The Saddle-Point Method in Differential Privacy

Wael Alghamdi [1] **Juan Felipe Gomez** [1] **Shahab Asoodeh** [2] **Flavio P. Calmon** [1] **Oliver Kosut** [3] **Lalitha Sankar** [3]

Abstract

We characterize the differential privacy guarantees of privacy mechanisms in the large-composition regime, i.e., when a privacy mechanism is sequentially applied a large number of times to sensitive data. Via exponentially tilting the privacy loss random variable, we derive a new formula for the privacy curve expressing it as a contour integral over an integration path that runs parallel to the imaginary axis with a free real-axis intercept. Then, using the method of steepest descent from mathematical physics, we demonstrate that the choice of *saddle-point* as the real-axis intercept yields closed-form accurate approximations of the desired contour integral. This procedure—dubbed the saddle-point accountant (SPA)—yields a constant-time accurate approximation of the privacy curve. Theoretically, our results can be viewed as a refinement of both Gaussian Differential Privacy and the moments accountant method found in Rényi Differential Privacy. In practice, we demonstrate through numerical experiments that the SPA provides a precise approximation of privacy guarantees competitive with purely numerical-based methods (such as FFT-based accountants), while enjoying closed-form mathematical expressions.

1. Introduction

Differential privacy (DP) is a widely adopted standard for privacy-preserving machine learning (ML). Differentially private mechanisms used in ML tasks typically operate in the *large-composition regime*, where mechanisms are sequentially applied many times to sensitive data. For example, when training neural networks using stochastic gradient descent, DP can be ensured by clipping and adding Gaussian noise to each gradient update (Abadi et al., 2016). Here, a DP mechanism (gradient clipping plus noise) is applied hundreds or thousands of times to the training data.

Quantifying the privacy loss after a large number of compositions of DP mechanisms is a central challenge in privacy-preserving ML. A key result by Murtagh & Vadhan (2016, Theorem 1.5) states that computing exact privacy parameters under composition is in general #P-complete, hence infeasible. This challenge has spurred several follow-up works on *privacy accounting*, e.g., (Dong et al., 2022; Koskela et al., 2020; Koskela & Honkela, 2021; Koskela et al., 2021; Gopi et al., 2021; Ghazi et al., 2022; Doroshenko et al., 2022), which compute upper bounds on the privacy budget parameters (ε, δ) in DP (see (6) for a formal definition).

The currently available accountants have several limitations. The accountants that have closed-form formulas—thereby attaining constant (in composition) runtimes—such as the moments accountant (Abadi et al., 2016; Mironov, 2017) and the CLT-based Gaussian-DP accountant (Bu et al., 2020), suffer from either overestimating or underestimating, respectively, the privacy parameters. On the other hand, convolution-based accountants, such as FFT-based approaches (Koskela et al., 2020; Gopi et al., 2021), while working well in practice, do not have constant runtimes, cannot generate the full privacy curve, and are limited by machine precision due to their purely numerical nature.[1] For example, existing implementations of the FFT-based approaches fail to estimate values of δ below 10^{-10} (Gopi et al., 2021, Appendix B) or 10^{-12} (Doroshenko et al., 2022, Appendix C).

We overcome these challenges by introducing a new approach for estimating DP parameters using complex analysis. Our approach is based on the method of steepest descent for integral approximation—a well-known method in mathematical physics (Jeffreys & Jeffreys, 1999). We derive the *saddle-point accountant* (SPA),[2] which:

[1]School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts, USA [2]Department of Computing and Software, McMaster University, Hamilton, Ontario, Canada [3]School of Electrical, Computer, and Energy Engineering, Arizona State University, Tempe, Arizona, USA. Correspondence to: Wael Alghamdi <alghamdi@g.harvard.edu>.

Proceedings of the 40^{th} International Conference on Machine Learning, Honolulu, Hawaii, USA. PMLR 202, 2023.

[1]Of course, this limitation can be alleviated by using custom implementations and arbitrary float-point precision libraries. Our point is that closed-form formulas do not have this limitation.

[2]We provide a Python implementation of the proposed SPA at https://github.com/Felipe-Gomez/saddlepoint_accountant

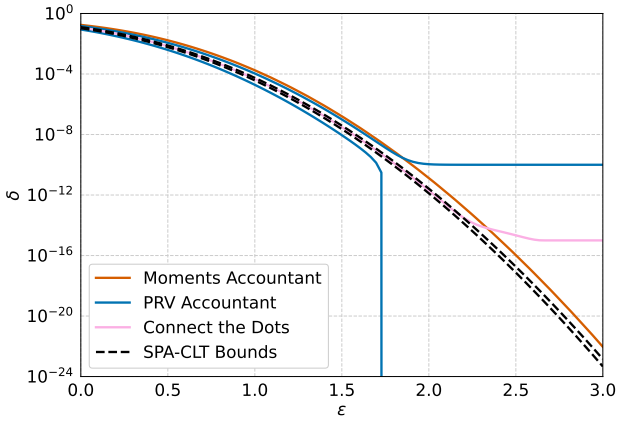

Figure 1. Accounting for the composition of 3000 subsampled Gaussian mechanisms, with noise scale $\sigma = 2$ and subsampling rate $\lambda = 0.01$. The remaining FFT discretization parameters are set[4] to $\varepsilon_{\text{error}} = 0.07$, $\delta_{\text{error}} = 10^{-10}$ for the PRV Accountant (Gopi et al., 2021), and discretization interval length of 2×10^{-4} for Connect the Dots (Doroshenko et al., 2022).

1) has a computable closed-form formula, hence enjoys constant runtime complexity in the number of compositions;

2) estimates the privacy parameters accurately and with provable error bounds; and

3) works for any value of δ, however small, thus describing the full range of (ε, δ) guarantees.

We illustrate the above properties of the SPA in Figure 1, which shows a comparison between the SPA and the state-of-the-art (SOTA) DP accountants when computing the (ε, δ) curve of a composition of 3000 subsampled Gaussian mechanisms. Only the moments accountant and the SPA are able to trace the whole privacy curve (see for example the region $\delta > 10^{-15}$). Further, the SPA upper and lower bounds have a narrow gap between them.

The SPA combines large-deviation and central-limit approaches for bounding expectations of sums of independent random variables, thereby attaining the best of both worlds. The large deviation approach uses the moment-generating function to approximate the probability of very unlikely events. The central limit theorem (CLT) approximates a random variable by a Gaussian with the same mean and variance. For DP accounting, the large deviation approach led to the moments accountant (Abadi et al., 2016); the CLT approach led to Gaussian-DP (Sommer et al., 2019; Dong et al., 2022). Both these accountant methods can be computed in constant time, but their accuracy is far less than the SOTA FFT accountant (Gopi et al., 2021). The saddle-point method can be viewed as a combination of two basic approaches: maintaining from large deviations the ability to handle very small values of δ, as well as the precise guarantees of the CLT. The resulting SPA achieves better accuracy than either approach on its own, while maintaining the optimality of the runtime complexity.

A brief overview of the SPA. Suppose that a DP mechanism has a privacy loss random variable whose cumulant-generating function $K(t)$ is finite for positive values of t (see Section 2 for precise definitions). Note that $K(t)$ is a familiar quantity used in DP accounting; for instance, it can be verified that the mechanism satisfies exactly $(t+1, K(t)/t)$-Rényi-DP for each $t > 0$ (Mironov, 2017). The SPA performs the following steps to estimate δ given ε:

1) Set $F(t) \triangleq K(t) - \varepsilon t - \log t - \log(t+1)$,
2) solve $F'(t) = 0$ over $t > 0$,
3) return $\delta(\varepsilon) \approx e^{F(t)} / \sqrt{2\pi F''(t)}$.

From this general workflow, it is clear that the SPA runs in *constant time* for n-fold self-composition; indeed, the cumulant-generating function for the composition is nK. Moreover, the root-finding in step 2 is similar to the one performed in the moments accountant (Abadi et al., 2016), which solves $K'(t) - \varepsilon = 0$ instead.

We refer to the approximation returned by this simple procedure as SPA-MSD.[5] The reason SPA-MSD approximates the privacy curve well is the following three steps. First, we express the privacy curve as the following contour integral:

$$\delta(\varepsilon) = \frac{1}{2\pi i} \int_{t-i\infty}^{t+i\infty} e^{F(z)}\, dz, \tag{1}$$

which holds *independently* of the choice of $t > 0$. Second, we apply the method of steepest descent, which uses a judicious choice of the integration path in the complex plane: the line parallel to the imaginary axis with real part corresponding to the *saddle-point* of the integrand, i.e., the unique point $t > 0$ for which $F'(t) = 0$. This approach leads to a new series expansion for δ given a fixed ε (see (26)), where the first term of this series corresponding to the approximation in step 3 above. This new expression for the privacy curve is our first main contribution.

Our experiments demonstrate that the SPA-MSD approximation is very accurate and can consistently achieve relative errors below 0.1% in ε for a fixed δ (see Figure 3). However, this approach *does not* provide a provable upper bound on the privacy curve—only an approximation. Consequently, we introduce another SPA, named SPA-CLT, where we first expand the K term in the integrand in (1) as an Edgeworth

[4] Decreasing these error parameters makes the PRV accountant more accurate, and we emphasize that the reason we include the PRV accountant in this plot is that it serves as a proxy for the ground truth.

[5] MSD stands for "method of steepest descent."

series (Hall, 2013), then apply the Berry-Esseen theorem to prove upper and lower bounds on the privacy curve. This procedure is equivalent to applying CLT to a tilted version of the privacy loss random variable.

The SPA-CLT amounts to replacing step 3 above by a slightly different approximation given in Proposition 5.3. This second approximation also enjoys constant runtime, yields provable and accurate upper bounds for the privacy curve even for very small values of δ, and is our second main contribution.

Finally, our third main contribution is an asymptotically tight DP composition theorem (see Theorem 4.1) which is useful in the error rate analysis for the SPA and is of independent interest.

The rest of the paper is organized as follows. Preliminaries on DP and the method of steepest descent are recalled in Section 2. We derive a contour-integral formula and an asymptotic expansion for the privacy curve in Section 3. This asymptotic expansion gives rise to heuristics for approximating the privacy curve, which leads to the the SPA-MSD method in Section 3.3. Then, we derive a tight composition theorem and the decay rate of the saddle-point in Section 4. In Section 5, we introduce the SPA-CLT (the second version of the SPA) and apply the results from Section 4 to derive rigorous bounds on the privacy curve. All proofs can be found in the Appendices.

2. Preliminaries

We collect in this section some of the required background on differential privacy, the method of steepest descent, and exponential tilting. We also prove a useful inequality on subsampling in Lemma 2.2, and we clarify our notation and assumptions.

2.1. Notation

For a random variable L, the moment-generating function (MGF) is denoted by $M_L(t) \triangleq \mathbb{E}[e^{tL}]$, and the cumulant-generating function (CGF) by $K_L(t) \triangleq \log M_L(t)$. The hockey-stick divergence (with parameter $\gamma \geq 1$) of a probability measure P from another Q is defined as

$$\mathsf{E}_\gamma(P \,\|\, Q) \triangleq \sup_{B \text{ Borel}} P(B) - \gamma Q(B). \tag{2}$$

If $X \sim P$ and $Y \sim Q$ are random variables, we also write $\mathsf{E}_\gamma(X \,\|\, Y) \triangleq \mathsf{E}_\gamma(P \,\|\, Q)$. The standard-normal cumulative distribution function is denoted by Φ. The Q function is defined by $Q(x) \triangleq 1 - \Phi(x)$. We also denote the function $q : \mathbb{R} \to (0, \infty)$ by

$$q(z) \triangleq Q(z) \cdot \sqrt{2\pi}\, e^{z^2/2}. \tag{3}$$

The (m, k)-th partial Bell polynomial is denoted by (with $\boldsymbol{x} = (x_1, \cdots, x_m)$)

$$B_{m,k}(\boldsymbol{x}) \triangleq \sum_{\substack{k_1+\cdots+k_m=k \\ 1\cdot k_1+\cdots+m\cdot k_m=m}} \binom{m}{k_1, \cdots, k_m} \prod_{j=1}^{m} \left(\frac{x_j}{j!}\right)^{k_j} \tag{4}$$

where the sum runs over nonnegative integers k_j, and the m-th complete Bell polynomial by

$$B_m(\boldsymbol{x}) \triangleq \sum_{k=1}^{m} B_{m,k}(\boldsymbol{x}). \tag{5}$$

The variance of a random variable X is denoted by σ_X^2. We will use the standard Bachmann-Landau notations $O, \Omega, \Theta, o, \omega$, and we will let $\asymp$ indicate the equivalence of order, i.e., $a_n \asymp b_n$ if and only if $a_n/b_n \to 1$ as $n \to \infty$. We will also write $f_n \sim \sum_{k\in\mathbb{N}} a_{n,k}$ to indicate an asymptotic expansion, i.e., the series might not converge but the first few partial sums approximate f_n well.

2.2. Differential Privacy

We review the basics of differential privacy, and derive a useful inequality for subsampling.

Definition 2.1 ((ε, δ)-DP (Dwork et al., 2006a;b; Zhu et al., 2022))**.** A mechanism (i.e., randomized algorithm) $\mathcal{M}$ is (ε, δ)-*differentially private* (DP) if, for every pair of neighboring datasets, denoted $D \simeq D'$, and event E,

$$\mathbb{P}[\mathcal{M}(D) \in E] - e^\varepsilon\, \mathbb{P}[\mathcal{M}(D') \in E] \leq \delta, \tag{6}$$

i.e., if

$$\sup_{D\simeq D'} \mathsf{E}_{e^\varepsilon}(\mathcal{M}(D) \,\|\, \mathcal{M}(D')) \leq \delta. \tag{7}$$

A pair of probability measures (P, Q) is called a *dominating pair* for $\mathcal{M}$ if, for every $\varepsilon \geq 0$, event E, and neighboring datasets $D \simeq D'$, the following inequality holds:

$$\mathbb{P}[\mathcal{M}(D) \in E] - e^\varepsilon\, \mathbb{P}[\mathcal{M}(D') \in E] \leq P(E) - e^\varepsilon Q(E). \tag{8}$$

If (8) is tight, i.e., if

$$\begin{aligned} \sup_{D\simeq D'} \mathbb{P}[\mathcal{M}(D) \in E] - e^\varepsilon\, \mathbb{P}[\mathcal{M}(D') \in E] \\ = P(E) - e^\varepsilon Q(E) \end{aligned} \tag{9}$$

for each fixed $\varepsilon \geq 0$, then (P, Q) is said to be a *tightly dominating pair*. For any dominating pair (P, Q) consisting of equivalent measures, we associate a *privacy loss random variable* (PLRV) that is defined as

$$L \triangleq \log \frac{dP}{dQ}(X), \qquad X \sim P. \tag{10}$$

It is not hard to see that a mechanism $\mathcal{M}$ having PLRV L will satisfy $(\varepsilon, \delta_L(\varepsilon))$-DP for every $\varepsilon \geq 0$, where we define the *privacy curve* (with $a^+ \triangleq \max(0, a)$)

$$\delta_L(\varepsilon) \triangleq \mathbb{E}\left[\left(1 - e^{\varepsilon - L}\right)^+\right]. \tag{11}$$

Composition of DP Mechanisms. The use of PLRVs facilitates DP accounting under composition. The *adaptive* composition of two mechanisms $\mathcal{M}_1$ and $\mathcal{M}_2$ is given by the mechanism

$$(\mathcal{M}_1 \circ \mathcal{M}_2)(D) \triangleq (\mathcal{M}_1(D), \mathcal{M}_2(D, \mathcal{M}_1(D))), \quad (12)$$

that is, $\mathcal{M}_2$ can look at both the dataset and the output of $\mathcal{M}_1$. Let $\mathcal{M}^{(n)} = \mathcal{M}_1 \circ \cdots \circ \mathcal{M}_n$ denote the adaptive composition of n, possibly distinct, mechanisms. We can form a PLRV for the composed mechanism that splits additively. In other words, $L^{(n)} \triangleq L_1 + \cdots + L_n$, where $L_1, \cdots, L_n$ are independent PLRVs for $\mathcal{M}_1, \cdots, \mathcal{M}_n$, respectively, is a PLRV for the composition $\mathcal{M}^{(n)}$ (Dong et al., 2022, Theorem 3.2). The ensuing privacy curve $\delta_{L^{(n)}}$ (as defined by (11)) gives a privacy guarantee for $\mathcal{M}^{(n)}$. Like all accounting methods cited herein, we focus on computing or approximating the curve $\delta_{L^{(n)}}$, which we refer to henceforth as the *composition curve*. We also denote the n-fold self-composition of a mechanism $\mathcal{M}$ by $\mathcal{M}^{\circ n}$, and in this case we may choose the L_j to be i.i.d.

Subsampling and DP-SGD. In the context of differentially-private stochastic gradient descent (DP-SGD), one applies a DP mechanism on a subset of the dataset. The fraction of the batch size over the size of the dataset is called the *subsampling rate*, denoted by λ. Subsampling is known to amplify the privacy guarantees (Balle et al., 2018). In this setting, with $\mathcal{M}_\lambda$ denoting the subsampled mechanism, one should bound both orders $\mathsf{E}_{e^\varepsilon}(\mathcal{M}(D) \,\|\, \mathcal{M}_\lambda(D'))$ and $\mathsf{E}_{e^\varepsilon}(\mathcal{M}_\lambda(D) \,\|\, \mathcal{M}(D'))$ to obtain the value of δ. In the following lemma, we show that in fact one order dominates. See Appendix A for the proof and further details on subsampling.

Lemma 2.2. *Fix a Borel probability measure P over $\mathbb{R}^n$ that is symmetric around the origin (i.e., $P(\mathcal{A}) = P(-\mathcal{A})$ for every Borel $\mathcal{A} \subset \mathbb{R}^n$), and fix constants $(s, \lambda, \gamma) \in \mathbb{R}^n \times [0,1] \times [1, \infty)$. Let $T_s P$ be the probability measure given by $(T_s P)(\mathcal{A}) = P(\mathcal{A} - s)$, and let $Q = (1-\lambda)P + \lambda T_s P$. We have the inequality $\mathsf{E}_\gamma(P \,\|\, Q) \le \mathsf{E}_\gamma(Q \,\|\, P)$, with equality if and only if $(\gamma - 1)\,\lambda\,\|s\|\,\mathsf{E}_\gamma(Q \,\|\, P) = 0$.*

Proof. See Appendix A. □

2.3. The Method of Steepest Descent

We give a brief overview of the method of steepest descent (see Appendix B for details). We need to compute

$$I_n = \frac{1}{2\pi i} \int_{t-i\infty}^{t+i\infty} e^{F_n(z)}\, dz \quad (13)$$

for a given F_n, provided that I_n is *independent* of the value of $t \in \mathbb{R}$. In a nutshell, the method of steepest descent is a powerful tool for choosing the best parameter t that renders the computation of I_n easiest. Namely, t is the *saddle-point* of F_n, defined as the unique solution to $F_n'(t_0) = 0$. Then, one would obtain the "asymptotic expansion":

$$I_n \sim \frac{e^{F_n(t_0)}}{\sqrt{2\pi F_n''(t_0)}} \left(1 + \sum_{m=2}^{\infty} \beta_{n,m}\right), \quad (14)$$

where we define the constants

$$\beta_{n,m} \triangleq \frac{(-1)^m B_{2m}(0, 0, F_n^{(3)}(t_0), \ldots, F_n^{(2m)}(t_0))}{2^m m! F_n''(t_0)^m}. \quad (15)$$

Recall that this does not mean that the above equation holds for I_n with equality for any particular n. Rather, it is a heuristic indicating the potential for the truncated expansion to give close approximations for the intended integral I_n.

In our application of the method of steepest descent to DP, we show in Theorem 3.1 that the privacy curve can be represented as the contour integral (13) for the choice of function

$$F_n(z) = K_{L^{(n)}}(z) - z\varepsilon - \log z - \log(1+z). \quad (16)$$

2.4. Exponential Tilting

An essential tool that we use for our theoretical analysis is exponential tilting of random variables, defined as follows.

Definition 2.3. The *exponential tilting* with parameter $t \in \mathbb{R}$ of a random variable L having a finite MGF at t is the random variable $\tilde{L}$ whose probability measure is given by

$$P_{\tilde{L}}(B) \triangleq \frac{1}{M_L(t)} \int_B e^{tx}\, dP_L(x) \quad (17)$$

for any Borel set B. If L has PDF p_L, then $\tilde{L}$ is given by its PDF

$$p_{\tilde{L}}(x) = \frac{e^{tx} p_L(x)}{M_L(t)}. \quad (18)$$

A simple key feature of exponential tilting, stated here without proof, is that it respects addition and independence.

Lemma 2.4. *For independent L_j, the exponential tilting of $L = L_1 + \cdots + L_n$ with parameter t is $\tilde{L} = \tilde{L}_1 + \cdots + \tilde{L}_n$, where $\tilde{L}_j$ is the exponential tilting of L_j with parameter t for each j. Further, $\tilde{L}_1, \ldots, \tilde{L}_n$ are independent too.*

2.5. Assumptions

We will require the PLRVs to have finite MGFs.

Assumption 2.5. The MGF $M_L(t)$ of the PLRV L is finite for every $t > 0$.

Under Assumption 2.5, both the MGF and CGF can be extended to be holomorphic functions over the half-plane $z \in (0, \infty) + i\mathbb{R} \subset \mathbb{C}$.

We impose the following technical assumption on the distribution of the PLRV so that Parseval's identity applies.

Assumption 2.6. The induced probability measure P_L by the PLRV L decomposes as a sum $P_L = Q_L + R_L$ for Q_L absolutely continuous with respect to the Lebesgue measure and discrete R_L. Further, with q_L denoting the PDF of Q_L, we assume that $x \mapsto e^{tx} q_L(x)^2$ is integrable for each $t > 0$.

For our error analysis, we will assume the following on the growth of the first three moments of a PLRV.

Assumption 2.7. With $\tilde{L} = \tilde{L}_1 + \cdots + \tilde{L}_n$ being the exponential tilting with parameter $t > 0$, and denoting

$$\mathrm{P}_t \triangleq \sum_{j=1}^{n} \mathbb{E}\left[\left|\tilde{L}_j - \mathbb{E}[\tilde{L}_j]\right|^3\right], \tag{19}$$

we assume that there are constants $\mathrm{KL}, \mathrm{V} > 0$, and P such that $t = o(n^{-1/3})$ yields the limit (as $n \to \infty$)

$$\frac{1}{n} \cdot (\mathbb{E}[\tilde{L}], \sigma^2_{\tilde{L}}, \mathrm{P}_t) \to (\mathrm{KL}, \mathrm{V}, \mathrm{P}). \tag{20}$$

A few remarks on the satisfiability of the above assumptions are in order.

Remark 2.8. Assumption 2.7 is automatically satisfied under Assumption 2.5 for *self-composition*.

Remark 2.9. It is worth noting that all of the above assumptions are satisfied by the usual continuous DP mechanisms, including both the subsampled Gaussian mechanism (because its PLRV is continuous with a PDF that decays super-exponentially) and the subsampled Laplace mechanism (because its PLRV's continuous part is bounded). See Appendix C for more details.

Remark 2.10. Although finiteness of the MGF rules out DP mechanisms whose PLRVs are infinite with nonzero probability (e.g., discrete mechanisms or compactly-supported mechanisms), our approach may be extended to encapsulate this case too. Specifically, the mass at infinity should be added to the value of δ directly. Indeed, we may rewrite the privacy curve δ_L via conditioning on the event $\{L < \infty\}$ as

$$\begin{aligned}\delta_L(\varepsilon) &= \mathbb{P}[L = \infty] + \mathbb{E}\left[(1 - e^{\varepsilon - L})^+ | L < \infty\right] \mathbb{P}[L < \infty] \\ &= \mathbb{P}[L = \infty] + \delta_Z(\varepsilon)\, \mathbb{P}[L < \infty],\end{aligned} \tag{21}$$

where Z is a random variable obtained from L via conditioning on the event $\{L < \infty\}$. Then, we may apply our methods on δ_Z and obtain results for the original curve δ_L in view of the relation (21).

Remark 2.11. It is not hard to see that the MGF M_L of the PLRV L is finite for any additive continuous mechanism with PDF of the form $e^{-g(x)}$ for a continuous g such that $g(x) \asymp \beta |x|^\alpha$ for some $\alpha, \beta > 0$. For such DP mechanisms, the MGF of the ensuing PLRV is finite at any $t > 0$ and for any subsampling rate $\lambda \in [0, 1]$.

3. New Representations of the Privacy Curve

In Theorem 3.1, we derive two new formulas for the privacy curve. Then, we apply the method of steepest descent to the contour-integral formula (23). This yields the asymptotic expansion (26) of the privacy curve, which is the basis for the SPA-MSD as given by Definition 3.6. Later, in Section 5, we derive rigorous bounds on a CLT-based approximation that is inspired by the approximations in the present section.

We assume that we have access to a PLRV L for mechanism $\mathcal{M}$ (see Definition 2.1). In most cases, the relevant variable is $L^{(n)} = L_1 + \cdots + L_n$, such that, as discussed in Section 2.2, $\delta_{L^{(n)}}$ is the composition curve for the adaptive composition $\mathcal{M}^{(n)} = \mathcal{M}_1 \circ \cdots \circ \mathcal{M}_n$ (and $L_1, \cdots, L_n$ are PLRVs for $\mathcal{M}_1, \cdots, \mathcal{M}_n$ that are independent). However, in this section we derive formulas for the privacy curve δ_L for any variable L. We note that for these formulas to be numerically computable, it suffices that the distribution of L be known to an extent that the derivatives of the MGF $M_L^{(k)}(t)$ can be computed.

3.1. The Privacy Curve as a Contour Integral

The privacy curve is defined in (11) as the expectation $\delta_L(\varepsilon) = \mathbb{E}[f(L)]$, where $f(x) = (1 - e^{\varepsilon - x})^+$. We want to transform this integral—via Parseval's identity—into the frequency domain. However, as $f \notin L^1(\mathbb{R})$, we cannot directly apply Parseval's identity. Nevertheless, exponentially tilting L, we may replace $f(x)$ by $e^{-tx} f(x)$, which decays fast. We carry out the details of this idea in Appendix E to obtain the following new formulas for δ_L.

Theorem 3.1. *If the PLRV L satisfies Assumption 2.5, then, for every $t > 0$, we may write the privacy curve as*

$$\delta_L(\varepsilon) = M_L(t)\, \mathbb{E}\left[e^{-t\tilde{L}} \left(1 - e^{\varepsilon - \tilde{L}}\right)^+\right] \tag{22}$$

for all $\varepsilon \geq 0$, where $\tilde{L}$ is the exponential tilting of L with parameter t (see Definition 2.3). If, in addition, L satisfies Assumption 2.6, then we also have the formula[6]

$$\delta_L(\varepsilon) = \frac{1}{2\pi i} \int_{t - i\infty}^{t + i\infty} e^{F_\varepsilon(z)}\, dz \tag{23}$$

for all $\varepsilon \geq 0$, where we define the exponent by[7]

$$F_\varepsilon(z) \triangleq K_L(z) - z\varepsilon - \log z - \log(1 + z). \tag{24}$$

Proof. See Appendix E. □

[6] The independence of formula (23) of t is not surprising, given Cauchy's integration theorem. More importantly, the theorem states that an integration path with real part t is actually equivalent to exponential tilting with parameter t.

[7] We use the principal branch of the complex logarithm, so F_ε is well defined and analytic over the half-plane $z \in (0, \infty) + i\mathbb{R}$.

The two formulas in (22)–(23) lead to two paths for approximating δ_L. The first is a direct application of the method of steepest descent, where F_ε is expanded around the saddle-point (see Section 2.3). The second simply approximates the expectation formula in (22) via the CLT, by replacing $\tilde{L}$ with a Gaussian. The first path (described next) leads to better approximations numerically, but the second path is more amenable to an error analysis (see Section 5).

3.2. The Privacy Curve in Terms of Bell Polynomials

As we have proved a formula in (23) for the privacy curve δ_L representing it as a contour integral like in (13), we can now apply the method of steepest descent to approximate it. Recall from Section 2.3 that the best choice for the real-axis intercept in (23) is the saddle-point.

Definition 3.2. The *saddle-point* associated with a PLRV L satisfying Assumption 2.5 and a privacy parameter ε satisfying $\varepsilon < \operatorname{ess\,sup} L$ is the unique $t_0 > 0$ such that $F'_\varepsilon(t_0) = 0$, or equivalently

$$K'_L(t_0) = \varepsilon + \frac{1}{t_0} + \frac{1}{t_0 + 1}. \tag{25}$$

Remark 3.3. The original moments accountant aims to solve $K'_L(t) = \varepsilon$, indicating the connection between the moments accountant and the SPA, introduced formally in Section 3.3.

Remark 3.4. We show in Appendix D that the saddle-point, as given by Definition 3.2, is indeed well-defined.

Applying the method of steepest descent to the contour integral in (23) with the choice of t being the saddle-point, we obtain the following asymptotic expansion for the privacy curve in terms of the derivatives of the MGF, connected via Bell polynomials (see Section 2.3).

Heuristic 3.5. *Let L be a PLRV satisfying Assumption 2.5. Then, for any $\varepsilon \in [0, \operatorname{ess\,sup} L)$, and with t_0 denoting the associated saddle-point, we have the asymptotic expansion*

$$\delta_L(\varepsilon) \sim \frac{e^{F_\varepsilon(t_0)}}{\sqrt{2\pi F''_\varepsilon(t_0)}} \left(1 + \sum_{m=2}^{\infty} \beta_{\varepsilon,m}\right), \tag{26}$$

where, with $B_k(x_1, \ldots, x_k)$ denoting the k-th Bell polynomial and $F_\varepsilon^{(k)}$ the k-th derivative, we denote the constants

$$\beta_{\varepsilon,m} \triangleq \frac{(-1)^m B_{2m}(0, 0, F_\varepsilon^{(3)}(t_0), \ldots, F_\varepsilon^{(2m)}(t_0))}{2^m m! F''_\varepsilon(t_0)^m}. \tag{27}$$

Further, with $B_{k,j}(x_1, \cdots, x_k)$ denoting the (k, j)-th partial Bell polynomial, the derivatives of F_ε are[8] *(for $k \geq 2$)*

$$F_\varepsilon^{(k)}(t_0) = (-1)^{k-1}(k-1)! \left(\frac{1}{t_0^k} + \frac{1}{(t_0+1)^k}\right) + \sum_{j=1}^{k} \frac{(-1)^{j-1}(j-1)!}{M_L(t_0)^j} B_{k,j}(M'_L(t_0), \cdots, M_L^{(k)}(t_0)). \tag{28}$$

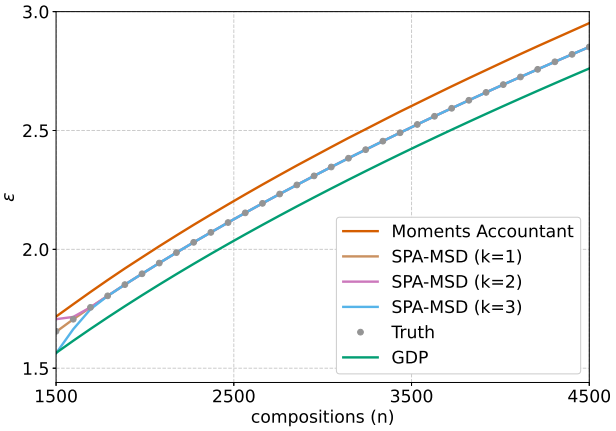

Figure 2. Privacy budget ε of the subsampled Gaussian mechanism after $1500 \leq n \leq 4500$ compositions using the proposed SPA-MSD (29) and the other closed-form accountants. We use the subsampling rate $\lambda = 0.01$, noise scale $\sigma = 2$, and $\delta = 10^{-15}$.

3.3. Application: The Saddle-Point Accountant

Based on the asymptotic expansion in (26), we can derive various approximations of δ_L depending on how many terms we keep. This leads to the following versions of the *saddle-point accountant* (SPA).

Definition 3.6. The *order-k method-of-steepest-descent saddle-point accountant* (SPA-MSD) for the mechanism $\mathcal{M}$ with PLRV L satisfying Assumption 2.5 is defined by

$$\delta^{(k)}_{L,\,\text{SP-MSD}}(\varepsilon) \triangleq \frac{e^{F_\varepsilon(t_0)}}{\sqrt{2\pi F''_\varepsilon(t_0)}} \left(1 + \sum_{m=2}^{k} \beta_{\varepsilon,m}\right) \tag{29}$$

when $\varepsilon < \operatorname{ess\,sup} L$, where $t_0 > 0$ is the saddle-point (i.e., $F'_\varepsilon(t_0) = 0$), and we set $\delta^{(k)}_{L,\,\text{SP-MSD}}(\varepsilon) = 0$ if $\varepsilon \geq \operatorname{ess\,sup} L$. Here, the $\beta_{\varepsilon,m}$ are as defined in (27).

The first SPA-MSD is

$$\delta^{(1)}_{L,\,\text{SP-MSD}}(\varepsilon) = \frac{e^{F_\varepsilon(t_0)}}{\sqrt{2\pi F''_\varepsilon(t_0)}}, \tag{30}$$

[8] The formula for $F_\varepsilon^{(k)}$ follows immediately by Faà di Bruno's formula for the derivatives of composition of functions.

which can be expanded using the definition of F_ε as

$$\delta^{(1)}_{L,\text{SP-MSD}}(\varepsilon) = \frac{e^{K_L(t_0)-\varepsilon t_0}}{\sqrt{2\pi}\sqrt{t_0(t_0+1)K_L''(t_0)+t_0^2+(1+t_0)^2}}. \quad (31)$$

The order-2 SPA-MSD is given by

$$\delta^{(2)}_{L,\text{SP-MSD}}(\varepsilon) = \frac{e^{F_\varepsilon(t_0)}}{\sqrt{2\pi F_\varepsilon''(t_0)}}\left(1+\frac{1}{8}\frac{F_\varepsilon^{(4)}(t_0)}{F_\varepsilon''(t_0)^2}\right), \quad (32)$$

and the order-3 SPA-MSD is given by

$$\delta^{(3)}_{L,\text{SP-MSD}}(\varepsilon) = \frac{e^{F_\varepsilon(t_0)}}{\sqrt{2\pi F_\varepsilon''(t_0)}}\times \left(1+\frac{1}{8}\frac{F_\varepsilon^{(4)}(t_0)}{F_\varepsilon''(t_0)^2}-\frac{5}{24}\frac{F_\varepsilon^{(3)}(t_0)^2}{F_\varepsilon''(t_0)^3}-\frac{1}{48}\frac{F_\varepsilon^{(6)}(t_0)}{F_\varepsilon''(t_0)^3}\right). \quad (33)$$

Empirical Accuracy of SPA-MSD. The expressions for the SPA-MSD displayed in (31)–(33) can traverse privacy curves that are virtually indistinguishable from the ground-truth. We illustrate this in Figure 2 for the subsampled Gaussian, where we estimate ε (for fixed $\delta = 10^{-15}$) under a varying number of compositions. In this experiment, SPA-MSD improves on the other closed-form accountants (which run in constant time). Hence, SPA-MSD can be seen a correction to both the large deviation method and the CLT-based method found in the Moments Accountant and Gaussian-DP, respectively. A zoomed-in version of this figure is given in Figure 3, which shows the relative errors. See Appendix L for the SPA-MSD pseudocode, Appendix M for computing the ground-truth in Figure 2, and Appendix N for more experiments.

4. Asymptotically Tight Composition Theorem

We show next that the lowest ε under composition cannot deviate from the mean of the PLRV by a large multiple of the standard deviation of the PLRV. This result is used afterwards to derive the asymptotic behavior of the saddle-point. The asymptotic behavior of the saddle-point, in turn, will be helpful in the next section to derive rigorous bounds on the SPA approximation error. We prove the following asymptotically tight DP composition theorem.

Theorem 4.1. *Let $\mathcal{M} = \mathcal{M}_1 \circ \cdots \circ \mathcal{M}_n$ have a PLRV $L = L_1 + \cdots + L_n$, where the L_j are PLRVs for the $\mathcal{M}_j$ that are independent. Assume that the L_j have finite absolute third moments, and $\mathrm{P}_0 = o(\sigma_L^3)$ as $n \to \infty$ (see (19)). Let $\delta \in (0, 1/2)$ be such that $\limsup \delta < 1/2$ (so δ is allowed to vary with n). If $\sigma_L/(-\Phi^{-1}(\delta)) \to \infty$ as $n \to \infty$, then $\mathcal{M}$ is $(\mathbb{E}[L] - \Phi^{-1}(\delta)\sigma_L, \delta\cdot(1+o(1)))$-DP. Conversely, this result is tight in the following sense. If $\delta_0 \in (0, 1/2)$ is fixed, $\sigma_L \to \infty$, and $\mathcal{M}$ is $(\mathbb{E}[L] + b\sigma_L, \delta_0 + o(1))$-DP, then we must have $\liminf b \geq -\Phi^{-1}(\delta_0)$.*

Proof. See Appendix F. □

A more compact way to state the constant-δ claim in the theorem is that, for any fixed $\delta \in (0, 1/2)$, we have

$$\delta_L(\mathbb{E}[L] - \Phi^{-1}(\delta)\sigma_L) \to \delta. \quad (34)$$

For example, Theorem 4.1 implies that $\delta_L(\varepsilon)$ is close to 10^{-10} if and only if ε is around $\mathbb{E}[L] + 6.4\sigma_L$ for all large n, since $-\Phi^{-1}(10^{-10}) \approx 6.4$. Thus, if one hopes to have a small value of δ, the only "interesting" values of ε, in the regime of high n, are those that are above $\mathbb{E}[L]$ by the derived multiple of σ_L.

4.1. Asymptotic Formula for the Saddle-Point

We re-parameterize $\varepsilon = \mathbb{E}[L] + b\sigma_L$, so b can be seen as the "Z-score" of ε, which is justified by Theorem 4.1. For this regime of values of ε, we prove the following asymptotic characterization of the saddle-point.

Theorem 4.2. *Let $L = L_1 + \cdots + L_n$ for independent L_j satisfying Assumption 2.5, and suppose that $(\mathbb{E}[L], \sigma_L^2) \asymp n \cdot (\mathrm{KL}, \mathrm{V})$ for some constants $\mathrm{KL}, \mathrm{V} > 0$. Let $\varepsilon = \mathbb{E}[L] + b\sigma_L$, where $b > 0$ satisfies $b = o(n^{1/6})$, and assume that $\varepsilon < \operatorname{ess\,sup} L$. Then, the value of the saddle-point (as given by Definition 3.2) satisfies the asymptotic relation*

$$t_0 \asymp \frac{b + \sqrt{b^2+4}}{2\sigma_L}. \quad (35)$$

Proof. See Appendix G. □

This asymptotic formula for the saddle-point will be useful in deriving the asymptotic rate of the approximation error of the SPA in the next section.

5. CLT Error Bound Analysis

While the approximations of Section 3 are often very precise (see Figure 2), they are merely *approximations*, and do not provide any hard guarantees on the (ε, δ)-DP of a given mechanism. In this section, we derive the alternative form of the SPA by applying the Berry-Esseen theorem to the saddle-point exponentially tilted PLRV, thereby obtaining upper and lower bounds on the achieved privacy parameters.

5.1. CLT Based Version of the SPA

We return to the expectation based formula for δ_L shown in Theorem 3.1, which can be rewritten as

$$\delta_L(\varepsilon) = e^{K_L(t)-\varepsilon t}\,\mathbb{E}\left[\bar{f}\left(\tilde{L}-\varepsilon, t\right)\right], \quad (36)$$

where

$$\bar{f}(x,t) \triangleq e^{-xt}\left(1-e^{-x}\right)^{+}, \tag{37}$$

with $t > 0$ varying freely and $\tilde{L}$ being the exponential tilting of L with parameter t. Here, $L = L_1 + \cdots + L_n$ for independent L_j satisfying Assumption 2.5. We will simply replace $\tilde{L}$ by a Gaussian with the same first two moments,[9] and choose t to be the saddle-point of L as per Definition 3.2. Thus, we introduce the following version of the SPA.

Definition 5.1. Under Assumption 2.5, the *CLT version of the saddle-point accountant* (SPA-CLT) is defined by

$$\delta_{L,\,\text{SP-CLT}}(\varepsilon) \triangleq e^{K_L(t_0)-\varepsilon t_0}\,\mathbb{E}\left[\bar{f}(Z-\varepsilon,t_0)\right] \tag{38}$$

if $\varepsilon < \operatorname{ess\,sup} L$, where $Z \sim \mathcal{N}\left(K_L'(t_0), K_L''(t_0)\right)$, and t_0 is the saddle-point for L as given by Definition 3.2. We define $\delta_{L,\,\text{SP-CLT}}(\varepsilon) = 0$ for $\varepsilon \geq \operatorname{ess\,sup} L$.

Remark 5.2. The approach giving rise to $\delta_{L,\,\text{SP-CLT}}$ can be seen as a series expansion of the $e^{K_L(z)}$ part of the integrand in Theorem 3.1, or equivalently as an (order-0) Edgeworth expansion (Hall, 2013) of the distribution of $\tilde{L}$. However, the Edgeworth expansion approach delineated herein is different from what can be found in the DP literature (Wang et al., 2022). Specifically, we apply the Edgeworth expansion on the *tilted* random variable $\tilde{L}$, whereas the approach of Wang et al. (2022) uses the Edgeworth expansion of the *non-tilted* version L. This distinction can yield very different approximations. We include a comparison between our approach and the standard CLT in Appendix H.

The following result expresses the CLT-based SPA in terms of easily computable functions. In what follows, we let $\delta_{L,\,\text{SP-CLT}}(\varepsilon;t)$ denote the same expression as in (38) but with t_0 replaced by a free $t > 0$, i.e.,

$$\delta_{L,\,\text{SP-CLT}}(\varepsilon;t) \triangleq e^{K_L(t)-\varepsilon t}\mathbb{E}\left[\bar{f}(Z-\varepsilon,t)\right] 1_{[0,\operatorname{ess\,sup} L)}(\varepsilon) \tag{39}$$

where $Z \sim \mathcal{N}(K_L'(t), K_L''(t))$. In particular $\delta_{L,\,\text{SP-CLT}}(\varepsilon) = \delta_{L,\,\text{SP-CLT}}(\varepsilon;t_0)$ for t_0 the saddle-point.

Proposition 5.3. *Suppose Assumption 2.5 holds. Fix any $t > 0$ and $\varepsilon \in [0, \operatorname{ess\,sup} L)$, and denote*

$$\begin{aligned}\gamma &\triangleq \frac{K_L'(t)-\varepsilon}{\sqrt{K_L''(t)}}, \quad \alpha \triangleq \sqrt{K_L''(t)}\,t - \gamma,\\ \beta &\triangleq \sqrt{K_L''(t)}\,(t+1) - \gamma.\end{aligned} \tag{40}$$

Then, we have that (with q as defined in (3))

$$\delta_{L,\,\text{SP-CLT}}(\varepsilon;t) = \frac{q(\alpha)-q(\beta)}{\sqrt{2\pi}} e^{K_L(t)-\varepsilon t-\gamma^2/2}. \tag{41}$$

Proof. See Appendix I.1. □

[9] It is not hard to see that the mean and variance of $\tilde{L}$ are given by $\mathbb{E}[\tilde{L}] = K_L'(t)$ and $\sigma_{\tilde{L}}^2 = K_L''(t)$.

Remark 5.4. It holds that $0 < q(z) < \min(1/z, \sqrt{\pi/2})$ for all $z > 0$, and $q(z) \asymp 1/z$ as $z \to \infty$ (NIS, Section 7.8).

While the two methods of approximation—the steepest descent as in Section 3.3, and the CLT approach in this section—lead to different approximations, these two approximations are closely related, as described by the following simple inequality.

Proposition 5.5. *Under Assumption 2.5, for any $t > 0$*

$$\delta_{L,\,\text{SP-CLT}}(\varepsilon;t) \leq \frac{e^{F_\varepsilon(t)}}{\sqrt{2\pi K_L''(t)}}. \tag{42}$$

Proof. See Appendix I.2. □

Note that the only difference between the right-hand side of (42) and $\delta^{(1)}_{L,\,\text{SP-MSD}}(\varepsilon)$ is that the denominator involves K_L'' instead of F_ε''.

5.2. Finite-Composition Error Bound

Using the Berry-Esseen theorem, we prove the following theorem for the error bounds on the approximation $\delta_{L,\,\text{SP-CLT}}$ for arbitrary tilts.

Theorem 5.6. *Suppose Assumption 2.5 holds. For any $t > 0$ and $\varepsilon \geq 0$, there is a $\zeta \in [-1,1]$ such that*

$$\begin{aligned}\delta_L(\varepsilon) = e^{K_L(t)-\varepsilon t}\,\mathbb{E}\left[e^{-t(Z-\varepsilon)}\left(1-e^{-(Z-\varepsilon)}\right)^{+}\right]\\ + \zeta\,\text{err}_{\text{SP}}(\varepsilon;t),\end{aligned} \tag{43}$$

where $Z \sim \mathcal{N}(K_L'(t), K_L''(t))$ and the error is defined by

$$\text{err}_{\text{SP}}(\varepsilon;t) \triangleq e^{K_L(t)-\varepsilon t}\frac{t^t}{(1+t)^{1+t}} \cdot \frac{1.12\,\text{P}_t}{K_L''(t)^{3/2}}. \tag{44}$$

Proof. See Appendix J. □

Note that omitting the ζ term in the right-hand side of (43) gives exactly $\delta_{L,\,\text{SP-CLT}}(\varepsilon;t)$ as per Definition 5.1. Thus, Theorem 5.6 can be equivalently restated as the following error bound for SPA-CLT: for each $t > 0$ and $\varepsilon \geq 0$,

$$|\delta_L(\varepsilon) - \delta_{L,\,\text{SP-CLT}}(\varepsilon;t)| \leq \text{err}_{\text{SP}}(\varepsilon;t). \tag{45}$$

5.3. Asymptotic Error Rate

While Theorem 5.6 holds for any positive value of t around which the MGF is finite, a natural choice of t is the saddle-point t_0 itself, defined as the solution to (25). We analyze the ensuing error rate for this particular choice of tilt next.

Specifically, we show that the error rate in approximating δ_L by $\delta_{L,\,\text{SP-CLT}}$ decays roughly at least as fast as $1/(\sqrt{n}\,e^{b^2/2})$ for the choice $\varepsilon = \mathbb{E}[L] + b\sigma_L$, and we characterize the constant term too.

Theorem 5.7. *Let $L = L_1 + \cdots + L_n$ for independent PLRVs $L_1, \cdots, L_n$ that satisfy Assumption 2.5. Suppose that Assumption 2.7 holds too. Let $\varepsilon = \mathbb{E}[L] + b\sigma_L$ for $b > 0$ satisfying $b = o(n^{1/6})$, and let t_0 be the saddle-point of L (see Definition 3.2). Then, as $n \to \infty$, we have*

$$\mathrm{err}_{\mathrm{SP}}(\varepsilon; t_0) \asymp \frac{1.12\sqrt{e}\,\mathrm{P}}{\mathrm{V}^{3/2} \cdot C(b)^{\tau} \cdot \sqrt{n}}, \tag{46}$$

where $\tau < 1$ satisfies $\tau \to 1$, and we define the term $C(b) \triangleq \exp\left((b^2 + b\sqrt{b^2+4})/4\right)$. Furthermore, writing $t_0 = \tau_0 \cdot \frac{b+\sqrt{b^2+4}}{2\sigma_L}$, we may take $\tau = (2-\tau_0)\tau_0$ in (46).

Proof. See Appendix K. □

Remark 5.8. In Appendix H, we illustrate the benefit of tilting the PLRV by comparing the error term in (46) with the corresponding standard CLT error (i.e., without tilting). For example, for a limiting value of $\delta = 10^{-10}$, the error term incurred by our tilting-based approach is roughly 9-orders of magnitude smaller than the standard approach without tilting.

5.4. Relative-Error Comparisons

The SPA-CLT approximation (41) and its error bound (44) can approximate the privacy parameters accurately. In Figure 3, we plot the relative error[10] in estimating ε given $\delta = 10^{-15}$ incurred by SPA-CLT (both for the approximation in (41) and the approximation $\pm$ the error term (44)), SPA-MSD (for comparison), and the other closed-form accountants. The setting is for the subsampled Gaussian mechanism, with the same parameters as in Figure 2. Here, SPA improves on both the moments accountant and Gaussian-DP.

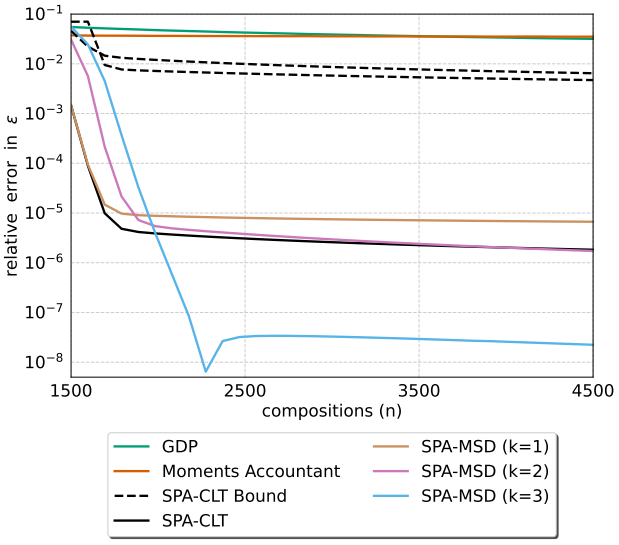

Figure 3. Accounting for the privacy budget ε, given $\delta = 10^{-15}$, for the subsampled Gaussian mechanism, with subsampling rate $\lambda = 0.01$, and noise scale $\sigma = 2$. We plot the relative error in estimating ε (i.e., $|1 - \hat{\varepsilon}(\delta)/\varepsilon(\delta)|$ for an estimate $\hat{\varepsilon}$) versus the number of compositions, n. the SPA outperforms the other closed-form accountants for this experiment.

6. Conclusion and Open Problems

We introduce a novel application of the method of steepest descent in DP. First, using the exponentially-tilted version of the PLRV, we derive new formulas for the privacy curve (Theorem 3.1). Inspired by the method of steepest descent, we fix the exponential tilt to be the saddle-point of the integrand's exponent. This amounts to solving the scalar equation

$$K_L'(t) = \varepsilon + \frac{1}{t} + \frac{1}{t+1}. \tag{47}$$

The ensuing closed-form formulas provide constant-runtime accurate approximations that can traverse the full privacy curve (e.g., for $\delta < 10^{-10}$). Our approach can be seen as a correction to both large-deviation methods (e.g., the moments accountant, via the additional $1/t + 1/(t+1)$ term) and CLT-based methods (e.g., Gaussian-DP, via preprocessing the PLRV with exponential tilting). This way, we retain the constant runtime of closed-form accountants without sacrificing accuracy, as demonstrated by our experiments.

The saddle-point approach leaves a few questions open. The relative-error plot in Figure 3 indicates that, while the SPA-CLT bounds achieve reasonable relative error, the original approximation given by SPA-CLT and SPA-MSD seem to be several orders of magnitudes more accurate than can be captured by the bounds we derive herein. Hence, it is an interesting future line of work to refine our bounds to further reveal the power of the saddle-point approximation. One promising path towards such a refinement might be through finding mechanism-specific bounds. Relatedly, such finer bounds would shed light on the question of "how large is large-enough n?" The additional experiments in Appendix N show that n might only need to be of moderate size for the SPA to provide tight guarantees, yet a more complete answer requires additional techniques.

Acknowledgements

We thank the anonymous referees for their careful critique, which helped improve the quality of the paper considerably. This material is based upon work supported by the National Science Foundation under Grant Nos. CAREER-1845852, FAI-2040880, CIF-1900750, SCH-2312666, CIF-1922971, and CIF-1901243; by NSERC Canada; and by the U.S. Department of Energy Award No. DE-SC0022158.

[10] We take the relative error of a privacy curve estimate $\hat{\varepsilon}(\delta)$, with a ground-truth of $\varepsilon(\delta)$, to be $|1 - \hat{\varepsilon}(\delta)/\varepsilon(\delta)|$.

References

NIST Digital Library of Mathematical Functions. http://dlmf.nist.gov/, Release 1.1.6 of 2022-06-30. F. W. J. Olver, A. B. Olde Daalhuis, D. W. Lozier, B. I. Schneider, R. F. Boisvert, C. W. Clark, B. R. Miller, B. V. Saunders, H. S. Cohl, and M. A. McClain, eds.

Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., and Zhang, L. Deep learning with differential privacy. In *Proc. ACM SIGSAC CCS*, pp. 308–318, 2016. doi: 10.1145/2976749.2978318.

Balle, B., Barthe, G., and Gaboardi, M. Privacy amplification by subsampling: Tight analyses via couplings and divergences. In *NeurIPS*, pp. 6280–6290, 2018.

Bu, Z., Dong, J., Long, Q., and Su, W. Deep learning with Gaussian differential privacy. *Harvard Data Science Review*, 2(3), sep 30 2020. https://hdsr.mitpress.mit.edu/pub/u24wj42y.

De, S., Berrada, L., Hayes, J., Smith, S. L., and Balle, B. Unlocking high-accuracy differentially private image classification through scale. *arXiv preprint arXiv:2204.13650*, 2022.

Dong, J., Roth, A., and Su, W. J. Gaussian Differential Privacy. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(1):3–37, 02 2022. ISSN 1369-7412. doi: 10.1111/rssb.12454. URL https://doi.org/10.1111/rssb.12454.

Doroshenko, V., Ghazi, B., Kamath, P., Kumar, R., and Manurangsi, P. Connect the dots: Tighter discrete approximations of privacy loss distributions. In *Privacy Enhancing Technologies Symposium (PETS)*, 2022.

Dwork, C., Kenthapadi, K., McSherry, F., Mironov, I., and Naor, M. Our data, ourselves: Privacy via distributed noise generation. In Vaudenay, S. (ed.), *EUROCRYPT*, pp. 486–503, 2006a.

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

Ghazi, B., Kamath, P., Kumar, R., and Manurangsi, P. Faster privacy accounting via evolving discretization. In Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., and Sabato, S. (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 7470–7483. PMLR, 17–23 Jul 2022.

Gopi, S., Lee, Y. T., and Wutschitz, L. Numerical composition of differential privacy. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.

Hall, P. *The bootstrap and Edgeworth expansion*. Springer Science & Business Media, 2013.

Jeffreys, H. and Jeffreys, B. *Methods of Mathematical Physics*. Cambridge University Press, 3rd edition, 1999. doi: 10.1017/CBO9781139168489.

Kasiviswanathan, S. P., Lee, H. K., Nissim, K., Raskhodnikova, S., and Smith, A. What can we learn privately? *SIAM J. Comput.*, 40(3):793–826, June 2011. ISSN 0097-5397. doi: 10.1137/090756090.

Koskela, A. and Honkela, A. Computing differential privacy guarantees for heterogeneous compositions using fft. *CoRR*, abs/2102.12412, 2021. URL https://arxiv.org/abs/2102.12412.

Koskela, A., Jälkö, J., and Honkela, A. Computing tight differential privacy guarantees using fft. In *International Conference on Artificial Intelligence and Statistics*, pp. 2560–2569. PMLR, 2020.

Koskela, A., Jälkö, J., Prediger, L., and Honkela, A. Tight differential privacy for discrete-valued mechanisms and for the subsampled gaussian mechanism using fft. In *International Conference on Artificial Intelligence and Statistics*, pp. 3358–3366. PMLR, 2021.

Mironov, I. Rényi differential privacy. In *Proc. Computer Security Found. (CSF)*, pp. 263–275, 2017.

Mironov, I., Talwar, K., and Zhang, L. Rényi differential privacy of the sampled gaussian mechanism. *arXiv preprint arXiv:1908.10530*, 2019.

Murtagh, J. and Vadhan, S. The complexity of computing the optimal composition of differential privacy. In *Proc. Int. Conf. Theory of Cryptography*, pp. 157–175, 2016.

Sommer, D. M., Meiser, S., and Mohammadi, E. Privacy loss classes: The central limit theorem in differential privacy. *Proceedings on Privacy Enhancing Technologies*, 2019(2):245–269, 2019.

Wang, H., Gao, S., Zhang, H., Shen, M., and Su, W. J. Analytical composition of differential privacy via the edgeworth accountant. *arXiv preprint arXiv:2206.04236*, 2022.

Zhu, Y. and Wang, Y.-X. Poission subsampled rényi differential privacy. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *ICML*, volume 97, pp. 7634–7642, 09–15 Jun 2019.

Zhu, Y., Dong, J., and Wang, Y.-X. Optimal accounting of differential privacy via characteristic function. In *International Conference on Artificial Intelligence and Statistics*, pp. 4782–4817. PMLR, 2022.

A. Subsampling: Proof of Lemma 2.2

Subsampling is a fundamental tool in the analysis of differentially-private mechanisms. Informally, subsampling entails applying a differentially-private mechanism to a small set of randomly sampled datapoints from a given dataset. There are several ways of formally defining the subsampling operator, see, e.g., (Balle et al., 2018). The most well-known one, Poisson subsampling, is parameterized by the subsampling rate $\lambda \in (0,1]$ which indicates the probability of selecting a datapoint. More formally, the subsampled datapoints from a dataset D can be expressed as $\{x \in D \,:\, B_x = 1\}$, where B_x is a Bernoulli random variable with parameter λ independent for each $x \in D$. Given any mechanism $\mathcal{M}$, we define the subsampled mechanism $\mathcal{M}_\lambda$ as the composition of $\mathcal{M}$ and the Poisson subsampling operator. Characterizing the privacy guarantees of subsampled mechanisms is the subject of "privacy amplification by subsampling" principle (Kasiviswanathan et al., 2011). This principle is well-studied particularly for characterizing the privacy guarantees of subsampled Gaussian mechanisms in the context of a variant of differential privacy, namely, Rényi differential privacy (Zhu & Wang, 2019; Abadi et al., 2016; Mironov et al., 2019). We can mirror their formulation to characterize ε and δ for the subsampled Gaussian mechanisms. Recall that a Gaussian mechanism satisfies $\mathcal{M}(D) = \mathcal{N}(f(D), \sigma^2 I_d)$ where f is a query function with ℓ_2-sensitivity 1. For the *subsampled* Gaussian, the optimal privacy curve (of a single composition) is

$$\delta_{\mathcal{M}_\lambda}(\varepsilon) = \max\left\{\mathsf{E}_{e^\varepsilon}(P\|Q), \mathsf{E}_{e^\varepsilon}(Q\|P)\right\}, \tag{48}$$

where $P = \mathcal{N}(0, \sigma^2 I_d)$ and $Q = (1-\lambda)P + \lambda P'$, and $P' \sim \mathcal{N}(\boldsymbol{e}_1, \sigma^2 I_d)$ where $\boldsymbol{e}_1$ is the first standard basis vector. In Lemma 2.2 (restated below for convenience), we show that the above maximum is always attained by $\mathsf{E}_{e^\varepsilon}(Q\|P)$ for any $\varepsilon \geq 0$, and that it holds for a larger family of DP mechanisms (including Gaussian and Laplace mechanisms). A similar ordering bound was proved by Mironov et al. (2019, Theorem 5) for the Rényi divergence.

Lemma A.1. *Fix a Borel probability measure P over $\mathbb{R}^n$ that is symmetric around the origin (i.e., $P(\mathcal{A}) = P(-\mathcal{A})$ for every Borel $\mathcal{A} \subset \mathbb{R}^n$), and fix constants $(s, \lambda, \gamma) \in \mathbb{R}^n \times [0,1] \times [1,\infty)$. Let $T_s P$ be the probability measure given by $(T_s P)(\mathcal{A}) = P(\mathcal{A} - s)$, and let $Q = (1-\lambda)P + \lambda T_s P$. We have the inequality $\mathsf{E}_\gamma(P\|Q) \leq \mathsf{E}_\gamma(Q\|P)$, with equality if and only if $(\gamma - 1)\,\lambda\,\|s\|\,\mathsf{E}_\gamma(Q\|P) = 0$.*

Proof. The case $\lambda = 0$ is clear, so assume $\lambda \in (0,1]$. Suppose for now that $\gamma \cdot (1-\lambda) < 1$. Denote $R \triangleq T_s P$, and consider the function $G : (0,\infty) \to [0,\infty)$ defined by

$$G(t) \triangleq t \cdot \mathsf{E}_{1+\frac{\gamma-1}{t}}(P\|R). \tag{49}$$

Since $\gamma' \mapsto \mathsf{E}_{\gamma'}(P\|R)$ is monotonically decreasing, we have that G is monotonically increasing. Note that $0 < \gamma\lambda + 1 - \gamma \leq \lambda$. Thus, plugging $t \in \{\gamma\lambda + 1 - \gamma, \lambda\}$ into G, we obtain

$$(\gamma\lambda + 1 - \gamma) \cdot \mathsf{E}_{\frac{\gamma\lambda}{\gamma\lambda+1-\gamma}}(P\|R) \leq \lambda \cdot \mathsf{E}_{\frac{\lambda-(1-\gamma)}{\lambda}}(P\|R). \tag{50}$$

Now, note that

$$\begin{aligned}
&(\gamma\lambda + 1 - \gamma) \cdot \mathsf{E}_{\frac{\gamma\lambda}{\gamma\lambda+1-\gamma}}(P\|R) \\
&\quad = (\gamma\lambda + 1 - \gamma) \cdot \sup_{\mathcal{A}} P(\mathcal{A}) - \frac{\gamma\lambda}{\gamma\lambda + 1 - \gamma} \cdot R(\mathcal{A}) && (51)\\
&\quad = \sup_{\mathcal{A}} P(\mathcal{A}) - \gamma \cdot ((1-\lambda)P(\mathcal{A}) + \lambda R(\mathcal{A})) && (52)\\
&\quad = \mathsf{E}_\gamma(P\|Q), && (53)
\end{aligned}$$

where the suprema are taken over all Borel sets $\mathcal{A} \subset \mathbb{R}^n$. In addition, by symmetry of P around the origin, we have that

$$\begin{aligned}
\mathsf{E}_{\gamma'}(P\|R) &= \sup_{\mathcal{A}} P(\mathcal{A}) - \gamma' P(\mathcal{A} - s) && (54)\\
&= \sup_{\mathcal{A}} P(-\mathcal{A}) - \gamma' P(-\mathcal{A} - s) && (55)\\
&= \sup_{\mathcal{A}} P(\mathcal{A}) - \gamma' P(\mathcal{A} + s) && (56)\\
&= \sup_{\mathcal{A}} P(\mathcal{A} - s) - \gamma' P(\mathcal{A}) && (57)\\
&= \mathsf{E}_{\gamma'}(R\|P). && (58)
\end{aligned}$$

Therefore,

$$\begin{aligned}
&\lambda \cdot \mathsf{E}_{\frac{\lambda-(1-\gamma)}{\lambda}}(P\|R) \\
&\quad = \lambda \cdot \mathsf{E}_{\frac{\lambda-(1-\gamma)}{\lambda}}(R\|P) && (59)\\
&\quad = \lambda \cdot \sup_{\mathcal{A}} R(\mathcal{A}) - \frac{\lambda - (1-\gamma)}{\lambda} \cdot P(\mathcal{A}) && (60)\\
&\quad = \sup_{\mathcal{A}} ((1-\lambda)P(\mathcal{A}) + \lambda R(\mathcal{A})) - \gamma P(\mathcal{A}) && (61)\\
&\quad = \mathsf{E}_\gamma(Q\|P). && (62)
\end{aligned}$$

We conclude from (50) the desired inequality $\mathsf{E}_\gamma(P\|Q) \leq \mathsf{E}_\gamma(Q\|P)$. In addition, the case $\gamma \cdot (1-\lambda) \geq 1$ follows immediately since then $\mathsf{E}_\gamma(P\|Q) = 0 \leq \mathsf{E}_\gamma(Q\|P)$. □

In light of this lemma, the privacy guarantee of a subsampled Gaussian mechanism is fully characterized by computing only $\mathsf{E}_{e^\varepsilon}((1-\lambda)P + \lambda T_s P\|P)$, where $P = \mathcal{N}(0, \sigma^2 I_d)$. Based on this result, for our numerical experiments, we only compute the saddle-point accountant with this order of P and Q.

B. The Method of Steepest Descent

We describe the general approach for the method of steepest descent. Our task is to compute the contour integral

$$I_n = \frac{1}{2\pi i} \int_{t-i\infty}^{t+i\infty} e^{F_n(z)}\, dz. \tag{63}$$

What we will obtain is an asymptotic expansion

$$I_n \sim \frac{e^{F_n(t_0)}}{\sqrt{2\pi F_n''(t_0)}} \left(1 + \sum_{m=2}^{\infty} \beta_{n,m}\right). \tag{64}$$

In a nutshell, the method of steepest descent is a powerful tool for choosing the best parameter t that renders the computation of I_n easiest. In particular, this choice of t is called the *saddle-point*, which is found as follows.

Here, F_n is holomorphic over a strip $(0,T) + i\mathbb{R}$ in the complex plane, the parameter $n \in \mathbb{N}$ is growing without bound, and $t \in (0,T) \subset \mathbb{R}$ is a free parameter. In particular, the value of the integral I_n is assumed to be independent of the parameter t. This could be satisfied for certain choices of F_n by virtue of its analyticity and in view of Cauchy's integral theorem. As we show in Theorem 3.1, computing the above contour integral amount to exactly computing the privacy parameter $\delta_{L^{(n)}}(\varepsilon)$ if we choose the function

$$F_n(z) = K_{L^{(n)}}(z) - z\varepsilon - \log z - \log(1+z). \tag{65}$$

Suppose that $F_n''(t) > 0$ over $t \in (0,T)$—in particular, F_n is strictly convex over the real interval $(0,T)$—and that there is a value $t_0 \in (0,T)$ solving the equation $F_n'(t_0) = 0$, which is then necessarily unique. Then, a second order Taylor expansion around t_0 yields that

$$F_n(z) = F_n(t_0) + \frac{(z-t_0)^2}{2} F_n''(t_0) + o(|z-t_0|^3). \tag{66}$$

Looking at the the values of the approximating quadratic $F_n(t_0) + \frac{(z-t_0)^2}{2} F_n''(t_0)$ for z near t_0 along the real axis (so $z = t$ for $t_0 \approx t \in \mathbb{R}$) and along the axis $t_0 + i\mathbb{R}$ (so $z = t_0 + is$ for $0 \approx s \in \mathbb{R}$), we see that this approximation has a local minimum at t_0 along the real axis and it has a local maximum at t_0 along the axis $t_0 + i\mathbb{R}$. Hence, t_0 is a saddle-point for the approximating quadratic. Further, as the integral we are concerned with runs along the contour $t + i\mathbb{R}$, we expect the value of I_n to come primarily from values $z \approx t_0$.

Now, consider the function

$$G_n(z) = F_n(t_0 + z) - F_n(t_0) - \frac{z^2}{2} F_n''(t_0). \tag{67}$$

We have that G_n is holomorphic over some vertical strip centered at the origin, and

$$G_n^{(k)}(0) = \begin{cases} 0, & 0 \le k \le 2 \\ F_n^{(k)}(t_0), & k \ge 3. \end{cases} \tag{68}$$

We assume for the next steps that G_n is an entire function. Thus, G_n has the expansion

$$G_n(z) = \sum_{k \ge 3} \frac{F_n^{(k)}(t_0)}{k!} z^k. \tag{69}$$

Furthermore, $e^{G_n(z)}$ has the power series expansion

$$e^{G_n(z)} = 1 + \sum_{k \ge 3} \alpha_{n,k} z^k, \tag{70}$$

where

$$\alpha_{n,k} = \frac{1}{k!} B_k(0, 0, F_\varepsilon^{(3)}(t_0), \ldots, F_\varepsilon^{(k)}(t_0)). \tag{71}$$

As we may write

$$F_n(t_0 + is) = F_n(t_0) + G_n(is) - \frac{F_n''(t_0)}{2} s^2, \tag{72}$$

we get the exact value of the integral

$$I_n = \frac{e^{F_n(t_0)}}{2\pi} \int_{-\infty}^{\infty} e^{-s^2 F_n''(t_0)/2} \left(1 + \sum_{k \ge 3} \alpha_{n,k} (is)^k\right) ds. \tag{73}$$

The derived steps thus far have all been justified rigorously. The final step, however, is a heuristic, where we truncate the power series expansion to obtain possible estimates of I_n. The point is that the derived expressions through this heuristic have the potential of being proved by other means to be indeed close approximations of I_n.

For instance, dropping the whole series beyond the constant term yields the basic saddle-point approximation

$$I_{n,1} \triangleq \frac{e^{F_n(t_0)}}{2\pi} \int_{-\infty}^{\infty} e^{-s^2 F_n''(t_0)/2}\, ds = \frac{e^{F_n(t_0)}}{\sqrt{2\pi F_n''(t_0)}}. \tag{74}$$

Note that this approximation is in fact exact if F_n is a quadratic, i.e., for computing the Gaussian integral. Keeping the terms $k \in \{3, \cdots, 2k^\star\}$, it is not hard to see that one obtains the $k^\star$-th estimate

$$I_{n,k^\star} \triangleq \frac{e^{F_n(t_0)}}{\sqrt{2\pi F_n''(t_0)}} \left(1 + \sum_{m=2}^{k^\star} \beta_{n,m}\right), \tag{75}$$

where we denote the constants

$$\beta_{n,m} \triangleq \frac{(-1)^m B_{2m}(0, 0, F_n^{(3)}(t_0), \ldots, F_n^{(2m)}(t_0))}{2^m m! F_n''(t_0)^m}. \tag{76}$$

Then one might say that I_n has the "asymptotic expansion"

$$I_n \sim \frac{e^{F_n(t_0)}}{\sqrt{2\pi F_n''(t_0)}} \left(1 + \sum_{m=2}^{\infty} \beta_{n,m}\right). \tag{77}$$

Recall that this does not mean that the above equation holds with equality for any particular n. Rather, it is a heuristic indicating the potential for the truncated expansion to give close approximations for the intended integral I_n.

C. Satisfiability of the Assumptions

We explain here how Assumption 2.6 is satisfied by the subsampled Gaussian and Laplace mechanisms. Note that by the Lebesgue decomposition theorem, the probability measure of the PLRV can always be decomposed into a sum of an absolutely continuous measure, a discrete measure, and a singular measure (such as the Cantor distribution). Thus, Assumption 2.6 requires the exclusion of singular components. This can be easily seen to be satisfied by the subsampled Gaussian and Laplace mechanisms. Further, Assumption 2.6 does not impose any requirement on the discrete part. Thus, we consider the continuous part here.

Note that the PLRV for the subsampled Gaussian mechanism (with subsampling rate λ, variance σ^2, and sensitivity s) is given by

$$L = \log\left(1 - \lambda + \lambda e^{(2sX - s^2)/(2\sigma^2)}\right), \tag{78}$$

where $X \sim (1-\lambda)\mathcal{N}(0,\sigma^2) + \lambda\mathcal{N}(s,\sigma^2)$. Hence,

$$\mathbb{P}\left[L \le z\right] = \mathbb{P}\left[X \le \frac{s}{2} + \frac{\sigma^2}{s}\log\left(\frac{e^z - (1-\lambda)}{\lambda}\right)\right] \tag{79}$$

if $z > \log(1-\lambda)$, and $\mathbb{P}[L \le z] = 0$ otherwise. So, L is continuous with PDF

$$p_L(z) = A\frac{e^{2z}}{g(z)^{3/2}}\cdot\left(\frac{g(z)}{\lambda}\right)^{-\frac{\sigma^2}{2s^2}\log\frac{g(z)}{\lambda}}\cdot 1_{(\log(1-\lambda),\infty)}(z), \tag{80}$$

where $g(z) = e^z - (1-\lambda)$ and we have the constant $A = \frac{\sigma}{s}\cdot\sqrt{\frac{\lambda}{2\pi}}\exp\left(-\frac{s^2}{8\sigma^2}\right)$. From this, we see that $p_L(z)$ decays superexponentially as $z \to \infty$. Further, it is continuous. Indeed, we only need to check continuity at $z = \log(1-\lambda)$. But this is immediate using, e.g., $y = \log\frac{g(z)}{\lambda}$ and taking $y \to -\infty$. These properties imply that Assumption 2.6 is satisfied by the subsampled Gaussian mechanism.

Finally, we note that the case of the subsampled Laplace mechanism is simpler. Indeed, taking the analogous expression for L as in (78), we see that L has only a discrete component and a continuous component. Further, the continuous part comes from values of X between 0 and s. This boundedness translates into the fact that the PDF of the continuous part of L is compactly supported, so Assumption 2.6 is satisfied in this case too.

D. Well-Definedness of the Saddle-Point

The well-definedness of the saddle-point, given $\varepsilon < \operatorname{ess\,sup} L$, follows from convexity of F_ε over the positive reals. Namely, we show that F_ε is *complex-differentiable* and that there is a unique positive real t_0 such that $F'_\varepsilon(t_0) = 0$. Let $K_L|_{\mathbb{R}}$ be the restriction of the CGF to the real axis. We have that $K_L|_{\mathbb{R}}$ is convex over $(0,\infty)$, and thus, $F_\varepsilon|_{\mathbb{R}}$ is strictly convex there. Thus, the minimum of F_ε over the positive reals is unique; further, the *real* derivative at this minimum vanishes. Nevertheless, finiteness of M_L over $(0,\infty)$ implies its analyticity over the half-plane $(0,\infty) + i\mathbb{R}$; in particular, the *complex* derivative of F_ε exists in the same half-plane. Hence, the function F_ε is complex-differentiable at t_0, and its derivative vanishes there, as required.

E. New Formulas for the Privacy Curve: Proof of Theorem 3.1

Before proving Theorem 3.1, we show the following general Parseval identity. For $f \in L^1(\mathbb{R})$, we denote the Fourier transform by

$$\hat{f}(\xi) \triangleq \int_{\mathbb{R}} f(x)e^{-ix\xi}\,dx. \tag{81}$$

Lemma E.1. *Let $P = Q + R$ be a Borel probability measure on $\mathbb{R}$, where Q is absolutely continuous with respect to the Lebesgue measure whose PDF is square-integrable and R is discrete. For any continuous function $f : \mathbb{R} \to \mathbb{R}$ such that $f \in L^1(\mathbb{R}) \cap L^2(\mathbb{R})$, $\hat{f} \in L^1(\mathbb{R})$, and $\mathbb{E}_{X\sim P}[|f(X)|] < \infty$, we have the Parseval identity*

$$\int_{\mathbb{R}} f(x)\,dP(x) = \frac{1}{2\pi}\int_{\mathbb{R}} \hat{f}(\xi)\phi_P(\xi)\,d\xi, \tag{82}$$

where $\phi_P(\xi) \triangleq \mathbb{E}_{X\sim P}[e^{i\xi X}]$ is the characteristic function.

Proof. Let q denote the PDF of Q. Suppose R is supported over $\{x_j\}_{j\in J}$, where J is at most countable, and write $r_j = R(\{x_j\})$. Then, we may write

$$\int_{\mathbb{R}} f(x)\,dP(x) = \int_{\mathbb{R}} f(x)\,dQ(x) + \int_{\mathbb{R}} f(x)\,dR(x) \tag{83}$$

$$= \int_{\mathbb{R}} f(x)q(x)\,dx + \sum_{j\in J} f(x_j)r_j. \tag{84}$$

Since $f, q \in L^1(\mathbb{R}) \cap L^2(\mathbb{R})$, we have the Parseval identity

$$\int_{\mathbb{R}} f(x)q(x)\,dx = \frac{1}{2\pi}\int_{\mathbb{R}} \hat{f}(\xi)\phi_Q(\xi)\,d\xi. \tag{85}$$

As we also have continuity of f and integrability of $\hat{f}$, we also have the Fourier inversion

$$f(x) = \frac{1}{2\pi}\int_{\mathbb{R}} \hat{f}(\xi)e^{ix\xi}\,d\xi \tag{86}$$

for *every* $x \in \mathbb{R}$. In particular, we have that

$$\sum_{j\in J} f(x_j)r_j = \frac{1}{2\pi}\int_{\mathbb{R}} \hat{f}(\xi)\phi_R(\xi)\,d\xi. \tag{87}$$

The desired result follows by $\phi_P = \phi_Q + \phi_R$. □

Now, we apply Lemma E.1 to derive Theorem 3.1.

Proof of Theorem 3.1. Expectations of functions of $\tilde{L}$ can be written in terms of L as $\mathbb{E}[f(\tilde{L})] = \mathbb{E}[e^{tL}f(L)]/M_L(t)$. Thus, the MGF of the tilted variable $\tilde{L}$ is given by

$$M_{\tilde{L}}(z) = \mathbb{E}[e^{z\tilde{L}}] = \frac{\mathbb{E}[e^{tL}e^{zL}]}{M_L(t)} = \frac{M_L(t+z)}{M_L(t)}. \tag{88}$$

Similarly, expectations of functions of L can be written in terms of $\tilde{L}$ as $\mathbb{E}[f(L)] = M_L(t)\,\mathbb{E}[e^{-t\tilde{L}}f(\tilde{L})]$. Thus, we can write the privacy curve δ_L in terms of the tilted variable $\tilde{L}$ as

$$\delta_L(\varepsilon) = \mathbb{E}\left[\left(1 - e^{\varepsilon - L}\right)^+\right] \tag{89}$$

$$= M_L(t)\,\mathbb{E}\left[e^{-t\tilde{L}}\left(1 - e^{\varepsilon - \tilde{L}}\right)^+\right]. \tag{90}$$

In other words, the formula in (22) holds.

Next, we use Assumption 2.6 to apply Parseval's identity (Lemma E.1) to the expectation in (90) to get the contour-integral formula in (23). Specifically, consider the function

$$f(x) = e^{-tx}\left(1 - e^{\varepsilon - x}\right)^+, \tag{91}$$

Note that f is bounded, continuous, and $f \in L^1(\mathbb{R}) \cap L^2(\mathbb{R})$. Further, we have the Fourier transform

$$\hat{f}(s) = \frac{e^{-(t+is)\varepsilon}}{(t+is)(t+1+is)} \in L^1(\mathbb{R}). \tag{92}$$

In addition, by Assumption 2.6, the probability measure P_L induced by L may be written as $P_L = Q_L + R_L$, where Q_L is absolutely continuous with respect to the Lebesgue measure whose PDF q_L satisfies that $x \mapsto e^{\tau x}q_L(x)^2$ is integrable for every $\tau > 0$ and R_L is discrete. Suppose R_L is supported over $\{x_j\}_{j\in J}$ with J at most countable, and write $r_{L,j} = R(\{x_j\})$. Then, by definition of exponential tilting, for every Borel set $B \subset \mathbb{R}$, we have that

$$P_{\tilde{L}}(B) = \frac{1}{M_L(t)}\int_B e^{tx}\,dP_L(x) \tag{93}$$

$$= \frac{1}{M_L(t)}\int_B e^{tx}\,dQ_L(x) + \frac{1}{M_L(t)}\int_B e^{tx}\,dR_L(x) \tag{94}$$

$$= \frac{1}{M_L(t)}\int_B e^{tx}q_L(x)\,dx + \frac{1}{M_L(t)}\sum_{\substack{j\in J\\ x_j\in B}} e^{tx_j}r_{L,j} \tag{95}$$

$$= \tilde{Q}(B) + \tilde{R}(B), \tag{96}$$

where we define the Borel measures

$$\tilde{Q}(B) \triangleq \frac{1}{M_L(t)}\int_B e^{tx}q_L(x)\,dx, \tag{97}$$

$$\tilde{R}(B) \triangleq \frac{1}{M_L(t)}\sum_{\substack{j\in J\\ x_j\in B}} e^{tx_j}r_{L,j}. \tag{98}$$

From these definitions, it is clear that $\tilde{R}$ is discrete and $\tilde{Q}$ is absolutely continuous with respect to the Lebesgue measure with PDF $\tilde{q}(x) \triangleq e^{tx}q_L(x)/M_L(t)$. Furthermore, by assumption on q_L, we have that $\tilde{q} \in L^2(\mathbb{R})$. Therefore, we may apply Parseval's identity (Lemma E.1) on f and $P_{\tilde{L}}$ to obtain

$$\mathbb{E}\left[f(\tilde{L})\right] = \frac{1}{2\pi}\int_{\mathbb{R}} \hat{f}(s)\phi_{P_{\tilde{L}}}(s)\,ds. \tag{99}$$

Next, applying the formula for $M_{\tilde{L}}$ in (88), we see that

$$\phi_{P_{\tilde{L}}}(s) = \mathbb{E}\left[e^{is\tilde{L}}\right] = M_{\tilde{L}}(is) = \frac{M_L(t+is)}{M_L(t)}. \tag{100}$$

Therefore, combining formulas (90) and (99), we get

$$\delta_L(\varepsilon) = M_L(t)\mathbb{E}\left[f(\tilde{L})\right] = \frac{1}{2\pi}\int_{\mathbb{R}} \hat{f}(s)M_L(t+is)\,ds. \tag{101}$$

Now, using the contour $\{z = t + is : -\infty < s < \infty\}$ oriented counter-clockwise, we see that (101) may be rewritten as the contour integral

$$\delta_L(\varepsilon) = \frac{1}{2\pi i}\int_{t-i\infty}^{t+i\infty} \hat{f}((z-t)/i)M_L(z)\,dz. \tag{102}$$

Finally, using the formula for $\hat{f}$ in (92), we deduce

$$\delta_L(\varepsilon) = \frac{1}{2\pi i}\int_{t-i\infty}^{t+i\infty} \frac{e^{-z\varepsilon}}{z(z+1)}M_L(z)\,dz \tag{103}$$

$$= \frac{1}{2\pi i}\int_{t-i\infty}^{t+i\infty} e^{F_\varepsilon(z)}\,dz, \tag{104}$$

where we define

$$F_\varepsilon(z) \triangleq K_L(z) - \varepsilon z - \log(z) - \log(1+z) \tag{105}$$

and we take the principal branch for the complex logarithm. This is precisely the desired formula for δ_L stated in (23), and the proof of the theorem is therefore complete. □

F. The Large-Composition Regime: Proof of Theorem 4.1

We may show this result using the standard Berry-Esseen approach. By the Berry-Esseen theorem, we have for $Z \sim$

$\mathcal{N}(\mathbb{E}[L], \sigma_L^2)$ that

$$\begin{aligned}
\delta_L(\varepsilon) &= \mathbb{E}\left[\left(1 - e^{\varepsilon - L}\right)^+\right] && (106)\\
&= \int_0^\infty \mathbb{P}\left[(1 - e^{\varepsilon - L})^+ > u\right] du && (107)\\
&= \int_0^\infty \mathbb{P}\left[1 - e^{\varepsilon - L} > u\right] du && (108)\\
&= \int_0^1 \mathbb{P}\left[1 - e^{\varepsilon - L} > u\right] du && (109)\\
&= \int_0^1 \mathbb{P}\left[L > \varepsilon - \log(1-u)\right] du && (110)\\
&= \delta_Z(\varepsilon) + \theta \cdot \frac{0.56\,\mathrm{P}_0}{\sigma_L^3} && (111)
\end{aligned}$$

where $|\theta| \le 1$. A direct computation yields that for any $\varepsilon \ge \mathbb{E}[L]$ (with $Z \sim \mathcal{N}(\mathbb{E}[L], \sigma_L^2)$)

$$\begin{aligned}
\delta_Z(\varepsilon) = \Phi&\left(\frac{\mathbb{E}[L] - \varepsilon}{\sigma_L}\right)\\
&- e^{\varepsilon - \mathbb{E}[L] + \sigma_L^2/2}\,\Phi\left(\frac{\mathbb{E}[L] - \sigma_L^2 - \varepsilon}{\sigma_L}\right).
\end{aligned} \tag{112}$$

Plugging in $\varepsilon = \mathbb{E}[L] - \Phi^{-1}(\delta)\sigma_L$, we obtain that

$$\begin{aligned}
&\delta_Z(\mathbb{E}[L] - \Phi^{-1}(\delta)\sigma_L)\\
&\quad = \delta - e^{-\Phi^{-1}(\delta)\sigma_L + \sigma_L^2/2}\Phi\left(\Phi^{-1}(\delta) - \sigma_L\right).
\end{aligned} \tag{113}$$

Using $\Phi(-x) = Q(x) = \frac{q(x)e^{-x^2/2}}{\sqrt{2\pi}}$ for $x > 0$, we obtain

$$\begin{aligned}
&\delta_Z(\mathbb{E}[L] - \Phi^{-1}(\delta)\sigma_L) && (114)\\
&\quad = \delta - \frac{q(\sigma_L - \Phi^{-1}(\delta))}{\sqrt{2\pi}} e^{-\Phi^{-1}(\delta)^2/2} && (115)\\
&\quad = \delta - \frac{q(\sigma_L - \Phi^{-1}(\delta))}{\sqrt{2\pi}} e^{-(-\Phi^{-1}(\delta))^2/2} && (116)\\
&\quad = \delta - \frac{q(\sigma_L - \Phi^{-1}(\delta))}{q(-\Phi^{-1}(\delta))}\Phi(\Phi^{-1}(\delta)) && (117)\\
&\quad = \delta \cdot \left(1 - \frac{q(\sigma_L - \Phi^{-1}(\delta))}{q(-\Phi^{-1}(\delta))}\right). && (118)
\end{aligned}$$

Note that $q(x) \asymp 1/x$ as $x \to \infty$. Since $\sigma_L/(-\Phi^{-1}(\delta)) \to \infty$ by assumption, we also have $\sigma_L - \Phi^{-1}(\delta)) \to \infty$. Thus, we obtain

$$\frac{q(\sigma_L - \Phi^{-1}(\delta))}{q(-\Phi^{-1}(\delta))} \asymp \frac{1}{-\Phi^{-1}(\delta)q(-\Phi^{-1}(\delta))} \cdot \frac{1}{\frac{\sigma_L}{-\Phi^{-1}(\delta)} - 1}. \tag{119}$$

As $\limsup\ \delta < 1/2$, we get that the term $-\Phi^{-1}(\delta)q(-\Phi^{-1}(\delta))$ is bounded away from 0. Therefore, we get that

$$\delta_Z(\mathbb{E}[L] - \Phi^{-1}(\delta)\sigma_L) = \delta \cdot (1 + o(1))\,. \tag{120}$$

From (111), and since $\mathrm{P}_0 = o(\sigma_L^3)$ by assumption, we conclude that

$$\delta_L(\mathbb{E}[L] - \Phi^{-1}(\delta)\sigma_L) = \delta \cdot (1 + o(1))\,. \tag{121}$$

In other words, $\mathcal{M}$ is $(\mathbb{E}[L] - \Phi^{-1}(\delta)\sigma_L, \delta \cdot (1 + o(1)))$-DP, as desired.

G. Asymptotic of the Saddle-Point: Proof of Theorem 4.2

We write $K = K_L$ for short. Consider the saddle-point equation (25):

$$K'(t) = \varepsilon + \frac{1}{t} + \frac{1}{1+t}. \tag{122}$$

The left-hand side strictly increases from $\mathbb{E}[L]$ to $\operatorname{ess\,sup} L$ over $t \in [0, \infty)$, whereas the right-hand side strictly decreases from ∞ to ε over the same interval. Hence, there exists a unique solution $t = t_0 > 0$, which we call the saddle-point.

We show first that $t_0 \to 0$ as $n \to \infty$. Suppose, for the sake of contradiction, that $t^* \triangleq \limsup_{n\to\infty} t_0 > 0$, and let $n_k \nearrow \infty$ be a sequence of indices such that the sequence of the n_k-th saddle points, denoted $t_0^{(k)}$, converge to t^*. Let $\rho_2 : (0,\infty) \to (0,\infty)$ be defined by $\rho_2(t) \triangleq (K'(t) - \mathbb{E}[L])/(t\sigma_L^2)$, so $\rho_2(t) \to 1$ as $t \to 0^+$ and

$$K'(t) = \mathbb{E}[L] + \sigma_L^2 t \rho_2(t). \tag{123}$$

Note that ρ_2 is a continuous function. Noting that $\varepsilon = \mathbb{E}[L] + b\sigma_L$, rearranging the saddle-point equation yields that

$$\frac{1 + \frac{\sigma_L^2}{\mathbb{E}[L]} t\rho_2(t)}{1 + b\frac{\sigma_L}{\mathbb{E}[L]}} = 1 + \frac{1}{\varepsilon t} + \frac{1}{\varepsilon \cdot (1+t)}. \tag{124}$$

Taking $t \in \{t_0^{(k)}\}_{k\in\mathbb{N}}$, letting $k \to \infty$, and recalling the assumptions that $(\mathbb{E}[L], \sigma_L^2) \asymp n \cdot (\mathrm{KL}, \mathrm{V})$ for $\mathrm{KL}, \mathrm{V} > 0$ and that $b = o(\sqrt{n})$, we infer from (124) that

$$\frac{\mathrm{V} t^* \rho_2(t^*)}{\mathrm{KL}} = 0. \tag{125}$$

Equality (125) contradicts that $\mathrm{V}, t^*, \rho_2(t^*), \mathrm{KL} > 0$. Thus, we must have that $t^* = 0$.

Consider the reparametrization $t = d/\sigma_L$, so d is a variable over $(0,\infty)$. The saddle-point equation can be rewritten as

$$\left(\rho_2(t) - \frac{b}{\sigma_L}\right) d^2 - \left(b + \frac{2}{\sigma_L}\right) d - \left(1 - \frac{\rho_2(t) d^3}{\sigma_L}\right) = 0. \tag{126}$$

We rewrite the saddle-point equation in this "quadratic" form since it closely approximates the quadratic $d^2 - bd -$

$1 = 0$ at the saddle-point. Indeed, let $d_0 > 0$ be such that $t_0 = d_0/\sigma_L$. We obtain from (126) the inequality $\frac{1}{2}d_0^2 - (b+1)d_0 - 1 \le 0$ for all large n. This latter inequality yields that

$$d_0 \le b + 1 + \sqrt{(b+1)^2 + 2} = o(n^{1/6}). \tag{127}$$

Hence, $\rho_2(t_0)d_0^3/\sigma_L \to 0$ as $n \to \infty$, i.e., the "constant" term in (126) approaches 1. Thus, for all large n, completing the square in (126) yields (denoting $t = t_0$, $\rho = \rho_2$, and $\sigma = \sigma_L$ for short)

$$d_0 = \frac{b + \frac{2}{\sigma} + \sqrt{\left(b + \frac{2}{\sigma}\right)^2 + 4\left(1 - \frac{\rho(t)d_0^3}{\sigma}\right)\left(\rho(t) - \frac{b}{\sigma}\right)}}{2\left(\rho(t) - \frac{b}{\sigma}\right)}. \tag{128}$$

Taking $n \to \infty$, we obtain

$$d_0 \asymp \frac{b + \sqrt{b^2+4}}{2}, \tag{129}$$

which gives the desired asymptotic formula for the saddle-point $t_0 = d_0/\sigma_L$.

H. Contrast between SPA and the Standard CLT

To illustrate the advantage of our tilting approach, we compare the asymptotic behavior of the error in Theorem 5.7 to that obtainable from non-tilted Berry-Esseen. Let $L = L_1 + \cdots + L_n$ for independent PLRVs $L_1, \cdots, L_n$ that satisfy Assumption 2.5. Suppose that Assumption 2.7 holds too.

By the Berry-Esseen theorem, we have for a Gaussian $Z \sim \mathcal{N}(\mathbb{E}[L], \sigma_L^2)$ that[11]

$$\delta_L(\varepsilon) = \mathbb{E}\left[\left(1 - e^{\varepsilon - L}\right)^+\right] \tag{130}$$

$$= \int_0^1 \mathbb{P}\left[L > \varepsilon - \log(1-u)\right] du \tag{131}$$

$$= \delta_Z(\varepsilon) + \theta \cdot \frac{0.56\ \mathrm{P}_0}{\sigma_L^3} \tag{132}$$

where $|\theta| \le 1$. By Assumption 2.7, the error term in the standard Berry-Esseen approach shown above satisfies

$$\mathrm{err}_{\text{Standard}}(\varepsilon) \triangleq \frac{0.56\ \mathrm{P}_0}{\sigma_L^3} \asymp \frac{0.56\ \mathrm{P}}{\mathrm{V}^{3/2}\cdot\sqrt{n}}. \tag{133}$$

Thus, the improvement our approach yields is asymptotically (see Theorem 5.7 for the definitions of $C(b)$ and τ)

$$\frac{\mathrm{err}_{\mathrm{SP}}(\varepsilon; t_0)}{\mathrm{err}_{\text{Standard}}(\varepsilon)} \asymp \frac{2\sqrt{e}}{C(b)^\tau}. \tag{134}$$

[11] Note that Z is not necessarily a PLRV associated to a Gaussian mechanism, since in general $\sigma_L^2 \neq 2\mathbb{E}[L]$.

Even for moderate values of b, the above ratio is very small (recall that we denote $\varepsilon = \mathbb{E}[L] + b\sigma_L$). For example, if $b \approx 6.4$ (so $\delta \approx 10^{-10}$ in the limit; see Theorem 4.1 on the high-composition regime), we obtain the limit of the ratio

$$\lim_{n\to\infty} \frac{\mathrm{err}_{\mathrm{SP}}(\varepsilon; t_0)}{\mathrm{err}_{\text{Standard}}(\varepsilon)} \approx 3 \times 10^{-9}. \tag{135}$$

In addition, in the complementary regime of $\delta \to 0$, e.g., when $\varepsilon = \mathbb{E}[L] + b\sigma_L$ with $b \ge \sqrt{\log n}$ (and still $b = o(n^{1/6})$), one has that the error term in the standard CLT *dominates* the approximation of δ:

$$\delta_Z(\varepsilon) = o\left(\mathrm{err}_{\text{Standard}}(\varepsilon)\right). \tag{136}$$

In contrast, in the same regime, our error term $\mathrm{err}_{\mathrm{SP}}(\varepsilon; t_0)$ is always vanishingly smaller than the approximation itself, i.e.,

$$\mathrm{err}_{\mathrm{SP}}(\varepsilon; t_0) = o\left(\delta_{L,\,\text{SP-CLT}}(\varepsilon)\right). \tag{137}$$

I. Proofs of Section 5.1

I.1. Proof of Proposition 5.3

Denote $K = K_L$ for short. The Gaussian expectation may be computed as

$$\begin{aligned}
&\mathbb{E}\left[\bar{f}\left(Z - \varepsilon, t\right)\right] \\
&= \exp\left(\frac{K''(t)t^2}{2} - (K'(t) - \varepsilon)t\right) \times \\
&\quad Q\left(\sqrt{K''(t)}\,t - \frac{K'(t) - \varepsilon}{\sqrt{K''(t)}}\right) \\
&\quad - \exp\left(\frac{K''(t)(t+1)^2}{2} - (K'(t) - \varepsilon)(t+1)\right) \times \\
&\quad Q\left(\sqrt{K''(t)}\,(t+1) - \frac{K'(t) - \varepsilon}{\sqrt{K''(t)}}\right).
\end{aligned} \tag{138}$$

Using $Q(z) = \frac{q(z)}{\sqrt{2\pi}} e^{-z^2/2}$ and the definitions of α, β, γ, we get

$$\mathbb{E}\left[\bar{f}\left(Z - \varepsilon, t\right)\right] = \frac{q(\alpha) - q(\beta)}{\sqrt{2\pi}}\, e^{-\gamma^2/2}. \tag{139}$$

Plugging this into the definition of $\delta_{L,\,\text{SP-CLT}}$ completes the proof.

I.2. Proof of Proposition 5.5

Let $Z \sim \mathcal{N}\left(K_L'(t), K_L''(t)\right)$ be the variable in the expectation in (38). Its PDF is upper bounded by $p_Z(z) \le$

$\frac{1}{\sqrt{2\pi K_L''(t)}}$. Thus

$$\mathbb{E}\left[e^{-t(Z-\varepsilon)}\left(1-e^{-(Z-\varepsilon)}\right)^+\right]$$
$$= \int_\varepsilon^\infty p_Z(z) e^{-t(z-\varepsilon)}\left(1-e^{-(z-\varepsilon)}\right) dz \tag{140}$$
$$\le \frac{1}{\sqrt{2\pi K_L''(t)}} \int_\varepsilon^\infty e^{-t(z-\varepsilon)}\left(1-e^{-(z-\varepsilon)}\right) dz \tag{141}$$
$$= \frac{1}{\sqrt{2\pi K_L''(t)}\, t(t+1)}. \tag{142}$$

Applying this bound to the definition of $\delta_{L,\,\text{SP-CLT}}(\varepsilon)$ in (38) completes the proof.

J. Error Bound for SPA-CLT: Proof of Theorem 5.6

Fix $t > 0$. Recall from (36) that

$$\delta_L(\varepsilon) = e^{K_L(t)-\varepsilon t}\,\mathbb{E}\left[\bar{f}\left(\tilde{L}-\varepsilon, t\right)\right] \tag{143}$$

where $\tilde{L}$ is the exponential tilting of L with parameter t, and

$$\bar{f}(x,t) = e^{-xt}(1-e^{-x})^+ \tag{144}$$

Note that $K_L'(t) = \mathbb{E}[\tilde{L}]$ and $K_L''(t) = \mathrm{Var}[\tilde{L}]$. We consider the function $\bar{f}(x,t)$. We show next that, for fixed t, $x \mapsto \bar{f}(x,t)$ is a unimodal function with a maximal value of $t^t/(t+1)^{t+1}$. Certainly $\bar{f}(x,t) \ge 0$ for all x. For $x > 0$ the derivative (with respect to x) is

$$\bar{f}'(x,t) = -te^{-tx}(1-e^{-x}) + e^{-tx}e^{-x} \tag{145}$$
$$= e^{-tx}\left[-t + (t+1)e^{-x}\right]. \tag{146}$$

Note that $-t + (t+1)e^{-x}$ is monotonically decreasing in x, which means that $\bar{f}(x,t)$ is increasing until $-t + (t+1)e^{-x} = 0$, and is subsequently decreasing. In particular, the maximal value of $\bar{f}$ is attained when

$$x = x_0 = -\log\frac{t}{t+1}. \tag{147}$$

Note that $x_0 > 0$. Thus, the maximal value of $\bar{f}$ is

$$f_{\max} \triangleq \bar{f}(x_0,t) = \bar{f}\left(-\log\frac{t}{t+1}, t\right) \tag{148}$$
$$= \left(\frac{t}{t+1}\right)^t\left(1-\frac{t}{t+1}\right) = \frac{t^t}{(t+1)^{t+1}}. \tag{149}$$

Thus, between $x = 0$ and $x = x_0$, $\bar{f}(x,t)$ is monotonically increasing from 0 to $f_{\max}$; then from $x = x_0$ to $x = \infty$, $\bar{f}(x,t)$ is monotonically decreasing from $f_{\max}$ to 0. Thus, there exist functions $f_1^{-1}(z)$, $f_2^{-1}(z)$ such that, for any $z \in (0, f_{\max})$, $\bar{f}(x,t) > z$ if and only if

$$f_1^{-1}(z) < x < f_2^{-1}(z).$$

Therefore,

$$\mathbb{E}[\bar{f}(\tilde{L}-\varepsilon,t)]$$
$$= \int_0^{f_{\max}} \mathbb{P}\left[\bar{f}(\tilde{L}-\varepsilon,t) > z\right] dz \tag{150}$$
$$= \int_0^{f_{\max}} \mathbb{P}\left[f_1^{-1}(z) < \tilde{L}-\varepsilon < f_2^{-1}(z)\right] dz. \tag{151}$$

In addition, we may apply the Berry-Esseen theorem to write

$$\sup_{x\in\mathbb{R}}\left|\mathbb{P}\left[\tilde{L} > x\right] - \mathbb{P}[Z > x]\right| \le \frac{0.56\,\mathrm{P}_t}{K_L''(t)^{3/2}} \tag{152}$$

where $Z \sim \mathcal{N}(K_L'(t), K_L''(t))$ and P_t is defined in the beginning of Section 5. Thus we have the upper bound

$$\delta_L(\varepsilon)$$
$$= e^{K_L(t)-\varepsilon t}\,\mathbb{E}\left[\bar{f}\left(\tilde{L}-\varepsilon,t\right)\right] \tag{153}$$
$$= e^{K_L(t)-\varepsilon t}\int_0^{f_{\max}} \mathbb{P}\left[f_1^{-1}(z) < \tilde{L}-\varepsilon < f_2^{-1}(z)\right] dz \tag{154}$$
$$\le e^{K_L(t)-\varepsilon t}\left(\frac{1.12 f_{\max}\mathrm{P}_t}{K_L''(t)^{3/2}} + \int_0^{f_{\max}} \mathbb{P}\left[f_1^{-1}(z) < Z-\varepsilon < f_2^{-1}(z)\right] dz\right) \tag{155}$$
$$= e^{K_L(t)-\varepsilon t}\left(\mathbb{E}\left[\bar{f}\left(Z-\varepsilon,t\right)\right] + \frac{1.12 f_{\max}\mathrm{P}_t}{K_L''(t)^{3/2}}\right) \tag{156}$$

Similarly, we have the lower bound

$$\delta_L(\varepsilon)$$
$$\ge e^{K_L(t)-\varepsilon t}\left(\mathbb{E}\left[\bar{f}\left(Z-\varepsilon,t\right)\right] - \frac{1.12 f_{\max}\mathrm{P}_t}{K_L''(t)^{3/2}}\right). \tag{157}$$

This completes the proof of the theorem.

K. Asymptotic of the SPA-CLT Approximation Error: Proof of Theorem 5.7

We write $K = K_L$ for short. Recall the definition of the error term in (44)

$$\mathrm{err}_{\mathrm{SP}}(\varepsilon; t_0) = e^{K(t_0)-\varepsilon t_0}\frac{t_0^{t_0}}{(1+t_0)^{1+t_0}}\cdot\frac{1.12\,\mathrm{P}_{t_0}}{K''(t_0)^{3/2}}. \tag{158}$$

From the characterization of the saddle-point in Theorem 4.2, we have that

$$t_0 \asymp \frac{b+\sqrt{b^2+4}}{2\sigma_L}. \tag{159}$$

By Assumption 2.7, we have that $\sigma_L^2 = K''(0) \asymp n\mathrm{V}$ as $n \to \infty$. Hence, $t_0 \asymp c/\sqrt{n}$ for $c = (b+\sqrt{b^2+4})/(2\mathrm{V}) = o(n^{1/6})$. Thus, by Assumption 2.7 again, $(K''(t_0), \mathrm{P}_{t_0}) \asymp n \cdot (\mathrm{V}, \mathrm{P})$. As we also have that $t_0 \to 0$, we conclude that

$$\frac{t_0^{t_0}}{(1+t_0)^{1+t_0}} \cdot \frac{1.12\,\mathrm{P}_{t_0}}{K''(t_0)^{3/2}} \asymp \frac{1.12\,\mathrm{P}}{\mathrm{V}^{3/2} \cdot \sqrt{n}}. \tag{160}$$

Thus, it only remains to analyze the asymptotic of $\exp\left(K(t_0) - \varepsilon t_0\right)$.

We use the following Taylor expansion of K around 0:

$$K(t_0) = t_0 \cdot \mathbb{E}[L] + \frac{t_0^2}{2} \cdot \sigma_L^2 + \frac{t_0^3}{6} \cdot K'''(\xi), \tag{161}$$

where $0 \le \xi \le t_0$. Using $\varepsilon = \mathbb{E}[L] + b\sigma_L$, and writing $t_0 = d_0/\sigma_L$ (so $d_0 \asymp (b+\sqrt{b^2+4})/2$ by (159)), we obtain

$$K(t_0) - \varepsilon t_0 = \frac{d_0^2}{2} - bd_0 + \frac{d_0^3 K'''(\xi)}{6\sigma_L^3}. \tag{162}$$

Now, note that $K'''(\xi) = \sum_{j=1}^n K_{L_j}'''(\xi)$. Thus, applying the triangle inequality, we obtain that $|K'''(\xi)| \le \mathrm{P}_\xi$. As $0 \le \xi \le t_0$, Assumption 2.7 yields that $|K'''(\xi)| = O(n)$. As $\sigma_L = \Theta(\sqrt{n})$, and $d_0 = o(n^{1/6})$, we infer that

$$\frac{d_0^3 K'''(\xi)}{6\sigma_L^3} \to 0 \tag{163}$$

as $n \to \infty$. Hence,

$$\exp\left(K(t_0) - \varepsilon t_0\right) \asymp \exp\left(\frac{d_0^2}{2} - bd_0\right). \tag{164}$$

Writing $d_0 = \tau_0 \cdot (b + \sqrt{b^2+4})/2$, so $\tau_0 > 0$ and $\tau_0 \to 1$ by (159), then collecting terms, we obtain

$$\frac{d_0^2}{2} - bd_0 = \frac{\tau_0^2}{2} - (2-\tau_0)\tau_0 \cdot \frac{b^2 + b\sqrt{b^2+4}}{4}. \tag{165}$$

Therefore, we obtain that

$$\exp\left(K(t_0) - \varepsilon t_0\right) \asymp \frac{\sqrt{e}}{C(b)^\tau} \tag{166}$$

where $\tau \triangleq (2-\tau_0)\tau_0 \to 1$. Putting the asymptotics shown above together, we conclude that

$$\mathrm{err}_{\mathrm{SP}}(\varepsilon; t_0) \asymp \frac{1.12\sqrt{e}\,\mathrm{P}}{\mathrm{V}^{3/2} \cdot C(b)^\tau \cdot \sqrt{n}}, \tag{167}$$

as desired.

L. Instantiation of the Saddle-point Accountant

The algorithm SADDLEPOINTACCOUNTANT (Algorithm 1), giving the workflow of the versions of the SPA, is presented here.

Algorithm 1 : SADDLEPOINTACCOUNTANT (SPA)

Input: A finite set $\mathcal{E} \subset [0,\infty)$ (values of ε), and tightly dominating distributions $(P_1, Q_1), \ldots, (P_n, Q_n)$.

Output: Four approximations $\delta_{L,\text{SP-MSD}}^{(k)}$, $1 \le k \le 3$, and $\delta_{L,\,\text{SP-CLT}}$ of the privacy curve δ_L, and an error bound so that $|\delta_L(\varepsilon) - \delta_{L,\,\text{SP-CLT}}(\varepsilon)| \le \mathrm{err}_{\mathrm{SP}}(\varepsilon)$.

$L_j \leftarrow \log \frac{dP_j}{dQ_j}(X_j)$ where $X_j \sim P_j$ $\quad j \in [n]$

$K_{L_j}(t) \leftarrow \log \mathbb{E}\left[e^{tL_j}\right]$ $\quad j \in [n]$

$L \leftarrow L_1 + \cdots + L_n$

$K_L \leftarrow K_{L_1} + \cdots + K_{L_n}$

for $\varepsilon \in \mathcal{E}$ **do**

  $t_0 \leftarrow$ positive solution to $K_L'(t_0) = \varepsilon + \frac{1}{t_0} + \frac{1}{t_0+1}$

  $F_\varepsilon(t) \leftarrow K_L(t) - \varepsilon t - \log t - \log(t+1)$

  $\beta_{\varepsilon,2} \leftarrow \frac{1}{8}\frac{F_\varepsilon^{(4)}(t_0)}{F_\varepsilon''(t_0)^2}$

  $\beta_{\varepsilon,3} \leftarrow -\frac{5}{24}\frac{F_\varepsilon^{(3)}(t_0)^2}{F_\varepsilon''(t_0)^3} - \frac{1}{48}\frac{F_\varepsilon^{(6)}(t_0)}{F_\varepsilon''(t_0)^3}$

  $\delta_{L,\,\text{SP-MSD}}^{(1)}(\varepsilon) \leftarrow \frac{e^{F_\varepsilon(t_0)}}{\sqrt{2\pi F_\varepsilon''(t_0)}}$

  $\delta_{L,\,\text{SP-MSD}}^{(2)}(\varepsilon) \leftarrow \frac{e^{F_\varepsilon(t_0)}}{\sqrt{2\pi F_\varepsilon''(t_0)}}\left(1 + \beta_{\varepsilon,2}\right)$

  $\delta_{L,\,\text{SP-MSD}}^{(3)}(\varepsilon) \leftarrow \frac{e^{F_\varepsilon(t_0)}}{\sqrt{2\pi F_\varepsilon''(t_0)}}\left(1 + \beta_{\varepsilon,2} + \beta_{\varepsilon,3}\right)$

  $\gamma \leftarrow \frac{K_L'(t_0) - \varepsilon}{\sqrt{K_L''(t_0)}}$

  $(\alpha, \beta) \leftarrow \left(\sqrt{K_L''(t_0)}\, t_0 - \gamma,\ \sqrt{K_L''(t_0)}\,(t_0+1) - \gamma\right)$

  $\delta_{L,\,\text{SP-CLT}}(\varepsilon) \leftarrow e^{K_L(t_0) - \varepsilon t_0 - \gamma^2/2}\frac{q(\alpha) - q(\beta)}{\sqrt{2\pi}}$

  $\tilde{L}_j \leftarrow$ exp. tilt of L_j with parameter t_0, $\quad j \in [n]$

  $\mathrm{P}_{t_0} \leftarrow \sum_{j\in[n]} \mathbb{E}\left[\left|\tilde{L}_j - K_{L_j}'(t_0)\right|^3\right]$

  $\mathrm{err}_{\mathrm{SP}}(\varepsilon) \leftarrow e^{K_L(t_0) - \varepsilon t_0}\frac{t_0^{t_0}}{(1+t_0)^{1+t_0}} \cdot \frac{1.12\,\mathrm{P}_{t_0}^{(n)}}{K_L''(t_0)^{3/2}}$

end for

Return: $\delta_{L,\text{SP-MSD}}^{(k)}, 1 \le k \le 3, \delta_{L,\,\text{SP-CLT}}, \mathrm{err}_{\mathrm{SP}}$.

M. Ground-Truth Curve Computation

We explain here how the ground-truth curve in Figure 2 is computed. Since the setting there is for self-composition, we employ that here too. So, let $L_1, \cdots, L_n$ be i.i.d. PLRVs for the subsampled Gaussian mechanism, and consider the PLRV $L = L_1 + \cdots + L_n$ for the composed mechanism.

Recall that the saddle-point accountant gives various approximations to the contour integral given in Theorem 3.1, which we copy here:

$$\delta_{L_1}(\varepsilon) = \frac{1}{2\pi i} \int_{t-i\infty}^{t+i\infty} e^{F_\varepsilon(z)}\, dz \tag{168}$$

where the function F_ε is defined as:

$$F_\varepsilon(z) = K_{L_1}(z) - \varepsilon z - \log z - \log(1+z). \tag{169}$$

After n compositions, the contour integral becomes:

$$\delta_L(\varepsilon) = \frac{1}{2\pi i} \int_{t-i\infty}^{t+i\infty} e^{nK_{L_1}(z) - \varepsilon z - \log z - \log(1+z)}\, dz. \tag{170}$$

Recall that this formula holds for any value of $t > 0$.

The ground-truth in (170) is then computed via standard numerical integration, which evidently is a time-consuming process, yet it is one that can produce a reference value to relatively compare accountants' accuracies.

Let $P = \mathcal{N}(0, \sigma^2)$, $Q = (1-\lambda)\mathcal{N}(0, \sigma^2) + \lambda\mathcal{N}(s, \sigma^2)$. The composed subsampled Gaussian has the PLRV $L = L_1 + \cdots + L_n$, where the L_j are independent and (see Lemma 2.2)

$$\begin{aligned} &L_j = \log \frac{dQ}{dP}(X) = \log\left(1 - \lambda + \lambda e^{s(2X-s)/(2\sigma^2)}\right), \\ &X \sim (1-\lambda)\mathcal{N}(0, \sigma^2) + \lambda\mathcal{N}(s, \sigma^2). \end{aligned} \tag{171}$$

In addition, the MGF of L_1 may be written as

$$M_{L_1}(z) = \mathbb{E}[e^{zL_1}] \tag{172}$$

$$= \mathbb{E}_{X\sim Q}\left[\left(\frac{dQ}{dP}(X)\right)^z\right] \tag{173}$$

$$= \mathbb{E}_{X\sim P}\left[\left(\frac{dQ}{dP}(X)\right)^{z+1}\right] \tag{174}$$

$$= \int_{-\infty}^{\infty} \left(1 - \lambda + \lambda e^{s(2x-s)/(2\sigma^2)}\right)^{z+1} dP(x). \tag{175}$$

Recall that the CGF is given by

$$K_{L_1}(z) = \log M_{L_1}(z). \tag{176}$$

Plugging in the log integral (176) into the contour integral (170), the contour integral can be directly computed using standard numerical libraries. We note that this calculation is very slow, as the integrand in (170) itself involves an integral over $\mathbb{R}$. Moreover, we numerically invert this function via bisection to obtain the curve described in Figure 2. This ground-truth curve was computed on a 64-core cluster using multi-processing to distribute the workload, and took a wall-time of 45 minutes. This amounts to a runtime of 48 CPU hours. In contrast, all other accountants run in the order of seconds on a commercial laptop.

N. Additional Numerical Experiments

We provide further experiments exploring the flexibility of the saddle-point accountant. We show that the SPA-MSD approximations can be accurate even in the moderate-composition regime, though the SPA-CLT bounds become loose for a small number of compositions. We demonstrate this using parameters used by a real-world application of DP on the image classification SGD algorithm in (De et al., 2022), which uses the subsampled Gaussian as the DP mechanism. In particular, we use the noise scale $\sigma = 9.4$ and subsampling rate $\lambda = 2^{14}/50000$, as these were the values that allowed a 40-layer Wide-ResNet to achieve a new SOTA accuracy of 81.4% on CIFAR-10 under $(\varepsilon = 8, \delta = 10^{-5})$-DP. This algorithm went up to $n = 2000$ compositions to achieve this SOTA.

First, we plot the (ε, δ)-curves at $n \in \{100, 250, 500, 2000\}$ compositions in Figure 4. We observe that the CLT bounds get tighter as the number of compositions increases, but the order-1 SPA-MSD remains consistently accurate for all presented compositions and values of δ.

Second, we demonstrate the accuracy of the order-1 SPA-MSD for all compositions less than 2000 in Figure 5, where we fix $\delta = 10^{-5}$, vary the number of compositions, and plot the resulting value of ε.

These two plots verify that the order-1 SPA-MSD is much more accurate than the CLT bounds suggest.

Finally, to demonstrate the applicability of the saddle-point accountant to other mechanisms beyond the subsampled Gaussian. In Figure 6 we present results for the composed Laplace mechanism. In particular, we borrow the mechanism parameters used in the PRV Accountant notebook showcasing their implementation, i.e., $n = 1000$ compositions with a sensitivity of $s = 0.01$ and a shape parameter $b = 1$. We compare against the Moments Accountant (using the Laplace cumulant), PRV Accountant and Connect the Dots implementation, in the same spirit as Figure 1. Note that we use $\varepsilon_{\text{error}} = 0.01$ for the PRV Accountant, which is an order of magnitude finer than the $\varepsilon_{\text{error}} = 0.1$ used in the linked notebook, for demonstration purposes. Once again, the SPA upper and lower bounds have a narrow gap between them.

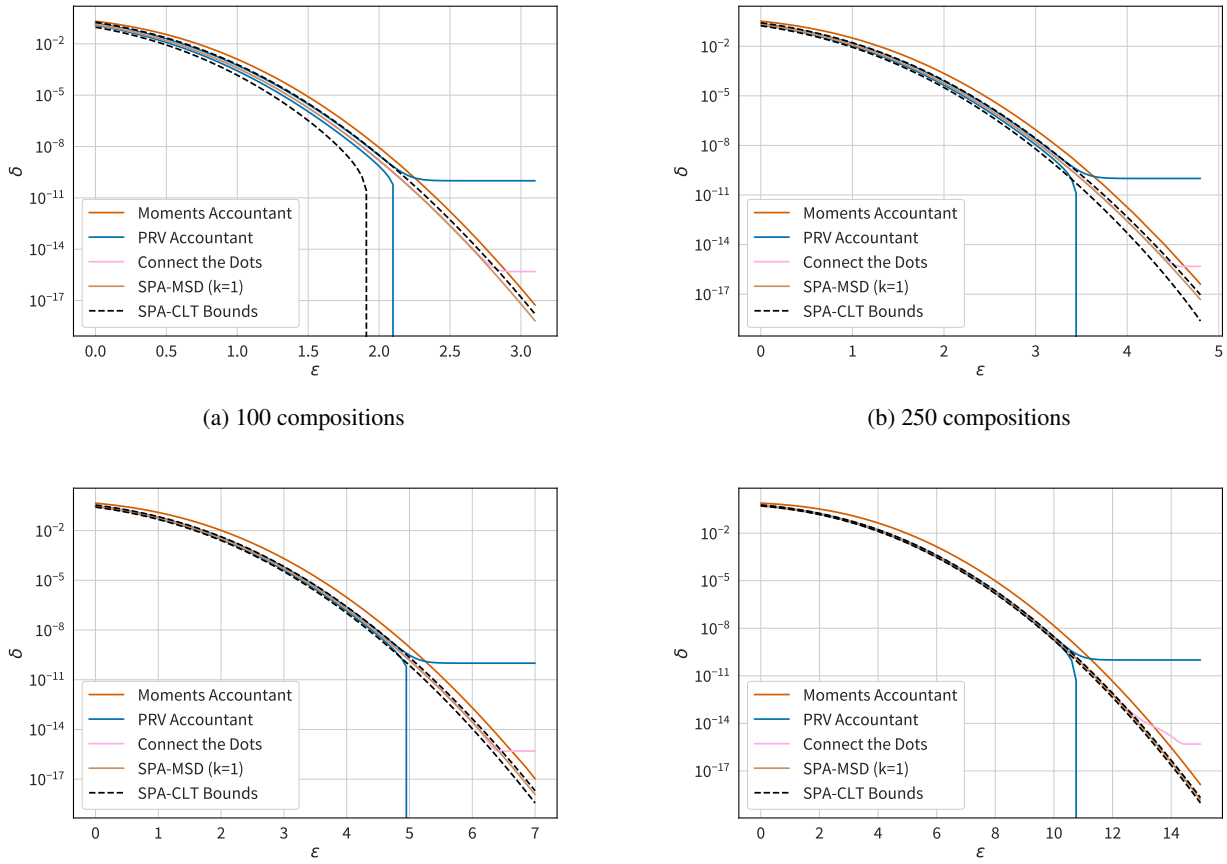

(a) 100 compositions

(b) 250 compositions

(c) 500 compositions

(d) 2000 compositions

Figure 4. Accounting for the composition of $n \in \{100, 250, 500, 2000\}$ subsampled Gaussian mechanisms, with noise scale $\sigma = 9.4$ and subsampling rate $\lambda = 2^{14}/50000$. The PRV Accountant (Gopi et al., 2021) discretization parameters are $\varepsilon_{\text{error}} = 0.1$, $\delta_{\text{error}} = 10^{-10}$. The Connect the Dots (Doroshenko et al., 2022) discretization interval length is 0.005.

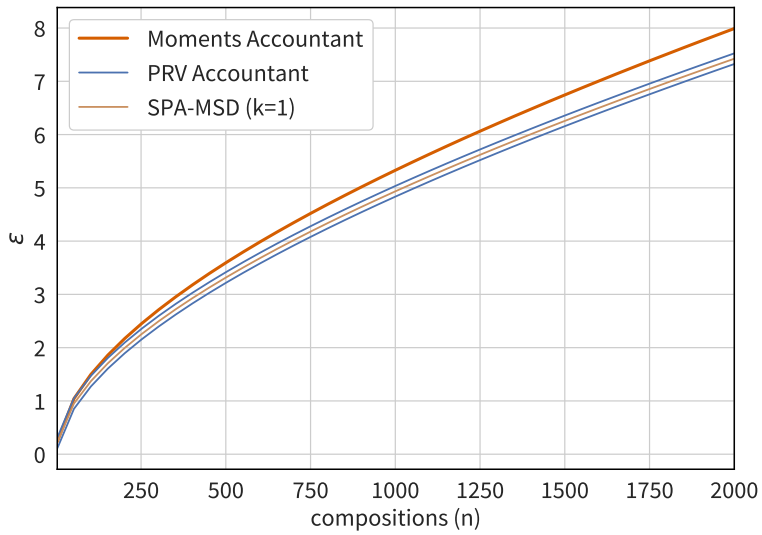

Figure 5. Privacy budget ε of the subsampled Gaussian mechanism after $1 \leq n \leq 2000$ compositions using the order-1 SPA-MSD, the Moments Accountant, and the PRV Accountant (Gopi et al., 2021). We use subsampling $\lambda = 2^{14}/50000$, noise scale $\sigma = 9.4$, and $\delta = 10^{-5}$. The discretization parameters for the PRV Accountant are $\varepsilon_{\text{error}} = 0.1, \delta_{\text{error}} = 10^{-10}$.

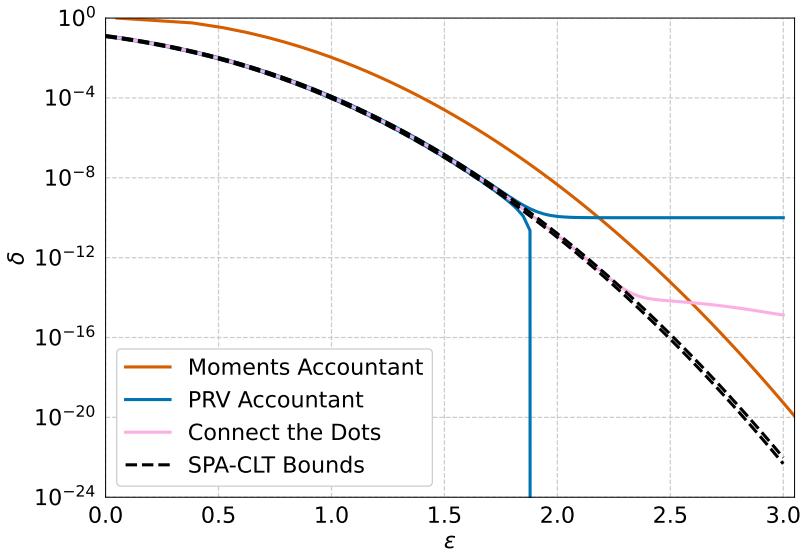

Figure 6. Accounting for the composition of 1000 Laplace mechanisms having shape parameter $b = 1$ and with sensitivity $s = 0.01$. The remaining FFT discretization parameters are set to $\varepsilon_{\text{error}} = 0.01, \delta_{\text{error}} = 10^{-10}$ for the PRV Accountant (Gopi et al., 2021), and discretization interval length of 2×10^{-4} for Connect the Dots (Doroshenko et al., 2022).