Tight Analysis of Privacy and Utility Tradeoff in Approximate Differential Privacy

Quan Geng Google Research
Wei Ding Google Research
Ruiqi Guo Google Research
Sanjiv Kumar Google Research

Abstract

We characterize the minimum noise amplitude and power for noise-adding mechanisms in (ϵ, δ)-differential privacy for single real-valued query function. We derive new lower bounds using the duality of linear programming, and new upper bounds by analyzing a special class of (ϵ, δ)-differentially private mechanisms, the *truncated Laplacian* mechanisms. We show that the multiplicative gap of the lower bounds and upper bounds goes to zero in various high privacy regimes, proving the tightness of the lower and upper bounds. In particular, our results close the previous constant multiplicative gap in the discrete setting. Numeric experiments show the improvement of the truncated Laplacian mechanism over the optimal Gaussian mechanism in all privacy regimes.

1 Introduction

Differential privacy, introduced by Dwork et al. (2006b), is a framework to quantify to what extent individual privacy in a statistical dataset is preserved while releasing useful aggregate information about the dataset. Differential privacy provides strong privacy guarantees by requiring the near-indistinguishability of whether an individual is in the dataset or not based on the released information. For more motivation and background of differential privacy, we refer the readers to the survey by Dwork (2008) and the book by Dwork and Roth (2014).

Since its introduction, differential privacy has spawned a large body of research in differentially private data-releasing mechanism design, and the noise-adding mechanism has been applied in many machine learning algorithms to preserve differential privacy, e.g., logistic regression (Chaudhuri and Monteleoni, 2008), empirical risk minimization (Chaudhuri et al., 2011; Wang et al., 2018), online learning (Jain et al., 2012), statistical risk minimization (Duchi et al., 2012), statistical learning (Dziugaite and Roy, 2018), deep learning (Shokri and Shmatikov, 2015; Abadi et al., 2016; Phan et al., 2016) distributed optimization (Agarwal et al., 2018), hypothesis testing (Sheffet, 2018), matrix completion (Jain et al., 2018), expectation maximization (Park et al., 2017), and principal component analysis (Chaudhuri et al., 2012; Ge et al., 2018).

The classic differential privacy is called ϵ-differential privacy, which imposes an upper bound e^ϵ on the multiplicative distance of the probability distributions of the randomized query outputs for any two neighboring datasets. The standard approach for preserving ϵ-differential privacy is adding a noise with the Laplacian distribution to the query output. Introduced by Dwork et al. (2006a), the approximate differential privacy is (ϵ, δ)-differential privacy, and the common interpretation of (ϵ, δ)-differential privacy is that it is ϵ-differential privacy "except with probability δ" (Mironov, 2017). The standard approach for preserving (ϵ, δ)-differential privacy is the Gaussian mechanism, which adds a Gaussian noise to the query output.

To fully make use of the differentially private mechanisms, it is important to understand the fundamental trade-off between privacy and utility (accuracy). For example, within the class of noise-adding mechanisms, given the privacy constraint ϵ and δ, we are interested in deriving the minimum amount of noise added to achieve the highest accuracy and utility while preserving the differential privacy. In the literature, there have been many works on optimal differential privacy mechanism design and characterizing the privacy and utility tradeoff in differential privacy. For a single count query function under ϵ-differential privacy, Ghosh et al. (2009) show that the geometric mechanism is universally optimal under a Bayesian framework, and Gupte and Sundararajan (2010) derived the optimal noise proba-

Proceedings of the 23[rd]International Conference on Artificial Intelligence and Statistics (AISTATS) 2020, Palermo, Italy. PMLR: Volume 108.

bility distributions under a minimax cost framework. Geng and Viswanath (2016b) show that the optimal noise distribution has a staircase-shaped probability density function for single real-valued query function under ϵ-differential privacy, and Geng et al. (2015) generalized the result to two-dimensional query functions. Soria-Comas and Domingo-Ferrer (2013) also independently derived the staircase-shaped noise probability distribution under a different optimization framework.

Geng and Viswanath (2016a) show that for a single integer-valued query function under (ϵ, δ)-differential privacy, the discrete uniform noise distribution and the discrete Laplacian noise distribution are asymptotically optimal within a constant multiplicative gap in the high privacy regions. Balle and Wang (2018) improved the classic analysis of the Gaussian mechanism for (ϵ, δ)-differential in the high privacy regime ($\epsilon \to 0$), and developed an optimal Gaussian mechanism whose variance is calibrated directly using the Gaussian cumulative density function instead of a tail bound approximation. Geng et al. (2019) derive the optimal noise-adding mechanism for single real-valued query function under $(0, \delta)$-differential privacy, and show that a uniform noise distribution with probability mass at the origin is optimal for a large class of cost functions.

1.1 Our Contributions

In this work, we characterize the minimum noise amplitude and power for noise-adding mechanisms in (ϵ, δ)-differential privacy for single real-valued query function. Our contributions are three-fold:

First, we analyze a new class of (ϵ, δ)-differentially private noise-adding mechanisms, *truncated Laplacian* mechanisms. Applying the truncated Laplacian mechanism, we derive new achievable upper bounds on minimum noise amplitude and noise power in (ϵ, δ)-differential privacy for single real-valued query function. The key insights from the new mechanisms design are that the noise probability density function shall decay as fast as possible while being ϵ-differentially private when the noise is small, and then sharply reduce to zero when the noise is big, to avoid a heavy tail distribution which would incur a high cost.

Second, we derive new lower bounds on the minimum noise amplitude and minimum noise power. The key technique is to discretize the continuous probability distribution and the loss function, and transform the continuous functional optimization problem to linear programming. Applying the lower bound result in Geng and Viswanath (2016a) for *integer-valued* query function, which is based on the duality of linear programming, we derive new lower bounds for *real-valued* query functions under (ϵ, δ)-differential privacy.

Third, we show that the multiplicative gap of the lower bounds and upper bounds goes to zero in various high privacy regimes, proving the tightness of the lower and upper bounds, and thus establish the optimality of the truncated Laplacian mechanism for minimizing the noise amplitude and noise power under (ϵ, δ)-differential privacy. In particular, our result closes the previous constant multiplicative gap between the lower bound and the upper bound (using discrete uniform distribution and discrete Laplacian distribution) in Geng and Viswanath (2016a).

Comprehensive numeric experiments show the improvement of the truncated Laplacian mechanism over the optimal Gaussian mechanism in Balle and Wang (2018) by significantly reducing the noise amplitude and noise power in all privacy regimes.

1.2 Organization

The paper is organized as follows. In Section 2, we give some preliminaries on differential privacy, and derive the (ϵ, δ)-differential privacy constraint on the additive noise probability distribution and define the minimum noise amplitude and noise power under (ϵ, δ)-differential privacy. Section 3 presents the truncated Laplacian mechanism for preserving (ϵ, δ)-differential privacy, and derives new upper bounds for minimum noise amplitude and noise power. Section 4 derives new lower bounds on the minimum noise magnitude and noise power. Section 5 shows that the multiplicative gap between the lower bounds and the upper bounds goes to zero in various privacy regimes, and thus proves the tightness of the new lower and upper bounds. Section 6 conducts comprehensive numeric experiments to compare the performance of the truncated Laplacian mechanism with the optimal Gaussian mechanisms, and demonstrates the improvement in all privacy regimes. Section 7 discusses some additional properties of the truncated Laplacian mechanism and concludes this paper.

2 Problem Formulation

In this section, we first give some preliminaries on differential privacy, and then define the minimum noise amplitude V_1^* and minimum noise power V_2^* for (ϵ, δ)-differentially private noise-adding mechanisms.

Consider a real-valued query function $q : \mathcal{D} \to \mathbb{R}$, where $\mathcal{D}$ is the set of all possible datasets. The real-valued query function q will be applied to a dataset, and the query output is a real number. Two datasets $D_1, D_2 \in \mathcal{D}$ are called neighboring datasets if they differ in at most one element, i.e., one is a proper

subset of the other and the larger dataset contains just one additional element Dwork (2008). A randomized query-answering mechanism $\mathcal{K}$ for the query function q will randomly output a number with probability distribution depending on query output $q(D)$, where D is the dataset.

Definition 1 ((ϵ, δ)-differential privacy (Dwork et al., 2006a)). *A randomized mechanism $\mathcal{K}$ gives (ϵ, δ)-differential privacy if for all data sets D_1 and D_2 differing on at most one element, and for any measurable set $S \subset Range(\mathcal{K})$,*

$$Pr[\mathcal{K}(D_1) \in S] \leq e^{\epsilon}\ Pr[\mathcal{K}(D_2) \in S] + \delta. \qquad (1)$$

The sensitivity of a real-valued query function measures how the query changes for neighboring datasets.

Definition 2 (Query Sensitivity). *The sensitivity of q is defined as*

$$\Delta := \max_{D_1, D_2 \in \mathcal{D}} |q(D_1) - q(D_2)|,$$

for all D_1, D_2 differing in at most one element.

A standard approach for preserving differential privacy is query-output independent noise-adding mechanisms, where a random noise is added to the query output. Given a dataset D, a query-output independent noise-adding mechanism $\mathcal{K}$ will release the query output $t = q(D)$ corrupted by an additive random noise X with probability distribution $\mathcal{P}$:

$$\mathcal{K}(D) = t + X.$$

We derive the differential privacy constraint on the noise probability distribution $\mathcal{P}$ in Lemma 1, which is essentially an extension of Equation (18) in Geng and Viswanath (2016a) to the continous case.

Lemma 1. *Given the query sensitivity Δ and privacy parameters ϵ and δ, the noise probability distribution $\mathcal{P}$ preserves (ϵ, δ)-differential privacy if and only if*

$$\mathcal{P}(S) - e^{\epsilon}\mathcal{P}(S + d) \leq \delta, \forall\ |d| \leq \Delta, \textit{measurable set } S \subset \mathbb{R}. \qquad (2)$$

Proof. The differential privacy constraint (1) on $\mathcal{K}$ is that for any $t_1, t_2 \in \mathbb{R}$ such that $|t_1 - t_2| \leq \Delta$ (corresponding to the query outputs for two neighboring datasets[1]),

$$\mathcal{P}(S - t_1) \leq e^{\epsilon}\mathcal{P}(S - t_2) + \delta, \forall \text{ measurable set } S \subset \mathbb{R}, \qquad (3)$$

where $\forall t \in \mathbb{R}$, $S + t$ is defined as the set $\{s + t \,|\, s \in S\}$.

Since (3) has to hold for any measurable set S and any $|t_1 - t_2| \leq \Delta$, equivalently, we have

$$\mathcal{P}(S) \leq e^{\epsilon}\mathcal{P}(S + d) + \delta, \forall\ |d| \leq \Delta, \text{measurable set } S \subset \mathbb{R}.$$

□

Let $\mathcal{P}_{\epsilon,\delta}$ denote the set of noise probability distributions satisfying the (ϵ, δ)-differential privacy constraint (2). Given $\mathcal{P} \in \mathcal{P}_{\epsilon,\delta}$, the expected noise amplitude and noise power are $\int_{x\in\mathbb{R}} |x|\mathcal{P}(dx)$ and $\int_{x\in\mathbb{R}} x^2\mathcal{P}(dx)$. The goal of this work is to characterize the minimum expected noise amplitude and noise power under (ϵ, δ)-differential privacy. More precisely, define

$$V_1^* := \inf_{\mathcal{P}\in\mathcal{P}_{\epsilon,\delta}} \int_{x\in\mathbb{R}} |x|\mathcal{P}(dx) \quad \text{(min noise amplitude)},$$

$$V_2^* := \inf_{\mathcal{P}\in\mathcal{P}_{\epsilon,\delta}} \int_{x\in\mathbb{R}} x^2\mathcal{P}(dx) \quad \text{(min noise power)}.$$

In this work, we characterize V_1^* and V_2^* in terms of Δ, ϵ, δ by deriving tight lower bounds V_1^{low}, V_2^{low} and upper bounds V_1^{upp}, V_2^{upp} such that $V_1^{low} \leq V_1^* \leq V_1^{upp}$ and $V_2^{low} \leq V_2^* \leq V_2^{upp}$.

In the next section, we present the new upper bounds V_1^{upp} and V_2^{upp}. The lower bounds V_1^{low} and V_2^{low} are presented in Section 4.

3 Upper Bound: Truncated Laplacian Mechanism

In this section, we present a new class of (ϵ, δ)-differentially private noise-adding mechanism, *truncated Laplacian* mechanism. Applying the truncated Laplacian mechanism, we derive new achievable (and tight) upper bounds V_1^{upp} and V_2^{upp} on minimum noise amplitude V_1^* and minimum noise power V_2^* in Theorem 2 and Theorem 3.

Before presenting the exact form of the truncated Laplacian mechanism, we first discuss some key ideas and insights behind the new mechanism design.

The standard Laplacian distribution for preserving ϵ-differential privacy has a symmetric probability density function $f(x) = \frac{\epsilon}{2\Delta}e^{-\frac{|x|\epsilon}{\Delta}}$. Note that for any $x \geq 0$, the probability density decay rate, $\frac{f(x)}{f(x+\Delta)}$, is exactly e^{ϵ}. Geng and Viswanath (2016b) show that the decay rate e^{ϵ} is optimal under ϵ-differential privacy. Indeed, if the decay rate is higher, it is no longer ϵ-differentially private; if the decay rate is lower, it will incur a higher cost. However, under (ϵ, δ)-differential privacy, Laplacian distribution is not optimal as it has a heavy tail distribution.

[1] In this work we impose no prior on the query function other than the query sensitivity Δ. For any $t_1, t_2 \in \mathbb{R}$ such that $|t_1 - t_2| \leq \Delta$, there may exist two neighboring datasets D_1 and D_2 with $q(D_1) = t_1$ and $q(D_2) = t_2$.

(ϵ, δ)-differential privacy relaxes the ϵ-differential privacy constraint, and it allows that for a set of points with a probability mass δ, the decay rate can exceed e^ϵ. The Gaussian mechanism is widely used in (ϵ, δ)-differential privacy, and for $x > 0$, its probability density decay rate is $\frac{f(x)}{f(x+\Delta)} = \frac{e^{-\frac{x^2}{\sigma^2}}}{e^{-\frac{(x+\Delta)^2}{\sigma^2}}} = e^{\frac{\Delta^2+2\Delta x}{\sigma^2}} = e^{\frac{\Delta^2}{\sigma^2}} e^{\frac{2\Delta}{\sigma^2}x}$, which is exponentially increasing with respect to x. When x is big, the decay rate can be very high. While the Gaussian mechanism addresses the long tail distribution to some extent by having higher decay rate for large x, the decay rate is smaller than e^ϵ when x is small.

Motivated by the observation that under (ϵ, δ)-differential privacy, the decay rate shall be as high as possible without exceeding e^ϵ, except for a set of points with a probability mass δ (for those there is no limit on the decay rate), we derive a symmetric truncated Laplacian distribution where the probability density decay rate is exactly e^ϵ, except for a set of points with probability mass δ where the decay rate is infinite.

Definition 3 (Truncated Laplacian Distribution)**.** *Given the privacy parameters $0 < \delta < \frac{1}{2}, \epsilon > 0$ and the query sensitivity $\Delta > 0$, the probability density function of the truncated Laplacian distribution $\mathcal{P}_{TLap}$ is defined as:*

$$f_{TLap}(x) := \begin{cases} Be^{-\frac{|x|}{\lambda}}, & \text{for } x \in [-A, A] \\ 0, & \text{otherwise} \end{cases} \tag{4}$$

where

$$\lambda := \frac{\Delta}{\epsilon},$$
$$A := \frac{\Delta}{\epsilon} \log(1 + \frac{e^\epsilon - 1}{2\delta}),$$
$$B := \frac{1}{2\lambda(1 - e^{-\frac{A}{\lambda}})} = \frac{1}{2\frac{\Delta}{\epsilon}(1 - \frac{1}{1+\frac{e^\epsilon-1}{2\delta}})}.$$

f_{TLap} is a valid probability density function, as $f_{\text{TLap}}(x) \geq 0$ and $\int_{x\in\mathbb{R}} f_{\text{TLap}}(x)dx = \int_0^A 2Be^{-\frac{|x|}{\lambda}}dx = 2\lambda B(1 - e^{-\frac{A}{\lambda}}) = 1$.

We discuss the key properties of the symmetric probability density function $f_{\text{TLap}}(x)$:

- The decay rate in $[0, A - \Delta]$ is exactly e^ϵ, i.e., $\frac{f_{\text{TLap}}(x)}{f_{\text{TLap}}(x+\Delta)} = e^\epsilon$, $\forall x \in [0, A - \Delta]$.
- The probability mass in the interval $[A - \Delta, A]$ is δ, i.e., $\mathcal{P}_{\text{TLap}}([A - \Delta, A]) = \delta$.

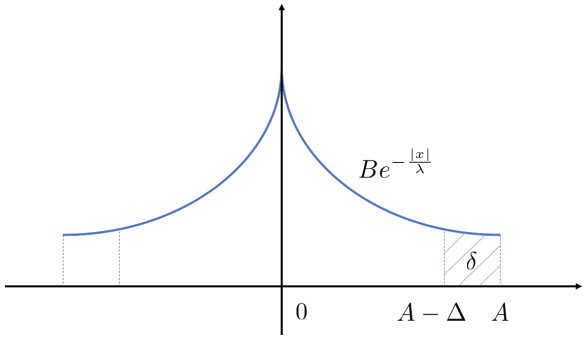

Figure 1: Noise probability density function f_{TLap} of the truncated Laplacian mechanism. f_{TLap} is a symmetric truncated exponentinal function with a probability mass δ in the last interval with length Δ in the support of f_{TLap}, i.e., the interval $[A - \Delta, A]$. The decay rate $\frac{f_{\text{TLap}}(x)}{f_{\text{TLap}}(x+\Delta)}$ is exactly e^ϵ for $x \in [0, A - \Delta)$. The parameters A and B are then derived by solving the equations that $\int_{x\in\mathbb{R}} f_{\text{TLap}}(x)dx = 1$ and $\int_{A-\Delta}^{A} f_{\text{TLap}}(x)dx = \delta$.

- The decay rate $\frac{f_{\text{TLap}}(x)}{f_{\text{TLap}}(x+\Delta)}$ is $+\infty$ for $x \in (A - \Delta, A]$, as $f_{\text{TLap}}(x) = 0$ for $x \in (A, +\infty)$.

Definition 4 (Truncated Laplacian mechanism)**.** *Given the query sensitivity Δ, and the privacy parameters ϵ, δ, the truncated Laplacian mechanism adds a noise with probability distribution $\mathcal{P}_{TLap}$ defined in* (4) *to the query output.*

Theorem 1. *The truncated Laplacian mechanism preserves (ϵ, δ)-differential privacy.*

Proof. Equivalently, we need to show that the truncated Laplacian distribution $\mathcal{P}_{\text{TLap}}$ defined in (4) satisfies the (ϵ, δ)-differential privacy constraint (2).

We are interested in maximizing $\mathcal{P}_{\text{TLap}}(S) - e^\epsilon \mathcal{P}_{\text{TLap}}(S + d)$ in (2) and show that the maximum over $S \subseteq \mathbb{R}$ is upper bounded by δ. Since $f_{\text{TLap}}(x)$ is symmetric and monotonically decreasing in $[0, +\infty)$, without loss of generality, we can assume $d \geq 0$ and thus $d \in [0, \Delta]$.

To maximize $\mathcal{P}_{\text{TLap}}(S) - e^\epsilon \mathcal{P}_{\text{TLap}}(S + d)$, S shall not contain points in $(-\infty, -\frac{\Delta}{2}]$, as

$$f_{\text{TLap}}(x) \leq f_{\text{TLap}}(x + d), \forall x \in (-\infty, -\frac{\Delta}{2}].$$

S shall not contain points in $[-\frac{\Delta}{2}, A - \Delta]$, as

$$f_{\text{TLap}}(x) \leq e^\epsilon f_{\text{TLap}}(x + d), \forall x \in [-\frac{\Delta}{2}, A - \Delta].$$

Therefore, $\mathcal{P}_{\text{TLap}}(S) - e^\epsilon \mathcal{P}_{\text{TLap}}(S + d)$ is maximized for some set $S \subseteq [A - \Delta, +\infty)$. Since

$f_{\text{TLap}}(x)$ is monotonically decreasing in $[A-\Delta, +\infty)$, $\mathcal{P}_{\text{TLap}}(S) - e^{\epsilon}\mathcal{P}_{\text{TLap}}(S+d)$ is maximized at $S = [A-\Delta, +\infty)$ and the maximum value is $\int_{A-\Delta}^{A-\Delta+d} f(x)dx \leq \int_{A-\Delta}^{A} f_{\text{TLap}}(x)dx = \delta$.

We conclude that $\mathcal{P}_{\text{TLap}}$ satisfies the (ϵ, δ)-differential privacy constraint (2). □

Next, we apply the truncated Laplacian mechanism to derive new upper bounds on the minimum noise amplitude V_1^* and noise power V_2^*.

Theorem 2 (Upper Bound on Minimum Noise Amplitude)**.**

$$V_1^* \leq V_1^{upp} := \frac{\Delta}{\epsilon}\left(1 - \frac{\log(1+\frac{e^{\epsilon}-1}{2\delta})}{\frac{e^{\epsilon}-1}{2\delta}}\right). \tag{5}$$

Proof. We can compute the expected noise amplitude for the truncated Laplacian distribution $\mathcal{P}_{\text{TLap}}$ defined in (4) via

$$\begin{aligned} V_1^{upp} &:= \int_{x\in\mathbb{R}} f_{\text{TLap}}(x)|x|dx = 2\int_0^A Be^{-\frac{x}{\lambda}}xdx \\ &= \frac{\Delta}{\epsilon}\left(1 - \frac{\log(1+\frac{e^{\epsilon}-1}{2\delta})}{\frac{e^{\epsilon}-1}{2\delta}}\right). \end{aligned}$$

Since the truncated Laplacian mechanism preserves (ϵ, δ)-differential privacy, this gives an upper bound on the minimum noise amplitude V_1^* under (ϵ, δ)-differential privacy. □

In Theorem 2, the upper bound V_1^{upp} is composed of two parts. The first part is $\frac{\Delta}{\epsilon}$, which is the noise amplitude of the Laplacian mechanism under ϵ-differential privacy. The second part reduces the noise by a portion of $\frac{\log(1+\frac{e^{\epsilon}-1}{2\delta})}{\frac{e^{\epsilon}-1}{2\delta}}$ due to the δ-relaxation in (ϵ, δ)-differential privacy.

We analyze the asympotic properties of V_1^{upp} in the high privacy regimes as $\epsilon \to 0, \delta \to 0$:

- Given ϵ, $\lim_{\delta\to 0} V_1^{upp} = \frac{\Delta}{\epsilon}$. The truncated Laplacian mechanism will be reduced to the standard Laplacian mechanism as $\delta \to 0$,.
- Given δ, $\lim_{\epsilon\to 0} V_1^{upp} = \frac{\Delta}{4\delta}$. Indeed, when $\epsilon \to 0$, $\frac{e^{\epsilon}-1}{2\delta} \to 0$, and thus[2]

$$\begin{aligned} V_1^{upp} &\approx \frac{\Delta}{\epsilon}\left(1 - \frac{\frac{e^{\epsilon}-1}{2\delta} - \frac{(\frac{e^{\epsilon}-1}{2\delta})^2}{2}}{\frac{e^{\epsilon}-1}{2\delta}}\right) \\ &= \frac{\Delta}{\epsilon}\frac{\frac{e^{\epsilon}-1}{2\delta}}{2} \approx \frac{\Delta}{\epsilon}\frac{\epsilon}{4\delta} = \frac{\Delta}{4\delta}. \end{aligned}$$

[2]Note that the additive error terms in the analysis below can be more precisely characterized using the Taylor expansion: $\log(1+x) = x + O(x^2)$, and $e^x = 1 + x + O(x^2)$.

  As $\epsilon \to 0$, the truncated Laplacian distribution is reduced to a uniform distribution in the interval $[-\frac{\Delta}{2\delta}, \frac{\Delta}{2\delta}]$ with probability density $\frac{\delta}{\Delta}$.

- In the regime $\delta = \epsilon \to 0$, the upper bound

$$\begin{aligned} V_1^{upp} &= \frac{\Delta}{\epsilon}\left(1 - \frac{\log(1+\frac{e^{\epsilon}-1}{2\epsilon})}{\frac{e^{\epsilon}-1}{2\epsilon}}\right) \\ &\approx \frac{\Delta}{\epsilon}\left(1 - \frac{\log(1+\frac{\epsilon}{2\epsilon})}{\frac{\epsilon}{2\epsilon}}\right) \\ &= \frac{\Delta}{\epsilon}(1 - 2\log\frac{3}{2}). \end{aligned} \tag{6}$$

  In Section 5, we show that the constant factor $(1 - 2\log\frac{3}{2})$ is actually tight and the upper bound V_1^{upp} matches the lower bound V_1^{low} defined in Theorem 4.

Theorem 3 (Upper Bound on Minimum Noise Power)**.** *Define*

$$V_2^{upp} := \frac{2\Delta^2}{\epsilon^2}\left(1 - \frac{\frac{1}{2}\log^2(1+\frac{e^{\epsilon}-1}{2\delta}) + \log(1+\frac{e^{\epsilon}-1}{2\delta})}{\frac{e^{\epsilon}-1}{2\delta}}\right). \tag{7}$$

We have

$$V_2^* \leq V_2^{upp}.$$

Proof. We can show that V_2^{upp} is the noise power under the truncated Laplacian mechanism, and thus V_2^{upp} is an upper bound for V_2^*. Please see the supplementary manuscript for the complete proof. □

It turns out that the upper bounds V_1^{upp} and V_2^{upp} in Theorem 2 and Theorem 3 are tight. We derive new lower bounds for V_1^* and V_2^* in the next section, and show that the multiplicative gap between the lower bounds and the upper bounds goes to zero in the high privacy regions in Section 5.

4 Lower Bound

In this section, we derive new lower bounds V_1^{low} and V_2^{low} on the minimum noise amplitude V_1^* and minimum noise power V_2^*, respectively. The key technique is to discretize the continuous probability distribution and the loss function, and transform the continuous functional optimization problem to linear programming, and then apply the discrete result from Geng and Viswanath (2016a).

Geng and Viswanath (2016a) derived lower bounds for an *integer-valued* query function under (ϵ, δ)-differential privacy. For integer-valued query functions, they formulate a linear programming problem with the objective

of minimizing the additive noise. They studied the dual problem and constructed a dual feasible solution which gives a lower bound. Extending this result to the continuous setting, we show a similar lower bound for *real-valued* query function under (ϵ, δ)-differential privacy.

First, we give a lower bound for (ϵ, δ)-differential privacy for *integer-valued* query function due to Geng and Viswanath (2016a).

Define

$$a := \frac{\delta + \frac{e^\epsilon - 1}{2}}{e^\epsilon},$$
$$b := e^{-\epsilon}.$$

To avoid integer rounding issues, assume that there exists an integer n such that $\sum_{k=0}^{n-1} ab^k = \frac{1}{2}$.

Lemma 2 (Theorem 8 in Geng and Viswanath (2016a)). *Consider a symmetric cost function $\mathcal{L}(\cdot) : \mathbb{Z} \to \mathbb{R}$, where $\mathbb{Z}$ denotes the set of all integers. Given the privacy parameters ϵ, δ and the discrete query sensitivity $\tilde{\Delta} \in \mathbb{Z}^+$, if a discrete probability distribution $\mathcal{P}$ satisfies*

$$\mathcal{P}(S) - e^\epsilon \mathcal{P}(S + d) \le \delta, \forall S \subseteq \mathbb{Z}, \forall d \in \mathbb{Z}, |d| \le \tilde{\Delta} \quad (8)$$

and the cost function $\mathcal{L}(\cdot)$ satisfies

$$\sum_{i=1}^{n-1} b^i \big(2\mathcal{L}(i\tilde{\Delta}) - \mathcal{L}(1 + (i-1)\tilde{\Delta}) - \mathcal{L}(1 + i\tilde{\Delta})\big) \ge \mathcal{L}(1), \quad (9)$$

then we have

$$\Sigma_{i \in \mathbb{Z}} \mathcal{L}(i)\mathcal{P}(i) \ge 2 \sum_{k=0}^{n-1} ab^k \mathcal{L}(1 + k\tilde{\Delta}). \quad (10)$$

Theorem 4 (Lower Bound on Minimum Noise Amplitude). *Define*

$$V_1^{low} := 2a \sum_{k=0}^{n-1} b^k k\Delta$$
$$= 2a \left(\frac{b - b^n}{(1-b)^2} - \frac{(n-1)b^n}{1-b} \right) \Delta. \quad (11)$$

We have

$$V_1^* \ge V_1^{low}.$$

Proof. Given $\mathcal{P} \in \mathcal{P}_{\epsilon,\delta}$, we can derive a lower bound on the cost by discretizing the probability distributions and applying the lower bound (10) for integer-valued query functions in Lemma 2.

We first discretize the probability distributions $\mathcal{P}$. Given a positive integer $N \ge 0$, define a discrete probability distribution $\tilde{\mathcal{P}}_N$ via

$$\tilde{\mathcal{P}}_N(i) := \mathcal{P}\big([\frac{\Delta}{2N}(2i-1), \frac{\Delta}{2N}(2i+1))\big), \forall i \in \mathbb{Z}.$$

For the noise cost function $|x|$, define the corresponding discrete cost function $\tilde{\mathcal{L}}_N$ via

$$\tilde{\mathcal{L}}_N(i) \triangleq \begin{cases} 0, & i = 0 \\ \frac{\Delta}{2N}(2i-1), & i \ge 1 \\ \tilde{\mathcal{L}}_N(-i), & i < 0. \end{cases}$$

It is ready to see that

$$\int_{x \in \mathbb{R}} |x| \mathcal{P}(dx) \ge \Sigma_{i \in \mathbb{Z}} \tilde{\mathcal{P}}_N(i) \tilde{\mathcal{L}}_N(i).$$

As the continuous probability distribution $\mathcal{P}$ satisfies (ϵ, δ)-differential privacy constraint (2) with the query sensitivity Δ, the discrete probability distribution $\tilde{\mathcal{P}}_N$ satisfies the discrete (ϵ, δ)-differential privacy constraint (8) with query sensitivity $\tilde{\Delta} = N$, i.e., $\tilde{\mathcal{P}}_N$ satisfies

$$\tilde{\mathcal{P}}_N(S) - e^\epsilon \tilde{\mathcal{P}}_N(S + d) \le \delta, \forall S \subseteq \mathbb{Z}, |d| \le N.$$

We can verify that the condition (9) in Lemma 2 holds for $\tilde{\mathcal{L}}_N$ and $\tilde{\mathcal{P}}_N$ with query sensitivity $\tilde{\Delta} = N$ when N is sufficiently large. Indeed, when $N \ge a + 2$,

$$\sum_{i=1}^{n-1} b^i [2\tilde{\mathcal{L}}_N(iN) - \tilde{\mathcal{L}}_N(1 + (i-1)N) - \tilde{\mathcal{L}}_N(1 + iN)]$$
$$- \tilde{\mathcal{L}}_N(1)$$
$$= \frac{\Delta}{2N}(\frac{N-2}{a} - 1) \ge 0.$$

The corresponding lower bound in (10) for $\tilde{\mathcal{L}}_N$ and $\tilde{\mathcal{P}}_N$ is

$$2 \sum_{k=0}^{n-1} ab^k \tilde{\mathcal{L}}_N(1 + kN) = 2 \sum_{k=0}^{n-1} ab^k \frac{\Delta}{2N}(2kN + 1)$$
$$= 2 \sum_{k=0}^{n-1} ab^k (k\Delta + \frac{\Delta}{2N}) = 2a\Delta \sum_{k=0}^{n-1} b^k k + \frac{\Delta}{2N}$$
$$\ge 2a\Delta \sum_{k=0}^{n-1} b^k k = 2a \left(\frac{b - b^n}{(1-b)^2} - \frac{(n-1)b^n}{1-b} \right) \Delta$$
$$= V_1^{low}$$

Therefore, for any $\mathcal{P} \in \mathcal{P}_{\epsilon,\delta}$, we have

$$\int_{x \in \mathbb{R}} |x| \mathcal{P}(dx) \ge \Sigma_{i \in \mathbb{Z}} \tilde{\mathcal{P}}_N(i) \tilde{\mathcal{L}}_N(i) \ge V_1^{low},$$

and thus $V_1^* \ge V_1^{low}$. □

Similarly, we derive the lower bound for the minimum noise power V_2^*.

Theorem 5 (Lower Bound on Minimum Noise Power)**.** *Define*

$$\begin{aligned} V_2^{low} &:= 2\sum_{k=0}^{n-1} ab^k k^2\Delta^2 \\ &= \frac{2a\Delta^2}{1-b}[-b + 2(\frac{b(1-b^{n-1})}{(1-b)^2} - \frac{(n-1)b^n}{1-b}) \\ &\quad - \frac{b^2(1-b^{n-2})}{1-b} - (n-1)^2 b^n]. \end{aligned} \tag{12}$$

We have

$$V_2^* \geq V_2^{low}.$$

Proof. The proof is similar to the proof of Theorem 4. Please see the supplementary manuscript for the complete proof. □

5 Tightness of the Lower and Upper Bounds

In this section, we compare the lower bounds V_1^{low}, V_2^{low} and the upper bounds V_1^{upp}, V_2^{upp} (derived from the truncated Laplacian mechanism) for the minimum noise amplitude and noise power under (ϵ, δ)-differential privacy. We show that they are close in the high privacy regions and the multiplicative gap goes to zero, which proves the tightness of these lower and upper bounds and thus establishes the near-optimality of the truncated Laplacian mechanism.

Theorem 6 (Tightness of Lower bound and Upper bound on Minimum Noise Amplitude)**.**

$$\begin{aligned} \lim_{\epsilon\to 0} \frac{V_1^{low}}{V_1^{upp}} &\geq 1 - 2\delta. \\ \lim_{\delta\to 0} \frac{V_1^{low}}{V_1^{upp}} &\geq \frac{\epsilon}{e^\epsilon - 1} = 1 - \frac{\epsilon}{2} + O(\epsilon^2). \\ \lim_{\epsilon=\delta\to 0} \frac{V_1^{low}}{V_1^{upp}} &= 1. \end{aligned}$$

Proof. 1. δ is fixed, and $\epsilon \to 0$:

When $\epsilon \to 0$, the upper bound $V_1^{upp} \to \frac{\Delta}{4\delta}$, and the lower bound $V_1^{low} \to 2\delta\frac{n(n-1)}{2}\Delta = 2\delta\frac{\frac{1}{2\delta}(\frac{1}{2\delta}-1)}{2}\Delta = (\frac{1}{4\delta} - \frac{1}{2})\Delta$. Therefore,

$$\lim_{\epsilon\to 0} \frac{V_1^{low}}{V_1^{upp}} \geq \frac{(\frac{1}{4\delta} - \frac{1}{2})\Delta}{\frac{\Delta}{4\delta}} = 1 - 2\delta.$$

Note that $1 - 2\delta \to 1$, as $\delta \to 0$, and thus the multiplicative gap between V_1^{low} and V_1^{upp} converges to zero.

2. ϵ is fixed, and $\delta \to 0$:

When $\delta \to 0$, the upper bound $V_1^{upp} \to \frac{\Delta}{\epsilon}$. For the lower bound V_1^{low}, we have

$$\begin{aligned} a &\to \frac{1-e^{-\epsilon}}{2}, \\ b^n &\to 0, \\ nb^n &\to 0, \end{aligned}$$

and thus $V_1^{low} \to \frac{\Delta}{e^\epsilon - 1}$ as $\delta \to 0$. Therefore,

$$\lim_{\delta\to 0} \frac{V_1^{low}}{V_1^{upp}} \geq \frac{\frac{\Delta}{e^\epsilon - 1}}{\frac{\Delta}{\epsilon}} = \frac{\epsilon}{e^\epsilon - 1} = 1 - \frac{\epsilon}{2} + O(\epsilon^2).$$

Therefore, the multiplicative gap between V_1^{low} and V_1^{upp} converges to zero as $\epsilon \to 0$.

3. $\epsilon = \delta \to 0$:

In this regime, $V_1^{upp} \approx \frac{\Delta}{\epsilon}(1 - 2\log\frac{3}{2})$ as shown in Section 3. For the lower bound V_1^{low}, since $\sum_{k=0}^{n-1} ab^k = \frac{1}{2}$, we have

$$a\frac{1-b^n}{1-b} = \frac{1}{2} \Rightarrow b^n = 1 - \frac{1-b}{2a}.$$

As $\epsilon = \delta \to 0$, $\frac{1-b}{2a} = \frac{1-e^{-\epsilon}}{2\frac{\delta + \frac{e^\epsilon - 1}{e^\epsilon}}{2}} \to \frac{1}{3}$, and thus

$$\begin{aligned} \lim_{\delta\to 0} b^n &= 1 - \frac{1}{3} = \frac{2}{3}, \\ n &= \Theta(\frac{\log\frac{3}{2}}{\delta}). \end{aligned}$$

Note that $a = \Theta(\frac{3}{2}\delta)$ as $\delta \to 0$.

Therefore, as $\epsilon = \delta \to 0$,

$$\begin{aligned} &2a\left(\frac{b-b^n}{(1-b)^2} - \frac{(n-1)b^n}{1-b}\right)\Delta \\ &\approx 2a(\frac{1-\frac{2}{3}}{\delta^2} - \frac{\frac{\log\frac{3}{2}}{\delta}\frac{2}{3}}{\delta})\Delta \\ &= 2a(\frac{1}{3\delta^2} - \frac{\frac{2}{3}\log(\frac{3}{2})}{\delta^2})\Delta \\ &\approx 2\frac{3}{2}\delta(\frac{1}{3\delta^2} - \frac{\frac{2}{3}\log(\frac{3}{2})}{\delta^2})\Delta \\ &= (1 - 2\log\frac{3}{2})\frac{\Delta}{\delta}. \end{aligned}$$

Therefore, V_1^* is lower bounded by $V_1^{low} \approx (1 - 2\log\frac{3}{2})\frac{\Delta}{\delta}$ in the regime $\epsilon = \delta \to 0$. Since it is also upper bounded by $V_1^{upp} \approx \frac{\Delta}{\epsilon}(1 - 2\log\frac{3}{2})$, we conclude that $\lim_{\epsilon=\delta\to 0} \frac{V_1^{low}}{V_1^{upp}} = 1$.

Note that our result closes the constant multiplicative gap in the discrete setting (see Equation (67) and (69) in Geng and Viswanath (2016a)). □

Similarly, we show that the lower bound and the upper bound on the minimum noise power are also tight.

Theorem 7 (Tightness of Lower bound and Upper bound on Minimum Noise Power)**.**

$$\lim_{\epsilon\to 0}\frac{V_2^{low}}{V_2^{upp}} \geq 1-3\delta+2\delta^2.$$
$$\lim_{\delta\to 0}\frac{V_2^{low}}{V_2^{upp}} \geq \frac{\epsilon^2(1+e^\epsilon)}{2(e^\epsilon-1)^2} = 1-\frac{\epsilon}{2}+O(\epsilon^2).$$
$$\lim_{\epsilon=\delta\to 0}\frac{V_2^{low}}{V_2^{upp}} = 1.$$

6 Comparison with the Optimal Gaussian Mechanism

In this section we conduct numeric experiments to compare the performance of the truncated Laplacian mechanisms with the optimal Gaussian mechanism described in Balle and Wang (2018).

A classic result on the Gaussian mechanism is that for any $\epsilon, \delta \in (0,1)$, adding a Gaussian noise with standard deviation $\sigma = \frac{\sqrt{2\log(1.25/\delta)}}{\epsilon}\Delta$ preserves (ϵ,δ)-differential privacy Dwork and Roth (2014). Balle and Wang (2018) developed the optimal Gaussian mechanism whose variance is calibrated directly using the Gaussian cumulative density function instead of a tail bound approximation.

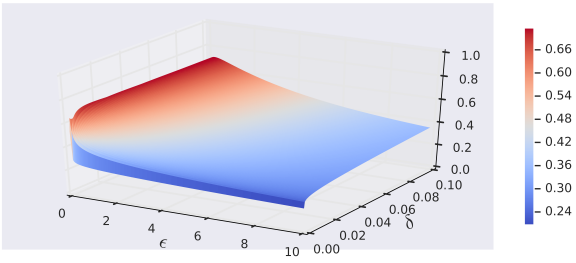

Figure 2: Ratio of the Noise Amplitude of the Truncated Laplacian Mechanism and the Optimal Gaussian Mechanism.

We plot the ratio of the noise amplitude of truncated Laplacian mechanism and the optimal Gaussian mechanism in Fig. 2, and plot the ratio of the noise power of truncated Laplacian mechanism and the optimal Gaussian mechanism in Fig. 3, where $\epsilon \in [10^{-4}, 10]$ and $\delta \in [10^{-6}, 0.1]$. Note that compared with the optimal Gaussian mechanism, the truncated Laplacian mechanism significantly reduces the noise amplitude and noise power in all privacy regimes. The improvement is not very surprising, as the truncated Laplacian mechanism universally improves the probability density decay rate (for both small and big noises) and thus leads to smaller noise amplitude and noise power in expectation.

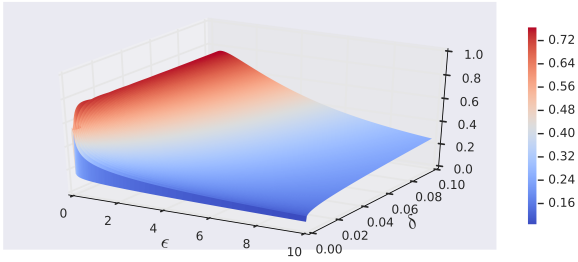

Figure 3: Ratio of the Noise Power of the Truncated Laplacian Mechanism and the Optimal Gaussian Mechanism.

7 Conclusion and Discussion

In this work, we characterize the minimum noise amplitude and noise power for noise-adding mechanisms in (ϵ,δ)-differential privacy for single real-valued query function. We derive new lower bounds using the duality of linear programming, and derive new upper bounds by proposing a new class of (ϵ,δ)-differentially private mechanisms, the *truncated Laplacian* mechanisms. We show that the multiplicative gap of the lower bounds and upper bounds goes to zero in various high privacy regimes, proving the tightness of the lower and upper bounds and thus establishing the optimality of the truncated Laplacian mechanism. In particular, our results close the previous constant multiplicative gap in Geng and Viswanath (2016a). Comprehensive numeric experiments show the improvement of the truncated Laplacian mechanism over the optimal Gaussian mechanism in Balle and Wang (2018) in all privacy regimes.

Note that the range of the truncated Laplacian noise is bounded between $[-A, A]$. Therefore, for two neighboring datasets, the randomized output ranges will have some non-overlapped set. While the truncated Laplacian mechanism can strictly preserve (ϵ,δ)-differential privacy, with a small probability up to δ (corresponding to the probability that the output is in the non-overlapped set), an adversary can distinguish the two neighboring datasets. To address this concern, one can improve over the truncated Laplacian mechanism and impose an arbitrarily light tail distribution over $[A, +\infty)$ to ensure that the output space is the same for all possible datasets.

Acknowledgement

This work was done when Quan Geng was with Google Research. Quan Geng is currently with Facebook, Inc.

References

Martin Abadi, Andy Chu, Ian Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, CCS '16, pages 308–318. ACM, 2016.

Naman Agarwal, Ananda Theertha Suresh, Felix Yu, Sanjiv Kumar, and Brendan McMahan. cpSGD: Communication-efficient and differentially-private distributed SGD. In *Advances in Neural Information Processing Systems*. 2018.

Borja Balle and Yu-Xiang Wang. Improving the Gaussian mechanism for differential privacy: Analytical calibration and optimal denoising. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 394–403. PMLR, 2018.

Kamalika Chaudhuri and Claire Monteleoni. Privacy-preserving logistic regression. In *Neural Information Processing Systems*, pages 289–296, 2008.

Kamalika Chaudhuri, Claire Monteleoni, and Anand D. Sarwate. Differentially private empirical risk minimization. *Journal of Machine Learning Research*, 12:1069–1109, 2011.

Kamalika Chaudhuri, Anand Sarwate, and Kaushik Sinha. Near-optimal differentially private principal components. In *Advances in Neural Information Processing Systems 25*, pages 989–997. 2012.

John Duchi, Michael Jordan, and Martin Wainwright. Privacy aware learning. In *Advances in Neural Information Processing Systems*, pages 1430–1438, 2012.

Cynthia Dwork. Differential Privacy: A Survey of Results. In *Theory and Applications of Models of Computation*, volume 4978, pages 1–19, 2008.

Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–407, 2014.

Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov, and Moni Naor. Our data, ourselves: privacy via distributed noise generation. In *Proceedings of the 24th annual international conference on The Theory and Applications of Cryptographic Techniques*, EUROCRYPT'06, pages 486–503. Springer-Verlag, 2006a.

Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography*, volume 3876 of *Lecture Notes in Computer Science*, pages 265–284. Springer Berlin / Heidelberg, 2006b.

Gintare Karolina Dziugaite and Daniel M Roy. Data-dependent pac-bayes priors via differential privacy. In *Advances in Neural Information Processing Systems 31*, pages 8440–8450. 2018.

Jason Ge, Zhaoran Wang, Mengdi Wang, and Han Liu. Minimax-optimal privacy-preserving sparse pca in distributed systems. In *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, volume 84 of *Proceedings of Machine Learning Research*, pages 1589–1598. PMLR, 09–11 Apr 2018.

Quan Geng and Pramod Viswanath. Optimal noise adding mechanisms for approximate differential privacy. *IEEE Transactions on Information Theory*, 62 (2):952–969, Feb 2016a.

Quan Geng and Pramod Viswanath. The optimal noise-adding mechanism in differential privacy. *IEEE Transactions on Information Theory*, 62(2):925–951, Feb 2016b.

Quan Geng, Peter Kairouz, Sewoong Oh, and Pramod Viswanath. The staircase mechanism in differential privacy. *IEEE Journal of Selected Topics in Signal Processing*, 9(7):1176–1184, Oct 2015.

Quan Geng, Wei Ding, Ruiqi Guo, and Sanjiv Kumar. Optimal Noise-Adding Mechanism in Additive Differential Privacy. In *Proceedings of the 22th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2019.

Arpita Ghosh, Tim Roughgarden, and Mukund Sundararajan. Universally utility-maximizing privacy mechanisms. In *Proceedings of the 41st annual ACM symposium on Theory of computing*, STOC '09, pages 351–360. ACM, 2009.

Mangesh Gupte and Mukund Sundararajan. Universally optimal privacy mechanisms for minimax agents. In *Symposium on Principles of Database Systems*, pages 135–146, 2010.

Prateek Jain, Pravesh Kothari, and Abhradeep Thakurta. Differentially private online learning. In *Proceedings of the 25th Annual Conference on Learning Theory*, volume 23 of *Proceedings of Machine Learning Research*, pages 24.1–24.34. PMLR, 25–27 Jun 2012.

Prateek Jain, Om Dipakbhai Thakkar, and Abhradeep Thakurta. Differentially private matrix completion revisited. In *Proceedings of the 35th International*

Conference on Machine Learning, volume 80 of *Proceedings of Machine Learning Research*, pages 2215–2224. PMLR, 10–15 Jul 2018.

Ilya Mironov. Rényi differential privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pages 263–275, Aug. 2017.

Mijung Park, James Foulds, Kamalika Chaudhuri, and Max Welling. DP-EM: Differentially Private Expectation Maximization. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pages 896–904. PMLR, 2017.

Ngoc-Son Phan, Yue Wang, Xintao Wu, and Dejing Dou. Differential privacy preservation for deep auto-encoders: an application of human behavior prediction. In *AAAI*, 2016.

Or Sheffet. Locally private hypothesis testing. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 4605–4614. PMLR, 10–15 Jul 2018.

Reza Shokri and Vitaly Shmatikov. Privacy-preserving deep learning. In *Proceedings of the 22Nd ACM SIGSAC Conference on Computer and Communications Security*, CCS '15, pages 1310–1321. ACM, 2015.

Jordi Soria-Comas and Josep Domingo-Ferrer. Optimal data-independent noise for differential privacy. *Information Sciences*, 250:200 – 214, 2013.

Di Wang, Marco Gaboardi, and Jinhui Xu. Empirical risk minimization in non-interactive local differential privacy revisited. In *Advances in Neural Information Processing Systems 31*, pages 973–982. 2018.