Differential Privacy at Risk: Bridging Randomness and Privacy Budget

Ashish Dandekar[1], **Debabrota Basu**[2], **Stéphane Bressan**[3]
[1] DI, École Normale Supérieure, Paris, France.
[2] Dept. of Computer Sci. and Engg., Chalmers University of Technology, Göteborg, Sweden.
[3] School of Computing, National University of Singapore, Singapore.
(ashishdandekar,debabrota.basu)@u.nus.edu

Abstract

The calibration of noise for a privacy-preserving mechanism depends on the sensitivity of the query and the prescribed privacy level. A data steward must make the non-trivial choice of a privacy level that balances the requirements of users and the monetary constraints of the business entity.

We study various sources of randomness that are involved in the design of a privacy-preserving mechanism, namely the explicit randomness induced by the noise distribution and the implicit randomness induced by the data-generation distribution. The study leads us to a probabilistic calibration of privacy-preserving mechanisms with quantifiable privacy guarantees. We instantiate it for the Laplace mechanism by providing analytical results.

We propose a cost model that bridges the gap between the privacy level and the compensation budget estimated by a GDPR compliant business entity. We illustrate a realistic scenario wherein the use of fine-tuning of the Laplace mechanism avoids the overestimation of the compensation budget. Additionally, the convexity of the proposed cost model leads to a unique fine-tuning of privacy level that minimises the compensation budget.

1 Introduction

(Dwork et al., 2014) quantify the privacy level ϵ in ϵ-*differential privacy* as an upper bound on the worst-case privacy loss incurred by a privacy-preserving mechanism. Generally, a privacy-preserving mechanism perturbs the results by adding the calibrated amount of random noise to them. The calibration of noise depends on the sensitivity of the query and the specified privacy level. In a real-world setting, a data steward must specify a privacy level that balances the requirements of the users and monetary constraints of the business entity. (Garfinkel et al., 2018) report the issues in deploying differential privacy as the privacy definition by the US census bureau. They highlight the lack of formal methods to choose the privacy level. They also report the empirical studies showing the loss in the utility due to the application of privacy-preserving mechanisms.

We address the dilemma of a data steward in two ways. Firstly, we propose a probabilistic quantification of privacy levels. Probabilistic quantification of privacy levels provide a portfolio to the data steward to take quantifiable risks under the desired utility of the data. We refer to the probabilistic quantification as *privacy at risk*. Secondly, we propose a cost model that links the privacy level to a monetary budget. This cost model helps the data steward to choose the privacy level constrained on the estimated budget and vice versa. In the end, we illustrate a realistic scenario that exemplifies how the data steward can avoid overestimation of the budget by using the proposed cost model by using privacy at risk.

The probabilistic quantification of privacy levels depends on two sources of randomness: the *explicit randomness* induced by the noise distribution and the *implicit randomness* induced by the data-generation distribution. Often, these two sources are coupled with each other. We require analytical forms of both the implicit and explicit sources of randomness as well as the analytical form of the query to derive a privacy guarantee. Computing the probabilistic quantification is generally a challenging task. Although we find multiple probabilistic privacy definitions in the literature (Machanavajjhala et al., 2008; Hall et al., 2012), we do not find the analytical quantification bridging the randomness and privacy level of a privacy-preserving mechanism. To the best of our knowledge, we are the first to analytically derive the probabilistic quantification, namely privacy at risk, for the widely used Laplace mechanism (Dwork et al., 2006b).

The privacy level proposed by the differential privacy framework is too abstract a quantity to be integrated in a business setting. We propose a cost model that maps the privacy level to a monetary budget. The corresponding cost model for the probabilistic quantification of privacy levels is a convex function of the privacy level. Hence, it leads to a unique probabilistic privacy level that minimises the cost. We illustrate a realistic scenario in a GDPR compliant business entity that needs an estimation of compensation budget that it needs to pay back to the stakeholders in an unfortunate event of a personal data breach. The illustration shows that the use of probabilistic privacy levels avoids overestimation of the compensation budget without sacrificing utility.

In conclusion, the benefits of the probabilistic quantification i.e. the privacy at risk are twofold. It not only quantifies

the privacy level for a given privacy-preserving mechanism but also facilitates decision-making in problems that focus on the privacy-utility trade-off and the compensation budget minimisation.

2 Background

We consider a universe of datasets $\mathcal{D}$. We explicitly mention when we consider that the datasets are sampled from a data-generation distribution $\mathcal{G}$ with support $\mathcal{D}$. Two datasets of equal cardinality x and y are said to be *neighbouring datasets* if they differ in one data point. A pair of neighbouring datasets is denoted by $x \sim y$. In this work, we focus on a specific class of queries called *numeric queries*. A numeric query f is a function that maps a dataset into a real-valued vector, i.e. $f : \mathcal{D} \to \mathbb{R}^k$. For instance, a sum query returns the sum of the values in a dataset.

In order to achieve a privacy guarantee, a *privacy-preserving mechanism*, which is a randomised algorithm, explicitly adds noise to the query from a given family of distributions. Thus, a privacy-preserving mechanism of a given family, $\mathcal{M}(f, \Theta)$, for the query f and the set of parameters Θ of the given noise distribution, is a function that maps a dataset into a real vector, i.e. $\mathcal{M}(f, \Theta) : \mathcal{D} \to \mathbb{R}^k$. We denote a privacy-preserving mechanism as $\mathcal{M}$, when the query and the parameters are clear from the context.

Definition 1. [Differential Privacy (Dwork et al., 2014).] *A privacy-preserving mechanism $\mathcal{M}$, equipped with a query f and with parameters Θ, is ϵ-differentially private if for all $Z \subseteq$ Range$(\mathcal{M})$ and $x, y \in \mathcal{D}$ such that $x \sim y$:*

$$\log\left(\left|\frac{\mathbb{P}(\mathcal{M}(f,\Theta)(x) \in Z)}{\mathbb{P}(\mathcal{M}(f,\Theta)(y) \in Z)}\right|\right) \leq \epsilon.$$

A privacy-preserving mechanism provides perfect privacy if it yields indistinguishable outputs for all neighbouring input datasets. The privacy level ϵ quantifies the privacy guarantee provided by ϵ-differential privacy. For a given query, a smaller value of ϵ provides higher privacy. A randomised algorithm that is ϵ-differentially private is also ϵ'-differential private for any $\epsilon' > \epsilon$.

In order to satisfy ϵ-differential privacy, the parameters of a privacy-preserving mechanism requires a calculated calibration. The amount of noise required to achieve a specified privacy level depends on the query. If the output of the query does not change drastically for two neighbouring datasets, then less noise is required to achieve a given privacy level. The measure of such fluctuations is called the *sensitivity* of the query. The parameters of a privacy-preserving mechanism are calibrated using the sensitivity of the query that quantifies the smoothness of a numeric query.

Definition 2. [Sensitivity.] *The sensitivity of a query $f : \mathcal{D} \to \mathbb{R}^k$ is defined as*

$$\Delta_f \triangleq \max_{\substack{x,y\in\mathcal{D}\\ x\sim y}} \|f(x) - f(y)\|_1.$$

The Laplace mechanism is a privacy-preserving mechanism that adds scaled noise sampled from a calibrated Laplace distribution to the numeric query.

Definition 3. [Laplace Mechanism (Dwork et al., 2006b).] *Given any function $f : \mathcal{D} \to \mathbb{R}^k$ and any $x \in \mathcal{D}$, the Laplace Mechanism is defined as*

$$\mathcal{L}_\epsilon^{\Delta_f}(x) \triangleq \mathcal{M}\left(f,\ \frac{\Delta_f}{\epsilon}\right)(x) = f(x) + (L_1, ..., L_k),$$

where L_i is drawn from $\mathrm{Lap}\left(\frac{\Delta_f}{\epsilon}\right)$ *and added to the i^{th} component of $f(x)$.*

Theorem 1. *[(Dwork et al., 2006b)] The Laplace mechanism, $\mathcal{L}_{\epsilon_0}^{\Delta_f}$, is ϵ_0-differentially private.*

3 Privacy at Risk: A Probabilistic Quantification of Randomness

The parameters of a privacy-preserving mechanism are calibrated using the privacy level and the sensitivity of the query. A data steward needs to choose appropriate privacy level for practical implementation. (Lee and Clifton, 2011) show that the choice of an actual privacy level by a data steward in regard to her business requirements is a non-trivial task. Recall that the privacy level in the definition of differential privacy corresponds to the worst case privacy loss. Business users are however used to taking and managing risks, if the risks can be quantified. For instance, (Jorion, 2000) defines *Value at Risk* that is used by risk analysts to quantify the loss in investments for a given portfolio and an acceptable confidence bound. Motivated by the formulation of *Value at Risk*, we propose to use the use of probabilistic privacy level. It provides us a finer tuning of an ϵ_0-differentially private privacy-preserving mechanism for a specified risk γ.

Definition 4. [Privacy at Risk.] *For a given data generating distribution $\mathcal{G}$, a privacy-preserving mechanism $\mathcal{M}$, equipped with a query f and with parameters Θ, satisfies ϵ-differential privacy with a* privacy at risk $0 \leq \gamma \leq 1$*, if for all $Z \subseteq$ Range$(\mathcal{M})$ and x, y sampled from $\mathcal{G}$ such that $x \sim y$:*

$$\mathbb{P}\left[\log\left|\frac{\mathbb{P}(\mathcal{M}(f,\Theta)(x) \in Z)}{\mathbb{P}(\mathcal{M}(f,\Theta)(y) \in Z)}\right| > \epsilon\right] \leq \gamma, \qquad (1)$$

where the outer probability is calculated with respect to the probability space $\mathrm{Range}(\mathcal{M} \circ \mathcal{G})$ *obtained by applying the privacy-preserving mechanism $\mathcal{M}$ on the data-generation distribution $\mathcal{G}$.*

If a privacy-preserving mechanism is ϵ_0-differentially private for a given query f and parameters Θ, the for any privacy level $\epsilon \geq \epsilon_0$, privacy at risk is 0. Our interest is to quantify the risk γ with which ϵ_0-differentially private privacy-preserving mechanism also satisfies a stronger ϵ-differential privacy, i.e. $\epsilon < \epsilon_0$.

Unifying Probabilistic and Random Differential Privacy. Interestingly, Equation 1 unifies the notions of probabilistic differential privacy and random differential privacy by accounting for both sources of randomness in a privacy-preserving mechanism. (Machanavajjhala et al., 2008) define probabilistic differential privacy that incorporates the explicit randomness of the noise distribution of the privacy-preserving mechanism whereas (Hall et al., 2012) define

random differential privacy that incorporates the implicit randomness of the data-generation distribution. In probabilistic differential privacy, the outer probability is computed over the sample space of $\mathrm{Range}(\mathcal{M})$ and all datasets are equally probable.

We do not only coalesce these two aspects but also extend them by providing analytical results connecting privacy level with risk for the Laplace mechanism.

3.1 The Case of Explicit Randomness

In this section, we study the effect of the explicit randomness induced by the noise sampled from Laplacian distribution. We provide a probabilistic quantification for fine tuning for the Laplace mechanism. We fine-tune the privacy level for a specified risk under by assuming that the sensitivity of the query is known a priori.

For a Laplace mechanism $\mathcal{L}_{\epsilon_0}^{\Delta_f}$ calibrated with sensitivity Δ_f and privacy level ϵ_0, we present the analytical formula relating privacy level ϵ and the risk γ_1 in Theorem 2. The proof is available in Appendix A.

Theorem 2. *The risk $\gamma_1 \in [0, 1]$ with which a Laplace Mechanism $\mathcal{L}_{\epsilon_0}^{\Delta_f}$ satisfies a privacy level $\epsilon \geq 0$ is given by*

$$\gamma_1 = \frac{\mathbb{P}(T \leq \epsilon)}{\mathbb{P}(T \leq \epsilon_0)}, \tag{2}$$

where T is a random variable dependent on the Laplace noise $\mathrm{Lap}(\frac{\Delta_f}{\epsilon_0})$, and follows the $\mathrm{BesselK}\left(k, \frac{\Delta_f}{\epsilon_0}\right)$ distribution.

Figure 1a shows the plot of the privacy level against risk for different values of k and for a Laplace mechanism $\mathcal{L}_{1.0}^{1.0}$. As the value of k increases, the amount of noise added in the output of numeric query increases. Therefore, for a specified risk, the privacy at risk level increases with the value of k.

The analytical formula representing γ_1 as a function of ϵ is bijective. We need to invert it to obtain the privacy level ϵ for a privacy at risk γ_1. However the analytical closed form for such an inverse function is not explicit. We use a numerical approach to compute privacy level for a given privacy at risk from the analytical formula of Theorem 2.

Result for a Real-valued Query. For the case $k = 1$, the analytical derivation is fairly straightforward because it only involves *Laplace* and *exponential distributions*, and does not require *gamma* and *BesselK-distribution*. In this case, we obtain an invertible closed-form of a privacy level for a specified risk. It is presented in Equation 3.

$$\epsilon = \ln\left(\frac{1}{1 - \gamma_1(1 - e^{-\epsilon_0})}\right) \tag{3}$$

Remarks on ϵ_0. For $k = 1$, Figure 1b shows the plot of privacy at risk level ϵ versus privacy at risk γ_1 for the Laplace mechanism $\mathcal{L}_{\epsilon_0}^{1.0}$. As the value of ϵ_0 increases, the probability of Laplace mechanism generating higher value of noise reduces. Therefore, we observe that for a fixed privacy at risk, privacy level increases with the value of ϵ_0. The same observation is made for $k > 1$.

3.2 The Case of Implicit Randomness

In this section, we study the effect of the implicit randomness induced by the data-generation distribution to provide a fine tuning for the Laplace mechanism. We fine-tune the risk for a specified privacy level without assuming that the sensitivity of the query.

If one takes into account randomness induced by the data-generation distribution, all pairs of neighbouring datasets are not equally probable. This leads to estimation of sensitivity of a query for a specified data-generation distribution. If we have access to an analytical form of the data-generation distribution and to the query, we could analytically derive the sensitivity distribution for the query. In general, we have access to the datasets, but not the data-generation distribution that generates them. We, therefore, statistically estimate sensitivity by constructing an empirical distribution. We call the sensitivity value obtained for a specified risk from the empirical cumulative distribution of sensitivity the *sampled sensitivity* (Definition 5). However, the value of sampled sensitivity is simply an estimate of the sensitivity for a specified risk. In order to capture this additional uncertainty introduced by the estimation from the empirical sensitivity distribution rather than the true unknown distribution, we compute a lower bound on the accuracy of this estimation. This lower bound yields a probabilistic lower bound on the specified risk. We refer to it as *empirical risk*. For a specified absolute risk γ_2, we denote by $\hat{\gamma_2}$ corresponding empirical risk.

For the Laplace mechanism $\mathcal{L}_{\epsilon}^{\Delta_{S_f}}$ calibrated with sampled sensitivity Δ_{S_f} and privacy level ϵ, we evaluate the empirical risk $\hat{\gamma_2}$. We present the result in Theorem 3. The proof is available in appendices.

Theorem 3. *Analytical bound on the empirical risk, $\hat{\gamma_2}$, for Laplace mechanism $\mathcal{L}_{\epsilon}^{\Delta_{S_f}}$ with privacy level ϵ and sampled sensitivity Δ_{S_f} for a query $f : \mathcal{D} \to \mathbb{R}^k$ is*

$$\hat{\gamma_2} \geq \gamma_2(1 - 2e^{-2\rho^2 n}) \tag{4}$$

where n is the number of samples used for estimation of the sampled sensitivity and ρ is the accuracy parameter. γ_2 denotes the specified absolute risk.

The error parameter ρ controls the closeness between the empirical cumulative distribution of the sensitivity to the true cumulative distribution of the sensitivity. Lower the value of the error, closer is the empirical cumulative distribution to the true cumulative distribution. Figure 2 shows the plot of number of samples as a function of the privacy at risk and the error parameter. We observe that as the value of the error reduces the number of samples in order to achieve the same privacy at risk exponentially increases. Let us now present the sketch of the derivation of the analytical bound on the empirical risk in Theorem 3.

Let, $\mathcal{G}$ denotes the data-generation distribution, either known apriori or constructed by subsampling the available data. We adopt the procedure of (Rubinstein and Aldà, 2017) to sample two neighbouring datasets with p data points each. We sample $p - 1$ data points from $\mathcal{G}$ that are common to both of these datasets and later two more data points. From

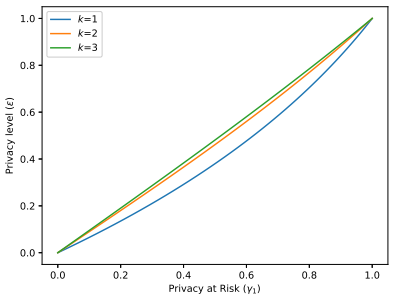

(a)

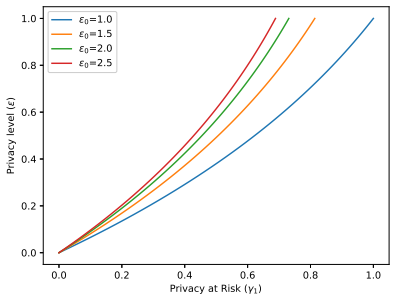

(b)

Figure 1: Privacy level ϵ for varying privacy at risk γ_1 for Laplace mechanism $\mathcal{L}_{\epsilon_0}^{1.0}$. In Figure 1a, we use $\epsilon_0 = 1.0$ and different values of k. In Figure 1b, for $k = 1$ and different values of ϵ_0.

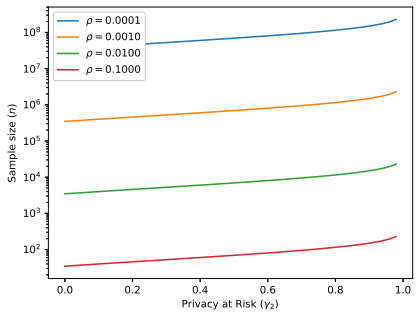

Figure 2: Number of samples n for varying privacy at risk γ_2 for different error parameter ρ.

those two points, we allot one data point to each of the two datasets.

Let, $S_f = \|f(x) - f(y)\|_1$ denotes the sensitivity random variable for a given query f, where x and y are two neighbouring datasets sampled from $\mathcal{G}$. Using n pairs of neighbouring datasets sampled from $\mathcal{G}$, we construct the empirical cumulative distribution, F_n, for the sensitivity random variable.

Definition 5. *For a given query f and for a specified risk γ_2, sampled sensitivity, Δ_{S_f}, is defined as the value of sensitivity random variable that is estimated using its empirical cumulative distribution function, F_n, constructed using n pairs of neighbouring datasets sampled from the data-generation distribution $\mathcal{G}$.*

$$\Delta_{S_f} \triangleq F_n^{-1}(\gamma_2)$$

If we knew analytical form of the data generation distribution, we could analytically derive the cumulative distribution function of the sensitivity, F, and find the sensitivity of the query as $\Delta_f = F^{-1}(1)$. Therefore, in order to have the sampled sensitivity close to the sensitivity of the query, we require the empirical cumulative distributions to be close to the cumulative distribution of the sensitivity. We use this insight to derive the analytical bound in the Theorem 3.

3.3 The Case of Explicit and Implicit Randomness

In this section, we study the combined effect of both explicit randomness induced by the noise distribution and implicit randomness in the data-generation distribution respectively. We do not assume the knowledge of the sensitivity of the query.

We estimate sensitivity using the empirical cumulative distribution of sensitivity. We construct the empirical distribution over the sensitivities using the sampling technique presented in the earlier case. Since we use the sampled sensitivity (Definition 5) to calibrate the Laplace mechanism, we estimate the *empirical risk* $\hat{\gamma_3}$.

For Laplace mechanism $\mathcal{L}_{\epsilon_0}^{\Delta_{S_f}}$ calibrated with sampled sensitivity Δ_{S_f} and privacy level ϵ_0, we present the analytical bound on the empirical sensitivity $\hat{\gamma_3}$ in Theorem 4 with proof in the appendix.

Theorem 4. *Analytical bound on the empirical risk $\hat{\gamma_3} \in [0, 1]$ to achieve a privacy level $\epsilon > 0$ for Laplace mechanism $\mathcal{L}_{\epsilon_0}^{\Delta_{S_f}}$ with sampled sensitivity Δ_{S_f} of a query $f : \mathcal{D} \to \mathbb{R}^k$ is*

$$\hat{\gamma_3} \geq \gamma_3(1 - 2e^{-2\rho^2 n}) \tag{5}$$

where n is the number of samples used for estimating the sensitivity, ρ is the accuracy parameter. γ_3 denotes the specified absolute risk.

The error parameter ρ controls the closeness between the empirical cumulative distribution of the sensitivity to the true cumulative distribution of the sensitivity. Figure 3 shows the dependence of the error parameter on the number of samples. In Figure 3a, we observe that the for a fixed number of samples and a privacy level, the privacy at risk decreases with the value of error parameter. For a fixed number of samples, smaller values of the error parameter reduce the probability of similarity between the empirical cumulative distribution of sensitivity and the true cumulative distribution. Therefore, we observe the reduction in the risk for a fixed privacy level. In Figure 3b, we observe that for a fixed value of error parameter and a fixed level of privacy level, the risk increases with the number of samples. For a fixed value of the error parameter, larger values of the sample size increase the probability of similarity between the empirical cumulative distribution of sensitivity and the true cumulative distribution. Therefore, we observe the increase in the risk for a fixed privacy level.

Effect of the consideration of implicit and explicit randomness is evident in the analytical expression for γ_3 in Equation 6. Proof is available in Appendix C. The privacy at risk is composed of two factors whereas the second term is a privacy at risk that accounts for inherent randomness. The first term takes into account the implicit randomness of the Laplace distribution along with a coupling coefficient η. We define η as the ratio of the true sensitivity of the query to its sampled sensitivity.

$$\gamma_3 \triangleq \frac{\mathbb{P}(T \leq \epsilon)}{\mathbb{P}(T \leq \eta\epsilon_0)} \cdot \gamma_2 \tag{6}$$

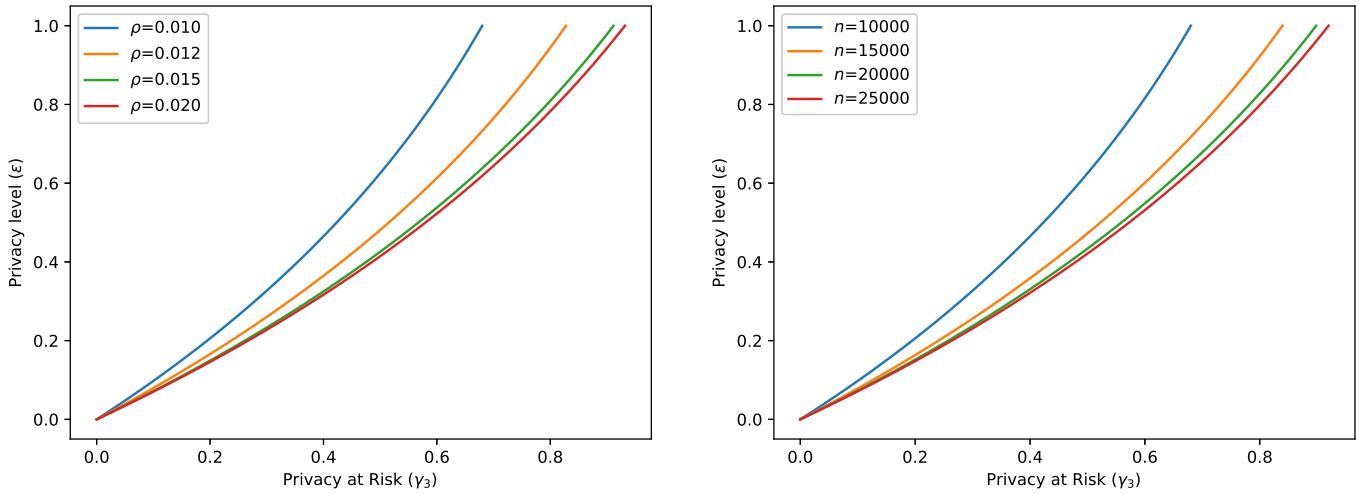

(a) (b)

Figure 3: Dependence of error and number of samples on the privacy at risk for Laplace mechanism $\mathcal{L}_{1.0}^{\Delta_{S_f}}$. For the figure on the left hand side, we fix the number of samples to 10000. For the Figure 3b we fix the error parameter to 0.01.

4 Minimising Compensation Budget for Privacy at Risk

Many service providers collect users' data to enhance user experience. In order to avoid misuse of this data, we require a legal framework that not only limits the use of the collected data but also proposes reparative measures in case of a data leak. General Data Protection Regulation (GDPR)[1] is such a legal framework.

Section 82 in GDPR states that any person who suffers from material or non-material damage as a result of a personal data breach has the right to demand compensation from the data processor. Therefore, every GDPR compliant business entity that either holds or processes personal data needs to secure a certain budget in the worst case scenario of the personal data breach. In order to reduce the risk of such an unfortunate event, the business entity may use privacy-preserving mechanisms that provide provable privacy guarantees while publishing their results. In order to calculate the compensation budget for a business entity, we devise a cost model that maps the privacy guarantees provided by differential privacy and privacy at risk to monetary costs. The discussions demonstrate the usefulness of probabilistic quantification of differential privacy in a business setting.

Cost Model for Differential Privacy. Let E be the compensation budget that a business entity has to pay to every stakeholder in case of a personal data breach when the data is processed without any provable privacy guarantees. Let E_ϵ^{dp} be the compensation budget that a business entity has to pay to every stakeholder in case of a personal data breach when the data is processed with privacy guarantees in terms of ϵ-differential privacy.

Privacy level, ϵ, in ϵ-differential privacy is the quantifier of indistinguishability of the outputs of a privacy-preserving mechanism when two neighbouring datasets are provided as inputs. When the privacy level is zero, the privacy-preserving mechanism outputs all results with equal probability. The indistinguishability reduces with increase in the privacy level. Thus, privacy level of zero bears the lowest risk of personal data breach and the risk increases with the privacy level. E_ϵ^{dp} needs to be commensurate to such a risk and, therefore, it needs to satisfy the following constraints.

1. For all $\epsilon \in \mathbb{R}^{\geq 0}$, $E_\epsilon^{dp} \leq E$.
2. E_ϵ^{dp} is a monotonically increasing function of ϵ.
3. As $\epsilon \to 0$, $E_\epsilon^{dp} \to E_{min}$ where E_{min} is the unavoidable cost that business entity might need to pay in case of personal data breach even after the privacy measures are employed.
4. As $\epsilon \to \infty$, $E_\epsilon^{dp} \to E$.

There are various functions that satisfy these constraints. In absence of any further constraints, we model E_ϵ^{dp} as defined in Equation 7.

$$E_\epsilon^{dp} \triangleq E_{min} + Ee^{-\frac{c}{\epsilon}} \tag{7}$$

E_ϵ^{dp} has two parameters, namely $c > 0$ and $E_{min} \geq 0$. c controls the rate of change in the cost as the privacy level changes and E_{min} is a privacy level independent bias. For this study, we use a simplified model with $c = 1$ and $E_{min} = 0$.

Cost Model for Privacy at Risk. Let, $E_{\epsilon_0}^{par}(\epsilon, \gamma)$ be the compensation that a business entity has to pay to every stakeholder in case of a personal data breach when the data is processed with an ϵ_0-differentially private privacy-preserving mechanism along with a probabilistic quantification of privacy level. Use of such a quantification allows use to provide a stronger a stronger privacy guarantee *viz.* $\epsilon < \epsilon_0$ for a specified privacy at risk at most γ for Thus, we calculate $E_{\epsilon_0}^{par}$

[1] https://eugdpr.org/

using Equation 8.

$$E^{par}_{\epsilon_0}(\epsilon, \gamma) \triangleq \gamma E^{dp}_{\epsilon} + (1-\gamma)E^{dp}_{\epsilon_0} \tag{8}$$

Existence of Minimum Compensation Budget. We want to find the privacy level, say ϵ_{min}, that yields the lowest compensation budget. We do that by minimising Equation 8 with respect to ϵ.

Lemma 1. *$E^{par}_{\epsilon_0}(\epsilon, \gamma)$ is a convex function of ϵ.*

By Lemma 1, there exists a unique ϵ_{min} that minimises the compensation budget for a specified parametrisation, say ϵ_0. Since the risk γ in Equation 8 is itself a function of privacy level ϵ, analytical calculation of ϵ_{min} is not possible in the most general case. When the output of the query is a real number, we derive the analytic form (Equation 3) to compute the risk under the consideration of explicit randomness. In such a case, ϵ_{min} is calculated by differentiating Equation 8 with respect to ϵ and equating it to zero. It gives us Equation 9 that we solve using any root finding technique such as Newton-Raphson method[(Press, 2007)] to compute ϵ_{min}.

$$\frac{1}{\epsilon} - \ln\left(1 - \frac{1-e^{\epsilon}}{\epsilon^2}\right) = \frac{1}{\epsilon_0} \tag{9}$$

Illustration. Suppose that the health centre in a university that complies to GDPR publishes statistics of its staff health checkup, such as obesity statistics, twice in a year. In January 2018, the health centre publishes that 34 out of 99 faculty members suffer from obesity. In July 2018, the health centre publishes that 35 out of 100 faculty members suffer from obesity. An intruder, perhaps an analyst working for an insurance company, checks the staff listings in January 2018 and July 2018, which are publicly available on website of the university. The intruder does not find any change other than the recruitment of John Doe in April 2018. Thus, with high probability, the intruder deduces that John Doe suffers from obesity. In order to avoid such a privacy breach, the health centre decides to publish the results using the Laplace mechanism. In this case, the Laplace mechanism operates on the count query.

In order to control the amount of noise, the health centre needs to appropriately set the privacy level. Suppose that the health centre decides to use the expected mean absolute error, defined in Equation 10, as the measure of *effectiveness* for the Laplace mechanism.

$$\mathbb{E}\left[|\mathcal{L}^1_{\epsilon}(x) - f(x)|\right] = \frac{1}{\epsilon} \tag{10}$$

Equation 10 makes use of the fact that the sensitivity of the count query is one. Suppose that the health centre requires the expected mean absolute error of at most two in order to maintain the quality of the published statistics. In this case, the privacy level has to be at least 0.5.

In order to compute the budget, the health centre requires an estimate of E. (Moriarty et al., 2012) show that the incremental cost of premiums for the health insurance with morbid obesity ranges between \$5467 to \$5530. With reference to this research, the health centre takes \$5500 as an estimate of E. For the staff size of 100 and the privacy level 0.5, the health centre uses Equation 7 in its simplified setting to compute the total budget of $\$74434.40$.

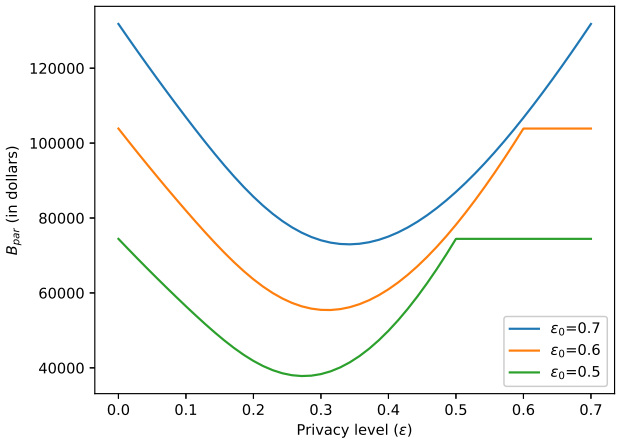

Figure 4: Variation in the budget for Laplace mechanism $\mathcal{L}^1_{\epsilon_0}$ under privacy at risk considering explicit randomness in the Laplace mechanism for the illustration in Section 4.

Is it possible to reduce this budget without degrading the effectiveness of the Laplace mechanism? We show that it is possible by fine-tuning the Laplace mechanism. Under the consideration of the explicit randomness introduced by the Laplace noise distribution, we show that ϵ_0-differentially private Laplace mechanism also satisfies ϵ-differential privacy with risk γ, which is computed using the formula in Theorem 2. Fine-tuning allows us to get a stronger privacy guarantee, $\epsilon < \epsilon_0$ that requires a smaller budget. In Figure 4, we plot the budget for various privacy levels. We observe that the privacy level 0.274, which is same as ϵ_{min} computed by solving Equation 9, yields the lowest compensation budget of $\$37805.86$. Thus, by using privacy at risk, the health centre is able to save $\$36628.532$ without sacrificing the quality of the published results.

Bounds on the Privacy at Risk. For a fixed budget, say B, re-arrangement of Equation 8 gives us an upper bound on the privacy level ϵ. We use the cost model with $c = 1$ and $E_{min} = 0$ to derive the upper bound. If we have a maximum permissible expected mean absolute error T, we use Equation 10 to obtain a lower bound on the privacy at risk level. Equation 11 illustrates the upper and lower bounds that dictate the permissible range of ϵ that a data publisher can promise depending on the budget and the permissible error constraints.

$$\frac{1}{T} \leq \epsilon \leq \left[\ln\left(\frac{\gamma E}{B - (1-\gamma)E^{dp}_{\epsilon_0}}\right)\right]^{-1} \tag{11}$$

Thus, the privacy level is constrained by the effectiveness requirement from below and by the monetary budget from above. (Hsu et al., 2014) calculate upper and lower bound on the privacy level in the differential privacy. They use a different cost model owing to the scenario of research study that compensates its participants for their data and releases the results in a differentially private manner. Their cost model is different than our GDPR inspired modelling.

5 Related Work

Researchers have proposed different privacy-preserving mechanisms to make different queries differentially private. These mechanisms can be broadly classified into two categories. In one category, the mechanisms explicitly add calibrated noise, such as Laplace noise in the work of (Dwork et al., 2006c) or Gaussian noise in the work of (Dwork et al., 2014), to the outputs of the query. In the other category, (Chaudhuri et al., 2011; Zhang et al., 2012; Acs et al., 2012; Hall et al., 2013) propose mechanisms that alter the query function so that the modified function satisfies differentially privacy. Privacy-preserving mechanisms in both of these categories perturb the original output of the query and make it difficult for a malicious data analyst to recover the original output of the query. These mechanisms induce randomness using the explicit noise distribution. Calibration of these mechanisms require the knowledge of the sensitivity of the query. (Nissim et al., 2007) consider the implicit randomness in the data-generation distribution to compute an estimate of the sensitivity. The authors propose the smooth sensitivity function that is an envelope over the local sensitivities for all individual datasets. Local sensitivity of a dataset is the maximum change in the value of the query over all of its neighboring datasets. In general, it is not easy to analytically estimate the smooth sensitivity function for a general query. (Rubinstein and Aldà, 2017) also study the inherent randomness in the data-generation algorithm. They do not use the local sensitivity. We adopt their approach of sampling the sensitivity from the empirical distribution of the sensitivity. They use order statistics to choose a particular value of the sensitivity. We use the risk, which provides a mediation tool for business entities to assess the actual business risks, on the sensitivity distribution to estimate the sensitivity.

In order to account for both sources of randomness, refinements of ϵ-differential privacy are proposed in order to bound the probability of occurrence of worst case scenarios. (Machanavajjhala et al., 2008) propose probabilistic differential privacy that considers upper bounds of the worst case privacy loss for corresponding confidence levels on the noise distribution. Definition of probabilistic differential privacy incorporates the explicit randomness induced by the noise distribution and bounds the probability over the space of noisy outputs to satisfy the ϵ-differential privacy definition. (Dwork and Rothblum, 2016) propose Concentrated differential privacy that considers the expected values of the privacy loss random variables for the corresponding. Definition of concentrated differential privacy incorporates the explicit randomness induced by the noise distribution but considering only the expected value of privacy loss satisfying ϵ-differential privacy definition instead of using the confidence levels limits its scope.

(Hall et al., 2013) propose random differential privacy that considers the privacy loss for corresponding confidence levels on the implicit randomness in the data-generation distribution. Definition of random differential privacy incorporates the implicit randomness induced by the data-generation distribution and bounds the probability over the space of datasets generated from the given distribution to satisfy the ϵ-differential privacy definition. (Dwork et al., 2006a) define approximate differential privacy by adding a constant bias to the privacy guarantee provided by the differential privacy. It is not a probabilistic refinement of the differential privacy.

(Kifer and Machanavajjhala, 2012) define Pufferfish privacy framework, and its variant by (Bassily et al., 2013), that considers randomness due to data-generation distribution as well as noise distribution. Despite the generality of their approach, the framework relies on the domain expert to define a set of *secrets* that they want to protect.

In this work, we consider the widely used Laplace mechanism proposed by (Dwork et al., 2006c). The Laplace mechanism adds Laplacian noise to the query output. (Xiao et al., 2011; Zhang et al., 2012; Acs et al., 2012) use Laplace mechanism by providing the calibration by computing sensitivity of the query.

(Ghosh and Roth, 2015; Chen et al., 2016) propose game theoretic methods that provide the means to evaluate the monetary cost of differential privacy. Our approach is inspired by the approach by the work of (Hsu et al., 2014). They model the cost under a scenario of a research study wherein the participants are reimbursed for their participation. Our cost modelling is driven by the scenario of securing a compensation budget in compliance with GDPR. Our requirement differs from the requirements for the scenario in the work of (Hsu et al., 2014). In our case, there is no monetary incentive for participants to share their data.

6 Conclusion and Future Works

In this paper, we provide a means to fine-tune the privacy level of a privacy-preserving mechanism by conducting a study of various sources of randomness. Such a fine-tuning leads to probabilistic quantification on privacy levels with quantified risks. We instantiate the analytical derivations for the Laplace mechanism. We propose a cost model that bridges the gap between the privacy level and the compensation budget estimated by a GDPR compliant business entity. We show the existence of a privacy level that yields the minimum compensation budget under the proposed fine-tuning.

Such a fine-tuning may be fully analytically computed in cases where the data-generation, or the sensitivity distribution, the noise distribution and the query are analytically known and take convenient forms. We are now looking at such convenient but realistic cases.

References

Acs, G., Castelluccia, C., and Chen, R. (2012). Differentially private histogram publishing through lossy compression. In *Data Mining (ICDM), 2012 IEEE 12th International Conference on*, pages 1–10. IEEE.

Askey, R. and Daalhuis, A. O. (2010). Generalized hypergeometric functions and meijer g-function. *NIST handbook of mathematical functions*, pages 403–418.

Bassily, R., Groce, A., Katz, J., and Smith, A. (2013). Coupled-worlds privacy: Exploiting adversarial uncertainty in statistical data privacy. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 439–448. IEEE.

Chaudhuri, K., Monteleoni, C., and Sarwate, A. D. (2011). Differentially private empirical risk minimization. *Journal of Machine Learning Research*, 12(Mar):1069–1109.

Chen, Y., Chong, S., Kash, I. A., Moran, T., and Vadhan, S. (2016). Truthful mechanisms for agents that value privacy. *ACM Transactions on Economics and Computation (TEAC)*, 4(3):13.

Dwork, C., Kenthapadi, K., McSherry, F., Mironov, I., and Naor, M. (2006a). Our data, ourselves: Privacy via distributed noise generation. In *Eurocrypt*, volume 4004, pages 486–503. Springer.

Dwork, C., McSherry, F., Nissim, K., and Smith, A. (2006b). Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference*, pages 265–284. Springer.

Dwork, C., McSherry, F., Nissim, K., and Smith, A. (2006c). *Calibrating Noise to Sensitivity in Private Data Analysis*, pages 265–284. Springer Berlin Heidelberg, Berlin, Heidelberg.

Dwork, C., Roth, A., et al. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407.

Dwork, C. and Rothblum, G. N. (2016). Concentrated differential privacy. *arXiv preprint arXiv:1603.01887*.

Garfinkel, S. L., Abowd, J. M., and Powazek, S. (2018). Issues encountered deploying differential privacy. *arXiv preprint arXiv:1809.02201*.

Ghosh, A. and Roth, A. (2015). Selling privacy at auction. *Games and Economic Behavior*, 91:334–346.

Hall, R., Rinaldo, A., and Wasserman, L. (2012). Random differential privacy. *Journal of Privacy and Confidentiality*, 4(2):43–59.

Hall, R., Rinaldo, A., and Wasserman, L. (2013). Differential privacy for functions and functional data. *Journal of Machine Learning Research*, 14(Feb):703–727.

Hsu, J., Gaboardi, M., Haeberlen, A., Khanna, S., Narayan, A., Pierce, B. C., and Roth, A. (2014). Differential privacy: An economic method for choosing epsilon. In *Computer Security Foundations Symposium (CSF), 2014 IEEE 27th*, pages 398–410. IEEE.

Jorion, P. (2000). Value at risk: The new benchmark for managing financial risk.

Kifer, D. and Machanavajjhala, A. (2012). A rigorous and customizable framework for privacy. In *Proceedings of the 31st ACM SIGMOD-SIGACT-SIGAI symposium on Principles of Database Systems*, pages 77–88. ACM.

Lee, J. and Clifton, C. (2011). How much is enough? choosing ε for differential privacy. In *International Conference on Information Security*, pages 325–340. Springer.

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

Massart, P. et al. (1990). The tight constant in the dvoretzky-kiefer-wolfowitz inequality. *The annals of Probability*, 18(3):1269–1283.

Moriarty, J. P., Branda, M. E., Olsen, K. D., Shah, N. D., Borah, B. J., Wagie, A. E., Egginton, J. S., and Naessens, J. M. (2012). The effects of incremental costs of smoking and obesity on health care costs among adults: a 7-year longitudinal study. *Journal of Occupational and Environmental Medicine*, 54(3):286–291.

Nissim, K., Raskhodnikova, S., and Smith, A. (2007). Smooth sensitivity and sampling in private data analysis. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 75–84. ACM.

Papoulis, A. and Pillai, S. U. (2002). *Probability, random variables, and stochastic processes*. Tata McGraw-Hill Education.

Press, W. H. (2007). *Numerical recipes 3rd edition: The art of scientific computing*. Cambridge university press.

Rubinstein, B. I. and Aldà, F. (2017). Pain-free random differential privacy with sensitivity sampling. In *International Conference on Machine Learning*, pages 2950–2959.

Xiao, X., Wang, G., and Gehrke, J. (2011). Differential privacy via wavelet transforms. *IEEE Transactions on Knowledge and Data Engineering*, 23(8):1200–1214.

Zhang, J., Zhang, Z., Xiao, X., Yang, Y., and Winslett, M. (2012). Functional mechanism: regression analysis under differential privacy. *Proceedings of the VLDB Endowment*, 5(11):1364–1375.

A Proof of Theorem 2 (Section 3.1)

Although a Laplace mechanism $\mathcal{L}_{\epsilon}^{\Delta_f}$ induces higher amount of noise on average than a Laplace mechanism $\mathcal{L}_{\epsilon_0}^{\Delta_f}$ for $\epsilon < \epsilon_0$, there is a non-zero probability that $\mathcal{L}_{\epsilon}^{\Delta_f}$ induces noise commensurate to $\mathcal{L}_{\epsilon_0}^{\Delta_f}$. This non-zero probability guides us to calculate the privacy at risk γ_1 for the privacy at risk level ϵ. In order to get an intuition, we illustrate the calculation of the overlap between two Laplace distributions as an estimator of similarity between the two distributions.

Definition 6. *[Overlap of Distributions, (Papoulis and Pillai, 2002)] The overlap, O, between two probability distributions P_1, P_2 with support $\mathcal{X}$ is defined as*

$$O = \int_{\mathcal{X}} \min[P_1(x), P_2(x)]\, dx.$$

Lemma 2. *The overlap O between two probability distributions,* $\mathrm{Lap}(\frac{\Delta_f}{\epsilon_1})$ *and* $\mathrm{Lap}(\frac{\Delta_f}{\epsilon_2})$*, such that $\epsilon_2 \leq \epsilon_1$, is given by*

$$O = 1 - (\exp(-\mu\epsilon_2/\Delta_f) - \exp(-\mu\epsilon_1/\Delta_f)),$$

where $\mu = \frac{\Delta_f \ln(\epsilon_1/\epsilon_2)}{\epsilon_1 - \epsilon_2}$.

Using the result in Lemma 2, we note that the overlap between two distributions with $\epsilon_0 = 1$ and $\epsilon = 0.6$ is 0.81. Thus, $\mathcal{L}_{0.6}^{\Delta_f}$ induces noise that is more than 80% times similar to the noise induced by $\mathcal{L}_{1.0}^{\Delta_f}$. Therefore, we can loosely say that at least 80% of the times a Laplace Mechanism $\mathcal{L}_{1.0}^{\Delta_f}$ will provide the same privacy as a Laplace Mechanism $\mathcal{L}_{0.8}^{\Delta_f}$.

Although the overlap between Laplace distributions with different scales offers an insight into the relationship between different privacy levels, it does not capture the constraint induced by the *sensitivity*. For a given query f, the amount of noise required to satisfy differential privacy is commensurate to the sensitivity of the query. This calibration puts a constraint on the noise that is required to be induced on a pair of neighbouring datasets. We state this constraint in Lemma 3, which we further use to prove that the Laplace Mechanism $\mathcal{L}_{\epsilon_0}^{\Delta_f}$ satisfies (ϵ, γ_1)-privacy at risk.

Lemma 3. *For a Laplace Mechanism $\mathcal{L}_{\epsilon_0}^{\Delta_f}$, the difference in the absolute values of noise induced on a pair of neighbouring datasets is upper bounded by the sensitivity of the query.*

Proof. Suppose that two neighbouring datasets x and y are given input to a numeric query $f : \mathcal{D} \to \mathbb{R}^k$. For any output $z \in \mathbb{R}^k$ of the Laplace Mechanism $\mathcal{L}_{\epsilon_0}^{\Delta_f}$,

$$\sum_{i=1}^{k} (|f(y_i) - z_i| - |f(x_i) - z_i|) \leq \sum_{i=1}^{k} (|f(x_i) - f(y_i)|) \leq \Delta_f.$$

We use triangular inequality in the first step and Definition 2 of sensitivity in the second step. □

We write $\mathrm{Exp}(b)$ to denote a random variable sampled from an *exponential distribution* with scale $b > 0$. We write $\mathrm{Gamma}(k, \theta)$ to denote a random variable sampled from a *gamma distribution* with shape $k > 0$ and scale $\theta > 0$.

Lemma 4. *[(Papoulis and Pillai, 2002)] If a random variable X follows Laplace Distribution with mean zero and scale b,* $|X| \sim \mathrm{Exp}(b)$.

Lemma 5. *[(Papoulis and Pillai, 2002)] If $X_1, ..., X_n$ are n i.i.d. random variables each following the Exponential Distribution with scale b,* $\sum_{i=1}^{n} X_i \sim$ *Gamma*(n, b).

Lemma 6. *If X_1 and X_2 are two i.i.d. Gamma(n, θ) random variables, the probability density function for the random variable $T = |X_1 - X_2|/\theta$ is given by*

$$P_T(t) = \frac{2^{2-n} t^{n-\frac{1}{2}} K_{n-\frac{1}{2}}(t)}{\sqrt{2\pi}\Gamma(n)}$$

where $K_{n-\frac{1}{2}}$ is the modified Bessel function of second kind.

Proof. Let X_1 and X_2 be two i.i.d. $\mathrm{Gamma}(n, \theta)$ random variables. Characteristic function of a Gamma random variable is given as

$$\phi_{X_1}(z) = \phi_{X_2}(z) = (1 - \iota z\theta)^{-n}.$$

Therefore,

$$\phi_{X_1 - X_2}(z) = \phi_{X_1}(z)\phi^*_{X_2}(z) = \frac{1}{(1 + (z\theta)^2)^n}$$

Probability density function for the random variable $X_1 - X_2$ is given by,

$$P_{X_1 - X_2}(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-\mathrm{i}zx} \phi_{X_1 - X_2}(z) dz = \frac{2^{1-n} |\frac{x}{\theta}|^{n-\frac{1}{2}} K_{n-\frac{1}{2}}(|\frac{x}{\theta}|)}{\sqrt{2\pi}\Gamma(n)\theta}$$

where $K_{n-\frac{1}{2}}$ is the Bessel function of second kind. Let, $T = |\frac{X_1 - X_2}{\theta}|$. Therefore,

$$P_T(t) = \frac{2^{2-n} t^{n-\frac{1}{2}} K_{n-\frac{1}{2}}(t)}{\sqrt{2\pi}\Gamma(n)}$$

We denote this probability distribution as $\mathrm{BesselK}(n, \theta)$. □

Lemma 7. *If X_1 and X_2 are two i.i.d. Gamma(n, θ) random variables and $|X_1 - X_2| \leq M$, then $T' = |X_1 - X_2|/\theta$ follows Truncated BesselK(n, θ, M) distribution with probability density function:*

$$P_{T'}(t') = \frac{P_T(t')}{P_T(T \leq M)},$$

where P_T is the probability density function of $\mathrm{BesselK}(n, \theta)$.

Lemma 8. *For Laplace Mechanism $\mathcal{L}_{\epsilon_0}^{\Delta_f}$ with query $f : \mathcal{D} \to \mathbb{R}^k$ and for any output $Z \subseteq Range(\mathcal{L}_{\epsilon_0}^{\Delta_f})$, $\epsilon \leq \epsilon_0$,*

$$\gamma_1 \triangleq \mathbb{P}\left[\log\left|\frac{\mathbb{P}(\mathcal{L}_{\epsilon_0}^{\Delta_f}(x) \in Z)}{\mathbb{P}(\mathcal{L}_{\epsilon_0}^{\Delta_f}(y) \in Z)}\right| \leq \epsilon\right] = \frac{\mathbb{P}(T \leq \epsilon)}{\mathbb{P}(T \leq \epsilon_0)},$$

where T follows $\text{BesselK}(k, \Delta_f/\epsilon_0)$.

Proof. Let, $x \in \mathcal{D}$ and $y \in \mathcal{D}$ be two datasets such that $x \sim y$. Let $f : \mathcal{D} \to \mathbb{R}^k$ be some numeric query. Let $\mathbb{P}_x(z)$ and $\mathbb{P}_y(z)$ denote the probabilities of getting the output z for Laplace mechanisms $\mathcal{L}_{\epsilon_0}^{\Delta_f}(x)$ and $\mathcal{L}_{\epsilon_0}^{\Delta_f}(y)$ respectively. For any point $z \in \mathbb{R}^k$ and $\epsilon \neq 0$,

$$\begin{aligned}\frac{\mathbb{P}_x(z)}{\mathbb{P}_y(z)} &= \prod_{i=1}^{k} \frac{\exp\left(\frac{-\epsilon_0|f(x_i)-z_i|}{\Delta_f}\right)}{\exp\left(\frac{-\epsilon_0|f(y_i)-z_i|}{\Delta_f}\right)} \\ &= \prod_{i=1}^{k} \exp\left(\frac{\epsilon_0(|f(y_i)-z_i| - |f(x_i)-z_i|)}{\Delta_f}\right) \\ &= \exp\left(\epsilon\left[\frac{\epsilon_0 \sum_{i=1}^{k}(|f(y_i)-z_i| - |f(x_i)-z_i|)}{\epsilon\Delta_f}\right]\right).\end{aligned} \tag{12}$$

By Definition 3,

$$(f(x)-z), (f(y)-z) \sim \text{Lap}(\Delta_f/\epsilon_0). \tag{13}$$

Application of Lemma 4 and Lemma 5 yields,

$$\sum_{i=1}^{k}(|f(x_i)-z_i|) \sim \text{Gamma}(k, \Delta_f/\epsilon_0). \tag{14}$$

Using Equations 13, 14, and Lemma 3, 7, we get

$$\left(\frac{\epsilon_0}{\Delta_f}\sum_{i=1}^{k}|(|f(y_i)-z| - |f(x_i)-z|)|\right) \sim \text{TruncatedBesselK}(k, \Delta_f/\epsilon_0, \Delta_f). \tag{15}$$

since, $\sum_{i=1}^{k}|(|f(y_i)-z| - |f(x_i)-z|)| \leq \Delta_f$. Therefore,

$$\mathbb{P}\left(\left[\frac{\epsilon_0}{\Delta_f}\sum_{i=1}^{k}|(|f(y_i)-z| - |f(x_i)-z|)|\right] \leq \epsilon\right) = \frac{\mathbb{P}(T \leq \epsilon)}{\mathbb{P}(T \leq \epsilon_0)}, \tag{16}$$

where T follows $\text{BesselK}(k, \Delta_f/\epsilon_0)$. Analytically,

$$\mathbb{P}(T \leq x) \propto \left({}_1F_2(\frac{1}{2}; \frac{3}{2}-k, \frac{3}{2}; \frac{x^2}{4})\sqrt{\pi}4^k x\right) - \left(2{}_1F_2(k; \frac{1}{2}+k, k+1; \frac{x^2}{4})x^{2k}\Gamma(k)\right)$$

where ${}_1F_2$ is the regularised generalised hypergeometric function as defined in (Askey and Daalhuis, 2010). From Equation 12 and 16,

$$\mathbb{P}\left[\log\left|\frac{\mathbb{P}(\mathcal{L}_{\epsilon_0}(x) \in S)}{\mathbb{P}(\mathcal{L}_{\epsilon_0}(y) \in S)}\right| \leq \epsilon\right] = \frac{\mathbb{P}(T \leq \epsilon)}{\mathbb{P}(T \leq \epsilon_0)}.$$

□

This completes the proof of Theorem 2.

Corollary 1. *Laplace Mechanism $\mathcal{L}_{\epsilon_0}^{\Delta_f}$ with $f : \mathcal{D} \to \mathbb{R}^k$ is (ϵ, δ)-probabilistically differentially private where*

$$\delta = \begin{cases} 1 - \frac{\mathbb{P}(T \leq \epsilon)}{\mathbb{P}(T \leq \epsilon_0)} & \epsilon \leq \epsilon_0 \\ 0 & \epsilon > \epsilon_0 \end{cases}$$

and T follows $\text{BesselK}(k, \Delta_f/\epsilon_0)$.

B Proof of Theorem 3 (Section 3.2)

Proof. Let, x and y be any two neighbouring datasets sampled from the data generating distribution $\mathcal{G}$. Let, Δ_{S_f} be the sampled sensitivity for query $f : \mathcal{D} \to \mathbb{R}^k$. Let, $\mathbb{P}_x(z)$ and $\mathbb{P}_y(z)$ denote the probabilities of getting the output z for Laplace mechanisms $\mathcal{L}_{\epsilon}^{\Delta_{S_f}}(x)$ and $\mathcal{L}_{\epsilon}^{\Delta_{S_f}}(y)$ respectively. For any point $z \in \mathbb{R}^k$ and $\epsilon \neq 0$,

$$\begin{aligned}\frac{\mathbb{P}_x(z)}{\mathbb{P}_y(z)} &= \prod_{i=1}^{k} \frac{\exp\left(\frac{-\epsilon|f(x_i)-z_i|}{\Delta_{S_f}}\right)}{\exp\left(\frac{-\epsilon|f(y_i)-z_i|}{\Delta_{S_f}}\right)} \\ &= \exp\left(\frac{\epsilon\sum_{i=1}^{k}(|f(y_i)-z_i| - |f(x_i)-z_i|)}{\Delta_{S_f}}\right) \\ &\leq \exp\left(\frac{\epsilon\sum_{i=1}^{k}|f(y_i)-f(x_i)|}{\Delta_{S_f}}\right) \\ &= \exp\left(\frac{\epsilon\|f(y)-f(x)\|_1}{\Delta_{S_f}}\right)\end{aligned} \tag{17}$$

We used triangle inequality in the penultimate step.

Using the trick in the work of (Rubinstein and Aldà, 2017), we define following events. Let, $B^{\Delta_{S_f}}$ denotes the set of pairs neighbouring dataset sampled from $\mathcal{G}$ for which the sensitivity random variable is upper bounded by Δ_{S_f}. Let, $C_\rho^{\Delta_{S_f}}$ denotes the set of sensitivity random variable values for which F_n deviates from the unknown cumulative distribution of S, F, at most by the accuracy value ρ. These events are defined in Equation 18.

$$\begin{aligned} B^{\Delta_{S_f}} &\triangleq \{x, y \sim \mathcal{G} \text{ such that } \|f(y)-f(x)\|_1 \leq \Delta_{S_f}\} \\ C_\rho^{\Delta_{S_f}} &\triangleq \left\{\sup_{\Delta}|F_S^n(\Delta) - F_S(\Delta)| \leq \rho\right\}\end{aligned} \tag{18}$$

$$\begin{aligned}\mathbb{P}(B^{\Delta_{S_f}}) &= \mathbb{P}(B^{\Delta_{S_f}}|C_\rho^{\Delta_{S_f}})\mathbb{P}(C_\rho^{\Delta_{S_f}}) \qquad (19)\\ &\quad + \mathbb{P}(B^{\Delta_{S_f}}|\overline{C\rho^{\Delta_{S_f}}})\mathbb{P}(\overline{C_\rho^{\Delta_{S_f}}}) \\ &\geq \mathbb{P}(B^{\Delta_{S_f}}|C_\rho^{\Delta_{S_f}})\mathbb{P}(C_\rho^{\Delta_{S_f}}) \\ &= F_n(\Delta_{S_f})\mathbb{P}(C_\rho^{\Delta_{S_f}}) \\ &\geq \gamma_2 \cdot (1 - 2e^{-2\rho^2 n}) \qquad (20)\end{aligned}$$

In the last step, we use the definition of the sampled sensitivity to get the value of the first term. The last term is obtained using DKW-inequality, as defined in (Massart et al., 1990),

where the n denotes the number of samples used to build empirical distribution of the sensitivity, F_n.

From Equation 17, we understand that if $\|f(y) - f(x)\|_1$ is less than or equals to the sampled sensitivity then the Laplace mechanism $\mathcal{L}_\epsilon^{\Delta_{S_f}}$ satisfies ϵ-differential privacy. Equation 20 provides the lower bound on the probability of the event $\|f(y) - f(x)\|_1 \leq \Delta_{S_f}$. Thus, combining Equation 17 and Equation 20 completes the proof. □

C Proof of Theorem 4 (Section 3.3)

Proof of Theorem 4 builds upon the ideas from the proofs for the rest of the two cases. In addition to the events defined in Equation 18, we define an additional event $A_{\epsilon_0}^{\Delta_{S_f}}$, defined in Equation 21, as a set of outputs of Laplace mechanism $\mathcal{L}_{\epsilon_0}^{\Delta_{S_f}}$ that satisfy the constraint of ϵ-differential privacy for a specified privacy at risk level ϵ.

$$A_{\epsilon_0}^{\Delta_{S_f}} \triangleq \left\{ z \sim \mathcal{L}_{\epsilon_0}^{\Delta_{S_f}} \ : \ \log \left| \frac{\mathcal{L}_{\epsilon_0}^{\Delta_{S_f}}(x)}{\mathcal{L}_{\epsilon_0}^{\Delta_{S_f}}(y)} \right| \leq \epsilon, x, y \sim \mathcal{G} \right\} \tag{21}$$

Corollary 2.

$$\mathbb{P}(A_{\epsilon_0}^{\Delta_{S_f}} | B^{\Delta_{S_f}}) = \frac{\mathbb{P}(T \leq \epsilon)}{\mathbb{P}(T \leq \eta\epsilon_0)}$$

where T follows $\text{BesselK}(k, \Delta_{S_f}/\epsilon_0)$ *and* $\eta = \frac{\Delta_f}{\Delta_{S_f}}$.

Proof. We provide the sketch of the proof. Proof follows from the proof of Lemma 8. For a Laplace mechanism calibrated with the sampled sensitivity Δ_{S_f} and privacy level ϵ_0, Equation 15 translates to,

$$\left(\frac{\epsilon_0}{\Delta_{S_f}} \sum_{i=1}^{k} |(|f(y_i) - z| - |f(x_i) - z|)| \right) \sim \text{Truncated BesselK}(k, \Delta_{S_f}/\epsilon_0, \Delta_f).$$

since, $\sum_{i=1}^{k} |(|f(y_i) - z| - |f(x_i) - z|)| \leq \Delta_f$. Using Lemma 7 and Equation 16,

$$\mathbb{P}(A_{\epsilon_0}^{\Delta_{S_f}}) = \frac{\mathbb{P}(T \leq \epsilon)}{\mathbb{P}(T \leq \eta\epsilon_0)}$$

where T follows $\text{BesselK}(k, \Delta_{S_f}/\epsilon_0)$ and $\eta = \frac{\Delta_f}{\Delta_{S_f}}$. □

For this case, we do not assume the knowledge of the sensitivity of the query. Using the empirical estimation presented in Section 3.2, if we choose the sampled sensitivity for privacy at risk $\gamma_2 = 1$, we obtain an approximation for η.

Lemma 9. *For a given value of accuracy parameter ρ,*

$$\frac{\Delta_f}{\Delta_{S_f}^*} = 1 + \mathcal{O}\left(\frac{\rho}{\Delta_{S_f}^*} \right)$$

where $\Delta_{S_f}^* = F_n^{-1}(1)$. $\mathcal{O}\left(\frac{\rho}{\Delta_{S_f}^*} \right)$ *denotes order of* $\frac{\rho}{\Delta_{S_f}^*}$, *i.e.,* $\mathcal{O}\left(\frac{\rho}{\Delta_{S_f}^*} \right) = k\frac{\rho}{\Delta_{S_f}^*}$ *for some* $k \geq 1$.

Proof. For a given value of accuracy parameter ρ and any $\Delta > 0$,

$$F_n(\Delta) - F(\Delta) \leq \rho$$

Since above inequality is true for any value of Δ, let $\Delta = F^{-1}(1)$. Therefore,

$$F_n(F^{-1}(1)) - F(F^{-1}(1)) \leq \rho$$
$$F_n(F^{-1}(1)) \leq 1 + \rho \tag{22}$$

Since a cumulative distribution function is 1-Lipschitz [(Papoulis and Pillai, 2002)],

$$|F_n(F_n^{-1}(1)) - F_n(F^{-1}(1))| \leq |F_n^{-1}(1) - F^{-1}(1)|$$
$$|F_n(F_n^{-1}(1)) - F_n(F^{-1}(1))| \leq |\Delta_{S_f}^* - \Delta_f|$$
$$\rho \leq \Delta_f - \Delta_{S_f}^*$$
$$1 + \frac{\rho}{\Delta_{S_f}^*} \leq \frac{\Delta_f}{\Delta_{S_f}^*}$$

where we used result from Equation 22 in step 3. Introducing $\mathcal{O}\left(\frac{\rho}{\Delta_{S_f}^*} \right)$ completes the proof. □

Lemma 10. *For Laplace Mechanism $\mathcal{L}_{\epsilon_0}^{\Delta_{S_f}}$ with sampled sensitivity Δ_{S_f} of a query $f : \mathcal{D} \to \mathbb{R}^k$ and for any $Z \subseteq Range(\mathcal{L}_\epsilon^{\Delta_{S_f}})$,*

$$\mathbb{P}\left[\log \left| \frac{\mathbb{P}(\mathcal{L}_{\epsilon_0}(x) \in Z)}{\mathbb{P}(\mathcal{L}_{\epsilon_0}(y) \in Z)} \right| \leq \epsilon \right] \geq \frac{\mathbb{P}(T \leq \epsilon)}{\mathbb{P}(T \leq \eta\epsilon_0)} \gamma_2 (1 - 2e^{-2\rho^2 n})$$

where n is the number of samples used to find sampled sensitivity, $\rho \in [0, 1]$ is a accuracy parameter and $\eta = \frac{\Delta_f}{\Delta_{S_f}}$. The outer probability is calculated with respect to support of the data-generation distribution $\mathcal{G}$.

Proof. The proof follows from the proof of Lemma 8 and Lemma 10. Consider,

$$\mathbb{P}(A_{\epsilon_0}^{\Delta_{S_f}}) \geq \mathbb{P}(A_{\epsilon_0}^{\Delta_{S_f}} | B^{\Delta_{S_f}}) \mathbb{P}(B^{\Delta_{S_f}} | C_\rho^{\Delta_{S_f}}) \mathbb{P}(C_\rho^{\Delta_{S_f}})$$
$$\geq \frac{\mathbb{P}(T \leq \epsilon)}{\mathbb{P}(T \leq \eta\epsilon_0)} \cdot \gamma_2 \cdot (1 - 2e^{-2\rho^2 n}) \tag{23}$$

The first term in the final step of Equation 23 follows from the result in Corollary 2 where T follows $\text{BesselK}(k, \frac{\Delta_{S_f}}{\epsilon_0})$. It is the probability with which the Laplace mechanism $\mathcal{L}_{\epsilon_0}^{\Delta_{S_f}}$ satisfies ϵ-differential privacy for a given value of sampled sensitivity. □

Probability of occurrence of event $A_{\epsilon_0}^{\Delta_{S_f}}$ calculated by accounting for both explicit and implicit sources of randomness gives the risk for privacy level ϵ. Thus, the proof of Lemma 10 completes the proof for Theorem 4.

Comparing the equations in Theorem 4 and Lemma 10, we observe that

$$\gamma_3 \triangleq \frac{\mathbb{P}(T \leq \epsilon)}{\mathbb{P}(T \leq \eta\epsilon_0)} \cdot \gamma_2 \tag{24}$$

The privacy at risk, as defined in Equation 24, is free from the term that accounts for the accuracy of sampled estimate. If we know cumulative distribution of the sensitivity, we do not suffer from the uncertainty of introduced by sampling from the empirical distribution.