Tailoring Differentially Private Bayesian Inference to Distance Between Distributions

Jiawen Liu[*], Mark Bun[**], Gian Pietro Farina[*], and Marco Gaboardi[*]

[*]University at Buffalo, SUNY.
{jliu223,gianpiet,gaboardi}@buffalo.edu
[**]Princeton University. mbun@cs.princeton.edu

Contents

Abstract

Bayesian inference is a statistical method which allows one to derive a *posterior* distribution, starting from a *prior* distribution and observed data. Several approaches have been explored in order to make this process differentially private. For example, [6], and [11] proved that, under specific conditions, sampling from the posterior distribution is already differentially private. [15], [9], designed differentially private mechanisms that output a representation of the full posterior distribution.

When the output of a differentially private mechanism is a probability distribution, accuracy is naturally measured by means of *probabilistic distances* measuring how far this distribution is from the original one. Some classical examples are total variation distance, Hellinger distance, χ^2-distance, KL-divergence, etc.

In this work, we explore the design space for differentially private bayesian inference, by considering different metrics and different algorithms we design a mechanism. We focus on two discrete models, the Beta-Binomial and the Dirichlet-multinomial models, and one probability distance, Hellinger distance. Our mechanism can be understood as a version of the exponential mechanism where the noise is calibrated to the smooth sensitivity of the utility function, rather than to its global sensitivity. In our setting, the utility function is the probability distance we want to use to measure accuracy. To show the usefulness of this mechanism we show an experimental analysis comparing it with mechanisms based on the Laplace mechanism.

1 Introduction

Data analysis techniques are broadly used in various applications in different areas to improve their services, including disease-medicine service, financial service, location service, social network and so on. In order to provide a better service, large of data are collected from users for analysis. As a consequence, the data privacy came to be a critical issue. Sensitivity information in data can be revealed through the analysis results. The key challenge here is to release a private analysis results, from which adversary cannot observe individual's sensitive information in data.

Plenty of work have been developed to solve this issue, guaranteeing the privacy in specific data analysis algorithms. They achieved the $\epsilon-$differential privacy by adopting either Laplace mechanism or achieved the $(\epsilon, \delta)-$differential privacy. But they are not giving better accuracy than the differential privacy mechanism itself. Here, we are proposing mechanism with better accuracy.

Our work is conducted under a Bayesian inference scenario, where the posterior distribution is the analysis result we obtained from the data. Publishing the posterior distribution inferred from a sensitive dataset can leak information about the individuals in the dataset. In order to guarantee differential privacy and to protect the individuals' data we can add noise to the posterior before releasing it. The amount of the noise that we need to introduced depends on the privacy parameter ϵ and the sensitivity of the inference to small changes in the data set. Sensitivity can be computed in many different ways based on which metric space we consider on the output set of the mechanism. In the literature on private Bayesian inference ([15, 13]), it is only measured with respect to the vector of numbers parametrizing the output distribution using, e.g. the ℓ_1 norm. A more natural approach which we explore here, is to measure sensitivity with respect to a metric on the space of inferred probability distributions. A re-loved question is that of how to measure accuracy. Again, this can be answered in different ways based on the metric imposed on the output space, and yet again only in few works in literature (e.g. [15]) distances between probability measures have been used for these purposes.

The question that this work aims at answering is whether an approach based on probability metrics can improve on the accuracy of approaches based on metrics over the numeric parameters of the distributions. We will see that in some cases this can happen.

Main contributions.

- We designed a differentially private Bayesian inference mechanism based on the standard exponential mechanism.
- We explored two ways to improve the accuracy: 1) calibrating noise to the sensitivity of a metric over distributions (e.g. Hellinger distance ($\mathcal{H}$), f-divergences, etc. . .). 2) Using a smooth upper bound on the local sensitivity and scale the noise to this smooth bound rather than global sensitivity, to improve the mechanism accuracy.
- A full proof on the newly designed mechanism is $(\epsilon, \delta)-$differential privacy is given in paper.
- We implemented the new proposed mechanism and other art-of-state mechanisms, comparing the performance in terms of accuracy and privacy.

Related Work.

A plentiful of data analysis algorithms have been studied to preserve differential privacy, including the subspace clustering algorithm [10], the gradient decedent algorithm in deep learning [1], logical regression [3], principle component analysis [4], probabilitic inference [12] and convergence in statistic estimation [2], etc.

In Bayesian Inference data analysis, mechanisms are proposed corresponded to maintain their differential privacy, focusing on 3 different goals: 1) Inherited differential privacy property of posterior sampling in Bayesian inference. [6], [15], [16] and [11]. 2) Data sampled and released from posterior distribution of Bayesian is differentially private [14], [7], [9]. 3) The inference process is differentially private and the posterior distribution released should be private itself, in the meantime, with good accuracy. The third topic where our work focus on is still very new.

2 Preliminaries

Bayesian Inference.

Given a prior belief $\Pr(\theta)$ on some parameter θ, and an observation $\mathbf{x}$, the posterior distribution on θ given $\mathbf{x}$ is computed as:

$$\Pr(\theta|\mathbf{x}) = \frac{\Pr(\mathbf{x}|\theta) \cdot \Pr(\theta)}{\Pr(\mathbf{x})}$$

where the expression $\Pr(\mathbf{x}|\theta)$ denotes the *likelihood* of observing $\mathbf{x}$ under a value of θ. Since we consider $\mathbf{x}$ to be fixed, the likelihood is a function of θ. For the same reason $\Pr(\mathbf{x})$ is a constant independent of θ. Usually in statistics the prior distribution $\Pr(\theta)$ is chosen so that it represents the initial belief on θ, that is, when no data has been observed. In practice though, prior distributions and likelihood functions are usually chosen so that the posterior belongs to the same *family* of distributions. In this case we say that the prior is conjugate to the likelihood function. Use of a conjugate prior simplifies calculations and allows for inference to be performed in a recursive fashion over the data. Then, we have:

$$\mathsf{bysInfer}(\Pr(\mathbf{x}|\theta), \Pr(\theta), \mathbf{x}) = \frac{\Pr(\mathbf{x}|\theta) \cdot \Pr(\theta)}{\Pr(\mathbf{x})}$$

Beta-binomial System.

In this work we will consider a specific instance of Bayesian inference and one of its generalizations. specifically, a Beta-binomial mode. We will consider the situation the underlying data is binomial distribution ($\sim$ *binomial*(θ)), where θ represents the parameter –informally called *bias*– of a Bernoulli distributed random variable. The prior distribution over $\theta \in [0,1]$ is going to be a beta distribution, $\mathsf{beta}(\alpha, \beta)$, with parameters $\alpha, \beta \in \mathbb{R}^+$, and with p.d.f:

$$\Pr(\theta) \equiv \frac{\theta^\alpha (1-\theta)^\beta}{\mathrm{B}(\alpha, \beta)}$$

where $\mathrm{B}(\cdot,\cdot)$ is the beta function. The data $\mathbf{x}$ will be a sequence of $n \in \mathbb{N}$ binary values, that is $\mathbf{x} = (x_1, \ldots x_n), x_i \in \{0,1\}$, and the likelihood function is:

$$\Pr(\mathbf{x}|\theta) \equiv \theta^{\Delta\alpha}(1-\theta)^{n-\Delta\alpha}$$

where $\Delta\alpha = \sum_{i=1}^{n} x_i$. From this, the inference can be easily derived:

$$\mathsf{bysInfer}(\mathit{binomial}(\theta), \mathsf{beta}(\alpha, \beta), \mathbf{x}) = \mathsf{beta}(\alpha + \Delta\alpha, \beta + n - \Delta\alpha)$$

Dirichlet-multinomial Systems.

The beta-binomial model can be immediately generalized to Dirichlet-multinomial, with underlying data multinomially distributed. The *bias* is represented by parameter $\boldsymbol{\theta}$, the vector of parameters of a categorically distributed random variable. The prior distribution over $\boldsymbol{\theta} \in [0,1]^k$ is given by a Dirichelet distribution,

$\mathsf{dirichlet}(\boldsymbol{\alpha})$, for $k \in \mathbb{N}$, and $\boldsymbol{\alpha} \in (\mathbb{R}^+)^k$, with p.d.f:

$$\Pr(\boldsymbol{\theta}) \equiv \frac{1}{\mathrm{B}(\boldsymbol{\alpha})} \cdot \prod_{i=1}^{k} \theta_i^{\alpha_i - 1}$$

where $\mathrm{B}(\cdot)$ is the generalized beta function. The data $\mathbf{x}$ will be a sequence of $n \in \mathbb{N}$ values coming from a universe $\mathcal{X}$, such that $\mid \mathcal{X} \mid= k$. The likelihood function will be:

$$\Pr(\mathbf{x}|\boldsymbol{\theta}) \equiv \prod_{a_i \in \mathcal{X}} \theta_i^{\Delta\alpha_i},$$

with $\Delta\alpha_i = \sum_{j=1}^{n} [x_j = a_i]$, where $[\cdot]$ represents Iverson bracket notation. Denoting by $\Delta\boldsymbol{\alpha}$ the vector $(\Delta\alpha_1, \dots \Delta\alpha_k)$ the inference over $\boldsymbol{\theta}$ turns out to be

$$\mathsf{bysInfer}(\mathit{multinomial}(\theta), \mathsf{dirichlet}(\boldsymbol{\alpha}), \mathbf{x}) = \mathsf{dirichlet}(\boldsymbol{\alpha} + \Delta\boldsymbol{\alpha})$$

where $+$ denotes the componentwise sum of vectors of reals.
Differential Privacy.

Definition 1. *$\epsilon-$differential privacy.*
A randomized mechanism $\mathcal{M} : \mathcal{X} \to \mathcal{Y}$ is differential privacy, iff for any adjacent[1] *input $\boldsymbol{x}, \boldsymbol{x'} \in \mathcal{X}$, a metric H over $\mathcal{Y}$ and a $B \subseteq H(\mathcal{Y})$, $\mathcal{M}$ satisfies:*

$$\Pr[H(\mathcal{M}(\boldsymbol{x})) \in B] = e^{\epsilon} \Pr[H(\mathcal{M}(\boldsymbol{x'})) \in B].$$

Definition 2. *$(\epsilon, \delta)-$differential privacy.*
A randomized mechanism $\mathcal{M} : \mathcal{X} \to \mathcal{Y}$ is differential privacy, iff for any $\boldsymbol{adj}(\boldsymbol{x}, \boldsymbol{x'}) \in \mathcal{X}$, a metric H over $\mathcal{Y}$ and a $B \subseteq H(\mathcal{Y})$, $\mathcal{M}$ satisfies:

$$\Pr[H(\mathcal{M}(\boldsymbol{x})) \in B] = e^{\epsilon} \Pr[H(\mathcal{M}(\boldsymbol{x'})) \in B] + \delta.$$

3 Technical Problem Statement and Motivations

We are interested in exploring mechanisms for privately releasing the full posterior distributions derived in section 2, as opposed to just sampling from them. It's worth noticing that the posterior distributions are fully characterized by their parameters, and the family (Beta, Dirichlet distribution) they belong to. Hence, in case of the Beta-Binomial model we are interested in releasing a private version of the pair of parameters $(\alpha', \beta') = (\alpha + \Delta\alpha, \beta + n - \Delta\alpha)$, and in the case of the Dirichlet-multinomial model we are interested in a private version of $\boldsymbol{\alpha}' = (\boldsymbol{\alpha} + \Delta\boldsymbol{\alpha})$. [15] and [13] have already attacked this problem by adding independent Laplacian noise to the parameters of the posteriors.

[1] Given $\mathbf{x}, \mathbf{x}'$ we say that $\mathbf{x}$ and $\mathbf{x}'$ are adjacent and we write, $\mathbf{adj}(\mathbf{x}, \mathbf{x}')$, iff $\sum_{i}^{n} [x_i = x'_i] \leq 1$.

That is, in the case of the Beta-Binomial system, the value released would be: $(\tilde{\alpha}, \tilde{\beta}) = (\alpha + \widetilde{\Delta\alpha}, \beta + n - \widetilde{\Delta\alpha})$ where $\widetilde{\Delta\alpha} \sim \mathsf{Lap}(\Delta\alpha, \frac{2}{\epsilon})$, and where $\mathsf{Lap}(\mu, \nu)$ denotes a Laplace random variable with mean μ and scale ν. This mechanism is ϵ-differentially private, and the noise is calibrated w.r.t. to a sensitivity of 2 which is derived by using ℓ_1 norm over the pair of parameters. Indeed, considering two adjacent data observations $\mathbf{x}, \mathbf{x}'$, that, from a unique prior, give rise to two posterior distributions, characterized by the pairs (α', β') and (α'', β'') then $|\alpha' - \alpha''| + |\beta' - \beta''| \leq 2$. This argument extends similarly to the Dirichelet-Multinomial system. Details are introduced in Sec. 4.

However, in previous works, the accuracy of the posterior was measured again with respect to ℓ_1 norm. That is, an upper bound was given on

$$\Pr[|\alpha - \tilde{\alpha}| + |\beta - \tilde{\beta}| \geq \gamma]$$

where $(\alpha, \beta), (\tilde{\alpha}, \tilde{\beta})$ are as defined above. This accuracy metric is meaningless when the results released are distributions rather than numerical values. In contrast, distribution metrics such as f-divergence, Hellinger distance, etc. come into mind overtly when we are measuring distance between distributions. This gives us motivation on exploring the design space of mechanisms by considering different norms (a distribution metric) to compute the sensitivity and provide guarantees on the accuracy.

Specifically, we will use the Hellinger distance $\mathcal{H}(\cdot, \cdot)$: Given two beta distributions $\mathsf{beta}(\alpha_1, \beta_1)$, and $\mathsf{beta}(\alpha_2, \beta_2)$ the following equality holds

$$\mathcal{H}(\mathsf{beta}(\alpha_1, \beta_1), \mathsf{beta}(\alpha_2, \beta_2)) = \sqrt{1 - \frac{\mathrm{B}(\frac{\alpha_1+\alpha_2}{2}, \frac{\beta_1+\beta_2}{2})}{\sqrt{\mathrm{B}(\alpha_1, \beta_1)\mathrm{B}(\alpha_2, \beta_2)}}}$$

Our choice to use Hellinger distance is motivated by three facts:

- It simplifies calculations in the case of the probabilistic models considered here.

- It also automatically yields bounds on the total variation distance, which represents also the maximum advantage an unbounded adversary can have in distingishing two distributions.

- The accuracy can be improved by using a smooth bound on Hellinger distance's local sensitivity. As shown in Fig. 1, taking advantage of the gap between the global sensitivity and local sensitivity, we can improve the accuracy by applying an upper bound on local sensitivity instead of using global sensitivity.

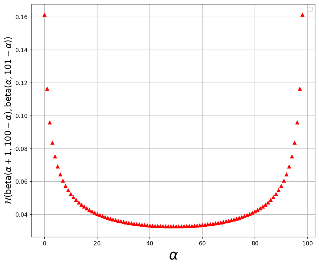

Figure 1: Sensitivity of $\mathcal{H}$

4 Mechanism Proposition

4.1 Laplace Mechanism Family

4.1.1 using ℓ_1 norm metric

Adding noise to the posterior distribution parameters directly, through Laplace mechanism with post-processing, $\mathsf{lapMech}(\mathsf{bysInfer}, \mathbf{x})$, producing an output:

$$\mathsf{lapMech}(\mathsf{bysInfer}, \mathbf{x}) = \mathsf{beta}(\alpha + \lfloor \Delta\alpha + Y \rceil_0^n, \beta + n - \lfloor \Delta\alpha + Y \rceil_0^n),$$

where $Y \sim \mathsf{Lap}(0, \frac{\Delta\mathsf{bysInfer}}{\epsilon})$, a Laplace distribution with location 0 and scale $\frac{\Delta\mathsf{bysInfer}}{\epsilon}$ in Beta-binomial model; and

$$\mathsf{lapMech}(\mathsf{bysInfer}, \mathbf{x}) = \mathsf{dirichlet}(\alpha_1 + \lfloor \Delta\alpha_1 + Y_1 \rceil_0^n, \cdots, \alpha_k + \lfloor n - \sum_{i=1}^{k-1} \lfloor \Delta\alpha_i + Y_i \rceil_0^n \rceil_0^n),$$

where $Y_i \sim \mathsf{Lap}(0, \frac{\Delta\mathsf{bysInfer}}{\epsilon})$ in Dirichlet-multinomial model. $\lfloor \cdot \rceil_0^n$ is taking the floor value and truncating into $[0, n]$ to make sure the noised posterior is valid.

We use $\mathsf{lapMech}(\mathbf{x}) = \mathsf{beta}(\tilde{\alpha}, \tilde{\beta})$ or $\mathsf{dirichlet}(\tilde{\boldsymbol{\alpha}})$ for short. Then release it as the private posterior distribution.

The sensitivity of the inference process w.r.t. ℓ_1 norm in $\mathsf{lapMech}$ is:

$$\Delta\mathsf{bysInfer} \equiv \max_{\mathbf{x}, \mathbf{x}' \in \{0,1\}^n, ||\mathbf{x} - \mathbf{x}'||_1 \leq 1} ||\mathsf{bysInfer}(\mathbf{x}) - \mathsf{bysInfer}(\mathbf{x}')||_1,$$

which is proportional to the dimensionality.

4.1.2 using improved ℓ_1 norm metric

Noise added to posterior distribution parameters are scaled to smaller sensitivity in this improved Laplace mechanism. Because in terms of two adjacent data sets $\mathbf{x}, \mathbf{x}'$, their posterior distributions by Bayesian inference – $\mathsf{bysInfer}(\mathbf{x}), \mathsf{bysInfer}(\mathbf{x}')$ – which parameter differs at most in 2 dimensions even though extended to Dirichlet-multinomial mode, i.e., $\Delta\mathsf{bysInfer} \leq 2$.

Then it is enough to use sensitivity 1 in 2 dimensions and 2 in higher dimensions, then the improved Laplace mechanism $\mathsf{ilapMech}(\mathsf{bysInfer}, \mathbf{x})$ is producing an output:

$$\mathsf{beta}(\alpha + \lfloor \Delta\alpha + Y \rceil_0^n, \beta + n - \lfloor \Delta\alpha + Y \rceil_0^n),$$

where $Y \sim \mathsf{Lap}(0, \frac{1}{\epsilon})$ in Beta-binomial model; and

$$\mathsf{dirichlet}(\alpha_1 + \lfloor \Delta\alpha_1 + Y_1 \rceil_0^n, \cdots, \alpha_k + \lfloor n - \sum_{i=1}^{k-1} \lfloor \Delta\alpha_i + Y_i \rceil_0^n \rceil_0^n),$$

where $Y_i \sim \mathsf{Lap}(0, \frac{2}{\epsilon})$ in Dirichlet-multinomial model.

Both Laplace mechanism and improved one are $\epsilon-$differential privacy[8].

4.2 Exponential Mechanism Family

In this section, we explore the exponential mechanism family: $\mathsf{expMech}(\cdot, \cdot, \cdot)$ by considering Hellinger distance metrics and different sensitivities.

Given a prior distribution $\boldsymbol{\beta}_{\text{prior}} = \mathsf{beta}(\alpha, \beta)$ and a sequence of n observations $\mathbf{x} \in \{0, 1\}^n$, we define the follwing set as candidate set where the mechanisms in this family sample from:

$$\mathcal{R}_{\text{post}} \equiv \{\mathsf{bysInfer}(\mathit{binomial}(\theta), \boldsymbol{\beta}_{\text{prior}}, \mathbf{x}) \mid \mathbf{x} \in \{0, 1\}^n\}.$$

For simplicity:

$$\mathcal{R}_{\text{post}} \equiv \{\mathsf{beta}(\alpha', \beta') \mid \alpha' = \alpha + \Delta\alpha, \beta' = \beta + n - \Delta\alpha\}$$

over all $\mathbf{x} \in \{0, 1\}^n$, where $\Delta\alpha$ is as defined in Section 2, and

$$\mathcal{R}_{\text{post}} \equiv \{\mathsf{dirichlet}(\boldsymbol{\alpha}') \mid \boldsymbol{\alpha}' = \boldsymbol{\alpha} + \Delta\boldsymbol{\alpha}\},$$

over all $\mathbf{x} \in \mathcal{X}^n$ in Dirichlet-multinomial model.

We don't explicitly parametrize the result by the prior and likelihood, which from now on we consider fixed and we denote them by $\boldsymbol{\beta}_{\text{prior}}$ and $\mathit{binomial}(\theta)$, $\mathit{multinomial}(\theta)$. We use $\mathsf{bysInfer}(\mathbf{x})$ to denote the Bayesian inference process when the other two parameters are fixed.

4.2.1 Standard Exponential Mechanism

Standard exponential mechanism $\mathsf{expMech}(x, u, \mathcal{R}_{\text{post}})$ samples an element from the candidate set $\mathcal{R}_{\text{post}} = \{r_1, r_2, \cdots r_n\}$ with probability proportional to $\exp(\frac{\epsilon u(x,r)}{2GS})$:

$$\underset{z \sim \mathsf{expMech}(x,u,\mathcal{R}_{\text{post}})}{Pr}[z = r] = \frac{exp(\frac{\epsilon u(x,r)}{2GS})}{\Sigma_{r' \in \mathcal{R}}\ exp(\frac{\epsilon u(x,r')}{2GS})},$$

where $u(x, r)$ is the Hellinger scoring function over candidates, $-\mathcal{H}(\mathsf{bysInfer}(\mathbf{x}), r)$, and GS is the global sensitivity of the scoring function (i.e. the global sensitivity of the Bayesian inference w.r.t the Hellinger distance), calculated by:

$$GS = \max_{\{|\mathbf{x},\mathbf{x}'|\leq 1; \mathbf{x},\mathbf{x}' \in \mathcal{X}^n\}} \max_{\{r \in \mathcal{R}\}} |\mathcal{H}(\mathsf{bysInfer}(\mathbf{x}), r) - \mathcal{H}(\mathsf{bysInfer}(\mathbf{x}'), r)|$$

Exponential mechanism is $\epsilon-$differential privacy[8].

4.2.2 Exponential Mechanism with Hellinger Metric and Local Sensitivity

Exponential mechanism with local sensitivity $\mathsf{expMech}^{local}(x, u, \mathcal{R}_{\text{post}})$ share the same candidate set and utility function as it with standard exponential mechanism. This outputs a candidate $r \in \mathcal{R}$ with probability proportional to $exp(\frac{\epsilon u(x,r)}{2LS(\mathbf{x})})$:

$$\Pr_{z \sim \mathsf{expMech}^{local}(x,u,\mathcal{R}_{\text{post}})}[z = r] = \frac{exp(\frac{\epsilon u(x,r)}{2LS(\mathbf{x})})}{\Sigma_{r' \in \mathcal{R}}\ exp(\frac{\epsilon u(x,r')}{2LS(\mathbf{x})})},$$

where $LS(\mathbf{x})$ is the local sensitivity of scoring function, calculated by:

$$LS(\mathbf{x}) = \max_{\mathbf{x}' \in \mathcal{X}^n : \mathbf{adj}(\mathbf{x},\mathbf{x}'), r \in \mathcal{R}} |\mathcal{H}(\mathsf{bysInfer}(\mathbf{x}'), r) - \mathcal{H}(\mathsf{bysInfer}(\mathbf{x}'), r)|.$$

The exponential mechanism with local sensitivity is non-differential privacy[8].

4.2.3 Exponential Mechanism with Hellinger Metric and Smoothed Sensitivity

Definition 3. *Given data set $\boldsymbol{x}$ of size n, we define $\gamma-$smooth sensitivity of Bayesian inference process w.r.t. the Hellinger distance $S(\boldsymbol{x})$ is:*

$$S(\boldsymbol{x}) = \max_{\boldsymbol{x}'' \in \{0,1\}^n} \left\{ \frac{1}{\frac{1}{LS(\boldsymbol{x}'')} + \gamma \cdot d(\boldsymbol{x}, \boldsymbol{x}'')} \right\}, \tag{1}$$

where d is the Hamming distance between two data sets.

Theorem 4.1. *Given data set $\boldsymbol{x}$, the $\gamma-$smooth sensitivity of Bayesian inference process w.r.t. the Hellinger distance $S(\boldsymbol{x})$ satisfying:*
For any $\boldsymbol{adj}(\boldsymbol{x}, \boldsymbol{x}')$:

$$\frac{1}{S(\boldsymbol{x})} - \frac{1}{S(\boldsymbol{x}')} \leq \gamma. \tag{2}$$

Proof. of Theorem. 4.1.

For $\mathbf{adj}(\mathbf{x}, \mathbf{x}')$ and arbitrary $\mathbf{x}'' \in \{0, 1\}^n$:

By Equation (1):

$$S(\mathbf{x}) = \max_{\mathbf{x}'' \in \{0,1\}^n} \left\{ \frac{1}{\frac{1}{LS(\mathbf{x}'')} + \gamma \cdot d(\mathbf{x}, \mathbf{x}'')} \right\}$$

$$\frac{1}{S(\mathbf{x})} = \min_{\mathbf{x}'' \in \{0,1\}^n} \left\{ \frac{1}{LS(\mathbf{x}'')} + \gamma \cdot d(\mathbf{x}, \mathbf{x}'') \right\}$$

Since $d(\mathbf{x}, \mathbf{x}'') \leq d(\mathbf{x}, \mathbf{x}') + d(\mathbf{x}', \mathbf{x}'') \leq 1 + d(\mathbf{x}', \mathbf{x}'')$:

$$\begin{aligned}
&\leq \min_{\mathbf{x}'' \in \{0,1\}^n} \left\{ \frac{1}{LS(\mathbf{x}'')} + \gamma \cdot (1 + d(\mathbf{x}', \mathbf{x}'')) \right\} \\
&= \min_{\mathbf{x}'' \in \{0,1\}^n} \left\{ \gamma + \frac{1}{LS(\mathbf{x}'')} + \gamma \cdot d(\mathbf{x}', \mathbf{x}'') \right\} \\
&= \gamma + \min_{\mathbf{x}'' \in \{0,1\}^n} \left\{ \frac{1}{LS(\mathbf{x}'')} + \gamma \cdot d(\mathbf{x}', \mathbf{x}'') \right\} \\
&= \gamma + \frac{1}{S(\mathbf{x}')}
\end{aligned}$$

$\Longrightarrow$

$$\frac{1}{S(\mathbf{x})} - \frac{1}{S(\mathbf{x}')} \leq \gamma.$$

□

Definition 4. $\mathsf{expMech}^{smoo}(\boldsymbol{x})$ *outputs a candidate* $r \in \mathcal{R}_{\text{post}}$ *with probability*

$$\underset{z \sim \mathsf{expMech}^{smoo}(\boldsymbol{x})}{Pr}[z = r] = \frac{exp\left(\frac{-\epsilon \cdot u(\boldsymbol{x}, r)}{4 \cdot S(\boldsymbol{x})}\right)}{\sum\limits_{r' \in \mathcal{R}_{\text{post}}} \exp\left(\frac{-\epsilon \cdot u(\boldsymbol{x}, r')}{4 \cdot S(\boldsymbol{x})}\right)}.$$

where $u(\boldsymbol{x}, r) = \mathcal{H}(\mathsf{bysInfer}(\boldsymbol{x}), r)$ *for* $r \in \mathcal{R}_{\text{post}}$ *and* $S(\boldsymbol{x})$ *is the* $\gamma-$*smooth sensitivity of* $\mathcal{H}(\mathsf{bysInfer}(\boldsymbol{x}), -)$*, calculated from Definition 3 by setting* $\gamma = 1$*.*

Mechanisms in Exponential mechanism family also extends to the Dirichlet-multinomial system $\mathsf{dirichlet}(\boldsymbol{\alpha})$ by rewriting the Hellinger distance as:

$$\mathcal{H}(\mathsf{dirichlet}(\boldsymbol{\alpha}_1), \mathsf{dirichlet}(\boldsymbol{\alpha}_2)) = \sqrt{1 - \frac{\mathrm{B}(\frac{\boldsymbol{\alpha}_1 + \boldsymbol{\alpha}_2}{2})}{\sqrt{\mathrm{B}(\boldsymbol{\alpha}_1)\mathrm{B}(\boldsymbol{\alpha}_2)}}},$$

and by replacing the $\mathcal{R}_{\text{post}}$ with set of posterior Dirichlet distributions candidates. Also, the smooth sensitivity $S(\mathbf{x})$ in (4.1) will be computed by letting $\mathbf{x}'$ range over all the elements in $\mathcal{X}^n$ adjacent to $\mathbf{x}$. Notice that $\mathcal{R}_{\text{post}}$ has $\binom{n+1}{m-1}$ elements in this case. We will denote by $\mathsf{expMech}^D_{\mathcal{H}}$ the mechanism for the Dirichlet-multinomial system.

5 Privacy Analysis

5.1 Privacy of Laplace Mechanism Family

Mechanisms in Laplace family are $\epsilon-$differential privacy by [8].

5.2 Privacy of Exponential Mechanism Family

5.2.1 $\mathsf{expMech}(,,)$ $\epsilon-$Differential Privacy

Standard exponential mechanism in $\epsilon-$differential privacy by [8].

5.2.2 $\mathsf{expMech}^{local}(,,)$ non-Differential Privacy

Exponential mechanism with local sensitivity non-differential privacy.

5.2.3 $\mathsf{expMech}^{smoo}$ $\epsilon-$Differential Privacy Proof

Lemma 5.1. $\mathsf{expMech}^{smoo}$ *is ϵ-differential privacy.*

Proof. of Lemma 5.1.

By Definition 1, to proof Lemma 5.1, we need to prove:

For any $\mathbf{adj}(\mathbf{x}, \mathbf{x'}) \in \mathcal{X}$ and any beta distribution r:

$$\underset{z\sim\mathsf{expMech}^{smoo}(\mathbf{x})}{Pr}[z = r] \leq e^{\epsilon} \underset{z\sim\mathsf{expMech}^{smoo}(\mathbf{x'})}{Pr}[z = r].$$

By definition 4:

$$\begin{aligned}
\underset{z\sim\mathsf{expMech}^{smoo}(\mathbf{x})}{Pr}[z = r] &= \frac{\exp\left(\frac{-\epsilon\cdot u(\mathbf{x},r)}{4\cdot S(\mathbf{x})}\right)}{\sum\limits_{r'\in\mathcal{R}_{\text{post}}} \exp\left(\frac{-\epsilon\cdot u(\mathbf{x},r')}{4\cdot S(\mathbf{x})}\right)} \\
&= \frac{\exp\left(\frac{-\epsilon\cdot(u(\mathbf{x},r)+u(\mathbf{x'},r)-u(\mathbf{x'},r))}{4\cdot S(\mathbf{x})}\right)}{\sum\limits_{r'\in\mathcal{R}_{\text{post}}} \exp\left(\frac{-\epsilon\cdot u(\mathbf{x},r')}{4\cdot S(\mathbf{x})}\right)} \\
&= \frac{\exp\left(\frac{-\epsilon\cdot(u(\mathbf{x'},r))}{4\cdot S(\mathbf{x})}\right)}{\sum\limits_{r'\in\mathcal{R}_{\text{post}}} \exp\left(\frac{-\epsilon\cdot u(\mathbf{x},r')}{4\cdot S(\mathbf{x})}\right)} \cdot \exp\left(\frac{\epsilon\cdot(u(\mathbf{x'},r) - u(\mathbf{x},r))}{4\cdot S(\mathbf{x})}\right)
\end{aligned}$$

Because $S(\mathbf{x}) \geq LS(\mathbf{x}) \geq (u(\mathbf{x}', r) - u(\mathbf{x}, r))$:

$$\leq \frac{exp\big(\frac{-\epsilon\cdot(u(\mathbf{x}',r))}{4\cdot S(\mathbf{x})}\big)}{\sum\limits_{r'\in\mathcal{R}_{\text{post}}} \exp\big(\frac{-\epsilon\cdot u(\mathbf{x},r')}{4\cdot S(\mathbf{x})}\big)} \cdot \exp\big(\frac{\epsilon}{4}\big)$$

$$= \exp\big(\frac{\epsilon}{4}\big) \cdot \frac{\exp\big(\frac{-\epsilon\cdot(u(\mathbf{x}',r))}{4\cdot S(\mathbf{x})}\big)}{\sum\limits_{r'\in\mathcal{R}_{\text{post}}} \exp\big(\frac{-\epsilon\cdot u(\mathbf{x},r')}{4\cdot S(\mathbf{x})}\big)} \exp\big(\frac{\epsilon \cdot (u(\mathbf{x}',r))}{4 \cdot S(\mathbf{x}')}\big) \exp\big(\frac{-\epsilon \cdot (u(\mathbf{x}',r))}{4 \cdot S(\mathbf{x}')}\big)$$

$$= \exp\big(\frac{\epsilon}{4}\big) \cdot \frac{\exp\big(\frac{-\epsilon\cdot(u(\mathbf{x}',r))}{4\cdot S(\mathbf{x}')}\big)}{\sum\limits_{r'\in\mathcal{R}_{\text{post}}} \exp\big(\frac{-\epsilon\cdot u(\mathbf{x},r')}{4\cdot S(\mathbf{x})}\big)} \exp\big(\frac{\epsilon \cdot (u(\mathbf{x}',r))}{4}\big(\frac{1}{S(\mathbf{x}')} - \frac{1}{S(\mathbf{x})}\big)\big)$$

Because $u(\mathbf{x}', r) = \mathcal{H}(\mathsf{bysInfer}(\mathbf{x}'), r) \leq 1$:

$$\leq \exp\big(\frac{\epsilon}{4}\big) \cdot \frac{\exp\big(\frac{-\epsilon\cdot(u(\mathbf{x}',r))}{4\cdot S(\mathbf{x}')}\big)}{\sum\limits_{r'\in\mathcal{R}_{\text{post}}} \exp\big(\frac{-\epsilon\cdot u(\mathbf{x},r')}{4\cdot S(\mathbf{x})}\big)} \exp\big(\frac{\epsilon}{4}\big(\frac{1}{S(\mathbf{x}')} - \frac{1}{S(\mathbf{x})}\big)\big)$$

Because the property of $\gamma-$smooth sensitivity: $\frac{1}{S(\mathbf{x}')} - \frac{1}{S(\mathbf{x})} \leq \gamma$:

$$\leq \exp\big(\frac{\epsilon}{4}\big) \cdot \frac{\exp\big(\frac{-\epsilon\cdot(u(\mathbf{x}',r))}{4\cdot S(\mathbf{x}')}\big)}{\sum\limits_{r'\in\mathcal{R}_{\text{post}}} \exp\big(\frac{-\epsilon\cdot u(\mathbf{x},r')}{4\cdot S(\mathbf{x})}\big)} \exp\big(\frac{\epsilon}{4} \cdot \gamma\big)$$

$$= \exp\big(\frac{\epsilon}{4} + \frac{\epsilon}{4} \cdot \gamma\big) \cdot \frac{\exp\big(\frac{-\epsilon\cdot(u(\mathbf{x}',r))}{4\cdot S(\mathbf{x}')}\big)}{\sum\limits_{r'\in\mathcal{R}_{\text{post}}} \exp\big(\frac{-\epsilon\cdot u(\mathbf{x},r')}{4\cdot S(\mathbf{x})}\big)}$$

Doing the same transformation in the denominator:

$$= \exp\big(\frac{\epsilon}{4} + \frac{\epsilon}{4} \cdot \gamma\big) \cdot \frac{\exp\big(\frac{-\epsilon\cdot(u(\mathbf{x}',r))}{4\cdot S(\mathbf{x}')}\big)}{\sum_{r'\in\mathcal{R}_{\text{post}}} \exp\big(\frac{-\epsilon\cdot(u(\mathbf{x},r')+u(\mathbf{x}',r')-u(\mathbf{x}',r'))}{4\cdot S(\mathbf{x})}\big)}$$

$$= \exp\big(\frac{\epsilon}{4} + \frac{\epsilon}{4} \cdot \gamma\big) \cdot \frac{\exp\big(\frac{-\epsilon\cdot(u(\mathbf{x}',r))}{4\cdot S(\mathbf{x}')}\big)}{\sum_{r'\in\mathcal{R}_{\text{post}}} \exp\big(\frac{-\epsilon\cdot(u(\mathbf{x}',r')}{4\cdot S(\mathbf{x})}\big) \exp\big(\frac{-\epsilon\cdot(u(\mathbf{x},r')-u(\mathbf{x}',r'))}{4\cdot S(\mathbf{x})}\big)}$$

Because $S(\mathbf{x}) \geq LS(\mathbf{x}) \geq (u(\mathbf{x}, r) - u(\mathbf{x}', r)) \implies \frac{-\epsilon\cdot(u(\mathbf{x},r')-u(\mathbf{x}',r'))}{4\cdot S(\mathbf{x})} \geq \frac{-\epsilon}{2}$:

$$\leq \exp\big(\frac{\epsilon}{4} + \frac{\epsilon}{4} \cdot \gamma\big) \cdot \frac{\exp\big(\frac{-\epsilon\cdot(u(\mathbf{x}',r))}{4\cdot S(\mathbf{x}')}\big)}{\sum_{r'\in\mathcal{R}_{\text{post}}} \exp\big(\frac{-\epsilon\cdot(u(\mathbf{x}',r')}{4\cdot S(\mathbf{x})}\big) \exp\big(\frac{-\epsilon}{4}\big)}$$

$$= \exp\big(\frac{\epsilon}{4} + \frac{\epsilon}{4} \cdot \gamma\big) \cdot \frac{\exp\big(\frac{-\epsilon\cdot(u(\mathbf{x}',r))}{4\cdot S(\mathbf{x}')}\big)}{\sum_{r'\in\mathcal{R}_{\text{post}}} \exp\big(\frac{-\epsilon\cdot(u(\mathbf{x}',r')}{4\cdot S(\mathbf{x})}\big) \exp\big(\frac{-\epsilon}{4}\big) \exp\big(\frac{\epsilon\cdot(u(\mathbf{x}',r'))}{4\cdot S(\mathbf{x}')}\big) \exp\big(\frac{-\epsilon\cdot(u(\mathbf{x}',r'))}{4\cdot S(\mathbf{x}')}\big)}$$

$$= \exp\big(\frac{\epsilon}{4} + \frac{\epsilon}{4} \cdot \gamma\big) \cdot \frac{\exp\big(\frac{-\epsilon\cdot(u(\mathbf{x}',r))}{4\cdot S(\mathbf{x}')}\big)}{\sum_{r'\in\mathcal{R}_{\text{post}}} \exp\big(\frac{-\epsilon\cdot(u(\mathbf{x}',r')}{4\cdot S(\mathbf{x}')}\big) \exp\big(\frac{-\epsilon}{2}\big) \exp\big(\frac{-\epsilon\cdot(u(\mathbf{x},r'))}{4}\big(\frac{1}{S(\mathbf{x})} - \frac{1}{S(\mathbf{x}')}\big)\big)}$$

Because $u(\mathbf{x}', r) = \mathcal{H}(\mathsf{bysInfer}(\mathbf{x}'), r) \leq 1 \implies \frac{-\epsilon \cdot (u(\mathbf{x}, r'))}{4} \geq \frac{-\epsilon}{4}$:

$$\leq \exp\big(\frac{\epsilon}{4} + \frac{\epsilon}{4} \cdot \gamma\big) \cdot \frac{\exp\big(\frac{-\epsilon \cdot (u(\mathbf{x}', r))}{4 \cdot S(\mathbf{x}')}\big)}{\sum_{r' \in \mathcal{R}_{\text{post}}} \exp\big(\frac{-\epsilon \cdot (u(\mathbf{x}', r')}{4 \cdot S(\mathbf{x}')}\big) \exp\big(\frac{-\epsilon}{4}\big) \exp\big(\frac{-\epsilon}{4}(\frac{1}{S(\mathbf{x})} - \frac{1}{S(\mathbf{x}')})\big)}$$

Because the property of $\gamma-$ smooth sensitivity: $\frac{1}{S(\mathbf{x})} - \frac{1}{S(\mathbf{x}')} \leq \gamma \implies \frac{-\epsilon}{4}(\frac{1}{S(\mathbf{x})} - \frac{1}{S(\mathbf{x}')}) \geq \frac{-\epsilon}{4} \cdot \gamma$

$$\begin{aligned}
&\leq \exp\big(\frac{\epsilon}{4} + \frac{\epsilon}{4} \cdot \gamma\big) \cdot \frac{\exp\big(\frac{-\epsilon \cdot (u(\mathbf{x}', r))}{4 \cdot S(\mathbf{x}')}\big)}{\sum_{r' \in \mathcal{R}_{\text{post}}} \exp\big(\frac{-\epsilon \cdot (u(\mathbf{x}', r')}{4 \cdot S(\mathbf{x}')}\big) \exp\big(\frac{-\epsilon}{4}\big) \exp\big(\frac{-\epsilon}{4} \cdot \gamma\big)} \\
&= \exp\big(\frac{\epsilon}{4} + \frac{\epsilon}{4} \cdot \gamma\big) \cdot \frac{\exp\big(\frac{-\epsilon \cdot (u(\mathbf{x}', r))}{4 \cdot S(\mathbf{x}')}\big)}{\sum_{r' \in \mathcal{R}_{\text{post}}} \exp\big(\frac{-\epsilon \cdot (u(\mathbf{x}', r')}{4 \cdot S(\mathbf{x}')}\big) \exp\big(\frac{-\epsilon}{4} + \frac{-\epsilon}{4} \cdot \gamma\big)} \\
&= e^{(\frac{\epsilon}{2} + \frac{\epsilon}{2} \cdot \gamma)} \cdot \frac{\exp\big(\frac{-\epsilon \cdot (u(\mathbf{x}', r))}{4 \cdot S(\mathbf{x}')}\big)}{\sum\limits_{r' \in \mathcal{R}_{\text{post}}} \exp\big(\frac{-\epsilon \cdot u(\mathbf{x}', r')}{4 \cdot S(\mathbf{x}')}\big)} \\
&= e^{(\frac{\epsilon}{2} + \frac{\epsilon}{2} \cdot \gamma)} \cdot \underset{z \sim \mathsf{expMech}^{smoo}(\mathbf{x}')}{Pr}[z = r]
\end{aligned}$$

Given $\gamma = 1$, $\epsilon-$differential privacy can be achieved.

□

6 Accuracy Analysis

6.1 Accuracy Bound for Baseline Mechanisms

6.1.1 Accuracy Bound for Laplace Mechanism

Given $Y \sim \mathsf{Lap}(0, b)$, we have[8]:$Pr[|Y| \geq t \cdot b] = e^{-t}$.
Based on this, we get:$Pr[|Y| \geq t] = e^{-\frac{t\epsilon}{\Delta \mathsf{bysInfer}}}$, where $Y \sim \mathsf{Lap}(0, \frac{\Delta \mathsf{bysInfer}}{\epsilon})$ in our setting.

Considering the post-processing (i.e., taking the floor value of Y) in $\mathsf{lapMech}$, we have:

$$Pr\big[\lfloor Y \rfloor = t\big] = Pr[t \leq Y < t + 1] = \frac{1}{2}(e^{-\frac{\epsilon(t)}{\Delta \mathsf{bysInfer}}} - e^{-\frac{\epsilon(t+1)}{\Delta \mathsf{bysInfer}}}).$$

when $t \geq 0$, and

$$Pr\big[\lfloor Y \rfloor = t\big] = Pr[t \leq Y < t + 1] = \frac{1}{2}(e^{\frac{\epsilon(t)}{\Delta \mathsf{bysInfer}}} - e^{\frac{\epsilon(t+1)}{\Delta \mathsf{bysInfer}}}).$$

when $t < 0$.

Let $\mathsf{beta}(\alpha, \beta)$ be the true posterior distribution, i.e., $\mathsf{bysInfer}(\mathbf{x}) = \mathsf{beta}(\alpha, \beta)$, and r_L be the posterior produced by Laplace mechanism, i.e., $r_L = \mathsf{beta}(\alpha + \lfloor Y \rfloor, \beta - \lfloor Y \rfloor)$. By applying Hellinger distance in our case, we get:

$$Pr\big[\mathcal{H}(\mathsf{bysInfer}(\mathbf{x}), r_L) = \mathcal{H}(\mathsf{beta}(\alpha, \beta), \mathsf{beta}(\alpha + t, \beta - t)\big] = \tfrac{1}{2}(e^{-\frac{\epsilon(t)}{\Delta \mathsf{bysInfer}}} - e^{-\frac{\epsilon(t+1)}{\Delta \mathsf{bysInfer}}})$$

when $t \geq 0$, and

$$Pr\big[\mathcal{H}(\mathsf{bysInfer}(\mathbf{x}), r_L) = \mathcal{H}(\mathsf{beta}(\alpha, \beta), \mathsf{beta}(\alpha + t, \beta - t)\big] = \tfrac{1}{2}(e^{\frac{\epsilon(t+1)}{\Delta \mathsf{bysInfer}}} - e^{\frac{\epsilon(t)}{\Delta \mathsf{bysInfer}}}).$$

when $t < 0$.

Unfolding the Hellinger distance formula($\mathcal{H}(\mathsf{beta}(\alpha, \beta), \mathsf{beta}(\alpha + t, \beta - t))$, we get:

case t is even

$$\begin{aligned} & Pr\left[\mathcal{H}(\mathsf{bysInfer}(\mathbf{x}), r_L)^2 = 1 - \prod_{k=0}^{\frac{t}{2}-1} \sqrt{1 - \frac{\frac{t}{2}}{a + k + \frac{t}{2}}} \cdot \prod_{k=1}^{\frac{t}{2}} \sqrt{1 - \frac{\frac{t}{2}}{\beta - k}}\ \right] \\ = \ & \frac{1}{2}(e^{-\frac{\epsilon(t)}{\Delta \mathsf{bysInfer}}} - e^{-\frac{\epsilon(t+1)}{\Delta \mathsf{bysInfer}}}) \end{aligned}$$

case t is odd

let $t = 2m + 1$

$$\begin{aligned} Pr\Bigg[\mathcal{H}(\mathsf{bysInfer}(\mathbf{x}), r_L)^2 = \quad & 1 - \frac{\Gamma(\alpha + \frac{1}{2})}{\Gamma(\alpha)} \cdot \frac{\Gamma(\beta - \frac{1}{2})}{\Gamma(\beta)} \\ & \cdot \prod_{k=0}^{m-1} \sqrt{(1 + \frac{1}{2(\alpha + k)})(1 + \frac{\frac{1}{2} - m}{(\alpha + m + k)})} \\ & \cdot \prod_{k=1}^{m} \sqrt{(1 + \frac{1}{2(\beta - \frac{1}{2} - k)})(1 + \frac{\frac{1}{2} - m}{(\beta - \frac{1}{2} - k)})}\ \Bigg] \\ = \quad & \frac{1}{2}(e^{-\frac{\epsilon(t)}{\Delta \mathsf{bysInfer}}} - e^{-\frac{\epsilon(t+1)}{\Delta \mathsf{bysInfer}}}) \end{aligned}$$

Given t with specific values ($1, 2, 3$ for example), we get following accuracy equations:

t = 0 $\underset{z \sim \mathsf{lapMech}(\mathbf{x})}{Pr}[\mathcal{H}(\mathsf{bysInfer}(\mathbf{x}), r_L) = 0] = 0.19673467014.$

t = 1 $\underset{z \sim \mathsf{lapMech}(\mathbf{x})}{Pr}[\mathcal{H}(\mathsf{bysInfer}(\mathbf{x}), r_L)^2 = 1 - \frac{\Gamma(\alpha + \frac{1}{2})}{\Gamma(\alpha)} \cdot \frac{\Gamma(\beta - \frac{1}{2})}{\Gamma(\beta)}] = 0.11932560927.$

$$\mathbf{t=2} \quad \underset{z\sim\mathsf{lapMech}(\mathbf{x})}{Pr}[\mathcal{H}(\mathsf{bysInfer}(\mathbf{x}), r_L)^2 = 1-\sqrt{1-\frac{\frac{t}{2}}{a+\frac{t}{2}}}\cdot\sqrt{1-\frac{\frac{t}{2}}{\beta}}] = 0.07237464051.$$

In all cases above, probability values are fixed even data size changes, only the distance are decreasing.

6.1.2 Accuracy Bound for Improved Laplace Mechanism

Accuracy bound for improved Laplace mechanism is obtained from the standard Laplace Mechanism by replacing the sensitivity of $\Delta\mathsf{bysInfer}$ with 1 in Beta-binomial model and 2 in Dirichlet-multinomial model.

Given t with the same specific values $(1, 2, 3$ as above), we get following accuracy equations:

$$\mathbf{t=0} \quad \underset{z\sim\mathsf{lapMech}(\mathbf{x})}{Pr}[\mathcal{H}(\mathsf{bysInfer}(\mathbf{x}), r_L) = 0] = 0.31606027941.$$

$$\mathbf{t=1} \quad \underset{z\sim\mathsf{lapMech}(\mathbf{x})}{Pr}[\mathcal{H}(\mathsf{bysInfer}(\mathbf{x}), r_L)^2 = 1-\frac{\Gamma(\alpha+\frac{1}{2})}{\Gamma(\alpha)}\cdot\frac{\Gamma(\beta-\frac{1}{2})}{\Gamma(\beta)}] = 0.11627207896.$$

$$\mathbf{t=2} \quad \underset{z\sim\mathsf{lapMech}(\mathbf{x})}{Pr}[\mathcal{H}(\mathsf{bysInfer}(\mathbf{x}), r_L)^2 = 1-\sqrt{1-\frac{\frac{t}{2}}{a+\frac{t}{2}}}\cdot\sqrt{1-\frac{\frac{t}{2}}{\beta}}] = 0.04277410743.$$

In all cases above, probability values are fixed even data size changes, only the distance are decreasing.

6.2 Accuracy Bound for $\mathsf{expMech}^{smoo}$

In Beta-binomial model, we have following formula for accuracy:

$$\underset{z\sim\mathsf{expMech}^{smoo}(\mathbf{x})}{Pr}[\mathcal{H}(\mathsf{bysInfer}(\mathbf{x}), z) = c] = \frac{\exp\left(\frac{-\epsilon c}{4S(\mathbf{x})}\right)}{\sum\limits_{r'\in\mathcal{R}_{\text{post}}}\exp\left(\frac{-\epsilon\cdot\mathcal{H}(\mathsf{BI}(\mathbf{x}),r')}{4\cdot S(\mathbf{x})}\right)}, \qquad (3)$$

$c = \mathcal{H}(\mathsf{beta}(\alpha,\beta), \mathsf{beta}(\alpha+t,\beta-t)$.

6.3 Accuracy Comparison between $\mathsf{expMech}^{smoo}$, $\mathsf{lapMech}$ and $\mathsf{ilapMech}$

For comparison with Laplace mechanism, we developed the Table 2.

In the case when $\alpha=\beta$, there are always 2 symmetric beta distributions with the same distance from the true posterior, then we have accuracy table as:

When ϵ and dimensions are fixed, the probability of getting the true posterior, or posterior with certain step from Laplace mechanism is fixed whatever the data size or prior changes.

By solving following equations:

c	0	$1-\frac{\Gamma(\alpha+\frac{1}{2})}{\Gamma(\alpha)}\cdot\frac{\Gamma(\beta-\frac{1}{2})}{\Gamma(\beta)}$	$1-\sqrt{1-\frac{\frac{t}{2}}{a+\frac{t}{2}}}\cdot\sqrt{1-\frac{\frac{t}{2}}{\beta}}$
$\underset{z\sim\mathsf{lapMech}(\mathbf{x})}{Pr}[\mathcal{H}(\mathsf{bysInfer}(\mathbf{x}),z)=c]$	0.19673467014	0.11932560927	0.07237464051
$\underset{z\sim\mathsf{ilapMech}(\mathbf{x})}{Pr}[\mathcal{H}(\mathsf{bysInfer}(\mathbf{x}),z)=c]$	0.31606027941	0.11627207896	0.04277410743
$\underset{z\sim\mathsf{expMech}^{smoo}(\mathbf{x})}{Pr}[\mathcal{H}(\mathsf{bysInfer}(\mathbf{x}),z)=c]$	$\frac{1}{\sum\limits_{r'\in\mathcal{R}_{\text{post}}}\exp\left(\frac{-\epsilon\cdot\mathcal{H}(\mathsf{BI}(\mathbf{x}),r')}{4\cdot S(\mathbf{x})}\right)}$	$\frac{\exp\left(\frac{-\epsilon\sqrt{1-\frac{\Gamma(\alpha+\frac{1}{2})}{\Gamma(\alpha)}\cdot\frac{\Gamma(\beta-\frac{1}{2})}{\Gamma(\beta)}}}{4S(\mathbf{x})}\right)}{\sum\limits_{r'\in\mathcal{R}_{\text{post}}}\exp\left(\frac{-\epsilon\cdot\mathcal{H}(\mathsf{BI}(\mathbf{x}),r')}{4\cdot S(\mathbf{x})}\right)}$	$\frac{\exp\left(\frac{-\epsilon\sqrt{1-\sqrt{1-\frac{\frac{t}{2}}{a+\frac{t}{2}}}\cdot\sqrt{1-\frac{\frac{t}{2}}{\beta}}}}{4S(\mathbf{x})}\right)}{\sum\limits_{r'\in\mathcal{R}_{\text{post}}}\exp\left(\frac{-\epsilon\cdot\mathcal{H}(\mathsf{BI}(\mathbf{x}),r')}{4\cdot S(\mathbf{x})}\right)}$

Table 1: Accuracy Comparison in Theoretical

c	0	$1-\frac{\Gamma(\alpha+\frac{1}{2})}{\Gamma(\alpha)}\cdot\frac{\Gamma(\beta-\frac{1}{2})}{\Gamma(\beta)}$	$1-\sqrt{1-\frac{\frac{t}{2}}{a+\frac{t}{2}}}\cdot\sqrt{1-\frac{\frac{t}{2}}{\beta}}$
$\underset{z\sim\mathsf{lapMech}(\mathbf{x})}{Pr}[\mathcal{H}(\mathsf{bysInfer}(\mathbf{x}),z)=c]$	0.19673467014	0.31606027941	0.19170024978
$\underset{z\sim\mathsf{ilapMech}(\mathbf{x})}{Pr}[\mathcal{H}(\mathsf{bysInfer}(\mathbf{x}),z)=c]$	0.31606027941	0.43233235838	0.1590461864
$\underset{z\sim\mathsf{expMech}^{smoo}(\mathbf{x})}{Pr}[\mathcal{H}(\mathsf{bysInfer}(\mathbf{x}),z)=c]$	$\frac{1}{\sum\limits_{r'\in\mathcal{R}_{\text{post}}}\exp\left(\frac{-\epsilon\cdot\mathcal{H}(\mathsf{BI}(\mathbf{x}),r')}{4\cdot S(\mathbf{x})}\right)}$	$\frac{2\exp\left(\frac{-\epsilon\sqrt{1-\frac{\Gamma(\alpha+\frac{1}{2})}{\Gamma(\alpha)}\cdot\frac{\Gamma(\beta-\frac{1}{2})}{\Gamma(\beta)}}}{4S(\mathbf{x})}\right)}{\sum\limits_{r'\in\mathcal{R}_{\text{post}}}\exp\left(\frac{-\epsilon\cdot\mathcal{H}(\mathsf{BI}(\mathbf{x}),r')}{4\cdot S(\mathbf{x})}\right)}$	$\frac{2\exp\left(\frac{-\epsilon\sqrt{1-\sqrt{1-\frac{\frac{t}{2}}{a+\frac{t}{2}}}\cdot\sqrt{1-\frac{\frac{t}{2}}{\beta}}}}{4S(\mathbf{x})}\right.}{\sum\limits_{r'\in\mathcal{R}_{\text{post}}}\exp\left(\frac{-\epsilon\cdot\mathcal{H}(\mathsf{BI}(\mathbf{x}),r')}{4\cdot S(\mathbf{x})}\right.}$

Table 2: Accuracy Comparison in Theoretical

- $$\frac{1}{\sum\limits_{r'\in\mathcal{R}_{\text{post}}}\exp\left(\frac{-\epsilon\cdot\mathcal{H}(\mathsf{BI}(\mathbf{x}),r')}{4\cdot S(\mathbf{x})}\right)}=0.19673467014 \ or \ 0.31606027941$$

  we can get the range where we can do better than $\mathsf{lapMech}$ or $\mathsf{ilapMech}$ on outputting the correct answer.

- $$\frac{\exp\left(\frac{-\epsilon\sqrt{1-\frac{\Gamma(\alpha+\frac{1}{2})}{\Gamma(\alpha)}\cdot\frac{\Gamma(\beta-\frac{1}{2})}{\Gamma(\beta)}}}{4S(\mathbf{x})}\right)}{\sum\limits_{r'\in\mathcal{R}_{\text{post}}}\exp\left(\frac{-\epsilon\cdot\mathcal{H}(\mathsf{BI}(\mathbf{x}),r')}{4\cdot S(\mathbf{x})}\right)}=0.11932560927 \ or \ 0.11627207896$$

  we can get the range where we can do better than $\mathsf{lapMech}$ or $\mathsf{ilapMech}$ on outputting the answer with one step from correct one.

- $$\frac{2\exp\left(\frac{-\epsilon\sqrt{1-\sqrt{1-\frac{\frac{t}{2}}{a+\frac{t}{2}}}\cdot\sqrt{1-\frac{\frac{t}{2}}{\beta}}}}{4S(\mathbf{x})}\right)}{\sum\limits_{r'\in\mathcal{R}_{\text{post}}}\exp\left(\frac{-\epsilon\cdot\mathcal{H}(\mathsf{BI}(\mathbf{x}),r')}{4\cdot S(\mathbf{x})}\right)}=0.07237464051 \ or \ 0.04277410743$$

  we can get the range where we can do better than $\mathsf{lapMech}$ or $\mathsf{ilapMech}$ on outputting the answer with two steps from the correct one.

- Following equations can be in the same way.

7 Experimental Evaluations

7.1 Efficiency Evaluation

The formula for computing the local sensitivity presented in Sec. 4: $LS(\mathbf{x}) = \max\limits_{\mathbf{x}' \in \mathcal{X}^n : \mathbf{adj}(\mathbf{x},\mathbf{x}'), r \in \mathcal{R}} |\mathcal{H}(\mathsf{bysInfer}(\mathbf{x}'), r) - \mathcal{H}(\mathsf{bysInfer}(\mathbf{x}'), r)|$ can be reduced to $\max\limits_{\{|\mathbf{x},\mathbf{x}'| \leq 1; \mathbf{x}' \in \mathcal{X}^n\}} \mathcal{H}(\mathsf{bysInfer}(\mathbf{x}), \mathsf{bysInfer}$ by applying the distance triangle property.

Specifically, the maximum value over $r \in R$ always achieves at $r = \mathsf{bysInfer}(\mathbf{x})$:

$$
\begin{aligned}
LS(\mathbf{x}) &= \max_{\{|\mathbf{x},\mathbf{x}'| \leq 1; \mathbf{x}' \in \mathcal{X}^n\}} \{\mathcal{H}(\mathsf{bysInfer}(\mathbf{x}), \mathsf{bysInfer}(\mathbf{x})) - \mathcal{H}(\mathsf{bysInfer}(\mathbf{x}'), \mathsf{bysInfer}(\mathbf{x}))|\} \\
&= \max_{\{|\mathbf{x},\mathbf{x}'| \leq 1; \mathbf{x}' \in \mathcal{X}^n\}} \{\mathcal{H}(\mathsf{bysInfer}(\mathbf{x}'), \mathsf{bysInfer}(\mathbf{x}))|\}.
\end{aligned}
$$

This equation is validated by an experimental result shown in Fig. 2. We calculate the $\max\limits_{\{|\mathbf{x},\mathbf{x}'| \leq 1; \mathbf{x}' \in \mathcal{X}^n\}}$ value for every candidate $r \in \mathcal{R}_{\text{post}}$. It is shown that maximum value taken when $r = \mathsf{bysInfer}(\mathbf{x})$.

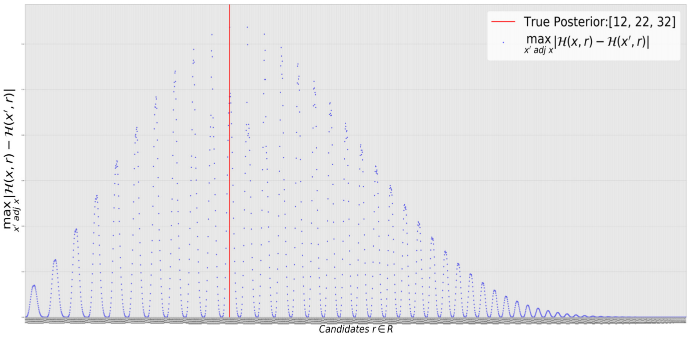

Figure 2: Experimental Results for Finding the Local Sensitivity Efficiently

7.2 Accuracy Evaluation

7.2.1 Theoretical Results

In Fig. 3 and 4, we plot on the x-axis the Hellinger distance from the true posterior and on the y-axis the theoretical probabilities of outputting the candidates with that distance under the different mechanisms. We consider *balanced* data sets, which means that in the Beta-Binomial model (Figure 4(a)) the datasets will consist of 50% 1s and the rest 0s, while for the Dirichelet-Multinomial (Figure 4(b)) the data will be split in the $k = 3$ bins with perecentages of: 33%, 33% and 34% in 3 dimensionality. Same concept in 4 dimensionality.

We consider 6 mechanisms in our comparison, including the Laplace mechanism, improved Laplace mechanism, standard exponential mechanism, non-private exponential mechanism (using local sensitivity) and two newly designed

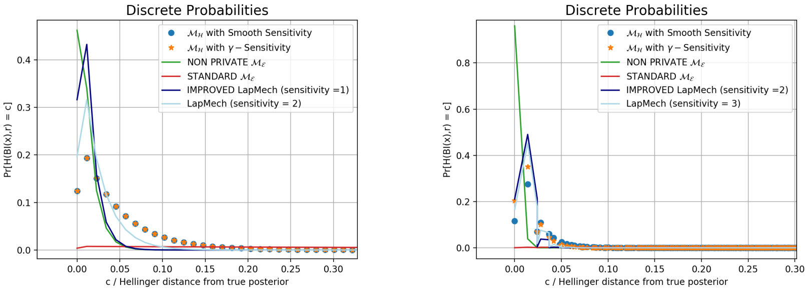

(a) 2 dimensions with data size 600

(b) 3 dimensions with data size 600

Figure 3: The theory probabilities of outputting candidates in certain distance from true posterior, with balanced data set and parameters $\epsilon = 1.0$

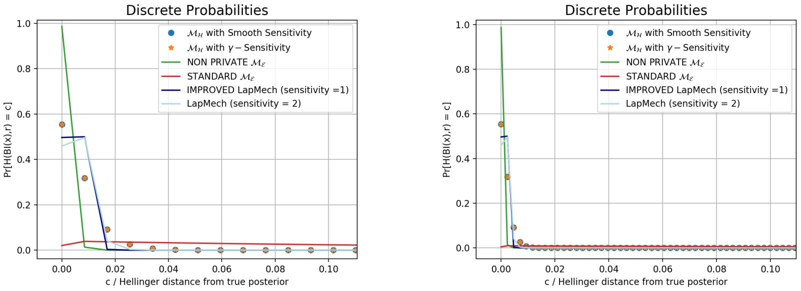

(a) 2 dimensions of data size 1000

(b) 2 dimensions of data size 10000

Figure 4: The theory probabilities of outputting candidates in certain distance from true posterior, with balanced data set and parameters $\epsilon = 5.0$

mechanisms (one with smooth sensitivity achieving $(\epsilon, \delta)-$dp and the other with γ sensitivity achieving $\epsilon-$dp).

In Fig. 3, candidates of smaller distance from true posterior can be outputted by $\mathsf{expMech}^{smoo}$ (in blue line) with larger probability than by baseline Laplace mechanism (in green line). This means $\mathsf{expMech}^{smoo}$ can produce good results with larger probability than baseline mechanism. However, the improved Laplace mechanism represented by red line can produce good results with probability higher than $\mathsf{expMech}^{smoo}$. It outperforms $\mathsf{expMech}^{smoo}$.

Increasing the privacy bound ϵ, we get theoretical results as in Fig. 4. In 2 dimensions, we can perform better than Laplace mechanisms.

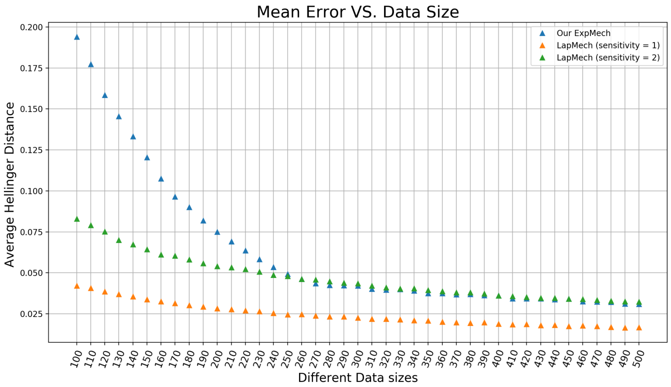

Figure 5: Increasing data size with prior $\mathsf{beta}(1, 1)$, balanced datasets and parameters $\epsilon = 1.0$

7.2.2 Experimental Results

In this section, we evaluate the accuracy of the mechanisms defined in Section (4) w.r.t. four variables, including data size, dimensions, data variance, prior distribution, and some combinations thereof. Every plot is an average over 1000 runs. In all the experiments we set $\epsilon = 1.0$.
In the following some of the plots show mean error as a function of the datasize while one is a whiskers-plot where the y-axis shows the average accuracy (or equivalently, the error) of the mechanisms, and the x-axis, instead shows different balanced priors used. The boxes extend from the lower to the upper quartile values of the data, with a line at the median. A notch on the box around the median is also drawn to give a rough guide to the significance of difference of medians; The whiskers extend from the box to show the range of the data. A blue box in the plots represents our newly designed exponential mechanism's behavior– where the sensitivity is calibrated w.r.t Hellinger distance– while the yellow box next to it represents the performance of a variation of the basic Laplace mechanism presented in Section (4) with the same settings: that is ϵ, δ, data, prior. The variation considered performs a postprocessing on the released parameters so that they are consistent. For instance when the sum of the noised parameters is greater than n we will truncate them so that they sum up to n.

Increasing data size with balanced datasets In Figures 5, 6(a) and 6(b) we still consider *balanced* data sets of observations. The results show that when the data size increases, the average errors of $\mathsf{expMech}^{smoo}$, Laplace mechanism and decrease. For small datasets, i.e with size less 300 in the case of Beta-Binomial models, both the baseline Laplace mechanisms and improved Laplace mechanism outperform $\mathsf{expMech}^{smoo}$. But for bigger data sets, that is, bigger

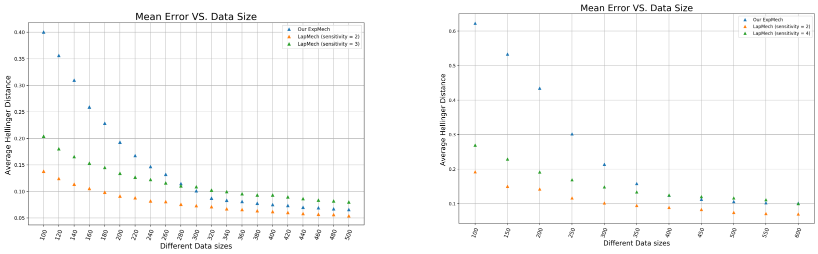

(a) Increasing data size with dirichlet$(1, 1, 1)$ prior distribution, balanced datasets and parameters $\epsilon = 1.0$

(b) Increasing data size with dirichlet$(1, 1, 1, 1)$ prior distribution, balanced datasets and parameters $\epsilon = 1.0$

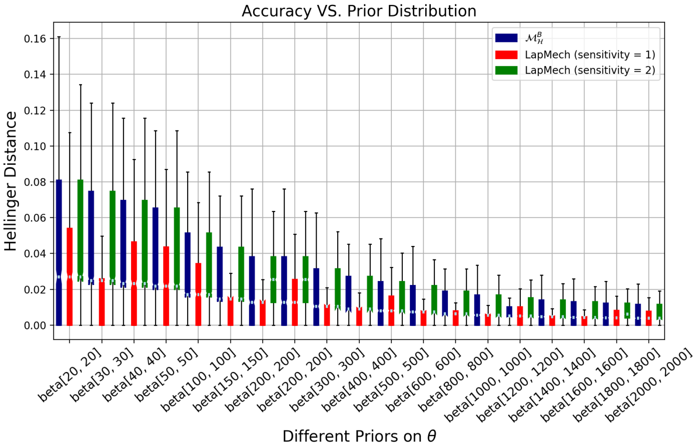

Figure 6: Observed data set is: $(50, 50)$, varying balanced priors

than 300, or as in Figure 5 where we considered data sets of the order of 15 thousands elements, the expMech^{smoo} outperforms the baseline Laplace mechanism, and asymptotically approaches the improved Laplace mechanism. Similar experimental tendencies were obtained for the Dirichlet-multinomial model (Figure 6(a) and 6(b)).

Fixed dataset varying balanced priors In Figure 6, we fix the data set to be $(50, 50)$, and the parameters the same as before: $\epsilon = 1.0$ and $\delta = 10^{-8}$. We studied the accuracy under different priors, where the priors considered are also balanced. Similar to the plots above, Figure 6 shows that in the beginning the baseline Laplace mechanism and improved Laplace mechanism performs better but the baseline approach is outperformed after a while, and very close to the improved Laplace mechanism.

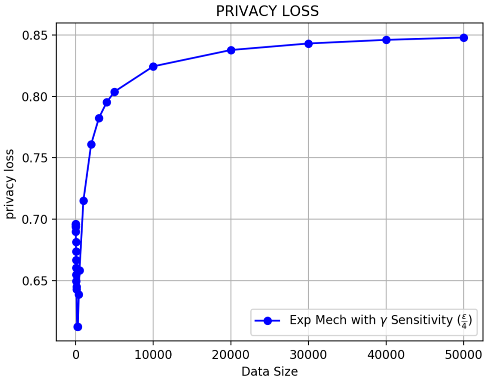

Figure 7: Actual privacy loss of different data size when privacy bound $\epsilon = 1.0$ in 2 dimensions, prior:$\mathsf{beta}(1,1)$ and balanced data

7.3 Privacy Evaluation

In order to see our privacy behavior, we study the accurate epsilon under concrete cases in this section. The (ϵ, δ) - differential privacy we proved in Sec. 4.2.3 is just an upper bound, we concrete ϵ should be smaller than upper bound in our exponential mechanism. We calculate the concrete privacy value in following ways wrt. the data size, and obtain plots in Fig. 7.

$\epsilon = 1.0$ is a privacy upper bound, we can observe that the concrete ϵ values are smaller than the upper bound. That is to say, we achieved a higher privacy level than expected.

8 Conclusion and Future Work

From what we have seen in the previous sections we can obtain some preliminary conclusions. That is, the probabiliy measure approach outperforms the ℓ_1-norm approach in the following cases:

1. $\mathsf{expMech}^{smoo}$ outperforms the baseline approach but not the improved one, for priors with small parameters.
2. When the prior parameters increase $\mathsf{expMech}^{smoo}$ is comparable with the improved baseline approach.

These results although very motivating, are still not enough for real world applications. Hence, we will continue our work in the follwoing directions:

1. For now, we just have a intuitive idea on the accuracy behavior of our mechanisms, and not a precise formula or bound on it. When do our

mechanisms perform better than the baseline mechanism and when they don't? How much influence will elements in Section 7 have on the accuracy? Are there any other important factors we missed? These are all questions w.r.t. the accuracy that we are going to explore next, and in a more principled and formal way.

2. Theorem 5.1 provides an upper bound on the privacy loss for $\mathsf{expMech}^{smoo}$ and $\mathsf{expMech}_{\mathcal{H}}^{D}$ but not necessarily a tight one. Indeed, experiments have shown that the actual privacy loss in the experiments can be smaller than ϵ. This means that we could improve accuracy, by adding less noise – that is noise proportional to a higher value of ϵ– but still achieve (ϵ, δ)-dp.

3. The choice of the Hellinger distance might seem quite ad-hoc. Hence, it is worth exploring other distances over distributions. An interesting class of probability metrics is the family of f-divergences [5].

References

[1] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *CCS 2016*, pages 308–318.

[2] Kamalika Chaudhuri and Daniel Hsu. Convergence rates for differentially private statistical estimation. In *ICML, 2012*, page 1327.

[3] Kamalika Chaudhuri and Claire Monteleoni. Privacy-preserving logistic regression. In *NIPS, 2009*, pages 289–296.

[4] Kamalika Chaudhuri, Anand Sarwate, and Kaushik Sinha. Near-optimal differentially private principal components. In *NIPS, 2012*, pages 989–997.

[5] I. Csiszár and P.C. Shields. Information theory and statistics: A tutorial. *Foundations and Trends® in Communications and Information Theory*, 1(4):417–528, 2004.

[6] Christos Dimitrakakis, Blaine Nelson, Aikaterini Mitrokotsa, and Benjamin IP Rubinstein. Robust and private bayesian inference. In *ALT, 2014*, pages 291–305.

[7] Christos Dimitrakakis, Blaine Nelson, Zuhe Zhang, Aikaterini Mitrokotsa, and Benjamin IP Rubinsein. Differential privacy in a bayesian setting through posterior sampling.

[8] Cynthia Dwork, Aaron Roth, et al. *The algorithmic foundations of differential privacy.* Now Publishers, Inc., 2014.

[9] James Foulds, Joseph Geumlek, Max Welling, and Kamalika Chaudhuri. On the theory and practice of privacy-preserving bayesian data analysis.

[10] Yining Wang, Yu-Xiang Wang, and Aarti Singh. Differentially private subspace clustering. In *NIPS, 2015*, pages 1000–1008.

[11] Yu-Xiang Wang, Stephen Fienberg, and Alex Smola. Privacy for free: Posterior sampling and stochastic gradient monte carlo. In *ICML, 2015*, pages 2493–2502.

[12] Oliver Williams and Frank McSherry. Probabilistic inference and differential privacy. In *NIPS, 2010*, pages 2451–2459.

[13] Yonghui Xiao and Li Xiong. Bayesian inference under differential privacy.

[14] Jun Zhang, Graham Cormode, Cecilia M. Procopiuc, Divesh Srivastava, and Xiaokui Xiao. Privbayes: Private data release via bayesian networks. pages 1423–1434.

[15] Zuhe Zhang, Benjamin IP Rubinstein, Christos Dimitrakakis, et al. On the differential privacy of bayesian inference. In *AAAI, 2016*, pages 2365–2371.

[16] Shijie Zheng. The differential privacy of bayesian inference. In *Bachelor's thesis, Harvard College, 2015.*