Journal of Privacy and Confidentiality (2016-2017) **7**, Number 3, 53–66

On the Meaning and Limits of Empirical Differential Privacy

Anne-Sophie Charest* and Yiwi Han †

1 Introduction

Differential privacy (DP) was introduced as a measure of confidentiality protection by Dwork et al. (2006b) and Dwork (2006). Designed to guarantee confidentiality even in a worst-case scenario, differential privacy protects the information of any individual in the database even against an adversary with complete knowledge of the rest of the dataset. This strong guarantee is achieved by limiting the influence that any one respondent can have on the released information.

Differentially-private mechanisms for various statistical tasks may now be found in the literature; see for e.g. Chaudhuri and Monteleoni (2009), Inan et al. (2010), Dwork and Smith (2010) and Smith (2011). However, the application of DP is rather complex for real datasets, as discussed (Dankar and El Emam, 2012), and is limited in practice. Several relaxations of DP have been proposed, including (ϵ, δ)-approximate differential privacy (Dwork et al., 2006a), (ϵ, δ)-probabilistic differential privacy (Machanavajjhala et al., 2008) and random differential privacy (Hall et al., 2011). While these relaxations often provide higher utility than DP, they are not that different from DP in that their application to a new statistical task requires the creation of an appropriate protection mechanism and a proof that it satisfies the privacy definition. These two tasks may be hard to achieve for data disseminators.

A significantly different approach is proposed (Abowd et al., 2013), where the goal is to conduct privacy-protected Bayesian mixed-effects modelling. Building on ideas presented in Abowd and Vilhuber (2008) and Machanavajjhala et al. (2008), the authors use a prior distribution diffuse enough to offer privacy protection of the posterior distribution, but they do not provide a guarantee that the method strictly satisfies DP, or any of its relaxation. Instead, they offer a convenient way to obtain "an empirical analogue of ϵ", an estimate of the DP protection of the statistical model used. The resulting privacy measure, termed Empirical Differential Privacy (EDP), is extremely attractive, as it can be computed easily for a large class of Bayesian models, even by users without the strong technical background required to produce DP mechanisms.

EDP has been applied to Bayesian linear mixed models in Abowd et al. (2013), Bayesian generalized linear mixed models in Schneider and Abowd (2015) and zero-inflated Poisson models in Schneider and Abowd (2015). However, little is known of its theoretical underpinnings. In this paper, we study carefully the meaning and limitations of EDP as a measure of privacy. We show that EDP is unfortunately not that

*Universitt́e Laval, Québec, Canada. mailto:anne-sophie.charest@met.ulaval.ca
†University of British Columbia, British Columbia, Canada. vickyhuw123@gmail.com

related to the original DP, that it is not entirely well-defined and that it can be very computationally demanding.

The rest of the paper is organized as follows. We define DP and EDP formally in Section 2. Section 3 discusses the meaning of EDP and explains why EDP can not simply be considered an empirical version of DP. In particular, we highlight the implicit conditioning in EDP, reframe EDP as a sensitivity measure on posterior distributions, and show that the EDP ϵ depends crucially on the choice of discretization used in the procedure. Section 4 considers the practical implementation of EDP and gives results of numerical experiments for two conjugate Bayesian models. We discuss the results in the last section.

2 Empirical Differential Privacy

This section introduces Empirical Differential Privacy (EDP), as proposed in Abowd et al. (2013). We first recall the definition of Differential Privacy (DP).

Differential Privacy: A randomized function κ is said to satisfy ϵ-differential privacy if and only if for all neighbouring datasets D_1 and D_2, and for all $S \subseteq \text{range}(\kappa)$,

$$e^{-\epsilon} \leq \frac{Pr[\kappa(D_1) \in S]}{Pr[\kappa(D_2) \in S]} \leq e^{\epsilon}. \tag{1}$$

For a matrix dataset, with rows associated to respondents and columns to variables, we usually consider two datasets to be neighbours if their entries are identical in all but one of the rows. Given an observed dataset D, one fixes a value of ϵ corresponding to the desired level of privacy, with smaller values of ϵ yielding greater privacy, and publishes the output $\kappa(D)$. Note that differential privacy is the property of the randomized function κ itself, not of the published output.

As noted before, EDP is proposed by Abowd et al. (2013) as an alternative to DP due to the difficulties with implementing DP in practice. The EDP procedure is easily applicable to Bayesian modelling of a sensitive dataset, and requires less technical work from the part of the user. We describe it below.

Empirical Differential Privacy (EDP): Let D be the original dataset, and D^{-i} be the same dataset without the i^{th} observation. Consider a parameter θ and a discretization of its posterior distribution, meaning bins $b_1, \ldots, b_B$ such that $P(\theta \in b_j|D) = 1/B, j = 1, \ldots, B$. Then, define

$$M_1 = \max_{i,j} \frac{P(\theta \in b_j|D^{-i})}{P(\theta \in b_j|D)} \qquad M_2 = \min_{i,j} \frac{P(\theta \in b_j|D^{-i})}{P(\theta \in b_j|D)} \tag{2}$$

Finally, compute $\epsilon = \max(|\log(M_1)|, |\log(M_2)|)$. This value is the level of EDP.

Note that the original definition actually considers several parameters simultaneously, and makes the implicit assumption that the domains of the posterior distributions of all these parameters are the same. For this paper, we will limit ourselves to a unique parameter of interest.

In practice, the EDP procedure is to compute the probabilities required to calculate M_1 and M_2 are estimated using samples from the appropriate posterior predictive distributions. Here are the details:

1. Obtain M draws from the posterior distribution of $\theta|D$.
   These are denoted $\theta_{Dm}, m = 1, \ldots, M$.

2. Compute the B-quantiles from these draws so that we obtain bins $b_1, \ldots, b_B$ such that
   $$\sum_{m=1}^{M} I(\theta_{Dm} \in b_j) = \sum_{m=1}^{M} I(\theta_{Dm} \in b_k) = \frac{M}{B} \quad \forall j, k \in \{1, \ldots, B\}.$$
   An estimate for $P(\theta \in b_j|D)$ is then $1/B$.

3. For each observation i, obtain M draws from the posterior distribution of $\theta|D^{-i}$.
   These are denoted $\theta_{D^{-i}m}, \;\; i = 1, \ldots n, \; m = 1, \ldots, M$.

4. Estimate $P(\theta \in b_j|D^{-i})$ as $\sum_{i=1}^{M} I(\theta_{D^{-i}m} \in b_j)/M$.

Note that the analyst must pick values for M and B; the example in Abowd et al. (2013) uses $M = 10,000$ and $B = 20$. We will discuss these choices in sections 3 and 4.

The proposed usage of EDP is to start with some prior for θ and compute the associated EDP-ϵ. Then, one may modify the prior for θ to reduce or increase ϵ, as desired. A more informative prior would reduce the impact of any single observation on the posterior distribution of θ, and thus lead to a higher degree of privacy as measured by EDP. Inversely,a less informative prior would lead to a smaller value of ϵ as measured by EDP. Through trial and error, one may obtain an output with the desired level of EDP.

Note that choosing the prior distribution in such a data-driven way goes against the philosophy of DP, and may leak some information about the dataset. In practice however, it may be a small enough risk for data providers to take.

3 The Meaning of EDP

EDP has been presented as an empirical version of DP. The term EDP itself is justified on page 89 of Abowd et al. (2013): "Because this is done after model estimation instead of before, we use the term 'empirical differential privacy'." EDP is presented throughout Abowd et al. (2013) as a convenient mechanism to estimate the DP protection of a Bayesian statistical model, with the value obtained from the EDP algorithm being

interpreted as an approximate value of differential privacy for releasing a sample from the posterior distribution. Careful analysis of the proposed methodology however reveals that EDP is very different from the original DP definition.

3.1 The Implicit Conditioning of EDP

A first important characteristic of the EDP procedure is its implicit conditioning on the observed dataset. DP is well-known to be a worst-case scenario, designed to protect any possible dataset, even one which is very unlikely to occur in practice. However, in EDP, we only consider the datasets $D, D^{-1}, \ldots, D^{-n}$, and are thus implicitly conditioning on the observed dataset D when computing the empirical ϵ.

To understand the impact of this conditioning, suppose that we define a conditional version of DP as follows:

Conditional Differential Privacy (CDP): Let D be the original dataset, and D^{-i} be the same dataset without the i^{th} observation. Then, a randomized function κ is said to satisfy ϵ-CDP for dataset D if $\forall i = 1, \ldots, n$

$$e^{-\epsilon} \leq \frac{Pr[\kappa(D^{-i}) \in S]}{Pr[\kappa(D) \in S]} \leq e^{\epsilon}.$$

We now illustrate the impact of this modification of DP in a simple setting.

Consider a dataset $X = (x_1, \ldots, x_n)$, where $x_i \in \{0, 1\}$ for $i = 1, \ldots, n$ are dichotomous variables. To release $x = \sum_{i=1}^{n} x_i$ while protecting the confidentiality of the respondents, one can publish an ϵ differentially-private version $\tilde{x}$ of this statistic. A possible, but not optimal, way to do so is to sample

$$\tilde{p} \sim \text{Beta}(\alpha_1 + x, \alpha_2 + n - x)$$
$$\tilde{x} \sim \text{Binomial}(n, \tilde{p})$$

with $\alpha_1 = \alpha_2 = (\exp(\epsilon) - 1)/n$ (Abowd and Vilhuber, 2008). This randomized function can be represented by an $n + 1$ by $n + 1$ transition matrix, as shown in Table 1 for $n = \tilde{n} = 5$ and $\alpha_1 = \alpha_2 = 0.5$, so that ϵ is $\log(11) \approx 2.397$.

One can verify the value of ϵ by computing the log ratios of posterior probabilities for neighbouring datasets. Indeed, the largest ratio is

$$\frac{P(\tilde{x} = 0|x = 4)}{P(\tilde{x} = 0|x = 5)} = \frac{0.010742}{0.000977} = 10.9948 \approx 11$$

Table 1 also allows to compute the CDP-ϵ for all possible datasets. The DP and CDP values are indeed equal at 2.4 if $x = 0, 1, 4$ or 5. However, the CDP-ϵ is only 1.466 if $x = 2$ or $x = 3$, and is thus smaller than the DP ϵ of 2.4. The smaller value

Table 1: **Transition probabilities for the beta-binomial synthesizer with $n = 5$ and $\alpha_1 = \alpha_2 = 0.5$, so that ϵ is $\log(11) \approx 2.397$.**

	$\tilde{x} = 0$	$\tilde{x} = 1$	$\tilde{x} = 2$	$\tilde{x} = 3$	$\tilde{x} = 4$	$\tilde{x} = 5$
x=0	0.715975	0.188415	0.066499	0.022166	0.005968	0.000977
x=1	0.339146	0.299247	0.199498	0.107422	0.043945	0.010742
x=2	0.139648	0.232747	0.250651	0.205078	0.125326	0.046549
x=3	0.046549	0.125326	0.205078	0.250651	0.232747	0.139648
x=4	0.010742	0.043945	0.107422	0.199498	0.299247	0.339146
x=5	0.000977	0.005968	0.022166	0.066499	0.188415	0.715975

is caused by the fact that the impact of the change of one observation on the output is less important for less extreme datasets. By conditioning on an observed dataset when computing ϵ, one thus risks overestimating the DP level of privacy.

For n larger than 5, many more CDP values are possible. Figure 1 shows dotplots of all possible CDP-ϵ for various sample sizes when the DP-ϵ is set to 2.4 or 5, with superimposed boxplots and means. As expected, the maximum CDP-ϵ is the DP-ϵ. Also note that the distributions of the CDP-ϵ are very skewed to the right. The risk of underestimating the DP ϵ by conditioning on the observed dataset D is thus large. For example, if $n = 200$ and $\epsilon = 5$ half of the possible datasets will lead to an estimate smaller than 1.60. This is more than 3 times smaller than the true DP-ϵ, corresponding to an estimated ratio of probabilities 30 times smaller than its true value.

3.2 EDP as a Measure of Sensitivity

Another crucial difference between DP and EDP is that while DP is a property of a randomized function, EDP is a property of a deterministic output. Indeed, EDP is concerned with a specific output, namely the posterior distribution of a parameter θ, and not a randomized function. In fact, it would be more appropriate to interpret EDP as a measure of sensitivity.

Consider a randomized function κ. Following the usual notation, the global sensitivity (Dwork et al., 2006b) of that function is defined as

$$GS_\kappa = \max_{D,D'} d(\kappa(D), \kappa(D')) \tag{3}$$

where D, D' are neighboring datasets, and d is a distance measure between outputs of κ, often the Euclidean distance. Similarly, the local sensitivity (Nissim et al., 2007) for a specific dataset D is given by

$$LS_\kappa = \max_{D'} d(\kappa(D), \kappa(D')) \tag{4}$$

EDP thus resembles a local sensitivity measure where the distance d between pos-

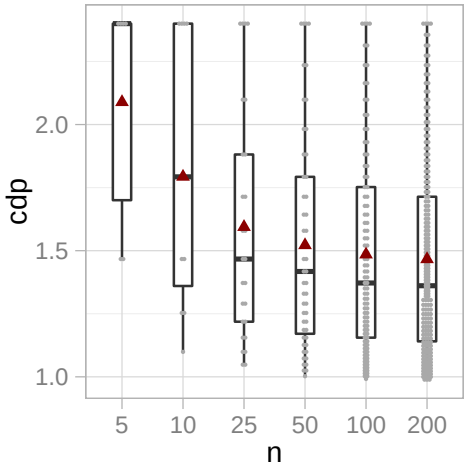

(a) DP-ϵ set to 2.4

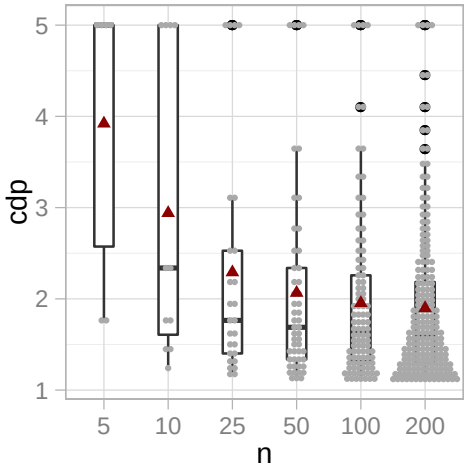

(b) DP-ϵ set to 5

Figure 1: Dotplots of all possible CDP-ϵ for various sample sizes when the DP ϵ is set to 2.4 or 5, with superimposed boxplots. Triangles indicate mean values. As expected, the maximum CDP-ϵ is the DP-ϵ. The distributions of the CDP-ϵ are very skewed. Except for a few extreme datasets, a conditional estimate of ϵ may thus be far from the true DP-ϵ.

terior distributions is taken to be $\max(|\log(M_1)|, \log(|M_2)|)$, with M_1 and M_2 as defined in equation (2). A slight difference is that in EDP neighbours of D are datasets

$D^{-i}, i = 1, \ldots, n$, i.e. datasets with one observation removed from D, whereas in equation (4) neighbours consist of datasets of size n with all but one records identical to those of D. We could thus think of EDP a one-sided measure of local sensitivity of the posterior distribution of a Bayesian model.

Measures such as global and local sensitivity are used to calculate how much noise is needed in a randomized function to satisfy ϵ-DP, but do not themselves provide us with an estimate of the DP protection offered by a randomization function. This is a second reason why EDP can not be interpreted as an empirical estimate of DP. [1]

3.3 The Impact of the Posterior Discretization

Now, even if we consider EDP as a sensitivity measure and not an empirical estimate of DP, its interpretation is problematic because it is not a well-defined measure. Indeed, the value of the EDP-ϵ depends on the number of bins B used in the discretization of the posterior distribution of the parameter of interest θ. The value $B = 20$ was used without justification in Abowd et al. (2013), but a different choice would have lead to a different estimate of ϵ. We illustrate the impact of B on two common conjugate models.

Beta-Binomial conjugate model

Consider data $X \sim$ Binomial(n,p). Suppose that we do not want to publish a private version of the observed data, but rather a sample from the posterior distribution of p. Given a prior distribution Beta(α_1, α_2) for p, the posterior distribution for p given $X = x$ is Beta$(\alpha_1 + x, \alpha_2 + n - x)$. For fixed values of n,x,α_1 and α_2, one may compute the EDP-ϵ by sampling from the posterior distribution of p, as described in Section 2. For this simple model however, exact values of M_1 and M_2 can also be calculated using the known posterior distribution, without having to sample from it. From these, one can easily obtain an exact value of the EDP-ϵ for a fixed B.

[1] A referee noted that the ϵ in DP can also be interpreted as a sensitivity measure, namely the global sensitivity of the conditional probability distribution defined on the output space. In that sense, both DP and EDP can be seen as sensitivity measures, but they are still not measuring the same quantity.

The exact EDP-ϵ are given in Table 2 for all possible true datasets and various B in the case where $n = 5, \alpha_1 = \alpha_2 = 0.5$. The EDP level clearly depends on the choice of B: as the number of bins increases, the associated EDP-ϵ increases. The difference in posterior distributions between two neighbouring datasets is more easily detected with finer bins.

Table 2: **True EDP-ϵ for all possible datasets for $n = 5, \alpha_1 = \alpha_2 = 0.5$ and various values of B.**

	$B = 5$	$B = 10$	$B = 20$	$B = 50$	$B = 100$	$B = 250$
x=0	4.01	5.41	6.81	8.64	10.03	11.87
x=1	2.16	2.49	2.97	3.61	4.08	4.70
x=2	1.61	1.97	2.31	2.74	3.05	3.44
x=3	1.61	1.97	2.31	2.74	3.04	3.44
x=4	2.16	2.49	2.97	3.61	4.08	4.70
x=5	4.01	5.41	6.81	8.64	10.03	11.87

One might be tempted to compare the results from Table 2 to the value $\log(11) \approx 2.4$, the DP-ϵ for the Beta-Binomial synthesizer when $n = 5$ and $\alpha_1 = \alpha_2 = 0.5$, to see how well EDP estimates DP. However, recall that both ϵ measure different things: the DP-ϵ corresponds to the transition matrix in Table 1, where we publish a confidential version of the original dataset, whereas the EDP-ϵ measures the sensitivity of the posterior distribution for p to a change in the true dataset.

Normal-Normal conjugate model

Let $D = \{y_1, \ldots, y_n\}$ be the observed dataset, and suppose a normal distribution with unknown mean μ and known variance σ^2 for y_i, $i = 1, \ldots, n$. Suppose a $N(\mu_0, \sigma_0^2)$ prior for μ. The posterior distribution for μ given D is then a normal distribution with parameters

$$\mu_1 = \frac{\sigma^2 \mu_0 + \sigma_0^2 \sum_{i=1}^n y_i}{\sigma^2 + n\sigma_0^2} \text{ and } \sigma_1^2 = \frac{\sigma_0^2 \sigma^2}{n\sigma_0^2 + \sigma^2}.$$

Once again, the posterior distribution of μ given D and D^{-i} is known,. To compute the exact EDP-ϵ for a given value of B, we simply calculate the B quantiles of the posterior distribution for μ given the real dataset D then calculate the probability associated with each bin under the different posterior distributions $\mu|D^{-i}$, $i = 1, \ldots, n$ to get M_1 and M_2, and infer the EDP-ϵ.

To test the impact of B, 100 datasets D were generated from a normal distribution with $\mu = 10$ and $\sigma^2 = 5$, for each of $n = 100$, $n = 1000$ and $n = 10,000$. The exact EDP-ϵ is was calculated with $\mu_0 = 10$ and $\sigma_0^2 = 3$ and $B \in \{10, 20, 50, 100, 250\}$ for each dataset. Figure 2 shows the distributions of the estimates in each case. Again, the EDP value increases with B. It also decreases with n, as expected since each observation

then has a smaller impact on the posterior distributions.

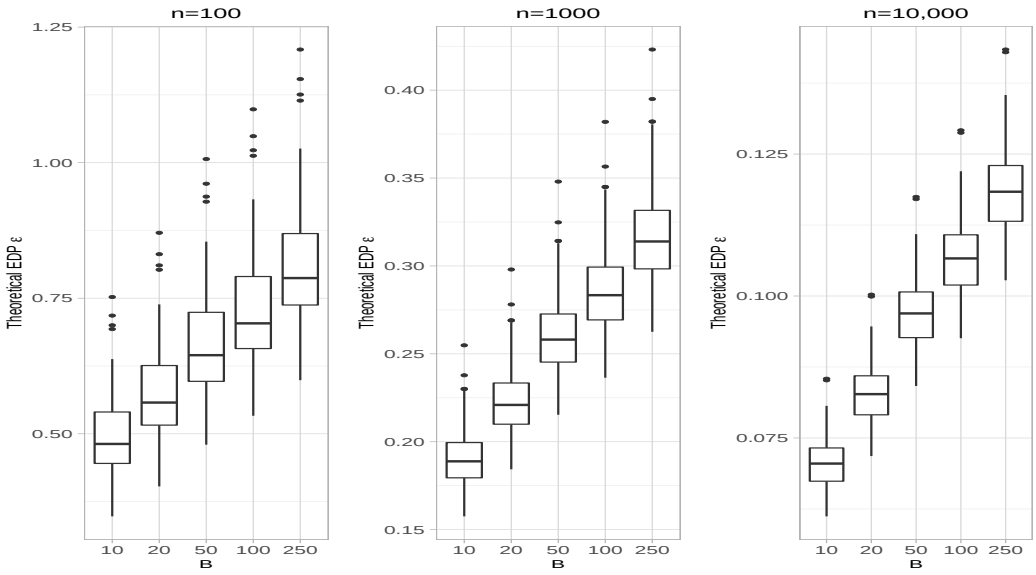

Figure 2: True EDP-ϵ for 100 datasets generated from a $N(10, 5)$ under the Normal-Normal model with prior $N(10, 3)$ for different sample sizes and discretizations. Again the value of ϵ increases with B. Note also that the EDP-ϵ decreases with n, as expected since each observation then has a smaller impact on the posterior distribution.

It is thus clear that the EDP-ϵ depends crucially on the discretization of the posterior distribution used in the procedure. Given the impact of B on the EDP value, we will for now on refer to the EDP ϵ for a given B as the B-EDP-ϵ.

4 EDP in Practice

In practice, EDP can be applied to more complicated Bayesian models, for which the posterior distribution may not have a closed form, but can be sampled from with an MCMC algorithm, as in Abowd et al. (2013), Schneider and Abowd (2015) and Schneider and Abowd (2015). The value of B-EDP-ϵ can then not be calculated exactly, but must then estimated using samples of size M from the posterior distributions of θ.

Using the Beta-Binomial and the Normal-Normal models, for which we can calculate the exact value of any B-EDP-ϵ, we now study the impact of the choice of M on the accuracy of the estimation of the B-EDP-ϵ.

Recall that Abowd et al. (2013) used $M = 10,000$ draws from the posterior distribution. That choice was justified by computing an empirical measure of the imprecision of the estimate from the EDP procedure. This measure is obtained by applying the

EDP procedure but using the original dataset D instead of D^{-i}, $i = 1, \ldots, n$. Since the posterior distributions are then equal, one should obtain $M_1 = M_2 = 1$, and thus zero as the estimate for ϵ. The authors obtained a value of 0.27, while their estimated EDP ϵ with B=20 was 3.27, and so the choice of $M = 10,000$ was deemed sufficient. We will show here that this procedure may be misleading to pick an appropriate value for M, and that larger values of M may be needed to estimate properly the B-EDP-ϵ.

Note that in practice, some of the bins created by the discretization of the posterior distribution $\theta|D$ may not contain any draw from $\theta|D^{-i}$, for some or several $i = 1, \ldots, n$. This potential problem was not discussed in the original EDP paper, but it happened frequently in our simulations, particularly for small M, large B and/or small values of B-EDP-ϵ. Simply using the maximum likelihood estimate of the probabilities for the bins would then lead to infinite estimates for B-EDP-ϵ. To alleviate this problem, we used additive smoothing (Johnson, 1932) with $\alpha = 0.01$ to obtain non-zero probability estimates for all the bins. The choice of α was seen to have little impact on the results in our setting.

Beta-Binomial model

Figure 3 shows estimates of EDP ϵ from 100 replications for the Beta-Binomial model for various values of B and M, still with $n = 5, \alpha_1 = \alpha_2 = 0.5$ and $x = 0$. Exact values, corresponding to the first row of Table 2, are illustrated with horizontal red dotted lines.

The percent relative bias of the estimate in each case is given in Table 3, under the heading RB(%). Table 3 also gives an empirical measure of imprecision of the EDP estimate. This measure is the average of the empirical measure computed as described above for 100 different samples from the posterior distribution given the original dataset D.

Table 3: Relative Bias (RB) (in %) and Empirical Measure of the Imprecision (EMI) of the estimate from the EDP procedure for the Beta-Binomial model for various values of M and B. Values are averages over 100 replications.

	$M = 10,000$		$M = 100,000$		$M = 1,000,000$	
	RB(%)	EMI	RB(%)	EMI	RB(%)	EMI
B=5	-0.05	0.039	-0.21	0.013	0.02	0.004
B=10	1.75	0.081	0.01	0.025	-0.02	0.008
B=20	-18.97	0.132	-0.85	0.041	0.22	0.013
B=50	-40.21	0.251	-13.79	0.079	2.69	0.025
B=100	-54.24	0.400	-31.14	0.122	19.07	0.039
B=250	-68.92	0.788	-49.51	0.221	8.67	0.069

It is clear that more than $M = 10,000$ draws are necessary to correctly estimate the value of B-EDP-ϵ for B larger than 10. More importantly, small values of the

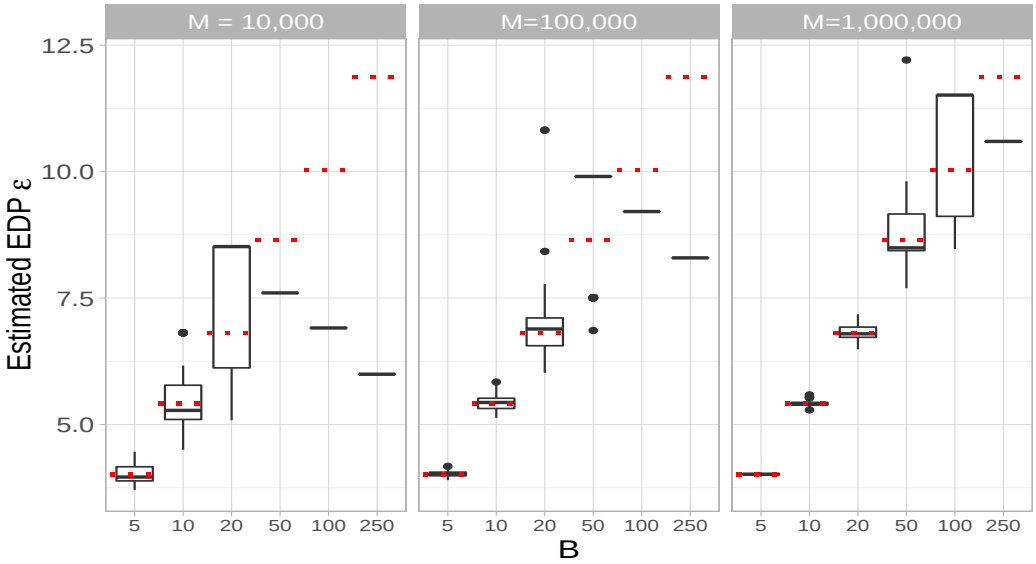

Figure 3: Estimates of EDP ϵ from 100 replications for the Beta-Binomial model with $n = 50, \alpha_1 = \alpha_2 = 0.5$ and $x = 0$, for various values of B and M. Real values, corresponding to the first row of Table 2, are illustrated with horizontal red dotted lines. With $M = 10,000$ posterior draws, only ϵ for $B = 5$ is relatively well-estimated. Increasing M improves the estimates, but even $M = 1,000,000$ is not sufficient for more than 20 bins.

empirical measure of imprecision do not always correspond to precise estimates of the EDP. Notice in particular that when $M = 1,000,000$ very small values are obtained even for $B = 100$ and $B = 250$, especially compared to the true values of 10.03 and 11.87, while the B-EDP-ϵ is not well estimated.

Similar results for larger values of n and different choices of x were obtained and are not shown here. We note however that with larger n, the true B-EDP-ϵ tends to be smaller, and thus a larger M is required for accurate estimation.

Normal-Normal model

Corresponding results for the Normal-Normal model are given in Figure 4 and Table 4. The same conclusions hold in this case, but an even larger M is required to get a relatively accurate of the B-EDP-ϵ.

Note that for large B the estimated relative bias for small M depends a lot on the details of the smoothing procedure. It is however clear that the bias of the estimate of the B-EDP-ϵ is very large. Also note that in some extreme cases the empirical measure of imprecision does indeed indicate that the value of M is not sufficient to estimate the B-EDP-ϵ, for example when $M = 10,000$ and $B = 100$ or $B = 250$.

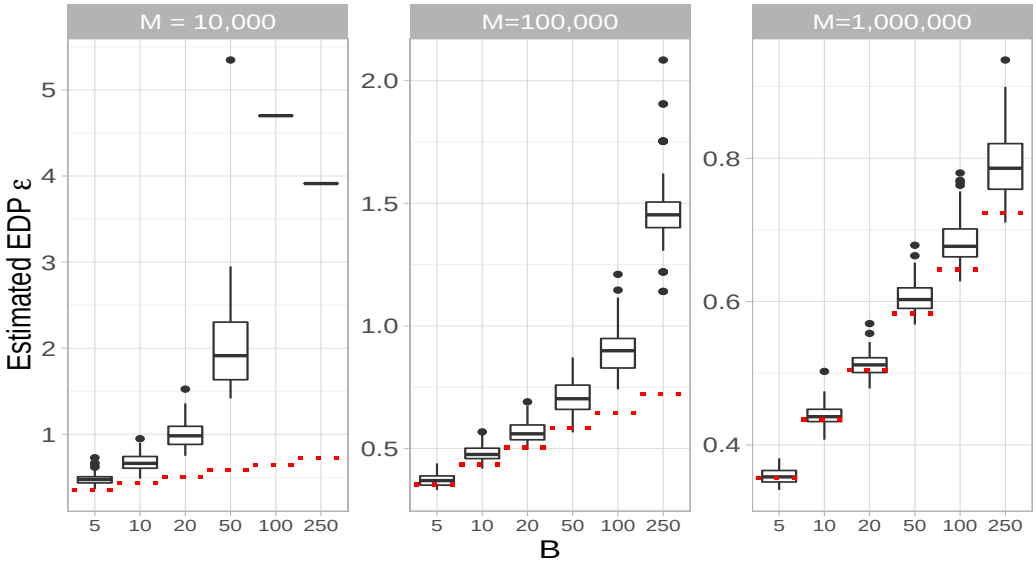

Figure 4: Estimates of B-EDP-ϵ from 100 replications for the Normal-Normal model with $\mu = 10$, $\sigma^2 = 5$, $\mu_0 = 10$, $\sigma_0^2 = 3$ and $n = 100$, for various values of B and M. Real values are illustrated with horizontal red dotted lines. Even for $B = 5$, estimates of the true B-EDP-ϵ are poor with $M = 100,000$.

Table 4: Relative Bias (RB) (in %) and Empirical Measure of the Imprecision (EMI) of the estimate from the EDP procedure for the Normal-Normal model for various values of M and B. Values are averages over 100 replications.

	$M = 10,000$		$M = 100,000$		$M = 1,000,000$	
	RB(%)	EMI	RB(%)	EMI	RB(%)	EMI
B=5	36.96	0.136	5.16	0.041	0.62	0.013
B=10	56.62	0.256	10.47	0.080	1.42	0.025
B=20	94.50	0.432	13.12	0.132	1.53	0.043
B=50	238.47	0.963	21.89	0.258	4.02	0.078
B=100	628.97	2.19	39.16	0.041	6.08	0.790
B=250	440.59	5.99	104.20	0.790	9.51	0.216

5 Conclusion

EDP has been proposed as a relaxation of DP for applications to Bayesian modeling of confidential data. Its simplicity of use makes it very attractive for privacy protection in practice. While it has been applied to Bayesian linear mixed models in Abowd et al. (2013), Bayesian generalized linear mixed models in Schneider and Abowd (2015) and

zero-inflated Poisson models in Schneider and Abowd (2015), this paper provides a first investigation of the meaning and limits of this relaxation of DP.

We showed in Section 3 that EDP can not simply be considered an empirical estimate of DP, in part because of the implicit conditioning in EDP, and that it should instead be considered a sensitivity measure on posterior distributions. Moreover, even when treated as such, EDP is not a well-defined criterion since its true value depends crucially on the number of bins used in the discretization part of the procedure, smaller bins leading to higher sensitivity measures.

By using simpler models for which the exact B-EDP-ϵ is known, we were also able to study the practical difficulty of estimating the EDP level from draws of the posterior distributions. We showed that to estimate accurately the exact B-EDP-ϵ many more random draws from posterior distributions are necessary than was previously thought. And note that both of our examples were conjugate models; even more draws may be needed when the posterior distribution itself is estimated using an MCMC algorithm. In addition, our results indicate that the tool proposed (Abowd et al., 2013) to choose the number of draws is not very effective.

Our results raise several questions about the future of EDP for privacy protection. We made it clear that EDP does not actually estimate DP, and should not be used as such for privacy protection. One may still want to use EDP as a measure of sensitivity with the goal of designing a randomized function κ which outputs a differentially-private posterior distribution. EDP could also be used directly as a sensitivity measure for statistical disclosure limitation, but it would need to be interpreted very differently than DP.

However, the dependence of the estimated value on the number of bins is problematic in any case, as a specific number of bins B will need to be chosen no matter how EDP is to be used. We know that larger B lead to larger measures of sensitivity, but more work will be needed to understand the meaning and consequences of any particular choice for B. The practical difficulties of estimating the B-EDP-ϵ also complicate the application of EDP in practice. If EDP is to continue to be used in practice, researchers will need to design better ways of estimating the probabilities involved in calculating ϵ. Since its value often depends crucially on the bins with smallest probabilities, importance sampling methods designed to estimate tail probabilities may be of use.

References

Abowd, J. M., Schneider, M. J., and Vilhuber, L. (2013). "Differential privacy applications to Bayesian and linear mixed model estimation." *Journal of Privacy and Confidentiality*, 5(1): 4.

Abowd, J. M. and Vilhuber, L. (2008). "How protective are synthetic data?" In *Privacy in Statistical Databases*, 239–246. Springer.

Chaudhuri, K. and Monteleoni, C. (2009). "Privacy-preserving logistic regression." In *Advances in Neural Information Processing Systems*, 289–296.

Dankar, F. K. and El Emam, K. (2012). "The application of differential privacy to health data." In *Proceedings of the 2012 Joint EDBT/ICDT Workshops*, 158–166. ACM.

Dwork, C. (2006). "Differential privacy." In *Automata, languages and programming*, 1–12. Springer.

Dwork, C., Kenthapadi, K., McSherry, F., Mironov, I., and Naor, M. (2006a). "Our data, ourselves: Privacy via distributed noise generation." In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, 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*, 265–284. Springer.

Dwork, C. and Smith, A. (2010). "Differential privacy for statistics: What we know and what we want to learn." *Journal of Privacy and Confidentiality*, 1(2): 2.

Hall, R., Rinaldo, A., and Wasserman, L. (2011). "Random differential privacy." *arXiv preprint arXiv:1112.2680*.

Inan, A., Kantarcioglu, M., Ghinita, G., and Bertino, E. (2010). "Private record matching using differential privacy." In *Proceedings of the 13th International Conference on Extending Database Technology*, 123–134. ACM.

Johnson, W. E. (1932). "I. Probability : The Deductive and Inductive Problems." *Mind*, 41(164): 409–423.

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

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*, 75–84. ACM.

Schneider, M. J. and Abowd, J. M. (2015). "A new method for protecting interrelated time series with Bayesian prior distributions and synthetic data." *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 178(4): 963–975.

Smith, A. (2011). "Privacy-preserving statistical estimation with optimal convergence rates." In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, 813–822. ACM.