

Bayesian Pseudo Posterior Mechanism under Asymptotic Differential Privacy

Terrance D. Savitsky

*Office of Survey Methods Research
U.S. Bureau of Labor Statistics
2 Massachusetts Ave NE
Washington, DC 20212, USA*

SAVITSKY.TERRANCE@BLS.GOV

Matthew R. Williams

*National Center for Science and Engineering Statistics
National Science Foundation
2415 Eisenhower Ave
Alexandria, VA 22314, USA*

MRWILLIA@NSF.GOV

Jingchen Hu

*Vassar College
124 Raymond Ave, Box 27
Poughkeepsie, NY 12604, USA*

JIHU@VASSAR.EDU

Editor: Sayan Mukherjee

Abstract

We propose a Bayesian pseudo posterior mechanism to generate record-level synthetic databases equipped with an (ϵ, π) -probabilistic differential privacy (pDP) guarantee, where π denotes the probability that any observed database exceeds ϵ . The pseudo posterior mechanism employs a data record-indexed, risk-based weight vector with weight values $\in [0, 1]$ that surgically downweight the likelihood contributions for high-risk records for model estimation and the generation of record-level synthetic data for public release. The pseudo posterior synthesizer constructs a weight for each datum record by using the Lipschitz bound for that record under a log-pseudo likelihood utility function that generalizes the exponential mechanism (EM) used to construct a formally private data generating mechanism. By selecting weights to remove likelihood contributions with non-finite log-likelihood values, we guarantee a finite local privacy guarantee for our pseudo posterior mechanism at every sample size. Our results may be applied to *any* synthesizing model envisioned by the data disseminator in a computationally tractable way that only involves estimation of a pseudo posterior distribution for parameters, θ , unlike recent approaches that use naturally-bounded utility functions implemented through the EM. We specify conditions that guarantee the asymptotic contraction of π to 0 over the space of databases, such that the form of the guarantee provided by our method is asymptotic. We illustrate our pseudo posterior mechanism on the sensitive family income variable from the Consumer Expenditure Surveys database published by the U.S. Bureau of Labor Statistics. We show that utility is better preserved in the synthetic data for our pseudo posterior mechanism as compared to the EM, both estimated using the same non-private synthesizer, due to our use of targeted downweighting.

Keywords: differential privacy, pseudo posterior distribution, pseudo posterior mechanism, synthetic data

1. Introduction

Privacy protection is an important research topic, which attracts attention from government statistical agencies and private companies alike. A commonly-used data privacy approach generates synthetic data from statistical models estimated on closely-held, private data for proposed release by statistical agencies (Rubin, 1993) and (Little, 1993). This approach replaces the closely-held (by the statistical agency) database with multiple synthetically generated record-level databases. The synthetic databases are released to the public who would use them to conduct any analyses of which they would conceive to be conducted on the real, confidential record-level data. The synthetic data approach replaces multiple queries performed on a summary statistic with the publication of the synthetic databases encoded with privacy protection, making this approach independent of the specific queries performed by users or putative intruders.

1.1 Differential Privacy

Our focus metric for measuring the relative privacy guarantee of our pseudo posterior synthesizing data mechanism introduced in the sequel is differential privacy (Dwork et al., 2006). We next provide a definition for differential privacy (McSherry and Talwar, 2007).

Definition 1 (Differential Privacy) *Let \mathbf{x} be a database in input space \mathcal{X}^n , where \mathcal{X}^n denotes a space of databases of size (number of observations) n . Let \mathcal{M} be a randomized mechanism such that $\mathcal{M}() : \mathcal{X}^n \rightarrow \mathcal{O}$. Then \mathcal{M} is ϵ -differentially private if*

$$\frac{Pr[\mathcal{M}(\mathbf{x}) \in O]}{Pr[\mathcal{M}(\mathbf{y}) \in O]} \leq \exp(\epsilon),$$

for all possible outputs $O = \text{Range}(\mathcal{M})$ under all possible pairs of data sets $\mathbf{x} \in \mathcal{X}^n$ where $\mathbf{y} \in \mathcal{X}^{n-1}$ differs from \mathbf{x} by deleting one record or datum (under a leave-one-out (LOO) distance definition).

Differential privacy is a property of the mechanism or data generating process and a mechanism that meets the definition above is guaranteed to be ϵ -differentially private, or ϵ -DP. Differential privacy is called a “formal” privacy guarantee because the ϵ -level or guarantee is independent of the behavior of a putative intruder seeking to re-identify the data and the guarantee is not lessened by the existence of other data sources that may contain information about the same respondents included in \mathcal{X}^n .

Differential privacy assigns a disclosure risk for a statistic to be released to the public, $f(\mathbf{x})$ (e.g., total employment for a state-industry) of any $\mathbf{x} \in \mathcal{X}^{n-1}$ based on the global sensitivity, $\Delta = \sup_{\mathbf{x} \in \mathcal{X}^n, \mathbf{y} \in \mathcal{X}^{n-1}: \delta(\mathbf{x}, \mathbf{y})=1} |f(\mathbf{x}) - f(\mathbf{y})|$, over the space of databases, \mathcal{X} , where $\delta(\mathbf{x}, \mathbf{y})$ denotes the number of records omitted from \mathbf{x} in database, \mathbf{y} . The distance metric, $\delta(\mathbf{x}, \mathbf{y})$ denotes the LOO distance such that \mathbf{x} differs from \mathbf{y} by a single record, which is equivalent to using a Hamming-1 distance in the case of count based statistics of binary data records. If the value of the statistic, f , expresses a high magnitude change after the deletion of a data record in \mathbf{y} , then the mechanism will be required to induce a relatively higher level of distortion to f . The more sensitive is a statistic to the change of a record, the higher its disclosure risk.

Our focus in this paper is where the mechanism, \mathcal{M} , is a model parameterized by θ from which replicate data are synthesized under an ϵ -DP guarantee. A common approach for generating parameter draws of θ under the statistical model for synthesizing data is the exponential mechanism (EM) of McSherry and Talwar (2007), which inputs a non-private mechanism for θ and generates θ in such a way that induces an ϵ -DP guarantee on the overall mechanism. The EM is conditioned on the availability of a global sensitivity over the space of databases, Δ_u for some utility function, $u(\mathbf{x}, \theta)$, defined on the space of databases and the space of parameters, globally.

Definition 2 (*Exponential Mechanism*) *The exponential mechanism releases values of θ from a distribution proportional to,*

$$\exp(u(\mathbf{x}, \theta)), \tag{1}$$

where $u(\mathbf{x}, \theta)$ is a utility function. Let

$\Delta_u = \sup_{\mathbf{x} \in \mathcal{X}^n} \sup_{\mathbf{y}: \delta(\mathbf{x}, \mathbf{y})=1} \sup_{\theta \in \Theta} |u(\mathbf{x}, \theta) - u(\mathbf{y}, \theta)|$ be the sensitivity, defined globally over $\mathbf{x} = (x_1, \dots, x_n) \in \mathcal{X}^n$, the σ -algebra of data sets, \mathbf{x} , governed by product measure, P_{θ_0} and the LOO distance metric, $\delta(\mathbf{x}, \mathbf{y}) = 1$. Then each draw of θ from the exponential mechanism is guaranteed to be $\epsilon = 2\Delta_u$ -DP.

This result is based on the following definition of differential privacy under utility function, $u(\mathbf{x}, \theta)$.

Definition 3 (*Differential Privacy under the Exponential Mechanism*) *A utility function, u , indexed by random parameters, θ , gives ϵ -differential privacy if for all databases, $\mathbf{x} \in \mathcal{X}^n$ and associated databases, $\mathbf{y} : \delta(\mathbf{x}, \mathbf{y}) = 1$, and all parameter values, $\theta \in \Theta$,*

$$Pr(u(\mathbf{x}, \theta) \in O) \leq \exp(\epsilon) \times Pr(u(\mathbf{y}, \theta) \in O), \tag{2}$$

where $O = \text{range}(u)$.

In order to set an arbitrary $\epsilon \neq 2\Delta_u$, we must modify the utility function $u(\mathbf{x}, \theta)$. The statistical agency owning the closely-held data will typically desire to determine ϵ as a matter of policy and not leave it to be $\epsilon = 2\Delta_u$. The simplest and most common approach is to rescale it: $u^*(\mathbf{x}, \theta) = \frac{\epsilon}{2\Delta_u} u(\mathbf{x}, \theta)$ (See McSherry and Talwar, 2007; Dwork et al., 2006, among many others).

The EM inputs a utility function and its sensitivity constructed as the supremum of the utility over the space of databases, \mathcal{X}^n , and simultaneously, the parameter space, Θ . Wasserman and Zhou (2010) and Snoke and Slavkovic (2018) construct utility functions (e.g., the Kolmogorov-Smirnov distance between the empirical distributions of the real and synthetic data sets) that are *naturally* bounded over all $\mathbf{x} \in \mathcal{X}^n$, resolving the challenge of using the potentially unbounded log-likelihood as the utility function. Although the use of a naturally bounded utility resolves the issue of truncating the data and parameter spaces, there is a *large*, and perhaps *intractable*, computational cost to the use of these naturally bounded utilities to draw samples of θ from the distribution constructed from the EM; for example, Snoke and Slavkovic (2018) must compute their *pMSE* utility statistic multiple times for each proposed value, $\hat{\theta}_l$ ($l = 1, \dots, L$), under a Metropolis-Hastings algorithm used

to draw samples under the EM. Furthermore, they assume the existence of some synthesizing distribution, $g(\hat{\theta})$, from which to draw synthetic data, is needed to compute their $pMSE$. In practice, g will be defined as the posterior predictive distribution, $g(\mathbf{X} \mid \mathbf{x}, \theta_l)$, which means the posterior distribution must be repeatedly estimated for *each* draw from of θ from the EM.

Dimitrakakis et al. (2017) utilize the log-likelihood as the utility under the EM such that the EM reduces to the the model posterior distribution, $\xi(\theta \mid \mathbf{x})$, as the mechanism, \mathcal{M} . They specify conditions under which the Bayesian posterior distribution is a formally private mechanism for generating synthetic data. They construct their posterior distribution from the private data, \mathbf{x} , and parameters, θ , used to generate the synthetic data. They show that if the log-likelihood is Lipschitz continuous with bound Δ over the space of databases, $\mathbf{x} \in \mathcal{X}^n$ (the space of databases of size, n) and the space of parameters, $\theta \in \Theta$, then the posterior mechanism achieves an $\epsilon = 2\Delta$ -DP guarantee for each posterior draw of θ , the model parameter(s). However, Dimitrakakis et al. (2017) acknowledge that computing a finite Δ , in practice, under the use of the log-likelihood is particularly difficult for an unbounded parameter space. They specify relatively simple Bayesian probability models where the Lipschitz bound is analytically available. Even in such simple model setting, Dimitrakakis et al. (2017) require truncation of the support of the prior distribution to achieve a finite Δ . Relatively simply-constructed differentially private Bayesian synthesizers are similarly proposed by Machanavajjhala et al. (2008); Abowd and Vilhuber (2008); McClure and Reiter (2012) and Bowen and Liu (2020). The utility performance to preserve the real data distribution in the simulated synthetic data of these simple posterior mechanisms under a truncated prior support may be severely compromised by truncation and over-smoothing (induced by simple, parametric prior distributions).

Hu et al. (2021) design a record-indexed weight $\alpha_i \in [0, 1]$, which is set to be inversely proportional to their construction for the identification risk probability of record, i ; a data record that expresses a relatively high probability of identification disclosure will receive a likelihood weight, α_i , that is closer to 0, while a data record with a low disclosure probability will receive a likelihood weight, α_i , that is closer to 1. The vector weights $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_n)$ are subsequently applied to the likelihood function of all n records to form the pseudo posterior,

$$\xi^{\boldsymbol{\alpha}}(\theta \mid \mathbf{x}, \gamma) \propto \left[\prod_{i=1}^n p(x_i \mid \theta)^{\alpha_i} \right] \xi(\theta \mid \gamma), \quad (3)$$

where θ denotes the model parameters, γ denotes the model hyperparameters and $\xi(\cdot)$ denotes the prior distribution. This construction employs a data record-indexed, risk-based weight vector with weights $\in [0, 1]$ to surgically downweight high-risk records in estimation of a pseudo posterior distribution for θ , subsequently used to generate and release a synthetic record-level database. Hu et al. (2021) show that this selective downweighting of records reduces the average of by-record risks as compared to an unweighted synthesis, while inducing only a minor reduction in utility. Their risk measure is based on a calculated probability of identification for a record. They cast a radius around the true data value for each record and count the number of record values that lie outside of the radius, which directly measures the extent that the target record is isolated and, therefore, easier for an intruder to discover by random guessing. While this risk measure appeals to intuition, it

is based on an assumption about the behavior of a putative intruder. By contrast, the DP framework makes no explicit assumptions about the behavior or knowledge of an intruder.

This paper extends Hu et al. (2021) and Dimitrakakis et al. (2017) by constructing weights $\alpha(\mathbf{x}) = (\alpha_1(\mathbf{x}), \dots, \alpha_n(\mathbf{x}))$ where $\alpha_i(\mathbf{x}) \propto 1/\sup_{\theta \in \Theta} |f(x_i | \theta)|$ that we show achieves a formal privacy guarantee that they were not able to achieve.

The remainder of the paper is organized as follows: Section 2 introduces a pseudo posterior mechanism using weights α plugged into Equation 3. We generalize Dimitrakakis et al. (2017) by establishing a direct functional association between a Lipschitz bound, Δ_α , for the pseudo posterior mechanism and a $(\epsilon = 2\Delta_\alpha)$ - DP guarantee. In Section 3, we describe the computation details to produce a matrix of (absolute values for) log-likelihoods estimated for the n records and S parameter draws taken from the unweighted posterior distribution and their subsequent use to formulate a vector of record-indexed weights, α , for a single observed database, \mathbf{x} . We then discuss the procedure to use the α to estimate the pseudo posterior distribution and the computation of the Lipschitz bound for the pseudo posterior mechanism based on the observed or local database. We call a Lipschitz bound constructed from a single (observed) database as a local Lipschitz. By contrast, we label a Lipschitz guarantee that represents a uniform bound over the space of databases as a global Lipschitz. Section 4 specifies formal conditions that guarantee the asymptotic contraction of a local Lipschitz bound to the global Lipschitz bound over the space of databases. We include a Monte Carlo simulation study that generates a collection of local databases and shows that the infimum and supremum of the local Lipschitz bounds collapse together to a global value as n approaches 1000. Section 5 focuses on our application to synthesizing the family income variable of a sample from the Consumer Expenditure surveys administered by the U.S. Bureau of Labor Statistics (BLS). This section presents the risk and utility curves of locally differentially private synthetic data generated under the proposed pseudo posterior mechanism, compared to the EM. We conclude with a discussion in Section 6.

2. Differential Privacy for the Pseudo Posterior

In this section, we specify the connection between achieving a global Lipschitz bound, Δ_α , under our pseudo posterior mechanism of Equation (3) with weights α and an ϵ - DP (or global DP) guarantee (over the space of databases). We further re-purpose a result from Wasserman and Zhou (2010) to extend a global DP guarantee for the mechanism generating parameters, θ , to the pseudo posterior predictive mechanism for generating synthetic data that is based on integrating with respect to the globally DP privacy guaranteed pseudo posterior distribution mechanism (used to generate the model parameters). After having shown that achievement of a global Lipschitz under our pseudo posterior mechanism produces a global DP privacy guarantee, we discuss constructing by-record weights used in our pseudo posterior mechanism that are designed to be inversely proportional to the (absolute value of) log-likelihood utilities computed over the parameter space. The log-likelihood for each record represents its relative risk of identification disclosure for the record since it governs the Lipschitz bound that defines the sensitivity. This construction of weights allows us to achieve a global Lipschitz (linked to a global DP guarantee) without data or parameter truncation.

2.1 Preliminaries

We begin by constructing the probability space, (Θ, β_Θ) , equipped with prior distribution, $\xi(\theta)$. Observe a database sequence, $\mathbf{x} = (x_1, \dots, x_n) \in \mathcal{X}^n$ under $x_1, \dots, x_n \stackrel{\text{ind}}{\sim} P_{\theta_0}$, for some $\theta_0 \in \Theta$, we formulate the pseudo likelihood,

$$p_\theta^\alpha(\mathbf{x}) = \prod_{i=1}^n p_{\theta_i}(x_i)^{\alpha_i(\mathbf{x})}, \quad (4)$$

for each $\theta \in \Theta$ and $\mathbf{x} \in \mathcal{X}^n$. The pseudo likelihood exponentiates likelihood contributions by $\alpha(\mathbf{x}) = (\alpha_1(\mathbf{x}), \dots, \alpha_n(\mathbf{x}))$, where $\alpha_i(\mathbf{x}) \in [0, 1]$ denote weights that are constructed to be inversely proportional to the local identification disclosure risk for each observed data set record. These weights are subsequently used to selectively downweight the likelihood contributions for records in proportion to the level identification disclosure risks that they express.

Given the prior and pseudo likelihood, we construct the pseudo posterior distribution,

$$\xi^\alpha(B | \mathbf{x}) = \frac{\int_{\theta \in B} p_\theta^\alpha(\mathbf{x}) d\xi(\theta)}{\phi^\alpha(\mathbf{x})} = \frac{\int_{\theta \in B} e^{-r_{n,\alpha}(\theta, \theta^*)} d\xi(\theta)}{\int_{\theta \in \Theta} e^{-r_{n,\alpha}(\theta, \theta^*)} d\xi(\theta)}, \quad (5)$$

where $\phi^\alpha(\mathbf{x}) \triangleq \int_{\theta \in \Theta} p_\theta^\alpha(\mathbf{x}) d\xi(\theta)$ normalizes the pseudo posterior distribution and $r_{n,\alpha}(\theta, \theta^*) = \sum_{i=1}^n \alpha_i \log \{p_{\theta_i^*}(x_i)/p_{\theta_i}(x_i)\}$, which is a generalization of the definition from Bhattacharya et al. (2019) that uses a fixed, scalar weight to now incorporate risk-adjusted, record-indexed weights, $(\alpha_i)_{i=1, \dots, n}$ where each $\alpha_i(\mathbf{x})$ depends on the closely-held data.

We formulate the α -weighted log-pseudo likelihood,

$$f_\theta^\alpha(\mathbf{x}) = \sum_{i=1}^n \alpha_i(\mathbf{x}) \log p_\theta(x_i), \quad (6)$$

that we use to construct a pseudo posterior mechanism.

2.2 Main Results

Our task is to specify assumptions that guarantee our pseudo posterior mechanism achieves an ϵ -DP guarantee. In particular, we extend Dimitrakakis et al. (2017) to show a direct relationship between the Lipschitz bound for the pseudo likelihood and the resulting ϵ -DP guarantee where both are a function of the record-indexed vector of weights, α , specified by the data provider. We present a collection of related results in this section with all of the associated proofs in Appendix A.

2.2.1 LINK THE GLOBAL LIPSCHITZ BOUND TO THE GLOBAL DP GUARANTEE

In this section and corresponding sections in Appendix A, we use the explicit notation $\alpha(\mathbf{x})$ to emphasize the dependence of the $\alpha_i \leq 1$ on the closely-held data, \mathbf{x} . We begin by extending the definition of DP from Dimitrakakis et al. (2017) to our α -weighted pseudo posterior mechanism.

Definition 4 (*Differential Privacy under the Pseudo Posterior Mechanism*)

$$\sup_{\mathbf{x} \in \mathcal{X}^n, \mathbf{y} \in \mathcal{X}^{n-1} : \delta(\mathbf{x}, \mathbf{y}) = 1} \sup_{B \in \beta_\Theta} \frac{\xi^{\alpha(\mathbf{x})}(B | \mathbf{x})}{\xi^{\alpha(\mathbf{y})}(B | \mathbf{y})} \leq e^\epsilon,$$

which limits the change in the pseudo posterior distribution over all sets, $B \in \beta_\Theta$ (i.e. β_Θ is the σ -algebra of measurable sets on Θ), from the inclusion of a single record (under the leave-one-out (LOO) distance, $\delta(\mathbf{x}, \mathbf{y}) = 1$, such that \mathbf{y} differs from \mathbf{x} by the omission of a *single* data record). Although the pseudo posterior distribution mass assigned to B depends on \mathbf{x} , the ϵ guarantee is defined as the supremum over all $\mathbf{x} \in \mathcal{X}^n$.

Our main assumption extends Dimitrakakis et al. (2017) to bound the *log-pseudo likelihood ratio*, uniformly, for all databases, $\mathbf{y} \in \mathcal{X}^{n-1}$ that are at a LOO distance (i.e. $\delta(\mathbf{x}, \mathbf{y}) = 1$), over all $\mathbf{x} \in \mathcal{X}^n$ and over all $\theta \in \Theta$. The uniform bound defines a maximum sensitivity in the log-pseudo likelihood from the inclusion of a record. Our intuition that the magnitude of this sensitivity for the log-pseudo likelihood ratio is directly tied to the resulting ϵ -DP guarantee of the pseudo posterior is confirmed in Theorem 5 for pseudo posterior draws of θ and in Lemma 6 for the subsequent generation of a synthetic database from a draw of θ .

Assumption 1 (*Lipschitz continuity*)

Fix some $\theta \in \Theta$ and define a collection of record indexed mappings $\alpha(\cdot) : \{\mathcal{X}_i \rightarrow [0, 1]\}_n$ for records $i = 1, \dots, n$ and construct the Lipschitz function of θ over the space of databases,

$$\ell^\alpha(\theta) \triangleq \inf \left\{ w : \left| f_\theta^{\alpha(\mathbf{x})}(\mathbf{x}) - f_\theta^{\alpha(\mathbf{y})}(\mathbf{y}) \right| \leq w, \forall \mathbf{x} \in \mathcal{X}^n, \mathbf{y} \in \mathcal{X}^{n-1} : \delta(\mathbf{x}, \mathbf{y}) = 1 \right\}.$$

Assumption 1 restricts Θ such that the Lipschitz function of θ is uniformly bounded from above,

$$\ell^\alpha(\theta) \leq \Delta_\alpha = \sup_{\theta \in \Theta} \{\ell^\alpha(\theta)\}.$$

Since $\alpha(\cdot)$ is a vector of record-indexed functions, $\alpha(\mathbf{x})$ and $\alpha(\mathbf{y})$ only differ for a single record j when \mathbf{x} and \mathbf{y} only differ in one record. Then $\alpha(\mathbf{y}) = \alpha(\mathbf{x}_{-j})$. We note that the subscripting of Δ with α is a notational device that denotes a Lipschitz bound computed using the log-pseudo likelihood, $f_\theta^{\alpha(\mathbf{x})}(\mathbf{x})$ as contrasted with Δ computed using the unweighted posterior mechanism. We further note that our general result simplifies to that of Dimitrakakis et al. (2017) by specifying $\alpha = \mathbf{1}$ and therefore $\ell^\alpha(\theta) \leq \Delta_\alpha \leq \Delta$.

We refer to Δ_α as “global” over the space of databases, $\mathbf{x} \in \mathcal{X}^n$ and it represents the sensitivity of the α -weighted pseudo likelihood of Equation (4) that we use as our utility function. The Lipschitz function of θ and α , $\ell^\alpha(\theta)$, is constructed using the pseudo log-likelihood, $f_\theta^{\alpha(\mathbf{x})}(\mathbf{x})$ that incorporates record-indexed weights, $\alpha(\mathbf{x})$, each of which is ≤ 1 . Selecting an α_i close to zero indicates strong downweighting of a highly sensitive record for an unweighted posterior mechanism (with a high magnitude log-likelihood ratio for some $\theta \in \Theta$), which will reduce the sensitivity of that record under our pseudo posterior mechanism. We see in our first two results that reducing the sensitivity of the log-likelihood ratio directly improves (i.e. reduces the value of ϵ) the ϵ -DP guarantee.

Our next result directly connects the (global) Lipschitz bound, Δ_α , for the log-pseudo likelihood of Assumption 1 to resulting DP guarantee, $\epsilon = 2\Delta_\alpha$, for each draw of θ from the pseudo posterior distribution.

Theorem 5 $\forall \mathbf{x} \in \mathcal{X}^n, \mathbf{y} \in \mathcal{X}^{n-1} : \delta(\mathbf{x}, \mathbf{y}) = 1, B \in \beta_\Theta$ (where β_Θ is the σ -algebra of measurable sets on Θ) under $\alpha(\cdot)$ with $\Delta_\alpha > 0$ satisfying Assumption 1:

$$\sup_{B \in \beta_\Theta} \sup_{\mathbf{x} \in \mathcal{X}^n, \mathbf{y} \in \mathcal{X}^{n-1} : \delta(\mathbf{x}, \mathbf{y}) = 1} \frac{\xi^{\alpha(\mathbf{x})}(B | \mathbf{x})}{\xi^{\alpha(\mathbf{y})}(B | \mathbf{y})} \leq \exp(2\Delta_\alpha), \quad (7)$$

i.e. the pseudo posterior $\xi^{\alpha(\mathbf{x})}(\cdot | \mathbf{x})$ is $2\Delta_\alpha$ -DP.

This result directly connects the global Lipschitz bound to the global DP guarantee and will allow us to control the DP guarantee, indirectly, by setting the record-indexed weights, $\alpha(\mathbf{x}) = (\alpha_1(\mathbf{x}), \dots, \alpha_n(\mathbf{x}))$, that determines the Lipschitz bound.

Our next result extends our DP guarantee from pseudo posterior draws of θ for models that satisfy Assumption 1 to draws of synthetic data, $\zeta = (\zeta_1, \dots, \zeta_m)$, constructed from the model pseudo posterior predictive distribution. The generation of synthetic data is the purpose for the pseudo posterior mechanism.

Lemma 6 Define $P^{\alpha(\mathbf{x})}(\zeta \in C | \mathbf{x}) = \int P(\zeta \in C | \theta, \mathbf{x}) d\xi^{\alpha(\mathbf{x})}(\theta | \mathbf{x})$ as the pseudo posterior predictive probability mass for ζ in set $C \in \mathcal{A}^n$ (the σ -algebra of sets for \mathcal{X}^n), constructed from our pseudo posterior model for θ that satisfies DP with expenditure, ϵ . Let $\zeta = (\zeta_1, \dots, \zeta_m)$ be m independent draws from $P^{\alpha(\mathbf{x})}(\zeta \in C | \mathbf{x})$. This defines a mechanism for ζ that satisfies DP with expenditure ϵ for any $m \leq n$.

We next formalize the method to construct our weighting scheme that characterizes our pseudo posterior mechanism.

Assumption 2 (*Risk-based Weighting for Pseudo Posterior Mechanism*)

Fix a value for n , the number of data records. Let $m(\cdot)$ be a monotonically decreasing scalar function $m : [0, \infty) \rightarrow [0, 1]$ such that $m(0) = 1$, and $m(\infty) = 0$. For every $\mathbf{x} \in \mathcal{X}^n$ choose a mapping $\alpha(\cdot)$ such that

$$\alpha_i = m \left(\sup_{\theta \in \Theta} |f_\theta(x_i)| \right), \quad (8)$$

where $f_\theta(x_i)$ is computed from the unweighted, non-differentially private posterior synthesizer. Under this procedure for selecting risk-based weights, α_i , $i = 1, \dots, n$, if $f_\theta(x_i)$ is non-finite for any x_i and value of $\theta \in \Theta$, α_i is set to $m(\infty) = 0$, which removes the contribution of database record, i , from the pseudo likelihood of Equation (4) used to formulate the pseudo posterior mechanism of Equation (5).

The mapping $m(\cdot)$ in Assumption 2 includes threshold ($m(z) = \mathbf{1}_{\{z < z^*\}}$) as well as smooth functions ($m(z) = (z + 1)^{-1}$), providing us the flexibility for how to implement the weighting in practice. Since we remove the likelihood contributions for all database records with non-finite log-likelihoods by setting their associated weights in our pseudo posterior mechanism to $m(\infty) = 0$, our mechanism is guaranteed to satisfy Assumption 1 with a finite $\Delta_\alpha < \infty$ and thus be globally differentially private. This is a non-asymptotic result at every n ; however we want to *estimate* the global Δ_α (and, therefore, ϵ), rather than simply knowing it exists.

We use Assumption 2 to implement our α -weighted pseudo posterior mechanism. Fix a database, \mathbf{x} , and compute a record-indexed vector of log-likelihood ratios, $|f_\theta(x_i)|$ and

linearly transform them to $|\tilde{f}_{\theta,i}| \in [0, 1]$ such that records with lower values for $|f_{\theta,i}|$, that indicate lower identification risks, produce values of $\tilde{f}_{\theta,i}$ near 0. We, next, set $\alpha_i = c \times (1 - \tilde{f}_i) + g$ where c and g may be used by the data provider to scale and shift the weights, respectively, restricted to $\alpha_i \in [0, 1]$, $\forall i \in (1, \dots, n)$ in order to achieve a desired Lipschitz bound, $\Delta_{\alpha, \mathbf{x}}$, for database, \mathbf{x} and the local DP privacy guarantee of $\epsilon_{\mathbf{x}} = 2\Delta_{\alpha, \mathbf{x}}$ (Hu et al. (2021) demonstrate the uses of c and g to fine tune the risk-utility trade-off in non-differentially private synthetic data settings). So the data provider indirectly controls the local privacy guarantee by formulating the weights. We discuss an asymptotic method in Section 4 that “discovers” a global Lipschitz bound and associated global ϵ of an (ϵ, π) -probabilistic DP guarantee from a local result. We show that the π , the probability of deviating from ϵ -DP, contracts onto 0 for a sufficiently large sample size, n .

In our application to a local data set \mathbf{x} we may want to use a weighting scheme $\alpha(\mathbf{x})$ which mildly violates the stated conditions in Assumption 1. In particular, we consider estimation of α_i which weakly depends on x_j for $i \neq j$, where this dependence attenuates asymptotically. For example, we use estimates of θ from an unweighted posterior distribution which weakly depend on all values of x_i . We see asymptotically that this dependence decays as θ collapses to a point, such that the results in this section apply to this case of weak dependence among the α_i except for minor updates to notation.

3. Computing a Local Lipschitz Bound

In this section, we describe the implementation algorithm to compute the pseudo likelihood weights, $\alpha = (\alpha_1, \dots, \alpha_n)$ for a local database, \mathbf{x} , from the *unweighted* synthesizer and the subsequent computation of the local Lipschitz bound, $\Delta_{\alpha, \mathbf{x}}$, associated with the pseudo posterior mechanism. In Section 3.1, we lay out the connection between the scalar-weighted pseudo posterior mechanism and the EM, with a discussion of the implications on the data utility of locally differentially private synthetic data generated under the two mechanisms.

1. Compute weights α

- (a) Let $|f_{\theta_s,i}|$ denote the absolute value of the log-likelihood computed from the unweighted pseudo posterior synthesizer for database record, $i \in (1, \dots, n)$ and MCMC draw, $s \in (1, \dots, S)$ of θ .
- (b) Compute the $S \times n$ matrix of by-record (absolute value of) log-likelihoods, $L = \{|f_{\theta_s,i}|\}_{i=1,\dots,n, s=1,\dots,S}$.
- (c) Compute the maximum over each $S \times 1$ column of L to produce the $n \times 1$ (database record-indexed) vector, $\mathbf{f} = (f_1, \dots, f_n)$. We use a linear transformation of each f_i to $\tilde{f}_i \in [0, 1]$ where values of \tilde{f}_i closer to 1 indicates relatively higher identification disclosure risk: $\tilde{f}_i = \frac{f_i - \min_j f_j}{\max_j f_j - \min_j f_j}$.
- (d) We formulate by-record weights, $\alpha = (\alpha_1, \dots, \alpha_n)$,

$$\alpha_i = c \times (1 - \tilde{f}_i) + g, \tag{9}$$

where c and g denote a scaling and a shift parameters, respectively, of the α_i used to tune the risk-utility trade-off. If we set scaling tuning parameter, $c = 1$ and

shift tuning parameter, $g = 0$, then each α_i is simply $(1 - \tilde{f}_i)$ such that the pseudo likelihood weights are solely a function of the record-indexed log likelihoods. As discussed in Hu et al. (2021), decreasing $c < 1$ will compress the distribution of the (α_i) while setting $g < 0$ will shift downward the distribution of the weights such that more weights will be close to 0. We use truncation to ensure each $\alpha_i \in [0, 1]$. These α satisfy a slightly weaker asymptotic form of Assumptions 1 and 2.

We will show in Section 5 the effects of different configurations of c and g on the risk and utility profiles of the differentially private synthetic data set for the CE sample, generated under our proposed α -weighted pseudo posterior mechanism.

2. Compute Lipschitz bound, $\Delta_{\alpha, \mathbf{x}}$

- (a) Use $\alpha = (\alpha_1, \dots, \alpha_n)$ to construct the pseudo likelihood of Equation 4 from which the pseudo posterior of Equation 5 is estimated. Draw $(\theta_s)_{s=1, \dots, S}$ from the α -weighted pseudo posterior distribution.
- (b) As earlier, compute the $S \times n$ matrix of log-pseudo likelihood values, $L^\alpha = \left\{ |f_{\theta_s, i}^\alpha| \right\}_{i=1, \dots, n, s=1, \dots, S}$
- (c) Compute $\Delta_{\alpha, \mathbf{x}} = \max_{s, i} |f_{\theta_s, i}^\alpha|$.

3. Draw synthetic data, ζ_ℓ , from the pseudo posterior distribution

- (a) Using the $(\theta_s)_{s=1, \dots, S}$ drawn from the α -weighted pseudo posterior distribution estimated in the earlier step, randomly sample $\ell = 1, \dots, (m = 20)$ parameter values and draw synthetic data value, $\zeta_{\ell, i} \stackrel{\text{ind}}{\sim} p_{\theta_\ell}(\cdot)$ for parameter draw $\ell \in (1, \dots, m)$ and database record $i \in (1, \dots, n)$. This step accomplishes a draw from the pseudo posterior predictive distribution.
- (b) Release the synthetic data, $\zeta = (\zeta_1, \dots, \zeta_m)$, in place of the closely-held real data, \mathbf{x} .

Our pseudo posterior mechanism *indirectly* sets the local DP guarantee, $2\Delta_{\alpha, \mathbf{x}}$ through the computation and subsequent scaling and shifting of the likelihood weights, α .

3.1 Exponential Mechanism Reduces to Scalar Weighting

Wasserman and Zhou (2010); Zhang et al. (2016); Snok and Slavkovic (2018) use the EM to generate synthetic data with privacy guarantees from a non-private mechanism. Suppose we start with a non-private mechanism, such as an unweighted posterior synthesizer,

$$\xi(\theta | \mathbf{x}, \gamma) \propto \left[\prod_{i=1}^n p(x_i | \theta) \right] \xi(\theta | \gamma). \quad (10)$$

Under the set-up of Zhang et al. (2016) that uses the log-likelihood function as the utility function, i.e. $u(\mathbf{x}, \theta) = \log(\prod_{i=1}^n p(x_i | \theta))$, the EM generates private samples from

$$\hat{\theta} \propto \exp\left(\frac{\epsilon \log(\prod_{i=1}^n p(x_i | \theta))}{2\Delta}\right) \xi(\theta | \gamma), \quad (11)$$

where the prior, $\xi(\theta | \gamma)$, is chosen as the “base” distribution as specified by McSherry and Talwar (2007) that ensures the EM produces a proper density function. Furthermore,

$$\begin{aligned} \exp\left(\frac{\epsilon \log\left(\prod_{i=1}^n p(x_i | \theta)\right)}{2\Delta}\right) \xi(\theta | \gamma) &= \exp\left(\log\left(\prod_{i=1}^n p(x_i | \theta)\right)^{\frac{\epsilon}{2\Delta}}\right) \xi(\theta | \gamma) \\ &= \left(\prod_{i=1}^n p(x_i | \theta)^{\frac{\epsilon}{2\Delta}}\right) \xi(\theta | \gamma), \end{aligned} \quad (12)$$

which demonstrates that the EM under a log-likelihood utility is equivalent to a risk-adjusted, scalar-weighted pseudo posterior synthesizer with scalar weight $\frac{\epsilon}{2\Delta}$, where $\alpha_i = \frac{\epsilon}{2\Delta}$, $\forall i \in (1, \dots, n)$. Wang et al. (2015) derived this same scalar-weighted result in their implementation of a gradient descent algorithm to sample the EM under a pseudo log-likelihood utility.

Using a scalar weight, $\alpha_i = \frac{\epsilon}{2\Delta}$, $\forall i \in (1, \dots, n)$, shown in Equation (12), we expect a resulting lower utility for synthetic data draws under this mechanism than we do under our α -weighted pseudo posterior shown in Equation (3), which uses a vector of record-indexed weights. The α -weighted pseudo posterior is more surgical and concentrates the downweighting to records with higher risk, whereas the EM must downweight all records the same amount. Downweighting all records the same amount will be conservative because the scalar weight is based on the *worst case* sensitivity, Δ , over the entire database of records and the parameter space, which is required to achieve a local DP privacy guarantee, and not tuned to the risk (\tilde{f}_i) of each record.

The re-casting of the EM as a scalar-weighted pseudo likelihood under a log-likelihood utility also provides insight into why our α -weighted pseudo posterior mechanism sets the ϵ -DP guarantee indirectly through specification of the vector of weights, $\alpha = (\alpha_1, \dots, \alpha_n)$, that determines Δ_α , which in turn, determines $\epsilon = 2\Delta_\alpha$. Since the commonly-used EM utilizes a single, scalar weight for all records, it is straightforward to directly set ϵ , but at a tremendous loss of efficiency in terms of risk-utility trade-off as compared to the α -weighted pseudo posterior mechanism. So our mechanism achieves a higher utility for an equivalent guarantee, ϵ .

We illustrate in Section 5 the reduction in utility of the local differentially private synthetic data set generated under the EM, compared to that generated under our proposed α -weighted pseudo posterior mechanism, at an equivalent privacy guarantee for both mechanisms.

4. Turning A Local Bound into A Global Bound

In this section we proceed to demonstrate that a local Lipschitz bound or sensitivity, $\Delta_{\alpha, \mathbf{x}}$, computed on observed database, \mathbf{x} , contracts on or becomes arbitrarily close to Δ_α , the global Lipschitz bound or supremum over the space of databases, \mathcal{X}^n , for sample size, n , sufficiently large.

4.1 Asymptotic Convergence of Local Lipschitz to Global Lipschitz

Although our DP result is non-asymptotic for every n , in the sense that we have earlier shown that a *finite* global Δ_α is guaranteed to exist under our α -weighted pseudo posterior

mechanism, we nevertheless do not *know* its value. We employ asymptotics to learn the global Lipschitz bound, Δ_{α} , to any degree of desired precision. We develop a contraction result for any α -weighted pseudo distribution to demonstrate under a set of conditions that convergence of the pseudo posterior distribution leads to asymptotic convergence of the local Lipschitz bound, $\Delta_{\alpha, \mathbf{x}}$, to the global bound, Δ_{α} in P_{θ_0} -probability for n sufficiently large.

Our asymptotic contraction of the local Lipschitz bound onto the global Lipschitz bound (that has a direct functional relationship to the global privacy guarantee, ϵ) does not provide a global ϵ -DP *guarantee* because there is the possibility of leakage of private information, π , at any fixed sample size such that our computed ϵ on a local database may be exceeded. Therefore, we employ our asymptotic result on the contraction of Lipschitz bounds to claim an (ϵ, π) -probabilistic DP guarantee where δ represents a probability that there are some databases in the space of databases for which ϵ is exceeded. Under our asymptotic contraction of local Lipschitz bounds to the global bound, we achieve that π contracts onto 0.

We formally introduce a definition for probabilistic differential privacy (pDP) that adapts the formulation of Machanavajjhala et al. (2008) to our α -weighted pseudo posterior mechanism.

Definition 7 (*Probabilistic Differential Privacy*) *Let $\epsilon > 0$ and $0 < \pi < 1$. We say that our pseudo posterior mechanism is (ϵ, π) -probabilistically differentially private (pDP) if $\forall \mathbf{x} \in \mathcal{X}^n$,*

$$Pr(\mathbf{x} \in \text{Disc}(\mathbf{x}, \epsilon)) \leq \pi,$$

where the probability is taken over $\mathbf{x} \in \mathcal{X}^n$ and $\text{Disc}(\mathbf{x}, \epsilon)$ denotes the *disclosure set*, $\{\mathbf{x} \in \mathcal{X}^n : \sup_{B \in \beta_{\Theta}} \log \left(\frac{\xi^{\alpha(\mathbf{x})}(B|\mathbf{x})}{\xi^{\alpha(\mathbf{y})}(B|\mathbf{y})} \right) > \epsilon, \forall \mathbf{y} : \delta(\mathbf{x}, \mathbf{y}) = 1\}$, the subspace of \mathcal{X}^n where our α -weighted pseudo posterior mechanism exceeds an ϵ -DP guarantee.

This definition constructs a probability for the event that there are *any* databases in the space of databases for which our pseudo posterior mechanism *exceeds* ϵ under the leave-one-out (LOO) distance. We recall that our vector weights, $\alpha = (\alpha_1, \dots, \alpha_n)$ determine Δ_{α} , which *indirectly* sets $\epsilon \leq 2\Delta_{\alpha}$. Our asymptotic result on the contraction of the local to global Lipschitz bound, presented in this section, reveals that π , which represents the (maximum) probability that ϵ -DP is exceeded, limits to 0 in P_{θ_0} -probability.

We verify our theoretical result by conducting a simulation study in Section 4.5 that demonstrates the contraction of the distribution for the local $\Delta_{\alpha, \mathbf{x}}$ for a relatively moderate sample sizes. Furthermore, we suggest a procedure for selecting a global ϵ that would result in a very small-to-negligible π .

4.2 Preliminaries

We next demonstrate the frequentist properties of our pseudo posterior Bayesian estimator. We generalize the result of Bhattacharya et al. (2019) developed for a fixed, scalar weight to our vector of record-indexed weights that depend on the closely-held data. Suppose $x_1, \dots, x_n \stackrel{\text{ind}}{\sim} P_{\theta_0}$ for $\theta_0 \in \Theta$. Under *frequentist* consistency, the $\mathbf{x} = (x_1, \dots, x_n)$ are random with respect to P_{θ_0} (for fixed θ_0), so taking probabilities and expectations with respect to P_{θ_0} requires us to address the dependence of α_i on \mathbf{x} to construct the contraction rate for

correctness and thoroughness. We drop the notation denoting the explicit dependence of $\alpha_i(x_i)$ for exposition of our consistency results in the sequel and just use α_i for readability when the context is clear.

Since our pseudo posterior formulation induces misspecification, we allow the true generating parameters, θ_0 , to lie outside the parameter space, Θ . We will show in the sequel that our model contracts on $\theta^* \in \Theta$ in P_{θ_0} -probability, where θ^* is the point that minimizes the Kullback-Liebler (KL) divergence from P_{θ_0} ; that is,

$$\theta^* := \arg \min_{\theta \in \Theta} D(p_\theta, p_{\theta_0}), \quad (13)$$

where $D(p, q) = \int p \log(p/q) d\mu$ for dominating measure, μ .

Our asymptotic result on the contraction in P_{θ_0} -probability relies on bounding the α -Rényi divergence measure,

$$D_{\theta_0, \alpha}^{(n)}(\theta, \theta^*) = \sum_{i=1}^n D_{\theta_0, \alpha, i}(\theta, \theta^*) = \sum_{i=1}^n \frac{1}{\alpha_i - 1} \log \{A_{\theta_0, \alpha, i}(\theta, \theta^*)\}, \quad (14)$$

where $A_{\theta_0, \alpha, i}(\theta, \theta^*) = \int \left(\frac{p_{\theta_i}}{p_{\theta_i^*}}\right)^{\alpha_i} p_{\theta_0, i} d\mu_i$ under dominating measure μ_i is defined as the α -affinity for record, x_i , such that $A_{\theta_0, \alpha}^{(n)}(\theta, \theta^*) = \prod_{i=1}^n A_{\theta_0, \alpha, i}(\theta, \theta^*)$, the α -affinity for the product measure space, where we have updated definitions from use of a scalar, α , to record-indexed α_i .

The posterior probability of the α -Rényi distance between $\theta \in \Theta$ and the point θ^* limits to 0 at a rate that is a function of n for any weighting scheme, $\alpha(\mathbf{x})$, where the construction of α depends on the observed data, \mathbf{x} , as does ours. We require the following two conditions to achieve contraction of the local $\Delta_{\alpha, \mathbf{x}}$ to the global Δ_α :

Assumption 3 (*Prior mass covering truth*) We construct a KL neighborhood of θ^* with radius, η , with,

$$B_n(\theta^*, \eta; \theta_0) = \left\{ \theta \in \Theta : \sum_{i=1}^n \int p_{\theta_0, i} \log(p_{\theta_i^*}/p_{\theta_i}) d\mu_i \leq n\eta^2, \right. \\ \left. \sum_{i=1}^n \int p_{\theta_0, i} \log^2(p_{\theta_i^*}/p_{\theta_i}) d\mu_i \leq n\eta^2 \right\}. \quad (15)$$

Restrict the prior, ξ , to place positive probability on this KL neighborhood,

$$\xi(B_n(\theta^*, \eta; \theta_0)) \geq e^{-n\tau_n^2}. \quad (16)$$

Assumption 4 (*Control size of α*) Let $A_n := \{i : \alpha_i < 1^-\}; i \in 1, \dots, n\}$ and $n_A := |A_n|$, where $|A_n|$ denotes the number of elements in A_n . Let

$Q_n := \{i : \alpha_i = \alpha^{(n)} \geq 1^-\}; i \in 1, \dots, n\}$ for some constant $\alpha^{(n)}$ and $n_Q := |Q_n|$.

$$\limsup_n |A_n| = \limsup_n n_A = \mathcal{O}\left(n^{\frac{1}{2}}\right), \text{ with } P_{\theta_0}\text{-probability } 1,$$

$$\limsup_n (1 - \alpha^{(n)}) = \mathcal{O}\left(n_Q^{-\frac{1}{2}}\right), \text{ with } P_{\theta_0}\text{-probability } 1,$$

such that for constants $C_1, C_3 > 0$ and n sufficiently large,

$$\begin{aligned} \sup_n |A_n| &\leq C_1 n^{\frac{1}{2}}, \\ \sup_n (1 - \alpha^{(n)}) &\leq C_3 \tau_n n_Q^{-\frac{1}{2}}. \end{aligned}$$

These two assumptions are required for consistency of our α -pseudo posterior mechanism at θ^* . The first assumption requires the prior to place some mass on a KL ball near θ^* as defined in Equation (13). The second assumption outlines a dyadic subgrouping of data records, where A_n contains those records whose likelihood contributions are downweighted to lessen the estimated identification disclosure risk (and improve privacy) for those records in the resulting synthetic data. The second subset of records, Q_n , contains those records that are minimally downweighted due to nearly zero values for identification disclosure risks. Since $\alpha_i \leq 1$, $\forall i \in (1, \dots, n)$, the constant value, $\alpha^{(n)}$, for all units in Q_n approaches 1 from the left. We show that the consistency result to θ^* for the synthesizer is dominated by the likelihood weighting for records in the downweighted set, A_n . Assumption 4 restricts the number of downweighted records (where $\alpha_i < 1^-$) to grow at a slower rate than the sample size, n , such that the downweighting becomes relatively more sparse. Our experience demonstrates that when weights are constructed based on disclosure risks, downweighting is confined to isolated records, which are sparse.

Theorem 8 (*Contraction of the α -pseudo posterior distribution*).

Let $\alpha = (\alpha_1 \in [0, 1], \dots, \alpha_n \in [0, 1])$. Define $\alpha_m := \max_{i \in A_n} \alpha_i \in [0, 1]$ and $\alpha_l := \min_{i \in A_n} \alpha_i \in [0, 1]$.

Let $D_{\theta_0, \alpha}^{(n_A)}(\theta, \theta^*) = \sum_{i \in A_n} D_{\theta_0, \alpha, i}$ and $D_{\theta_0, 1^-}^{(n_Q)}(\theta, \theta^*) = \sum_{i \in Q_n} D_{\theta_0, 1^-, i}$. Let θ^* be as defined in Equation (13). Assume that τ_n satisfies $n\tau_n^2 \geq 2$ and suppose Assumptions 3 and 4 hold. Let $C_1^* = \sqrt{2 + C_1^2 + C_3^2} \geq \sqrt{2}$. Then for any $D \geq 2$ and $t > 0$,

$$\xi^\alpha \left(\frac{1}{n} \left[(1 - \alpha_m) D_{\theta_0, \alpha}^{(n_A)}(\theta, \theta^*) + (1 - \alpha^{(n)}) D_{\theta_0, 1^-}^{(n_Q)}(\theta, \theta^*) \right] \geq (D + 3t)\tau_n^2 | \mathbf{x} \right) \leq e^{-tn\tau_n^2}, \quad (17)$$

hold with P_{θ_0} -probability at least $1 - [(\alpha_l^2 + 2)(C_1^*)^2 / \alpha_m^2 \times 2 / \{(D + t - 1)^2 n \tau_n^2\}]$.

Since $(1 - \alpha^{(n)}) = \mathcal{O}(n_Q^{-1/2})$, while $n_A = \mathcal{O}(n^{1/2})$, the first term dominates with increasing n , so that the $(1 - \alpha_m)^{-1}$ is the dominating penalty on the τ_n contraction rate of the α -pseudo posterior onto θ^* . Even though the downweighting becomes relatively more sparse due to Assumption 4, it is the maximum value of α_i for $i \in A_n$ on the set of downweighted records that penalizes the rate. We observe that the rate of contraction is injured by factor, $(1 - \alpha_m)^{-1}$. Since $\alpha_i \leq 1^-$, $\forall i \in A_n$, our result generalizes Bhattacharya et al. (2019) to allow a tempering of a *portion* of the posterior distribution and there is a penalty to be paid in terms of contraction rate for the tempering. Since we induce the misspecification through the weights, α , the distance of the point of contraction, θ^* from the true generating parameters, θ_0 , and the contraction rate on this point are *both* impacted by the induced misspecification. The requirement for increasing sparsity in the number of downweighted record likelihood contributions, however, ensures that θ^* will be relatively close to θ_0 that produces a high utility for our (pseudo posterior) estimator.

If we plug in for τ_n , we see that our contraction of Θ to θ^* occurs at a rate that is of $\mathcal{O}(n^{-1/2})$.

4.3 Contraction of Local Lipschitz bound onto Global bound

Asymptotically Theorem 8 guarantees that the space $\theta \in \Theta$ collapses onto θ^* for n sufficiently large. The space of databases, \mathcal{X}^n , drawn under this distribution collapses unto a single distribution, $\tilde{x} \stackrel{\text{ind}}{\sim} P^\alpha(x|\theta^*)$ with density $p^\alpha(x|\theta^*) \propto \exp(f_{\theta^*}^\alpha(x))$. The term \tilde{x} denotes the risk-corrected version of \mathbf{x} under which high disclosure risk records are less likely to be drawn due to their downweighting. High-risk records are isolated relative to other records located in regions of the unweighted generating distribution, P_{θ_0} , of low-probability mass, such as the tails. Since the contraction of the pseudo posterior distribution induces the collapsing of the parameter space to a point and the space of databases to a single distribution (conditioned on θ^*) for large n , this result guarantees that the local Lipschitz bound, $\Delta_{\alpha,\mathbf{x}}$ and the the global bound Δ_α contract together for n sufficiently large. First, we revisit the different forms of the Δ bound:

$$\Delta_\alpha = \sup_{\theta \in \Theta} \left\{ \inf \left\{ w : \left| f_\theta^{\alpha(x_i)}(x_i) \right| \leq w, \forall x_i \in \mathbf{x}, \forall \mathbf{x} \in \mathcal{X}^n \right\} \right\} \quad (18)$$

$$= \sup_{\theta \sim \xi^\alpha(\theta|\mathbf{x})} \left\{ \inf \left\{ w : \left| f_\theta^{\alpha(x_i)}(x_i) \right| \leq w, \forall x_i \in \mathbf{x}, \forall \mathbf{x} \in \mathcal{X}^n \right\} \right\} \quad (19)$$

$$(20)$$

The formulation of Δ_α comes directly from Assumption 1. This is a bound of the α -weighted log-likelihood over the full support of both θ and \mathbf{x} . Note that we can replace the sup over Θ in Equation 18 with the sup over infinite draws from the pseudo posterior to achieve the same result in Equation 19 because the support is still the entire space, Θ , for any finite n .

The local $\Delta_{\alpha,\mathbf{x}}$ in Equation 21 is a random quantity based on only the observed values x_i from n draws from the generating distribution $P(x|\theta_0)$, and a fixed finite number S of draws θ_s from the posterior $\xi^\alpha(\theta|\mathbf{x})$. Finally, $\Delta_{\alpha,\mathbf{x}}^\infty$ in Equation 22 expresses the bound of the log-likelihood based on the observed database values evaluated at the limiting point θ^* .

$$\Delta_{\alpha,\mathbf{x}} = \max_{\theta_s \sim \xi^\alpha(\theta|\mathbf{x})} \left\{ \inf \left\{ w : \left| f_{\theta_s}^{\alpha(x_i)}(x_i) \right| \leq w, i \in \{1, \dots, n\}, x_i \sim P(x|\theta_0) \right\} \right\} \quad (21)$$

$$\Delta_{\alpha,\mathbf{x}}^\infty = \inf \left\{ w : \left| f_{\theta^*}^{\alpha(x_i)}(x_i) \right| \leq w, i \in \{1, \dots, \infty\}, x_i \sim P(x|\theta_0) \right\} \quad (22)$$

Then based on Theorem 8 both $P_{\theta_0}\{|\Delta_{\alpha,\mathbf{x}} - \Delta_{\alpha,\mathbf{x}}^\infty| > 0\} \rightarrow 0$ and $P_{\theta_0}\{|\Delta_\alpha - \Delta_{\alpha,\mathbf{x}}^\infty| > 0\} \rightarrow 0$ because the pseudo posterior degenerates to a point mass at θ^* . Thus

$$P_{\theta_0}\{|\Delta_{\alpha,\mathbf{x}} - \Delta_\alpha| > 0\} \rightarrow 0 \quad (23)$$

Assumption 2 ensures a formal privacy guarantee since $\Delta_\alpha < \infty$, by construction, and the asymptotic result guarantees that the local $\Delta_{\alpha,\mathbf{x}}$ will get arbitrarily close to the global Δ_α where $\epsilon = 2\Delta_\alpha$. For a large n , then, $\Delta_{\alpha,\mathbf{x}} \rightarrow \Delta_\alpha$ becomes independent of $\mathbf{x} \in \mathcal{X}^n$, where we recall that $\epsilon = 2\Delta_\alpha$. This contraction of the local Lipschitz bound onto a global value that determines the privacy guarantee, ϵ , indicates that π of our (ϵ, π) -pDP guarantee in Equation 13 contracts onto 0 at $\mathcal{O}(n^{-1/2})$ rate at which $\Delta_{\alpha,\mathbf{x}}$ contracts onto Δ_α . To speed convergence and add stability for $\Delta_{\alpha,\mathbf{x}}$ for finite n , we consider employing a threshold M for

the α -weighted log-likelihood, such that $\left|f_{\theta}^{\alpha(x_i)}(x_i)\right| > M$ is replaced by 0, through setting $\alpha(x_i) = 0$.

To make intuitive the rate of contraction of π to 0 at $\mathcal{O}(n^{-1/2})$, we conduct a Monte Carlo simulation study next to develop a distribution of local Lipschitz bounds from which we compute the local Lipschitz, $\Delta_{\alpha, \mathbf{x}}$, each at an increasing sequence of sample sizes, n . We reveal that the distribution over local Lipschitz bounds contracts together onto a single global value, demonstrating the local-to-global contraction as n increases. In particular, the use of the M threshold greatly stabilizes and speeds convergence.

4.4 Asymptotic Differential Privacy Guarantee

We have noted that Theorem 8 induces the contraction of $\Delta_{\alpha, \mathbf{x}}$ computed on database, \mathbf{x} , to the global Lipschitz bound, Δ_{α} . This contraction is driven by the collapsing of the parameter space, Θ , to a point, θ^* , asymptotically in data size, n .

Our implementation for computing the by-record Lipschitz bound, $\sup_{\theta \in \Theta} |f(x_i | \theta)|$, on a database relies on this asymptotic convergence. For implementation on a database we evaluate $\sup_{\theta \in \xi^{\alpha}(\theta | \mathbf{x})} |f(x_i | \theta)|$; that is, we compute the supremum of the absolute value of the log-likelihood over the subset of Θ that receives positive posterior mass. This subset shrinks to a point for n sufficiently large, making our treatment conservative for large n . Our de facto truncation of Θ to that subset receiving positive posterior measure is similar to the probabilistic Lipschitz condition of Assumption 2 in Dimitrakakis et al. (2017); only, in our implementation of differential privacy guarantee is asymptotic such that we rely on the shrinking of the size of Θ with increasing n to state our privacy guarantee. Our use of a threshold M speeds this convergence, making it useful for moderate sample sizes. This indirect truncation of Θ is much simpler to implement compared to specifying a meaningful truncated prior distribution in high-dimension.

4.5 Simulation Study

We next utilize a Monte Carlo simulation study by fixing a sample size, n , and repeatedly generating a count data sample from a Poisson generating model. We proceed to compute the local Lipschitz bound for each sample database for the α -weighted pseudo posterior mechanism and also the unweighted posterior synthesizer to provide a comparison. This procedure gives us a distribution of the local Lipschitz bounds across databases of size n . We repeat this process for an increasing sequence of sample sizes, $\mathbf{n} = 100 * 4^{(0,1,2,3)} = (100, 400, 1600, 6400)$.

In addition to computing the local Lipschitz bounds at each n for the α -weighted pseudo posterior mechanism, we introduce an extension to our pseudo posterior mechanism that truncates the weight, α_i , for each likelihood contribution in the following procedure:

1. Compute weights, α , for local database, \mathbf{x} , using the procedure of Section 3. We first compute f_i (the maximum of absolute log-likelihood values for record, i , over the sampled values of θ_s) for each database record, $i \in (1, \dots, n)$, from the unweighted posterior mechanism. Then using the linear transform, $\alpha_i = 1 - \tilde{f}_i$, where $\tilde{f}_i = \frac{f_i - \min_j f_j}{\max_j f_j - \min_j f_j}$.

2. We add a step to truncate the weight for any record whose weighted log-pseudo likelihood value is greater than some threshold, M , to 0, completely removing the likelihood contribution for record i . We accomplish this truncation by forming a weighted absolute log-pseudo likelihood for each record, i , as $\alpha_i \times f_i$. If $\alpha_i \times f_i > M$, we set final weight, $\alpha_i^* = 0$; otherwise we leave $\alpha_i^* = \alpha_i$ unchanged. The motivation for this method is to more tightly control or correct the local Lipschitz to M that we will observe in the sequel speeds convergence. We choose M based on oracle information based on experience with databases of similar types.

The use of a threshold, M , to truncate weights is a stricter implementation from our weight-setting procedure of Assumption 2. We recall that this assumption guarantees the existence of a global Lipschitz because for every database it sets the weight for a record with a non-finite absolute log likelihood to 0. In this stricter implementation, we set $\alpha_i^* = 0$ if its weighted absolute log-pseudo likelihood is $> M$, where we choose M based on oracle information based on experience with databases of similar types.

Using the means model for Poisson distributed data, $y \sim Pois(\mu)$ (with $\mu = 100$) our simulation procedure is, as follows.

1. For sample size, $n \in \{100, 400, 1600, 6400\}$, repeat the following Monte Carlo procedure to generate a distribution of local Lipschitz bounds:
2. For $r = 1, \dots, 400$:
 - Generate $\mathbf{y}_r \sim Pois(\mu)$, each of size n .
 - Compute the *local* Lipschitz bound, $\Delta_{\alpha, \mathbf{y}_r}$, for the unweighted, α -weighted, and M -truncation-weighted pseudo posterior mechanisms.
 - Construct the distribution of $\Delta_{\alpha, \mathbf{y}_r}$ and note the maximum of the distribution and difference between the maximum and minimum values of the distribution of the local Lipschitz bounds at each sample size, n .
3. Assess contraction of the $\max_r \Delta_{\alpha, \mathbf{y}_r}$ to a single (global) value and whether the minimum and maximum values collapse together.

To assess the *contraction* of the maximum point in the distribution of *local* Lipschitz bounds to the *global* Lipschitz bound, we repeat the simulation above using sample sizes $\mathbf{n} = (100, 400, 1600, 6400)$. Figure 1 compares the distributions across the $R = 400$ replications for the unweighted (labeled as “Unweighted”), α -weighted (labeled as “Weighted”) that does *not* use truncation of weights and the truncated weighted at M (labeled as “Weighted-M”) mechanisms, from left-to-right. The distribution of local Lipschitz bounds for the Unweighted mechanism increases (or drifts) with larger sample sizes. The Weighted mechanism (that includes no weight truncation) shows a pronounced decrease in drift in the maximum Lipschitz of local databases over the increasing sample sizes as compared to the Unweighted mechanism, though even at sample size, 6400, there is still a small, though decreasing drift of the maximum Lipschitz. By contrast, the Weighted-M mechanism, under setting $M = 3.5$, demonstrates rapid contraction of both the minimum and maximum local Lipschitz values onto M . This is still a probabilistic formal privacy result because the local Lipschitz values are not strictly bounded below M due to sampling variability of θ .

The maximum of the distribution of local Lipschitz bounds at each sample size is slightly larger than M , indicating that our guarantee is probabilistic. While both the Weighted and Weighted-M local Lipschitz bounds contract at $\mathcal{O}(n^{-1/2})$, the multiplicative constant of the contraction rate is much smaller for Weighted-M because of the truncation to an asymptotic global Lipschitz of M defined by the owner of the closely-held data.

Figure 2 presents the distributions for the averages of the mean parameter, μ , over the $R = 400$ Monte Carlo iterations. We see there is some utility loss relative to Unweighted and Weighted under use of Weighted-M, though the resulting utility is still relatively robust. The deterioration of the utility for Weighted-M as n increases (as represented by the growing dissimilarity of the pseudo posterior distribution for μ to that under Unweighted) is a conservative result because we use the same $M = 3.5$ for all sample sizes. Yet, the DP guarantee is based on the space of databases at a particular sample size, n , and M will be set based on agency experience with a particular class of data (e.g., monthly survey responses) that all have very similar values for n .

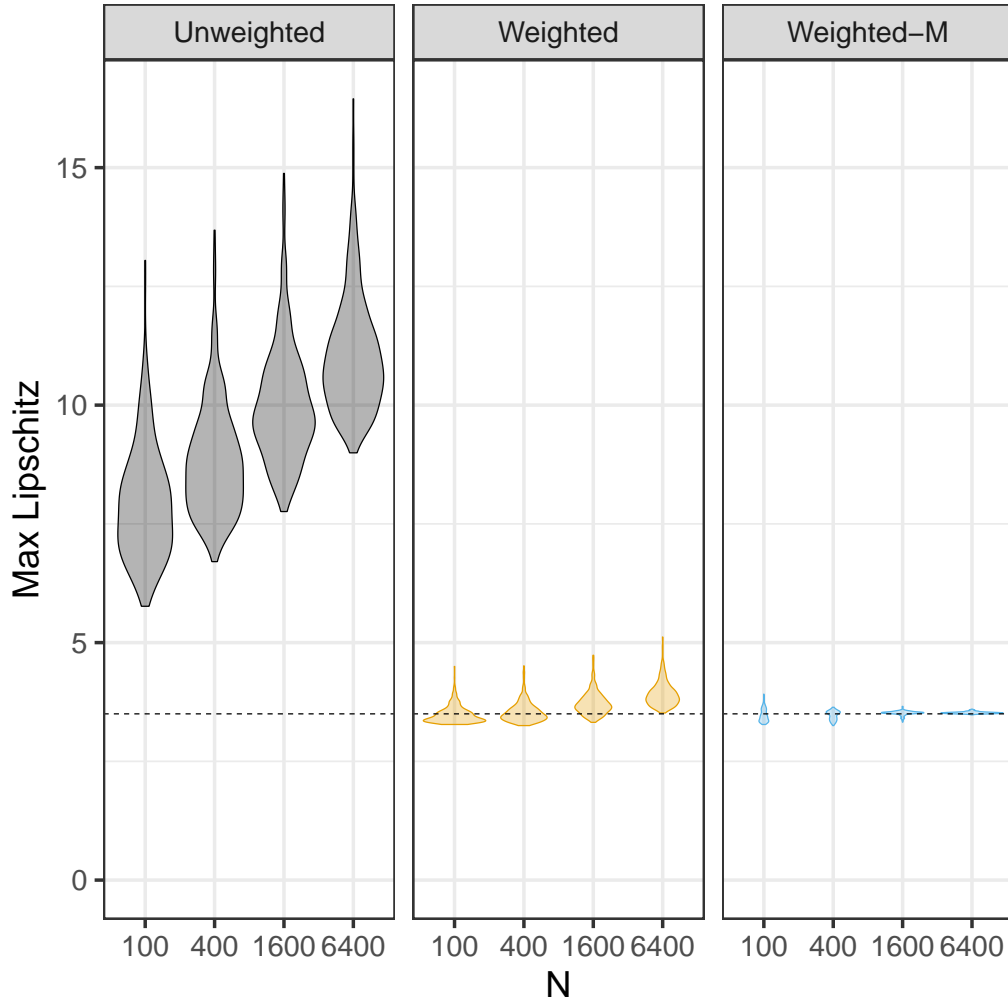


Figure 1: Distribution of the maximum observed Lipschitz bound $\Delta_{\mathbf{y}}$ for each of sample sizes (100, 400, 1600, 6400) from $R = 400$ realizations of pseudo posterior samples of (left to right) unweighted, weighted, and weighted-M (M -truncated weights).

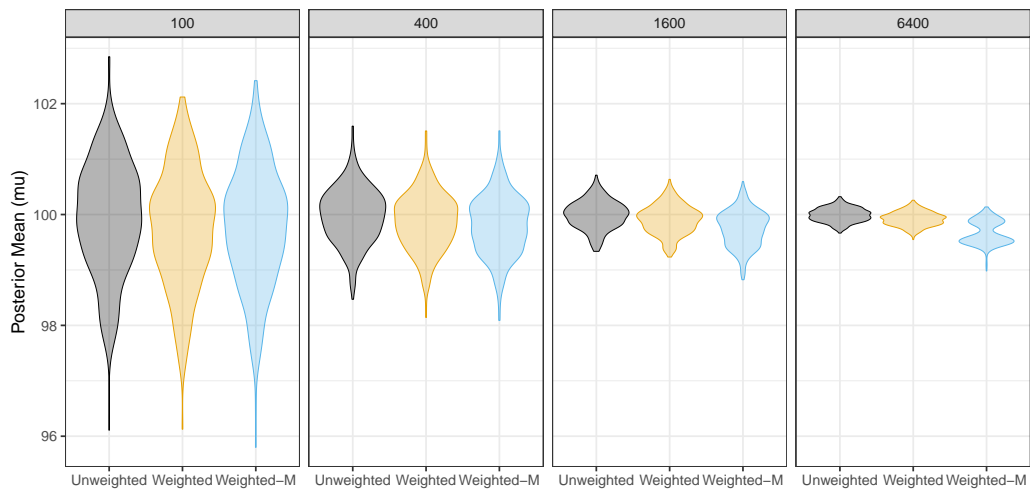


Figure 2: Distributions of the average of mean parameter μ for each of sample size (100, 400, 1600, 6400) from $R = 400$ realizations of pseudo posterior samples of (left to right) unweighted, weighted, and weighted-M (M -truncated weights).

4.6 Recommendation for Setting a Global ϵ from a weighted- M Lipschitz

Although we have demonstrated an $\mathcal{O}(n^{-1/2})$ contraction rate of π (the probability of exceeding ϵ - global DP), in theory, and have further illustrated this convergence in our above simulation study, it is difficult in practice to discover at what sample size under a specific synthesizer that one may declare the local Lipschitz to be global. For typically used sample sizes > 1000 we suggest to take the Weighted-M Lipschitz and employ a multiplicative “factor of safety”, $s \in (1, 1.05)$, to develop an upper, global bound that, in turn, determines ϵ because the contraction is extremely rapid. If the sample size is < 1000 , we recommend to set multiplicative $s' \in (1.05, 1.10)$.

5. Application to the CE Sample

We introduce the CE sample of consumer units (CU) or households in Section 5.1, where our goal is to synthesize a highly-skewed continuous variable, family income, under a local DP guarantee provided by our α -weighted pseudo posterior mechanism. In Section 5.2, we present risk and utility profiles of synthetic data drawn from our α -weighted pseudo posterior mechanism, along with comparisons to the EM, the non-differentially private risk-weighted synthesizer of Hu et al. (2021) and the unweighted posterior mechanism. Section 5.3 presents privacy and utility results with different scaling and shifting, (c, g) , configurations for vector weights in Equation (9) to sketch out a risk-utility curve for our α -weighted pseudo posterior mechanism that we compare to that of the EM. A risk-utility curve provides the Bureau of Labor Statistics (BLS) options for selecting a risk-utility setting that matches their policy objectives.

5.1 The CE Sample and Unweighted Synthesizer

Our application of the α -weighted pseudo posterior mechanism focuses on providing privacy protection for a family income variable published by the CE. The CE is administered by the BLS with the purpose of providing income and expenditure patterns indexed by geographic domains to support policy-making by State and Federal governments. The description of the CE sample included here closely follows that in Hu et al. (2021). The CE contain data on expenditures, income, and tax statistics about CUs across the U.S. The CE public-use microdata (PUMD)¹ is publicly available record-level data, published by the CE. The CE PUMD has undergone masking procedures to provide privacy protection of survey respondents. Notably, the family income variable has undergone top-coding, a popular Statistical Disclosure Limitation (SDL) procedure that may result in reduced utility and insufficient privacy protection (An and Little, 2007; Hu et al., 2021).

The CE sample in our application contains $n = 6208$ CUs, coming from the 2017 1st quarter CE Interview Survey. It includes the family income variable, which is highly right-skewed and deemed sensitive; see Figure 3 for its density plot. The CE sample also contains 10 categorical variables, listed in Table 1. These categorical variables are deemed insensitive and used as predictors in building a flexible synthesizer for the synthesis of the sensitive family income variable.

1. For for information about CE PUMD, visit <https://www.bls.gov/cex/pumd.htm>.




Figure 3: Density plot of Family Income in the CE sample.

Variable	Description
Gender	Gender of the reference person; 2 categories
Age	Age of the reference person; 5 categories
Education Level	Education level of the reference person; 8 categories
Region	Region of the CU; 4 categories
Urban	Urban status of the CU; 2 categories
Marital Status	Marital status of the reference person; 5 categories
Urban Type	Urban area type of the CU; 3 categories
CBSA	2010 core-based statistical area (CBSA) status; 3 categories
Family Size	Size of the CU; 11 categories
Earners	Earners status of the reference person; 2 categories
Family Income	Imputed and reported income before tax of the CU; approximate range: (-7K, 1,800K)

Table 1: Variables used in the CE sample. Data taken from the 2017 Q1 Consumer Expenditure Surveys.

To generate partially synthetic databases for the CE sample with synthetic family income, we use an unweighted, non-private synthesizer: a flexible, parametric finite mixture synthesizer. This finite mixture synthesizer has been shown to produce synthetic data characterized by a high utility, but also with an unacceptable level of disclosure risk (Hu et al., 2021). We leave the details of the synthesizer in the Appendix B for brevity and direct interested readers to Hu et al. (2021) for further information.

5.2 Risk and Utility Comparisons

To generate synthetic data and compare results, we apply four synthesizers: 1) the unweighted, non-(locally) private synthesizer, labeled “Unweighted”; 2) the locally private synthesizer under the α -weighted pseudo posterior mechanism, labeled “DPweighted”, with configuration $(c, g) = (0.7, 0.0)$; 3) the locally private synthesizer under the EM, labeled “EMweighted”, which is designed to privacy target, ϵ , achieved by “DPweighted”; 4) and the weighted, though non-(locally) private pseudo posterior synthesizer proposed by Hu et al. (2021), labeled “Countweighted”, that utilizes their method for measuring the by-record disclosure risk (based on an assumption about the behavior of an intruder). We use α_c to denote the risk-adjusted record-indexed weights calculated in the Countweighted method. The labels are used throughout the remainder of this paper when presenting various risk and utility results.

We first look at the risk profiles of the four synthesizers. Figure 4 plots the distributions of the Lipschitz bounds, Δ_{x_i} ’s, for each of the four synthesizers computed by taking the maximum of the S log-likelihood ratios for each record, $i = 1, \dots, (n = 6208)$ over the S draws of θ from its posterior distribution. The maximum value of the (Δ_{x_i}) over all of the records is denoted as $\Delta_{\mathbf{x}}$, the Lipschitz bound for the mechanism.

The Unweighted, non-private synthesizer clearly has the highest maximum $\Delta_{\mathbf{x}}$ with $\Delta_{Unweighted} = 78.7$. The other non-private Countweighted synthesizer achieves a much lower maximum $\Delta_{\mathbf{x}}$ with $\Delta_{\alpha_c, Countweighted} = 11.17$. The large reduction in the Countweighted synthesizer owes to the positive correlation between by-record weights, α_c , where each α_{ci} is computed as the probability that the value for each target record is relatively isolated from that of other records used in the Countweighted synthesizer, on the one hand, with the by-record log-pseudo likelihood ratio bounds used for the DPweighted mechanism, on the other hand. DPweighted denotes our α -pseudo posterior mechanism. The two locally private synthesizers both achieve even lower maximum $\Delta_{\mathbf{x}}$: $\Delta_{\alpha, DPweighted} = 10.1$, $\Delta_{\alpha, EMweighted} = 10.2$, indicating the best risk profiles. The EMweighted mechanism was estimated by setting the scalar with a target $\epsilon_{\mathbf{x}} = 2\Delta_{\alpha, \mathbf{x}}$, the local privacy guarantee (expenditure) achieved by our DPweighted mechanism with Lipschitz $\Delta_{\alpha, \mathbf{x}}$. Our intent is to compare the utility performances between the two private mechanisms (DPweighted and EMweighted) where each achieves an equivalent privacy guarantee. It bears mention that while the DPweighted under the pseudo posterior mechanism and the EMweighted under the EM achieve similar maximum local Lipschitz bounds, which governs the local DP guarantee, the EM tends to produce notably lower risk for most records than the DPweighted mechanism. The lower record-indexed risk for EMweighted as compared to DPweighted is evident in the flattened shape of the violin plot for EMweighted. The EM sets the scalar weight based on the risk of the worst case over all records because the same level of downweighting must be applied




Figure 4: Violin plots of the distribution of the Lipschitz bounds, $\Delta_{\mathbf{x}}$'s, for synthetic data generated the four synthesizers. The corresponding maximum $\Delta_{\mathbf{x}}$ values are: $\Delta_{Unweighted} = 78.7$, $\Delta_{\alpha, DPweighted} = 10.1$, $\Delta_{EMweighted} = 10.2$, $\Delta_{\alpha_c, Countweighted} = 11.17$.

to all records in contrast with the by-record weighting under of our α -weighted pseudo posterior mechanism in DPweighted.

Figure 5 and Figure 6 show a collection of violin plots of the distribution (obtained from re-sampling) for each of the mean and the 90th quantile statistics, respectively, estimated on the synthetic data generated under each of our four synthesizers and also on the closely-held confidential (real) data for comparison, labeled “Data”. These figures allow us to compare the utility performances across our synthesizers by the examination of how well the real data distribution for each statistic is reproduced by the synthetic database for each of our synthesizers. For the synthesizers, a set of $m = 20$ synthetic databases were generated and the distribution for each statistic was estimated on each databases (under re-sampling). The resulting barycenter of the individual distributions in the Wasserstein space of measures was computed by averaging the quantiles over the M databases (Srivastava et al., 2015). Our privacy guarantees apply to *each* synthetic draw from our mechanism, so the total privacy expenditure is that for each database shown in Figure 4 multiplied by m . We compute utilities over $m = 20$ synthetic databases to fully capture the uncertainty in the synthetic data generation process from the (pseudo) posterior predictive distributions. Generating multiple synthetic databases are also standard practice in the research and practice of synthetic data using Bayesian synthesizers (Reiter and Raghunathan, 2007). We note that the distribution of each statistic for a single synthetic database is very similar.

The DPweighted synthesizer outperforms the EMweighted and Countweighted mechanisms in utility preservation. First, especially evident in Figure 6, DPweighted (the α -weighted pseudo posterior mechanism) provides better estimates than EMweighted (the




Figure 5: Violin plots of the mean estimation of the confidential CE sample and the four synthesizers.




Figure 6: Violin plots of the 90th quantile estimation of the confidential CE sample and the four synthesizers.

scalar-weighted EM). The notably deteriorated utility preservation of the EM derives from the setting that scalar weight applied to all records based on the highest risk records as earlier discussed. Since both mechanisms achieve the same maximum Lipschitz bound $\Delta_{\mathbf{x}}$, which governs the local DP guarantee, these results indicate that the EM has to compromise a large amount of the utility to achieve a similar local DP guarantee compared to the α -weighted pseudo posterior mechanism.

Second, while the non-private Unweighted synthesizer and the locally private DPweighted synthesizer provide equally good estimates for both the mean and the 90th quantile, the much greater Lipschitz bound of the Unweighted synthesizer shown in Figure 4 indicates a much worse balance for the utility-risk trade-off as compared to DPweighted. The third minor point is that the Countweighted synthesizer, albeit non-locally private, achieves only a slightly higher maximum Lipschitz bound compared to our private DPweighted synthesizer. However, its utility preservation is worse, especially evident in Figure 6 for the 90th quantile estimation.

In summary, our private DPweighted mechanism outperforms the other three synthesizers to achieve a highly satisfactory risk-utility trade-off balance. We next explore different scaling and shift configurations of (c, g) , introduced in Section 3, to sketch out the risk-utility curves for DPweighted and EMweighted.

5.3 Mapping DP Risk and Utility Curves

We conclude by applying the scaling parameter, c , and the shift parameter, g , to the distribution of weights, α , used in our α -weighted pseudo posterior mechanism in order to enumerate the risk-utility curve to support the choice of α (and, hence, $\Delta_{\alpha, \mathbf{x}}$, and $\epsilon_{\mathbf{x}}$). Having such a risk-utility curve would allow the BLS (or, more generally, the owner of the closely-held private database) to discover the setting configuration that best represents their policy goal for the level of privacy protection sought. We compare the risk-utility mapping produced by the α -weighted pseudo posterior mechanism to that of the EM, which we recall reduces to a scalar-weighted pseudo posterior under use of the log-likelihood as the utility measure. As discussed in Hu et al. (2021), applying a scaling constant, $c < 1$, will induce a compression in the distribution of the weights while apply a scaling $g < 0$ will induce a downward shift in the distribution of the record-indexed weights. We apply the scaling and shifting in a manner that uses truncation to ensure each of the resulting weights are restricted to lie in $[0, 1]$.

Each violin plot in Figure 7 presents a distribution of the 90th quantile for a synthetic database generated under a particular (c, g) configuration. The sequence of plots from left-to-right are ordered from less scaling and shifting (with a relatively higher or looser level for the privacy guarantee) to more scaling and shifting (with a relatively lower or tighter level for the privacy guarantee). The specific local sensitivity or Lipschitz value, $\Delta_{\alpha, \mathbf{x}}$, associated with each configuration are shown in Table 5.3, where we recall that the associated local privacy guarantee is $\epsilon_{\mathbf{x}} = 2 \times \Delta_{\alpha, \mathbf{x}} \times (m = 20)$, where the multiplication by $m = 20$ derives from our use of multiple posterior draws to generate multiple synthetic databases.

The accompanying Table 5.3 demonstrates a nearly 80% reduction in the level for the local DP guarantee of the α -weighted pseudo posterior mechanism over the range of configurations.




Figure 7: Violin plots of the 90th quantile estimation of: 1) the confidential CE sample; 2) the unweighted, non-private synthesizer; and overlapping violin plots of the 90th quantile estimation of the synthesizer under the α -weighted pseudo posterior mechanism compared to the synthesizer under the EM with equivalent $\Delta_{\alpha, x}$ values, for the following (c, g) configurations: 3) $(c, g) = (0.7, 0.0)$; 4) $(c, g) = (0.6, 0.0)$; 5) $(c, g) = (0.6, -0.1)$; 6) $(c, g) = (0.5, 0.0)$; 7) $(c, g) = (0.5, -0.1)$; 8) $(c, g) = (0.5, -0.3)$.

Figure 7 plots the distribution of the 90–th quantile for the generated synthetic data under each of the Unweighted (“UW”), Exponential (“EM”) and α -weighted pseudo posterior (“DP”) mechanisms at a sequence of (scaling, shifting), (c, g) , combinations. The local sensitivity/Lipschitz, $\Delta_{\alpha, x}$, is lower as one traverses left-to-right, indicating a stronger local privacy guarantee on the right-hand side. This sequence of plots demonstrates a much flatter or reduced deterioration of the 90–th quantile distribution for the DPweighted mechanism, the α -weighted pseudo posterior mechanism, as compared to the EMweighted mechanism. The superior result for DPweighted is not surprising due to the greater flexibility of DPweighted to concentrate downweighting to high-risk records versus the application of a scalar weight based on the highest risk record to all records under EMweighted.

6. Conclusion

This paper adapts the α -weighted pseudo posterior synthesizer and converts any non-private posterior synthesizer to a formally private mechanism. Our pseudo posterior mechanism provides a much higher utility than the EM for equivalent local privacy guarantee, ϵ_x , due to its surgical downweighting of high-risk records (as opposed to the scalar down-

(c, g)	$\Delta_{\alpha, \mathbf{x}}$ value
(0.7, 0.0)	10.10
(0.6, 0.0)	8.16
(0.6, -0.1)	7.30
(0.5, 0.0)	6.09
(0.5, -0.1)	5.71
(0.5, -0.3)	2.25

Table 2: Table of values of the Lipschitz bound $\Delta_{\alpha, \mathbf{x}}$, of the synthesizer under the α -weighted pseudo posterior mechanism, for a series of (c, g) configurations. $\Delta_{Unweighted} = 78.7$.

weighting imposed by the EM). The construction for the α -weighted pseudo posterior mechanism utilizes the log-pseudo likelihood to develop the local Lipschitz bound. We provide an asymptotic result on the contraction of a local Lipschitz to a global bound that guarantees an (ϵ, π) -probabilistic DP guarantee where π contracts onto 0 for n sufficiently large. We are able to increase the rate of contraction by truncating the weight, $\alpha_i^* = 0$ if the weighted log likelihood contribution, $\alpha_i \times f_\theta(x_i) > M$, where M becomes the targeted global point of contraction.

Our α -weighted pseudo posterior mechanism has the feature that it accommodates any synthesizer model formulated by the statistical agency and offers a simple weighting scheme that guarantees a pDP result. The simple weighting allows the posterior sampling scheme devised for the non-private synthesizer to be utilized for synthesis with minor modification for the α -weighted pseudo posterior mechanism.

Appendix A. Proofs for Theoretical Results in Sections 2 and 4

A.1 Proof for Theorem 5

We begin by stating an enabling result that connects the global Lipschitz bound, Δ_α , to the KL divergence between the posterior densities (given \mathbf{x} versus \mathbf{y}) from the inclusion of a database record.

Theorem 9 $\forall \mathbf{x} \in \mathcal{X}^n, \mathbf{y} \in \mathcal{X}^{n-1} : \delta(\mathbf{x}, \mathbf{y}) = 1$ and $\alpha(\cdot)$ with $\Delta_\alpha > 0$ satisfying Assumption 1,

$$\sup_{\mathbf{x} \in \mathcal{X}^n, \mathbf{y} \in \mathcal{X}^{n-1} : \delta(\mathbf{x}, \mathbf{y}) = 1} D_{KL} \left[(\xi^{\alpha(\mathbf{x})}(\cdot | \mathbf{x}) \parallel \xi^{\alpha(\mathbf{y})}(\cdot | \mathbf{y})) \right] \leq 2\Delta_\alpha, \quad (24)$$

where $D_{KL}((P \parallel Q)) = \int_{\mathcal{X}^n} \ln \frac{dP}{dQ} dP$.

Proof

$$\begin{aligned} D_{KL} \left[(\xi^{\alpha(\mathbf{x})}(\cdot | \mathbf{x}) \parallel \xi^{\alpha(\mathbf{y})}(\cdot | \mathbf{y})) \right] &= \int_{\Theta} \ln \frac{d\xi^{\alpha(\mathbf{x})}(\theta | \mathbf{x})}{d\xi^{\alpha(\mathbf{y})}(\theta | \mathbf{y})} d\xi^{\alpha(\mathbf{x})}(\theta | \mathbf{x}) \\ &= \int_{\Theta} \ln \frac{p_{\theta}^{\alpha(\mathbf{x})}(\mathbf{x})}{p_{\theta}^{\alpha(\mathbf{y})}(\mathbf{y})} d\xi^{\alpha(\mathbf{x})}(\theta | \mathbf{x}) + \int_{\Theta} \ln \frac{\phi^{\alpha(\mathbf{y})}(\mathbf{y})}{\phi^{\alpha(\mathbf{x})}(\mathbf{x})} d\xi^{\alpha(\mathbf{x})}(\theta | \mathbf{x}) \\ &\leq \int_{\Theta} \left| \ln \frac{p_{\theta}^{\alpha(\mathbf{x})}(\mathbf{x})}{p_{\theta}^{\alpha(\mathbf{y})}(\mathbf{y})} \right| d\xi^{\alpha(\mathbf{x})}(\theta | \mathbf{x}) + \int_{\Theta} \ln \frac{\phi^{\alpha(\mathbf{y})}(\mathbf{y})}{\phi^{\alpha(\mathbf{x})}(\mathbf{x})} d\xi^{\alpha(\mathbf{x})}(\theta | \mathbf{x}) \\ &\leq \Delta_\alpha + \left| \ln \frac{\phi^{\alpha(\mathbf{y})}(\mathbf{y})}{\phi^{\alpha(\mathbf{x})}(\mathbf{x})} \right| \end{aligned} \quad (25)$$

From Assumption 1, $p_{\theta}^{\alpha(\mathbf{x})}(\mathbf{x}) \leq \exp(\Delta_\alpha) p_{\theta}^{\alpha(\mathbf{y})}(\mathbf{y}), \forall \theta \in \Theta$, so

$$\phi^{\alpha(\mathbf{y})}(\mathbf{y}) = \int_{\Theta} p_{\theta}^{\alpha(\mathbf{y})}(\mathbf{y}) d\xi(\theta) \leq \exp(\Delta_\alpha) \int_{\Theta} p_{\theta}^{\alpha(\mathbf{x})}(\mathbf{x}) d\xi(\theta) = \exp(\Delta_\alpha) \phi^{\alpha(\mathbf{x})}(\mathbf{x}), \quad (26)$$

which gives

$$\sup_{\mathbf{x} \in \mathcal{X}^n, \mathbf{y} \in \mathcal{X}^{n-1} : \delta(\mathbf{x}, \mathbf{y}) = 1} D_{KL} \left[(\xi^{\alpha(\mathbf{x})}(\cdot | \mathbf{x}) \parallel \xi^{\alpha(\mathbf{y})}(\cdot | \mathbf{y})) \right] \leq 2\Delta_\alpha. \quad (27)$$

■

A.1.1 PROOF OF THEOREM 5

From Assumption 1, $\frac{p_{\theta}^{\alpha(\mathbf{x})}(\mathbf{x})}{p_{\theta}^{\alpha(\mathbf{y})}(\mathbf{y})} \leq \exp(\Delta_\alpha)$. From Theorem 9, we show $\phi^{\alpha(\mathbf{y})}(\mathbf{y}) \leq \exp(\Delta_\alpha) \phi^{\alpha(\mathbf{x})}(\mathbf{x})$.

Then, $\forall \mathbf{x} \in \mathcal{X}^n$ and for each $\mathbf{x}, \forall \mathbf{y} \in \mathcal{X}^{n-1} : \delta(\mathbf{x}, \mathbf{y} = 1)$,

$$\begin{aligned} \xi^{\alpha(\mathbf{x})}(B | \mathbf{x}) &= \frac{\int_B \frac{p_{\theta}^{\alpha(\mathbf{x})}(\mathbf{x})}{p_{\theta}^{\alpha(\mathbf{y})}(\mathbf{y})} p_{\theta}^{\alpha(\mathbf{y})}(\mathbf{y}) d\xi(\theta)}{\phi^{\alpha(\mathbf{y})}(\mathbf{y})} \cdot \frac{\phi^{\alpha(\mathbf{y})}(\mathbf{y})}{\phi^{\alpha(\mathbf{x})}(\mathbf{x})} \\ &\leq \exp(2\Delta_\alpha) \xi^{\alpha(\mathbf{y})}(B | \mathbf{y}). \end{aligned} \quad (28)$$

A.2 Proof for Lemma 6

$$\begin{aligned}
 P^{\alpha(\mathbf{x})}(\zeta \in C \mid \mathbf{x}) &= \int P(\zeta \in C \mid \mathbf{x}, \theta) d\xi^{\alpha(\mathbf{x})}(\theta \mid \mathbf{x}) \\
 &= \int P(\zeta \in C \mid \theta) d\xi^{\alpha(\mathbf{x})}(\theta \mid \mathbf{x}) \\
 &= \int P(\zeta \in C \mid \theta) \frac{d\xi^{\alpha(\mathbf{x})}(\theta \mid \mathbf{x})}{d\xi^{\alpha(\mathbf{y})}(\theta \mid \mathbf{y})} d\xi^{\alpha(\mathbf{y})}(\theta \mid \mathbf{y}) \\
 &\leq e^\epsilon \int P(\zeta \in C \mid \theta) d\xi^{\alpha(\mathbf{y})}(\theta \mid \mathbf{y}) \\
 &= e^\epsilon P^{\alpha(\mathbf{y})}(\zeta \in C \mid \mathbf{y}).
 \end{aligned} \tag{29}$$

A.3 Proof of Theorem 8

Let us the define the following subset of $\theta \in \Theta$,

$$U_n = \left\{ \theta \in \Theta : \left[(1 - \alpha_m) D_{\theta_0, \alpha}^{(n_A)}(\theta, \theta^*) + (1 - \alpha^{(n)}) D_{\theta_0, 1^-}^{(n_Q)}(\theta, \theta^*) \right] \geq (D + 3t)n\tau_n^2 \right\},$$

which is the restricted set for which we will bound the pseudo posterior distribution, $\xi^\alpha(U_n \mid \mathbf{x})$, from above to achieve the result of Theorem 8. We begin with the statement and proof of Lemma 10 that extends Lemma 8.1 of Ghosal et al. (2000) to our α -pseudo posterior in order to provide a concentration inequality to probabilistically (in P_{θ_0} -probability) bound the denominator of the α -pseudo posterior distribution, $\xi^\alpha(U_n \mid \mathbf{x})$, from below.

A.3.1 ENABLING LEMMA

Lemma 10 (*Concentration Inequality*) *Suppose Assumption 3 holds. Define $\alpha_m = \max_{i \in A_n} \alpha_i$ and $\alpha_l = \min_{i \in A_n} \alpha_i$. For every $\tau_n > 0$ and measure Π on the set $B_n(\theta^*, \xi; \theta_0)$, we have for every $C_1^* = \sqrt{2 + C_1^2 + C_3^2}$, and n sufficiently large,*

$$P_{\theta_0} \left\{ \int_{\theta \in B_n} e^{-r_n, \alpha(\theta, \theta^*)} \xi(d\theta) \leq e^{-\alpha_m(D+t)n\tau_n^2} \right\} \leq \frac{(1 + \alpha_l^2)(C_1^*)^2}{\alpha_m^2} \times \frac{1}{(D + t - 1)^2 n \tau_n^2}, \tag{30}$$

where the above probability is taken with the respect to P_{θ_0} .

Proof The proof follows that of Savitsky and Toth (2016) by bounding the probability expression on left-hand side of Equation (30). We construct an α -weighted empirical distribution that we will need for the proof with,

$$\mathbb{P}_{n, \alpha} = \frac{1}{n} \sum_{i=1}^n \alpha_i \delta(x_i), \tag{31}$$

where $\delta(x_i)$ denotes the Dirac delta function with probability mass 1 at x_i . We construct the associated scaled and centered empirical process, $\mathbb{G}_{n, \alpha} = \sqrt{n}(\mathbb{P}_{n, \alpha} - P_{\theta_0})$. The usual equally-weighted empirical distribution, $\mathbb{P}_n = \frac{1}{n} \sum_{i=1}^n \delta(x_i)$ and associated, $\mathbb{G}_n =$

$\sqrt{n}(\mathbb{P}_n - P_{\theta_0})$ may be viewed as special cases. We may define the associated expectation functionals with respect to the α -weighted empirical distribution by $\mathbb{P}_{n,\alpha}g = \frac{1}{n} \sum_{i=1}^n \alpha_i g(x_i)$.

Using Jensen's inequality,

$$\begin{aligned} & \log \int_{\theta \in B_n} \prod_{i=1}^n \left[\frac{p_{\theta_i}}{p_{\theta_i^*}}(X_i) \right]^{\alpha_i} \xi(d\theta) \\ & \geq \sum_{i=1}^n \int_{\theta \in B_n} \alpha_i \log \frac{p_{\theta_i}}{p_{\theta_i^*}} \xi(d\theta) \\ & = n \mathbb{P}_{n,\alpha} \int_{\theta \in B_n} \log \frac{p_{\theta}}{p_{\theta^*}} \xi(d\theta) \end{aligned} \quad (32)$$

We may use the above to now bound the left-hand side of Equation (30)

$$P_{\theta_0} \left\{ \int_{\theta \in B_n} e^{-r_{n,\alpha}(\theta, \theta^*)} \xi(d\theta) \leq e^{-\alpha_m(D+t)n\tau_n^2} \right\} \quad (33a)$$

$$\leq P_{\theta_0} \left\{ n \mathbb{P}_{n,\alpha} \int_{\theta \in B_n} \log \frac{p_{\theta}}{p_{\theta^*}} \xi(d\theta) \leq -\alpha_m(D+t)n\tau_n^2 \right\} \quad (33b)$$

$$= P_{\theta_0} \left\{ \mathbb{G}_{n,\alpha} \int_{\theta \in B_n} \log \frac{p_{\theta}}{p_{\theta^*}} \xi(d\theta) \leq -\alpha_m(D+t)n\tau_n^2 - \sqrt{n}P_{\theta_0} \log \frac{p_{\theta}}{p_{\theta^*}} \xi(d\theta) \right\} \quad (33c)$$

$$\leq P_{\theta_0} \left\{ \mathbb{G}_{n,\alpha} \int_{\theta \in B_n} \log \frac{p_{\theta}}{p_{\theta^*}} \xi(d\theta) \leq -\alpha_m(D+t)\sqrt{n}\tau_n^2 - \sqrt{n}\tau_n^2 \right\} \quad (33d)$$

$$= P_{\theta_0} \left\{ \mathbb{G}_{n,\alpha} \int_{\theta \in B_n} \log \frac{p_{\theta}}{p_{\theta^*}} \xi(d\theta) \leq -\alpha_m(D+t-1)\sqrt{n}\tau_n^2 \right\}, \quad (33e)$$

where the bound in Equation (33d) uses the prior mass result from Assumption 3. We proceed to use Chebyshev to bound the resultant probability, as follows:

$$\begin{aligned} & P_{\theta_0} \left\{ \mathbb{G}_{n,\alpha} \int_{\theta \in B_n} \log \frac{p_{\theta}}{p_{\theta^*}} \xi(d\theta) \leq -\alpha_m(D+t-1)\sqrt{n}\tau_n^2 \right\} \\ & \leq \frac{\int_{\theta \in B_n} \left[\mathbb{E}_{P_{\theta_0}} \left(\mathbb{G}_{n,\alpha} \log \frac{p_{\theta}}{p_{\theta^*}} \right)^2 \right] \xi(d\theta)}{\alpha_m^2(D+t-1)^2 n \tau_n^4}, \end{aligned} \quad (34)$$

where we have applied Fubini to the right side of Equation (34) to move the expectation through the integral. We now proceed to further bound the expression in brackets on the

right-hand side of Equation (34) from above. We may decompose the expectation, as follows

$$\mathbb{E}_{P_{\theta_0}} \left(\mathbb{G}_{n,\alpha} \log \frac{p_\theta}{p_{\theta^*}} \right)^2 \leq n \mathbb{E}_{P_{\theta_0}} \left(\mathbb{P}_{n,\alpha} \log \frac{p_\theta}{p_{\theta^*}} - \mathbb{P}_n \log \frac{p_\theta}{p_{\theta^*}} \right)^2 + \mathbb{E}_{P_{\theta_0}} \left(\mathbb{G}_{n,\alpha} \log \frac{p_\theta}{p_{\theta^*}} \right)^2 \quad (35)$$

We first bound the second term on the right,

$$\mathbb{E}_{P_{\theta_0}} \left(\mathbb{G}_{n,\alpha} \log \frac{p_\theta}{p_{\theta^*}} \right)^2 \quad (36a)$$

$$\leq \mathbb{E}_{P_{\theta_0}} \left(\sqrt{n} \mathbb{P}_{n,\alpha} \log \frac{p_\theta}{p_{\theta^*}} \right)^2 \quad (36b)$$

$$\leq \mathbb{E}_{P_{\theta_0}} \left(\frac{1}{\sqrt{n}} \sum_{i=1}^n \log \frac{p_\theta}{p_{\theta^*}} \right)^2 \quad (36c)$$

$$\leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{P_{\theta_0}} \left(\log \frac{p_\theta}{p_{\theta^*}} \right)^2 \quad (36d)$$

$$\leq \frac{1}{n} \times n \tau_n^2 = \tau_n^2, \quad (36e)$$

where we use independence of the X_i to establish the fourth equation and Assumption 3 to achieve the fifth equation.

We proceed to further simplify the bound in the first term on the right in Equation (35):

$$n\mathbb{E}_{P_{\theta_0}} \left(\mathbb{P}_{n,\alpha} \log \frac{p_\theta}{p_{\theta^*}} - \mathbb{P}_n \log \frac{p_\theta}{p_{\theta^*}} \right)^2 \quad (37a)$$

$$= n\mathbb{E}_{P_{\theta_0}} \left(\frac{1}{n} \sum_{i=1}^n (\alpha_i - 1) \log \frac{p_{\theta_i}}{p_{\theta_i^*}} \right)^2 \quad (37b)$$

$$= \frac{1}{n} \sum_{i,j=1}^n \mathbb{E}_{P_{\theta_0}} \left[(\alpha_i - 1) (\alpha_j - 1) \log \frac{p_{\theta_i}}{p_{\theta_i^*}} (X_i) \log \frac{p_{\theta,j}}{p_{\theta^*,j}} (X_j) \right] \quad (37c)$$

$$= \frac{1}{n} \sum_{i=j=1}^n \mathbb{E}_{P_{\theta_0}} \left[(\alpha_i - 1)^2 \log \frac{p_{\theta_i}}{p_{\theta_i^*}} (X_i)^2 \right] \quad (37d)$$

$$+ \frac{1}{n} \sum_{i \neq j=1}^n \mathbb{E}_{P_{\theta_0}} \left| \left[(\alpha_i - 1) (\alpha_j - 1) \log \frac{p_{\theta_i}}{p_{\theta_i^*}} (X_i) \log \frac{p_{\theta,j}}{p_{\theta^*,j}} (X_j) \right] \right|$$

$$\leq \frac{1}{n} \left\{ (1 - \alpha_l)^2 \sum_{i \neq j=1}^n \mathbb{E}_{P_{\theta_0}} \left[\log \frac{p_{\theta_i}}{p_{\theta_i^*}} (X_i)^2 \right] \right\} + \frac{1}{n} (1 - \alpha_l)^2 \sum_{i \neq j \in A_n} \left| \mathbb{E}_{P_{\theta_0}} \log \frac{p_{\theta_i}}{p_{\theta_i^*}} (X_i) \log \frac{p_{\theta,j}}{p_{\theta^*,j}} (X_j) \right| \quad (37e)$$

$$+ \frac{1}{n} (1 - \alpha^{(n)})^2 \sum_{i \neq j \in Q_n} \left| \mathbb{E}_{P_{\theta_0}} \log \frac{p_{\theta_i}}{p_{\theta_i^*}} (X_i) \log \frac{p_{\theta,j}}{p_{\theta^*,j}} (X_j) \right| \leq \frac{1}{n} \left\{ (1 - \alpha_l)^2 n \tau_n^2 \right\} + \frac{1}{n} (1 - \alpha_l)^2 (C_1^2 n - C_1 \sqrt{n}) \tau_n^2 + n_Q \frac{C_3^2 \tau_n^2}{n_Q} \quad (37f)$$

$$= \left\{ (1 - \alpha_l)^2 \tau_n^2 \right\} + (1 - \alpha_l)^2 C_1^2 \tau_n^2 + C_3^2 \tau_n^2, \quad (37g)$$

for sufficiently large n . The bound in Equation (37f) results from the restriction of θ to $B_n(\theta^*, \eta; \theta_0)$ and also from Assumption 4 that regulates the growth of the number of $\alpha_i < 1^-$ and the magnitude of $(1 - \alpha^{(n)})$.

We may now bound the expectation on the right-hand size of Equation (34),

$$\mathbb{E}_{P_{\theta_0}} \left(\mathbb{G}_{n,\alpha} \log \frac{p_\theta}{p_{\theta^*}} \right)^2 \leq \left\{ (1 - \alpha_l)^2 \tau_n^2 \right\} (1 - \alpha_l)^2 C_1^2 \tau_n^2 + \tau_n^2 \quad (38a)$$

$$\leq \left\{ (1 - 2\alpha_l + \alpha_l^2) \tau_n^2 + (1 - 2\alpha_l + \alpha_l^2) C_1^2 \tau_n^2 + C_3^2 \eta n^2 + \tau_n^2 \right\} \quad (38b)$$

$$\leq (2 + C_1^2 + C_3^2) \tau_n^2 + (1 + C_1^2) \alpha_l^2 \tau_n^2 \leq (1 + \alpha_l)^2 (C_1^*)^2 \tau_n^2 \quad (38c)$$

for n sufficiently large, where we set $C_1^* := \sqrt{C_1^2 + C_3^2 + 2}$. This concludes the proof. \blacksquare

A.3.2 PROOF OF THEOREM 8

We begin by constructing the α -pseudo posterior distribution on the set, U_n ,

$$\xi^\alpha(U_n | \mathbf{x}) = \frac{\int_{U_n} e^{-r_{n,\alpha}(\theta, \theta^*)} \xi(d\theta)}{\int_{\Theta} e^{-r_{n,\alpha}(\theta, \theta^*)} \xi(d\theta)}. \quad (39)$$

We next bound the numerator from above in P_{θ_0} -probability.

$$\mathbb{E}_{P_{\theta_0}} \int_{U_n} e^{-r_{n,\alpha}(\theta, \theta^*)} \xi(d\theta) \quad (40a)$$

$$= \int_{U_n} A_{\theta_0, \alpha}^{(n)}(\theta, \theta^*) \xi(d\theta) \quad (40b)$$

$$= \int_{U_n} e^{-\sum_{i=1}^n (1-\alpha_i) D_{\theta_0, \alpha, i}} \xi(d\theta) \quad (40c)$$

$$\leq \int_{U_n} e^{-(1-\alpha_m) \sum_{i \in A_n} D_{\theta_0, \alpha, i} - (1-\alpha^{(n)}) \sum_{i \in Q_n} D_{\theta_0, 1^-, i}} \xi(d\theta) \quad (40d)$$

$$\leq e^{-(D+3t)n\tau_n^2}, \quad (40e)$$

where we use Fubini to switch the order of expectation and integration in Equation (40b). We achieve the bound in Equation (40d) since $D_{\theta_0, \alpha, i} > 0$, $\forall i \in (1, \dots, n)$ and Bhattacharya et al. (2019) shows that $D_{\theta_0, 1^-}^{(n)}(\theta, \theta^*)$ is finite and contracts on the KL divergence. The final bound uses the definition of U_n .

We proceed to use the Markov inequality and the definition for U_n to achieve the numerator bound with respect to P_{θ_0} -probability,

$$P_{\theta_0} \left\{ \int_{U_n} e^{-r_{n,\alpha}(\theta, \theta^*)} \xi(d\theta) \geq e^{-(D+2t)n\tau_n^2} \right\} \quad (41a)$$

$$\leq \frac{e^{-(D+3t)n\tau_n^2}}{e^{-(D+2t)n\tau_n^2}} = e^{-tn\tau_n^2} \leq \frac{(1 + \alpha_l^2)(C_1^*)^2}{\alpha_m^2(D-1+t)^2n\tau_n^2}. \quad (41b)$$

We, next, turn to bounding the denominator of Equation (39), from below. Since,

$$\int_{\theta \in \Theta} e^{-r_{n,\alpha}(\theta, \theta^*)} \xi(d\theta) \geq \int_{\theta \in B_n} e^{-r_{n,\alpha}(\theta, \theta^*)} \xi(d\theta),$$

we may use the result of Lemma 10 in,

$$P_{\theta_0} \left\{ \int_{\theta \in \Theta} e^{-r_{n,\alpha}(\theta, \theta^*)} \xi(d\theta) \geq e^{-\alpha_m(D+t)n\tau_n^2} \right\} > 1 - \frac{(1 + \alpha_l^2)(C_1^*)^2}{\alpha_m^2(D-1+t)^2n\tau_n^2}. \quad (42)$$

Finally, combining the results of Equations (39), (41) and (42): With probability at least $1 - [2/(D+t-1)^2 n \tau_n^2 \times (1 + \alpha_l^2 (C_1^*)^2)/\alpha_m^2]$,

$$\begin{aligned} \xi^\alpha \left(\left[(1 - \alpha_m) D_{\theta_0, \alpha}^{(n_A)}(\theta, \theta^*) + (1 - \alpha^{(n)}) D_{\theta_0, 1^-}^{(n_Q)}(\theta, \theta^*) \right] \geq (D + 3t) n \tau_n^2 | \mathbf{x} \right) &\leq \\ &e^{-(D+2t)n\tau_n^2} e^{\alpha_m(D+t)n\tau_n^2} \\ &\leq e^{-tn\tau_n^2} \end{aligned}$$

Appendix B. Unweighted, Non-private Synthesizer

Our description of the unweighted, non-private synthesizer follows closely of that in Hu et al. (2021). To simulate partially synthetic data for the CE sample, where only the sensitive, continuous family income variable is synthesized, we propose using a flexible, parametric finite mixture synthesizer.

Equation (43) and Equation (44) present the first two levels of the hierarchical parametric finite mixture synthesizer: y_i is the logarithm of the family income for CU i , and \mathbf{x}_i is the $R \times 1$ predictor vector for CU i . The finite mixture utilizes a hyperparameter for the maximum number of mixture components (i.e., clusters), K , that is to set to be over-determined to permit the flexible clustering of CUs. A subset of CUs that are assigned to cluster, k , employ the same generating parameters for y , (β_k^*, σ_k^*) , that we term a ‘‘location’’. Locations, (β^*, σ^*) , and the $n \times 1$ vector of cluster indicators, $z_i \in (1, \dots, K)$, are all sampled for each CU, $i \in (1, \dots, n)$.

$$y_i | \mathbf{X}_i, z_i, \mathbf{B}^*, \sigma^* \sim \text{Normal}(y_i | \mathbf{x}_i' \beta_{z_i}^*, \sigma_{z_i}^*), \quad (43)$$

$$z_i | \pi \sim \text{Multinomial}(1; \pi_1, \dots, \pi_K), \quad (44)$$

where the $K \times R$ matrix of regression locations, $\mathbf{B}^* = (\beta_1^*, \dots, \beta_K^*)'$, denote cluster-indexed regression coefficients for R predictors. The (π_1, \dots, π_K) are, in turn, assigned a sparsity inducing Dirichlet distribution with hyperparameters specified as α/K for $\alpha \in \mathbb{R}^+$. We next describe our prior specification.

We induce sparsity in the number of clusters with,

$$(\pi_1, \dots, \pi_K) \sim \text{Dirichlet} \left(\frac{\alpha}{K}, \dots, \frac{\alpha}{K} \right), \quad (45)$$

$$\alpha \sim \text{Gamma}(a_\alpha, b_\alpha). \quad (46)$$

We specify multivariate Normal priors for each regression coefficient vector of coefficient locations, β_k^* ,

$$\beta_k^* \stackrel{\text{iid}}{\sim} \text{MVN}_R(\mathbf{0}, \text{diag}(\sigma_\beta) \times \Omega_\beta \times \text{diag}(\sigma_\beta)), \quad (47)$$

where the $R \times R$ correlation matrix, Ω_β , receives a uniform prior over the space of $R \times R$ correlation matrices, and each component of σ_β receives a Student-t prior with 3 degrees of freedom,

$$\sigma_k^* \stackrel{\text{iid}}{\sim} t(3, 0, 1). \quad (48)$$

We proceed to describe how to generate partially synthetic data for the CE sample. To implement the finite mixture synthesizer, we first generate sample values of $(\boldsymbol{\pi}^{(l)}, \beta^{*,(l)}, \sigma^{*,(l)})$

from the posterior distribution at MCMC iteration l . Second, for CU i , we generate cluster assignments, $z_i^{(l)}$, from its full conditional posterior distribution given in Hu et al. (2021) using the posterior samples of $\boldsymbol{\pi}^{(l)}$. Lastly, we generate synthetic family income for CU i , $y_i^{*,(l)}$, from Equation (43) given \boldsymbol{x}_i , and samples of $z_i^{(l)}$, $\boldsymbol{\beta}^{*,(l)}$ and $\boldsymbol{\sigma}^{*,(l)}$. We perform these draws for all n CUs, and obtain a partially synthetic database, $\boldsymbol{Z}^{(l)}$ at MCMC iteration l . We repeat this process for m times, creating m independent partially synthetic databases $\boldsymbol{Z} = (\boldsymbol{Z}^{(1)}, \dots, \boldsymbol{Z}^{(m)})$.

References

- J. Abowd and L. Vilhuber. How protective are synthetic data? In J. Domingo-Ferrer and Y. Saygin, editors, *Privacy in Statistical Databases*, volume 5262 of *Lecture Notes in Computer Science*, pages 239–246. Springer, 2008.
- D. An and R. J. A. Little. Multiple imputation: an alternative to top coding for statistical disclosure control. *Journal of the Royal Statistical Society, Series A*, 170:923–940, 2007.
- A. Bhattacharya, D. Pati, and Y. Yang. Bayesian fractional posteriors. *The Annals of Statistics*, 47(1):39–66, 2019.
- C. M. Bowen and F. Liu. Comparative study of differentially private data synthesis methods. *Statistical Science*, 35:280–307, 2020.
- C. Dimitrakakis, B. Nelson, Z. Zhang, A. Mitrokotsa, and B. I. P. Rubinstein. Differential privacy for bayesian inference through posterior sampling. *J. Mach. Learn. Res.*, 18(1): 343–381, January 2017. ISSN 1532-4435.
- C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the Third Conference on Theory of Cryptography, TCC’06*, pages 265–284, Berlin, Heidelberg, 2006. Springer-Verlag. ISBN 3-540-32731-2, 978-3-540-32731-8. doi: 10.1007/11681878_14.
- S. Ghosal, J. K. Ghosh, and A. W. Van Der Vaart. Convergence rates of posterior distributions. *Ann. Statist.*, pages 500–531, 2000.
- J. Hu, T. D. Savitsky, and M. R. Williams. Risk-Efficient Bayesian Data Synthesis for Privacy Protection. *Journal of Survey Statistics and Methodology*, 04 2021. ISSN 2325-0984. doi: 10.1093/jssam/smab013. URL <https://doi.org/10.1093/jssam/smab013>.
- R. J. A. Little. Statistical analysis of masked data. *Journal of Official Statistics*, 9:407–426, 1993.
- A. Machanavajjhala, D. Kifer, J. Abowd, J. Gehrke, and L. Vilhuber. Privacy: Theory meets practice on the map. In *2008 IEEE 24th International Conference on Data Engineering*, pages 277–286, 2008.
- A. Machanavajjhala, D. Kifer, J. Abowd, J. Gehrke, and L. Vilhuber. Privacy: Theory meets practice on the map. In *Proceedings of the 2008 IEEE 24th International Conference on Data Engineering*, pages 277–286. IEEE Computer Society, 2008.

- D. McClure and J. P. Reiter. Differential privacy and statistical disclosure risk measures: An illustration with binary synthetic data. *Transactions on Data Privacy*, 5:535–552, 2012.
- M. McSherry and K. Talwar. Mechanism design via differential privacy. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*, pages 94–103, 2007.
- J. P. Reiter and T. E. Raghunathan. The multiple adaptations of multiple imputation. *Journal of the American Statistical Association*, 102:1462–1471, 2007.
- D. B. Rubin. Discussion statistical disclosure limitation. *Journal of Official Statistics*, 9: 461–468, 1993.
- T. D. Savitsky and D. Toth. Bayesian Estimation Under Informative Sampling. *Electronic Journal of Statistics*, 10(1):1677–1708, 2016.
- J. Snoké and A. Slavkovic. pMSE mechanism: Differentially private synthetic data with maximal distributional similarity. In J. Domingo-Ferrer and F. Montes, editors, *Privacy in Statistical Databases*, volume 11126 of *Lecture Notes in Computer Science*, pages 138–159. Springer, 2018.
- S. Srivastava, V. Cevher, Q. Dinh, and D. Dunson. WASP: Scalable Bayes via barycenters of subset posteriors. In *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics*, pages 912–920, 2015.
- Yu-Xiang Wang, Stephen E. Fienberg, and Alex Smola. Privacy for free: Posterior sampling and stochastic gradient monte carlo, 2015.
- L. Wasserman and S. Zhou. A statistical framework for differential privacy. *Journal of the American Statistical Association*, 105:375–389, 2010.
- Z. Zhang, B. I. P. Rubinstein, and C. Dimitrakakis. On the differential privacy of Bayesian inference. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*, pages 2365–2371. AAAI, 2016.