Mechanisms for Global Differential Privacy under Bayesian Data Synthesis

Terrance D. Savitsky [1] Matthew R. Williams [2]
Monika (Jingchen) Hu [3]

[1] U.S. Bureau of Labor Statistics (Office of Survey Methods Research)

[2] RTI International

[3] Vassar College (Mathematics and Statistics Department)

FCSM Conference, October 23, 2024

Outline

- Background
- Five microdata synthesizers
- Simulation study and SDR application
- Application
- Concluding remarks

Replicate data $\mathbf{x}^*$ given (|) observed data, $\mathbf{x}$ (Hu, 2019)

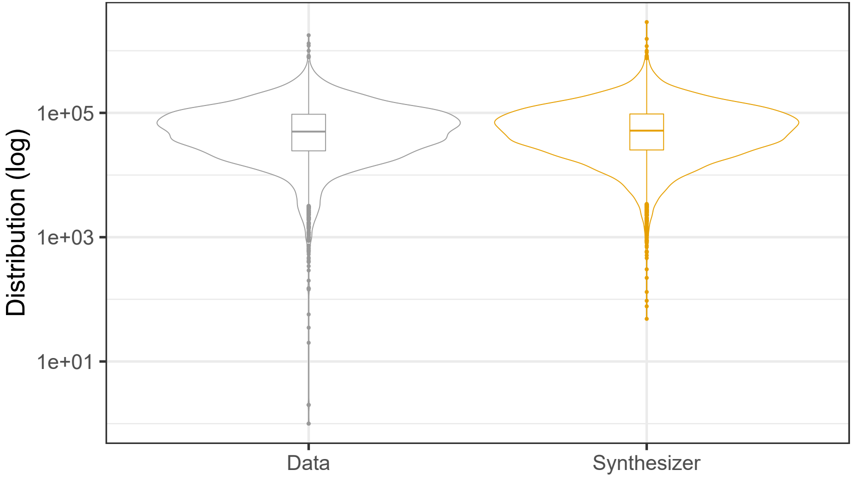

Differential Privacy under $\mathcal{M} = \xi(\theta \mid \mathbf{x})$ (Dimitrakakis et al., 2017)

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

- ϵ bounds the **change** in the probability measure ξ
  - from the inclusion of a **single record** $\delta(\mathbf{x}, \mathbf{x}^{'}) = 1$,
  - over **all possible outcomes**, $B \in \beta_\Theta$ – sets in the space of measurable sets of Θ.
  - over **all possible data sets** $\mathbf{x}, \mathbf{x}^{'} \in \mathcal{X}^n$ of size n.

Outline

Synthesizer #1: Weighted pseudo posterior (Savitsky et al., 2022)

- The mechanism $\mathcal{M}$ is the pseudo posterior:

$$\xi^{\boldsymbol{\alpha}(\boldsymbol{y})}(\theta \mid \boldsymbol{x}) \propto \prod_{i=1}^{n} p(x_i \mid \theta)^{\alpha_i} \times \xi(\theta) \tag{1}$$

- Fit any Bayesian synthesizer to confidential data $\boldsymbol{x}$
- Formulate privacy weight α_i and estimate a pseudo posterior
  - Downweight each likelihood by $\alpha_i \in [0, 1]$
  - Higher disclosure risk, lower α_i
- Calculate Lipschitz where $f_\theta(\boldsymbol{x})$ is the log-likelihood:

$$\sup_{\mathbf{x},\mathbf{x}' \in \mathcal{X}^n : \delta(\mathbf{x},\mathbf{x}')=1} \sup_{\theta \in \Theta} |\alpha(\mathbf{x}) f_\theta(\mathbf{x}) - \alpha(\mathbf{x}^{'}) f_\theta(\mathbf{x}^{'})| \leq \Delta_\alpha$$

$$\max_{i \in 1,\ldots,n} \sup_{\theta \in \Theta} |\alpha_i \times f_\theta(x_i)| \leq \Delta_{\boldsymbol{\alpha},\mathbf{x}}$$

- Each posterior draw with $\epsilon_{\boldsymbol{x}} = 2\Delta_{\boldsymbol{\alpha},\mathbf{x}}$ produces one synthetic $\boldsymbol{x}^*$

Synthesizer #2: Weighted-e pseudo posterior (Savitsky et al., 2022)

- In addition to the observation-indexed privacy weight α_i in Weighted
- Savitsky *et al.* (2022) introduce a truncation of each weight: If a record's log-likelihood contribution $> \epsilon/2$, set final weight $\alpha_i^* = 0$
- Truncation induces a rapid contraction of $\epsilon_{\boldsymbol{x}}$ to global ϵ
- As with the Weighted, Weighted-e also achieves aDP

Asymptotic Differential Privacy

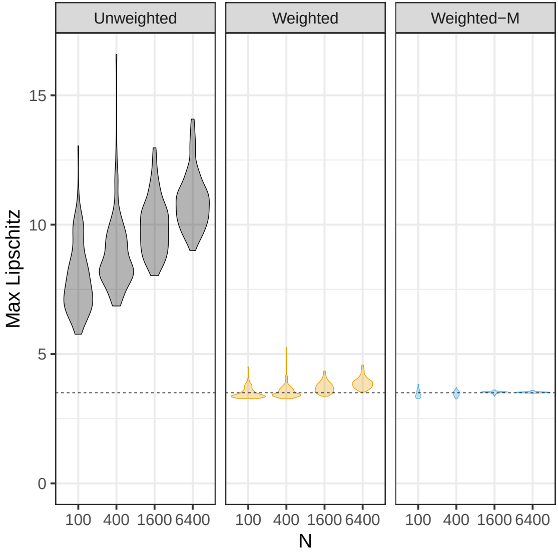

Weighted-e Produces Slightly Worse Utility

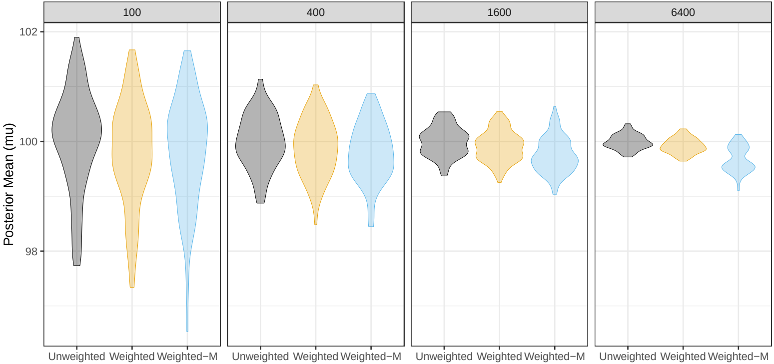

Figure: Distributions of the average of mean parameter μ for each of sample size $(100, 400, 1600, 6400)$ from $R = 100$ realizations.

Synthesizers #3 & #4: Censored (likelihood)

- Censor the log-likelihood at a target threshold, $\epsilon/2$

$$p_c^{\boldsymbol{\alpha}}(x_i \mid \theta) = \begin{cases} \exp(\epsilon/2), & p(x_i \mid \theta)^{\boldsymbol{\alpha}} > \exp(\epsilon/2), \\ \exp(-\epsilon/2), & p(x_i \mid \theta)^{\boldsymbol{\alpha}} < \exp(-\epsilon/2), \\ p(x_i \mid \theta)^{\boldsymbol{\alpha}}, & \text{otherwise}, \end{cases}$$

  for use in

$$\xi_c^{\boldsymbol{\alpha}}(\theta \mid \mathbf{x}) \propto \prod_{i=1}^{n} p_c^{\boldsymbol{\alpha}}(x_i \mid \theta)\xi(\theta),$$

- Embeds weights inside the censoring mechanism; labeled as Censor_w and it achieves DP (Hu et al., 2022)
- Censored unweighted (posterior) the censoring mechanism; labeled as Censor_uw and it achieves DP
- Censoring offers a practical, low-dimensional alternative to truncating the parameter space to achieve DP

Synthesizer #5: Perturbed histogram (Wasserman and Zhou, 2010)

- Under the assumption of a bounded and continuous univariate variable
  1. Discretize it into a histogram with a selected number of bins
  2. Adding Laplace noise to the histogram to achieve DP
  3. Simulate microdata from the private histogram under DP
- Labeled as PH and it achieves DP if bounded data

Outline

Simulation design

- Monte Carlo simulation under repeated sampling
- For $r = 1, \cdots, R = 100$, simulate a local database $\mathbf{x}_r$ of size $n = 2000$ from Beta(0.5, 3)
- For each local database $\mathbf{x}_r$, we fit and create a synthetic dataset for each of the synthesizers
  1. Weighted: asymptotic DP (aDP)
  2. Weighted-e: aDP with faster convergence
  3. Censor_w: DP
  4. Censor_uw: DP
  5. Perturbed histogram: DP if bounded data
- We experiment with $\epsilon \in \{5, 4, 3\}$

Simulation: privacy comparison results

- Violin plots of Lipschitz bounds over $R = 100$ replicates
- A dashed horizontal line at $\epsilon/2$ is included in each panel

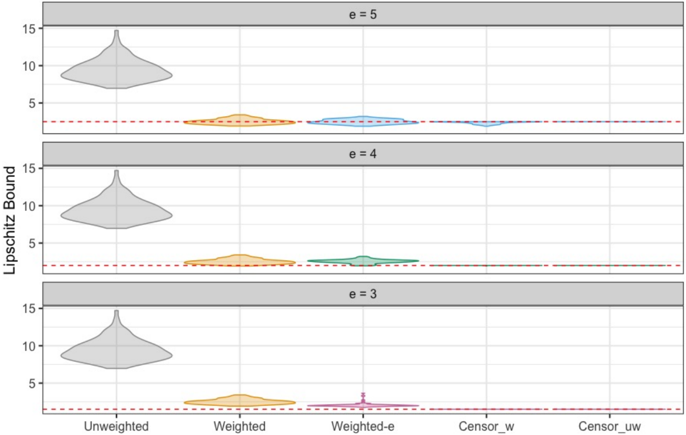

Simulation: global utility comparison results

- Violin plots of ECDF utility - maximum record-level difference over $R = 100$ replicates
- Smaller the ECDF, higher the utility

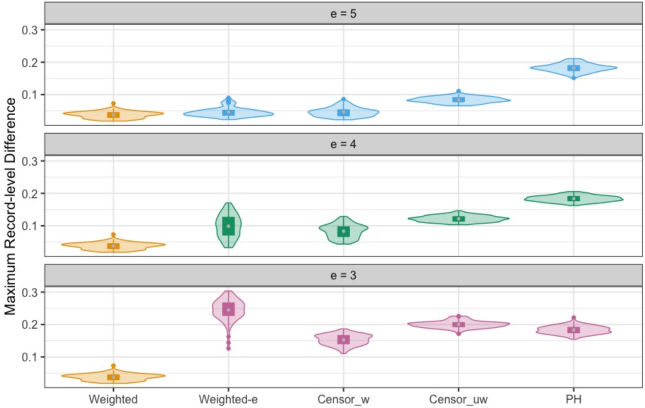

Simulation: analysis-specific utility comparison results

- Violin plots of median over $R = 100$ replicates at $\epsilon = 5$
- A dashed horizontal line at the analytical median from Beta(0.5, 3) is included

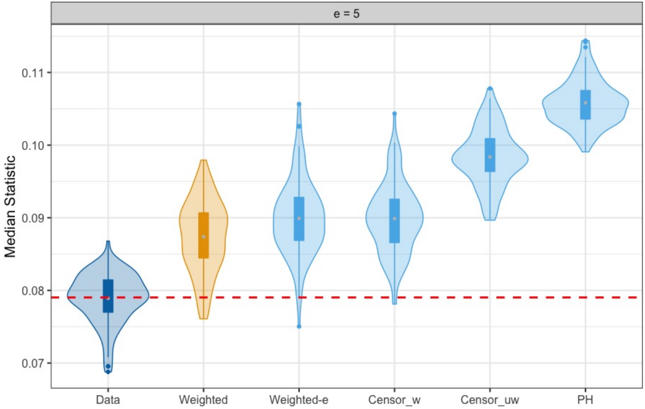

Fine tuning with downscaling - $\tilde{\alpha}_i = c_1 \times \alpha_i,\ c_1 < 1$

- Violin plots of Lipschitz bounds over $R = 100$ replicates without downscaling (top) and with downscaling (bottom) at $\epsilon = 4$
- A dashed horizontal line at $\epsilon/2$ is included

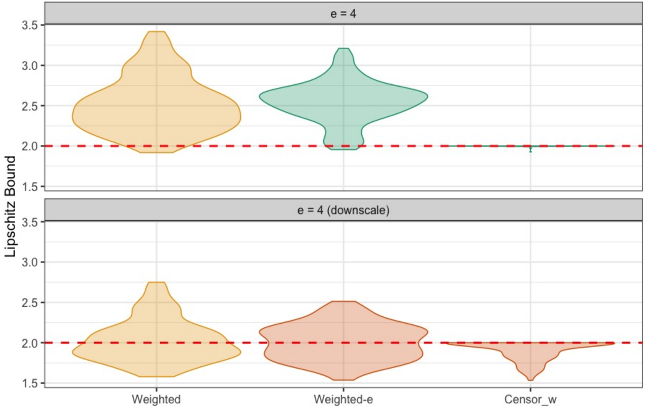

Improved utility with downscaling

- Violin plots of ECDF utility - average record-level squared difference over $R = 100$ replicates without downscaling (top) and with downscaling (bottom) at $\epsilon = 4$

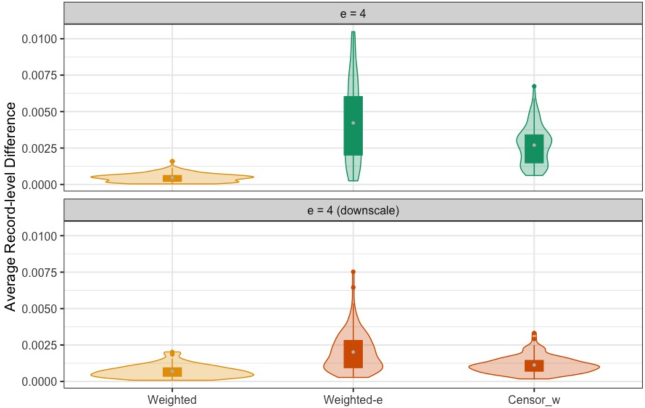

Outline

Application

- Survey of Doctoral Recipients public use file from 2017
- $n = 1601$ respondents who have positions at a 4-year college or university in the field of mathematics and statistics
- Variables: salary, gender, age, and the number of working weeks
- Synthesizer: a beta regression after transformation of salary
- We fit and create a synthetic dataset for each of the synthesizers
  1. Weighted: asymptotic DP (aDP)
  2. Weighted-e: aDP with faster convergence
  3. Censor_w: DP
  4. Censor_uw: DP
  5. Perturbed histogram: DP if bounded data
- Target $\epsilon = 5$

SDR application: utility and privacy results

	Data	Weighted	Weighted-e	Censor_w	Censor_uw	PH
Lipschitz	NA	2.62	2.61	2.50	2.50	NA
Privacy ϵ	NA	5.24	5.22	5.00	5.00	5.00
max-ECDF	NA	**0.0656**	0.1020	<u>0.0968</u>	0.1350	0.1310
avg-ECDF	NA	**0.0011**	<u>0.0025</u>	0.0026	0.0039	0.0057
Mean	91019	<u>92994</u>	**89525**	88581	88840	93654
Median	80000	**80135**	<u>76451</u>	75642	75211	91180
15th Q	51000	**44303**	40574	<u>41142</u>	38107	30108
90th Q	150000	162351	**158037**	<u>158426</u>	162686	163872

In the utility rows, the best performing synthesizer is in bold and the second best is underlined

Outline

Background

Five microdata synthesizers

Simulation study and SDR application

Application

Concluding remarks

Summary

- We propose a stronger, non-asymptotic DP mechanism through the censoring of log-likelihood
- It offers a practical, low-dimensional alternative to truncating the parameter space to achieve DP
- Weighted and Censor_w are recommended given their efficient balance of utility-risk trade-off
  - Weighted demonstrates superior utility preservation at the cost of an aDP guarantee
  - Censor_w provides a stronger, non-asymptotic DP guarantee at the price of slightly reduced utility performance

References I

Dimitrakakis, C., Nelson, B., Zhang, Z., Mitrokotsa, A. and Rubinstein, B. I. P. (2017), 'Differential privacy for bayesian inference through posterior sampling', *Journal of Machine Learning Research* **18**(1), 343–381.

Hu, J. (2019), 'Bayesian estimation of attribute and identification disclosure risks in synthetic data', *Transactions on Data Privacy* **12**, 61–89.

Hu, J., Williams, M. R. and Savitsky, T. D. (2022), 'Mechanisms for global differential privacy under bayesian data synthesis', *Statistica Sinica* .
URL: *https://api.semanticscholar.org/CorpusID:248665552*

Savitsky, T. D., Williams, M. R. and Hu, J. (2022), 'Bayesian pseudo posterior mechanism under asymptotic differential privacy', *Journal of Machine Learning Research* **23**, 1–37.

Wasserman, L. and Zhou, S. (2010), 'A statistical framework for differential privacy', *Journal of the American Statistical Association* **105**, 375–389.

The Paper Covering This Presentation