Journal of Privacy and Confidentiality
Vol. 15 (2) 2025

Submitted June 2024
Published August 2025

# BAYESIAN METHODS TO IMPROVE THE ACCURACY OF DIFFERENTIALLY PRIVATE MEASUREMENTS OF CONSTRAINED PARAMETERS

RYAN JANICKI, SCOTT H. HOLAN, KYLE M. IRIMATA, JAMES LIVSEY, AND ANDREW RAIM

Center for Statistical Research and Methodology, U.S. Census Bureau, 4600 Silver Hill Road, Washington, DC, 20233, USA
*e-mail address*: Ryan.Janicki@census.gov

Department of Statistics, University of Missouri, 146 Middlebush Hall, Columbia, MO, 65211, USA

Office of the Associate Director for Research and Methodology, U.S. Census Bureau, 4600 Silver Hill Road, Washington, DC, 20233, USA

Center for Statistical Research and Methodology, U.S. Census Bureau, 4600 Silver Hill Road, Washington, DC, 20233, USA

Center for Statistical Research and Methodology, U.S. Census Bureau, 4600 Silver Hill Road, Washington, DC, 20233, USA

Center for Statistical Research and Methodology, U.S. Census Bureau, 4600 Silver Hill Road, Washington, DC, 20233, USA

Abstract. Formal disclosure avoidance techniques are necessary to ensure that published data can not be used to identify information about individuals. The addition of statistical noise to unpublished data can be implemented to achieve *differential privacy*, which provides a formal mathematical privacy guarantee. However, the infusion of noise results in data releases which are less precise than if no noise had been added, and can lead to some of the individual data points being nonsensical. Examples of this are estimates of population counts which are negative, or estimates of the ratio of counts which violate known constraints. A straightforward way to guarantee that published estimates satisfy these known constraints is to specify a statistical model and incorporate a prior on census counts and ratios which properly constrains the parameter space. We utilize rejection sampling methods for drawing samples from the posterior distribution and we show that this implementation produces estimates of population counts and ratios which maintain formal privacy, are more precise than the original unconstrained noisy measurements, and are guaranteed to satisfy prior constraints.

*Key words and phrases:* Bayesian analysis, Differential privacy, Inequality constraints, Privacy protection.

www.journalprivacyconfidentiality.org
DOI:10.29012/jpc.936

## 1. Introduction

National statistical agencies such as the U.S. Census Bureau are tasked with collecting information about a country's population, housing characteristics, and business establishments, and disseminating data products for public use. An important consideration that needs to be made when deciding which datasets can be released is protection of individual respondent data. Federal law requires all collected data to be kept confidential. In particular, Title 13 of the U.S. Code prohibits the U.S. Census Bureau from disclosing any "information reported by, or on behalf of, any particular respondent" and from creating "any publication whereby the data furnished by any particular establishment or individual under this title can be identified." [1] (U.S. Census Bureau, 2021).

Disclosure avoidance (DA) is the process of protecting the confidentiality of an individual respondent's personal data. Historically, statistical agencies have used ad hoc methods such as partial or full data suppression of some data tables for small geographic areas to avoid indirect disclosure. More sophisticated methods, such as data swapping, which swaps records for some households with similar characteristics, along with top and bottom coding of sensitive data, have also been used. See McKenna (2018) for a historical review of DA techniques used by the U.S. Census Bureau.

With easy access to increasingly powerful computing resources and advanced data science tools it has become easier to reconstruct and re-identify information about individuals from aggregated public-use data products. Outside parties can now conceivably take published tables from the U.S. Census Bureau and link to information from commercial databases. This database reconstruction has necessitated more sophisticated DA approaches. See McKenna (2019) for examples of successful re-identification attacks using publicly available data.

Recently, in a series of papers beginning with Dwork (2006), a new DA framework based on *differential privacy* (DP) has been developed. Much of the appeal of DP is that it gives a rigorous definition of what is meant by privacy and provides mathematically provable guarantees of privacy protection. Roughly speaking, a privacy protection mechanism that provides differential privacy ensures that the addition or removal of a single record from a database does not dramatically affect the outcome of any analysis (Dwork, 2008). DP has quickly become the gold standard for privacy protection algorithms and has been adopted by the U.S. Census Bureau for its 2020 decennial data products (U.S. Census Bureau, 2021). A comprehensive review of DP can be found in Dwork and Roth (2014) and a review of historical privacy protection methods at the U.S. Census Bureau as well as an overview of methods used in the 2020 decennial census can be found in Abowd and Hawes (2023).

While DP represents a great advancement in the field of disclosure avoidance, it also introduces serious consequences from the perspective of data users. With any DA methodology there is an inherent trade-off between privacy protection and the utility of data. DP, in particular, makes use of noise infusion into confidential data to ensure privacy, which results in published data products that are necessarily less precise than if no DA mechanism had been employed. Noise infusion also has the unfortunate effect of causing estimates to sometimes violate known constraints. For example, DP can result in estimates of population counts which are negative, and estimates of ratios which violate logical constraints. Bayesian modeling has been proposed as a potential method for reducing the impact of DP noise (Charest, 2011). Some examples include Bayesian inference for proportions (Li and Reiter, 2022), the use of Bayesian methods for exact inference (Gong, 2022) or linear

[1] Title 13 U.S. Code, Sections 8–9; Title 13 U.S. Code, Section 141.

regression (Bernstein and Sheldon, 2019). Recent work on addressing issues with bounded data protected by differential privacy can be found in Kazan and Reiter (2024).

In this article we propose a Bayesian model-based solution for producing estimates which are guaranteed to satisfy all known constraints, which are more precise than the original noisy measurements, and which maintain all privacy protections guaranteed by the original DA algorithm. We model the noisy measurements after implementation of DP, using the known noise infusion mechanism, and incorporate all knowledge about logical constraints into the prior distribution on population counts. We provide a brief overview of differential privacy in Section 2 and describe the main results in Section 3. A data example using 2010 decennial census tables is given in Section 4. Concluding remarks are made in Section 5.

## 2. Background

Let $\boldsymbol{X} \subset \mathcal{X}$ denote a dataset (or database) which is a collection of individual records obtained through a survey sample or a census, and let $x_i$ be the number of records in $\boldsymbol{X}$ of type $i$. For example, $\boldsymbol{X}$ might contain information on a collection of records' age, race, sex and geography. Then $i$ would index the possible age by race by sex cross classifications at each geographic region, and $x_i$ would be the number of individuals in $\boldsymbol{X}$ which belong to cell $i$.

The statistical agency that collected the data $\boldsymbol{X}$ would like to release tabulations (or queries) of $\boldsymbol{X}$ to the public, such as the number of persons at different geographies with certain characteristics of interest, but to also ensure that information about any one record in $\boldsymbol{X}$ can not be identified. The privacy of individual records can be achieved by applying DA techniques to either $\boldsymbol{X}$ or tabulations of $\boldsymbol{X}$ prior to dissemination. DA which satisfies differential privacy is a principled way to ensure protection of individual respondents.

Differential privacy, which was introduced in Dwork (2006), requires a definition of distance between databases. For any two databases $\boldsymbol{X}$ and $\boldsymbol{X}'$ in $\mathcal{X}$ we can define their distance using the $l^1$ norm

$$\left\|\boldsymbol{X} - \boldsymbol{X}'\right\|_1 = \sum_{i=1}^{|\mathcal{X}|} \left|x_i - x_i'\right|, \tag{2.1}$$

where $|\mathcal{X}|$ denotes the number of unique records in $\mathcal{X}$. Roughly speaking, if an algorithm provides privacy protections, then the outputs should be similar when applied to similar databases, so that any one individual record is not overly influential and can not easily be recovered. The following formal definition of differential privacy can be found in Dwork and Roth (2014).

**Definition 2.1.** A randomized algorithm $\mathcal{M}$ is $(\varepsilon, \delta)$-differentially private if for all $\mathcal{S} \subseteq \text{Range}(\mathcal{M})$ and for all databases $\boldsymbol{X}, \boldsymbol{X}'$ such that $\|\boldsymbol{X} - \boldsymbol{X}'\| \leq 1$,

$$P\left(\mathcal{M}(\boldsymbol{X}) \in \mathcal{S}\right) \leq \exp\{\varepsilon\} P\left(\mathcal{M}(\boldsymbol{X}') \in \mathcal{S}\right) + \delta,$$

for $\varepsilon > 0$ and $\delta \geq 0$.

The privacy budget is given by the parameters $\varepsilon$ and $\delta$ and determines the amount of privacy guarantee. Small values of $\varepsilon$ and $\delta$ provide greater privacy protection at the expense of less accurate published data, while larger values of $\varepsilon$ and $\delta$ result in more accurate data in exchange for weaker privacy guarantees.

Let $Y$ be a tabulation of $\boldsymbol{X}$ that a statistical agency would like to publish. For example, $Y$ could represent the number of households by relationship for the population under 18 years in a county in the U.S. Additional examples of tabulations published by the U.S. Census Bureau are given in Section 4. The tabulation $Y$ cannot be released to the public without first applying DA techniques. A simple privacy protection algorithm which achieves differential privacy is adding statistical noise to $Y$ and releasing this noisy version of $Y$. The noisy measurement is denoted by $Z$, and can be generated as

$$Z = Y + \varepsilon, \tag{2.2}$$

where $\varepsilon$ is sampled from a noise-generating (probability) distribution. Two of the most used distributions for $\varepsilon$ are the Gaussian distribution, which has probability density function

$$f\left(x; \sigma^2\right) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^2}x^2}$$

and the Laplace distribution, which has density function

$$f\left(x; \lambda\right) = \frac{1}{2\lambda} e^{-\frac{|x|}{\lambda}}.$$

The Laplace mechanism applied in this way preserves $(\varepsilon, 0)$-differential privacy while the Gaussian mechanism preserves $(\varepsilon, \delta)$-differential privacy (Dwork and Roth, 2014).

The addition of statistical noise from a Laplace or Gaussian distribution guarantees that differential privacy will be satisfied. However, it has the unfortunate effect of reducing the utility of the data as the noisy measurements, $Z$, are less precise than the tabulations $Y$. The noisy measurements may also violate certain constraints that the unperturbed tabulations are known to satisfy. For example, if $Y$ is a count tabulation, then $Y$ must be nonnegative. If $Y$ is the ratio of two count tabulations, there may be a relationship between the numerator and the denominator that must be taken into account. Publishing these noisy measurements as official data products could affect how users interpret the data and may result in lowered confidence in the quality of the estimates. In the next section we introduce model-based methods for improving the quality of noisy (differentially private) measurements obtained by adding statistical noise to tabulations by incorporating constraints into a prior distribution.

## 3. Modeling set up

Let $\boldsymbol{Y} \in \mathbb{R}^m$ be a vector of tabulations of the database $\boldsymbol{X}$, and let $\boldsymbol{Z}$ be a privacy-protected measurement of $\boldsymbol{Y}$ obtained by independently adding noise to each component of $\boldsymbol{Y}$. Let $f$ denote the noise generating distribution. Then

$$\boldsymbol{Z} \mid \boldsymbol{Y}, \boldsymbol{\theta} \sim \prod_{i=1}^{m} f\left(Z_i; Y_i, \boldsymbol{\theta}\right) \tag{3.1}$$

where $\boldsymbol{\theta}$ is a vector of parameters determined by the DA algorithm which will be fully known to the analyst. In many situations we will have prior knowledge about logical constraints that must be satisfied by the vector $\boldsymbol{Y}$, but that are not necessarily respected by the vector of noisy measurements, $\boldsymbol{Z}$. For example, if $\boldsymbol{Y}$ is a vector of counts, it follows that each component must be nonnegative. Another potential issue occurs in tables consisting of ratios. Some examples are given in Section 4.

Let $p$ denote the number of known inequality constraints that must be satisfied by the components of $\boldsymbol{Y}$. We can summarize this information in terms of a vector of lower bounds, $\boldsymbol{l} \in \mathbb{R}^p$, a vector of upper bounds, $\boldsymbol{u} \in \mathbb{R}^p$, and a constraint matrix $\boldsymbol{D} \in \mathbb{R}^{p \times m}$,

$$\boldsymbol{l} \leq \boldsymbol{D}\boldsymbol{Y} \leq \boldsymbol{u}, \tag{3.2}$$

where the inequalities are to be interpreted componentwise. A straightforward way to incorporate the constraints is to use a prior distribution on $\boldsymbol{Y}$ with support implied by the inequalities in (3.2). For our work we used an improper distribution

$$\pi(\boldsymbol{Y}) \propto I(\boldsymbol{l} \leq \boldsymbol{DY} \leq \boldsymbol{u}), \tag{3.3}$$

where $I(\cdot)$ is the indicator function. Combining (3.1) and (3.3) results in a posterior distribution

$$\boldsymbol{Y} \mid \boldsymbol{Z}, \boldsymbol{\theta} \propto \prod_{i=1}^{m} f(Z_i; Y_i, \boldsymbol{\theta}) I(\boldsymbol{l} \leq \boldsymbol{DY} \leq \boldsymbol{u}). \tag{3.4}$$

Since the prior is improper when either the upper or lower bound is infinite, it does need to be verified that the expression in (3.4) is integrable. Fortunately, in most practical applications, the noise will be additive so that $f$ is location invariant and (3.4) will be proper.

Samples drawn from (3.4) will be guaranteed to satisfy (3.2). The problem becomes more of a computational one as it is not straightforward to sample from (3.4) efficiently. Neither the posterior distribution (3.4) nor any of its full conditional distributions belong to standard parametric families for any choice of $f$, so there is no known way to directly sample from (3.4). Instead, indirect methods involving proposal distributions and accept/reject algorithms such as the Metropolis Hastings method (Metropolis et al., 1953) must be used. Gibbs sampling under inequality constraints was studied in Gelfand et al. (1992), and can be done with univariate cross-sectional Gibbs sampling which draws from the unconstrained model and only retains samples which lie in the constrained space. Another approach is to sample from the unrestricted posterior distribution and either reject samples which fall outside the constrained space, or project onto the constrained space (Dunson and Neelon, 2003). The main issue with these methods is efficiency as accept/reject algorithms can result in most proposals being rejected if not properly tuned.

For the special case when $f$ is Gaussian, the posterior distribution is a multivariate truncated Gaussian distribution. Sampling from a multivariate truncated Gaussian distribution has been studied by several authors, including Li and Ghosh (2015), Ma et al. (2020), and Ghosal and Ghosh (2022). An efficient implementation of the algorithm in Ma et al. (2020) is in the R package `tmvmixnorm` (Ma et al., 2020).

We are not aware of custom samplers for other multivariate truncated distributions. If the noise mechanism is some other distribution, we propose using a multivariate truncated Gaussian that is close to the target distribution as a proposal. The Gaussian distribution can be chosen by a variety of methods such as moment matching or minimizing a distance. When a Laplace distribution is used, there is a simple, closed-form expression which minimizes the Kullback-Leibler divergence from a Gaussian distribution. Using an approximating Gaussian distribution was found to be an effective solution in Irimata et al. (2022).

Samples generated from (3.4) are forced to satisfy all logical constraints and estimates. Importantly, model estimates of $\mathbf{Y}$ made using the posterior samples (e.g. the posterior mean) will also maintain all privacy protection guarantees as the noisy measurements, $\boldsymbol{Z}$. Using posterior inference for estimation of $\boldsymbol{Y}$ can be thought of as a postprocessing of the noisy measurements, and postprocessing maintains the same level of privacy protections as the preprocessed noisy measurements (Dwork and Roth, 2014, Proposition 2.1).

In addition, the posterior mean will generally be more precise than the noisy measurements as we are utilizing additional information about the unknown parameter $\mathbf{Y}$ through the constraints in the prior distribution. Note that because the noise generating mechanism, including all parameters, is completely known, and the constraints are based on accurate

| | Alabama | Alaska | Arizona | Arkansas | California |
|---|---|---|---|---|---|
| Total: | 3.02 | 3.21 | 3.19 | 3.00 | 3.45 |
| Under 18 years: | 0.87 | 1.07 | 1.01 | 0.90 | 1.05 |
| 18 years and over: | 2.14 | 2.14 | 2.18 | 2.10 | 2.40 |

Table 1. Average family size by age: 2010 published decennial census state-level total population tabulations. Available at data.census.gov.

prior information, there is no possibility of model misspecification. The constraints imposed in the prior effectively reduce the parameter space. Intuitively, if it is known that the true data-generating process is a submodel of some working model, then precision should be increased by using the submodel rather than the full model. An asymptotic result supporting this idea can be found in Altham (1984). We also verify this empirically in the next section.

## 4. Example

This project was motivated by the production needs of the Social, Economic, and Housing Statistics (SEHSD) and Population (POP) Divisions of the U.S. Census Bureau. SEHSD and POP produce the Supplemental Detailed Housing Characteristic (S-DHC) tables which contain information about characteristics of persons, households, and person-household joins (tables which combine person data and household data). There are 8 tables included in the S-DHC. These tables are Average Household Size by Age (PH1), Household Type for the Population in Households (PH2), Households by Relationship For the Population Under 18 years (PH3), Population in Families by Age (PH4), Average Family Size by Age (PH5), Family Type and Age For Own Children Under 18 years (PH6), Total Population in Occupied Housing Units by Tenure (PH7), and Average Household Size of Occupied Housing Units by Tenure (PH8). These tables will be published at the nation and state levels of geography using 2020 census data, and 6 of the eight tables will be iterated by major race category and Hispanic origin. Further detail about S-DHC and other 2020 decennial census data products can be found at https://www.census.gov/programs-surveys/decennial-census/decade/2020/planning-management/process/disclosure-avoidance/newsletters/update-2020-census-data-products.html.

In this section we give an example using the 2010 version of the PH5 table, Average Family Size by Age, Race and Ethnicity in states in the U.S. The race and ethnicity iterations are White alone; Black or African American alone; Asian alone; American Indian and Alaska Native alone; Native Hawaiian and Other Pacific Islander alone; Some Other Race alone; Two or More Races; Hispanic or Latino; White alone, not Hispanic or Latino; and unattributed. The estimates that are produced in this table are the ratio of number of persons 18 and under in families to the number of family households, the number of persons over 18 in families to the number of family households, and the total number of persons in families to the number of family households. Table 1 shows the published 2010 state-level ratios for total population for five states.

Three noisy measurements are generated for each geography and each race iteration for the PH5 table: the total population under 18 in families, the total population over 18 in families, and the number of family households. Denote the true counts as $Y_{18-}$, $Y_{18+}$, and $Y_{\mathrm{FHH}}$, and the noisy measurements as $Z_{18-}$, $Z_{18+}$, and $Z_{\mathrm{FHH}}$. The published values in PH5 are

$$\frac{Z_{18-}}{Z_{\text{FHH}}},\ \frac{Z_{18+}}{Z_{\text{FHH}}},\ \text{and}\ \frac{Z_{18-}+Z_{18+}}{Z_{\text{FHH}}}, \tag{4.1}$$

which are estimates of the ratios

$$\frac{Y_{18-}}{Y_{\text{FHH}}},\ \frac{Y_{18+}}{Y_{\text{FHH}}},\ \text{and}\ \frac{Y_{18-}+Y_{18+}}{Y_{\text{FHH}}}, \tag{4.2}$$

respectively.

The constraints that must be satisfied are $Y_{18-} \geq 0$, $Y_{18+} \geq 0$, $Y_{\text{FHH}} \geq 1$ (we only consider areas with at least one occupied household), and $Y_{18+} + Y_{18-} \geq 2Y_{\text{FHH}}$ (the universe is family households). We have an additional constraint that is specific to our application that $Y_{18+} + Y_{18-} \leq \kappa Y_{\text{FHH}}$, where $\kappa$ is a positive integer. This constraint is due to the privacy algorithm used by the U.S. Census Bureau for person-household join tables which truncates the family household universe to households with at most $\kappa$ individuals. Let $\boldsymbol{Y}^\intercal = (Y_{18-}, Y_{18+}, Y_{\text{FHH}})$ and $\boldsymbol{Z}^\intercal = (Z_{18-}, Z_{18+}, Z_{\text{FHH}})$. For this problem, the constraints in (3.2) are

$$\boldsymbol{l} = \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix},\ \boldsymbol{D} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 1 & -2 \\ -1 & -1 & \kappa \end{bmatrix},\ \boldsymbol{u} = \begin{bmatrix} \infty \\ \infty \\ \infty \\ \infty \\ \infty \end{bmatrix}. \tag{4.3}$$

When all of the noisy measurements $Z_{18-}$, $Z_{18+}$ and $Z_{\text{FHH}}$ are large relative to the amount of noise added, the ratios in (4.1) will typically be sensible, accurate estimates of the true ratios in (4.2). However, when one or more of the true counts which make up the true ratios in (4.2) is very small, the noisy measurements of these counts can be negative, resulting in a ratio which is negative. Furthermore, if $Y_{\text{FHH}}$ is close to zero, the ratios of noisy measurements "blow up" and appear as an impossibly large number.

The preliminary DA algorithm uses a privacy-loss budget which results in 90% margins of error of 200 and a truncation level of 10. This margin of error implies a value of $\rho = 0.016371$ when using zero-concentrated differential privacy (zCDP). See Bun and Steinke (2016) for a discussion of zCDP and its relationship to classical DP. The 90% margin of error of 200 is equivalent to a variance parameter of $\sigma^2 = 14{,}782$ when using a Gaussian noise distribution, or a scale parameter of $\lambda = 86.86$ when using a Laplace noise distribution. We performed two experiments based on these parameter settings. We first generated a set of noisy measurements by adding independent Gaussian noise with variance $\sigma^2 = 14{,}782$ to the true 2010 census counts described above. We then drew 10,000 samples from the posterior distribution (3.4) using the correctly-specified Gaussian likelihood and constraints as in (4.3). All computational work was done using `R` (R Core Team, 2023) and samples were drawn from the multivariate truncated normal distribution using the `rtmvn` function in the `tmvmixnorm` package (Ma et al., 2020).

We then repeated this experiment, but instead added independent Laplace noise with the scale parameter set to 86.86. The posterior distribution in (3.4) is then a truncated multivariate Laplace distribution. Since there is no way to directly sample from this distribution, we instead used the Metropolis-Hastings algorithm to draw samples (Metropolis et al., 1953). We used a multivariate truncated Gaussian distribution as the proposal distribution with the variance set to $\sigma^2 = \pi 86.86^2/2$; it is easy to verify that this minimizes the Kullback-Leibler distance to the truncated multivariate Laplace distribution (Irimata et al., 2022). The computations were completed in just over 17 minutes on a Linux machine

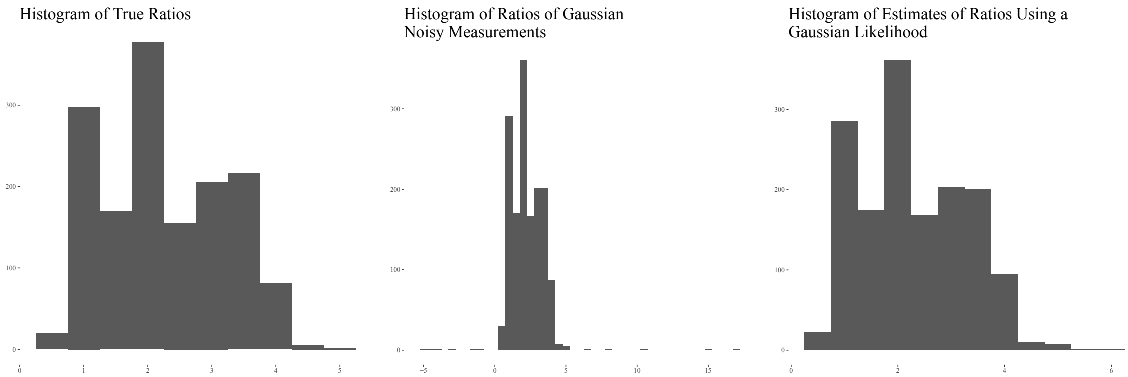


FIGURE 1. Histograms of true ratios, ratios of noisy measurements, and model-based predictions of the ratios.

using a single processor. The truncated Gaussian distribution appeared to be a good proposal distribution in this example as the acceptance rate was 82%.

Figure 1 shows the histogram of the true ratios, the histogram of the ratio of noisy measurements, and the histogram of the predicted ratios using Model (3.4) with a Gaussian likelihood. The true ratios are all between 0 and 6. The histogram of the ratio of noisy measurements has heavier tails with many ratios smaller than 0 or larger than 10 while the histogram of model-based predictions is more in line with the histogram of true ratios.

Table 2 presents summary measures of the noisy measurements and the model-based estimates for each experiment. Recall that the true value of each ratio must be between 0 and 10. The metrics from the 10,000 samples are: the minimum value of all ratios (MIN), the maximum value of all ratios (MAX), the percent of the estimates which violate the constraints in (4.3) (BAD%), the root mean squared error of the ratios (RMSE), given by $\left(\sum_{i=1}^{10{,}000}(\widehat{Y}_i - Y_i)^2/10,000\right)^{1/2}$, where $\widehat{Y}_i$ is the estimate and $Y_i$ is the true count, the average coverage rate of the interval estimates (COV), which were obtained from the 5th and 95th percentiles of the posterior distribution, and the average length of the interval estimates (LEN).

When the noisy measurements are generated from the Gaussian distribution, confidence intervals for the ratio of noisy measurements can be calculated using Fieller's method (Fieller, 1954). There is no analogous method for obtaining confidence intervals for the ratio of Laplace variables, so the COV and LEN columns for the noisy measurements for this experiment are omitted. The interval estimates for the predicted values were taken to be the 5th and 95th percentiles of the samples taken from the posterior distribution.

In both experiments, only 1.4% of the ratios violated the constraints in (4.3). However, the consequences of these violations become clear when looking at the results in Table 2. In both the Gaussian and Laplace experiments there are estimates of average number of persons in family households at the state level that are either negative or unreasonably large. Publishing such numbers could result in lack of confidence from data users. The model-based post-processing of the noisy measurements eliminated all nonsensical ratio estimates and drastically reduced the root mean squared error and length of the coverage intervals. Also, the MH algorithm used with the Laplace mechanism had a high acceptance

| Mechanism | Estimate | MIN | MAX | BAD% | RMSE | COV | LEN |
|---|---|---|---|---|---|---|---|
| Gauss | NM | -5.2 | 17.2 | 1.4 | 0.7 | 87.5 | 1.2 |
| | MB | 0.5 | 6.2 | 0.0 | 0.2 | 89.3 | 0.3 |
| Laplace | NM | -6.6 | 12.5 | 1.4 | 0.6 | NA | NA |
| | MB | 0.5 | 6.2 | 0.0 | 0.2 | 86.7 | 0.3 |

Table 2. Comparison of Noisy Measurements (NM) and Model-based predictions (MB) when the noisy measurements are generated using either a Gaussian mechanism or a Laplace mechanism. The metrics shown are the maximum value (MAX), minimum value (MIN), the percent which are outside the constrained region (BAD%), root mean squared error (RMSE), coverage rate (COV) and interval length (LEN).

rate and the precision of the ratios was as good as the experiment using the Gaussian mechanism, although the coverage rate of the intervals was slightly below the nominal rate.

We also compared the results to those obtained by simple constrained optimization, that is, finding the closest point to the vector of noisy measurements which satisfies the given constraints. We used the `solve.QP` function in the `R` package `quadprog` (Turlach and Weingessel, 2019) to solve this optimization problem. The RMSE of the estimates made using constrained optimization was 0.54, which is a clear improvement over the noisy measurements, but far worse than the model-based predictions. This is due to the fact that the constrained optimization problem is often solved at a point on the boundary, whereas the model-based estimates are forced to be in the interior of the parameter space. Another issue with constrained optimization is that there is no obvious way to measure the uncertainty or to construct a valid interval.

## 5. Conclusion

The analysis shown in Section 4 was for a person-household join table which requires a rather large privacy budget compared to other decennial census products, and also requires a truncation of the housing universe to households containing a pre-specified maximum number of persons. We gave an example which produced a set of noisy measurements for this table at the U.S. state level, and post-processed these noisy measurements using a model through posterior inference using a prior distribution which incorporates known constraints. We demonstrated that this model-based procedure results in estimates which are more precise than the noisy measurements and belong to the constrained parameter space. As a byproduct of our approach, a measure of uncertainty is provided. In principle, this measure of uncertainty could be incorporated downstream in subsequent analyses conducted by data users. However, the best way to do this would be analysis specific and is a subject of future research.

In the future we would like to produce estimates at substate geographies, such as county, tract, and block group levels. Generating noisy measurements for these geographies requires a much larger privacy budget and would result in ratios which more often violate the constraints. Future research is needed to determine whether this procedure still results in publishable estimates at this higher level of noise. Future work is also needed to determine whether covariate information can be utilized to further improve the quality of estimates.

## Acknowledgments

This article is released to inform interested parties of ongoing research and to encourage discussion. The views expressed on statistical issues are those of the authors and not those of the U.S. Census Bureau. The authors thank the associate editor and two referees for their careful review of this paper and helpful comments.

## References

Abowd, J.M., Hawes, M.B., 2023. Confidentiality protection in the 2020 us census of population and housing. Annual Review of Statistics and Its Application 10, 119–144. doi:https://doi.org/10.1146/annurev-statistics-010422-034226.

Altham, P.M.E., 1984. Improving the precision of estimation by fitting a model. Journal of the Royal Statistical Society: Series B (Methodological) 46, 118–119. doi:10.1111/j.2517-6161.1984.tb01283.x.

Bernstein, G., Sheldon, D.R., 2019. Differentially private bayesian linear regression, in: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché Buc, F., Fox, E., Garnett, R. (Eds.), Advances in Neural Information Processing Systems, Curran Associates, Inc. URL: https://proceedings.neurips.cc/paper_files/paper/2019/file/f90f2aca5c640289d0a29417bcb63a37-Paper.pdf.

Bun, M., Steinke, T., 2016. Concentrated differential privacy: Simplifications, extensions, and lower bounds, in: Hirt, M., Smith, A. (Eds.), Theory of Cryptography, Springer Berlin Heidelberg, Berlin, Heidelberg. pp. 635–658. doi:http://dx.doi.org/10.1007/978-3-662-53641-4_24.

Charest, A.S., 2011. How can we analyze differentially-private synthetic datasets? Journal of Privacy and Confidentiality 2. doi:10.29012/jpc.v2i2.589.

Dunson, D.B., Neelon, B., 2003. Bayesian inference on order-constrained parameters in generalized linear models. Biometrics 59, 286–295. doi:10.1111/1541-0420.00035.

Dwork, C., 2006. Differential privacy, in: Bugliesi, M., Preneel, B., Sassone, Vladimiroand Wegener, I. (Eds.), Automata, Languages and Programming, Springer Berlin Heidelberg, Berlin, Heidelberg. pp. 1–12. doi:http://dx.doi.org/10.1007/11787006_1.

Dwork, C., 2008. Differential privacy: A survey of results, in: Agrawal, M., Du, D., Duan, Z., Li, A. (Eds.), Theory and Applications of Models of Computation, Springer Berlin Heidelberg, Berlin, Heidelberg. pp. 1–19. doi:http://dx.doi.org/10.1007/978-3-540-79228-4_1.

Dwork, C., Roth, A., 2014. The algorithmic foundations of differential privacy. Foundations and Trends in Theoretical Computer Science 9, 211–407. doi:10.1561/0400000042.

Fieller, E.C., 1954. Some problems in interval estimation. Journal of the Royal Statistical Society: Series B (Methodological) 16, 175–185. doi:10.1111/j.2517-6161.1954.tb00159.x.

Gelfand, A.E., Smith, A.F.M., Lee, T.M., 1992. Bayesian analysis of constrained parameter and truncated data problems using Gibbs sampling. Journal of the American Statistical Association 87, 523–532. doi:10.1080/01621459.1992.10475235.

Ghosal, R., Ghosh, S.K., 2022. Bayesian inference for generalized linear model with linear inequality constraints. Computational Statistics & Data Analysis 166, 107335. doi:https://doi.org/10.1016/j.csda.2021.107335.

Gong, R., 2022. Exact inference with approximate computation for differentially private data via perturbations. Journal of Privacy and Confidentiality 12. doi:10.29012/jpc.797.

Irimata, K.M., Raim, A.M., Janicki, R., Livsey, J.A., Holan, S.H., 2022. Evaluation of Bayesian Hierarchical Models of Differentially Private Data Based on an Approximate Data Model. Technical Report. U. S. Census Bureau. URL: https://www.census.gov/library/working-papers/2022/adrm/RRS2022-05.html. Research Report Series (Statistics #2022-05), Center for Statistical Research and Methodology, US Bureau of the Census, Washington, DC.

Kazan, Z., Reiter, J.P., 2024. Bayesian inference under differential privacy: Prior selection considerations with application to univariate Gaussian data and regression. arXiv:2405.13801.

Li, L., Reiter, J.P., 2022. Bayesian inference for estimating subset proportions using differentially private counts. Journal of Survey Statistics and Methodology 10, 785–803. doi:10.1093/jssam/smab060.

Li, Y., Ghosh, S.K., 2015. Efficient sampling methods for truncated multivariate normal and student-t distributions subject to linear inequality constraints. Journal of Statistical Theory and Practice 9, 712–732. doi:10.1080/15598608.2014.996690.

Ma, T.F.R., Ghosh, S.K., Li, Y., 2020. tmvmixnorm: Sampling from Truncated Multivariate Normal and t Distributions. URL: https://CRAN.R-project.org/package=tmvmixnorm. r package version 1.1.1.

McKenna, L., 2018. Disclosure Avoidance Techniques Used for the 1970 through 2010 Decennial Censuses of Population and Housing. Technical Report. U. S. Census Bureau. URL: https://www.census.gov/library/working-papers/2018/adrm/ces-wp-18-47.html.

McKenna, L., 2019. U. S. Census Bureau Reidentification Studies. Technical Report. U. S. Census Bureau. URL: https://www.census.gov/library/working-papers/2019/adrm/CED-WP-2019-008.html.

Metropolis, N., Rosenbluth, A.W., Rosenbluth, M.N., Teller, A.H., Teller, E., 1953. Equation of state calculations by fast computing machines. The Journal of Chemical Physics 21, 1087–1092. doi:10.1063/1.1699114.

R Core Team, 2023. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing. Vienna, Austria. URL: https://www.R-project.org/.

Turlach, B.A., Weingessel, A., 2019. quadprog: Functions to Solve Quadratic Programming Problems. URL: https://CRAN.R-project.org/package=quadprog. r package version 1.5-8.

U.S. Census Bureau, 2021. Disclosure Avoidance for the 2020 Census: An Introduction. U.S. Government Publishing Office, Washington, DC. URL: https://www2.census.gov/library/publications/decennial/2020/2020-census-disclosure-avoidance-handbook.pdf.