Bayesian Inference Under Differential Privacy: Prior Selection Considerations with Application to Univariate Gaussian Data and Regression

Zeki Kazan
Department of Statistical Science
Duke University
Durham, NC 27708
zekican.kazan@duke.edu

Jerome P. Reiter
Department of Statistical Science
Duke University
Durham, NC 27708
jreiter@duke.edu

Abstract

We describe Bayesian inference for the mean and variance of bounded data protected by differential privacy and modeled as Gaussian. Using this setting, we demonstrate that analysts can and should take the constraints imposed by the bounds into account when specifying prior distributions. Additionally, we provide theoretical and empirical results regarding what classes of default priors produce valid inference for a differentially private release in settings where substantial prior information is not available. We discuss how these results can be applied to Bayesian inference for regression with differentially private data.

1 Introduction

Differential privacy (DP) [11] is the current gold standard for protecting individual privacy when performing data releases involving confidential data. A common method of ensuring a release satisfies DP is to add calibrated random noise to sufficient statistics of the data. Once the noisy statistics are released, analysts in theory can use them to perform statistical tasks such as the estimation of unknown parameters of interest, prediction of new data values, and quantification of the uncertainty around these quantities. Performing these tasks in a principled way, however, can be nontrivial and is an area of active research.

Bayesian methods are natural for statistical inference after a DP release. By sampling from posterior distributions, analysts can perform estimation and prediction tasks automatically with "built-in" uncertainty quantification. As examples, Bernstein and Sheldon (2018) [4] and Ju et al. (2022) [18] use Gibbs sampling to estimate differentially private posterior distributions, the former for exponential families and the latter for likelihoods satisfying a record additivity condition. Gong (2022) [15] uses approximate Bayesian computation for settings where a perturbation mechanism is applied. Other works investigate Bayesian estimation of subset proportions [27], linear regression coefficients [5], and Poisson factorization [35]. These methods all employ variations of the following scheme: repeatedly impute plausible values of the nonprivate statistics based on their private counterparts, and obtain samples of the parameters of interest via a traditional nonprivate Bayesian analysis.

DP is related to Bayesian methods more generally. For example, releasing a sample from a posterior distribution can be used to achieve DP [8, 14, 17, 30, 39, 41, 42]. Another line of work examines Bayesian semantics for DP and their relationship to disclosure risk assessment [11, 20, 21, 22, 24, 25, 28]. We also note that researchers have developed non-Bayesian methods for inference under DP, including interval estimates [2, 9, 10, 12, 19, 36] and significance tests [1, 2, 3, 6, 23, 34].

Existing methods for Bayesian inference under DP generally presume that the analyst sets an informative prior distribution. In practice, however, the analyst may not have sufficient prior knowledge

Preprint. Under review.

to set an accurate, informative prior. Additionally, methods in the literature typically rely on prior distributions with unbounded support. Often, however, the confidential data are assumed to be constrained in some way; for example, values lie within an interval $[a, b]$ [5, 10, 36, 38, 40]. Arguably, samples from the posterior distribution should reflect this constraint as well. Yet, existing methods for Bayesian inference under DP can produce estimates or predictions outside the feasible range. This may lead analysts to resort to ad hoc fixes such as clipping estimates or predictions to be at the bounds of the interval, which can lead to undesirable and unpredictable behavior.

We address issues resulting from ignoring constraints in Bayesian inference under DP and the use of default prior distributions under DP. Our primary contributions include:

1. We illustrate the incorporation of constraints into a differentially private, Bayesian inference procedure and demonstrate their effect on estimation and prediction tasks.
2. We present a first-of-its-kind examination of default prior distributions analysts might consider in lieu of highly informative proper priors, and present theoretical results demonstrating the need for caution when certain weakly informative priors are used in a DP analysis.

Throughout, we use the univariate Gaussian setting as a motivating case study. This setting is among the most important and commonly encountered in statistics but is relatively understudied in the DP literature, with only a few works [9, 10, 12, 19] providing frequentist methods for inference with uncertainty quantification. General Bayesian DP inference methods such as those in [4, 15, 18] can be applied to this setting but require approximations or substantial computation time. As part of our contributions, we propose a Gibbs sampling procedure for the univariate Gaussian setting without these drawbacks.

2 Background and Case Study

In this section, we briefly review the definition of DP and several of its properties. We then describe the DP univariate Gaussian setting and present our novel Gibbs sampling scheme.

2.1 Differential Privacy

As described by [11], a release mechanism $\mathcal{M}$ satisfies ε-DP if for all data sets x, x' that differ in only one row and all $S \subseteq Range(\mathcal{M})$,

$$P[\mathcal{M}(x) \in S] \leq e^{\varepsilon} P[\mathcal{M}(x') \in S]. \tag{1}$$

DP has many desirable properties [11]. Here, we make use of *post-processing.* For any data set x, if $\mathcal{M}(x)$ satisfies ε-DP and g is not a function of x, then $g(\mathcal{M}(x))$ satisfies ε-DP. We also make use of *composition.* For any data set x, if $m_1 = \mathcal{M}_1(x)$ satisfies ε_1-DP and $m_2 = \mathcal{M}_2(x, m_1)$ satisfies ε_2-DP, then $\mathcal{M}(x) = (m_1, m_2)$ satisfies $(\varepsilon_1 + \varepsilon_2)$-DP.

For any function f applied to any data set x, the sensitivity is defined as $\Delta f = \max_{x,x'} |f(x) - f(x')|$ for all data sets x, x' that differ in only one row. The sensitivity of f is used in the *Laplace Mechanism,* which is a common method for achieving DP. If a function $f(x)$ has sensitivity Δf, then the mechanism $\mathcal{M}(x) = \mathsf{Lap}(f(x), \Delta f/\varepsilon)$ satisfies ε-DP.

2.2 The Univariate Gaussian Setting

Let $Y_1, \ldots, Y_n$ be the confidential, scalar data values for n individuals. We assume that all Y_i are contained in some publicly known interval $[a, b]$. The analyst models the data as independent and identically distributed with $Y_i \sim \mathcal{N}(\mu, \sigma^2)$ for all i. Gaussian models are reasonable and commonly employed when the observed data are symmetric with a range not near the bounds. For example, quantities like proportions and standardized test scores have known ranges. Other quantities like commute times and stock market returns have known lower bounds and effective upper bounds.

The sufficient statistics are $\bar{Y} = \sum_{i=1}^{n} Y_i/n$ and $S^2 = \sum_{i=1}^{n}(Y_i - \bar{Y})^2/(n-1)$. As shown in [10], $\bar{Y}$ and S^2 have sensitivities $(b-a)/n$ and $(b-a)^2/n$, respectively. To achieve ε-DP, the data curator can release $\bar{Y}^* \sim \mathsf{Lap}(\bar{Y}, (b-a)/(\varepsilon_1 n))$ and $S^{2*} \sim \mathsf{Lap}(S^2, (b-a)^2/(\varepsilon_2 n))$, where $\varepsilon_1 + \varepsilon_2 = \varepsilon$. We make the simplifying assumption that $a = 0$ and $b = 1$. We can do so without loss of generality, as described in Theorem 1, which is proved in Appendix A.

Theorem 1. *Let $Y_1, \ldots, Y_n \in [a, b]$ and let $\tilde{Y}_i = (Y_i - a)/(b - a) \in [0, 1]$. Let $\bar{Y}$ and S^2 be the sufficient statistics for $\{Y_i\}$ and let $\tilde{\bar{Y}}$ and $\tilde{S}^2$ be the sufficient statistics for $\{\tilde{Y}_i\}$. Suppose each statistic is released via the Laplace Mechanism under ε-DP and denote the DP statistics $\bar{Y}^*$, S^{2*}, $\tilde{\bar{Y}}^*$, and $\tilde{S}^{2*}$. Then, $\bar{Y}^* \stackrel{d}{=} (b-a)\tilde{\bar{Y}}^* + a$ and $S^{2*} \stackrel{d}{=} (b-a)^2 \tilde{S}^{2*}$.*

For now, we suppose the analyst wishes to incorporate prior beliefs into their analysis via an informative prior distribution. A conjugate choice for the prior in the public setting is

$$\sigma^2 \sim \mathsf{IG}\left(\nu_0/2, \nu_0\sigma_0^2/2\right), \qquad \mu \mid \sigma^2 \sim \mathcal{N}\left(\mu_0, \sigma^2/\kappa_0\right) \tag{2}$$

with analyst-specified hyperparameters $(\nu_0, \kappa_0, \mu_0, \sigma_0^2)$.

We now present a Gibbs sampler to estimate this model, building on the ideas in [4, 5]. Our innovation is with respect to the full conditional for S^2, for which we derive an exact form. As long as $\varepsilon_2 < 2$, the full conditional for S^2 is a distribution we call a truncated gamma mixture (TGM). [1]

Definition 1. *A random variable $X \sim \mathsf{TGM}(\alpha, \beta, \lambda, \tau)$ for $\alpha > 0$, $\beta > \lambda \geq 0$, and $\tau \in \mathbb{R}$ if it has probability density function*

$$p(x) = \begin{cases} \frac{(\beta+\lambda)^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-(\beta+\lambda)x}, & \text{if } \tau \leq 0; \\ \pi_1 \frac{(\beta-\lambda)^\alpha}{\gamma(\alpha,(\beta-\lambda)\tau)} x^{\alpha-1} e^{-(\beta-\lambda)x}, & \text{if } \tau > 0 \text{ and } x \leq \tau; \\ \pi_2 \frac{(\beta+\lambda)^\alpha}{\Gamma(\alpha,(\beta+\lambda)\tau)} x^{\alpha-1} e^{-(\beta+\lambda)x} & \text{if } \tau > 0 \text{ and } x > \tau, \end{cases} \tag{3}$$

where

$$\pi_1 = \frac{e^{-\lambda\tau} \frac{\gamma(\alpha,(\beta-\lambda)\tau)}{(\beta-\lambda)^\alpha}}{e^{-\lambda\tau} \frac{\gamma(\alpha,(\beta-\lambda)\tau)}{(\beta-\lambda)^\alpha} + e^{\lambda\tau} \frac{\Gamma(\alpha,(\beta+\lambda)\tau)}{(\beta+\lambda)^\alpha}}, \qquad \pi_2 = \frac{e^{\lambda\tau} \frac{\Gamma(\alpha,(\beta+\lambda)\tau)}{(\beta+\lambda)^\alpha}}{e^{-\lambda\tau} \frac{\gamma(\alpha,(\beta-\lambda)\tau)}{(\beta-\lambda)^\alpha} + e^{\lambda\tau} \frac{\Gamma(\alpha,(\beta+\lambda)\tau)}{(\beta+\lambda)^\alpha}}. \tag{4}$$

Here, $\gamma(\alpha, x)$ is the lower incomplete gamma function and $\Gamma(\alpha, x)$ is the upper incomplete gamma function. We provide an algorithm for sampling from the TGM in Appendix B.

Derivations for all full conditionals for the Gibbs sampler are provided in Appendix C. The full conditionals for μ and σ^2 are identical to those of the nonprivate setting. To obtain the full conditional for $\bar{Y}$, we follow the strategy of [4] based on the Bayesian LASSO [33] and use the fact that $\bar{Y}^* \sim \mathsf{Lap}(\bar{Y}, 1/(\varepsilon_1 n))$ is equal in distribution to $\bar{Y}^* \sim \mathcal{N}(\bar{Y}, \omega^2)$ for $\omega^2 \sim \mathsf{Exp}(\varepsilon_1^2 n^2/2)$. The full conditional for $\bar{Y}^*$ is Gaussian, the full conditional for $1/\omega^2$ is inverse-Gaussian, and the full conditional for S^2 is a TGM. The parameters of the full conditionals are provided in Appendix D.

A major advantage of this sampler is its runtime. If T is the number of Gibbs iterations, then the computational complexity of this sampler is $\mathcal{O}(T)$; the runtime does not scale with the sample size, n. The Gibbs sampler of [4] for exponential families applied to the univariate Gaussian setting also has computational complexity $\mathcal{O}(T)$. This sampler, however, requires a Gaussian approximation for S^2, which, particularly when σ^2 is near zero or n is small, can lead to unreliable results. The sampler of [18] does not require any approximations; however, it has computational complexity $\mathcal{O}(nT)$, making it potentially infeasible for large n. Our proposed sampler achieves both the exact sampling of [18] and the $\mathcal{O}(T)$ computational complexity of [4].

3 Incorporating Constraints

By incorporating constraints on data values, analysts can reduce the uncertainty in estimates of the parameters of interest "for free." They need not make additional assumptions; they merely determine how the conditions on the data values constrain the model used for analysis.

3.1 Constraints in the Univariate Gaussian Setting

In the univariate Gaussian setting from Section 2.2, the constraint that $Y_i \in [0, 1]$ affects the model in two ways. Firstly, the bound constrains the parameters μ and σ^2, since bounded random variables

[1] For the distribution of S^2 to be a TGM, we require $\sigma^2 < (n-1)/(2n\varepsilon_2)$. We show in Corollary 1 that in this setting we can guarantee $\sigma^2 \leq 1/4$, which implies $\sigma^2 < (n-1)/(2n\varepsilon_2)$ automatically when $\varepsilon_2 < 2$. Within the Gibbs sampler, we reject and resample any draws for which $\sigma^2 \geq (n-1)/(2n\varepsilon_2)$.

have bounded expectations and variances. Secondly, the bound constrains the sufficient statistics $\bar{Y}$ and S^2, since samples of bounded random variables have bounded means and variances. We demonstrate how we establish these bounds in a manner that can be easily incorporated into the Gibbs sampler from Section 2.2. See Appendix E for proofs.

First, Theorem 2 establishes conditional bounds for $\mu \mid \sigma^2$ and $\sigma^2 \mid \mu$ that can be incorporated into the prior. That is, since the data values are bounded, we may replace the unbounded normal-inverse gamma prior from (2) with a prior of the same form truncated to be within the feasible region. This yields a posterior of the same form as in Section 2.2 truncated to the feasible region.

Theorem 2. *Let $Y_1, \ldots, Y_n \sim \mathcal{N}(\mu, \sigma^2)$ and suppose each $Y_i \in [0, 1]$. If μ is known, then $\sigma^2 \in [0, \mu(1-\mu)]$. If σ^2 is known, then $\mu \in [1/2 - \sqrt{1/4 - \sigma^2}, 1/2 + \sqrt{1/4 - \sigma^2}]$.*

Corollary 1. *If $Y_1, \ldots, Y_n \sim \mathcal{N}(\mu, \sigma^2)$ and $Y_i \in [0, 1]$, then $\sigma^2 \leq 1/4$.*

Theorem 3 establishes bounds for $\bar{Y} \mid S^2$ and $S^2 \mid \bar{Y}$ that can be incorporated into the likelihood. This yields a posterior of the same form as in Section 2.2 truncated to the feasible region.

Theorem 3. *Let $Y_1, \ldots, Y_n \sim \mathcal{N}(\mu, \sigma^2)$ and suppose each $Y_i \in [0, 1]$. If $\bar{Y}$ is known, then $S^2 \in [0, n/(n-1)\bar{Y}(1-\bar{Y})]$. If S^2 is known, then $\bar{Y} \in [1/2 - \sqrt{1/4 - (n-1)/n \cdot S^2}, 1/2 + \sqrt{1/4 - (n-1)/n \cdot S^2}]$.*

The full conditionals under these constraints are described in Appendix D.

3.2 Example: Constraints for the Blood Lead Levels Dataset

We now demonstrate the effect of incorporating these constraints using genuine data. All computations are performed on the $[0, 1]$ scale, but all quantities are converted back to the original scale for clarity of presentation. Interval estimates are those with highest posterior density (HPD).[2]

Example 1. *Researchers sampled the blood lead levels of $n = 43$ policemen assigned to outdoor work in Egypt to examine the effect of leaded gasoline on exposed individuals [31]. The sample mean was $\bar{Y} = 32.08\,\mu g/dL$, and the sample variance was $S^2 = 16.98^2\,\mu g^2/dL^2$. An expert advises that blood lead levels in this region are reasonably bounded above by $100\,\mu g/dL$. A realization of the Laplace Mechanism with $\varepsilon_1 = \varepsilon_2 = 0.25$ yields noisy statistics $\bar{Y}^* = 34.30\,\mu g/dL$ and $S^{2*} = 47.16^2\,\mu g^2/dL^2$. An analyst wishes to estimate μ, the average blood lead level of all policemen assigned to outdoor work in Egypt. She believes blood lead levels are reasonably approximated by a Gaussian distribution. To reflect prior knowledge, she uses* (2) *with $(\mu_0 = 12.5, \sigma_0^2 = 3.8^2, \kappa_0 = 1, \nu_0 = 1)$, roughly representing the equivalent information of a single "prior data point."*

Figure 1 presents the joint posterior distribution for (μ, σ^2) with and without constraints accounted for. For this particular release, S^{2*} is larger than the nonprivate S^2 and is near the upper bound on σ^2 determined by Theorem 2. Thus, more than 10% of samples from the unconstrained analysis have infeasibly large σ^2. Additionally, the unconstrained analysis produces samples with μ either negative or above the upper bound. The constrained analysis respects the bounds and produces only reasonable samples for μ and σ^2.

In both analyses, the analyst's posterior mode is around $\hat{\mu} \approx 16\,\mu\text{g/dL}$ and $\hat{\sigma}^2 \approx 5^2\,\mu\text{g}^2/\text{dL}^2$. The analyses, however, have substantially different amounts of uncertainty. The 95% HPD interval for σ^2 is $[1.2^2, 50.8^2]$ without constraints and $[1.0^2, 39.3^2]$ with constraints. Accounting for constraints allows the analyst to rule out the region with infeasibly high σ^2, yielding a tighter interval estimate. The 95% HPD interval for μ is $[3.4, 48.2]$ without constraints and $[1.9, 42.0]$ with constraints. Similarly to the above, the tighter interval in the constrained analysis is due primarily to ruling out μ in the region where σ^2 is infeasibly large. In both analyses, the posterior concentration is highest near the prior values, μ_0 and σ_0^2; there is much less density near the released noisy values. This indicates that the posterior inference is likely quite sensitive to the analyst's prior distribution.

We emphasize that the constrained posterior is not merely a truncated version of the unconstrained posterior; truncation in this way would yield inaccurate point and interval estimates. Instead, by truncating within the Gibbs sampler, we have mathematically coherent estimates.

[2] HPD intervals are preferred over quantile-based intervals when the posterior distribution is skewed [26]. HPD intervals can be computed with standard tools, such as the `HDInterval` package in `R` [29].

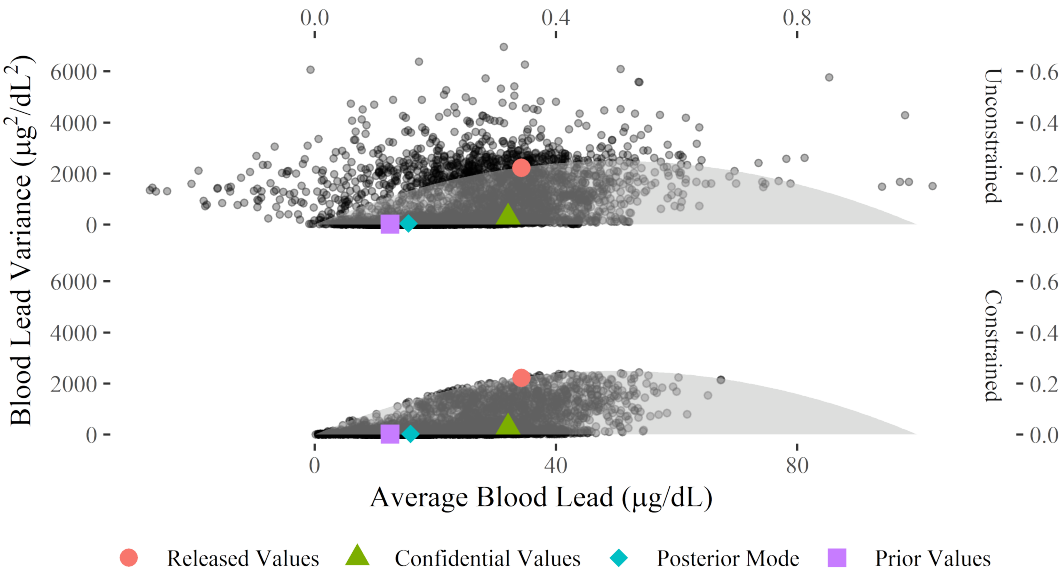


Figure 1: Joint posterior distribution for (μ, σ^2) in Example 1. Here, $(\bar{Y}^* = 34.30, S^{2*} = 47.17^2)$ is represented by the red circle, the unreleased $(\bar{Y} = 32.08, S^2 = 16.96^2)$ is represented by the green triangle, the analyst's posterior mode is represented by the blue diamond, and $(\mu_0 = 12.5, S^2 = 3.8^2)$ is represented by the purple square. Upper and lower panels display the posterior when constraints are and are not accounted for, respectively. The shaded area represents the feasible region for (μ, σ^2) from Theorem 2. The main and secondary axes present the original and $[0, 1]$ scale, respectively. Plots based on 5,000 Gibbs sampler iterations.

4 Default Prior Choices

As evident in Example 1, the results of a Bayesian analysis under DP can be sensitive to the choice of prior distribution. This is unsurprising, since the addition of noise yields data with less information about model parameters than in the public setting, making it difficult to overwhelm the information in the prior distribution. One potential fix—used in examples in the Bayesian DP inference literature—is to increase the prior variance, yielding a weakly informative prior. One also could examine the limit in which the prior variance is infinite. We show, however, that this strategy may lead to improper limiting posterior distributions, even in settings where the nonprivate posterior is proper. In such cases, weakly informative priors may yield unreliable inference and should be avoided [13].

4.1 Default Priors in the Univariate Gaussian Setting

For the unconstrained analysis in the univariate Gaussian setting, a weakly informative prior of the form in (2) is created by making κ_0 and ν_0 very small. In the limit as these hyperparameters go to zero, this produces the default prior $p(\mu, \sigma^2) \propto (\sigma^2)^{-1}$. In the nonprivate setting, this is called the independent Jeffrey's prior [37]. While it does not correspond to a proper probability distribution, the posterior distribution it produces is not only proper, but also frequentist matching.[3] In the private setting, unfortunately, the posterior distribution produced by $p(\mu, \sigma^2) \propto (\sigma^2)^{-1}$ is not a proper probability distribution. The following result, proved in Appendix F, demonstrates this.

Theorem 4. *For confidential data* $Y_1, \ldots, Y_n \overset{iid}{\sim} \mathcal{N}(\mu, \sigma^2)$ *where sufficient statistics* $\bar{Y}^* \sim \textsf{Lap}(\bar{Y}, 1/(\varepsilon_1 n))$ *and* $S^{2*} \sim \textsf{Lap}(S^2, 1/(\varepsilon_2 n))$ *are released, if an analyst has prior* $p(\mu, \sigma^2) \propto (\sigma^2)^{-1}$*, then their posterior* $p(\mu, \sigma^2 \mid \bar{Y}^*, S^{2*})$ *is not a proper probability distribution.*

An alternative strategy for finding a default prior for this setting is to focus on the likelihood's Laplace portion. We show in Appendix G that for a Laplace likelihood in the nonprivate setting,

[3] In general, Bayesian credible intervals do not have the property that a 95% interval will cover the true value 95% of the time under repeated sampling. Frequentist matching priors, however, produce credible intervals identical to corresponding frequentist confidence intervals and so 95% intervals have 95% coverage (see [7]).

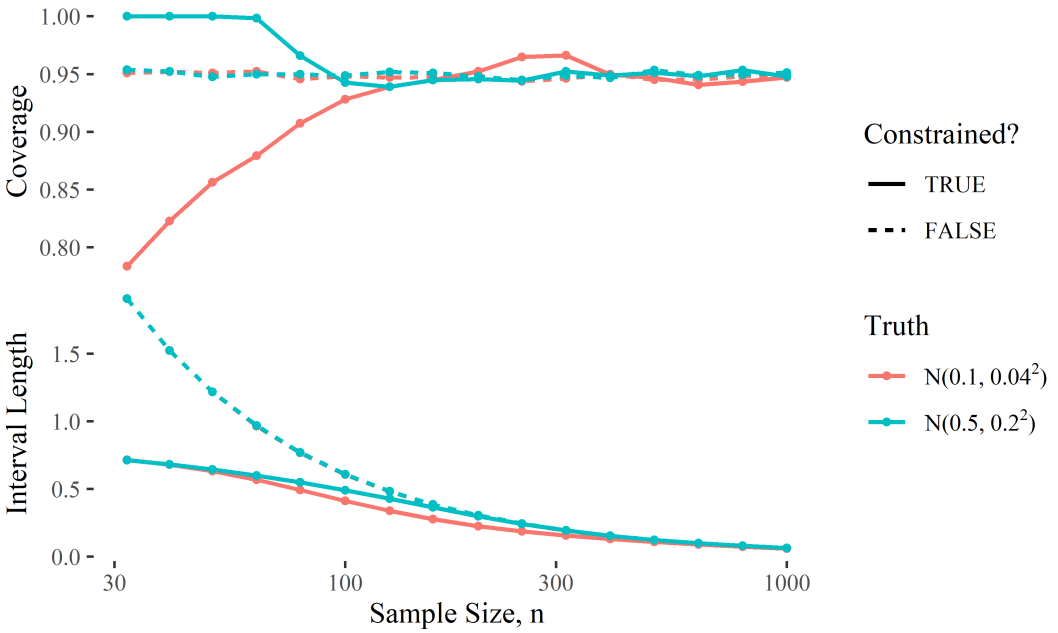


Figure 2: The empirical coverage rate (top) and average length (bottom) of 95% HPD intervals for μ for different sample sizes. Results based on 10,000 simulated datasets $Y_i \in [0, 1]$ released with $\varepsilon_1 = \varepsilon_2 = 0.1$ and analyzed with prior $p(\mu, \sigma^2) \propto 1$. Analyses with constraints accounted for are solid lines and not accounted for are dashed lines. Data generating model is either $\mathcal{N}(\mu = 0.1, \sigma^2 = 0.04^2)$ (red) or $\mathcal{N}(\mu = 0.5, \sigma^2 = 0.2^2)$ (blue). Each Gibbs sampler is run for 20,000 iterations.

a uniform prior produces a proper, frequentist matching posterior distribution. A uniform prior on (μ, σ^2) thus may be a reasonable choice in the private setting. Theorem 5, proved in Appendix F, demonstrates that $p(\mu, \sigma^2) \propto 1$ does indeed produce a proper posterior distribution.

Theorem 5. *For confidential data $Y_1, \ldots, Y_n \overset{iid}{\sim} \mathcal{N}(\mu, \sigma^2)$ where $n > 3$ and sufficient statistics $\bar{Y}^* \sim \mathsf{Lap}(\bar{Y}, 1/(\varepsilon_1 n))$ and $S^{2*} \sim \mathsf{Lap}(S^2, 1/(\varepsilon_2 n))$ are released, if an analyst has prior $p(\mu, \sigma^2) \propto 1$, then their posterior $p(\mu, \sigma^2 \mid \bar{Y}^*, S^{2*})$ is a proper probability distribution.*

Similarly, we can determine posterior propriety for the constrained analysis. Since a constrained uniform distribution has bounded support, it is a proper prior distribution and thus produces a proper posterior distribution. Likewise, it can be verified empirically that the posterior produced by $p(\mu, \sigma^2) \propto (\sigma^2)^{-1}$ also is improper in the constrained analysis.

We use simulation experiments to assess the properties of the posterior distributions based on priors that use or disregard the constraints on the data values. For each of 10,000 simulated datasets of a given size n, we create a 95% HPD interval based on the Gibbs sampler in Appendix D. Figure 2 displays the results. Without enforcing constraints, the 95% credible interval has approximately the nominal 95% coverage rate regardless of n, and the average interval length is linearly related to n on the log-log scale.[4] When enforcing the constraints, the results are similar to those for the unconstrained analysis when n is relatively large. In these cases, the credible intervals for μ are far from the boundary, so that the constraints have little effect. When n is relatively small, the 95% credible interval has lower than 95% coverage rate when the true μ is close to the boundary and near 100% coverage when μ is in the middle of the range. In fact, it can be shown that the prior $p(\mu, \sigma^2) \propto 1$ over the region where $\mu \in [0, 1]$ and $\sigma^2 \in (0, \mu(1 - \mu)]$ induces the marginal distribution $\mu \sim \mathsf{Beta}(2, 2)$. This Beta distribution places more probability density in the center of the distribution than near 0 or 1, resulting in the observed over and under coverage, depending on the true value of μ. This suggests that $p(\mu, \sigma^2) \propto 1$ is most appropriate when an analyst believes μ more likely is near the center of the distribution than the tails. That is, while the prior is "non-informative" in the unconstrained analysis, it is not so in the constrained analysis.

[4] A regression of log average interval length on log n has a slope of -1 (with $R^2 > 0.999$), indicating that average interval length decays proportionally to $1/n$. A similar result holds for the RMSE; see Appendix H.

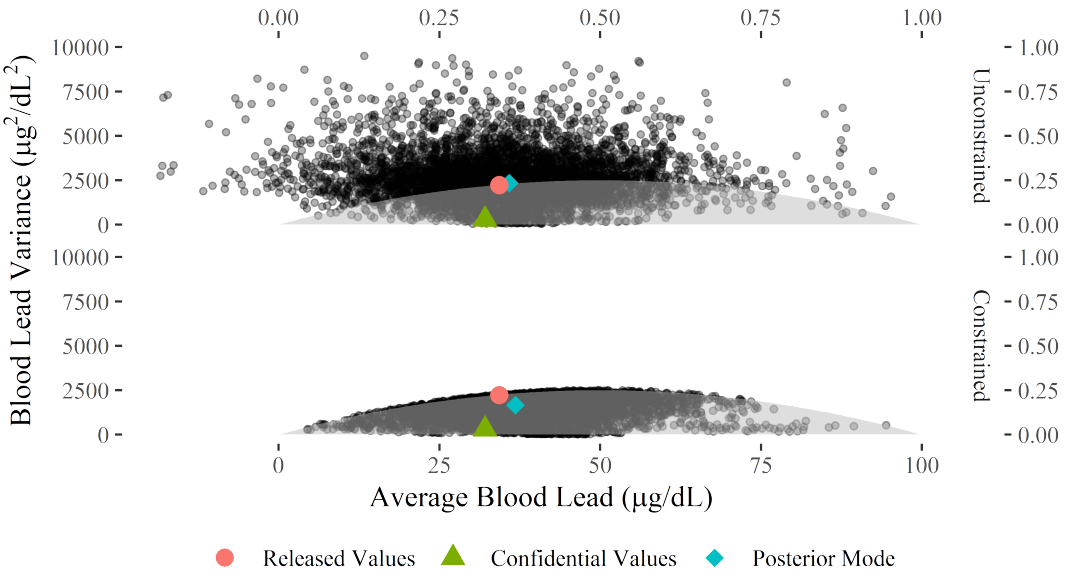


Figure 3: Plot of the joint posterior distribution for (μ, σ^2) in Example 1 under a uniform prior. The point $(\bar{Y}^* = 34.30, S^2 = 47.17^2)$ is represented by the red circle, the unreleased point $(\bar{Y} = 32.08, S^2 = 16.96^2)$ is represented by the green triangle, and the analyst's posterior mode is represented by the blue diamond. The upper and lower panels provide the posterior when constraints are and are not accounted for, respectively. The shaded area represents the feasible region for (μ, σ^2) from Theorem 2. The main axes and secondary axes present the original and $[0, 1]$ scale, respectively. This plot is based on 5,000 Gibbs iterations.

Enforcing the constraint, however, does have advantages. For all points in Figure 2, the average interval length for the constrained analysis is less than 0.75. Meanwhile, the average interval length without constraints is greater than 1 for all $n \leq 50$, indicating that the credible intervals must include values that are not within $[0, 1]$. Such intervals often include all of $[0, 1]$, yielding intervals that are useless in practice. Indeed, the calibrated coverage rate for the unconstrained analysis is of dubious merit, as these credible intervals are inflated by permitting probability mass on infeasible regions of support. Additionally, when we estimate μ with its posterior mode, the RMSEs are uniformly lower in the constrained analysis than in the unconstrained; see Figure 8 in Appendix H. When $\mu = 0.5$ the constrained analysis delivers a "win-win" result: it has a higher coverage rate with a shorter average interval length and a lower RMSE than the unconstrained analysis.

The results described above use a total ε of 0.2. When ε is larger, the constrained and unconstrained analyses are more similar; see Figure 9 in Appendix H for additional simulations with total $\varepsilon = 2$. When ε is larger, the unconstrained analysis no longer offers approximately exact coverage: the 95% posterior credible intervals cover more than 95% of the time.

4.2 Default Priors for the Blood Lead Levels Example

Returning to the setting of Example 1, we examine the effect of replacing the informative prior used in Section 3.2 with the default prior $p(\mu, \sigma^2) \propto 1$. Figure 3 presents a plot analogous to Figure 1 with the new prior. In the unconstrained analysis, the posterior mode is approximately equal to the released statistics, but more than 50% of posterior draws are outside of the feasible region. In the constrained analysis, as discussed above, the posterior mode for μ is shifted towards the center of the distribution. A similar effect is observed for σ^2; the prior has more mass closer to zero and so the posterior mode is smaller than S^{2*}.

Analysts also might be interested in using the Bayesian model for prediction. Figure 4 plots the posterior predictive distribution for the uniform prior, which is computed via $Y_{\text{new}}^{(t)} \sim \mathcal{N}(\mu^{(t)}, \sigma^{2(t)})$ for each Gibbs iteration t, under the unconstrained and constrained analyses.[5] The constrained pos-

[5]For the constrained analysis, the distribution is replaced by a truncated Gaussian.

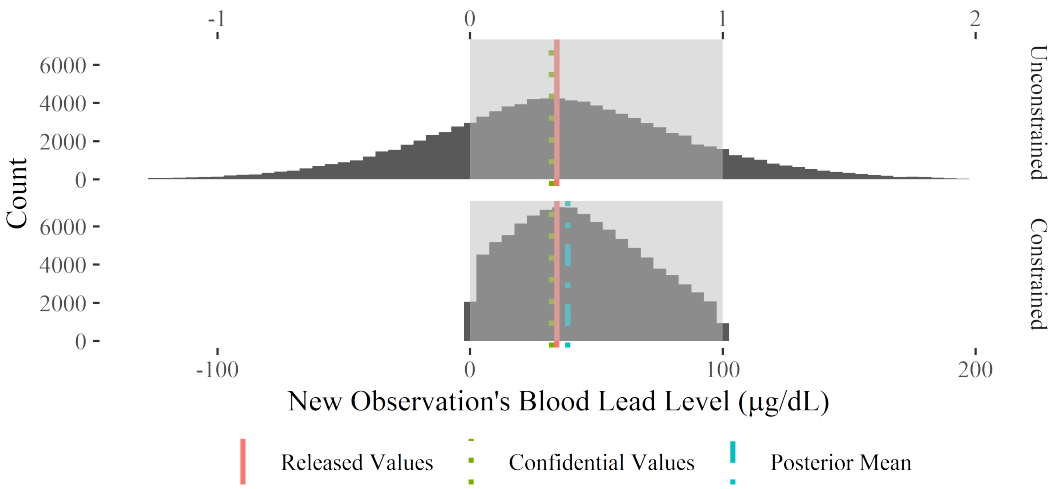


Figure 4: Posterior predictive distribution of a new observation given the posterior draws from Example 1 and prior $p(\mu, \sigma^2) \propto 1$. The value $\bar{Y}^* = 34.30$ is represented by the solid red line, the (unreleased) value $\bar{Y} = 32.08$ is represented by the dotted green line, and the analyst's posterior mode is represented by the blue dot-dashed line. Upper and lower panels provide the posterior predictive distributions without and without accounting for constraints, respectively. The shaded area represents the feasible region for a new observation. The main axis and secondary axis present the original and $[0, 1]$ scale, respectively. Plots based on 100,000 Gibbs sampler iterations.

terior predictive distribution only includes values in the feasible region of 0 to 100 μg/dL, with the distribution centered around the posterior mode. The unconstrained posterior predictive distribution, meanwhile, generates 24% of predictive draws as negative values and 10% as values greater than 100 μg/dL. To obtain reasonable predictions with the draws from the unconstrained analysis, an analyst must resort to ad hoc methods. They might, for example, re-code all negative predictions to 0 and all predictions greater than 100 to 100. This would, among other effects, lead to 24% of predictions being exactly zero, which is not reasonable scientifically. Alternatively, the analyst could sample predictions from a truncated Gaussian distribution. The variance of this distribution is still substantially inflated by the samples of overly large values of $\sigma^{2(t)}$. Thus, predictions have larger variability under this ad hoc approach than under the theoretically principled constrained analysis.

5 Non-Univariate Gaussian Application

The strategies we discuss can be applied to any Bayesian analysis under DP. To demonstrate, we consider the DP Bayesian linear regression method of [5]. In particular, we adapt an example from [5] using data from [**?**] to illustrate how the strategies for enforcing constraints in Section 3 can be used off-the-shelf to supplement an existing method.

The data in [5**?**] include the response variable cirrhosis rate, $\{y_i\}$, explanatory variable drinking rate, $\{x_i\}$, and sample size $n = 46$. The authors rescale both variables to lie in $[0, 1]$ and assume the range of both quantities is public information. Letting $\mathbf{Y} \in \mathbb{R}^n$ be the vector of responses and $\mathbf{X} \in \mathbb{R}^{n \times 2}$ be a design matrix with a column of ones, the authors' model is a simple linear regression of the form below with the following informative, conjugate prior.

$$\mathbf{Y} \sim \mathcal{N}_n(\mathbf{X}\boldsymbol{\theta}, \sigma^2 \mathbf{I}_n), \qquad \boldsymbol{\theta} \mid \sigma^2 \sim \mathcal{N}_2\left(\begin{bmatrix}1\\0\end{bmatrix}, \frac{\sigma^2}{4}\mathbf{I}_2\right), \quad \sigma^2 \sim \mathsf{IG}\left(\frac{40}{2}, \frac{1}{2}\right), \tag{5}$$

where $\mathbf{I}_p$ is the $p \times p$ identity matrix. Their proposed method involves applying the Laplace Mechanism to each element of the sufficient statistics $\mathbf{X}^\top\mathbf{X}$, $\mathbf{X}^\top\mathbf{Y}$ and $\mathbf{Y}^\top\mathbf{Y}$. To ensure a valid regression solution exists, the authors enforce that the matrix $\mathbf{Z}^\top\mathbf{Z}$, where $\mathbf{Z} = [\mathbf{X} \;\; \mathbf{Y}]$, is positive semi-definite (PSD) by projecting sampled values to the nearest PSD matrix. They apply their `Gibbs-SS-Noisy` method, which also adds Laplace noise to the matrix of fourth sample moments of $\mathbf{X}$.

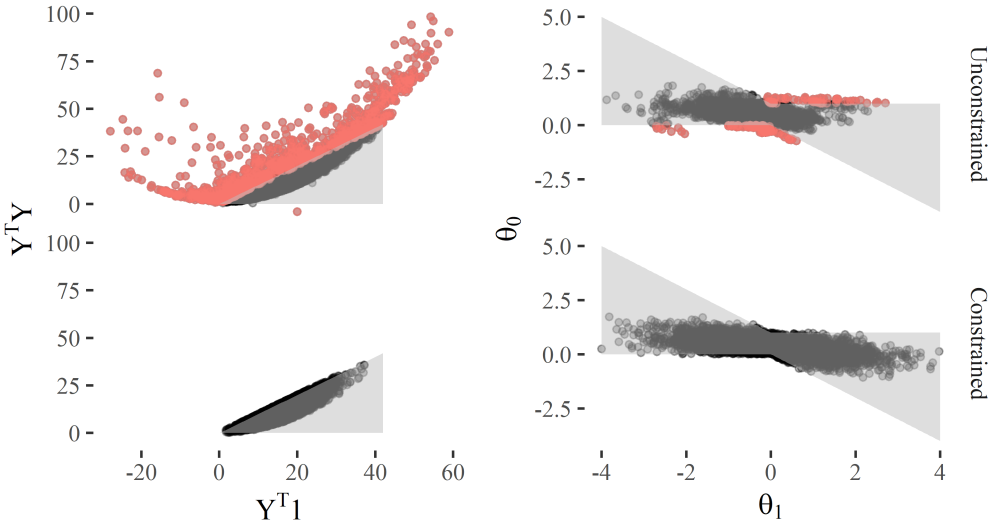


Figure 5: Plot of posterior draws from the linear regression method of [5]. The left panels represent draws of the imputed sufficient statistics $\mathbf{Y}^\top\mathbf{1}$ and $\mathbf{Y}^\top\mathbf{Y}$, while the right panels represent draws of the parameters θ_1 and θ_0. The upper and lower panels provide the posterior when constraints are and are not accounted for, respectively. The shaded areas represent the feasible regions; points outside the feasible regions are colored in red. This plot is based on 10,000 Gibbs iterations.

Since it is known that $x_i, y_i \in [0, 1]$ for all i, we can exploit the resulting structure in the linear regression to constrain parameter updates in the Gibbs sampler. In particular, letting $\mathbf{X} = [\mathbf{1}_n \ \mathbf{x}_1]$, we show in Appendix I that $0 \leq \mathbf{x}_1^\top\mathbf{x}_1 \leq \mathbf{x}_1^\top\mathbf{1} \leq n$ and $0 \leq \mathbf{Y}^\top\mathbf{Y} \leq \mathbf{Y}^\top\mathbf{1} \leq n$. We also show that $\mathbf{x}_1^\top\mathbf{Y} \leq \min\{\mathbf{x}_1^\top\mathbf{1}, \mathbf{Y}^\top\mathbf{1}\}$ and that $0 \leq \sigma^2 \leq 1/4$. Letting the regression coefficients be $\boldsymbol{\theta} = [\theta_0 \ \theta_1]^\top$, we know that (θ_1, θ_0) must lie in the region depicted in the right panels of Figure 5. We create a constrained Gibbs sampler by drawing each iterate from the unconstrained full conditional distribution and rejecting and resampling whenever all constraints are not satisfied. We also reject and resample if $\mathbf{Z}^\top\mathbf{Z}$ is not PSD.

Figure 5 demonstrates some of the effects of enforcing these constraints for one particular set[6] of released noisy values with $\varepsilon = 0.1$ for each of the 11 queries. The shaded triangles in the left panels represent the constraints that $0 \leq \mathbf{Y}^\top\mathbf{Y} \leq \mathbf{Y}^\top\mathbf{1} \leq n$. We find that around 20% of the posterior draws without accounting for constraints do not satisfy this inequality, as indicated by the red points. The shaded regions in the right panels represent the feasible region for $\boldsymbol{\theta}$. In this case, less than 5% of the posterior draws did not satisfy the constraint in $\boldsymbol{\theta}$. But notably, when all relevant constraints are enforced, the posterior distribution is more spread out throughout the feasible region. In this case, not enforcing the constraints leads artificially to an estimated model with too little uncertainty about its parameter estimates and thus underestimated uncertainty around downstream predictions.

This regression example, as well as univariate Gaussian case, suggest that incorporating constraints in Bayesian DP inference, thereby respecting the actual support of the parameters, can result in potentially more accurate representations of posterior distributions. Additionally, the examples indicate that informative prior distributions can have significant impact on DP posterior inferences. The use of weakly informative or default prior distributions may be advantageous, although analysts must take care to ensure posterior distributions are not improper, as demonstrated in Section 4. Finally, these results suggest there is scope for further development of default prior distributions to obtain theoretically valid and accurate posterior distributions with desirable frequentist properties.

[6]The addition of Laplace noise to the fourth sample moments of $\mathbf{X}$ adds additional uncertainty which is not accounted for in the procedure in [5]. Because of this, our rejection sampling approach struggles when the released noisy moments are far from their nonprivate counterparts, leading to the Gibbs sampling procedure being unable to find a draw for the next Gibbs iterate satisfying the constraints and getting "stuck." For demonstrative purposes, we choose a seed where this issue is not encountered.

Acknowledgments and Disclosure of Funding

This research was supported by NSF grant SES-2217456.

References

[1] Daniel Alabi and Salil Vadhan. Hypothesis testing for differentially private linear regression. *Advances in Neural Information Processing Systems*, 35:14196–14209, 2022.

[2] Jordan Awan and Aleksandra Slavković. Differentially private uniformly most powerful tests for binomial data. *Advances in Neural Information Processing Systems*, 31, 2018.

[3] Andrés F Barrientos, Jerome P Reiter, Ashwin Machanavajjhala, and Yan Chen. Differentially private significance tests for regression coefficients. *Journal of Computational and Graphical Statistics*, 28(2):440–453, 2019.

[4] Garrett Bernstein and Daniel R Sheldon. Differentially private Bayesian inference for exponential families. *Advances in Neural Information Processing Systems*, 31, 2018.

[5] Garrett Bernstein and Daniel R Sheldon. Differentially private Bayesian linear regression. *Advances in Neural Information Processing Systems*, 32, 2019.

[6] Simon Couch, Zeki Kazan, Kaiyan Shi, Andrew Bray, and Adam Groce. Differentially private nonparametric hypothesis testing. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pages 737–751, 2019.

[7] Gauri Sankar Datta and Trevor J Sweeting. Probability matching priors. In D K Dey and C R Rao, editors, *Bayesian Thinking*, volume 25 of *Handbook of Statistics*, pages 91–114. Elsevier, 2005.

[8] Christos Dimitrakakis, Blaine Nelson, Zuhe Zhang, Aikaterini Mitrokotsa, and Benjamin IP Rubinstein. Differential privacy for Bayesian inference through posterior sampling. *Journal of Machine Learning Research*, 18(11):1–39, 2017.

[9] Vito D'Orazio, James Honaker, and Gary King. Differential privacy for social science inference. *Sloan Foundation Economics Research Paper*, (2676160), 2015.

[10] Wenxin Du, Canyon Foot, Monica Moniot, Andrew Bray, and Adam Groce. Differentially private confidence intervals. *arXiv preprint arXiv:2001.02285*, 2020.

[11] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography: Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006. Proceedings 3*, pages 265–284. Springer, 2006.

[12] Cecilia Ferrando, Shufan Wang, and Daniel Sheldon. Parametric bootstrap for differentially private confidence intervals. In *International Conference on Artificial Intelligence and Statistics*, pages 1598–1618. PMLR, 2022.

[13] Andrew Gelman. Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper). *Bayesian Analysis*, 1(3):515–534, 2006.

[14] Joseph Geumlek, Shuang Song, and Kamalika Chaudhuri. Rényi differential privacy mechanisms for posterior sampling. *Advances in Neural Information Processing Systems*, 30, 2017.

[15] Ruobin Gong. Exact inference with approximate computation for differentially private data via perturbations. *Journal of Privacy and Confidentiality*, 12(2), 2022.

[16] Peter D Hoff. *A First Course in Bayesian Statistical Methods*, volume 580. Springer, 2009.

[17] Jack E Jewson, Sahra Ghalebikesabi, and Chris C Holmes. Differentially private statistical inference through beta-divergence one posterior sampling. *Advances in Neural Information Processing Systems*, 36, 2024.

[18] Nianqiao Ju, Jordan Awan, Ruobin Gong, and Vinayak Rao. Data augmentation mcmc for Bayesian inference from privatized data. *Advances in neural information processing systems*, 35:12732–12743, 2022.

[19] Vishesh Karwa and Salil Vadhan. Finite sample differentially private confidence intervals. *arXiv preprint arXiv:1711.03908*, 2017.

[20] Shiva P Kasiviswanathan and Adam Smith. On the'semantics' of differential privacy: A Bayesian formulation. *Journal of Privacy and Confidentiality*, 6(1), 2014.

[21] Zeki Kazan and Jerome P Reiter. Prior-itizing privacy: A Bayesian approach to setting the privacy budget in differential privacy. *arXiv preprint arXiv:2306.13214*, 2023.

[22] Zeki Kazan and Jerome P Reiter. Assessing statistical disclosure risk for differentially private, hierarchical count data, with application to the 2020 us decennial census. *Statistica Sinica*, 34:to appear, 2024.

[23] Zeki Kazan, Kaiyan Shi, Adam Groce, and Andrew P Bray. The test of tests: A framework for differentially private hypothesis testing. In *International Conference on Machine Learning*, pages 16131–16151. PMLR, 2023.

[24] Daniel Kifer, John M Abowd, Robert Ashmead, Ryan Cumings-Menon, Philip Leclerc, Ashwin Machanavajjhala, William Sexton, and Pavel Zhuravlev. Bayesian and frequentist semantics for common variations of differential privacy: Applications to the 2020 census. *arXiv preprint arXiv:2209.03310*, 2022.

[25] Daniel Kifer and Ashwin Machanavajjhala. Pufferfish: A framework for mathematical privacy definitions. *ACM Transactions on Database Systems (TODS)*, 39(1):1–36, 2014.

[26] John K Kruschke. Bayesian approaches to testing a point ("null") hypothesis. In John K Kruschke, editor, *Doing Bayesian Data Analysis (Second Edition)*, pages 335–358. Academic Press, Boston, second edition edition, 2015.

[27] Linlin Li and Jerome P Reiter. Bayesian inference for estimating subset proportions using differentially private counts. *Journal of Survey Statistics and Methodology*, 10(3):785–803, 2022.

[28] David McClure and Jerome P Reiter. Differential privacy and statistical disclosure risk measures: An investigation with binary synthetic data. *Trans. Data Priv.*, 5(3):535–552, 2012.

[29] Mike Meredith and John Kruschke. *HDInterval: Highest (Posterior) Density Intervals*, 2020. R package version 0.2.2.

[30] Kentaro Minami, Hitomi Arai, Issei Sato, and Hiroshi Nakagawa. Differential privacy without sensitivity. *Advances in Neural Information Processing Systems*, 29, 2016.

[31] Wael Mortada, Mohamed Sobh, M M El-Defrawy, and S E Farahat. Study of lead exposure from automobile exhaust as a risk for nephrotoxicity among traffic policemen. *American Journal of Nephrology*, 21(4):274–279, 2001.

[32] Frederick Novomestky and Saralees Nadarajah. *truncdist: Truncated Random Variables*, 2016. R package version 1.0-2.

[33] Trevor Park and George Casella. The Bayesian lasso. *Journal of the American Statistical Association*, 103(482):681–686, 2008.

[34] Víctor Peña and Andrés F Barrientos. Differentially private hypothesis testing with the subsampled and aggregated randomized response mechanism. *Statistica Sinica*, 34:to appear, 2024.

[35] Aaron Schein, Zhiwei Steven Wu, Alexandra Schofield, Mingyuan Zhou, and Hanna Wallach. Locally private Bayesian inference for count models. In *International Conference on Machine Learning*, pages 5638–5648. PMLR, 2019.

[36] Or Sheffet. Differentially private ordinary least squares. In *International Conference on Machine Learning*, pages 3105–3114. PMLR, 2017.

[37] Dongchu Sun and James O Berger. Objective Bayesian analysis for the multivariate normal model. *Bayesian Statistics*, 8:525–562, 2007.

[38] Yu-Xiang Wang. Revisiting differentially private linear regression: optimal and adaptive prediction & estimation in unbounded domain. *arXiv preprint arXiv:1803.02596*, 2018.

[39] Yu-Xiang Wang, Stephen Fienberg, and Alex Smola. Privacy for free: Posterior sampling and stochastic gradient Monte Carlo. In *International Conference on Machine Learning*, pages 2493–2502. PMLR, 2015.

[40] Jun Zhang, Zhenjie Zhang, Xiaokui Xiao, Yin Yang, and Marianne Winslett. Functional mechanism: regression analysis under differential privacy. *Proceedings of the VLDB Endowment*, 5(11):1364–1375, 2012.

[41] Wanrong Zhang and Ruqi Zhang. Dp-fast mh: Private, fast, and accurate Metropolis-Hastings for large-scale Bayesian inference. In *International Conference on Machine Learning*, pages 41847–41860. PMLR, 2023.

[42] Zuhe Zhang, Benjamin Rubinstein, and Christos Dimitrakakis. On the differential privacy of Bayesian inference. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, 2016.

A Re-scaling The Data

Throughout the document, we assume that the data $Y_1, \ldots, Y_n$ are contained in the interval $[a, b]$ for known, public a and b. In this section, we prove that, without loss of generality, we may consider data on the interval $[0, 1]$ by re-scaling the data as follows to obtain $\tilde{Y}_i \in [0, 1]$,

$$\tilde{Y}_i = \frac{Y_i - a}{b - a}. \tag{6}$$

We can convert back to the original scale via the relationship $Y_i = (b - a)\tilde{Y}_i + a$.

We now show that the sufficient statistics released via the Laplace Mechanism on the $[0, 1]$ scale are re-scaled versions of the sufficient statistics released via the Laplace Mechanism on the $[a, b]$ scale. Thus no information is lost from scaling before the DP release and re-scaling afterwards.

Theorem 1. *Let $Y_1, \ldots, Y_n \in [a, b]$ and let $\tilde{Y}_i = (Y_i - a)/(b - a) \in [0, 1]$. Let $\bar{Y}$ and S^2 be the sufficient statistics for $\{Y_i\}$ and let $\tilde{\bar{Y}}$ and $\tilde{S}^2$ be the sufficient statistics for $\{\tilde{Y}_i\}$. Suppose each statistic is released via the Laplace Mechanism under ε-DP and denote the DP statistics $\bar{Y}^*$, S^{2*}, $\tilde{\bar{Y}}^*$, and $\tilde{S}^{2*}$. Then, $\bar{Y}^* \stackrel{d}{=} (b - a)\tilde{\bar{Y}}^* + a$ and $S^{2*} \stackrel{d}{=} (b - a)^2\tilde{S}^{2*}$.*

Proof. To begin, note that we may relate the sample means on the two scales as follows

$$\bar{Y} = \frac{1}{n}\sum_{i=1}^{n} Y_i = \frac{1}{n}\sum_{i=1}^{n}\Big[(b - a)\tilde{Y}_i + a\Big] = (b - a)\frac{1}{n}\sum_{i=1}^{n}\tilde{Y}_i + a = (b - a)\tilde{\bar{Y}} + a. \tag{7}$$

Similarly, we may relate the sample variances on the two scales as follows.

$$\begin{aligned} S^2 &= \frac{1}{n-1}\sum_{i=1}^{n}(Y_i - \bar{Y})^2 && (8)\\ &= \frac{1}{n-1}\sum_{i=1}^{n}([(b - a)\tilde{Y}_i + a] - [(b - a)\tilde{\bar{Y}} + a])^2 && (9)\\ &= (b - a)^2\frac{1}{n-1}\sum_{i=1}^{n}(\tilde{Y}_i - \tilde{\bar{Y}})^2 && (10)\\ &= (b - a)^2\tilde{S}^2. && (11) \end{aligned}$$

We now consider the sensitivities of each of the four statistics. By Lemmas 11 and 12 in [10], the sensitivity of $\bar{Y}$ is $(b-a)/n$ and the sensitivity of S^2 is $(b-a)^2/n$. Similarly, the sensitivity of $\tilde{\bar{Y}}$ is $1/n$ and the sensitivity of $\tilde{S}^2$ is $1/n$. Thus, when released under ε-DP via the Laplace Mechanism, the released statistics have distribution

$$\bar{Y}^* \sim \mathsf{Lap}\left(\bar{Y}, \frac{b-a}{\varepsilon n}\right), \qquad S^{2*} \sim \mathsf{Lap}\left(S^2, \frac{(b-a)^2}{\varepsilon n}\right) \tag{12}$$

$$\tilde{\bar{Y}}^* \sim \mathsf{Lap}\left(\tilde{\bar{Y}}, \frac{1}{\varepsilon n}\right), \qquad \tilde{S}^{2*} \sim \mathsf{Lap}\left(\tilde{S}^2, \frac{1}{\varepsilon n}\right). \tag{13}$$

Recall that if X has a Laplace distribution with location parameter m and scale parameter b, i.e., $X \sim \mathsf{Lap}(m, b)$, then if $k > 0$ and $c \in \mathbb{R}$, it follows that $kX + c \sim \mathsf{Lap}(km + c, kb)$. By this property, $(b-a)\tilde{\bar{Y}}^* + a$ has location parameter $(b-a)\tilde{\bar{Y}} + a = \bar{Y}$ and scale parameter $(b-a)/(\varepsilon n)$. Similarly, $(b-a)^2\tilde{S}^{2*}$ has location parameter $(b-a)^2\tilde{S}^2 = S^2$ and scale parameter $(b-a)^2/(\varepsilon n)$. The result follows. □

By this result, it is equivalent to release $\bar{Y}^*$ and S^{2*} on the $[a, b]$ scale directly via the Laplace distribution or to re-scale the data to the $[0, 1]$ scale, release $\tilde{\bar{Y}}^*$ and $\tilde{S}^{2*}$, and then use the relationships in Theorem 1 to convert back to the $[a, b]$ scale.

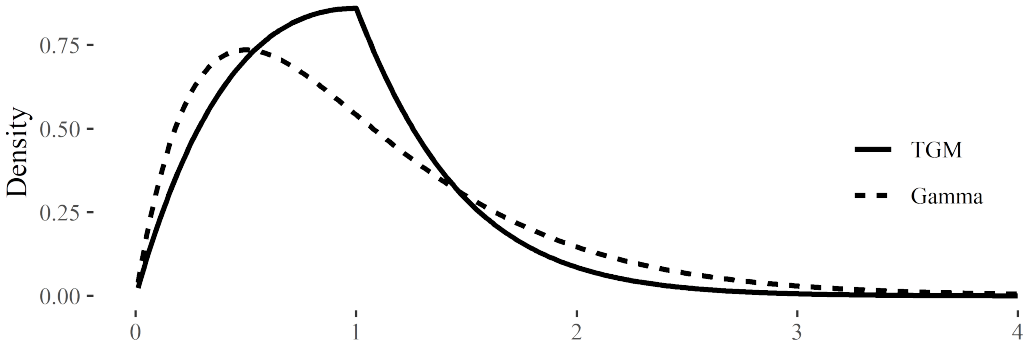


Figure 6: Comparison of the probability density functions of $\mathsf{TGM}(\alpha = 2, \beta = 2, \lambda = 1, \tau = 1)$ (the solid line) and $\mathsf{Gam}(\alpha = 2, \beta = 2)$ (the dashed line).

B The Truncated Gamma Mixture Distribution

This section provides additional details related to the truncated gamma mixture distribution (abbreviated TGM). The TGM distribution has four parameters: α is analogous to the gamma distribution's scale parameter, β is analogous to the gamma distribution's rate parameter, λ determines how far the TGM distribution is from a gamma distribution, and τ determines the point of truncation. When $\tau > 0$, the distribution is a mixture of $\mathsf{Gam}_{(0,\tau]}(\alpha, \beta - \lambda)$ and $\mathsf{Gam}_{(\tau,\infty)}(\alpha, \beta + \lambda)$ with the mixture weights in (4). When $\tau \leq 0$, the distribution is equivalent to $\mathsf{Gam}(\alpha, \beta + \lambda)$.

Figure 6 compares the probability density of the TGM distribution to the probability density of the gamma distribution with the same α and β. We see that the shapes are similar, but the rates are different on each side of τ. Notably, the distribution is continuous but is clearly not differentiable at τ.

Algorithm 1 provides a straightforward procedure for sampling from the TGM distribution. It assumes one is able to sample from a truncated gamma distribution, which can be done via standard tools such as the `rtrunc` function from the `truncdist` package in `R` [32].

Algorithm 1: Sample $X \sim \mathsf{TGM}(\alpha, \beta, \lambda, \tau)$

Input: α, β, λ, τ
if $\tau \leq 0$ **then**
  Sample $X \sim \mathsf{Gam}(\alpha, \beta + \lambda)$
else
  Compute π_1 via (4)
  Sample $U \sim \mathsf{Unif}(0, 1)$
  if $U \leq \pi_1$ **then**
    Sample $X \sim \mathsf{Gam}(\alpha, \beta - \lambda)$ truncated to $(0, \tau]$
  else
    Sample $X \sim \mathsf{Gam}(\alpha, \beta + \lambda)$ truncated to (τ, ∞)
Output: X

We now prove that Algorithm 1 correctly samples from the TGM distribution.

Theorem 6. *X sampled via Algorithm 1 has distribution $X \sim \mathsf{TGM}(\alpha, \beta, \lambda, \tau)$.*

Proof. If $\tau \leq 0$, then $X \sim \mathsf{Gam}(\alpha, \beta + \lambda)$ and so has PDF

$$p(x) = \frac{(\beta + \lambda)^\alpha}{\Gamma(\alpha)} x^{\alpha - 1} e^{-(\beta + \lambda)x}. \tag{14}$$

Thus, the distribution of X is as desired for $\tau \leq 0$.

We now consider the case where $\tau > 0$. Let $F_1(x)$ and $f_1(x)$ be the CDF and PDF, respectively, of $\mathsf{Gam}(\alpha, \beta - \lambda)$ truncated to $(0, \tau]$. Let $F_2(x)$ and $f_2(x)$ be the CDF and PDF, respectively, of $\mathsf{Gam}(\alpha, \beta + \lambda)$ truncated to (τ, ∞). By the law of total probability, the CDF of X is

$$P[X \le x] = P[X \le x \mid U \le \pi_1]\, P[U \le \pi_1] + P[X \le x \mid U > \pi_1]\, P[U > \pi_1] \tag{15}$$

$$= \pi_1\, F_1(x) + \pi_2\, F_2(x). \tag{16}$$

This function is differentiable everywhere except $x = \tau$. When $x \neq \tau$, we may take the derivative with respect to x to obtain the PDF

$$p(x) = \frac{\mathrm{d}}{\mathrm{d}x}\left[\pi_1\, F_1(x) + \pi_2\, F_2(x)\right] \tag{17}$$

$$= \pi_1 f_1(x) + \pi_2 f_2(x) \tag{18}$$

$$= \pi_1 \frac{(\beta-\lambda)^\alpha}{\gamma(\alpha, (\beta-\lambda)\tau)} x^{\alpha-1} e^{-(\beta-\lambda)x} \mathbf{1}[x \le \tau] + \pi_2 \frac{(\beta+\lambda)^\alpha}{\Gamma(\alpha, (\beta+\lambda)\tau)} x^{\alpha-1} e^{-(\beta+\lambda)x} \mathbf{1}[x > \tau]. \tag{19}$$

Thus, we have shown that $X \sim \mathsf{TGM}(\alpha, \beta, \lambda, \tau)$ for all $X \neq \tau$. Since the event $X = \tau$ occurs with probability zero, this completes the proof. □

The TGM distribution primarily arises in the following setting.

Theorem 7. *If $\tau \mid \mu \sim$ Lap$(\mu, 1/\lambda)$ and $\mu \sim$ Gam(α, β) for $\beta > \lambda$, then $\mu \mid \tau \sim$ TGM$(\alpha, \beta, \lambda, \tau)$.*

Proof. By Bayes' Theorem, the desired distribution is

$$p(\mu \mid \tau) = \frac{p(\tau \mid \mu)\, p(\mu)}{\int_0^\infty p(\tau \mid \mu)\, p(\mu)\, \mathrm{d}\mu}. \tag{20}$$

We begin by investigating the numerator of (20), which has the following form.

$$p(\tau \mid \mu)\, p(\mu) = \frac{\lambda}{2} e^{-\lambda|\tau-\mu|} \frac{\beta^\alpha}{\Gamma(\alpha)} \mu^{\alpha-1} e^{-\beta\mu} \tag{21}$$

$$= \frac{\lambda\beta^\alpha}{2\Gamma(\alpha)} \mu^{\alpha-1} \left[e^{-\lambda\tau} e^{-(\beta-\lambda)\mu} \mathbf{1}[\mu \le \tau] + e^{\lambda\tau} e^{-(\beta+\lambda)\mu} \mathbf{1}[\mu > \tau]\right] \tag{22}$$

$$= \frac{\lambda\beta^\alpha}{2\Gamma(\alpha)} \left[e^{-\lambda\tau}\, \mu^{\alpha-1} e^{-(\beta-\lambda)\mu} \mathbf{1}[\mu \le \tau] + e^{\lambda\tau}\, \mu^{\alpha-1} e^{-(\beta+\lambda)\mu} \mathbf{1}[\mu > \tau]\right] \tag{23}$$

$$= \frac{\lambda\beta^\alpha}{2\Gamma(\alpha)} \Bigg[e^{-\lambda\tau} \frac{\gamma(\alpha, (\beta-\lambda)\tau)}{(\beta-\lambda)^\alpha} \frac{(\beta-\lambda)^\alpha}{\gamma(\alpha, (\beta-\lambda)\tau)} \mu^{\alpha-1} e^{-(\beta-\lambda)\mu} \mathbf{1}[\mu \le \tau]$$
$$+ e^{\lambda\tau} \frac{\Gamma(\alpha, (\beta+\lambda)\tau)}{(\beta+\lambda)^\alpha} \frac{(\beta+\lambda)^\alpha}{\Gamma(\alpha, (\beta+\lambda)\tau)} \mu^{\alpha-1} e^{-(\beta+\lambda)\mu} \mathbf{1}[\mu > \tau]\Bigg]. \tag{24}$$

The denominator of (20) is then as follows.

$$\int_0^\infty p(\tau \mid \mu)\, p(\mu)\, \mathrm{d}\mu = \frac{\lambda\beta^\alpha}{2\Gamma(\alpha)} \Bigg[e^{-\lambda\tau} \frac{\gamma(\alpha, (\beta-\lambda)\tau)}{(\beta-\lambda)^\alpha} \int_0^\tau \frac{(\beta-\lambda)^\alpha}{\gamma(\alpha, (\beta-\lambda)\tau)} \mu^{\alpha-1} e^{-(\beta-\lambda)\mu}\, \mathrm{d}\mu$$
$$+ e^{\lambda\tau} \frac{\Gamma(\alpha, (\beta+\lambda)\tau)}{(\beta+\lambda)^\alpha} \int_\tau^\infty \frac{(\beta+\lambda)^\alpha}{\Gamma(\alpha, (\beta+\lambda)\tau)} \mu^{\alpha-1} e^{-(\beta+\lambda)\mu}\, \mathrm{d}\mu\Bigg] \tag{25}$$

$$= \frac{\lambda\beta^\alpha}{2\Gamma(\alpha)} \left[e^{-\lambda\tau} \frac{\gamma(\alpha, (\beta-\lambda)\tau)}{(\beta-\lambda)^\alpha} + e^{\lambda\tau} \frac{\Gamma(\alpha, (\beta+\lambda)\tau)}{(\beta+\lambda)^\alpha}\right]. \tag{26}$$

The second equality follows by recognizing the integrands as the PDFs of truncated gamma distributions. Thus, the distribution of $\mu \mid \tau$ is

$$p(\mu \mid \tau) = \frac{e^{-\lambda\tau} \frac{\gamma(\alpha, (\beta-\lambda)\tau)}{(\beta-\lambda)^\alpha}}{e^{-\lambda\tau} \frac{\gamma(\alpha, (\beta-\lambda)\tau)}{(\beta-\lambda)^\alpha} + e^{\lambda\tau} \frac{\Gamma(\alpha, (\beta+\lambda)\tau)}{(\beta+\lambda)^\alpha}} \frac{(\beta-\lambda)^\alpha}{\gamma(\alpha, (\beta-\lambda)\tau)} \mu^{\alpha-1} e^{-(\beta-\lambda)\mu} \mathbf{1}[\mu \le \tau]$$
$$+ \frac{e^{\lambda\tau} \frac{\Gamma(\alpha, (\beta+\lambda)\tau)}{(\beta+\lambda)^\alpha}}{e^{-\lambda\tau} \frac{\gamma(\alpha, (\beta-\lambda)\tau)}{(\beta-\lambda)^\alpha} + e^{\lambda\tau} \frac{\Gamma(\alpha, (\beta+\lambda)\tau)}{(\beta+\lambda)^\alpha}} \frac{(\beta+\lambda)^\alpha}{\Gamma(\alpha, (\beta+\lambda)\tau)} \mu^{\alpha-1} e^{-(\beta+\lambda)\mu} \mathbf{1}[\mu > \tau], \tag{27}$$

which we recognize as the PDF of a $\mathsf{TGM}(\alpha, \beta, \lambda, \tau)$. □

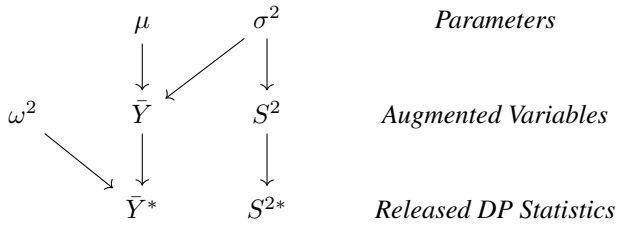


Figure 7: A graphical model representing the structure of the univariate Gaussian setting.

C Derivation of Full Conditionals for Univariate Gaussian Gibbs Sampler

In this section, we derive the full conditionals for the Gibbs sampler in Section 2.2. Recall that if $Y_1, \ldots, Y_n \overset{iid}{\sim} \mathcal{N}(\mu, \sigma^2)$, then the distributions of the sufficient statistics are

$$\bar{Y} \mid \mu, \sigma^2 \sim \mathcal{N}\left(\mu, \frac{\sigma^2}{n}\right), \qquad S^2 \mid \sigma^2 \sim \mathsf{Gam}\left(\frac{n-1}{2}, \frac{n-1}{2\sigma^2}\right), \tag{28}$$

where $(\bar{Y} \perp\!\!\!\perp S^2) \mid \mu, \sigma^2$. Assuming that the $Y_i \in [0, 1]$, then the statistics released via the Laplace Mechanism have distributions

$$\bar{Y}^* \mid \bar{Y} \sim \mathsf{Lap}\left(\bar{Y}, \frac{1}{\varepsilon_1 n}\right), \qquad S^{2*} \mid S^2 \sim \mathsf{Lap}\left(S^2, \frac{1}{\varepsilon_2 n}\right). \tag{29}$$

The release mechanism is such that $(\bar{Y}^* \perp\!\!\!\perp S^{2*}) \mid \bar{Y}, S^2$. It is convenient to replace the above distribution of $\bar{Y}^*$ with the following equivalent formulation, as used in [4, 5, 33].

$$\bar{Y}^* \mid \bar{Y}, \omega^2 \sim \mathcal{N}\left(\bar{Y}, \omega^2\right), \qquad \omega^2 \sim \mathsf{Exp}\left(\frac{\varepsilon_1^2 n^2}{2}\right). \tag{30}$$

Finally, we take the conjugate prior from the public setting

$$\sigma^2 \sim \mathsf{IG}\left(\frac{\nu_0}{2}, \frac{\nu_0 \sigma_0^2}{2}\right), \qquad \mu \mid \sigma^2 \sim \mathcal{N}\left(\mu_0, \frac{\sigma^2}{\kappa_0}\right). \tag{31}$$

Figure 7 presents the graphical model corresponding to this likelihood, which we can factor as follows.

$$\begin{aligned} p(\bar{Y}^*, S^{2*}, \bar{Y}, S^2, \omega^2 \mid \mu, \sigma^2) &= p(\bar{Y}^* \mid \bar{Y}, \omega^2)\, p(S^{2*} \mid S^2)\, p(\bar{Y} \mid \mu, \sigma^2)\, p(S^2 \mid \sigma^2)\, p(\omega^2) && (32) \\ &= \mathcal{N}(\bar{Y}^*; \bar{Y}, \omega^2)\, \mathsf{Lap}\left(S^{2*}; S^2, \frac{1}{\varepsilon_2 n}\right) \mathcal{N}\left(\bar{Y}; \mu, \frac{\sigma^2}{n}\right) \\ &\quad \mathsf{Gam}\left(S^2; \frac{n-1}{2}, \frac{n-1}{2\sigma^2}\right) \mathsf{Exp}\left(\omega^2; \frac{\varepsilon_1^2 n^2}{2}\right) && (33) \end{aligned}$$

We now examine each full conditional. The full conditional for μ is

$$p(\mu \mid \bar{Y}^*, S^{2*}, \bar{Y}, S^2, \omega^2, \sigma^2) \propto p(\bar{Y} \mid \mu, \sigma^2)\, p(\mu \mid \sigma^2) = \mathcal{N}\left(\bar{Y}; \mu, \frac{\sigma^2}{n}\right) \mathcal{N}\left(\mu; \mu_0, \frac{\sigma^2}{\kappa_0}\right), \tag{34}$$

which we recognize as the full conditional from the public setting, as derived in Chapter 5 of [16]. Similarly, the full conditional for σ^2 is

$$\begin{aligned} p(\sigma^2 \mid \bar{Y}^*, S^{2*}, \bar{Y}, S^2, \omega^2, \mu) &\propto p(\bar{Y} \mid \mu, \sigma^2)\, p(S^2 \mid \sigma^2)\, p(\sigma^2) && (35) \\ &= \mathcal{N}\left(\bar{Y}; \mu, \frac{\sigma^2}{n}\right) \mathsf{Gam}\left(S^2; \frac{n-1}{2}, \frac{n-1}{2\sigma^2}\right) \mathsf{IG}\left(\frac{\nu_0}{2}, \frac{\nu_0 \sigma_0^2}{2}\right). && (36) \end{aligned}$$

This is the same full conditional as in the public setting, as derived in Chapter 6 of [16]. Thus, the full conditionals have the form

$$\mu \mid \sigma^2, \omega^2, \bar{Y}, S^2, \bar{Y}^*, S^{2*} \sim \mathcal{N}\left(\frac{n\bar{Y} + \kappa_0\mu_0}{n+\kappa_0}, \frac{\sigma^2}{n+\kappa_0}\right) \tag{37}$$

$$\sigma^2 \mid \mu, \omega^2, \bar{Y}, S^2, \bar{Y}^*, S^{2*} \sim \mathsf{IG}\left(\frac{n+\nu_0}{2}, \frac{\nu_0\sigma_0^2 + (n-1)S^2 + n(\bar{Y}-\mu)^2}{2}\right). \tag{38}$$

We next examine the full conditional for $\bar{Y}$, which is

$$p(\bar{Y} \mid \bar{Y}^*, S^{2*}, S^2, \omega^2, \mu, \sigma^2) \propto p(\bar{Y}^* \mid \bar{Y}, \omega^2)\, p(\bar{Y} \mid \mu, \sigma^2) = \mathcal{N}(\bar{Y}^*; \bar{Y}, \omega^2)\, \mathcal{N}\left(\bar{Y}; \mu, \frac{\sigma^2}{n}\right). \tag{39}$$

We recognize this is equivalent to a Gaussian model with known variance and a Gaussian prior on the mean. By Chapter 5 of [16], the full conditional for $\bar{Y}$ is

$$\bar{Y} \mid \mu, \sigma^2, \omega^2, S^2, \bar{Y}^*, S^{2*} \sim \mathcal{N}\left(\frac{\frac{\bar{Y}^*}{\omega^2} + \frac{n\mu}{\sigma^2}}{\frac{1}{\omega^2} + \frac{n}{\sigma^2}}, \frac{1}{\frac{1}{\omega^2} + \frac{n}{\sigma^2}}\right). \tag{40}$$

The full conditional for ω^2 is as follows, expressed in terms of its inverse, $1/\omega^2$.

$$p(1/\omega^2 \mid \bar{Y}^*, S^{2*}, \bar{Y}, S^2, \mu, \sigma^2) \propto p(\bar{Y}^* \mid \bar{Y}, 1/\omega^2)\, p(1/\omega^2) \tag{41}$$

$$= \mathcal{N}(\bar{Y}^*; \bar{Y}, (1/\omega^2)^{-1})\, \mathsf{IG}\left(1/\omega^2; 1, \frac{\varepsilon_1^2 n^2}{2}\right). \tag{42}$$

We recognize this form from the Bayesian LASSO [33], which yields the distribution

$$1/\omega^2 \mid \mu, \sigma^2, \bar{Y}, S^2, \bar{Y}^*, S^{2*} \sim \mathsf{InvGaus}\left(\frac{\varepsilon_1 n}{|\bar{Y}^* - \bar{Y}|}, \varepsilon_1^2 n^2\right). \tag{43}$$

Finally, the full conditional for S^2 is

$$p(S^2 \mid \bar{Y}^*, S^{2*}, \bar{Y}, \omega^2, \mu, \sigma^2) \propto p(S^{2*} \mid S^2)\, p(S^2 \mid \sigma^2) \tag{44}$$

$$= \mathsf{Lap}\left(S^{2*}; S^2, \frac{1}{\varepsilon_2 n}\right) \mathsf{Gam}\left(S^2; \frac{n-1}{2}, \frac{n-1}{2\sigma^2}\right). \tag{45}$$

By Theorem 7 when $(n-1)/(2\sigma^2) > \varepsilon_2 n$, the distribution is

$$S^2 \mid \mu, \sigma^2, \omega^2, \bar{Y}, \bar{Y}^*, S^{2*} \sim \mathsf{TGM}\left(\frac{n-1}{2}, \frac{n-1}{2\sigma^2}, n\varepsilon_2, S^{2*}\right). \tag{46}$$

Note that by Corollary 1 of the main text, which is repeated in Section E below, since $Y_i \in [0,1]$ it follows that $\sigma^2 \leq 1/4$. Thus, we may use this full conditional when

$$\varepsilon_2 < \frac{n-1}{n} \cdot \frac{1}{2\sigma^2} \leq 2\,\frac{n-1}{n}. \tag{47}$$

If $\varepsilon_2 > 2(n-1)/n$, then the full conditional for S^2 is not necessarily a TGM without additional assumptions. If reasonable for the application, the analyst could impose the additional constraint that $\sigma^2 < (n-1)/(n\varepsilon_2)$, which is sufficient to ensure the full conditional for S^2 is of the form in (46). Alternatively, if n is large, an analyst could use an approximation of the full conditional for S^2, such as the proposals of [4, 5].

D List of Full Conditionals

In this section, we summarize the full conditional under different prior distributions for the DP univariate Gaussian setting with likelihood given by $Y_i \overset{iid}{\sim} \mathcal{N}_{[0,1]}(\mu, \sigma^2)$, $\bar{Y}^* \sim \mathsf{Lap}(\bar{Y}, 1/(\varepsilon_1 n))$, and $S^{2*} \sim \mathsf{Lap}(S^2, 1/(\varepsilon_2 n))$. Subscripts on a distribution denote truncation of the distribution to an interval.

If the prior is $\sigma^2 \sim \mathsf{IG}\left(\nu_0/2, \nu_0\sigma_0^2/2\right)$ and $\mu \mid \sigma^2 \sim \mathcal{N}\left(\mu_0, \sigma^2/\kappa_0\right)$ with no constraints enforced, then

$$\mu \mid \sigma^2, \omega^2, \bar{Y}, S^2, \bar{Y}^*, S^{2*} \sim \mathcal{N}\left(\frac{n\bar{Y} + \kappa_0\mu_0}{n + \kappa_0}, \frac{\sigma^2}{n + \kappa_0}\right) \tag{48}$$

$$\sigma^2 \mid \mu, \omega^2, \bar{Y}, S^2, \bar{Y}^*, S^{2*} \sim \mathsf{IG}_{\left(0, \frac{n-1}{2n\varepsilon_2}\right)}\left(\frac{n + \nu_0}{2}, \frac{\nu_0\sigma_0^2 + (n-1)S^2 + n(\bar{Y} - \mu)^2}{2}\right) \tag{49}$$

$$\bar{Y} \mid \mu, \sigma^2, \omega^2, S^2, \bar{Y}^*, S^{2*} \sim \mathcal{N}\left(\frac{\frac{\bar{Y}^*}{\omega^2} + \frac{n\mu}{\sigma^2}}{\frac{1}{\omega^2} + \frac{n}{\sigma^2}}, \frac{1}{\frac{1}{\omega^2} + \frac{n}{\sigma^2}}\right) \tag{50}$$

$$1/\omega^2 \mid \mu, \sigma^2, \bar{Y}, S^2, \bar{Y}^*, S^{2*} \sim \mathsf{InvGaus}\left(\frac{\varepsilon_1 n}{|\bar{Y}^* - \bar{Y}|}, \varepsilon_1^2 n^2\right) \tag{51}$$

$$S^2 \mid \mu, \sigma^2, \omega^2, \bar{Y}, \bar{Y}^*, S^{2*} \sim \mathsf{TGM}\left(\frac{n-1}{2}, \frac{n-1}{2\sigma^2}, n\varepsilon_2, S^{2*}\right). \tag{52}$$

If the prior is $\sigma^2 \sim \mathsf{IG}\left(\nu_0/2, \nu_0\sigma_0^2/2\right)$ and $\mu \mid \sigma^2 \sim \mathcal{N}\left(\mu_0, \sigma^2/\kappa_0\right)$ with constraints enforced, then

$$\mu \mid \sigma^2, \omega^2, \bar{Y}, S^2, \bar{Y}^*, S^{2*} \sim \mathcal{N}_{\left[1/2 - \sqrt{1/4 - \sigma^2}, 1/2 + \sqrt{1/4 - \sigma^2}\right]}\left(\frac{n\bar{Y} + \kappa_0\mu_0}{n + \kappa_0}, \frac{\sigma^2}{n + \kappa_0}\right) \tag{53}$$

$$\sigma^2 \mid \mu, \omega^2, \bar{Y}, S^2, \bar{Y}^*, S^{2*} \sim \mathsf{IG}_{\left[0, \min\{\mu(1-\mu), \frac{n-1}{2n\varepsilon_2}\right)}\left(\frac{n + \nu_0}{2}, \frac{\nu_0\sigma_0^2 + (n-1)S^2 + n(\bar{Y} - \mu)^2}{2}\right) \tag{54}$$

$$\bar{Y} \mid \mu, \sigma^2, \omega^2, S^2, \bar{Y}^*, S^{2*} \sim \mathcal{N}_{\left[1/2 - \sqrt{1/4 - \frac{n-1}{n}S^2}, 1/2 + \sqrt{1/4 - \frac{n-1}{n}S^2}\right]}\left(\frac{\frac{\bar{Y}^*}{\omega^2} + \frac{n\mu}{\sigma^2}}{\frac{1}{\omega^2} + \frac{n}{\sigma^2}}, \frac{1}{\frac{1}{\omega^2} + \frac{n}{\sigma^2}}\right) \tag{55}$$

$$1/\omega^2 \mid \mu, \sigma^2, \bar{Y}, S^2, \bar{Y}^*, S^{2*} \sim \mathsf{InvGaus}\left(\frac{\varepsilon_1 n}{|\bar{Y}^* - \bar{Y}|}, \varepsilon_1^2 n^2\right) \tag{56}$$

$$S^2 \mid \mu, \sigma^2, \omega^2, \bar{Y}, \bar{Y}^*, S^{2*} \sim \mathsf{TGM}_{\left(0, \frac{n}{n-1}\bar{Y}(1-\bar{Y})\right]}\left(\frac{n-1}{2}, \frac{n-1}{2\sigma^2}, n\varepsilon_2, S^{2*}\right). \tag{57}$$

If the prior is $p(\mu, \sigma^2) \propto 1$ with no constraints enforced, then

$$\mu \mid \sigma^2, \omega^2, \bar{Y}, S^2, \bar{Y}^*, S^{2*} \sim \mathcal{N}\left(\bar{Y}, \frac{\sigma^2}{n}\right) \tag{58}$$

$$\sigma^2 \mid \mu, \omega^2, \bar{Y}, S^2, \bar{Y}^*, S^{2*} \sim \mathsf{IG}_{\left(0, \frac{n-1}{2n\varepsilon_2}\right)}\left(\frac{n-2}{2}, \frac{(n-1)S^2 + n(\bar{Y} - \mu)^2}{2}\right) \tag{59}$$

$$\bar{Y} \mid \mu, \sigma^2, \omega^2, S^2, \bar{Y}^*, S^{2*} \sim \mathcal{N}\left(\frac{\frac{\bar{Y}^*}{\omega^2} + \frac{n\mu}{\sigma^2}}{\frac{1}{\omega^2} + \frac{n}{\sigma^2}}, \frac{1}{\frac{1}{\omega^2} + \frac{n}{\sigma^2}}\right) \tag{60}$$

$$1/\omega^2 \mid \mu, \sigma^2, \bar{Y}, S^2, \bar{Y}^*, S^{2*} \sim \mathsf{InvGaus}\left(\frac{\varepsilon_1 n}{|\bar{Y}^* - \bar{Y}|}, \varepsilon_1^2 n^2\right) \tag{61}$$

$$S^2 \mid \mu, \sigma^2, \omega^2, \bar{Y}, \bar{Y}^*, S^{2*} \sim \mathsf{TGM}\left(\frac{n-1}{2}, \frac{n-1}{2\sigma^2}, n\varepsilon_2, S^{2*}\right). \tag{62}$$

If the prior is $p(\mu, \sigma^2) \propto 1$ with constraints enforced, then

$$\mu \mid \sigma^2, \omega^2, \bar{Y}, S^2, \bar{Y}^*, S^{2*} \sim \mathcal{N}_{\left[1/2-\sqrt{1/4-\sigma^2}, 1/2+\sqrt{1/4-\sigma^2}\right]}\left(\bar{Y}, \frac{\sigma^2}{n}\right) \tag{63}$$

$$\sigma^2 \mid \mu, \omega^2, \bar{Y}, S^2, \bar{Y}^*, S^{2*} \sim \mathsf{IG}_{[0, \min\{\mu(1-\mu), \frac{n-1}{2n\varepsilon_2}\})}\left(\frac{n-2}{2}, \frac{(n-1)S^2 + n(\bar{Y}-\mu)^2}{2}\right) \tag{64}$$

$$\bar{Y} \mid \mu, \sigma^2, \omega^2, S^2, \bar{Y}^*, S^{2*} \sim \mathcal{N}_{\left[1/2-\sqrt{1/4-\frac{n-1}{n}S^2}, 1/2+\sqrt{1/4-\frac{n-1}{n}S^2}\right]}\left(\frac{\frac{\bar{Y}^*}{\omega^2} + \frac{n\mu}{\sigma^2}}{\frac{1}{\omega^2} + \frac{n}{\sigma^2}}, \frac{1}{\frac{1}{\omega^2} + \frac{n}{\sigma^2}}\right) \tag{65}$$

$$1/\omega^2 \mid \mu, \sigma^2, \bar{Y}, S^2, \bar{Y}^*, S^{2*} \sim \mathsf{InvGaus}\left(\frac{\varepsilon_1 n}{|\bar{Y}^* - \bar{Y}|}, \varepsilon_1^2 n^2\right) \tag{66}$$

$$S^2 \mid \mu, \sigma^2, \omega^2, \bar{Y}, \bar{Y}^*, S^{2*} \sim \mathsf{TGM}_{(0, \frac{n}{n-1}\bar{Y}(1-\bar{Y})]}\left(\frac{n-1}{2}, \frac{n-1}{2\sigma^2}, n\varepsilon_2, S^{2*}\right). \tag{67}$$

E Constraints for the Univariate Gaussian Setting

In this section, we provide proofs for results in Section 3 of the main text regarding constraints on parameters and statistics in the univariate Gaussian setting. The result regarding constraints on the parameters μ and σ^2 is restated below.

Theorem 2. *Let $Y_1, \ldots, Y_n \sim \mathcal{N}(\mu, \sigma^2)$ and suppose each $Y_i \in [0, 1]$. If μ is known, then $\sigma^2 \in [0, \mu(1-\mu)]$. If σ^2 is known, then $\mu \in [1/2 - \sqrt{1/4 - \sigma^2}, 1/2 + \sqrt{1/4 - \sigma^2}]$.*

Proof. Since $Y_i \in [0, 1]$, it follows that $Y_i^2 \leq Y_i$. and so by monotonicity of expectation,

$$\sigma^2 = Var[Y_i] = E[Y_i^2] - E[Y_i]^2 \leq E[Y_i] - E[Y_i]^2 = \mu - \mu^2. \tag{68}$$

Thus, since Y_i must have non-negative variance, if μ is known then $\sigma^2 \in [0, \mu(1-\mu)]$. On the other hand, if σ^2 is known then by (68), μ must be such that $\mu^2 - \mu + \sigma^2 \leq 0$. Applying the quadratic formula, we find that $\mu^2 - \mu + \sigma^2$ has roots $\mu = 1/2 \pm \sqrt{1/4 - \sigma^2}$. Recognizing that $\mu^2 - \mu + \sigma^2 \leq 0$ when μ is between these roots yields the desired result. □

The corollary of Theorem 2 is restated below.

Corollary 1. *If $Y_1, \ldots, Y_n \sim \mathcal{N}(\mu, \sigma^2)$ and $Y_i \in [0, 1]$, then $\sigma^2 \leq 1/4$.*

Proof. By Theorem 2 $\sigma^2 \leq \mu(1-\mu)$ and since $f(\mu) = \mu(1-\mu)$ attains its maximum value at $\mu = 1/2$, it follows that $f(\mu) \leq f(1/2) = 1/4$ for all μ. The result follows. □

The result regarding constraints on the sufficient statistics $\bar{Y}$ and S^2 is restated below.

Theorem 3. *Let $Y_1, \ldots, Y_n \sim \mathcal{N}(\mu, \sigma^2)$ and suppose each $Y_i \in [0, 1]$. If $\bar{Y}$ is known, then $S^2 \in [0, n/(n-1)\bar{Y}(1-\bar{Y})]$. If S^2 is known, then $\bar{Y} \in [1/2 - \sqrt{1/4 - (n-1)/n \cdot S^2}, 1/2 + \sqrt{1/4 - (n-1)/n \cdot S^2}]$.*

Proof. We begin by expressing S^2 in a more convenient form, as follows.

$$(n-1)S^2 = \sum_{i=1}^n (Y_i - \bar{Y})^2 = \sum_{i=1}^n Y_i^2 + n\bar{Y}^2 - 2\bar{Y}\sum_{i=1}^n Y_i = \sum_{i=1}^n Y_i^2 - n\bar{Y}^2. \tag{69}$$

Since $Y_i \in [0, 1]$, we have that $Y_i \leq Y_i^2$ and so

$$(n-1)S^2 = \sum_{i=1}^n Y_i^2 - n\bar{Y}^2 \leq \sum_{i=1}^n Y_i - n\bar{Y}^2 = n\bar{Y}(1-\bar{Y}). \tag{70}$$

Since $(n-1)S^2$ is a sum of squares, it is non-negative and so if $\bar{Y}$ is known, $S^2 \in [0, n/(n-1)\bar{Y}(1-\bar{Y})]$. On the other hand, if S^2 is known then $\bar{Y}$ must be such that $\bar{Y}^2 - \bar{Y} + (n-1)/n \cdot S^2 \leq 0$. By an analogous argument to the proof of Theorem 2, it follows that $\bar{Y} \in [1/2 - \sqrt{1/4 - (n-1)/n \cdot S^2}, 1/2 + \sqrt{1/4 - (n-1)/n \cdot S^2}]$. □

F Determining Whether the Posterior Distribution is Proper

In this section, we prove that the posterior distribution in the differentially private univariate Gaussian setting is not a proper probability distribution when $p(\mu, \sigma^2) \propto (\sigma^2)^{-1}$, but is a proper probability distribution when $p(\mu, \sigma^2) \propto 1$. The results are reproduced below.

Theorem 4. *For confidential data $Y_1, \ldots, Y_n \overset{iid}{\sim} \mathcal{N}(\mu, \sigma^2)$ where sufficient statistics $\bar{Y}^* \sim \mathsf{Lap}(\bar{Y}, 1/(\varepsilon_1 n))$ and $S^{2*} \sim \mathsf{Lap}(S^2, 1/(\varepsilon_2 n))$ are released, if an analyst has prior $p(\mu, \sigma^2) \propto (\sigma^2)^{-1}$, then their posterior $p(\mu, \sigma^2 \mid \bar{Y}^*, S^{2*})$ is not a proper probability distribution.*

Proof. To begin, note that under the prior $p(\mu, \sigma^2) \propto (\sigma^2)^{-1}$, for observed $\bar{Y}^* = \bar{y}^*$ and $S^{2*} = s^{2*}$ the posterior distribution, if proper, should be

$$p(\mu, \sigma^2 \mid \bar{y}^*, s^{2*}) = \frac{p(\bar{y}^*, s^{2*} \mid \mu, \sigma^2)\, p(\mu, \sigma^2)}{\int_0^\infty \int_{-\infty}^\infty p(\bar{y}^*, s^{2*} \mid \mu, \sigma^2)\, p(\mu, \sigma^2)\, \mathrm{d}\mu\, \mathrm{d}\sigma^2} \tag{71}$$

$$= \frac{p(\bar{y}^*, s^{2*} \mid \mu, \sigma^2)\, (\sigma^2)^{-1}}{\int_0^\infty \int_{-\infty}^\infty p(\bar{y}^*, s^{2*} \mid \mu, \sigma^2)\, (\sigma^2)^{-1}\, \mathrm{d}\mu\, \mathrm{d}\sigma^2}. \tag{72}$$

The posterior is a proper probability distribution if the integral in the denominator of (72) is finite. We now examine this integral. Introducing the latent variables $\bar{Y} = \bar{y}$ and $S^2 = s^2$ and exploiting the conditional independence structure, this integral is equivalent to

$$\int_0^\infty \int_{-\infty}^\infty p(\bar{y}^*, s^{2*} \mid \mu, \sigma^2)\, (\sigma^2)^{-1}\, \mathrm{d}\mu\, \mathrm{d}\sigma^2 \tag{73}$$

$$= \int_0^\infty \int_{-\infty}^\infty \int_0^\infty \int_{-\infty}^\infty p(\bar{y}^*, s^{2*}, \bar{y}, s^2 \mid \mu, \sigma^2)\, (\sigma^2)^{-1}\, \mathrm{d}\bar{y}\, \mathrm{d}s^2\, \mathrm{d}\mu\, \mathrm{d}\sigma^2 \tag{74}$$

$$= \int_0^\infty (\sigma^2)^{-1} \int_{-\infty}^\infty \int_0^\infty \int_{-\infty}^\infty p(\bar{y}^* \mid \bar{y})\, p(s^{2*} \mid s^2)\, p(\bar{y} \mid \mu, \sigma^2)\, p(s^2 \mid \sigma^2)\ \mathrm{d}\bar{y}\, \mathrm{d}s^2\, \mathrm{d}\mu\, \mathrm{d}\sigma^2 \tag{75}$$

$$= \int_0^\infty (\sigma^2)^{-1} \left[\int_0^\infty p(s^{2*} \mid s^2)\, p(s^2 \mid \sigma^2)\, \mathrm{d}s^2 \right] \left[\int_{-\infty}^\infty \int_{-\infty}^\infty p(\bar{y}^* \mid \bar{y})\, p(\bar{y} \mid \mu, \sigma^2)\, \mathrm{d}\bar{y}\, \mathrm{d}\mu \right] \mathrm{d}\sigma^2. \tag{76}$$

We examine the double integral in the second bracket. If we exchange the order of integration,

$$\int_{-\infty}^\infty \int_{-\infty}^\infty p(\bar{y}^* \mid \bar{y})\, p(\bar{y} \mid \mu, \sigma^2)\, \mathrm{d}\mu\, \mathrm{d}\bar{y} = \int_{-\infty}^\infty p(\bar{y}^* \mid \bar{y}) \int_{-\infty}^\infty \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(\bar{y}-\mu)^2}{2\sigma^2}}\, \mathrm{d}\mu\, \mathrm{d}\bar{y}, \tag{77}$$

since we recognize the quantity under the inner integral as the density function of $\mathcal{N}(\bar{y}, \sigma^2)$, the integral is equal to 1. Thus,

$$\int_{-\infty}^\infty \int_{-\infty}^\infty p(\bar{y}^* \mid \bar{y})\, p(\bar{y} \mid \mu, \sigma^2)\, \mathrm{d}\mu\, \mathrm{d}\bar{y} = \int_{-\infty}^\infty p(\bar{y}^* \mid \bar{y})\, \mathrm{d}\bar{y} = \int_{-\infty}^\infty \frac{\varepsilon_1 n}{2} e^{-\varepsilon_1 n |\bar{y}^* - \bar{y}|}\, \mathrm{d}\bar{y} = 1, \tag{78}$$

since we recognize the quantity under the integral as the density function of $\mathsf{Lap}(\bar{y}^*, 1/(\varepsilon_1 n))$. By Fubini's Theorem, since this integral is finite and the quantity being integrated is nonnegative, it is equivalent to the quantity in brackets above. Thus, the quantity of interest reduces to

$$\int_0^\infty \int_{-\infty}^\infty p(\bar{y}^*, s^{2*} \mid \mu, \sigma^2)\, (\sigma^2)^{-1}\, \mathrm{d}\mu\, \mathrm{d}\sigma^2 = \int_0^\infty (\sigma^2)^{-1} \int_0^\infty p(s^{2*} \mid s^2)\, p(s^2 \mid \sigma^2)\, \mathrm{d}s^2\, \mathrm{d}\sigma^2. \tag{79}$$

If $s^{2*} > 0$, then we bound the inner integral below as follows.

$$\int_0^\infty p(s^{2*} \mid s^2)\, p(s^2 \mid \sigma^2)\, \mathrm{d}s^2 = \frac{\varepsilon_2 n \left(\frac{n-1}{2\sigma^2}\right)^{\frac{n-1}{2}}}{2\,\Gamma(\frac{n-1}{2})} \int_0^\infty e^{-\varepsilon_2 n |s^{2*} - s^2|} \left(s^2\right)^{\frac{n-1}{2}-1} e^{-\frac{n-1}{2\sigma^2} s^2}\, \mathrm{d}s^2 \tag{80}$$

$$= \frac{\varepsilon_2 n \left(\frac{n-1}{2\sigma^2}\right)^{\frac{n-1}{2}}}{2\,\Gamma(\frac{n-1}{2})} \left[e^{-\varepsilon_2 n s^{2*}} \int_0^{s^{2*}} \left(s^2\right)^{\frac{n-1}{2}-1} e^{-(\frac{n-1}{2\sigma^2} - \varepsilon_2 n) s^2}\, \mathrm{d}s^2 \right.$$
$$\left. + e^{\varepsilon_2 n s^{2*}} \int_{s^{2*}}^\infty \left(s^2\right)^{\frac{n-1}{2}-1} e^{-(\frac{n-1}{2\sigma^2} + \varepsilon_2 n) s^2}\, \mathrm{d}s^2 \right] \tag{81}$$

$$\geq \frac{\varepsilon_2 n \left(\frac{n-1}{2\sigma^2}\right)^{\frac{n-1}{2}}}{2\,\Gamma(\frac{n-1}{2})} \left[e^{-\varepsilon_2 n s^{2*}} \int_0^{s^{2*}} \left(s^2\right)^{\frac{n-1}{2}-1} e^{-(\frac{n-1}{2\sigma^2} + \varepsilon_2 n) s^2}\, \mathrm{d}s^2 \right.$$
$$\left. + e^{-\varepsilon_2 n s^{2*}} \int_{s^{2*}}^\infty \left(s^2\right)^{\frac{n-1}{2}-1} e^{-(\frac{n-1}{2\sigma^2} + \varepsilon_2 n) s^2}\, \mathrm{d}s^2 \right] \tag{82}$$

$$= \frac{\varepsilon_2 n \left(\frac{n-1}{2\sigma^2}\right)^{\frac{n-1}{2}}}{2\,\Gamma(\frac{n-1}{2})} e^{-\varepsilon_2 n s^{2*}} \int_0^\infty \left(s^2\right)^{\frac{n-1}{2}-1} e^{-(\frac{n-1}{2\sigma^2} + \varepsilon_2 n) s^2}\, \mathrm{d}s^2 \tag{83}$$

$$= \frac{\varepsilon_2 n}{2} \left(\frac{\frac{n-1}{2\sigma^2}}{\frac{n-1}{2\sigma^2} + \varepsilon_2 n} \right)^{\frac{n-1}{2}} \tag{84}$$

$$= \frac{\varepsilon_2 n}{2} \left(\frac{\frac{n-1}{2\varepsilon_2 n}}{\frac{n-1}{2\varepsilon_2 n} + \sigma^2} \right)^{\frac{n-1}{2}}. \tag{85}$$

The inequality follows since $e^x \geq e^{-x}$ for $x > 0$ and since $e^{-(a-b)x} \geq e^{-(a+b)x}$ for $a > b > 0$ and $x > 0$. If $s^{2*} \leq 0$, then the same bound holds by an analogous argument (with the integral from 0 to s^{2*} omitted). Thus, the quantity of interest is bounded as follows.

$$\int_0^\infty \int_{-\infty}^\infty p(\bar{y}^*, s^{2*} \mid \mu, \sigma^2)\, (\sigma^2)^{-1}\, \mathrm{d}\mu\, \mathrm{d}\sigma^2 \geq \frac{\varepsilon_2 n}{2} \int_0^\infty (\sigma^2)^{-1} \left(\frac{\frac{n-1}{2\varepsilon_2 n}}{\frac{n-1}{2\varepsilon_2 n} + \sigma^2} \right)^{\frac{n-1}{2}} \mathrm{d}\sigma^2 \tag{86}$$

$$\geq \frac{\varepsilon_2 n}{2} \int_0^1 (\sigma^2)^{-1} \left(\frac{\frac{n-1}{2\varepsilon_2 n}}{\frac{n-1}{2\varepsilon_2 n} + \sigma^2} \right)^{\frac{n-1}{2}} \mathrm{d}\sigma^2 \tag{87}$$

$$\geq \frac{\varepsilon_2 n}{2} \int_0^1 (\sigma^2)^{-1} \left(\frac{\frac{n-1}{2\varepsilon_2 n}}{\frac{n-1}{2\varepsilon_2 n} + 1} \right)^{\frac{n-1}{2}} \mathrm{d}\sigma^2 \tag{88}$$

$$= \frac{\varepsilon_2 n}{2} \left(\frac{\frac{n-1}{2\varepsilon_2 n}}{\frac{n-1}{2\varepsilon_2 n} + 1} \right)^{\frac{n-1}{2}} \int_0^1 (\sigma^2)^{-1}\, \mathrm{d}\sigma^2. \tag{89}$$

Since this integral diverges, the quantity of interest diverges. Thus, (72) is not a proper probability distribution. □

Theorem 5. *For confidential data $Y_1, \ldots, Y_n \overset{iid}{\sim} \mathcal{N}(\mu, \sigma^2)$ where $n > 3$ and sufficient statistics $\bar{Y}^* \sim$ Lap$(\bar{Y}, 1/(\varepsilon_1 n))$ and $S^{2*} \sim$ Lap$(S^2, 1/(\varepsilon_2 n))$ are released, if an analyst has prior $p(\mu, \sigma^2) \propto 1$, then their posterior $p(\mu, \sigma^2 \mid \bar{Y}^*, S^{2*})$ is a proper probability distribution.*

Proof. To begin, note that under the prior $p(\mu, \sigma^2) \propto 1$, for observed $\bar{Y}^* = \bar{y}^*$ and $S^{2*} = s^{2*}$ the posterior distribution, if proper, should be

$$p(\mu, \sigma^2 \mid \bar{y}^*, s^{2*}) = \frac{p(\bar{y}^*, s^{2*} \mid \mu, \sigma^2)\, p(\mu, \sigma^2)}{\int_0^\infty \int_{-\infty}^\infty p(\bar{y}^*, s^{2*} \mid \mu, \sigma^2)\, p(\mu, \sigma^2)\, \mathrm{d}\mu\, \mathrm{d}\sigma^2} \tag{90}$$

$$= \frac{p(\bar{y}^*, s^{2*} \mid \mu, \sigma^2)}{\int_0^\infty \int_{-\infty}^\infty p(\bar{y}^*, s^{2*} \mid \mu, \sigma^2)\, \mathrm{d}\mu\, \mathrm{d}\sigma^2}. \tag{91}$$

The posterior is a proper probability distribution if the integral in the denominator of (91) is finite. By analogy to the proof of Theorem 4, this integral is equivalent to

$$\int_0^\infty \int_{-\infty}^\infty p(\bar{y}^*, s^{2*} \mid \mu, \sigma^2)\, \mathrm{d}\mu\, \mathrm{d}\sigma^2 \tag{92}$$

$$= \int_0^\infty \left[\int_0^\infty p(s^{2*} \mid s^2)\, p(s^2 \mid \sigma^2)\, \mathrm{d}s^2 \right] \left[\int_{-\infty}^\infty \int_{-\infty}^\infty p(\bar{y}^* \mid \bar{y})\, p(\bar{y} \mid \mu, \sigma^2)\, \mathrm{d}\bar{y}\, \mathrm{d}\mu \right] \mathrm{d}\sigma^2. \tag{93}$$

By (78, the double integral in the second bracket is equal to 1. Thus, the quantity of interest reduces to

$$\int_0^\infty \int_{-\infty}^\infty p(\bar{y}^*, s^{2*} \mid \mu, \sigma^2)\, \mathrm{d}\mu\, \mathrm{d}\sigma^2 = \int_0^\infty \int_0^\infty p(s^{2*} \mid s^2)\, p(s^2 \mid \sigma^2)\, \mathrm{d}s^2\, \mathrm{d}\sigma^2. \tag{94}$$

Exchanging the order of integration yields

$$\int_0^\infty p(s^{2*} \mid s^2) \int_0^\infty p(s^2 \mid \sigma^2)\, \mathrm{d}\sigma^2\, \mathrm{d}s^2. \tag{95}$$

Since $n > 3$, the inner integral is the kernel of an inverse-gamma distribution and so

$$\int_0^\infty p(s^2 \mid \sigma^2)\, \mathrm{d}\sigma^2 = \int_0^\infty \frac{\left(\frac{n-1}{2\sigma^2}\right)^{\frac{n-1}{2}}}{\Gamma\left(\frac{n-1}{2}\right)} (s^2)^{\frac{n-1}{2}-1} e^{-\frac{(n-1)s^2}{2\sigma^2}}\, \mathrm{d}\sigma^2 \tag{96}$$

$$= \frac{\left(\frac{n-1}{2}\right)^{\frac{n-1}{2}}}{\Gamma\left(\frac{n-1}{2}\right)} (s^2)^{\frac{n-1}{2}-1} \int_0^\infty (\sigma^2)^{-\frac{n-3}{2}-1} e^{-\frac{(n-1)s^2}{2\sigma^2}}\, \mathrm{d}\sigma^2 \tag{97}$$

$$= \frac{\left(\frac{n-1}{2}\right)^{\frac{n-1}{2}}}{\Gamma\left(\frac{n-1}{2}\right)} (s^2)^{\frac{n-1}{2}-1} \frac{\Gamma\left(\frac{n-3}{2}\right)}{\left(\frac{(n-1)s^2}{2}\right)^{\frac{n-3}{2}}} \tag{98}$$

$$= \frac{\frac{n-1}{2}\Gamma\left(\frac{n-3}{2}\right)}{\Gamma\left(\frac{n-1}{2}\right)} \tag{99}$$

$$= \frac{n-1}{n-3}. \tag{100}$$

Thus,

$$\int_0^\infty p(s^{2*} \mid s^2) \int_0^\infty p(s^2 \mid \sigma^2)\, \mathrm{d}\sigma^2\, \mathrm{d}s^2 = \frac{n-1}{n-3} \int_0^\infty \frac{\varepsilon_2 n}{2} e^{-\varepsilon_2 n |s^{2*} - s^2|}\, \mathrm{d}s^2 \tag{101}$$

$$= \frac{n-1}{n-3} \left(\frac{1}{2} + \frac{1}{2} \mathrm{sgn}\left(s^{2*}\right) \left(1 - e^{-\varepsilon_2 n |s^{2*}|}\right) \right), \tag{102}$$

since we recognize the quantity under the integral as the density function of $\mathsf{Lap}(s^{2*}, 1/(\varepsilon_2 n))$, we plug in the form of its CDF at zero. By Fubini's Theorem, since this integral is finite and the quantity being integrated is nonnegative, it is equivalent to the original quantity of interest. Thus,

$$\int_0^\infty \int_{-\infty}^\infty p(\bar{y}^*, s^{2*} \mid \mu, \sigma^2)\, (\sigma^2)^0\, \mathrm{d}\mu\, \mathrm{d}\sigma^2 = \frac{n-1}{n-3} \left(\frac{1}{2} + \frac{1}{2} \mathrm{sgn}\left(s^{2*}\right) \left(1 - e^{-\varepsilon_2 n |s^{2*}|}\right) \right) < \infty \tag{103}$$

and so the posterior is a proper probability distribution. □

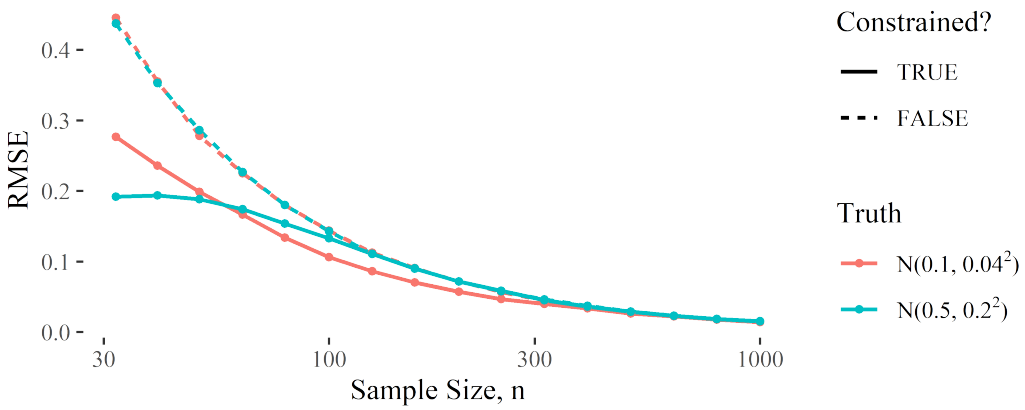


Figure 8: Root mean square error (RMSE) for estimating μ with the posterior mode for different sample sizes for the simulation in Section 4 of the main text. Results based on 10,000 simulated datasets $Y_i \in [0, 1]$ released with $\varepsilon_1 = \varepsilon_2 = 0.1$ and analyzed with prior $p(\mu, \sigma^2) \propto 1$. Analyses with constraints accounted for are solid lines and not accounted for are dashed lines. Data generating model is either $\mathcal{N}(\mu = 0.1, \sigma^2 = 0.04^2)$ (red) or $\mathcal{N}(\mu = 0.5, \sigma^2 = 0.2^2)$ (blue). Each Gibbs sampler is run for 20,000 iterations.

G Nonprivate Laplace-Uniform Posterior Distribution

In this section, we demonstrate that if $X \sim \mathsf{Lap}(\xi, 1/\lambda)$ and λ is known, then a uniform prior on ξ yields a proper, frequentist matching posterior. The main result is as follows.

Theorem 8. *Let $X \sim \mathsf{Lap}(\xi, 1/\lambda)$, where λ is known. Then the prior $p(\xi) \propto 1$ yields a proper posterior distribution that is frequentist matching.*

Proof. Applying properties of the Laplace distribution, if $X \sim \mathsf{Lap}(\xi, 1/\lambda)$, then $\xi - X \sim \mathsf{Lap}(0, 1/\lambda)$. Note that this is a pivotal quantity, since its distribution does not depend on any unknown parameters. Let $\ell_{\alpha/2}$ and $\ell_{1-\alpha/2}$ be the $\alpha/2$ and $(1 - \alpha/2)$ quantiles of $\mathsf{Lap}(\xi, 1/\lambda)$. Then,

$$P[X + \ell_{\alpha/2} \leq \xi \leq X + \ell_{1-\alpha/2}] = P[\ell_{\alpha/2} \leq \xi - X \leq \ell_{1-\alpha/2}] = 1 - \alpha \tag{104}$$

and so $[X + \ell_{\alpha/2}, X + \ell_{1-\alpha/2}]$ is a $(1 - \alpha)$ confidence interval for ξ.

In the Bayesian setting, if $p(\xi) \propto 1$ then

$$p(\xi \mid X) \propto p(X \mid \xi)\, p(\xi) \propto \exp\{-\lambda|X - \xi|\} = \exp\{-\lambda|\xi - X|\}.$$

Thus, the posterior distribution is $(\xi \mid X) \sim \mathsf{Lap}(X, 1/\lambda)$. Again applying properties of the Laplace distribution, $(\xi - X \mid X) \sim \mathsf{Lap}(0, 1/\lambda)$. Thus,

$$P[X + \ell_{\alpha/2} \leq \xi \leq X + \ell_{1-\alpha/2} \mid X] = P[\ell_{\alpha/2} \leq \xi - X \leq \ell_{1-\alpha/2} \mid X] = 1 - \alpha \tag{105}$$

and so $[X + \ell_{\alpha/2}, X + \ell_{1-\alpha/2}]$ is a $(1 - \alpha)$ credible interval. Since this interval is the same as the confidence interval above, this posterior distribution is frequentist matching. □

H Additional Simulation Studies

Figure 8 presents the RMSE from estimating μ with the posterior mode in the simulation study from Section 4 used to make Figure 2 of the main text. The RMSE is uniformly lower in the constrained analysis than in the unconstrained analysis. For the constrained analysis, a regression of log RMSE on log n has a slope of -1 (with $R^2 > 0.999$), indicating that RMSE decays proportionally to $1/n$

Figure 9 presents the results of a simulation study identical to that of Figure 2 and Figure 8 of the main text, but with ε increased by a factor of 10. There is little practical difference between the constrained and unconstrained analysis at this larger ε.

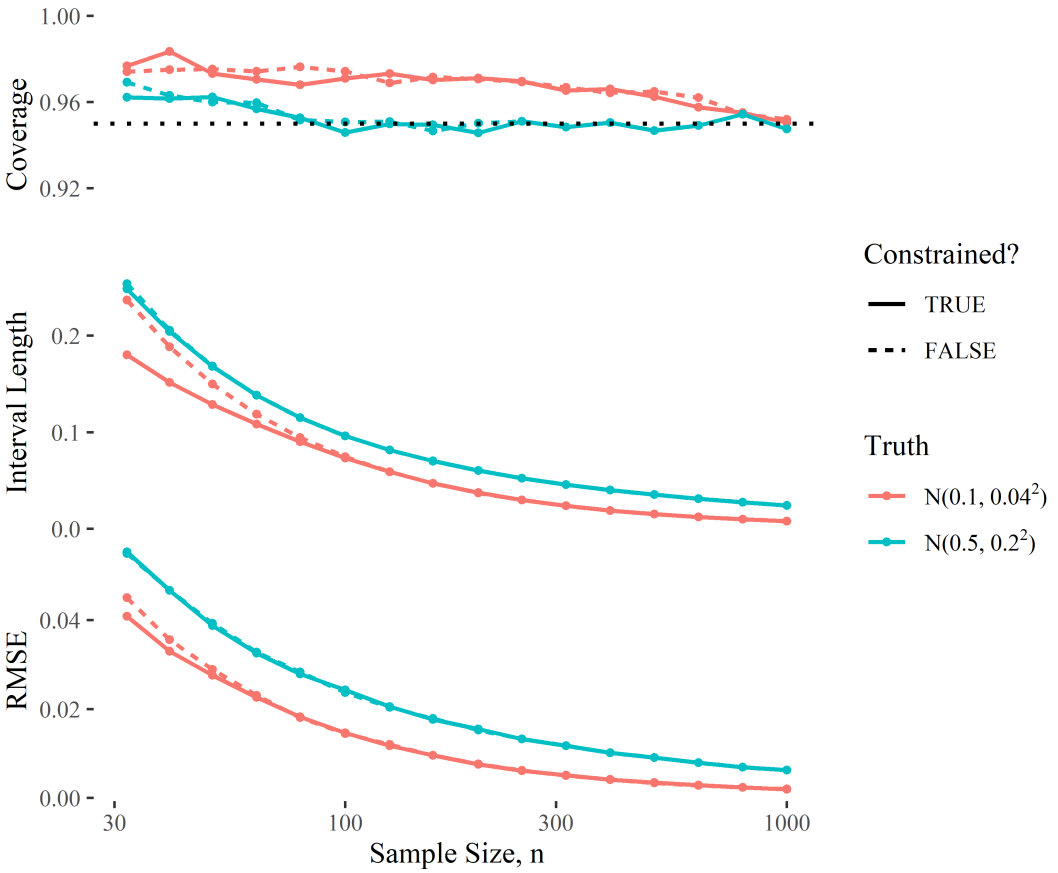


Figure 9: Empirical coverage rate of 95% HPD intervals for μ (top), average 95% HPD interval length for μ (middle), and root mean square error for estimating μ (bottom) for different sample sizes. Results based on 10,000 simulated datasets $Y_i \in [0,1]$ released with $\varepsilon_1 = \varepsilon_2 = 1$ and analyzed with prior $p(\mu, \sigma^2) \propto 1$. Analyses with constraints accounted for are solid lines and not accounted for are dashed lines. Data generating model is either $\mathcal{N}(\mu = 0.1, \sigma^2 = 0.04^2)$ (red) or $\mathcal{N}(\mu = 0.5, \sigma^2 = 0.2^2)$ (blue). Each Gibbs sampler is run for 20,000 iterations.

I Constraints in Simple Linear Regression

In this section, we derive constraints on the model proposed by [5] in the simple linear regression setting. Letting $\mathbf{X} = [\mathbf{1}_n \ \mathbf{x}_1]$, the authors release $\mathbf{X}^\top\mathbf{X}$, $\mathbf{X}^\top\mathbf{Y}$, and $\mathbf{Y}^\top\mathbf{Y}$, which include the following elements.

$$\mathbf{X}^\top\mathbf{X} = \begin{bmatrix} \mathbf{1}_n^\top\mathbf{1}_n & \mathbf{1}^\top\mathbf{x}_1 \\ \mathbf{x}_1^\top\mathbf{1}_n & \mathbf{x}_1^\top\mathbf{x}_1 \end{bmatrix}, \qquad \mathbf{X}^\top\mathbf{Y} = \begin{bmatrix} \mathbf{1}_n^\top\mathbf{Y} \\ \mathbf{x}_1^\top\mathbf{Y} \end{bmatrix}. \tag{106}$$

Note that $\mathbf{1}_n^\top\mathbf{1}_n = n$ and as sums of squares, $\mathbf{x}_1^\top\mathbf{x}_1 = \sum_{i=1}^n x_i^2 \geq 0$ and $\mathbf{Y}^\top\mathbf{Y} = \sum_{i=1}^n y_i^2 \geq 0$. The requirement that $x_i, y_i \in [0,1]$ for all $i \in \{1, \ldots, n\}$ implies the following additional constraints.

1. Since $x_i \leq 1$ for all i, it follows that $\mathbf{x}_1^\top\mathbf{1}_n = \mathbf{1}^\top\mathbf{x}_1 = \sum_{i=1}^n x_i \leq n$.
2. Since $y_i \leq 1$ for all i, it follows that $\mathbf{1}_n^\top\mathbf{Y} = \sum_{i=1}^n y_i \leq n$.
3. Since $x_i, y_i \geq 0$ for all i, it follows that $\mathbf{x}_1^\top\mathbf{Y} = \sum_{i=1}^n x_i y_i \geq 0$.
4. Since $x_i \leq 1$ for all i, it follows that $\mathbf{x}_1^\top\mathbf{Y} = \sum_{i=1}^n x_i y_i \leq \sum_{i=1}^n y_i = \mathbf{1}_n^\top\mathbf{Y}$.
5. Since $y_i \leq 1$ for all i, it follows that $\mathbf{x}_1^\top\mathbf{Y} = \sum_{i=1}^n x_i y_i \leq \sum_{i=1}^n x_i = \mathbf{1}_n^\top\mathbf{x}_1$.
6. Since $0 \leq x_i \leq 1$ for all i, $x_i^2 \leq x_i$ and so $\mathbf{x}_1^\top\mathbf{x}_1 = \sum_{i=1}^n x_i^2 \leq \sum_{i=1}^n x_i = \mathbf{1}_n^\top\mathbf{x}_1$.

7. Since $0 \leq y_i \leq 1$ for all i, $y_i^2 \leq y_i$ and so $\mathbf{Y}^\top \mathbf{Y} = \sum_{i=1}^n y_i^2 \leq \sum_{i=1}^n y_i = \mathbf{1}_n^\top \mathbf{Y}$.

The form of the model implies constraints on the parameters θ_0, θ_1, and σ^2. We can rewrite the model in the following form, where for each $i \in \{1, \ldots, n\}$,

$$y_i \sim \mathcal{N}(\theta_0 + \theta_1 x_i, \sigma^2). \tag{107}$$

For the regression coefficients, θ_0 and θ_1, we observe that by monotonicity of expectation, since $y_i \in [0, 1]$, it follows that $E[y_i] = \theta_0 + \theta_1 x_i \in [0, 1]$. Using this fact, for a given value of θ_1, we obtain the following constraints on θ_0.

1. When $\theta_1 \geq 0$, since $x_i \geq 0$ for all i, it follows that $\theta_0 \leq \theta_0 + \theta_1 x_i \leq 1$.
2. When $\theta_1 \geq 0$, since $x_i \leq 1$ for all i, it follows that $\theta_0 + \theta_1 \geq \theta_0 + \theta_1 x_i \geq 0$ and so $\theta_0 \geq -\theta_1$.
3. When $\theta_1 < 0$, since $x_i \geq 0$ for all i, it follows that $\theta_0 \geq \theta_0 + \theta_1 x_i \geq 0$.
4. When $\theta_1 < 0$, since $x_i \leq 1$ for all i, it follows that $\theta_0 + \theta_1 \leq \theta_0 + \theta_1 x_i \leq 1$ and so $\theta_0 \leq 1 - \theta_1$.

We summarize these constraints as follows.

$$\begin{cases} 0 \leq \theta_0 \leq 1 - \theta_1, & \text{if } \theta_1 < 0; \\ -\theta_1 \leq \theta_0 \leq 1, & \text{if } \theta_1 \geq 0. \end{cases} \tag{108}$$

Finally, we use a similar argument to Corollary 1 to obtain a constraint on σ^2. By monotonicity of expectation and since $y_i \in [0, 1]$ implies $y_i^2 \leq y_i$, for all i,

$$\sigma^2 = Var[y_i] = E[y_i^2] - E[y_i]^2 \leq E[y_i] - E[y_i]^2 \leq 1/4. \tag{109}$$

The final inequality is obtained by observing that the quadratic $f(x) = x - x^2$ is minimized at $x = 1/2$ and has minimum $f(1/2) = 1/4$. Since the variance of y_i must be nonnegative, it follows that $\sigma^2 \in [0, 1/4]$.