

Differentially Private Methods for Managing Model Uncertainty in Linear Regression Models

Víctor Peña

*Department d'Estadística i Investigació Operativa
Universitat Politècnica de Catalunya
Barcelona, Spain*

VICTOR.PENA.PIZARRO@UPC.EDU

Andrés F. Barrientos

*Department of Statistics
Florida State University
Tallahassee, FL 32306, USA*

ABARRIENTOS@FSU.EDU

Editor: Moritz Hardt

Abstract

In this work, we propose differentially private methods for hypothesis testing, model averaging, and model selection for normal linear models. We propose Bayesian methods based on mixtures of g -priors and non-Bayesian methods based on likelihood-ratio statistics and information criteria. The procedures are asymptotically consistent and straightforward to implement with existing software. We focus on practical issues such as adjusting critical values so that hypothesis tests have adequate type I error rates and quantifying the uncertainty introduced by the privacy-ensuring mechanisms.

Keywords: Confidential Data, Regression, Bayesian Methods, Information Criteria

1. Introduction

Differential privacy (Dwork et al., 2006) is a formal framework for quantifying the privacy of randomized algorithms. Its theoretical properties have been studied extensively (see e.g. Dwork et al. (2014)) and it has been adopted by companies such as Google, Apple, and Microsoft (Garfinkel et al., 2018), as well as institutions like the U.S. Census Bureau (Abowd, 2018).

In this article, we develop differentially private methods for normal linear models. We propose differentially private hypothesis tests for comparing nested models (in Section 4) as well as methods for model averaging and selection (in Section 5). We consider Bayesian methods based on mixtures of g -priors (Liang et al., 2008) and non-Bayesian methods that are built upon likelihood-ratio statistics and information criteria.

In Bayesian hypothesis testing and model selection, prior distributions must be chosen carefully because their effect does not vanish as the sample size grows (Bayarri et al., 2012). Our work is based on mixtures of g -priors because, when combined with right-Haar priors on the common parameters, they satisfy a list of appealing criteria proposed in Bayarri et al. (2012). They are also conveniently implemented in the R package `library(BAS)` (Clyde, 2020).

From a non-Bayesian perspective, we work with likelihood-ratio tests and information criteria. We make this choice because of their theoretical properties, intuitive appeal, and ease of use. The class of information criteria we consider includes the Akaike Information Criterion (AIC; Akaike (1974)) and the Bayesian Information Criterion (BIC; Schwarz (1978)), among others. Although we focus on normal linear models, information criteria are useful in model averaging and selection problems for more complex models. We recommend the monograph Claeskens et al. (2008) for an extensive overview of the approach.

We enforce differential privacy with well-established techniques. For hypothesis testing, we use the *subsample and aggregate* technique (Nissim et al., 2007; Smith, 2011), which consists in splitting the data into subgroups and releasing perturbed averages. For model averaging and selection, we use *sufficient-statistic perturbation*, which consists in releasing a noisy version of a sufficient statistic (see, for example, McSherry and Mironov (2009); Vu and Slavkovic (2009) and Bernstein and Sheldon (2019)).

1.1 Related Work

There is a growing literature on differentially private methods for linear regression models. For example, Amitai and Reiter (2018) estimate quantiles and posterior tail probabilities of coefficients, Barrientos et al. (2019) test the significance of individual regression coefficients, and Ferrando et al. (2022) propose methods for point and interval estimation that can be applied to the normal linear model. Lei et al. (2018) consider the problem of model selection based on information criteria, but do not consider model averaging or Bayesian approaches, and Bernstein and Sheldon (2019) propose a method for sampling from posterior distributions on regression coefficients, but do not consider hypothesis testing, model averaging or selection.

Our methods for hypothesis testing involve data-splitting and censoring. Both operations can induce bias in the outputs. Evans et al. (2020) considers the effects of censoring in estimates that are asymptotically normal and proposes strategies to correct the bias induced by censoring. Covington et al. (2021) uses bags of little bootstraps (Kleiner et al., 2012) to find unbiased estimates and valid confidence intervals. Alternatively, Ferrando et al. (2022) use the parametric bootstrap for bias correction. In hypothesis testing, the bias induced by data-splitting and censoring leads to inappropriate critical values for rejecting null hypotheses. We address this issue by simulating the distribution of the differentially private test statistics under the null hypothesis and then find corrected critical values that are adequately calibrated.

In general, our Bayesian methodology draws from the objective Bayesian literature for hypothesis testing, averaging, and selection, especially from Liang et al. (2008) and Bayarri et al. (2012). In these references, the classes of priors we consider here are proposed after showing that they satisfy a list of conceptually appealing criteria.

1.2 Main Contributions

In Section 4, we argue that working on a logarithmic scale is natural for defining differentially private Bayes factors. We show that the methods are asymptotically consistent under regularity conditions that are similar to the ones needed for consistency when there are no privacy constraints. In Section 4.2, we describe a simple procedure to quantify the effects

of the privacy-ensuring mechanisms. Then, in Section 4.3, we study the effects of censoring and data-splitting.

In Section 5, we use sufficient-statistic perturbation to define methods for model averaging and selection. If we took a naive approach to the problem, the variance of the perturbation term would grow exponentially in the number of predictors. With sufficient-statistic perturbation, the variance of the perturbation term grows quadratically in the number of predictors. The methods for model selection are consistent under conditions that are similar to those needed for consistency without privacy constraints. In Section 5.1, we propose a strategy that quantifies the uncertainty introduced by the mechanisms that is analogous to the one pursued in Section 4.2.

From a practical point of view, we give guidelines for maximizing the statistical utility of the methods in finite samples. We illustrate the performance of our methods in Sections 4.4 and 5.2. We include additional results from the simulation studies in the Appendix.

The proofs of the Propositions stated in the main text can be found in the Appendix A.

1.3 Notation

Extrema appear often in Section 4 because the methods are based on censored statistics. Our notation for them is $a \vee b = \max(a, b)$ and $a \wedge b = \min(a, b)$. The censored statistics are of the form $T^c = (T \vee a) \wedge b$: that is, $T^c = a$ whenever $T \leq a$ and $T^c = b$ whenever $T \geq b$ (with the understanding that $a < b$).

We use the following notation for probability distributions. The p -variate normal distribution with mean μ and covariance matrix Σ is $N_p(\mu, \Sigma)$, the Laplace distribution with location parameter μ and scale parameter b is $L_1(\mu, b)$, and the $(p \times p)$ -dimensional Wishart distribution with degree of freedom $n > p - 1$ and positive-definite scale matrix V is $W_p(n, V)$.

In the case of matrices, all of them are assumed to be full-rank unless otherwise stated. Our notation for basic matrix operations and special matrices is as follows. The matrix transpose of A is A' , the perpendicular projection operator onto the column space of A is $P_A = A(A'A)^{-1}A'$, and the upper Cholesky factor of A is $A^{1/2}$. The $(n \times n)$ -dimensional zero matrix is $0_{n \times n}$ and the $(n \times n)$ -dimensional identity matrix is I_n . For vectors, the usual p -norm on \mathbb{R}^d is $\|\cdot\|_p$. The zero vector is $0_n = (0, 0, \dots, 0)'$ and the vector of ones is $1_n = (1, 1, \dots, 1)'$. The expectation of a random variable X is $\mathbb{E}(X)$ and the variance is $\text{Var}(X)$.

2. Brief Review of Differential Privacy

Differential privacy (Dwork et al., 2006, 2014) is a probabilistic property of randomized algorithms. Intuitively, differential privacy limits how much we can learn about individual entries in a data set.

In the literature, randomized algorithms that ensure differential privacy are referred to as *mechanisms*. Conceptually, mechanisms are functions \mathcal{M} that take data D as inputs and output a random $\mathcal{M}(D)$ that is, in some sense, private. In the definition of differential privacy, a key notion is that of *neighboring data sets*. Two data sets D and \tilde{D} are neighbors, which is denoted by $D \sim \tilde{D}$, if they only differ in one row.

Definition 1 ((ε, δ)-differential privacy) *Let $\varepsilon > 0$ and $0 \leq \delta \leq 1$. A mechanism \mathcal{M} satisfies (ε, δ)-differential privacy if, for all $D \sim \tilde{D}$ any \mathcal{M} -measurable set S , we have $P[\mathcal{M}(D) \in S] \leq \exp(\varepsilon)P[\mathcal{M}(\tilde{D}) \in S] + \delta$.*

When $\delta = 0$, we retrieve ε -differential privacy, which is the most popular formalization of privacy in the literature. When ε is small, the privacy of \mathcal{M} increases; in such cases, the probability distributions of $\mathcal{M}(D)$ and $\mathcal{M}(\tilde{D})$ are forced to be similar, so the output of \mathcal{M} is not very sensitive to small changes in D . On the other hand, as ε increases, the distributions of $\mathcal{M}(D)$ and $\mathcal{M}(\tilde{D})$ are allowed to be more different, which decreases the privacy of the output. When $0 < \delta \leq 1$, Definition 1 allows $\mathcal{M}(D)$ to release outputs that lead to a “high privacy loss” with probability δ (see Section 2 in Dwork et al. (2014) for details).

In this article, $\mathcal{M}(D)$ are perturbed versions of confidential statistics $T(D)$. The scale of the perturbation depends on the global sensitivity of $T(D)$, a concept we define below.

Definition 2 (Global sensitivity) *Let T be a statistic mapping data to \mathbb{R}^d . The global sensitivity of T is defined as $\Delta_p = \sup_{D \sim \tilde{D}} \|T(D) - T(\tilde{D})\|_p$.*

A key property of differential privacy is that transformations of (ε, δ)-differentially private statistics are (ε, δ)-differentially private. In the literature, this is known as the *post-processing* property of differential privacy. We use this property frequently; for example, we use it to find posterior probabilities of hypotheses given Bayes factors.

3. Brief Review of Bayes Factors and Information Criteria

In this section, we review basic facts about Bayes factors and information criteria that are relevant for our purposes. This is not meant to be a comprehensive review; we refer the reader to Berger and Pericchi (2001), Liang et al. (2008), and Bayarri et al. (2012) for further background on Bayesian methods and Claeskens et al. (2008) for information criteria.

3.1 Hypothesis Testing

In Section 4, we introduce differentially private methods for hypothesis testing. We cover Bayesian approaches that are based on Bayes factors and non-Bayesian approaches that are based on likelihood-ratio tests and information criteria. Now, we review some core concepts in Bayesian and non-Bayesian testing that are helpful for contextualizing our work.

Let $y = (y_1, \dots, y_n)$ be a vector collecting independent and identically distributed observations from a statistical model with sampling density $f(y | \theta) = \prod_{i=1}^n f(y_i | \theta)$, for $\theta \in \Theta$. Our goal is testing $H_0 : \theta \in \Theta_0$ against $H_1 : \theta \in \Theta_1$, where Θ_0 and Θ_1 are disjoint subsets of Θ .

In the Bayesian paradigm, all unknowns have probability distributions associated to them, including the hypotheses H_0 and H_1 and the parameter θ . Before observing any data, the uncertainty in the hypotheses is reflected in the prior probabilities $P(H_0)$ and $P(H_1) = 1 - P(H_0)$. The uncertainty in θ is usually expressed conditionally through the prior distributions $\pi(\theta | H_0)$ and $\pi(\theta | H_1)$. Upon observing data, the uncertainty about θ

and the hypotheses is updated in the posterior distribution, which is simply the conditional distribution of these unknowns given the data.

A key quantity in Bayesian hypothesis testing is the Bayes factor, which is defined as

$$B_{10} = \int_{\theta \in \Theta_1} \pi(\theta | H_1) f(y | \theta) \nu(d\theta) / \int_{\theta \in \Theta_0} \pi(\theta | H_0) f(y | \theta) \nu(d\theta)$$

for a dominating measure $\nu(\cdot)$.

Bayes factors are important in Bayesian testing for several reasons. For example, if $o_{10} = P(H_1)/P(H_0)$ are the prior odds of the hypotheses, then the posterior probability of H_1 is

$$P(H_1 | y) = o_{10} B_{10} / (1 + o_{10} B_{10}),$$

which depends on the data only through the Bayes factor B_{10} . Bayes factors can be motivated as ratios of integrated likelihoods that quantify the evidence in favor of H_1 relative to H_0 (Berger et al., 1999). There have been efforts to categorize the strength of evidence in favor or against H_1 based on the magnitude of B_{10} alone. Perhaps the most popular categorization is “Jeffreys’ scale of evidence” (Jeffreys (1939), see Table 1).

Table 1: Jeffreys’ scale of evidence for Bayes factors (Jeffreys, 1939)

Bayes factor	Interpretation
$B_{10} < 1$	H_0 supported
$1 < B_{10} < 10^{1/2}$	Evidence against H_0 , but not worth more than a bare mention
$10^{1/2} < B_{10} < 10$	Evidence against H_0 substantial
$10 < B_{10} < 10^{3/2}$	Evidence against H_0 strong
$10^{3/2} < B_{10} < 10^2$	Evidence against H_0 very strong
$B_{10} > 10^2$	Evidence against H_0 decisive

Logarithms of Bayes factors are naturally symmetric about zero. To see this, assume that H_0 and H_1 are equally likely *a priori*. Then, the posterior probability of H_1 is the standard logistic function in $\log B_{10}$: that is, $P(H_1 | D) = 1/[1 + \exp(-\log B_{10})]$. Changing the sign of $\log B_{10}$ leads to the complement $1 - P(H_1 | D)$, and $P(H_1 | D) = 1/2$ if and only if $\log B_{10} = 0$.

In Bayesian hypothesis testing, the priors $\pi(\theta | H_0)$ and $\pi(\theta | H_1)$ must be chosen carefully. Vague priors on θ , which are commonplace in estimation problems, can lead to posterior probabilities that overwhelmingly support H_0 no matter what the data are. This phenomenon is known in the literature as Lindley’s paradox (see Lindley (1957) and Robert (2014)), and it can occur when H_1 represents a larger set than H_0 . In normal linear regression problems, mixtures of g -priors have been studied carefully in Liang et al. (2008) and Bayarri et al. (2012). They avoid Lindley’s paradox by having a fixed scale matrix and they have other desirable properties such as invariance of Bayes factors with respect to changes of measurement units and large-sample consistency. In Sections 4 and 5, we work with priors within this class.

From a non-Bayesian perspective, we propose working with likelihood-ratio tests and information criteria. The likelihood ratio of H_1 to H_0 is defined as

$$\Lambda_{10} = \max_{\theta \in \Theta_1} f(y | \theta) / \max_{\theta \in \Theta_0} f(y | \theta).$$

The likelihood ratio Λ_{10} is similar to the Bayes factor B_{10} , but instead of averaging the likelihoods with weights given by $\pi(\theta | H_1)$ and $\pi(\theta | H_0)$, the likelihoods are maximized under H_1 and H_0 . Likelihood ratio tests are standard within the field of statistics and their properties are well-characterized (see, for example, Lehmann and Romano (2005)). Under mild conditions, the transformed log-likelihood ratio $2 \log \Lambda_{10}$ is asymptotically distributed as chi-squared under H_0 and consistent under H_1 .

Finally, we define the class of information criteria

$$I_{10} = n^{-\rho/2} \Lambda_{10},$$

which encompasses AIC for $\rho = 2p/\log n$, BIC for $\rho = p$, and the likelihood ratio statistic Λ_{10} for $\rho = 0$. For a fixed ρ , the information criterion I_{10} can be interpreted as a penalized likelihood ratio, where ρ acts as a penalty for model complexity. For the normal linear model, likelihood-ratio tests and AIC fail to be consistent under H_0 ; BIC, on the other hand, is consistent under H_0 and H_1 . We refer the reader to Chapter 4 in the monograph Claeskens et al. (2008) for a more general version of this result and a discussion on related issues.

3.2 Model Averaging and Selection

In Section 5, we focus on model averaging and selection. The context is a regression problem where there is an outcome variable $Y \in \mathbb{R}^n$ and p predictors collected in a design matrix $X \in \mathbb{R}^{n \times p}$. We do not know which variables in X , if any, should be included in our model for Y given X . This type of uncertainty is often referred to as *model uncertainty*. For an introduction to the topic with a strong Bayesian flavor, we recommend Draper (1995).

Model uncertainty can be parameterized through a binary vector $\gamma \in \{0, 1\}^p$ that indicates active predictors: $\gamma_i = 0$ if the i th predictor is not active and $\gamma_i = 1$ if it is. Conceptually, we assume that there is a true model generating the data identified by $T \in \{0, 1\}^p$. Given finite data D , we do not know what the truth (T) is. From a Bayesian perspective, we can put a prior on γ and find posterior probabilities to quantify this uncertainty; from a non-Bayesian perspective, we can find a point estimate of γ or average our uncertainty over it with rules inspired by Bayesian procedures.

For each model, which we identify by its active predictors in $\gamma \in \{0, 1\}^p$, we can compute a Bayes factor or information criterion relative to the null model, which does not include any active predictors. We denote these null-based Bayes factors and information criteria B_{γ_0} and I_{γ_0} , respectively. With these, we can perform model selection after maximizing over γ or we can find model-averaged estimates with weights proportional to B_{γ_0} or I_{γ_0} .

4. Hypothesis Testing

In this section, we describe differentially private methods for testing a null hypothesis H_0 against an alternative H_1 . The methodology described here can be applied in general, but in Section 4.1 we focus on nested linear regression models, for which we have found theoretical guarantees.

Our methods are based on the subsample and aggregate technique (Nissim et al., 2007). It consists in splitting the data into M disjoint subgroups, computing statistics within the subgroups, and averaging the results. The output is made differentially private by adding a perturbation term η . The variance of η is increasing in the global sensitivity (i.e., the range) of the statistics involved.

We apply the subsample and aggregate technique as follows. First, we split the data into M disjoint subgroups of sample size b_1, b_2, \dots, b_M ($\sum_{i=1}^M b_i = n$) and compute censored statistics $T_i^c = (T_i \vee L) \wedge U \in [L, U]$ for $i \in \{1, 2, \dots, M\}$, where T_i are raw statistics computed from confidential data. After censoring the T_i , we know that they have global sensitivity $\Delta = U - L$. Finally, we release the noisy average $\sum_{i=1}^M T_i^c / M + \eta$, where η is a random perturbation term that ensures (ϵ, δ) -differential privacy. If $\delta = 0$, then $\eta \sim L_1(0, \Delta / (M\epsilon))$; if $0 < \delta \leq 1$, then $\eta \sim N_1(0, a\Delta^2 / (2M\epsilon))$, where a is a constant that can be computed with Algorithm 1 in Balle and Wang (2018).

The confidential statistics T_i are logarithms of Bayes factors or logarithms of information criteria. We justify working on a logarithmic scale for Bayes factors in the subsequent paragraphs. A similar argument can be used to justify this choice for information criteria.

Let $B_{10,i}^c$ be the censored Bayes factor in the i th subgroup. Naively, one could release the noisy average $\sum_{i=1}^M B_{10,i}^c / M + \eta$. Unfortunately, that approach has undesirable properties. Under both the Laplace and analytic Gaussian mechanisms, η is supported on \mathbb{R} and symmetric around zero, whereas B_{10} is always non-negative and shows equal support to H_0 and H_1 when $B_{10} = 1$. If there are no privacy constraints, Bayes factors satisfy $B_{01} = B_{10}^{-1}$, but in general $\sum_{i=1}^M B_{01,i}^c / M + \eta \neq (\sum_{i=1}^M B_{10,i}^c / M + \eta)^{-1}$.

Alternatively, we propose working on a logarithmic scale, defining

$$\log \tilde{B}_{10} = \sum_{i=1}^M \log B_{10,i}^c / M + \eta, \quad \log B_{10,i}^c = (\log B_{10,i} \vee L) \wedge U.$$

Logarithms of Bayes factors are supported on \mathbb{R} and, as we argued in Section 3, they have a natural symmetry around zero, so it is sensible to add a zero mean perturbation term on that scale. After exponentiating, we obtain a geometric mean of censored Bayes factors with a multiplicative perturbation:

$$\tilde{B}_{10} = \exp(\eta) \left(\prod_{i=1}^M B_{10,i}^c \right)^{1/M}.$$

Since the distribution of η is symmetric, $(\tilde{B}_{10})^{-1}$ is equal in distribution to \tilde{B}_{01} , and it is exactly equal to \tilde{B}_{01} when $\eta = 0$ (i.e., when there are no privacy constraints). Geometric means of Bayes factors have appeared in the objective Bayesian literature in geometric intrinsic Bayes factors (Berger and Pericchi, 1996).

The statistic \tilde{B}_{10} is based on censored Bayes factors that are supported on $[L, U]$. However, the support of \tilde{B}_{10} is not $[L, U]$ after introducing the perturbation term η . This issue can be solved by censoring \tilde{B}_{10} , defining

$$B_{10}^* = (\tilde{B}_{10} \vee L^*) \wedge U^*, \tag{1}$$

where $L^* = \exp(L)$ and $U^* = \exp(U)$. With B_{10}^* , we can define the differentially private posterior probability of H_1 given D as $P^*(H_1 | D) = [1 - P(H_0)]B_{10}^*/\{P(H_0) + [1 - P(H_0)]B_{10}^*\}$.

Following the same reasoning, we can define a differentially private information criterion

$$\begin{aligned} I_{10}^* &= (\tilde{I}_{10} \vee L^*) \wedge U^* \\ \tilde{I}_{10} &= \exp(\eta) \left(\prod_{i=1}^M I_{10,i}^c \right)^{1/M} \\ I_{10,i}^c &= (b_i^{-\rho/2} \Lambda_{10,i} \vee L) \wedge U. \end{aligned} \tag{2}$$

4.1 Nested Linear Regression Models

Consider the normal linear model

$$Y = X_0\beta_0 + X\beta + \sigma W, \quad W \sim N_n(0_n, I_n),$$

where $X_0 \in \mathbb{R}^{n \times p_0}$ and $X \in \mathbb{R}^{n \times p}$ are full-rank and $n > p + p_0$. In this section, we present differentially private methods for testing $H_0 : \beta = 0_p$ against $H_1 : \beta \neq 0_p$. The set of predictors in X_0 is common to H_0 and H_1 and it can be, for instance, an intercept 1_n .

For the Bayesian methods, we define our priors after reparameterizing the model. We rewrite the model as $Y = X_0\psi + V\beta + \sigma W$ for $V = (I_n - P_{X_0})X$ and $\psi = \beta_0 + (X_0'X_0)^{-1}X_0'X\beta$. In this parameterization, the common predictors X_0 are orthogonal to V , which is specific to H_1 . If X_0 is an intercept 1_n , the reparameterization simply centers the predictors in X .

Our prior specification is

$$\pi(\psi, \sigma^2) \propto 1/\sigma^2, \quad \pi(\beta | \sigma^2, H_1) = \int_0^\infty N_p(\beta | 0_p, g\sigma^2(V'V)^{-1}) \pi(g) \nu(dg).$$

where ν is an appropriate dominating measure. We allow the prior measure on g to depend on n and p , but not on Y .

The prior on (ψ, σ^2) is the right-Haar prior for this problem, which is improper, whereas $\beta | \sigma^2, H_1$ is a mixture of g -priors. This class of priors has strong theoretical support (see Liang et al. (2008) and Bayarri et al. (2012) for details). For example, it leads to Bayes factors and posterior probabilities of hypotheses that are invariant with respect to invertible linear transformations of the design matrix V , such as changes of units. This property would not be satisfied if we had chosen a diagonal covariance matrix for $\beta | g, \sigma^2, H_1$.

The prior distribution for β_0 and σ^2 is improper, so the marginal distributions can only be defined up to arbitrary multiplicative constants. We use the same constants for both H_0 and H_1 , which is justified by the principle of invariance (Berger et al., 1998) and predictive matching arguments (Bayarri et al., 2012).

When the data are not confidential, we can report the Bayes factor (Liang et al., 2008):

$$B_{10} = \int_0^\infty (g+1)^{(n-p-p_0)/2} [1 + g(1 - R^2)]^{-(n-p_0)/2} \pi(g) \nu(dg), \tag{3}$$

where $R^2 = Y'P_V Y / Y'(I_n - P_{X_0})Y$ or, from a non-Bayesian perspective, we can report the information criterion

$$I_{10} = n^{-\rho/2} \Lambda_{10} = n^{-\rho/2} (1 - R^2)^{-n/2}. \quad (4)$$

If there are privacy constraints, B_{10} or I_{10} cannot be released directly. We propose applying the subsample and aggregate technique to release differentially private versions of B_{10} and I_{10} . That is, we split up the data set into M disjoint subgroups of sample sizes b_1, b_2, \dots, b_M ($\sum_{i=1}^M b_i = n$), define priors $\pi_i(g_i)$ and penalties ρ_i for $i \in \{1, 2, \dots, M\}$, and release B_{10}^* and I_{10}^* as defined in Equations (1) and (2), respectively.

Proposition 1 below states that B_{10}^* and I_{10}^* are consistent under some regularity conditions. Consistency does not follow from Liang et al. (2008) for various reasons. First, there is a growing number of subgroups and the sample sizes within the subgroups grow to infinity. Another difference is that the Bayes factors are censored and there is a perturbation term η . Our result does not follow directly from Smith (2011), either. Smith (2011) studies the asymptotic behavior of averages $\sum_{i=1}^M T_i / M$ released with the subsample and aggregate method. If T_i^* is the ‘‘true value’’ of T_i , Lemma 6 in Smith (2011) assumes that $\sqrt{b_i}(T_i - T_i^*)$ is asymptotically normal, $\mathbb{E}(T_i) - T_i^* \in O(1/b_i)$, and $\mathbb{E}\{[\sqrt{b_i}(T_i - T_i^*)]^3\} \in O(1)$. Furthermore, the results in Smith (2011) are for independent and identically distributed data. In our case, it is unclear whether the asymptotic conditions are satisfied (if they are, it would require proof), and we would need additional assumptions on, at least, the design matrices, the priors $\pi_i(g_i)$, and the penalties ρ_i . Instead of verifying the conditions in Smith (2011), our proof uses a union bound and tail inequalities for R^2 .

Proposition 1 *Under the regularity conditions listed below, B_{10}^* and I_{10}^* are consistent under H_0 , $B_{10}^* \rightarrow_P 0$ and $I_{10}^* \rightarrow_P 0$; under H_1 , $(B_{10}^*)^{-1} \rightarrow_P 0$ and $(I_{10}^*)^{-1} \rightarrow_P 0$.*

1. **Well-specified model:** *The raw confidential data are generated from the normal linear model described in this section.*
2. **Growth of M and b_i :** *$\lim_{n \rightarrow \infty} M = \infty$ and $\lim_{n \rightarrow \infty} \inf_{i \in 1:M} b_i = \infty$. If $p \geq 2$, $\lim_{n \rightarrow \infty} \sup_{i \in 1:M} M/b_i = 0$ and, if $p = 1$, $\lim_{n \rightarrow \infty} \sup_{i \in 1:M} M \sqrt{\log b_i / b_i} = 0$.*
3. **Censoring limits:** *$\lim_{n \rightarrow \infty} L = -\infty$, $\lim_{n \rightarrow \infty} U = \infty$.*
4. **Privacy parameters:** *The privacy parameters ε and δ are such that $\lim_{n \rightarrow \infty} (U - L) / (M\varepsilon) = \lim_{n \rightarrow \infty} a(U - L)^2 / (M\varepsilon) = 0$.*
5. **Design matrices:** *Under H_1 , $\lim_{n \rightarrow \infty} \inf_{i \in 1:M} \beta_T' X_i' X_i \beta_T / (\sigma_T^2 b_i) > 0$, where β_T and σ_T^2 are the fixed true values of β and σ^2 , respectively.*
6. **Priors on g_i and penalties ρ_i :**

$$\lim_{n \rightarrow \infty} \sup_{i \in 1:M} \int_0^\infty b_i^{p/2} (g_i + 1)^{-p/2} \pi_i(g_i) \nu(dg_i) < \infty, \quad \rho_i \leq \max(p, \log b_i)$$

$$\lim_{n \rightarrow \infty} \inf_{i \in 1:M} \int_{b_i}^\infty b_i^{p/2} (g_i + 1)^{-p/2} \pi_i(g_i) \nu(dg_i) > 0, \quad \rho_i \geq p.$$

Assumption 1 states that the model for the confidential data is well-specified. We explicitly state this assumption because it is not necessarily satisfied in Section 5. Assumption 2 implies that both the number of subgroups M and the sample sizes within the subgroups b_i grow to infinity, and it imposes conditions on the growth of M relative to the growth of b_i . Assumption 3 assumes that the censoring limits increase with the sample size. Assumption 4 is related to the privacy parameters ε and δ . It holds for fixed ε and δ , but it also includes cases where ε and δ depend on n . However, the dependency has to be such that, ultimately, η converges to zero in probability: the privacy parameters cannot decrease too fast, relative to the censoring limits. Assumption 5 is a regularity condition on design matrices that is commonly made in the literature (see Equation (22) in Liang et al. (2008) or Equation (17) in Bayarri et al. (2012)). Finally, when combined with the previous assumptions, Assumption 6 is sufficient for establishing consistency. It imposes conditions on the prior on g_i or, from a non-Bayesian perspective, the model complexity penalties ρ_i . Priors such as Zellner's g -prior with $g_i = b_i$, the robust prior proposed in Bayarri et al. (2012), and Zellner-Siow (Zellner and Siow, 1980) satisfy the conditions. The differentially private version of BIC satisfies the conditions on ρ_i , but AIC and the likelihood ratio statistic do not. When there are no privacy constraints, both AIC and the likelihood ratio statistic fail to be consistent under H_0 (Claeskens et al., 2008).

Data-splitting is a commonly used strategy for analyzing big data (see, for instance, Chen et al. (2021) for a review on the topic). Proposition 1 establishes consistency of mixtures of g -priors and information criteria if the data are split into independent subgroups, with or without privacy constraints (the latter corresponds to the case where $\varepsilon \rightarrow \infty$). To the best of our knowledge, there are no equivalent results in the literature.

The Bayes factors and information criteria we work with are misspecified, since the observed values are perturbed averages of censored statistics. Therefore, Proposition 1 can be of interest to those studying the properties of Bayesian procedures under model misspecification models, in the spirit of Rossell and Rubio (2019).

4.2 Quantifying the Uncertainty Introduced by the Mechanism

Given a differentially private Bayes factor, likelihood ratio or information criterion, we can find a confidence interval that quantifies the uncertainty introduced by the privacy-ensuring mechanism.

Let $\tilde{T} = \sum_{i=1}^n T_i^c/M + \eta$ be a differentially private statistic, and let $\eta_{1-\alpha/2}$ be the $(1-\alpha/2)$ -th quantile of the perturbation term η . Then, $\tilde{\mathcal{C}}_\alpha = \tilde{T} \pm \eta_{1-\alpha/2} = [\tilde{l}_\alpha, \tilde{u}_\alpha]$ is a $1-\alpha$ confidence interval for $\sum_{i=1}^n T_i^c/M$. The interval can be shortened because $\sum_{i=1}^n T_i^c/M \in [L, U]$, so $\mathcal{C}_{1-\alpha} = [(\tilde{l}_\alpha \vee L) \wedge U, (\tilde{u}_\alpha \vee L) \wedge U]$ is also a $1-\alpha$ confidence interval for $\sum_{i=1}^n T_i^c/M$.

Confidence intervals for transformations $f(\tilde{T})$ can be found by applying f to the endpoints of $\mathcal{C}_{1-\alpha}$. For Bayes factors of H_1 to H_0 , the transformation is $f(x) = \exp(x)$; for posterior probabilities of H_1 , it is $f(x) = [1 - P(H_0)] \exp(x) / \{P(H_0) + [1 - P(H_0)] \exp(x)\}$.

4.3 Effects of Censoring and Data-splitting

The subsample and aggregate technique requires the specification of censoring limits $L < U$ and a number of subgroups M . These choices affect the performance of the methods, which is well-acknowledged in the literature (see Evans et al. (2020), Covington et al. (2021), and

Ferrando et al. (2022)). These articles focus on bias correction to obtain unbiased point estimates and valid confidence intervals. Our context is different: since we are concerned with hypothesis testing, the bias correction takes the form of correcting critical values to ensure appropriate type I errors.

Differentially private likelihood ratios Λ_{10}^* are not distributed as their non-private counterparts. Therefore, the critical values for rejecting H_0 with Λ_{10}^* must be bias-corrected to guarantee adequate type I error rates.

We can find corrected critical values for Λ_{10}^* via simulation. In the linear regression model, Λ_{10}^* depends on the data only through R_i^2 , which under H_0 are distributed as beta with shape parameters $p/2$ and $(b_i - p - p_0)/2$. We can repeatedly simulate R_i^2 under H_0 , compute Λ_{10}^* , and estimate the correct critical value with empirical quantiles. In contexts other than linear regression models, the adjustments can be done approximately since, under mild conditions, Wilks' theorem establishes that the transformed likelihood ratio statistic $2 \log \Lambda_{10}$ is asymptotically distributed as chi-squared under H_0 .

In the remainder of this section, we quantify the effects of censoring and splitting the data in terms of bias-variance trade-offs. In Appendix B, we present the results of a simulation study that aims to quantify these trade-offs. Our conclusion is consistent with what we report here.

Given the symmetric nature of logarithms of Bayes factors (explained in Section 3), we specify symmetric censoring limits $L = -C$ and $U = C$ for $C > 0$. Information criteria with penalties as specified in Proposition 1 (a prime example being BIC) are consistent under H_0 and H_1 , so it is sensible to specify symmetric censoring limits for their logarithms. On the other hand, log-likelihood ratios are non-negative for nested models (which is the case we consider in Section 4), so $L = 0$ is a natural lower bound for them. In what follows, we assume that the confidential statistics are censored with the limits we have described in this paragraph.

Let $\tilde{T} = \sum_{i=1}^n T_i^c/M + \eta$ be an ε -differentially private statistic released by the subsample and aggregate technique (similar arguments can be made for (ε, δ) -differentially private statistics with $\delta > 0$). Then,

$$\text{Var}(\tilde{T}) = \left[\sum_{i=1}^n \text{Var}(T_i^c) + 2\Delta^2/\varepsilon^2 \right] / M^2.$$

The variance of \tilde{T} is clearly decreasing in the number of subgroups M , and it is increasing in the sensitivity $\Delta = U - L$ because $\text{Var}(T_i^c)$ is increasing in Δ .

The bias of the differentially private \tilde{T} with respect to \mathcal{T} , which is the statistic we would obtain without censoring or splitting the data, can be decomposed into the bias induced by censoring plus the bias induced by splitting the data:

$$\underbrace{\mathbb{E}(\tilde{T} - \mathcal{T})}_{\text{total bias}} = \underbrace{\mathbb{E} \left\{ \sum_{i=1}^M T_i^c/M - \sum_{i=1}^M T_i/M \right\}}_{\text{censoring bias}} + \underbrace{\mathbb{E} \left\{ \sum_{i=1}^M T_i/M - \mathcal{T} \right\}}_{\text{data-splitting bias}}.$$

The censoring bias is decreasing in the sensitivity Δ and the number of subgroups M . In the case of logarithms of Bayes factors and information criteria, where the censoring is

symmetric about zero, the censoring bias is

$$\frac{1}{M} \sum_{i=1}^M \{P(T_i \leq -C) [-C - \mathbb{E}(T_i | T_i \leq -C)] + P(T_i \geq C) [C - \mathbb{E}(T_i | T_i \geq C)]\}.$$

The censoring bias would be zero if the T_i were symmetric about zero. However, the distributions of logarithms of Bayes factors and information criteria tend to be asymmetric because they are consistent: under H_0 , logarithms of Bayes factors and information criteria tend to be negative, favoring the null; under H_1 , they tend to be positive, favoring the alternative. If the distribution of the T_i is not symmetric about zero, the censoring bias decreases in $\Delta = 2C$ and M . In the case of log-likelihood ratio statistics, the censoring bias is

$$\frac{1}{M} \sum_{i=1}^M P(T_i \geq U) [U - \mathbb{E}(T_i | T_i \geq U)],$$

which is decreasing in the upper censoring limit U and the number of subgroups M .

The data-splitting bias is harder to characterize generally but, through examples and simulations, we observe that it is typically increasing in the number of subsets. In the case that the subgroups are balanced ($b_i = b$ for all $i \in \{1, 2, \dots, M\}$) the data-splitting bias is

$$\mathbb{E} \left\{ \sum_{i=1}^M T_i / M - \mathcal{T} \right\} = \mathbb{E}(T_i - \mathcal{T}),$$

where T_i is the same statistic as \mathcal{T} but computed with fewer observations. If T_i and \mathcal{T} were functions like the sample mean, this term would be zero, but for our statistics, this bias will be non-zero (in general). Intuitively, if the number of subgroups is large and T_i is based on a smaller sample, the bias will tend to increase. This phenomenon occurs empirically in Section 4.4 and Appendix B, where the bias increases as the number of subgroups increase. For a more analytical argument, the example below considers the case where T_i and \mathcal{T} are log-likelihood ratios. In this case, we are able to find the data-splitting bias under the null hypothesis. This example is relevant because our information criteria and Bayes factors depend on the data only through log-likelihood ratios.

Example 1 Assume the subgroups are balanced ($b_i = b$ for all $i \in \{1, 2, \dots, M\}$), so we have

$$T_i = \log \Lambda_{10,i} = -b \log(1 - R_i^2)/2, \quad \mathcal{T} = \log \Lambda_{10} = -n \log(1 - R^2)/2.$$

Under H_0 , $1 - R_i^2 \sim \text{Beta}((b-p-p_0)/2, p/2)$ and $1 - R^2 \sim \text{Beta}((n-p-p_0)/2, p/2)$. Using properties of the beta distribution, the bias under H_0 is

$$\mathbb{E}(T_i - \mathcal{T}) = \frac{b}{2} \left[\psi \left(\frac{b-p_0}{2} \right) - \psi \left(\frac{b-p-p_0}{2} \right) \right] - \frac{n}{2} \left[\psi \left(\frac{n-p_0}{2} \right) - \psi \left(\frac{n-p-p_0}{2} \right) \right],$$

where ψ is the digamma function (Abramowitz et al., 1964). Using properties of the digamma function, it is straightforward to see that the bias is positive and increasing as the number of subgroups increases (i.e., as b decreases). The bias is zero, as it should be, when there is no data-splitting and $b = n$.

In conclusion, we observe two bias-variance tradeoffs:

1. **Censoring:** Increasing the sensitivity $\Delta = U - L$ decreases the bias but increases the variance of the differentially private statistic.
2. **Data-splitting:** Increasing the number of subgroups M decreases the variance of the differentially private statistic. The relationship between M and the total bias is slightly more complicated: increasing the number of subgroups decreases the censoring bias, but it increases the data-splitting bias. In our empirical evaluations, we systematically observe that the total bias of the statistic is increasing in M .

4.4 Application: High School and Beyond Survey

We analyze a random sample of 200 students from the High School and Beyond survey, which was conducted by the National Center of Education Statistics. We obtained the data from Diez et al. (2012). In R, they are available as `data(hsb2)` in `library(openintro)`.

We consider two hypothesis tests. Both have `math` scores as their outcome variable Y . In the first one, we test if the variable `gender` (which, in this data set, can take on the values `male` or `female`) is predictive of `math` scores. Under the null hypothesis (H_{01}), the model contains only an intercept, whereas under the alternative (H_{11}), the model includes an intercept and `gender` as a predictor. Then, we test if `read` scores are predictive of `math` scores when `science` scores are already in the model. In the null hypothesis (H_{02}), the model includes an intercept and `science` scores. Under the alternative hypothesis (H_{12}), the model contains an intercept, `science` scores, and `read` scores.

In Figure 1, we show the distribution of the differentially private posterior probabilities $P^*(H_{11} | D)$ and $P^*(H_{12} | D)$ for different combinations of ε and M , simulating random data splits and perturbation terms $\eta \sim L_1(0, \Delta/(M\varepsilon))$. Our prior probabilities on the hypotheses are $P(H_{01}) = P(H_{02}) = 0.5$, and the prior on the regression coefficients is Zellner’s g -prior with g_i equal to the sample sizes of the subgroups b_i . We censor logarithms of Bayes factors at $L = \log((1-0.99)/0.99)$ and $U = \log((1-0.01)/0.01)$, which is equivalent to censoring posterior probabilities of hypotheses at 0.01 and 0.99, respectively. Each error bar has been computed with 10^4 simulations.

As expected, increasing ε and M reduces the variability in $P^*(H_{11} | D)$ and $P^*(H_{12} | D)$. Additionally, M induces a conservative bias, in the sense that the posterior probabilities are shrunk towards 0.5. For example, if $M = 10$ and ε is large, $P^*(H_{11} | D) \approx 0.25$ and $P^*(H_{12} | D) \approx 0.70$, which are more conservative than the confidential answers $P(H_{11} | D) \approx 0.07$ and $P(H_{12} | D) \approx 0.99$.

In Figure 2, we show the distributions of $2 \log \Lambda_{10,1}^*$ and $2 \log \Lambda_{10,2}^*$ for $L = 0$, $U = 7$, $\varepsilon = 1$, and different values of $0 < \delta < 1$. We treat the data as fixed and quantify the uncertainty introduced by splitting the data randomly and the perturbation term η , which in this case is normally distributed. The non-private test of H_{01} against H_{11} does not reject H_{01} at the 0.01 or 0.05 significance levels, whereas the non-private test of H_{02} against H_{12} rejects H_{02} at the 0.01 and 0.05 significance levels. In fact, $2 \log \Lambda_{10,2}$ is greater than U .

As we argued in Section 4.3, the critical values for rejecting the null hypothesis have to be corrected to take the censoring and data-splitting into account. When δ is small and $M \in \{2, 5\}$, the corrected critical values (gray + and Δ) are larger than the uncorrected

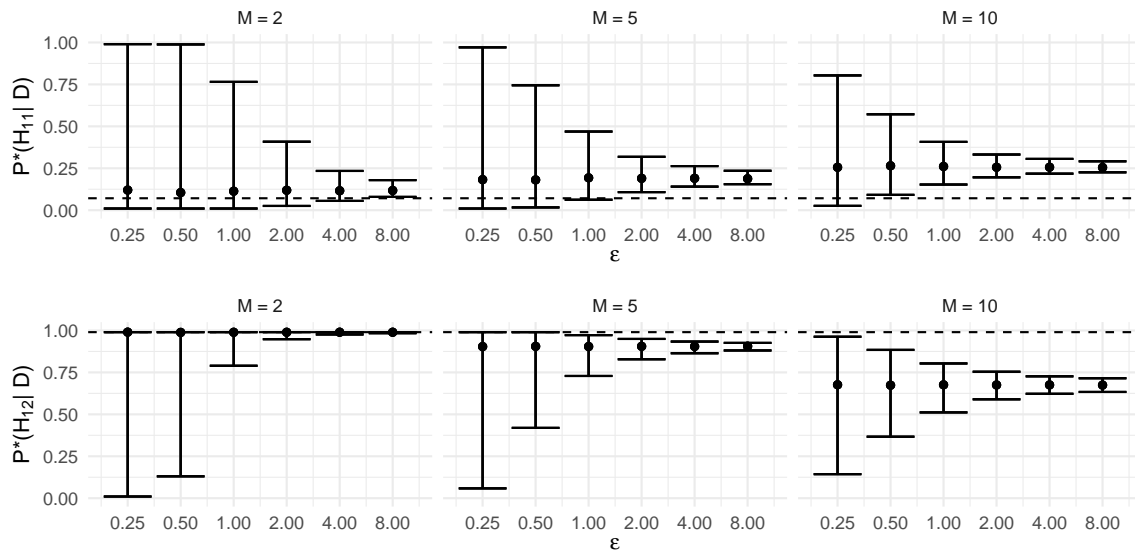


Figure 1: Distribution of $P^*(H_{11} | D)$ and $P^*(H_{12} | D)$ as a function of ε and M . The lower endpoint of the error bars is the first quartile of the distribution, the midpoint is the median, and the upper endpoint is the third quartile. The dashed lines are the true confidential posterior probabilities.

ones (black + and Δ), so if we did not correct the critical values, we would have inflated type I errors. For larger δ and $M = 10$, the corrected critical points are lower than the uncorrected ones, showing that correcting critical values can increase power. The distribution of $2 \log \Lambda_{10,1}^*$ is generally below the corrected critical values, which coincides with the confidential decision of not rejecting the null. In the case of $2 \log \Lambda_{10,2}^*$, corrected tests would reject the null most of the time, in agreement with the confidential answer, especially when $M > 2$ and $\delta \geq 0.25$.

In Section B of the Appendix, we repeat the simulation study with different censoring limits. Decreasing the sensitivity $U - L$ decreases the variance, but it can introduce bias and lead to tests that are not powerful. In the next subsection, we give simple guidelines to set the censoring limits.

4.5 Guidelines

In Section 4.4, we observed a bias-variance tradeoff when choosing the number of subgroups M : as M increases, the variability introduced by η decreases at the expense of introducing some bias.

Before obtaining a private statistic, users can simulate mean squared errors to decide the value of M that is most appropriate in their application. The simulations can be run for different values of R^2 . Another option is using the sampling distribution of R^2 , which is known, and average our uncertainty over it.

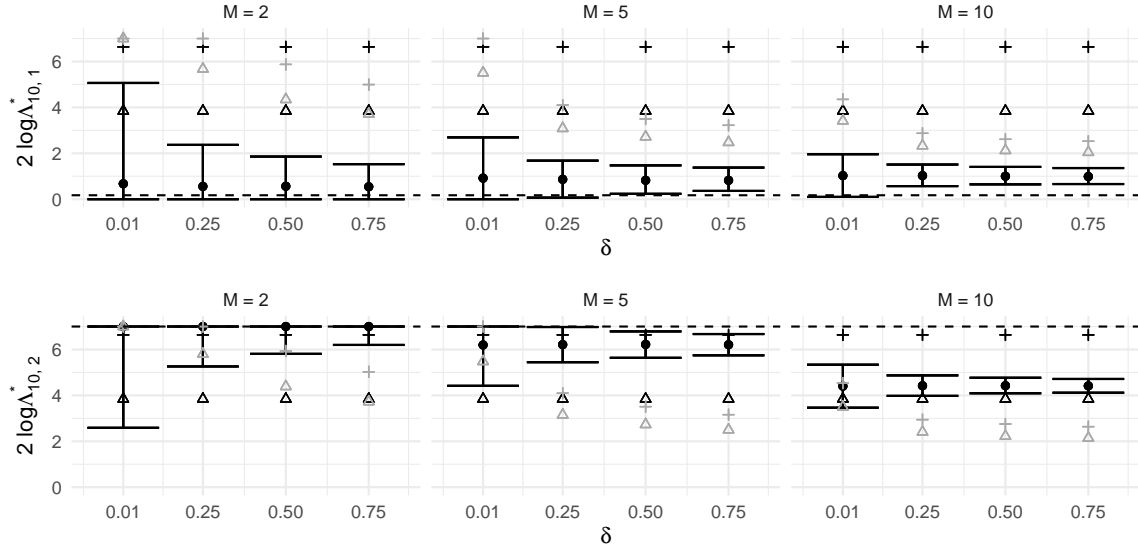


Figure 2: Distribution of $2 \log \Lambda_{10,1}^*$ and $2 \log \Lambda_{10,2}^*$ as a function of δ and M . The gray $+$ and Δ are corrected critical values at the 0.01 and 0.05 significance levels, respectively. The black $+$ and Δ are uncorrected critical values. Dashed horizontal lines are $2 \log \Lambda_{10,1}$ (first row) and $U = 7$ (second row).

The bounds L and U can be values of the confidential statistic that users consider small or large enough. Stringent censoring decreases the variance of the output, but it can introduce substantial bias if the true posterior probability or likelihood ratio lies outside the censoring limits. In Section 4.4, we chose L and U that were equivalent to censoring posterior probabilities at 0.01 and 0.99, respectively, since those values indicate strong support against or in favor of hypotheses. Practitioners can also use Jeffreys' scale of evidence (Table 1) to censor Bayes factors instead of posterior probabilities. In the case of likelihood ratios, we recommend setting the lower censoring bound L to 0, which is natural in this case, and let the upper censoring value be greater than the critical value for rejection of the non-private test to avoid loss of power. Setting U to approximately twice the critical value performed well in our applications.

5. Model Averaging and Selection

In this section, we work with the normal linear model as defined in Section 4, but instead of comparing two nested models, we consider all the 2^p models that can be constructed by taking subsets of columns of X . We express our model uncertainty through a binary vector $\gamma = (\gamma_1, \gamma_2, \dots, \gamma_p)' \in \{0, 1\}^p$ such that $\gamma_j = 0$ if and only if $\beta_j = 0$. We use the notation $|\gamma| = \sum_{j=1}^p \gamma_j$ for the number of active coefficients in a model.

Given a model indexed by γ , we write the linear model as

$$Y = X_0 \beta_0 + V_\gamma \beta_\gamma + \sigma W, \quad W \sim N_n(0_n, I_n),$$

where $\beta_\gamma \in \mathbb{R}^{|\gamma| \times 1}$ is a vector including the β_j such that $\gamma_j = 1$ and $V_\gamma \in \mathbb{R}^{n \times |\gamma|}$ is a matrix with the active variables in γ . Just as we did in Section 4, we parameterize the model so that X_0 and V_γ are orthogonal. For convenience, we denote the null model where none of the variables are active as $\gamma = 0$. The matrix X_0 represents a set of predictors we are sure to include in our model. Usually, X_0 contains an intercept 1_n , but it can be empty as well. If X_0 is empty, the methods in this section would be defined in analogous manner after replacing $I_n - P_{X_0}$ by I_n .

From a Bayesian perspective, our prior specification on the regression coefficients β_γ is the same we had in Section 4: we put mixtures of g -priors on $\beta_\gamma \mid \sigma^2, \gamma$ and the right-Haar prior $\pi(\psi, \sigma^2) \propto 1/\sigma^2$ on the common parameters. If γ is not the null model, the expression for the non-private Bayes factor of model γ to the null model, denoted $B_{\gamma 0}$, is identical to the one in Equation (3) after substituting p by $|\gamma|$ and R^2 by $R_\gamma^2 = Y'P_{V_\gamma}Y/Y'(I - P_{X_0})Y$. If γ is the null model, we have $B_{\gamma 0} = 1$.

Given a prior distribution $\pi(\gamma)$ on γ , the posterior probability of the model identified by γ given the data D is defined as

$$P(\gamma \mid D) = P(\gamma) B_{\gamma 0} / \sum_{\tilde{\gamma} \in \{0,1\}^p} P(\tilde{\gamma}) B_{\tilde{\gamma} 0},$$

which depends on the data only through the Bayes factors $B_{\gamma 0}$. We assume that the prior $P(\gamma)$ can depend on p , but not on n or the design matrix. Most common choices for $P(\gamma)$, like a uniform prior $P(\gamma) = 2^{-p}$ or the hierarchical uniform prior recommended by Scott and Berger (2010) satisfy the condition.

From a non-Bayesian perspective, the information criteria $I_{\gamma 0}^*$ are as in Equation 4 after substituting R^2 by R_γ^2 and ρ by $\rho_{|\gamma|}$, where $\rho_{|\gamma|}$ is an increasing function in $|\gamma|$ such as $\rho_{|\gamma|} = |\gamma|$.

We could use the method in Section 4.1 to release differentially private versions of all the $B_{\gamma 0}$ or $I_{\gamma 0}$. However, if we released those 2^p statistics, the variance of the private statistics would increase exponentially in p . Instead, we propose working with a perturbed version of a sufficient statistic whose dimension increases quadratically in p .

Let $Z = (I - P_{X_0})Y \in \mathbb{R}^{n \times 1}$ be the null-centered outcome variable and $V \in \mathbb{R}^{n \times p}$ be the design matrix of the full model that includes all p predictors, which we collect in a data matrix $D = [V; Z]$. Assuming $Z'Z > 0$ almost surely, we define

$$G = D'D = \begin{bmatrix} V'V & V'Z \\ Z'V & Z'Z \end{bmatrix} = \begin{bmatrix} V'V & V'Y \\ Y'V & Y'(I - P_{X_0})Y \end{bmatrix}.$$

The Gram matrix G is a sufficient statistic for the normal linear model (see, for example, Seber and Lee (2012)). As a consequence, all of the R_γ^2 , $B_{\gamma 0}$, and $I_{\gamma 0}$ can be constructed by taking appropriate subsets of G .

We propose releasing a noisy version of the sufficient statistic G , a technique known in the differential privacy literature as *sufficient-statistic perturbation* (see, for example, McSherry and Mironov (2009); Vu and Slavkovic (2009) and Bernstein and Sheldon (2019)).

We construct a differentially private version of G by adding a random perturbation term, defining $G^* = G + E$, where E is a random perturbation matrix that ensures differential privacy. To ensure ε -differential privacy, we use the Laplace mechanism. For (ε, δ) -differential

privacy with $0 < \varepsilon < 1$ and $0 < \delta \leq 1$, we use Algorithm 2 in Sheffet (2019), which we refer to as the Wishart mechanism.

To establish the parameters of the distribution of E , we assume that there are lower and upper bounds for the data: that is, there are l and u such that each entry d_{ij} in D is within the interval $[l, u]$. Since the response and predictors are centered, $l < 0$ and $u > 0$. The entries of $G = D'D$ are of the form $\sum_{i=1}^n d_{ij}d_{ij'}$. Here, d_{1j} is the first row and j th column of D . If we replace this entry by another one, say $\tilde{d}_{1j'}$, then the maximum absolute difference in the entries of G is $|d_{1j}d_{1j'} - \tilde{d}_{1j}\tilde{d}_{1j'}|$. When computing the sensitivity Δ_1 , we only consider the entries where $j \geq j'$ because G is symmetric. Therefore, the sensitivity Δ_1 can be upper-bounded as follows:

$$\begin{aligned} \Delta_1 &= \sup_{D \sim \tilde{D}} \|D'D - \tilde{D}'\tilde{D}\|_1 \\ &= \sup_{D \sim \tilde{D}} \sum_{j \geq j'} |d_{1j}d_{1j'} - \tilde{d}_{1j}\tilde{d}_{1j'}| \\ &= \sum_{j \geq j'} \sup_{D \sim \tilde{D}} |d_{1j}d_{1j'} - \tilde{d}_{1j}\tilde{d}_{1j'}| \\ &\leq (p+1)(p+2) \sup_{d_{11}, d_{11'} \in (u, l)} |d_{11}d_{11'}| \\ &\leq (p+1)(p+2)(l^2 \vee u^2). \end{aligned}$$

For the Laplace mechanism, we define the perturbation term E as a symmetric random matrix of the form

$$E = \begin{bmatrix} e_{11} & e_{12} & \dots & e_{1,p+1} \\ e_{12} & e_{22} & \dots & e_{2,p+1} \\ \vdots & \vdots & \ddots & \vdots \\ e_{1,p+1} & e_{2,p+1} & \dots & e_{p+1,p+1} \end{bmatrix},$$

where $e_{jj'} \stackrel{\text{iid}}{\sim} L_1(0, \Delta_1/\varepsilon)$, for $j \geq j'$, and $e_{j'j} = e_{jj'}$. For the Wishart mechanism (i.e., Algorithm 2 in Sheffet, 2019), we need a uniform bound on the Euclidean norm of the rows of D . Given the assumption $l < d_{ij} < u$, we can use $\sqrt{(p+1)(l^2 \vee u^2)}$ as a uniform bound. In this case, given the bound, we define the random perturbation as $E = M - \mathbb{E}(M)$, where $M \sim W_{p+1}(k, (p+1)(l^2 \vee u^2)I_{p+1})$ and $k = \lfloor p+1 + 28 \log(4/\delta)/\varepsilon^2 \rfloor$ with $0 < \varepsilon < 1$ and $0 < \delta \leq 1$.

For both mechanisms, the variances of the entries in E can be high, especially for small values of ε or δ . This, in turn, can lead to outputs that overestimate the number of active predictors. To avoid this issue, we propose post-processing G^* in two ways: hard-thresholding off-diagonal elements as in Bickel et al. (2008) and adding a constant r to the diagonal elements.

We propose thresholding the off-diagonal entries of G^* at e_λ , the λ -th percentile of E_{ij} . More formally, we define G^{**} to be the matrix with typical element $G_{ij}^* \mathbb{1}(i = j \text{ or } |G_{ij}^*| \geq e_\lambda)$. This choice can be justified as follows: if an off-diagonal entry of $G_{ij}^* = G_{ij} + E_{ij}$ is not an extreme value in the distribution of E_{ij} , it is likely that G_{ij}^* is essentially E_{ij} and G_{ij} is nearly zero.

Adding a constant r to the diagonal elements of G^* can be seen as a ridge-type of regularization that can reduce the variability of the output. It can also be used to guarantee

that the output is positive-definite because, after adding the Laplace perturbation terms, G^* may not be positive-definite. In fact, the hard-thresholded matrix G^{**} need not be positive-definite even if G^* is positive-definite (Bickel et al., 2008). Given a symmetric matrix A , which here can be either G^* or G^{**} , the matrix $A_r = A + rI_{p+1}$ is positive-definite as long as $r > -\text{eig}_{\min}(A)$, where $\text{eig}_{\min}(\cdot)$ is a function returning the minimum eigenvalue of a matrix.

In essence, we propose releasing a differentially private Gram matrix G^* , which we then transform to obtain a better estimate of G . Then, we use the formulas we would use if we had access to G after plugging in its differentially private estimate.

In our applications (in Section 5.2), we consider methods that hard-threshold G^* and methods that do not. That is, we compare methods that are based on $G_r^{**} = G^{**} + rI_{p+1}$ to methods based on $G_r^* = G^* + rI_{p+1}$, where G^* is not hard-thresholded. In our experience, hard-thresholding is helpful when the ground truth is sparse, but it can be detrimental when most predictors are active. We justify this argument in Section 5.2.

Proposition 2 establishes model-selection consistency under some assumptions. More precisely, we show that the differentially private Bayes factor of any model γ to the true model, which can be expressed as $B_{\gamma T}^* = B_{\gamma 0}^*/B_{T 0}^*$, converges to zero in probability for any $\gamma \neq T$. The differentially private information criteria $I_{\gamma T}^*$ are also consistent under the assumptions listed below. We proved the result by characterizing the asymptotic behavior of G_r^* , G_r^{**} , and $R_{\gamma}^{2,*}$, and then bounding the Bayes factors and information criteria above and below. Just as we had in Section 4, the proof covers ε and (ε, δ) differentially private methods. The proof does not follow from Liang et al. (2008) for several reasons, one of them being that we are not assuming that the response given the covariates is normal, since we assume that the data are bounded.

Proposition 2 *Let $T \in \{0, 1\}^p$ be a vector indexing the truly active predictors. As $n \rightarrow \infty$, and under the regularity conditions listed below, $B_{\gamma T}^* \rightarrow_P 0$ and $I_{\gamma T}^* \rightarrow_P 0$ for any $\gamma \neq T$.*

1. **Boundedness:** *The data D are within the interval $[l, u]$ for finite l and u .*
2. **Regression mean and variance:** *$\mathbb{E}(Y \mid X_0, V) = X_0\psi + V_T\beta_T$ and $\text{Var}(Y \mid X_0, V) = \sigma_T^2 I_n$, where $V_T \in \mathbb{R}^{n \times p_T}$ is a matrix that contains the truly active predictors.*
3. **Privacy parameters:** *The privacy parameters ε and δ are such that the matrix perturbation term $E/n \rightarrow_P 0$.*
4. **Regularization parameters:** *λ is fixed and r is so that $\lim_{n \rightarrow \infty} r/n = 0$.*
5. **Design matrices:** *$\lim_{n \rightarrow \infty} V'V/n = S_1$, where S_1 is symmetric and positive-definite.*
6. **Priors on g and penalties ρ :** *for all $1 \leq |\gamma| \leq p$, the prior $\pi(g)$ satisfies*

$$\lim_{n \rightarrow \infty} \int_0^\infty n^{p/2} (g+1)^{-|\gamma|/2} \pi(g) \nu(dg) < \infty$$

$$\lim_{n \rightarrow \infty} \int_n^\infty n^{p/2} (g+1)^{-|\gamma|/2} \pi(g) \nu(dg) > 0,$$

or from a non-Bayesian perspective, $\rho_{|\gamma|}$ is increasing in $|\gamma|$ and satisfies $|\gamma| \leq \rho_{|\gamma|} \leq |\gamma| \vee \log n$.

The framework here is different to the one in Section 4. We assume that the confidential data are bounded (Assumption 1) and do not assume that the distribution of the outcome given the predictors is normal. In other words, Bayes factors and information criteria are misspecified beyond the addition of the perturbation term E . Nonetheless, Proposition 2 shows that the methods are consistent. Our setup is also distinct to the one adopted in Lei et al. (2018), where it is simultaneously assumed that the response is normal (Assumption 1 in Lei et al. (2018)) and bounded (Assumption 4). Assumption 2 requires that the regression mean be well-specified and that the covariance of the response given the predictors be spherical. Assumption 3 forces the perturbation matrix E to be so that $E/n \rightarrow_P 0$. As we had in Section 4, it suffices to let ε and δ be fixed for it to hold. Assumption 4 imposes conditions on the regularization parameters. The case of a non-thresholded matrix G_r^* is included as $\lambda = 0$. Assumption 5 is similar to the regularity condition on design matrices in Proposition 1. Finally, Assumption 6 is essentially the same as the assumptions on the priors and penalties in 2. Just as we had in Section 4, the differentially private Bayes factors with Zellner’s g -prior with $g = n$, the robust prior, Zellner-Siow, and BIC are all consistent, and so is BIC.

In the common scenario where X_0 is an intercept 1_n , the methods described in this section can be conveniently implemented in R with the `bas.lm` function in `library(BAS)` (Clyde, 2020). Given a set of predictors and an outcome variable, the `bas.lm` function enumerates Bayes factors for small to moderate p and samples from the model space for large p . The function outputs other statistics of interest such as posterior inclusion probabilities and model-averaged estimates. To use `bas.lm` for our problem, we need to generate a synthetic data set \mathcal{D} (containing both centered predictors and outcome) whose sufficient statistic $\mathcal{D}'\mathcal{D}$ is equal to a fixed Gram matrix \mathcal{G} , which can be G_r^* or G_r^{**} . Proposition 3 below shows how to obtain such a synthetic data set $\mathcal{D} = [\mathcal{V}; \mathcal{Z}]$ given a matrix \mathcal{G} .

Proposition 3 *Let $\mathcal{U} \in \mathbb{R}^{n \times (p+1)}$ be a full-rank matrix and define $\mathcal{M} = (I_n - 1_n 1_n'/n)\mathcal{U}$. Given a matrix \mathcal{G} , we can generate a synthetic data set $\mathcal{D} = [\mathcal{V}; \mathcal{Z}]$ with the formula $\mathcal{D} = \mathcal{M}(\mathcal{M}'\mathcal{M})^{-1/2}\mathcal{G}^{1/2}$. The synthetic data set \mathcal{D} satisfies the identities $\mathcal{D}'\mathcal{D} = \mathcal{G}$, $\mathcal{V}'1_n = 0_p$, and $\mathcal{Z}'1_n = 0$.*

Proposition 3 guarantees that the outputs we obtain from running `bas.lm` on the synthetic data \mathcal{D} are identical to what we would find by taking subsets of \mathcal{G} directly. In the proposition, the matrix \mathcal{U} is arbitrary; in practice, its entries can be simulated by sampling independently from the uniform distribution.

5.1 Quantifying the Uncertainty Introduced by the Mechanism

With the non-thresholded methods, the private statistics are of the form $G_r^* = G + E + rI_{p+1}$. Since r and the distribution of E are both known, we can define a confidence set for the non-private G given G_r^* . With such a set, it is possible to find confidence regions for summaries of interest $T(G)$ like least-squares estimates or inclusion probabilities.

To define a $1 - \alpha$ confidence set for G , we first find $\mathcal{E}_{1-\alpha}$ such that $P(E \in \mathcal{E}_{1-\alpha}) = 1 - \alpha$. For a matrix norm $\|\cdot\|$, define $\mathcal{E}_{1-\alpha} = \{E : \|E - \mathbb{E}(E)\| \leq q_{1-\alpha}\}$ where $P[\|E - \mathbb{E}(E)\| \leq$

$q_{1-\alpha}] = 1 - \alpha$. Then, $\tilde{\mathcal{C}}_{1-\alpha} = \{G_r^* - rI_{p+1} - E : E \in \mathcal{E}_{1-\alpha}\}$ is a $1 - \alpha$ confidence set for G . Since G is symmetric and positive-definite, we can intersect $\tilde{\mathcal{C}}_{1-\alpha}$ with the set of symmetric positive-definite matrices \mathcal{S}_{++} to define a $1 - \alpha$ confidence region $\mathcal{C}_{1-\alpha} = \tilde{\mathcal{C}}_{1-\alpha} \cap \mathcal{S}_{++}$ whose volume is at most that of $\tilde{\mathcal{C}}_{1-\alpha}$. The confidence set $\mathcal{C}_{1-\alpha}$ can be transformed into $T(\mathcal{C}_{1-\alpha})$ to produce confidence sets for summaries of interest $T(G)$.

We can approximate $T(\mathcal{C}_{1-\alpha})$ with a rejection sampler. First, simulate $E_1, E_2, \dots, E_{n_{\text{sim}}}$ from the appropriate mechanism and compute $\|E_i - \mathbb{E}(E)\|$ for $i \in \{1, 2, \dots, n_{\text{sim}}\}$. Then, approximate the $((1 - \alpha) \times 100)$ th percentile $q_{1-\alpha}$ with its empirical version $\hat{q}_{1-\alpha} = \inf\{q : \sum_{i=1}^{n_{\text{sim}}} \mathbb{1}(\|E_i - \mathbb{E}(E)\| \leq q) / n_{\text{sim}} \geq 1 - \alpha\}$ and define $\hat{\mathcal{C}}_{1-\alpha} = \{E_i : \|E_i - \mathbb{E}(E)\| \leq \hat{q}_{1-\alpha}, i = 1, \dots, n_{\text{sim}}\}$. After that, we can find $T(\hat{\mathcal{C}}_{1-\alpha})$; that is, compute $T(G_r^* - rI_{p+1} - E_i)$ for $E_i \in \hat{\mathcal{C}}_{1-\alpha}$ and only keep those such that $G_r^* - rI_{p+1} - E_i \in \mathcal{S}_{++}$. This computational strategy to construct $T(\mathcal{C}_{1-\alpha})$ is an application of the parametric bootstrap: the quantity to be inferred is the unknown confidential summary $T(G)$, and the only source of randomness is the noise injected into G to make it differentially private.

In general, the confidence set need not be an interval, but we can summarize the confidence set with a histogram. To do so, we define the bins of the histogram as $\mathcal{B}_k = [t_{k-1}, t_k]$ with $\min\{T(\hat{\mathcal{C}}_{1-\alpha})\} = t_0 < t_1 < \dots < t_{K-1} < t_K = \max\{T(\hat{\mathcal{C}}_{1-\alpha})\}$ and their corresponding relative frequencies $\#\mathcal{B}_k$ using the number of elements of $T(\hat{\mathcal{C}}_{1-\alpha})$ that fall in \mathcal{B}_k , $k = 1, \dots, K$. We denote the histogram summarizing $T(\hat{\mathcal{C}}_{1-\alpha})$ as $\text{Hist}(T, \hat{\mathcal{C}}_{1-\alpha}) = \{(\mathcal{B}_k, \#\mathcal{B}_k)\}_{k=1}^K$. We can also report histograms in a density scale; that is, $\text{Hist}(T, \hat{\mathcal{C}}_{1-\alpha}) = \{(\mathcal{B}_k, d_k)\}_{k=1}^K$, where

$$d_k = \frac{\#\mathcal{B}_k}{|T(\hat{\mathcal{C}}_{1-\alpha})|(t_k - t_{k-1})}$$

and $|T(\hat{\mathcal{C}}_{1-\alpha})|$ denotes the cardinality of $T(\hat{\mathcal{C}}_{1-\alpha})$. While $\text{Hist}(T, \hat{\mathcal{C}}_{1-\alpha})$ displays the distribution of the estimator of $T(G)$ constrained to $T(\hat{\mathcal{C}}_{1-\alpha})$, its support, denoted by $[t_0, t_K]$, corresponds to an $1 - \alpha$ confidence interval. This is because $t_0 = \min T(\hat{\mathcal{C}}_{1-\alpha})$ and $t_K = \max T(\hat{\mathcal{C}}_{1-\alpha})$. Therefore, analysts can directly report $[t_0, t_K]$ as the confidence interval, if desired. Additionally, analysts can use the histogram to assess whether $[t_0, t_K]$ provides a good representation of the set $T(\hat{\mathcal{C}}_{1-\alpha})$. They can also consider if $T(\hat{\mathcal{C}}_{1-\alpha})$ would be better represented by the union of non-adjacent intervals.

5.2 Empirical Evaluations

We evaluate the performance of the methods described in this section in a simulation study and an application. The simulation study is similar to the one in Liang et al. (2008), whereas the real data set is a subset of the March 2000 Current Population Survey that was analyzed in Barrientos et al. (2019).

We implement the methods with the R package `BAS` (Clyde, 2020). In the simulation study, we found Bayes factors with the Zellner-Siow prior (ZS) and information criteria with BIC. The prior distribution on the model space $\pi(\gamma)$ is the hierarchical uniform prior proposed in Scott and Berger (2010). From a non-Bayesian perspective, $\pi(\gamma)$ acts as a function that weighs the information criteria. The results with Zellner-Siow and BIC are almost identical. We report the outputs based on the former here and show the results with the latter in Appendix C.

We compare the results we obtain by hard-thresholding and not thresholding the Gram matrix G^* . In all cases, we add a regularization parameter r to the diagonal entries of G^* . For the Laplace mechanism, we set r to be the 99-th percentile of $\text{eig}_{\min}(E)$, which we find via simulation. For the Wishart mechanism, we use the analytical expression in Remark 2 of Sheffet (2019).

5.2.1 SIMULATION STUDY

We simulate data from a normal linear model with p predictors, where p is set to 2, 6, or 9. The sample size n (in thousands) varies from 5 to 10,000. The number of active predictors in the true model $|T|$ depends on the value of p and ranges from 0 (null model is true) to p (full model is true). Specifically, if $p = 2$, we set $|T| \in \{0, 1, 2\}$; if $p = 6$, we set $|T| \in \{0, 3, 6\}$; and if $p = 9$, we set $|T| \in \{0, 4, 9\}$. The predictors are independently drawn from the uniform distribution on $(-2, 2)$. Following Hastie et al. (2017), we define the signal-to-noise ratio (SNR) as the variance of the regression mean (which is random, since we are simulating predictors and β) divided by σ^2 . In our simulations, we assume that the intercept is zero and β is a p -dimensional vector equal to $b[1, \dots, 1]'$. We use optimization to find σ^2 and b such that $\text{SNR} = 0.5$ and the response falls within $(-2, 2)$ with high probability. For each combination of $|T|$ and n , we simulate 1,000 data sets. All the data sets we simulated are such that the response falls in $(-2, 2)$. We consider $\varepsilon \in \{0.5, 0.9\}$ and, in the case of the Wishart mechanism, we set $\delta = 1/n$.

We assess the performance of the methods by tracking Monte Carlo averages of predictive mean squared errors and the posterior probability of the true model. We define the predictive mean squared error as $\text{PMSE} = n^{-1} \|V_T \beta_T - V \beta^*\|_2^2$, where V_T is a design matrix containing truly active predictors, β_T is the true value of β , and β^* is the differentially private model-averaged posterior expectation.

Figure 3 displays PMSEs for different values of n , p , ε , and $|T|$. As expected, the PMSEs for both private and non-private approaches decrease as the sample size increases. We observe that the PMSEs of the differentially private methods are smaller when ε is 0.9 compared to when ε is 0.5, and they are always higher than the non-private PMSEs. This is expected, as larger values of ε should lead to greater statistical utility. In most cases, the methods based on the Laplace mechanism have a lower PMSE than those based on the Wishart mechanism. However, the Wishart mechanism seems to perform better in the case when p is either 6 or 9, $|T|$ is zero, and the sample size is small. Although this is slightly less evident in the figure, upon closer inspection, we can see that methods based on hard-thresholding tend to have a slightly lower PMSE when $|T|$ is zero but cease to be advantageous when $|T|$ is large and the sample size is small.

Figure 4 displays the posterior probabilities of the true model, for the same values of n , p , ε , and $|T|$ that we used in Figure 3. The results are consistent with what we observed for PMSEs. We also observe that, although all probabilities increase as the sample size increases, the rate at which they increase depends on p . For higher values of p , the rate of increase is lower because the computational complexity of the problem increases with the dimension of G . The number of entries in G increases quadratically in p , and so does the variance of the perturbation term added to ensure differential privacy. This fact affects the convergence rate of the posterior probabilities.

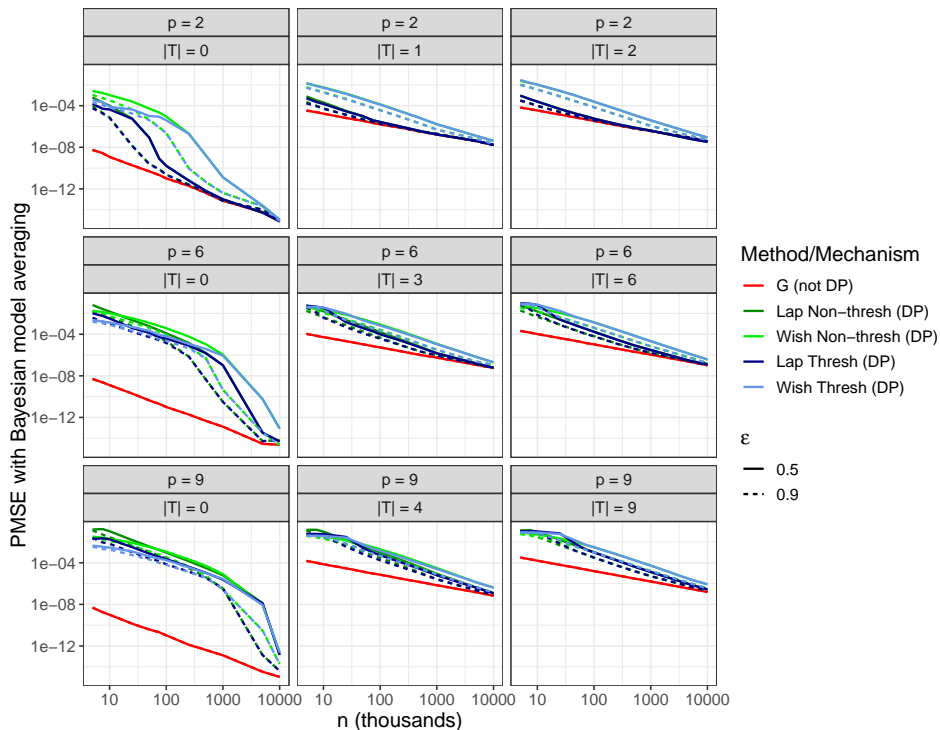


Figure 3: Simulation study: Sample size (x -axis) against $\log(\text{PMSE})$ (y -axis) with Zellner-Siow prior.

5.2.2 APPLICATION: CURRENT POPULATION SURVEY

The data set includes $n = 49,436$ heads of households with non-negative incomes. We consider 6 predictors: age in years (β_1), age squared (β_2), marital status (β_3), sex (β_4), education (β_5), and race (β_6). All predictors are numeric or binary except for education, which is an ordinal variable. To reduce the number of coefficients in the model, we treat education as numeric, ranging from 1 (for less than 1st grade) to 16 (for doctoral degree). The binary predictors are: marital status (1: civilian spouse present; 0: otherwise), sex (1: male; 0: female), and race (1: white; 0: otherwise). The response variable is income.

In this application, the non-private inclusion probabilities are all close to one. To provide a more challenging benchmark for our methods, we permute the rows for marital status and education in the design matrix to artificially make the inclusion probabilities for β_3 and β_5 close to zero. The predictors and the response are centered and rescaled to the interval $(-0.5, 0.5)$.

Figure 5 displays the posterior expected values of β_1 , β_3 , and β_4 with the Zellner-Siow prior and $\epsilon = 0.9$.

We use the histograms described in Section 5.1 to define approximate 95% confidence sets for $T(G) = \mathbb{E}(\beta_j | G)$. Our choice of matrix norm is the Frobenius norm. Specifically, we run our procedure 250 times and, for each run and a fixed collection of bins $\mathcal{B}_1, \dots, \mathcal{B}_K$, we summarize each $T(\hat{\mathcal{C}}_{1-\alpha})$ with its corresponding histogram $\text{Hist}(T, \hat{\mathcal{C}}_{0.95}) = \{(\mathcal{B}_k, d_k)\}_{k=1}^K$.

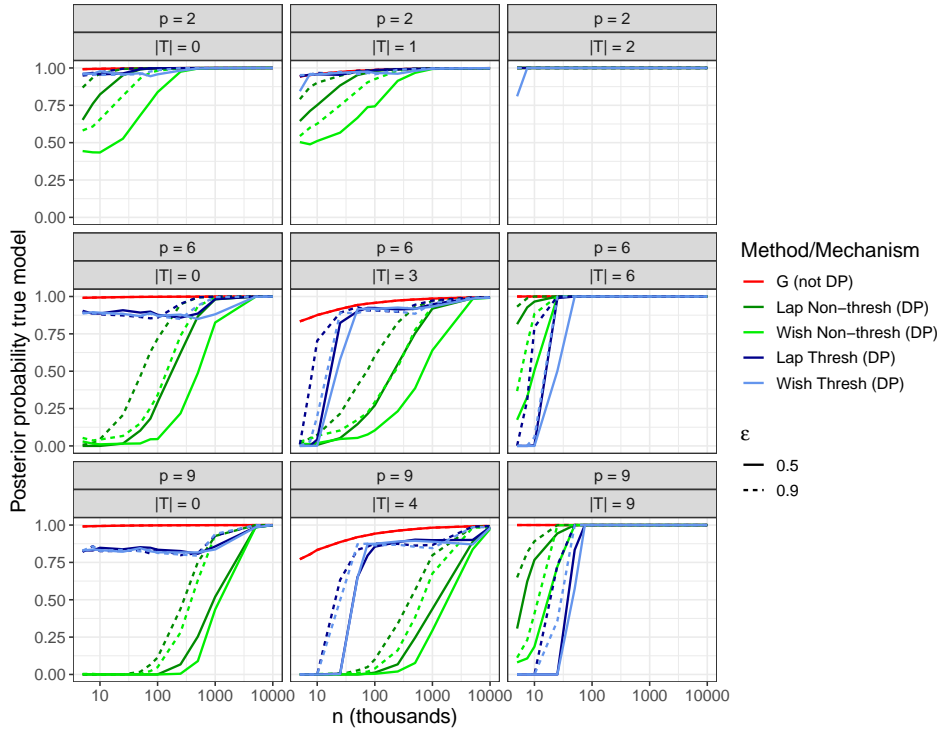


Figure 4: Simulation study: Sample size (x -axis) against posterior probability of the true model (y -axis) with Zellner-Siow prior.

If we let $d_1^{(l)}, \dots, d_K^{(l)}$ be the densities values of the histogram associated with the l -th run, $l = 1, \dots, 250$, we define average histograms as

$$\overline{\text{Hist}} \left(T(\cdot) = \mathbb{E}(\beta_j \mid \cdot), \hat{C}_{0.95} \right) = \left\{ \left(\mathcal{B}_k, \bar{d}_k = \frac{1}{250} \sum_{l=1}^{250} d_k^{(l)} \right) \right\}_{k=1}^K, \quad j \in \{1, 3, 4\}.$$

The results are displayed in Figure 5. In all cases, the differentially private methods are close to the non-private answers. We can also see that the histograms are useful for quantifying the uncertainty introduced by the mechanism, since their spread increases when the thresholded and non-thresholded methods do not agree in their estimates.

5.3 Guidelines

Thresholded methods tend to perform best when the true number of predictors is small. On the other hand, when most predictors are active, non-thresholded methods tend to outperform thresholded methods. The root cause behind this phenomenon is that thresholding shrinks the elements in G^* , which promotes sparsity.

In practice, we recommend that users run analyses with both thresholded and non-thresholded methods. This can be done without affecting the privacy budget of the analyst

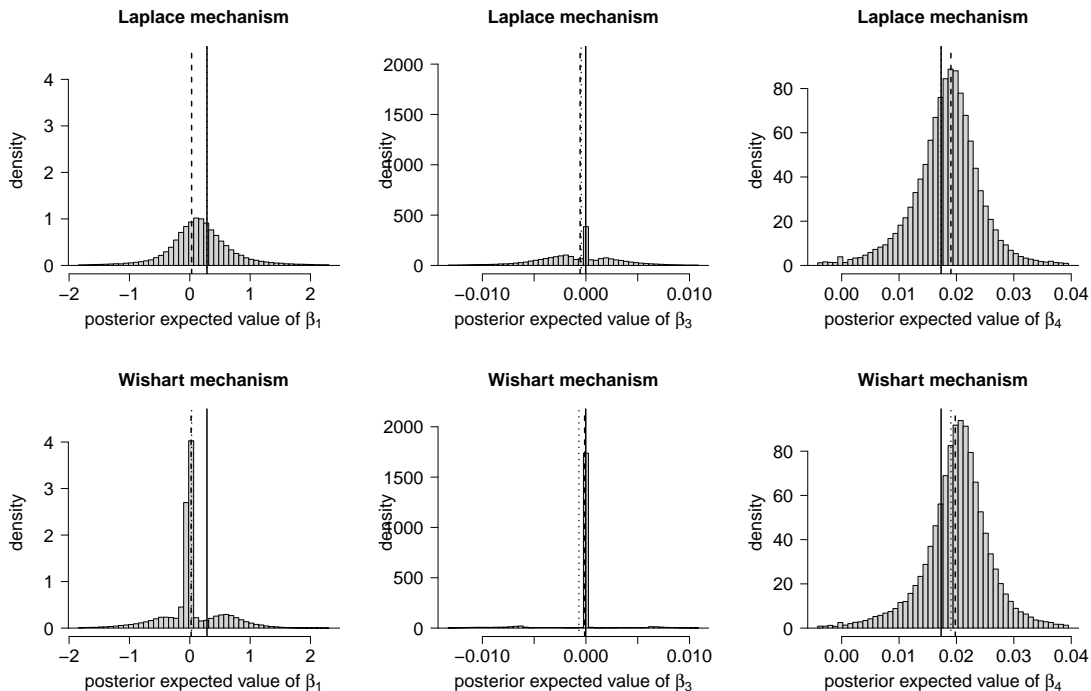


Figure 5: Current population survey: Posterior expectations of coefficients β_1, β_3 , and β_4 with Zellner-Siow prior and $\varepsilon = 0.9$. The vertical solid lines are the non-private posterior expectations, whereas the dashed lines and dotted lines are averaged posterior expectations estimated with non-thresholded methods and thresholded methods, respectively.

because both G_r^* and the thresholded matrix G_r^{**} are post-processed versions of the same differentially private matrix G^* .

Finally, we recommend reporting confidence sets whenever possible, since we find them to be a valuable tool for quantifying the uncertainty introduced by the mechanisms.

6. Conclusions and Future Work

In this article, we proposed differentially private methods for hypothesis testing, model averaging, and model selection for normal linear models. Under regularity conditions, the methods are consistent. The regularity conditions we have imposed are similar to the conditions used in the literature for establishing consistency of non-differentially private methods.

Our methods for hypothesis testing are based on data-splitting and censoring statistics. We have studied the effects of these operations on the performance of the methods. In the case of data-splitting, increasing the number of subsets reduces the variance of the differentially private statistics, but it adds bias. In the case of censoring, more stringent

censoring reduces the variance, but it can lead to substantial bias if the true, confidential statistic lies outside the uncensored range.

The methods we proposed for model averaging and selection are based on a perturbed sufficient statistic. If we suspect that the ground truth is sparse, we recommend hard-thresholding the perturbed sufficient statistic; however, if most predictors are active, hard-thresholding can lead to underfitting.

The methodology proposed here could be extended in a number of ways. It would be useful to extend the methods to generalized linear models through the framework proposed in Li and Clyde (2018). The implementation for hypothesis testing is straightforward, but our methods for model uncertainty, which are based on sufficient statistics, cannot be applied directly. This obstacle can be overcome using approximate sufficient statistics, as proposed in Huggins et al. (2017). This approach has been used successfully in estimation problems under differential privacy constraints in Kulkarni et al. (2021). It would also be interesting to extend the methods to survival models because health records are confidential. In this case, the extension could adapt the framework proposed in Castellanos et al. (2021).

In our work, we have used off-the-shelf techniques for establishing differential privacy. While we observe that our proposals can be useful in practice, it might be possible to design more efficient mechanisms that are specifically tailored to the tasks we considered. This is another interesting avenue for future research.

Acknowledgments

The authors would like to thank the feedback from two anonymous reviewers that greatly improved the presentation and contents of the article. The research of the second author was supported by the National Science Foundation National Center for Science and Engineering Statistics [49100420C0002 and 49100422C0008] and the Test Resource Management Center (TRMC) within the Office of the Secretary of Defense (OSD), contract #FA807518D0002.

Appendix

In Section A, we include the proofs of the propositions in the main text, as well as auxiliary results that are helpful for proving them. In Section B we study the effects of censoring and data-splitting in the High School and Beyond Survey application we considered in Section 4. Finally, in Section C we include additional plots for the simulation study and the application in Section 4 of the main text.

Appendix A. Proofs

We use the notation $a \lesssim b$ and $a \gtrsim b$ to denote $\exists K < \infty : a \leq Kb$ and $\exists K' < \infty : a \geq K'B$, respectively, with the understanding that, if we are taking limits that depend on n , K and K' do not depend on n . Unless stated otherwise, vector norms are Euclidean norms; that is, $\|x\|^2 = \sum_{i=1}^p x_i^2$ for $x \in \mathbb{R}^p$. We use the $O_p(1)$ notation in our proofs, which we define below.

Definition 1 *Let X_n be a random variable taking values in \mathbb{R}^n . Then $X_n \in O_p(1)$ if, for all $\epsilon > 0$, there exist $n_\epsilon, M_\epsilon < \infty$ so that for all $n \geq n_\epsilon$, $P[\|X_n\|^2 > M_\epsilon^2] < \epsilon$.*

A.1 Auxiliary results for Proposition 3

Proposition 2 [Birgé (2001), Laurent and Massart (2000)] *Let $X \sim \chi_m^2(\lambda)$, $Y \sim \chi_n^2$ and $x \geq 0$. Then,*

$$\begin{aligned} P(X \leq m + \lambda - 2\sqrt{(m + 2\lambda)x}) &\leq \exp(-x) \\ P(Y \geq n + 2\sqrt{nx} + 2x) &\leq \exp(-x) \end{aligned}$$

Proposition 3 *Let $R_i^2 \sim \text{Beta}(p/2, (b_i - p - p_0)/2)$ for $b_i > p_0 + p + 2$. Then, if $p \geq 2$:*

$$P(R_i^2 > k_i) \leq \frac{2(1 - k_i)^{(b_i - p - p_0)/2}}{(b_i - p - p_0) B(p/2, (b_i - p - p_0)/2)}.$$

If $p = 1$:

$$P(R_i^2 > k_i) \leq \frac{1}{B(1/2, (b_i - 1 - p_0)/2)} \left(\frac{(1 - k_i)^{b_i - p_0 - 2} \log(1/k_i)}{b_i - p_0 - 2} \right)^{1/2},$$

where $B(\cdot, \cdot)$ is the Beta function.

Proof *This result is straightforward to prove, but we could not find it in the literature. Assuming $p \geq 2$:*

$$\begin{aligned} P(R_i^2 > k_i) &= \frac{1}{B(p/2, (b_i - p - p_0)/2)} \int_{k_i}^1 x^{p/2-1} (1-x)^{(b_i - p - p_0)/2-1} dx \\ &\stackrel{(if\ p \geq 2)}{\leq} \frac{1}{B(p/2, (b_i - p - p_0)/2)} \int_{k_i}^1 (1-x)^{(b_i - p - p_0)/2-1} dx \\ &= \frac{2(1 - k_i)^{(b_i - p - p_0)/2}}{(b_i - p - p_0) B(p/2, (b_i - p - p_0)/2)}. \end{aligned}$$

If $p = 1$, using the Cauchy-Schwarz inequality:

$$\begin{aligned}
P(R_i^2 > k_i) &= \frac{1}{B(1/2, (b_i - p_0 - 1)/2)} \int_{k_i}^1 x^{-1/2} (1-x)^{(b_i - p_0 - 1)/2 - 1} dx \\
&\leq \frac{1}{B(1/2, (b_i - p_0 - 1)/2)} \left(\int_{k_i}^1 x^{-1} dx \right)^{1/2} \left(\int_{k_i}^1 (1-x)^{b_i - p_0 - 3} dx \right)^{1/2} \\
&= \frac{1}{B(1/2, (b_i - p_0 - 1)/2)} \left(\frac{(1-k_i)^{b_i - p_0 - 2} \log(1/k_i)}{b_i - p_0 - 2} \right)^{1/2},
\end{aligned}$$

as required. ■

A.2 Proof of Proposition 1 in main text

First of all, note that

$$\log \tilde{B}_{10} = \sum_{i=1}^M \log B_{10,i}^c / M + \eta, \quad \log \tilde{I}_{10} = \sum_{i=1}^M \log I_{10,i}^c / M + \eta,$$

and recall that $\log B_{10}^* = (\log \tilde{B}_{10} \vee L) \wedge U$ and $\log I_{10}^* = (\log \tilde{I}_{10} \vee L) \wedge U$. Given Assumption 3, the censoring is asymptotically irrelevant. In other words, for any fixed $a \in \mathbb{R}$, there exists a finite N such that $n \geq N$ implies that $P(\log B_{10}^* > a) = P(\log \tilde{B}_{10} > a)$ and $P(\log I_{10}^* > a) = P(\log \tilde{I}_{10} > a)$. For that reason, we show consistency by studying the asymptotic behavior of \tilde{B}_{10} and \tilde{I}_{10} . Asymptotically, the effect of η is irrelevant as well since, by Assumption 4, $\eta \rightarrow_P 0$.

It remains to show that averages of censored logarithms of Bayes factors and information criteria are consistent. We prove consistency by cases. First, we show consistency when H_0 is true. Then, we show consistency when H_1 is true.

Case H_0 is true: We prove consistency for Bayes factors first. The proof for information criteria is very similar.

Using the inequality $(1 + x/n)^n \leq \exp(x)$, we can bound the non-private Bayes factor as follows:

$$\begin{aligned}
B_{10,i} &= \int_0^\infty (g_i + 1)^{-p/2} \left[1 + \frac{g_i R_i^2}{1 + g_i(1 - R_i^2)} \right]^{(b_i - p_0)/2} \pi_i(g_i) \nu(dg_i) \\
&\leq \exp\left(\frac{b_i - p_0}{2} \frac{R_i^2}{1 - R_i^2}\right) \int_0^\infty (g_i + 1)^{-p/2} \pi_i(g_i) \nu(dg_i) \\
&= \exp\left(\frac{b_i - p_0}{2} \frac{R_i^2}{1 - R_i^2}\right) I_{\pi_i}.
\end{aligned}$$

Let $a \in \mathbb{R}$ be a fixed constant. By Assumption 3, there exists a finite N so that $n \geq N$ implies $L < a < U$. Using a union bound, for $n \geq N$:

$$\begin{aligned} P\left(\frac{1}{M} \sum_{i=1}^M \log B_{10,i}^c > a\right) &\leq \sum_{i=1}^M P(\log B_{10,i}^c > a) \\ &\leq \sum_{i=1}^M P(R_i^2 > k_i), \end{aligned}$$

where $k_i = t_i/(1+t_i)$ for $t_i = 2[a - \log I_{\pi_i}]/(b_i - p_0)$, which is positive for b_i large enough since $\lim_{n \rightarrow \infty} \sup_{i \in 1:M} \int_0^\infty b_i^{p/2} (g_i + 1)^{-p/2} \pi_i(g_i) \nu(dg_i) < \infty$ by Assumption 6.

Let $p \geq 2$. Under H_0 , $R_i^2 \sim \text{Beta}(p/2, (b_i - p - p_0)/2)$. We show that $\sum_{i=1}^M P(R_i^2 > k_i)$ goes to zero with the help of Proposition 3 and the inequality for the beta function $1/B(x, y) \leq (x+y)^{x+y-1}/(x^{x-1}y^{y-1})$ (see e.g. Equation 4 in Grenié and Molteni (2015)). We also use the inequality $(1+x/k)^{-k} \leq \exp(-x)(1-x^2/k)^{-1}$ for $x = b_i(a - \log I_{\pi_i})/(b_i - p_0)$ and $k = b_i/2$, which is valid for $|x| \leq k$. The condition is satisfied when the sample size is large enough by Assumption 6.

Putting it all together:

$$\begin{aligned} \sum_{i=1}^M P(R_i^2 > k_i) &\leq \sum_{i=1}^M \frac{2(1-k_i)^{(b_i-p-p_0)/2}}{(b_i-p-p_0) B(p/2, (b_i-p-p_0)/2)} \\ &\lesssim \sum_{i=1}^M (b_i-p-p_0)^{p/2-1} (1-k_i)^{b_i/2} \\ &= \sum_{i=1}^M \frac{(b_i-p-p_0)^{p/2-1}}{[1+2[a-\log I_{\pi_i}]/(b_i-p_0)]^{b_i/2}} \\ &\lesssim \sum_{i=1}^M (b_i-p-p_0)^{p/2-1} I_{\pi_i} \\ &\lesssim \sup_{i \in 1:M} \frac{M}{b_i-p-p_0} b_i^{p/2} I_{\pi_i} \\ &\rightarrow 0. \end{aligned}$$

In the last steps, we used Assumptions 2 and 6. We can use essentially the same argument to show consistency for information criteria under H_0 when $p \geq 2$. In this case, $k_i = 1 - b_i^{-\rho_i/b_i} \exp(-2a/b_i)$ and

$$\begin{aligned} P\left(\sum_{i=1}^M \log I_{10,i}^c/M > a\right) &\lesssim \sum_{i=1}^M P(R_i^2 > k_i) \\ &\lesssim \sum_{i=1}^M b_i^{(p-\rho_i)/2-1} \\ &\lesssim \sup_{i \in 1:M} \frac{M}{b_i} b_i^{(p-\rho_i)/2} \\ &\rightarrow_{\rho_i \geq p} 0. \end{aligned}$$

Let $p = 1$. We can use the same strategy we used for $p \geq 2$, but with the appropriate bounds and assumptions. Once again, we can use the inequalities $1/B(x, y) \leq (x + y)^{x+y-1}/(x^{x-1}y^{y-1})$, Proposition 3, and assumptions as needed.

More explicitly:

$$\begin{aligned}
 \sum_{i=1}^M P(R_i^2 > k_i) &\leq \sum_{i=1}^M \frac{1}{B(1/2, (b_i - 1 - p_0)/2)} \left(\frac{(1 - k_i)^{b_i - p_0 - 2} \log(1/k_i)}{b_i - p_0 - 2} \right)^{1/2} \\
 &\lesssim \sum_{i=1}^M \left[(1 - k_i)^{b_i - p_0 - 2} \log(1/k_i) \right]^{1/2} \\
 &\lesssim \sum_{i=1}^M I_{\pi_i} \log(1/k_i)^{1/2} \\
 &\lesssim \sum_{i=1}^M I_{\pi_i} [\log(b_i / \log b_i)]^{1/2} \\
 &\lesssim \sum_{i=1}^M b_i^{-1/2} [\log(b_i / \log b_i)]^{1/2} \\
 &\lesssim M \sup_{i \in 1:M} \sqrt{\log(b_i / \log b_i) / b_i} \\
 &\rightarrow 0,
 \end{aligned}$$

as required. Using the same argument, we can prove consistency under H_0 for information criteria when $p = 1$:

$$\begin{aligned}
 P \left(\sum_{i=1}^M \log I_{10,i}^c / M > a \right) &\lesssim \sum_{i=1}^M P(R_i^2 > k_i) \\
 &\lesssim \sum_{i=1}^M (1 - k_i)^{(b_i - 3)/2} \log(1/k_i)^{1/2} \\
 &\lesssim_{\rho_i \geq 1} b_i^{-\rho_i - \rho_i/b_i} \sqrt{\log b_i} \\
 &\lesssim \frac{M \sqrt{\log b_i}}{b_i} b_i^{1 - \rho_i(1 - 1/b_i)} \\
 &\rightarrow 0,
 \end{aligned}$$

which completes the proof.

Case H_1 is true: First, we show that

$$\lim_{n \rightarrow \infty} P(\forall i \in 1 : M, R_i^2 \geq b_i^{-0.5}) = 1.$$

Under H_1 , $R_i^2 \stackrel{\text{ind.}}{\sim}$ noncentralBeta($p/2, (b_i - p - p_0)/2, \lambda_i$), where $\lambda_i = \beta' X_i' X_i \beta / \sigma^2$ is the noncentrality parameter of the noncentral beta distribution (type I). In other words, we can write

$$R_i^2 = \frac{X_i}{X_i + Y_i}; \quad X_i \sim \chi_p^2(\lambda_i) \perp Y_i \sim \chi_{b_i - p - p_0}^2.$$

Let $a_i = b_i^{-0.5}/(1 - b_i^{-0.5})$. Then,

$$R_i^2 > b_i^{-0.5} \Leftrightarrow \frac{X_i}{Y_i} > a_i$$

Let $k_i = a_i[(b_i - p - p_0) + 2\sqrt{(b_i - p - p_0)b_i^{0.5} + 2b_i^{0.5}}]$. Then,

$$P(X_i/Y_i > a_i) \geq P(X_i > k_i)P(Y_i < k_i/a_i).$$

Provided that $\lambda_i + p > k_i$, Proposition 2 ensures that

$$P(X_i > k_i) \geq 1 - \exp\left(-\frac{(p + \lambda_i - k_i)^2}{4(p + 2\lambda_i)}\right).$$

The condition $\lambda_i + p > k_i$ is satisfied for a large enough sample size for all $i \in \{1, 2, \dots, M\}$ under Assumption 5. On the other hand, also by Proposition 2,

$$P(Y_i < k_i/a_i) \geq 1 - \exp(-b_i^{0.5}),$$

Thus, under Assumption 2:

$$\begin{aligned} P\{\forall i \in 1 : M, R_i^2 \geq b_i^{-0.5}\} &\geq \prod_{i=1}^M \{P(X_i > k_i)P(Y_i < k_i/a_i)\} \\ &\geq \inf_{i \in 1:M} \left\{ \left[1 - \exp\left(-\frac{(p + \lambda_i - k_i)^2}{4(p + 2\lambda_i)}\right) \right]^{b_i} [1 - \exp(-b_i^{0.5})]^{b_i} \right\}^{M/b_i} \\ &\rightarrow 1, \end{aligned}$$

as required.

Conditioning on the high probability event $\{\forall i \in 1 : M, R_i^2 \geq b_i^{-0.5}\}$:

$$\begin{aligned} \frac{1}{M} \sum_{i=1}^M \log B_{10,i}^c &\geq \frac{1}{M} \sum_{i=1}^M \log \int_0^\infty \frac{(g_i + 1)^{(b_i - p - p_0)/2}}{[1 + g_i(1 - b_i^{-0.5})]^{(b_i - p_0)/2}} \pi_i(g_i) \nu(dg_i) \\ &\gtrsim \log \inf_{i \in 1:M} \left[\frac{b_i + 1}{b_i(1 - b_i^{-0.5}) + 1} \right]^{(b_i - p_0)/2} \int_{b_i}^\infty (g_i + 1)^{-p/2} \pi_i(g_i) \nu(dg_i) \\ &\rightarrow \infty \end{aligned}$$

and, with high probability,

$$\begin{aligned} \frac{1}{M} \sum_{i=1}^M \log I_{10,i}^c &\gtrsim \inf_{i \in 1:M} \log \left(\frac{b_i^{-\rho_i/2}}{(1 - b_i^{-0.5})^{b_i/2}} \right) \\ &\rightarrow \infty, \end{aligned}$$

which completes the proof.

A.3 Auxiliary results for Proposition 2

Proposition 4 (*Hanson-Wright inequality; see e.g. Rudelson et al. (2013)*). Let W be a random vector in \mathbb{R}^n . Let $\{W_i\}_{i=1}^n$ be the components of W . Let $\mathbb{E}(W_i) = 0$ and $0 < K < \infty$ be a constant so that $\mathbb{E}[\exp(W_i^2/K^2)] < 2$. Let A be a $n \times n$ real matrix. Then, for a constant $c > 0$,

$$P[|W'AW - \mathbb{E}(W'AW)| > t] \leq 2 \exp \left\{ -c \frac{t}{K^2} \min \left(\frac{t}{K^2 \|A\|_F^2}, \frac{1}{\|A\|_{op}} \right) \right\},$$

where $\|A\|_F = \sqrt{\text{tr}(A'A)}$ is the Frobenius norm of A and $\|A\|_{op} = \sup_{\|x\| \neq 0} \|Ax\|/\|x\|$.

Proposition 5 (*Convergence of G_r^* and G_r^{**}*) Under the assumptions of Proposition 2 in the main text, $G_r^*/n \rightarrow_P \bar{G}_\infty$ and $G_r^{**}/n \rightarrow_P \bar{G}_\infty$, a constant symmetric matrix.

Proof Let $G^*/n = G/n + E/n$. The entries E_{ij}/n converge to zero in probability by Assumption 3.

We show that $G/n \rightarrow_P \bar{G}_\infty$ with \bar{G}_∞ symmetric. Recall that

$$G/n = \begin{bmatrix} V'V/n & V'Z/n \\ Z'V/n & Z'Z/n \end{bmatrix}.$$

We establish the convergence of G/n to \bar{G}_∞ by establishing the convergence of the blocks. All the probability statements, expectations, and variances in this proof are given X_0 and V .

By Assumption 5, $\lim_{n \rightarrow \infty} V'V/n = S_1$. Then,

$$\begin{aligned} \lim_{n \rightarrow \infty} \mathbb{E}(V'Z/n) &= \lim_{n \rightarrow \infty} V'V_T \beta_T/n = S_2 \beta_T \\ \lim_{n \rightarrow \infty} \text{Var}(V'Z/n) &= \lim_{n \rightarrow \infty} \frac{\sigma_T^2}{n^2} V'V = 0_{p \times p} \\ V'Z/n &\rightarrow_P S_2 \beta_T, \end{aligned}$$

where S_2 is a submatrix of S_1 .

Let us study the asymptotic behavior of $Z'Z/n$. We can write

$$\begin{aligned} Z'Z/n &= Q/n + 2\beta_T' V_T' Y/n - \beta_T' V_T' V_T \beta_T/n, \\ Q &= [Y - \mathbb{E}(Y)]'(I_n - P_{X_0})[Y - \mathbb{E}(Y)]. \end{aligned}$$

On the one hand,

$$\begin{aligned} \lim_{n \rightarrow \infty} \mathbb{E}(2\beta_T' V_T' Y/n - \beta_T' V_T' V_T \beta_T/n) &= \lim_{n \rightarrow \infty} \beta_T' V_T' V_T \beta_T/n = \beta_T' S_3 \beta_T \\ \lim_{n \rightarrow \infty} \text{Var}(2\beta_T' V_T' Y/n - \beta_T' V_T' V_T \beta_T/n) &= \lim_{n \rightarrow \infty} \frac{4\sigma_T^2}{n^2} \beta_T' V_T' V_T \beta_T = 0 \\ 2\beta_T' V_T' Y/n - \beta_T' V_T' V_T \beta_T/n &\rightarrow_P \beta_T' S_3 \beta_T, \end{aligned}$$

where S_3 is a submatrix of S_1 .

We can show that $Q/(n-p_0) \rightarrow_P \sigma_T^2$ with the Hanson-Wright inequality, which in turn implies that $Q/n \rightarrow_P \sigma_T^2$. We have $\mathbb{E}[Q/(n-p_0)] = \sigma_T^2$. Define $W = Y - \mathbb{E}(Y)$. Then, $\mathbb{E}(W) = 0_n$ and W_i^2 are uniformly bounded since both Y and $\mathbb{E}(Y)$ are by Assumption 1. Therefore, we can pick a finite constant K satisfying $\mathbb{E}[\exp(W_i^2/K^2)] < 2$.

Applying the Hanson-Wright inequality, for any given $\epsilon > 0$,

$$P[|Q/(n-p_0) - \sigma_T^2| > \epsilon] \leq 2e^{-2c\epsilon \min[\epsilon(K^2\|(I-P_{X_0})/(n-p_0)\|_F)^{-1}, (\|(I-P_{X_0})/(n-p_0)\|_{op})^{-1}]} \rightarrow_{n \rightarrow \infty} 0,$$

because $\|(I-P_{X_0})/(n-p_0)\|_{op} = \|(I-P_{X_0})/(n-p_0)\|_F = 1/\sqrt{n-p_0} \rightarrow 0$ as $n \rightarrow \infty$. This implies $Q/n \rightarrow_P \sigma_T^2$ and $Z'Z/n \rightarrow_P \sigma_T^2 + \beta_T' S_3 \beta_T$.

We have shown that

$$G/n \rightarrow_P G/n = \begin{bmatrix} S_1 & S_2 \beta_T \\ \beta_T' S_2' & \sigma_T^2 + \beta_T' S_3 \beta_T \end{bmatrix} = \bar{G}_\infty.$$

Therefore, we have $G^*/n \rightarrow_P \bar{G}_\infty$.

In the case of the non-thresholded matrix G_r^* , we have established that $G_r^*/n = G^*/n + r/n I_{p+1} \rightarrow_P \bar{G}_\infty$ since $r/n \rightarrow 0$ by Assumption 4. In the case of G_r^{**} , there is an indicator that can hard-threshold off-diagonal elements. Let $G_{ij}^{**}/n = G_{ij}^*/n \mathbb{1}(i=j \text{ or } |G_{ij}^*|/n \geq e_\lambda/n)$ be the (i, j) -th entry of G^{**} . On the one hand, $\lim_{n \rightarrow \infty} e_\lambda/n = 0$ because the variance of E_{ij} is finite and does not depend on n . For $i=j$ the indicator is equal to 1. For $i \neq j$:

$$\begin{aligned} E[\mathbb{1}(|G_{ij}^*|/n \geq e_\lambda/n)] &= P[|G_{ij}^*|/n \geq e_\lambda/n] \rightarrow 1 \\ \text{Var}[\mathbb{1}(|G_{ij}^*|/n \geq e_\lambda/n)] &= P[|G_{ij}^*|/n \geq e_\lambda/n] - P[|G_{ij}^*|/n \geq e_\lambda/n]^2 \rightarrow 0 \\ \mathbb{1}(|G_{ij}^*|/n \geq e_\lambda/n) &\rightarrow_P 1. \end{aligned}$$

By Slutsky's lemma, we have that $G_{ij}^{**}/n \rightarrow \bar{G}_{\infty, ij}$ for all i, j . This is enough to show that $G^*/n \rightarrow \bar{G}_\infty$. Finally, since $G_r^{**} = G^{**} + r I_{p+1}$ and $r/n \rightarrow 0$ by Assumption 4, we have $G_r^{**}/n \rightarrow_P \bar{G}_\infty$, as required. ■

Proposition 6 (Convergence of noisy R^2) *Given G_r^* or G_r^{**} and a model-indexing vector $\gamma \in \{0, 1\}^p$, we can construct a differentially private version of R_γ^2 , denoted $R_\gamma^{2,*}$. Let $T \in \{0, 1\}^p$ be the index of the true model. Under the regularity conditions stated in Proposition 5 and assuming that $Z'Z$ is not equal to zero almost surely, the following are true:*

1. *If γ does not nest T (i.e. if there exists i such that $T_i = 1$ and $\gamma_i = 0$), then $R_\gamma^{2,*} \rightarrow R_{\gamma, \infty}^2$ and $R_T^{2,*} \rightarrow R_{T, \infty}^2$ with $R_{\gamma, \infty}^2 < R_{T, \infty}^2$.*
2. *If the $T = 0_p$ is the null model, then $nR_\gamma^{2,*}$ is in $O_p(1)$.*
3. *If γ nests T (i.e. if $T_i = 1$ implies $\gamma_i = 1$), then $[(1 - R_T^{2,*})/(1 - R_\gamma^{2,*})]^{(n-p_0)/2}$ is in $O_p(1)$.*

Proof We prove the three statements separately.

Proof of statement 1. The proofs of $R_\gamma^2 \rightarrow_P R_{\gamma,\infty}^2$ and $R_\gamma^{2,*} \rightarrow_P R_{\gamma,\infty}^2$ are a direct consequence of Proposition 5. Note that

$$R_\gamma^2 = \frac{Z'V_\gamma(V_\gamma'V_\gamma)^{-1}V_\gamma'Z}{Z'Z} = \frac{Z'P_{V_\gamma}Z/n}{Z'Z/n}.$$

On the one hand, $n/Z'Z \rightarrow_P 1/(\sigma_T^2 + \beta_T'S_3\beta_T)$. Then, $V_\gamma'Z$ are subvectors of $V'Z$, so $V_\gamma'Z/n$ converges to a subvector of $S_2\beta_T$. Similarly, $V_\gamma'V_\gamma/n$ converges to a submatrix of S_1 and, since we assume that V is full-rank, $V_\gamma'V_\gamma$ is invertible for any γ . This implies that R_γ^2 converges in probability to some constant $R_{\gamma,\infty}^2$, which has to be between zero and one because $V_\gamma(V_\gamma'V_\gamma)^{-1}V_\gamma'$ is a projection matrix and $Z'P_XZ \leq Z'Z$ for any projection matrix P_X .

The convergence of the noisy $R_\gamma^{2,*}$ to $R_{\gamma,\infty}^2$ can be established after noting that $R_\gamma^{2,*}$ can be constructed by taking submatrices of G_r^* or G_r^{**} and multiplying and dividing terms as needed. We can invoke Proposition 5 and Slutsky's lemma and conclude that $R_\gamma^{2,*} \rightarrow_P R_{\gamma,\infty}^2$, as required.

It remains to show that if γ does not nest the true model T , $R_{\gamma,\infty}^2 < R_{T,\infty}^2$. To see this, note that

$$\lim_{n \rightarrow \infty} \mathbb{E}(Z'P_{V_\gamma}Z/n) = \lim_{n \rightarrow \infty} \beta_T'X_T'P_{V_\gamma}X_T\beta_T/n, \quad \lim_{n \rightarrow \infty} \text{Var}(Z'P_{V_\gamma}Z/n) = 0,$$

and $\beta_T'X_T'P_{V_\gamma}X_T\beta_T < \beta_T'X_T'X_T\beta_T = \beta_T'X_T'P_{V_T}X_T\beta_T$, which implies $R_{\gamma,\infty}^2 < R_{T,\infty}^2$.

Proof of statement 2. If T is the null model, we show that $nR_\gamma^{2,*}$ is in $O_p(1)$. It is useful to write

$$nR_\gamma^{2,*} = \frac{\frac{1}{\sqrt{n}}(Z'V)_\gamma^* \left[\frac{(V'V)_\gamma^*}{n} \right]^{-1} (V'Z)_\gamma^* \frac{1}{\sqrt{n}}}{(Z'Z)^*/n}.$$

By Proposition 5, we know that the denominator converges in probability to a constant. It is enough to show that the numerator is in $O_p(1)$. The matrix $[(V'V)_\gamma^*/n]^{-1}$ is in $O_p(1)$ by Proposition 5. It remains to show that $(V'Z)_\gamma^*/\sqrt{n}$ is in $O_p(1)$. Let $w^* = (V'Z) + E_2$ and k_λ the appropriate threshold for the off-diagonal elements of G^{**} . Then,

$$\begin{aligned} \left\| \frac{1}{\sqrt{n}}(V'Z)_\gamma^* \right\|^2 &\leq \frac{1}{n} \|w^*\|^2 \mathbb{1}(|w^*| > k_\lambda)^2 \\ &\leq \frac{1}{n} \|w^*\|^2 \\ &\leq \frac{1}{n} \|V'Z\|^2 + \frac{1}{n} \|E_2\|^2. \end{aligned}$$

It suffices to show that both the error E_2/\sqrt{n} and the non-private $V'Z/\sqrt{n}$ are in $O_p(1)$. Each of the E_{ij} in E has a finite variance that does not depend on n . Therefore, each E_{ij}/\sqrt{n} has a variance that goes to zero and, since E_2 has p elements, $\|E_2/\sqrt{n}\|$ is in $O_p(1)$. It only remains to show that $V'Z/\sqrt{n}$ is in $O_p(1)$, which we show using the Hanson-Wright inequality.

First, note that $\|V'Z/\sqrt{n}\|^2 = Z'VV'Z/n$. Then, since we are assuming that the true model is the null model:

$$\mathbb{E}(Z'VV'Z/n) = \frac{\sigma^2}{n} \text{tr}(V'V) \rightarrow_{n \rightarrow \infty} c > 0.$$

The limit is a positive constant because $V'V/n$ converges to a symmetric positive-definite matrix by Assumption 5, which also implies $\lim_{n \rightarrow \infty} \|VV'/n\|_F < \infty$ and $\lim_{n \rightarrow \infty} \|VV'/n\|_{op} < \infty$. From here, we can apply the Hanson-Wright inequality to establish that $|Z'VV'Z/n - \mathbb{E}(Z'VV'Z/n)|$ is in $O_p(1)$ and, since $\mathbb{E}(Z'VV'Z/n)$ converges to a constant, we conclude that $Z'VV'Z/n$ is in $O_p(1)$, as required.

Proof of statement 3. Finally, we show that if γ nests T , $[(1 - R_T^{2,*})/(1 - R_\gamma^{2,*})]^{(n-p_0)/2}$ is in $O_p(1)$. Consider the non-private $[(1 - R_T^2)/(1 - R_\gamma^2)]^{(n-p_0)/2}$. We can write

$$\begin{aligned} \left[\frac{1 - R_T^2}{1 - R_\gamma^2} \right]^{(n-p_0)/2} &= \left[1 + \frac{2}{n - p_0} \frac{Z'(P_{V_\gamma} - P_{V_T})Z}{2Z'(I_n - P_{V_\gamma})Z/(n - p_0)} \right]^{(n-p_0)/2} \\ &\leq \exp \left(\frac{Z'(P_{V_\gamma} - P_{V_T})Z}{2Z'(I_n - P_{V_\gamma})Z/(n - p_0)} \right). \end{aligned}$$

The denominator $Z'(I_n - P_{V_\gamma})Z/(n - p_0)$ converges to a constant. This fact follows directly given the asymptotic behavior of $Z'P_{V_\gamma}Z/n$ and $Z'Z/n$ we just described in the proof of the first statement. The same is true for the private version of the statistic, using the argument we used for $R_\gamma^{2,*}$.

The numerator $Z'(P_{V_\gamma} - P_{V_T})Z$ is in $O_p(1)$, which can be shown using the Hanson-Wright inequality. The expectation is $\mathbb{E}[Z'(P_{V_\gamma} - P_{V_T})Z] = \sigma^2(|\gamma| - |T|)$. Let $Q_\gamma = Z'(P_{V_\gamma} - P_{V_T})Z - \sigma^2(|\gamma| - |T|)$. Applying the Hanson-Wright inequality

$$P(Q_\gamma > M) \leq 2 \exp \left\{ -c \frac{M}{K^2} \min \left(\frac{M}{K^2(|\gamma| - |T|)}, 1 \right) \right\},$$

where K is a constant that can be chosen in a similar way as we did in Proposition 5. The right-hand side can be made arbitrarily close to zero by increasing M , which establishes that Q_γ and $Z'(P_{V_\gamma} - P_{V_T})Z$ are in $O_p(1)$. The same is true for the private version of the statistic, using an argument which is similar to the one we used for showing that $nR_\gamma^{2,*}$ is in $O_p(1)$ when the true model is the null model. Both $(V'Z)_\gamma^*/\sqrt{n}$ and $(V'Z)_T^*/\sqrt{n}$ converge to their private versions $(V'Z)_\gamma/\sqrt{n}$ and $(V'Z)_T/\sqrt{n}$ because the error goes to zero and the indicator converges (see Proposition 5 and the earlier proof for $nR_\gamma^{2,*}$ for more detailed versions of these arguments). The private $(V'V)_\gamma^*$ also converge to their non-private versions (see e.g. Proposition 5). Therefore, the private $(Z'(P_{V_\gamma} - P_{V_T})Z)^*$ has the same asymptotic behavior as the private $Z'(P_{V_\gamma} - P_{V_T})Z$, which we have shown to be in $O_p(1)$. This completes the proof. ■

A.4 Proof of Proposition 2 in main text

We consider two cases: one where the true model is the null model and another one where the true model is not the null model.

True model is the null model: Let γ be a model that is not the null model. Then, we have

$$B_{\gamma 0}^* \leq \exp\left(\frac{n-p_0}{2} \frac{R_\gamma^{2,*}}{1-R_\gamma^{2,*}}\right) \int_0^\infty (g+1)^{-|\gamma|/2} \pi(g) \nu(dg).$$

The integral converges to zero by Assumption 6 and the exponential term is in $O_p(1)$ because, by Proposition 6, we know that $nR_\gamma^{2,*}$ is in $O_p(1)$. Therefore, for any model γ which is not the null model, $B_{\gamma 0}^* \rightarrow_P 0$. A similar argument works for information criteria. In such case,

$$I_{\gamma 0}^* \leq n^{-\rho|\gamma|/2} \exp\left(\frac{n}{2} \frac{R_\gamma^{2,*}}{1-R_\gamma^{2,*}}\right) \rightarrow_P 0.$$

True model is not the null model: We study the asymptotic behavior of $B_{\gamma T}^*$, where T is the true model and γ is a model that is not the true model. We split this case into two subcases: one where γ nests the true model, and another one where γ does not nest the true model.

First, note that, since we are working with null-based Bayes factors,

$$B_{\gamma T}^* = \frac{B_{\gamma 0}^* (1 - R_T^{2,*})^{(n-p_0)/2}}{B_{T 0}^* (1 - R_T^{2,*})^{(n-p_0)/2}}$$

we will bound the numerator and denominator separately and put our bounds together. First, we bound the numerator:

$$\begin{aligned} B_{\gamma 0}^* (1 - R_T^{2,*})^{(n-p_0)/2} &= \int_0^\infty (g+1)^{-|\gamma|/2} \left[\frac{1 + g(1 - R_T^{2,*}) - R_T^{2,*}}{1 + g(1 - R_\gamma^{2,*})} \right]^{(n-p_0)/2} \pi(g) \nu(dg) \\ &\leq \left(\frac{1 - R_T^{2,*}}{1 - R_\gamma^{2,*}} \right)^{(n-p_0)/2} \int_0^\infty (g+1)^{-|\gamma|/2} \pi(g) \nu(dg). \end{aligned}$$

Then, we bound the denominator:

$$\begin{aligned} B_{T 0}^* (1 - R_T^{2,*})^{(n-p_0)/2} &\geq \int_n^\infty (g+1)^{-|T|/2} \left[1 - \frac{R_T^{2,*}}{1 + g(1 - R_T^{2,*})} \right]^{(n-p_0)/2} \pi(g) \nu(dg) \\ &\geq \left[1 - \frac{R_T^{2,*}}{1 + n(1 - R_T^{2,*})} \right]^{(n-p_0)/2} \int_n^\infty (g+1)^{-|T|/2} \pi(g) \nu(dg). \end{aligned}$$

Putting the bounds together and using Assumption 6:

$$\begin{aligned} B_{\gamma T}^* &\leq \left(\frac{1 - R_T^{2,*}}{1 - R_\gamma^{2,*}} \right)^{(n-p_0)/2} \left(1 - \frac{R_T^{2,*}}{1 + n(1 - R_T^{2,*})} \right)^{-(n-p_0)/2} \frac{\int_0^\infty (g+1)^{-|\gamma|/2} \pi(g) \nu(dg)}{\int_n^\infty (g+1)^{-|T|/2} \pi(g) \nu(dg)} \\ &\lesssim n^{(|T|-|\gamma|)/2} \left(\frac{1 - R_T^{2,*}}{1 - R_\gamma^{2,*}} \right)^{(n-p_0)/2} \exp(R_\gamma^{2,*}/(1 - R_\gamma^{2,*})). \end{aligned}$$

When γ nests T , $n^{(|T|-|\gamma|)/2}$ goes to zero and the remaining terms are in $O_p(1)$ by Proposition 6, so $B_{\gamma T}^*$ converges to zero in probability. When γ does not nest T , we show that

$$A_n = n^{(|T|-|\gamma|)/2} [(1 - R_T^{2,*}) / (1 - R_\gamma^{2,*})]^{(n-p_0)/2} \rightarrow_P 0$$

Let $R_n = (1 - R_T^{2,*}) / (1 - R_\gamma^{2,*})$, which converges in probability to a constant less than one. Taking logarithms

$$\log \left\{ n^{(|T|-|\gamma|)/2} R_n^{(n-p_0)/2} \right\} = \frac{n-p_0}{2} [(|T| - |\gamma|) \log n / (n-p_0) + \log R_n],$$

which diverges to $-\infty$ in probability, so $A_n \rightarrow_P 0$. The remaining term in the upper bound for $B_{\gamma T}^*$ is in $O_p(1)$. Therefore, we have shown that $B_{\gamma T}^* \rightarrow_P 0$.

The proof for information criteria is essentially the same, but we do not need to bound integrals. Indeed,

$$I_{\gamma T}^* = n^{(\rho_{|T|} - \rho_{|\gamma|})/2} \left(\frac{1 - R_T^{2,*}}{1 - R_\gamma^{2,*}} \right)^{n/2},$$

and we can use the same arguments we used for $B_{\gamma T}^*$ to show that $I_{\gamma T}^*$ is consistent.

A.5 Proof of Proposition 3 in main text

Let $\mathcal{U} \in \mathbb{R}^{n \times (p+1)}$ be a full-rank matrix and define $\mathcal{M} = (I_n - P_{X_0})\mathcal{U}$. Given a Gram matrix \mathcal{G} , we can generate a synthetic data set $\mathcal{D} = [\mathcal{V}; \mathcal{Z}]$ with the formula $\mathcal{D} = \mathcal{M}(\mathcal{M}'\mathcal{M})^{-1/2}\mathcal{G}^{1/2}$. In other words, we have $\mathcal{D}'\mathcal{D} = \mathcal{G}$:

$$\mathcal{D}'\mathcal{D} = \mathcal{G}^{1/2}\mathcal{M}(\mathcal{M}'\mathcal{M})^{-1/2}\mathcal{M}'\mathcal{M}(\mathcal{M}'\mathcal{M})^{-1/2}\mathcal{G}^{1/2} = \mathcal{G}^{1/2}\mathcal{G}^{1/2} = \mathcal{G}.$$

The synthetic data are also centered (the same way that V and Z are centered in our construction of D). This is true because \mathcal{D} is pre-multiplied by $I_n - P_{X_0}$, so it is orthogonal to the span of $X_0 = \mathbf{1}_n$.

Appendix B. Effects of censoring and data-splitting: High School and Beyond Survey

In this section, we revisit the High School and Beyond Survey data set (Section 4.4) to study the effects of setting different censoring limits. For concreteness, we restrict our attention to the test of H_{02} against H_{12} ; that is, the hypothesis test where we wish to know if `read` scores are predictive of `math` scores when `science` scores are already a covariate in the model. The results for the test of H_{01} against H_{11} are similar. The simulation setup is the same as described in Section B. The only changes are the censoring limits.

In Figure 6, we display the results for the differentially private posterior probability of H_{12} for different values of ε and M . We censor the posterior probability of H_{12} at $[0.35, 0.65]$, $[0.25, 0.75]$, $[0.01, 0.99]$, and $[0.001, 0.999]$. The true, non-private posterior probability of H_{12} is near 1. For all censoring limits, increasing the number of subgroups decreases the variance of the output, but it induces a bias that shrinks the probability to 0.5. More stringent censoring limits reduce the variance as well, but come at the cost of potential

bias: for instance, censoring the posterior probability at $[0.35, 0.65]$ is clearly too stringent, since the true posterior probability is much higher than the upper limit.

In Figure 7, we display the analogous result for likelihood ratios. In this case, the lower censoring limit is set to $L = 0$, which is a natural lower bound for likelihood ratios. The gray $+$ and Δ are calibrated critical values for rejection of H_{02} at significance at levels 0.01 and 0.05, respectively. The upper censoring limits U are set to 1, 2, 7, and 10. We arrive at the same conclusions we reached with the Bayesian analysis. The true, non-private likelihood ratio is above 10, and if we censor at a much lower value (such as 1 or 2) the test fails to be powerful. The test performs best if the lower bound is 10, in which case the test is quite powerful, especially for $M \geq 5$ and $\delta \geq 0.25$.

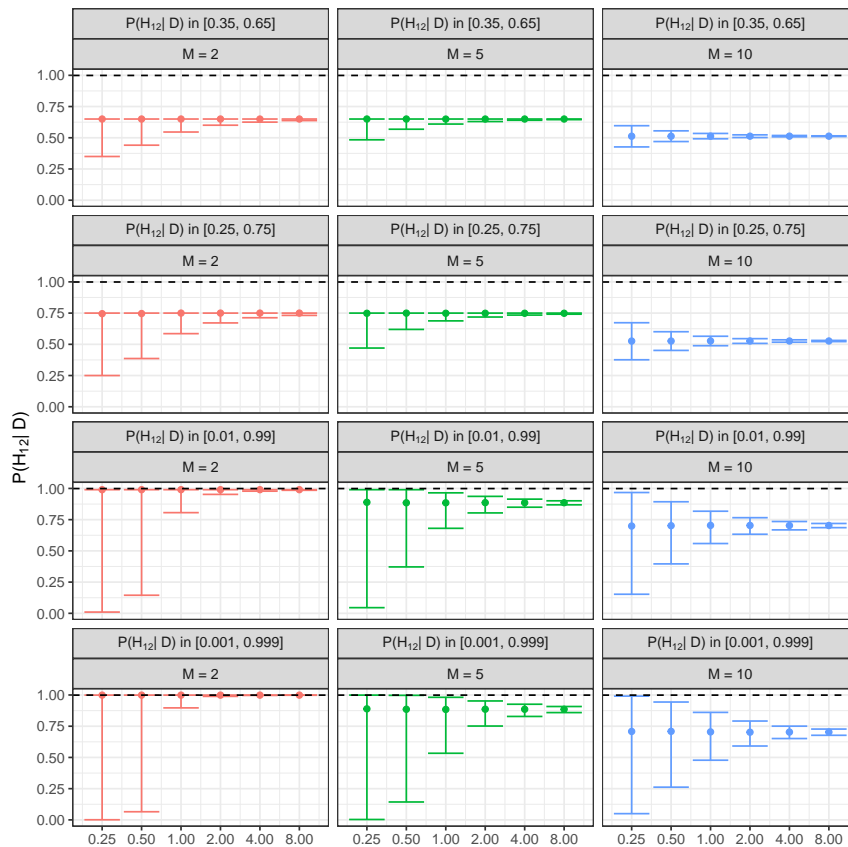


Figure 6: Distribution of $P^*(H_{12} | D)$ as a function of ε , M , and censoring limits. The lower endpoint of the error bars is the first quartile of the distribution, the midpoint is the median, and the upper endpoint is the third quartile. The dashed lines are the non-private posterior probabilities.

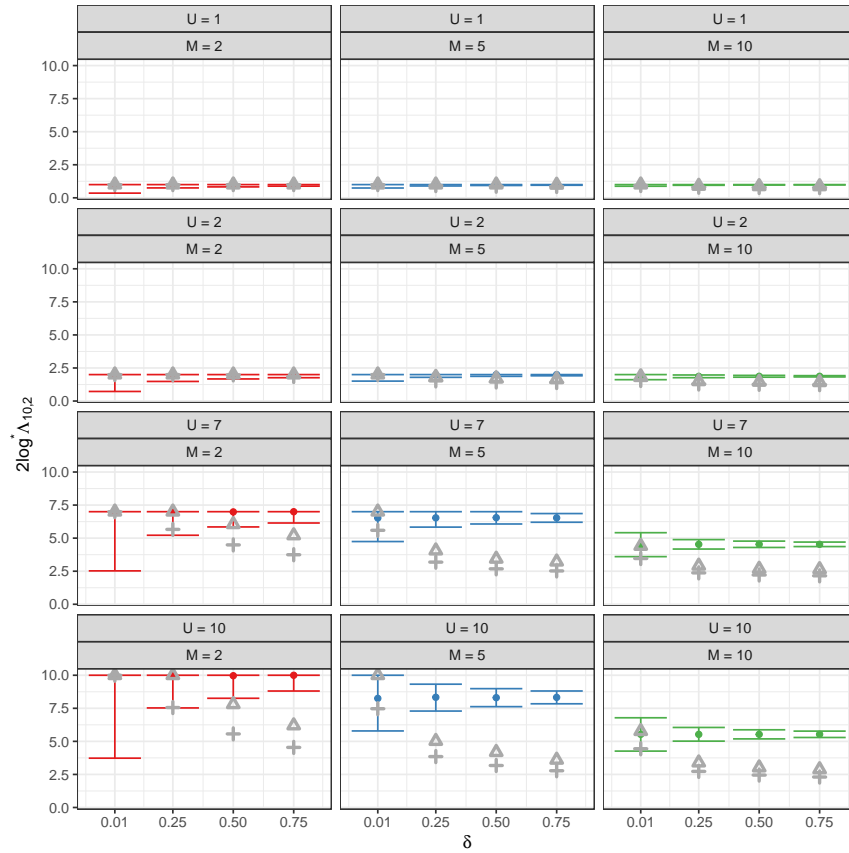


Figure 7: Distribution of $2 \log \Lambda_{10,1}^*$ and $2 \log \Lambda_{10,2}^*$ as a function of δ , M , and censoring upper limit U . The gray $+$ and Δ are corrected critical values at the 0.01 and 0.05 significance levels, respectively.

Appendix C. Additional Plots for Simulation Study

In this section, we include additional plots for the simulation study in Section 5.2. We show results with BIC combined with least-squares estimates, as well as results for additional values of ε that were not included in the main text. The interpretation of the plots is the same: thresholded methods perform best when $|T|$ is small, and non-thresholded methods perform best when $|T|$ is large.

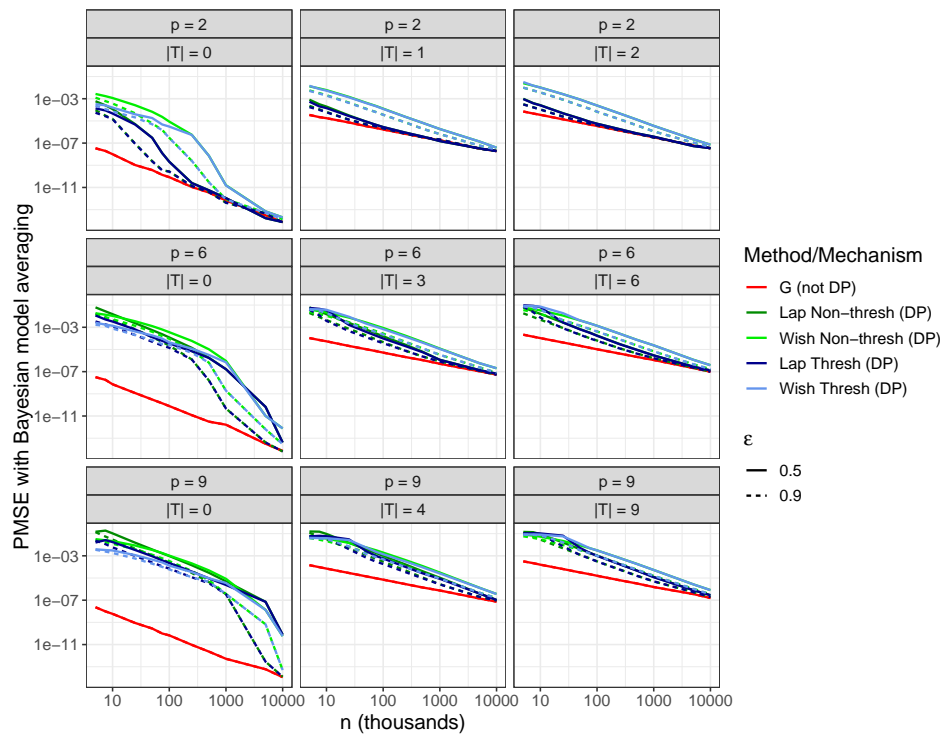


Figure 8: Simulation study: Sample size (x -axis) against $\log(\text{PMSE})$ (y -axis) with BIC prior.

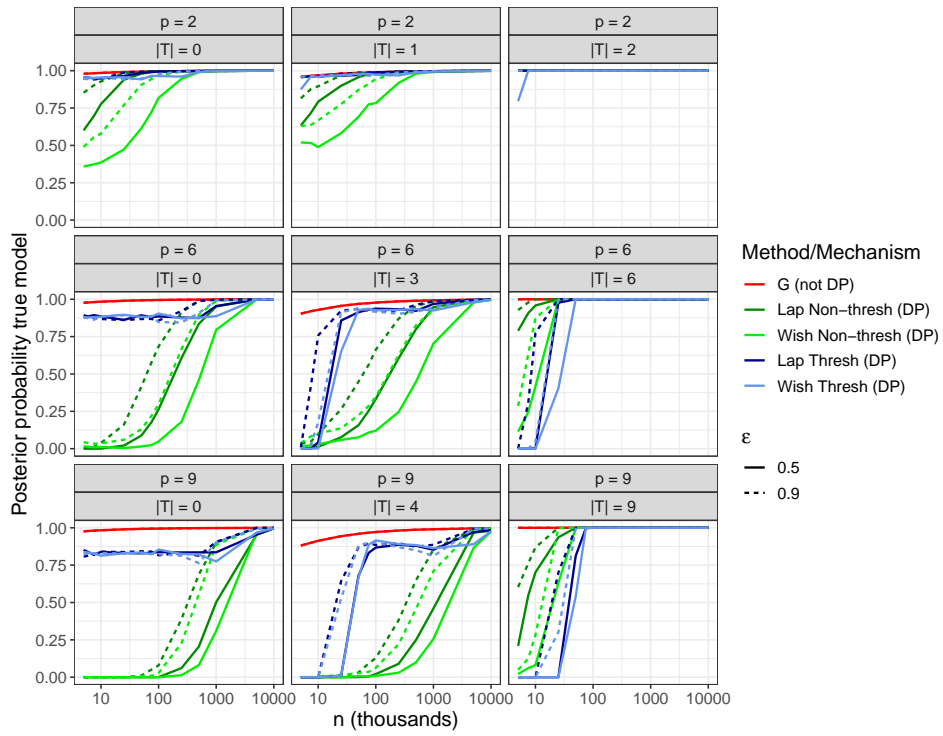


Figure 9: Simulation study: Sample size (x -axis) against posterior probability of the true model (y -axis) with BIC prior.

References

- John M Abowd. The US census bureau adopts differential privacy. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2867–2867, 2018.
- Milton Abramowitz, Irene A Stegun, et al. *Handbook of mathematical functions*, volume 55. Dover New York, 1964.
- Hirotsugu Akaike. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723, 1974.
- Gilad Amitai and Jerome Reiter. Differentially private posterior summaries for linear regression coefficients. *Journal of Privacy and Confidentiality*, 8(1), 2018.
- Borja Balle and Yu-Xiang Wang. Improving the Gaussian mechanism for differential privacy: Analytical calibration and optimal denoising. In *International Conference on Machine Learning*, pages 394–403. PMLR, 2018.
- 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.
- Maria J Bayarri, James O Berger, Anabel Forte, Gonzalo García-Donato, et al. Criteria for Bayesian model choice with application to variable selection. *The Annals of Statistics*, 40(3):1550–1577, 2012.
- James O Berger and Luis R Pericchi. The intrinsic Bayes factor for model selection and prediction. *Journal of the American Statistical Association*, 91(433):109–122, 1996.
- James O Berger and Luis R Pericchi. Objective Bayesian methods for model selection: Introduction and comparison. *Lecture Notes-Monograph Series*, pages 135–207, 2001.
- James O Berger, Luis R Pericchi, and Julia A Varshavsky. Bayes factors and marginal distributions in invariant situations. *Sankhyā: The Indian Journal of Statistics, Series A*, pages 307–321, 1998.
- James O Berger, Brunero Liseo, and Robert L Wolpert. Integrated likelihood methods for eliminating nuisance parameters. *Statistical science*, pages 1–22, 1999.
- Garrett Bernstein and Daniel R Sheldon. Differentially private Bayesian linear regression. *Advances in Neural Information Processing Systems*, 32:525–535, 2019.
- Peter J Bickel, Elizaveta Levina, et al. Covariance regularization by thresholding. *The Annals of Statistics*, 36(6):2577–2604, 2008.
- Lucien Birgé. An alternative point of view on Lepski’s method. *Lecture Notes-Monograph Series*, pages 113–133, 2001.

- Maria Eugenia Castellanos, Gonzalo Garcia-Donato, and Stefano Cabras. A model selection approach for variable selection with censored data. *Bayesian Analysis*, 16(1):271–300, 2021.
- Xueying Chen, Jerry Q Cheng, and Min-ge Xie. Divide-and-conquer methods for big data analysis. *arXiv preprint arXiv:2102.10771*, 2021.
- Gerda Claeskens, Nils Lid Hjort, et al. Model selection and model averaging. *Cambridge Books*, 2008.
- Merlise Clyde. *BAS: Bayesian Variable Selection and Model Averaging using Bayesian Adaptive Sampling*, 2020. R package version 1.5.5.
- Christian Covington, Xi He, James Honaker, and Gautam Kamath. Unbiased statistical estimation and valid confidence intervals under differential privacy. *arXiv preprint arXiv:2110.14465*, 2021.
- David M Diez, Christopher D Barr, and Mine Cetinkaya-Rundel. *OpenIntro statistics*. OpenIntro, 4 edition, 2012.
- David Draper. Assessment and propagation of model uncertainty. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1):45–70, 1995.
- Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pages 265–284. Springer, 2006.
- Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–407, 2014.
- Georgina Evans, Gary King, Margaret Schwenzfeier, and Abhradeep Thakurta. Statistically valid inferences from privacy protected data. *URL: GaryKing.org/dp*, 2020.
- 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.
- Simson L Garfinkel, John M Abowd, and Sarah Powazek. Issues encountered deploying differential privacy. In *Proceedings of the 2018 Workshop on Privacy in the Electronic Society*, pages 133–137, 2018.
- L. Grenié and G. Molteni. Inequalities for the beta function. *Mathematical Inequalities & Applications*, 18:1427–1442, 2015.
- Trevor Hastie, Robert Tibshirani, and Ryan J Tibshirani. Extended comparisons of best subset selection, forward stepwise selection, and the LASSO. *arXiv preprint arXiv:1707.08692*, 2017.
- Jonathan Huggins, Ryan P Adams, and Tamara Broderick. PASS-GLM: polynomial approximate sufficient statistics for scalable Bayesian GLM inference. *Advances in Neural Information Processing Systems*, 30, 2017.

- Harold Jeffreys. *The theory of probability*. Oxford University Press, 1939.
- Ariel Kleiner, Ameet Talwalkar, Purnamrita Sarkar, and Michael Jordan. The big data bootstrap. *arXiv preprint arXiv:1206.6415*, 2012.
- Tejas Kulkarni, Joonas Jälkö, Antti Koskela, Samuel Kaski, and Antti Honkela. Differentially private bayesian inference for generalized linear models. In *International Conference on Machine Learning*, pages 5838–5849. PMLR, 2021.
- Beatrice Laurent and Pascal Massart. Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics*, pages 1302–1338, 2000.
- Erich L Lehmann and Joseph P Romano. *Testing statistical hypotheses*. Springer-Verlag New York, 3 edition, 2005.
- Jing Lei, Anne-Sophie Charest, Aleksandra Slavkovic, Adam Smith, and Stephen Fienberg. Differentially private model selection with penalized and constrained likelihood. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 2018.
- Yingbo Li and Merlise A Clyde. Mixtures of g-priors in generalized linear models. *Journal of the American Statistical Association*, 113(524):1828–1845, 2018.
- Feng Liang, Rui Paulo, German Molina, Merlise A Clyde, and Jim O Berger. Mixtures of g-priors for Bayesian variable selection. *Journal of the American Statistical Association*, 103(481):410–423, 2008.
- Dennis V Lindley. A statistical paradox. *Biometrika*, 44(1/2):187–192, 1957.
- Frank McSherry and Ilya Mironov. Differentially private recommender systems: Building privacy into the netflix prize contenders. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 627–636, 2009.
- Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. Smooth sensitivity and sampling in private data analysis. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 75–84. ACM, 2007.
- Christian P Robert. On the Jeffreys-Lindley paradox. *Philosophy of Science*, 81(2):216–232, 2014.
- David Rossell and Francisco Javier Rubio. Additive Bayesian variable selection under censoring and misspecification. *arXiv preprint arXiv:1907.13563*, 2019.
- Mark Rudelson, Roman Vershynin, et al. Hanson-Wright inequality and sub-gaussian concentration. *Electronic Communications in Probability*, 18, 2013.
- Gideon Schwarz. Estimating the dimension of a model. *Annals of Statistics*, 6(2):461–464, 1978.
- James G Scott and James O Berger. Bayes and empirical-Bayes multiplicity adjustment in the variable-selection problem. *The Annals of Statistics*, pages 2587–2619, 2010.

- George AF Seber and Alan J Lee. *Linear regression analysis*. John Wiley & Sons, 2012.
- Or Sheffet. Old techniques in differentially private linear regression. In *Algorithmic Learning Theory*, pages 789–827. PMLR, 2019.
- Adam Smith. Privacy-preserving statistical estimation with optimal convergence rates. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 813–822, 2011.
- Duy Vu and Aleksandra Slavkovic. Differential privacy for clinical trial data: Preliminary evaluations. In *2009 IEEE International Conference on Data Mining Workshops*, pages 138–143. IEEE, 2009.
- Arnold Zellner and Aloysius Siow. Posterior odds ratios for selected regression hypotheses. *Trabajos de estadística y de investigación operativa*, 31(1):585–603, 1980.