Published in final edited form as:

Comput Stat. 2017 September ; 32(3): 867–888. doi:10.1007/s00180-017-0710-x.

Frequentist Standard Errors of Bayes Estimators

DongHyuk Lee¹, Raymond J. Carroll^{1,2}, and Samiran Sinha¹

¹Department of Statistics, Texas A&M University, 3143 TAMU, College Station, TX 77843-3143, USA

²School of Mathematical and Physical Sciences, University of Technology Sydney, Broadway NSW 2007, Australia

Summary

Frequentist standard errors are a measure of uncertainty of an estimator, and the basis for statistical inferences. Frequentist standard errors can also be derived for Bayes estimators. However, except in special cases, the computation of the standard error of Bayesian estimators requires bootstrapping, which in combination with Markov chain Monte Carlo (MCMC) can be highly time consuming. We discuss an alternative approach for computing frequentist standard errors of Bayesian estimators, including importance sampling. Through several numerical examples we show that our approach can be much more computationally efficient than the standard bootstrap.

Keywords

Bootstrap; Importance sampling; Markov chain; Posterior distribution; Standard error; Tail probability

1 Introduction

Suppose that $f(\cdot|\theta)$ is the data generating density and $\pi(\theta)$ is the prior distribution for the parameter θ . Let D be the observed data. Then the posterior distribution of θ is

$$\pi(\theta|D) = K_{\pi} f(D|\theta) \pi(\theta),$$

where K_{π} denotes the normalizing constant. There are several posterior summaries, such as the mean, $m(D) = E(\theta|D) = \int \theta \pi(\theta|D) d\theta$, the posterior median $\tilde{m}(D)$, which satisfies $\int_{-\infty}^{\tilde{m}(D)} \pi(\theta|D) d\theta = 0.5$, the α^{th} quantile $q_{\alpha}(D)$, that satisfies $\int_{-\infty}^{q_{\alpha}(D)} \pi(\theta|D) d\theta = \alpha$ for any $\alpha \in (0, 1)$, and the posterior mode $m_{\phi}(D) = \arg \max_{\theta} \pi(\theta|D)$. Suppose that by $s(D)$ we refer to any summary of the posterior distribution. Throughout this article we assume that D consists of $(X_1, \dots, X_n)$ independently and identically distributed observations. The goal of this paper is to discuss approaches of computing the frequentist standard error of $s(D)$.

Under a large sample, the observed data dominates the prior information in a Bayesian framework, and under standard regularity conditions, the posterior distribution of finite dimensional model parameters converges to the Gaussian distribution with the maximum likelihood estimator and the inverse of the Fisher Information matrix as the asymptotic mean and asymptotic variance, respectively. This asymptotic connection indicates that the Bayesian philosophy of integrating the observed data and the prior knowledge can be seen as a general procedure that encompasses the frequentist procedure as a special case. Therefore, as pointed out by a referee, frequentist standard error of a Bayes estimator is a way of assessing uncertainty of the general procedure. Particularly, for a large sample, the frequentist variance of the posterior mean converges to the inverse of the Fisher's information matrix. From the Bayesian perspective, frequentist standard errors can be used for comparing uncertainty of estimators under different priors (Efron, 2015). Although the posterior standard error (the standard deviation of the posterior distribution) is a measure of uncertainty of the posterior distribution, Efron (2015) argued that in a Bayesian paradigm, accuracy of a Bayes estimator, such as posterior mean could be judged based on the posterior distribution given that the prior distribution of the parameter reflects the truth in some degree. Therefore, finding accuracy of a Bayes estimator in an objective way among different subjective and objective priors is important. In fact, Berger (2006) discussed that the “pseudo-Bayes procedures” where subjective, objective, or a mixture of subjective and objective priors are used, often fail to provide any guidance on the performance of true subjective or objective Bayesian analysis. He then pointed out the necessity of validating these Bayesian approaches, and frequentist standard error of a posterior summary can be seen as a measure of such validation. Although Bayes factor is a way of comparing Bayesian procedures, many practitioners still want to compare estimators based on a frequentist uncertainty measure. Therefore, despite an apparent lack of coherence for incorporating a frequentist comparisons among Bayes procedures, it provides a measure of comparing uncertainties of the estimators. Of course, we should not use this measure solely to elicit the optimal prior for a Bayes procedure as for a proper comparison one should consider consistency, posterior convergence rate, along with the uncertainty of the estimator.

Efron (2015) proposed methods for computing frequentist standard errors of the posterior mean of a function of a parameter. In particular, he derived the approximate frequentist standard deviation of the posterior mean of a parameter based on the delta method. Suppose that T is the sufficient statistic. Following our notations, his formula for the approximate standard deviation of $\hat{t} = E\{t(\theta) | D\} = E\{t(\theta) | T\}$, the posterior mean of $\kappa(\theta)$, a function of θ , is $\left[\text{cov}\{t(\theta), \alpha_T(\theta) | T\}^T V_\theta \text{cov}\{t(\theta), \alpha_T(\theta) | T\} \right]^{1/2}$, where $\alpha_T(\theta) = \partial \log \{f_\theta(T)\} / \partial T$ denotes the gradient of $\log \{f_\theta(T)\}$ with respect to T , the sufficient statistic for θ , $f_\theta(T)$ is the density for the sufficient statistic T , and V_θ denotes the variance of the sufficient statistic. For application of this method it is critical that V_θ is readily available. Secondly, one key component of the delta method is the gradient of $\hat{t}$ with respect to T , and here this gradient is expressed as the posterior covariance $\text{cov}\{\kappa(\theta), \alpha_T(\theta) | T\}$. Expressing $\partial \hat{t} / \partial T$ as $\text{cov}\{\kappa(\theta), \alpha_T(\theta) | T\}$ critically relies on the fact that $\hat{t}$ is a posterior expectation. This posterior covariance is easy to estimate from a sample from the posterior distribution of θ . Therefore,

when V_{θ} is available and $\hat{\tau}$ is a posterior expectation, then Efron's formula is easy to apply and it is computationally fast.

In a special case with the exponential family of distributions where θ is considered to be the natural or canonical parameter vector, along with an uninformative prior for θ , he showed that the standard error of the posterior mean of $\kappa(\theta) = \theta$ can be computed without running the MCMC step to generate posterior samples for computing $\text{cov}\{\theta, \alpha_T(\theta)|T\}$. In lieu of the MCMC sampling, he used a parametric bootstrap resampling technique (Efron, 2012) to compute the posterior covariance term. Although the proposed method is applicable to only posterior means and when V_{θ} is easily available, the main advantage is that this method, when it is applicable, is much faster than the regular bootstrap procedure.

Inspired by this work we propose a general method of efficiently computing the frequentist standard error not only of the posterior mean but also of any posterior summary, $s(D)$. Our method is applicable for data generated from any parametric model, not necessarily from an exponential family of distributions. The proposed method relies on the bootstrap idea. Usually, the standard error of an estimator can be computed by the bootstrap method (Efron and Tibshirani, 1986), where the standard error is estimated by the standard deviation of the Bayes estimators obtained from a large number of bootstrap samples. On the other hand, the Bayes estimator for a bootstrap sample is usually calculated by drawing a large number of Markov chain Monte Carlo (MCMC) samples, which is often time consuming, and consequently drawing posterior samples for each of the bootstrap data can be a prohibitively time consuming task.

The main aim of this article is to reduce this computation time. To do so, the MCMC method will be used once to draw samples from the posterior distribution of the parameters given the original data. Then use these posterior samples along with the importance sampling idea to compute the posterior summary for each bootstrap data. The details are discussed in the following sections. To make it clearer, we want to re-state that in the proposed method, we do need bootstrap sampling, but we bypass the MCMC sampling for each bootstrap data by a clever use of the importance sampling method. Here is a brief description of the importance sampling method in a few words. Suppose that we are interested in estimating $\theta = \int g(x) f(x) dx$, where $f(x)$ is a density. With another density $h(x)$, we can re-write $\theta = \int g(x) \omega(x) h(x) dx$, where $\omega(x) = f(x)/h(x)$ is called the importance weight. Then

the importance sampling estimator of θ is $\hat{\theta} = m^{-1} \sum_{i=1}^m g(x_i) \omega(x_i)$, where $x_1, \dots, x_m$ are iid from $h(x)$. This technique is quite useful for efficient estimation of tail probabilities, and is used for drawing bootstrap samples, specially for estimating standard error of small probabilities (Efron and Tibshirani, 1994, pp. 349). However, in this paper we are using importance sampling technique to compute estimators based on a bootstrap re-sampled data. Basically in the proposed approach bootstrap samples are drawn using standard bootstrap resampling technique and then importance sampling is used to compute the Bayes estimators. Although importance sampling idea has been used in many other contexts, including but not limited to the simulated maximum likelihood estimation, computer graphics, modelling stock market data, modelling linear and nonlinear dynamic processes

(Liang, 2002), to the best of our knowledge, the use of this technique in the present context seems to be the new.

A brief outline of the remainder of the manuscript is as follows. In Section 2 we provide the background information. The main idea related to the posterior mean is discussed in Section 3, while Section 4 considers posterior quantiles and the posterior mode. Section 5 describes the results of two simulation studies and a real data examples. Section 6 contains conclusions.

2 Background

To motivate this research first we consider three commonly used models.

Logistic regression model

Suppose that $Y_1, \dots, Y_n$ are independently drawn from the Bernoulli(p_i) distribution, where $p_i = \text{pr}(Y_i = 1|X_i) = \{1 + \exp(-\alpha - \beta X_i)\}^{-1}$ with a scalar covariate X_i . Assume priors $\alpha \sim \text{Normal}(a, \sigma^2)$ and $\beta \sim \text{Normal}(b, \tau^2)$, and let D denote the observed data $\{(X_i, Y_i), i = 1, \dots, n\}$. Then the posterior distribution of α and β is

$$\begin{aligned} \pi(\alpha, \beta|D) &\propto \prod_{i=1}^n \left\{ \frac{1}{1 + \exp(-\alpha - \beta X_i)} \right\}^{Y_i} \left\{ \frac{\exp(-\alpha - \beta X_i)}{1 + \exp(-\alpha - \beta X_i)} \right\}^{(1-Y_i)} \\ &\times \frac{\exp\left\{-\frac{(\alpha - a)^2}{2\sigma^2}\right\}}{\sqrt{2\pi\sigma^2}} \\ &\times \frac{\exp\left\{-\frac{(\beta - b)^2}{2\tau^2}\right\}}{\sqrt{2\pi\tau^2}}. \end{aligned}$$

For computing any posterior summary for $\pi(\alpha, \beta|D)$, usually we draw posterior samples from $\pi(\alpha, \beta|D)$ using the MCMC method. So, a numerical method is must for computing frequentist standard errors of any summary of the posterior distribution. In the simulation section, for illustration, we apply the proposed method on this model.

Linear measurement error model

Now, we consider the following simple linear regression problem, where using the observed data $D = \{(Y_i, W_i), i = 1, \dots, n\}$, we want to fit $Y_i = \alpha + X_i\beta + \epsilon_i$, where X_i is unobserved but we observed its surrogate variable W_i , and $\epsilon_i \sim \text{Normal}(0, \sigma_\epsilon^2)$. The observed surrogate W_i is associated with the true X_i through the classical additive measurement error model $W_i = X_i + U_i$, where $U_i \sim \text{Normal}(0, \sigma_u^2)$ and σ_u^2 is considered to be known for simplicity. We further assume that measurement error is nondifferential such that Y_i is conditionally independent of W_i given the true X_i (Carroll et al., 2006, pp. 36), and $X_i \sim \text{Normal}(\mu_x, \sigma_x^2)$. It is well-known that the simple linear regression of Y on W will cause an attenuation towards 0 by the multiplicative factor $\sigma_x^2 / (\sigma_x^2 + \sigma_u^2)$. One of the corrections for attenuation is

the method of moments. That is, the resulting estimator $\hat{\beta} = \hat{\beta}_w \hat{\sigma}_w^2 / (\hat{\sigma}_w^2 - \sigma_u^2)$, where $\hat{\beta}_w$ is the OLS estimator ignoring measurement error, $\hat{\sigma}_w^2$ is the sample variance of the observed W , and σ_u^2 is the variance of U (Carroll et al., 2006, Section 3.4.1; Fuller, 1987, Section 2.5). In addition, it is well-known that $\hat{\beta}$ has no finite moments, because the denominator term $\hat{\sigma}_w^2 - \sigma_u^2$ can get arbitrarily close to zero (Fuller, 1987). Therefore, Bayesian calculations are an attractive alternative.

We attempt to use a Bayesian inference for the parameters $\theta = (\alpha, \beta, \mu_x, \sigma_x^2, \sigma_\epsilon^2)$ in which α and β are the main parameters of interest. Assigning normal priors, $\text{Normal}(0, \sigma_\alpha^2)$, $\text{Normal}(0, \sigma_\beta^2)$, $\text{Normal}(0, \sigma_\mu^2)$ for α, β, μ_x , respectively and inverse gamma priors $\text{IG}(\delta_x, \lambda_x)$, $\text{IG}(\delta_\epsilon, \lambda_\epsilon)$ for $\sigma_x^2, \sigma_\epsilon^2$, respectively (Carroll et al., 2006, Section 9.4), the joint posterior distribution of θ and the latent variable $X = (X_1, \dots, X_n)$ is

$$\pi(\theta, X|D) \propto (\sigma_\epsilon^2)^{-n/2-\delta_\epsilon-1} (\sigma_x^2)^{-n/2-\delta_x-1} \exp \left\{ -\frac{\sum_{i=1}^n (Y_i - \alpha - X_i \beta)^2 / 2 + \lambda_\epsilon}{\sigma_\epsilon^2} - \frac{\sum_{i=1}^n (W_i - X_i)^2 / 2 + \lambda_u}{\sigma_u^2} - \frac{\sum_{i=1}^n (X_i - \mu_x)^2 / 2 + \lambda_x}{\sigma_x^2} - \frac{\alpha^2}{\sigma_\alpha^2} - \frac{\beta^2}{\sigma_\beta^2} - \frac{\mu_x^2}{\sigma_\mu^2} \right\}.$$

Due to the conjugacy of the prior distributions, it is easy to apply the Gibbs sampler to draw posterior samples from $\pi(\theta, X|D)$. Specifically, the conditional posterior distributions of α and β given other parameters and the latent variable X are normal distributions so that we can easily obtain their posterior summaries. However, it is not an easy problem to find the variances of their posterior summaries mainly because they are dependent on the unobserved X . Thus a numerical method is required.

Weibull regression model

Suppose that $T_1, \dots, T_n$ are independently drawn from the Weibull(α, λ_i) distribution whose density is $g(t|\alpha, \lambda) = \alpha t^{\alpha-1} \exp\{\lambda - \exp(\lambda) t^\alpha\}$ (Ibrahim et al., 2001, eq. 2.2.1). Let $C_1, \dots, C_n$ be the corresponding censoring times whose distribution does not include any information about parameters α and λ_i (non-informative censoring) and $i_1, \dots, i_n$ be the censoring indicator where $i_i = 1$ if $T_i \leq C_i$ (observed) and $i_i = 0$ if $T_i > C_i$ (censored). In this example, let $D = \{Y_i, X_i, i = 1, \dots, n\}$, where $Y_i = \min(T_i, C_i)$, and X_i is the covariate for the i^{th} individual. We regress the parameter λ_i on covariates X_i i.e., $\lambda_i = X_i' \beta$. Assigning a normal prior, $\text{Normal}(\mu_0, \Sigma_0)$, for β and a gamma prior, $\text{Gamma}(\alpha_0, \kappa_0)$, for α , the posterior distribution of α and β is

$$\pi(\alpha, \beta|D) \propto \alpha^{\alpha_0+d-1} \exp \left[\sum_{i=1}^n \left\{ \Delta_i X_i' \beta + \Delta_i (\alpha - 1) \log(Y_i) - Y_i^\alpha \exp(X_i' \beta) \right\} - \kappa_0 \alpha - \frac{1}{2} (\beta - \mu_0)' \Sigma^{-1} (\beta - \mu_0) \right],$$

where $d = \sum_{i=1}^n \Delta_i$ (Ibrahim et al., 2001, eq. 2.2.4). Likewise in the logistic regression example, we need not only to draw posterior samples from $\pi(\alpha, \beta|D)$ using the MCMC method to evaluate any posterior summary, we also necessitate a numerical procedures for computing frequentist standard errors of those posterior summaries. We use this model to analyze the Melanoma data set in Section 5.3, and compute uncertainty measures using the proposed approach.

These examples show that even for these well researched models, posterior summaries may not have an explicit expression that is easy to compute. Additionally, the computation of the standard error of the posterior summaries requires extra numerical work.

3 Standard errors of posterior means

In this section we concentrate only on the posterior mean and its standard error calculations. In Section 4, we provide recipes for efficiently calculating frequentist standard errors of other types of Bayes estimators. For any generic vector a , we shall use $a^{\otimes 2}$ to denote aa^T .

The frequentist standard error of the posterior mean of θ , $\hat{\theta} = m^*(D) = E(\theta|D) = E_{\pi(\cdot|D)}(\theta)$, where $\pi(\cdot|D)$ is the posterior distribution, is

$$\sqrt{\text{var}_F(\hat{\theta})} = \sqrt{\int \{m^*(D)\}^{\otimes 2} dF(D|\theta) - \{\int m^*(D) dF(D|\theta)\}^{\otimes 2}}.$$

Suppose that one draws B random samples each of size m from the posterior distribution $\pi(\theta|D)$. Denote the b^{th} sample as $(\theta_{b1}, \dots, \theta_{bm})$, $b = 1, \dots, B$. Define $\hat{\theta}_b = \sum_{j=1}^m \theta_{bj}/m$, and $\bar{\theta} = \sum_{b=1}^B \hat{\theta}_b/B$. It is obvious that the variance among $\hat{\theta}_1, \dots, \hat{\theta}_B$ does not estimate $\text{var}_F(\hat{\theta})$ as $(B-1)^{-1} \sum_{b=1}^B (\hat{\theta}_b - \bar{\theta})^2 \rightarrow (1/m) \text{var}_{\pi(\cdot|D)}(\theta)$ almost surely as $B \rightarrow \infty$, where $\text{var}_{\pi(\cdot|D)}(\theta)$ denotes the posterior variance of θ . One obvious approach to estimate $\text{var}_F(\hat{\theta})$ is to adopt the bootstrap idea. In the bootstrap world, instead of $\text{var}_F(\hat{\theta})$ we target estimating $\text{var}_{\hat{F}}(\hat{\theta})$, where the observed data are treated as the entire population. In the bootstrap method, we draw B bootstrap samples with replacement from the original data, calculate the posterior mean for each bootstrap sample, and then take the variance of the B posterior means. Let $D^{(b)}$ be the b^{th} bootstrap data, and $\pi(\theta|D^{(b)})$ be the corresponding posterior distribution. Define $\hat{\theta}^{(b)} = E(\theta|D^{(b)}) = E_{\pi(\cdot|D^{(b)})}(\theta)$ as the posterior mean of θ for the b^{th} bootstrap data. Further define $\bar{\theta}^{(\cdot)} = \sum_{b=1}^B \hat{\theta}^{(b)}/B$. Then

$$(B-1)^{-1} \sum_{b=1}^B \left(\hat{\theta}^{(b)} - \bar{\theta}^{(\cdot)} \right)^2 \rightarrow \text{var}_{\hat{F}}(\hat{\theta}) \quad \text{as } B \rightarrow \infty.$$

In practice, $\hat{\theta}^{(b)}$ is estimated by the Monte Carlo estimator $\hat{\theta}_{mc}^{(b)} = \sum_{j=1}^M \theta_j^{(b)} / M$, where $\theta_1^{(b)}, \dots, \theta_M^{(b)}$ are M random draws from $\pi(\theta|D^{(b)})$, and $\hat{\theta}_{mc}^{(b)} \rightarrow \hat{\theta}^{(b)}$ almost surely as $M \rightarrow \infty$.

Also, define $\bar{\theta}_{mc}^{(\cdot)} = B^{-1} \sum_{b=1}^B \hat{\theta}_{mc}^{(b)}$. Then as $M \rightarrow \infty$,

$$(B-1)^{-1} \sum_{b=1}^B \left(\hat{\theta}_{mc}^{(b)} - \bar{\theta}_{mc}^{(\cdot)} \right)^2 \rightarrow (B-1)^{-1} \sum_{b=1}^B \left(\hat{\theta}^{(b)} - \bar{\theta}^{(\cdot)} \right)^2.$$

Hence $\sum_{b=1}^B \left(\hat{\theta}_{mc}^{(b)} - \bar{\theta}_{mc}^{(\cdot)} \right)^2 / (B-1)$ will be used as the estimator of $\text{var}_{\hat{\theta}}(\hat{\theta})$. In the following paragraph we describe how we estimate $\hat{\theta}^{(1)}, \dots, \hat{\theta}^{(B)}$ without having numerically computing B posterior distributions using B MCMC chains thereby saving lots of computation time.

Suppose that using MCMC method we have drawn $\theta_1, \dots, \theta_M$ from $\pi(\theta|D)$, the posterior distribution of θ given the entire data D . Suppose that in the b^{th} bootstrap sample, X_i occurs $r_i^{(b)}$ times, where $0 \leq r_i^{(b)} \leq n$, but $\sum_{i=1}^n r_i^{(b)} = n$. Then the posterior distribution of θ given the b^{th} bootstrap data $D^{(b)}$ is

$$\pi(\theta|D^{(b)}) = \frac{\prod_{i=1}^n f^{r_i^{(b)}}(X_i|\theta) \pi(\theta)}{\int \prod_{i=1}^n f^{r_i^{(b)}}(X_i|\theta) \pi(\theta) d\theta},$$

so

$$\hat{\theta}^{(b)} = \int \theta \pi(\theta|D^{(b)}) d\theta = \frac{\int \theta \prod_{i=1}^n f^{r_i^{(b)}}(X_i|\theta) \pi(\theta) d\theta}{\int \prod_{i=1}^n f^{r_i^{(b)}}(X_i|\theta) \pi(\theta) d\theta} = \frac{G_1^{(b)}}{G_0^{(b)}},$$

where $G_s^{(b)} = \int \theta^s \prod_{i=1}^n f^{r_i^{(b)}}(X_i|\theta) \pi(\theta) d\theta$ for $s = 0$ and 1 . Next, we can re-write

$$\begin{aligned} G_s^{(b)} &= \frac{1}{K_\pi} \int \theta^s \frac{\prod_{i=1}^n f^{r_i^{(b)}}(X_i|\theta) \pi(\theta)}{\prod_{i=1}^n f(X_i|\theta) \pi(\theta)} K_\pi \prod_{i=1}^n f(X_i|\theta) \pi(\theta) d\theta \\ &= \frac{1}{K_\pi} \int \theta^s \omega^{(b)}(\theta) K_\pi \prod_{i=1}^n f(X_i|\theta) \pi(\theta) d\theta, \end{aligned}$$

where the importance weight

$\omega^{(b)}(\theta) = \prod_{i=1}^n f^{r_i^{(b)}}(X_i|\theta) / \prod_{i=1}^n f(X_i|\theta) = \prod_{i=1}^n f^{(r_i^{(b)}-1)}(X_i|\theta)$. Hence $\hat{\theta}^{(b)}$ can be estimated by

$$\hat{\theta}_{is}^{(b)} = \frac{\sum_{j=1}^M \theta_j \omega^{(b)}(\theta_j)}{\sum_{j=1}^M \omega^{(b)}(\theta_j)}, \quad (1)$$

where $\theta_1, \dots, \theta_M$ are MMCMC samples drawn from $\pi(\theta|D)$, the posterior distribution of θ given the original data D . Importantly, under regularity conditions, $\hat{\theta}_{is}^{(b)} \rightarrow \hat{\theta}^{(b)}$ almost surely as $M \rightarrow \infty$ (see the Appendix). Next define $\bar{\theta}_{is}^{(\cdot)} = B^{-1} \sum_{b=1}^B \hat{\theta}_{is}^{(b)}$. As M gets large,

$$(B-1)^{-1} \sum_{b=1}^B \left(\hat{\theta}_{is}^{(b)} - \bar{\theta}_{is}^{(\cdot)} \right)^2 \rightarrow (B-1)^{-1} \sum_{b=1}^B \left(\hat{\theta}^{(b)} - \bar{\theta}^{(\cdot)} \right)^2.$$

Hence we use $\sum_{b=1}^B \left(\hat{\theta}_{is}^{(b)} - \bar{\theta}_{is}^{(\cdot)} \right)^2 / (B-1)$ to estimate $\text{var}_{\hat{F}}(\hat{\theta})$. The above procedure can be summarized in the following steps.

Step 1

Draw MMCMC samples from $\pi(\theta|D)$, and call them $(\theta_1, \dots, \theta_M)$.

Step 2

Draw B bootstrap samples with replacement from D , and each bootstrap sample consists of n observations. For the b^{th} sample we obtain $(r_1^{(b)}, \dots, r_n^{(b)})$, with $0 \leq r_i^{(b)} \leq n$ and $\sum_{i=1}^n r_i^{(b)} = n$, where $r_i^{(b)}$ is the number of times X_i appears in the b^{th} bootstrap sample, $b = 1, \dots, B$.

Step 3

Compute $\hat{\theta}_{is}^{(b)} = \sum_{j=1}^M \theta_j \omega^{(b)}(\theta_j) / \sum_{j=1}^M \omega^{(b)}(\theta_j)$ with $\omega^{(b)}(\theta_j) = \prod_{i=1}^n f^{(r_i^{(b)}-1)}(X_i|\theta_j)$ for $b = 1, \dots, B$, and $\bar{\theta}_{is}^{(\cdot)} = \sum_{b=1}^B \hat{\theta}_{is}^{(b)} / B$.

Step 4

Compute $(B-1)^{-1} \sum_{b=1}^B \left(\hat{\theta}_{is}^{(b)} - \bar{\theta}_{is}^{(\cdot)} \right)^2$.

One of the main concerns of importance sampling is the behavior of the importance weights that have influence on the efficiency of the estimator. The following remark gives an intuitive justification that our choice $\pi(\theta|D)$ as the trial distribution provides a bounded importance weight with high probability.

Remark 1

Note that $\omega^{(b)}(\theta) = \exp \left[\sum_{i=1}^n \left(r_i^{(b)} - 1 \right) \log \{f(X_i|\theta)\} \right] = \exp \left\{ \ell^{(b)}(\theta) - \ell(\theta) \right\}$, where $\ell^{(b)}(\theta) = \sum_{i=1}^n r_i^{(b)} \log f(X_i|\theta) + \log \{\pi(\theta)\}$ and $\ell(\theta) = \sum_{i=1}^n \log \{f(X_i|\theta)\} + \log \{\pi(\theta)\}$, and θ is drawn from the posterior distribution $\pi(\theta|D)$. Now,

$$\ell^{(b)}(\theta) - \ell(\tilde{\theta}) \leq \ell^{(b)}(\theta) - \ell(\theta) \leq \ell^{(b)}(\tilde{\theta}_{(b)}) - \ell(\theta),$$

where $\tilde{\theta}_{(b)}$ is the posterior mode based on the b^{th} bootstrap data set and $\tilde{\theta}$ is the posterior mode based on the original data. Then under certain regularity conditions, posterior distribution $\pi(\theta|D)$ has the asymptotic normal distribution having mean $\tilde{\theta}$ and the variance is minus the inverse Hessian of the log posterior evaluated at $\tilde{\theta}$ for large n (Theorem 3.1 of Carlin and Louis, 2008).

4 Other Bayes estimators**4.1 Posterior quantile**

Here we broadly discuss the standard error calculation of posterior quantiles that include the posterior median and credible intervals as special cases. The α^{th} quantile is defined as

$q_\alpha(D) = F_{\pi(\theta|D)}^{-1}(\alpha)$, where $F_{\pi(\theta|D)}(r) = \int_{-\infty}^r \pi(\theta|D) d\theta$. To estimate the frequentist standard error of $q_\alpha(D)$, we may apply the regular bootstrap method by calculating the α^{th} quantile for each of the B posterior distributions, that means one needs to draw posterior samples from $\pi(\theta|D^{(b)})$ using MCMC technique for each $b = 1, \dots, B$. Instead of doing this for multiple bootstrap data sets, here we can also apply the importance sampling idea. For a trial density $h(\theta)$, we have

$$F_{\pi(\theta|D^{(b)})}(r) = \int_{-\infty}^{\infty} I(\theta \leq r) \pi(\theta|D^{(b)}) d\theta = \int_{-\infty}^{\infty} I(\theta \leq r) \frac{\pi(\theta|D^{(b)})}{h(\theta)} h(\theta) d\theta = \int_{-\infty}^{\infty} I(\theta \leq r) \omega^{(b)}(\theta) h(\theta) d\theta,$$

where $\omega^{(b)}(\theta) = \pi(\theta|D^{(b)})/h(\theta)$. The distribution function can be estimated by

$$\hat{F}_{\pi(\theta|D^{(b)})}(r) = \frac{\sum_{j=1}^M I(\theta_j \leq r) \omega^{(b)}(\theta_j)}{\sum_{j=1}^M \omega^{(b)}(\theta_j)}, \quad (2)$$

where $\theta_1, \dots, \theta_M$ are drawn from $h(\theta)$. We shall evaluate $\hat{F}_{\pi(\theta|D^{(b)})}(r)$ for a grid of values of r . Next, the estimated α^{th} quantile is defined as $a_{\alpha, is}^{(b)} = \inf \left\{ r : \hat{F}_{\pi(\theta|D^{(b)})}(r) \geq \alpha \right\}$. Note that we shall use the same set of $\theta_1, \dots, \theta_M$ drawn from $h(\theta)$, for each bootstrap data set thereby saving considerable computation time.

When α takes a moderate value in the range of 0.2 to 0.8, the importance sampling estimates are reasonable if $\pi(\theta|D)$ is used as the trial distribution. For more extreme values of α , (smaller than 0.2 or larger than 0.8), we recommend the following trial distribution for efficient estimation of the α^{th} quantile. To be more specific, without any loss of generality,

write $\theta = (\theta_1, \theta_2^T)^T$, and suppose that we are interested in estimating the α^{th} quantile of θ_1 based on the b^{th} bootstrap data. Take $h(\theta) = h_1(\theta_1)h_2(\theta_2)$, where h_1 denotes the uniform density over $[l, u]$ for given values of l and u , and h_2 is taken as the posterior distribution of θ_2 given the data D , that means, $h_2(\theta_2) = \int \pi(\theta|D) d\theta_1$. Although there is no optimum choice of l or u , based on our computing experiences, we recommend $l = q_{0.5}(D) - 6 \times sd_{\theta_1}(D)$ and $u = q_{0.5}(D) + 6 \times sd_{\theta_1}(D)$, where $q_{\alpha}(D)$ and $sd_{\theta_1}(D)$ denote the α^{th} quantile and the posterior standard deviation of θ_1 given the entire data D .

Suppose that $(\theta_{11}, \dots, \theta_{1M})$ are M random draws from $h_1(\theta_1)$, and $(\theta_{21}, \dots, \theta_{2M})$ are M random draws from $\pi(\theta_2|D)$. The later sample is obtained by simply discarding the first component from each of the $MMCMC$ samples drawn from $\pi(\theta|D) \equiv \pi(\theta_1, \theta_2|D)$.

Computation of the importance weight $\omega^{(b)}(\theta)$ at $\theta = \theta_j = (\theta_{1j}, \theta_{2j}^T)^T$, for any $j = 1, \dots, M$, requires $h_2(\theta_{2j}) = \int \pi(\theta_1^*, \theta_{2j}^T|D) d\theta_1^* = \kappa^{-1} \int \prod_{i=1}^n f(X_i|\theta_1^*, \theta_{2j}^T) \pi(\theta_1^*, \theta_{2j}^T) d\theta_1^*$, where κ is the normalizing constant that does not depend on θ_j . In order to save computation time, instead of targeting to evaluate $h_2(\theta_{2j})$ separately, we consider directly evaluating $\omega^{(b)}(\theta_j)$, and

$$\begin{aligned} \omega^{(b)}(\theta_j) &= \frac{\pi(\theta_{1j}, \theta_{2j}^T|D^{(b)})}{h_1(\theta_{1j}) \int_{\theta_{1,\min}-\varepsilon}^{\theta_{1,\max}+\varepsilon} \pi(\theta_1^*, \theta_{2j}^T|D) d\theta_1^*} = \frac{\kappa_b^{-1} \prod_{i=1}^n f^{(b)}(X_i|\theta_{1j}, \theta_{2j}^T) \pi(\theta_{1j}, \theta_{2j}^T)}{h_1(\theta_{1j}) \kappa^{-1} \int_{\theta_{1,\min}-\varepsilon}^{\theta_{1,\max}+\varepsilon} \prod_{i=1}^n f(X_i|\theta_1^*, \theta_{2j}^T) \pi(\theta_1^*, \theta_{2j}^T) d\theta_1^*} \\ &= \left[h_1(\theta_{1j}) \frac{\kappa_b^{-1}}{\kappa^{-1}} \int_{\theta_{1,\min}-\varepsilon}^{\theta_{1,\max}+\varepsilon} \left\{ \frac{f(X_i|\theta_1^*, \theta_{2j}^T)}{f^{(b)}(X_i|\theta_{1j}, \theta_{2j}^T)} \right\} \left\{ \frac{\pi(\theta_1^*, \theta_{2j}^T)}{\pi(\theta_{1j}, \theta_{2j}^T)} \right\} d\theta_1^* \right]^{-1}, \end{aligned} \quad (3)$$

where κ_b is the normalizing constant for the b^{th} bootstrap data $D^{(b)}$, and $\theta_{1,\min}$ and $\theta_{1,\max}$ denote the observed minimum and maximum values of θ_1 in the posterior samples drawn from $\pi(\theta_1, \theta_2|D)$. To cover the entire domain of θ_1 , we extend the range of the integration by adding and subtracting a small number $\varepsilon > 0$. In all our computations, we used $\varepsilon = 0.1 \times \text{IQR}$, where IQR stands for the inter quartile range of the posterior distribution of θ_1 given the original data D . Importantly, we do not need to evaluate κ and κ_b for estimating

$F_{\pi(\theta|D^{(b)})}^{(r)}$ as they are independent of θ_j ; so they get canceled from the normalized weight. Finally, we recommend to use Gauss-Legendre quadrature to determine the above integral in (3). Also to reduce the computational burden, once $\omega^{(b)}(\theta_j)$ is calculated some b , then we compute $\omega^{(b')}(\theta_j)$ using the following formula

$$\omega^{(b')}(\theta_j) = \omega^{(b)}(\theta_j) \pi \left(\theta_{1j}, \theta_{2j}^T | D^{(b')} \right) / \pi \left(\theta_{1j}, \theta_{2j}^T | D^{(b)} \right), \text{ for any } b' \neq b \text{ as}$$

$$\omega^{(b')}(\theta_j) = \frac{\pi(\theta_{1j}, \theta_{2j}^T | \mathbf{D}^{(b')})}{h_1(\theta_{1j}) \int_{\theta_{1,min}-\varepsilon}^{\theta_{1,max}+\varepsilon} \pi(\theta_1^*, \theta_{2j}^T | \mathbf{D}) d\theta_1^*} = \underbrace{\frac{\pi(\theta_{1j}, \theta_{2j}^T | \mathbf{D}^{(b)})}{h_1(\theta_{1j}) \int_{\theta_{1,min}-\varepsilon}^{\theta_{1,max}+\varepsilon} \pi(\theta_1^*, \theta_{2j}^T | \mathbf{D}) d\theta_1^*}}_{\omega^{(b)}(\theta_j)} \times \frac{\pi(\theta_{1j}, \theta_{2j}^T | \mathbf{D}^{(b')})}{\pi(\theta_{1j}, \theta_{2j}^T | \mathbf{D}^{(b)})}. \quad (4)$$

4.2 Posterior mode

Here we do not apply the importance sampling idea but use another approach for time efficient computation. The posterior mode is defined as $\hat{\theta}_{mode} = \arg \max_{\theta} \pi(\theta | \mathbf{D})$. The

variance of $\hat{\theta}_{mode}$, $var_F(\hat{\theta}_{mode})$ can be estimated by $\sum_{b=1}^B \left(\hat{\theta}_{mode}^{(b)} - \bar{\theta}_{mode}^{(\cdot)} \right)^2 / (B-1)$,

where $\hat{\theta}_{mode}^{(b)}$ denotes the posterior mode for the b^{th} bootstrap sample, and

$\bar{\theta}_{mode}^{(\cdot)} = \sum_{b=1}^B \hat{\theta}_{mode}^{(b)} / B$. Since this standard bootstrap method could be time consuming as it requires to solve a set of gradient equations for each of the B bootstrap data sets, we propose the following alternative approach of estimating that variance.

Under sufficient smoothness conditions, $\hat{\theta}_{mode}$ will satisfy $S(\hat{\theta}_{mode} | \mathbf{D}) = 0$, where $S(\theta | \mathbf{D}) = \partial \log \{ \pi(\theta | \mathbf{D}) \} / \partial \theta = \partial \log \{ f(\mathbf{D} | \theta) \} / \partial \theta + \partial \log \{ \pi(\theta) \} / \partial \theta = 0$. Suppose that as $n \rightarrow \infty$, $\hat{\theta}_{mode} \rightarrow \theta_{mode}$. Then

$$\begin{aligned} 0 &= S(\hat{\theta}_{mode} | \mathbf{D}) \\ &= \frac{\partial}{\partial \theta} \log \{ f(\mathbf{D} | \hat{\theta}_{mode}) \} \\ &\quad + \frac{\partial}{\partial \theta} \log \{ \pi(\hat{\theta}_{mode}) \} \approx \left[\frac{\partial}{\partial \theta} \log \{ f(\mathbf{D} | \theta_{mode}) \} + \frac{\partial}{\partial \theta} \log \{ \pi(\theta_{mode}) \} \right] \\ &\quad + \left[\frac{\partial^2}{\partial \theta^2} \log \{ f(\mathbf{D} | \theta_{mode}) \} + \frac{\partial^2}{\partial \theta^2} \log \{ \pi(\theta_{mode}) \} \right] (\hat{\theta}_{mode} - \theta_{mode}). \end{aligned}$$

Thus, with $A = E \left[\frac{\partial^2}{\partial \theta^2} \log \{ f(\mathbf{D} | \theta_{mode}) \} + \frac{\partial^2}{\partial \theta^2} \log \{ \pi(\theta_{mode}) \} \right]$, we have

$(\hat{\theta}_{mode} - \theta_{mode}) \approx A^{-1} [\partial \log \{ f(\mathbf{D} | \theta_{mode}) \} / \partial \theta + \partial \log \{ \pi(\theta_{mode}) \} / \partial \theta]$, and consequently the variance can be obtained by the sandwich formula,

$$var_F(\hat{\theta}_{mode}) = A^{-1} var \left[\frac{\partial}{\partial \theta} \log \{ f(\mathbf{D} | \theta_{mode}) \} + \frac{\partial}{\partial \theta} \log \{ \pi(\theta_{mode}) \} \right] A^{-T}.$$

Here $\mathbf{A}$ can be estimated by $\hat{\mathbf{A}} = \partial^2 \log \{f(\mathbf{D}|\hat{\theta}_{mode})\} / \partial \theta^2 + \partial^2 \log \{\pi(\hat{\theta}_{mode})\} / \partial \theta^2$. The middle term of the variance formula is $var[\partial \log \{f(\mathbf{D}|\theta_{mode})\} / \partial \theta]$ that can be estimated by

$$\widehat{var} \left[\frac{\partial}{\partial \theta} \log \{f(\mathbf{D}|\theta_{mode})\} \right] = (B-1)^{-1} \sum_{b=1}^B \left[\frac{\partial}{\partial \theta} \log \{f(\mathbf{D}^{(b)}|\hat{\theta}_{mode})\} - \frac{1}{B} \sum_{b'=1}^B \frac{\partial}{\partial \theta} \log \left\{ f\left(\mathbf{D}^{(b')}|\hat{\theta}_{mode}\right) \right\} \right]^2,$$

and in particular, for fast computation we use

$\partial \log \{f(\mathbf{D}^{(b)}|\hat{\theta}_{mode})\} / \partial \theta = \sum_{i=1}^n r_i^{(b)} \partial \log \{f(X_i|\hat{\theta}_{mode})\} / \partial \theta$. Finally, $var_F(\hat{\theta}_{mode})$ is estimated by $\hat{\mathbf{A}}^{-1} \widehat{var}[\partial \log \{f(\mathbf{D}|\theta_{mode})\} / \partial \theta] \hat{\mathbf{A}}^{-T}$.

5 Numerical studies

In order to assess and compare the performances of the methods, we conducted simulation studies and real data analysis for the motivating examples described in Section 2. Specifically, we provide simulation results for the logistic regression model. Next, the linear measurement error model is illustrated using a simulated data set. Third, we present an analysis of real data set using the Weibull regression model. Finally, following a referee's request, we consider an application of the proposed method to a vector autoregressive (VAR) model.

5.1 Logistic regression model

We generated 500 data sets, and each simulated data set consists of $n = 500$ observations, denoted by $\{(X_i, Y_i), i = 1, \dots, n\}$. We drew X from Normal(0, 1) distribution and the response variable Y was simulated from a Bernoulli distribution with the success probability $\text{pr}(Y = 1|X) = \exp(a + \beta X) / \{1 + \exp(a + \beta X)\}$. The true values of a and β were -2.5 and 1 , respectively. That makes the proportion of success around 10%. For the Bayesian inference of the parameters a and β we used the same Normal(0, 2) priors for both of them. Then for the MCMC computation, we used 15,000 iterations with the first 5,000 samples were used as burn-in samples.

For each data set, we estimated the posterior mean of a and β . We also calculated standard errors of the posterior means for each data set. Let $\hat{\alpha}_j$ and $\hat{\beta}_j$ be the posterior mean based on the j^{th} data set, for $j = 1, \dots, 500$. For each data set, we computed the frequentist standard error of the estimator based on 1) the regular bootstrap method and 2) the proposed importance sampling based approach. For the j^{th} data set, we drew $B = 500$ bootstrap

samples with replacement. Suppose that $(\hat{\alpha}_{mcmc,j}^{(b)}, \hat{\beta}_{mcmc,j}^{(b)})$ denotes the posterior means for the b^{th} bootstrap data, for $b = 1, \dots, 500$, and these posterior means were calculated by applying the MCMC method to each bootstrap data separately. The regular bootstrap standard error for $\hat{\alpha}_j$ and $\hat{\beta}_j$ are now expressed as

$sd_{1,j}(\alpha) = \sqrt{(1/499) \sum_{b=1}^{500} \left(\hat{\alpha}_{mcmc,j}^{(b)} - \sum_{b'=1}^{500} \hat{\alpha}_{mcmc,j}^{(b')} / 500 \right)^2}$ and
 $sd_{1,j}(\beta) = \sqrt{(1/499) \sum_{b=1}^{500} \left(\hat{\beta}_{mcmc,j}^{(b)} - \sum_{b'=1}^{500} \hat{\beta}_{mcmc,j}^{(b')} / 500 \right)^2}$, respectively. Next, we computed the proposed importance sampling based standard error,
 $sd_{2,j}(\alpha) = \sqrt{(1/499) \sum_{b=1}^{500} \left(\hat{\alpha}_{is,j}^{(b)} - \sum_{b'=1}^{500} \hat{\alpha}_{is,j}^{(b')} / 500 \right)^2}$ and
 $sd_{2,j}(\beta) = \sqrt{(1/499) \sum_{b=1}^{500} \left(\hat{\beta}_{is,j}^{(b)} - \sum_{b'=1}^{500} \hat{\beta}_{is,j}^{(b')} / 500 \right)^2}$, where $(\hat{\alpha}_{is,j}^{(b)}, \hat{\beta}_{is,j}^{(b)})$ denotes the posterior means for the b^{th} bootstrap data based on the importance sampling idea. Our goal is to illustrate that instead of using the regular bootstrap idea that is way more time consuming, one can simply use the importance sampling based method to estimate the frequentist standard error of the Bayes estimators. We wanted to show that proposed method is computationally far more time efficient, and on the other hand, the standard error calculated using the proposed method is close to the standard error calculated based on the regular bootstrap method. We, once again, point out that the regular bootstrap approach requires enumeration of B MCMC chains, one for each of the B bootstrap data sets, while the proposed approach requires enumeration of only one MCMC chain. In the appendix, we compare the computational complexity of the two approaches.

Figure 1 shows a scatter plot of two standard errors (sd_1 and sd_2) for 500 data sets for the intercept and slope parameter. The figure reveals that the two estimates of the standard error are in good agreement as the points are well dispersed around the 45 degree line. Table 1 shows the computation time (in sec) for the two methods, and clearly the proposed importance sampling based approach is computationally far more superior than the regular bootstrap method.

Since the logistic regression belongs to the class of the generalized linear models, we are able to apply Efron (2015)'s method to evaluate the standard deviation of the posterior mean for the intercept and slope parameters. Let $\theta = (\alpha, \beta)^T$. From Equation (3.1) of Efron (2015),

$$f_{\theta}(T) = \exp \left[\theta^T T - \sum_{j=1}^n \log \{1 + \exp(\alpha + \beta X_j)\} \right], \text{ where } T = \left(\sum_{j=1}^n Y_j, \sum_{j=1}^n X_j Y_j \right)^T$$

is the sufficient statistic for θ . Then, $E(T) = \left(\sum_{j=1}^n p_j, \sum_{j=1}^n X_j p_j \right)^T$ and

$$\text{var}(T) = V_{\theta} = \sum_{j=1}^n p_j (1 - p_j) (1, X_j)^T (1, X_j), \text{ where}$$

$p_j = P(Y=1|X_j) = \exp(\alpha + \beta X_j) / \{1 + \exp(\alpha + \beta X_j)\}$, the success probability given $X = X_j$. Due to the numerical instability of the “conversion factors” we are not able to apply his method that completely avoids MCMC sampling, and this issue has been acknowledged in Efron (2015). However, we apply his general approach for calculating the standard deviation of the posterior mean that is summarized in the following steps.

Step 1—Draw M MCMC samples $(\theta_1, \dots, \theta_M)$ from $\pi(\theta|D)$.

Step 2—Estimate $\text{cov}(\theta, \theta|T)$ by $\widehat{\text{cov}} = \sum_{j=1}^M (\theta_j - \bar{\theta})(\theta_j - \bar{\theta})^T / M$, where $\bar{\theta} = \sum_{j=1}^M \theta_j / M$. Then we obtain $sd_3 = [\widehat{\text{cov}}^T V_{\hat{\theta}} \widehat{\text{cov}}]^{(1/2)}$.

Now we compare sd_3 with the gold standard approach, sd_1 , in Figure 2. In terms of computation time, Efron's approach is much much faster than any other procedure (Table 1). However, Efron's approach is applicable when V_{θ} is easily available, and his method can compute standard error for posterior mean only, not for any quantiles.

Next we calculated the standard error of the first quartile, third quartile, the 2.5th percentile, and the 97.5th percentile of the posterior distribution of α and β based on 1) the regular bootstrap method and 2) the importance sampling based method. We particularly considered 2.5th and 97.5th percentiles as they are often used for constructing credible intervals. Figures 3 and 4 show the standard errors computed using the two approaches for each of these summary statistics for the simulated data sets. We want to point out that for the 2.5th and 97.5th percentiles we used the trial distribution that is described in Section 4.1 and it involves with a slightly more computation than the scenario where $\pi(\alpha, \beta|D)$ is used as a trial distribution (see Table 1). However, despite of being more computationally involved, overall this approach is more time efficient (see Table 1) than the regular bootstrap method where one needs to run MCMC method on each bootstrap data set separately. We also need to keep in mind that this time comparison is heavily depended on the number of MCMC iterations used in the computation, and the time gain will be more if more MCMC iterations are used for the posterior inference. For a fair comparison, every core computation was conducted using FORTRAN 90 within an R script. That is, generation of random samples from the posterior distribution $\pi(\theta|D)$ and evaluation of the importance weight $\omega^{(b)}(\theta)$ in Sections 3, 4.1 were programmed in FORTRAN. Although there are a number of presumably optimized programs or R packages for Bayesian computing, we decide to write our own code for fair comparison across the methods.

The computational complexity of the proposed method and the regular bootstrap method using MCMC simulations are of the same order, and according to the Bachman-Landau notation it is $O(BMn)$, where B, M, n denote the number of bootstrap samples, the number of MCMC iterations, and the sample size, respectively. In Appendix B, we have explained the computational complexity for this example through algorithms, and similar algorithms can be written for other examples. Although the computational complexity of the regular bootstrap method and the proposed method are of the same order, by avoiding MCMC simulations the computation of posterior summary is much faster in the latter method than the former approach.

5.2 Linear measurement error model

Next, we revisit the linear measurement error model. We first note that the joint distribution of the observed Y and W , $f_{Y,W}(y, w)$ is an exponential family. Since

$$f_{Y,W}(y, w) = \int f(w, x, y) dx, \text{ where } f(y, w, x) \text{ is the joint density of } W, X, Y,$$

$$f_{Y,W}(Y, W) = h(Y, W) c(\theta) \exp \left[-\frac{1}{2} \left\{ \frac{1}{\sigma_\epsilon^2} - \frac{\beta^2/\sigma_\epsilon^4}{\beta^2/\sigma_\epsilon^2 + 1/\sigma_u^2 + 1/\sigma_x^2} \right\} Y^2 \right. \\ \left. - \left\{ \frac{\alpha\beta^2/\sigma_\epsilon^4 + \beta\mu_x/(\sigma_\epsilon^2\sigma_x^2)}{\beta^2/\sigma_\epsilon^2 + 1/\sigma_u^2 + 1/\sigma_x^2} - \frac{\alpha}{\sigma_\epsilon^2} \right\} Y + \frac{1/\sigma_u^4}{2(\beta^2/\sigma_\epsilon^2 + 1/\sigma_u^2 + 1/\sigma_x^2)} W^2 \right. \\ \left. - \frac{\alpha B/(\sigma_\epsilon^2\sigma_u^2) - \mu_x/(\sigma_u^2\sigma_x^2)}{\beta^2/\sigma_\epsilon^2 + 1/\sigma_u^2 + 1/\sigma_x^2} W + \frac{\beta/(\sigma_\epsilon^2\sigma_u^2)}{\beta^2/\sigma_\epsilon^2 + 1/\sigma_u^2 + 1/\sigma_x^2} YW \right],$$

where $h(Y, W) = \exp(-W^2/2\sigma_u^2)$ does not depend on θ since σ_u^2 is known and

$$c(\theta) = (2\pi)^{-1} \left\{ \sigma_\epsilon^2 \sigma_u^2 \sigma_x^2 \left(\beta^2/\sigma_\epsilon^2 + 1/\sigma_u^2 + 1/\sigma_x^2 \right) \right\}^{-1/2}$$

$$\exp \left\{ -\alpha^2/2\sigma_\epsilon^2 - \mu_x^2/2\sigma_x^2 + \left(\alpha^2\beta^2/2\sigma_\epsilon^4 + \mu_x^2/2\sigma_x^4 - \alpha\beta\mu_x/2\sigma_\epsilon^2\sigma_x^2 \right) / \left(\beta^2/\sigma_\epsilon^2 + 1/\sigma_u^2 + 1/\sigma_x^2 \right) \right\}$$

is a function of θ . Therefore, $T = (Y^2, Y, W^2, W, YW)$ is a sufficient statistic for the natural

parameter $\eta = (\eta_1, \dots, \eta_5)$, where $\eta_1 = 1/\sigma_\epsilon^2 - (\beta^2/\sigma_\epsilon^4) / (\beta^2/\sigma_\epsilon^2 + 1/\sigma_u^2 + 1/\sigma_x^2)$,

$$\eta_2 = \left\{ \alpha\beta^2/\sigma_\epsilon^4 + \beta\mu_x/(\sigma_\epsilon^2\sigma_x^2) \right\} / \left(\beta^2/\sigma_\epsilon^2 + 1/\sigma_u^2 + 1/\sigma_x^2 \right) - \alpha/\sigma_\epsilon^2,$$

$$\eta_3 = 1/\sigma_u^2 - (1/\sigma_u^4) / \left(\beta^2/\sigma_\epsilon^2 + 1/\sigma_u^2 + 1/\sigma_x^2 \right) \quad \eta_4 = \left\{ \alpha\beta/(\sigma_\epsilon^2\sigma_u^2) - \mu_x/(\sigma_u^2\sigma_x^2) \right\} / \left(\beta^2/\sigma_\epsilon^2 + 1/\sigma_u^2 + 1/\sigma_x^2 \right)$$

, and $\eta_5 = \left\{ \beta/(\sigma_\epsilon^2\sigma_u^2) \right\} / \left(\beta^2/\sigma_\epsilon^2 + 1/\sigma_u^2 + 1/\sigma_x^2 \right)$. In order to apply Efron (2015)'s method,

we need to find the variance covariance matrix V_η of T , which is a very difficult if not impossible task. Therefore, we applied our approach to compute the frequentist standard error for the posterior summaries of α and β .

We generated a single data set comprising of $D = \{(Y_i, W_i, X_i), i = 1, \dots, n = 1,000\}$ under the true model $Y_i = \alpha + \beta X_i + \epsilon_i$, $\alpha = 0.23$, $\beta = 0.47$, and $W_i = X_i + U_i$, where $\epsilon_i \sim \text{Normal}[0, (0.5)^2]$, $U_i \sim \text{Normal}[0, (0.5)^2]$ and $X_i \sim \text{Normal}(0.5, 1)$. We analyzed the data according to the method described in Sections 3 and 4, without using X in the analysis. We applied Gibbs sampling to draw samples from the posterior distribution of the parameters, and used $M = 10,000$ iterations after the first 5,000 samples as burn-in samples. For the prior distributions, we set $\sigma_\alpha^2 = \sigma_\beta^2 = \sigma_\mu^2 = 10,000$ and $\delta_x = \delta_\epsilon = \lambda_x = \lambda_\epsilon = 1$. Then we drew $B = 500$ bootstrap samples with replacement and we evaluated sd_1 and sd_2 as described in Section 5.1. Table 2 shows the frequentist standard errors corresponding to the posterior summaries of α and β , along with the computation time. The results show the advantages of the proposed method over the regular bootstrap method in terms of computational time.

5.3 Weibull regression model

We now analyze a subset of the E1684 melanoma clinical trial data (Example 1.2 and 2.2 of Ibrahim et al., 2001) to determine the frequentist standard errors of posterior summaries from the Weibull model. This was a phase III clinical trial conducted by Eastern Cooperative Oncology Group (ECOG) with chemotherapy of interferon alpha-2b in melanoma patient and can be found at "<http://merlot.stat.uconn.edu/~mhchen/survbook/>". The data set contains observed time measured in year, (right) censoring indicator and chemotherapy treatment indicator for each of 255 patients. The purpose of this clinical study was to examine the

treatment effect on the survival times (Y). Among the possible models for this objective, we fit a Weibull regression model on the survival times (Y) using chemotherapy as a covariate (X) according to Example 2.2 in Ibrahim et al. (2001). Following Ibrahim et al. (2001), we used a Gamma(1, 0.001) prior for α and a Normal($(0, 0)^T$, $10^4 I_2$) prior for β , where I_2 denotes the 2×2 identity matrix, for the Weibull regression model described in Section 2. Here we also generated $B = 500$ bootstrap data sets to calculate standard errors for the posterior summaries of parameters.

Table 3 shows the posterior estimates of β_0 , β_1 and α , corresponding frequentist standard errors, and computing times. Instead of presenting only posterior means as done in Table 2.2 of Ibrahim et al. (2001), we extend that table to include other posterior summaries and the frequentist uncertainty of the estimates. Moreover, following the method described in Section 5, we are able to calculate the standard errors more time efficiently.

Furthermore, it is worth to note that it is difficult to apply Efron (2015)'s approach for calculating frequentist standard deviation of posterior mean to the Weibull model because it is not an exponential family of distributions. Secondly, the joint density of the above model is

$$f(D|\alpha, \beta_0, \beta_1) = \exp \left[\sum_{i=1}^n \{ \Delta_i \log \alpha + \Delta_i (\beta_0 + X_i \beta_1) + \Delta_i (\alpha - 1) \log(Y_i) - Y_i^\alpha \exp(\beta_0 + X_i \beta_1) \} \right]$$

so that it is also hard to calculate V_θ the variance of the sufficient statistic, where $\theta = (\alpha, \beta_0, \beta_1)$. Hence, we are not able to apply his method in this context.

5.4 Vector autoregressive model (VAR)

In the previous examples, we discussed the frequentist standard errors of posterior summaries for parameters themselves. We now discuss a more complicated case where the main interest is a function of parameters. Suppose that we have a p -dimensional time series data y_s , $s = 1, \dots, S$, and assume that the data follows a vector autoregression (VAR) model.

The VAR model with lag L is $y'_s = \mu' + \sum_{j=1}^L y'_{s-j} B_j + \epsilon'_s$, where μ is an $1 \times p$ vector, B_j is a $p \times p$ coefficient matrix, $\epsilon_1, \dots, \epsilon_S$ are iid $N(0, \Sigma)$, and the covariance Σ is an unknown $p \times p$ positive definite matrix. Instead of focusing our attention on the elements of parameter

matrices $B = (B'_1, \dots, B'_L)'$ and Σ , it is more of interest to estimate the impact of changing an element of y_s on the future value y_{s+k} . These effects are called impulse responses (Stock and Watson, 2001), and they are defined as nonlinear functions of the parameters matrices B and Σ .

The likelihood function of (μ, B, Σ) is (μ, B, Σ) is

$$L(\Phi, \Sigma) = (2\pi)^{-Sp/2} |\Sigma|^{-S/2} \exp \left[-tr \left\{ (Y - X\Phi) \Sigma^{-1} (Y - X\Phi)' \right\} / 2 \right], \text{ where } Y =$$

$(y_1, \dots, y_S)'$, $\Phi = (\mu', B')'$, $X = (x_1, \dots, x_S)'$, and $x_s = (1, y'_{s-1}, \dots, y'_{s-L})'$. Note that Y is $S \times p$ matrix, X is $S \times (Lp + 1)$ matrix, and Φ is $(Lp + 1) \times p$ matrix. Here we consider the impulse response to orthogonalized errors $U = \epsilon' \Psi^{-1}$, where Ψ is the Cholesky matrix for Σ , i.e., $\Sigma = \Psi' \Psi$. That is, the impulse responses Z_k of y_{s+k} based on the structural shock $\epsilon'_s \Psi^{-1}$

is $\mathbf{Z}_k = \mathbf{\Psi} \mathbf{H}_k$, where $\mathbf{H}_j = \sum_{i=1}^j \mathbf{B}_j \mathbf{H}_{j-i}$, and $\mathbf{B}_i = 0$ for i larger than lag L and $\mathbf{B}_0 = \mathbf{I}$ (Sims, 1980; Ni et al., 2007).

For the computational purpose, we consider conjugate priors for $(\mathbf{\Phi}, \mathbf{\Sigma})$. That is,

$\pi(\mathbf{\Sigma}) \propto |\mathbf{\Sigma}|^{-(p+1)/2}$, the Jeffreys prior, and

$\pi(\phi) \propto |\mathbf{M}_0|^{-1/2} \exp \left\{ -(\phi - \phi_0)' \mathbf{M}_0^{-1} (\phi - \phi_0) / 2 \right\}$, where $\phi = \text{vec}(\mathbf{\Phi})$. Next,

following Ni et al. (2007), the conditional density of ϕ given $\mathbf{\Sigma}, \mathbf{D}$ is $\mathcal{N}(\mathbf{m}, \mathbf{V})$ and the conditional density of $\mathbf{\Sigma}$ given $\mathbf{\Phi}, \mathbf{D}$ is inverse Wishart $(S(\mathbf{\Phi}), \mathbf{M})$, where

$$\mathbf{m} = \hat{\phi}_{mle} + \left\{ \mathbf{M}_0^{-1} + \mathbf{\Sigma}^{-1} \otimes (\mathbf{X}' \mathbf{X}) \right\}^{-1} \mathbf{M}_0^{-1} (\phi_0 - \hat{\phi}_{mle}), \mathbf{V} = \left\{ \mathbf{M}_0^{-1} + \mathbf{\Sigma}^{-1} \otimes (\mathbf{X}' \mathbf{X}) \right\}^{-1},$$

$$\hat{\phi}_{mle} = \text{vec}(\hat{\mathbf{\Phi}}_{mle}), \hat{\mathbf{\Phi}}_{mle} = (\mathbf{X}' \mathbf{X})^{-1} \mathbf{X}' \mathbf{Y}, S(\mathbf{\Phi}) = (\mathbf{Y} - \mathbf{X}\mathbf{\Phi})'(\mathbf{Y} - \mathbf{X}\mathbf{\Phi}), \text{ and } \mathbf{D} = \{\mathbf{y}_1, \dots, \mathbf{y}_S\}$$

is the observed data. Since the impulse response is a function of $\mathbf{B}$ and $\mathbf{\Sigma}$, we rewrite $\mathbf{Z}_k = \mathbf{Z}(\mathbf{B}, \mathbf{\Sigma}, k)$. We take the posterior mean as a Bayes estimator of the impulse response, and it

is $(\hat{\mathbf{Z}}_k)_{(i,j)} = \int \{\mathbf{Z}(\theta, k)\}_{(i,j)} \pi(\theta | \mathbf{D}) d\theta$, where $\theta = (\mathbf{B}, \mathbf{\Sigma})$, and $(\mathbf{Z}_s)_{(i,j)}$ is the (i, j) element of $\mathbf{Z}_s$. Now, we apply the proposed method to calculate the frequentist standard error of $\hat{\mathbf{Z}}_k$. In

this complex example, it is nearly impossible to find the variance-covariance matrix of the sufficient statistics of θ , therefore it is not possible to apply Efron's approach.

For illustration purpose, we generated a data set from the following VAR(1) model with $p = 2$,

$$\mathbf{y}'_s = \begin{bmatrix} -0.7 & 1.3 \end{bmatrix} + \begin{bmatrix} 0.7 & 0.3 \\ 0.2 & 0.6 \end{bmatrix} \mathbf{y}'_{s-1} + \boldsymbol{\epsilon}_s, \boldsymbol{\epsilon}_s \stackrel{iid}{\sim} N(0, \mathbf{\Sigma}), \mathbf{\Sigma} = \begin{bmatrix} 1 & 0.5 \\ 0.5 & 1 \end{bmatrix}, s=1, \dots, S,$$

with $S = 1,000$. Since the time series data is no more independent, we used the moving block bootstrap (MBB), where we divided the series into N overlapping blocks of length ℓ to preserve the dependence structure of the original dataset (Kreiss and Lahiri, 2012). Then we chose b blocks out of N blocks to make the bootstrap observations $\mathbf{y}_1^*, \dots, \mathbf{y}_S^*$.

We fit a VAR(2) model to the simulated dataset. As in previous examples, we used $M = 10000$ iterations after the burn-in samples. We imposed noninformative priors for $\mathbf{\Phi}$, where $\phi_0 = 0$ and $\mathbf{M}_0 = 20\mathbf{I}$. Then we drew $B = 500$ MBB samples with 15% of the total dataset as a block length (ℓ).

Figure 5 show the point estimate (posterior mean) and the 95% confidence band based on the frequentist standard error of the posterior mean for the impulse responses of $\mathbf{y}_2$ to $\mathbf{y}_1$ and $\mathbf{y}_1$ to $\mathbf{y}_2$, respectively. The confidence bands based on sd_1 and sd_2 are similar, but computationally the second approach (sd_2) was about 5.6 times faster than the first approach (sd_1). In Table 4, we also report the numerical values of the standard errors at each time lag, and the results do not show any appreciable difference between sd_1 and sd_2 .

6 Conclusions

In this paper we have discussed numerical approaches for efficient computation of standard errors for posterior summaries. The main theme of the paper is to use bootstrap samples but avoid using full blown MCMC based inference for each of the bootstrap data. The methods rely on the importance sampling idea, and are broadly applicable. The R code for our computation is available at <https://stat.tamu.edu/~sinha/research.html>.

It is well-known that the presence of outliers results in a poor performance in a bootstrap approach because they are more frequent in bootstrap samples than the original dataset if we consider the classical nonparametric bootstrap (Salibian-Barrera and Zamar, 2002; Willems and Van Aelst, 2005; Huber and Ronchetti, 2009). Therefore, the performance of our proposed method can be affected by the outliers in the data as we have used the classical nonparametric bootstrap with replacement. However, this may be overcome by considering robust bootstrap methods for drawing samples (Singh, 1998; Hu and Hu, 2000; Salibian-Barrera and Zamar, 2002), or a combination of a robust bootstrap method and a robust Bayesian method, possibly with a flat-tailed prior (Berger et al., 1994; Marín, 2000).

Acknowledgments

The authors would like to thank two referees for their thoughtful comments that lead to a much improved manuscript. Carroll's research was supported by a grant from the National Cancer Institute (U01-CA057030). Lee and Sinha's research was supported by NIH grant R03CA176760.

Appendix

A Proof of the convergence result of Section 3

Here we discuss the convergence of (1). Suppose that $\omega^{(b)}(\theta)$ and $\theta\omega^{(b)}(\theta)$ are integrable functions of θ with respect to the posterior distribution of the original data $\pi(\theta|\mathbf{D})$ so that $G_s^{(b)} = \int \theta^s \omega^{(b)}(\theta) \pi(\theta|\mathbf{D}) d\theta / K_\pi = E_{\pi(\cdot|\mathbf{D})} \{ \theta^s \omega^{(b)}(\theta) \} / K_\pi$ is finite for all b and $s = 0, 1$. Therefore, as $M \rightarrow \infty$, from the ergodic theorem (Jones, 2004; Robert and Casella, 2005), with probability 1,

$$\begin{aligned} \frac{1}{M} \sum_{j=1}^M \omega^{(b)}(\theta_j) &\rightarrow E_{\pi(\cdot|\mathbf{D})} \{ \omega^{(b)}(\theta) \} = K_\pi G_0^{(b)}, \\ \frac{1}{M} \sum_{j=1}^M \theta_j \omega^{(b)}(\theta_j) &\rightarrow E_{\pi(\cdot|\mathbf{D})} \{ \theta \omega^{(b)}(\theta) \} = K_\pi G_1^{(b)}. \end{aligned}$$

From Remark 1 in Section 3, $\omega^{(b)}(\theta) = \exp\{\ell(\theta) - \ell(\theta^*)\}$ implies $\omega^{(b)}(\theta)$ is positive for all θ . Therefore, $\omega^{(b)}(\theta)$ and $\theta\omega^{(b)}(\theta)$, and consequently

$$\hat{\theta}_{is}^{(b)} = \frac{\sum_{j=1}^M \theta_j^s \omega^{(b)}(\theta_j)}{\sum_{j=1}^M \omega^{(b)}(\theta_j)} \rightarrow \frac{G_1^{(b)}}{G_0^{(b)}} = \hat{\theta}^{(b)}$$

with probability 1 as $M \rightarrow \infty$.

B Computational complexity of the two approaches for the logistic regression example

Algorithm 1

Full Bootstrap method for the logistic regression model in Section 5.1

```

for  $b = 1$  to  $B$  do
    Draw a bootstrap sample
    Initialize  $\alpha_0^{(b)}$  and  $\beta_0^{(b)}$ 

    for  $m = 1$  to  $M + \text{burn}$  do
        Propose  $\alpha^{cand}, \beta^{cand} \sim q(\alpha, \beta \mid \alpha_{m-1}^{(b)}, \beta_{m-1}^{(b)})$ 

        Calculate
        
$$r = \min \left[ 1, \frac{\pi(\alpha^{cand}, \beta^{cand} \mid D^{(b)}) q(\alpha_{m-1}^{(b)}, \beta_{m-1}^{(b)} \mid \alpha^{cand}, \beta^{cand})}{\pi(\alpha_{m-1}^{(b)}, \beta_{m-1}^{(b)} \mid D^{(b)}) q(\alpha^{cand}, \beta^{cand} \mid \alpha_{m-1}^{(b)}, \beta_{m-1}^{(b)})} \right]$$


        Generate  $u \sim U(0, 1)$ 

        if  $u < r$  then
             $\alpha_m^{(b)} = \alpha^{cand}$ 
             $\beta_m^{(b)} = \beta^{cand}$ 
        else
             $\alpha_m^{(b)} = \alpha_{m-1}^{(b)}$ 
             $\beta_m^{(b)} = \beta_{m-1}^{(b)}$ 
        end if
    end for

    Find averages  $\hat{\alpha}^{(b)}, \hat{\beta}^{(b)}$  for the  $b^{\text{th}}$  bootstrap sample:
    
$$\hat{\alpha}^{(b)} = \sum_{j=1}^M \alpha_j^{(b)} \Big/ M \text{ and } \hat{\beta}^{(b)} = \sum_{j=1}^M \beta_j^{(b)} \Big/ M$$


end for

Evaluate standard deviations  $sd_1(\alpha)$  and  $sd_1(\beta)$  as in Section 5.1

```

Algorithm 2

Proposed method for the logistic regression model in Section 5.1

```

for  $b = 1$  to  $B$  do
    Draw a bootstrap sample, and obtain  $(r_1^{(b)}, \dots, r_n^{(b)})$ 

    for  $m = 1$  to  $M$  do
        evaluate  $\omega^{(b)}(\alpha_m, \beta_m) = \prod_{i=1}^n f^{(r_i^{(b)} - 1)}(X_i, Y_i \mid \alpha_m, \beta_m)^{\dagger}$ 
    end for

```

Find averages $\hat{\alpha}^{(b)}, \hat{\beta}^{(b)}$ for the b^{th} bootstrap sample ^{*}:

$$\hat{\alpha}^{(b)} = \sum_{j=1}^M \alpha_j \omega^{(b)}(\alpha_m, \beta_m) \bigg/ \sum_{j=1}^M \omega^{(b)}(\alpha_m, \beta_m) \text{ and}$$

$$\hat{\beta}^{(b)} = \sum_{j=1}^M \beta_j \omega^{(b)}(\alpha_m, \beta_m) \bigg/ \sum_{j=1}^M \omega^{(b)}(\alpha_m, \beta_m)$$

end for

Evaluate standard deviations $sd_2(\alpha)$ and $sd_2(\beta)$ as in Section 5.1

[†] $(\alpha_1, \beta_1), \dots, (\alpha_M, \beta_M)$ are from $\pi(\alpha, \beta | D)$.

^{*} If we are interested in the q^{th} quantile we will compute $\hat{\alpha}_q^{(b)}, \hat{\beta}_q^{(b)}$ based on equation 2 in Section 4.1 at this step.

References

- Berger J. The case for objective bayesian analysis. *Bayesian analysis*. 2006; 1:385–402.
- Berger JO, Moreno E, Pericchi LR, Bayarri MJ, Bernardo JM, Cano JA, De la Horra J, Martín J, Ríos-Insúa D, Betrò B, et al. An overview of robust bayesian analysis. *Test*. 1994; 3:5–124.
- Carlin, BP., Louis, TA. *Bayesian Methods for Data Analysis*. CRC Press; 2008.
- Carroll, R., Ruppert, D., Stefanski, L., Crainiceanu, C. *Measurement Error in Nonlinear Models: A Modern Perspective*. Second Edition. CRC Press; 2006.
- Efron B. Bayesian inference and the parametric bootstrap. *Annals of Applied Statistics*. 2012; 6:1971–1997. [PubMed: 23843930]
- Efron B. Frequentist accuracy of Bayesian estimates. *Journal of the Royal Statistical Society, Series B*. 2015; 77:617–646.
- Efron B, Tibshirani R. Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Statistical Science*. 1986; 1:54–75.
- Efron, B., Tibshirani, R. *An introduction to the bootstrap*. CRC press; 1994.
- Fuller, W. *Measurement Error Models*. John Wiley & Sons; New York: 1987.
- Hu F, Hu J. A note on breakdown theory for bootstrap methods. *Statistics & probability letters*. 2000; 50:49–53.
- Huber, P., Ronchetti, E. *Robust Statistics*. John Wiley & Sons; New York: 2009.
- Ibrahim, J., Chen, M., Sinha, D. *Bayesian Survival Analysis*. Springer; 2001.
- Jones GL. On the markov chain central limit theorem. *Probability Surveys*. 2004; 1:299–320.
- Kreiss J-P, Lahiri S. Bootstrap methods for time series. *Handbook of Statistics: Time Series Analysis: Methods and Applications*. 2012; 30
- Liang F. Dynamically weighted importance sampling in Monte Carlo computation. *Journal of the American Statistical Association*. 2002; 97:807–821.
- Marín, JM. *Robust Bayesian Analysis*. Springer; 2000. A robust version of the dynamic linear model with an economic application.; p. 373-383.
- Ni S, Sun D, Sun X. Intrinsic bayesian estimation of vector autoregression impulse responses. *Journal of Business & Economic Statistics*. 2007; 25:163–176.
- Robert, C., Casella, G. *Monte Carlo Statistical Methods*. Springer; New York: 2005.
- Salibian-Barrera M, Zamar RH. Bootstrapping robust estimates of regression. *The Annals of Statistics*. 2002; 30:556–582.
- Sims CA. Macroeconomics and reality. *Econometrica*. 1980; 48:1–48.
- Singh K. Breakdown theory for bootstrap quantiles. *The annals of Statistics*. 1998; 26:1719–1732.
- Stock JH, Watson MW. Vector autoregressions. *The Journal of Economic Perspectives*. 2001; 15:101–115.
- Willems G, Van Aelst S. Fast and robust bootstrap for LTS. *Computational Statistics & Data Analysis*. 2005; 48:703–715.

Figure 1.

Frequentist standard errors of posterior means of the intercept (α) and the slope (β) of the logistic regression model from the 500 simulated data sets in Section 5.1 based on the regular bootstrap method (Y-axis) and the proposed importance sampling based method (X-axis).

Figure 2.

Frequentist standard errors of posterior means of the intercept (α) and the slope (β) of the logistic regression model from the 500 simulated data sets in Section 5.1 based on the regular bootstrap method (Y-axis) and the approach proposed in Efron (2015) (X-axis).

Figure 3.

Frequentist standard errors of Q_1 , Q_3 , 2.5th percentile, and 97.5th percentile of the posterior distribution of α in the logistic regression model from the 500 simulated data set in Section 5.1. Regular bootstrap standard errors are presented along the Y-axis while importance sampling based standard errors are presented along the X-axis.

Figure 4.

Frequentist standard errors of Q_1 , Q_3 , 2.5th percentile, and 97.5th percentile of the posterior distribution of β in the logistic regression model from the 500 simulated data set in Section 5.1. Regular bootstrap standard errors are presented along the Y-axis while importance sampling based standard errors are presented along the X-axis.

Figure 5.

The estimated impulse responses (solid line) and its 95% (pointwise) confidence band of y_2 to a shock in y_1 (left panel) and vice versa (right panel) referenced in Section 5.4. The bold dotted line is based on regular bootstrap approach (sd_1) while the circled solid line is based on importance sampling based approach (sd_2).

Table 1

Average computing time ($\pm$ standard deviation of 500 simulated data sets) measured in seconds for calculating standard errors of posterior summaries in logistic regression model from Section 5.1 based on the 1) regular bootstrap method, 2) the importance sampling based approach, and 3) the method proposed in Efron (2015). Here Q_1 and Q_3 denote the first and third quartiles, and 2.5th and 97.5th denote the 2.5th percentile and 97.5th percentile of the posterior distribution, respectively.

Method	Time to calculate			Computational complexity
	Mean	Q_1 (Q_3)	2.5 th (97.5 th)	
1	247.78 $\pm$ 6.70	247.78 $\pm$ 6.70	247.78 $\pm$ 6.70	$O(BMn)$
2	46.65 $\pm$ 0.41	51.66 $\pm$ 0.40	120.44 $\pm$ 0.93	$O(BMn)$
3	4.18 $\pm$ 0.6			$O(Mn)$

Table 2

The frequentist standard errors and computing times for α and β of the linear measurement error model in Section 5.2. Here sd_1 and sd_2 denote the standard errors based on the regular bootstrap method and the importance sampling based approach.

Parameter		Mean	Posterior		
			Q_2	2.5 th	97.5 th
α	sd_1	0.028	0.027	0.028	0.027
	sd_2	0.027	0.029	0.028	0.030
β	sd_1	0.034	0.034	0.033	0.036
	sd_2	0.030	0.032	0.035	0.031
Computation time in second	sd_1	233.06	233.06	233.06	233.06
	sd_2	22.99	26.36	24.16	24.16
Computational complexity	sd_1	$O(BMn)$	$O(BMn)$	$O(BMn)$	$O(BMn)$
	sd_2	$O(BMn)$	$O(BMn)$	$O(BMn)$	$O(BMn)$

Table 3

Posterior summaries and the corresponding frequentist standard errors of β_0 , β_1 , and α used in the Weibull model for analyzing the E1684 melanoma data given in Section 5.3. Here sd_1 and sd_2 denote the standard errors based on the regular bootstrap method and the importance sampling based approach.

Parameter		Posterior			
		Mean	Q_2	2.5 th	97.5 th
β_0		-1.103	-1.101	-1.710	-0.586
	sd_1	0.278	0.278	0.295	0.266
	sd_2	0.255	0.265	0.252	0.261
β_1		-0.256	-0.256	-0.585	0.090
	sd_1	0.177	0.178	0.180	0.179
	sd_2	0.169	0.176	0.163	0.183
α		0.791	0.793	0.688	0.891
	sd_1	0.038	0.039	0.035	0.043
	sd_2	0.037	0.038	0.034	0.038
Computation time in sec	sd_1	151.77	151.77	151.77	151.77
	sd_2	37.31	43.59	85.75	85.75
Computational complexity	sd_1	$O(BMn)$	$O(BMn)$	$O(BMn)$	$O(BMn)$
	sd_2	$O(BMn)$	$O(BMn)$	$O(BMn)$	$O(BMn)$

Table 4

The frequentist standard errors of the estimated the impulse responses at each time lag. Here sd_1 and sd_2 denote the standard errors based on the regular bootstrap method and the importance sampling based approach.

Time lag	y_2 to y_1		y_1 to y_2	
	sd_1	sd_2	sd_1	sd_2
1	0.0241	0.0222	0.0404	0.0403
2	0.0388	0.0374	0.0338	0.0334
3	0.0442	0.0434	0.0304	0.0298
4	0.0431	0.0428	0.0279	0.0272
5	0.0401	0.0402	0.0262	0.0256
6	0.0369	0.0372	0.0253	0.0248
7	0.0339	0.0342	0.0247	0.0244
8	0.0311	0.0315	0.0243	0.0242
9	0.0286	0.0290	0.0239	0.0240
10	0.0263	0.0267	0.0234	0.0236
11	0.0242	0.0246	0.0227	0.0232
12	0.0222	0.0226	0.0220	0.0225
13	0.0204	0.0208	0.0212	0.0218
14	0.0188	0.0191	0.0203	0.0210
15	0.0173	0.0176	0.0194	0.0201