

Simulation-based Calibration of Uncertainty Intervals under Approximate Bayesian Estimation

Terrance D. Savitsky

*Office of Survey Methods Research
U.S. Bureau of Labor Statistics
2 Massachusetts Ave NE
Washington, DC 20212, USA*

SAVITSKY.TERRANCE@BLS.GOV

Julie Gershunskaya

*Office of Employment and Unemployment Statistics
U.S. Bureau of Labor Statistics
2 Massachusetts Ave NE
Washington, DC 20212, USA*

GERSHUNSKAYA.JULIE@BLS.GOV

Editor: Edo Airoldi

Abstract

The mean field variational Bayes (VB) algorithm implemented in Stan is relatively fast and efficient, making it feasible to produce model-estimated official statistics on a rapid timeline. Yet, while consistent point estimates of parameters are achieved for continuous data models, the mean field approximation often produces inaccurate uncertainty quantification to the extent that parameters are correlated a posteriori. In this paper, we propose a simulation procedure that calibrates uncertainty intervals for model parameters estimated under approximate algorithms to achieve nominal coverages. Our procedure detects and corrects biased estimation of both first and second moments of approximate marginal posterior distributions induced by any estimation algorithm that produces consistent first moments under specification of the correct model. The method generates replicate data sets using parameters estimated in an initial model run. The model is subsequently re-estimated on each replicate data set, and we use the empirical distribution over the re-samples to formulate calibrated confidence intervals of parameter estimates of the initial model run that are guaranteed to asymptotically achieve nominal coverage. We demonstrate the performance of our procedure in Monte Carlo simulation study and apply it to real data from the Current Employment Statistics survey.

Keywords: Small area estimation, Approximate Bayesian computation, Variational Bayes, Parametric bootstrap, Uncertainty quantification, Bayesian hierarchical modeling, Survey data estimation

1. Introduction

The Bureau of Labor Statistics (BLS) publishes employment estimates for geography-by-industry domains on a monthly basis in the Current Employment Statistics (CES) survey. Both a point and associated variance estimate is constructed for each domain from underlying survey data (linked to that domain) acquired from business establishments. These survey-based domain-level point and variance estimates contain both sampling and measure-

ment sources of error. It is common practice to conduct small domain estimation modeling using a Bayesian hierarchical formulation (Savitsky and Gershunskaya, 2022) that borrows information among correlated domains and outputs smoothed domain point and variance estimates. The Bayesian hierarchical model used by BLS to produce CES domain estimates are richly parameterized (with the number of parameters on the order of the number of domains) to capture the correlation structures; for example, the model parameterizes domain-level random effects under a mixture or clustering prior distribution to capture dependence in the response variable of interest among the domains.

We utilize Stan (Carpenter et al., 2017) to estimate our Bayesian hierarchical model for its expressive syntax and flexibility of prior specifications. Stan selects a (“No U-turns”) version of Hamiltonian Monte Carlo algorithm (HMC) to conduct numerical integration. HMC is a version of Markov Chain Monte Carlo (MCMC) that samples all parameters, simultaneously, under the joint posterior distribution in a single sweep in each MCMC iteration. By contrast, a Gibbs sampler blocks parameters and samples parameter full conditional distributions sequentially to build the joint distribution. The joint sampling of all parameters on each iteration produces better mixing when model parameters are a posteriori highly correlated as compared to Gibbs sampling.

We are, however, unable to use HMC for our clustering model introduced in the sequel because the discrete cluster assignment parameter vector must be integrated out (since HMC does not sample discrete parameter spaces), which induces label switching. The label switching prevents convergence under HMC estimation.

Stan provides an implementation of variational Bayes (VB) in their Automatic Differentiation Variational Inference (ADVI) algorithm that uses optimization instead of the numerical integration of MCMC that resolves label switching for our clustering model configuration. (A brief over of the mean field variational formulation underlying ADVI is provided in Appendix A). BLS utilize Stan’s ADVI to produce the CES monthly smoothed estimators from a Bayesian hierarchical model (that we review in the sequel). The ADVI algorithm produces very accurate smoothed point estimates for parameters under the BLS model for the CES (Gershunskaya and Savitsky, 2019), but the resulting variances or scales of the estimated marginal posterior distributions are inaccurate. BLS desires to publish both the smoothed point estimate, $\mathbb{E}(\theta_i | \mathbf{y})$ of the posterior distribution for domain, $i \in (1, \dots, N)$ for data set, $\mathbf{y} = (y_1, \dots, y_N)$, its associated variance, $V(\theta_i | \mathbf{y})$, along with associated uncertainty intervals for any level of misspecification γ . The ADVI estimation algorithm often produces inaccurate scale estimates of the posterior distributions for $\{\theta_i\}_{i=1}^N$ because it uses a mean-field algorithm that assumes the model parameters are *a posteriori* independent, which is typically not the case.

We introduce a simulation-based procedure in this paper that corrects or calibrates the scale of the marginal posterior parameter distributions and outputs confidence intervals that achieve nominal coverage under any misspecification probability, γ . Our simulation-based calibration method may be used to calibrate uncertainty quantification under *any* estimation algorithm for the marginal posterior distributions (and not just ADVI) under two conditions: 1. The specified estimation model is asymptotically correct (for data generation); 2. The estimation algorithm produces consistent point estimates, though those point estimates may asymptotically express a constant bias. For example, if the algorithm tends to produce an overly small posterior variance for θ_i under an asymptotically correct model specification,

then if we generate a replicate data set indexed by α , $\mathbf{y}^{(\alpha)}$, using posterior draw $\theta_i^{(\alpha)}$ as “truth” and re-estimate the model on $\mathbf{y}^{(\alpha)}$, then the estimated posterior variance of $\theta_i^{(\alpha)}$ from this model will also be downwardly biased. The estimation error propagates from the real data to replicate data because the error is induced by the estimation algorithm.

Let index $i \in (1, \dots, N)$ index a domain. Our simulation-based calibration procedure takes a draw, $\alpha \in (1, \dots, A)$, of $(\theta_i^{(\alpha)})_{i=1}^N$ and generates a replicate data set, $\mathbf{y}^{(\alpha)} = (y_1^{(\alpha)}, \dots, y_n^{(\alpha)})$. We re-estimate the model on each $\mathbf{y}^{(\alpha)}$ from which we compute the approximate posterior means and variances, $(m(\theta_i^{(\alpha)}), v(\theta_i^{(\alpha)}))$ from $\mathbf{y}^{(\alpha)}$. We construct a pivot statistic indexed for each re-sample, α , from $(m(\theta_i^{(\alpha)}), v(\theta_i^{(\alpha)}))$ by treating the point estimate, $\theta_i^{(\alpha)}$, as “truth”. We use the distribution of this pivot statistic to over the $\alpha \in (1, \dots, A)$ replications to construct a mean (shift) and variance (scale) adjustment. Each adjusted mean and variance for domain i are subsequently used to formulate a calibrated confidence interval for θ_i that replaces the approximate posterior distributions from the ADVI algorithm.

Our simulation-based procedure introduced in the sequel may be applied to any posterior estimation algorithm that satisfies the two conditions stated above. So, although we illustrate our procedure to correct the joint posterior distribution estimated under ADVI, the procedure as described may be applied to *any* approximate estimation algorithm that produces consistent point estimation but incorrect uncertainty quantification.

1.1 Literature Review

Our calibration procedure directly operates on the marginal posterior distributions for model parameters and outputs corrected variances and associated confidence intervals that are asymptotically guaranteed to achieve nominal coverage.

By contrast, conformal inference (Lei et al., 2018) focuses on generating prediction intervals for Y_{n+1} , rather than parameter confidence sets. Conformal prediction intervals assume data, $(Y_i, X_i) \stackrel{\text{iid}}{\sim} \mathbb{P}$, $i = 1, \dots, n$, are exchangeable and that the point estimate, $\hat{\mu}(X_i)$, is symmetric. Under these two assumptions, conformal inference outputs prediction intervals that achieve nominal coverage under any misspecification γ for finite samples. Remarkably, this result makes no assumptions about (the consistency of) $\hat{\mu}$ or distribution, \mathbb{P} (though the resulting intervals will be relatively longer in proportion to the bias in $\hat{\mu}$).

Fong and Holmes (2021) construct conformal prediction intervals under a Bayesian modeling setup. The posterior predictive distribution, $g_i := g(Z_1, \dots, Z_{n+1}; Z_i) = p(Y_i | X_i, Z_1, \dots, Z_{n+1}) = \int f_\theta(Y_i | X_i) \times \pi(\theta | Z_1, \dots, Z_{n+1}) d\theta$, where $Z_i = (X_i, Y_i)$. In their set-up the posterior predictive distribution for each i is used as the “conforming score”. Score g_i is exchangeable because the $\pi(\theta | Z_1, \dots, Z_{n+1}) \propto \pi(\theta) \times \prod_{i=1}^{n+1} f_\theta(Y_i | X_i)$.

The resulting rank of g_{n+1} , $\pi(y) = \frac{1}{n+1} \sum_{i=1}^{n+1} \mathbf{1}(g_i \leq g_{n+1})$ is uniformly distributed such that one may form the conformal predictive set, $C_\gamma(X_{n+1}) = \{y \in \mathbb{R} : \pi(y) > \gamma\}$, where y is sampled on a grid and the resulting intervals are guaranteed to achieve nominal coverage. In the typical case the model must be re-fit for each value of y , though Fong and Holmes (2021) devised an importance re-sampling scheme to re-weight the posterior draws for parameters, θ , estimated from the first n data points.

A key assumption to produce conformal prediction intervals under a Bayesian hierarchical model implementation (besides exchangeability) is that valid parameter samples from the model posterior distribution, $\pi(\theta \mid Z_1, \dots, Z_n)$, are available. In our set-up, however, we don't have samples from a valid posterior distribution due to our use of the computationally scalable ADVI estimation algorithm. Even more, our focus is not on prediction of the next datum, but on estimation of the marginal parameter posterior distributions.

Chen et al. (2018) develop a simulation-based calibration procedure that is similar to ours, though their focus is on frequentist estimation of global parameters under variational inference. They utilize a non-parametric bootstrap procedure since their inferential interest is on model *global* parameters. By contrast, our procedure uses a parametric-based re-sampling procedure because our inferential interest is in *local* parameters, θ_i , indexed by observation / domain (such that the re-sampling with replacement procedure for the non-parametric bootstrap would disallow inference on local parameters indexed by domain i). Chen et al. (2018) form a pivot statistic, but construct their variance estimate for each re-sampled data set by using a sandwich estimator that requires analytical computation of the gradient and Hessian. By contrast, we develop a numerical variance adjustment factor that is applied to the estimated variance, $v(\theta_i)$, from the original run under the approximate estimation algorithm that does not require access to analytical gradients and Hessians.

The remainder of this paper presents our target model formulation in Section 2 that motivates our calibration adjustments of marginal posterior distributions. We introduce our resampling-based calibration adjustment method in Section 3 followed by simulation studies in Section 4 that apply our calibration method to our models of Section 2. We discuss the theoretical properties of our two proposed calibration methods that provide context on simulation study results in Section 5. We subsequently apply our simulation-based calibration procedure to correct uncertainty quantification on a real data application from the Current Employment Statistics survey in Section 6 followed by a concluding discussion in Section 7.

2. Model Formulations to Apply Simulation-based Calibration

We introduce a domain-level model used in the Current Employment Statistics (CES) survey that estimates total employment in small domains defined by intersections of industry and geography. This application motivates our development of the resampling-based calibration algorithm introduced in the next section. We subsequently use the calibration procedure we develop to correct bias in the mean and variance estimates from approximate marginal posterior distributions. The approximate marginal posterior distributions are estimated under our domain-level model from a computationally scalable optimization algorithm.

The domain-indexed CES total employment point estimate, y_i , for domain $i \in 1, \dots, N$, is constructed as a survey weighted statistic of underlying responses for sampled business establishments in that domain. These survey-weighted domain statistics are referred to as direct estimates and are not model-based. An accompanying statistic, v_i , representing the variance of y_i (with respect to the distribution governing the taking of samples from a fixed population) is also constructed.

Our inferential interest is extracting local parameters, θ_i , that represent latent “true” domain point estimates. Our model is a so-called small domain estimation formulation

where correlations between domains with relatively large respondent sample sizes and those with relatively small respondent sample sizes are used to improve estimation quality in the latter (as compared to the direct point estimate, y_i). (See Rao and Molina (2015) and Savitsky et al. (2023) for a background on small domain estimation.)

We developed and refined our total employment model for CES domains by estimating the model on CES employment historical series and subsequently testing its performance (e.g., mean-squared prediction error) against a “gold standard” of census/population-based employment levels that become available to researchers several months *after* the actual publication of the CES estimates. The results showed satisfactory performance of the model fitted point estimates, as compared to direct sample based estimates or estimates based on legacy CES models.

We estimate our CES model using the Variational Bayes (VB) ADVI algorithm implemented in Stan. The VB algorithm is relatively fast, making it feasible for use under a tight production timeline to produce monthly survey-based domain point estimates and variances.

Based on this research, we are satisfied with the model fitted point estimates of employment. Our model framework specifies a joint distribution for domain-level point estimate and variance statistics. The variance statistic measures the variation of the point estimate with respect to the distribution governing the taking of samples from a fixed population. The variance statistic is noisy and our model produces a “true” variance of the point estimate. We co-model the point estimate and variance because they are correlated and because the variance provides information about the quality of the point estimate for regulating model-induced shrinkage. That said, we are interested in the variance of the “true” point estimate for each domain, θ_i (for domain $i \in 1, \dots, N$), from the model, which is based on the model credibility intervals for θ_i .

Yet, the credibility intervals for the estimated true point estimate are based on draws from mean field variational algorithm that assumes the joint posterior distribution is fully factored. So, we would *a priori* expect the credibility intervals to be too optimistic and induce undercoverage.

Our chosen model may be viewed as a generalization of the classical area-level Fay-Herriot (FH) model. Thus, we start by formulating the FH model (Section 2.1) and discussing its shortcomings to motivate the extension to our model formulation. A modification of the FH model that allows co-modeling of direct point estimates and their variances is given in Section 2.2. Finally, we formulate the “production version” of the model in Section 2.3.

2.1 The Fay-Herriot model

The Fay-Herriot (FH) model’s name comes from Fay and Herriot (1979) paper, where the authors adapted the James-Stein estimator for making estimates of income for small places. The hierarchical formulation of the model is also familiar to many from the classical eight schools example of Rubin (1981) and it is used as an introductory example to hierarchical modeling by Gelman et al. (2013).

Let y_i be a survey estimate of target “true” point estimate θ_i for domain i (in the case of CES, it represents a relative change in employment from month $t - 1$ to month t in domain

i); for each domain, $i = 1, \dots, N$, assume

$$y_i | \theta_i \stackrel{\text{ind}}{\sim} N(\theta_i, v_i), \quad (1)$$

$$\theta_i | \beta, \tau_u^2 \stackrel{\text{ind}}{\sim} N(\mathbf{x}_i^T \beta, \tau_u^2). \quad (2)$$

Note that variances $v_i, i = 1, \dots, N$, in Equation 1 are assumed to be *fixed and known*.

In practice, true variances of survey estimates are not known, and v_i represents some *estimate* of the variance. Since direct survey estimates of variances contain noise, the traditional approach is to use some sort of *preprocessing* to smooth variances (for example, based on a generalized variance function, or GVF) before they are used in the modeling.

2.2 Co-modeling point estimates and their variances

A more efficient approach seeks to remove noise from the direct variance estimates by fitting them, simultaneously, with the point estimates in a single model (Maiti et al. (2014), Sugawara et al. (2015), Gershunskaya and Savitsky (2019)). For each domain i we consider data vector (y_i, v_i) of direct survey estimates. Assume the following model for $i = 1, \dots, N$:

$$y_i | \theta_i, \sigma_i^2 \stackrel{\text{ind}}{\sim} N(\theta_i, \sigma_i^2), \quad (3)$$

$$\theta_i | \beta, \tau_u^2 \stackrel{\text{ind}}{\sim} N(\mathbf{x}_i^T \beta, \tau_u^2). \quad (4)$$

$$v_i | a, \sigma_i^2 \stackrel{\text{ind}}{\sim} G\left(\frac{a n_i^*}{2}, \frac{a n_i^*}{2 \sigma_i^2}\right), \quad (5)$$

$$\sigma_i^2 | \gamma \stackrel{\text{ind}}{\sim} IG(2, \exp(z_i^T \gamma)), \quad (6)$$

where $G(\alpha, \beta)$ and $IG(\alpha, \beta)$ denote the gamma and inverse gamma distributions with shape parameter, α and scale parameter β , respectively.

The shape and scale parameters of the gamma distribution depend on an *unknown* parameter a and domain sample size n_i (we use standardized values for sample sizes,

$$n_i^* = \left(n_i - \left\{ \min_i n_i - 1 \right\} \right) / \left(\max_i n_i - \min_i n_i \right) \in [0, 1]).$$

Note that $a n_i^*$ controls the variation in the likelihood around latent true variance σ_i^2 , where domains with relatively larger values for n_i^* will have smaller variation, which accords with the notion that the observed variances v_i are of relatively higher quality for domains with larger sample. Parameter a allows the data to estimate a uniform scale of this sample size effect. Equation 6 denotes the prior distribution for the true variance σ_i^2 ; the true variance has an inverse gamma prior distribution where the mean depends on vector of covariates z_i through $\exp(z_i^T \gamma)$, where γ are unknown parameters. We label this model as “FHV”.

2.3 Relaxing linearity assumptions using finite mixtures

We now introduce our main model from Savitsky and Gershunskaya (2022). Here, we assume that Equation 3 holds and relax the parametric assumptions of Equations 4, 5,

and 6 by specifying finite mixtures in place of each of their respective distributions,

$$\theta_i | \pi, \mu, \beta, \tau_u^2 \stackrel{\text{iid}}{\sim} \sum_{k=1}^K \pi_k N(\mu_k + \mathbf{x}_i^T \beta, \tau_u^2) \quad (7)$$

$$v_i | a, \sigma_i^2, \mathbf{b}, \pi \stackrel{\text{ind}}{\sim} \sum_{k=1}^K \pi_k G\left(\frac{an_i^*}{2}, \frac{an_i^*}{2b_k\sigma_i^2}\right) \quad (8)$$

$$\sigma_i^2 | \gamma, \pi \stackrel{\text{ind}}{\sim} \sum_{k=1}^K \pi_k IG(2, \exp(z_i^T \gamma_k)). \quad (9)$$

The maximum number of clusters, K , is over-specified and denotes a truncated model. The actual number of discovered clusters, $K^* \leq K$, is governed by the following prior formulations,

$$\pi_k | \alpha \sim D\left(\frac{\alpha}{K}, \dots, \frac{\alpha}{K}\right) \quad (10)$$

$$\alpha \sim G(a_\alpha = 1, b_\alpha = 1) \quad (11)$$

$$\gamma_k \stackrel{\text{ind}}{\sim} N_P(\mu_\gamma = 0, \Sigma_\gamma = \mathbb{I}_P) \quad (12)$$

$$b_k \stackrel{\text{ind}}{\sim} G\left(\frac{\beta_b}{2} - 1, \frac{\beta_b}{2}\right), \quad (13)$$

where $D(\cdot)$ denotes the Dirichlet distribution, and \mathbb{I}_P denotes the $P \times P$ identity matrix. Equation 10 encourages sparsity in the number of clusters/components discovered and makes the setup equivalent to a truncated Dirichlet process. We set hyperparameter, $\beta_b = 4$ in the prior for b_k and P denotes the number of predictors in z_i .

The formulation is completed by utilizing the specification of a matrix of continuous predictors that allows the prior distributions for any two domains to cluster together to be influenced by predictors, \mathbf{g}_i , so that two domains, i_1 and i_2 , that have similar predictor values, \mathbf{g}_{i_1} and \mathbf{g}_{i_2} would be assigned a higher probability to cluster together, *a priori*. We accomplish this setup by placing a further prior distribution on predictors, \mathbf{g}_i with,

$$\mathbf{g}_i | \mu_g, \Sigma_g, \pi \stackrel{\text{ind}}{\sim} \sum_{k=1}^K \pi_k \mathbf{N}_l(\mu_{gk}, \Sigma_{gk}), \quad (14)$$

where we influence the co-clustering of domains by incorporating mixture parameters, (μ_{gk}, Σ_{gk}) , into the set of clustered parameters. Using predictors to influence the prior probability of co-clustering parameterizes a predictor-dependent clustering formulation and is another way to use predictors beyond a regression relationship with the response variable. We do *not* believe the predictors, \mathbf{g}_i , are random and are not interested in the fitted values of parameters μ_{gk}, Σ_{gk} , but fitting \mathbf{g}_i and including the μ_{gk} and Σ_{gk} in the clusters along with μ_k, b_k, γ_k helps better identify the clustering structure to give sharper estimates. Placing a probability model on \mathbf{g}_i is a device that allows the use of the simple truncated mixture to parameterize a predictor-dependent clustering model (Müller et al., 2011).

Finally, parameters b_k allow the data to estimate a bias in the observed variances v_i (for those domains i linked to cluster k). The bias b_k multiplicatively offsets the estimated true variance, σ_i^2 . (See Savitsky and Gershunskaya (2022) for more details.) We label our resulting model with “CFHV” to denote a clustering version of the Fay-Herriot that also co-models the variance.

3. Resampling-based Calibration

We use the ADVI estimation algorithm (Kucukelbir et al. (2017)) implemented in Stan to estimate the mixture model described in Section 2.3. This algorithm implements a mean field variational Bayes approximation for the posterior distribution characterized by a fully factorized normal distribution (on a transformed parameter space such that each transformed parameter lies on the real line). It has been reported in the literature (Yao et al. (2018)) that the mean field approximation may produce inadequate measures of uncertainty for model parameters, even if the resulting posterior means are unbiased estimates of parameters.

In this section we consider a re-sampling, bootstrap-like procedure (that we will refer to as the “bootstrap” procedure in the sequel) to evaluate and re-calibrate potential bias in the second moment of the estimation algorithm that improves the estimation accuracy of parameter variances. BLS publishes both point and variance estimates for CES domains, where each variance connotes a measure of estimation quality and its publication is facilitated by our proposed re-calibration procedure. We illustrate how to use our re-calibration procedure to also correct bias in the first moment, though use of approximate estimation algorithms like ADVI are typically predicated on producing high quality point estimates (from first moments). We further demonstrate how to use re-calibrated first and second moments to adjust confidence intervals that ensure the nominal coverage at any level of misspecification γ .

The bootstrap procedure is now described in the following sequence of steps:

1. Estimate a hierarchical model (e.g., of Section 2.3) using the ADVI algorithm on the given data. Let $m(\theta_i)$ and $v(\theta_i)$ denote the approximate posterior mean and variance, respectively, of parameter θ_i for domain i obtained from the initial model estimation using ADVI (on the untransformed parameter space in Stan). Extract random *draws* of model parameters $\theta_i^{(\alpha)}$ indexed by $\alpha = (1, \dots, A)$ from the approximate posterior distribution of model parameters θ_i from the initial model run with ADVI. Use these draws, $\theta_i^{(\alpha)}$ to, in turn, generate draws of replicate data $y_i^{(\alpha)}$ from the approximate posterior predictive distribution of y_i conditioned on the parameter draws. This procedure constructs A re-sampled data sets for each domain, $i \in 1, \dots, n$.

In each data set, obtained from draw α , we treat the associated $y_i^{(\alpha)}$ as “observed” data, and $\theta_i^{(\alpha)}$ as “bootstrap true” parameters for $i = (1, \dots, n)$.

2. Estimate the model *again* under the same ADVI algorithm for *each* re-sampled data set $(y_i^{(\alpha)})_{i=1}^n$, indexed by $\alpha = 1, \dots, A$. We obtain approximate posterior means, $m(\theta_i^{(\alpha)})$ and variances, $v(\theta_i^{(\alpha)})$ for respective parameters from each model estimation on data set α .
3. Use model estimations on the A re-sampled data sets to adjust the first moment, as follows:

The bias of the model fitted approximate posterior mean $m(\theta_i)$ of parameter θ_i is expressed with,

$$a_i = m(\theta_i) - \bar{m}(\theta_i^{(\alpha)}),$$

where $m(\theta_i)$ is the approximate posterior mean from the initial run and $\bar{m}(\theta_i^{(\alpha)}) = A^{-1} \sum_{\alpha=1}^A m(\theta_i^{(\alpha)})$ is the bootstrap mean, obtained as the average of approximate posterior means, $m(\theta_i^{(\alpha)})$, over the A replicate runs. The first order bias adjusted posterior mean for θ_i is

$$\tilde{m}(\theta_i) = m(\theta_i) + a_i. \tag{15}$$

4. Use the A model estimations to formulate a second moment adjustment that we next derive from properties of the pivot statistic,

$$T_i = \frac{m(\theta_i) - \theta_i}{\sqrt{v(\theta_i)}}. \tag{16}$$

If $m(\theta_i)$ and $v(\theta_i)$ are consistent estimates of the domain point estimate and variance, respectively, then the pivot is distributed as a t-statistic with $(0, 1)$ mean and variance, respectively, that contracts on a standard normal distribution. We write that $T_i \sim (0, 1)$ as a shorthand notation that denotes this statistic is distributed under a symmetric distribution with mean 0 and variance 1. We next use this standardization of the pivot statistic to derive our adjustment to the estimated variance.

We know that $v(\theta_i)$ is *not* a consistent estimator of the true variance (of the numerator of the pivot) when estimated from the ADVI algorithm. We define variance scale adjustment, c_i , under an assumption that the variance correction factor is of a multiplicative scaling formulation,

$$\text{Var}(\theta_i) = v(\theta_i)c_i,$$

where $\text{Var}(\theta_i)$ denote the true (posterior) variance of the numerator in Equation 16. We estimate the pivot statistic of Equation 16 from its distribution over the A bootstrap samples (each indexed by $\alpha \in 1, \dots, A$) using,

$$T_i^{(\alpha)} = \frac{m(\theta_i^{(\alpha)}) - \theta_i^{(\alpha)}}{\sqrt{v(\theta_i^{(\alpha)})}}. \tag{17}$$

- (a) We first focus on correcting or adjusting variance of the bootstrap runs, $v(\theta_i^\alpha)$, where the variance of the original model run, $v(\theta_i)$, is treated as the truth. The bootstrap variances may be corrected back to the original model run using, $v(\theta_i) = v(\theta_i^\alpha)c_i$. Thus, we divide pivot T_i^α by $\sqrt{c_i}$ to correct the bootstrap variances in the pivot statistic back to the bootstrap truth.
- (b) However, since we draw our bootstrap samples from an approximate posterior distribution with *biased* variance $v(\theta_i)$, we must again divide pivot T_i^α by $\sqrt{c_i}$, to correct back to the *true* posterior variance, $\text{Var}(\theta_i)$.

We use the standardization property of the pivot under use of the correct variance to derive our adjustment from,

$$\text{Var}(c_i^{-1}T_i^{(\alpha)}) = 1.$$

which gives us $c_i^2 = \text{Var}(T_i^{(\alpha)})$ for the form of c_i . This form for c_i , in turn, produces, $c_i = \sqrt{A^{-1} \sum_{\alpha=1}^A (T_i^{(\alpha)} - \bar{T}_i)^2}$, where $\bar{T}_i = A^{-1} \sum_{\alpha=1}^A T_i^{(\alpha)}$.

5. *Adjusted confidence intervals (II): the pivotal approach*

The adjusted pivot after performing the variance correction using the estimated c_i in the previous step is,

$$\tilde{T}_i^{(\alpha)} = \frac{T_i^{(\alpha)} - \bar{T}_i}{c_i} \sim (0, 1),$$

which we invert to produce the calibrated level- γ confidence intervals. Here γ denotes the misspecification probability for target $1 - \gamma$ confidence intervals and $\tilde{t}_{i,\gamma}$ is the γ quantile statistic of the empirical distribution, $\tilde{T}_i^{(\alpha)}$ indexed by α .

$$C_i(\gamma) = \left[\tilde{m}(\theta_i) + \sqrt{v(\theta_i)c_i\tilde{t}_{i,\gamma}}, \tilde{m}(\theta_i) + \sqrt{v(\theta_i)c_i\tilde{t}_{i,1-\gamma}} \right], \quad (18)$$

6. *Adjusted confidence intervals (II): direct re-scaling approach.* We can obtain adjusted confidence intervals for fitted θ_i by rescaling each draw θ_i^* from the approximate posterior distribution of θ_i produced by the ADVI algorithm to match the first two moments:

$$\tilde{\theta}_i^* = \frac{\theta_i^* - m(\theta_i)}{\sqrt{v(\theta_i)}} \sqrt{v(\theta_i)c_i} + \tilde{m}(\theta_i), \quad (19)$$

where we use the same c_i derived from standardizing the bootstrap indexed pivot to scale adjust the approximate posterior variance, $v(\theta_i)$. The adjusted confidence intervals are obtained by computing percentiles over the adjusted draws $\tilde{\theta}_i^*$.

4. Simulation study

4.1 Calibration under the Fay Herriott (FH) Model

In the first part of our simulation study, we generate data from the relatively simple Fay-Herriot model and then fit the data using ADVI, evaluate and adjust the variance to produce revised confidence intervals using the two methods to construct intervals proposed in Section 3. We consider $N = 150$ domains and let true model parameters be $\beta = 1$, $\tau_u^2 = 1$, $\sigma_i^2 = 1$; we generate covariates as $x_i \stackrel{\text{ind}}{\sim} \text{Unif}(0, 2)$, random effects as $u_i \stackrel{\text{ind}}{\sim} N(0, \tau_u^2)$, and random errors as $\epsilon_i \stackrel{\text{ind}}{\sim} N(0, \sigma_i^2)$. True domain values are $\theta_i = \mathbf{x}_i^T \beta + u_i$ and “direct” domain estimates are $y_i = \theta_i + \epsilon_i$. For this setup, variances of y_i are assumed to be observed and measured exactly ($v_i = \sigma_i = 1$.)

We run the model and extract $S = 200$ *simulation* data sets from the posterior distribution of θ_i and posterior predictive distribution of y_i , where S denotes the number of Monte Carlo simulation iterations (that allows us to assess coverage of our rescaled estimator). Next, we perform the re-sampling algorithm $A = 500$ times within *each* of the $S = 200$ Monte Carlo simulation data sets.

To assess the coverage properties of the confidence intervals (CIs), we compute the number of times, over $S = 200$, that the CI - before and after adjustment - covers the true parameter value. We implement the pivotal and rescaling based adjustments solely to adjust the second moment and not the first moment of the model estimation algorithm. So, we use the posterior means, $m(\theta_i^{(\alpha)})$ (in place of $\tilde{m}(\theta_i^{(\alpha)})$) to form revised confidence intervals. Our experience reveals that the ADVI produces accurate point estimates for our

	Orig Fitted	Rescaled	Pivot
Coverage	0.531	0.493	0.492
Length	1.019	0.935	0.933

Table 1: Coverage properties for model fitted $m(\theta_i)$, 50% nominal, over 150 domains and $S = 200$ simulation runs, using $A = 500$ re-samples

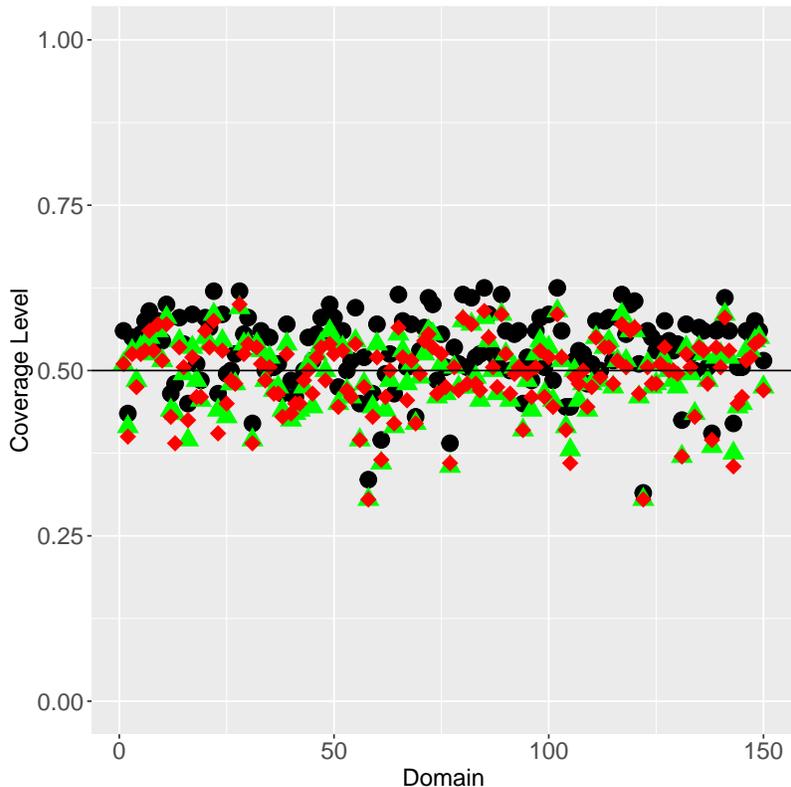


Figure 1: Domain coverages (50% nominal) of fitted point estimates vs number of respondents. Black dots represent originally fitted model; green triangles represent adjusted coverages using rescaling method; red diamonds represent adjusted coverages based on the pivotal method.

CES models of Section 2.3. Including the first order bias correction in this case where ADVI produces relatively accurate point estimates induces added noise in the calibrated confidence intervals.

Results for 50% nominal coverage, averaged over all 150 domains for $S = 200$ simulation runs, are presented in Table 1. The original fit produces slight overcoverage; both rescaled and pivotal based adjustment methods give nearly nominal coverages. The lengths of the adjusted intervals are shorter than the original lengths.

In Figure 1, we plot coverages for *each* domain. The black dots represent coverages from the original fit, the green triangles and red diamonds correspond to the rescaled and

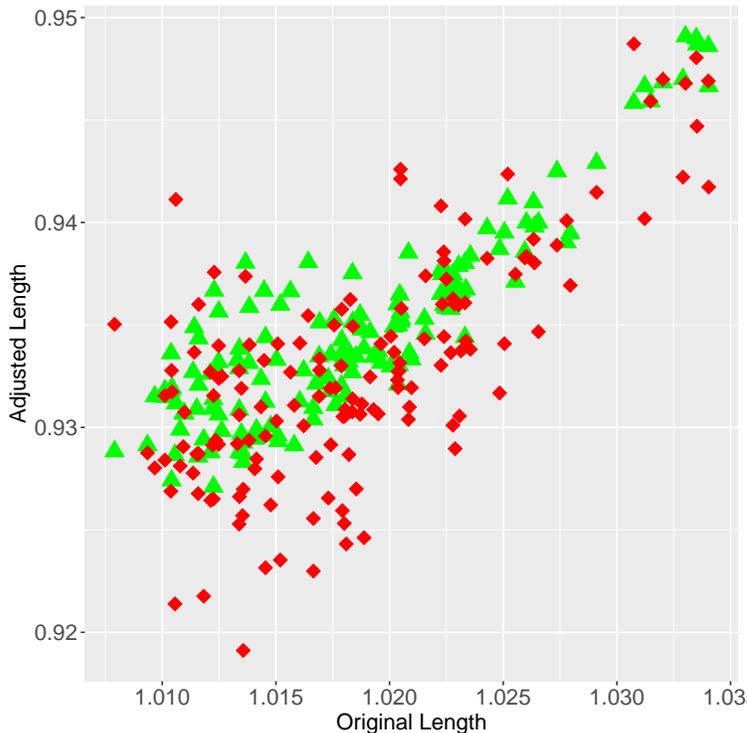


Figure 2: Lengths of the adjusted vs the original (50% nominal) confidence intervals. Green triangles represent adjusted coverages using rescaling method; red diamonds represent adjusted coverages based on the pivotal method.

pivotal based adjustment methods, respectively. The adjustments bring coverages closer to the nominal 50% line.

In Figure 2, we plot the length of the adjusted CIs against the lengths of the original CIs. The lengths of CIs from the rescaled and pivotal based adjustment methods (green triangles and red diamonds, respectively) are shorter than the original CIs. The pivotal based CIs are somewhat noisier but a little shorter than the rescaled-based CIs.

To check the performance of the method under a different scenario, we conducted additional simulations, where the data are generated by setting the true value of σ_i^2 parameter to 2. The variation among domains remain essentially unchanged. We include results of this sensitivity study in Appendix B.

4.2 Calibration under the clustering CFHV model

In the second part of our simulation study, we evaluate the proposed calibration procedure for a more complicated CFHV model, given in Section 2.3. For this part of the study, simulation data sets are generated as follows: We first estimate the mixture model of Section 2.3 on an available *real* data set and extract $S = 200$ random draws from the posterior distribution of model parameters. These can be viewed as “true” values of the model parameters for data set s , $s = 1, \dots, S$. We use the same initial run to draw “observed”

	Orig Fitted	Rescaled	Pivot
Coverage	0.625	0.559	0.549
Length	0.618	0.537	0.528

Table 2: Coverage properties for model fitted $m(\theta_i)$, 50% nominal, over 166 domains and $S = 200$ simulation runs, using $A = 500$ re-samples

data $(y_i^{(s)}, v_i^{(s)})$ from the posterior predictive distribution of (y_i, v_i) . These data, along with covariates, domain labels, number of respondents from the original data set form our S simulated data sets. As noted, the data set is based on real domains of various sizes with the number of respondents ranging from 1 to 247 over the total of 166 domains.

For each simulation data set s , we generate $A = 500$ re-sampled data sets obtained from a model estimation run and perform the calibration procedure as described in Section 3.

We use posterior draws from respective $A = 500$ model estimation runs (for each data set, $s \in 1, \dots, S$) to construct CIs for each model parameter and compute the number of times, over $S = 200$, that the CI - before and after adjustment - covers the true parameter value.

We present coverage properties (for 50% nominal coverage) of the confidence intervals for model fitted $m(\theta_i)$ in Table 2. The results in the table are averaged over all domains. One notices the overcoverage of the original (“Orig Fitted”) intervals obtained from the model posterior draws. The adjustment using the rescaling and pivotalbased approaches brings the coverage closer to the nominal.

To complement the overall results presented in Table 2, we present plots showing the distribution of the coverages over domains. We show coverage properties for *each* domain of the 50% confidence intervals in Figure 3 computed over the S simulation runs - both before adjustment and after adjustment under the rescaling and pivotal based approaches. Coverage for point estimates is plotted on the vertical axis vs the number of respondents in a given domain, on the horizontal axis where the domains are ordered according to an increasing number of sampled establishments (sampled in a domain) from left-to-right. Black dots represent the original model results, green triangles show results from the rescaling method, and red diamonds represent the pivotal based adjusted coverages. We observe overcoverage of the original model fit for most of domains. The rescaling and pivotal based approaches produce similar results for dispersion of domain-indexed coverages around the truth 50% nominal, though the pivotal based coverage results are slightly more compactly distributed along the horizontal line representing the nominal 50% coverage.

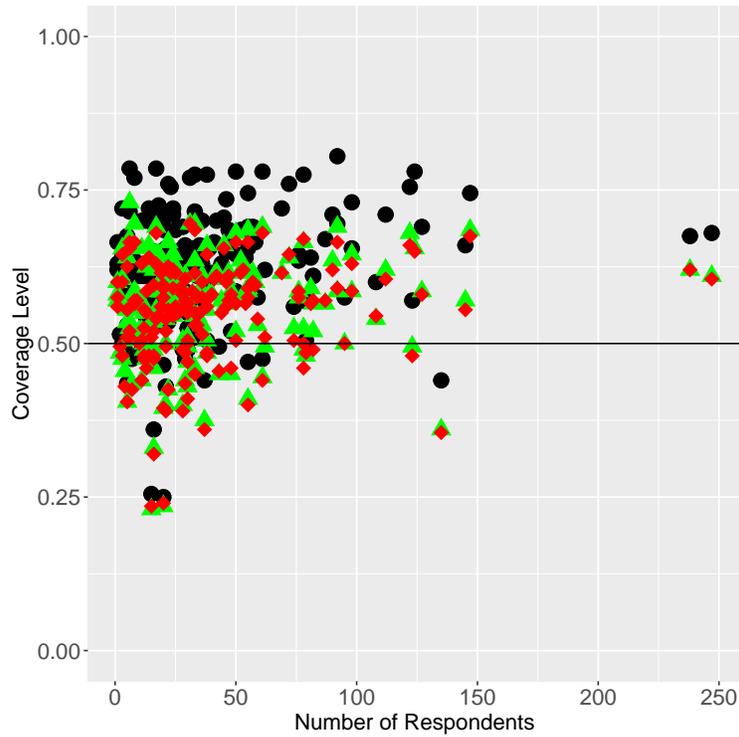


Figure 3: Domain coverages (50% nominal) of fitted point estimates vs number of respondents. Black dots represent originally fitted model; green triangles represent adjusted coverages using rescaling method; red diamonds represent adjusted coverages based on the pivotal method. Domains are ordered along the horizontal axis by increasing number of sampled establishments from left-to-right.

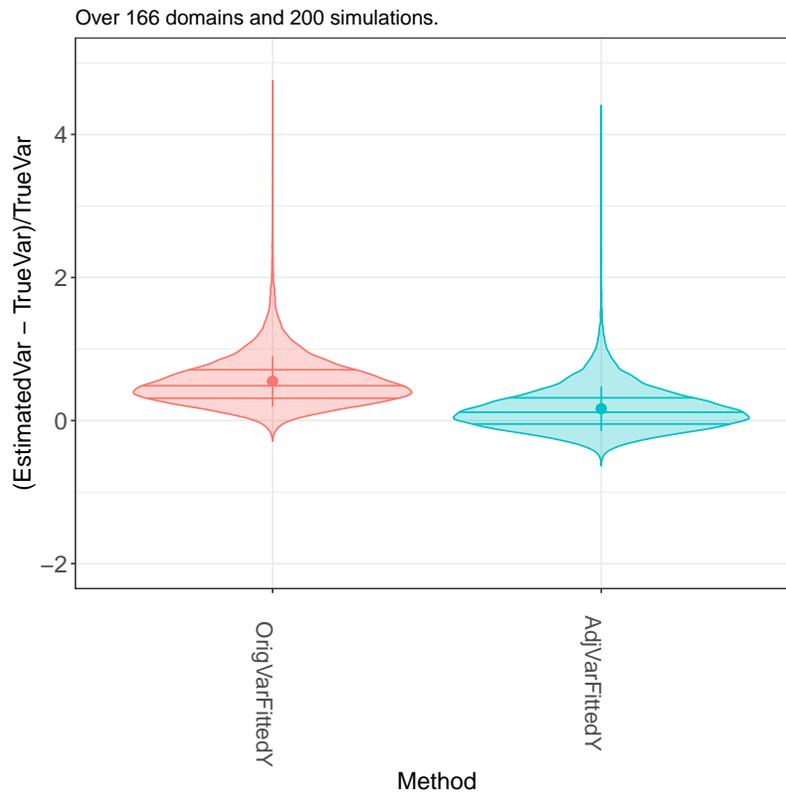


Figure 4: Distributions, over 166 domains and 200 simulation runs, of residuals of model posterior variances $v(\theta_i)$ and adjusted variances $\tilde{v}(\theta_i)$, relative to respective true variances of θ_i . ("Original" - model approximate posterior variances; "Adjusted" - adjusted variances).

The violin plots in Figure 4 show distributions over domains and simulation runs of relative residuals of $v(\theta_i)$, the model approximate posterior variances of θ_i , and also the adjusted variances $\tilde{v}(\theta_i)$ under the pivotal calibration method. The original model posterior variances of θ_i are upwardly biased and the bias is corrected in the calibrated version of the variances.

In Figure 5, we plot lengths of the adjusted confidence intervals against the lengths of the original CIs. Each symbol corresponds a domain. The adjusted intervals are shorter than the original ones, with the pivotal based method being slightly shorter than the rescaling based.

Although our motivation for using a variance calibrated version of ADVI derives from the inability to estimate the clustering / mixture model in HMC, it bears mention that another motivation for using some optimization-based approximation (that may not be ADVI) to the joint posterior distribution may be to reduce computation time for estimation; however, in our case, if we set the number of mixtures in our CFHV model to be $K = 1$ to avoid label switching and allow HMC estimation, HMC computes much faster than ADVI + calibration. The $K = 1$ model reduces to a relatively simple parametric model where an HMC estimation run on our small domain estimation model takes a matter of minutes (versus many hours to conduct all of the ADVI re-estimations to derive the calibration scale, c). So, if one is motivated by the goal to reduce computation time in their choice of an optimization-based approximation and subsequent calibration of uncertainty intervals, our experience suggests that a computation benefit would require a relatively complex, semi-parametric formulation and/or a relatively larger data set.

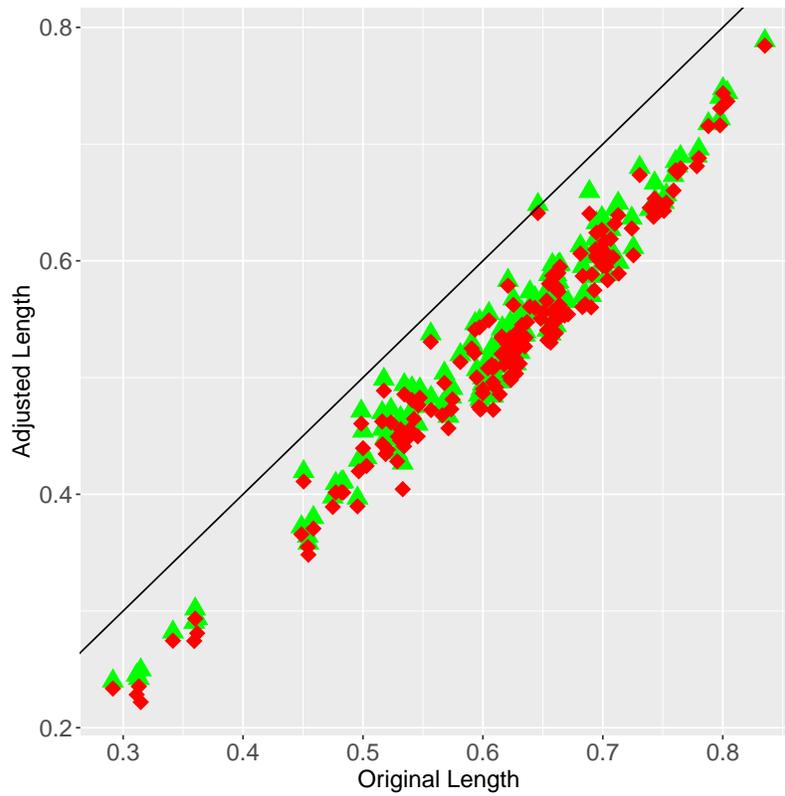


Figure 5: Lengths of the adjusted vs the original (50% nominal) confidence intervals. Green triangles represent adjusted coverages using rescaling method; red diamonds represent adjusted coverages based on the pivotal method.

5. Convergence of Adjusted Uncertainty Intervals

We next discuss a convergence result of our corrected γ -level uncertainty intervals to nominal coverage at $\mathcal{O}(N^{-1})$ rate where we recall that N denotes the number of domains. We leverage existing theoretical results from the bootstrap literature to assess the convergence properties of uncertainty intervals produced from our simulation-based calibration procedure. To leverage these results, we require two conditions: Firstly, the specified *model* must be asymptotically correct for data generation because our simulation-based calibration procedure uses a parametric bootstrap since our interest focuses on inference for model parameters. (See Ghosal and van der Vaart (2017) for assumptions that guaranty the consistency of a Bayesian model. Two important assumptions are: 1. Bounded entropy that regulates the complexity of the model; 2. The joint prior distribution is required to place positive mass on the truth.) Secondly, we require first-order consistency of the *estimation* algorithm, although the converged value may express a constant or fixed bias (e.g., If the (unknown) true value for parameter θ_i is denoted by $\theta_{0,i}$, the estimation algorithm may contract onto fixed point, $\theta_{a,i} \neq \theta_{0,i}$).

We compare the convergence of the adjusted marginal credibility interval of the pivotal method to that obtained with the scale correction method. The pivotal method is constructed from a re-sampled distribution, while the scale correction method performs a scale adjustment to each parameter approximate marginal posterior distribution.

We formulate a pivot statistic for each re-sampled data set, $\alpha \in (1, \dots, A)$ with scale adjustment factor c_i^2 multiplying the model posterior variance, $v(\theta_i^{(\alpha)})$ in the denominator of the pivot indexed by re-sample iteration α . The quantiles of the resulting distribution of the pivot statistic over the re-sampled iterations are used to form a confidence interval for θ_i under our pivotal method. So, these intervals are based on the re-sampling/empirical distribution (of the pivot statistic), not the approximate model posterior distribution (of θ_i) obtained from the scalable estimation algorithm, such as ADVI. We replace the approximate posterior distribution with the empirical distribution to formulate a calibrated confidence interval for θ_i .

Efron and Tibshirani (1993) discuss that the convergence of the pivotal method intervals under a bootstrap re-sampling procedure are “second order accurate”; that is $\mathbb{P}(\theta_i \in C_i(\gamma)) \rightarrow 1 - \gamma$ at an $\mathcal{O}(1/N)$ rate. Our re-sampling procedure that draws replicate data from the model posterior predictive distribution may be viewed as equivalent to the bootstrap procedure conditioned on the correctness of the underlying model used to estimate the posterior predictive distribution.

Efron and Tibshirani (1993) notes that the variance of the distribution of the pivot statistic (that we will refer to as “pivot distribution”) is independent of θ_i such one does not have to estimate the variance by plugging in $\hat{\theta}_i$. Therefore, the variance of the pivot distribution converges to the truth at a relatively faster (higher order) $\mathcal{O}(1/N)$ rate (versus the slower $\mathcal{O}(1/\sqrt{N})$ rate for the percentiles of the underlying bootstrap distribution that require plugging in $\hat{\theta}_i$, which converges to θ_i at $\mathcal{O}(1/\sqrt{N})$). The convergence of the higher distribution moments (e.g., skewness and kurtosis) are $\mathcal{O}(1/N^{3/2})$ and $\mathcal{O}(1/N)$, respectively. Therefore, the resulting pivotal based confidence intervals (realized from the distribution governed by repeatedly drawing samples from the parametric model) converge to nominal coverage at a second order rate without requiring convergence to a normal distribution. So,

this rate will be faster than the convergence of intervals that require convergence of the estimated distribution to a normal distribution for intervals to achieve nominal coverage. Chen et al. (2018) demonstrates this same result for intervals based on the pivot statistic when using variational inference to obtain point estimates for θ_i .

Our re-scaling option, by contrast with the pivotal method, uses the variance correction factor, c_i , to directly perform a re-scaling (variance correction) of the draws taken from the approximate posterior distribution for θ_i . The convergence of the adjusted posterior distribution to the truth (required for convergence of the adjusted credibility intervals to achieve nominal coverage) requires that the adjusted posterior distribution converges to a normal distribution. Even where the adjusted posterior distribution converges to a normal distribution, we expect the convergence of the associated credibility intervals to be at a slower first order ($\mathcal{O}(1/\sqrt{N})$) rate.

The faster convergence of the pivotal based confidence intervals as compared to the re-scaled credibility intervals explains why their domain-indexed ($i \in (1, \dots, N)$) coverages are more tightly clustered around the $1 - \gamma = 50\%$ target for the former as compared to the those of the latter in Figure 3.

Our motivating application on estimation of CES point estimates, θ_i , determines our choice of the ADVI (mean-field variational Bayes) algorithm estimation. BLS use ADVI to computationally scale the model estimations to support the monthly production process for CES employment estimates. Using the re-scaling option to calibrate the uncertainty intervals requires that the underlying ADVI produces approximate posterior distributions that converge to a normal distribution. Wang and Titterington (2004) show that the mean field variational Bayes approximate posterior distribution *does* achieve asymptotic normality for parametric models of the exponential family. However, Zhang and Gao (2020) demonstrate that the distribution error has two parts: 1. Convergence of the true posterior distribution; 2. Variational Bayes error. For nonparametric and high dimensional models, a specific prior formulation must be used such that the first error source dominates the second and asymptotic normality is achieved. However, if a prior distribution of this specific form is not used, the approximate posterior distributions will not converge to a normal distribution. In this case, the second order correction will not resolve errors in higher moments and the resulting adjusted credibility intervals will not achieve nominal coverage, no matter how large is the sample size.

Even though it may be the case that most model formulations will produce approximate posterior distributions that converge to normal distributions under use of mean-field variational Bayes, there is yet some possibility that such will not be the case for a specific model formulation suggests that one may not universally rely on uncertainty intervals formed from the adjusted (calibrated) posterior distribution. Replacing the approximate posterior distribution with use of the empirical distribution over the re-samples avoids this issue of convergence to normality that is tied to the approximation algorithm and may be universally relied on. It bears mention that there is a computational cost to re-estimate the model for each re-sampled data set and there is some minimum (e.g., $A = 100$) re-samples required to control Monte Carlo estimation error to estimate the empirical distribution over the re-sampled parameter estimates (though the replicate model re-estimations may be parallelized if computational resources permit).

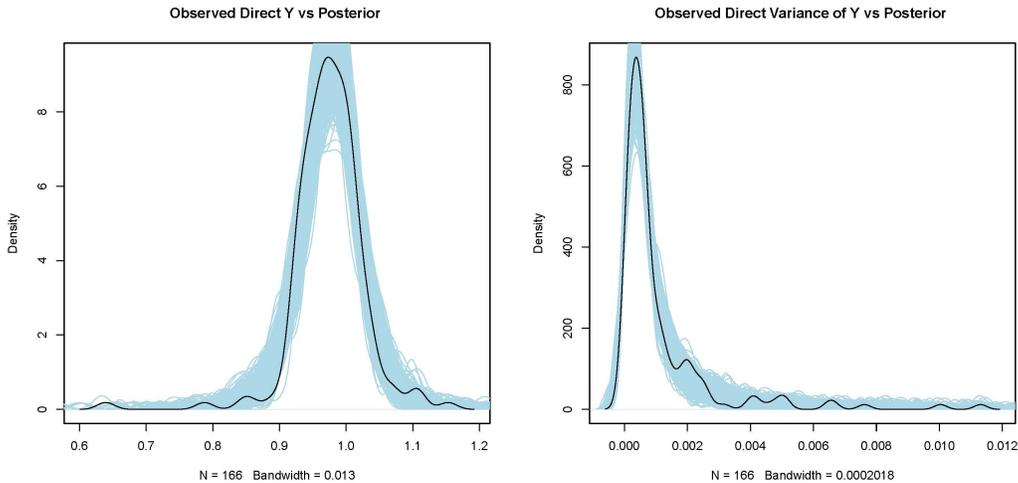


Figure 6: Posterior predictive check for (a) direct estimates y_i and (b) observed variances of v_i of direct estimates y_i . The black line in each plot represents the distribution of the real data.

Lastly, our simulation-based calibration procedure may be used with *any* estimation algorithm that propagates its error over multiple data estimations to calibrate confidence intervals that achieve nominal coverage. While we have focused on ADVI, there are other computationally scalable algorithms that may be selected to estimate an approximation to the model posterior distribution (Vehtari et al., 2020). Convergence of the approximate posterior distributions to (true) normal distributions would have to be verified for each estimation algorithm and class of models if one desires to re-scale the approximate posterior distribution to calibrate the uncertainty intervals.

6. Application of simulation-based calibration to real data

6.1 Model fit checking and calibration procedure

We now apply the calibration procedure to real data from the Current Employment Statistics (CES) survey, where we use our area-level mixture model described in Section 2.3 to obtain better estimates of employment levels in 166 domains within the “Leisure” industrial sector. As a first step, we fit the model and extract draws from the model posterior predictive distribution of y_i and v_i .

Before performing simulation-based calibration, we conduct posterior predictive checks to assess the correctness of the model formulation based on how well it preserves the distribution properties of the real data. We examine the marginal distributions over domains $i \in (1, \dots, N)$ for replicate draws of each of $(y_i^{(\alpha)})_{i=1}^N$ and $(v_i^{(\alpha)})_{i=1}^N$ as well as examine the joint distributions, $(y_i^{(\alpha)}, v_i^{(\alpha)})_{i=1}^N$. We only expect our simulation-based calibration procedure to work if the specified model well reproduces the distribution properties of the real data.

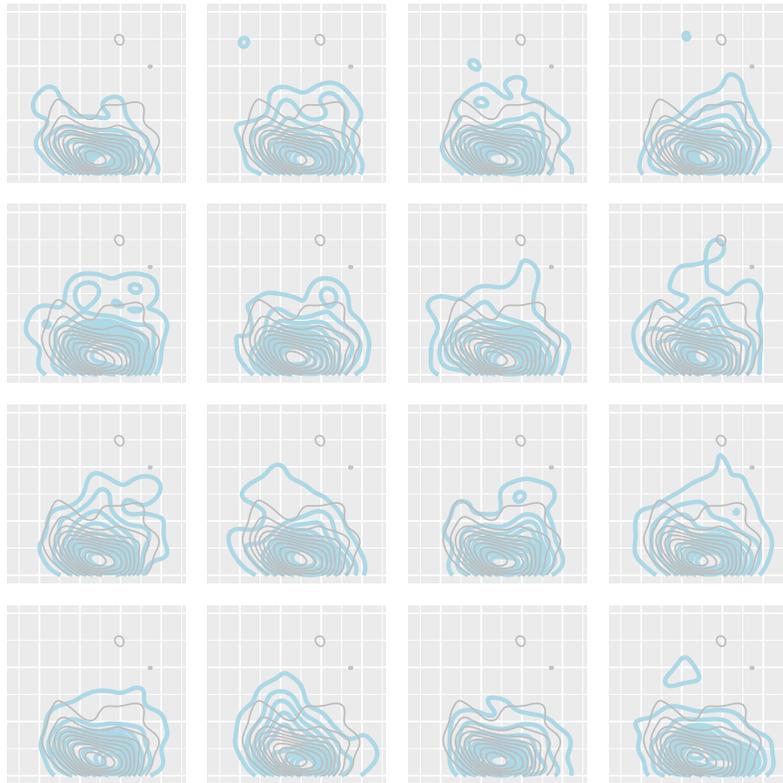


Figure 7: Posterior predictive check for pairs (y_i, v_i) , based on 16 random draws from the posterior predictive distribution. Each of the 16 panels represents a single vector draw (in light blue) from the posterior predictive distribution, overlaid by the observed distribution (gray) over domains.

Plots in Figure 6 demonstrate that the model provides reasonable fit for observed y_i and v_i . Next, we briefly check if the model also preserves correlation between vectors y_i and v_i . For this, we use a contour plot to show the density of two-dimensional vector (y_i, v_i) over the set of domains. On the same plot, we include the analogous density of vector $(y_i^{(d)}, v_i^{(d)})$ obtained by randomly drawing from the posterior predictive distribution of (y_i, v_i) . The plot panels in Figure 7 show $D = 16$ sets of such contour plots corresponding to $D = 16$ draws from the posterior predictive distribution overlaid by the observed distribution of (y_i, v_i) . The contour plots demonstrate that correlation between y_i and v_i in the simulated data is similar to the observed data. We, therefore, conclude that the model describes the data reasonably well.

We now save $A = 100$ random draws $(y_i^{(\alpha)}, v_i^{(\alpha)})$, $\alpha = 1, \dots, A$, from the posterior predictive distribution of (y_i, v_i) and compute calibration adjustments as described in Section 3.

There is a practical concern that it would not be feasible to perform calibration in the tight real-time monthly CES production environment. To overcome this difficulty, we conjecture that the CES calibration adjustments would be similar for each month in a given industry because the underlying model formulations and data structures are highly similar. If this is the case, our plan for the real-time application of the method would be to derive a set of adjustments once a year based on past year historical data, before the start of a new production cycle for the current year. These historical-based monthly adjustments would be averaged over the 12 months of the previous year, and the average values would be applied in the upcoming monthly production cycle.

To test the proposed procedure, we save another $B = 100$ sets of draws from our initial model run, in addition to the $A = 100$ sets already used to derive the calibration adjustments. We then run the model once for each of the B data sets and apply our average adjustments to check coverage properties. The procedure is summarized below.

1. Use historical CES monthly data and fit the model for each month of the year. Call it the “initial” model run.
2. Produce $A = 100$ re-sampled data sets using draws from the initial model run. Derive calibration adjustments from these data sets, as described in Section 3. Run the calibration procedure for each month of a year (to obtain calibrated intervals for each month).
3. Produce additional $B = 100$ “test” data sets using B additional draws from the initial model run. Fit the model once on each of the $b = 1, \dots, B$ test data sets using the same algorithm as in the initial model fit. For each model run, b , we use the associated $(y_i^{(b)}, v_i^{(b)})$ as “observed” data.
4. Apply calibration adjustments to estimates obtained from each of the B model runs.
5. For each b , construct confidence intervals for model fitted posterior means $m(\theta_i^{(b)})$. Compute the number of times (over B) that respective true values $\theta_i^{(b)}$ are covered by the interval. Perform the coverage computation for the originally fitted and adjusted estimates.

Month	Orig Fitted	Rescaled	Pivot
Jan	0.577 (0.750)	0.523 (0.668)	0.503 (0.642)
Feb	0.572 (0.582)	0.523 (0.520)	0.500 (0.500)
Mar	0.589 (0.531)	0.537 (0.474)	0.516 (0.455)
Apr	0.569 (0.663)	0.517 (0.591)	0.497 (0.566)
May	0.597 (0.511)	0.544 (0.456)	0.522 (0.437)
Jun	0.580 (0.573)	0.529 (0.511)	0.509 (0.491)
Jul	0.590 (0.585)	0.538 (0.521)	0.516 (0.500)
Aug	0.576 (0.710)	0.520 (0.633)	0.501 (0.608)
Sep	0.590 (0.469)	0.538 (0.419)	0.516 (0.402)
Oct	0.568 (0.661)	0.516 (0.589)	0.498 (0.565)
Nov	0.590 (0.608)	0.540 (0.542)	0.518 (0.520)
Dec	0.604 (0.542)	0.549 (0.483)	0.529 (0.464)

Table 3: Coverage properties of model fitted θ_i , 50% nominal, over 166 domains and 100 simulation runs

We tested the average adjustment factors on our $B = 100$ “test” data sets and evaluated for each of the 12 months. Results are presented in Section 6.2.

6.2 Results

We use draws from the VB posterior distribution of θ_i to construct credibility intervals; the adjusted confidence intervals are constructed using the rescaling and pivotal approaches of Section 3. For each domain i , we used adjustments averaged over 12 months and we test results separately for each month.

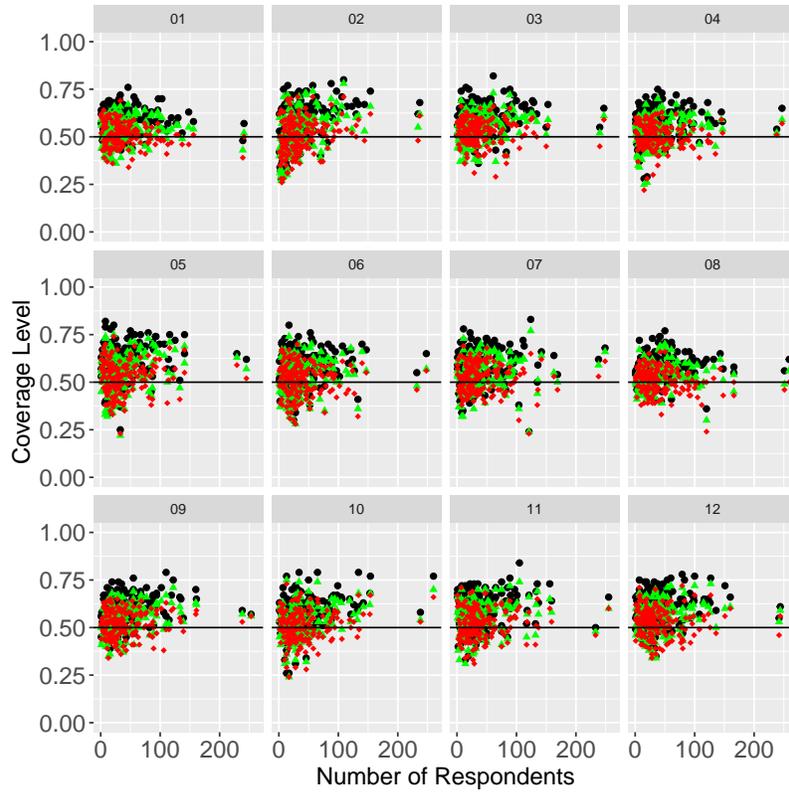


Figure 8: Domain coverages (50% nominal) of fitted point estimates vs number of respondents for each of the 12 months. Black dots represent originally fitted model; green triangles represent adjusted coverages using rescaling method; red diamonds represent adjusted coverages based on the pivotal method. Domains in each panel are ordered along the horizontal axis by increasing number of sampled establishments from left-to-right.

In Table 3, we show monthly test results of coverages, averaged over domains. Respective average CI lengths are shown in parenthesis. We note the overcoverage of the original model CIs, ranging from 57 to 60 percent covered, for the nominal 50%. The rescaled CIs provide improved coverage (from 52 to 55 percent); the pivotal method gives best close to the nominal results and the shortest CI lengths.

To complement the results averaged over domains shown in Table 3, we present monthly coverages for each domain in Figure 8. Percent covered for each domain in a given month is plotted against the number of respondents. The black horizontal reference line corresponds to the 50% nominal level. Black dots represent coverages corresponding to the original model fit. Similar to the simulations results of Section 4, we observe overcoverage across the domains in the originally fitted model results. The rescaling method (green triangles) corrects the overcoverage to a certain degree. Red diamonds corresponding to the pivotal based adjustment method lie closer around the horizontal reference line, which is consistent with our conclusion that the pivotal based adjustment gives the best performing method among the alternatives considered.

7. Discussion

This paper formulates a parametric re-sampling procedure to calibrate the outputs of a scalable, approximate Bayesian estimation algorithm where the algorithm produces accurate point estimates (e.g., approximate posterior means) but *incorrect* uncertainty quantification. Our simulation-based calibration procedure produces asymptotically correct marginal posterior variances. Our procedure requires that a scalable algorithm provide consistent point estimation (though there can be an asymptotically constant bias) under a Bayesian model that asymptotically contracts on the true generating model. It would be expected that practitioners would only use a scalable algorithm to the extent that it produced similarly high quality point estimation as do Markov Chain Monte Carlo (MCMC) algorithms.

Bayesian hierarchical models are favored for small domain estimation in BLS for their ability to allow the data to estimate complex correlation structures among domains that regulates shrinking to produce more accurate domain point estimates. A fast turnaround monthly BLS data publication cycle requires use of computationally scalable estimation algorithms. BLS seeks to publish both smoothed domain point estimates, θ_i , as well as their variances (of the θ_i) so that users may assess the quality of each point estimate. We use Stan’s (Carpenter et al., 2017) ADVI variational Bayes estimation algorithm and achieve high quality point estimation, but the variances are incorrect such that uncertainty intervals do not achieve nominal coverage.

Our application of the simulation-based calibration procedure requires re-estimation of the model on the order of $A = 100$ times, which reduces the scalability of the estimation algorithm (unless one possesses the computational resources to parallelize the A model runs). Yet, we show that under a regularly produced data set for a collection of domains that computation of the mean and variance adjustment factors may be performed once and applied for each month throughout the year because the model formulation and data structures that together determine the calibration adjustments are largely unchanged over the months. This procedure produces uncertainty intervals on a monthly basis that provide much better coverage than unadjusted intervals.

Our derivation of a calibration scale adjustment factor for each domain, c_i , by standardizing the distribution of a pivot statistic formed over replicate model runs provides a fast and easy numerical procedure.

Finally, an alternative simulation-based procedure calibrates the first and second moments using replicate data drawn from the *prior* predictive distribution (as contrasted with our approach that calibrates data drawn from the approximate *posterior* predictive distribution) (Talts et al., 2020). Calibrating from the prior predictive does not require an initial model estimation using the approximation algorithm as does our algorithm. Yet, we find that while calibrating from the prior predictive achieves nominal coverage for simple models like the Fay-Herriot of Section 2.1, it performs poorly (fails to achieve nominal coverage) for more complex models that co-estimate both the point estimates and variances as does our FHV and CFHV models of Sections 2.2 and 2.3. We speculate that there is an implicit condition that the distribution of replicate data generated from the prior predictive distribution express some minimal mass covering the true data generating distribution, which is overly difficult to achieve in practice for a complex hierarchical model. Even more, in the case of the FH where calibration from the prior predictive is feasible, it nevertheless produces longer interval lengths than calibrating from the posterior predictive where both algorithms achieve nominal coverage.

Acknowledgments

The authors would like to thank Andrew Gelman, Professor of Statistics in the Departments of Statistics and Political Science at Columbia University for suggesting the pursuit of a re-sampling based calibration procedure and for his coaching throughout the project development.

Appendix A. Overview of Mean Field Variational Bayes (MFVB)

A.1 The Mean Field Assumption

Let $\mathbf{z} = (z_1, \dots, z_m)$ denote “local” or latent variables or model parameters. The goal of variational Bayes (VB) is to approximate an intractable posterior $p(\mathbf{z}|\mathbf{x})$ with a simpler distribution $q(\mathbf{z})$. In the Mean Field formulation, we assume the joint parameter posterior distribution is fully independent, and so can be factorized as:

$$q(\mathbf{z}) = \prod_{i=1}^m q_i(z_i) \quad (20)$$

A.2 The Evidence Lower Bound (ELBO)

The ELBO, denoted as $\mathcal{L}(q) = \mathbb{E}_{q(\mathbf{z})}[\ln p(\mathbf{x}, \mathbf{z})] - \mathbb{E}_{q(\mathbf{z})}[\ln q(\mathbf{z})]$, is the objective function maximized during optimization. It is derived from the log marginal likelihood of the data (the evidence) using Jensen’s Inequality:

$$\ln p(\mathbf{x}) = \mathcal{L}(q) + \text{KL}(q(\mathbf{z})\|p(\mathbf{z}|\mathbf{x})) \quad (21)$$

	Orig Fitted	Rescaled	Pivot
Coverage	0.550	0.508	0.507
Length	0.929	0.849	0.847

Table 4: Coverage properties for model fitted $m(\theta_i)$, 50% nominal, over 150 domains and $S = 200$ simulation runs, using $A = 500$ re-samples, for the $\sigma_i^2 = 2$ setup.

Because the KL divergence is non-negative, $\mathcal{L}(q)$ serves as a lower bound on the log-evidence: $\ln p(\mathbf{x}) \geq \mathcal{L}(q)$. So, by maximizing $\mathcal{L}(q)$ for estimation of $q(\mathbf{z})$ one achieves the closest approximation within the class of $q(\mathbf{z}) = \prod_{i=1}^m q_i(z_i)$ to the target distribution, $p(\mathbf{z}|\mathbf{x})$.

A.3 Practical Implementation via ADVI

As implemented in ADVI, $q(\mathbf{z})$, is specified as Gaussian with each of the factorized distributions, $q(z_i)$ specified as normal. Stan implements ADVI by first transforming all parameters to the real line (and then back transforms them, applying the Jacobian). While the Bayesian central limit theorem asserts that the Gaussian joint distribution will provide a good approximation for a sufficiently large sample size under mild regularity conditions Ghosal and van der Vaart (2017), the factorization will tend to mis-estimate the uncertainty quantification to the extent that parameters are correlated *a posteriori*.

Appendix B. Sensitivity analysis

To explore the sensitivity of the method to the underlying model parameters to the scale of the data, we used the same simulation setup as described in Section 4.1 but now set true model parameter σ_i^2 to 2 (as opposed to original $\sigma_i^2 = 1$.) The results are given in Table 4 and Figures 9 and 10. We note that these results are very similar to the original simulation results, in terms of the coverage and variation among domains.

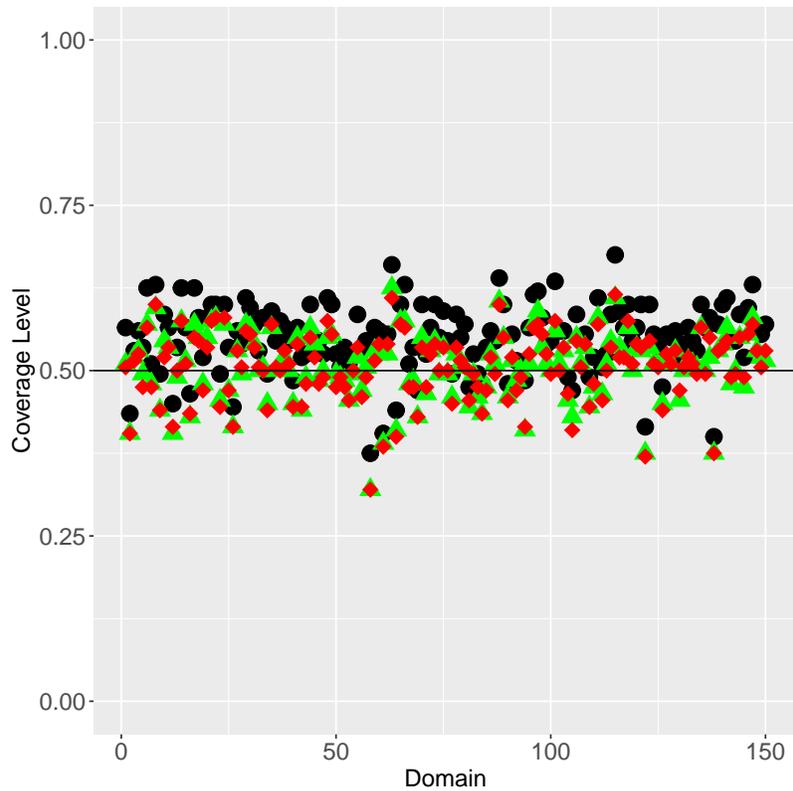


Figure 9: Domain coverages (50% nominal) of fitted point estimates vs number of respondents. Black dots represent originally fitted model; green triangles represent adjusted coverages using rescaling method; red diamonds represent adjusted coverages based on the pivotal method.

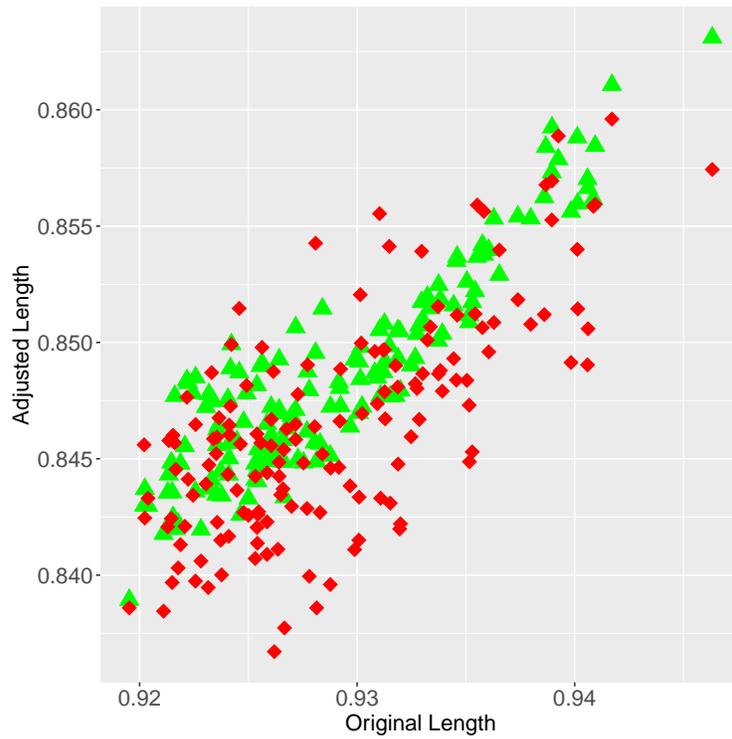


Figure 10: Lengths of the adjusted vs the original (50% nominal) confidence intervals. Green triangles represent adjusted coverages using rescaling method; red diamonds represent adjusted coverages based on the pivotal method.

References

- Bob Carpenter, Andrew Gelman, Matthew D. Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Marcus Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. Stan: A probabilistic programming language. *Journal of Statistical Software*, 76(1):1–32, 2017. doi: 10.18637/jss.v076.i01.
- Yen-Chi Chen, Y. Samuel Wang, and Elena A. Erosheva. On the use of bootstrap with variational inference: Theory, interpretation, and a two-sample test example. *The Annals of Applied Statistics*, 12(2):846 – 876, 2018. doi: 10.1214/18-AOAS1169. URL <https://doi.org/10.1214/18-AOAS1169>.
- Bradley Efron and Robert J Tibshirani. *An Introduction to the Bootstrap*. Chapman and Hall, New York, 1993. ISBN 978-0-412-04231-7.
- R.E. Fay and R.A. Herriot. Estimates of income for small places: An application of james-stein procedures to census data. *Journal of the American Statistical Association*, 85: 398–409, 1979. URL <https://doi.org/10.1080/01621459.1979.10482505>.
- Edwin Fong and Chris C Holmes. Conformal bayesian computation. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, volume 34, pages 18268–18279. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/file/97785e0500ad16c18574c64189ccf4b4-Paper.pdf.
- Andrew Gelman, John B Carlin, Hal S Stern, David B Dunson, Aki Vehtari, and Donald B Rubin. *Bayesian Data Analysis*. Chapman and Hall/CRC, London, 3rd edition, 2013.
- Julie Gershunskaya and Terrance D Savitsky. Model-Based Screening for Robust Estimation in the Presence of Deviations from Linearity in Small Domain Models. *Journal of Survey Statistics and Methodology*, 8(2):181–205, 04 2019. ISSN 2325-0984. doi: 10.1093/jssam/smz004. URL <https://doi.org/10.1093/jssam/smz004>.
- Subhashis Ghosal and Aad van der Vaart. *Fundamentals of Nonparametric Bayesian Inference*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2017. doi: 10.1017/9781139029834.
- Alp Kucukelbir, David M. Blei, Andrew Gelman, Rajesh Ranganath, and Dustin Tran. Automatic differentiation variational inference. *Journal of Machine Learning Research*, 18:1–45, January 2017. ISSN 1532-4435. Publisher Copyright: © 2017 Alp Kucukelbir, Dustin Tran, Rajesh Ranganath, Andrew Gelman and David M. Blei.
- Jing Lei, Max G’Sell, Alessandro Rinaldo, Ryan J. Tibshirani, and Larry Wasserman. Distribution-free predictive inference for regression. *Journal of the American Statistical Association*, 113(523):1094–1111, 2018. doi: 10.1080/01621459.2017.1307116. URL <https://doi.org/10.1080/01621459.2017.1307116>.

- Tapabrata Maiti, Hao Ren, and Samiran Sinha. Prediction Error of Small Area Predictors Shrinking Both Means and Variances. *Scandinavian Journal of Statistics*, 41(3):775–790, September 2014. URL <https://ideas.repec.org/a/bla/scjsta/v41y2014i3p775-790.html>.
- Peter Müller, Fernando Quintana, and Gary L. Rosner. A product partition model with regression on covariates. *Journal of Computational and Graphical Statistics*, 20(1):260–278, 2011. doi: 10.1198/jcgs.2011.09066. URL <https://doi.org/10.1198/jcgs.2011.09066>. PMID: 21566678.
- JNK Rao and Isabel Molina. *Small area estimation*. Wiley Series in Survey Methodology. John Wiley & Sons, 2015.
- Donald B. Rubin. Estimation in parallel randomized experiments. *Journal of Educational Statistics*, 6(4):377–401, 1981. doi: 10.3102/10769986006004377. URL <https://doi.org/10.3102/10769986006004377>.
- Terrance D. Savitsky and Julie Gershunskaya. Bayesian Nonparametric Joint Model For Domain Point Estimates and Variances Under Biased Observed Variances. *Journal of Survey Statistics and Methodology*, 02 2022. ISSN 2325-0984. doi: 10.1093/jssam/smac003. URL <https://doi.org/10.1093/jssam/smac003>. smac003.
- Terrance D. Savitsky, Julie Gershunskaya, and Mark Crankshaw. Joint point and variance estimation under a hierarchical Bayesian model for survey count data. *The Annals of Applied Statistics*, 17(3):2002 – 2018, 2023. doi: 10.1214/22-AOAS1704. URL <https://doi.org/10.1214/22-AOAS1704>.
- Shonosuke Sugawara, Hiromasa Tamae, and Tatsuya Kubokawa. Bayesian estimators for small area models shrinking both means and variances. *Scandinavian Journal of Statistics*, 44:150 – 167, 2015.
- Sean Talts, Michael Betancourt, Daniel Simpson, Aki Vehtari, and Andrew Gelman. Validating bayesian inference algorithms with simulation-based calibration, 2020.
- Aki Vehtari, Andrew Gelman, Tuomas Sivula, Pasi Jylänki, Dustin Tran, Swupnil Sahai, Paul Blomstedt, John P. Cunningham, David Schiminovich, and Christian P. Robert. Expectation propagation as a way of life: A framework for bayesian inference on partitioned data. *Journal of Machine Learning Research*, 21(17):1–53, 2020. URL <http://jmlr.org/papers/v21/18-817.html>.
- Bo Wang and D. M. Titterton. Convergence and asymptotic normality of variational bayesian approximations for exponential family models with missing values. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, UAI '04, page 577–584, Arlington, Virginia, USA, 2004. AUAI Press. ISBN 0974903906.
- Yuling Yao, Aki Vehtari, Daniel Simpson, and Andrew Gelman. Yes, but did it work?: Evaluating variational inference. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of*

Machine Learning Research, pages 5581–5590. PMLR, 10–15 Jul 2018. URL <https://proceedings.mlr.press/v80/yao18a.html>.

Fengshuo Zhang and Chao Gao. Convergence rates of variational posterior distributions. *The Annals of Statistics*, 48(4):2180 – 2207, 2020. doi: 10.1214/19-AOS1883. URL <https://doi.org/10.1214/19-AOS1883>.