

CALIBRATED PERCENTILE DOUBLE BOOTSTRAP FOR ROBUST LINEAR REGRESSION INFERENCE

Daniel McCarthy¹, Kai Zhang², Lawrence D. Brown³

Richard Berk³, Andreas Buja³, Edward I. George³ and Linda Zhao³

¹*Emory University*, ²*University of North Carolina at Chapel Hill*
and ³*University of Pennsylvania*

Abstract: We consider inference for the parameters of a linear model when the covariates are random and the relationship between response and covariates is possibly non-linear. Conventional inference methods such as z intervals perform poorly in these cases. We propose a double bootstrap-based calibrated percentile method, `perc-cal`, as a general-purpose CI method which performs very well relative to alternative methods in challenging situations such as these. The superior performance of `perc-cal` is demonstrated by a thorough, full-factorial design synthetic data study as well as a data example involving the length of criminal sentences. We also provide theoretical justification for the `perc-cal` method under mild conditions. The method is implemented in the R package ‘perccal’, available through CRAN and coded primarily in C++, to make it easier for practitioners to use.

Key words and phrases: Confidence intervals, Edgeworth expansion, resampling, second-order correctness.

1. Introduction

In many applied settings, practitioners would like to make interpretable statements such as the expected average difference in a response variable, Y , associated with a unit difference in a covariate of interest, X_j , controlling for all other predictors. In situations like these, practitioners often run linear regressions despite the fact that the true but unobservable relationship between Y and X_j 's may be non-linear. Doing so may be sensible when the utility of being able to make more interpretable statements such as the one above outweighs the cost of possible model bias, which may be hard to discern (particularly in multivariate settings). An important challenge is how practitioners can produce valid inference upon their estimates of the true population-level best linear approximation for the relationship between predictor and response in these settings. We denote

the target of interest by β (for more detail, please see Section 2.4). Our aim with this paper is to help practitioners perform better inference for β in these situations.

Buja et al. (2015) call much-needed attention to this issue, showing that when the relationship between response Y and covariates $\vec{\mathbf{X}} = (1, X_1, \dots, X_p)^T$ is truly non-linear with noise that is possibly heteroskedastic, and when $\vec{\mathbf{X}}$ is itself random, standard linear model theory standard errors are asymptotically invalid. They show that the “sandwich estimator” of standard error does provide asymptotically correct inference for the population slopes, even when non-linearity and heteroskedasticity are present, and $\vec{\mathbf{X}}$ is random.

While the sandwich estimator may provide asymptotically valid inference, practitioners will also be understandably interested in better understanding how the finite sample performance of various methods of inference compare, as well as the asymptotic properties of those methods. We will show that in our setting, empirical coverage of population regression slopes deteriorates considerably for all traditional confidence interval methods. The primary contribution of this paper is to shine new light on these issues, proposing and studying an inference method which is convincingly superior to the sandwich estimator, and making a very fast implementation of this proposed method accessible to practitioners.

We propose a double bootstrap-based calibrated percentile method, `perc-cal`. The seminal work of Peter Hall shows the advantages of double bootstrap approach in classical settings involving population means. Population slopes as defined here are a more complex, non-linear object. For example, Hall (1992) studies univariate data without model misspecification. The methods he uses, then, need to be augmented with additional material about Edgeworth expansions that is adapted from Jensen (1989). For the first time, we prove in Section 3 and in the appendix that even when Y and $\vec{\mathbf{X}}$ have a non-linear joint distribution, and $\vec{\mathbf{X}}$ is random, that under relatively mild regularity conditions the rate of coverage error of `perc-cal` for two-sided confidence intervals of the best linear population slopes between a response variable Y and p -dimensional covariates $\vec{\mathbf{X}}$ is $\mathcal{O}(n^{-2})$. In contrast, conventional methods achieve a rate of coverage error of $\mathcal{O}(n^{-1})$. We then show in a Monte Carlo study that `perc-cal` performs better than traditional confidence interval methods, including the BCa method (Efron (1987)), and other Sandwich-based estimators discussed in Cribari-Neto, Souza and Vasconcellos (2007) and MacKinnon (2013). Our study is similar in structure to the simulation study that was performed in Gonçalves and White (2005), but modified to study a very wide variety of misspecified mean functions. We fol-

low up this synthetic simulation study with a data example involving a criminal sentencing dataset, and show that `perc-cal` once again performs satisfactorily. We have released an R package, ‘`perccal`’ (McCarthy (2016)), available through CRAN and coded primarily in C++, so that practitioners may benefit from a fast implementation of `perc-cal` for their own analyses.

We argue that the combination of theoretical and empirical justification presented in this paper supports the claim that `perc-cal` is a reliable confidence interval methodology that performs well in general, even in the presence of relatively severe model misspecification. The remainder of the paper is organized as follows. Section 2 provides a review of confidence interval methods. Section 3 presents the theoretical results. Section 4 compares the performance of `perc-cal` with that of other often more commonly used confidence interval estimators in synthetic and real data settings. Section 5 provides concluding remarks, and an Appendix gives all of the proofs.

2. Literature Review

2.1. Review of bootstrap confidence intervals

There is a very wide variety of bootstrap methods that have been proposed in the literature to compute $(1 - \alpha)$ confidence intervals. These methods include Efron’s percentile method (Efron (1981)), Hall’s percentile approach (Hall (1992)), and Hall’s percentile- t method (Hall (1988)). Other forms of bootstrap CIs include symmetric CIs (Hall (1992)) and short bootstrap confidence intervals (Hall (1992)). In general, performance of these methods depends upon the properties of the data generating process and/or the sample size. We are primarily interested in confidence interval methods that assume much less about the true underlying data generating process, which is usually unknown and often not well behaved in applications, making these methods less relevant to the work which follows.

Hall advocates the use of pivotal bootstrap statistics because they have higher asymptotic accuracy when the limiting distributions are indeed pivotal (Hall (1992, p. 83)). We emphasize that Hall’s preference for pivotal bootstrap statistics, and much of the discussion regarding the relative merits of various confidence interval methods, are based on the *asymptotic* properties of these methods. When the sample size is small, these asymptotic considerations do not necessarily reflect the empirical performance of these methods. For example, Hall cautioned that “our criticism of the percentile method and our preference for

percentile- t lose much of their force when a stable estimate of σ^2 is not available" (Hall (1988)). Simulation studies that reinforce this include Scholz (2007).

Another class of confidence intervals may be formed by replacing the standard error estimator in the standard \mathbf{z} or \mathbf{t} interval with a so-called 'Sandwich' estimator (White (1980)), or one of the many extensions of the Sandwich estimator (Cribari-Neto, Souza and Vasconcellos (2007)). A comprehensive review of Sandwich estimators can be found in MacKinnon (2013), and we will compare these methods with our proposed method in Section 4.

2.2. Review of iterative bootstrap confidence intervals

The idea of the iterative bootstrap (or double-bootstrap) was first introduced in Efron (1983). The improvement on coverage probability of CIs was first analyzed in Hall (1986) and later discussed in more detail in Loh (1987), Hall and Martin (1988), Hall, Martin and Schucany (1989) and Martin (1990a,b). A comprehensive review can be found in Section 3.11 in Hall (1992) (see also Efron and Tibshirani (1994)). In general, the iterative bootstrap provides more accurate coverage probability at the cost of more computing.

To fix ideas, in this section we shall introduce the proposed double-bootstrap confidence interval method in a univariate case with generic notations. We will extend this procedure to the regression setting in Section 2.4. We assume that we observe $Z_1, \dots, Z_m \stackrel{iid}{\sim} F$ for some distribution F . Let $\theta = \theta(F)$ be a parameter of our interest. We will estimate θ through the empirical distribution $\hat{F}(z) = (1/m) \sum_{i=1}^m I(Z_i \leq z)$. The estimator is denoted by $\hat{\theta} = \theta(\hat{F}) = \theta(Z_1, \dots, Z_m)$. The construction of the confidence interval is illustrated in Figure 1 and is described as follows.

1. For chosen bootstrap sample size B_1 , obtain bootstrap samples $(\mathbf{Z}_1^*, \dots, \mathbf{Z}_{B_1}^*)$. Each \mathbf{Z}_j^* consists of m i.i.d. samples with replacement from \hat{F} . For chosen bootstrap sample size B_2 , obtain double bootstrap samples corresponding to all bootstrap samples, $(\mathbf{Z}_{1,1}^{**}, \dots, \mathbf{Z}_{1,B_2}^{**}, \mathbf{Z}_{2,1}^{**}, \dots, \mathbf{Z}_{2,B_2}^{**}, \dots, \mathbf{Z}_{B_1,1}^{**}, \dots, \mathbf{Z}_{B_1,B_2}^{**})$ in the same manner as in the first-level bootstrap. Denote the empirical distributions by \hat{F}_j^* 's, $j = 1, \dots, B_1$, and $\hat{F}_{j,k}^{**}$'s, $j = 1, \dots, B_1, k = 1, \dots, B_2$, respectively.
2. Obtain parameter estimates corresponding to the observed sample, $\hat{\theta} = \theta(\hat{F})$, all bootstrap samples, $(\hat{\theta}_1^*, \dots, \hat{\theta}_{B_1}^*)$ with $\hat{\theta}_j^* = \theta(\hat{F}_j^*)$ and all double bootstrap samples corresponding to all bootstrap samples, $(\hat{\theta}_{1,1}^{**}, \dots, \hat{\theta}_{1,B_2}^{**}, \hat{\theta}_{2,1}^{**}, \dots, \hat{\theta}_{2,B_2}^{**}, \dots, \hat{\theta}_{B_1,1}^{**}, \dots, \hat{\theta}_{B_1,B_2}^{**})$ with $\hat{\theta}_{j,k}^{**} = \theta(\hat{F}_{j,k}^{**})$.

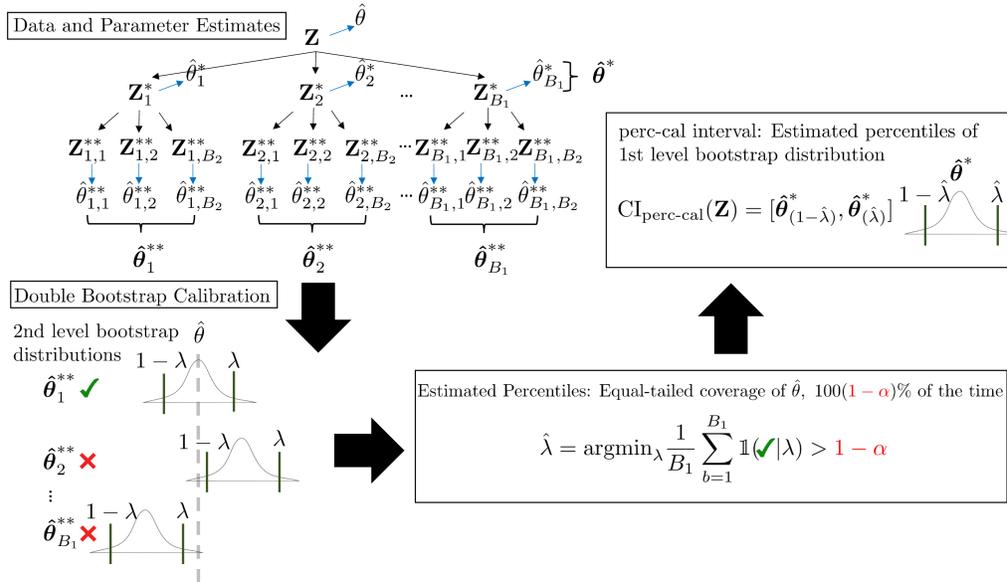


Figure 1. perc-cal diagram.

3. Form B_1 double-bootstrap histograms $\hat{\theta}_1^{**}, \dots, \hat{\theta}_{B_1}^{**}$, where each histogram $\hat{\theta}_j^{**}$ is comprised of all B_2 double bootstrap estimates $(\hat{\theta}_{j,1}^{**}, \dots, \hat{\theta}_{j,B_2}^{**})$ corresponding to the j th bootstrap sample and estimate, \mathbf{Z}_j^* and $\hat{\theta}_j$, respectively, $j \in \{1, 2, \dots, B_1\}$.
4. Find the largest $\hat{\lambda}$ such that $1/2 < \hat{\lambda} < 1$ and that $\hat{\theta}$ lies in the $1 - \hat{\lambda}$ percentile and the $\hat{\lambda}$ percentile of the histograms $1 - \alpha$ proportion of the time.
5. $\hat{\theta}$ lies between the $(1 - \hat{\lambda}, \hat{\lambda})$ percentiles of the second-level bootstrap distributions $1 - \alpha$ proportion of the time. Therefore our **perc-cal** $(1 - \alpha)$ interval for θ is equal to the $(1 - \hat{\lambda}, \hat{\lambda})$ percentiles of the first-level bootstrap distribution, $[\hat{\theta}_{(1-\hat{\lambda})}^*, \hat{\theta}_{(\hat{\lambda})}^*]$.

For a $(1 - \alpha)$ left-sided **perc-cal** confidence interval for θ , the only change in the procedure is in Step 4, where one uses the histograms to find the smallest $\hat{\lambda}$ such that $\hat{\theta}$ lies below the $\hat{\lambda}$ percentile of the histograms $1 - \alpha$ percent of the time. In what follows, we shall refer the two-sided **perc-cal** interval as $\mathcal{I}_2 = [\hat{\theta}_{(1-\hat{\lambda})}^*, \hat{\theta}_{(\hat{\lambda})}^*]$ and the one-sided **perc-cal** interval as $\mathcal{I}_1 = (\infty, \hat{\theta}_{(\hat{\lambda})}^*]$.

A similar double-bootstrap confidence interval is the double-bootstrap- t method which uses the second-level bootstrap to calibrate the coefficient of the bootstrap

standard deviation estimate. In practice, both methods can be applied. Hall commented in his book (Hall (1992)) that “either of the two percentile methods could be used, although the ‘other percentile method’ seems to give better results in simulations, for reasons that are not clear to us.” Here “the ‘other percentile method’” refers to confidence intervals \mathcal{I}_1 and \mathcal{I}_2 . Our simulation studies in Section 4 demonstrate the same phenomenon. We have observed through simulation that the performance of double-bootstrap- t can be erratic at times due to the instability which arises by relying upon a statistic which has a double bootstrap estimated standard error in its denominator. Some samples will inevitably have very small or even degenerate bootstrap standard errors, making this statistic very large. This issue is particularly acute when sample sizes are small.

Research on optimizing the trade-off between the number of simulations, B_1 and B_2 , in double-bootstrap and the CI accuracy can be found in Beran (1987, 1988), Booth and Hall (1994), Booth and Presnell (1998), Lee and Young (1999), among many others. In particular, Lee and Young (1999) study the asymptotic convergence rate of the coverage probability involving B_1 and B_2 , and suggest an adaptive method to optimize the choice of B_2 as a function of B_1 . They note that B_1 should be set to a larger number (for example, $B_1 = 1,000$), and B_2 equal to a lesser value. In all simulations which follow, we set $B_1 = B_2 = 2,000$. Since the computation of `perc-cal` is reasonably efficient as discussed in Section 5, we do not optimize the number of bootstrap samples further but note that further performance gains for robust linear regression inference are a promising area for future research (in particular, with respect to B_2).

2.3. Review of bootstrap applications in conventional linear models

Bootstrap in linear models is studied in Section 4.3 in Hall (1992). Hall refers the fixed design case the “regression model” and the random design case the “correlation model.” Bootstrap estimation and confidence intervals for the slopes, as well as simultaneous confidence bands, are described.

Since the seminal paper of Freedman (1981), the bootstrap has been widely used in regression models because of its robustness to the sample distributions. A review of bootstrap methods in economics can be found in MacKinnon (2006). Gonçalves and White (2005) consider bootstrapping the sandwich estimator for the standard error when the observations are dependent and heterogeneous. Bootstrap applications under other types of model misspecifications are recently considered in Kline and Santos (2012) and Spokoiny and Zhilova (2014). In this paper, we focus on a different case when observations of $(Y, \vec{\mathbf{X}})$ are i.i.d. but the

joint distribution is assumption-lean – we elaborate upon this further in the next section.

2.4. The assumption-lean framework and double-bootstrap applications

Conventional linear models assume $\mathbf{E}(Y|\vec{\mathbf{X}}) = \vec{\mathbf{X}}\boldsymbol{\beta}$ for some $(p+1)$ -vector $\boldsymbol{\beta}$ as regression coefficients, so that Y depends on $\vec{\mathbf{X}}$ only through a linear function. While this is commonly assumed in the bootstrap literature, we may not want to require it when performing inference in data settings because the true relationship may not be linear. Moreover, as first noted in White (1980), a non-linear relationship between Y and $\vec{\mathbf{X}}$ and randomness in $\vec{\mathbf{X}}$ can lead to serious bias in the estimation of standard errors. Buja et al. (2015) reviewed this problem, and proposed an “assumption-lean” framework for inference in regressions. In this framework, no assumption is made regarding the relationship between Y and $\vec{\mathbf{X}}$. The only assumptions are on the existence of certain moments of the joint distribution of $\vec{\mathbf{V}} = (X_1, \dots, X_p, Y)^T$. This consideration makes the model very general and thus widely applicable. Readers are referred to Buja et al. (2015) for more details.

Even though a linear relationship between $\mathbf{E}(Y|\vec{\mathbf{X}})$ and $\vec{\mathbf{X}}$ is not assumed in an assumption-lean framework, the slope coefficients that are estimated are always well-defined through a population least-squares consideration: the population least-squares coefficients $\boldsymbol{\beta}$ minimize squared error risk over all possible linear combinations of $\vec{\mathbf{X}}$:

$$\boldsymbol{\beta} = \operatorname{argmin}_{\mathbf{b}} \mathbf{E} \|Y - \mathbf{b}^T \vec{\mathbf{X}}\|_2^2 = \mathbf{E}(\vec{\mathbf{X}}\vec{\mathbf{X}}^T)^{-1} \mathbf{E}(Y\vec{\mathbf{X}}). \quad (2.1)$$

This definition of coefficients $\boldsymbol{\beta}$ is meaningful in addition to being well defined under minimal assumptions: $\boldsymbol{\beta}$ provides us with the best linear approximation from $\vec{\mathbf{X}}$ to Y , whether or not $\vec{\mathbf{X}}$ and Y are linearly related to one another. This setup allows for situations including random $\vec{\mathbf{X}}$, non-Normality, non-linearity and heteroskedasticity and we show later that the proposed `perc-cal` method provides better empirical coverage of the true population least-squares coefficients $\boldsymbol{\beta}$ on average over a wide variety of data generating processes, even if those data generating processes involve random $\vec{\mathbf{X}}$, non-linearity in $\mathbf{E}(Y|\vec{\mathbf{X}})$ and/or heteroskedasticity. In contrast, previous research on the double bootstrap has studied functions that are linear or approximately linear.

To estimate $\boldsymbol{\beta}$, denote the i.i.d. observations of $\vec{\mathbf{V}}$ by $\vec{\mathbf{V}}_1, \dots, \vec{\mathbf{V}}_n$ and denote the $n \times (p+1)$ matrix with rows $\vec{\mathbf{V}}_1, \dots, \vec{\mathbf{V}}_n$ by \mathbf{V} . Denote the dis-

tribution of \mathbf{V} by G . The multivariate empirical distribution of $\vec{\mathbf{V}}$ is then $\hat{G}(\vec{\mathbf{V}}) = \hat{G}(x_1, \dots, x_p, y) = (1/n) \sum_{i=1}^n I(X_{1,i} \leq x_1, \dots, X_{p,i} \leq x_p, Y_i \leq y)$. The least squares estimate for β defined in (2.1) is

$$\hat{\beta} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{Y}, \quad (2.2)$$

where \mathbf{X} is the $n \times (p+1)$ matrix and \mathbf{Y} is the $n \times 1$ vector containing the i.i.d. observations of $\vec{\mathbf{X}}$ and Y respectively. Note that each estimate $\hat{\beta}_j$ of β_j can be written as a function of \hat{G} , $\hat{\beta}_j = \beta_j(\hat{G})$.

The `perc-cal` confidence intervals for each of the slopes β_j in β are constructed similarly as described in Section 2.2. We use the pairs bootstrap first proposed by Freedman (1981) and create B_1 i.i.d. bootstrap samples, \mathbf{V}_k^* , where each matrix \mathbf{V}_k^* consists n i.i.d. samples with replacement from \hat{G} . From these samples, one can create the empirical distribution \hat{G}_k^* . To find the proper calibration of \mathcal{I}_1 or \mathcal{I}_2 (as defined in Section 2.2), we sample B_2 i.i.d. pairs bootstraps $\mathbf{V}_{k,h}^{**}$ with empirical distributions $\hat{G}_{k,h}^{**}$. The other steps in the construction are identical to those in Section 2.2 with $\theta(\cdot)$ replaced by $\beta_j(\cdot)$ with respective empirical distributions as arguments. We note here that although confidence intervals can be built through empirical process theory van der Vaart (1998), the accuracy is usually not as good as the proposed `perc-cal` method, as will be discussed in the next section. While Delaigle, Hall and Jamshidi (2015) studies bootstrap confidence bands under a nonparametric regression setting, we are unaware of any existing general bootstrap theory that yields general results for iterated bootstrap confidence intervals for a non-linear function of expectations of non-linear functions.

3. Asymptotic Theory

In this section, we discuss the theoretical properties of `perc-cal`. The following theorem describes the accuracy on the coverage probability of `perc-cal` confidence intervals.

Theorem 1. *Consider n i.i.d. observations of the $(p+1)$ -dimensional random vector $\vec{\mathbf{V}} = (X_1, \dots, X_p, Y)^T$. Denote the vector of Y and the continuous X_j 's by $\vec{\mathbf{V}}_C$ and that of the discrete X_j 's by $\vec{\mathbf{V}}_D$. Suppose that*

1. $\exists C_0 > 0$, such that the minimal eigenvalue of the covariance matrix $\text{Var}(\vec{\mathbf{V}})$ is larger than C_0 .
2. The moment generating function of the random variable $\|\vec{\mathbf{V}}\|_2$ exists for all $t \in \mathbb{R}$: $\mathbf{E}[e^{t\|\vec{\mathbf{V}}\|_2}] < \infty$ for all $t \in \mathbb{R}$.

3. Conditional on every single value of $\vec{\mathbf{V}}_D$, the joint distribution of $\vec{\mathbf{V}}_C$ is absolutely continuous with respect to the Lebesgue measure.

Consider the population least squares parameter defined in (2.1) whose estimate is the sample least squares defined in (2.2). Then for each $1 \leq j \leq p$, the $(1 - \alpha)$ **perc-cal** CI for β_j described in Section 2.2 and Section 2.4 have the coverage probabilities

$$\mathbf{P}(\beta_j \in \mathcal{I}_1) = 1 - \alpha + O(n^{-1}), \quad (3.1)$$

$$\mathbf{P}(\beta_j \in \mathcal{I}_2) = 1 - \alpha + O(n^{-2}). \quad (3.2)$$

The proof of Theorem 1 is in Section A in the Appendix. It uses techniques of the Edgeworth expansion as in Section 3.11.3 of Hall (1992), with a focus on the assumption-lean regression model setting where the dependence between Y and X_j 's are not necessarily linear. We are not aware of prior investigation of the performance of the double bootstrap confidence intervals under this situation. Moreover, the results in Martin (1990a) and Hall (1992) accommodate only X_j 's satisfying Cramèr's condition, which excludes distributions such as the Poisson distribution. Discrete distributions for X_j are often encountered in models involving categorical predictors, and the data example we study in Section 4 involves such covariates. Through a conditioning argument in Jensen (1989), we are able to show that the same performance is enjoyed by a wider class of mixed discrete and continuous X_j 's.

Our results show that the coverage probability is $1 - \alpha + O(n^{-2})$. Note that for other construction approaches of confidence intervals such as from the sandwich estimator or from the empirical process theory, the resulting one-sided confidence intervals often have a coverage probability of $1 - \alpha + O(n^{-1/2})$, and two-sided ones often have a coverage probability of $1 - \alpha + O(n^{-1})$ (see Section 3.5.4 and Section 3.5.5 in Hall (1992)). Thus, the double bootstrap method provides better coverage.

4. Numerical Studies

In this section, we study the performance of **perc-cal** compared to alternative (often more common) methods for forming confidence intervals, including other double bootstrap methods. We first compare **perc-cal** to these other methods using simulated data under a very wide variety of true data generating processes. We then illustrate our approach in a data example. We will see that **perc-cal** performs very satisfactorily in general.

4.1. Synthetic simulation study

4.1.1. Design: synthetic simulation

We compared the performance of `perc-cal` with ten other methods that are commonly used for constructing confidence intervals:

1. Standard normal interval: `z` (Efron and Tibshirani (1994)).
- 2-6. Five sandwich variants: `sand1`, `sand2`, `sand3`, `sand4` and `sand5` (MacKinnon (2013) provides a review of these methods, denoted there by H1, H2, H3, HC4 and HC5).
7. Hall's Studentized interval: `stud` (Hall (1988)).
8. Hall's "bootstrap-t" method: `boot-t` (Efron and Tibshirani (1994)).
9. Efron's BCa interval: `BCa` (Efron (1987)).
10. Single percentile method: `perc` (Efron and Tibshirani (1994)).

We considered a very wide range of underlying true data generating models, to obtain a more general understanding for how these confidence interval methods compare against one another in a wide variety of data settings, for large sample sizes as well as small. The data generating models represent a full factorial design of the following 48 factors, after excluding non-denerate combinations (i.e., combinations for which the conditional mean is not finite):

- Simple regression - one predictor, $Y = \beta_0 + \beta_1 X + \epsilon$.
- Sample size $n = 32, 64, 128, 256$.
- Relationships between Y and X: (1) $Y = X + e$; (2) $Y = \exp(X) + e$; (3) $Y = X^3 + e$.
- Distribution of X: (1) $X \sim \mathcal{N}(0, 1)$, (2) $X \sim \exp(\mathcal{N}(0, 1))$.
- Noise: $\epsilon \sim$ (1) $N(0, 1)$; (2) $|X| * \mathcal{N}(0, 1)$; (3) $\exp(\mathcal{N}(0, 1))$.

In each of the above cases, we used 2,000 first and second-level bootstrap samples for all bootstrap methods ($B_1 = B_2 = 2,000$). We obtained empirical coverage figures for the slope coefficient in the regression, β_1 . Results were averaged over 500 replications to reduce the empirical standard error of the resulting intervals to below 1.5% on average across scenarios and methods. We present

Empirical Coverage of True Slope Coefficient, 90% Target Coverage

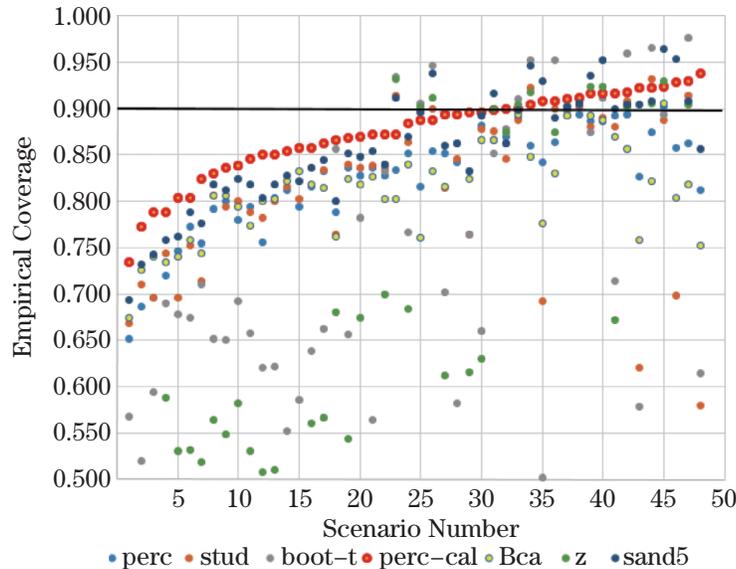


Figure 2. Scatterplot of empirical coverage proportion of methods by scenario number – 90% target coverage.

the results for a target coverage of 90%, without loss of generality (results for a target coverage of 95% are qualitatively the same).

4.1.2. Results

To more easily visualize the performance of many methods under many different scenarios, we begin with a coverage scatterplot in Figure 2. We rank sort the 48 scenarios by the empirical coverage proportion of `perc-cal`, in ascending order. This ordered list determines the scenario number associated with each scenario (i.e., scenario number 1 represents the scenario in which `perc-cal`'s empirical coverage was the smallest, while scenario number 48 represents the scenario in which `perc-cal`'s empirical coverage was the largest). On the x-axis, we provide the scenario number. On the y-axis, we provide the empirical coverage proportion of β_1 using all methods. We exclude `sand1`, `sand2`, `sand3`, and `sand4` but include `sand5`, because `sand5` has a better empirical coverage proportion than the other sandwich estimators. We add a horizontal line to the graph at the desired target coverage level of 90%.

In general, none of the methods was “perfect” in the sense of always providing coverage at or above the target level of coverage. All noticeably undercover in particular cases and in these cases, `perc-cal`'s relative performance is generally

noticeably strong. Specifically, `perc-cal` achieves coverage less than target in 31 of the 48 scenarios. In 27 of those 31 below-target scenarios, `perc-cal` achieved a higher empirical coverage than all other alternative methods. In the remaining 4 of the 31 below-target scenarios where `perc-cal` was not the best-performing method, its absolute performance was close to target, averaging to 88.7%, and only one scenario with under 88% coverage. While `perc-cal` also had the highest empirical average coverage across all scenarios, we believe this robustness to challenging scenarios is more important to practitioners, who may take comfort in the fact that when `perc-cal` undercovers it almost always outperforms alternative methods, and in all other scenarios, it achieves or exceeds the desired coverage.

Across scenarios, while `perc-cal` provided the most consistent empirical coverage, `sand5` was itself generally superior to the other alternative methods, including `BCa`, which has favorable asymptotic properties. `perc-cal` achieved a mean absolute deviation from target coverage of only 3.8 percentage points, versus 5.8 and 8.9 percentage points for `sand5` and `BCa`, respectively. Restricting our attention to just the scenarios in which `perc-cal` achieved empirical coverage below 90%, the corresponding mean absolute deviation statistics were 5, 7.7, and 10.4, respectively. While it is perhaps no surprise that traditional methods such as `z` interval fare so poorly because they typically assume a fixed- X setting, the relative performance of these methods which do not make such assumptions is more interesting.

`perc-cal`'s improved coverage came at the cost of modestly longer interval lengths – across these 48 scenarios, `perc-cal` had an average interval length of 1.39, at the upper end (but not the top) of other methods – excluding the poor performance of `z` and `boot-t`, these other methods had interval lengths between 0.99 and 1.49, averaging to 1.16. While these other methods have interval lengths that are approximately 16% smaller on average, it would not be acceptable to a practitioner for this shortness to come at the expense of falling below desired target coverage. Only when target coverage is achieved do considerations like average interval length become a primary concern, and 16% longer intervals seems like a reasonable price for the “insurance” provided by `perc-cal`.

4.2. Data example: criminal sentencing dataset

4.2.1. Design: data example

We turn now to an example of how well `perc-cal` performs in practice, on data. In this section, we compare `perc-cal` to other methods on a criminal

sentencing dataset. This dataset contains information regarding criminal sentencing in a large state from January 1st 2002 to December 31st 2007. There are a total of 119,983 offenses in the dataset, stemming from a variety of crimes – murder (899 cases), violent crime (80,402 cases), sex-related crime (7,496 cases), property-related crime (92,743 cases), firearm-related crime (15,326 cases) and drug crime (93,506 cases). An individual offense can involve multiple types of crime, and an offender’s case can involve multiple charges of each type of crime.

Our modeling objective is to form marginal confidence intervals for the slope coefficients of a linear regression. The response variable of our regression is the number of days of jail time an offender must serve (log-transformed), which we predict with the following 8 covariates:

1. **race**: Binary variable for the race of the offender (1 if white, 0 if non-white).
2. **seriousness**: A numerical variable scaled to lie between 0 and 10 indicating the severity of the crime. A larger number denotes a more serious crime.
3. **age**: Age of offender at the time of the offense.
4. **race**: The percent of the neighborhood that is not of Caucasian ethnicity in the offender’s home zip code.
5. **in-state**: Binary variable for whether the offender committed the crime in his/her home state (1 if in offenders home state, 0 otherwise).
6. **juvenile**: Binary variable for whether the offender had at any point committed a crime as a juvenile (1 if yes, 0 otherwise). 18% of all offenses involved offenders who had committed a crime as a child.
7. **prior-jaildays**: Number of days of jail time the offender had previously served.
8. **age-firstcrime**: The age of the offender when the offender was charged with his/her first crime as an adult.

This is truly a random X setting because the predictors themselves are stochastic, coming to us from an unknown distribution. \vec{X} is stochastic, the relationship between \vec{X} and Y is unknown and possibly non-linear, and error may have heteroskedastic variance. These results are not meant to be a complete study of the issue, but rather are presented to illustrate the potential of our methodology.

We first ran a linear regression upon the full dataset containing all 119,983 offenses. We treated the coefficients as if this were a population-level regression. We then proceeded as if we did not have the full dataset and instead only had the ability to observe random subsets of the dataset of size 500 – large, but not so large that all coefficient estimates are over-powered. We studied the empirical coverage performance of confidence intervals formed using the methods in the simulation exercise over repeated realizations which are obtained through random subsamples of size 500.

4.2.2. Results: data example

Linear regression across the full dataset has an R^2 of 16.9% with 6 of the 8 predictors coming up as significant. We then took repeated random subsamples of size 500 from this population of offenses and treated these subsamples as if they were the observed dataset. Presupposing that each crime represents an *iid* draw, this framework allows us to compare and contrast the empirical coverage performance of confidence interval methods.

In Figure 3, we present the empirical coverage for each of our predictors when we form 90% confidence intervals. The y-axis of the plot below represents the empirical coverage over 10,000 realizations for each of the methods in question (i.e., 90% empirical coverage for a particular method implies that 9,000 of the 10,000 realizations had confidence intervals for that method which contained the true but unknown population-level parameters). Along the x-axis, we have the predictors listed above. We include a bold horizontal line at the target level of empirical coverage of 90%. The standard error associated with the coverages presented below average to 0.002 across scenarios, predictors and methods.

There are a number of inferences that we can draw from the above chart:

- All methods generally perform as expected, with empirical coverage proportions generally falling between 85% and 92%.
- **perc-cal** is the only method that consistently achieves empirical coverage over 90%. All other methods, including **sand5**, were unable to do so.
- **prior-jaildays** appears to be the predictor with the most disappointing empirical coverage. All methods except for **perc-cal** do not achieve 90% empirical coverage. The average empirical coverage of **prior-jaildays** for all non-**perc-cal** methods was 87.0%.
- There is also considerable disparity in the ability of various methods to cover the coefficients associated with the intercept term and the **in-state**

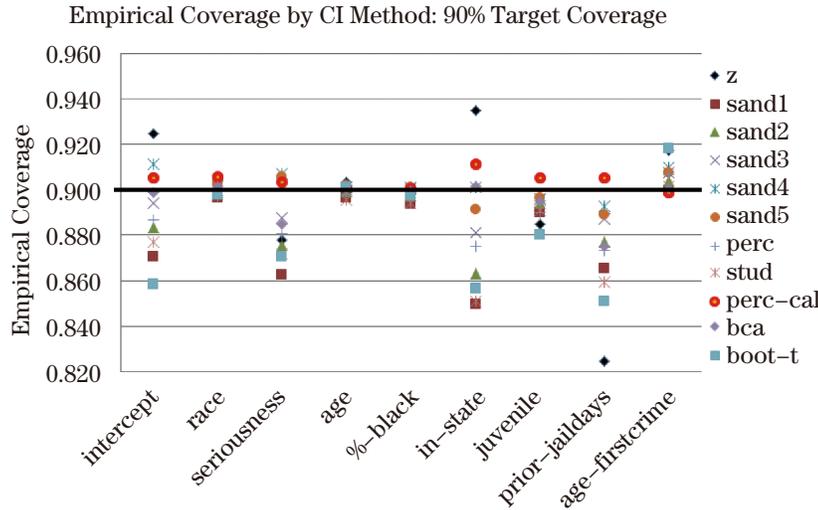


Figure 3. Scatterplot of coverage proportion of methods – 90% Target Coverage.

covariate. Although the BCa method has near-90% empirical coverage of the `in-state` super-population coefficient, its coverage is less satisfactory for the `seriousness` and `prior-jaildays` covariates.

When we plot the relationship of jail length (log transformed) against prior total jail length in Figure 4, adjusted for all of the other covariates in the super-population, we see an almost bi-modal relationship.

It is clear from the plot in Figure 4 that the highly misspecified relationship between Y and \vec{X} is likely to be driving the large disparity (and general deterioration) in coverage performance across the various non-`perc-cal` confidence interval methods. Overall, these results support the notion that `perc-cal` is a good all-purpose confidence interval method, and that all other methods, while performing well for some of the covariates, do not perform well for all of the covariates as was the case for `perc-cal`. The results assuming target coverage of 95% are qualitatively the same as the results presented above.

5. Discussion and Concluding Remarks

If `perc-cal` performs so well relative to alternative more popular CI methods, why is it not used more in practice? We believe the use of double bootstrap methods in general have not been widely adopted primarily because of their computational cost. Although it is true that double bootstrap methods in general and `perc-cal` in particular require more computation, the computational burden

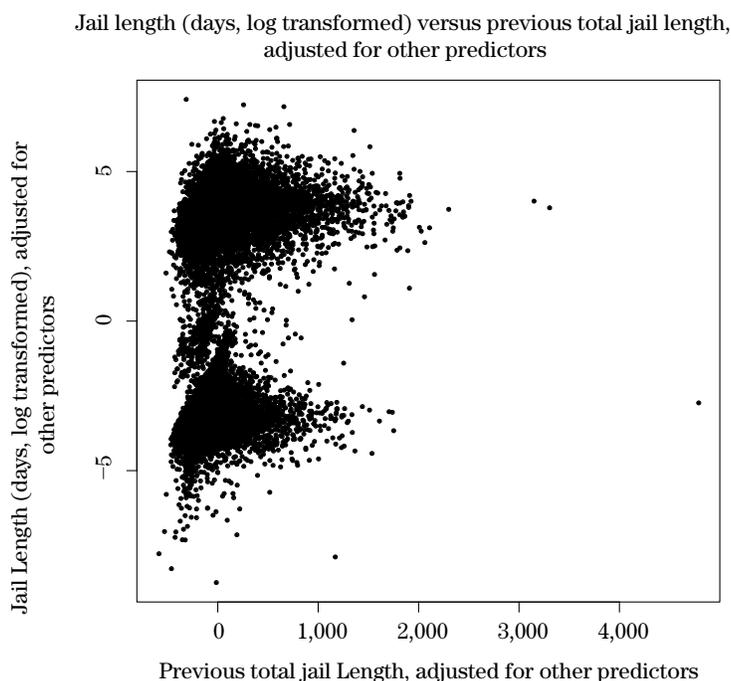


Figure 4. Jail length (log transformed) versus previous total jail length, adjusted for other predictors.

of these procedures is far less problematic than in the past because of current computational advances. For example, the rise of grid computing has greatly facilitated parallel computation. Because `perc-cal` is trivially parallelizable, it is relatively straightforward to compute all second-level bootstrap calculations in parallel, allowing researchers to compute `perc-cal` at a computational “cost” that is on the order of a single bootstrap. Furthermore, the perceived computational cost of double bootstrap methods may be inflated due to the inefficiency with which the calculations are carried out in popular statistical programming languages, most notably R – the very same calculations are orders of magnitude faster in lower level languages, such as C++. The rising popularity and adoption of packages integrating R with C++ (Eddelbuettel et al. (2011)) can greatly reduce the cost of double bootstrap methods for practitioners performing data analysis in R who do not know C++. In the spirit of this, the R package we have created allows users to compute `perc-cal` intervals in R efficiently using C++ code via `Rcpp`. We are optimistic that the use of double bootstrap methods will only increase further as the cost of computing declines further over the next 10 years.

We have restricted our attention to equal-tailed intervals for all methods considered here. It is natural and certainly possible to extend our approach to compute the shortest *unequal*-tailed interval, even if other methods cannot or would not, because of the symmetry of the asymptotic distribution underlying those alternative methods. At the same time, this advantage should not be overstated – for example, one may be forced so far into the tails of the bootstrap distribution that a considerably larger number of first and second-level bootstrap samples are required. Because this is not the focus of our paper, we do not pursue it further here.

The asymptotic theory we developed, examined, and compared the more traditional percentile and “bootstrap-t” methods to their double bootstrap analogs in our “assumption lean” setting. We did not study the asymptotic properties of alternative confidence interval methods in our setting. Although it would be interesting to do so, there are a very wide range of methods in the literature, making systematic theoretical study impractical.

In summary, randomness in \vec{X} , non-linearity in the relationship between Y and \vec{X} , and heteroskedasticity “conspire” against classical inference in a regression setting (Buja et al. (2015)), particularly when the sample size is small. We have shown that, in theory, the percentile-calibrated method `perc-cal` provides very satisfactory empirical coverage – the asymptotic rate of coverage error under mild regularity conditions for a two-sided confidence interval of the best linear approximation between Y and \vec{X} is $\mathcal{O}(1/n^2)$. Furthermore, `perc-cal` performs very well in practice, both in synthetic and data settings. We believe that `perc-cal` is a good general-purpose CI method and merits consideration when confidence intervals are needed in applied settings by practitioners.

Acknowledgment

Lawrence D. Brown and Daniel McCarthy were partially supported by NSF grant DMS-1512084. Edward George is partially supported by the NSF grant DMS-1406563. Kai Zhang is partially supported by NSF grant DMS-1309619, NSF DMS-1613112 and IIS-1633212, and is also partially supported by the NSF under Grant DMS-1127914 to the Statistical and Applied Mathematical Sciences Institute. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Science Foundation.

Appendix

A. Proof of Theorem 1

The proof consists of two parts. The first part shows the existence of the Edgeworth expansion of the pivoting quantity. The second part derives the asymptotic order of the error term by using the first terms in this expansion and the Cornish-Fisher expansion, which can be regarded as the inverse of the Edgeworth expansion.

Part I: Existence of Edgeworth Expansion.

Under three assumptions on the joint distribution in which part of the variables can be discrete, the validity of the Edgeworth expansion was shown in Jensen (1989) through a conditioning argument. In this section, we shall show the existence of the Edgeworth expansion of the pivoting quantity for the confidence intervals by checking these three assumptions.

We note first that by (2.1), $\beta_{(p+1) \times 1}$ can be written as a smooth function of the moments $\mathbf{E}(X_j Y)$'s and $\mathbf{E}(X_{j_1}^{k_1} X_{j_2}^{k_2})$'s, where $1 \leq j \leq p$, $1 \leq j_1 < j_2 \leq p$, $k_1, k_2 \geq 0$, and $k_1 + k_2 \leq 2$. Moreover, by the Central Limit Theorem, the asymptotic distribution of the least square estimate $\hat{\beta}$ in (2.2) is given by

$$\sqrt{n}(\hat{\beta} - \beta) \rightarrow \mathcal{N}\left(\mathbf{0}, \mathbf{E}(\vec{\mathbf{X}}\vec{\mathbf{X}}^T)^{-1} \mathbf{E}\{(Y - \vec{\mathbf{X}}^T \beta)^2 \vec{\mathbf{X}}\vec{\mathbf{X}}^T\} \mathbf{E}(\vec{\mathbf{X}}\vec{\mathbf{X}}^T)^{-1}\right). \quad (\text{A.1})$$

Starting with the general case when $p \geq 4$, we note that the $(p+1) \times 1$ vector of the asymptotic variances of each $\sqrt{n}(\hat{\beta}_j - \beta_j)$, $1 \leq j \leq p$, is again a smooth function of the moments $\mathbf{E}(X_{j_1}^{k_1} X_{j_2}^{k_2} X_{j_3}^{k_3} X_{j_4}^{k_4})$'s, $\mathbf{E}(X_{j_5}^{k_5} X_{j_6}^{k_6} X_{j_7}^{k_7} Y)$'s, and $\mathbf{E}(X_{j_8}^{k_8} X_{j_9}^{k_9} Y^2)$'s, where $1 \leq j_1 < j_2 < j_3 < j_4 \leq p$, $1 \leq j_5 < j_6 < j_7 \leq p$, $1 \leq j_8 < j_9 \leq p$, $k_1, k_2, \dots, k_9 \geq 0$, $k_1 + k_2 + k_3 + k_4 \leq 4$, $k_5 + k_6 + k_7 \leq 3$, and $k_8 + k_9 \leq 2$. These moments for the asymptotic mean and variance of $\hat{\beta}$ can all be consistently estimated by their corresponding sample moments.

We now collect all related monomials of X_j and Y into a random vector \mathbf{W}_p such that

$$\mathbf{W}_p = (\{X_{j_1}^{k_1} X_{j_2}^{k_2} X_{j_3}^{k_3} X_{j_4}^{k_4}\}, \{X_{j_5}^{k_5} X_{j_6}^{k_6} X_{j_7}^{k_7} Y\}, \{X_{j_8}^{k_8} X_{j_9}^{k_9} Y^2\})^T, \quad (\text{A.2})$$

where $1 \leq j_1 < j_2 < j_3 < j_4 \leq p$, $1 \leq j_5 < j_6 < j_7 \leq p$, $1 \leq j_8 < j_9 \leq p$, $k_1, k_2, \dots, k_9 \geq 0$, $k_1 + k_2 + k_3 + k_4 \leq 4$, $k_5 + k_6 + k_7 \leq 3$, and $k_8 + k_9 \leq 2$. For $p \geq 4$, the dimension of \mathbf{W}_p is

$$d_p = \binom{p}{4} \binom{8}{4} + \binom{p}{3} \binom{6}{3} + \binom{p}{2} \binom{4}{2} = \frac{1}{12} p(p-1)(35p^2 - 135p + 166).$$

For $1 \leq p \leq 3$, we can take a similar approach to write out \mathbf{W}_p .

1. For $p = 1$, $\mathbf{W}_1 = (\{X_1^{k_1}\}, \{X_1^{k_2}Y\}, \{X_1^{k_3}Y^2\})^T$ where $0 \leq k_1 \leq 4$, $0 \leq k_2 \leq 3$, and $0 \leq k_3 \leq 2$. The dimension of \mathbf{W}_1 is $d_1 = 12$.
2. For $p = 2$, $\mathbf{W}_2 = (\{X_1^{k_1}X_2^{k_2}\}, \{X_1^{k_3}X_2^{k_4}Y\}, \{X_1^{k_5}X_2^{k_6}Y^2\})^T$ where $k_1 + k_2 \leq 4$, $k_3 + k_4 \leq 3$, and $k_5 + k_6 \leq 2$. The dimension of \mathbf{W}_2 is $d_2 = 31$.
3. For $p = 3$, $\mathbf{W}_3 = (\{X_1^{k_1}X_2^{k_2}X_3^{k_3}\}, \{X_1^{k_4}X_2^{k_5}X_3^{k_6}Y\}, \{X_{j_7}^{k_7}X_{j_8}^{k_8}Y^2\})^T$, where $1 \leq j_7 < j_8 \leq 3$, $k_1, k_2, \dots, k_9 \geq 0$, $k_1 + k_2 + k_3 \leq 4$, $k_4 + k_5 + k_6 \leq 3$, and $k_7 + k_8 \leq 2$. The dimension of \mathbf{W}_3 is $d_3 = 73$.

With \mathbf{W}_p defined, we can write $\beta_{(p+1) \times 1} = \beta(\mathbf{E}(\mathbf{W}_p))$. We can also write the vector of the asymptotic variances of $\sqrt{n}(\hat{\beta}_j - \beta_j)$, $1 \leq j \leq p$, as $\sigma_{(p+1) \times 1}^2 = \sigma^2(\mathbf{E}(\mathbf{W}_p))$. From the n i.i.d. samples of $\vec{\mathbf{V}} = (X_1, \dots, X_p, Y)^T$, we can form the \mathbf{W}_p 's as $\mathbf{W}_{p,1}, \dots, \mathbf{W}_{p,n}$. The estimates of β and σ^2 can thus be written as $\hat{\beta} = \beta(\bar{\mathbf{W}}_p)$ and $\hat{\sigma}^2 = \sigma^2(\bar{\mathbf{W}}_p)$, respectively, where $\bar{\mathbf{W}}_p = (1/n) \sum_{i=1}^n \mathbf{W}_{p,i}$.

To check the three assumptions in Jensen (1989), we first denote the random vector of discrete variables in \mathbf{W}_p by \mathbf{W}_D with dimension d_D and the random vector of the rest of the variables in \mathbf{W}_p by \mathbf{W}_C with dimension d_C . We shall use the same subscripts D and C for other related quantities to distinguish discrete variables from others. Note that in \mathbf{W}_C , some variables are products of discrete and continuous variables. The distributions of these product variables may not be absolutely continuous with respect to the Lebesgue measure. Nonetheless, the characteristic function of \mathbf{W}_C still exists, and the proof in Jensen (1989) still holds with this relaxation. Thus, we continue to check these three conditions.

Assumptions 1(i), 1(ii), 1(iii), 3(i), and 3(ii) pertain to the cumulants of \mathbf{W}_p . Under the existence of the moment generating function of $\vec{\mathbf{V}}$, these assumptions are satisfied.

To check Assumptions 2(i) and 2(ii), note that the pivoting quantity for the confidence intervals for β is $(\hat{\beta} - \beta) \circ \sigma^{\circ(-1)}$, where \circ denotes the Hadamard product. Let $\tilde{\mathbf{W}}_p = \bar{\mathbf{W}}_p - \mathbf{E}(\mathbf{W}_p)$ be the centered version of $\bar{\mathbf{W}}_p$. The function $g(\cdot)$ in Jensen (1989) for regression can be written as

$$g(\tilde{\mathbf{W}}_p) = \{\beta(\tilde{\mathbf{W}}_p + \mathbf{E}(\mathbf{W}_p)) - \beta(\mathbf{E}(\mathbf{W}_p))\} \circ \sigma(\mathbf{E}(\mathbf{W}_p))^{\circ(-1)}. \tag{A.3}$$

It is easy to see that $g(\mathbf{0}) = \mathbf{0}$. Furthermore, it can be shown that, since the distribution of $\vec{\mathbf{V}}$ is non-degenerate, the derivatives of g with respect to $\tilde{\mathbf{W}}_C$ exist, are continuous in a neighborhood of $\mathbf{0}$, and the $d_C \times (p + 1)$ matrix of the derivatives has full rank.

To check Assumptions 1(iv) and 1(v), we first note that the derivatives of the characteristic functions are bounded by appropriate moments, which all

exist under the assumptions of Theorem 1. Therefore, we only need to show that the characteristic function of \mathbf{W}_p is bounded away from 1. To show this, we decompose this characteristic function by conditioning on \mathbf{W}_D ,

$$\mathbf{E}(e^{it^T \mathbf{W}_p}) = \sum_{\mathbf{w}_D} e^{it_D^T \mathbf{w}_D} \mathbf{E}(e^{it_C^T \mathbf{W}_C} | \mathbf{W}_D = \mathbf{w}_D), \quad (\text{A.4})$$

where the dimensions of \mathbf{t} , \mathbf{t}_D , and \mathbf{t}_C are d_p , d_D , and d_C respectively. To verify Assumptions 1(iv) and 1(v), it now suffices to show that $|\mathbf{E}[e^{it_C^T \mathbf{W}_C} | \mathbf{W}_D = \mathbf{w}_D]| < 1$ for each \mathbf{w}_D . This proof is patterned after the arguments in Section 2.4 of Hall (1992).

Denote the joint density of $\vec{\mathbf{V}}_C$ condition on w_D by $f_{w_D}(\mathbf{x}_C, y) : \mathbb{R}^{p_C} \rightarrow \mathbb{R}$. We first approximate this density through simple functions. Given $\epsilon > 0$, let $f_1(\mathbf{x}_C, y) = \sum_{m=1}^M c_m I((\mathbf{x}_C, y) \in \mathbf{S}_m)$, where c_m 's are appropriate constants, \mathbf{S}_m 's are appropriate rectangular prisms, and $1 \leq m \leq M$ for some appropriate $M > 0$ such that $\int_{\mathbb{R}^{p_C}} |f_{w_D}(\mathbf{x}_C, y) - f_1(\mathbf{x}_C, y)| d\mathbf{x}_C dy < \epsilon$. We can then focus on showing that $\lim_{\|\mathbf{t}_C\|_2 \rightarrow \infty} |\int_{\mathbf{S}_m} e^{it_C^T \mathbf{w}_C} d\mathbf{x}_C dy| \rightarrow 0$ for each m .

Let $\mathbf{t}_C = \mathbf{t}_C(u)$, where u is an index that diverges to infinity. Through a subsequence argument, we can assume without loss of generality that for some $1 \leq h \leq d_C$, and with $s(h, \tilde{h}, u) = \mathbf{t}_{\tilde{h}}(u)/\mathbf{t}_h(u)$, the limit $s(h, \tilde{h}) = \limsup_{u \rightarrow \infty} s(h, \tilde{h}, u)$ exists for $1 \leq \tilde{h} \leq d_C$, and $|s(h, \tilde{h})| \leq 1$. We now can take $\mathbf{s} = \{s(h, \tilde{h}, u)\}$ and write the real part of the integral $\int_{\mathbf{S}_m} e^{it_C^T \mathbf{w}_C} d\mathbf{x}_C dy$ as $\int_{\mathbf{S}_m} \cos(t_h \mathbf{s}^T \mathbf{w}_C) d\mathbf{x}_C dy$. Note that each entry in \mathbf{w}_C is a monomial of x_j 's and y . Thus, $\mathbf{s}^T \mathbf{w}_C$ is a polynomial of x_j 's and y . We can now take transformations of variables to show that $\int_{\mathbf{S}_m} \cos(t_h \mathbf{s}^T \mathbf{w}_C) d\mathbf{x}_C dy = O(t_h^{-1})$. Similarly, the imaginary part of the integral can be shown to be $O(t_h^{-1})$. These facts entail that the magnitude of the conditional characteristic function is strictly less than 1 when $\|\mathbf{t}_C\|_2$ is large, which completes the proof.

Part II: The Asymptotic Accuracy of Double Bootstrap CIs

With the existence of the Edgeworth expansion, we develop the asymptotic accuracy of the two-sided double-bootstrap CI for regression. In this section, we use θ_0 to denote a generic β_j and use $\hat{\theta}$ to denote the corresponding $\hat{\beta}_j$. We show here only the proof for the two-sided `perc-cal` confidence intervals \mathcal{I}_1 . The one-sided case for \mathcal{I}_1 is proved in a similar (and easier) manner. The techniques used in this proof are patterned after those in Section 3.11 in Hall (1992) but are reorganized for readability and included so that our analysis is self-contained.

Consider the distribution of $\hat{A}(\bar{\mathbf{W}}^*) = (\hat{\theta}^* - \hat{\theta})/\hat{\sigma}$, where $\bar{\mathbf{W}}^*$ is the bootstrap version of $\bar{\mathbf{W}}_p$. For any $0 < \gamma < 1$, the quantile estimate \hat{v}_γ satisfies

$$\mathbf{P}\left(\sqrt{n}\frac{\hat{\theta}^* - \hat{\theta}}{\hat{\sigma}} \leq \hat{v}_\gamma \mid \mathbf{W}_{p,1}, \dots, \mathbf{W}_{p,n}\right) = \gamma. \tag{A.5}$$

Due to the existence of the Edgeworth expansion as described in Part I, we can write \hat{v}_γ in the standard normal quantile z_γ through the Cornish-Fisher expansion:

$$\hat{v}_\gamma = z_\gamma + n^{-1/2}\hat{p}_1(z_\gamma) + n^{-1}\hat{p}_2(z_\gamma) + O_p(n^{-3/2}) \tag{A.6}$$

where \hat{p}_1 and \hat{p}_2 are polynomials whose coefficients are sample estimates of those of p_1 and p_2 , and the coefficients of p_1 and p_2 depend only on the moments of \mathbf{W}_p . Given the condition that all moments of $\vec{\mathbf{V}}$ exist, all of these estimates are root- n consistent.

Now consider the quantile \hat{w}_λ in the bootstrap distribution of $\hat{\theta}^*$ such that

$$\mathbf{P}(\hat{\theta}^* \leq \hat{w}_\lambda \mid \mathbf{W}_{p,1}, \dots, \mathbf{W}_{p,n}) = \lambda. \tag{A.7}$$

By comparing (A.7) and (A.5) with $\gamma = \lambda$, we see

$$\hat{w}_\lambda = \hat{\theta} + n^{-1/2}\hat{\sigma}\hat{v}_\lambda = \hat{\theta} + n^{-1/2}\hat{\sigma}\{z_\lambda + n^{-1/2}\hat{p}_1(z_\lambda) + n^{-1}\hat{p}_2(z_\lambda) + O_p(n^{-3/2})\} \tag{A.8}$$

Thus, by Proposition 3.1 in Hall (1992), we have

$$\begin{aligned} &\mathbf{P}(\theta_0 \in (-\infty, \hat{w}_\lambda)) \\ &= \mathbf{P}(\theta_0 \leq \hat{\theta} + n^{-1/2}\hat{\sigma}\{z_\lambda + n^{-1/2}\hat{p}_1(z_\lambda) + n^{-1}\hat{p}_2(z_\lambda) + O_p(n^{-3/2})\}) \\ &= \lambda + n^{-1/2}r_1(z_\lambda)\phi(z_\lambda) + n^{-1}r_2(z_\lambda)\phi(z_\lambda) + O(n^{-3/2}) \end{aligned} \tag{A.9}$$

where ϕ is the density of the standard normal distribution, and r_1 and r_2 are even and odd polynomials whose coefficients can be root- n consistently estimated.

Let $\xi = 2(1 - \alpha/2 - \lambda)$ and $\lambda = 1 - \alpha/2 + \xi/2$. To find a proper λ for the perc-cal interval \mathcal{I}_2 is now to find ξ such that

$$\mathbf{P}(\theta_0 \in (\hat{w}_{1-\lambda}, \hat{w}_\lambda)) = \mathbf{P}(\theta_0 \in (\hat{w}_{\alpha/2-\xi/2}, \hat{w}_{1-\alpha/2+\xi/2})) = 1 - \alpha. \tag{A.10}$$

Note that the coverage probability of a two-sided CI can be written as

$$\begin{aligned} &\mathbf{P}(\theta_0 \in (\hat{w}_{1-\lambda}, \hat{w}_\lambda)) \\ &= \mathbf{P}(\theta_0 \leq \hat{w}_\lambda) - \mathbf{P}(\theta_0 \leq \hat{w}_{1-\lambda}) \\ &= 2\lambda - 1 + 2n^{-1}r_2(z_\lambda)\phi(z_\lambda) + O(n^{-2}). \\ &= 1 - \alpha + \xi + 2n^{-1}r_2(z_{1-\alpha/2+\xi/2})\phi(z_{1-\alpha/2+\xi/2}) + O(n^{-2}) \end{aligned} \tag{A.11}$$

The cancellation of the $O(n^{-1/2})$ term due to that $-z_{1-\lambda} = z_\lambda$ and that r_1 is an even polynomial is crucial for the improvement in double-bootstrap. To achieve the accuracy of the coverage in Theorem 1, we would like to choose ξ such that

$$\xi = -2n^{-1}r_2(z_{1-\alpha/2+\xi/2})\phi(z_{1-\alpha/2+\xi/2}) + O(n^{-2}) \tag{A.12}$$

Now consider the second-level bootstrap, in which we calibrate $\hat{\xi}$ for $\hat{\lambda} = 1 - \alpha/2 + \hat{\xi}/2$ in the `perc-cal` intervals. Through a similar argument for the first-level bootstrap, we see that the calibrated $\hat{\xi}$ satisfies that

$$\hat{\xi} = -2n^{-1}\hat{r}_2(z_{1-\alpha/2+\hat{\xi}/2})\phi(z_{1-\alpha/2+\hat{\xi}/2}) + O_p(n^{-2}) \quad (\text{A.13})$$

so that

$$\hat{\xi} - \xi = O_p(n^{-3/2}). \quad (\text{A.14})$$

Finally, consider the coverage probability of the double-bootstrap CI $(\hat{w}_{\alpha/2-\hat{\xi}/2}, \hat{w}_{1-\alpha/2+\hat{\xi}/2})$. By (A.8), the Taylor expansion

$$z_{\gamma+\epsilon} = z_{\gamma} + \epsilon\phi(z_{\gamma})^{-1} + O(\epsilon^2), \quad (\text{A.15})$$

and the derivations for (3.36) in Hall (1992), we have

$$\begin{aligned} & \mathbf{P}(\theta_0 \in (-\infty, \hat{w}_{1-\alpha/2+\hat{\xi}/2})) \\ &= \mathbf{P}\left(\sqrt{n}\frac{\hat{\theta} - \theta_0}{\hat{\sigma}} > -z_{1-\alpha/2+\hat{\xi}/2} - n^{-1/2}\hat{p}_1(z_{1-\alpha/2+\hat{\xi}/2}) - n^{-1}\hat{p}_2(z_{1-\alpha/2+\hat{\xi}/2}) \right. \\ & \quad \left. + \dots\right) \\ &= \mathbf{P}\left(\sqrt{n}\frac{\hat{\theta} - \theta_0}{\hat{\sigma}} > -z_{1-\alpha/2+\xi/2} - n^{-1/2}\hat{p}_1(z_{1-\alpha/2+\xi/2}) - \frac{1}{2}(\hat{\xi} - \xi)\phi(z_{1-\alpha/2})^{-1} \right. \\ & \quad \left. - n^{-1}\hat{p}_2(z_{1-\alpha/2+\xi/2}) + O_p(n^{-2})\right) \\ &= \mathbf{P}\left(\sqrt{n}\frac{\hat{\theta} - \theta_0}{\hat{\sigma}} > -z_{1-\alpha/2+\xi/2} - n^{-1/2}\hat{p}_1(z_{1-\alpha/2+\xi/2}) - n^{-1}\hat{p}_2(z_{1-\alpha/2+\xi/2}) \right. \\ & \quad \left. + \dots + \frac{1}{2}(\hat{\xi} - \xi)\phi(z_{1-\alpha/2})^{-1}\right) + O(n^{-2}) \\ &= \mathbf{P}(\theta_0 < \hat{w}_{1-\alpha/2+\xi/2}) + n^{-3/2}bz_{1-\alpha/2}\phi(z_{1-\alpha/2}) + O(n^{-2}) \end{aligned} \quad (\text{A.16})$$

where the constant b is defined through

$$\mathbf{E}\left(\sqrt{n}\frac{\hat{\theta} - \theta_0}{\hat{\sigma}}n^{3/2}\frac{\hat{\xi} - \xi}{2}\right) = b + O(n^{-1}). \quad (\text{A.17})$$

The $O(n^{-1})$ term is derived as in equation (3.35) in Hall (1992). Similarly,

$$\mathbf{P}(\theta_0 \in (-\infty, \hat{w}_{\alpha/2-\hat{\xi}/2})) = \mathbf{P}(\theta_0 \in (-\infty, \hat{w}_{\alpha/2-\xi/2})) - n^{-3/2}bz_{\alpha/2}\phi(z_{\alpha/2}) + O(n^{-2}) \quad (\text{A.18})$$

Now

$$\mathbf{P}(\theta_0 \in (\hat{w}_{\alpha/2-\hat{\xi}/2}, \hat{w}_{1-\alpha/2+\hat{\xi}/2}))$$

$$\begin{aligned}
&= \mathbf{P}(\theta_0 \in (-\infty, \hat{w}_{1-\alpha/2+\hat{\xi}/2})) - \mathbf{P}(\theta_0 \in (-\infty, \hat{w}_{\alpha/2-\hat{\xi}/2})) \\
&= \mathbf{P}(\theta_0 \in (-\infty, \hat{w}_{1-\alpha/2+\xi/2})) + n^{-3/2}bz_{1-\alpha/2}\phi(z_{1-\alpha/2}) + O(n^{-2}) \quad (\text{A.19}) \\
&\quad - (\mathbf{P}(\theta_0 \in (-\infty, \hat{w}_{\alpha/2-\xi/2})) - n^{-3/2}bz_{\alpha/2}\phi(z_{\alpha/2}) + O(n^{-2})) \\
&= 1 - \alpha + O(n^{-2}),
\end{aligned}$$

which concludes our proof.

References

- Beran, R. (1987). Prepivoting to reduce level error of confidence sets. *Biometrika* **74**, 457–468.
- Beran, R. (1988). Prepivoting test statistics: A bootstrap view of asymptotic refinements. *Journal of the American Statistical Association* **83**, 687–697.
- Booth, J. and Presnell, B. (1998). Allocation of monte carlo resources for the iterated bootstrap. *Journal of Computational and Graphical Statistics* **7**, 92–112.
- Booth, J. G. and Hall, P. (1994). Monte Carlo approximation and the iterated bootstrap. *Biometrika* **81**, 331–340.
- Buja, A., Berk, R., Brown, L., George, E., Pitkin, E., Traskin, M., Zhao, L. and Zhang, K. (2015). Models as approximations: A conspiracy of random regressors and model deviations against classical inference in regression. *arXiv:1404.1578 [stat.ME]*.
- Cribari-Neto, F., Souza, T. C. and Vasconcellos, K. L. (2007). Inference under heteroskedasticity and leveraged data. *Communications in Statistics-Theory and Methods* **36**, 1877–1888.
- Delaigle, A., Hall, P. and Jamshidi, F. (2015). Confidence bands in non-parametric errors-in-variables regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **77**, 149–169.
- Eddelbuettel, D., François, R., Allaire, J., Chambers, J., Bates, D. and Ushey, K. (2011). Rcpp: Seamless r and c++ integration. *Journal of Statistical Software* **40**, 1–18.
- Efron, B. (1981). Nonparametric standard errors and confidence intervals. *Canadian Journal of Statistics* **9**, 139–158.
- Efron, B. (1983). Estimating the error rate of a prediction rule: improvement on cross-validation. *Journal of the American Statistical Association* **78**, 316–331.
- Efron, B. (1987). Better bootstrap confidence intervals. *Journal of the American Statistical Association* **82**, 171–185.
- Efron, B. and Tibshirani, R. J. (1994). *An Introduction to the Bootstrap*, **57**. CRC press.
- Freedman, D. A. (1981). Bootstrapping regression models. *The Annals of Statistics* **9**, 1218–1228.
- Gonçalves, S. and White, H. (2005). Bootstrap standard error estimates for linear regression. *Journal of the American Statistical Association* **100**, 970–979.
- Hall, P. (1986). On the bootstrap and confidence intervals. *The Annals of Statistics*, 1431–1452.
- Hall, P. (1988). Theoretical comparison of bootstrap confidence intervals. *The Annals of Statistics*, 927–953.
- Hall, P. (1992). *The Bootstrap and Edgeworth Expansion*. Springer.
- Hall, P. and Martin, M. A. (1988). On bootstrap resampling and iteration. *Biometrika* **75**, 661–671.

- Hall, P., Martin, M. A. and Schucany, W. R. (1989). Better nonparametric bootstrap confidence intervals for the correlation coefficient. *Journal of Statistical Computation and Simulation* **33**, 161–172.
- Jensen, J. L. (1989). Validity of the formal edgeworth expansion when the underlying distribution is partly discrete. *Probability Theory and Related Fields* **81**, 507–519.
- Kline, P. and Santos, A. (2012). Higher order properties of the wild bootstrap under misspecification. *Journal of Econometrics* **171**, 54–70.
- Lee, S. M. S. and Young, G. A. (1999). The effect of Monte Carlo approximation on coverage error of double-bootstrap confidence intervals. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **61**, 353–366.
- Loh, W.-Y. (1987). Calibrating confidence coefficients. *Journal of the American Statistical Association* **82**, 155–162.
- MacKinnon, J. G. (2006). Bootstrap methods in econometrics. *Economic Record* **82**, 2–18.
- MacKinnon, J. G. (2013). Thirty years of heteroskedasticity-robust inference. In *Recent Advances and Future Directions in Causality, Prediction, and Specification Analysis*, pp. 437–461. Springer.
- Martin, M. A. (1990a). On bootstrap iteration for coverage correction in confidence intervals. *Journal of the American Statistical Association* **85**, 1105–1118.
- Martin, M. A. (1990b). On the double bootstrap. Technical report, Stanford University, Stanford, CA, Technical Report.
- McCarthy, D. (2016). perccal: Implementing double bootstrap linear regression confidence intervals using the ‘perc-cal’ method. R package version 1.0. <https://CRAN.R-project.org/package=perccal>.
- Scholz, F. (2007). The bootstrap small sample properties. *Boeing Computer Services, Research and Technology, Tech. Rep.*
- Spokoiny, V. and Zhilova, M. (2014). Bootstrap confidence sets under a model misspecification. *arXiv preprint arXiv:1410.0347*.
- van der Vaart, A. W. (1998). *Asymptotic Statistics*. Cambridge University Press, Cambridge.
- White, H. (1980). A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica: Journal of the Econometric Society*, 817–838.

Goizueta Business School, Emory University, 1300 Clifton Road NE, Room 517 Atlanta, GA 30322, USA.

E-mail: daniel.mccarthy@emory.edu

Department of Statistics and Operations Research, 306 Hanes Hall, CB# 3260, Chapel Hill, NC, 27599, USA.

E-mail: zhangk@email.unc.edu

Department of Statistics, The Wharton School, University of Pennsylvania, Philadelphia, PA 19104, USA.

E-mail: lbrown@wharton.upenn.edu

Department of Statistics, The Wharton School, University of Pennsylvania, Philadelphia, PA 19104, USA.

E-mail: berkr@sas.upenn.edu

Department of Statistics, The Wharton School, University of Pennsylvania, Philadelphia, PA

19104, USA.

E-mail: buja.at.wharton@gmail.com

Department of Statistics, The Wharton School, University of Pennsylvania, Philadelphia, PA
19104, USA.

E-mail: edgerge@wharton.upenn.edu

Department of Statistics, The Wharton School, University of Pennsylvania, Philadelphia, PA
19104, USA.

E-mail: lzhao@wharton.upenn.edu

(Received December 2016; accepted January 2017)