



Annual Review of Economics

Bootstrap Methods in Econometrics

Joel L. Horowitz

Department of Economics, Northwestern University, Evanston, Illinois 60208, USA;
email: joel-horowitz@northwestern.edu

Annu. Rev. Econ. 2019. 11:193–224

The *Annual Review of Economics* is online at
economics.annualreviews.org

<https://doi.org/10.1146/annurev-economics-080218-025651>

Copyright © 2019 by Annual Reviews.
All rights reserved

Keywords

resampling, confidence interval, hypothesis test, asymptotic refinement

Abstract

The bootstrap is a method for estimating the distribution of an estimator or test statistic by resampling one's data or a model estimated from the data. Under conditions that hold in a wide variety of econometric applications, the bootstrap provides approximations to distributions of statistics, coverage probabilities of confidence intervals, and rejection probabilities of hypothesis tests that are more accurate than the approximations of first-order asymptotic distribution theory. The reductions in the differences between true and nominal coverage or rejection probabilities can be very large. In addition, the bootstrap provides a way to carry out inference in certain settings where obtaining analytic distributional approximations is difficult or impossible. This article explains the usefulness and limitations of the bootstrap in contexts of interest in econometrics. The presentation is informal and expository. It provides an intuitive understanding of how the bootstrap works. Mathematical details are available in the references that are cited.

1. INTRODUCTION

The bootstrap is a method for estimating the distribution of an estimator or test statistic by re-sampling one's data or a model estimated from the data. It amounts to treating the data as if they were the population for the purpose of evaluating the distribution of interest. Under mild regularity conditions, the bootstrap yields an approximation to the distribution of an estimator or test statistic that is at least as accurate as and often more accurate than the approximation obtained from first-order asymptotic theory. Thus, the bootstrap provides a way to substitute computation for mathematical analysis if calculating the asymptotic distribution of an estimator or statistic is difficult, and it often provides a practical way to improve upon first-order approximations.

Improvements on first-order approximations are called asymptotic refinements. The bootstrap's ability to provide asymptotic refinements is important because first-order asymptotic theory often gives poor approximations to the distributions of test statistics with the sample sizes available in applications. Therefore, the nominal probability that a test based on an asymptotic critical value rejects a true null hypothesis or the nominal coverage probability of a confidence interval can be very different from the true rejection probability (RP) or coverage probability. The information matrix test of White (1982) is a well-known example of a test in which large finite-sample errors in the RP can occur when asymptotic critical values are used (Horowitz 1994, Kennan & Neumann 1988, Orme 1990, Taylor 1987). The bootstrap often provides a tractable way to reduce or eliminate finite-sample errors in the RPs of statistical tests.

The bootstrap was introduced by Efron (1979) and was the object of much research in the ensuing 25 years. Most of the theory of the bootstrap and methods for implementing it were developed during this period. The results of this research are synthesized in books by Beran & Ducharme (1991), Davison & Hinkley (1997), Efron & Tibshirani (1993), Hall (1992), Mammen (1992), and Shao & Tu (1995). Hall (1994), Horowitz (1997, 2001), Jeong & Maddala (1993), and Vinod (1993) provide reviews with an econometric orientation. In addition to topics covered by these books and reviews, this article summarizes several recent extensions of bootstrap methods. These include methods for obtaining confidence bands in nonparametric estimation, methods for carrying out nonasymptotic inference in parametric models, and methods for high-dimensional models.

This article explains and illustrates the usefulness and limitations of the bootstrap in contexts of interest in econometrics. The presentation is informal and expository. Its aim is to provide an intuitive understanding of how the bootstrap works and a feeling for its practical value in econometrics. It does not provide a mathematically detailed or rigorous treatment of the theory of the bootstrap. Mathematically rigorous treatments are available in the books by Beran & Ducharme (1991) and Hall (1992) and in references that are cited below.

Although the bootstrap is often very accurate, it can be inaccurate and misleading if it is used incorrectly. Examples of incorrect uses include inference about a parameter that is on the boundary of the parameter set, inference about the maximum or minimum of random variables, and inference in the presence of weak instruments. The difference between correct and incorrect use is not always intuitive. A further objective of this article is to explain the difference between correct and incorrect use.

The variety of different bootstrap methods and applications and the associated literature are very large. A given bootstrap method typically works well in some settings but not in others. It is not possible to provide an exhaustive review of bootstrap methods or the bootstrap literature in one article, and this article does not attempt to do so. Instead, the article focuses on a few methods and settings that arise frequently in applied econometrics. Owing to space limitations, this article does not present Monte Carlo or other illustrations of the numerical performance

of the bootstrap. Such illustrations are provided by Horowitz (2001) and other references cited throughout the article.

The remainder of this article is divided into eight sections. Section 2 explains the bootstrap sampling procedure and gives conditions under which the bootstrap distribution of a statistic is a consistent estimator of the statistic's asymptotic distribution. Section 3 explains when and why the bootstrap provides asymptotic refinements. This section concentrates on data that are independent random samples from a distribution and statistics that either are smooth functions of sample moments or can be approximated with asymptotically negligible error by such functions (the smooth function model). Section 4 extends the results of Section 3 to statistics that do not satisfy the assumptions of the smooth function model. Section 5 describes the use of the bootstrap in nonparametric estimation of a conditional mean or quantile function. Section 6 describes the multiplier and wild bootstraps. Section 7 presents results on the application of the bootstrap to statistics based on the least absolute shrinkage and selection operator (LASSO) and adaptive LASSO parameter estimates. Section 8 discusses the application of the bootstrap to dependent data, and Section 9 presents concluding comments.

2. THE BOOTSTRAP WITH INDEPENDENTLY AND IDENTICALLY DISTRIBUTED DATA

Much research in empirical economics is based on data that are a random sample from a distribution. This type of data includes panel data based on randomly sampled individuals, where the data consist of observations of individuals and other relevant variables over time. To minimize the complexity of the discussion in this section, let the data consist of a random sample of size n from the distribution of a possibly vector-valued random variable \mathbf{X} . If the data are used in a model that has dependent and explanatory variables, then both types of variables are included among the components of \mathbf{X} . Denote the data by $\{\mathbf{X}_i : i = 1, \dots, n\}$. Let F_0 denote the distribution function of \mathbf{X} and F_n denote the empirical distribution function of the sample. That is, for any vector \mathbf{z} whose dimension is the same as that of \mathbf{X} , we have

$$F_n(\mathbf{z}) = \frac{1}{n} \sum_{i=1}^n I(\mathbf{X}_i \leq \mathbf{z}),$$

where $I(\cdot)$ is the indicator function, and the inequality holds component by component. Suppose that the object of interest is a statistic $T_n(\mathbf{X}_1, \dots, \mathbf{X}_n)$. For example, if X is a scalar random variable, T_n might be the t -statistic for testing the null hypothesis $H_0 : E(X) = 0$, where E is the expectation operator, and $E(X)$ is the population mean of X . In this case, we have

$$T_n(X_1, \dots, X_n) = \bar{X}/s_n,$$

where \bar{X} is the sample average, and s_n is the sample standard deviation:

$$\bar{X} = \frac{1}{n} \sum_{i=1}^n X_i; \quad s_n^2 = \frac{1}{n-1} \sum_{i=1}^n (X_i - \bar{X})^2.$$

In most applications, the distribution of X and, therefore, of T_n is unknown. It is necessary to estimate the distribution of T_n to obtain a critical value for testing H_0 . Asymptotic distribution theory provides one way of estimating this distribution. In this simple example, T_n is asymptotically distributed as $N(0, 1)$. Moreover, the $N(0, 1)$ distribution provides a very accurate approximation to

the unknown exact finite sample distribution of T_n with samples of the sizes typically encountered in empirical economics. In more complicated settings, however, the asymptotic approximation may be inaccurate or the asymptotic distribution of T_n hard to compute.

2.1. The Nonparametric Bootstrap and Bootstrap Consistency

The nonparametric bootstrap estimates the distribution of T_n by treating the empirical distribution of the data as if it were the population distribution. In other words, the nonparametric bootstrap conditions on the data and pretends that they were drawn from a population whose distribution function is F_n instead of F_0 . In this pretend world, \mathbf{X} has a known distribution (the empirical distribution of the data), and the exact finite-sample distribution of T_n can be computed with any desired accuracy by simulation. The steps of this procedure are

1. Generate a bootstrap sample $\{\mathbf{X}_i^* : i = 1, \dots, n\}$ by sampling the original data randomly with replacement.
2. Compute $T_n^* = T_n(\mathbf{X}_1^*, \dots, \mathbf{X}_n^*)$.
3. Use the results of many repetitions of steps 1 and 2 to compute the empirical probability of the event $T_n^* \leq \tau$ for any τ . This probability is the proportion of repetitions in which the event $T_n^* \leq \tau$ occurs.

Brown (2000) and Hall (1992, appendix II) discuss simulation methods that take advantage of techniques for reducing Monte Carlo sampling variation. Regardless of the simulation method that is used, however, the essential characteristic of the nonparametric bootstrap is treating F_n as if it, instead of F_0 , were the distribution function of the population from which the data were sampled.

Now let $G_n(\cdot, F)$ be the distribution function of T_n when the data are sampled randomly from a population whose distribution function is F . Thus, the true but unknown distribution function of T_n is $P(T_n \leq \tau) = G_n(\tau, F_0)$. The bootstrap approximates $G_n(\tau, F_0)$ by $G_n(\tau, F_n)$. The Glivenko-Cantelli theorem ensures that $|F_n(\mathbf{z}) - F_0(\mathbf{z})| \rightarrow 0$ almost surely uniformly over \mathbf{z} . Therefore, intuition suggests that $G_n(\tau, F_n)$ should be close to $G_n(\tau, F_0)$ if n is large, and if G_n is continuous at F_0 in the sense that $G_n(\tau, F)$ is close to $G_n(\tau, F_0)$ whenever F is close to F_0 . Let $G_\infty(\tau, F_0)$ denote the asymptotic distribution function of T_n . Since by definition $G_n(\tau, F_0)$ converges to $G_\infty(\tau, F_0)$ as $n \rightarrow \infty$, one can hope that $G_n(\tau, F_n) \rightarrow G_\infty(\tau, F_0)$ occurs as $n \rightarrow \infty$. In other words, the bootstrap estimate of the distribution function of T_n converges to the asymptotic distribution function of T_n as $n \rightarrow \infty$. This property of the bootstrap is called consistency. The formal definition of bootstrap consistency is

Definition 1: Let P_n denote the probability distribution of the sample $\{\mathbf{X}_i : i = 1, \dots, n\}$. The bootstrap estimator $G_n(\cdot, F_n)$ is consistent if, for each $\varepsilon > 0$ and F_0 ,

$$\lim_{n \rightarrow \infty} P_n \left[\sup_{\tau} |G_n(\tau, F_n) - G_\infty(\tau, F_0)| > \varepsilon \right] = 0.$$

Beran & Ducharme (1991) give general conditions under which the bootstrap is consistent, but these conditions are hard to check. Mammen (1992) gives less general but much more useful conditions for bootstrap consistency. Specifically, Mammen (1992) gives necessary and sufficient conditions for the bootstrap to consistently estimate the distribution of a linear functional of F_0 . These conditions are important because they are often easy to check, and many econometric estimators and test statistics are asymptotically equivalent to linear functionals of some F_0 . Gill (1989) and Hall (1990) give related theorems. Mammen's (1992) theorem is

Theorem 1 (Mammen 1992): Let $\{\mathbf{X}_i : i = 1, \dots, n\}$ be a random sample from a population. For a sequence of functions g_n and sequences of numbers t_n and σ_n , define $\bar{g}_n = n^{-1} \sum_{i=1}^n g_n(\mathbf{X}_i)$ and $T_n = (\bar{g}_n - t_n)/\sigma_n$. For the bootstrap sample $\{\mathbf{X}_i^* : i = 1, \dots, n\}$, define $\bar{g}_n^* = n^{-1} \sum_{i=1}^n g_n(\mathbf{X}_i^*)$ and $T_n^* = (\bar{g}_n^* - \bar{g}_n)/\sigma_n$. Let $G_n(\tau) = P(T_n \leq \tau)$ and $G_n^*(T_n^*) = P_n^*(T_n^* \leq \tau)$, where P_n^* is the probability distribution induced by bootstrap sampling conditional on the original data $\{\mathbf{X}_i\}$. Then $G_n^*(\cdot)$ consistently estimates $G_n(\cdot)$ if and only if $T_n \rightarrow^d N(0, 1)$.

Bootstrap sampling is conditional on the original data $\{\mathbf{X}_i\}$. Therefore, in bootstrap sampling and estimation, the original data and any functions of the data are nonstochastic. Only functions of the bootstrap sample $\{\mathbf{X}_i^*\}$ are random in bootstrap sampling.

Maximum likelihood estimators, generalized method of moments estimators, and other commonly encountered extremum estimators and test statistics are asymptotically linear and asymptotically normal (or are asymptotically chi-square quadratic forms of asymptotic normal statistics) under the usual regularity conditions (e.g., Amemiya 1985, chapter 4). Therefore, the nonparametric bootstrap estimates their asymptotic distributions consistently and can be applied to them. Indeed, the bootstrap is often more accurate than asymptotic distribution theory, as is discussed in Section 3.

There are, however, important settings in which a statistic is neither asymptotically linear nor normal and for which the nonparametric bootstrap is inconsistent. These settings are discussed in the sections below.

2.1.1. Manski's maximum score estimator for a binary response model. In Manski's (1975, 1985) model, the binary dependent variable Y , vector of explanatory variables \mathbf{X} , and vector of constant coefficients β are related by

$$Y = \begin{cases} 1 & \text{if } \beta' \mathbf{X} - U \geq 0 \\ -1 & \text{if } \beta' \mathbf{X} - U < 0 \end{cases}$$

$$P(U \leq 0 | \mathbf{X}) = 0.5.$$

Let $\{Y_i, \mathbf{X}_i : i = 1, \dots, n\}$ be a random sample from the distribution of (Y, \mathbf{X}) and $\|\cdot\|$ denote the ℓ_2 norm. For any scalar v , define

$$\text{sgn}(v) = \begin{cases} 1 & \text{if } v \geq 0 \\ -1 & \text{if } v < 0. \end{cases}$$

The maximum score estimator is

$$\hat{\mathbf{b}} = \arg \max_{\|\mathbf{b}\|=1} \sum_{i=1}^n Y_i \text{sgn}(\mathbf{b}' \mathbf{X}_i).$$

Manski (1975, 1985) gives conditions under which $\hat{\mathbf{b}}$ estimates β up to an unidentified scale parameter, but the asymptotic distribution of $\hat{\mathbf{b}}$ after suitable centering and scaling is that of the maximum of a multidimensional Gaussian process with quadratic drift (Cavanagh 1987, Kim & Pollard 1990). The nonparametric bootstrap does not estimate this distribution consistently (Abrevaya & Huang 2005).

2.1.2. Estimation when a parameter is on the boundary of the parameter set. This situation arises in moment inequality models and estimation under shape restrictions, among other settings, where it is not known whether one or more inequality constraints is binding. The following simple example illustrates the problem. Let $\{X_i : i = 1, \dots, n\}$ be a random sample from the $N(\mu, 1)$ distribution, and suppose that it is known that $\mu \geq 0$. Let $\bar{X} = n^{-1} \sum_{i=1}^n X_i$ be the sample average of X . The maximum likelihood estimate of μ is $m = \max(0, \bar{X})$. The exact finite-sample distribution of m is given by

$$P[n^{1/2}(m - \mu) \leq z] = \begin{cases} \Phi(z) & \text{if } m \geq 0, \\ 0 & \text{if } m < 0, \end{cases}$$

where Φ is the standard normal distribution function. This is a censored normal distribution with censoring at the boundary point $m = 0$. If $\mu > 0$, then the probability of censoring approaches 0 as $n \rightarrow \infty$, and the limiting distribution of $n^{1/2}(m - \mu)$ is $N(0, 1)$. However, if $\mu = 0$, then the limiting distribution is

$$P[n^{1/2}m \leq z] = \begin{cases} \Phi(z) & \text{if } z \geq 0. \\ 0 & \text{if } z < 0. \end{cases}$$

It follows from Theorem 1 that the nonparametric bootstrap estimates the distribution of $n^{1/2}(m - \mu)$ consistently if $\mu > 0$. However, asymptotic normality of $n^{1/2}(m - \mu)$ is necessary for consistency according to Theorem 1. Therefore, nonparametric bootstrap does not estimate the distribution of $n^{1/2}(m - \mu)$ consistently if $\mu = 0$. Andrews (2000) describes several alternative resampling methods for estimating the distribution of $n^{1/2}(m - \mu)$ consistently when $\mu = 0$. Andrews & Barwick (2012), Andrews & Han (2009), Bugni (2010, 2016), and Bugni et al. (2015, 2017) discuss bootstrap methods for moment inequality models.

2.1.3. Distribution of the maximum of a sample. The following example, which is based on the work of Bickel & Freedman (1981), is a simplified version of a situation that occurs in empirical models of auctions and search models, among others. Bowlus et al. (2001), Donald & Paarsch (1996), Flinn & Heckman (1982), and Heckman et al. (1997) describe such models. Chernozhukov & Hong (2004) and Hirano & Porter (2003) discuss some of the theoretical issues that are involved.

Let $\{X_i : i = 1, \dots, n\}$ be a random sample from the uniform distribution on $[0, \theta_0]$, where θ_0 is a positive constant whose true but unknown value is $\theta_0 = 1$. The maximum likelihood estimator of θ_0 is $\hat{\theta}_n = \max(X_1, \dots, X_n)$. Define $T_n = n(\hat{\theta}_n - 1)$. As $n \rightarrow \infty$, we have $P(T_n \leq -z) = e^{-z}$ for any $z \geq 0$. Moreover, we have $P(T_n = 0) = 0$ for all n . Let $\{X_i^* : i = 1, \dots, n\}$ be a bootstrap sample that is obtained by sampling the data $\{X_i : i = 1, \dots, n\}$ randomly with replacement. The bootstrap analog of θ_0 is $\hat{\theta}_n^*$ because the bootstrap samples the empirical distribution of the data, and $\hat{\theta}_n^*$ is the upper bound of the support of this distribution. The bootstrap estimator of $\hat{\theta}_n$ is $\theta_n^* = \max(X_1^*, \dots, X_n^*)$. The bootstrap analog of T_n is $T_n^* = n(\theta_n^* - \hat{\theta}_n)$. Let P_n^* denote probability under bootstrap sampling conditional on the original data $\{X_i\}$. Then we have $P_n^*(T_n^* = 0) = 1 - (1 - n^{-1})^n \rightarrow 1 - e^{-1}$ as $n \rightarrow \infty$. Thus, the nonparametric bootstrap does not estimate the distribution of T_n consistently.

2.2. The Parametric and Residual Bootstraps

The distribution of the random variable in a model is sometimes assumed to be known up to a finite-dimensional parameter. This happens, for example, in maximum likelihood estimation and

in ordinary least squares (OLS) estimation of a normal linear model with a fixed design. Let \mathbf{U} be the random variable or vector in a model, and let $F(\mathbf{u}, \boldsymbol{\theta}) = P(\mathbf{U} \leq \mathbf{u})$ be the cumulative distribution function of \mathbf{U} , where F is a known function, and $\boldsymbol{\theta}$ is a finite-dimensional parameter whose true but unknown value is $\boldsymbol{\theta}_0$. Let $\hat{\boldsymbol{\theta}}_n$ be a consistent estimator of $\boldsymbol{\theta}_0$. If $F(\mathbf{u}, \boldsymbol{\theta})$ is a continuous function of $\boldsymbol{\theta}$ in a neighborhood of $\boldsymbol{\theta}_0$ for all \mathbf{u} , then bootstrap samples can be drawn randomly from the distribution whose cumulative distribution function is $F(\mathbf{u}, \hat{\boldsymbol{\theta}}_n)$.

If the distribution of \mathbf{U} is not known, then it may be possible to draw bootstrap samples from a model's residuals. As an example, consider the linear model

$$Y_i = \boldsymbol{\theta}'_0 \mathbf{X}_i + U_i; \quad i = 1, \dots, n, \tag{1}$$

where the \mathbf{X}_i s are fixed in repeated samples and the U_i s are independently and identically distributed random variables with means of zero and finite variances. Let $\hat{\boldsymbol{\theta}}_n$ be the OLS estimator of $\boldsymbol{\theta}_0$, and let $\hat{U}_i = Y_i - \hat{\boldsymbol{\theta}}_n' \mathbf{X}_i$ ($i = 1, \dots, n$) denote the OLS residuals. Bootstrap analogs Y_i^* of the Y_i s can be generated from the equation

$$Y_i^* = \hat{\boldsymbol{\theta}}_n' \mathbf{X}_i + U_i^*; \quad i = 1, \dots, n,$$

where the U_i^* s are sampled randomly with replacement from the \hat{U}_i s. This is called the residual bootstrap because it samples the residuals of the estimated model, not the original data. The residual bootstrap sample is $\{Y_i^*, \mathbf{X}_i : i = 1, \dots, n\}$. The bootstrap analog of $\boldsymbol{\theta}_0$ is $\hat{\boldsymbol{\theta}}_n$ and is non-stochastic in bootstrap sampling, although it is random in sampling from the original population. Let $\boldsymbol{\theta}_n^*$ be the OLS estimator of $\hat{\boldsymbol{\theta}}_n$ based on the bootstrap sample. Then the bootstrap distribution of $n^{1/2}(\boldsymbol{\theta}_n^* - \hat{\boldsymbol{\theta}}_n)$ consistently estimates the population distribution of $n^{1/2}(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0)$. Specifically, we have

$$\lim_{n \rightarrow \infty} \sup_{\mathbf{z}} |P_n^*[n^{1/2}(\boldsymbol{\theta}_n^* - \hat{\boldsymbol{\theta}}_n) \leq \mathbf{z}] - P[n^{1/2}(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0) \leq \mathbf{z}]| = o_p(0),$$

where P_n^* is the probability induced by bootstrap sampling of the U_i^* s conditional on the original data $\{Y_i, \mathbf{X}_i\}$.

The residual approach can also be used in nonparametric mean and quantile regression. In a nonparametric mean regression, for example, $\boldsymbol{\theta}'_0 \mathbf{X}_i$ in Equation 1 is replaced by $g(\mathbf{X}_i) = E(Y|\mathbf{X}_i)$, and $\hat{\boldsymbol{\theta}}_n' \mathbf{X}_i$ is replaced by a nonparametric estimate of $g(\mathbf{X}_i)$. Denote this estimate by $\hat{g}(\mathbf{X}_i)$. Denote the residuals of the nonparametric model by $\tilde{U}_i = Y_i - \hat{g}(\mathbf{X}_i)$, and let $\hat{U}_i = \tilde{U}_i - n^{-1} \sum_{j=1}^n \tilde{U}_j$ denote residuals that are centered to have a sample average of zero. Residual bootstrap samples $\{Y_i^*, \mathbf{X}_i : i = 1, \dots, n\}$ are generated by setting

$$Y_i^* = \hat{g}(\mathbf{X}_i) + U_i^*,$$

where the U_i^* s are sampled randomly from the \hat{U}_i s.

2.3. Subsampling

The nonparametric bootstrap draws samples from the empirical distribution of the data, not the population distribution. The nonparametric bootstrap is inconsistent in certain settings because the two distributions are not the same. Although the empirical distribution of the data converges to the population distribution almost surely as the sample size increases, the distribution of the bootstrap analog of a statistic may not converge to the population distribution of the statistic.

This problem can be avoided by sampling the true population distribution instead of the empirical distribution of the data. It is not possible to draw repeated samples of size n from the population distribution, but it is possible to draw repeated samples of size $m < n$. This is done by drawing samples of size m randomly without replacement from the estimation data. Each such sample is a subsample of the original data and is a random sample of size m from the population distribution of the data.

This method of subsampling is called nonreplacement subsampling. It was originally proposed by Politis & Romano (1994), who show that it consistently estimates the distribution of a statistic under conditions that are much weaker than those required for consistency of the bootstrap estimator. Politis et al. (1997) extend the nonreplacement subsampling method to heteroskedastic time series. Politis et al. (1999) provide a detailed description of nonreplacement subsampling and examples of its application. Bertail et al. (1999) provide a data-based method for choosing a tuning parameter that occurs in nonreplacement subsampling.

To describe the nonreplacement subsampling method, let $t_n = t_n(\mathbf{X}_1, \dots, \mathbf{X}_n)$ be an estimator of the population parameter θ , and set $T_n = \rho_n(t_n - \theta)$, where ρ_n is a normalizing factor that is chosen so that $G_n(\tau, F_0) = P(T_n \leq \tau)$ converges to a nondegenerate limit $G_\infty(\tau, F_0)$ at points τ where the latter function is continuous. For example, if θ is a population mean, then we have $t_n = \bar{X}$ (the sample average) and $\rho_n = n^{1/2}$. Let $\{\mathbf{X}_{i_j} : j = 1, \dots, m\}$ be a subset of $m < n$ observations taken from the original sample $\{\mathbf{X}_i : i = 1, \dots, n\}$. Define $N_{nm} = \binom{n}{m}$ to be the total number of subsets of size m that can be formed. Let t_{mk} denote the estimator t_m evaluated at the k th subset. The nonreplacement subsampling method estimates $G_n(\tau, F_0)$ by

$$G_{nm}(\tau) = N_{nm}^{-1} \sum_{k=1}^{N_{nm}} I[\rho_m(t_{mk} - t_n) \leq \tau].$$

The intuition behind this method is as follows. Each subsample $\{\mathbf{X}_{i_j} : j = 1, \dots, m\}$ is a random sample of size m from the population distribution whose cumulative distribution function is F_0 . Therefore, $G_m(\tau, F_0)$ is the exact sampling distribution of $\rho_m(t_m - \theta)$. Moreover, we have

$$G_m(\tau, F_0) = EI[\rho_m(t_m - \theta) \leq \tau]. \tag{2}$$

The quantity on the right-hand side of Equation 2 cannot be calculated in an application because F_0 and θ are unknown. $G_{nm}(\tau)$ is the estimator of $G_m(\tau, F_0)$ that is obtained by replacing the population expectation with the average over subsamples and θ with t_n . If n is large, but m/n is small, then random fluctuations in t_n are small relative to those in t_m . Accordingly, the sampling distributions of $\rho_m(t_m - t_n)$ and $\rho_m(t_m - \theta)$ are close. Similarly, if N_{nm} is large, then the average over subsamples is a good approximation to the population average. These ideas are formalized by Politis & Romano (1994), who show that if $\rho_m/\rho_n \rightarrow 0$ and $m/n \rightarrow 0$ hold true as $m, n \rightarrow \infty$, and if T_n has a well-behaved asymptotic distribution, then the nonreplacement subsampling method consistently estimates this distribution. The nonreplacement subsampling method also consistently estimates asymptotic critical values for T_n and asymptotic confidence intervals for t_n . Andrews & Guggenberger (2010) show that nonreplacement subsampling is not uniformly consistent in a certain sense. Andrews & Guggenberger (2009) describe methods for overcoming this problem.

In practice, N_{nm} is likely to be very large, which makes G_{nm} hard to compute. This problem can be overcome by replacing the average over all N_{nm} subsamples with the average over a random sample of subsamples (Politis & Romano 1994). These can be obtained by sampling the data $\{\mathbf{X}_i : i = 1, \dots, n\}$ randomly without replacement.

The nonreplacement subsampling method enables one to consistently estimate the asymptotic distributions of statistics under very weak conditions. However, the bootstrap is typically more accurate than nonreplacement subsampling when the former is consistent. Under conditions that are satisfied in most applications of the bootstrap, the errors made by the nonparametric, parametric, and residual bootstrap estimators of a distribution are at most $O_p(n^{-1/2})$ and can be much less. In contrast, the error made by the nonreplacement subsampling estimator is $O_p(n^{-1/3})$ (Bugni 2010, Politis & Romano 1994). Bertail (1997), Hall & Jing (1996), and Politis et al. (1999) describe extrapolation methods for improving this rate, although the improved rate is slower than that of the bootstrap under similar conditions. Thus, the bootstrap estimator of $G_n(\tau, F_0)$ is more accurate than the nonreplacement subsampling estimator in most applications in econometrics. The subsampling method is useful, however, if characteristics of the sampled population or the statistic of interest cause the bootstrap to be inconsistent, or if checking the consistency of the bootstrap is difficult.

Another method of subsampling, called the m out of n bootstrap, consists of drawing $m < n$ observations randomly with replacement from the estimation sample. Except for the size of the bootstrap sample, the m out of n bootstrap is identical to the standard bootstrap. The m out of n bootstrap has properties similar to those of nonreplacement subsampling. Swanepoel (1986) gives conditions under which the m out of n bootstrap consistently estimates the distribution of the maximum of a sample. Andrews (2000) gives conditions under which it consistently estimates the distribution of a parameter on the boundary of the parameter set. Bickel et al. (1997) provide a detailed discussion of the consistency and rates of convergence of the m out of n bootstrap and of an extrapolation method to increase the rate of convergence. Chung & Lee (2001) describe a method for improving the accuracy of the m out of n bootstrap to that of the conventional bootstrap in certain situations.

3. ASYMPTOTIC REFINEMENTS

The term asymptotic refinements refers to the ability of the bootstrap to provide approximations to the distributions of statistics that are more accurate than the approximations of conventional asymptotic distribution theory. This section explains how the bootstrap provides asymptotic refinements for a large class of statistics that are important in applied research. I continue to assume that the data are an independent random sample from a distribution. Throughout this section and in Sections 4 and 5, the term bootstrap refers to the nonparametric bootstrap of Section 2.1 and the parametric and residual bootstraps of Section 2.2, unless otherwise stated.

The class of statistics discussed in this section is called the smooth function model. To define this model, let $\{\mathbf{X}_i : i = 1, \dots, n\}$ be an independent random sample from the distribution of the random variable or vector \mathbf{X} . As in Section 2, it is not necessary to distinguish between dependent and explanatory variables in this section. Accordingly, the components of \mathbf{X} and the \mathbf{X}_i s include any dependent variables. Let $\mathbf{Z}(\cdot)$ be a possibly vector-valued function on the support of \mathbf{X} . Define $\boldsymbol{\theta} = E\mathbf{Z}(\mathbf{X})$, $\mathbf{Z}_i = \mathbf{Z}(\mathbf{X}_i)$, and $\bar{\mathbf{Z}} = n^{-1} \sum_{i=1}^n \mathbf{Z}_i$. Let $H(\mathbf{z})$ be a smooth scalar-valued function whose argument has the same dimension as \mathbf{Z} . Smooth means that $H(\mathbf{z})$ has bounded partial derivatives of sufficiently high order with respect to any combination of the components of \mathbf{z} . The required order of the derivatives depends on the bootstrap application (e.g., bias reduction, a symmetrical confidence interval) and is not specified here. Let s_n^2 be a consistent estimator of $Var\{n^{1/2}[H(\bar{\mathbf{Z}}) - H(\boldsymbol{\theta})]\}$. The statistics in the smooth function model include

1. $H(\bar{\mathbf{Z}})$, which estimates $H(\boldsymbol{\theta})$, and
2. $n^{1/2}[H(\bar{\mathbf{Z}}) - H(\boldsymbol{\theta})]$ and $n^{1/2}[H(\bar{\mathbf{Z}}) - H(\boldsymbol{\theta})]/s_n$, which can be used to form confidence intervals for and test hypotheses about $H(\boldsymbol{\theta})$.

These statistics are all smooth functions H of the sample moments \bar{Z} . Many estimators and test statistics used in applied econometrics either are smooth functions of sample moments or can be approximated by such functions with an approximation error that decreases very rapidly as n increases and is negligible. The OLS estimator of the slope coefficients in a linear regression model and the t statistic for testing a hypothesis about a coefficient are exact functions of sample moments. Maximum likelihood and generalized method of moments estimators of the parameters of nonlinear models can be approximated with asymptotically negligible error by smooth functions of sample moments if the log-likelihood function or moment conditions have sufficiently many derivatives with respect to the unknown parameters.

Some test statistics do not satisfy the assumptions of the smooth function model. Quantile estimators, such as the least absolute deviations (LAD) estimator, do not satisfy the assumptions of the smooth function model because their objective functions are not sufficiently smooth. Non-parametric density and mean-regression estimators and semiparametric estimators that require kernel or other forms of smoothing also do not fit within the smooth function model. Bootstrap methods for such estimators are discussed in Sections 4 and 5.

3.1. Bias Reduction

This section explains how the bootstrap can be used to reduce the finite-sample bias of an estimator. To minimize the complexity of the discussion, it is assumed that θ and Z in the smooth function model are scalars. However, the method outlined in this section applies to any estimator that satisfies the assumptions of the smooth function model.

To begin, define $\mu = H(\theta)$. Observe that $m_n = H(\bar{Z})$ is a consistent estimator of μ , but m_n is biased if H is a nonlinear function. That is, we have $m_n \rightarrow^{a.s.} \mu$ but $E(m_n) \neq \mu$. A Taylor series expansion of $H(\bar{Z})$ about $\bar{Z} = \theta$ yields

$$H(\bar{Z}) = H(\theta) + H'(\theta)(\bar{Z} - \theta) + (1/2)H''(\theta)(\bar{Z} - \theta)^2 + r_n,$$

where r_n is a remainder term that satisfies $E(r_n) = O(n^{-2})$. Therefore, the bias of m_n is

$$E(m_n - \mu) = (1/2)H''(\theta)E(\bar{Z} - \theta)^2 + O(n^{-2}).$$

The leading term of the bias is

$$B_n \equiv (1/2)H''(\theta)E(\bar{Z} - \theta)^2 = O(n^{-1}).$$

Therefore, the bias of m_n through $O(n^{-1})$ is B_n .

Now consider the bootstrap. Let $\{X_i^* : i = 1, \dots, n\}$ be a bootstrap sample that is obtained by sampling the X_i s randomly with replacement. Define $\bar{Z}^* = n^{-1} \sum_{i=1}^n Z(X_i^*)$. The bootstrap analog of μ is m_n , and the bootstrap estimator of m_n is $m_n^* = H(\bar{Z}^*)$. A Taylor series expansion of $H(\bar{Z}^*)$ about $\bar{Z}^* = \bar{Z}$ gives the bootstrap bias through $O(n^{-1})$ almost surely as

$$B_n^* = (1/2)H''(\bar{Z})E^*(m_n^* - m_n)^2,$$

where E^* denotes the expectation under bootstrap sampling (i.e., the expectation relative to the empirical distribution of the estimation data and conditional on the estimation data). Because the

distribution that the bootstrap samples is known, B_n^* can be computed with arbitrary accuracy by Monte Carlo simulation. Thus, B_n^* is a feasible estimator of the bias of m_n .

The differences between B_n and B_n^* are that \bar{Z} replaces θ and E^* replaces E in B_n^* . Moreover, we have $E(B_n^*) = B_n + O(n^{-2})$. Therefore, through $O(n^{-1})$, the bootstrap bias estimate B_n^* provides the same bias reduction that would be obtained if the infeasible population value B_n could be used. This is the source of the bootstrap's ability to reduce the bias of m_n . The resulting bias-corrected estimator of μ is $m_{n,Corr} = m_n - B_n^*$. It satisfies $E(m_{n,Corr} - \mu) = O(n^{-2})$. Thus, the bias of $m_{n,Corr}$ is $O(n^{-2})$, whereas the bias of m_n is $O(n^{-1})$.

3.2. Asymptotic Refinements to the Distributions of Test Statistics

This section explains the bootstrap's ability to provide approximations to the distributions of test statistics that are more accurate than the approximations of conventional asymptotic distribution theory, such as the asymptotic normal approximation. I assume that the estimators and test statistics of interest belong to the smooth function model. Specifically, I work with statistics of the form

$$\Delta_n = n^{1/2}[H(\bar{\mathbf{Z}}) - H(\boldsymbol{\theta})] \quad \text{and} \quad t_n = n^{1/2}[H(\bar{\mathbf{Z}}) - H(\boldsymbol{\theta})]/s_n,$$

where $\bar{\mathbf{Z}}$ and $\boldsymbol{\theta}$, respectively, are the sample average and population mean of the random variable or vector \mathbf{Z} ; s_n^2 is a consistent estimator of the variance of the asymptotic distribution of Δ_n ; and H has sufficiently many derivatives with respect to the components of its argument. As is explained in Section 3, many estimators and test statistics used in applied econometrics either are smooth functions of sample moments or can be approximated by such functions with a negligible error. We add a further assumption, which is called the Cramér condition.

Cramér Condition: Let $\boldsymbol{\tau}$ be a vector of constants with the same dimension as \mathbf{Z} . Let $i = \sqrt{-1}$. \mathbf{Z} satisfies the Cramér condition if

$$\lim_{\|\boldsymbol{\tau}\| \rightarrow \infty} \sup |E \exp(i\boldsymbol{\tau}' \mathbf{Z})| < 1.$$



The Cramér condition is satisfied if \mathbf{Z} is continuously distributed (that is, has a conventional probability density) but not if \mathbf{Z} is discrete. $\mathbf{Z} = \mathbf{Z}(\mathbf{X})$ is a function of the observed random variable \mathbf{X} . The Cramér condition is satisfied even if some components of \mathbf{X} are discretely distributed if $\mathbf{Z}(\mathbf{X})$ is continuously distributed.

Under the assumptions of the smooth function model and the Cramér condition, Δ_n , t_n , and their bootstrap analogs have higher-order asymptotic expansions called Edgeworth expansions. Let $G_n(\cdot, F_0)$ and $G_\infty(\cdot, F_0)$, respectively, denote the distribution and asymptotic distribution functions of either Δ_n or t_n when \mathbf{X} is sampled from a population whose cumulative distribution function is F_0 . The Edgeworth expansions of Δ_n and t_n are

$$G_n(\tau, F_0) = G_\infty(\tau, F_0) + n^{-1/2}g_1(\tau, \kappa_1) + n^{-1}g_2(\tau, \kappa_2) + n^{-3/2}g_3(\tau, \kappa_3) + O(n^{-2}) \quad 3.$$

uniformly over τ , where the κ s are cumulants through order 4 of the distribution of Δ_n or t_n ; G_∞ is the asymptotic distribution function of Δ_n or t_n ; g_1 and g_3 are even functions of their first arguments; g_2 is an odd function of its first argument; and g_1, g_2 , and g_3 are differentiable functions

of their second arguments. The Edgeworth expansions for the bootstrap analogs of Δ_n and t_n are

$$G_n(\tau, F_n) = G_\infty(\tau, F_n) + n^{-1/2}g_1(\tau, \kappa_{n1}) + n^{-1}g_2(\tau, \kappa_{n2}) + n^{-3/2}g_3(\tau, \kappa_{n3}) + O(n^{-2}) \quad 4.$$

almost surely uniformly over τ , where F_n is the empirical distribution of \mathbf{X} , and the κ_n s are cumulants of the empirical distribution of Δ_n or t_n .

To evaluate the accuracy of the bootstrap distribution function $G_n(\tau, F_n)$ as an approximation to the population distribution function $G_n(\tau, F_0)$, subtract Equation 3 from Equation 4 to obtain

$$G_n(\tau, F_0) - G_n(\tau, F_n) = G_\infty(\tau, F_0) - G_\infty(\tau, F_n) + n^{-1/2}[g_1(\tau, \kappa_1) - g_1(\tau, \kappa_{n1})] + n^{-1}[g_2(\tau, \kappa_2) - g_2(\tau, \kappa_{n2})] + O(n^{-3/2}) \quad 5.$$

almost surely uniformly over τ . The leading term on the right-hand side of Equation 5 is $G_\infty(\tau, F_0) - G_\infty(\tau, F_n)$, whose size is $O(n^{-1/2})$ almost surely because $F_n - F_0 = O(n^{-1/2})$ almost surely uniformly over the support of F_0 . Thus, the bootstrap makes an error of size $O(n^{-1/2})$ almost surely, which is the same as the size of the error made by conventional asymptotic approximations. In terms of the rate of convergence of the approximation error to zero, the bootstrap has the same accuracy as conventional approximations.

Now focus on t_n . This statistic is asymptotically pivotal, meaning that $G_\infty(\tau, F_0)$ does not depend on F_0 . Most test statistics are asymptotically pivotal, but most estimators are not. Usually, G_∞ for an asymptotically pivotal statistic is the standard normal distribution function. Asymptotic chi-square statistics are quadratic forms of asymptotically normal statistics and have properties that are straightforward modifications of the properties of asymptotic normal statistics.

When a statistic is asymptotically pivotal, the first term on the right-hand side of Equation 5 is zero. The second term is $O(n^{-1})$ almost surely because cumulants can be written as smooth functions of moments and, therefore, $\kappa_1 - \kappa_{n1} = O(n^{-1/2})$. Thus, the error of the bootstrap approximation to the distribution function of an asymptotically pivotal statistic converges to zero more rapidly than does the conventional asymptotic approximation. In this sense, the bootstrap is more accurate than conventional asymptotic approximations.

The bootstrap approximation to the symmetrical distribution function $P(|t_n| \leq \tau)$ is even more accurate. Because G_∞ does not depend on F_0 or F_n , g_1 and g_3 are even functions, and g_2 is an odd function,

$$P(|t_n| \leq \tau) = G_n(\tau, F_0) - G_n(-\tau, F_0) = \frac{2}{n}g_2(\tau, \kappa_2) + O(n^{-2}), \quad 6.$$

uniformly over τ . The bootstrap analog is

$$P^*(|t_n^*| \leq \tau) = G_n(\tau, F_n) - G_n(-\tau, F_n) = \frac{2}{n}g_2(\tau, \kappa_{n2}) + O(n^{-2}) \quad 7.$$

almost surely, where t_n^* is the bootstrap analog of t_n . Because $\kappa_{n2} - \kappa_2 = O(n^{-1/2})$ almost surely, it follows from Equations 6 and 7 that

$$P^*(|t_n^*| \leq \tau) - P(|t_n| \leq \tau) = O(n^{-3/2}) \quad 8.$$

almost surely. In contrast, the error made by conventional asymptotic approximations such as the normal approximation is $O(n^{-1})$.

In summary, the error in the bootstrap approximation to a one-sided distribution function of an asymptotically pivotal statistic is almost surely $O(n^{-1})$. The error in the bootstrap approximation to a symmetrical distribution function is almost surely $O(n^{-3/2})$. In contrast, the errors made by

conventional asymptotic approximations to one-sided and symmetrical distribution functions are $O(n^{-1/2})$ and $O(n^{-1})$, respectively.

Now suppose that the asymptotically pivotal statistic t_n is a statistic for testing a null hypothesis H_0 about the sampled population. Let $z_{1-\alpha}$ denote the $1 - \alpha$ quantile of the distribution of $|t_n|$. Thus, we have $P(|t_n| \leq z_{1-\alpha}) = 1 - \alpha$. A symmetrical test based on t_n rejects H_0 at the α level if $|t_n| > z_{1-\alpha}$. However, $z_{1-\alpha}$ is unknown in most settings. Let $z_{1-\alpha}^*$ be the $1 - \alpha$ quantile of the bootstrap distribution of $|t_n^*|$. Then we have $P^*(|t_n^*| \leq z_{1-\alpha}^*) = 1 - \alpha$. The quantity $z_{1-\alpha}^*$ can be estimated with any desired accuracy through Monte Carlo simulation. Therefore, $z_{1-\alpha}^*$ can be treated as known. A simulation algorithm is given at the end of this section. A feasible test rejects H_0 if $|t_n| > z_{1-\alpha}^*$. If H_0 is correct, then the probability that it is rejected is $P(|t_n| > z_{1-\alpha}^*)$. The difference between this probability and the nominal RP of $1 - \alpha$ [the error in the RP (ERP)] is

$$P(|t_n| > z_{1-\alpha}^*) - P^*(|t_n^*| > z_{1-\alpha}^*) = P(|t_n| \leq z_{1-\alpha}^*) - P^*(|t_n^*| \leq z_{1-\alpha}^*).$$

This difference is not given by Equation 8 because $z_{1-\alpha}^*$ is a function of the original data and, therefore, a random variable relative to the population distribution of the data. It is not a random variable relative to the bootstrap distribution. A lengthy argument based on an Edgeworth expansion of the distribution of $|t_n| - z_{1-\alpha}^*$ (see, for example, Hall 1992) shows that, if H_0 is correct, then we have

$$P(|t_n| > z_{1-\alpha}^*) - \alpha = O(n^{-2}).$$

Thus, the ERP based on the bootstrap critical value $z_{1-\alpha}^*$ is $O(n^{-2})$. In contrast, the ERP based on the asymptotic critical value [e.g., the $1 - \alpha/2$ quantile of the $N(0, 1)$ distribution] is $O(n^{-1})$. The ERP based on the bootstrap critical value converges to zero more rapidly than the ERP based on the asymptotic critical value. In samples of practical size, a test based on the bootstrap critical value usually has a smaller ERP than a test based on the asymptotic critical value. The reduction in the ERP can be dramatic (see, for example, Horowitz 1994, 1998a).

A one-sided upper-tailed test rejects H_0 if t_n is too much greater than zero. A one-sided lower-tailed test rejects H_0 if t_n is too much less than zero. The bootstrap versions of the two tests have similar properties, so we consider only the upper-tailed test. Define $\tilde{z}_{1-\alpha}^*$ as the $1 - \alpha$ quantile of the bootstrap distribution of t_n^* , that is, $P^*(t_n^* \leq \tilde{z}_{1-\alpha}^*) = 1 - \alpha$. A one-sided upper-tailed test based on t_n and the bootstrap critical value rejects H_0 if $t_n > \tilde{z}_{1-\alpha}^*$. Then, if H_0 is correct, we have

$$P(t_n > \tilde{z}_{1-\alpha}^*) - \alpha = O(n^{-1}).$$

The ERP in a one-tailed test with a bootstrap critical value is $O(n^{-1})$. In contrast, the error with a conventional asymptotic critical value is $O(n^{-1/2})$ (Hall 1992).

There are, however, circumstances in which the ERP with a bootstrap critical value is $O(n^{-3/2})$. Hall (1992) shows that this is true for a one-sided t test of a hypothesis about a slope (but not intercept) coefficient in a homoskedastic, linear, mean-regression model. Davidson & MacKinnon (1999) show that it is true whenever t_n is asymptotically independent of $g_2(\tilde{z}_{1-\alpha/2}, \kappa_{n2})$, where $\tilde{z}_{1-\alpha/2}$ is the asymptotic $1 - \alpha$ quantile of t_n . Many familiar test statistics satisfy this condition. Pretorius & Swanepoel (2018) describe a method that combines bootstrap estimates with results obtained from analytic Cornish-Fisher expansions, which are inversions of Edgeworth expansions. The method of Pretorius & Swanepoel (2018) achieves an ERP of $O(n^{-3/2})$ and, in some cases, $O(n^{-2})$. However, its need for analytic Cornish-Fisher expansions makes it intractable for most statistics of interest in econometrics.

The results above also apply to symmetrical and one-sided confidence intervals for a parameter. In the smooth function model, define $\mu = H(\theta)$. Then we have

$$t_n = n^{1/2}[H(\bar{\mathbf{Z}}) - \mu]/s_n.$$

A symmetrical $1 - \alpha$ confidence interval for μ based on the bootstrap critical value $z_{1-\alpha}^*$ is

$$H(\bar{\mathbf{Z}}) - n^{-1/2}s_n z_{1-\alpha}^* \leq \mu \leq H(\bar{\mathbf{Z}}) + n^{-1/2}s_n z_{1-\alpha}^*.$$

A one-sided upper confidence interval is

$$\mu \leq H(\bar{\mathbf{Z}}) + n^{-1/2}s_n z_{1-\alpha}^*.$$

A one-sided lower confidence interval can be defined similarly. The difference between the nominal and true coverage probabilities [error in the coverage probability (ECP)] of the symmetrical confidence interval is $O(n^{-2})$. It is $O(n^{-1})$ for the one-sided interval with the qualifications noted in the previous paragraph. The ECPs based on conventional asymptotic critical values are $O(n^{-1})$ and $O(n^{-1/2})$, respectively, for symmetrical and one-sided confidence intervals.

Tests based on statistics that are asymptotically chi-square distributed behave like symmetrical, two-tailed tests. Therefore, their ERP under a correct H_0 is $O(n^{-1})$ with conventional asymptotic critical values and $O(n^{-2})$ with bootstrap critical values.

The results presented in this section show that the bootstrap reduces the ERPs of hypothesis tests and the ECPs of confidence intervals based on smooth, asymptotically pivotal statistics. These include many asymptotically normal and asymptotically chi-square test statistics that are used for testing hypotheses about the parameters of econometric models. Models that satisfy the required smoothness conditions include linear and nonlinear mean-regression models, error-components mean-regression models for panel data, logit and probit models that have at least one continuously distributed explanatory variable, and tobit models. The smoothness conditions are also satisfied by parametric sample-selection models in which the selection equation is a logit or probit model with at least one continuously distributed explanatory variable. Asymptotically pivotal statistics based on quantile regression models do not satisfy the smoothness conditions. Bootstrap methods for such statistics are discussed in Section 4. Statistics based on nonparametric estimators also have different properties. These are discussed in Section 5.

The ability of the bootstrap to provide asymptotic refinements for smooth, asymptotically pivotal statistics provides a powerful argument for using these statistics in applications of the bootstrap whenever possible. The bootstrap may be applied to statistics that are not asymptotically pivotal, but it does not provide higher-order approximations to their distributions. Estimators of the structural parameters of econometric models (e.g., slope and intercept parameters, including regression coefficients; standard errors; covariance matrix elements; and autoregressive coefficients) usually are not asymptotically pivotal. The asymptotic distributions of centered structural parameter estimators are often normal, with means of zero, but have variances that depend on the unknown population distribution of the data. The errors of bootstrap estimates of the distributions of statistics that are not asymptotically pivotal converge to zero at the same rate as the errors made by first-order asymptotic approximations.

Bootstrap hypothesis tests and confidence intervals based on the asymptotically pivotal statistic t_n have smaller ERPs and ECPs than tests and confidence intervals based on conventional asymptotic critical values because the bootstrap estimate of the distribution of t_n is a form of low-order Edgeworth expansion that accounts for skewness and kurtosis of the distribution.

Normal and chi-square distributional approximations do not account for skewness or kurtosis. The bootstrap based on an asymptotically pivotal statistic should be used to obtain hypothesis tests or confidence intervals, not standard errors. Standard errors are not asymptotically pivotal, and the bootstrap does not provide asymptotic refinements for them. More importantly, standard errors are related in a simple way to quantiles of the normal distribution and, therefore, to confidence intervals based on the asymptotic normal distribution of t_n . However, the bootstrap estimate of the distribution of t_n is non-normal, and there is no simple relation between standard errors and quantiles of non-normal distributions. The use of the bootstrap to estimate standard errors when an asymptotically pivotal statistic is available fails to take advantage of the bootstrap's ability to provide asymptotic refinements and results that are numerically more accurate than those provided by asymptotic normal or chi-square approximations.

The following is an algorithm for Monte Carlo computation of symmetrical bootstrap critical values and confidence intervals. The algorithm for one-sided critical values and confidence intervals is similar and not given.

MONTE CARLO PROCEDURE FOR COMPUTING THE BOOTSTRAP CRITICAL VALUE FOR TESTING A HYPOTHESIS ABOUT OR FORMING A SYMMETRICAL CONFIDENCE INTERVAL FOR A PARAMETER μ

Step 1: Use the estimation data to compute a consistent, asymptotically normal estimate $\hat{\mu}_n$ of μ and its standard error s_n .

Step 2: Generate a bootstrap sample of size n by sampling the empirical distribution of the data randomly with replacement, sampling a parametric estimate of the distribution of the data, or using the residual bootstrap. Estimate μ and its standard error from the bootstrap sample. Call the results μ_n^* and s_n^* . The bootstrap version of the asymptotically pivotal statistic t_n is $t_n^* = n^{1/2}(\mu_n^* - \hat{\mu}_n)/s_n^*$.

Step 3: Use the results of many repetitions of Step 2 to compute the empirical distribution of $|t_n^*|$. Set $z_{1-\alpha}^*$ equal to the $1 - \alpha$ quantile of this distribution.

Step 4: Reject the hypothesis $H_0 : \mu = \mu_0$ at the $1 - \alpha$ level if $|n^{1/2}(\hat{\mu}_n - \mu_0)/s_n| > z_{1-\alpha}^*$.

Step 5: A symmetrical $1 - \alpha$ confidence interval for μ is

$$\hat{\mu}_n - n^{-1/2}s_n z_{1-\alpha}^* \leq \mu \leq \hat{\mu}_n + n^{-1/2}s_n z_{1-\alpha}^*.$$

4. NONSMOOTH STATISTICS

Some important estimators and test statistics do not satisfy the assumptions of the smooth function model. Quantile estimators and Manski's (1975, 1985) maximum score estimator for a binary response model are two examples. The objective function for quantile estimation has cusps, so its first derivative is discontinuous. The objective function for maximum score estimation is a step function, so its first derivative is either zero or infinity. This section outlines the properties of bootstrap methods for quantile and maximum score estimators.

4.1. Quantile Estimators

In this section, I restrict attention to the LAD estimator of a linear median-regression model. Similar ideas and methods apply to nonlinear quantile models and estimators for other quantiles.

The model is

$$Y = \beta' \mathbf{X} + U; \quad P(U \leq 0 | \mathbf{X}) = 0.5,$$

where β is a vector of constant parameters, \mathbf{X} is a vector of explanatory variables, and U is an unobserved random variable. Let $\{Y_i, \mathbf{X}_i : i = 1, \dots, n\}$ be an independent random sample from the distribution of (Y, \mathbf{X}) . The LAD estimator of β is

$$\begin{aligned} \hat{\mathbf{b}} &= \arg \min_{\mathbf{b}} \sum_{i=1}^n |Y_i - \mathbf{b}' \mathbf{X}_i| \\ &= \arg \min_{\mathbf{b}} \sum_{i=1}^n (Y_i - \mathbf{b}' \mathbf{X}_i) [2I(Y_i - \mathbf{b}' \mathbf{X}_i \geq 0) - 1]. \end{aligned} \tag{9}$$

Bassett & Koenker (1978) and Koenker & Bassett (1978) give conditions under which $n^{1/2}(\hat{\mathbf{b}} - \beta)$ is asymptotically normal with mean zero. The bootstrap estimates the distribution of $n^{1/2}(\hat{\mathbf{b}} - \beta)$ consistently (De Angelis et al. 1993, Hahn 1995), but the nonsmoothness of the LAD objective function causes the difference between bootstrap approximation to the distribution function of $n^{1/2}(\hat{\mathbf{b}} - \beta)$ and the true distribution function to converge to zero slowly. Depending on how the bootstrap is implemented and certain other conditions, the rate is either $O(n^{-1/4})$ or $O(n^{-2/5})$ (De Angelis et al. 1993). In addition, the Edgeworth expansion of the distribution of $n^{1/2}(\hat{\mathbf{b}} - \beta)$ is very complicated. It is not known whether bootstrap sampling applied to Equation 9 provides asymptotic refinements for hypothesis tests and confidence intervals.

Horowitz (1998b) proposes smoothing the cusp in the LAD objective function to make it differentiable. To state the smoothed objective function, let $H(\cdot)$ be a function with sufficiently many continuous derivatives such that $H(v) = 1$ if $v \geq 1$ and $H(v) = 0$ if $v \leq -1$. For example, $H(v)$ might be the cumulative distribution function of a random variable whose support is $[-1, 1]$. Let $\{b_n\}$ be a decreasing sequence of positive constants that converges to zero as $n \rightarrow \infty$. The smoothed LAD (SLAD) estimator is

$$\tilde{\mathbf{b}} = \arg \min_{\mathbf{b}} \sum_{i=1}^n (Y_i - \mathbf{b}' \mathbf{X}_i) \left[2H\left(\frac{Y_i - \mathbf{b}' \mathbf{X}_i}{b_n}\right) - 1 \right]. \tag{10}$$

The objective function in Equation 10 is identical to the LAD objective function in Equation 9 if $|Y_i - \mathbf{b}' \mathbf{X}_i| \geq b_n$, but it smooths the discontinuities in the indicator function at $Y_i - \mathbf{b}' \mathbf{X}_i = 0$. The interval over which smoothing occurs decreases as $b_n \rightarrow 0$. This is necessary to make $\tilde{\mathbf{b}}$ a consistent estimator of β .

Because the objective function in Equation 10 is differentiable, standard Taylor series methods can be used to obtain the asymptotic distribution of $n^{1/2}(\tilde{\mathbf{b}} - \beta)$. Horowitz (1998b) gives conditions under which $n^{1/2}(\tilde{\mathbf{b}} - \beta) \rightarrow^d N(0, V)$ holds true, where V is a covariance matrix. Let \hat{V} be a consistent estimator of V ; \hat{V}_{jj} denote the (j, j) component of \hat{V} ; and \tilde{b}_j and β_j denote the j th components of $\tilde{\mathbf{b}}$ and β , respectively. Then a t statistic for testing the hypothesis $H_0 : \beta_j = \beta_{0j}$ is

$$t_n = \frac{n^{1/2}(\tilde{b}_j - \beta_{0j})}{\hat{V}_{jj}^{1/2}}.$$

A bootstrap version of t_n can be obtained by replacing the original sample $\{Y_i, \mathbf{X}_i : i = 1, \dots, n\}$ with the bootstrap sample $\{Y_i^*, \mathbf{X}_i^* : i = 1, \dots, n\}$ in Equation 10. Let \tilde{b}_j^* denote the estimate of β_j obtained from the bootstrap sample and V_{jj}^* denote the bootstrap analog of \hat{V}_{jj} . The bootstrap

analog of t_n is

$$t_n^* = \frac{n^{1/2}(\tilde{b}_j^* - \tilde{b}_j)}{(V_{jj}^*)^{1/2}}.$$

Horowitz (1998b) shows that t_n and t_n^* have Edgeworth expansions that are identical almost surely through $O[(nb_n)^{-1}]$. Therefore, arguments similar to those in Section 3.2 show that the bootstrap provides asymptotic refinements for hypothesis tests and confidence intervals based on the SLAD estimator. For example, consider a symmetrical t test of H_0 . Let $z_{1-\alpha}^*$ be the $1 - \alpha$ quantile of the bootstrap distribution of $|t_n^*|$. Then we have

$$P(|t_n| > z_{1-\alpha}^*) = \alpha + o[(nb_n)^{-1}]$$

if H_0 is true. In contrast, the asymptotic normal approximation makes an error of $O[(nb_n)^{-1}]$. Use of the bootstrap critical value reduces the approximation error because the bootstrap captures a higher-order term in the Edgeworth expansion of $|t_n|$, whereas the asymptotic normal approximation drops this term. The bootstrap also provides asymptotic refinements for one-sided hypothesis tests and confidence intervals and for asymptotic chi-square tests of hypotheses about several components of β . In addition, the bootstrap provides asymptotic refinements for a smoothed version of Powell's (1984, 1986) censored LAD estimator.

4.2. The Maximum Score Estimator

Manski's (1975, 1985) maximum score estimator applies to the binary response model $Y = I(\beta'X + U \geq 0)$, where Y is an observed random variable, X is an observed random vector, U is an unobserved random variable satisfying $P(U \leq 0|X) = 0.5$, and β is an unknown vector of constant parameters to be estimated. The vector β is identified only up to scale, so a scale normalization is needed. In this section, scale normalization will consist of setting $|\beta_1| = 1$, where β_1 is the first component of β . Assume that the first component of X is continuously distributed. Define the set $B = \{\mathbf{b} : |b_1| = 1\}$. Let $\{Y_i, \mathbf{X}_i : i = 1, \dots, n\}$ be an independent random sample of (Y, X) . The maximum score estimator of β can be written as

$$\hat{\mathbf{b}}_{MS} = \arg \max_{\mathbf{b} \in B} \sum_{i=1}^n (2Y_i - 1)I(\mathbf{b}'\mathbf{X}_i \geq 0). \quad 11.$$

Manski (1975, 1985) gives conditions under which $\hat{\mathbf{b}}_{MS} \rightarrow \beta$ almost surely. Cavanagh (1987) and Kim & Pollard (1990) show that $\hat{\mathbf{b}}_{MS}$ converges to β at the rate $n^{-1/3}$, and that $n^{1/3}(\hat{\mathbf{b}}_{MS} - \beta)$ has a complicated, non-normal asymptotic distribution. Abrevaya & Huang (2005) show that the nonparametric bootstrap does not provide a consistent estimate of this distribution.

The maximum score estimator converges slowly and has a complicated limiting distribution because it is obtained by maximizing a step function. Several authors have proposed bootstrap sampling procedures and/or modifications of the maximum score objective function that make the bootstrap consistent (Cattaneo et al. 2018; Hong & Li 2015; Horowitz 1992, 2002; Patra et al. 2018). The method of Horowitz (1992, 2002), which provides asymptotic refinements, is described in this section. Horowitz (1992) proposes replacing the indicator function on the right-hand side of Equation 11 with a differentiable function. Let H and b_n be defined as in Equation 10. The

resulting smoothed maximum score (SMS) estimator is

$$\hat{\mathbf{b}}_{SMS} = \arg \max_{\mathbf{b} \in B} \sum_{i=1}^n (2Y_i - 1)H\left(\frac{\mathbf{b} \cdot \mathbf{X}_i}{b_n}\right).$$

Horowitz (1992) gives conditions under which $\hat{\mathbf{b}}_{SMS}$ converges to β at a rate that is at least as fast as $n^{-2/5}$, and $(nb_n)^{1/2}(\hat{\mathbf{b}}_{SMS} - \beta)$ is asymptotically normally distributed. Asymptotic normality makes it possible to form asymptotically pivotal t statistics for testing hypotheses about β and obtain asymptotic refinements with the bootstrap (Horowitz 2002). Let t_n be a t statistic, and let t_n^* be the bootstrap analog that is obtained by SMS estimation based on the bootstrap sample $\{Y_i^*, \mathbf{X}_i^* : i = 1, \dots, n\}$. Now consider a symmetrical t test of the hypothesis $H_0 : \beta_j = \beta_{0j}$. Let $z_{1-\alpha}^*$ be the $1 - \alpha$ quantile of the bootstrap distribution of $|t_n^*|$. Then we have

$$P(|t_n| > z_{1-\alpha}^*) = \alpha + o[(nb_n)^{-1}]$$

if H_0 is true. In contrast, the asymptotic normal approximation makes an error of $O[(nb_n)^{-1}]$. As in SLAD estimation, the bootstrap captures a higher-order term in the Edgeworth expansion of $|t_n|$, whereas the asymptotic normal approximation drops this term. Also as in SLAD estimation, the bootstrap also provides asymptotic refinements for one-sided hypothesis tests and confidence intervals and for asymptotic chi-square tests of hypotheses about several components of β .

5. NONPARAMETRIC ESTIMATION

This section is concerned with inference about the unknown function g in the nonparametric mean-regression model

$$Y = g(\mathbf{X}) + \varepsilon; \quad E(\varepsilon|\mathbf{X}) = 0 \tag{12}$$

and the nonparametric quantile regression model

$$Y = g(\mathbf{X}) + \varepsilon; \quad P(\varepsilon \leq 0|\mathbf{X}) = \tau; \quad 0 < \tau < 1. \tag{13}$$

Nonparametric estimators of g and their properties are described by Fan & Gijbels (1996), Fan et al. (1994), Härdle (1990), and Yu & Jones (1998), among many others.

Any method for testing a hypothesis about g or constructing a confidence interval or band for g based on a nonparametric estimate must deal with the problem of asymptotic bias. Let \hat{g} denote a nonparametric estimate of g . The expected value of \hat{g} does not equal g , the asymptotic normal distribution of the scaled estimate is not centered at g , and the true coverage probability of an asymptotic confidence interval for g (or acceptance probability of a hypothesis test) that is constructed from the normal distribution in the usual way is less than the nominal probability. This problem is usually overcome by undersmoothing or explicit bias reduction. For example, suppose that \hat{g} is a kernel estimator. Then undersmoothing consists of making the bias asymptotically negligible by using a bandwidth whose rate of convergence is faster than the asymptotically optimal rate. In explicit bias reduction, an estimate of the asymptotic bias is used to construct an asymptotically unbiased estimate of g . Most explicit bias reduction methods involve some form of oversmoothing, that is, using a bandwidth whose rate of convergence is slower than the asymptotically optimal rate. Hall (1992) and Horowitz (2001) describe the ability of the bootstrap to provide asymptotic refinements for hypothesis tests and confidence intervals based on nonparametric estimators with undersmoothing or explicit bias correction.

Methods based on undersmoothing or oversmoothing require a bandwidth whose rate of convergence is faster or slower than the asymptotically optimal rate. However, there are no effective empirical ways to choose these bandwidths. Hall & Horowitz (2013) and Horowitz & Krishnamurthy (2019) describe bootstrap methods for overcoming this problem in the models in Equations 12 and 13, respectively. These methods use bandwidths chosen by standard empirical methods such as cross validation or a plug-in rule. Instead of under- or oversmoothing, the methods use the bootstrap to estimate the bias of \hat{g} . The bootstrap estimate of the bias can be obtained using the procedure in Section 3.1. However, in this nonparametric setting, the bootstrap bias estimate has stochastic noise that is comparable in size to the bias itself. Hall & Horowitz and Horowitz & Krishnamurthy give conditions under which combining a suitable quantile of the distribution of the bootstrap bias estimate with \hat{g} provides pointwise confidence intervals and uniform confidence bands with asymptotic coverage probabilities that equal or exceed the nominal probabilities. The methods of Hall & Horowitz and Horowitz & Krishnamurthy use the bootstrap to select critical values that are larger than those obtained from the usual asymptotic normal approximations and, therefore, yield wider confidence intervals and bands that achieve the desired coverage probabilities asymptotically without the need for bias correction through undersmoothing or explicit bias estimation. The same methods can also be used to form hypothesis tests. In contrast to the use of the bootstrap in Sections 3 and 4, the objective in this case is not to obtain asymptotic refinements. Rather, the bootstrap is used to overcome a problem in constructing first-order asymptotic confidence bands and hypothesis tests.

Hall & Horowitz (2013) and Horowitz & Krishnamurthy (2019) provide the details of the methods, which are lengthy, although not difficult to implement. As an illustration, the method for obtaining a pointwise confidence band for the mean-regression model in Equation 12 with a scalar covariate X with compact support \mathcal{R} and a homoskedastic ε consists of the following steps.

Step 1: Form an estimate \hat{g} of g by using a kernel-based method such as a local polynomial estimator. Use a standard empirical method such as cross-validation or plug-in to choose the bandwidth. Form an estimate $\hat{\sigma}^2$ of $\sigma^2 = E(U^2)$ by using a standard method such as that of Rice (1984). Let $\mathcal{S}(x)\hat{\sigma}^2$ denote an estimate of the variance of $\hat{g}(x)$. In local polynomial estimation, $\mathcal{S}(x)$ is a known function of the X_i s in the data. Hall & Horowitz (2013) provide details.

Step 2: Compute residuals $\tilde{\varepsilon}_i = Y_i - \hat{g}(X_i)$ ($i = 1, \dots, n$). Set $\bar{\varepsilon} = n^{-1} \sum_{i=1}^n \tilde{\varepsilon}_i$ and compute centered residuals $\hat{\varepsilon}_i$, defined by $\hat{\varepsilon}_i = \tilde{\varepsilon}_i - \bar{\varepsilon}$.

Step 3: Form the residual bootstrap sample $\{Y_i^*, X_i : i = 1, \dots, n\}$, where $Y_i^* = \hat{g}(X_i) + \varepsilon_i^*$ and the ε_i^* s are obtained by sampling the $\hat{\varepsilon}_i$ s randomly with replacement. The X_i s are the same as in the original data and are not resampled.

Step 4: Use the estimation methods and bandwidth of Step 1 with the bootstrap sample from Step 3 to obtain bootstrap estimates of g and σ^2 . Denote these by g^* and σ^{*2} . This step uses the bandwidth obtained in Step 1. It does not use the bootstrap sample to obtain a new bandwidth. Form the bootstrap pointwise $1 - \alpha$ confidence band

$$B^*(\alpha) = \{(y, x) : g^*(x) - \mathcal{S}(x)\sigma^* z_{1-\alpha/2} \leq \hat{g}(x) \leq g^*(x) + \mathcal{S}(x)\sigma^* z_{1-\alpha/2}\},$$

where $z_{1-\alpha/2}$ is the $1 - \alpha/2$ quantile of the $N(0, 1)$ distribution.

Step 5: Repeat Step 4 B times, where B is a large positive integer (e.g., $B = 1000$). Let $B_b^*(\alpha)$ denote the confidence band obtained on the b th out of B repetition. For $x \in \text{supp}(X)$,

compute the bootstrap coverage probability

$$\pi^*(x, \alpha) = B^{-1} \sum_{b=1}^B I[(x, \hat{g}(x) \in \mathcal{B}_b^*(\alpha)].$$

Step 6: Let $1 - \alpha_0$ be the desired coverage probability of a pointwise confidence band. Define $\beta^*(x, \alpha_0)$ to be the solution in α of $\pi^*(x, \alpha) = 1 - \alpha_0$. For $\xi \in (0, .5]$, let $\alpha_\xi^*(\alpha_0)$ denote the ξ -level quantile of points in the set $\{\beta^*(x, \alpha_0) : x \in \mathcal{R}\}$. In practice, it suffices to find the quantile over a closely spaced grid of discrete points in \mathcal{R} . Asymptotically, the confidence band

$$\mathcal{B}[\alpha_\xi(\alpha_0)] = \{(y, x) : \hat{g}(x) - \mathcal{S}(x)\hat{\sigma}z_{1-\alpha_\xi(\alpha_0)/2} \leq g(x) \leq \hat{g}(x) + \mathcal{S}(x)\hat{\sigma}z_{1-\alpha_\xi(\alpha_0)/2}\}$$

covers all but a fraction $1 - \xi$ of points in $x \in \mathcal{R}$. The exceptional points are in regions where $g(\mathbf{x})$ has sharp peaks or troughs that cause the bias of \hat{g} to be unusually large. These regions are typically visible in a plot of \hat{g} . Hall & Horowitz (2013) provide a theoretical analysis of this phenomenon and Monte Carlo illustrations of the numerical performance of  method for the mean regression model in Equation 12. Horowitz & Krishnamurthy (2014) illustrate the performance of the method for the quantile regression model in Equation 13.

6. THE MULTIPLIER AND WILD BOOTSTRAPS

The multiplier and wild bootstraps are methods for generating bootstrap samples that do not consist of resampling the original data or residuals, as in Sections 2–5. Rather, the multiplier and wild bootstraps combine the data with random variables drawn from a known distribution to form a bootstrap sample. The multiplier and wild bootstraps provide ways to deal with issues such as heteroskedasticity of unknown form in fixed-design regression models or random-design models in which one conditions on the covariates. They also provide methods for obtaining nonasymptotic bounds on the ERPs of certain hypothesis tests.

6.1. The Wild Bootstrap

The wild bootstrap enables one to carry out accurate inference in regression models with heteroskedasticity of unknown form. This section focuses on the linear model

$$Y_i = \beta' \mathbf{X}_i + \varepsilon_i; \quad i = 1, \dots, n, \tag{14}$$

where the \mathbf{X}_i s are fixed in repeated samples or random but inference is conditional on the observed \mathbf{X}_i s; $E(\varepsilon_i) = 0$ for all $i = 1, \dots, n$; $E(\varepsilon_i^2) = \sigma^2(\mathbf{X}_i)$; and $\sigma^2(\cdot)$ is an unknown function. Asymptotic inference about β can be carried out by using the heteroskedasticity-consistent covariance matrix estimator (HCCME) of Eicker (1963, 1967) and White (1980), but this estimator can be severely biased downward in finite samples, with the consequence that the true finite-sample probability of rejecting a correct null hypothesis can be significantly larger than the nominal probability (Chesher & Jewitt 1987). If the \mathbf{X}_i s are treated as random, then the bootstrap can be implemented for the model in Equation 14 by sampling the data $\{Y_i, \mathbf{X}_i : i = 1, \dots, n\}$ randomly with replacement. However, this method can be inaccurate because it does not impose the moment condition $E(\varepsilon_i | \mathbf{X}_i) = 0$.

The wild bootstrap imposes the moment restriction. It was introduced by Liu (1988) following a suggestion of Wu (1986). Mammen (1993) establishes the ability of the wild bootstrap

to provide asymptotic refinements for the model in Equation 14. Cao-Abad (1991), Härdle & Mammen (1993), and Härdle & Marron (1991) use the wild bootstrap in nonparametric regression. To describe the method for a linear model, let $\hat{\mathbf{b}}$ be the OLS estimate of the vector of slope coefficients, $\boldsymbol{\beta}$, based on data $\{Y_i, \mathbf{X}_i : i = 1, \dots, n\}$ in the model

$$Y_i = \boldsymbol{\beta}'\mathbf{X}_i + \varepsilon_i; E(\varepsilon_i) = 0; E(\varepsilon_i^2) = \sigma_i^2. \tag{15}$$

The variances σ_i^2 are unknown and not necessarily equal. The wild bootstrap generates bootstrap samples $\{Y_i^*, \mathbf{X}_i : i = 1, \dots, n\}$ from

$$Y_i^* = \hat{\mathbf{b}}'\mathbf{X}_i + \varepsilon_i^*. \tag{16}$$

As in the residual bootstrap method in Section 2.2, the wild bootstrap uses the \mathbf{X}_i s from the original data. The \mathbf{X}_i s are not resampled. The ε_i^* s are generated by either of the following two methods:

1. Let $\hat{\varepsilon}_i = Y_i - \hat{\mathbf{b}}'\mathbf{X}_i$ ($i = 1, \dots, n$) be the OLS residuals from the model in Equation 15. For each $i = 1, \dots, n$, let F_i correspond to the unique two-point distribution that satisfies $E(\eta_i) = 0$; $E(\eta_i^2|F_i) = \hat{\varepsilon}_i^2$; and $E(\eta_i^3) = \hat{\varepsilon}_i^3$; where η_i is a random variable with the cumulative distribution function F_i . Then we have $\eta_i = (1 - \sqrt{5})\hat{\varepsilon}_i$ with probability $(1 + \sqrt{5})/(2\sqrt{5})$ and $\eta_i = (1 + \sqrt{5})\hat{\varepsilon}_i/2$ with probability $1 - (1 + \sqrt{5})/(2\sqrt{5})$. Set $\varepsilon_i^* = \eta_i$ for each i . Mammen (1993) provides a detailed discussion of the properties of this method.
2. The second method is an example of the multiplier bootstrap, meaning that the ε_i^* s are multiples of transformations of the residuals $\hat{\varepsilon}_i$ and independent random variables. Specifically, let U_i ($i = 1, \dots, n$) be random variables that are independent of each other and the OLS residuals such that $E(U_i) = 0$ and $E(U_i^2) = 1$. One possibility is $U_i \sim N(0, 1)$. Let $f(\hat{\varepsilon}_i)$ be a transformation of the OLS residuals, possibly $f(\hat{\varepsilon}_i) = \hat{\varepsilon}_i$. Set $\varepsilon_i^* = U_i f(\hat{\varepsilon}_i)$. Davidson & Flachaire (2008) discuss properties of this method.

Regardless of the method used to generate ε_i^* , implementation of the wild bootstrap proceeds as follows.

Step 1: Generate a bootstrap sample $\{Y_i^*, \mathbf{X}_i : i = 1, \dots, n\}$ from Equation 16. Estimate $\boldsymbol{\beta}$ by OLS using this sample. Compute the resulting bootstrap t statistic, t_n^* , by using the HCCME or a variant such as that of MacKinnon & White (1985).

Step 2: Obtain the empirical distribution of t_n^* by repeating Step 1 many times. Obtain the critical value of t_n^* . Use this critical value with the t statistic from the original data to test hypotheses about and form confidence intervals for the components of $\boldsymbol{\beta}$.

Horowitz (2001) and Davidson & Flachaire (2008) provide examples of the numerical performance of the wild bootstrap.

6.2. Nonasymptotic Inference in Maximum Likelihood Estimation

This section describes the use of the multiplier bootstrap to carry out a likelihood ratio test of a finite-dimensional parameter $\boldsymbol{\theta} \in \mathbb{R}^d$ for some finite integer $d \geq 1$. The hypothesis is $H_0 : \boldsymbol{\theta} = \boldsymbol{\theta}_0$ for some known $\boldsymbol{\theta}_0$. Let the data $\{\mathbf{X}_i : i = 1, \dots, n\}$ be an independently but not necessarily identically distributed random sample from some population. As in Section 2, the components of \mathbf{X}_i include any dependent variables, as well as explanatory variables. Let $f_i(\cdot, \boldsymbol{\theta})$ denote the probability density function of the i th observation. Technically, f_i is a Radon-Nikodym density, so the components of \mathbf{X}_i may be discretely distributed or have probability distributions that are continuous in some regions and discrete in others, but this section uses ordinary probability density notation

to minimize the complexity of the discussion. The log-likelihood function is

$$\log L(\boldsymbol{\theta}) = \sum_{i=1}^n \log f_i(\mathbf{X}_i, \boldsymbol{\theta}).$$

The maximum likelihood estimate of $\boldsymbol{\theta}$ is

$$\hat{\boldsymbol{\theta}} = \arg \max_{\boldsymbol{\theta} \in \Theta} \log L(\boldsymbol{\theta}),$$

where $\Theta \subset \mathbb{R}^d$ is the parameter set. Under $H_0 : \boldsymbol{\theta} = \boldsymbol{\theta}_0$, we have

$$\boldsymbol{\theta}_0 = \arg \max_{\boldsymbol{\theta} \in \Theta} E[\log L(\boldsymbol{\theta})].$$

The likelihood ratio statistic for testing H_0 is

$$LR = 2[\log L(\hat{\boldsymbol{\theta}}) - \log L(\boldsymbol{\theta}_0)].$$

To construct the multiplier bootstrap version of LR , let $\{U_i : i = 1, \dots, n\}$ be scalar random variables that are independent of each other and the \mathbf{X}_i s such that $E(U_i) = E(U_i^2) = 1$ and $E \exp(U_i) < \infty$ for all $i = 1, \dots, n$. The multiplier bootstrap version of the log-likelihood function is

$$\log L^*(\boldsymbol{\theta}) = \sum_{i=1}^n U_i \log f_i(\mathbf{X}_i, \boldsymbol{\theta}).$$

The multiplier bootstrap parameter estimate is

$$\boldsymbol{\theta}^* = \arg \max_{\boldsymbol{\theta} \in \Theta} \log L^*(\boldsymbol{\theta}).$$

Let E^* denote the expectation with respect to the distribution of the U_i s with the \mathbf{X}_i s held constant. Then, under H_0 , we have

$$\hat{\boldsymbol{\theta}} = \arg \max_{\boldsymbol{\theta} \in \Theta} E^*[\log L^*(\boldsymbol{\theta})],$$

and the multiplier bootstrap version of the likelihood ratio statistic is

$$LR^* = 2[\log L(\boldsymbol{\theta}^*) - \log L(\hat{\boldsymbol{\theta}})].$$

Let P_n^* denote the probability distribution induced by sampling the U_i s while holding the \mathbf{X}_i s constant. Let $z_{1-\alpha}^*$ denote the $1 - \alpha$ quantile of the distribution of LR^* under P_n^* probability. Then we have

$$z_{1-\alpha}^* = \inf_{z \geq 0} \{z : P_n^*(LR^* > z) \leq \alpha\}.$$

The quantity $z_{1-\alpha}^*$ can be estimated with any desired accuracy by repeated sampling of the U_i s and computation of LR^* . Spokoiny & Zhilova (2015) give conditions under which we have

$$|P(LR > z_{1-\alpha}^*) - \alpha| \leq C \left(\frac{d + \nu}{n} \right)^{1/8} \tag{17}$$

for any ν such that $\alpha \leq 1 - 8e^{-\nu}$, where C is a constant. The right-hand side of Equation 17 is small only if n is very large. Nonetheless, Equation 17 is important because it is a finite-sample result showing that a version of the bootstrap can be used to carry out nonasymptotic inference with a familiar and frequently used statistic. In addition, the right-hand side of Equation 17 is a worst-case bound that accommodates the most extreme distributions of $\{\mathbf{X}_i : i = 1, \dots, n\}$. Spokoiny & Zhilova (2015) provide Monte Carlo evidence indicating that the bound is much tighter than Equation 17 suggests in less extreme cases.

6.3. Inference with the Multiplier Bootstrap in High-Dimensional Settings

Let $\{\mathbf{X}_i : i = 1, \dots, n\}$ be independent $p \times 1$ random vectors with means of zero and finite covariance matrices. Define the $p \times 1$ random vector

$$\mathbf{X} = n^{-1/2} \sum_{i=1}^n \mathbf{X}_i,$$

and let $X^{(j)}$ denote the j th component of \mathbf{X} . Chernozhukov et al. (2013) consider the problem of estimating the probability distribution of

$$Z = \max_{1 \leq j \leq p} X^{(j)}$$

when p may be greater than n . This problem arises in testing multiple hypotheses and in certain high-dimensional estimation methods, among other applications.

If p is fixed, and certain other conditions are satisfied, then a standard central limit theorem shows that \mathbf{X} is asymptotically multivariate normal. The asymptotic distribution of Z can be obtained from this result. However, standard central limit theorems do not apply if $p > n$.

Chernozhukov et al. (2013) define a multiplier bootstrap version of Z . Let $\{e_i : i = 1, \dots, n\}$ denote a sequence of $N(0, 1)$ random variables that are independent of each other and the \mathbf{X}_i s. Let $X^{*(j)}$ be the j th component of the $p \times 1$ random vector

$$\mathbf{X}^* = n^{-1/2} \sum_{i=1}^n \mathbf{X}_i e_i.$$

Define the following multiplier bootstrap version of Z :

$$Z^* = \max_{1 \leq j \leq p} n^{-1/2} X^{*(j)}.$$

Let $z_{1-\alpha}^*$ denote the $1 - \alpha$ quantile of the distribution of Z^* under sampling of the e_i s while holding the \mathbf{X}_i s fixed at the values in the data. Chernozhukov et al. (2013) give conditions under which $|P(Z \leq z_{1-\alpha}^*) - (1 - \alpha)| \rightarrow 0$ as $n \rightarrow \infty$ even if $p > n$. Chernozhukov et al. (2013) also show that

their version of the multiplier bootstrap provides consistent estimates of the distributions of statistics related to Z . Finally, Chernozhukov et al. (2013) give examples of the application of their results to certain problems in high-dimensional estimation and testing.

7. THE LASSO

The LASSO is a method for parameter estimation in settings in which the number of unknown parameters is comparable to or may exceed the sample size. The LASSO was introduced by Tibshirani (1996). It has been the object of much subsequent research and has generated a vast literature. Bühlmann & van de Geer (2011) synthesize the results of this research. The LASSO is applicable to a wide variety of models. In this section, however, I discuss only the homoskedastic linear model

$$Y_i = \sum_{j=1}^p \beta_j X_{ij} + \varepsilon_i; \quad i = 1, \dots, n, \tag{18}$$

where $\{\beta_j : j = 1, \dots, p\}$ are constant parameters to be estimated, $\{X_{ij} : i = 1, \dots, n; j = 1, \dots, p\}$ are covariates that are fixed in repeated samples, and $\{\varepsilon_i : i = 1, \dots, n\}$ are unobserved random variables that are independently and identically distributed with means of zero and finite variances. The model in Equation 18 is assumed to be sparse in the sense that most of the β_j s are zero but a relatively small number are nonzero. It is possible that $p > n$, in which case the β_j s cannot be estimated by OLS.

The LASSO consists of estimating the β_j s by solving the penalized least squares problem

$$\hat{\mathbf{b}} \equiv (\hat{b}_1, \dots, \hat{b}_p)' = \arg \min_{\mathbf{b}} \sum_{i=1}^n \left(Y_i - \sum_{j=1}^p b_j X_{ij} \right)^2 + \lambda \sum_{j=1}^p |b_j|,$$

where $\lambda > 0$ is a constant penalization parameter. This estimator has a complicated, non-normal asymptotic distribution even if $p < n$ (Knight & Fu 2000). The bootstrap does not estimate this distribution consistently (Chatterjee & Lahiri 2011).

The adaptive LASSO (ALASSO) (Zou 2006) is a variant of the LASSO that overcomes these problems. The ALASSO estimates $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)'$ in two steps. The first step consists of obtaining a $n^{-1/2}$ -consistent initial estimate $\tilde{\mathbf{b}}$ of $\boldsymbol{\beta}$. For example, $\tilde{\mathbf{b}}$ can be the OLS estimate if $p < n$ and the LASSO estimate if $p > n$. Denote the initial estimate by $\tilde{\mathbf{b}} = (\tilde{b}_1, \dots, \tilde{b}_p)'$. The second step estimate is

$$\hat{\mathbf{b}}_{AL} = \arg \min_{\mathbf{b}} \sum_{i=1}^n \left(Y_i - \sum_{\substack{j=1 \\ \tilde{b}_j \neq 0}}^p b_j X_{ij} \right)^2 + \lambda \sum_{\substack{j=1 \\ \tilde{b}_j \neq 0}}^p \frac{|b_j|}{|\tilde{b}_j|},$$

where $\hat{b}_{AL,j} = 0$ if $\tilde{b}_j = 0$. The ALASSO estimate is asymptotically normal. In particular, let the number of nonzero components of $\boldsymbol{\beta}$ be fixed. Define $A_+ = \{j : \beta_j \neq 0\}$, $\boldsymbol{\beta}_+ = \{\beta_j : j \in A_+\}$, and $\hat{\mathbf{b}}_{AL+} = \{\hat{b}_{AL,j} : j \in A_+\}$. Then we have $n^{1/2}(\hat{\mathbf{b}}_{AL+} - \boldsymbol{\beta}_+) \rightarrow^d N(0, V_+)$, where V_+ is the covariance matrix of the OLS estimate of $\boldsymbol{\beta}_+$ obtained from a model that contains only explanatory variables X_{ij} for which $j \in A_+$. This property is called oracle efficiency.

To implement the bootstrap with the ALASSO, let $\bar{\varepsilon}_i = Y_i - \hat{\mathbf{b}}'_{AL} \mathbf{X}_i$ ($i = 1, \dots, n$) be the ALASSO residuals, $\bar{\varepsilon} = n^{-1} \sum_{i=1}^n \bar{\varepsilon}_i$, and $\hat{\varepsilon}_i = \bar{\varepsilon}_i - \bar{\varepsilon}$ be the centered residuals. Define

$$Y_i^* = \sum_{j=1}^p \hat{b}_{AL,j} X_{ij} + \varepsilon_i^*; \quad i = 1, \dots, n,$$

where the ε_i^* s are drawn randomly with replacement from the $\hat{\varepsilon}_i$ s. The bootstrap sample is $\{Y_i^*, \mathbf{X}_i : i = 1, \dots, n\}$. Let \mathbf{b}_{AL}^* denote the ALASSO estimate of $\boldsymbol{\beta}$ based on the bootstrap sample. Chatterjee & Lahiri (2013) give conditions under which the bootstrap provides asymptotic refinements for the ALASSO. As an illustration of their results, let \mathbf{c} be a $p \times 1$ vector of constants. Let t_n and t_n^* , respectively, be t statistics for testing the hypothesis $\mathbf{c}'\boldsymbol{\beta} = 0$ and its bootstrap analog. Specifically, we have

$$t_n = \frac{n^{1/2} \mathbf{c}'(\hat{\mathbf{b}}_{AL} - \boldsymbol{\beta})}{s_n}$$

and

$$t_n^* = \frac{n^{1/2} \mathbf{c}'(\mathbf{b}_{AL}^* - \hat{\mathbf{b}}_{AL})}{s_n^*},$$

where s_n is a consistent estimate of the standard error of $n^{1/2} \mathbf{c}'\hat{\mathbf{b}}_{AL}$, and s_n^* is its bootstrap analog. Let P^* denote probability under bootstrap sampling of the ε_i^* s conditional on the

X_{ij} s. Chatterjee & Lahiri (2013) give conditions under which we have

$$\sup_{\tau} |P(t_n \leq \tau) - P^*(t_n^* \leq \tau)| = o_p(n^{-1/2}).$$

This is more accurate than the error made by the asymptotic normal approximation, which is $O(n^{-1/2})$. Chatterjee & Lahiri (2013) also give conditions under which the bootstrap estimates the distribution of a modified t statistic with an accuracy of $O_p(n^{-1})$.

8. TIME SERIES DATA

Bootstrap sampling with time series data must capture the dependence of the data-generation process (DGP) in a suitable way. Methods for doing this depend on what one assumes about the form of the dependence and can be complex if one assumes little. This section summarizes several methods for implementing the bootstrap with time series data. The details of some methods are lengthy. More thorough presentations are provided by Härdle et al. (2003) and Horowitz (2001) and their references, as well as the references cited in this section.

Bootstrap sampling can be carried out relatively easily if there is a parametric model, such as an autoregressive moving average (ARMA) model, that reduces the DGP to a transformation of independent random variables. For example, suppose that the time series $\{X_t : t = 1, \dots, T\}$ is generated by the stationary, invertible, finite-order ARMA model

$$A(L, \boldsymbol{\alpha})X_t = B(L, \boldsymbol{\beta})U_t, \tag{19}$$

where A and B are known functions, L is the backshift operator, $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are vectors of parameters, and $\{U_t\}$ is a sequence of independently and identically distributed random variables with means

of zero. Let $\hat{\alpha}$ and $\hat{\beta}$ be $T^{1/2}$ -consistent, asymptotically normal estimators of α and β , and let $\{\hat{U}_t\}$ be the residuals of the estimated model in Equation 19, centered so that $T^{-1} \sum_{t=1}^T \hat{U}_t = 0$. Then a bootstrap sample $\{X_t^*\}$ can be generated as

$$A(L, \hat{\alpha})X_t^* = B(L, \hat{\beta})U_t^*,$$

where $\{U_t^*\}$ is an independent random sample from the empirical distribution of the centered residuals $\{\hat{U}_t\}$. If the distribution of U_t is assumed to belong to a known parametric family (e.g., the normal distribution), then $\{U_t^*\}$ can be generated by independent sampling from the estimated distribution. Bose (1988) provides a rigorous discussion of the use of the bootstrap with autoregressions. Bose (1990) discusses moving average models.

Another possibility is that the data are generated by a stationary, linear process. That is, the DGP has the form

$$X_i - \mu = \sum_{j=1}^{\infty} \alpha_j (X_{i-j} - \mu) + U_i, \tag{20}$$

where $\mu = E(X_i)$ for all i and $\{U_i\}$ is a sequence of independently and identically distributed random variables. Assume that $\sum_{j=1}^{\infty} \alpha_j^2 < \infty$ and that all of the roots of the power series $1 - \sum_{j=1}^{\infty} \alpha_j z^j$ are outside of the unit circle. Bühlmann (1997, 1998), Kreiss (1988, 1992), and Paparoditis (1996) proposed approximating Equation 20 by an AR(p) model in which $p = p(n)$ increases with increasing sample size n . Let $\{a_{nj} : j = 1, \dots, p\}$ denote least squares or Yule-Walker estimates of the coefficients of the approximating process, and let $\{\hat{U}_{nj}\}$ denote the centered residuals. This procedure, which is called the sieve bootstrap, consists of generating bootstrap samples according to the process

$$\hat{X}_i - m = \sum_{j=1}^p a_{nj} (\hat{X}_{i-j} - m) + \hat{U}_j, \tag{21}$$

where $m = n^{-1} \sum_{i=1}^n X_i$, and the \hat{U}_j s are sampled randomly with replacement from the U_{nj} s. Bühlmann (1997), Kreiss (1992, 2000), and Paparoditis (1996) give conditions under which this procedure consistently estimates the distributions of sample averages, sample autocovariances and autocorrelations, and the regression coefficients a_{nj} , among other statistics.

Choi & Hall (2000) investigate the ability of the sieve bootstrap to provide asymptotic refinements to the coverage probability of a one-sided confidence interval for the mean of a linear statistic when the X_i s are scalar random variables. A linear statistic has the form

$$\theta_n = (n - q + 1)^{-1} \sum_{i=1}^{n-q+1} G(X_i, \dots, X_{i+q-1}),$$

where $q \geq 1$ is a fixed integer, and G is a known function. Define $\theta = E[G(X_1, \dots, X_q)]$. Choi & Hall (2000) consider the problem of finding a one-sided confidence interval for θ . They give conditions under which the difference between the true and nominal coverage probabilities of this interval is $O(n^{-1+\varepsilon})$ for any $\varepsilon > 0$. This is only slightly larger than the difference of $O(n^{-1})$ that is available with independent data.

A third possibility is that the DGP is a Markov process or can be approximated by such a process. The class of Markov and approximate Markov processes contains autoregressive conditionally heteroskedastic (ARCH), generalized ARCH (GARCH), and many other processes that

are important in applications. The assumption that the DGP is a Markov or approximate Markov process is weaker and, therefore, more general than the assumption that the DGP belongs to a finite-dimensional parametric family or is linear.

When the DGP is a stationary Markov process, the bootstrap can be implemented by estimating the Markov transition density nonparametrically. Bootstrap samples are generated by the stochastic process implied by the estimated transition density. Call this procedure the Markov bootstrap (MB). The MB was proposed by Rajarshi (1990), who gives conditions under which it consistently estimates the asymptotic distribution of a statistic. Datta & McCormick (1995) give conditions under which the error in the MB estimator of the distribution function of a normalized sample average is almost surely $o(n^{-1/2})$.

Horowitz (2003) investigates the ability of the MB to provide asymptotic refinements for confidence intervals and tests based on Studentized statistics of the form

$$t_n = T^{1/2}[H(m) - H(\mu)]/s_T,$$

where H is a smooth function, $\mu = E(X_1)$, $m = T^{-1} \sum_{t=1}^T X_t$, and s_T^2 is a consistent estimator of the variance of the asymptotic distribution of $T^{1/2}[H(m) - H(\mu)]$. Thus, for example, a symmetrical $1 - \alpha$ confidence interval for $H(\mu)$ is

$$H(m) - z_{1-\alpha} s_T \leq H(\mu) \leq H(m) + z_{1-\alpha} s_T,$$

where $z_{1-\alpha}$ is a critical value. Horowitz (2003) gives conditions under which the difference between the true and nominal coverage probabilities of a symmetrical confidence interval is $O(n^{-3/2+\varepsilon})$ for any $\varepsilon > 0$. The difference for a one-sided confidence interval is $O(n^{-1+\varepsilon})$. In contrast, the asymptotic normal approximation makes errors of size $O(n^{-1})$ and $O(n^{-1/2})$, respectively, for symmetrical and one-sided confidence intervals.

The block bootstrap is a method that makes even weaker assumptions about the DGP. This method consists of dividing the data into blocks and sampling the blocks randomly with replacement. The blocks may be nonoverlapping (Carlstein 1986, Hall 1985) or overlapping (Hall 1985, Künsch 1989, Politis & Romano 1993). Andrews (2004) and Politis & Romano (1993) describe variants of these blocking schemes. Details of the various blocking schemes are presented in the references above and by Härdle et al. (2003), Hall & Horowitz (1996), and Lahiri (2003). The root mean-square errors of bootstrap estimators of distribution functions are smaller with overlapping blocks than with nonoverlapping ones. This suggests that overlapping blocks are preferred for applications, although the differences between the numerical results obtained with the two types of blocking are often very small (Andrews 2004). The rates of convergence of the errors made with overlapping and nonoverlapping blocks are the same.

Regardless of the blocking method that is used, the block length must increase with increasing sample size n to make bootstrap estimators of moments and distribution functions consistent. The block length that minimizes mean square estimation error depends on what is being estimated. Hall et al. (1995) show that, with either overlapping or nonoverlapping blocks, the optimal block length is $l \propto n^r$, with $r = 1/3$ for estimating bias or variance, $r = 1/4$ for estimating a one-sided distribution function [e.g., $P(t_n \leq \tau)$], and $r = 1/5$ for estimating a symmetrical distribution function [e.g., $P(|t_n| \leq \tau)$]. The results obtained with the block bootstrap can be sensitive to the choice of block length. Hall et al. (1995) and Lahiri (2003) describe data-based methods for choosing the block length in applications.

Block bootstrap sampling does not exactly replicate the dependence structure of the original DGP. For example, if nonoverlapping blocks are used, then bootstrap observations that belong to

the same block are deterministically related, whereas observations that belong to different blocks are independent. This dependence structure is unlikely to be present in the original DGP. As a result, the finite-sample covariance matrices of the asymptotic forms of parameter estimators obtained from the original sample and from the bootstrap sample are different. The practical consequence of this difference is that asymptotic refinements cannot be obtained by applying the usual formulae for test statistics to the block bootstrap sample. It is necessary to develop special formulae for the bootstrap versions of test statistics. These formulae contain factors that correct for the differences between the asymptotic covariances of the original sample and bootstrap versions of test statistics without distorting the higher-order terms of asymptotic expansions that produce refinements. Härdle et al. (2003), Hall & Horowitz (1996), and Lahiri (2003), among the other references cited in this section, describe the appropriate versions of test statistics and explain why they are needed.

The asymptotic refinements provided by the block bootstrap depend in a complicated way on what is being estimated. Andrews (2004), Härdle et al. (2003), Hall & Horowitz (1996), and Hall et al. (1995), among the other references cited in this section, provide details and numerical illustrations. The estimation errors made by the block bootstrap converge more slowly than the errors made by methods that make stronger assumptions about the DGP but more rapidly than the errors made by asymptotic normal approximations.

9. CONCLUSIONS

The bootstrap consistently estimates the asymptotic distributions of econometric estimators and test statistics under conditions that are sufficiently general to accommodate most applications. To achieve consistency, however, the bootstrap sampling procedure must be matched to the application. Often, it suffices to sample one's data randomly with replacement, but there are important cases that require sampling procedures that are more complex and/or modifications of the estimation procedure. This article reviews several examples of such cases. When used properly, the bootstrap provides a way to substitute computation for mathematical analysis if analytical calculation of the asymptotic distribution of an estimator or test statistic is difficult or impossible. The bootstrap is more accurate than subsampling in settings where the bootstrap is consistent and the relative accuracy of the two methods is known.

Under conditions that are stronger than those required for consistency but still general enough to accommodate many econometric applications, the bootstrap provides a more accurate approximation to the finite-sample distribution of an estimator or test statistic than does first-order asymptotic theory. The approximations of first-order asymptotic theory are often inaccurate with samples of the sizes encountered in applications. As a result, the errors in the rejection probabilities of hypothesis tests and the coverage probabilities of confidence intervals based on first-order approximations can be very large. The bootstrap can provide dramatic reductions in these errors. In many cases of practical importance, the bootstrap essentially eliminates finite-sample errors in rejection and coverage probabilities.

This article emphasizes the need for care in applying the bootstrap and the importance of asymptotically pivotal statistics for obtaining asymptotic refinements. However, even if asymptotic refinements are not available or not desired, the bootstrap should be used to obtain confidence intervals and critical values for hypothesis tests, not standard errors. Standard errors are important because they bear a simple relation to quantiles of the normal distribution, but bootstrap distributional approximations are non-normal. There is no simple relation between standard errors and quantiles of bootstrap distributions.

DISCLOSURE STATEMENT

The author is not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

LITERATURE CITED

- Abrevaya J, Huang J. 2005. On the bootstrap of the maximum score estimator. *Econometrica* 73:1175–204
- Amemiya T. 1985. *Advanced Econometrics*. Cambridge, MA: Harvard Univ. Press
- Andrews DWK. 2000. Inconsistency of the bootstrap when a parameter is on the boundary of the parameter space. *Econometrica* 68:399–405
- Andrews DWK. 2004. The block-block bootstrap: improved asymptotic refinements. *Econometrica* 72:673–700
- Andrews DWK, Barwick PL. 2012. Inference for parameters defined by moment inequalities: a recommended moment selection procedure. *Econometrica* 80:2805–26
- Andrews DWK, Guggenberger P. 2009. Hybrid and size-corrected subsampling methods. *Econometrica* 77:721–62
- Andrews DWK, Guggenberger P. 2010. Asymptotic size and a problem with subsampling and with the m out of n bootstrap. *Econom. Theory* 26:426–68
- Andrews DWK, Han S. 2009. Invalidity of the bootstrap and the m out of n bootstrap for confidence interval endpoints defined by moment inequalities. *Econom. J.* 12:S172–99
- Bassett G, Koenker R. 1978. Asymptotic theory of least absolute error regression. *J. Am. Stat. Assoc.* 73:618–21
- Beran R, Ducharme GR. 1991. *Asymptotic Theory for Bootstrap Methods in Statistics*. Montreal, Can.: Publ. CRM
- Bertail P. 1997. Second-order properties of an extrapolated bootstrap without replacement under weak assumptions. *Bernoulli* 3:149–79
- Bertail P, Politis DN, Romano JP. 1999. On subsampling estimators with unknown rate of convergence. *J. Am. Stat. Assoc.* 94:569–79
- Bickel PJ, Freedman DA. 1981. Some asymptotic theory for the bootstrap. *Ann. Stat.* 9:1196–217
- Bickel PJ, Götze F, van Zwet WR. 1997. Resampling fewer than n observations: gains, losses, and remedies for losses. *Stat. Sin.* 7:1–32
- Bose A. 1988. Edgeworth correction by bootstrap in autoregressions. *Ann. Stat.* 16:1709–22
- Bose A. 1990. Bootstrap in moving average models. *Ann. Inst. Stat. Math.* 42:753–68
- Bowlus A, Neumann G, Kiefer N. 2001. Equilibrium search models and the transition from school to work. *Int. Econ. Rev.* 42:317–43
- Brown BW. 2000. Simulation variance reduction for bootstrapping. In *Simulation-Based Econometrics: Methods and Applications*, ed. R Mariano, T Schuermann, M Weeks, pp. 437–73. Cambridge, UK: Cambridge Univ. Press
- Bugni FA. 2010. Bootstrap inference in partially identified models defined by moment inequalities. *Econometrica* 78:735–53
- Bugni FA. 2016. Comparison of inferential methods in partially identified models in terms of error in coverage probability. *Econom. Theory* 32:187–242
- Bugni FA, Canay IA, Shi X. 2015. Specification tests for partially identified models defined by moment conditions. *J. Econom.* 185:259–82
- Bugni FA, Canay IA, Shi X. 2017. Inference for subvectors and other functions of partially identified parameters in moment inequality models. *Quant. Econ.* 8:1–38
- Bühlmann P. 1997. Sieve bootstrap for time series. *Bernoulli* 3:123–48
- Bühlmann P. 1998. Sieve bootstrap for smoothing nonstationary time series. *Ann. Stat.* 26:48–83
- Bühlmann P, van de Geer S. 2011. *Statistics for High Dimensional Data*. Berlin: Springer
- Cao-Abad R. 1991. Rate of convergence for the wild bootstrap in nonparametric regression. *Ann. Stat.* 19:2226–31
- Carlstein E. 1986. The use of subseries methods for estimating the variance of a general statistic from a stationary time series. *Ann. Stat.* 14:1171–79
- Cattaneo M, Jansson M, Nagasawa K. 2018. *Bootstrap-based inference for cube root consistent estimators*. Work. Pap., Dep. Stat., Univ. Mich., Ann Arbor

- Cavanagh CL. 1987. *Limiting behavior of estimators defined by optimization*. Unpublished manuscript, Dep. Econ., Harvard Univ., Cambridge, MA
- Chatterjee A, Lahiri SN. 2011. Bootstrapping Lasso estimators. *J. Am. Stat. Assoc.* 106:608–25
- Chatterjee A, Lahiri SN. 2013. Rates of convergence of the adaptive Lasso estimators to the oracle distribution and higher order refinements by the bootstrap. *Ann. Stat.* 41:1232–59
- Chernozhukov V, Chetverikov D, Kato K. 2013. Gaussian approximations and multiplier bootstrap for maxima of sums of high-dimensional random vectors. *Ann. Stat.* 41:2786–819
- Chernozhukov V, Hong H. 2004. Likelihood estimation and inference in a class of nonregular econometric models. *Econometrica* 72:1445–80
- Chesher A, Jewitt I. 1987. The bias of a heteroskedasticity consistent covariance matrix estimator. *Econometrica* 55:1217–22
- Choi E, Hall P. 2000. Bootstrap confidence regions computed from autoregressions of arbitrary order. *J. R. Stat. Soc. B* 62:461–77
- Chung K-H, Lee SMS. 2001. Optimal bootstrap sample size in construction of percentile confidence bounds. *Scand. J. Stat.* 28:225–39
- Datta S, McCormick WP. 1995. Some continuous Edgeworth expansions for Markov chains with applications to bootstrap. *J. Multivar. Anal.* 52:83–106
- Davidson R, Flachaire E. 2008. The wild bootstrap, tamed at last. *J. Econom.* 146:162–69
- Davidson R, MacKinnon JG. 1999. The size distortion of bootstrap tests. *Econom. Theory* 15:361–76
- Davison AC, Hinkley DV. 1997. *Bootstrap Methods and Their Application*. Cambridge, UK: Cambridge Univ. Press
- De Angelis D, Hall P, Young GA. 1993. Analytical and bootstrap approximations to estimator distributions in L^1 regression. *J. Am. Stat. Assoc.* 88:1310–16
- Donald SG, Paarsch HJ. 1996. Identification, estimation, and testing in empirical models of auctions within the independent private values paradigm. *Econom. Theory* 12:517–67
- Efron B. 1979. Bootstrap methods: another look at the jackknife. *Ann. Stat.* 7:1–26
- Efron B, Tibshirani RJ. 1993. *An Introduction to the Bootstrap*. New York: Chapman & Hall
- Eicker F. 1963. Asymptotic normality and consistency of the least squares estimators for families of linear regressions. *Ann. Math. Stat.* 34:447–56
- Eicker F. 1967. Limit theorems for regression with unequal and dependent errors. In *Proceedings of the 5th Berkeley Symposium on Mathematical Statistics and Probability*, ed. L LeCam, J Neyman, pp. 59–82. Berkeley, CA: Univ. Calif. Press
- Fan J, Gijbels I. 1996. *Local Polynomial Modelling and Its Applications*. New York: Chapman & Hall
- Fan J, Hu T-C, Truong YK. 1994. Robust nonparametric function estimation. *Scand. J. Stat.* 21:433–46
- Flinn C, Heckman J. 1982. New methods for analyzing structural models of labor force dynamics. *J. Econom.* 18:115–68
- Gill RD. 1989. Non- and semi-parametric maximum likelihood estimators and the von Mises method (Part 1). *Scand. J. Stat.* 16:97–128
- Hahn J. 1995. Bootstrapping the quantile regression estimators. *Econom. Theory* 11:105–21
- Hall P. 1985. Resampling a coverage process. *Stoch. Process Appl.* 19:259–69
- Hall P. 1990. Asymptotic properties of the bootstrap for heavy-tailed distributions. *Ann. Probab.* 18:1342–60
- Hall P. 1992. *The Bootstrap and Edgeworth Expansion*. Berlin: Springer
- Hall P. 1994. Methodology and theory for the bootstrap. In *Handbook of Econometrics*, Vol. 4, ed. RF Engle, DF McFadden, pp. 2341–81. Amsterdam: Elsevier
- Hall P, Horowitz JL. 1996. Bootstrap critical values for tests based on generalized-method-of-moments estimators. *Econometrica* 64:891–916
- Hall P, Horowitz JL. 2013. A simple bootstrap method for constructing nonparametric confidence bands for functions. *Ann. Stat.* 41:1892–921
- Hall P, Horowitz JL, Jing B-Y. 1995. On blocking rules for the bootstrap with dependent data. *Biometrika* 82:561–74
- Hall P, Jing B-Y. 1996. On sample reuse methods for dependent data. *J. R. Stat. Soc. B* 58:727–37
- Härdle W. 1990. *Applied Nonparametric Regression*. Cambridge, UK: Cambridge Univ. Press

- Härdle W, Horowitz JL, Kreiss J-P. 2003. Bootstrap methods for time series. *Int. Stat. Rev.* 71:435–59
- Härdle W, Mammen E. 1993. Comparing nonparametric versus parametric regression fits. *Ann. Stat.* 21:1926–47
- Härdle W, Marron JS. 1991. Bootstrap simultaneous error bars for nonparametric regression. *Ann. Stat.* 19:778–96
- Heckman JJ, Smith J, Clements N. 1997. Making the most out of programme evaluations and social experiments: accounting for heterogeneity in programme impacts. *Rev. Econ. Stud.* 64:487–535
- Hirano K, Porter JR. 2003. Asymptotic efficiency in parametric structural models with parameter dependent support. *Econometrica* 71:1307–38
- Hong H, Li J. 2015. *The numerical bootstrap*. Work. Pap., Dep. Econ., Stanford Univ., Stanford, CA
- Horowitz JL. 1992. A smoothed maximum score estimator for the binary response model. *Econometrica* 60:505–31
- Horowitz JL. 1994. Bootstrap-based critical values for the information-matrix test. *J. Econom.* 61:395–411
- Horowitz JL. 1997. Bootstrap methods in econometrics: theory and numerical performance. In *Advances in Economics and Econometrics: Theory and Applications*, Vol. 3, ed. DM Kreps, KF Wallis, pp. 188–222. Cambridge, UK: Cambridge Univ. Press
- Horowitz JL. 1998a. Bootstrap methods for covariance structures. *J. Hum. Resour.* 33:39–61
- Horowitz JL. 1998b. Bootstrap methods for median regression models. *Econometrica* 66:1327–51
- Horowitz JL. 2001. The bootstrap. In *Handbook of Econometrics*, Vol. 5, ed. JL Heckman, E Leamer, pp. 3159–228. Amsterdam: Elsevier
- Horowitz JL. 2002. Bootstrap critical values for tests based on the smoothed maximum score estimator. *J. Econom.* 111:141–67
- Horowitz JL. 2003. Bootstrap methods for Markov processes. *Econometrica* 71:1049–82
- Horowitz JL, Krishnamurthy A. 2017. A bootstrap method for constructing pointwise and uniform confidence bands for conditional quantile functions. *Stat. Sin.* In press
- Jeong G, Maddala GS. 1993. A perspective on application of bootstrap methods in econometrics. In *Handbook of Statistics*, Vol. 11, ed. GS Maddala, CR Rao, HD Vinod, pp. 573–610. Amsterdam: North-Holland
- Kennan J, Neumann GR. 1988. *Why does the information matrix test reject so often?* Work. Pap. 88–4, Dep. Econ., Univ. Iowa, Iowa City
- Kim J, Pollard D. 1990. Cube root asymptotics. *Ann. Stat.* 18:191–219
- Knight K, Fu W. 2000. Asymptotics for LASSO-type estimators. *Ann. Stat.* 28:1356–78
- Koenker R, Bassett G. 1978. Regression quantiles. *Econometrica* 46:33–50
- Kreiss J-P. 1988. *Asymptotic statistical inference for a class of stochastic processes*. Work. Pap., Univ. Hamburg, Ger.
- Kreiss J-P. 1992. Bootstrap procedures for AR(∞) processes. In *Bootstrapping and Related Techniques*, ed. KH Jöckel, G Rothe, W Sender, pp. 107–13. Lect. Notes Econ. Math. Syst. 376. Berlin: Springer
- Kreiss J-P. 2000. *Residual and wild bootstrap for infinite order autoregressions*. Work. Pap., Inst. Math. Stoch., Tech. Univ. Braunschweig, Ger.
- Künsch HR. 1989. The jackknife and the bootstrap for general stationary observations. *Ann. Stat.* 17:1217–41
- Lahiri SN. 2003. *Resampling Methods for Dependent Data*. Berlin: Springer
- Liu RY. 1988. Bootstrap procedures under some non-i.i.d. models. *Ann. Stat.* 16:1696–708
- MacKinnon JG, White H. 1985. Some heteroskedasticity consistent covariance matrix estimators with improved finite sample properties. *J. Econom.* 29:305–25
- Mammen E. 1992. *When Does Bootstrap Work? Asymptotic Results and Simulations*. Berlin: Springer
- Mammen E. 1993. Bootstrap and wild bootstrap for high dimensional linear models. *Ann. Stat.* 21:255–85
- Manski CF. 1975. Maximum score estimation of the stochastic utility model of choice. *J. Econom.* 3:205–28
- Manski CF. 1985. Semiparametric analysis of discrete response: asymptotic properties of the maximum score estimator. *J. Econom.* 27:313–34
- Orme C. 1990. The small-sample performance of the information-matrix test. *J. Econom.* 46:309–31
- Paparoditis E. 1996. Bootstrapping autoregressive and moving average parameter estimates of infinite order vector autoregressive processes. *J. Multivar. Anal.* 57:277–96
- Patra RK, Seijo E, Sen B. 2018. A consistent bootstrap procedure for the maximum score estimator. *J. Econom.* 205:488–507

- Politis DN, Romano JP. 1993. The stationary bootstrap. *J. Am. Stat. Assoc.* 89:1303–13
- Politis DN, Romano JP. 1994. Large sample confidence regions based on subsamples under minimal assumptions. *Ann. Stat.* 22:2031–50
- Politis DN, Romano JP, Wolf M. 1997. Subsampling for heteroskedastic time series. *J. Econom.* 81:281–317
- Politis DN, Romano JP, Wolf M. 1999. *Subsampling*. Berlin: Springer
- Powell JL. 1984. Least absolute deviations estimation for the censored regression model. *J. Econom.* 25:303–25
- Powell JL. 1986. Censored regression quantiles. *J. Econom.* 32:143–55
- Pretorius C, Swanepoel JWH. 2018. On the asymptotic theory of new bootstrap confidence bounds. *Ann. Stat.* 46:438–56
- Rajarshi MB. 1990. Bootstrap in Markov-sequences based on estimates of transition density. *Ann. Inst. Stat. Math.* 42:253–68
- Rice J. 1984. Bandwidth choice for nonparametric regression. *Ann. Stat.* 12:1215–30
- Shao U, Tu D. 1995. *The Jackknife and Bootstrap*. Berlin: Springer
- Spokoiny V, Zhilova M. 2015. Bootstrap confidence sets under model misspecification. *Ann. Stat.* 43:2653–75
- Swanepoel JWH. 1986. A note on proving that the (modified) bootstrap works. *Commun. Stat. Theory Methods* 15:3193–203
- Taylor LW. 1987. The size bias of White's information matrix test. *Econ. Lett.* 24:63–67
- Tibshirani R. 1996. Regression shrinkage and selection via the Lasso. *J. R. Stat. Soc. B* 58:267–88
- Vinod HD. 1993. Bootstrap methods: applications in econometrics. In *Handbook of Statistics*, Vol. 11, ed. GS Maddala, CR Rao, HD Vinod, pp. 629–61. Amsterdam: North-Holland
- White H. 1980. A heteroscedasticity-consistent covariance matrix estimator and a direct test for heteroscedasticity. *Econometrica* 48:817–38
- White H. 1982. Maximum likelihood estimation of misspecified models. *Econometrica* 50:1–25
- Wu CFJ. 1986. Jackknife, bootstrap and other resampling methods in regression analysis. *Ann. Stat.* 14:1261–95
- Yu K, Jones MC. 1998. Local linear quantile regression. *J. Am. Stat. Assoc.* 93:228–37
- Zou H. 2006. The adaptive Lasso and its oracle properties. *J. Am. Stat. Assoc.* 101:1418–29