INCONSISTENCY OF THE BOOTSTRAP WHEN A PARAMETER IS ON THE BOUNDARY OF THE PARAMETER SPACE

BY

DONALD W. K. ANDREWS

COWLES FOUNDATION PAPER NO. 994

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

2000

http://cowles.econ.yale.edu/

Econometrica, Vol. 68, No. 2 (March, 2000), 399–405

NOTES AND COMMENTS

INCONSISTENCY OF THE BOOTSTRAP WHEN A PARAMETER IS ON THE BOUNDARY OF THE PARAMETER SPACE

BY DONALD W. K. ANDREWS[1]

1. INTRODUCTION

RESEARCHERS IN ECONOMETRICS sometimes view the bootstrap as a panacea for statistical inference. Indeed, it does have widespread applicability; e.g., see Hall (1992), Efron and Tibshirani (1993), and Hall and Horowitz (1996). Nevertheless, there are situations where the bootstrap is not consistent. In this note, we provide such a counterexample to the bootstrap. This counterexample is quite simple, but it generalizes to a wide variety of estimation problems that are of importance in econometric applications. The counterexample should serve as a useful reminder that the bootstrap is not a universal solution to problems of statistical inference.

We consider the maximum likelihood estimator of the mean of a sample of iid normal random variables with mean μ and variance one (denoted $N(\mu, 1)$) when the mean is restricted to be nonnegative. The maximum likelihood estimator in this case is just the maximum of the sample mean and zero. When the true mean is zero, the bootstrap is not asymptotically correct to first order. This is true of the nonparametric bootstrap based on the empirical distribution function, as well as the parametric bootstrap based on the restricted or unrestricted maximum likelihood estimator.

The above counterexample to the bootstrap generalizes to a wide variety of estimation problems that have considerable relevance in applications. For example, in models with random coefficients, it is often the case that the estimated variances of some of the random coefficients are small and, hence, the true variances of some of the random coefficients may be zero. If any of the coefficient variances are zero, the bootstrap is not consistent. More generally, the bootstrap is not consistent if the parameter is on a boundary of the parameter space defined by linear or nonlinear *inequality* or mixed *inequality/equality* constraints. If a parameter is on a boundary defined by linear or nonlinear *equality* constraints, then the bootstrap is consistent.

We provide four alternatives to the bootstrap that are asymptotically correct to first order in this context. The first method is based on testing. The second is a parametric bootstrap procedure that uses an estimator that shrinks towards the boundary of the parameter space. The third is a subsample method of Wu (1990) and Politis and Romano (1994). The fourth is a "rescaled" bootstrap procedure. We note that none of these methods is consistent when the true mean is of the form $\mu_n = \mu/n^{1/2}$. Thus, these methods may not work as well as desired when the true mean is small, but not zero. We show that it is impossible to derive a consistent scheme when the mean equals $\mu_n = \mu/n^{1/2}$.

[1] The author gratefully acknowledges the research support of the National Science Foundation via Grant Numbers SBR-9410675 and SBR-9730277. The author thanks Peter Bickel and Joe Romano for references and Dale Stahl and two referees for helpful comments.

In Andrews (1997), we consider two types of bootstrap percentile confidence intervals in the above example. The first is what Hall (1992) calls a bootstrap percentile confidence interval. The second is what Efron and Tibshirani (1993) call a bootstrap percentile confidence interval. We find that they both have asymptotic coverage probability that exceeds the nominal asymptotic level when the true value of the mean μ equals zero.

On the other hand, consider the level $\alpha \in (0, 1/2)$ one-sided bootstrap test of H_0 : $\mu = 0$ versus H_1 : $\mu > 0$ that rejects H_0 when $n^{1/2}\hat{\mu}_n > \hat{t}_{1-\alpha}$, where $\hat{t}_{1-\alpha}$ is the $1-\alpha$ quantile of $n^{1/2}(\hat{\mu}_n^* - \hat{\mu}_n)$ conditional on $\hat{F}_n$. This test has the correct asymptotic null rejection rate; see the end of Section 4.

The remainder of this paper is organized as follows. Section 2 discusses counterexamples to the bootstrap that are already in the literature. Section 3 establishes the inconsistency of the bootstrap in the example described above. Section 4 considers alternatives to the bootstrap for the example above.

2. RESULTS IN THE LITERATURE

The literature contains a number of examples in which the bootstrap of Efron (1979) does not consistently estimate the true distribution of a statistic correctly to first order. These examples are all nonstandard in some way or other. The example considered in this paper is very simple and close to being standard.

Bickel and Freedman (1981) provide two counterexamples to the (nonparametric) bootstrap. Their first counterexample is a U-statistic of degree two in which the kernel $\omega(x, x)$ does not satisfy the condition $\int \omega^2(x, x)\, dF(x) < \infty$, where F denotes the true distribution of the data. Their second example is the largest order statistic from an iid sample of uniform $(0, \theta)$ random variables. This example is extended in Bickel, Götze, and van Zwet (1997, Example 3).

Other counterexamples to the (nonparametric) bootstrap in the literature include: extrema for unbounded distributions (see Athreya and Fukuchi (1994) and Deheuvels, Mason, and Shorack (1993)); the sample mean in the case of infinite variance random variables (see Babu (1984) and Athreya (1987)); some versions of Hodges' super-efficient estimator, viz., those with the rescaling parameter $b \in (0, 1)$ (see Beran (1982, 1997) and Putter and van Zwet (1996)); Stein's shrinkage estimator (see Beran (1997)); degenerate U and V statistics (see Bretagnolle (1983)); estimators of the eigenvalues of a covariance matrix whose eigenvalues are not distinct (see Beran and Srivastava (1985)); nondifferentiable functions of the empirical distribution function (see Dümbgen (1993)); the distribution of the square of a sample average when the population mean equals zero (see Datta (1995)); and the nonparametric kernel estimator of the mode of a smooth unimodal density when the smoothing parameter (for both the estimator and the bootstrap) is chosen to be optimal for the estimation problem (see Romano (1988)). A counterexample to the parametric bootstrap is given in Sriram (1993) for critical branching processes with immigration estimated by maximum likelihood. Beran (1997) gives necessary and sufficient conditions for an "intuitive" parametric bootstrap to be consistent and for a nonparametric bootstrap to be consistent (for cases in which the underlying random variables have finite support).

The counterexample to the bootstrap introduced in this paper, based on a parameter being on the boundary of the parameter space, seems simpler and more relevant to economic applications than most of the counterexamples just listed.

Another counterexample to the bootstrap, which is of relevance to economic problems, is the failure of the residual-based bootstrap in the estimation of the autoregressive coefficient in a first order autoregressive model when the true coefficient equals unity (see Basawa, Mallik, McCormick, Reeves, and Taylor (1991)). Unlike the counterexample considered in the present paper, however, this counterexample does not involve the standard nonparametric bootstrap applied in an iid context.

3. THE COUNTEREXAMPLE

We now analyze the counterexample described in the Introduction. Let $\{X_i: i \geq 1\}$ be a sequence of independent identically distributed (iid) $N(\mu, 1)$ random variables. Suppose the parameter space for μ is $R^+ := \{y: y \geq 0\}$, where $:=$ denotes equality by definition. The maximum likelihood estimator of μ in this case is $\hat{\mu}_n := \max\{\overline{X}_n, 0\}$, where $\overline{X}_n := (1/n)\sum_{i=1}^n X_i$. It is easy to see that

$$(1) \qquad n^{1/2}(\hat{\mu}_n - \mu) \xrightarrow{d} \begin{cases} Z & \text{if } \mu > 0 \\ \max\{Z, 0\} & \text{if } \mu = 0 \end{cases} \quad \text{as } n \to \infty, \quad \text{where} \quad Z \sim N(0,1).$$

First, we consider the standard nonparametric bootstrap. Let $\{X_i^*: i \leq n\}$ be iid with $X_i^* \sim \hat{F}_n$, where $\hat{F}_n(x) := (1/n)\sum_{i=1}^n 1(X_i \leq x)$. The bootstrap maximum likelihood estimator $\hat{\mu}_n^*$ is defined by $\hat{\mu}_n^* := \max\{\overline{X}_n^*, 0\}$, where $\overline{X}_n^* := (1/n)\sum_{i=1}^n X_i^*$. In cases where the bootstrap is asymptotically valid, the bootstrap distribution of $n^{1/2}(\hat{\mu}_n^* - \hat{\mu}_n)$ is used to approximate the distribution of $n^{1/2}(\hat{\mu}_n - \mu)$. Asymptotic validity of the bootstrap requires that with probability one the asymptotic distribution of $n^{1/2}(\hat{\mu}_n^* - \hat{\mu}_n)$ conditional on $\{\hat{F}_n: n \geq 1\}$ equals the asymptotic distribution of $n^{1/2}(\hat{\mu}_n - \mu)$. We show that this does not hold in the present example.

Suppose $\mu = 0$. Let $A_c := \{\liminf_{n \to \infty} n^{1/2}\overline{X}_n < -c\}$ for $0 < c < \infty$. By the law of the iterated logarithm, $P(A_c) = 1$. For $\omega \in A_c$, consider a subsequence $\{n_k: k \geq 1\}$ of $\{n: n \geq 1\}$ such that $n_k^{1/2}\overline{X}_{n_k}(\omega) \leq -c$ for all k. Then,

$$\begin{aligned}
(2) \qquad & n_k^{1/2}(\hat{\mu}_{n_k}^* - \hat{\mu}_{n_k}(\omega)) \\
&= \max\left\{n_k^{1/2}\left(\overline{X}_{n_k}^* - \overline{X}_{n_k}(\omega)\right) + n_k^{1/2}\overline{X}_{n_k}(\omega), 0\right\} - \max\left\{n_k^{1/2}\overline{X}_{n_k}(\omega), 0\right\} \\
&\leq \max\left\{n_k^{1/2}\left(\overline{X}_{n_k}^* - \overline{X}_{n_k}(\omega)\right) - c, 0\right\} \\
&\xrightarrow{d} \max\{Z - c, 0\} \text{ as } k \to \infty \text{ conditional on } \{\hat{F}_n: n \geq 1\} \\
&\leq \max\{Z, 0\},
\end{aligned}$$

where the last inequality is strict with positive probability and the convergence in distribution holds by a triangular array central limit theorem. So, along the subsequence $\{n_k\}$, $n_k^{1/2}(\hat{\mu}_{n_k}^* - \hat{\mu}_{n_k}(\omega)) \nrightarrow_d \max\{Z, 0\}$ as $k \to \infty$ conditional on $\{\hat{F}_{n_k}: k \geq 1\}$. Hence, $n^{1/2}(\hat{\mu}_n^* - \hat{\mu}_n(\omega)) \nrightarrow_d \max\{Z, 0\}$ as $n \to \infty$ conditional on $\{\hat{F}_n: n \geq 1\}$. This is true for all $\omega \in A_c$. We conclude that with probability one (with respect to the randomness in $\{\hat{F}_n: n \geq 1\}$), the bootstrap distribution is not consistent.

Note that the bootstrap also is not correct when $\mu = 0$ for sample paths $\omega \in B_c := \{\limsup_{n \to \infty} n^{1/2}\overline{X}_n > c\}$ for any $0 < c < \infty$ and sample sizes $\{n_m: m \geq 1\}$ for which $n_m^{1/2}\overline{X}_{n_m}(\omega) \geq c$ for all m. In this case, we have

$$\begin{aligned}
(3) \qquad n_m^{1/2}(\hat{\mu}_{n_m}^* - \hat{\mu}_{n_m}(\omega)) &= \max\left\{n_m^{1/2}\left(\overline{X}_{n_m}^* - \overline{X}_{n_m}(\omega)\right), -n_m^{1/2}\overline{X}_{n_m}(\omega)\right\} \\
&\leq \max\left\{n_m^{1/2}\left(\overline{X}_{n_m}^* - \overline{X}_{n_m}(\omega)\right), -c\right\} \\
&\xrightarrow{d} \max\{Z, -c\} \text{ as } m \to \infty \text{ conditional on } \{\hat{F}_n: n \geq 1\} \\
&\leq \max\{Z, 0\},
\end{aligned}$$

where the last inequality is strict with positive probability. Note that $P(B_c) = 1$ for all $0 < c < \infty$. Thus, the bootstrap is incorrect both when $n^{1/2}\overline{X}_n(\omega)$ is negative for n large and when $n^{1/2}\overline{X}_n(\omega)$ is positive for n large. In both cases, the bootstrap distribution is too small (i.e., has too much mass to the left) when $\mu = 0$.

One can see why the bootstrap fails when $\mu = 0$ by inspecting equations (2) and (3) and utilizing the fact that $n^{1/2}(\overline{X}_n^* - \overline{X}_n(\omega))$ and $n^{1/2}(\overline{X}_n - \mu)$ have the same $N(0,1)$ distribution asymptotically. When $\overline{X}_n(\omega) = -c$ for $c > 0$, then $n^{1/2}(\hat{\mu}_n^* - \hat{\mu}_n(\omega)) = 0$ whenever $\overline{X}_n^* - \overline{X}_n(\omega) < c$, whereas $n^{1/2}(\hat{\mu}_n - \mu) = 0$ whenever $\overline{X}_n - \mu < 0$. Because $c > 0$, $n^{1/2}(\hat{\mu}_n^* - \hat{\mu}_n(\omega))$ has a higher probability of equalling zero than does $n^{1/2}(\hat{\mu}_n - \mu)$. Alternatively, when $\overline{X}_n(\omega) = c$ for $c > 0$, then $n^{1/2}(\hat{\mu}_n^* - \hat{\mu}_n(\omega)) = \max\{n^{1/2}(\overline{X}_n^* - \overline{X}_n(\omega)), -c\}$, whereas $n^{1/2}(\hat{\mu}_n - \mu) = \max\{n^{1/2}(\overline{X}_n - \mu), 0\}$. Because $-c < 0$, the distribution of $n^{1/2}(\hat{\mu}_n^* - \hat{\mu}_n(\omega))$ is to the left of that of $n^{1/2}(\hat{\mu}_n - \mu)$.

We note that consistency of the nonparametric bootstrap cannot be rescued by using the distribution of $n^{1/2}(\hat{\mu}_n^* - \overline{X}_n)$, rather than that of $n^{1/2}(\hat{\mu}_n^* - \hat{\mu}_n)$, to approximate the distribution of $n^{1/2}(\hat{\mu}_n - \mu)$. This follows because (3) holds with $\hat{\mu}_{n_m}(\omega)$ replaced by $\overline{X}_{n_m}(\omega)$.

Next, we consider a parametric bootstrap. Let $\{X_i^*: i \leq n\}$ be iid with $X_i^* \sim N(\hat{\mu}_n, 1)$. The analysis of equation (3) is exactly the same as with the nonparametric bootstrap (except that the asymptotic distribution actually holds in finite samples as well). Thus, this parametric bootstrap is not consistent. An alternative parametric bootstrap takes $\{X_i^*: i \leq n\}$ to be iid with $X_i^* \sim N(\overline{X}_n, 1)$. In this case, the analysis of both (2) and (3) is exactly the same as with the nonparametric bootstrap (except that the asymptotic distribution in both equations holds in finite samples as well). Thus, this parametric bootstrap also is not consistent.

The above counterexample to the nonparametric bootstrap generalizes to any estimation problem in which the true parameter is on the boundary of the parameter space and the parameter space is not locally equal to (or locally approximated by) a linear subspace (as defined in Andrews (1999)). The primary case where the latter occurs is with linear or nonlinear equality constraints. If the parameter is on a boundary defined by linear or nonlinear inequality constraints or mixed inequality/equality constraints, then the bootstrap is not consistent. For brevity, we do not provide the details. See Andrews (1998, 1999) for general results providing the asymptotic distribution of extremum estimators, including maximum likelihood estimators, minimum distance estimators, etc., when the true parameter is on the boundary of the parameter space. Such results are needed to demonstrate the inconsistency of the bootstrap in more general cases than the simple example provided above.

Bickel and Freedman (1981, Sec. 6) list three conditions for the bootstrap distribution of a statistic to be consistent in iid contexts. The first is weak convergence of the statistic when $X_i \sim G$ for all distributions G in a neighborhood of the true distribution F. The second is uniform weak convergence over distributions G in a neighborhood of the true distribution F. The third is continuity of the mapping from the underlying distribution G to the asymptotic distribution of the statistic. Bickel and Freedman provide two counterexamples to the bootstrap (described in the Introduction) that violate the second condition, viz., uniformity. The counterexample given above violates the third condition, viz., continuity.

4. ALTERNATIVES TO THE BOOTSTRAP

We now suggest four methods for obtaining consistent estimators of the asymptotic distribution of the normalized maximum likelihood estimator, $n^{1/2}(\hat{\mu}_n - \mu)$, in the iid

$N(\mu, 1)$ counterexample given above. These methods are designed to be consistent whether or not the true parameter is on the boundary. The methods generalize to the problem of an arbitrary extremum estimator when the true parameter may be on the boundary of the parameter space; see Andrews (1999).

We note that none of the four methods consistently estimates the asymptotic distribution of $n^{1/2}(\hat{\mu}_n - \mu_n)$ when the true mean is of the form $\mu_n = \mu/n^{1/2}$ for some $\mu > 0$. In fact, it is impossible to consistently estimate the asymptotic distribution of $n^{1/2}(\hat{\mu}_n - \mu_n)$ (or of $n^{1/2}\hat{\mu}_n$) in this case. The reason is that the asymptotic distribution is $\max\{Z, -\mu\}$ (or $\max\{Z + \mu, 0\}$ respectively) and consistent estimation of this term would imply consistent estimation of μ, which is not possible when $\mu_n = \mu/n^{1/2}$.

The first method is based on tests. Let $\{\eta_n \colon n \geq 1\}$ be a sequence of positive random variables (possibly constants) that satisfies

$$P\left(\lim_{n\to\infty} \eta_n = 0 \text{ and } \lim_{n\to\infty}\inf \eta_n (n/(2\ln\ln n))^{1/2} > 1\right) = 1. \tag{4}$$

If $\hat{\mu}_n \leq \eta_n$, then we estimate the asymptotic distribution of $n^{1/2}(\hat{\mu}_n - \mu)$ to be $\max\{Z, 0\}$. Otherwise, we estimate the asymptotic distribution to be Z. (Note that the η_n's could be chosen to be the critical values for a sequence of one-sided tests of H_0: $\mu = 0$ versus H_1: $\mu > 0$ whose significance levels converge to zero as $n \to \infty$ at a rate such that (4) holds.)

This estimator of the asymptotic distribution is strongly consistent, because

$$\begin{aligned} &P\left(\limsup_{n\to\infty}(\hat{\mu}_n - \eta_n) \leq 0\right) \\ &= P\left(\limsup_{n\to\infty}\left(\max\left\{(2n\ln\ln n)^{-1/2}\sum_{i=1}^{n} X_i, 0\right\} - \eta_n (n/(2\ln\ln n))^{1/2}\right) \leq 0\right) \\ &= \begin{cases} 0 & \text{if } \mu > 0, \\ 1 & \text{if } \mu = 0, \end{cases} \end{aligned} \tag{5}$$

by the law of the iterated logarithm. Equation (5) also holds with the $\limsup_{n\to\infty}$ replaced by $\liminf_{n\to\infty}$.

This method of estimating the asymptotic distribution can be generalized to the case of an arbitrary extremum estimator with a parameter space that is defined by linear or nonlinear inequality constraints by specifying a criterion for each inequality constraint to assess whether it is binding or not. The method can be applied when the data are iid, as well as when the data exhibit temporal dependence, including stochastic and deterministic time trends. See Andrews (1999) for details.

The second method is a parametric bootstrap procedure in which the parameter estimator used to generate the bootstrap, $\tilde{\mu}_n$, shrinks to the boundary of the parameter space. This is an application of an idea of Beran (1997, Corollary 2.1). Define $\tilde{\mu}_n = \overline{X}_n 1(\overline{X}_n > \eta_n)$, where $\{\eta_n \colon n \geq 1\}$ are as above. Let $\{X_i^* \colon i \leq n\}$ be iid with $X_i^* \sim N(\tilde{\mu}_n, 1)$. By equation (5), when $\mu > 0$, $\tilde{\mu}_n = \overline{X}_n$ for n sufficiently large with probability one. And when $\mu = 0$, $\tilde{\mu}_n = 0$ for n sufficiently large with probability one. In consequence, this parametric bootstrap is asymptotically correct for all $\mu \geq 0$.

The third method is a subsample method introduced by Wu (1990) and extended by Politis and Romano (1994) (also see Bickel, Götze, and van Zwet (1997)) to cover cases where the statistic of interest has *some* asymptotic distribution, not necessarily normal, such as that which arises when the true parameter is on the boundary of the parameter space. The method is applicable in iid contexts, as well as in stationary time series contexts; see Politis and Romano (1994). A random subsampling variant of the procedure is also available; see Politis and Romano (1994, Sec. 2.2).

The method is as follows. Let $D_1, \ldots, D_{N_n}$ denote the $N_n = \binom{n}{m}$ subsets of size m of $\{X_i: i = 1, \ldots, n\}$, ordered in any fashion. Let $\hat{\mu}_{m,j}$ denote the statistic $\hat{\mu}_n$ computed using the data set D_j rather than $\{X_i: i = 1, \ldots, n\}$. The empirical distribution of $\{m^{1/2}(\hat{\mu}_{m,j} - \hat{\mu}_n): j = 1, \ldots, N_n\}$ is used to estimate the distribution of $n^{1/2}(\hat{\mu}_n - \mu)$. The empirical distribution function, denoted $L_n(x)$, is defined by

$$L_n(x) = \frac{1}{N_n} \sum_{j=1}^{N_n} 1\left(m^{1/2}(\hat{\mu}_{m,j} - \hat{\mu}_n) \leq x\right).$$

Provided $m \to \infty$ and $m/n \to 0$, $L_n(x) \to_p F_\mu(x)$ as $n \to \infty$ for all $x \in R$, where $F_\mu(x)$ denotes the asymptotic distribution function of $n^{1/2}(\hat{\mu}_n - \mu)$; see Politis and Romano (1994, Theorem 2.1). Under stronger conditions on m, the convergence holds with probability one.

The fourth method is a variant of the bootstrap, called a *rescaled* bootstrap, in which bootstrap samples of size m $(< n)$, rather than n, are employed. This method has been used previously as a means of fixing the bootstrap in the U-statistic counterexample of Bickel and Freedman (1981) by Bretagnolle (1983), in the largest order statistic example of Bickel and Freedman (1981) by Swanepoel (1986), and in the sample mean with infinite variance random variables counterexample of Babu (1984); see Arcones (1990), who attributes the idea to an unpublished paper of Athreya. See Bickel, Götze, and van Zwet (1997) for further applications and analysis of this method.

The idea is to use the bootstrap distribution of $m^{1/2}(\hat{\mu}_m^* - \hat{\mu}_n)$ to estimate the distribution of $n^{1/2}(\hat{\mu}_n - \mu)$, where $\hat{\mu}_m^* := \max\{\overline{X}_m^*, 0\}$, $\overline{X}_m^* := (1/m)\Sigma_{i=1,\ldots,m} X_i^*$, and $\{X_i^*: i \leq m\}$ are iid with $X_i^* \sim \hat{F}_n$. This variant of the bootstrap is consistent with probability one if $m \to \infty$ and $m(\ln\ln n)/n \to 0$ as $n \to \infty$. The reason is that

$$\begin{aligned}
(6) \qquad & m^{1/2}(\hat{\mu}_m^* - \hat{\mu}_n) \\
&= \max\left\{m^{1/2}\left(\overline{X}_m^* - \overline{X}_n\right) + m^{1/2}(\overline{X}_n - \mu), -m^{1/2}\mu\right\} - m^{1/2}(\hat{\mu}_n - \mu) \\
&= \max\left\{m^{1/2}\left(\overline{X}_m^* - \overline{X}_n\right) + o(1), -m^{1/2}\mu\right\} + o(1) \\
&\xrightarrow{d} \begin{cases} Z & \text{if } \mu > 0 \\ \max\{Z, 0\} & \text{if } \mu = 0 \end{cases} \quad \text{as } n \to \infty \text{ conditional on } \{\hat{F}_n: n \geq 1\},
\end{aligned}$$

where the second equality holds with probability one by the law of the iterated logarithm and the convergence in distribution holds by the central limit theorem for triangular arrays of row-wise iid random variables.

Lastly, we note that the reason the one-sided test described in the Introduction has correct null rejection rate is that if $\mu = 0$ and $\overline{X}_n < 0$, then $\hat{\mu}_n = 0$, $\hat{t}_{1-\alpha} \geq 0$, and the test does not reject; whereas if $\mu = 0$ and $\overline{X}_n \geq 0$, then $\hat{\mu}_n = \overline{X}_n$, $n^{1/2}(\hat{\mu}_n^* - \hat{\mu}_n) = \max\{n^{1/2}(\overline{X}_n^* - \overline{X}_n), -\overline{X}_n\}$, $\hat{t}_{1-\alpha}$ converges in probability to the standard normal $1 - \alpha$ quantile, and the test rejects H_0 with asymptotic probability α.

Cowles Foundation for Research in Economics, 30 Hillhouse Ave., P.O. Box 208281, New Haven, CT 06520-8281, U.S.A.

Manuscript received June, 1998; final revision received November, 1998.

REFERENCES

ANDREWS, D. W. K. (1997): "A Simple Counterexample to the Bootstrap," unpublished manuscript, Cowles Foundation for Research in Economics, Yale University.

——— (1998): "Estimation When a Parameter Is on a Boundary: Part II," unpublished manuscript, Cowles Foundation for Research in Economics, Yale University.

——— (1999): "Estimation When a Parameter Is on a Boundary," *Econometrica*, 67, 1341–1383.

ARCONES, M. (1990): "On the Asymptotic Theory of the Bootstrap," Ph.D. dissertation, The City University of New York.

ATHREYA, K. B. (1987): "Bootstrap of the Mean in the Infinite Variance Case," *Annals of Statistics*, 15, 724–731.

ATHREYA, K. B., AND J. FUKUCHI (1994): "Bootstrapping Extremes of I.I.D. Random Variables," *Proceedings of the Conference on Extreme Value Theory*, NIST.

BABU, J. (1984): "Bootstrapping Statistics with Linear Combinations of Chi-squares as Weak Limit," *Sankhya*, 46, 86–93.

BASAWA, I. V., A. K. MALLIK, W. P. MCCORMICK, J. H. REEVES, AND R. L. TAYLOR (1991): "Bootstrapping Unstable First-order Autoregressive Processes," *Annals of Statistics*, 19, 1098–1101.

BERAN, R. (1982): "Estimated Sampling Distributions: The Bootstrap and Competitors," *Annals of Statistics*, 10, 212–225.

——— (1997): "Diagnosing Bootstrap Success," *Annals of the Institute of Statistical Mathematics*, 49, 1–24.

BERAN, R., AND M. S. SRIVASTAVA (1985): "Bootstrap Tests and Confidence Regions for Functions of a Covariance Matrix," *Annals of Statistics*, 13, 95–115.

BICKEL, P. J., AND D. FREEDMAN (1981): "Some Asymptotic Theory for the Bootstrap," *Annals of Statistics*, 9, 1196–1217.

BICKEL, P. J., F. GÖTZE, AND W. R. VAN ZWET (1997): "Resampling Fewer than n Observations: Gains, Losses, and Remedies for Losses," *Statistica Sinica*, 7, 1–31.

BRETAGNOLLE, J. (1983): "Lois Limites du Bootstrap de Certaines Fonctionnelles," *Annales de L'institut Henri Poincaré*, Sec. B, 19, 281–296.

DATTA, S. (1995): "On a Modified Bootstrap for Certain Asymptotically Non-normal Statistics," *Statistics and Probability Letters*, 24, 91–98.

DEHEUVELS, P., D. MASON, AND G. SHORACK (1993): "Some Results on the Influence of Extremes on the Bootstrap," *Annales de l'Institut Henri Poincaré*, 29, 83–103.

DÜMBGEN, L. (1993): "On Nondifferentiable Functions and the Bootstrap," *Probability Theory and Related Fields*, 95, 125–140.

EFRON, B. (1979): "Bootstrap Methods: Another Look at the Jackknife," *Annals of Statistics*, 7, 1–26.

EFRON, B., AND R. TIBSHIRANI (1993): *An Introduction to the Bootstrap*. New York: Chapman and Hall.

FREEDMAN, D. (1981): "Bootstrapping Regression Models," *Annals of Statistics*, 9, 1218–1228.

HALL, P. (1992): *The Bootstrap and Edgeworth Expansion*. New York: Springer-Verlag.

HALL, P., AND J. L. HOROWITZ (1996): "Bootstrap Critical Values for Tests Based on Generalized-Method-of-Moment Estimators," *Econometrica*, 64, 891–916.

POLITIS, D. N., AND J. P. ROMANO (1994): "Large Sample Confidence Regions Based on Subsamples under Minimal Assumptions," *Annals of Statistics*, 22, 2031–2050.

PUTTER, H., AND W. R. VAN ZWET (1996): "Resampling: Consistency of Substitution Estimators," *Annals of Statistics*, 24, 2297–2318.

ROMANO, J. (1988): "Bootstrapping the Mode," *Annals of the Institute of Statistical Mathematics*, 40, 565–586.

SRIRAM, T. N. (1993): "Invalidity of the Bootstrap for Critical Branching Processes with Immigration," *Annals of Statistics*, 22, 1013–1023.

SWANEPOEL, J. W. H. (1986): "A Note on Proving That the (Modified) Bootstrap Works," *Communications in Statistics—Theory and Methods Section*, 15, 3193–3203.

WU, C. F. J. (1990): "On the Asymptotic Properties of the Jackknife Histogram," *Annals of Statistics*, 18, 1438–1452.