

Double-bootstrap methods that use a single double-bootstrap simulation

BY JINYUAN CHANG

*School of Statistics, Southwestern University of Finance and Economics, Chengdu,
Sichuan 611130, China*

jinyuan.chang@unimelb.edu.au

AND PETER HALL

*Department of Mathematics and Statistics, The University of Melbourne, Parkville,
Victoria 3010, Australia*

halpstat@ms.unimelb.edu.au

SUMMARY

We show that, when the double bootstrap is used to improve performance of bootstrap methods for bias correction, techniques based on using a single double-bootstrap sample for each single-bootstrap sample can produce third-order accuracy for much less computational expense than is required by conventional double-bootstrap methods. However, this improved level of performance is not available for the single double-bootstrap methods that have been suggested to construct confidence intervals or distribution estimators.

Some key words: Bias correction; Bias estimation; Confidence intervals; Distribution estimation; Edgeworth expansion; Second-order correctness; Third-order correctness.

1. INTRODUCTION

Double-bootstrap methods that use a single simulation at the second bootstrap level have been studied in at least one context for more than a decade. An early contribution was made by [White \(2000\)](#), although in the setting of diagnosing the overuse of a dataset, rather than speeding up Monte Carlo simulation for general applications of the bootstrap. [Davidson & Mackinnon \(2002\)](#), the same authors in a number of subsequent papers accessible via an unpublished 2006 technical report available from the second author, and [Davidson & Mackinnon \(2007\)](#) introduced the concept independently and explored its applications. [Giacomini et al. \(2013\)](#) christened the technique the warp-speed double-bootstrap method, nomenclature that we shall use here, too. [Giacomini et al. \(2013\)](#) demonstrated that this approach is asymptotically consistent. All this work is for distribution estimation and its application to constructing confidence intervals and hypothesis tests.

In statistics the conventional double bootstrap is used in two main classes of problems: (i) to improve the effectiveness of bias correction and (ii) to improve the coverage accuracy of confidence intervals. In problem (i), an application of the double bootstrap reduces the order of magnitude of bias by the factor $O(n^{-1})$, and in problem (ii) it reduces coverage error by the factor $O(n^{-1/2})$ for one-sided confidence intervals, and $O(n^{-1})$ for two-sided intervals. In the setting of problem (i), it is not clear whether there exists a version of warp-speed methodology for bias correction, and whether, should it exist, it successfully reduces the order of magnitude of bias.

Call these questions 1 and 2, respectively. In problem (ii), it is unclear whether the warp-speed double bootstrap is as effective as the conventional double bootstrap, in the sense of offering the above levels of improved accuracy; we shall refer to this as question 3. In the present paper we show that the answers to questions 1 and 2 are positive, but that the answer to question 3 is negative. In particular, the warp-speed bootstrap does not reduce the order of magnitude of coverage error of a confidence interval.

There is an extensive literature on conventional double-bootstrap methods, particularly in the context of improving the coverage accuracy of single-bootstrap methods. The first mention of the double bootstrap in this setting apparently was by Hall (1986), followed quickly by contributions of Beran (1987, 1988); see also Hall & Martin (1988). The approach suggested by Hall (1992, Ch. 3) allows general multiple bootstrap methods to be developed together, so that different settings do not require separate treatment. However, details of properties of the technique seem to be very problem-specific. Efron (1983) was the first to use the double bootstrap in any setting; in that paper his work was in the context of estimating the error rate of classifiers. Research on optimizing the trade-off between the numbers of simulations in the first and second stages of the conventional double bootstrap, in the context of distribution estimation and constructing confidence intervals, includes that of Booth & Hall (1994), Booth & Presnell (1998) and Lee & Young (1999).

It has become conventional to assess performance of the bootstrap in terms of Edgeworth expansions, not least because that approach enables theoretical properties to be developed in the very broad context addressed by Bhattacharya & Ghosh (1978). The resulting approximations are valid, in absolute rather than relative terms, uniformly in the tails. An alternative approach, based on large deviation probabilities, is valid in relative terms; see, e.g., Hall (1990). However, it requires either more stringent assumptions or specialized methods that, at least at present, are not available in the context of the models used by Bhattacharya & Ghosh (1978). In the setting of absolute rather than relative accuracy, arbitrarily far out into the tails, the results in this paper take the result of consistency, demonstrated by Giacomini et al. (2013), much further.

2. MODEL AND METHODOLOGY FOR BIAS CORRECTION

2.1. Model

Let $\theta = f(\mu)$ be a parameter expressible as a known function, f , of a p -variate mean, μ , and let \bar{X} denote an unbiased estimator of $\mu = (\mu_1, \dots, \mu_p)^\top$. Our estimator of θ is the same function of a sample mean, \bar{X} :

$$\hat{\theta} = f(\bar{X}). \quad (1)$$

The smooth function f maps a point x in p -variate Euclidean space to a point on the real line. We do not insist that \bar{X} be a mean of n , say, independent and identically distributed random p -vectors, since it might be the case that $\bar{X} = (\bar{X}_1, \dots, \bar{X}_p)^\top$, with

$$\bar{X}_j = \frac{1}{n_j} \sum_{i=1}^{n_j} X_{ji},$$

where X_{ji} , for $i = 1, \dots, n_j$, are independent for each i , $E(X_{ji}) = \mu_j$ for each j , and the n_j are not all equal. Nevertheless, in mathematical terms we shall assume that the n_j are all functions of an integer parameter n , and that each $n_j \asymp n$; that is, each ratio n_j/n is bounded away from zero and infinity as $n \rightarrow \infty$.

These issues are related to dependence relationships among the random variables X_{ji} , which should be reflected in resampling methodology. In our theoretical work we shall suppose the following condition.

Condition 1. Either (i) each $n_j = n$ and the vectors $(X_{1i}, \dots, X_{pi})^T$, for $i \geq 1$, are independent and identically distributed; or (ii) the X_{ji} are totally independent, for $i = 1, \dots, n_j$ and $j = 1, \dots, p$, and in this case, for each $j \in \{1, \dots, p\}$ the variables X_{j1}, X_{j2}, \dots are identically distributed, and $n_j \asymp n$.

Conditions 1(i) and (ii) can be generalized, for example to hybrid cases where, for positive integers p_1, \dots, p_r that satisfy $\sum_{j=1}^r p_j = p$, and defining $q_j = \sum_{k=1}^j p_k$, the vectors $V_{ji} = (X_{q_{j+1},i}, \dots, X_{q_{j+i},i})^T$, for $j = 0, \dots, r - 1$ and $i \geq 1$, are completely independent, and for each j the vectors V_{ji} , for $i \geq 1$, are identically distributed. Bootstrap methods that reflect these properties can be constructed readily, and theory providing authoritative support in this setting can be developed, but for the sake of brevity, in our theoretical work we shall restrict attention to cases where Condition 1 holds.

2.2. Bias correction

Bias-corrected estimators of θ , based on the conventional bootstrap and the double bootstrap, respectively, are given by

$$\hat{\theta}^{bc} = 2\hat{\theta} - E(\hat{\theta}^* | \mathcal{X}), \quad \hat{\theta}^{bcc} = 3\hat{\theta} - 3E(\hat{\theta}^* | \mathcal{X}) + E(\hat{\theta}^{**} | \mathcal{X}). \quad (2)$$

Here $\mathcal{X} = \{X_{ji} : i = 1, \dots, n_j, j = 1, \dots, p\}$ denotes the original dataset, $\hat{\theta}^*$ is the version of $\hat{\theta}$ computed from a resample \mathcal{X}^* drawn randomly, with replacement, from \mathcal{X} , in a manner that reflects appropriately the dependence structure, and $\hat{\theta}^{**}$ is the version of $\hat{\theta}$ computed from \mathcal{X}^{**} , which in turn is drawn randomly with replacement from \mathcal{X}^* , again reflecting dependence.

Monte Carlo approximations to the quantities $\hat{\theta}^{bc}$ and $\hat{\theta}^{bcc}$ in (2) are given respectively by

$$\tilde{\theta}^{bc} = 2\hat{\theta} - \frac{1}{B} \sum_{b=1}^B \hat{\theta}_b^*, \quad \tilde{\theta}^{bcc} = 3\hat{\theta} - \frac{3}{B} \sum_{b=1}^B \hat{\theta}_b^* + \frac{1}{BC} \sum_{b=1}^B \sum_{c=1}^C \hat{\theta}_{bc}^{**}, \quad (3)$$

where $\hat{\theta}_b^*$ denotes the b th out of B independent and identically distributed, conditional on \mathcal{X} , versions of $\hat{\theta}^*$, computed from respective resamples \mathcal{X}_b^* drawn by sampling randomly, with replacement, from the data in \mathcal{X} , and $\hat{\theta}_{bc}^{**}$ is the c th out of C independent and identically distributed, conditional on \mathcal{X} and \mathcal{X}^* , versions of $\hat{\theta}^{**}$, and is computed from a resample \mathcal{X}_{bc}^{**} drawn by sampling randomly, with replacement, from \mathcal{X}_b^* .

2.3. Bootstrap algorithms

Reflecting the model at (1), we can express $\hat{\theta}_b^*$ and $\hat{\theta}_{bc}^{**}$ in (3) as $\hat{\theta}_b^* = f(\bar{X}_b^*)$ and $\hat{\theta}_{bc}^{**} = f(\bar{X}_{bc}^{**})$, where $\bar{X}_b^* = (\bar{X}_{b1}^*, \dots, \bar{X}_{bp}^*)^T$, $\bar{X}_{bc}^{**} = (\bar{X}_{bc1}^{**}, \dots, \bar{X}_{bcp}^{**})^T$, \bar{X}_{bj}^* denotes the mean of data in the resample $\mathcal{X}_{bj}^* = \{X_{bj1}^*, \dots, X_{bjn_j}^*\}$, \bar{X}_{bcj}^{**} is the mean of data in the re-resample $\mathcal{X}_{bcj}^{**} = \{X_{bcj1}^{**}, \dots, X_{bcjn_j}^{**}\}$ drawn by sampling with replacement from \mathcal{X}_{bj}^* , the resampling operations at the first bootstrap level are undertaken by resampling the vectors $X_i = (X_{1i}, \dots, X_{pi})^T$ randomly, with replacement, if Condition 1(i) holds, or by resampling the X_{ji} randomly and completely independently, conditional on \mathcal{X} and with replacement, if Condition 1(ii) obtains, and resampling at the second bootstrap level is undertaken analogously.

2.4. *Main conclusions*

Theorem 1 shows that if $C \rightarrow \infty$, no matter how slowly, as n and B diverge, then the asymptotic distribution of the Monte Carlo simulation error incurred when constructing $\tilde{\theta}^{\text{bcc}}$ at (3) is the same as it would be if $C = \infty$. In particular, not only is the error of order $(nB)^{-1/2}$, the large-sample limiting distribution of the relevant asymptotically normal random variable, which has standard deviation proportional to $(nB)^{-1/2}$, and which describes in relative detail the accuracy of Monte Carlo bootstrap simulation, is identical to the limiting distribution that would arise if $C = \infty$.

Moreover, if C is held fixed then the order of magnitude, $(nB)^{-1/2}$, remains unchanged, but the standard deviation of the large-sample limiting distribution referred to above changes by a constant factor. This result is critical. It demonstrates the relatively small gains that are to be achieved by taking C to be large, and argues in favour of taking $C = 1$, for example. This is the analogue, for bias correction, of the warp-speed bootstrap for distribution estimation when constructing confidence intervals.

Therefore the order of magnitude of Monte Carlo simulation error in $\tilde{\theta}^{\text{bcc}}$ is unchanged even if C is held fixed. Incidentally, the order of magnitude, $(nB)^{-1/2}$, should be compared with that of the uncorrected bias that remains after applying the bias correction that leads to $\tilde{\theta}^{\text{bcc}}$; it is n^{-3} . Therefore, unless B is of order n^5 or larger, for the regular bootstrap, the orders of magnitude involving B , discussed above, dominate the error in the bias correction.

3. MODEL AND METHODOLOGY FOR CONSTRUCTING CONFIDENCE INTERVALS

3.1. *Model*

As in § 2.1 we shall assume that the parameter θ can be represented as $f(\mu)$, where the function $f: \mathbb{R}^p \rightarrow \mathbb{R}$ is known, and $\mu = E(X)$ is an unknown p -vector of parameters, estimated by $\bar{X} = n^{-1} \sum_{i=1}^n X_i$ where $\mathcal{X} = \{X_1, \dots, X_n\}$ is a random sample of data vectors. Here and below we use Condition 1(i) for the data, but only minor modifications are needed if Condition 1(ii) is employed instead.

In such cases, provided that f is sufficiently smooth and $\hat{\theta}$ is given by (1), the asymptotic variance, σ_n^2 , of $\hat{\theta}$ is estimated root- n consistently by $n^{-1} \hat{\sigma}^2$, where

$$\hat{\sigma}^2 = \sum_{j_1=1}^p \sum_{j_2=1}^p f_{j_1 j_2}(\bar{X}) \frac{1}{n} \sum_{i=1}^n (X_{j_1 i} - \bar{X}_{j_1})(X_{j_2 i} - \bar{X}_{j_2}).$$

Here, given a p -vector $x = (x_1, \dots, x_p)^\top$, and integers j_1, \dots, j_r between 1 and p ; and assuming that f has r well-defined derivatives with respect to each variable; we put

$$f_{j_1 \dots j_r}(x) = (\partial/\partial x_{j_1}) \cdots (\partial/\partial x_{j_r}) f(x).$$

The above definitions of $\hat{\theta}$ and $\hat{\sigma}$ are used in (4) below.

3.2. *Bootstrap algorithms*

Let R , referred to as the root by [Giacomini et al. \(2013\)](#), be given by either of the formulae

$$R = n^{1/2} (\hat{\theta} - \theta), \quad R = n^{1/2} (\hat{\theta} - \theta) / \hat{\sigma}. \quad (4)$$

Here $\hat{\theta}$ and $\hat{\sigma}$ are estimators of parameters θ and σ computed from the random sample \mathcal{X} , and σ^2 denotes the asymptotic variance of $n^{1/2} \hat{\theta}$. The warp-speed bootstrap of [Giacomini et al. \(2013\)](#),

closely related to suggestions by [White \(2000\)](#) and [Davidson & Mackinnon \(2002, 2007\)](#), can be defined as follows.

As in § 2, let \mathcal{X}_b^* , for $b = 1, \dots, B$, be drawn randomly, with replacement, from \mathcal{X} , and be independent conditional on \mathcal{X} . Draw \mathcal{X}_b^{**} , denoting a single double-bootstrap resample, by sampling randomly, with replacement, from \mathcal{X}_b^* for $b = 1, \dots, B$, in such a manner that these re-resamples are independent, conditional on \mathcal{X} and $\mathcal{X}_1^*, \dots, \mathcal{X}_B^*$. In the context of § 2, \mathcal{X}_b^{**} would be one of the resamples $\mathcal{X}_{b1}^{**}, \dots, \mathcal{X}_{bC}^{**}$ that were drawn by resampling from \mathcal{X}_b^* , but on the present occasion we require only one of these resamples.

Let $\hat{\theta}_b^*$ and $\hat{\theta}_b^{**}$ denote the versions of $\hat{\theta}$ computed from \mathcal{X}_b^* and \mathcal{X}_b^{**} , respectively, instead of \mathcal{X} , and write $\hat{\sigma}_b^*$ and $\hat{\sigma}_b^{**}$ for the corresponding versions of $\hat{\sigma}$. If R is given by one of the formulae at (4), define

$$R_b^* = n^{1/2} (\hat{\theta}_b^* - \hat{\theta}), \quad R_b^* = n^{1/2} (\hat{\theta}_b^* - \hat{\theta}) / \hat{\sigma}_b^*, \tag{5}$$

$$R_b^{**} = n^{1/2} (\hat{\theta}_b^{**} - \hat{\theta}_b^*), \quad R_b^{**} = n^{1/2} (\hat{\theta}_b^{**} - \hat{\theta}_b^*) / \hat{\sigma}_b^{**}, \tag{6}$$

in the respective cases, and put

$$\hat{F}_B^*(x) = \frac{1}{B} \sum_{b=1}^B I(R_b^* \leq x), \quad \tilde{F}_B^*(x) = \frac{1}{B} \sum_{b=1}^B I(R_b^{**} \leq x). \tag{7}$$

Then \hat{F}_B^* is the conventional single-bootstrap Monte Carlo approximation to the distribution function F of R , and the limit of \hat{F}_B^* , as $B \rightarrow \infty$, is the conventional single-bootstrap approximation to F . The function \tilde{F}_B^* is a short-cut, warp-speed, double-bootstrap approximation to F .

Given a nominal coverage level $\alpha \in (0, 1)$ of a confidence interval, define $x = \hat{x}_\alpha^*$ to be the solution of the equation $\tilde{F}_B^*(x) = \alpha$, and similarly let \hat{x}_α be the solution of $\hat{F}_B^*(x) = \alpha$. If R is given by either of the expressions in (4), consider the respective confidence intervals,

$$\mathcal{I}_{b\alpha}^* = (\hat{\theta}_b^* - n^{-1/2} \hat{x}_\alpha^*, \infty), \quad \mathcal{I}_{b\alpha} = (\hat{\theta}_b^* - n^{-1/2} \hat{\sigma}_b^* \hat{x}_\alpha^*, \infty), \tag{8}$$

which are bootstrap versions of the respective intervals

$$\mathcal{I}_\alpha = (\hat{\theta} - n^{-1/2} \hat{x}_\alpha, \infty), \quad \mathcal{I}_\alpha = (\hat{\theta} - n^{-1/2} \hat{\sigma} \hat{x}_\alpha, \infty).$$

In either case, our estimator of the probability p_α that the interval \mathcal{I}_α covers θ is given by

$$\hat{p}_{B\alpha} = \frac{1}{B} \sum_{b=1}^B I(\hat{\theta} \in \mathcal{I}_{b\alpha}^*).$$

We take the final interval to be $\mathcal{I}_{\hat{\beta}_{B\alpha}}$, where $\beta = \hat{\beta}_{B\alpha}$ denotes the solution of $\hat{p}_{B\beta} = \alpha$.

Earlier warp-speed bootstrap methodology is a little ambiguous in the percentile- t setting, i.e., in the context of the second definition in each of (4)–(6), where the technique is not completely clear from the algorithms of [White \(2000\)](#), [Davidson & Mackinnon \(2002\)](#) and [Giacomini et al. \(2013, pp. 570–1\)](#). In particular it is unclear from [Giacomini et al. \(2013\)](#) when, or whether, the estimator $\hat{\sigma}$ should be replaced by its single- or double-bootstrap forms, $\hat{\sigma}^*$ and $\hat{\sigma}^{**}$, for example in (5)–(8). The choices we have made are appropriate, however, and in particular the algorithm would not be second-order accurate, or third-order accurate in the case of the double bootstrap, if we were to use simply $\hat{\sigma}$ in those instances.

In § 5.2 we shall show that in the percentile- t case, using the case $B = \infty$ as a benchmark, the approach suggested above produces quantile estimators that are identical to those obtained using the standard single-bootstrap method, up to an error of order $n^{-3/2}$. In particular, they do not reduce the $O(n^{-1})$ coverage error of single-bootstrap methods. Similar results hold for percentile-method bootstrap procedures.

4. NUMERICAL PROPERTIES

4.1. Bias correction

Here we report the results of a simulation study comparing the performances of six different bootstrap methods for bias correction: the single bootstrap, the conventional double bootstrap, and the suggested alternative method involving only $C = 1, 2, 5$ or 10 double-bootstrap replications. The data were of two types, either the exponential distribution, with density $2^{-1}e^{-x/2}$ on the positive half-line, or the log-normal distribution. These two distributions both have nonzero skewness and nonzero kurtosis, making them challenging for the bootstrap. The parameter of interest also took two forms, both of them nonlinear: either $\theta = f(\mu) = \mu^3$ or $\theta = \sin(\mu)$, where μ was the population mean. In such cases there is a term with order n^{-2} in the bias expansion, which cannot be eliminated by the single bootstrap but can be removed by the double bootstrap. This is reflected in our simulation results, which show that the double bootstrap provides better bias correction than the single bootstrap method.

Sample size, n , was chosen in steps of 20 between 20 and 80; the number of simulations, B , in the first bootstrap step was set equal to n^2 , for each of the bootstrap methods; and the number of simulations, C , for the second bootstrap step in the conventional double bootstrap was taken to be the integer part of $10 B^{1/2}$, which we write as $\lfloor 10 B^{1/2} \rfloor$. The choice of $B^{1/2}$ here was suggested by Booth & Hall (1994) in the context of confidence intervals, and gives an expression for C that is orders of magnitude larger than obtained using relatively small, fixed C . For example, when $n = 20$ the value of $C = \lfloor 10 B^{1/2} \rfloor$ is between 20 and 200 times the values $C = 1, 2, 5$ or 10 used to simulate the alternative approach to double-bootstrap methods; when $n = 80$ the respective factors are 80 to 800.

From equation (3),

$$\frac{1}{B} \sum_{b=1}^B \hat{\theta}_b^* - \hat{\theta}, \quad \frac{3}{B} \sum_{b=1}^B \hat{\theta}_b^* - \frac{1}{BC} \sum_{b=1}^B \sum_{c=1}^C \hat{\theta}_{bc}^{**} - 2\hat{\theta},$$

provide the estimates of the true bias of $\hat{\theta}$, i.e., $E(\hat{\theta}) - \theta$, via the single and double bootstraps, respectively. Empirical approximations to bias, computed by averaging over the results of 5000 Monte Carlo trials in each case, are reported in Tables 1–2 in the Supplementary Material, and the ratios of such approximations and true bias are graphed in Fig. 1. For the values of B used in our analysis, there is little to choose between performance when using $C = 1$ and $C = \lfloor 10 B^{1/2} \rfloor$.

4.2. Confidence intervals

In this section we illustrate the coverage performance of bootstrap confidence intervals, with nominal coverage 0.9, for the population means of the two distributions considered in § 4.1, i.e., the exponential and log-normal distributions. Sample size n was taken equal to 20 and 40 in each case; B was increased from 200 to 700 in steps of 100, as indicated on the horizontal axis of each panel; and one-sided and two-sided equal-tailed bootstrap confidence intervals were considered, each using either the percentile or percentile- t bootstrap, implemented via the single bootstrap,

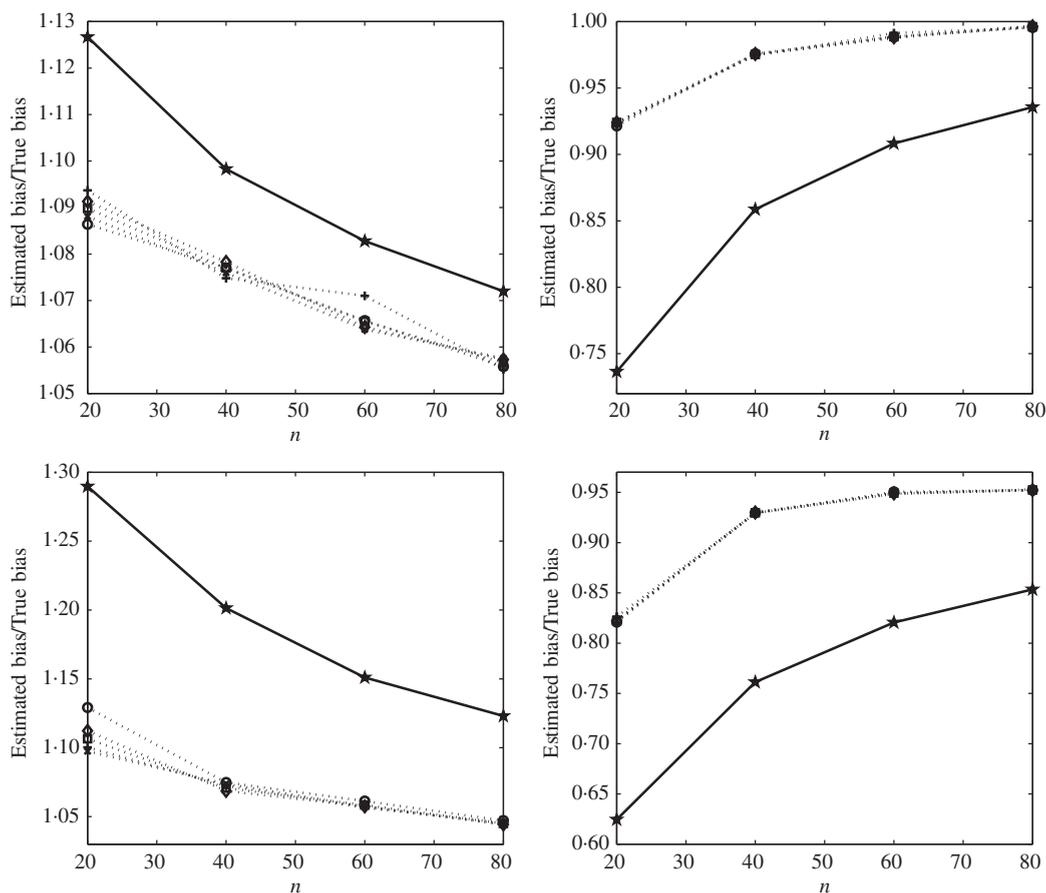


Fig. 1. Performance of bootstrap methods for bias correction. First and second rows show results for the exponential distribution, and the log-normal distribution, respectively; left- and right-hand panels show results for $\theta = \mu^3$ and $\theta = \sin(\mu)$, respectively. In each panel the graphs represent single-bootstrap method (—★—) and conventional double-bootstrap methods with $C = 1$ (···+···), $C = 2$ (···○···), $C = 5$ (···×···), $C = 10$ (···◇···) and $C = \lfloor 10 B^{1/2} \rfloor$ (···□···), respectively.

the conventional double bootstrap, $C = \lfloor 10 B^{1/2} \rfloor$, and the warp-speed bootstrap, i.e., the double bootstrap with $C = 1$. This choice of C was suggested by Lee & Young (1999). To provide a perspective different from that in § 4.1, in the present section we graph coverage as a function of B for fixed n , rather than as a function of n for fixed B as in § 4.1. Results in the two settings can of course be expressed in the same way; the conclusions do not alter.

Results for sample size $n = 20$, with each point on each graph based on 5000 Monte Carlo simulations, are presented in Fig. 2. For each confidence interval type, the conventional double-bootstrap method gives greater coverage accuracy than the single-bootstrap and warp-speed bootstrap. Results for $n = 40$ are similar, and are reported in the Supplementary Material.

5. THEORETICAL PROPERTIES

5.1. Bias correction

Our main regularity condition, in addition to the model assumptions (1) and Condition 1, is as follows.

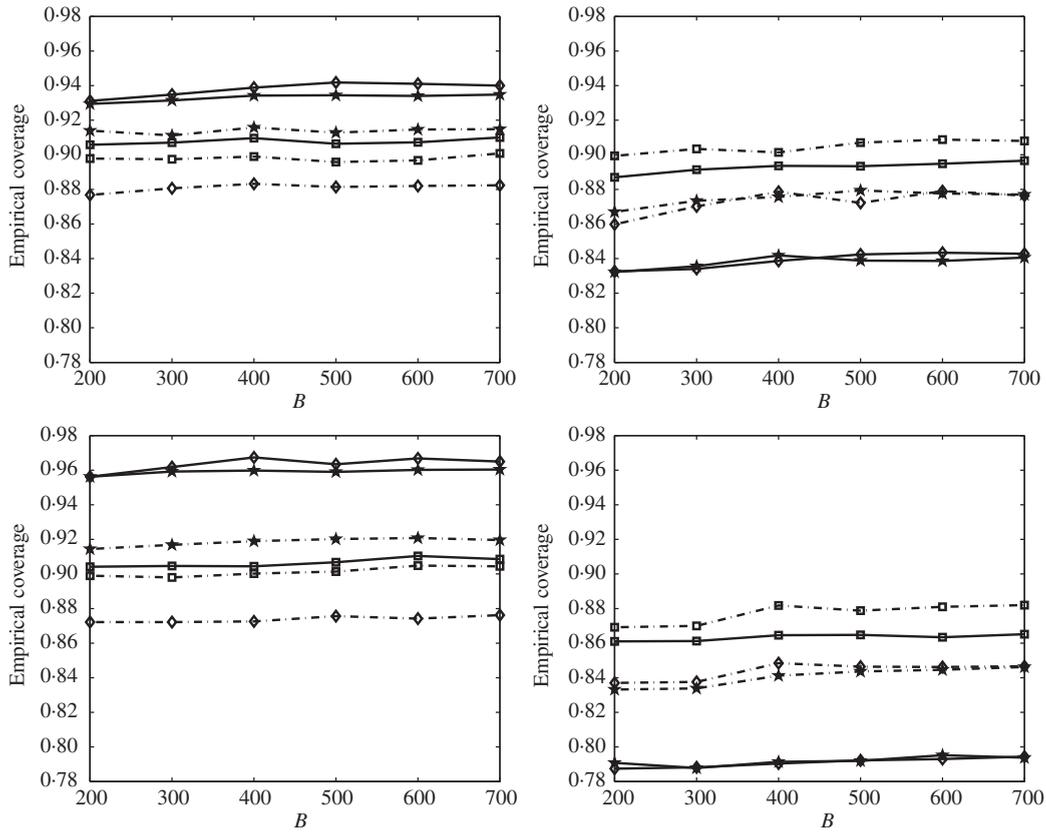


Fig. 2. Performance of bootstrap methods for confidence intervals when $n = 20$. First and second rows show results for the exponential distribution, and the log-normal distribution, respectively; left- and right-hand panels show results for one-sided and two-sided equal-tailed confidence intervals, respectively. In each panel the graphs represent single-bootstrap percentile ($- \star -$), single-bootstrap percentile- t ($- \cdot \star \cdot -$), conventional double-bootstrap percentile ($- \square -$), conventional double-bootstrap percentile- t ($- \cdot \square \cdot -$), warp-speed percentile ($- \diamond -$) and warp-speed percentile- t methods ($- \cdot \diamond \cdot -$).

Condition 2. (i) The function $f(x)$ is differentiable six times with respect to any combination of the p components of x ; and those derivatives, as well as f itself, are uniformly bounded; and (ii) the data X_{ji} have at least six finite moments, and $E(X_{ji}^6)$ is bounded uniformly in i and j .

Condition 2 can be generalized, but, for example, if we relax significantly the condition of boundedness of f and its derivatives, in Condition 2(i), then we need to strengthen the assumption about tails of the distributions of the X_{ji} , in Condition 2(ii). We shall define

$$\tau^2 = E \left[\left\{ \sum_{j=1}^p (X_{j1} - \mu_j) f_j(\mu) \right\}^2 \right].$$

In Theorem 1, we decompose the bias-corrected estimators $\tilde{\theta}^{bc}$, based on the single bootstrap, and $\tilde{\theta}^{bcc}$, based on the double bootstrap, as follows:

$$\tilde{\theta}^{bc} = U^{bc} + V^{bc}, \quad \tilde{\theta}^{bcc} = U^{bcc} + V^{bcc}. \tag{9}$$

Here U^{bc} and U^{bcc} are the ideal versions of $\tilde{\theta}^{bc}$ and $\tilde{\theta}^{bcc}$, respectively, that we would obtain if we were to do an infinite number of simulations, i.e., if we were to take $B = C = \infty$; and V^{bc} and V^{bcc} denote error terms arising from doing only a finite number of Monte Carlo simulations. Part (d) of Theorem 1 shows that the error terms V^{bc} in the case of the single bootstrap, and V^{bcc} for the double bootstrap, both equal $O_p\{(nB)^{-1/2}\}$, and that this is the exact order, regardless of the selection of C in the second bootstrap stage. Although the Monte Carlo error terms in the single bootstrap and the double bootstrap share the same convergence rate, equations (10) show that the double bootstrap provides a higher degree of accuracy, in terms of bias correction, than the single bootstrap if we take $B = C = \infty$. Part (d) also implies that if B is sufficiently large, or more precisely if $n^5 = O(B)$, then the Monte Carlo error is of the same order as, or an order smaller than, the deterministic remainders in (10). These are the main theoretical findings of Theorem 1.

THEOREM 1. *Assume that the data are generated according to either of the models at Condition 1, that Condition 2 holds, and that $B = B(n) \rightarrow \infty$ as $n \rightarrow \infty$. Then: (a) equations (9) hold, where U^{bc} and U^{bcc} are functions of \mathcal{X} alone, and in particular do not involve \mathcal{X}^* or \mathcal{X}^{**} , and satisfy*

$$E(U^{bc}) = \theta + O(n^{-2}), \quad E(U^{bcc}) = \theta + O(n^{-3}); \tag{10}$$

and V^{bc} and V^{bcc} are functions of both \mathcal{X} and \mathcal{X}^ , and also of \mathcal{X}^{**} , in the case of V^{bcc} , and satisfy $E(V^{bc} | \mathcal{X}) = E(V^{bcc} | \mathcal{X}) = 0$. (b) Both U^{bc} and U^{bcc} equal $\hat{\theta} + O_p(n^{-1})$, and both satisfy the same central limit theorem as $\hat{\theta}$. (c) In particular, both U^{bc} and U^{bcc} are asymptotically normally distributed with mean θ and a variance, σ_n^2 say, which has the property that $n\sigma_n^2$ is bounded as $n \rightarrow \infty$. (d) Conditional on \mathcal{X} , V^{bc} and V^{bcc} are asymptotically normally distributed with zero means and variances of size $(nB)^{-1}$, and if $C = C(n) \rightarrow \infty$ as $n \rightarrow \infty$ then the ratio of the variances converges to 1 as n diverges. In the case of Condition 1(i) the asymptotic variances of V^{bc} and V^{bcc} , both conditional on \mathcal{X} and unconditionally, are $(Bn)^{-1} \tau^2$ and $(4 + C^{-1})(Bn)^{-1} \tau^2$, respectively.*

In connection with part (d) it can be shown that, if C diverges, no matter how slowly, as n increases, the asymptotic distribution of the error is the same as it would be if $C = \infty$. If σ_n is as in part (c) then, under Condition 1(i), there exists a positive constant c such that $n\sigma_n^2 = c + o(1)$ as $n \rightarrow \infty$. However, this is not necessarily correct under Condition 1(ii), since in that setting we do not require the ratios n_j/n to converge. In the context of Condition 1(i), formulae for U^{bc} and U^{bcc} are given in the Supplementary Material.

The orders of magnitude of the remainders in (10) are exact when skewness and kurtosis are nonzero. It follows from part (b) of Theorem 1 that, in the case $B = C = \infty$, $\tilde{\theta}^{bc}$ and $\tilde{\theta}^{bcc}$ satisfy identical central limit theorems, and in particular both have the same asymptotic variances.

5.2. Distribution estimation and confidence interval construction

We shall assume that X , which represents a generic p -vector $X_i = (X_{1i}, \dots, X_{pi})^T$, where $i = 1, \dots, n$ and Condition 1(i) holds, satisfies the following multivariate version of Cramér’s continuity condition (Hall, 1992):

$$\limsup_{\|t\| \rightarrow \infty} |E\{\exp(it^T X)\}| < 1. \tag{11}$$

On this occasion, i denotes $\sqrt{-1}$. For brevity we shall treat in detail only the percentile- t case, evidenced by the second formula in each of (4)–(6), and discuss the percentile method briefly below Theorem 2.

Let Φ and ϕ denote the standard normal distribution and density functions, respectively. Assume that an unknown scalar parameter θ can be written as $\theta = f(\mu)$, where $\mu = E(X)$, and that our estimator of θ is $\hat{\theta} = f(\bar{X})$, as at (1), where $\bar{X} = n^{-1} \sum_{i=1}^n X_i$. Methods of [Bhattacharya & Ghosh \(1978\)](#) can be used to prove that, under conventional assumptions such as those in Theorem 2 below,

$$\begin{aligned} G(x) &\equiv \text{pr}\{n^{1/2}(\hat{\theta} - \theta)/\hat{\sigma} \leq x\} \\ &= \Phi(x) + \sum_{j=1}^3 n^{-j/2} Q_j(x) \phi(x) + n^{-2} A_n(x), \end{aligned} \quad (12)$$

where Q_j is a polynomial of degree $3j - 1$, and is an even or odd function according as j is odd or even, respectively; and the remainder $A_n(x)$ satisfies

$$\sup_{n \geq 1} \sup_{-\infty < x < \infty} |A_n(x)| < \infty.$$

The coefficients of Q_j are rational polynomials in moments of the distribution of X .

For simplicity in this section we take $B = \infty$, which is the ideal case where there is no error generated from Monte Carlo approximation. Inverting the Edgeworth expansion at (12) we obtain a Cornish–Fisher expansion:

$$x_\alpha \equiv G^{-1}(\alpha) = z_\alpha + n^{-1/2} Q_1^{\text{cf}}(z_\alpha) + n^{-1} Q_2^{\text{cf}}(z_\alpha) + n^{-3/2} Q_3^{\text{cf}}(z_\alpha) + O(n^{-2}), \quad (13)$$

where $z_\alpha = \Phi^{-1}(\alpha)$, the functions Q_1^{cf} , Q_2^{cf} and Q_3^{cf} are Cornish–Fisher polynomials and for example are given by $Q_1^{\text{cf}} = -Q_1$ and $Q_2^{\text{cf}}(x) = Q_1(x) Q_1'(x) - \frac{1}{2} x Q_1(x)^2 - Q_2(x)$, and the remainder in (13) is of the stated order, uniformly in $\alpha \in [\alpha_1, \alpha_2]$, whenever $0 < \alpha_1 < \alpha_2 < 1$.

The conventional percentile- t bootstrap estimator of G is \hat{G} , defined by

$$\hat{G}(x) = \text{pr}\{n^{1/2}(\hat{\theta}^* - \hat{\theta})/\hat{\sigma}^* \leq x \mid \mathcal{X}\}$$

and satisfying an empirical version of the Edgeworth expansion at (13):

$$\hat{x}_\alpha \equiv \hat{G}^{-1}(\alpha) = z_\alpha + n^{-1/2} \hat{Q}_1^{\text{cf}}(z_\alpha) + n^{-1} \hat{Q}_2^{\text{cf}}(z_\alpha) + n^{-3/2} \hat{Q}_3^{\text{cf}}(z_\alpha) + O_p(n^{-2}), \quad (14)$$

where \hat{Q}_k^{cf} is derived from empirical Edgeworth polynomials $\hat{Q}_1, \dots, \hat{Q}_k$ in the standard way, discussed above (13); and \hat{Q}_k is derived from the Edgeworth polynomial, Q_k , on replacing moments of the distribution of X , appearing in coefficients of Q_k , by the same respective moments of the distribution of X^* , conditional on \mathcal{X} , with X^* drawn by sampling, randomly and with replacement, from \mathcal{X} . The coefficients of \hat{Q}_k depend on moments of X^* , conditional on \mathcal{X} , through rational polynomials in those conditional moments.

If we knew the sampling distribution of X , and wished to construct an upper one-sided confidence interval for θ , we would employ the Studentized confidence interval $(\hat{\theta} - n^{-1/2} \hat{\sigma} x_\alpha, \infty)$, where x_α is as at (13); if we were to use the percentile- t bootstrap method, it would be $(\hat{\theta} - n^{-1/2} \hat{\sigma} \hat{x}_\alpha, \infty)$, where \hat{x}_α is as at (14); and if we were to employ the warp-speed bootstrap method, it would be $(\hat{\theta} - n^{-1/2} \hat{\sigma} \hat{x}_{\hat{\beta}_\alpha}, \infty)$, as discussed in §3.2, where $\hat{\beta}_\alpha$ denotes the limit, as $B \rightarrow \infty$, of the quantity $\hat{\beta}_{B\alpha}$ introduced there. However, we shall show in Theorem 2 that $\hat{x}_{\hat{\beta}_\alpha} = \hat{x}_\alpha + O_p(n^{-3/2})$, and so the endpoints of standard percentile- t and warp-speed bootstrap confidence intervals differ only to order $n^{-3/2}$. This signals that conventional arguments, based on Edgeworth expansions, can be used to prove that the standard percentile- t confidence

interval, and its warp-speed bootstrap variant, have identical coverage error up to and including terms of order n^{-1} , and of course that can be done under the assumptions of Theorem 2. Since the coverage error of the percentile- t interval is genuinely of order n^{-1} (Hall, 1986), then it follows that the warp-speed bootstrap does not improve on that accuracy.

THEOREM 2. *Assume that Condition 1(i) applies; that the function f , in the definition $\theta = f(\mu)$, has five bounded derivatives; and that (11) holds, $E(\|X\|^K) < \infty$ for sufficiently large $K > 0$, and $B = \infty$. Then $\hat{x}_{\hat{\beta}_\alpha} = \hat{x}_\alpha + O_p(n^{-3/2})$.*

The appropriate number of moments that should be assumed for general Edgeworth or Cornish–Fisher expansions, even in relatively simple, non-bootstrap cases, is awkward to determine. For example, the argument of Bhattacharya & Ghosh (1978) requires at least six moments in the case of the Studentized mean, whereas it is known that three moments are sufficient; see, e.g., Hall (1987). Even if we were to develop, in full detail, a proof of Theorem 2 based on the methods of Bhattacharya & Ghosh (1978), the number of moments we would need to assume would be unduly generous, and instead refer to the number as simply K . We choose not to provide such a detailed development here. However, the number of derivatives is relatively easy to address, and the theorem provides detail in that respect.

Let

$$\tilde{F}^*(x) = \text{pr}\{n^{1/2}(\hat{\theta}^{**} - \hat{\theta}^*)/\hat{\sigma}^{**} \leq x \mid \mathcal{X}\},$$

which is the limit of $\tilde{F}_B^*(x)$, defined in (7), as $B \rightarrow \infty$. Then $\hat{x}_{\hat{\beta}_\alpha}$ is the solution of $\tilde{F}^*(x) = \alpha$. Our focus on the case $B = \infty$ deserves comment. In the early days of the bootstrap, $B = \infty$ was seen as the statistical bootstrap method, and the case of finite B was interpreted as a Monte Carlo approximation to the bootstrap. Indeed, taking $B < \infty$ was viewed more as an issue to be addressed in computational or numerical terms, rather than statistical ones. Reflecting this, for about eight years from the mid 1980s considerable effort was spent developing efficient computational methods for undertaking bootstrap resampling. However, by the early 1990s computers had become so fast that this area of research had largely disappeared. This remains the case today; taking B in the thousands, without using numerical devices to increase simulation efficiency, is now the rule rather than the exception. The difference between such large values of B , and using the mathematical ideal value $B = \infty$, is particularly small.

ACKNOWLEDGEMENT

This research was supported by the Australian Research Council. We are grateful to three referees and the editor for helpful suggestions. Chang is also affiliated with the University of Melbourne.

SUPPLEMENTARY MATERIAL

Supplementary material available at *Biometrika* online includes theoretical proofs of Theorems 1 and 2, and additional simulation results for §§ 4.1 and 4.2.

REFERENCES

- BERAN, R. (1987). Prepivoting to reduce level error in confidence sets. *Biometrika* **74**, 457–68.
 BERAN, R. (1988). Prepivoting test statistics: A bootstrap view of asymptotic refinements. *J. Am. Statist. Assoc.* **83**, 687–97.
 BHATTACHARYA, R. N. & GHOSH, J. K. (1978). On the validity of the formal Edgeworth expansion. *Ann. Statist.* **6**, 434–51.

- BOOTH, J. G. & HALL, P. (1994). Monte Carlo approximation and the iterated bootstrap. *Biometrika* **81**, 331–40.
- BOOTH, J. G. & PRESNELL, B. (1998). Allocation of Monte Carlo resources for the iterated bootstrap. *J. Comp. Graph. Statist.* **7**, 92–112.
- DAVIDSON, R. & MACKINNON, J. G. (2002). Fast double bootstrap tests of nonnested linear regression models. *Economet. Rev.* **21**, 417–27.
- DAVIDSON, R. & MACKINNON, J. G. (2007). Improving the reliability of bootstrap tests with the fast double bootstrap. *Comp. Statist. Data Anal.* **51**, 3259–81.
- EFRON, B. (1983). Estimating the error rate of a prediction rule: Improvement on cross-validation. *J. Am. Statist. Assoc.* **78**, 316–31.
- GIACOMINI, R., POLITIS, D. N. & WHITE, H. (2013). A warp-speed method for conducting Monte Carlo experiments involving bootstrap estimators. *Economet. Theory* **29**, 567–89.
- HALL, P. (1986). On the bootstrap and confidence intervals. *Ann. Statist.* **14**, 1431–52.
- HALL, P. (1987). Edgeworth expansion for Student's t statistic under minimal moment conditions. *Ann. Prob.* **15**, 920–31.
- HALL, P. (1990). On the relative performance of bootstrap and Edgeworth approximations of a distribution function. *J. Mult. Anal.* **35**, 108–29.
- HALL, P. (1992). *The Bootstrap and Edgeworth Expansion*. New York: Springer.
- HALL, P. & MARTIN, M. A. (1988). On bootstrap resampling and iteration. *Biometrika* **75**, 661–71.
- LEE, S. M. S. & YOUNG, G. A. (1999). The effect of Monte Carlo approximation on coverage error of double-bootstrap confidence intervals. *J. R. Statist. Soc. B* **61**, 353–66.
- WHITE, H. (2000). A reality check for data snooping. *Econometrica* **68**, 1097–126.

[Received April 2014. Revised September 2014]