

Bootstrap-Based Inference for Cube Root Consistent Estimators*

Matias D. Cattaneo[†]

Michael Jansson[‡]

Kenichi Nagasawa[§]

April 24, 2017

Abstract

This note proposes a consistent bootstrap-based distributional approximation for cube root consistent estimators such as the maximum score estimator of [Manski \(1975\)](#) and the isotonic density estimator of [Grenander \(1956\)](#). In both cases, the standard nonparametric bootstrap is known to be inconsistent. Our method restores consistency of the nonparametric bootstrap by altering the shape of the criterion function defining the estimator whose distribution we seek to approximate. This modification leads to a generic and easy-to-implement resampling method for inference that is conceptually distinct from other available distributional approximations based on some form of modified bootstrap. We offer simulation evidence showcasing the performance of our inference method in finite samples. An extension of our methodology to general M-estimation problems is also discussed.

Keywords: cube root asymptotics, bootstrapping, maximum score estimation, isotonic density estimation.

JEL: C12, C14, C21.

*We thank Andres Santos for very thoughtful discussions. Cattaneo gratefully acknowledges financial support from the National Science Foundation through grant SES-1459931, and Jansson gratefully acknowledges financial support from the National Science Foundation through grant SES-1459967 and the research support of CREATES (funded by the Danish National Research Foundation under grant no. DNRF78).

[†]Department of Economics and Department of Statistics, University of Michigan.

[‡]Department of Economics, University of California at Berkeley and CREATES.

[§]Department of Economics, University of Michigan.

1 Introduction

In a seminal paper, [Kim and Pollard \(1990\)](#) studied a class of M-estimators exhibiting cube root asymptotics. These estimators not only have a non-standard rate of convergence, but also have the property that rather than being Gaussian their limiting distributions are of the [Chernoff \(1964\)](#) type. More precisely, the limiting distribution of $\sqrt[3]{n}$ times their estimation error is that of the maximizer of a Gaussian process (where n denotes the sample size). In fact, in leading examples of cube root consistent estimators such as the maximum score estimator of [Manski \(1975\)](#), the covariance kernel of the Gaussian process characterizing the limiting distribution depends on an infinite-dimensional nuisance parameter. As a consequence, whereas it is customary to conduct inference using analytical “plug-in” covariance matrix estimators in the standard \sqrt{n} -normal case (i.e., using a finite-dimensional nuisance parameter estimator), resampling-based distributional approximations offer the most attractive approach to inference in estimation problems exhibiting cube root asymptotics. The purpose of this note is to propose an easy-to-implement bootstrap-based distributional approximation applicable in such cases.

As does the standard nonparametric bootstrap, the method proposed herein employs bootstrap samples of size n from the empirical distribution function. But unlike the nonparametric bootstrap, which is inconsistent in general (e.g., [Abrevaya and Huang, 2005](#); [Léger and MacGibbon, 2006](#)), our method offers a consistent distributional approximation for $\sqrt[3]{n}$ -consistent estimators and therefore has the advantage that it can be used to construct asymptotically valid inference procedures. Consistency is achieved by altering the shape of the criterion function defining the estimator whose distribution we seek to approximate: heuristically, the method is designed to ensure that the bootstrap version of a certain empirical process has a mean which resembles the large sample version of its population counterpart. The latter is quadratic in the problems we study and known up to the value of a certain matrix. As a consequence, the only ingredient needed to implement the proposed “reshapement” of the objective function is a consistent estimator of the unknown matrix entering the quadratic mean of the empirical process. Such estimators turn out to be generically available and easy to compute.

Several alternative resampling-based distributional approximations for cube root consistent estimators have been proposed in the literature. The best known alternative is probably subsampling

(Politis and Romano, 1994), whose applicability to cube root asymptotic problems was pointed out by Delgado, Rodriguez-Poo, and Wolf (2001). A related method is the rescaled bootstrap (Dümbgen, 1993), whose validity in cube root asymptotic M -estimation problems was established recently by Hong and Li (2017). In addition, case-specific smooth bootstrap methods have been proposed for leading examples such as maximum score estimation (Patra, Seijo, and Sen, 2015) and isotonic density estimation (Kosorok, 2008; Sen, Banerjee, and Woodroffe, 2010). Each of these methods can also be viewed as offering a “robust” alternative to the nonparametric bootstrap but, unlike our proposed approach, they all achieve consistency by modifying the distribution used to generate the bootstrap counterpart of the estimator whose distribution is being approximated. In contrast, our method achieves consistency by means of an analytic modification to the objective function used to construct the standard nonparametric bootstrap distributional approximation, and hence is conceptually distinct from the modifications employed in the existing literature.

The note proceeds as follows. Section 2 is heuristic in nature and serves the purpose of outlining the main idea underlying our approach. The heuristics of Section 2 are then made rigorous in Section 3. Section 4 considers the maximum score estimator, illustrating the implications of our main results for that estimator and investigating the small sample properties of our proposed inference procedure in a simulation experiment. Two distinct extensions of our results, to M -estimators exhibiting an arbitrary rate of convergence and to the Grenander (1956) estimator of an isotonic density, respectively, are provided in Section 5. Simulation evidence for the latter example is also provided. All derivations and proofs have been collected in the supplemental appendix.

2 Cube Root Asymptotics

Suppose $\theta_0 \in \Theta \subseteq \mathbb{R}^d$ is an estimand admitting the characterization

$$\theta_0 = \operatorname{argmax}_{\theta \in \Theta} M(\theta), \quad M(\theta) = \mathbb{E}[m(\mathbf{z}, \theta)],$$

where \mathbf{z} is a random vector of which a random sample $\{\mathbf{z}_i : 1 \leq i \leq n\}$ is available, and where m is a known function. Studying estimation problem of this kind for non-smooth m , Kim and Pollard

(1990) found that M -estimators such as

$$\hat{\boldsymbol{\theta}}_n = \operatorname{argmax}_{\boldsymbol{\theta} \in \Theta} M_n(\boldsymbol{\theta}), \quad M_n(\boldsymbol{\theta}) = \frac{1}{n} \sum_{i=1}^n m(\mathbf{z}_i, \boldsymbol{\theta}), \quad (1)$$

often exhibit cube root asymptotics:

$$\sqrt[3]{n}(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0) \rightsquigarrow \operatorname{argmax}_{\mathbf{s} \in \mathbb{R}^d} \{\mathcal{Q}(\mathbf{s}) + \mathcal{G}(\mathbf{s})\}, \quad (2)$$

where \rightsquigarrow denotes weak convergence, $\mathcal{Q}(\mathbf{s})$ is a quadratic form given by

$$\mathcal{Q}(\mathbf{s}) = -\frac{1}{2} \mathbf{s}' \mathbf{V}_0 \mathbf{s}, \quad \mathbf{V}_0 = -\frac{\partial^2 M(\boldsymbol{\theta}_0)}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}'},$$

and \mathcal{G} is a non-degenerate zero-mean Gaussian process with $\mathcal{G}(0) = 0$.

Whereas the matrix \mathbf{V}_0 governing the shape of \mathcal{Q} is finite-dimensional, the covariance kernel of \mathcal{G} in (2) typically involves infinite-dimensional unknown quantities. As a consequence, the limiting distribution of $\hat{\boldsymbol{\theta}}_n$ tends to be more difficult to approximate than conventional Gaussian distributions, implying in turn that basing inference on $\hat{\boldsymbol{\theta}}_n$ is more challenging under cube root asymptotics than in the more familiar case where $\hat{\boldsymbol{\theta}}_n$ is (\sqrt{n} -consistent and) asymptotically normally distributed.

As a candidate method of approximating the distribution of $\hat{\boldsymbol{\theta}}_n$, consider the standard nonparametric bootstrap. To describe it, let $\{\mathbf{z}_i^* : 1 \leq i \leq n\}$ denote a random sample from the empirical distribution of $\{\mathbf{z}_i : 1 \leq i \leq n\}$ and let

$$\hat{\boldsymbol{\theta}}_n^* = \operatorname{argmax}_{\boldsymbol{\theta} \in \Theta} M_n^*(\boldsymbol{\theta}), \quad M_n^*(\boldsymbol{\theta}) = \frac{1}{n} \sum_{i=1}^n m(\mathbf{z}_i^*, \boldsymbol{\theta}),$$

be the natural bootstrap analogue of $\hat{\boldsymbol{\theta}}_n$. Then the nonparametric bootstrap approximation to the distribution of $\hat{\boldsymbol{\theta}}_n$ is given by

$$\mathbb{P}[\sqrt[3]{n}(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0) \leq \mathbf{t}] \approx \mathbb{P}^*[\sqrt[3]{n}(\hat{\boldsymbol{\theta}}_n^* - \hat{\boldsymbol{\theta}}_n) \leq \mathbf{t}],$$

where \mathbb{P}^* denotes a probability computed under the bootstrap distribution conditional on the data.

As is well documented, however, this distributional approximation is inconsistent under cube root asymptotics (e.g., [Abrevaya and Huang, 2005](#); [Léger and MacGibbon, 2006](#)).

For the purpose of giving a heuristic, yet constructive, explanation of the inconsistency of the nonparametric bootstrap, it is helpful to recall that for $\hat{\boldsymbol{\theta}}_n$ defined in (1), a proof of (2) can be based on the representation

$$\sqrt[3]{n}(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0) = \operatorname{argmax}_{\mathbf{s} \in \mathbb{R}^d} \{Q_n(\mathbf{s}) + G_n(\mathbf{s})\}, \quad (3)$$

where

$$G_n(\mathbf{s}) = n^{2/3}[M_n(\boldsymbol{\theta}_0 + \mathbf{s}/n^{1/3}) - M_n(\boldsymbol{\theta}_0) - M(\boldsymbol{\theta}_0 + \mathbf{s}/n^{1/3}) + M(\boldsymbol{\theta}_0)]$$

is a zero-mean random process, while

$$Q_n(\mathbf{s}) = n^{2/3}[M(\boldsymbol{\theta}_0 + \mathbf{s}/n^{1/3}) - M(\boldsymbol{\theta}_0)]$$

is a non-random function that is correctly centered in the sense that $\operatorname{argmax}_{\mathbf{s} \in \mathbb{R}^d} Q_n(\mathbf{s}) = \mathbf{0}$. In cases where m is non-smooth but M is smooth, Q_n and G_n are usually asymptotically quadratic and asymptotically Gaussian, respectively, in the sense that

$$Q_n(\mathbf{s}) \rightarrow \mathcal{Q}(\mathbf{s}) \quad (4)$$

and

$$G_n(\mathbf{s}) \rightsquigarrow \mathcal{G}(\mathbf{s}). \quad (5)$$

Under regularity conditions ensuring among other things that the convergence in (4) and (5) is suitably uniform in \mathbf{s} , the proof of (2) can then be completed by applying a continuous mapping-type theorem for the argmax functional to the representation in (3).

Similarly to (3), the bootstrap analogue of $\hat{\boldsymbol{\theta}}_n$ admits a representation of the form

$$\sqrt[3]{n}(\hat{\boldsymbol{\theta}}_n^* - \hat{\boldsymbol{\theta}}_n) = \operatorname{argmax}_{\mathbf{s} \in \mathbb{R}^d} \{Q_n^*(\mathbf{s}) + G_n^*(\mathbf{s})\}, \quad (6)$$

where

$$G_n^*(\mathbf{s}) = n^{2/3}[M_n^*(\hat{\boldsymbol{\theta}}_n + \mathbf{s}/n^{1/3}) - M_n^*(\hat{\boldsymbol{\theta}}_n) - M_n(\hat{\boldsymbol{\theta}}_n + \mathbf{s}/n^{1/3}) + M_n(\hat{\boldsymbol{\theta}}_n)]$$

and

$$Q_n^*(\mathbf{s}) = n^{2/3}[M_n(\hat{\boldsymbol{\theta}}_n + \mathbf{s}/n^{1/3}) - M_n(\hat{\boldsymbol{\theta}}_n)].$$

Under mild conditions, G_n^* satisfies the following bootstrap counterpart of (5):

$$G_n^*(\mathbf{s}) \rightsquigarrow_{\mathbb{P}} \mathcal{G}(\mathbf{s}), \tag{7}$$

where $\rightsquigarrow_{\mathbb{P}}$ denotes weak convergence in probability. On the other hand, although Q_n^* is non-random under the bootstrap distribution and satisfies $\operatorname{argmax}_{\mathbf{s} \in \mathbb{R}^d} Q_n^*(\mathbf{s}) = \mathbf{0}$, it turns out that $Q_n^*(\mathbf{s}) \not\rightarrow_{\mathbb{P}} \mathcal{Q}(\mathbf{s})$ in general. In other words, and as explained in more detail in Section 5, the natural bootstrap counterpart of (4) typically fails and, as a partial consequence, so does the natural bootstrap counterpart of (2); that is, $\sqrt[3]{n}(\hat{\boldsymbol{\theta}}_n^* - \hat{\boldsymbol{\theta}}_n) \not\rightarrow_{\mathbb{P}} \operatorname{argmax}_{\mathbf{s} \in \mathbb{R}^d} \{\mathcal{Q}(\mathbf{s}) + \mathcal{G}(\mathbf{s})\}$ in general.

To the extent that the implied inconsistency of the bootstrap can be attributed to the fact that the shape of Q_n^* fails to replicate that of Q_n , it seems plausible that a consistent bootstrap-based distributional approximation can be obtained by basing the approximation on

$$\tilde{\boldsymbol{\theta}}_n^* = \operatorname{argmax}_{\boldsymbol{\theta} \in \Theta} \tilde{M}_n^*(\boldsymbol{\theta}), \quad \tilde{M}_n^*(\boldsymbol{\theta}) = \frac{1}{n} \sum_{i=1}^n \tilde{m}_n(\mathbf{z}_i^*, \boldsymbol{\theta}),$$

where \tilde{m}_n is a suitably “reshaped” version of m satisfying two properties. First, \tilde{G}_n^* should be asymptotically equivalent to G_n^* , where

$$\tilde{G}_n^*(\mathbf{s}) = n^{2/3}[\tilde{M}_n^*(\hat{\boldsymbol{\theta}}_n + \mathbf{s}/n^{1/3}) - \tilde{M}_n^*(\hat{\boldsymbol{\theta}}_n) - \tilde{M}_n(\hat{\boldsymbol{\theta}}_n + \mathbf{s}/n^{1/3}) + \tilde{M}_n(\hat{\boldsymbol{\theta}}_n)], \quad \tilde{M}_n(\boldsymbol{\theta}) = \frac{1}{n} \sum_{i=1}^n \tilde{m}_n(\mathbf{z}_i, \boldsymbol{\theta}),$$

is the counterpart of G_n^* associated with \tilde{m}_n . Second, \tilde{Q}_n^* should be asymptotically quadratic, where

$$\tilde{Q}_n^*(\mathbf{s}) = n^{2/3}[\tilde{M}_n(\hat{\boldsymbol{\theta}}_n + \mathbf{s}/n^{1/3}) - \tilde{M}_n(\hat{\boldsymbol{\theta}}_n)]$$

is the counterpart of Q_n^* associated with \tilde{m}_n .

Accordingly, let

$$\tilde{m}_n(\mathbf{z}, \boldsymbol{\theta}) = m(\mathbf{z}, \boldsymbol{\theta}) - M_n(\boldsymbol{\theta}) - \frac{1}{2}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_n)' \tilde{\mathbf{V}}_n (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_n),$$

where $\tilde{\mathbf{V}}_n$ is an estimator of \mathbf{V}_0 . Then

$$\sqrt[3]{n}(\tilde{\boldsymbol{\theta}}_n^* - \hat{\boldsymbol{\theta}}_n) = \operatorname{argmax}_{\mathbf{s} \in \mathbb{R}^d} \left\{ \tilde{Q}_n^*(\mathbf{s}) + \tilde{G}_n^*(\mathbf{s}) \right\},$$

where, by construction, $\tilde{G}_n^*(\mathbf{s}) = G_n^*(\mathbf{s})$ and $\tilde{Q}_n^*(\mathbf{s}) = -\mathbf{s}' \tilde{\mathbf{V}}_n \mathbf{s} / 2$. Because $\tilde{G}_n^* = G_n^*$, $\tilde{G}_n^*(\mathbf{s}) \rightsquigarrow_{\mathbb{P}} \mathcal{G}(\mathbf{s})$ whenever (7) holds. In addition, $\tilde{Q}_n^*(\mathbf{s}) \rightarrow_{\mathbb{P}} \mathcal{Q}(\mathbf{s})$ provided $\tilde{\mathbf{V}}_n \rightarrow_{\mathbb{P}} \mathbf{V}_0$. As a consequence, it seems plausible that the distributional approximation

$$\mathbb{P}[\sqrt[3]{n}(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0) \leq \mathbf{t}] \approx \mathbb{P}^*[\sqrt[3]{n}(\tilde{\boldsymbol{\theta}}_n^* - \hat{\boldsymbol{\theta}}_n) \leq \mathbf{t}]$$

is consistent if $\tilde{\mathbf{V}}_n \rightarrow_{\mathbb{P}} \mathbf{V}_0$.

3 Main Result

To make the heuristics of the previous section precise, we impose the following condition.

Condition CRA (Cube Root Asymptotics) (i) For every $\delta > 0$, $\sup_{\|\boldsymbol{\theta} - \boldsymbol{\theta}_0\| > \delta} M(\boldsymbol{\theta}) < M(\boldsymbol{\theta}_0)$.

Also, the class $\mathcal{M} = \{m(\cdot, \boldsymbol{\theta}) : \boldsymbol{\theta} \in \boldsymbol{\Theta}\}$ is manageable for the envelope $\bar{M}(\cdot) = \sup_{m \in \mathcal{M}} |m(\cdot)|$, and $\mathbb{E}[\bar{M}(\mathbf{z})^4] < \infty$.

(ii) $\boldsymbol{\theta}_0$ is an interior point of $\boldsymbol{\Theta}$.

(iii) M is twice continuously differentiable near $\boldsymbol{\theta}_0$, with \mathbf{V}_0 positive definite.

(iv) For $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2$ near $\boldsymbol{\theta}_0$, $\mathbb{E}[|m(\mathbf{z}, \boldsymbol{\theta}_1) - m(\mathbf{z}, \boldsymbol{\theta}_2)|] = O(\|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|)$.

(v) For all $\mathbf{s}, \mathbf{t} \in \mathbb{R}^d$,

$$H(\mathbf{s}, \mathbf{t}) = \lim_{\boldsymbol{\theta} \rightarrow \boldsymbol{\theta}_0, \delta \downarrow 0} \frac{1}{\delta} \mathbb{E}[\{m(\mathbf{z}, \boldsymbol{\theta} + \delta \mathbf{s}) - m(\mathbf{z}, \boldsymbol{\theta})\} \{m(\mathbf{z}, \boldsymbol{\theta} + \delta \mathbf{t}) - m(\mathbf{z}, \boldsymbol{\theta})\}]$$

exists, and $H(\mathbf{s}, \mathbf{s}) + H(\mathbf{t}, \mathbf{t}) - 2H(\mathbf{s}, \mathbf{t}) > 0$ for all $\mathbf{s} \neq \mathbf{t}$.

(vi) For R near zero, the classes $\mathcal{D}_R = \{m(\cdot, \boldsymbol{\theta}) - m(\cdot, \boldsymbol{\theta}_0) : \|\boldsymbol{\theta} - \boldsymbol{\theta}_0\| \leq R\}$ are uniformly

manageable for the envelopes $\bar{D}_R(\cdot) = \sup_{D \in \mathcal{D}_R} |D(\cdot)|$, $\mathbb{E}[\bar{D}_R(\mathbf{z})^2] = O(R)$, and for every $\epsilon > 0$ there is a constant K such that $\mathbb{E}[\bar{D}_R(\mathbf{z})^2 \mathbf{1}\{\bar{D}_R(\mathbf{z}) > K\}] < \epsilon R$.

Condition CRA is similar to, but slightly stronger than, assumptions (iii)-(vii) of the main theorem of [Kim and Pollard \(1990\)](#), to which the reader is referred for a discussion of these assumptions as well as a definition of the term (uniformly) manageable. To be specific, parts (ii)-(iv) and (vi) are identical to their counterparts in [Kim and Pollard \(1990\)](#), part (v) is a locally uniform (with respect to $\boldsymbol{\theta}$ near $\boldsymbol{\theta}_0$) version of its counterpart in [Kim and Pollard \(1990\)](#), while (i) can be thought of as replacing the high level condition $\hat{\boldsymbol{\theta}}_n \rightarrow_{\mathbb{P}} \boldsymbol{\theta}_0$ with more primitive conditions that imply it for (approximate M -estimators) $\hat{\boldsymbol{\theta}}_n$ satisfying

$$M_n(\hat{\boldsymbol{\theta}}_n) \geq \sup_{\boldsymbol{\theta} \in \Theta} M_n(\boldsymbol{\theta}) - o_{\mathbb{P}}(n^{-2/3}). \quad (8)$$

In the case of both (i) and (v), the purpose of strengthening the assumptions of [Kim and Pollard \(1990\)](#) is to be able to analyze the bootstrap.

Our main result is the following.

Theorem 1 *Suppose Condition CRA holds and that $\hat{\boldsymbol{\theta}}_n$ satisfies (8). If $\tilde{\mathbf{V}}_n \rightarrow_{\mathbb{P}} \mathbf{V}_0$ and if $\tilde{M}_n^*(\tilde{\boldsymbol{\theta}}_n^*) \geq \sup_{\boldsymbol{\theta} \in \Theta} \tilde{M}_n^*(\boldsymbol{\theta}) - o_{\mathbb{P}}(n^{-2/3})$, then*

$$\sup_{\mathbf{t} \in \mathbb{R}^d} \left| \mathbb{P}^*[\sqrt[3]{n}(\tilde{\boldsymbol{\theta}}_n^* - \hat{\boldsymbol{\theta}}_n) \leq \mathbf{t}] - \mathbb{P}[\sqrt[3]{n}(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0) \leq \mathbf{t}] \right| \rightarrow_{\mathbb{P}} 0. \quad (9)$$

Under the conditions of the theorem, it follows from [Kim and Pollard \(1990\)](#) that (2) holds, with \mathcal{G} having covariance kernel H . Mimicking the derivation of that result, the proof of the theorem proceeds by establishing the following bootstrap counterpart of (2):

$$\sqrt[3]{n}(\tilde{\boldsymbol{\theta}}_n^* - \hat{\boldsymbol{\theta}}_n) \rightsquigarrow_{\mathbb{P}} \operatorname{argmax}_{\mathbf{s} \in \mathbb{R}^d} \{ \mathcal{Q}(\mathbf{s}) + \mathcal{G}(\mathbf{s}) \}. \quad (10)$$

The theorem offers a valid bootstrap-based distributional approximation for $\hat{\boldsymbol{\theta}}_n$. To implement the approximation, only a consistent estimator of $\mathbf{V}_0 = -\partial^2 M(\boldsymbol{\theta}_0) / \partial \boldsymbol{\theta} \partial \boldsymbol{\theta}'$ is needed. A generic

estimator based on numerical derivatives is $\tilde{\mathbf{V}}_n^{\text{ND}}$, the matrix whose element (k, l) is given by

$$\tilde{\mathbf{V}}_{n,kl}^{\text{ND}} = -\frac{1}{4\epsilon_n^2} \left[M_n(\hat{\boldsymbol{\theta}}_n + \mathbf{e}_k\epsilon_n + \mathbf{e}_l\epsilon_n) - M_n(\hat{\boldsymbol{\theta}}_n + \mathbf{e}_k\epsilon_n - \mathbf{e}_l\epsilon_n) \right. \\ \left. - M_n(\hat{\boldsymbol{\theta}}_n - \mathbf{e}_k\epsilon_n + \mathbf{e}_l\epsilon_n) + M_n(\hat{\boldsymbol{\theta}}_n - \mathbf{e}_k\epsilon_n - \mathbf{e}_l\epsilon_n) \right],$$

where \mathbf{e}_k is the k^{th} unit vector in \mathbb{R}^d and where $\epsilon_n = o(1)$ is a tuning parameter. Conditions under which this estimator is consistent are given in the following lemma.

Lemma 1 *Suppose Condition CRA holds and that $\hat{\boldsymbol{\theta}}_n$ satisfies (8). If $\epsilon_n \rightarrow 0$ and $n\epsilon_n^3 \rightarrow \infty$, then $\tilde{\mathbf{V}}_n^{\text{ND}} \rightarrow_{\mathbb{P}} \mathbf{V}_0$.*

The proof of the lemma goes beyond consistency and develops a Nagar-type mean squared error (MSE) expansion for $\tilde{\mathbf{V}}_n^{\text{ND}}$ under the additional assumption that M is four times continuously differentiable near $\boldsymbol{\theta}_0$. The approximate MSE (AMSE) can be minimized by choosing ϵ_n proportional to $n^{-1/7}$, the optimal factor of proportionality being a functional of the covariance kernel H and the fourth order derivatives of M evaluated at $\boldsymbol{\theta}_0$. For details, see the supplemental appendix (Theorem A.2 and Section A.5), which also contains a brief discussion of alternative generic estimators of \mathbf{V}_0 .

4 The Maximum Score Estimator

Arguably the most prominent econometric example of an estimator exhibiting cube root asymptotics is the maximum score estimator of [Manski \(1975\)](#). To describe a version of this estimator, suppose $\{\mathbf{z}_i : 1 \leq i \leq n\}$ is a random sample of $\mathbf{z} = (y, \mathbf{x}')'$ generated by the binary response model

$$y = \mathbb{1}(\mathbf{x}'\boldsymbol{\beta}_0 + \varepsilon \geq 0), \quad F_{\varepsilon|\mathbf{x}}(0|\mathbf{x}) = 1/2, \quad (11)$$

where $\boldsymbol{\beta}_0 \in \mathbb{R}^{d+1}$ is an unknown parameter of interest, $\mathbf{x} \in \mathbb{R}^{d+1}$ is a vector of covariates, and $F_{\varepsilon|\mathbf{x}}(\cdot|\mathbf{x})$ is the conditional cumulative distribution function of the unobserved error term ε given \mathbf{x} . Following [Abrevaya and Huang \(2005\)](#), we normalize the (unidentified) scale of $\boldsymbol{\beta}_0$ by setting the first element of the parameter vector equal to unity. In other words, we employ the parameterization $\boldsymbol{\beta}_0 = (1, \boldsymbol{\theta}_0)'$, where $\boldsymbol{\theta}_0 \in \mathbb{R}^d$ is unknown. Partitioning \mathbf{x} conformably with $\boldsymbol{\beta}_0$ as $\mathbf{x} = (x_1, \mathbf{x}_2)'$, a

maximum score estimator of $\boldsymbol{\theta}_0 \in \Theta \subseteq \mathbb{R}^d$ is any $\hat{\boldsymbol{\theta}}_n^{\text{MS}}$ satisfying (8) for

$$m(\mathbf{z}, \boldsymbol{\theta}) = m^{\text{MS}}(\mathbf{z}, \boldsymbol{\theta}) = (2y - 1)\mathbb{1}(x_1 + \mathbf{x}'_2\boldsymbol{\theta} \geq 0).$$

Regarded as a member of the class of M -estimators exhibiting cube root asymptotics, the maximum score estimator is representative in a couple of respects. First, under easy-to-interpret primitive conditions the estimator is covered by the results of Section 3. To state a set of such conditions, let $f_{x_1|\mathbf{x}_2}(\cdot|\mathbf{x}_2)$ denote the density function of x_1 given \mathbf{x}_2 .

Condition MS (Maximum Score) (i) $(y, \mathbf{x}')'$ satisfies (11) and $0 < \mathbb{P}(y = 1|\mathbf{x}) < 1$. Also,

$F_{\varepsilon|x_1, \mathbf{x}_2}(\varepsilon|x_1, \mathbf{x}_2)$ is differentiable with respect to ε and x_1 , with continuous and bounded derivatives.

(ii) The support of \mathbf{x} is not contained in any proper linear subspace of \mathbb{R}^{d+1} , $\mathbb{E}[\|\mathbf{x}_2\|^6] < \infty$, and $f_{x_1|\mathbf{x}_2}(\cdot|\mathbf{x}_2)$ is continuous, bounded, and everywhere positive.

(iii) $\boldsymbol{\beta}_0 = (1, \boldsymbol{\theta}'_0)'$, $\boldsymbol{\theta}_0$ is an interior point of $\Theta \subseteq \mathbb{R}^d$, and Θ is compact.

(iv) M^{MS} is twice continuously differentiable near $\boldsymbol{\theta}_0$, with $\mathbf{V}_0^{\text{MS}} = -\partial^2 M^{\text{MS}}(\boldsymbol{\theta}_0)/\partial\boldsymbol{\theta}\partial\boldsymbol{\theta}'$ positive definite, where $M^{\text{MS}}(\boldsymbol{\theta}) = \mathbb{E}[(2y - 1)\mathbb{1}(x_1 + \mathbf{x}'_2\boldsymbol{\theta} \geq 0)]$.

Second, in addition to the generic estimator $\tilde{\mathbf{V}}_n^{\text{ND}}$ discussed above, the maximum score estimator admits an example-specific consistent estimator of the \mathbf{V}_0 associated with it. Let

$$\tilde{\mathbf{V}}_n^{\text{MS}} = -\frac{1}{nh_n^2} \sum_{i=1}^n (2y_i - 1) \dot{K} \left(\frac{x_{1i} + \mathbf{x}'_{2i} \hat{\boldsymbol{\theta}}_n^{\text{MS}}}{h_n} \right) \mathbf{x}_{2i} \mathbf{x}'_{2i},$$

where $h_n = o(1)$ is a bandwidth and $\dot{K}(u) = dK(u)/du$, where K is a kernel function. As defined, $\tilde{\mathbf{V}}_n^{\text{MS}}$ is simply minus the second derivative, evaluated at $\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}_n^{\text{MS}}$, of the criterion function associated with the smoothed maximum score estimator of Horowitz (1992). The estimator $\tilde{\mathbf{V}}_n^{\text{MS}}$ is consistent under mild conditions on h_n , provided the kernel satisfies the following.

Condition K (Kernel) (i) $\int_{\mathbb{R}} K(u)du = 1$ and $\lim_{|u| \rightarrow \infty} |uK(u)| = 0$.

(ii) $\int_{\mathbb{R}} |u\dot{K}(u)|du + \int_{\mathbb{R}} \dot{K}(u)^2 du < \infty$ and, for some $B(\cdot)$ with $\int_{\mathbb{R}} B(u)^2 du < \infty$,

$$|\dot{K}(u_1) - \dot{K}(u_2)| \leq B(u_1)|u_1 - u_2| \quad \forall u_1, u_2 \in \mathbb{R}.$$

Our results about the maximum score estimator can be summarized as follows.

Theorem 2 *Suppose Condition MS holds.*

(i) *Then Condition CRA holds for $m = m^{\text{MS}}$.*

Suppose also that $\hat{\boldsymbol{\theta}}_n^{\text{MS}}$ satisfies (8) and that Condition K holds.

(ii) *If $h_n \rightarrow 0$ and $nh_n^3 \rightarrow \infty$, then $\tilde{\mathbf{V}}_n^{\text{MS}} \rightarrow_{\mathbb{P}} \mathbf{V}_0^{\text{MS}}$.*

As in the case of Lemma 1, the proof of part (ii) of the theorem goes beyond consistency and develops a Nagar-type MSE expansion for $\tilde{\mathbf{V}}_n^{\text{MS}}$ under some the additional assumptions. The AMSE can be minimized by choosing h_n proportional to $n^{-1/7}$; for details, including a characterization of the optimal factor of proportionality and a simple rule-of-thumb choice thereof based on a Gaussian reference model, see the supplemental appendix (Theorem A.3 and Section A.5).

To investigate the finite sample properties of our proposed bootstrap-based inference procedures, we conducted a Monte Carlo experiment. Following Horowitz (2002), and to allow for a comparison with his bootstrap-based inference method for the smoothed maximum score estimator, we generate data from a model of the form (11) with $d = 1$, where

$$\mathbf{x} = (x_1, x_2)' \sim \mathcal{N} \left(\left(\begin{array}{c} 0 \\ 1 \end{array} \right), \left(\begin{array}{cc} 1 & 0 \\ 0 & 1 \end{array} \right) \right), \quad \mathbf{x} \perp \varepsilon,$$

and where ε can take three distinct distributions. Specifically, DGP 1 sets $\varepsilon \sim \text{Logistic}(0, 1)/\sqrt{2\pi^2/3}$, DGP 2 sets $\varepsilon \sim \mathcal{T}_3/\sqrt{3}$, where \mathcal{T}_k denotes a Student's t distribution with k degrees of freedom, and DGP 3 sets $\varepsilon \sim (1 + 2(x_1 + x_2)^2 + (x_1 + x_2)^4)\text{Logistic}(0, 1)/\sqrt{\pi^2/48}$. The parameter is $\boldsymbol{\theta}_0 = 1$ in all cases.

The Monte Carlo experiment employs a sample size $n = 1,000$ with $B = 2,000$ bootstrap replications and $S = 2,000$ simulations. For each of the three DGPs, we implement the standard non-parametric bootstrap, m -out-of- n bootstrap, and our proposed method using the two estimators $\tilde{\mathbf{V}}_n^{\text{MS}}$ and $\tilde{\mathbf{V}}_n^{\text{ND}}$ of \mathbf{V}_0 . We report empirical coverage for nominal 95% confidence intervals and their average interval length. For the case of our proposed procedures, we investigate their performance using both (i) a grid of fixed tuning parameter values (bandwidth/derivative step) around the MSE-optimal choice and (ii) infeasible and feasible AMSE-optimal choices of the

tuning parameter.

Table 1 presents the main results, which are consistent across all three simulation designs. First, as expected, the standard nonparametric bootstrap (labeled “Standard”) does not perform well, leading to confidence intervals with an average 64% empirical coverage rate. Second, the m -out-of- n bootstrap (labeled “m-out-of-n”) performs somewhat better for small subsamples, but leads to very large average interval length of the resulting confidence intervals. Our proposed methods, on the other hand, exhibit excellent finite sample performance in this Monte Carlo experiment. Results employing the example-specific plug-in estimator $\tilde{\mathbf{V}}_n^{\text{MS}}$ are presented under the label “Plug-in” while results employing the generic numerical derivative estimator $\tilde{\mathbf{V}}_n^{\text{ND}}$ are reported under the label “Num Deriv”. Empirical coverage appears stable across different values of the tuning parameters for each method, with better performance in the case of $\tilde{\mathbf{V}}_n^{\text{MS}}$. We conjecture that $n = 1,000$ is too small for the numerical derivative estimator $\tilde{\mathbf{V}}_n^{\text{ND}}$ to lead to as good inference performance as $\tilde{\mathbf{V}}_n^{\text{MS}}$ (e.g., note that the MSE-optimal choice ϵ_{MSE} is greater than 1). Nevertheless, empirical coverage of confidence intervals constructed using our proposed bootstrap-based method is close to 95% in all cases except when $\tilde{\mathbf{V}}_n^{\text{ND}}$ is used with either the infeasible asymptotic choice ϵ_{AMSE} or its estimated counterpart $\hat{\epsilon}_{\text{AMSE}}$, and with an average interval length of at most half that of any of the m -out-of- n competing confidence intervals. In particular, confidence intervals based on $\tilde{\mathbf{V}}_n^{\text{MS}}$ implemented with the feasible bandwidth \hat{h}_{AMSE} perform quite well across the three DGPs considered.

In sum, applying the bootstrap-based inference methods proposed in this note to the case of the Maximum Score estimator of Manski (1975) lead to confidence intervals with very good coverage and length properties in the simulation designs considered.

5 Extensions

The scope of some of the main insights of this note extends beyond the $\sqrt[3]{n}$ -consistent M -estimators covered by Theorem 1. This section briefly discusses two possible extensions. First, we indicate how our main results can be generalized to M -estimators exhibiting an arbitrary rate of convergence. Second, we illustrate how the idea of reshaping can be used to achieve consistency of a bootstrap-based approximation to the distribution of another prominent $\sqrt[3]{n}$ -consistent estimator, the isotonic density estimator, which is not of M -estimator type.

5.1 M -estimators with an arbitrary rate of convergence

For the purposes of explaining the inconsistency of the nonparametric bootstrap, the unusual rate of convergence of $\hat{\boldsymbol{\theta}}_n$ is a bit of a “red herring”. Accordingly, suppose $\hat{\boldsymbol{\theta}}_n$ satisfies (1) and that, for some increasing function r_n of n (such as $\sqrt[3]{n}$ or \sqrt{n}), $r_n(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0)$ has a non-degenerate limiting distribution. Then, in perfect analogy with Section 2, the shape of that limiting distribution can usually be anticipated with the help of the representation

$$r_n(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0) = \operatorname{argmax}_{\mathbf{s} \in \mathbb{R}^d} \{Q_n(\mathbf{s}) + G_n(\mathbf{s})\},$$

where

$$G_n(\mathbf{s}) = r_n^2 [M_n(\boldsymbol{\theta}_0 + \mathbf{s}/r_n) - M_n(\boldsymbol{\theta}_0) - M(\boldsymbol{\theta}_0 + \mathbf{s}/r_n) + M(\boldsymbol{\theta}_0)]$$

and

$$Q_n(\mathbf{s}) = r_n^2 [M(\boldsymbol{\theta}_0 + \mathbf{s}/r_n) - M(\boldsymbol{\theta}_0)].$$

Specifically, assuming the functions G_n and Q_n satisfy convergence properties of the form (5) and (4), respectively, it stands to reason that

$$\hat{\mathbf{s}}_n = r_n(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0) \rightsquigarrow \operatorname{argmax}_{\mathbf{s} \in \mathbb{R}^d} \{\mathcal{Q}(\mathbf{s}) + \mathcal{G}(\mathbf{s})\} = \mathcal{S}.$$

In particular, for example, if $r_n = \sqrt{n}$ and $\mathcal{G}(\mathbf{s}) = \mathbf{s}'\dot{\mathcal{G}}$ with $\dot{\mathcal{G}}$ some zero-mean Gaussian vector, then \mathcal{S} becomes the usual normal distribution of \sqrt{n} -consistent (parametric) M -estimators.

The bootstrap analogue of $\hat{\boldsymbol{\theta}}_n$ satisfies

$$r_n(\hat{\boldsymbol{\theta}}_n^* - \hat{\boldsymbol{\theta}}_n) = \operatorname{argmax}_{\mathbf{s} \in \mathbb{R}^d} \{Q_n^*(\mathbf{s}) + G_n^*(\mathbf{s})\},$$

where

$$G_n^*(\mathbf{s}) = r_n^2 [M_n^*(\hat{\boldsymbol{\theta}}_n + \mathbf{s}/r_n) - M_n^*(\hat{\boldsymbol{\theta}}_n) - M_n(\hat{\boldsymbol{\theta}}_n + \mathbf{s}/r_n) + M_n(\hat{\boldsymbol{\theta}}_n)]$$

and

$$Q_n^*(\mathbf{s}) = r_n^2 [M_n(\hat{\boldsymbol{\theta}}_n + \mathbf{s}/r_n) - M_n(\hat{\boldsymbol{\theta}}_n)].$$

Again, G_n^* will satisfy (7) under mild conditions. On the other hand, because

$$Q_n^*(\mathbf{s}) = Q_n(\hat{\mathbf{s}}_n + \mathbf{s}) + G_n(\hat{\mathbf{s}}_n + \mathbf{s}) - Q_n(\hat{\mathbf{s}}_n) - G_n(\hat{\mathbf{s}}_n),$$

we would expect

$$\begin{aligned} Q_n^*(\mathbf{s}) &\rightsquigarrow \mathcal{Q}(\mathcal{S} + \mathbf{s}) + \mathcal{G}(\mathcal{S} + \mathbf{s}) - \mathcal{Q}(\mathcal{S}) - \mathcal{G}(\mathcal{S}) \\ &= \mathcal{Q}(\mathbf{s}) - \mathcal{S}'\mathbf{V}_0\mathbf{s} + \mathcal{G}(\mathcal{S} + \mathbf{s}) - \mathcal{G}(\mathcal{S}) = \mathcal{Q}^*(\mathbf{s}). \end{aligned}$$

Unlike \mathcal{Q} , the process \mathcal{Q}^* is random in general. Indeed, \mathcal{Q}^* coincides with \mathcal{Q} if and only if \mathcal{G} is of the form $\mathcal{G}(\mathbf{s}) = \mathbf{s}'\dot{\mathcal{G}}$ for some zero-mean Gaussian vector $\dot{\mathcal{G}}$. As a consequence, the bootstrap analogue of (4) fails unless $\mathcal{G}(\mathbf{s}) = \mathbf{s}'\dot{\mathcal{G}}$. Because $\mathcal{S} = \mathbf{V}_0^{-1}\dot{\mathcal{G}}$ when $\mathcal{G}(\mathbf{s}) = \mathbf{s}'\dot{\mathcal{G}}$, the limiting distribution of $r_n(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0)$ is Gaussian whenever the bootstrap analogue of (4) holds. In perfect qualitative agreement with Fang and Santos (2016, Corollary 3.1), we therefore find that Gaussianity of the limiting distribution of $r_n(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0)$ is a “heuristically necessary” condition for consistency of the nonparametric bootstrap.

The rate of convergence of $\hat{\boldsymbol{\theta}}_n$ is also a “red herring” in our main constructive results. In particular, the reshaped function \tilde{m}_n employed in the construction of $\tilde{\boldsymbol{\theta}}_n^*$ makes no use of the fact that $\hat{\boldsymbol{\theta}}_n$ was assumed to be cube root consistent. Moreover, because

$$\sup_{\mathbf{t} \in \mathbb{R}^d} \left| \mathbb{P}^*[r(\tilde{\boldsymbol{\theta}}_n^* - \hat{\boldsymbol{\theta}}_n) \leq \mathbf{t}] - \mathbb{P}[r(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0) \leq \mathbf{t}] \right| = \sup_{\mathbf{t} \in \mathbb{R}^d} \left| \mathbb{P}^*[\tilde{\boldsymbol{\theta}}_n^* - \hat{\boldsymbol{\theta}}_n \leq \mathbf{t}] - \mathbb{P}[\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0 \leq \mathbf{t}] \right|, \quad \forall r > 0,$$

the factor $\sqrt[3]{n}$ has been included in the statement of the bootstrap consistency result (9) itself solely to facilitate its interpretation. As a consequence, also in the more general setting of the current discussion one would expect that if $\tilde{\mathbf{V}}_n \rightarrow_{\mathbb{P}} \mathbf{V}_0$, then our bootstrap-based approximation will be consistent in the sense that

$$\sup_{\mathbf{t} \in \mathbb{R}^d} \left| \mathbb{P}^*[r_n(\tilde{\boldsymbol{\theta}}_n^* - \hat{\boldsymbol{\theta}}_n) \leq \mathbf{t}] - \mathbb{P}[r_n(\hat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0) \leq \mathbf{t}] \right| \rightarrow_{\mathbb{P}} 0.$$

5.2 The Grenander estimator

Another prominent example of a cube root consistent estimator is the isotonic density estimator of Grenander (1956). The asymptotic properties of Grenander estimator have been studied by Prakasa Rao (1969), Groeneboom (1985), and Kim and Pollard (1990), among others. More recently, inconsistency of standard bootstrap-based approximations to the distribution of Grenander estimator has been pointed out by Kosorok (2008) and Sen, Banerjee, and Woodroffe (2010), among others.

To describe the Grenander estimator, suppose $\{z_i : 1 \leq i \leq n\}$ is a random sample of a continuously distributed nonnegative random variable z , whose Lebesgue density f is continuous and non-increasing on $[0, \infty)$. For a given evaluation point $x_0 \in (0, \infty)$, the Grenander estimator $\hat{f}_n(x_0)$ of $f(x_0)$ is the left derivative at x_0 of the least concave majorant (LCM) of $\hat{F}_n(\cdot) = n^{-1} \sum_{i=1}^n \mathbb{1}(x_i \leq \cdot)$, the empirical distribution function.

Although not an M -estimator of the form considered in this note, the Grenander estimator can be handled by adapting some of the ideas used to obtain the results for M -estimators. Assuming f is differentiable at x_0 with strictly negative derivative $f'(x_0)$ and letting \mathcal{W} denote a two-sided Wiener process with $\mathcal{W}(0) = 0$, it is well known that

$$\sqrt[3]{n}(\hat{f}_n(x_0) - f(x_0)) \rightsquigarrow \sqrt[3]{|4f'(x_0)f(x_0)|} \operatorname{argmax}_{s \in \mathbb{R}} \{\mathcal{W}(s) - s^2\},$$

a result that can be obtained by using empirical process methods similar to those used when deriving (2); for a textbook treatment, see van der Vaart and Wellner (1996, Example 3.2.14).

A natural bootstrap analogue of $\hat{f}_n(x_0)$ is given by $\hat{f}_n^*(x_0)$, the left derivative at x_0 of the LCM of $\hat{F}_n^*(\cdot) = n^{-1} \sum_{i=1}^n \mathbb{1}(x_i^* \leq \cdot)$, where $\{x_i^* : 1 \leq i \leq n\}$ denotes a random sample from the empirical distribution of $\{x_i : 1 \leq i \leq n\}$. Kosorok (2008) and Sen, Banerjee, and Woodroffe (2010) proved, among other things, that the distributional approximation

$$\mathbb{P}[\sqrt[3]{n}(\hat{f}_n(x_0) - f(x_0)) \leq t] \approx \mathbb{P}^*[\sqrt[3]{n}(\hat{f}_n^*(x_0) - \hat{f}_n(x_0)) \leq t]$$

is inconsistent. Once again, this inconsistency can be attributed to the fact that the bootstrap approximation uses an estimator based on a process whose mean function, under the bootstrap

distribution, fails to replicate that of its population counterpart. To be specific, the empirical distribution function \hat{F}_n , upon which $\hat{f}_n(x_0)$ is based, is an unbiased estimator of the cumulative distribution function F . Around x_0 , that function admits a quadratic approximation of the form

$$F(x) \approx F(x_0) + f(x_0)(x - x_0) + \frac{1}{2}f'(x_0)(x - x_0)^2.$$

The bootstrap mean of the function function \hat{F}_n^* upon which $\hat{f}_n^*(x_0)$ is based is given by \hat{F}_n . Unlike F , the function \hat{F}_n does not admit a quadratic approximation around x_0 , a fact to which [Sen, Banerjee, and Woodroffe \(2010, p. 1968\)](#) attribute the inconsistency of the bootstrap. Adapting the logic used to motivate the functional form of \tilde{m}_n , a reshaped version of \hat{F}_n^* is given by

$$\tilde{F}_n^*(x) = \hat{F}_n^*(x) - \hat{F}_n(x) + \hat{F}_n(x_0) + \hat{f}_n(x_0)(x - x_0) + \frac{1}{2}\tilde{f}'_n(x_0)(x - x_0)^2,$$

where $\tilde{f}'_n(x_0)$ is an estimator of $f'(x_0)$. Letting $\tilde{f}_n^*(x_0)$ denote the left derivative at x_0 of the LCM of \tilde{F}_n^* , the hope is of course that the approximation

$$\mathbb{P}[\sqrt[3]{n}(\hat{f}_n(x_0) - f(x_0)) \leq t] \approx \mathbb{P}^*[\sqrt[3]{n}(\tilde{f}_n^*(x_0) - \hat{f}_n(x_0)) \leq t]$$

is consistent under mild conditions on $\tilde{f}'_n(x_0)$. In perfect analogy with [Theorem 1](#), it turns out that this is indeed the case.

Theorem 3 *Suppose that f is differentiable at x_0 with strictly negative derivative $f'(x_0)$. If $\tilde{f}'_n(x_0) \rightarrow_{\mathbb{P}} f'(x_0)$, then*

$$\sup_{t \in \mathbb{R}} \left| \mathbb{P}^*[\sqrt[3]{n}(\tilde{f}_n^*(x_0) - \hat{f}_n(x_0)) \leq t] - \mathbb{P}[\sqrt[3]{n}(\hat{f}_n(x_0) - f(x_0)) \leq t] \right| \rightarrow_{\mathbb{P}} 0.$$

The consistency requirement $\tilde{f}'_n(x_0) \rightarrow_{\mathbb{P}} f'(x_0)$ associated with the Grenander estimator is mild, being met by standard nonparametric estimators. One obvious choice, and one we have found to work well in simulations, is a standard kernel density estimator with its corresponding plug-in MSE-optimal bandwidth selector; for details, see the supplemental appendix. Another possibility, also discussed in the supplemental appendix, is to use an estimator based on numerical derivatives.

We investigate the finite sample properties of confidence intervals for $f(x_0)$ constructed using the bootstrap-based distributional approximation whose consistency was established in Theorem 3. We employ the DGPs and simulation setting previously considered in Sen, Banerjee, and Woodroffe (2010). This, as in the case of the Maximum Score estimator discussed previously, allows for a direct comparison with other bootstrap-based inference methods and their numerical performance already studied in previous work available in the literature.

We estimate $f(x_0)$ at the evaluation point $x_0 = 1$ using a random sample of observations, where three distinct distributions are considered: DGP 1 sets $x \sim \text{Exponential}(1)$, DGP 2 sets $x \sim |\text{Normal}(0, 1)|$, and DGP 3 sets $x \sim |\mathcal{T}_3|$. As in the case of the Maximum Score example, the Monte Carlo experiment employs a sample size $n = 1,000$ with $B = 2,000$ bootstrap replications and $S = 2,000$ simulations, and compares three types of bootstrap-based inference procedures: the standard non-parametric bootstrap, m -out-of- n bootstrap, and our proposed method using two distinct estimators of $f'(x_0)$ (plug-in and numerical derivative).

Table 2 presents the numerical results for this example. We continue to report empirical coverage for nominal 95% confidence intervals and their average interval length. For the case of our proposed procedures, we again investigate their performance using both (i) a grid of fixed tuning parameter value (derivative step/bandwidth) and (ii) infeasible and feasible AMSE-optimal choice of tuning parameter. Also in this case, the numerical evidence is very encouraging. Our proposed bootstrap-based inference method leads to confidence intervals with excellent empirical coverage and average interval length, outperforming both the standard nonparametric bootstrap (which is inconsistent) and the m -out-of- n bootstrap (which is consistent). In particular, in this example, the plug-in method employs an off-the-shelf kernel derivative estimator, which in this case leads to confidence intervals that are very robust (i.e., insensitive) to the choice of bandwidth. Furthermore, when the corresponding feasible off-the-shelf MSE-optimal bandwidth is used, the resulting confidence intervals continue to perform excellently. Finally, the generic numerical derivative estimator also leads to very good performance of bootstrap-based infeasible and feasible confidence intervals.

In sum, this example provides a second numerical illustration of the very good finite sample performance of inference based on our proposed bootstrap-based distributional approximation for cube root consistent estimators.

6 Conclusion

We introduced a new bootstrap-based distributional approximation for M-estimators having cube root asymptotic distributions. Our method employs the standard nonparametric bootstrap but with a carefully chosen reshaping of the objective function to ensure a valid distributional approximation. We applied our results to two leading examples of $\sqrt[3]{n}$ -consistent estimators, Maximum Score and Isotonic Density, and in both cases simulation evidence showed excellent performance in terms of empirical coverage and average interval length of the resulting confidence intervals estimators. We also discussed how our main ideas could be applied to general parametric M-estimators with arbitrary convergence rates.

References

- ABREVAYA, J., AND J. HUANG (2005): “On the Bootstrap of the Maximum Score Estimator,” *Econometrica*, 73(4), 1175–1204.
- CHERNOFF, H. (1964): “Estimation of the Mode,” *Annals of the Institute of Statistical Mathematics*, 16(1), 31–41.
- DELGADO, M. A., J. M. RODRIGUEZ-POO, AND M. WOLF (2001): “Subsampling Inference in Cube Root Asymptotics with an Application to Manski’s Maximum Score Estimator,” *Economics Letters*, 73(2), 241–250.
- DÜMBGEN, L. (1993): “On Nondifferentiable Functions and the Bootstrap,” *Probability Theory and Related Fields*, 95, 125–140.
- FANG, Z., AND A. SANTOS (2016): “Inference on Directionally Differentiable Functions,” arXiv:1404.3763.
- GRENANDER, U. (1956): “On the Theory of Mortality Measurement: Part II,” *Scandinavian Actuarial Journal*, 39(2), 125–153.
- GROENEBOOM, P. (1985): “Estimating a Monotone Density,” in *Proceedings of the Berkeley Conference in Honor of Jerzy Neyman and Jack Kiefer*, ed. by L. M. Le Cam, R. A. Olshen, and C.-S. Cheng. Hayward, CA: Institute of Mathematical Statistics, 539–555.
- HONG, H., AND J. LI (2017): “The Numerical Delta Method and Bootstrap,” Working Paper, Stanford University.
- HOROWITZ, J. L. (1992): “A Smoothed Maximum Score Estimator for the Binary Response Model,” *Econometrica*, 60(3), 505–531.
- (2002): “Bootstrap Critical Values for Tests Based on the Smoothed Maximum Score Estimator,” *Journal of Econometrics*, 111(2), 141–167.
- KIM, J., AND D. POLLARD (1990): “Cube Root Asymptotics,” *Annals of Statistics*, 18(1), 191–219.

- KOSOROK, M. R. (2008): “Bootstrapping the Grenander Estimator,” in *Beyond Parametrics in Interdisciplinary Research: Festschrift in Honor of Professor Pranab K. Sen*, pp. 282–292. Institute of Mathematical Statistics.
- LÉGER, C., AND B. MACGIBBON (2006): “On the Bootstrap in Cube Root Asymptotics,” *Canadian Journal of Statistics*, 34(1), 29–44.
- MANSKI, C. F. (1975): “Maximum Score Estimation of the Stochastic Utility Model of Choice,” *Journal of econometrics*, 3(3), 205–228.
- PATRA, R. K., E. SELJO, AND B. SEN (2015): “A Consistent Bootstrap Procedure for the Maximum Score Estimator,” arXiv:1105.1976.
- POLITIS, D. N., AND J. P. ROMANO (1994): “Large Sample Confidence Regions Based on Subsamples under Minimal Assumptions,” *Annals of Statistics*, 22(4), 2031–2050.
- PRAKASA RAO, B. (1969): “Estimation of a Unimodal Density,” *Sankhya: The Indian Journal of Statistics, Series A*, pp. 23–36.
- SEN, B., M. BANERJEE, AND M. WOODROOFE (2010): “Inconsistency of Bootstrap: The Grenander Estimator,” *Annals of Statistics*, 38(4), 1953–1977.
- VAN DER VAART, A. W., AND J. A. WELLNER (1996): *Weak Convergence and Empirical Processes*. Springer.

Table 1: Simulations, Maximum Score Estimator, 95% Confidence Intervals.

(a) $n = 1,000$, $S = 2,000$, $B = 2,000$

	DGP 1			DGP 2			DGP 3		
	h, ϵ	Coverage	Length	h, ϵ	Coverage	Length	h, ϵ	Coverage	Length
Standard		0.639	0.475		0.645	0.480		0.640	0.247
m-out-of-n									
$m = \lceil n^{1/2} \rceil$		0.998	1.702		0.997	1.754		0.999	1.900
$m = \lceil n^{2/3} \rceil$		0.979	1.189		0.979	1.223		0.985	0.728
$m = \lceil n^{4/5} \rceil$		0.902	0.824		0.894	0.839		0.904	0.447
Plug-in: $\tilde{\mathbf{V}}_n^{\text{MS}}$									
$0.7 \cdot h_{\text{MSE}}$	0.434	0.941	0.501	0.406	0.947	0.513	0.105	0.904	0.256
$0.8 \cdot h_{\text{MSE}}$	0.496	0.946	0.503	0.464	0.952	0.516	0.120	0.917	0.260
$0.9 \cdot h_{\text{MSE}}$	0.558	0.951	0.506	0.522	0.951	0.518	0.135	0.930	0.267
$1.0 \cdot h_{\text{MSE}}$	0.620	0.954	0.510	0.580	0.952	0.522	0.150	0.941	0.273
$1.1 \cdot h_{\text{MSE}}$	0.682	0.959	0.515	0.638	0.955	0.526	0.165	0.948	0.281
$1.2 \cdot h_{\text{MSE}}$	0.744	0.961	0.522	0.696	0.959	0.532	0.180	0.958	0.288
$1.3 \cdot h_{\text{MSE}}$	0.806	0.962	0.531	0.754	0.960	0.539	0.195	0.966	0.296
h_{AMSE}	0.385	0.938	0.499	0.367	0.941	0.510	0.119	0.917	0.260
\hat{h}_{AMSE}	0.446	0.947	0.509	0.415	0.949	0.518	0.155	0.941	0.275
Num Deriv: $\tilde{\mathbf{V}}_n^{\text{ND}}$									
$0.7 \cdot \epsilon_{\text{MSE}}$	0.980	0.912	0.431	0.904	0.891	0.422	0.203	0.864	0.216
$0.8 \cdot \epsilon_{\text{MSE}}$	1.120	0.922	0.442	1.033	0.897	0.432	0.232	0.888	0.228
$0.9 \cdot \epsilon_{\text{MSE}}$	1.260	0.929	0.460	1.163	0.909	0.448	0.261	0.904	0.238
$1.0 \cdot \epsilon_{\text{MSE}}$	1.400	0.939	0.484	1.292	0.919	0.469	0.290	0.917	0.248
$1.1 \cdot \epsilon_{\text{MSE}}$	1.540	0.943	0.514	1.421	0.928	0.497	0.319	0.928	0.257
$1.2 \cdot \epsilon_{\text{MSE}}$	1.680	0.948	0.549	1.550	0.932	0.531	0.348	0.939	0.265
$1.3 \cdot \epsilon_{\text{MSE}}$	1.820	0.955	0.590	1.679	0.935	0.568	0.377	0.947	0.274
ϵ_{AMSE}	0.483	0.878	0.410	0.476	0.871	0.412	0.216	0.877	0.221
$\hat{\epsilon}_{\text{AMSE}}$	0.518	0.877	0.414	0.513	0.884	0.418	0.368	0.932	0.269

Notes:

(i) Panel **Standard** refers to standard nonparametric bootstrap, Panel **m-out-of-n** refers to m -out-of- n nonparametric bootstrap with subsample m , Panel **Plug-in: $\tilde{\mathbf{V}}_n^{\text{MS}}$** refers to our proposed bootstrap-based implemented using the example-specific plug-in drift estimator, and Panel **Num Deriv: $\tilde{\mathbf{V}}_n^{\text{ND}}$** refers to our proposed bootstrap-based implemented using the generic numerical derivative drift estimator.

(ii) Column “ h, ϵ ” reports tuning parameter value used or average across simulations when estimated, and Columns “Coverage” and “Length” report empirical coverage and average length of bootstrap-based 95% percentile confidence intervals, respectively.

(iii) h_{MSE} and ϵ_{MSE} correspond to the simulation MSE-optimal choices, h_{AMSE} and ϵ_{AMSE} correspond to the AMSE-optimal choices, and \hat{h}_{AMSE} and $\hat{\epsilon}_{\text{AMSE}}$ correspond to the ROT feasible implementation of \hat{h}_{AMSE} and $\hat{\epsilon}_{\text{AMSE}}$ described in the supplemental appendix.

Table 2: Simulations, Isotonic Density Estimator, 95% Confidence Intervals.

(a) $n = 1,000$, $S = 2,000$, $B = 2,000$

	DGP 1			DGP 2			DGP 3		
	h, ϵ	Coverage	Length	h, ϵ	Coverage	Length	h, ϵ	Coverage	Length
Standard		0.828	0.146		0.808	0.172		0.821	0.155
m-out-of-n									
$m = \lceil n^{1/2} \rceil$		1.000	0.438		0.995	0.495		0.998	0.452
$m = \lceil n^{2/3} \rceil$		0.989	0.314		0.979	0.360		0.989	0.328
$m = \lceil n^{4/5} \rceil$		0.953	0.235		0.937	0.274		0.948	0.248
Plug-in: $\tilde{\mathbf{V}}_n^{\text{ID}}$									
$0.7 \cdot h_{\text{MSE}}$	0.264	0.955	0.157	0.202	0.947	0.183	0.209	0.957	0.165
$0.8 \cdot h_{\text{MSE}}$	0.302	0.954	0.157	0.231	0.946	0.182	0.239	0.952	0.165
$0.9 \cdot h_{\text{MSE}}$	0.339	0.951	0.156	0.260	0.944	0.181	0.269	0.949	0.164
$1.0 \cdot h_{\text{MSE}}$	0.377	0.949	0.154	0.289	0.941	0.180	0.299	0.948	0.163
$1.1 \cdot h_{\text{MSE}}$	0.415	0.940	0.151	0.318	0.938	0.178	0.329	0.944	0.161
$1.2 \cdot h_{\text{MSE}}$	0.452	0.934	0.147	0.347	0.934	0.176	0.359	0.939	0.158
$1.3 \cdot h_{\text{MSE}}$	0.490	0.922	0.142	0.376	0.928	0.173	0.389	0.935	0.155
h_{AMSE}	0.380	0.949	0.154	0.300	0.940	0.180	0.333	0.943	0.161
\hat{h}_{AMSE}	0.364	0.950	0.155	0.290	0.941	0.180	0.401	0.930	0.154
Num Deriv: $\tilde{\mathbf{V}}_n^{\text{ND}}$									
$0.7 \cdot \epsilon_{\text{MSE}}$	0.726	0.954	0.158	0.527	0.947	0.183	0.554	0.952	0.165
$0.8 \cdot \epsilon_{\text{MSE}}$	0.830	0.956	0.159	0.602	0.947	0.182	0.633	0.950	0.164
$0.9 \cdot \epsilon_{\text{MSE}}$	0.933	0.956	0.160	0.678	0.944	0.181	0.712	0.949	0.163
$1.0 \cdot \epsilon_{\text{MSE}}$	1.037	0.956	0.159	0.753	0.942	0.180	0.791	0.948	0.162
$1.1 \cdot \epsilon_{\text{MSE}}$	1.141	0.955	0.159	0.828	0.940	0.179	0.870	0.946	0.161
$1.2 \cdot \epsilon_{\text{MSE}}$	1.244	0.956	0.160	0.904	0.936	0.177	0.949	0.943	0.160
$1.3 \cdot \epsilon_{\text{MSE}}$	1.348	0.960	0.163	0.979	0.935	0.176	1.028	0.940	0.159
ϵ_{AMSE}	0.927	0.956	0.160	0.731	0.943	0.180	0.812	0.948	0.162
$\hat{\epsilon}_{\text{AMSE}}$	0.888	0.956	0.159	0.708	0.943	0.181	0.978	0.942	0.159

Notes:

(i) Panel **Standard** refers to standard nonparametric bootstrap, Panel **m-out-of-n** refers to m -out-of- n nonparametric bootstrap with subsample m , Panel **Plug-in: $\tilde{\mathbf{V}}_n^{\text{ID}}$** refers to our proposed bootstrap-based implemented using the example-specific plug-in drift estimator, and Panel **Num Deriv: $\tilde{\mathbf{V}}_n^{\text{ND}}$** refers to our proposed bootstrap-based implemented using the generic numerical derivative drift estimator.

(ii) Column “ h, ϵ ” reports tuning parameter value used or average across simulations when estimated, and Columns “Coverage” and “Length” report empirical coverage and average length of bootstrap-based 95% percentile confidence intervals, respectively.

(iii) h_{MSE} and ϵ_{MSE} correspond to the simulation MSE-optimal choices, h_{AMSE} and ϵ_{AMSE} correspond to the AMSE-optimal choices, and \hat{h}_{AMSE} and $\hat{\epsilon}_{\text{AMSE}}$ correspond to the ROT feasible implementation of \hat{h}_{AMSE} and $\hat{\epsilon}_{\text{AMSE}}$ described in the supplemental appendix.