

NBER WORKING PAPER SERIES

A SCORE BASED APPROACH TO WILD BOOTSTRAP INFERENCE

Patrick M. Kline
Andres Santos

Working Paper 16127
<http://www.nber.org/papers/w16127>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2010

We thank Justin McCrary, Nicolas Lepage-Saucier, Graham Elliott, Michael Jansson as well as the editor Jason Abrevaya and two anonymous referees for comments that helped greatly improve the paper. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2010 by Patrick M. Kline and Andres Santos. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

A Score Based Approach to Wild Bootstrap Inference
Patrick M. Kline and Andres Santos
NBER Working Paper No. 16127
June 2010, Revised May 2011
JEL No. C01,C12

ABSTRACT

We propose a generalization of the wild bootstrap of Wu (1986) and Liu (1988) based upon perturbing the scores of M-estimators. This "score bootstrap" procedure avoids recomputing the estimator in each bootstrap iteration, making it substantially less costly to compute than the conventional nonparametric bootstrap, particularly in complex nonlinear models. Despite this computational advantage, in the linear model, the score bootstrap studentized test statistic is equivalent to that of the conventional wild bootstrap up to order $O_p(n^{-1})$. We establish the consistency of the procedure for Wald and Lagrange Multiplier type tests and tests of moment restrictions for a wide class of M-estimators under clustering and potential misspecification. In an extensive series of Monte Carlo experiments we find that the performance of the score bootstrap is comparable to competing approaches despite its computational savings.

Patrick M. Kline
Department of Economics
UC, Berkeley
508-1 Evans Hall #3880
Berkeley, CA 94720
and NBER
pkline@econ.berkeley.edu

Andres Santos
Department of Economics
9500 Gilman Drive
La Jolla, CA 92093-0508
a2santos@ucsd.edu

1 Introduction

The bootstrap of Efron (1979) has become a standard tool for conducting inference with economic data. Among the numerous variants of the original bootstrap, the so-called “wild” bootstrap of Wu (1986) and Liu (1988) has been found to yield dramatic improvements in the ability to control the size of Wald tests of OLS regression coefficients in small samples (Mammen (1993), Horowitz (1997, 2001), Cameron, Gelbach, and Miller (2008)).

Originally proposed as an alternative to the residual bootstrap of Freedman (1981), the wild bootstrap has often been interpreted as a procedure that resamples residuals in a manner that captures any heteroscedasticity in the underlying errors. Perhaps for this reason, the applications and extensions of the wild bootstrap have largely been limited to linear models where residuals are straightforward to obtain; see for example Hardle and Mammen (1993) for nonparametric regression, You and Chen (2006) for partially linear regression, Davidson and MacKinnon (2010) for IV regression and Cavaliere and Taylor (2008) for unit root inference.

We propose a new variant of the wild bootstrap (the “score” bootstrap) which perturbs the fitted score contributions of an M-estimator with i.i.d. weights conditional on a fixed Hessian. In the linear model, our score bootstrap procedure is numerically equivalent to the conventional wild bootstrap for unstudentized statistics and higher order equivalent for studentized ones. However, in contrast to the wild bootstrap, our approach is easily adapted to estimators without conventional residuals and avoids recomputing the estimator in each bootstrap iteration. As a result, the score bootstrap possesses an important advantage over existing bootstraps in settings where the model is computationally expensive to estimate or poorly behaved in a subset of the bootstrap draws. For example, computational problems often arise in small samples even in simple probit or logit models where, for some nonparametric bootstrap draws, the estimator cannot be computed.¹

The score bootstrap is closely related to several existing bootstrap procedures in the literature. Most notably, it bears a close relationship to the estimating equation bootstrap of Hu and Zidek (1995) who propose resampling score contributions in the linear model conditional on a fixed Hessian. Hu and Kalbfleisch (2000) generalize this approach to nonlinear models by resampling both score and Hessian contributions evaluated at the estimated parameter vector. In the case of score or Lagrange Multiplier tests, our approach can be interpreted as a wild bootstrap analogue to their pairs resampling procedure. Also related is the generalized bootstrap of Chatterjee and Bose (2005)

¹See Kaido and Santos (2011) for a recent application of the score bootstrap to estimating the asymptotic distribution of set estimators, an area where computational considerations are particularly important.

which perturbs the objective function of an M-estimator with i.i.d. weights. This approach is closer to the weighted bootstrap (e.g. Barbe and Bertail (1995), Ma and Kosorok (2005)) than the wild bootstrap and, in contrast to the score bootstrap, requires reoptimization of the estimator under each perturbation of the criterion function. Finally, our procedure has an interpretation as a variant of the k-step bootstrap procedure of Davidson and MacKinnon (1999a) which involves taking a finite number of Newton steps towards optimization of an M-estimator in a bootstrap sample. Andrews (2002) showed that this procedure yields an Edgeworth refinement depending on the number of optimization steps taken. Like the conventional nonparametric bootstrap however, the k-step procedure may be difficult to compute if, in some bootstrap samples, the Hessian is poorly behaved or of less than full rank, problems which the score bootstrap avoids.

We provide results establishing the consistency of the score bootstrap for a broad class of test statistics under weak regularity conditions and in the presence of potential misspecification. Our framework is shown to encompass Wald and Lagrange Multiplier (LM) tests as well as tests of moment restrictions. To assess the empirical relevance of these theoretical results, we conduct an extensive series of Monte Carlo experiments comparing the performance of several different bootstrap procedures in settings with clustered data. Our focus on clustered data is motivated by the prevalence of such settings in applied work and a large literature (e.g. Bertrand, Duflo, and Mullainathan (2004), Wooldridge (2003), Donald and Lang (2007), Cameron, Gelbach, and Miller (2008)) finding that asymptotic cluster robust methods often perform poorly in small samples. We find that variants of our proposed score based bootstrap substantially outperform analytical cluster robust methods. The performance of these procedures is also comparable to that of competing bootstrap methods, despite their large difference in computational cost.

The remainder of the paper is structured as follows: Section 2 reviews the wild bootstrap, while Section 3 introduces the score bootstrap and establishes its higher order equivalence. In Section 4 we develop the consistency of the score bootstrap under weak regularity conditions and illustrate its applicability to a variety of settings. Our simulation study is contained in Section 5, while Section 6 briefly concludes. All proofs are contained in the Appendix.

2 Wild Bootstrap Review

We begin by reviewing the wild bootstrap and the reasons for its consistency in the context of a linear model. A careful examination of the arguments justifying its validity provides us with the intuition necessary for developing the score bootstrap and its extension to M-estimation problems.

While there are multiple approaches to implementing the wild bootstrap, for expository purposes we focus on the original methodology developed in Liu (1988). Suppose $\{Y_i, X_i\}_{i=1}^n$ is an i.i.d. sequence of random variables, with $Y_i \in \mathbf{R}$, $X_i \in \mathbf{R}^m$ and satisfying the linear relationship:

$$Y_i = X_i' \beta_0 + \epsilon_i . \quad (1)$$

Letting $\hat{\beta}$ denote the OLS estimate of β_0 and $e_i \equiv (Y_i - X_i' \hat{\beta})$ the implied residual, the wild bootstrap generates new residuals of the form $\epsilon_i^* \equiv W_i e_i$ for some randomly generated i.i.d. sequence $\{W_i\}_{i=1}^n$ that is independent of $\{Y_i, X_i\}_{i=1}^n$ and satisfies $E[W_i] = 0$ and $E[W_i^2] = 1$. Common choices of distributions for W_i include the Standard Normal, Rademacher,² and the two-point distribution advocated in Mammen (1993).³ Under these assumptions on $\{W_i\}_{i=1}^n$ it follows that:

$$E[\epsilon_i^* | \{Y_i, X_i\}_{i=1}^n] = 0 \quad E[(\epsilon_i^*)^2 | \{Y_i, X_i\}_{i=1}^n] = e_i^2 . \quad (2)$$

Hence ϵ_i^* is mean independent of $\{Y_i, X_i\}_{i=1}^n$ and, in addition, captures the pattern of heteroscedasticity found in the original sample. This property, originally noted in Wu (1986), enables the wild bootstrap to remain consistent even in the presence of heteroscedasticity or model misspecification.⁴

The wild bootstrap resampling scheme consists of generating dependent variables $\{Y_i^*\}_{i=1}^n$ by

$$Y_i^* \equiv X_i' \hat{\beta} + \epsilon_i^* \quad (3)$$

and then conducting OLS on the sample $\{Y_i^*, X_i\}_{i=1}^n$ in order to obtain a bootstrap estimate $\hat{\beta}^*$. The distribution of $\sqrt{n}(\hat{\beta}^* - \hat{\beta})$ conditional on $\{Y_i, X_i\}_{i=1}^n$ (but not on $\{W_i\}_{i=1}^n$) is then used as an estimate of the unknown distribution of $\sqrt{n}(\hat{\beta} - \beta_0)$. Since the former distribution can be computed through simulation, the wild bootstrap provides a simple way to obtain critical values for inference.

We review why the wild bootstrap is consistent by drawing from arguments in Mammen (1993). First, observe that standard OLS algebra and the relationships in (1) and (3) imply that:

$$\sqrt{n}(\hat{\beta} - \beta_0) = H_n^{-1} \frac{1}{\sqrt{n}} \sum_{i=1}^n X_i \epsilon_i \quad \sqrt{n}(\hat{\beta}^* - \hat{\beta}) = H_n^{-1} \frac{1}{\sqrt{n}} \sum_{i=1}^n X_i \epsilon_i^* , \quad (4)$$

where $H_n \equiv n^{-1} \sum_i X_i X_i'$. When $\hat{\beta}$ is viewed as the maximum likelihood estimator of a normal model, H_n is the Hessian of the likelihood, while $\sum_i X_i \epsilon_i$ is the gradient (or score) evaluated at the true parameter value β_0 . Since both the full sample score contributions ($\{X_i \epsilon_i\}_{i=1}^n$) and their bootstrap analogues ($\{X_i \epsilon_i^*\}_{i=1}^n$) are properly centered, the expressions in (4) can be expected

²A Rademacher random variable puts probability one half on the values one and negative one.

³Here W_i equals $\frac{1-\sqrt{5}}{2}$ with probability $\frac{\sqrt{5}+1}{2\sqrt{5}}$ and $\frac{\sqrt{5}+1}{2}$ with probability $1 - \frac{\sqrt{5}+1}{2\sqrt{5}}$.

⁴We refer to misspecification in model (1) as $E[\epsilon_i | X_i] \neq 0$ but $E[\epsilon_i X_i] = 0$.

to converge to a normal limit. Therefore, consistency of the wild bootstrap hinges on whether these limits are the same or, equivalently, whether the asymptotic variances agree. However, since $E[W_i^2] = 1$ and $\{W_i\}_{i=1}^n$ is independent of $\{Y_i, X_i\}_{i=1}^n$, we may write:

$$E\left[\left(\frac{1}{\sqrt{n}} \sum_{i=1}^n X_i \epsilon_i\right) \left(\frac{1}{\sqrt{n}} \sum_{i=1}^n X_i \epsilon_i\right)'\right] = E[X_i X_i' \epsilon_i^2] \quad (5)$$

$$E\left[\left(\frac{1}{\sqrt{n}} \sum_{i=1}^n X_i \epsilon_i^*\right) \left(\frac{1}{\sqrt{n}} \sum_{i=1}^n X_i \epsilon_i^*\right)' \mid \{Y_i, X_i\}_{i=1}^n\right] = \frac{1}{n} \sum_{i=1}^n X_i X_i' e_i^2, \quad (6)$$

which implies, by standard arguments, that the second moments indeed agree asymptotically. As a result, $\sqrt{n}(\hat{\beta} - \beta_0)$ and $\sqrt{n}(\hat{\beta}^* - \hat{\beta})$ converge in distribution to the same normal limit and the consistency of the wild bootstrap is immediate.

While the ability of the wild bootstrap to asymptotically match the first two moments of the full sample score provides the basis for establishing its validity, it does not elucidate why it often performs better than a normal approximation. Improvements occur when the bootstrap is able to additionally match higher moments of the score. If, for example, $E[W_i^3] = 1$, then the third moments match asymptotically and the wild bootstrap provides a refinement over the normal approximation to a studentized statistic by providing a skewness correction (Liu (1988)). The additional requirement that $E[W_i^3] = 1$ is satisfied, for example, by the weights proposed in Mammen (1993), as well as for $W_i = (V_i - 2)$ with V_i following a Gamma distribution with mean 2 and variance 1. Alternatively, the Rademacher distribution, which satisfies $E[W_i] = E[W_i^3] = 0$ and $E[W_i^2] = E[W_i^4] = 1$, is able to match the first four moments for symmetric distributions and can in such cases provide an additional refinement (Liu (1988); Davidson and Flachaire (2008)).

3 The Score Bootstrap

The wild bootstrap resampling scheme is often interpreted as a means of generating a set of bootstrap residuals mimicking the heteroscedastic nature of the true errors. However, the residuals only influence the limiting distribution of the OLS estimator through the score. Thus, we may alternatively view the wild bootstrap as creating a set of bootstrap score contributions ($\{X_i \epsilon_i^*\}_{i=1}^n$) that mimic the heteroscedastic nature of the true score contributions ($\{X_i \epsilon_i\}_{i=1}^n$). In this section, we develop the implications of this observation, which provides the basis for our procedure.

The relationship between the wild bootstrap and the score is transparent from the discussion of its consistency in Section 2. Since $\epsilon_i^* = e_i W_i$, we learn from (4) that the wild bootstrap may be interpreted as a perturbation of the score contributions ($\{X_i(Y_i - X_i' \beta)\}_{i=1}^n$) evaluated at the

estimated parameter value ($\hat{\beta}$) that leaves the Hessian ($\frac{1}{n} \sum_i X_i X_i'$) unchanged.⁵ More precisely, a numerically equivalent way to implement the wild bootstrap is given by the following algorithm:

- STEP 1: Obtain the OLS estimate $\hat{\beta}$ and generate the fitted score contributions $\{X_i(Y_i - X_i'\hat{\beta})\}_{i=1}^n$.
- STEP 2: Using random weights $\{W_i\}_{i=1}^n$ independent of $\{Y_i, X_i\}_{i=1}^n$ and satisfying $E[W_i] = 0$ and $E[W_i^2] = 1$, construct a new set of perturbed score contributions $\{X_i(Y_i - X_i'\hat{\beta})W_i\}_{i=1}^n$.
- STEP 3: Multiply the perturbed score by the inverse Hessian to obtain $H_n^{-1} \frac{1}{\sqrt{n}} \sum_i (Y_i - X_i'\hat{\beta}) X_i W_i$ and use its distribution conditional on $\{Y_i, X_i\}_{i=1}^n$ as an estimate of the distribution of $\sqrt{n}(\hat{\beta} - \beta_0)$.

Unlike the residual based view of the wild bootstrap, the score interpretation is easily generalized to more complex nonlinear models. One may simply perturb the fitted score contributions of such a model while keeping the Hessian unchanged and, provided $E[W_i] = 0$ and $E[W_i^2] = 1$, the first two moments of the perturbed score and true score will match asymptotically. Under the appropriate regularity conditions, this moment equivalence will suffice for establishing the consistency of the proposed bootstrap. For obvious reasons, we term this procedure a “score bootstrap”.

To fix ideas, the following example shows how this intuition may be applied in a nonlinear model.

Example 3.1. Consider a standard probit model in which $\epsilon_i \sim N(0, 1)$ and (Y_i, X_i) satisfy:

$$Y_i = 1\{X_i'\beta_0 \geq \epsilon_i\}, \quad (7)$$

where $1\{\cdot\}$ is the indicator function. Suppose we wish to approximate the distribution of $\sqrt{n}(\hat{\beta} - \beta_0)$, where $\hat{\beta}$ is the maximum likelihood estimate of β_0 . The log-likelihood which $\hat{\beta}$ maximizes is then:

$$L_n(\beta) \equiv \frac{1}{n} \sum_{i=1}^n \{Y_i \log(\Phi(X_i'\beta)) + (1 - Y_i)(1 - \Phi(X_i'\beta))\}, \quad (8)$$

for Φ the cdf of a standard normal random variable. Thus, the score for the model is given by:

$$S_n(\beta) \equiv \frac{1}{n} \sum_{i=1}^n s(Y_i, X_i, \beta) \quad s(y, x, \beta) \equiv \frac{\phi(x'\beta)(y - \Phi(x'\beta))}{\Phi(x'\beta)(1 - \Phi(x'\beta))} x, \quad (9)$$

where ϕ is the derivative of Φ . The principles derived from the linear model then suggest bootstrapping the distribution of $\sqrt{n}(\hat{\beta} - \beta_0)$ by: (i) Obtaining fitted score contributions $\{s(Y_i, X_i, \hat{\beta})\}_{i=1}^n$; (ii) Perturbing them by random weights $\{W_i\}_{i=1}^n$ to obtain $\{s(Y_i, X_i, \hat{\beta})W_i\}_{i=1}^n$; (iii) Multiplying the perturbed score by the inverse Hessian ($[\frac{1}{n} \sum_i \nabla_{\beta} s(Y_i, X_i, \hat{\beta})]^{-1} \frac{1}{\sqrt{n}} \sum_i s(Y_i, X_i, \hat{\beta})W_i$) and employing its distribution conditional on $\{Y_i, X_i\}_{i=1}^n$ to approximate that of $\sqrt{n}(\hat{\beta} - \beta_0)$. ■

⁵In contrast, the weighted bootstrap perturbs the score and the Hessian (Barbe and Bertail (1995)).

3.1 Higher Order Equivalence

In the linear model, the wild and score bootstrap statistics for $\sqrt{n}(\hat{\beta} - \beta_0)$ are numerically equivalent. However, in most instances the statistic of interest is studentized, since only in this context is a refinement over an analytical approximation available (Liu (1988), Horowitz (2001)). In accord with the perturbed score interpretation, it is natural to simply employ the sample variance of the perturbed score contributions for studentization. For this reason, we define the bootstrap statistics:

$$T_n^{*w} \equiv (H_n^{-1} \Sigma_n^*(\hat{\beta}^*) H_n^{-1})^{-\frac{1}{2}} \sqrt{n}(\hat{\beta}^* - \hat{\beta}) \quad T_n^{*s} \equiv (H_n^{-1} \Sigma_n^*(\hat{\beta}) H_n^{-1})^{-\frac{1}{2}} H_n^{-1} \frac{1}{\sqrt{n}} \sum_{i=1}^n X_i \epsilon_i^*, \quad (10)$$

where $\Sigma_n^*(\beta) \equiv \frac{1}{n} \sum_i X_i X_i' (Y_i^* - X_i' \beta)^2$ and T_n^{*w} and T_n^{*s} are the studentized wild and score bootstrap statistics respectively. It is important to note that in the computation of T_n^{*s} , the full sample estimator $\hat{\beta}$ is used in obtaining the standard errors, and hence calculation of $\hat{\beta}^*$ and its implied residuals is unnecessary. As a result, the score bootstrap is computationally simpler to implement than the wild bootstrap which requires obtaining bootstrap residuals.

While for the statistics in (4) the wild and score bootstraps are numerically equivalent, such a relationship fails to hold for the studentized versions. An important concern then is whether this difference is of importance and, in particular, whether the refinement of the wild bootstrap over a normal approximation (Liu (1988)) is lost due to this discrepancy. Somewhat surprisingly, the differences between the wild and score bootstrap are asymptotically negligible even at higher order. Specifically, the wild and score bootstrap statistics are asymptotically equivalent up to a higher order than that of the refinement the wild bootstrap possesses over the normal approximation. As a result, under appropriate regularity conditions, the score bootstrap not only remains consistent despite not recomputing the estimator but can in addition be expected to obtain a refinement over an analytical approximation in precisely the same instances as the wild bootstrap.

In order to establish the higher order equivalence of T_n^{*s} and T_n^{*w} , we impose the following:

Assumption 3.1. (i) $\{Y_i, X_i\}_{i=1}^n$ are i.i.d. $E[X_i \epsilon_i] = 0$, and $E[X_i X_i']$, $E[X_i X_i' \epsilon_i^2]$ are full rank; (ii) The moments $E[\|X_i\|^4]$, $E[\epsilon_i^4]$ and $E[\|X_i\|^4 \epsilon_i^4]$ are finite; (iii) $\{W_i\}_{i=1}^n$ are i.i.d., independent of $\{Y_i, X_i\}_{i=1}^n$ with $E[W_i] = 0$, $E[W_i^2] = 1$ and $E[W_i^4] < \infty$.

Let P^* and E^* denote probability and expectation conditional on $\{Y_i, X_i\}_{i=1}^n$ (but not $\{W_i\}_{i=1}^n$). Under Assumption 3.1 we can then establish the higher order equivalence of T_n^{*w} and T_n^{*s} .

Lemma 3.1. Under Assumption 3.1, $T_n^{*w} = T_n^{*s} + O_{p^*}(n^{-1})$ almost surely.

If the conditions for an Edgeworth expansion of the bootstrap statistics T_n^{*w} and T_n^{*s} are satisfied, then Lemma 3.1 implies that they can be expected to disagree only in terms of order n^{-1} or smaller; see Chapter 2.7 in Hall (1992) for such arguments.⁶ Therefore, in settings where the wild bootstrap obtains the traditional Edgeworth refinement of order $n^{-\frac{1}{2}}$ over a normal approximation, the score bootstrap should as well. Kline and Santos (2011) show that such a refinement may be available even in models in which ϵ_i is not mean independent of X_i but merely uncorrelated with it.

The higher order equivalence of T_n^{*w} and T_n^{*s} is at first glance unexpected since the score bootstrap appears to violate the usual plug-in approach of the standard bootstrap. However, this only introduces a smaller order error due to the residuals $\{\epsilon_i^*\}_{i=1}^n$ being mean independent of $\{X_i\}_{i=1}^n$ under the bootstrap distribution. Importantly, the higher order equivalence would fail to hold if the residuals $\{\epsilon_i^*\}$ were sampled in a manner under which they were merely uncorrelated with $\{X_i\}_{i=1}^n$ under the bootstrap distribution.

Remark 3.1. The bootstrap estimator $\hat{\beta}^*$ acquired from running OLS in the sample $\{Y_i^*, X_i\}$ may easily be obtained from the score bootstrap procedure by the equality:

$$\hat{\beta}^* = \hat{\beta} + H_n^{-1} \frac{1}{n} \sum_{i=1}^n X_i \epsilon_i^*. \quad (11)$$

Note that the right hand side of equation (11) is a single Newton-Raphson step towards the wild bootstrap estimator $\hat{\beta}^*$ starting from $\hat{\beta}$. Thus, there is a close connection between our approach and the k-step bootstrap procedure studied by Davidson and MacKinnon (1999a) and Andrews (2002). However, as Lemma 3.1 suggests and our Monte Carlos confirm, in the linear model, computation of $\hat{\beta}^*$ may be avoided while still obtaining a refinement over an analytical approximation. ■

4 Inference

We turn now to establishing the validity of a score bootstrap procedure for estimating the critical values of a large class of tests. Building on our earlier discussion we consider test statistics based upon the fitted parametric scores of M-estimators, using perturbations of those scores to estimate their sampling distribution. Since this approach does not depend upon resampling of residuals, we do not distinguish between dependent and exogenous variables and instead consider a random vector $Z_i \in \mathcal{Z} \subseteq \mathbf{R}^m$ which may contain both.

⁶More precisely, Lemma 3.1 is not sufficient for showing the equivalence of the first two terms in the Edgeworth expansions. Such an equivalence can be established if $P(P^*(\|T_n^{*w} - T_n^{*s}\| > (n^{\frac{1}{2}} \log n)^{-1}) > n^{-\frac{1}{2}}) = o(n^{-\frac{1}{2}})$ and the Edgeworth expansion is valid in the bootstrap sample with probability $1 - o(n^{-\frac{1}{2}})$ (Lemma 5 Andrews (2002)).

We focus on test statistics G_n that are quadratic forms of a vector valued statistic T_n :

$$G_n \equiv T_n' T_n . \quad (12)$$

Under the null hypothesis, the underlying statistic T_n is required to be asymptotically pivotal and allow for a linear expansion. More precisely, we require that under the null hypothesis:

$$T_n = (A_n(\theta_0)\Sigma_n(\theta_0)A_n(\theta_0)')^{-\frac{1}{2}}S_n(\theta_0) + o_p(1) \quad S_n(\theta) \equiv A_n(\theta)\frac{1}{\sqrt{n}}\sum_{i=1}^n s(Z_i, \theta) , \quad (13)$$

where $A_n(\theta)$ is a $r \times k$ matrix, $s(z, \theta)$ is a $k \times 1$ vector, $\Sigma_n(\theta)$ is the sample covariance matrix of $s(Z_i, \theta)$ and θ_0 is an unknown parameter vector. In accord with the terminology we employed for the linear model, we refer to $A_n(\theta)$ as the inverse of the Hessian, $\sum_i s(Z_i, \theta)$ as the score of the model, and to $\{s(Z_i, \theta)\}_{i=1}^n$ as the score contributions. Under appropriate regularity conditions, T_n is therefore asymptotically normally distributed with identity covariance matrix and hence G_n is asymptotically Chi-squared distributed with degrees of freedom equal to the dimension of T_n . Though we only consider asymptotically pivotal statistics, our results readily extend to unstudentized ones as well.

The bootstrap statistics employed to estimate the distributions of G_n and T_n are given by:

$$G_n^* \equiv T_n'^* T_n^* \quad T_n^* \equiv (A_n(\hat{\theta})\Sigma_n^*(\hat{\theta})A_n(\hat{\theta})')^{-\frac{1}{2}}S_n^*(\hat{\theta}) \quad S_n^*(\theta) \equiv A_n(\theta)\frac{1}{\sqrt{n}}\sum_{i=1}^n s(Z_i, \theta)W_i \quad (14)$$

where $\Sigma_n^*(\theta)$ is the sample covariance matrix of $s(Z_i, \theta)W_i$ and $\hat{\theta}$ is a consistent estimator for θ_0 . As discussed in the previous section, implementation of the score bootstrap only requires calculation of the full sample estimator $\hat{\theta}$; no additional optimization is needed in each bootstrap iteration.

Remark 4.1. An alternative to perturbing the fitted score contributions by random weights is to instead resample them with replacement. Specifically, for $\tilde{s}_i \equiv s(Z_i, \hat{\theta})$, we may consider drawing from $\{\tilde{s}_i\}_{i=1}^n$ with replacement and approximating the distribution of G_n by that of \tilde{G}_n , where:

$$\tilde{G}_n \equiv \tilde{T}_n' \tilde{T}_n \quad (A_n(\hat{\theta})\tilde{\Sigma}_n(\hat{\theta})A_n(\hat{\theta})')^{-\frac{1}{2}}\tilde{S}_n(\hat{\theta}) \quad \tilde{S}_n(\theta) \equiv A_n(\theta)\frac{1}{\sqrt{n}}\sum_{i=1}^n \tilde{s}_i \quad (15)$$

for $\tilde{\Sigma}_n(\hat{\theta})$ the sample covariance matrix of $\{\tilde{s}_i\}_{i=1}^n$. In the linear model, this procedure corresponds to that of Hu and Zidek (1995). For nonlinear problems, Hu and Kalbfleisch (2000) propose a closely related approach that additionally resamples the Hessian. Specifically, when the inverse Hessian takes the form $A_n^{-1}(\hat{\theta}) = \sum_i a(Z_i, \hat{\theta})$, their bootstrap procedure samples with replacement from both $\{\tilde{s}_i\}_{i=1}^n$ and $\{a(Z_i, \hat{\theta})\}_{i=1}^n$ to obtain a new score and inverse Hessian. This may be thought of as an approximation to the traditional nonparametric (“pairs”) bootstrap. Like the pairs bootstrap however, their procedure may encounter computational difficulties when the Hessian is poorly behaved in some bootstrap draws, a problem which becomes more likely in small samples when some of the covariates are discrete. ■

4.1 Bootstrap Consistency

We establish the consistency of the bootstrap under the following set of assumptions:

Assumption 4.1. (i) $\hat{\theta} \xrightarrow{p} \theta_0$ with $\hat{\theta}, \theta_0 \in \Theta \subset \mathbf{R}^p$ and Θ a compact set; (ii) The limit point θ_0 satisfies $E[s(Z_i, \theta_0)s(Z_i, \theta_0)'] < \infty$ and the matrix $A(\theta_0)E[s(Z_i, \theta_0)s(Z_i, \theta_0)']A(\theta_0)'$ is invertible.

Assumption 4.2. (i) Under the null hypothesis T_n satisfies (13) and θ_0 is such that $E[s(Z_i, \theta_0)] = 0$; (ii) Under the alternative hypothesis $G_n \xrightarrow{p} \infty$.

Assumption 4.3. (i) $\{Z_i\}_{i=1}^n$ is i.i.d. (ii) $\sup_{\theta \in \Theta} \|A_n(\theta) - A(\theta)\|_F = o_p(1)$ with $A(\theta)$ continuous.

Assumption 4.4. (i) $\{W_i\}_{i=1}^n$ is an i.i.d. sample, independent of $\{Z_i\}_{i=1}^n$ satisfying $E[W_i] = 0$ and $E[W_i^2] = 1$; (ii) For $\text{conv}(\Theta)$ the convex hull of Θ , $s(z, \theta)$ is continuously differentiable in $\theta \in \text{conv}(\Theta)$ and $\sup_{\theta \in \text{conv}(\Theta)} \|\nabla s(z, \theta)\|_F \leq F(z)$ for some function $F(z)$ with $E[F^2(Z_i)] < \infty$.

In Assumption 4.1 we require $\hat{\theta}$ to converge in probability to some parameter vector $\theta_0 \in \Theta$ whose value may depend upon the distribution of Z_i . The compactness of the parameter space Θ is employed to verify the perturbed scores form a Donsker class. This restriction may be relaxed at the expense of a more complicated argument that exploits the consistency of $\hat{\theta}$ for a local analysis. Though in the notation we suppress such dependence, it is important to note that θ_0 may take different values under the null and alternative hypotheses. In Assumptions 4.3(ii) and 4.4(ii), $\|\cdot\|_F$ denotes the Frobenius norm. Assumptions 4.2 and 4.3, in turn enable us to establish the asymptotic behavior of G_n under the null and alternative hypotheses. Assumption 4.4(i) imposes the only requirements on the random weights $\{W_i\}_{i=1}^n$, which are the same conditions imposed for inference on the linear model in previous wild bootstrap studies. Assumption 4.4(ii) allows us to establish that the empirical process induced by functions of the form $ws(z, \theta)$ is asymptotically tight. Differentiability is not necessary for this end, but we opt to impose it due to its ease of verification and wide applicability.⁷ We note, however, that estimation of $A(\theta_0)$ may be more challenging in the non-differentiable case as this quantity usually depends on the population Hessian.

Assumptions 4.1-4.4 are sufficient for establishing the consistency of the proposed score bootstrap procedure under the null hypothesis.

Theorem 4.1. Let F_n and F_n^* be the cdfs of G_n and of G_n^* conditional on $\{Z_i\}_{i=1}^n$ and suppose that Assumptions 4.1, 4.2, 4.3 and 4.4 hold. If the null hypothesis is true, it then follows that:

$$\sup_{c \in \mathbf{R}} |F_n(c) - F_n^*(c)| = o_p(1) .$$

⁷For non-differentiable settings, the relevant condition is that $\mathcal{F} \equiv \{ws(z, \theta) : \theta \in \Theta\}$ be a Donsker class.

Theorem 4.1 justifies the use of quantiles from the distribution of G_n^* conditional on $\{Z_i\}_{i=1}^n$ as critical values. In order to control the size of the test at level α , we may employ:

$$\hat{c}_{1-\alpha} \equiv \inf\{c : P(G_n^* \leq c \mid \{Z_i\}_{i=1}^n) \geq 1 - \alpha\} . \quad (16)$$

While difficult to compute analytically, $\hat{c}_{1-\alpha}$ may easily be calculated via simulation. Employing a random number generator, B samples $\{\{W_{i1}\}_{i=1}^n, \dots, \{W_{iB}\}_{i=1}^n\}$ may be created independently of the data and used to construct B statistics $\{G_{n1}^*, \dots, G_{nB}^*\}$. Provided B is sufficiently large, the empirical $1 - \alpha$ quantile of $\{G_{n1}^*, \dots, G_{nB}^*\}$ will yield an accurate approximation to $\hat{c}_{1-\alpha}$.

While Theorem 4.1 implies that the critical value $\hat{c}_{1-\alpha}$ in conjunction with the test statistic G_n delivers size control, it does not elucidate the behavior of the test under the alternative hypothesis. As in other bootstrap procedures, the test is consistent due to the bootstrap statistic G_n^* being properly centered even under the alternative. As a result, $\hat{c}_{1-\alpha}$ converges in probability to the $1 - \alpha$ quantile of a Chi-squared distribution with r degrees of freedom, while G_n diverges to infinity. Therefore, under the alternative hypothesis, G_n is larger than $\hat{c}_{1-\alpha}$ with probability tending to one and the test rejects asymptotically. We summarize these findings in the following corollary:

Corollary 4.1. *Under Assumptions 4.1, 4.2, 4.3 and 4.4, it follows that under the null hypothesis:*

$$\lim_{n \rightarrow \infty} P(G_n \geq \hat{c}_{1-\alpha}) = 1 - \alpha ,$$

for any $0 < \alpha < 1$. Under the same assumptions, if the alternative hypothesis is instead true, then:

$$\lim_{n \rightarrow \infty} P(G_n \geq \hat{c}_{1-\alpha}) = 1 .$$

4.2 Parameter Tests

A principal application of the proposed bootstrap is in obtaining critical values for parametric hypothesis tests. We consider a general M-estimation framework in which the parameter of interest θ_M is the unique minimizer of some non-stochastic but unknown function $Q : \Theta \rightarrow \mathbf{R}$:

$$\theta_M = \arg \min_{\theta \in \Theta} Q(\theta) . \quad (17)$$

We examine the classic problem of conducting inference on a function of θ_M . Specifically, for some known and differentiable mapping $c : \Theta \rightarrow \mathbf{R}^l$ with $l \leq p$, the hypothesis we study is:

$$H_0 : c(\theta_M) = 0 \quad H_1 : c(\theta_M) \neq 0 . \quad (18)$$

Standard tests for this hypothesis include the Wald and Lagrange Multiplier (LM) tests. Intuitively, the Wald test examines whether the value of the function c evaluated at an unrestricted estimator $\hat{\theta}_M$ is statistically different from zero. In contrast, the LM test instead checks whether the first order condition of an estimator $\hat{\theta}_{M,R}$ computed imposing the null hypothesis is statistically different from zero. Therefore, in the nomenclature of Assumption 4.1(i), $\hat{\theta}$ equals $\hat{\theta}_M$ for the Wald test and $\hat{\theta}_{M,R}$ for the LM test. Similarly, if $\theta_{M,R}$ denotes the minimizer of Q over Θ subject to $c(\theta) = 0$, then θ_0 equals θ_M and $\theta_{M,R}$ under the Wald and LM test respectively.

We proceed to illustrate the details of the score bootstrap in this setting for both generalized method of moments (GMM) and maximum likelihood (ML) estimators. We focus on the analytical expressions $A_n(\theta)$ and $s(z, \theta)$ take in those specific settings and provide references for primitive conditions that ensure Assumptions 4.1, 4.2, 4.3 and 4.4 hold.

4.2.1 ML Estimators

For an ML estimator, the criterion function Q and its sample analogue Q_n are of the form:

$$Q_n(\theta) \equiv \frac{1}{n} \sum_{i=1}^n q(Z_i, \theta) \quad Q(\theta) \equiv E[q(Z_i, \theta)] , \quad (19)$$

where $q : \mathcal{Z} \times \Theta \rightarrow \mathbf{R}$ is the log-likelihood. If q is twice differentiable in θ , then we may define the Hessian $H_n(\theta) \equiv n^{-1} \sum_i \nabla^2 q(Z_i, \theta)$. For notational convenience, it is also helpful to denote the gradient of the function c evaluated at θ by $C(\theta) \equiv \nabla c(\theta)$.

Example 4.1. (Wald) The relevant Wald statistic is the studentized quadratic form of $\sqrt{n}c(\hat{\theta}_M)$, which under both the null and alternative hypothesis satisfies the asymptotic expansion:

$$\sqrt{n}(c(\hat{\theta}_M) - c(\theta_M)) = -C(\theta_M)H_n^{-1}(\theta_M) \frac{1}{\sqrt{n}} \sum_{i=1}^n \nabla q(Z_i, \theta_M) + o_p(1) . \quad (20)$$

Therefore, the Wald statistic fits the formulation in (13) with $A_n(\theta) = -C(\theta)H_n^{-1}(\theta)$ and $s(z, \theta) = \nabla q(z, \theta)$. Under the alternative hypothesis, G_n diverges to infinity since $c(\theta_M) \neq 0$. Refer to Section 3.2 in Newey and McFadden (1994) for a formal justification of these arguments. ■

Example 4.2. (LM) In this setting, the LM statistic is the normalized quadratic form of:

$$C(\hat{\theta}_{M,R})H_n^{-1}(\hat{\theta}_{M,R}) \frac{1}{\sqrt{n}} \sum_{i=1}^n \nabla q(Z_i, \hat{\theta}_{M,R}) . \quad (21)$$

Moreover, under conditions stated in Chapter 12.6.2 in Wooldridge (2002), we additionally have:

$$C(\hat{\theta}_{M,R})H_n^{-1}(\hat{\theta}_{M,R}) \frac{1}{\sqrt{n}} \sum_{i=1}^n \nabla q(Z_i, \hat{\theta}_{M,R}) = C(\theta_{M,R})H_n^{-1}(\theta_{M,R}) \frac{1}{\sqrt{n}} \sum_{i=1}^n \nabla q(Z_i, \theta_{M,R}) + o_p(1) , \quad (22)$$

under the null hypothesis. Thus, the LM statistic also fits the general formulation in (13) with $A_n(\theta) = C(\theta)H_n^{-1}(\theta)$ and $s(z, \theta) = \nabla q(z, \theta)$. Under the alternative, $G_n \xrightarrow{p} \infty$ provided $\theta_{M,R}$ is not a local minimizer of Q , $C(\theta_{M,R})E[\nabla^2 q(Z_i, \theta_{M,R})]$ is full rank and Assumption 4.1(ii) holds. ■

4.2.2 GMM Estimators

In the context of GMM estimation, the criterion function Q and its sample analogue Q_n are:

$$Q_n(\theta) \equiv \left[\frac{1}{n} \sum_{i=1}^n q(Z_i, \theta)' \right] \Omega_n \left[\frac{1}{n} \sum_{i=1}^n q(Z_i, \theta) \right] \quad Q(\theta) \equiv E[q(Z_i, \theta)'] \Omega E[q(Z_i, \theta)] , \quad (23)$$

where $q : \mathcal{Z} \times \Theta \rightarrow \mathbf{R}^k$ is a known function and Ω_n, Ω are positive definite matrices such that $\Omega_n \xrightarrow{p} \Omega$. Assuming q is differentiable in θ , let $D_n(\theta) \equiv n^{-1} \sum_i \nabla q(Z_i, \theta)$ and $B_n(\theta) \equiv D_n(\theta)' \Omega_n D_n(\theta)$. As in the discussion of ML estimators, we also denote $C(\theta) \equiv \nabla c(\theta)$.

Example 4.3. (Wald) The Wald statistic for the hypothesis in (18) is given by the studentized quadratic form of $\sqrt{nc}(\hat{\theta}_M)$. In the present context we therefore obtain an expansion of the form:

$$\sqrt{n}(c(\hat{\theta}_M) - c(\theta_M)) = -C(\theta_M)B_n^{-1}(\theta_M)D_n(\theta_M)' \Omega_n \frac{1}{\sqrt{n}} \sum_{i=1}^n q(Z_i, \theta_M) + o_p(1) , \quad (24)$$

which implies $A_n(\theta) = -C(\theta)B_n^{-1}(\theta)D_n(\theta)' \Omega_n$ and $s(z, \theta) = q(z, \theta)$ and Assumption 4.2(i) is satisfied provided $E[q(Z_i, \theta_M)] = 0$.⁸ Primitive conditions under which Assumptions 4.1-4.4 hold in this context can be found in Section 3.3 of Newey and McFadden (1994). ■

Example 4.4. (LM) In this setting, the LM test statistic is the studentized quadratic form of:

$$C(\hat{\theta}_{M,R})B_n^{-1}(\hat{\theta}_{M,R})D_n(\hat{\theta}_{M,R})' \Omega_n \frac{1}{\sqrt{n}} \sum_{i=1}^n q(Z_i, \hat{\theta}_{M,R}) , \quad (25)$$

which, as shown in Section 9.1 of Newey and McFadden (1994), is asymptotically equivalent to:

$$C(\theta_{M,R})B_n^{-1}(\theta_{M,R})D_n(\theta_{M,R})' \Omega_n \frac{1}{\sqrt{n}} \sum_{i=1}^n q(Z_i, \theta_{M,R}) \quad (26)$$

under the null hypothesis. Hence, $A_n(\theta) = C(\theta)B_n^{-1}(\theta)D_n(\theta)' \Omega_n$ and $s(z, \theta) = q(z, \theta)$. ■

⁸Notice this is trivially satisfied in a just identified system. The extension to overidentified models in which $E[q(Z_i, \theta_M)] \neq 0$ but $E[\nabla q(Z_i, \theta_M)'] \Omega E[q(Z_i, \theta_M)] = 0$ can be accomplished by letting $s(z, \theta)$ depend on n and setting $s_n(z, \theta) = D_n(\theta)' \Omega_n q(z, \theta)$. Though straightforward to establish, we do not pursue such an extension.

4.3 Moment Restrictions

An additional application of the bootstrap procedure we consider is for testing the hypothesis:

$$H_0 : E[m(Z_i, \theta_M)] = 0 \quad H_1 : E[m(Z_i, \theta_M)] \neq 0 , \quad (27)$$

where $m : \mathcal{Z} \times \Theta \rightarrow \mathbf{R}^l$ is a known function and θ_M is the minimizer of some unknown non-stochastic $Q : \Theta \rightarrow \mathbf{R}$. Such restrictions arise, for example, in tests of proper model specification and hypotheses regarding average marginal effects in nonlinear models. As in Section 4.2, the specific nature of the bootstrap statistic is dependent on whether Q is as in (19) (ML) or as in (23) (GMM). For brevity, we focus on the former, though the extension to GMM can be readily derived following manipulations analogous to those in Example 4.3.

The Wald test statistic for the hypothesis in (27) is based on the studentized plug-in estimator:

$$\frac{1}{\sqrt{n}} \sum_{i=1}^n m(Z_i, \hat{\theta}_M) , \quad (28)$$

where $\hat{\theta}_M$ is in this case the unconstrained minimizer of Q_n on Θ . Hence, in this setting θ_0 equals θ_M and $\hat{\theta}$ equals $\hat{\theta}_M$ in the notation of Assumption 4.1(i). Obtaining an expansion for T_n as in (13) is straightforward provided m and q are once and twice continuously differentiable in θ respectively. Defining the gradient $M_n(\theta) \equiv n^{-1} \sum_i \nabla m(Z_i, \theta)$ and Hessian $H_n(\theta) \equiv n^{-1} \sum_i \nabla^2 q(Z_i, \theta)$, standard arguments imply that under the null hypothesis:

$$\frac{1}{\sqrt{n}} \sum_{i=1}^n m(Z_i, \hat{\theta}_M) = \frac{1}{\sqrt{n}} \sum_{i=1}^n m(Z_i, \theta_M) - M_n(\theta_M) H_n^{-1}(\theta_M) \frac{1}{\sqrt{n}} \sum_{i=1}^n \nabla q(Z_i, \theta_M) + o_p(1) ; \quad (29)$$

see Newey (1985a) for primitive conditions for (29). Thus, in this setting $s(z, \theta)$ and $A_n(\theta)$ are:

$$s(z, \theta) = \begin{pmatrix} m(z, \theta) \\ \nabla q(z, \theta) \end{pmatrix} \quad A_n(\theta) = \begin{bmatrix} I & - M_n(\theta) H_n^{-1}(\theta) \end{bmatrix} . \quad (30)$$

Moreover, if θ_M is an interior point of Θ , then $E[\nabla q(Z_i, \theta_M)] = 0$ because θ_M minimizes Q on Θ . Therefore, $G_n \xrightarrow{p} \infty$ under the alternative hypothesis due to $E[m(Z_i, \theta_M)] \neq 0$.

Remark 4.2. Similar manipulations may be employed to show the score bootstrap can be applied to Wald tests in two stage parametric estimation problems. Unlike the nonparametric bootstrap, however, such a procedure would require an analytical derivation of the effect of the first stage estimator on the influence function of the second estimator. ■

4.3.1 ML Specification Tests

A prominent application of hypotheses as in (27) is in model specification testing. In particular, this setting encompasses moment based specification tests (“m-tests”) for maximum likelihood models, as considered in White (1982, 1994), Newey (1985b) and Tauchen (1985).⁹ Computations are simplified for ML models by the generalized information matrix equality, which implies:

$$E[\nabla^2 q(Z_i, \theta_M)] = -E[\nabla q(Z_i, \theta_M) \nabla q(Z_i, \theta_M)'] \quad E[\nabla m(Z_i, \theta_M)] = -E[m(Z_i, \theta_M) \nabla q(Z_i, \theta_M)'] .$$

For example, as noted in Chesher (1984) and Newey (1985b), computation of the Wald test statistic for the null hypothesis in (27) can be performed through the auxiliary regression:

$$1 = m(Z_i, \hat{\theta}_M)' \delta + \nabla q(Z_i, \hat{\theta}_M)' \gamma + \epsilon_i . \quad (31)$$

Equation (31) is often termed the Outer Product of Gradient (OPG) regression form for moment tests; see Chapter 15.2 in Davidson and MacKinnon (2004) or Chapter 8.2.2 in Cameron and Trivedi (2005). If R^2 is the uncentered R -squared of the regression in (31), then under the generalized information matrix equality result in (31) the Wald test statistic is asymptotically equivalent to:

$$G_n = nR^2 . \quad (32)$$

The calculation of the score bootstrap simplifies in an analogous fashion. Under a uniform law of large numbers, we obtain that $A_n(\hat{\theta}_M)$ as defined in (30) satisfies,

$$A_n(\hat{\theta}_M) = \left[I : -\frac{1}{n} \sum_{i=1}^n m(Z_i, \hat{\theta}_M) \nabla q(Z_i, \hat{\theta}_M)' \times \left[\frac{1}{n} \sum_{i=1}^n \nabla q(Z_i, \hat{\theta}_M) \nabla q(Z_i, \hat{\theta}_M)' \right]^{-1} \right] + o_p(1) , \quad (33)$$

under the null hypothesis. As a result, the score bootstrap has a simple interpretation in terms of the multivariate regression of the moments $m(Z_i, \hat{\theta}_M)$ on $\nabla q(Z_i, \hat{\theta}_M)$:

$$\begin{aligned} m^{(1)}(Z_i, \hat{\theta}_M) &= \nabla q(Z_i, \hat{\theta}_M)' \beta_1 + \epsilon_{1,i} \\ \vdots &= \vdots \\ m^{(l)}(Z_i, \hat{\theta}_M) &= \nabla q(Z_i, \hat{\theta}_M)' \beta_l + \epsilon_{l,i} \end{aligned} , \quad (34)$$

where $m^{(j)}(Z_i, \hat{\theta}_M)$ is the j^{th} component of $m(Z_i, \hat{\theta}_M)$. Letting $e_{j,i} \equiv m^{(j)}(Z_i, \hat{\theta}_M) - \nabla q(Z_i, \hat{\theta}_M)' \hat{\beta}_j$ be the fitted residual of the j^{th} regression and $e_i = (e_{1,i}, \dots, e_{l,i})'$, we obtain that:

$$S_n^*(\hat{\theta}_M) = \frac{1}{\sqrt{n}} \sum_{i=1}^n e_i W_i . \quad (35)$$

⁹A bootstrap construction for the Information Matrix Equality test was also developed in Horowitz (1994).

Therefore, G_n^* is simply the Wald test for the null hypothesis that the mean of $e_i W_i$ equals zero.

In summary, if the generalized information matrix equality holds, then in testing (27) we may employ the following simple algorithm:

STEP 1: Run the regression in (31) and compute the uncentered R -squared to obtain G_n as in (32).

STEP 2: Regress $\{m(Z_i, \hat{\theta}_M)\}_{i=1}^n$ on $\{\nabla q(Z_i, \hat{\theta}_M)\}_{i=1}^n$ to generate residual vectors $\{e_i\}_{i=1}^n$.

STEP 3: Using random weights $\{W_i\}_{i=1}^n$ independent of $\{Y_i, X_i\}_{i=1}^n$ with $E[W_i] = 0$ and $E[W_i^2] = 1$, perturb the original residual vectors $\{e_i\}_{i=1}^n$ to obtain a new set of residual vectors $\{e_i W_i\}_{i=1}^n$.

STEP 4: Let G_n^* be the Wald test statistic for the null that $E[e_i W_i] = 0$ calculated using $\{e_i W_i\}_{i=1}^n$. To control size at level α , reject if G_n is larger than the $1 - \alpha$ quantile of G_n^* conditional on $\{Z_i\}_{i=1}^n$.

4.4 Clustered Data

Theorem 4.1 and Corollary 4.1 may be applied to clustered data provided clusters are i.i.d. with the same number of observations. Extensions to settings where the clusters are unbalanced or there is heteroskedasticity across them are feasible, essentially requiring an extension of Theorem 4.1 to independent but not identically distributed observations.

Let Z_{ic} denote observation number i in cluster c , J be the total number of observations per cluster, n be the total number of clusters and $Z_c = \{Z_{1c}, \dots, Z_{Jc}\}$. Following (13), we consider test statistics of the general form $\tilde{G}_n \equiv \tilde{T}_n' \tilde{T}_n$, where \tilde{T}_n satisfies:

$$\tilde{T}_n = (A_n(\theta_0) \tilde{\Sigma}_n(\theta_0) A_n(\theta_0)')^{-\frac{1}{2}} \tilde{S}_n(\theta_0) + o_p(1) \quad \tilde{S}_n(\theta) \equiv A_n(\theta) \frac{1}{\sqrt{n}} \sum_{c=1}^n \frac{1}{\sqrt{J}} \sum_{i=1}^J \tilde{s}(Z_{ic}, \theta), \quad (36)$$

$A_n(\theta)$ is again a $r \times m$ matrix, $\tilde{s}(z, \theta)$ maps each (Z_{ic}, θ) into a $m \times 1$ vector and $\tilde{\Sigma}_n(\theta)$ is a robust covariance matrix that allows for arbitrary correlation within cluster. The Wald and LM test statistics, as well as the moment restriction tests previously discussed all extend to this setting when observations are allowed to be dependent within clusters.

The applicability of Theorem 4.1 and Corollary 4.1 to the present context is immediate once we notice that we may define $s(z, \theta)$, mapping each (Z_c, θ) into a $m \times 1$ vector, to be given by:

$$s(Z_c, \theta) = \frac{1}{\sqrt{J}} \sum_{i=1}^J \tilde{s}(Z_{ic}, \theta). \quad (37)$$

The statistics \tilde{T}_n and $\tilde{S}_n(\theta)$ are then special cases of T_n and $S_n(\theta)$ as considered in (13) but with Z_c in place of Z_i . Hence, equations (13) and (37) indicate that the relevant bootstrap statistic should

perturb the data at the level of the cluster rather than the individual observation. We thus define:

$$\tilde{G}_n^* \equiv \tilde{T}_n^{*'} \tilde{T}_n^* \quad \tilde{T}_n^* \equiv (A_n(\hat{\theta}) \tilde{\Sigma}_n^*(\hat{\theta}) A_n(\hat{\theta})')^{-\frac{1}{2}} \tilde{S}_n^*(\hat{\theta}) \quad \tilde{S}_n^*(\theta) \equiv A_n(\theta) \frac{1}{\sqrt{n}} \sum_{c=1}^n \frac{W_c}{\sqrt{J}} \sum_{i=1}^J \tilde{s}(Z_{ic}, \theta)$$

where $\tilde{\Sigma}_n^*(\theta)$ is a robust bootstrap covariance matrix for $s(Z_{ic}, \theta) W_c$.

Given these definitions, it is readily apparent that \tilde{G}_n^* , \tilde{T}_n^* and $\tilde{S}_n^*(\theta)$ are themselves special cases of the bootstrap statistics G_n^* , T_n^* and $S_n^*(\theta)$. The consistency of the proposed score bootstrap then follows immediately provided the clusters are i.i.d., the number of clusters tends to infinity and $s(z, \theta)$ as defined in (37) satisfies Assumption 4.1(ii), 4.2(i) and 4.4(ii).

Corollary 4.2. *Under Assumptions 4.1, 4.2, 4.3 and 4.4, it follows that under the null hypothesis:*

$$\lim_{n \rightarrow \infty} P(\tilde{G}_n \geq \hat{c}_{1-\alpha}) = 1 - \alpha ,$$

for any $0 < \alpha < 1$. Under the same assumptions, if the alternative hypothesis is instead true, then:

$$\lim_{n \rightarrow \infty} P(\tilde{G}_n \geq \hat{c}_{1-\alpha}) = 1 .$$

5 Simulation Evidence

To assess the small sample behavior of the score bootstrap we conduct a series of Monte Carlo experiments examining the performance of bootstrap Wald and LM tests of hypotheses regarding the parameters of a linear model estimated by OLS and a nonlinear probit model estimated by maximum likelihood. We also examine the performance of a test for the presence of intra cluster correlation in the probit model. Because small sample issues often arise in settings with dependent data, we work with hierarchical data generating processes (DGPs) exhibiting dependence of micro-units i within independent clusters c . We consider balanced panels with 20 observations per cluster and sampling designs ranging from 5 to 200 independent clusters.¹⁰

In order to allow a comparison of the wild and score bootstraps with the traditional nonparametric block bootstrap, we consider a setting with continuous regressors so that the block bootstrap distribution may be computed in small samples. It is important to note, however, that in many studies the regressor of interest will have discrete or binary support, in which case the statistic of interest will be undefined in bootstrap samples where only one value of the regressor is sampled. Moreover, even in bootstrap draws where the regressor of interest does exhibit variation, the Hessian may not be full rank. In such settings the traditional resampling based bootstrap will not be viable and the case for consideration of the wild and score bootstraps will be much stronger.

¹⁰In unreported results we found our results to be insensitive to variation in the number of observations per cluster.

5.1 Designs

As pointed out by Chesher (1995), symmetric Monte Carlo designs are likely to yield an overly optimistic assessment of the ability of testing procedures to control size. For this reason we study the performance of our proposed bootstrap procedures under a variety of different designs meant to reflect realistic features of microeconomic datasets. Throughout, the linear model we examine is given by:

$$Y_{ic} = X_{ic} + D_c + \eta_c + \epsilon_{ic} , \quad (38)$$

where the regressors (X_{ic}, D_c) and cluster level error (η_c) are generated according to:

$$X_{ic} = X_c + \xi_{ic} \quad \eta_c = (1 + D_c + X_c)v_c . \quad (39)$$

The regressor of interest is D_c , which varies only at the cluster level. Note that the cluster level random effect η_c exhibits heteroscedasticity with respect to D_c and X_{ic} . The designs are:

Design I: (baseline) We let $(X_c, D_c, \xi_{ic}, \epsilon_{ic})$ be normally distributed with identity covariance matrix, and v_c independent of other variables with a t -distribution with six degrees of freedom.

Design II: (skewed regressor) Design I is modified to generate D_c according to a mixture between a $N(0, 1)$ with probability 0.9 and a $N(2, 9)$ with probability 0.1 as in Horowitz (1997). This yields a regressor with occasional “outliers” and substantial skew and kurtosis in its marginal distribution.

Design III: (misspecification) The model estimated is still (38), but the DGP is modified to:

$$Y_{ic} = X_{ic} + D_c + .1D_c^2 + \eta_c + \epsilon_{ic} , \quad (40)$$

and other features remain as in Design I. Hence, the quadratic term in the regressor of interest is ignored in estimation which yields a skewed reduced form regression error. Note that $E[D_c^3] = E[X_{ic}D_c^2] = 0$ which ensures the population regression coefficient on D_c is still one.

To study the performance of the score bootstrap in a nonlinear model we consider probit estimation of the following data generating process:

$$Y_{ic} = 1\{X_{ic} + D_c + \eta_c + \epsilon_{ic} \geq 0\} \quad X_{ic} = X_c + \xi_{ic} . \quad (41)$$

This is essentially a latent variable representation of the model in (38) without heteroscedasticity in the group error η_c . We consider the following two designs for our probit analysis:

Design IV: (baseline probit) In (41), we let $(X_c, D_c, \xi_{ic}) \sim N(0, I_3/16)$ and $(\eta_c, \epsilon_{ic}) \sim N(0, I_2/2)$.¹¹

¹¹Though the DGP contains a cluster level random effect, the marginal model for the outcome given covariates is a standard probit ensuring that conventional maximum likelihood estimation is consistent.

Design V: (skew probit) We modify Design IV by generating D_c according to a mixture distribution as in Design II, so that the regressor of interest is heavily skewed.

Finally, we illustrate the methods of Section 4.3 by testing for the presence of intra cluster correlation in the probit model. Specifically, we test the hypothesis:

$$H_0 : E\left[\sum_{i=1}^{20} \nu_{ic} \sum_{j \neq i} \nu_{jc}\right] = 0 \quad H_1 : E\left[\sum_{i=1}^{20} \nu_{ic} \sum_{j \neq i} \nu_{jc}\right] \neq 0 , \quad (42)$$

where $\nu_{ic} = [Y_{ic} - p_{ic}] / \sqrt{p_{ic}(1 - p_{ic})}$ is a generalized residual and $p_{ic} = \Phi(X_{ic} + D_c)$ is the conditional probability that Y_{ic} equals one given D_c and X_{ic} .¹² Note that under the probit model $E[\nu_{ic}] = 0$ and $E[\nu_{ic}^2] = 1$. A test of H_0 examines whether within cluster dependence is present in the data, the finding of which might suggest the presence of an unmodeled cluster level random effect. In order to ensure the null hypothesis is true, we employ designs IV and V but set $\eta_c = 0$ almost surely and change the variance of ϵ_{ic} to equal one.

5.2 Results

Table 1 provides empirical false rejection rates from 10,000 Monte Carlo repetitions of Wald and LM tests of the null that the population least squares coefficient on D_c in (38) is one. All tests have a nominal size of 5% and are studentized using a recentered variance estimator.¹³ Bootstrap tests are computed via bootstrap p-values based on 199 repetitions, and reject the null hypothesis whenever the p-value is less than 0.05. Stata code for our Monte Carlo experiments is available online.¹⁴

We consider implementations of the score bootstrap using both Rademacher weights and the skew correcting weights suggested by Mammen (1993). For comparison with the various score bootstraps we also compute the empirical rejection rates of Wald and LM tests based upon analytical clustered standard errors, the original wild bootstrap of Liu (1988), and the pairs-based block bootstrap. Following the results of Davidson and MacKinnon (1999b) on the value of “imposing the null hypothesis” on bootstrap tests, we include in our exercise a variant of the wild bootstrap studied in Cameron, Gelbach, and Miller (2008), which perturbs the restricted score contributions obtained from estimates constraining the coefficient on D_c to equal one, a procedure we term “Wild2”. We also examine the performance of an analogous score bootstrap, which we term “Score2,” that works

¹²See McCall (1994) and Card and Hyslop (2005) for further examples of the use of generalized residual correlations as specification diagnostics.

¹³We also make a finite sample degrees of freedom correction of $n/(n - 1)$ to all variance estimators.

¹⁴URL: <http://www.econ.berkeley.edu/~pkline/papers/score.zip>

Table 1: EMPIRICAL REJECTION RATES - LINEAR MODEL

Wald Tests	Normal Regressor				Mixture Regressor				Misspecification						
	$n = 5$	$n = 10$	$n = 20$	$n = 50$	$n = 200$	$n = 5$	$n = 10$	$n = 20$	$n = 50$	$n = 200$	$n = 5$	$n = 10$	$n = 20$	$n = 50$	$n = 200$
	Analytical	0.381	0.216	0.130	0.085	0.059	0.443	0.322	0.255	0.168	0.087	0.383	0.219	0.132	0.087
Pairs	0.013	0.046	0.041	0.039	0.045	0.018	0.082	0.084	0.068	0.039	0.013	0.046	0.042	0.040	0.046
Wild Rademacher	0.178	0.098	0.066	0.057	0.051	0.229	0.191	0.143	0.086	0.053	0.179	0.100	0.067	0.056	0.051
Wild Mammen	0.175	0.123	0.090	0.069	0.054	0.229	0.185	0.149	0.107	0.063	0.177	0.125	0.093	0.071	0.053
Wild2 Rademacher	0.057	0.059	0.055	0.052	0.051	0.077	0.083	0.070	0.056	0.050	0.062	0.060	0.056	0.052	0.050
Wild2 Mammen	0.051	0.024	0.027	0.036	0.048	0.066	0.035	0.041	0.036	0.032	0.050	0.027	0.026	0.035	0.046
Score Rademacher	0.267	0.184	0.127	0.092	0.063	0.325	0.301	0.268	0.203	0.107	0.269	0.189	0.130	0.093	0.062
Score Mammen	0.374	0.248	0.165	0.104	0.063	0.437	0.363	0.310	0.222	0.114	0.375	0.253	0.165	0.105	0.062
Score2 Rademacher	0.042	0.058	0.053	0.055	0.051	0.047	0.082	0.077	0.063	0.052	0.044	0.059	0.053	0.054	0.051
Score2 Mammen	0.043	0.036	0.059	0.063	0.053	0.046	0.045	0.067	0.068	0.054	0.043	0.037	0.058	0.061	0.052
LM Tests	Normal Regressor				Mixture Regressor				Misspecification						
Analytical	$n = 5$	$n = 10$	$n = 20$	$n = 50$	$n = 200$	$n = 5$	$n = 10$	$n = 20$	$n = 50$	$n = 200$	$n = 5$	$n = 10$	$n = 20$	$n = 50$	$n = 200$
Pairs	0.157	0.090	0.065	0.056	0.052	0.143	0.088	0.054	0.037	0.039	0.157	0.092	0.067	0.055	0.050
Score Rademacher	0.020	0.041	0.039	0.039	0.047	0.021	0.039	0.045	0.044	0.035	0.020	0.041	0.040	0.038	0.048
Score Mammen	0.105	0.101	0.079	0.065	0.053	0.121	0.145	0.110	0.074	0.054	0.098	0.104	0.080	0.064	0.054
	0.080	0.031	0.035	0.044	0.050	0.096	0.053	0.049	0.040	0.034	0.078	0.032	0.035	0.043	0.049

Table 2: EMPIRICAL REJECTION RATES, PROBIT

Wald Tests	Normal Regressor					Mixture Regressor				
	$n = 5$	$n = 10$	$n = 20$	$n = 50$	$n = 200$	$n = 5$	$n = 10$	$n = 20$	$n = 50$	$n = 200$
Analytical	0.326	0.167	0.104	0.070	0.051	0.308	0.170	0.100	0.065	0.058
Pairs	n.a.	0.055	0.060	0.055	0.048	n.a.	0.052	0.060	0.052	0.050
Score Rademacher	0.186	0.128	0.090	0.065	0.049	0.167	0.125	0.085	0.060	0.056
Score Mammen	0.306	0.177	0.108	0.070	0.049	0.279	0.176	0.105	0.065	0.057
Score2 Rademacher	0.140	0.090	0.069	0.059	0.050	0.140	0.113	0.092	0.071	0.058
Score2 Mammen	0.159	0.096	0.070	0.057	0.050	0.160	0.121	0.092	0.070	0.058
LM Tests	Normal Regressor					Mixture Regressor				
	$n = 5$	$n = 10$	$n = 20$	$n = 50$	$n = 200$	$n = 5$	$n = 10$	$n = 20$	$n = 50$	$n = 200$
Analytical	0.171	0.106	0.080	0.062	0.050	0.158	0.110	0.080	0.060	0.055
Pairs ¹⁵	n.a.	0.083	0.082	0.064	0.053	n.a.	0.076	0.081	0.066	0.053
Score Rademacher	0.081	0.079	0.066	0.057	0.049	0.075	0.082	0.069	0.057	0.054
Score Mammen	0.061	0.023	0.038	0.050	0.048	0.058	0.025	0.041	0.053	0.053

with an estimator’s restricted score but employs an unrestricted variance estimate. Details of the various procedures are described in the Implementation Appendix.

The standard clustered Wald test severely over-rejects in samples with few clusters, with performance further degrading when the regressor of interest is generated according to a mixture distribution. Mild misspecification of the sort captured by Design III has little effect on the rejection rates of any of the procedures. A conventional pairs bootstrap of the Wald test yields dramatic improvements in size control though its performance degrades somewhat when the regressor of interest exhibits outliers. Wild bootstrapping the Wald test yields improvements over analytical methods but under performs relative to pairs regardless of whether Mammen or Rademacher weights are used. As suggested by our theoretical results, the score bootstrap yields improvements over analytical methods but is somewhat outperformed by the wild bootstrap particularly in the skewed regressor design. Our variants of the score and wild bootstrap Wald tests that work with restricted residuals perform much better than their unrestricted counterparts. Both Wald2 and Score2 yield performance on par with Pairs even under the relatively difficult skewed regressor design.

In contrast to the Wald tests, the LM tests appear to perform well across a range of sample sizes, regardless of the distribution of the regressors. It is only in samples with very few clusters

¹⁵We were unable to compute the LM statistic in the majority of pairs draws with 5 clusters.

Table 3: EMPIRICAL REJECTION RATES, m-TEST (PROBIT)

Wald Tests	Normal Regressor					Mixture Regressor				
	$n = 5$	$n = 10$	$n = 20$	$n = 50$	$n = 200$	$n = 5$	$n = 10$	$n = 20$	$n = 50$	$n = 200$
Analytical	0.767	0.448	0.257	0.136	0.076	0.772	0.439	0.253	0.142	0.072
Score Rademacher	0.441	0.394	0.233	0.130	0.073	0.437	0.386	0.229	0.134	0.072
Score Mammen	0.370	0.237	0.223	0.128	0.069	0.353	0.234	0.218	0.134	0.068

that the analytical LM test yields significant overrejection. Score bootstrapping the LM statistic with Mammen weights largely removes these distortions as does application of the nonparametric pairs bootstrap.

Table 2 examines the performance of Wald and LM tests in the probit model. Here both Wald and LM tests tend to overreject when asymptotic critical values are used. Use of the pairs bootstrap corrects for this overrejection though in small samples we were sometimes unable to compute the bootstrap distribution.¹⁶ Score bootstrapping the Wald test yields improvements over analytical clustered standard errors but substantial overrejection remains in small samples. Use of the restricted score variant of the test yields smaller improvements than were found with OLS. Score bootstrapping the LM test with Mammen weights, on the other hand, yields size control roughly on par with the pairs bootstrap.

Table 3 examines the performance of tests for intra cluster correlation of the generalized residuals in the probit model, as in (42).¹⁷ Because the information matrix equality holds under both DGPs we use the outer product version of the test described in 4.3.1 generalized to allow for clustering. We see that the analytical m-test procedure overrejects substantially in small samples. Both score bootstraps partially correct this problem, though they significantly overreject as well. With 200 clusters, the analytical and bootstrap approaches appear to work equally well.

Table 4: COMPUTATIONAL TIME (IN SECONDS)

	Score LM	Score2	Wild2	Pairs
OLS	13	15	116	278
Probit	18	21	n.a.	30,718

Finally, to illustrate the dramatic computational advantages of the score bootstrap relative to

¹⁶We discarded simulations for which we were unable to compute an estimate in some bootstrap draws.

¹⁷A description of the implementation of this test can be found in the Implementation Appendix

the wild bootstrap and pairs resampling, Table 4 presents the time elapsed in conducting bootstrap Wald tests using 9,999 bootstrap repetitions of the Score LM, Score2, Wild2, and pairs bootstrap procedures on a simulated dataset with twenty clusters. These computations were performed in Stata/SE 11.1 on a single core of a 2.3 Ghz Quad Core AMD Opteron Processor running Linux.

For OLS, the score bootstrap yields nearly an order of magnitude improvement in computational time over the Wild bootstrap and more than a twenty fold improvement over pairs. For the probit model the results are even more striking. The score bootstrap is more than 1,000 times faster than nonparametric pairs resampling.

6 Conclusion

The score bootstrap provides a substantial computational advantage over the wild and pairs bootstraps and may easily be applied to estimators that lack conventional residuals. Our Monte Carlo experiments suggest these computational advantages come at little cost in terms of performance. Particularly when applied to LM test statistics, the score bootstrap tends to yield substantial improvements over traditional asymptotic testing procedures in small sample environments and exhibits performance comparable to more computationally expensive bootstrap procedures.

IMPLEMENTATION DETAILS

We provide here implementation details for the various bootstrap procedures discussed in the Monte Carlo study of Section 5. We restrict our discussion to the linear model and to the test of hypothesis (42), as the generalization to the probit model is straightforward but notationally intensive. Throughout, the linear model we consider is given by:

$$Y_{ic} = X'_{ic}\beta_0 + \epsilon_{ic} , \quad (43)$$

where $\beta_0 \in \mathbf{R}^m$ and (Y_{ic}, X_{ic}) denotes observation i in cluster c with J the total number of observations per cluster and n the total number of clusters. We examine hypotheses of the form:

$$H_0 : R\beta_0 = r \quad H_1 : R\beta_0 \neq r , \quad (44)$$

where R is a $d \times m$ matrix and r a $d \times 1$ column vector. In our Monte Carlo simulations, R is a vector that selects the coefficient corresponding to D_c (in (38)) and $r = 1$.

Wald Tests

A number of bootstrap procedures we consider provide approximations to the distribution of:

$$T_n \equiv nJ(R\hat{\beta}_u - r)'(RH_n^{-1}\Sigma_n(\hat{\beta}_u)H_n^{-1}R')^{-1}(R\hat{\beta}_u - r) , \quad (45)$$

where $\hat{\beta}_u$ is the OLS estimator, $H_n \equiv \frac{1}{nJ} \sum_{c=1}^n \sum_{i=1}^J X_{ic}X'_{ic}$ and for $e_{ic}(\beta) \equiv (Y_{ic} - X'_{ic}\beta)$ we let:

$$\Sigma_n(\beta) \equiv \frac{1}{(n-1)} \sum_{c=1}^n \left(\frac{1}{J} \sum_{i=1}^J \{X_{ic}e_{ic}(\beta) - \overline{Xe(\beta)}\} \right) \left(\frac{1}{J} \sum_{i=1}^J \{X_{ic}e_{ic}(\beta) - \overline{Xe(\beta)}\} \right)' , \quad (46)$$

and $\overline{Xe(\beta)} \equiv \frac{1}{nJ} \sum_{c=1}^n \sum_{i=1}^J X_{ic}e_{ic}(\beta)$. Tests employing analytical critical values reject at level α if T_n is larger than the $1 - \alpha$ quantile of a χ_d^2 random variable. Note that in Section 5, $d = 1$ and hence the analytical critical value is approximately 3.841.

Wald - Pairs Cluster Bootstrap

Let $X_c = \{X_{ic}\}_{i=1}^J$ and $Y_c = \{Y_{ic}\}_{i=1}^J$. The pairs cluster bootstrap draws with replacement n observations from $\{Y_c, X_c\}_{c=1}^n$ – here an “observation” is an entire cluster. For $\{\tilde{Y}_c, \tilde{X}_c\}_{c=1}^n$ the bootstrap sample, we then compute the OLS estimator $\tilde{\beta}_u$, $\tilde{H}_n \equiv \frac{1}{nJ} \sum_{c=1}^n \sum_{i=1}^J \tilde{X}_{ic}\tilde{X}'_{ic}$ and:

$$\tilde{T}_n \equiv nJ(R\tilde{\beta}_u - R\hat{\beta}_u)'(R\tilde{H}_n^{-1}\tilde{\Sigma}_n(\tilde{\beta}_u)\tilde{H}_n^{-1}R')^{-1}(R\tilde{\beta}_u - R\hat{\beta}_u) , \quad (47)$$

where $\tilde{\Sigma}_n(\beta)$ is computed as $\Sigma_n(\beta)$ with $\{\tilde{Y}_c, \tilde{X}_c\}_{c=1}^n$ in place of $\{Y_c, X_c\}_{c=1}^n$. The pairs cluster bootstrap test then rejects the null hypothesis at level α if:

$$P(\tilde{T}_n \geq T_n | \{Y_c, X_c\}_{c=1}^n) < \alpha . \quad (48)$$

In practice, the probability in (48) can be accurately approximated by simulation of \tilde{T}_n . That is we: (i) Draw B samples $\{\tilde{Y}_c, \tilde{X}_c\}_{c=1}^n$ and obtain for sample b a statistic $\tilde{T}_{n,b}$; (ii) Reject the null hypothesis if the proportion of bootstrap draws $\{\tilde{T}_{n,b}\}_{b=1}^B$ that is larger than T_n is less than α .

Wald - Score Cluster Bootstrap

The score bootstraps we study are constructed employing the score and variance statistics:

$$S_n^*(\beta) \equiv \frac{1}{nJ} \sum_{c=1}^n W_c \sum_{i=1}^J X_{ic} e_{ic}(\beta) \quad (49)$$

$$\Sigma_n^{*s}(\beta) \equiv \frac{1}{(n-1)} \sum_{c=1}^n \left(\frac{W_c}{J} \sum_{i=1}^J \{X_{ic} e_{ic}(\beta) - \overline{X e(\beta)^*}\} \right) \left(\frac{W_c}{J} \sum_{i=1}^J \{X_{ic} e_{ic}(\beta) - \overline{X e(\beta)^*}\} \right)' , \quad (50)$$

where $\overline{X e(\beta)^*} \equiv \frac{1}{nJ} \sum_{c=1}^n W_c (\sum_{i=1}^J X_{ic} e_{ic}(\beta))$, and as in Section 4.4 $\{W_c\}_{c=1}^n$ is a sample of random weights satisfying $E[W_c] = 0$ and $E[W_c^2] = 1$. The score bootstrap statistics we consider are then:

$$T_{n,1}^{*s} \equiv (RH_n^{-1} S_n^*(\hat{\beta}_u))' (RH_n^{-1} \Sigma_n^{*s}(\hat{\beta}_u) H_n^{-1} R')^{-1} (RH_n^{-1} S_n^*(\hat{\beta}_u)) \quad (51)$$

$$T_{n,2}^{*s} \equiv (RH_n^{-1} S_n^*(\hat{\beta}_r))' (RH_n^{-1} \Sigma_n^{*s}(\hat{\beta}_u) H_n^{-1} R')^{-1} (RH_n^{-1} S_n^*(\hat{\beta}_r)) \quad (52)$$

for $\hat{\beta}_r$ the OLS estimate of (43) restricted to satisfy $R\hat{\beta}_r = r$. Here, $T_{n,1}^{*s}$ corresponds to ‘‘Score Wald’’ in Section 5 and $T_{n,2}^{*s}$ to ‘‘Score2 Wald’’. Both procedures reject at level α if:

$$P(T_{n,i}^{*s} \geq T_n \mid \{Y_c, X_c\}_{c=1}^n) < \alpha \quad (53)$$

for $i \in \{1, 2\}$. As is the case for the pairs cluster bootstrap, the probability in (53) need not be computed analytically but may be approximated through simulation. To this end: (i) Draw B samples of $\{W_c\}_{c=1}^n$ and employ each b sample to compute $T_{n,i,b}^{*s}$ for either $i \in \{1, 2\}$; (ii) Reject the null hypothesis if the proportion of draws $\{T_{n,i,b}^{*s}\}_{b=1}^B$ that is larger than T_n is smaller than α .

Wald - Wild Cluster Bootstrap

The wild bootstrap procedures require the generation of new dependent variables:

$$Y_{ic}^*(\beta) \equiv X_{ic}' \beta + e_{ic}(\beta) W_c , \quad (54)$$

where $\{W_c\}_{c=1}^n$ is independent of $\{Y_c, X_c\}_{c=1}^n$ and satisfy $E[W_c] = 0$ and $E[W_c^2] = 1$. For $Y_c^*(\beta) \equiv \{Y_{ic}^*(\beta)\}_{i=1}^J$, let $\hat{\beta}_2^*$ and $\hat{\beta}_1^*$ denote the OLS estimator on the bootstrap samples $\{Y_c^*(\hat{\beta}_r), X_c\}_{c=1}^n$ and $\{Y_c^*(\hat{\beta}_u), X_c\}_{c=1}^n$ respectively – i.e. $\hat{\beta}_2^*$ differs from $\hat{\beta}_1^*$ in that for the former the null hypothesis is imposed in the bootstrap distribution. Similarly, let $\Sigma_{n,1}^{*w}(\beta)$ and $\Sigma_{n,2}^{*w}(\beta)$ denote the analogues of $\Sigma_n(\beta)$ (in (46)) but respectively employing $\{Y_c^*(\hat{\beta}_u), X_c\}_{c=1}^n$ and $\{Y_c^*(\hat{\beta}_r), X_c\}_{c=1}^n$ in place of

$\{Y_c, X_c\}_{c=1}^n$. The two wild bootstrap statistics we consider are then:

$$T_{n,1}^{*w} \equiv (RH_n^{-1}S_n^*(\hat{\beta}_u))'(RH_n^{-1}\Sigma_{n,1}^{*w}(\hat{\beta}_1^*)H_n^{-1}R')^{-1}(RH_n^{-1}S_n^*(\hat{\beta}_u)) \quad (55)$$

$$T_{n,2}^{*w} \equiv (RH_n^{-1}S_n^*(\hat{\beta}_r))'(RH_n^{-1}\Sigma_{n,2}^{*w}(\hat{\beta}_2^*)H_n^{-1}R')^{-1}(RH_n^{-1}S_n^*(\hat{\beta}_r)) , \quad (56)$$

where $T_{n,1}^{*w}$ corresponds to ‘‘Wild’’ in Section 5 and $T_{n,2}^{*w}$ to ‘‘Wild2’’. Both wild bootstrap procedures then reject the null hypothesis at level α whenever:

$$P(T_{n,i}^{*w} \geq T_n \mid \{Y_c, X_c\}_{c=1}^n) < \alpha \quad (57)$$

for $i \in \{1, 2\}$. As for the score bootstrap, we may: (i) Draw B samples of $\{W_c\}_{c=1}^n$ and compute $T_{n,i,b}^{*w}$ for either $i \in \{u, \emptyset\}$; (ii) Reject the null hypothesis if the proportion of computed bootstrap statistics $\{T_{n,i,b}^{*w}\}_{b=1}^B$ that is larger than T_n is smaller than α .

Lagrange Multiplier (LM) Tests

Alternatively, we may also test the hypothesis in (44) employing a Lagrange Multiplier test. Let

$$S_n(\beta) \equiv \frac{1}{nJ} \sum_{c=1}^n \sum_{i=1}^J X_{ic}e_{ic}(\beta) , \quad (58)$$

which constitutes the full sample analogue of $S_n^*(\beta)$ as in (49). The LM test statistic is then:

$$L_n \equiv nJ(RH_n^{-1}S_n(\hat{\beta}_r))'(RH_n^{-1}\Sigma_n(\hat{\beta}_r)H_n^{-1}R')^{-1}(RH_n^{-1}S_n(\hat{\beta}_r)) , \quad (59)$$

which converges in distribution to a χ_d^2 random variable under the null hypothesis. Therefore, the analytical LM test rejects at level α when L_n is larger than the $1 - \alpha$ quantile of χ_d^2 .

LM - Pairs Cluster Bootstrap

As in the Wald pairs cluster bootstrap, let $\{\tilde{Y}_c, \tilde{X}_c\}_{c=1}^n$ denote a bootstrap sample drawn with replacement from $\{Y_c, X_c\}_{c=1}^n$. For $\tilde{\beta}_r$ the OLS estimator on $\{\tilde{Y}_c, \tilde{X}_c\}_{c=1}^n$ restricted to satisfy $R\tilde{\beta}_r = r$, $\tilde{H}_n \equiv \frac{1}{nJ} \sum_{c=1}^n \sum_{i=1}^J \tilde{X}_{ic}\tilde{X}'_{ic}$ and $\tilde{\Sigma}_n(\beta)$ and $\tilde{S}_n(\beta)$ the analogues of $\Sigma_n(\beta)$ (in (46)) and $S_n(\beta)$ (in (58)) but employing $\{\tilde{Y}_c, \tilde{X}_c\}_{c=1}^n$ in place of $\{Y_c, X_c\}_{c=1}^n$, define the bootstrap LM statistic:

$$\tilde{L}_n \equiv nJ(R(\tilde{H}_n^{-1}\tilde{S}_n(\tilde{\beta}_r) - H_n^{-1}S_n(\hat{\beta}_r)))'(R\tilde{H}_n^{-1}\tilde{\Sigma}_n(\tilde{\beta}_r)\tilde{H}_n^{-1}R')^{-1}(R(\tilde{H}_n^{-1}\tilde{S}_n(\tilde{\beta}_r) - H_n^{-1}S_n(\hat{\beta}_r))) . \quad (60)$$

The pairs cluster bootstrap for the Lagrange multiplier test then rejects at level α if:

$$P(\tilde{L}_n \geq L_n \mid \{Y_c, X_c\}_{c=1}^n) < \alpha . \quad (61)$$

In practice, we: (i) Draw B samples with replacement $\{\tilde{Y}_c, \tilde{X}_c\}_{c=1}^n$ and for each sample b compute the bootstrap statistic $\tilde{L}_{n,b}$; (ii) Reject the null hypothesis if the proportion of bootstrap statistics $\{\tilde{L}_{n,b}\}_{b=1}^B$ larger than L_n is smaller than α .

LM - Score Cluster Bootstrap

Given a sample $\{W_c\}_{c=1}^n$ of random weights, independent of $\{Y_c, X_c\}_{c=1}^n$ satisfying $E[W_c] = 0$ and $E[W_c^2] = 1$, and for $S_n^*(\beta)$ as in (49) and $\Sigma_n^{*s}(\beta)$ as in (50), define:

$$L_n^* \equiv nJ(RH_n^{-1}S_n^*(\hat{\beta}_r))'(RH_n^{-1}\Sigma_n^{*s}(\hat{\beta}_r)H_n^{-1}R')^{-1}(RH_n^{-1}S_n^*(\hat{\beta}_r)) . \quad (62)$$

The score bootstrap Lagrange Multiplier procedure then rejects at level α whenever:

$$P(L_n^* \geq L_n \mid \{Y_c, X_c\}_{c=1}^n) < \alpha . \quad (63)$$

We approximate this decision through simulation by: (i) Drawing B samples $\{W_c\}_{c=1}^n$ and employing each sample b to obtain a score bootstrap LM statistics $L_{n,b}^*$; (ii) Rejecting the null hypothesis if the proportion of $\{L_{n,b}^*\}_{b=1}^B$ larger than L_n is smaller than α .

Intra-Cluster Correlation Test (Probit)

The score bootstrap examined in Table 3 follows the discussion of Section 4.3.1. Specifically:

STEP 1 Obtain a probit estimate of model (41) and let $\hat{\theta}_M = (\hat{\theta}_0, \hat{\theta}_X, \hat{\theta}_D)$ where $\hat{\theta}_X$ is the estimated coefficient for X_{ic} and $\hat{\theta}_D$ for D_c and $\hat{\theta}_0$ the intercept. As in Section 4.4 also let $Z_c \equiv \{Y_{ic}, X_{ic}, D_c\}_{i=1}^{20}$.

STEP 2: Employing the probit estimate, construct the fitted prediction $\hat{p}_{ic} \equiv \Phi(\hat{\theta}_0 + X_{ic}\hat{\theta}_X + D_c\hat{\theta}_D)$ and let $\hat{\nu}_{ic} = (Y_{ic} - \hat{p}_{ic})/\sqrt{\hat{p}_{ic}(1 - \hat{p}_{ic})}$ denote the fitted generalized residual.

STEP 3: Following the notation of Section 4.3.1, we may then define the statistics:

$$m(Z_c, \hat{\theta}_M) \equiv \sum_{i=1}^{20} \hat{\nu}_{ic} \sum_{j \neq i} \hat{\nu}_{jc} \quad \nabla q(Z_c, \hat{\theta}_M) \equiv \sum_{i=1}^{20} \frac{(Y_{ic} - \hat{p}_{ic})\phi(\hat{\theta}_0 + X_{ic}\hat{\theta}_X + D_c\hat{\theta}_D)}{\hat{p}_{ic}(1 - \hat{p}_{ic})} X_{ic} , \quad (64)$$

and obtain the cluster level quantities $\{m(Z_c, \hat{\theta}_M)\}_{c=1}^n$ and $\{\nabla q(Z_c, \hat{\theta}_M)\}_{c=1}^n$.

STEP 4: As in Section 4.3.1, (i) Regress $\{m(Z_c, \hat{\theta}_M)\}_{c=1}^n$ on $\{\nabla q(Z_c, \hat{\theta}_M)\}_{c=1}^n$ and obtain residuals $\{e_c\}_{c=1}^n$; (ii) Perturb the cluster level residuals to obtain $\{W_c e_c\}_{c=1}^n$; (iii) For M_n and M_n^* the Wald test statistic for $E[e_c] = 0$ and $E[e_c W_c] = 0$ respectively, the bootstrap test then rejects (42) if:

$$P(M_n \geq M_n^* \mid \{Z_c\}_{c=1}^n) < \alpha . \quad (65)$$

APPENDIX

PROOF OF LEMMA 3.1: First notice that by Markov's inequality, $E[W_i^2] = 1$ and the i.i.d. assumption

$$P^*(\|\frac{1}{\sqrt{n}} \sum_{i=1}^n X_i \epsilon_i^* \| > C) \leq \frac{1}{nC^2} E^*[(\sum_{i=1}^n X_i \epsilon_i^*)'(\sum_{i=1}^n X_i \epsilon_i^*)] = \frac{1}{nC^2} \sum_{i=1}^n X_i' X_i e_i^2. \quad (66)$$

Let $H \equiv E[X_i X_i']$ and $\Sigma \equiv E[X_i X_i' e_i^2]$. Since $n^{-1} \sum_i X_i' X_i e_i^2 \xrightarrow{a.s.} \Sigma$ and $H_n \xrightarrow{a.s.} H$, we obtain from (66),

$$\|\sqrt{n}(\hat{\beta}^* - \hat{\beta})\| \leq \|H_n^{-1}\|_F \times \|\frac{1}{\sqrt{n}} \sum_{i=1}^n X_i \epsilon_i^*\| = O_{p^*}(1) \quad a.s., \quad (67)$$

where $\|\cdot\|_F$ denotes the Frobenius norm. Next observe that for $\|\cdot\|_o$ the operator norm, we have:

$$\begin{aligned} & \| (H_n^{-1} \Sigma_n^*(\hat{\beta}^*) H_n^{-1})^{-1} - (H_n^{-1} \Sigma_n^*(\hat{\beta}) H_n^{-1})^{-1} \|_o \\ & \leq \| (H_n^{-1} \Sigma_n^*(\hat{\beta}) H_n^{-1})^{-1} \|_o \times \| H_n^{-1} (\Sigma_n^*(\hat{\beta}) - \Sigma_n^*(\hat{\beta}^*)) H_n^{-1} \|_o \times \| (H_n^{-1} \Sigma_n^*(\hat{\beta}^*) H_n^{-1})^{-1} \|_o. \end{aligned} \quad (68)$$

Let $X_i^{(k)}$ denote the k^{th} element of the vector X_i . Arguing as in (66), it is straightforward to show that $n^{-\frac{1}{2}} \sum_i X_i^{(k)} X_i^{(l)} X_i^{(s)} \epsilon_i^* = O_{p^*}(1)$ almost surely for any indices k, l, s . Therefore, since $\|\cdot\|_o \leq \|\cdot\|_F$ and $E[\|X_i\|^4] < \infty$, we conclude from (67) and direct calculation that we must have:

$$\begin{aligned} \|\Sigma_n^*(\hat{\beta}) - \Sigma_n^*(\hat{\beta}^*)\|_o & \leq \|\frac{1}{n} \sum_{i=1}^n X_i X_i' \{(Y_i^* - X_i' \hat{\beta})^2 - (Y_i^* - X_i' \hat{\beta}^*)^2\}\|_F \\ & = \|\frac{1}{n} \sum_{i=1}^n X_i X_i' \{2\epsilon_i^* (X_i'(\hat{\beta} - \hat{\beta}^*)) + (X_i'(\hat{\beta} - \hat{\beta}^*))^2\}\|_F = O_{p^*}(n^{-1}) \quad a.s.. \end{aligned} \quad (69)$$

Moreover, since $E[(\epsilon_i^*)^k] = E[W_i^k] e_i^k$, we also obtain from the i.i.d. assumption and $E[\|X_i\|^4 e_i^4] < \infty$ that:

$$\begin{aligned} E^*[\|\frac{1}{n} \sum_{i=1}^n X_i X_i' \{(\epsilon_i^*)^2 - e_i^2\}\|_F^2] & = \sum_{l=1}^m \sum_{s=1}^m E^*[(\frac{1}{n} \sum_{i=1}^n X_i^{(l)} X_i^{(s)} \{(\epsilon_i^*)^2 - e_i^2\})^2] \\ & = \frac{1}{n} \sum_{l=1}^m \sum_{s=1}^m \frac{1}{n} \sum_{i=1}^n (X_i^{(l)} X_i^{(s)})^2 \{(E[W_i^4] - 1) e_i^4\} = o_{a.s.}(1). \end{aligned} \quad (70)$$

Therefore, since $n^{-1} \sum_i X_i X_i' e_i^2 \xrightarrow{a.s.} \Sigma$ and $H_n^{-1} \xrightarrow{a.s.} H^{-1}$, results (69) and (70) establish:

$$\|H_n^{-1} \Sigma_n^*(\hat{\beta}) H_n^{-1} - H^{-1} \Sigma H^{-1}\|_F = o_{p^*}(1) \quad \|H_n^{-1} \Sigma_n^*(\hat{\beta}^*) H_n^{-1} - H^{-1} \Sigma H^{-1}\|_F = o_{p^*}(1) \quad (71)$$

almost surely. Next, for any normal matrix A , let $\xi(A)$ denote its smallest eigenvalue. Since Corollary III.2.6 in Bhatia (1997) implies that $|\xi(A) - \xi(B)| \leq \|A - B\|_F$, it then follows from (71) that:

$$\xi(H_n^{-1} \Sigma_n^*(\hat{\beta}) H_n^{-1}) = \xi(H^{-1} \Sigma H^{-1}) + o_{p^*}(1) \quad \xi(H_n^{-1} \Sigma_n^*(\hat{\beta}^*) H_n^{-1}) = \xi(H^{-1} \Sigma H^{-1}) + o_{p^*}(1) \quad (72)$$

almost surely. However, since for any normal matrix A , $\|A^{-1}\|_o = 1/\xi(A)$, result (72) and Assumption 3.1(i) imply $\|H_n^{-1} \Sigma_n^*(\hat{\beta}^*) H_n^{-1}\|_o = O_{p^*}(1)$ and $\|H_n^{-1} \Sigma_n^*(\hat{\beta}) H_n^{-1}\|_o = O_{p^*}(1)$ almost surely. Hence,

$$\|(H_n^{-1} \hat{\Sigma}_n^*(\hat{\beta}^*) H_n^{-1})^{-1} - (H_n^{-1} \hat{\Sigma}_n^*(\hat{\beta}) H_n^{-1})^{-1}\|_o = O_{p^*}(n^{-1}) \quad a.s. \quad (73)$$

as a result of (68) and (69). In turn, (73) yields that $\|(H_n^{-1} \hat{\Sigma}_n^*(\hat{\beta}^*) H_n^{-1})^{-\frac{1}{2}} - (H_n^{-1} \hat{\Sigma}_n^*(\hat{\beta}) H_n^{-1})^{-\frac{1}{2}}\|_o = O_{p^*}(n^{-1})$ almost surely, and the claim of the Lemma then follows by result (67). ■

Lemma 6.1. *Let $\{W_i\}_{i=1}^n$ be an i.i.d. sample independent of $\{Z_i\}_{i=1}^n$ satisfying $E[W_i^2] = 1$. If Assumptions 4.1, 4.3(i) and 4.4(ii) hold, then the class $\mathcal{F} = \{s(z, \theta)s(z, \theta)'w^2 : \theta \in \Theta\}$ is Glivenko-Cantelli.*

PROOF: By Assumption 4.4(ii), $s(z, \theta)w$ is continuous in $\theta \in \Theta$, and hence so is $s(z, \theta)s(z, \theta)'w^2$. Let $s^{(l)}(z, \theta)$ be the l^{th} component of the vector $s(z, \theta)$. By the mean value theorem and Assumption 4.4(ii):

$$|s^{(l)}(z, \theta)| \leq |s^{(l)}(z, \theta) - s^{(l)}(z, \theta_0)| + |s^{(l)}(z, \theta_0)| \leq F(z)\|\theta - \theta_0\| + |s^{(l)}(z, \theta_0)|. \quad (74)$$

Hence, for $D = \text{diam}(\Theta)$ it then follows that $|s^{(l)}(z, \theta)s^{(j)}(z, \theta)w^2| \leq w^2(F(z)D + |s^{(l)}(z, \theta_0)|)(F(z)D + |s^{(j)}(z, \theta_0)|)$, which is integrable for all $1 \leq l \leq j \leq k$ by Assumption 4.1(i)-(ii) and 4.4(ii). We conclude that \mathcal{F} has an integrable envelope, and the Lemma follows by Example 19.8 in van der Vaart (1999). ■

Lemma 6.2. *Under Assumptions 4.1(i), 4.3(i) and 4.4(i)-(ii), $\mathcal{F} \equiv \{ws(z, \theta) : \theta \in \Theta\}$ is Donsker.*

PROOF: Let $\|\cdot\|_o$ and $\|\cdot\|_F$ denote the operator and Frobenius norms. Using $\|\cdot\|_o \leq \|\cdot\|_F$, Assumption 4.4(ii) and the mean value theorem, we obtain that for some $\bar{\theta}$ a convex combination of θ_1 and θ_2 :

$$\begin{aligned} & \|ws(z, \theta_1) - ws(z, \theta_2)\| \\ &= |w| \times \|\nabla s(z, \bar{\theta})(\theta_1 - \theta_2)\| \leq |w| \times \|\nabla s(z, \bar{\theta})\|_o \times \|\theta_1 - \theta_2\| \leq |w| \times F(z) \times \|\theta_1 - \theta_2\|. \end{aligned} \quad (75)$$

Hence, the class \mathcal{F} is Lipschitz in $\theta \in \Theta$. For $s^{(l)}(z, \theta)$ the l^{th} component of the vector $s(z, \theta)$, let $\mathcal{F}^l \equiv \{ws^{(l)}(z, \theta) : \theta \in \Theta\}$ and note that Theorem 2.7.11 in van der Vaart and Wellner (1996) implies:

$$N_{[\cdot]}(2\epsilon\|\tilde{F}\|_{L^2}, \mathcal{F}^l, \|\cdot\|_{L^2}) \leq N(\epsilon, \Theta, \|\cdot\|), \quad (76)$$

where $\tilde{F}(w, z) = |w|F(z)$. Let $D = \text{diam}(\Theta)$ and $M^2 = E[\tilde{F}^2(W_i, Z_i)]$ and notice that Assumptions 4.4(i)-(ii) imply $M < \infty$. Since by (75), the diameter of \mathcal{F}^l under $\|\cdot\|_{L^2}$ is less than or equal to MD :

$$\begin{aligned} \int_0^\infty \sqrt{\log N_{[\cdot]}(\epsilon, \mathcal{F}^l, \|\cdot\|_{L^2})} d\epsilon &\leq \int_0^{MD} \sqrt{\log N_{[\cdot]}(\epsilon, \mathcal{F}^l, \|\cdot\|_{L^2})} d\epsilon = 2M \int_0^{\frac{D}{2}} \sqrt{\log N_{[\cdot]}(2Mu, \mathcal{F}^l, \|\cdot\|_{L^2})} du \\ &\leq 2M \int_0^{\frac{D}{2}} \sqrt{\log N(u, \Theta, \|\cdot\|)} du \leq 2M \int_0^{\frac{D}{2}} \sqrt{p \log(D/u)} du < \infty \end{aligned} \quad (77)$$

where the first equality follows by the change of variables $u = \epsilon/2M$, the second inequality from (76) and the third by $N(u, \Theta, \|\cdot\|) \leq (\text{diam}(\Theta)/u)^p$. Since $E[\tilde{F}^2(Z_i, W_i)] < \infty$, (77) and Theorem 2.5.6 in van der Vaart and Wellner (1996) implies \mathcal{F}^l is Donsker for $1 \leq l \leq k$, and the claim of the Lemma follows. ■

Lemma 6.3. *Suppose Assumptions 4.1, 4.2, 4.3 and 4.4(ii) hold. If the null hypothesis is true, it then follows that $G_n \xrightarrow{L} \mathcal{X}_r^2$. On the other hand, if the alternative hypothesis is true, then $G_n \xrightarrow{P} \infty$.*

PROOF: We first study the limiting behavior of G_n under the null hypothesis. For this purpose, notice that Assumption 4.3(ii) implies that $A_n(\theta_0) = A(\theta_0) + o_p(1)$, while Lemma 6.1 applied to $W_i = 1$ with probability one yields $\Sigma_n(\theta_0) = \Sigma(\theta_0) + o_p(1)$ for $\Sigma(\theta_0) = E[s(Z_i, \theta_0)s(Z_i, \theta_0)']$. Therefore, we conclude:

$$A_n(\theta_0)\Sigma_n(\theta_0)A_n(\theta_0)' = A(\theta_0)\Sigma(\theta_0)A(\theta_0)' + o_p(1). \quad (78)$$

It follows that $A_n(\theta_0)\Sigma_n(\theta_0)A_n(\theta_0)'$ is then invertible with probability tending to one by Assumption 4.1(ii). Therefore, by Assumptions 4.2(i), 4.1(ii), 4.3(i) and the central limit theorem we conclude that:

$$T_n = (A_n(\theta_0)\Sigma_n(\theta_0)A_n(\theta_0)')^{-1}A_n(\theta_0)\frac{1}{\sqrt{n}}\sum_{i=1}^n s(Z_i, \theta_0) + o_p(1) \xrightarrow{L} N(0, I). \quad (79)$$

Hence, by the continuous mapping theorem and (79), $G_n \xrightarrow{L} \mathcal{X}_r$, which establishes the first claim of the Lemma. The second claim of the Lemma was assumed in Assumption 4.2(ii). ■

PROOF OF THEOREM 4.1: Let $\Sigma(\theta) = E[s(Z_i, \theta)s(Z_i, \theta)']$. As argued in (74), the matrix $s(z, \theta)s(z, \theta)'$ has an integrable envelope. Hence, since $s(z, \theta)s(z, \theta)'$ is continuous in θ for all z by Assumption 4.4(ii), the dominated convergence theorem implies $\Sigma(\theta)$ is continuous in θ . Therefore, by Lemma 6.1 and Assumption 4.1(i), we obtain $\Sigma_n^*(\hat{\theta}) = \Sigma(\theta_0) + o_p(1)$. In addition, $A_n(\hat{\theta}) = A(\theta_0) + o_p(1)$ by Assumption 4.3(ii) and hence Assumption 4.1(ii) and $\sup_{\theta \in \Theta} \|n^{-\frac{1}{2}}\sum_i s(Z_i, \theta)W_i\| = O_p(1)$ by Lemma 6.2 imply:

$$\begin{aligned} & (A_n(\hat{\theta})\Sigma_n^*(\hat{\theta})A_n(\hat{\theta})')^{-\frac{1}{2}}A_n(\hat{\theta})\frac{1}{\sqrt{n}}\sum_{i=1}^n s(Z_i, \hat{\theta})W_i \\ &= (A(\theta_0)\Sigma(\theta_0)A(\theta_0)')^{-\frac{1}{2}}A(\theta_0)\frac{1}{\sqrt{n}}\sum_{i=1}^n s(Z_i, \hat{\theta})W_i + o_p(1) \\ &= (A(\theta_0)\Sigma(\theta_0)A(\theta_0)')^{-\frac{1}{2}}A(\theta_0)\frac{1}{\sqrt{n}}\sum_{i=1}^n s(Z_i, \theta_0)W_i + o_p(1), \end{aligned} \quad (80)$$

where the second equality follows by Assumption 4.1(i) and Lemma 6.2. Let BL_c be the set of Lipschitz real valued functions whose Lipschitz constant and level are less than c . For two random variables Y, V :

$$\|Y - V\|_{BL_1} \equiv \sup_{f \in BL_1} |E[f(Y)] - E[f(V)]|, \quad (81)$$

metrizes weak convergence, see for example Theorem 1.12.4 in van der Vaart and Wellner (1996). Define:

$$\bar{T}_n^* \equiv (A(\theta_0)\Sigma(\theta_0)A(\theta_0)')^{-\frac{1}{2}}A(\theta_0)\frac{1}{\sqrt{n}}\sum_{i=1}^n s(Z_i, \theta_0)W_i. \quad (82)$$

Using that all $f \in BL_1$ are bounded in level and Lipschitz constant by one, we obtain for any $\eta > 0$:

$$\sup_{f \in BL_1} |E[f(\bar{T}_n^*) - f(T_n^*)|\{Z_i\}_{i=1}^n]| \leq \eta P(|\bar{T}_n^* - T_n^*| \leq \eta|\{Z_i\}_{i=1}^n) + 2P(|\bar{T}_n^* - T_n^*| > \eta|\{Z_i\}_{i=1}^n). \quad (83)$$

However, by the law of iterated expectations and (80), we have that $P(|\bar{T}_n^* - T_n^*| > \eta|\{Z_i\}_{i=1}^n)$ converges to zero in mean, and hence in probability. As a result, since η is arbitrary, result (83) in fact implies:

$$\sup_{f \in BL_1} |E[f(\bar{T}_n^*)|\{Z_i\}_{i=1}^n] - E[f(T_n^*)|\{Z_i\}_{i=1}^n]| = o_p(1). \quad (84)$$

Let $T_\infty^* \sim N(0, I)$. Since $\|\cdot\|_{BL_1}$ metrizes weak convergence, Assumptions 4.3(i) and 4.4(i) together with Lemma 2.9.5 in van der Vaart and Wellner (1996) in turn let us conclude that:

$$\sup_{f \in BL_1} |E[f(\bar{T}_n^*)|\{Z_i\}_{i=1}^n] - E[f(T_\infty^*)]| = o_p(1). \quad (85)$$

For any $M > 0$, define the map $g_M : \mathbf{R}^r \rightarrow \mathbf{R}$ to be given by $g_M(a) = \min\{a'a, M\}$ and notice that for any $a, b \in \mathbf{R}^r$ we have $|g_M(a) - g_M(b)| \leq 2\sqrt{M}\|a - b\|$ and $g_M(a) \leq M$ so that for $M \geq 4$ we have $g_M \in BL_M$. As a result, for any $f \in BL_1$, $f \circ g_M \in BL_M$ and $M^{-1}f \circ g_M \in BL_1$, which implies:

$$\sup_{f \in BL_1} |E[f(g_M(T_n^*))|\{Z_i\}_{i=1}^n] - E[f(g_M(T_\infty^*))]| \leq M \sup_{f \in BL_1} |E[f(T_n^*)|\{Z_i\}_{i=1}^n] - E[f(T_\infty^*)]| = o_p(1), \quad (86)$$

where the final result follows by (84) and (85). Since $G_n^* = T_n^{*'}T_n^*$ and every $f \in BL_1$ is bounded by one,

$$\sup_{f \in BL_1} |E[f(G_n^*) - f(g_M(T_n^*))|\{Z_i\}_{i=1}^n]| \leq 2P(T_n^{*'}T_n^* > M|\{Z_i\}_{i=1}^n). \quad (87)$$

By (80) and the continuous mapping theorem, $T_n^{*'}T_n^* \xrightarrow{L} \mathcal{X}_r^2$ unconditionally and hence is asymptotically tight. For an arbitrary $\eta > 0$ it then follows by Markov's inequality that for M sufficiently large:

$$\limsup_{n \rightarrow \infty} P(2P(T_n^{*'}T_n^* > M|\{Z_i\}_{i=1}^n) > \eta) \leq \limsup_{n \rightarrow \infty} \frac{2}{\eta} P(T_n^{*'}T_n^* > M) < \eta. \quad (88)$$

Similarly, let $G_\infty^* \sim \mathcal{X}_r^2$ and notice that by selecting M appropriately large we may also obtain:

$$\sup_{f \in BL_1} |E[f(G_\infty^*) - f(g_M(T_\infty^*))]| \leq 2P(T_\infty^{*'}T_\infty^* > M) < \eta. \quad (89)$$

Since η is arbitrary, results (86), (87), (88) and (89) in turn allow us to conclude that:

$$\sup_{f \in BL_1} |E[f(G_n^*)|\{Z_i\}_{i=1}^n] - E[f(G_\infty^*)]| = o_p(1), \quad (90)$$

which establishes the weak convergence of the distribution of G_n^* conditional on $\{Z_i\}_{i=1}^n$ to that of G_∞^* in probability. Letting F be the cdf of G_∞^* , we obtain by the Portmanteau theorem, G_∞^* having a continuous distribution, result (90) and Lemma 6.3 that for any $c \in \mathbf{R}$, $F_n^*(c) = F(c) + o_p(1)$ and $F_n(c) = F(c) + o(1)$. The Theorem follows since convergence is uniform in $c \in \mathbf{R}$ by Lemma 2.11 in van der Vaart (1999). ■

Lemma 6.4. *Let $F_n : \mathbf{R} \rightarrow [0, 1]$, $F : \mathbf{R} \rightarrow [0, 1]$ be monotonic, $\sup_{c \in \mathbf{R}} |F_n(c) - F(c)| = o_p(1)$ and define:*

$$c_\alpha \equiv \inf\{c : F(c) \geq \alpha\} \quad c_{n,\alpha} \equiv \inf\{c : F_n(c) \geq \alpha\}.$$

If F is strictly increasing at c_α , it then follows that $c_{n,\alpha} = c_\alpha + o_p(1)$.

PROOF: Fix $\epsilon > 0$. Since by hypothesis F is strictly increasing at c_α it follows by definition of c_α :

$$F(c_\alpha - \epsilon) < \alpha < F(c_\alpha + \epsilon). \quad (91)$$

Moreover, since $F_n(c_\alpha + \epsilon) > \alpha$ implies that $c_{n,\alpha} \leq c_\alpha + \epsilon$ and $F_n(c_\alpha - \epsilon) < \alpha$ implies that $c_{n,\alpha} > c_\alpha - \epsilon$,

$$\lim_{n \rightarrow \infty} P(|c_\alpha - c_{n,\alpha}| \leq \epsilon) \geq \lim_{n \rightarrow \infty} P(F_n(c_\alpha - \epsilon) < \alpha < F_n(c_\alpha + \epsilon)) = 1, \quad (92)$$

where the final equality follows from (91) and $\sup_c |F_n(c) - F(c)| = o_p(1)$ by hypothesis. ■

PROOF OF COROLLARY 4.1: Let F denote the cdf of a \mathcal{X}_r^2 random variable and $c_{1-\alpha}$ be its $1 - \alpha$ quantile. As argued following (90), $\sup_c |F_n^*(c) - F(c)| = o_p(1)$, and hence by Lemma 6.4 it follows that

$\hat{c}_{1-\alpha} = c_{1-\alpha} + o_p(1)$ provided $0 < \alpha < 1$. The first claim of the Corollary then follows by Lemma 6.3 and the continuous mapping theorem.

For the second claim of the Corollary, observe that the bootstrap statistic $S_n^*(\hat{\theta})$ remains properly centered. In fact, (90) was established without appealing to Assumption 4.2(i). Therefore, $\hat{c}_{1-\alpha} = c_{1-\alpha} + o_p(1)$ under the alternative hypothesis as well. However, under the alternative hypothesis $G_n \xrightarrow{P} \infty$ by Lemma 6.3 and therefore the second claim of the Corollary follows. ■

PROOF OF COROLLARY 4.2: Given the definitions, this is a special case of Corollary 4.1. ■

References

- ANDREWS, D. W. K. (2002): “Higher-Order Improvements of a Computationally Attractive k-Step Bootstrap for Extremum Estimators,” *Econometrica*, 70(1), 119–162.
- BARBE, P., AND P. BERTAIL (1995): *The Weighted Bootstrap*. Springer-Verlag, New York.
- BERTRAND, M., E. DUFLO, AND S. MULLAINATHAN (2004): “How Much Should We Trust Differences in Differences Estimates?,” *Quarterly Journal of Economics*, 119(1), 249–275.
- BHATIA, R. (1997): *Matrix Analysis*. Springer, New York.
- CAMERON, A. C., J. B. GELBACH, AND D. L. MILLER (2008): “Bootstrap-Based Improvements for Inference with Clustered Errors,” *Review of Economics and Statistics*, 90, 414–427.
- CAMERON, A. C., AND P. K. TRIVEDI (2005): *Microeconometrics - Methods and Applications*. Cambridge University Press, New York.
- CARD, D., AND D. R. HYSLOP (2005): “Estimating the Effects of a Time-Limited Earnings Subsidy for Welfare-Leavers,” *Econometrica*, 73(6), 1723–1770.
- CAVALIERE, G., AND A. M. R. TAYLOR (2008): “Bootstrap Unit Root Tests for Time Series with Nonstationary Volatility,” *Econometric Theory*, 24, 43–71.
- CHATTERJEE, S., AND A. BOSE (2005): “Generalized Bootstrap for Estimating Equations,” *Annals of Statistics*, 33, 414–436.
- CHESHER, A. (1984): “Testing for Neglected Heterogeneity,” *Econometrica*, 52(4), 865–872.
- (1995): “A Mirror Image Invariance for M-estimators,” *Econometrica*, 63(1), 207–211.
- DAVIDSON, R., AND E. FLACHAIRE (2008): “The Wild Bootstrap, Tamed at Last,” *Journal of Econometrics*, 146, 162–169.

- DAVIDSON, R., AND J. G. MACKINNON (1999a): “Bootstrap Testing in Nonlinear Models,” *International Economic Review*, 40, 487–508.
- (1999b): “The Size Distortion of Bootstrap Tests,” *Econometric Theory*, 15, 361–376.
- DAVIDSON, R., AND J. G. MACKINNON (2004): *Econometric Theory and Methods*. Oxford University Press, New York, Oxford.
- DAVIDSON, R., AND J. G. MACKINNON (2010): “Wild Bootstrap Tests for IV Regression,” *Journal of Business and Economic Statistics*, 28, 128–144.
- DONALD, S. G., AND K. LANG (2007): “Inference with Differences-in-Differences and Other Panel Data,” *Review of Economics and Statistics*, 89, 221233.
- EFRON, B. (1979): “Bootstrap Methods: Another Look at the Jackknife,” *The Annals of Statistics*, 7(1), 1–26.
- FREEDMAN, D. A. (1981): “Bootstrapping Regression Models,” *The Annals of Statistics*, 9(6), 1218–1228.
- HALL, P. (1992): *The Bootstrap and Edgeworth Expansion*. Springer-Verlag, New York.
- HARDLE, W., AND E. MAMMEN (1993): “Comparing Nonparametric Versus Parametric Regression Fits,” *The Annals of Statistics*, 21(4), 1926–1947.
- HOROWITZ, J. L. (1994): “Bootstrap-based Critical Values for the Information Matrix Test,” *Journal of Econometrics*, 61, 395–411.
- (1997): “Bootstrap Methods in Econometrics: Theory and Numerical Performance,” in *Advances in Economics and Econometrics: Theory and Applications, Seventh World Congress*, ed. by D. M. Kreps, and K. F. Wallis, vol. 3. Cambridge University Press.
- (2001): “The Bootstrap,” in *Handbook of Econometrics*, ed. by J. J. Heckman, and E. Leamer, vol. 5, chap. 52. Elsevier.
- HU, F., AND J. D. KALBFLEISCH (2000): “The Estimating Function Bootstrap,” *The Canadian Journal of Statistics*, 28(3), 449–481.
- HU, F., AND J. V. ZIDEK (1995): “A Bootstrap Based on the Estimating Equations of the Linear Model,” *Biometrika*, 82(2), 263–275.
- KAIDO, H., AND A. SANTOS (2011): “Asymptotically Efficient Estimation of Models Defined by Convex Moment Inequalities,” Working paper, University of California - San Diego.

- KLINE, P., AND A. SANTOS (2011): “Higher Order Properties of the Wild Bootstrap Under Misspecification,” Working paper, University of California - Berkeley.
- LIU, R. Y. (1988): “Bootstrap Procedures under some Non-I.I.D. Models,” *The Annals of Statistics*, 16(4), 1696–1708.
- MA, S., AND M. R. KOSOROK (2005): “Robust Semiparametric M-estimation and the Weighted Bootstrap,” *Journal of Multivariate Analysis*, 96, 190–217.
- MAMMEN, E. (1993): “Bootstrap and Wild Bootstrap for High Dimensional Linear Models,” *The Annals of Statistics*, 21(1), 255–285.
- MCCALL, B. P. (1994): “Specification Diagnostics for Duration Models: A Martingale Approach,” *Journal of Econometrics*, 60, 293–312.
- NEWKEY, W. K. (1985a): “Generalized Method of Moments Specification Testing,” *Journal of Econometrics*, 29, 229–256.
- (1985b): “Maximum Likelihood Specification Testing and Conditional Moment Tests,” *Econometrica*, 53(5), 1047–1070.
- NEWKEY, W. K., AND D. L. MCFADDEN (1994): “Large Sample Estimation and Hypothesis Testing,” in *Handbook of Econometrics*, ed. by R. F. Engle, and D. L. McFadden, vol. IV, pp. 2113–2245. Elsevier Science B.V.
- TAUCHEN, G. (1985): “Diagnostic Testing and Evaluation of Maximum Likelihood Models,” *Journal of Econometrics*, 30, 415–443.
- VAN DER VAART, A. (1999): *Asymptotic Statistics*. Cambridge University Press, New York.
- VAN DER VAART, A. W., AND J. A. WELLNER (1996): *Weak Convergence and Empirical Processes: with Applications to Statistics*. Springer, New York.
- WHITE, H. (1982): “Maximum Likelihood Estimation of Misspecified Models,” *Econometrica*, 50, 1–25.
- (1994): *Estimation, Inference, and Specification Analysis*. Cambridge University Press, New York.
- WOOLDRIDGE, J. (2003): “Jackknife, Bootstrap, and other Resampling Methods in Regression Analysis,” *American Economic Review*, 93(2), 133–138.
- WOOLDRIDGE, J. M. (2002): *Econometric Analysis of Cross Section and Panel Data*. The MIT Press, Cambridge.

WU, C. F. J. (1986): “Jackknife, Bootstrap, and other Resampling Methods in Regression Analysis,” *Annals of Statistics*, 14(4), 1261–1295.

YOU, J., AND G. CHEN (2006): “Wild Bootstrap Estimation in Partially Linear Models with Heteroscedasticity,” *Statistics and Probability Letters*, 76, 340–348.