

Article

Bootstrap Tests for Overidentification in Linear Regression Models

Russell Davidson ^{1,2,†} and James G. MacKinnon ^{3,†,*}

¹ Department of Economics and CIREQ, McGill University, Montréal, Québec H3A 2T7, Canada;
E-Mail: russell.davidson@mcgill.ca

² AMSE-GREQAM, Centre de la Vieille Charité, 13236 Marseille cedex 02, France

³ Department of Economics, Queen's University, Kingston, Ontario K7L 3N6, Canada

† These authors contributed equally to this work.

* Author to whom correspondence should be addressed; E-Mail: jgm@econ.queensu.ca;
Tel.: +1-613-533-2293; Fax: +1-613-533-6668.

Academic Editor: Kerry Patterson

Received: 13 July 2015 / Accepted: 24 November 2015 / Published: 9 December 2015

Abstract: We study the finite-sample properties of tests for overidentifying restrictions in linear regression models with a single endogenous regressor and weak instruments. Under the assumption of Gaussian disturbances, we derive expressions for a variety of test statistics as functions of eight mutually independent random variables and two nuisance parameters. The distributions of the statistics are shown to have an ill-defined limit as the parameter that determines the strength of the instruments tends to zero and as the correlation between the disturbances of the structural and reduced-form equations tends to plus or minus one. This makes it impossible to perform reliable inference near the point at which the limit is ill-defined. Several bootstrap procedures are proposed. They alleviate the problem and allow reliable inference when the instruments are not too weak. We also study their power properties.

Keywords: Sargan test; Basman test; Anderson-Rubin test; weak instruments

JEL classifications: C10; C12; C15; C30

1. Introduction

In recent years, there has been a great deal of work on the finite-sample properties of estimators and tests for linear regression models with endogenous regressors when the instruments are weak. Much of this work has focused on the case in which there is just one endogenous variable on the right-hand side, and numerous procedures for testing hypotheses about the coefficient of this variable have been studied. See, among many others, Staiger and Stock [1], Stock, Wright, and Yogo [2], Kleibergen [3], Moreira [4,5], Andrews, Moreira, and Stock [6], and Davidson and MacKinnon [7,8]. However, the closely related problem of testing overidentifying restrictions when the instruments are weak has been less studied. Staiger and Stock [1] provides some asymptotic results, and Guggenberger, Kleibergen, Mavroeidis, and Chen [9] derives certain asymptotic properties of tests for overidentifying restrictions in the context of hypotheses about the coefficients of a subset of right-hand side variables.

In the next section, we discuss the famous test of Sargan [10] and other asymptotic tests for overidentification in linear regression models estimated by instrumental variables (IV) or limited information maximum likelihood (LIML). We show that the test statistics are all functions of six quadratic forms defined in terms of the two endogenous variables of the model, the linear span of the instruments, and its orthogonal complement. In fact, they can be expressed as functions of a certain ratio of sums of squared residuals and are closely related to the test proposed by Anderson and Rubin [11]. In Section 3, we analyze the properties of these overidentification test statistics. We use a simplified model with only three parameters, which is nonetheless capable of generating statistics with exactly the same distributions as those generated by a more general model, and characterize the finite-sample distributions of the test statistics, not by providing closed-form solutions, but by providing unique recipes for simulating them in terms of just eight mutually independent random variables with standard distributions as functions of the three parameters.

In Section 4, we make use of these results in order to derive the limiting behavior of the statistics as the instrument strength tends to zero, as the correlation between the disturbances in the structural and reduced-form equations tends to unity, and as the sample size tends to infinity under the assumptions of weak-instrument asymptotics. Since the results of Section 4 imply that none of the tests we study is robust with weak instruments, Section 5 discusses a number of bootstrap procedures that can be used in conjunction with any of the overidentification tests. Some of these procedures are purely parametric, while others make use of resampling.

Then, in Section 6, we investigate by simulation the finite-sample behavior of the statistics we consider. Simulation evidence and theoretical analysis concur in strongly preferring a variant of a likelihood-ratio test to the more conventional forms of Sargan test. In Section 7, we look at the performance of bootstrap tests, finding that the best of them behave very well if the instruments are not too weak. However, as our theory suggests, they improve very little over tests based on asymptotic critical values in the neighborhood of the singularity that occurs where the instrument strength tends to zero and the correlation of the disturbances tends to one.

In Section 8, we analyze the power properties of the two main variants of bootstrap test. We obtain analytical results that generalize those of Section 3. Using those analytical results, we conduct extensive simulation experiments, mostly for cases that allow the bootstrap to yield reliable inference. We find

that bootstrap tests based on IV estimation seem to have a slight power advantage over those based on LIML, at the cost of slightly greater size distortion under the null when the instruments are not too weak. Section 9 presents a brief discussion of how both the test statistics and the bootstrap procedures can be modified to take account of heteroskedasticity and clustered data. Finally, some concluding remarks are made in Section 10.

2. Tests for Overidentification

Although the tests for overidentification that we deal with are applicable to linear regression models with any number of endogenous right-hand side variables, we restrict attention in this paper to a model with just one such variable. We do so partly for expositional convenience and partly because this special case is of particular interest and has been the subject of much research in recent years. The model consists of just two equations,

$$\mathbf{y}_1 = \beta \mathbf{y}_2 + \mathbf{Z}\boldsymbol{\gamma} + \mathbf{u}_1, \text{ and} \quad (1)$$

$$\mathbf{y}_2 = \mathbf{W}\boldsymbol{\pi} + \mathbf{u}_2. \quad (2)$$

Here \mathbf{y}_1 and \mathbf{y}_2 are n -vectors of observations on endogenous variables, \mathbf{Z} is an $n \times k$ matrix of observations on exogenous variables, and \mathbf{W} is an $n \times l$ matrix of instruments such that $\mathcal{S}(\mathbf{Z}) \subset \mathcal{S}(\mathbf{W})$, where the notation $\mathcal{S}(\mathbf{A})$ means the linear span of the columns of the matrix \mathbf{A} . The disturbances are assumed to be homoskedastic and serially uncorrelated. We assume that $l > k + 1$, so that the model is overidentified.

The parameters of this model are the scalar β , the k -vector $\boldsymbol{\gamma}$, the l -vector $\boldsymbol{\pi}$, and the 2×2 contemporaneous covariance matrix of the disturbances u_{1i} and u_{2i} :

$$\boldsymbol{\Sigma} \equiv \begin{bmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{bmatrix}. \quad (3)$$

Equation (1) is the structural equation we are interested in, and Equation (2) is a reduced-form equation for the second endogenous variable \mathbf{y}_2 .

The model (1) and (2) implicitly involves one identifying restriction, which cannot be tested, and $q \equiv l - k - 1$ overidentifying restrictions. These restrictions say, in effect, that if we append q regressors all belonging to $\mathcal{S}(\mathbf{W})$ to Equation (1) in such a way that the equation becomes just identified, then the coefficients of these q additional regressors are zero.

The most common way to test the overidentifying restrictions is to use a Sargan test (Sargan, [10]), which can be computed in various ways. The easiest is probably to estimate Equation (1) by instrumental variables (IV), using the l columns of \mathbf{W} as instruments, so as to obtain the IV residuals

$$\hat{\mathbf{u}}_1 = \mathbf{y}_1 - \hat{\beta}_{IV}\mathbf{y}_2 - \mathbf{Z}\hat{\boldsymbol{\gamma}}_{IV}, \quad (4)$$

where $\hat{\beta}_{IV}$ and $\hat{\boldsymbol{\gamma}}_{IV}$ denote the IV estimates of β and $\boldsymbol{\gamma}$, respectively. The IV residuals $\hat{\mathbf{u}}_1$ are then regressed on \mathbf{W} . The explained sum of squares from this regression divided by the IV estimate of σ_1^2 is the test statistic, and it is asymptotically distributed as $\chi^2(q)$.

The numerator of the Sargan statistic can be written as

$$(\mathbf{y}_1 - \mathbf{Z}\hat{\gamma}_{IV} - \hat{\beta}_{IV}\mathbf{y}_2)^\top \mathbf{P}_W (\mathbf{y}_1 - \mathbf{Z}\hat{\gamma}_{IV} - \hat{\beta}_{IV}\mathbf{y}_2), \tag{5}$$

where $\hat{\beta}_{IV}$ and $\hat{\gamma}_{IV}$ denote the IV estimates of β and γ , respectively, and $\mathbf{P}_W \equiv \mathbf{W}(\mathbf{W}^\top \mathbf{W})^{-1} \mathbf{W}^\top$ projects orthogonally onto $S(\mathbf{W})$. We define \mathbf{P}_Z similarly, and let $\mathbf{M}_W \equiv \mathbf{I} - \mathbf{P}_W$ and $\mathbf{M}_Z \equiv \mathbf{I} - \mathbf{P}_Z$. Since \mathbf{Z} is orthogonal to the IV residuals,

$$\mathbf{y}_1 - \mathbf{Z}\hat{\gamma}_{IV} - \hat{\beta}_{IV}\mathbf{y}_2 = \mathbf{M}_Z (\mathbf{y}_1 - \mathbf{Z}\hat{\gamma}_{IV} - \hat{\beta}_{IV}\mathbf{y}_2) = \mathbf{M}_Z (\mathbf{y}_1 - \hat{\beta}_{IV}\mathbf{y}_2).$$

Then, since $\mathbf{P}_W \mathbf{M}_Z = \mathbf{M}_Z \mathbf{P}_W = \mathbf{P}_W - \mathbf{P}_Z = \mathbf{M}_Z - \mathbf{M}_W$, the numerator of the Sargan statistic can also be written as

$$(\mathbf{y}_1 - \hat{\beta}_{IV}\mathbf{y}_2)^\top (\mathbf{M}_Z - \mathbf{M}_W) (\mathbf{y}_1 - \hat{\beta}_{IV}\mathbf{y}_2). \tag{6}$$

Similarly, the denominator is just

$$\frac{1}{n} (\mathbf{y}_1 - \mathbf{Z}\hat{\gamma}_{IV} - \hat{\beta}_{IV}\mathbf{y}_2)^\top (\mathbf{y}_1 - \mathbf{Z}\hat{\gamma}_{IV} - \hat{\beta}_{IV}\mathbf{y}_2) = \frac{1}{n} (\mathbf{y}_1 - \hat{\beta}_{IV}\mathbf{y}_2)^\top \mathbf{M}_Z (\mathbf{y}_1 - \hat{\beta}_{IV}\mathbf{y}_2). \tag{7}$$

In addition to being the numerator of the Sargan statistic, expression (6) is the numerator of the Anderson-Rubin, or AR, statistic for the hypothesis that $\beta = \hat{\beta}_{IV}$; see Anderson and Rubin [11]. The denominator of this same AR statistic is

$$\frac{1}{n-l} (\mathbf{y}_1 - \hat{\beta}_{IV}\mathbf{y}_2)^\top \mathbf{M}_W (\mathbf{y}_1 - \hat{\beta}_{IV}\mathbf{y}_2), \tag{8}$$

which may be compared to the right-hand side of Equation (7). We see that the Sargan statistic estimates σ_1^2 under the null hypothesis, and the AR statistic estimates it under the alternative.

Of course, AR statistics are usually calculated for the hypothesis that β takes on a specific value, say β_0 , rather than $\hat{\beta}_{IV}$. Since by definition $\hat{\beta}_{IV}$ minimizes the numerator (5), it follows that the numerator of the AR statistic is always no smaller than the numerator of the Sargan statistic. Even though the AR statistic is not generally thought of as a test of the overidentifying restrictions, it could be used as such a test, because it will always reject if the restrictions are sufficiently false, since that will cause the noncentrality parameter to be large.

It seems natural to modify the Sargan statistic by using (8) instead of (7) as the denominator, and this was done by Basmann [12]. The usual Sargan statistic can be written as

$$S = \frac{SSR_0 - SSR_1}{SSR_0/n} = n(1 - \zeta(\hat{\beta}_{IV})) \tag{9}$$

and the Basmann statistic as

$$S' = \frac{SSR_0 - SSR_1}{SSR_1/(n-l)} = (n-l)(\zeta^{-1}(\hat{\beta}_{IV}) - 1), \tag{10}$$

where SSR_0 is the sum of squared residuals from regressing $\mathbf{y}_1 - \hat{\beta}_{IV}\mathbf{y}_2$ on \mathbf{Z} , SSR_1 is the SSR from regressing $\mathbf{y}_1 - \hat{\beta}_{IV}\mathbf{y}_2$ on \mathbf{W} , and $\zeta(\hat{\beta}_{IV}) \equiv SSR_1/SSR_0$. Observe that both test statistics are simply monotonic functions of $\zeta(\hat{\beta}_{IV})$, the ratio of the two sums of squared residuals.

Another widely used test statistic for overidentification is the likelihood ratio, or LR, statistic associated with the LIML estimate $\hat{\beta}_{\text{LIML}}$. This statistic is simply $n \log \kappa(\hat{\beta}_{\text{LIML}})$, where

$$\kappa(\beta) = \frac{(\mathbf{y}_1 - \beta \mathbf{y}_2)^\top \mathbf{M}_Z (\mathbf{y}_1 - \beta \mathbf{y}_2)}{(\mathbf{y}_1 - \beta \mathbf{y}_2)^\top \mathbf{M}_W (\mathbf{y}_1 - \beta \mathbf{y}_2)}. \tag{11}$$

The LIML estimator $\hat{\beta}_{\text{LIML}}$ minimizes $\kappa(\beta)$ with respect to β . Since $\hat{\kappa} \equiv \kappa(\hat{\beta}_{\text{LIML}})$ is just $\zeta^{-1}(\hat{\beta}_{\text{LIML}})$, we see that the LR statistic is $-n \log \zeta(\hat{\beta}_{\text{LIML}})$. It is easy to show that, under both conventional and weak-instrument asymptotics, and under the null hypothesis, $\hat{\kappa} - 1 = O_p(n^{-1})$ as the sample size n tends to infinity. Therefore, the LR statistic is asymptotically equivalent to the linearized likelihood ratio statistic

$$\text{LR}' \equiv (n - l)(\hat{\kappa} - 1) = (n - l)\hat{\lambda} = (n - l)(\zeta^{-1}(\hat{\beta}_{\text{LIML}}) - 1) \tag{12}$$

where $\hat{\lambda} \equiv \hat{\kappa} - 1$. We define LR' as $(n - l)\hat{\lambda}$ rather than as $n\hat{\lambda}$ by analogy with (10). In what follows, it will be convenient to analyze LR' rather than LR.

We have seen that the Sargan statistic (9), the Basman statistic (10), and the two likelihood ratio statistics LR and LR' are all monotonic functions of the ratio of SSRs $\zeta(\hat{\beta})$ for some estimator $\hat{\beta}$. Both the particular function of $\zeta(\hat{\beta})$ that is used and the choice of $\hat{\beta}$ affect the finite-sample properties of an asymptotic test. For a bootstrap test, however, it is only the choice of $\hat{\beta}$ that matters. This follows from the fact that it is only the rank of the actual test statistic in the ordered list of the actual and bootstrap statistics that determines a bootstrap P value; see Section 6 below and Davidson and MacKinnon [13]. Therefore, for any given bootstrap data-generating process (DGP) and any estimator $\hat{\beta}$, bootstrap tests based on any monotonic transformation of $\zeta(\hat{\beta})$ yield identical results.

3. Analysis Using a Simpler Model

It is clear from (6), (7), and (11) that all the statistics we have considered for testing the overidentifying restrictions depend on \mathbf{y}_1 and \mathbf{y}_2 only through their projections $\mathbf{M}_Z \mathbf{y}_1$ and $\mathbf{M}_Z \mathbf{y}_2$. We see also that $\zeta(\beta)$ is homogeneous of degree zero with respect to $\mathbf{M}_Z \mathbf{y}_1$ and $\mathbf{M}_Z \mathbf{y}_2$ separately, for any β . Thus the statistics depend on the scale of neither \mathbf{y}_1 nor \mathbf{y}_2 . Moreover, the matrix \mathbf{Z} plays no essential role. In fact, it can be shown—see Davidson and MacKinnon [7], Section 3—that the distributions of the test statistics generated by the model (1) and (2) for sample size n are identical to those generated by the simpler model

$$\mathbf{y}_1 = \beta \mathbf{y}_2 + \mathbf{u}_1, \text{ and} \tag{13}$$

$$\mathbf{y}_2 = \mathbf{W}\boldsymbol{\pi} + \mathbf{u}_2, \tag{14}$$

where the sample size is now $n - k$, the matrix \mathbf{W} has $l - k$ columns, and $\sigma_1 = \sigma_2 = 1$.

In the remainder of the paper, we deal exclusively with the simpler model (13) and (14). By doing so, we avoid the additional algebra made necessary by the presence of \mathbf{Z} in Equation (1) without changing the results in any essential way. Of course, \mathbf{y}_1 , \mathbf{y}_2 , and \mathbf{W} in the simpler model are not the same as in the original model, although ρ is unchanged. For the original model, n and l in our results below would have to be replaced by $n - k$ and $l - k$, and \mathbf{y}_1 and \mathbf{y}_2 would have to be replaced by $\mathbf{M}_Z \mathbf{y}_1$ and $\mathbf{M}_Z \mathbf{y}_2$.

It is well known—see Mariano and Sawa [14]—that all the test statistics depend on the data generated by (13) and (14) only through the six quadratic forms

$$\begin{aligned}
 P_{11} &\equiv \mathbf{y}_1^\top \mathbf{P}_W \mathbf{y}_1, & P_{12} &\equiv \mathbf{y}_1^\top \mathbf{P}_W \mathbf{y}_2, & P_{22} &\equiv \mathbf{y}_2^\top \mathbf{P}_W \mathbf{y}_2, \\
 M_{11} &\equiv \mathbf{y}_1^\top \mathbf{M}_W \mathbf{y}_1, & M_{12} &\equiv \mathbf{y}_1^\top \mathbf{M}_W \mathbf{y}_2, & \text{and } M_{22} &\equiv \mathbf{y}_2^\top \mathbf{M}_W \mathbf{y}_2.
 \end{aligned}
 \tag{15}$$

This is also true for the general model (1) and (2), except that the projection matrix \mathbf{P}_W must be replaced by $\mathbf{P}_W - \mathbf{P}_Z = \mathbf{P}_W \mathbf{M}_Z$.

In this section and the next, we make the additional assumption that the disturbances \mathbf{u}_1 and \mathbf{u}_2 are normally distributed. Since the quadratic forms in (15) depend on the instruments only through the projections \mathbf{P}_W and \mathbf{M}_W , it follows that their joint distribution depends on \mathbf{W} only through the number of instruments l and the norm of the vector $\mathbf{W}\boldsymbol{\pi}$. We can therefore further simplify Equation (14) as

$$\mathbf{y}_2 = a\mathbf{w} + \mathbf{u}_2, \tag{16}$$

where the vector $\mathbf{w} \in \mathcal{S}(\mathbf{W})$ is normalized to have length unity, which implies that $a^2 = \boldsymbol{\pi}^\top \mathbf{W}^\top \mathbf{W} \boldsymbol{\pi}$. Thus the joint distribution of the six quadratic forms depends only on the three parameters β , a , and ρ , and on the dimensions n and l ; for the general model (1) and (2), the latter would be $n - k$ and $l - k$.

The above simplification was used in Davidson and MacKinnon [7] in the context of tests of hypotheses about β , and further details can be found there. The parameter a determines the strength of the instruments. In weak-instrument asymptotics, $a = O(1)$, while in conventional strong-instrument asymptotics, $a = O(n^{1/2})$. Thus, by treating a as a parameter of order unity, we are in the context of weak-instrument asymptotics; see Staiger and Stock [1]. The square of the parameter a is often referred to as the (scalar) concentration parameter; see Phillips [15] (p. 470), and Stock, Wright, and Yogo [2].

Davidson and MacKinnon [7] show that, under the assumption of normal disturbances, the six quadratic forms (15) can be expressed as functions of the three parameters β , a , and ρ and eight mutually independent random variables, the distributions of which do not depend on any of the parameters. Four of these random variables, which we denote by x_1 , x_2 , z_P , and z_M , are standard normal, and the other four, which we denote by t_{11}^P , t_{22}^P , t_{11}^M , and t_{22}^M , are respectively distributed as χ_{l-2}^2 , χ_{l-1}^2 , χ_{n-l}^2 , and χ_{n-l-1}^2 . In terms of these eight variables, we make the definitions

$$\begin{aligned}
 Q_{11} &\equiv x_1^2 + z_P^2 + t_{11}^P, & Q_{12} &\equiv x_1 x_2 + z_P \sqrt{t_{22}^P}, & Q_{22} &\equiv x_2^2 + t_{22}^P, \\
 N_{11} &\equiv t_{11}^M, & N_{12} &\equiv z_M \sqrt{t_{11}^M}, & \text{and } N_{22} &\equiv z_M^2 + t_{22}^M.
 \end{aligned}
 \tag{17}$$

These quantities have simple interpretations: $Q_{ij} = \mathbf{u}_i^\top \mathbf{P}_W \mathbf{u}_j$, and $N_{ij} = \mathbf{u}_i^\top \mathbf{M}_W \mathbf{u}_j$, for $i = 1, 2$.

Theorem 1. The Basman statistic S' defined in (10) can be expressed as

$$S' = \frac{(n - l)(P_{11}P_{22} - P_{12}^2)}{M_{11}P_{22} - 2P_{12}M_{12} + P_{12}^2M_{22}/P_{22}}, \tag{18}$$

and the statistic $\text{LR}' \equiv (n - l)\hat{\lambda}$ defined in (12) can be expressed as

$$\text{LR}' = \frac{(n - l)(P_{11}M_{22} - 2P_{12}M_{12} + P_{22}M_{11} - \Delta^{1/2})}{2(M_{11}M_{22} - M_{12}^2)}, \tag{19}$$

where

$$\begin{aligned} P_{11} &= Q_{11}, & P_{12} &= ax_1 + \rho Q_{11} + rQ_{12}, & M_{11} &= N_{11}, \\ P_{22} &= a^2 + 2a(\rho x_1 + rx_2) + \rho^2 Q_{11} + 2r\rho Q_{12} + r^2 Q_{22}, \\ M_{22} &= \rho^2 N_{11} + 2r\rho N_{12} + r^2 N_{22}, & \text{and } M_{12} &= \rho N_{11} + rN_{12}, \end{aligned} \tag{20}$$

with $r = \sqrt{1 - \rho^2}$, and

$$\Delta = (P_{11}M_{22} - 2P_{12}M_{12} + P_{22}M_{11})^2 - 4(M_{11}M_{22} - M_{12}^2)(P_{11}P_{22} - P_{12}^2).$$

Proof: See the Appendix.

Remark 1. By substituting the relations (20) and (17) into (18) and (19), realizations of S' and LR' can be generated as functions of realizations of the eight independent random variables. Since these have known distributions, it follows that Theorem 1 provides a complete characterization of the distributions of the statistics S' and LR' in finite samples for arbitrary values of the parameters β , a , and ρ . Although it seems to be impossible to obtain closed-form solutions for their distribution functions, it is of course possible to approximate them arbitrarily well by simulation.

4. Particular Cases

In this section, we show that no test of the overidentifying restrictions is robust to weak instruments. In fact, the distributions of S' and LR' have a singularity at the point in the parameter space at which $a = 0$ and $\rho = \pm 1$, or, equivalently, $a = r = 0$. By letting $n \rightarrow \infty$ for arbitrary values of a and r , we may obtain asymptotic results for both weak and strong instruments. In this way, we show that the singularity is present both in finite samples and under weak-instrument asymptotics. For conventional asymptotics, it is necessary to let both a and n tend to infinity, and in that case the singularity becomes irrelevant.

Like Theorem 1, the theorems in this section provide finite-sample results. Asymptotic results follow from these by letting the sample size n tend to infinity.

Theorem 2. When the parameter a is zero, the linearized likelihood ratio statistic LR' , in terms of the random variables defined in (17), is

$$\frac{(n - l)(Q_{11}N_{22} - 2Q_{12}N_{12} + Q_{22}N_{11} - \Delta^{1/2})}{2(N_{11}N_{22} - N_{12}^2)}, \tag{21}$$

where the discriminant Δ has become

$$(Q_{11}N_{22} - Q_{22}N_{11})^2 + 4(Q_{12}N_{11} - Q_{11}N_{12})(Q_{12}N_{22} - Q_{22}N_{12}), \tag{22}$$

independent of r . When $r = 0$, LR' is

$$(n - l)(Q_{11} - x_1^2)/N_{11}, \tag{23}$$

independent of a .

Proof. See the Appendix.

Remark 2.1. As with Theorem 1, the expressions (21)–(23) constitute a complete characterization of the distribution of LR' in the limits when $a = 0$ and when $r = 0$.

Remark 2.2. The limiting value (21) for $a = 0$ is independent of r . Thus the distribution of LR' in the limit of completely irrelevant instruments is independent of all the model parameters.

Remark 2.3. The limiting value (23) for $r = 0$ is independent of a , and it tends to a χ^2_{l-1} variable as $n \rightarrow \infty$.

Remark 2.4. The singularity mentioned above is a consequence of the fact that the limit at $a = r = 0$ is ill-defined, since LR' converges to two different random variables as $r \rightarrow 0$ for $a = 0$ and as $a \rightarrow 0$ for $r = 0$. These random variables are quite different and have quite different distributions.

Remark 2.5. Simulation results show that the limit to which LR' converges as it approaches the singularity with $a = \alpha r$ for some positive α depends on α . In other words, the limit depends on the direction from which the singularity is reached.

Remark 2.6. The limit as $r \rightarrow 0$ may appear to be of little interest. However, it may have a substantial impact on bootstrap inference. Davidson and MacKinnon [7] show that a and r are not consistently estimated with weak instruments, and so, even if the true value of r is far from zero, the estimate needed for the construction of a bootstrap DGP (see Section 5) may be close to zero.

Remark 2.7. In the limit of strong instruments, the distribution of LR' is the same as its distribution when $r = 0$, as shown by the following corollary.

Corollary 2.1. The limit of LR' as $a \rightarrow \infty$, which is the limit when the instruments are strong, is

$$(n - l)(Q_{11} - x_1^2)/N_{11}. \tag{24}$$

Proof of Corollary 2.1. Isolate the coefficients of powers of a in (19), and perform a Taylor expansion for small $1/a$. The result (24) follows easily. \square

As $n \rightarrow \infty$, $N_{11}/n \rightarrow 1$, and so the result (24) confirms that the asymptotic distribution with strong instruments is just χ^2_{l-1} . In fact, Guggenberger *et al.* [9] demonstrate that, for our simple model, the asymptotic distribution is bounded by χ^2_{l-1} however weak or strong the instruments may be. Some of the simulation results in Section 6 for finite samples are in accord with this result.

Theorem 3. The Basmann statistic S' in the limit with $a = 0$ becomes

$$(n - l) \frac{(Q_{11}Q_{22} - Q_{12}^2)(\rho^2Q_{11} + 2r\rho Q_{12} + r^2Q_{22})}{\rho^2D_0 + 2r\rho D_1 + r^2D_2}, \tag{25}$$

where

$$\begin{aligned} D_0 &= Q_{12}^2N_{11} - 2Q_{11}Q_{12}N_{12} + Q_{11}^2N_{22}, \\ D_1 &= Q_{12}Q_{22}N_{11} - N_{12}(Q_{11}Q_{22} + Q_{12}^2) + Q_{11}Q_{12}N_{22}, \text{ and} \\ D_2 &= Q_{22}^2N_{11} - 2Q_{12}Q_{22}N_{12} + Q_{12}^2N_{22}. \end{aligned}$$

The limit with $r = 0$ is

$$\frac{(n - l)(Q_{11} - x_1^2)(a^2 + 2ax_1 + Q_{11})}{N_{11}(a + x_1)^2}. \tag{26}$$

Proof. The proof is omitted. It involves tedious algebra similar to that in the proof of Theorem 2; see the Appendix. \square

Remark 3.1. The expression (25) for $a = 0$ depends on r , unlike the analogous expression for LR'. When $r \rightarrow 0$ with $a = 0$, it is easy to see that S' tends to the limit

$$(n - l) \frac{Q_{11}(Q_{11}Q_{22} - Q_{12}^2)}{Q_{12}^2 N_{11} - 2Q_{11}Q_{12}N_{12} + Q_{11}^2 N_{22}}. \tag{27}$$

Remark 3.2. Similarly, the expression (26) for $r = 0$ depends on a . Its limit as $a \rightarrow 0$ is

$$(n - l) \frac{Q_{11}(Q_{11} - x_1^2)}{N_{11}x_1^2}, \tag{28}$$

which is quite different from (27), where the order of the limits is inverted. As with LR', the limit as the singularity is approached depends on the direction of approach.

Remark 3.3. As expected, the limit of S' as $a \rightarrow \infty$ is the same as that of LR'.

The fact that the test statistics S' and LR' depend on the parameters a and ρ indicates that these statistics are not robust to weak instruments. Passing to the limit as $n \rightarrow \infty$ with weak-instrument asymptotics does not improve matters. Of the six quadratic forms on which everything depends, only the M_{ij} depend on n . Their limiting behavior is such that $M_{11}/n \rightarrow 1$, $M_{22}/n \rightarrow 1$, and $M_{12}/n \rightarrow \rho$ as $n \rightarrow \infty$. In contrast, the P_{ij} do not depend on n , and they do depend on a and ρ .

5. Bootstrap Tests

Every test statistic has a distribution which depends on the DGP that generated the sample from which it is computed. The “true” DGP that generated an observed realization of the statistic is in general unknown. However, according to the bootstrap principle, one can perform inference by replacing the unknown DGP by an estimate of it, which is called the bootstrap DGP. Because what we need for inference is the distribution of the statistic under DGPs that satisfy the null hypothesis, the bootstrap DGP must necessarily impose the null. This requirement by itself does not lead to a unique bootstrap DGP, and we will see in this section that, for an overidentification test, there are several plausible choices.

If the observed value of a test statistic τ is $\hat{\tau}$, and the rejection region is in the upper tail, then the bootstrap P value is the probability, under the bootstrap distribution of the statistic, that τ is greater than $\hat{\tau}$. To estimate this probability, one generates a large number, say B , of realizations of the statistic using the bootstrap DGP. Let the j th realization be denoted by τ_j^* . Then the simulation-based estimate of the bootstrap P value is just the proportion of the τ_j^* greater than $\hat{\tau}$:

$$\hat{p}^*(\hat{\tau}) = \frac{1}{B} \sum_{j=1}^B \mathbb{I}(\tau_j^* > \hat{\tau}),$$

where $\mathbb{I}(\cdot)$ is the indicator function, which is equal to 1 when its argument is true and 0 otherwise. If this fraction is smaller than α , the level of the test, then we reject the null hypothesis. See Davidson and MacKinnon [13].

5.1. Parametric Bootstraps

The DGPs contained in the simple model defined by Equations (13) and (16) are characterized by just three parameters, namely, β , a , and ρ . Since the value of β does not affect the distribution of the overidentification test statistics (see Lemma 1 in the Appendix), the bootstrap DGP for a parametric bootstrap (assuming normally distributed disturbances) is completely determined by the values of a and ρ that characterize it.

The test statistic $\hat{\tau}$ itself may be any one of the overidentification statistics we have discussed. The model that is actually estimated in order to obtain $\hat{\tau}$ is not the simple model, but rather the full model given by Equations (1) and (2). The parameters of this model include some whose values do not interest us for the purpose of defining a bootstrap DGP: β , since it has no effect on the distribution of the statistic, and γ , since the matrix Z plays no role in the simple model, from which the bootstrap DGP is taken. There remain π , ρ , σ_1 , and σ_2 .

For Equation (16), the parameter a was defined as the square root of $\pi^T W^T W \pi / \sigma_2^2$. However, this definition assumes that the vector w has unit length and that all the variables are scaled so that the variance of the disturbances u_2 is 1. A more general definition of a that does not depend on these assumptions is

$$a = \sqrt{\pi^T W^T W \pi / \sigma_2^2} \tag{29}$$

It follows from (29) that, in order to estimate a , it is necessary also to estimate σ_2^2 .

Since the parameter ρ is the correlation of the disturbances, which are not observed, any estimate of ρ must be based on the residuals from the estimation of Equations (1) and (2). Let these residuals be denoted by \ddot{u}_1 and \ddot{u}_2 . Then the obvious estimators of the parameters of the covariance matrix are

$$\ddot{\sigma}_1^2 = n^{-1} \ddot{u}_1^T \ddot{u}_1, \quad \ddot{\sigma}_2^2 = n^{-1} \ddot{u}_2^T \ddot{u}_2, \quad \text{and} \quad \ddot{\rho} = n^{-1} \ddot{u}_1^T \ddot{u}_2 / (\ddot{\sigma}_1^2 \ddot{\sigma}_2^2)^{1/2},$$

and the obvious estimator of a is given by

$$\ddot{a}^2 = n \ddot{\pi}^T W^T W \ddot{\pi} / \ddot{u}_2^T \ddot{u}_2, \tag{30}$$

where $\ddot{\pi}$ estimates π . For \ddot{u}_1 , there are two obvious choices, the IV residuals and the LIML residuals from (1). For \ddot{u}_2 , the obvious choice is the vector of OLS residuals from (2), possibly scaled by a factor of $(n/(n-l))^{1/2}$ to take account of the lost degrees of freedom in the OLS estimation. However, this obvious choice is not the only one, because, if we treat the model (1) and (2) as a system, the system estimator of π that comes with the IV estimator of β is the three-stage least squares (3SLS) estimator, and the one that comes with the LIML estimator of β is the full-information maximum likelihood (FIML) estimator. These system estimators give rise to estimators not only of π , but also of u_2 , that differ from those given by OLS.

The system estimators of π can be computed without actually performing a system estimation, by running the regression

$$y_2 = W\pi + \varphi \ddot{u}_1 + \text{residuals}; \tag{31}$$

see Davidson and MacKinnon [7], where this matter is discussed in greater detail. If $\ddot{\mathbf{u}}_1$ is the vector of IV residuals, then the corresponding estimator $\ddot{\pi}$ is the 3SLS estimator; if it is the vector of LIML residuals, then $\ddot{\pi}$ is the FIML estimator.

For the purpose of computation, it is worth noting that all these estimators can be expressed as functions of the six quadratic forms (15). A short calculation shows that the estimators of a^2 and ρ based on IV residuals, scaled OLS residuals, and the OLS estimator of π are

$$\hat{a}^2 = (n - l) \frac{P_{22}}{M_{22}} \quad \text{and} \quad \hat{\rho} = \frac{n^{-1}(M_{12} - \hat{b}M_{22})}{\hat{\sigma}_1 \hat{\sigma}_2}, \tag{32}$$

where $\hat{b} = P_{12}/P_{22}$ is the difference between the IV estimator of β and the true β of the DGP, and where

$$\hat{\sigma}_1^2 = n^{-1}(Q_{11} + N_{11} - 2\hat{b}(P_{12} + M_{12}) + \hat{b}^2(P_{22} + M_{22})), \quad \text{and} \quad \hat{\sigma}_2^2 = M_{22}/(n - l).$$

From (17), we can express \hat{a}^2 as

$$\frac{(n - l)(a^2 + 2a(\rho x_1 + r x_2) + \rho^2 Q_{11} + 2r\rho Q_{12} + r^2 Q_{22})}{\rho^2 N_{11} + 2r\rho N_{12} + r^2 N_{22}}.$$

The weak-instrument asymptotic limit of this expression replaces the denominator divided by $n - l$ by 1. The expectation of the numerator without the factor of $n - l$ is $a^2 + \rho^2 l + r^2 l = a^2 + l$. Consequently, it may be preferable to reduce bias in the estimation of a^2 by setting $\hat{a}^2 = (n - l) \max(0, P_{22}/M_{22} - l)$; see Davidson and MacKinnon [7].

The system estimator of the vector π , which we write as $\ddot{\pi}$, is computed by running regression (31). A tedious calculation shows that

$$\ddot{\pi}^\top \mathbf{W}^\top \mathbf{W} \ddot{\pi} = P_{22} + \frac{\ddot{P}_{11} \ddot{M}_{12}^2}{\ddot{M}_{11}^2} - 2 \frac{\ddot{P}_{12} \ddot{M}_{12}}{\ddot{M}_{11}}, \tag{33}$$

with

$$\begin{aligned} \ddot{P}_{11} &= \ddot{\mathbf{u}}_1^\top \mathbf{P}_W \ddot{\mathbf{u}}_1 = P_{11} - 2\ddot{b}P_{12} + \ddot{b}^2 P_{22}, \\ \ddot{P}_{12} &= \ddot{\mathbf{u}}_1^\top \mathbf{P}_W \mathbf{y}_2 = P_{12} - \ddot{b}P_{22}, \\ \ddot{M}_{11} &= \ddot{\mathbf{u}}_1^\top \mathbf{M}_W \ddot{\mathbf{u}}_1 = M_{11} - 2\ddot{b}M_{12} + \ddot{b}^2 M_{22}, \quad \text{and} \\ \ddot{M}_{12} &= \ddot{\mathbf{u}}_1^\top \mathbf{M}_W \mathbf{y}_2 = M_{12} - \ddot{b}M_{22}, \end{aligned}$$

where \ddot{b} is $\hat{b} = \hat{\beta}_{IV} - \beta$ and $\ddot{\mathbf{u}}_1$ is the vector of IV residuals if the structural Equation (1) is estimated by IV, and \tilde{b} is $\tilde{b} = \hat{\beta}_{LIML} - \beta$ and $\tilde{\mathbf{u}}_1$ is the vector of LIML residuals if (1) is estimated by LIML. Then, if σ_2^2 is estimated using the residuals $\mathbf{y}_2 - \mathbf{W} \ddot{\pi}$, we find that

$$\ddot{\sigma}_2^2 = \frac{1}{n - l} \left(M_{22} + \frac{\ddot{M}_{12}^2}{\ddot{M}_{11}^2} \ddot{P}_{11} \right). \tag{34}$$

Using the results (33) and (34) in (30) allows us to write

$$\ddot{a}^2 = \frac{(n - l)(P_{22} \ddot{M}_{11}^2 + \ddot{P}_{11} \ddot{M}_{12}^2 - 2\ddot{P}_{12} \ddot{M}_{11} \ddot{M}_{12})}{\ddot{M}_{11}^2 M_{22} + \ddot{M}_{12}^2 \ddot{P}_{11}}.$$

For the parameter ρ , more calculation shows that

$$\ddot{\rho} = \ddot{M}_{12} \left[\frac{\ddot{P}_{11} + \ddot{M}_{11}}{\ddot{M}_{11}^2 M_{22} + \ddot{M}_{12}^2 \ddot{P}_{11}} \right]^{1/2}.$$

We consider four different bootstrap DGPs. The simplest, which we call the IV-R bootstrap, uses the IV and OLS estimates of a and ρ given in (32). The IV-ER bootstrap is based on 3SLS estimation of the two-equation system, that is, on IV estimation of (1) and OLS estimation of regression (31) with \ddot{u}_1 the vector of IV residuals. Similarly, the LIML-ER bootstrap relies on FIML estimation of the system, that is, on LIML for (1) and on regression (31) with \ddot{u}_1 the LIML residuals. Finally, we also define the F(1)-ER bootstrap, which is the same as LIML-ER except that $\hat{\beta}_{\text{LIML}}$ is replaced at every step by the modified LIML estimator of Fuller [16] with $\eta = 1$ which is discussed in Section 6.

It is plain that, the closer the bootstrap DGP is to the true DGP, the better will be bootstrap inference; see Davidson and MacKinnon [17]. We may therefore expect that IV-ER should perform better than IV-R, and that LIML-ER should perform better than IV-ER. Between LIML-ER and F(1)-ER, there is no obvious reason *a priori* to expect that one of them would outperform the other. But, whatever the properties of these bootstraps may be when the true DGP is not in the neighborhood of the singularity at $a = 0, \rho = 1$, we cannot expect anything better than some improvement over inference based on asymptotic critical values, rather than truly reliable inference, in the neighborhood of the singularity.

5.2. Resampling

Any parametric bootstrap risks being unreliable if the strong assumptions used to define the null hypothesis are violated. Most practitioners would therefore prefer a more robust bootstrap method. The strongest assumption we have made so far is that the disturbances are normally distributed. It is easy to relax this assumption by using a bootstrap DGP based on resampling, in which the bivariate normal distribution is replaced by the joint empirical distribution of the residuals. The discussion of the previous subsection makes it clear that several resampling bootstraps can be defined, depending on the choice of residuals that are resampled.

The most obvious resampling bootstrap DGP in the context of IV estimation is

$$y_1^* = \hat{\beta}_{\text{IV}} y_2^* + Z \hat{\gamma}_{\text{IV}} + \hat{u}_1^* \tag{35}$$

$$y_2^* = W \hat{\pi} + \hat{u}_2^*, \tag{36}$$

where y_1^* and y_2^* are n -vectors of bootstrap observations, \hat{u}_1^* and \hat{u}_2^* are n -vectors of bootstrap disturbances with typical elements \hat{u}_{1i}^* and \hat{u}_{2i}^* , respectively, and $\hat{\pi}$ is the OLS estimate from Equation (2). The bootstrap disturbances are drawn in pairs from the bivariate empirical distribution of the structural residuals \hat{u}_{1i}^{IV} and the rescaled reduced-form residuals $(n/(n-l))^{1/2} \hat{u}_{2i}^{\text{OLS}}$. That is,

$$\begin{bmatrix} \hat{u}_{1i}^* \\ \hat{u}_{2i}^* \end{bmatrix} \sim \text{EDF} \left(\begin{pmatrix} \hat{u}_{1i}^{\text{IV}} \\ (n/(n-l))^{1/2} \hat{u}_{2i}^{\text{OLS}} \end{pmatrix} \right). \tag{37}$$

Here EDF stands for “empirical distribution function”.

The rescaling of the reduced form residuals \hat{u}_{2i}^{OLS} in (37) ensures that the distribution of the \hat{u}_{2i}^* has variance equal to the unbiased OLS variance estimator. The fact that the bootstrap disturbances are drawn in pairs ensures that the bootstrap DGP mimics the correlation between the two sets of residuals. Of course, whether that correlation consistently estimates the correlation between the actual disturbances depends on whether or not $\hat{\rho}$ can be estimated consistently. Under weak-instrument asymptotics, $\hat{\rho}$ cannot be estimated consistently; see Davidson and MacKinnon [7].

Since all of the overidentification test statistics are invariant to the values of β and γ , we may replace the bootstrap DGP for \mathbf{y}_1^* given by (35) by

$$\mathbf{y}_1^* = \hat{\mathbf{u}}_1^*. \tag{38}$$

The bootstrap statistics generated by (38) and (36) are identical to those generated by (35) and (36). We will refer to the bootstrap DGP given by (36)–(38), as the IV-R resampling bootstrap. It is a semiparametric bootstrap, because it uses parameter estimates of the reduced-form equation, but it does not assume a specific functional form for the joint distribution of the disturbances. The empirical distribution of the residuals has a covariance matrix which is exactly that used to estimate a and ρ by the IV-R parametric bootstrap; hence our nomenclature.

The IV-ER resampling bootstrap draws pairs from the joint EDF of the IV residuals $\hat{\mathbf{u}}_1^{IV}$ from Equation (1) and the residuals $\mathbf{y}_2 - \mathbf{W}\hat{\pi}_{IV}$ computed by running regression (31) with $\hat{\mathbf{u}}_1^{IV}$ replacing $\hat{\mathbf{u}}_1$. It also uses the resulting estimator $\hat{\pi}_{IV}$ in (36) instead of the OLS estimator $\hat{\pi}$. Note that the residuals $\mathbf{y}_2 - \mathbf{W}\hat{\pi}_{IV}$ are *not* the residuals from regression (31), but rather those residuals plus $\hat{\varphi}\hat{\mathbf{u}}_1^{IV}$.

The LIML-ER resampling bootstrap is very similar to the IV-ER one, except that it uses $\hat{\mathbf{u}}_1^{LIML}$ both directly and in regression (31). Formally, the resampling procedure draws pairs from the bivariate empirical distribution of

$$\begin{bmatrix} \hat{u}_{i1}^{LIML} \\ \hat{u}_{i2}^{LIML} \end{bmatrix} = \begin{bmatrix} y_{i1} - \hat{\beta}_{LIML}y_{2i} - \mathbf{Z}_i\hat{\gamma}_{LIML} \\ y_{i2} - \mathbf{W}_i\hat{\pi}_{LIML} \end{bmatrix}. \tag{39}$$

Similarly, for the F(1)-ER resampling bootstrap, the structural Equation (1) is estimated by Fuller’s method, and the residuals from this are used both for resampling and in regression (31); see the next section.

A word of caution is advisable here. Although the values of overidentification test statistics are invariant to β , thereby allowing us to use (38) instead of (35) in the bootstrap DGP, the residuals from which we resample in (37) and (39) do depend on the estimate of β , as does the estimate of π if it is based on any variant of Equation (31). Nevertheless, the test statistics depend on the estimate of β only through the residuals and the estimate of π .

6. Performance of Asymptotic Tests: Simulation Results

The discussion in Section 4 was limited to the statistics S' and LR' . For bootstrap tests, it is enough to consider just these two, since all other statistics mentioned in Section 2 are monotonic transforms of them. Of course, the different versions of the Sargan test and the LR test have different properties when used with (strong-instrument) asymptotic critical values. In this section, we therefore present some Monte Carlo results on the finite-sample performance of five test statistics, including S , S' , LR , and LR' .

The fifth test statistic we examine is based on the estimator proposed by Fuller [16]. Like the IV and LIML estimators, Fuller's estimator is a K -class estimator for model (1) and (2). It takes the form

$$\begin{bmatrix} \hat{\beta} \\ \hat{\gamma} \end{bmatrix} = (\mathbf{X}^\top(\mathbf{I} - K\mathbf{M}_W)\mathbf{X})^{-1}\mathbf{X}^\top(\mathbf{I} - K\mathbf{M}_W)\mathbf{y}_1. \quad (40)$$

Setting $K = \hat{\kappa}$, the minimized value of the variance ratio (11), in Equation (40) gives the LIML estimator, while setting $K = 1$ gives the IV estimator. Fuller's estimator sets $K = \hat{\kappa} + \eta/(n - l)$ for some fixed number $\eta > 0$ independent of the sample size n . We set $\eta = 1$. With this choice, Fuller's estimator $\hat{\beta}_F$ has all moments (except when the sample size is very small) and is approximately unbiased. The corresponding test statistic is simply $-n \log \zeta(\hat{\beta}_F)$, which has the same form as the LR statistic. We will refer to this as the LRF test.

The DGPs used for our simulations all belong to the simplified model (13) and (16). The disturbances are generated according to the relations

$$\mathbf{u}_1 = \mathbf{v}_1 \quad \text{and} \quad \mathbf{u}_2 = \rho\mathbf{v}_1 + r\mathbf{v}_2,$$

where \mathbf{v}_1 and \mathbf{v}_2 are n -vectors with independent standard normal elements, and $r \equiv (1 - \rho^2)^{1/2}$. Of course, it is quite unnecessary to generate simulated samples of n observations, as it is enough to generate the six quadratic forms (15) as functions of eight mutually independent random variables, using the relations (17) and (20). The sample size n affects only the degrees of freedom of the two χ^2 random variables t_{11}^M and t_{22}^M that appear in (17). Although any DGP given by (13) and (16) involves no explicit overidentifying restrictions, the test statistics are computed for the model (1) and (2), for which there are $q \equiv l - k - 1$ of them, with $k = 0$ in the experiments.

The first group of experiments is intended to provide guidance on the appropriate sample size to use in the remaining experiments. Our objective is to mimic the common situation in which the sample size is reasonably large and the instruments are quite weak. Since our simulation DGPs embody weak-instrument asymptotics, we should not expect any of the test statistics to have the correct size asymptotically. However, for any given a and ρ , the rejection frequency converges as $n \rightarrow \infty$ to that given by the asymptotic distribution of the statistic used; these asymptotic distributions were discussed at the end of Section 4. In the experiments, we use sample sizes of 20, 28, 40, 56, and so on, up to 1810. Each of these numbers is larger than its predecessor by approximately $\sqrt{2}$. Each experiment used 10^6 replications.

The results of four sets of experiments are presented in Figure 1, in which we plot rejection frequencies in the experiments for a nominal level of 5%. In the top two panels, $a = 2$, so that the instruments are very weak. In the bottom two panels, $a = 8$, so that they are reasonably strong. Recall that the concentration parameter is a^2 . In the two panels on the left, $\rho = 0.5$, so that there is moderate correlation between the structural and reduced form disturbances. In the two panels on the right, $\rho = 0.9$, so that there is strong correlation. Note that the vertical axis differs across most of the panels.




Figure 1. Rejection frequencies for asymptotic tests as functions of n for $q = 8$. **(a)** $a = 2, \rho = 0.5$; **(b)** $a = 2, \rho = 0.9$; **(c)** $a = 8, \rho = 0.5$; **(d)** $a = 8, \rho = 0.9$. Note that the scale of the vertical axis differs across panels.

It is evident that the various asymptotic tests can perform very differently and that their performance varies greatly with the sample size. The Basmann test S' always rejects more often than LR'. This must be the case, because the two test statistics are exactly the same function of ζ — compare Equations (10) and (12) — but the LIML estimator minimizes ζ^{-1} and the IV estimator does not. Thus $\zeta^{-1}(\hat{\beta}_{LIML})$ must be closer to unity than $\zeta^{-1}(\hat{\beta}_{IV})$.

The Sargan (S) and Basmann (S') tests perform almost the same for large samples but very differently for small ones, with the latter much more prone to overreject than the former. This could have been predicted from Equations (9) and (10). If it were not for the different factors of n and $n - l$, S' would always be greater than S , because $\zeta^{-1} - 1 > 1 - \zeta$. For $\zeta \approx 1$, the difference between $\zeta^{-1} - 1$ and $1 - \zeta$

is $O(\zeta^{-2})$. Therefore, under the null hypothesis, the difference between the two test statistics is $O(n^{-1})$. This explains their rapid convergence as n increases.

For $a = 2$, the LR test and its linearized version LR' perform somewhat differently in small samples but almost identically once $n \geq 200$, with LR rejecting more often than LR'. The difference between the two LIML-based statistics is much smaller than the difference between the two IV-based statistics for two reasons. First, because the LIML estimator minimizes ζ^{-1} , the two LR test statistics are being evaluated at values of ζ that are closer to unity than the two IV-based statistics. Second, $(\zeta^{-1} - 1) + \log(\zeta)$ is almost exactly half the magnitude of $(\zeta^{-1} - 1) - (1 - \zeta)$. Based on the fact that $(\zeta^{-1} - 1) > -\log(\zeta)$, we might expect LR' to reject more often than LR. However, the difference between $n - l$ and n evidently more than offsets this, at least for $l = 9$, the case in the figure.

In this case, the Fuller variant of the LR test performs somewhat differently from both LR and LR' for all sample sizes. In contrast, for $a = 8$, LR and LRF are so similar that we did not graph LR to avoid making the figure unreadable. LR, LR', and LRF perform almost identically, and very well indeed, for large sample sizes, even though they overreject severely for small sample sizes.

As expected, all of the rejection frequencies seem to be converging to constants as $n \rightarrow \infty$. Moreover, in every case, it appears that the (interpolated) results for $n = 400$ are very similar to the results for larger values up to $n = 1810$. Accordingly, we used $n = 400$ in all the remaining experiments.

In the second group of experiments, the number of overidentifying restrictions q is varied. The four panels in Figure 2 correspond to those of Figure 1. In most cases, performance deteriorates as q increases. Sometimes, rejection frequencies seem to be converging, but by no means always. In the remaining experiments, we somewhat arbitrarily set $q = 8$. Choosing a smaller number would generally have resulted in smaller size distortions.

In the third group of experiments, the results of which are shown in Figure 3, we set $n = 400$ and $q = 8$, and we vary ρ between 0.0 and 0.99 at intervals of 0.01 for four values of a . The vertical axis is different in each of the four panels, because the tests all perform much better as a increases. For clarity, rejection frequencies for LR are not shown in the figure, because they always lie between those for LR' and LRF. They are very close to those for LR' when a is small, and they are very close to those for LRF when a is large.

For the smaller values of a , all of the tests can either overreject or underreject, with rejection frequencies increasing in ρ . The Sargan and Basmann tests overreject very severely when a is small and ρ is large. The LR', LR, and LRF tests underreject severely when a is small and ρ is not large, but they overreject slightly when a is large. Based on Figure 1 and on the analysis of the previous section, we expect that this slight overrejection vanishes for larger samples.




Figure 2. Rejection frequencies for asymptotic tests as functions of q for $n = 400$. **(a)** $a = 2, \rho = 0.5$; **(b)** $a = 2, \rho = 0.9$; **(c)** $a = 8, \rho = 0.5$; **(d)** $a = 8, \rho = 0.9$. Note that the scale of the vertical axis differs across panels.

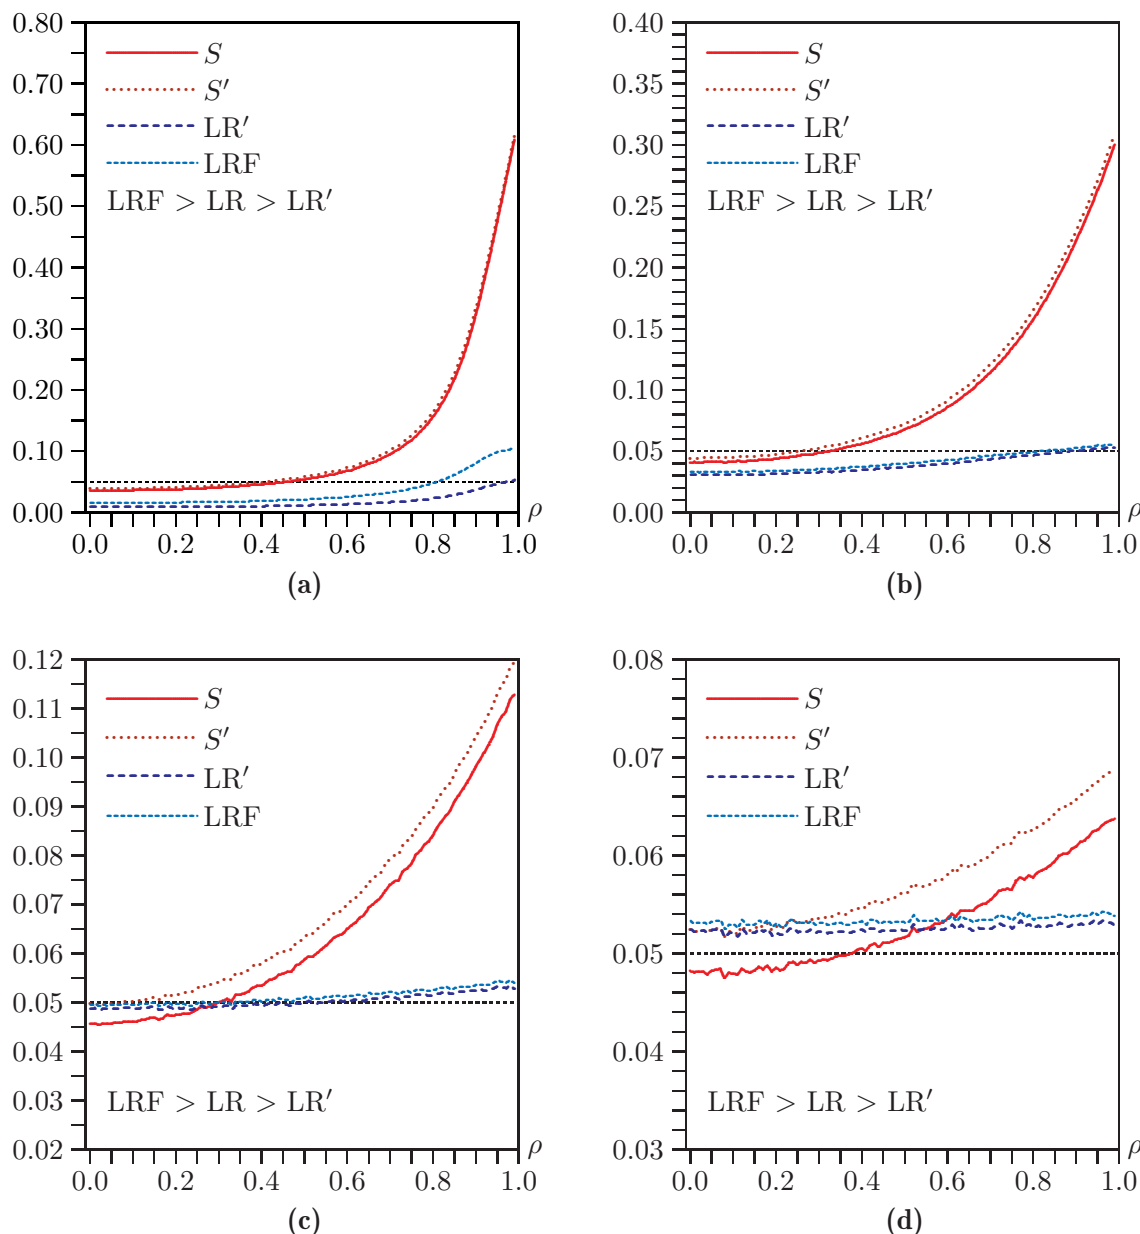


Figure 3. Rejection frequencies for asymptotic tests as functions of ρ for $q = 8$ and $n = 400$. (a) Very weak instruments: $a = 2$; (b) Weak instruments: $a = 4$; (c) Moderately strong instruments: $a = 8$; (d) Very strong instruments: $a = 16$. Note that the scale of the vertical axis differs across panels.

Although the performance of all the overidentification tests is quite poor when a is small, it is worth noting that the Sargan tests are not as unreliable as t tests of the hypothesis that β has a specific value, and the LR tests for overidentification are not as unreliable as LR tests for that hypothesis; see Davidson and MacKinnon [7,8].

6.1. Near the Singularity

From Figures 1–3, we see that the rejection probabilities of all the tests vary considerably with the parameters a and ρ as they vary in the neighborhood of the singularity at $a = 0, \rho = 1$. Further insight

into this phenomenon is provided by Figures 4 and 5. These are contour plots of rejection frequencies near the singularity for tests at the 0.05 level with a and ρ on the horizontal and vertical axes, respectively. Figure 4 is for the Basmann statistic S' , and Figure 5 is for the LR' statistic. Both figures are for the case dealt with in Figure 3, for which $n = 400$ and $q = 8$. The rejection frequencies are, once again, estimated using 10^6 replications.




Figure 4. Contours of rejection frequencies for S' tests with $q = 8$ and $n = 400$.




Figure 5. Contours of rejection frequencies for likelihood ratio (LR)' tests with $q = 8$ and $n = 400$.

It is clear from these figures that rejection frequencies tend to be greatest as the singularity is approached by first setting $r = 0$ and then letting a tend to zero. In this limit, S' is given by expression (28), and LR' is given by expression (23). For extremely small values of a , S' actually underrejects. But, as a rises to values that are still very very small, rejection frequencies soar, sometimes to over 0.80. In contrast, LR' underrejects severely for small values of a , values which do not have to be nearly as small as in the case of S' . In much of the figure, however, the rejection frequencies for LR' are just a little greater than 0.05.

The 95% quantile of the distribution of expression (28) has the huge value of 16,285, as estimated from 9,999,999 independent realizations. In contrast, recall that the 95% quantile of the χ_q^2 distribution for $q = 8$ is 15.5073. Since the distribution of S' for arbitrary a and ρ is stochastically bounded by that of (28), S' is boundedly pivotal. However, basing inference on the distribution of (28) is certain to be *extremely* conservative.

7. Performance of Bootstrap Tests

In principle, any of the bootstrap DGPs discussed in the previous section can be combined with any of the test statistics discussed in Section 2. However, there is no point considering both S and S' , or both LR and LR' , because in each case one test statistic is simply a monotonic transformation of the other. If both the statistics in each pair are bootstrapped using the same bootstrap DGP, they must therefore yield identical results.

All of our experiments involve 100,000 replications for each set of parameter values, and the bootstrap tests mostly use $B = 399$. This is a smaller number than should generally be used in practice, but it is perfectly satisfactory for simulation experiments, because experimental randomness in the bootstrap p values tends to average out across replications. Although the disturbances of the true DGPs are taken to be normally distributed, the bootstrap DGPs we investigate in the main experiments are resampling ones, because we believe they are the ones that will be used in practice.

Figures 6–8 present the results of a large number of Monte Carlo experiments. Figure 6 concerns Sargan tests, Figure 7 concerns LR tests, and Figure 8 concerns Fuller LR tests. Each of the figures shows rejection frequencies as a function of ρ for 34 values of ρ , namely, 0.00, 0.03, 0.06, . . . , 0.99. The four panels correspond to $a = 2, 4, 6,$ and 8 . Note that the scale of the vertical axis often differs across panels within each figure and across figures for panels corresponding to the same value of a . It is important to keep this in mind when interpreting the results.

As we have already seen, for small and moderate values of a , Sargan tests tend to overreject severely when ρ is large and to underreject modestly when it is small. It is evident from Figure 6 that, for $a = 2$, using either the IV-R or IV-ER bootstrap improves matters only slightly. However, both these methods do provide a more and more noticeable improvement as a increases. For $a = 8$, the improvement is very substantial. If we were increasing n as well as a , it would be natural to see this as evidence of an asymptotic refinement.

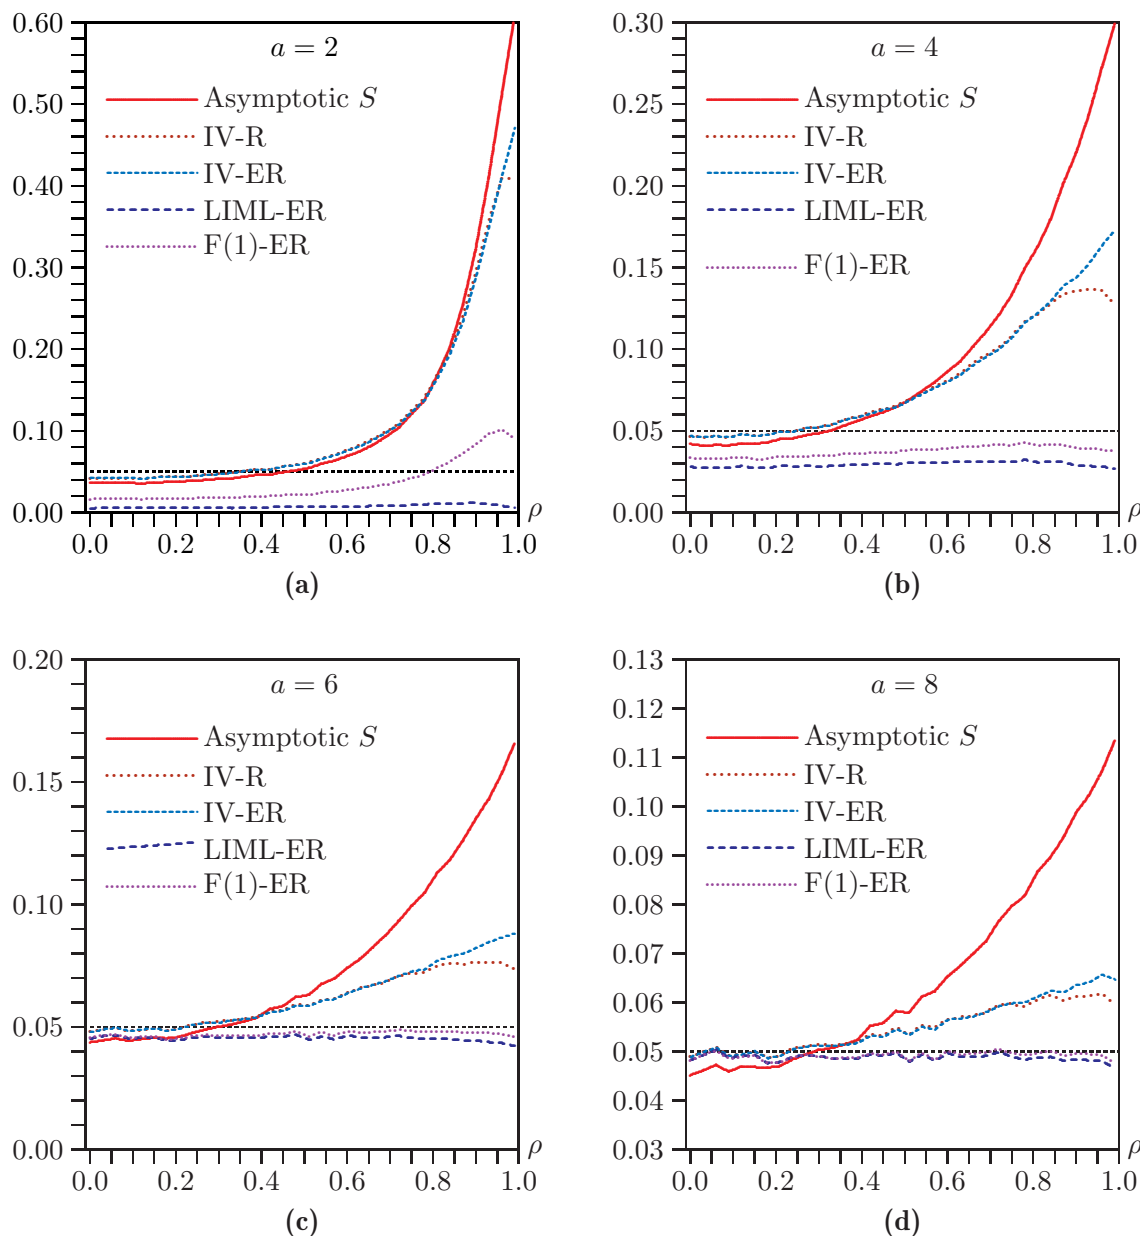


Figure 6. Rejection frequencies for Sargan tests as functions of ρ for $q = 8$ and $n = 400$. (a) $a = 2$; (b) $a = 4$; (c) $a = 8$; (d) $a = 16$. Note that the scale of the vertical axis differs across panels.

There seems to be no advantage to using IV-ER rather than IV-R. In fact, the latter always works a bit better when ρ is very large. This result is surprising in the light of the findings of Davidson and MacKinnon [7,8] for bootstrapping t tests on β . However, the bootstrap methods considered in those papers imposed the null hypothesis that $\beta = \beta_0$, while the ones considered here do not. Apparently, this makes a difference.

Using the LIML-ER and F(1)-ER bootstraps with the Sargan statistic yields entirely different results. The former underrejects very severely for all values of ρ when a is small, but the extent of the underrejection drops rapidly as a increases. The latter always underrejects less severely than LIML-ER (it actually overrejects for large values of ρ when $a = 2$), and it performs surprisingly well for $a \geq 6$.

Of course, it may seem a bit strange to bootstrap a test statistic based on IV estimation using a bootstrap DGP based on LIML or its Fuller variant.




Figure 7. Rejection frequencies for LR tests as functions of ρ for $q = 8$ and $n = 400$. (a) $a = 2$; (b) $a = 4$; (c) $a = 8$; (d) $a = 16$. Note that the scale of the vertical axis differs across panels.

In Figure 7, we see that, in contrast to the Sargan test, the LR test generally underrejects, often very severely when both ρ and a are small. Its performance improves rapidly as a increases, however, and it actually overrejects slightly when ρ and a are both large. All of the bootstrap methods improve matters, and the extent of the improvement increases with a . For $a = 8$, all the bootstrap methods work essentially perfectly. For small values of a , the IV-R bootstrap actually seems to be the best in many cases, although it does lead to modest overrejection when ρ is large.




Figure 8. Rejection frequencies for Fuller LR tests as functions of ρ for $q = 8$ and $n = 400$. (a) $a = 2$; (b) $a = 4$; (c) $a = 8$; (d) $a = 16$. Note that the scale of the vertical axis differs across panels.

In Figure 8, we see that the Fuller LR test never underrejects as much as the LR test, and it actually overrejects quite severely when ρ is large and $a = 2$. However, that is the only case in which it overrejects much. This is the only test for which its own bootstrap DGP, namely, F(1)-ER, is arguably the best one to use. Except when the asymptotic test already works perfectly, using that bootstrap method almost always improves the performance of the test. The bottom two panels of Figure 8 look very similar to the corresponding panels of Figure 7, except that the bootstrapped Fuller test tends to underreject just a bit. It is evident that, as a increases, the LR test and its Fuller variant become almost indistinguishable.

Figures 6–8 provide no clear ranking of tests and bootstrap methods. There seems to be a preference for the LR and Fuller LR tests, and for the LIML-ER and F(1)-ER bootstrap DGPs. In no case does any combination of those tests and those bootstrap DGPs overreject anything like as severely as the

Sargan test bootstrapped using IV-R or IV-ER. Provided the instruments are not very weak, any of these combinations should yield reasonably accurate, but perhaps somewhat conservative, inferences in most cases.

The rather mixed performance of the bootstrap tests can be understood by using the concept of “bootstrap discrepancy,” which is a function of the nominal level of the test, say α . The bootstrap discrepancy is simply the actual rejection rate for a bootstrap test at level α minus α itself. Davidson and MacKinnon [18] shows that the bootstrap discrepancy at level α is a conditional expectation of the random variable

$$q(\alpha) \equiv R(Q(\alpha, \mu^*), \mu) - \alpha, \tag{41}$$

where $R(\alpha, \mu)$ is the probability, under the DGP μ , that the test statistic is in the rejection region for nominal level α , and $Q(\alpha, \mu)$ is the inverse function that satisfies the equation

$$R(Q(\alpha, \mu), \mu) = \alpha = Q(R(\alpha, \mu), \mu).$$

Thus $Q(\alpha, \mu)$ is the true level- α critical value of the asymptotic test under μ . The random element in (41) is μ^* , the bootstrap DGP. If $\mu^* = \mu$, then we see clearly that $q(\alpha) = 0$, and the bootstrap discrepancy vanishes. For more detail, see Davidson and MacKinnon [18].




Figure 9. Contours of rejection frequencies for instrumental variables (IV)-R bootstrap Sargan tests.

Suppose now that the true DGP μ_0 is near the singularity. The bootstrap DGP can reasonably be expected also to be near the singularity, but most realizations are likely to be farther away from the singularity than μ_0 itself. If μ_0 were actually at the singularity, then any bootstrap DGP would necessarily be farther away. If the statistic used is S , then we see from Figure 4 that rejection frequencies fall as

the DGP moves away from the singularity in most, but not all, directions. Thus, for most such bootstrap DGPs, $Q(\alpha, \mu^*)$ is smaller than $Q(\alpha, \mu_0)$ for any α , and so the probability mass $R(Q(\alpha, \mu^*), \mu_0)$ in the distribution generated by μ_0 is greater than α . This means that $q(\alpha)$ is positive, and so the bootstrap test overrejects. However, if the statistic used is LR, the reverse is the case, as we see from Figure 5, and the bootstrap test underrejects. This is just what we see in Figures 6 to 8.

Figures 9 and 10 are contour plots similar to Figures 4 and 5, but they are for bootstrap rather than asymptotic tests. The IV-R parametric bootstrap is used for the Sargan test in Figure 9, and the LIML-ER parametric bootstrap is used for the LR test in Figure 10. In both cases, there are 100,000 replications, and $B = 199$. Figure 9 looks remarkably like Figure 4, with low rejection frequencies for extremely small values of a , then a ridge where rejection frequencies are very high for slightly larger values of a . The ridge is not quite as high as the one in Figure 4, and the rejection frequencies diminish more rapidly as a increases.




Figure 10. Contours of rejection frequencies for limited information maximum likelihood (LIML)-ER bootstrap LR tests.

Similarly, Figure 10 looks like Figure 5, but the severe underrejection in the far left of the figure occurs over an even smaller region, and there is an area of modest overrejection nearby. Both these size distortions can be explained by Figure 5. When a is extremely small, the estimate used by the bootstrap DGP tends on average to be larger, so the bootstrap critical values tend, on average, to be overestimates. This leads to underrejection. However, there is a region where a is not quite so small in which the bootstrap DGP uses estimates of a that are sometimes too small and sometimes too large. The former causes overrejection, the latter underrejection. Because of the curvature of the rejection probability function, the net effect is modest overrejection; see Davidson and MacKinnon [17]. This is actually the

case for most of the parameter values shown in the figure, but the rejection frequencies are generally not much greater than 0.05.

8. Power Considerations

Overidentification tests are performed in order to check whether some of the assumptions for the two-equation model (1) and (2) to be correctly specified are valid. Those assumptions are not valid if the DGP for Equation (1) is actually

$$y_1 = Z\gamma_1 + W_1\delta + \beta y_2 + u_1, \tag{42}$$

where the columns of the matrix W_1 are in the span of the columns of the matrix W and are linearly independent of those of Z . As in Section 3, we can eliminate Z from the model, replacing all other variables and the disturbances by their projections onto the orthogonal complement of the span of the columns of Z . The simpler model of Equations (13) and (16) becomes

$$y_1 = \beta y_2 + \delta w_p + v_1 \tag{43}$$

$$y_2 = aw_1 + u_2 = aw_1 + \rho v_1 + r v_2. \tag{44}$$

The vector $W\pi$ is now written as aw_1 instead of aw , and the vector $W_1\delta$ is written as δw_p . As before, we make the normalizations that $\|w_1\|^2 = 1$ and $a^2 = \pi^T W^T W \pi$. In addition, we normalize so that $\|w_p\|^2 = 1$ and $\delta^2 = \delta^T W_1^T W_1 \delta$.

The disturbance vector u_1 of (13) is written as v_1 . In the rightmost expression of Equations (44), the vector u_2 has been replaced by $\rho v_1 + r v_2$, where $r \equiv (1 - \rho^2)^{1/2}$. Since we are assuming, as in Section 3, that the disturbance vectors u_1 and u_2 are normally distributed, the vectors v_1 and v_2 are independent $N(0, I)$. Further, since w_1 and w_p are not in general orthogonal, we write

$$w_p = \theta w_1 + t w_2,$$

where $w_2^T w_1 = 0$, and $t = (1 - \theta^2)^{1/2}$.

The Basman statistic S' is still given by Equation (18), which is simply an algebraic consequence of the definition (10). Since the DGP for y_2 is unchanged, the quantities denoted in (10) by P_{22} and M_{22} are the same under the alternative as under the null. Since the DGP for $M_W y_1$ is also the same under the null and the alternative, so are M_{11} and M_{12} . Thus only P_{11} and P_{12} differ from the expressions for them in Equations (20). It is easy to check that neither the numerator nor the denominator of S' in (18) depends on β under the alternative, and so in our computations we set $\beta = 0$ without loss of generality.

In order to analyze the asymptotic power of the Sargan test in Basman form, we seek to express its limiting asymptotic distribution as a chi-squared variable that is non-central under the alternative. As usual, in order for the non-centrality parameter (NCP) to have a finite limit, we invoke a Pitman drift. With our normalization of w_p , this just means that δ is constant as the sample size n tends to infinity. Again, we cannot expect to find a limiting chi-squared distribution with weak-instrument asymptotics, and so our asymptotic construction supposes that $a \rightarrow \infty$ as $n \rightarrow \infty$.

Under the null and the alternative, the denominator of (18), divided by $(n - l)P_{22}$, is simply an estimate of the variance of v_1 . For the purposes of the asymptotic analysis of the simpler model, it can

therefore be replaced by 1. The quantity of which the limiting distribution is expected to be chi-squared is therefore $P_{11} - P_{12}^2/P_{22}$. Recall that this is just the numerator of both the S and S' statistics.

With $\beta = 0$, we compute as follows:

$$\begin{aligned} P_{11} &= \mathbf{y}_1^\top \mathbf{P}_W \mathbf{y}_1 = \delta^2 + 2\delta\theta x_1 + 2\delta t z_1 + \mathbf{v}_1^\top \mathbf{P}_W \mathbf{v}_1, \\ P_{12} &= \mathbf{y}_1^\top \mathbf{P}_W \mathbf{y}_2 = a(x_1 + \delta\theta) + O_p(1), \text{ and} \\ P_{22} &= \mathbf{y}_2^\top \mathbf{P}_W \mathbf{y}_2 = a^2 + 2a(\rho x_1 + r x_2) + O_p(1), \end{aligned}$$

where the symbol $O_p(1)$ means of order unity as $a \rightarrow \infty$. As before, we let $x_i = \mathbf{w}_1^\top \mathbf{v}_i$ for $i = 1, 2$, and we also let $z_i = \mathbf{w}_2^\top \mathbf{v}_i$. Thus the limit as $a \rightarrow \infty$ of $P_{11} - P_{12}^2/P_{22}$ is

$$\delta^2 + 2\delta\theta x_1 + 2\delta t z_1 + \mathbf{v}_1^\top \mathbf{P}_W \mathbf{v}_1 - (x_1 + \delta\theta)^2 = \mathbf{v}_1^\top \mathbf{P}_W \mathbf{v}_1 - x_1^2 + \delta^2 t^2 + 2\delta t z_1. \tag{45}$$

In Equation (17), we introduced the quantity Q_{11} , equal to $\mathbf{v}_1^\top \mathbf{P}_W \mathbf{v}_1$ and distributed as χ_l^2 . It was expressed as the sum of three mutually independent random variables, x_1^2 , z_P^2 , and t_{11}^P . Now we separate out both the terms x_1^2 and z_1^2 to obtain

$$Q_{11} = x_1^2 + z_1^2 + z_P^2 + t_{11}^{P_0}, \tag{46}$$

where all four random variables above are independent, with x_1 , z_1 , and z_P standard normal, and $t_{11}^{P_0}$ distributed as χ_{l-3}^2 . The random variable $t_{11}^{P_0}$ is not to be confused with t_{11}^P in Equations (17), which is distributed as χ_{l-2}^2 . It is legitimate to write Q_{11} in this way because it can be constructed as the sum of the squares of the l independent $N(0,1)$ variables $\mathbf{w}_j^\top \mathbf{v}_1$, where the \mathbf{w}_j form an arbitrary orthonormal basis of the span of the columns of \mathbf{W} .

Using (46), the right-hand side of Equation (45) can be written as

$$z_P^2 + t_{11}^{P_0} + (z_1 + \delta t)^2.$$

This is the sum of three independent random variables. The first is χ_1^2 , the second is χ_{l-3}^2 , and the last is noncentral $\chi_1^2(\delta^2 t^2)$. It follows that, when a^2 and the sample size both tend to infinity, which implies that the instruments are not weak, the numerator of the test statistic is distributed as $\chi_{l-1}^2(\delta^2 t^2)$. Note that, if $\theta = 1$, so that $\mathbf{w}_p = \mathbf{w}_1$, the NCP $\delta^2 t^2$ vanishes.

For the general model (1) and (2), with DGP given by Equation (42), it can be shown that the NCP is

$$\frac{1}{\sigma_1^2} \delta^\top \mathbf{W}_1^\top \mathbf{M}_Z \mathbf{W}_1 \delta - \frac{1}{\sigma_1^2} \frac{(\boldsymbol{\pi}^\top \mathbf{W}^\top \mathbf{M}_Z \mathbf{W}_1 \delta)^2}{\boldsymbol{\pi}^\top \mathbf{W}^\top \mathbf{M}_Z \mathbf{W} \boldsymbol{\pi}}. \tag{47}$$

For the simpler model given by Equations (43) and (44), the first term here collapses to δ^2 and the second term, which arises because β has to be estimated, collapses to $-\theta^2 \delta^2$. Therefore, expression (47) as a whole corresponds to $\delta^2 t^2$ for the simpler model.

8.1. Finite-Sample Concerns

The asymptotic result that S' follows the $\chi_{l-1}^2(\delta^2 t^2)$ distribution strongly suggests that S , LR, and LR' must do so as well, because all these statistics are asymptotically equivalent. In fact, a more

tedious calculation than that in Equations (45) and (46) shows that the limiting distribution of LR' as both n and a tend to infinity is the same as for S' , namely $\chi_{l-1}^2(\delta^2 t^2)$. Because these results are only asymptotic, however, it is necessary to resort to simulation to investigate behavior under the alternative in finite samples.

Under the null, we were able in Section 3 to express all six quantities, the P_{ij} and the M_{ij} , for $i, j = 1, 2$, in terms of eight independent random variables. Under the alternative, we require ten of these variables. For the M_{ij} , there is no need to change the expressions for them in (20), where we use the three variables t_{11}^M, t_{22}^M , and z_M , distributed respectively as $\chi_{n-l}^2, \chi_{n-l-1}^2$, and $N(0,1)$. These represent the projections of v_1 and v_2 onto the orthogonal complement of the span of the instruments. For the P_{ij} , however, we decompose as follows:

$$\begin{aligned}
 Q_{11} &= x_1^2 + z_1^2 + z_P^2 + t_{11}^{P_0}, \\
 Q_{12} &= x_1 x_2 + z_1 z_2 + z_P \sqrt{t_{22}^{P_0}}, \text{ and} \\
 Q_{22} &= x_2^2 + z_2^2 + t_{22}^{P_0}.
 \end{aligned}$$

Here $x_i, z_i, i = 1, 2$, and z_P are standard normal, t_{11}^P is χ_{l-3}^2 , and t_{22}^P is χ_{l-2}^2 , all seven variables being mutually independent. We can simulate both S' and LR' very cheaply, by drawing ten random variables, independently of either the sample size n or the degree of overidentification $l-1$, because all the statistics are deterministic functions of the P_{ij} and the M_{ij} , and, of course, n and l . The relations in (20) hold except those for P_{11} and P_{12} . These are replaced by

$$P_{11} = Q_{11} + \delta^2 + 2\delta\theta x_1 + 2\delta t z_1,$$

and

$$P_{12} = ax_1 + \rho Q_{11} + rQ_{12} + \delta(a\theta + \rho\theta x_1 + \rho t z_1 + r\theta x_2 + r t z_2).$$

These equations differ from the corresponding ones in (20) only by terms proportional to a positive power of δ .

8.2. Simulation Evidence

Since we have seen that the LR' test often has much better finite-sample properties than the S' test, even when both are bootstrapped, it is important to see whether the superior performance of LR' comes at the expense of power. In this section, we employ simulation methods to do so.

Given the considerable size distortion of the asymptotic tests for most of that part of the parameter space considered in Section 7, we limit attention to parametric bootstrap tests. In this, we follow Horowitz and Savin [19], which argues that the best way to proceed, as long as the rejection probability of a test is far removed from its nominal level, is to consider a bootstrap test. But that proposition is based on the assumption that the bootstrap discrepancy is small enough to be ignored, which is not the case for the overidentification tests we have considered in the neighborhood of the singularity. Because of that, and because it is unreasonable to expect that there is much in the way of usable power near the singularity, it is primarily of interest to investigate power for situations in which the instruments are not too weak.

As before, all the simulation results are presented graphically. These results are based on 200,000 replications with 399 bootstrap repetitions. The same random variables are used for every set of parameter values. These experiments would have been extremely computationally demanding without the theoretical results of Section 6 and the first part of this section, which allow us to calculate everything very cheaply after we have generated and stored $200,000 \times 10$ plus $200,000 \times 399 \times 8$ random variables. The first set of random variables is used to calculate the actual test statistics and the estimates of a and ρ , and the second set is used to calculate the bootstrap statistics.




Figure 11. Power of bootstrap tests as functions of δ with $n = 400$, $q = 8$, $\rho = 0.5$, $\theta = 0$, and $t = 1$. (a) $a = 2$; (b) $a = 4$; (c) $a = 8$; (d) $a = 6$.

We report results only for S' bootstrapped using the IV parameter estimates and for LR' bootstrapped using the LIML estimates. Recall from Section 2 that the former results apply to S as well as S' , and the latter apply to LR as well as LR' , because the test statistics in each pair are monotonically related.

Figure 11 shows power functions for $q = 8$, $\rho = 0.5$, and four values of a . When $a = 2$, LR' rejects much less frequently than S' , both under the null and under the alternative. Both power functions level out as δ becomes large, and it appears that neither test rejects with probability one as $\delta \rightarrow \infty$. As a increases, the two power functions converge, and both tests do seem to reject with probability one for large δ .

The top two panels of Figure 12 are comparable to the top two panels of Figure 11, but with $q = 2$. When $a = 2$, S' now rejects less often that it did before, but LR' rejects more often. When $a = 4$, LR' rejects very much more often than it did before, and the two power functions are quite close. We also obtained results for $a = 6$, $a = 8$, and $a = 16$, which are not shown. For $a = 6$, the power functions for S' and LR' are extremely similar, and for $a \geq 8$ they are visually indistinguishable.




Figure 12. Power of bootstrap tests as functions of δ with $n = 400$, $q = 2$, $\theta = 0$, and $t = 1$.
(a) $a = 2$, $\rho = 0.5$; **(b)** $a = 4$, $\rho = 0.5$; **(c)** $a = 4$, $\rho = 0.1$; **(d)** $a = 4$, $\rho = 0.9$.

The bottom two panels of Figure 12 are comparable to the top right panel, except that $\rho = 0.1$ or $\rho = 0.9$ instead of $\rho = 0.5$. It is evident that the shapes of the power functions depend on ρ , but for most

values of δ the dependence is moderate. This justifies our use of $\rho = 0.5$ in most of the experiments. Using other values of ρ would not change the main results.

When one power function is always above another, as is the case in all the panels of Figures 11 and 12, it is difficult to conclude that one test is genuinely more powerful than the other. Perhaps greater power is just an artifact of greater rejection frequencies whether or not the null hypothesis is true.

One way to compare such tests is to graph rejection frequencies under the alternative against rejection frequencies under the null. Each point on such a “size-power curve” corresponds to some nominal level for the bootstrap test, with levels running from 0 to 1. The abscissa is the rejection frequency when the DGP satisfies the null, the ordinate the rejection frequency when the DGP belongs to the alternative. For a level of 0, the test never rejects, since bootstrap P values cannot be negative. If the level is 1, the test always rejects. As the nominal level increases from 0 to 1, we expect power (on the vertical axis) to increase more rapidly than the rejection frequency under the null (on the horizontal axis). See Davidson and MacKinnon [20].

The top two panels of Figure 13 show size-power curves for $q = 2$, $a = 4$, and four values of δ . Perhaps surprisingly, the curves for LR' in the left-hand panel look remarkably similar to the ones for S' in the right-hand panel. The apparently greater power of S' , which is evident in the top right panel of Figure 12, seems to be almost entirely accounted for by its greater tendency to reject under the null.

The bottom two panels of Figure 13 show size-power curves for $q = 2$, $\delta = 4$, and four values of a . It is clear that power increases with a , but at a decreasing rate. As $a \rightarrow \infty$, the curves converge to the one given by asymptotic theory, where the distribution under the null is central χ^2_{l-1} and the one under the alternative is noncentral $\chi^2_{l-1}(\delta^2 t^2)$. This curve is graphed in the figure and labelled $a = \infty$.

The asymptotic result that the test statistics follow the $\chi^2_{l-1}(\delta^2 t^2)$ distribution suggests that only the product $\delta t = \delta(1 - \theta^2)^{1/2}$ influences power, and that, in particular, there should be no power beyond the level of the test when $\theta = 1$. In finite samples, things turn out to be more complicated, as can be seen from Figure 14, which plots power against θ for $\delta = 4$. The top two panels show results for $a = 2$ and $a = 4$. The S' test has substantial power when $\theta = 1$ and $a = 2$, which presumably reflects its tendency to overreject severely under the null when the instruments are weak. Those panels also show, once again, that S' can reject far more often than LR' when the instruments are weak. This is much less evident in the bottom two panels, which show results for larger values of a , namely, 6 and 8.

One surprising feature of Figure 14 is that, in all cases, power initially increases as θ increases from 0, even though $\delta(1 - \theta^2)^{1/2}$ declines. This is true even for quite large values of a , such as $a = 16$, although, of course, it is not true for extremely large values.

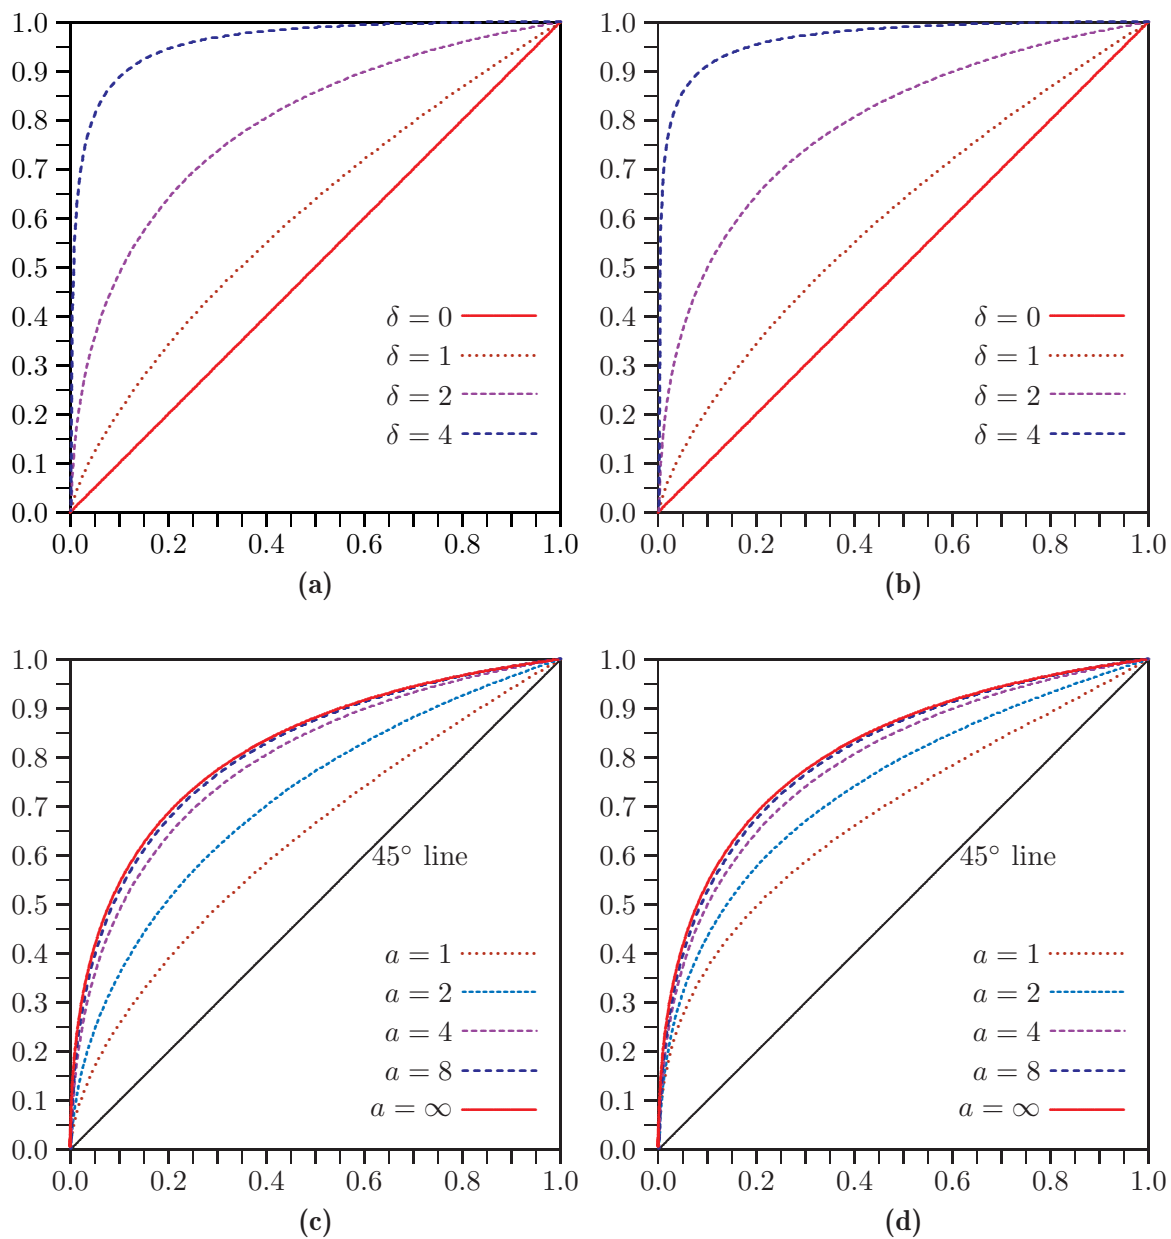


Figure 13. Size-power curves for $q = 2$, $n = 400$, $\rho = 0.5$, and $\theta = 0$. (a) LR': $a = 4$ and several values of δ ; (b) S' : $a = 4$ and several values of δ ; (c) LR': $\delta = 2$ and several values of a ; (d) S' : $\delta = 2$ and several values of a .




Figure 14. Power as a function of θ for $n = 400$, $\rho = 0.5$, and $\delta = 4$. (a) $q = 2$; (b) $q = 8$; (c) $q = 2$; (d) $q = 8$

9. Relaxing the IID Assumption

The resampling bootstraps that we looked at in Section 7 do not implicitly make the assumption that the disturbances are normal. They do, however, assume that the disturbances are pairwise IID. If instead the disturbances are heteroskedastic, then the covariance matrix of their bivariate distribution may be different for each observation. In that case, all the test statistics we have studied have distributions that depend on the pattern of heteroskedasticity, and so they are no longer approximately pivotal for the model (1) and (2) under either weak-instrument or strong-instrument asymptotics.

Andrews, Moreira, and Stock [21] proposes heteroskedasticity-robust versions of test statistics for tests about the value of β that are robust to weak instruments. Note that, although Andrews, Moreira, and Stock [6] is based on the 2004 paper and has almost the same title, it does not contain this material.

However, this work cannot be applied here, because, as we have seen, the overidentification tests are not robust to weak instruments.

The role of the denominators of the statistics S , S' , and LR' is simply to provide non-robust estimates of the scale of the numerators. In order to make those statistics robust to heteroskedasticity, we have to provide robust measures instead. The numerators of all three statistics can be written as

$$\hat{u}_1^\top P_W \hat{u}_1 = \hat{u}_1^\top W (W^\top W)^{-1} W^\top \hat{u}_1, \tag{48}$$

where the vector \hat{u}_1 denotes either $y_1 - Z\hat{\gamma}_{IV} - \hat{\beta}_{IV}y_2$, in the case of S and S' , or $y_1 - Z\hat{\gamma}_{LIML} - \hat{\beta}_{LIML}y_2$, in the case of LR' . Expression (48) is a quadratic form in the l -vector $W^\top \hat{u}_1$. The usual estimate of the covariance matrix of that vector is $W^\top \hat{\Omega} W$, where $\hat{\Omega} = \text{diag } \hat{u}_{1i}^2$. Thus the heteroskedasticity-robust variant of all three test statistics is the quadratic form

$$\hat{u}_1^\top W (W^\top \hat{\Omega} W)^{-1} W^\top \hat{u}_1. \tag{49}$$

There would be no point in using a heteroskedasticity-robust statistic along with a bootstrap DGP that imposed homoskedasticity. The natural way to avoid doing so is to use the wild bootstrap. In Davidson and MacKinnon [8], the wild bootstrap is shown to have good properties when used with tests about the value of β . The disturbances of the wild bootstrap DGP are given by

$$\begin{bmatrix} u_{1i}^* \\ u_{2i}^* \end{bmatrix} = \begin{bmatrix} \hat{u}_{1i} \nu_i^* \\ \hat{u}_{2i} \nu_i^* \end{bmatrix}, \tag{50}$$

where ν_i^* is an auxiliary random variable with expectation 0 and variance 1. The easiest choice for the distribution of the ν_i^* is the Rademacher distribution, which sets ν_i^* to +1 or -1, each with probability one half. This is also probably the best choice in most cases; see Davidson and Flachaire [22].

The IID assumption can, of course, be relaxed in other ways. In particular, it would be easy to modify the test statistic (49) to allow for clustered data by replacing the middle matrix with one that resembles the middle matrix for the usual cluster robust covariance matrix. We could then use a variant of the wild cluster bootstrap of Cameron, Gelbach, and Miller [23] that allows for simultaneity. The Rademacher random variable associated with each cluster, the analog of ν_i^* in Equation (50), would then multiply the residuals for all observations within that cluster for both equations.

10. Concluding Remarks

We have shown that the well-known Sargan test for overidentification in a linear simultaneous-equations model estimated by instrumental variables often overrejects severely when the instruments are weak. In the same circumstances, the likelihood ratio test often underrejects severely. In Sections 3 and 4, we provided a finite-sample analysis that explains these facts. The distributions of the test statistics we consider have a singularity when the concentration parameter vanishes and the absolute value of the correlation between the disturbances of the structural and reduced-form equations tends to one. Thus it can be risky to use asymptotic tests in this situation. In addition, we proposed a new test based on Fuller’s modified LIML estimator, which often outperforms the ordinary LR test.

We have also proposed four bootstrap methods which can be applied to all three of these tests. Although bootstrapping does not help much when the instruments are extremely weak, especially

when the disturbances of the two equations are highly correlated, it does help substantially when the instruments are only moderately weak. In particular, using a bootstrap DGP based on Fuller's estimator generally leads to much more accurate inferences than simply using asymptotic theory in this case.

There is a cost in terms of power to using a bootstrap test based on any version of the likelihood ratio statistic relative to a test based on the conventional Sargan or Basman statistics. This cost generally seems to be very modest, except when the instruments are very weak.

Acknowledgments

This research was supported, in part, by grants from the Social Sciences and Humanities Research Council of Canada, the Canada Research Chairs program (Chair in Economics, McGill University), and the Fonds de recherche du Québec—Société et culture (FRQSC). For comments that have led us to improve the paper, we thank some anonymous referees. We are also grateful for helpful remarks made by seminar participants at the University of Exeter and the Tinbergen Institute, Amsterdam, as well as at meetings of the Canadian Econometric Study Group.

Author Contributions

The two authors contributed equally to the paper as a whole.

Conflicts of Interest

The authors declare no conflict of interest.

Appendix: Proofs

We have remarked several times that all of the overidentification test statistics are invariant to the parameter β . Since the proof of Theorem 1 depends on this result for two of those statistics, it is worth stating the following lemma.

Lemma 1. Neither the Basman statistic S' defined in (10), nor the linearized likelihood ratio statistic LR' defined in (12), depends on the value of the parameter β .

Proof of Lemma 1. Observe that, from (13),

$$P_{11} = \mathbf{y}_1^\top \mathbf{P}_W \mathbf{y}_1 = \mathbf{u}_1^\top \mathbf{P}_W \mathbf{u}_1 + 2\beta \mathbf{y}_2^\top \mathbf{P}_W \mathbf{u}_1 + \beta^2 P_{22},$$

where, since \mathbf{y}_2 does not depend on β , all dependence on β in this equation is explicit. By a similar argument,

$$P_{12} = \mathbf{y}_2^\top \mathbf{P}_W \mathbf{u}_1 + \beta P_{22},$$

and it then follows that

$$P_{11}P_{22} - P_{12}^2 = P_{22}\mathbf{u}_1^\top \mathbf{P}_W \mathbf{u}_1 - (\mathbf{y}_2^\top \mathbf{P}_W \mathbf{u}_1)^2,$$

which does not depend on β . Similar calculations show that the denominator in (18) and the coefficients A and B in Equation (54) do not depend on β , and so neither do the statistics S' and LR' . \square

Proof of Theorem 1. Let $\hat{\mathbf{u}}_1^{IV}$ denote $\mathbf{y}_1 - \hat{\beta}_{IV}\mathbf{y}_2$. The Basman statistic (10) can be expressed as

$$S' = (n - l) \frac{(\hat{\mathbf{u}}_1^{IV})^\top \mathbf{P}_W \hat{\mathbf{u}}_1^{IV}}{(\hat{\mathbf{u}}_1^{IV})^\top \mathbf{M}_W \hat{\mathbf{u}}_1^{IV}}.$$

The IV estimator $\hat{\beta}_{IV}$ satisfies the estimating equation $\mathbf{y}_2^\top \mathbf{P}_W (\mathbf{y}_1 - \hat{\beta}_{IV}\mathbf{y}_2) = 0$, so that $\hat{\beta}_{IV} = P_{12}/P_{22}$. Therefore

$$\begin{aligned} (\hat{\mathbf{u}}_1^{IV})^\top \mathbf{P}_W \hat{\mathbf{u}}_1^{IV} &= P_{11} - P_{12}^2/P_{22}, \text{ and} \\ (\hat{\mathbf{u}}_1^{IV})^\top \mathbf{M}_W \hat{\mathbf{u}}_1^{IV} &= M_{11} - 2P_{12}M_{12}/P_{22} + P_{12}^2M_{22}/P_{22}^2. \end{aligned}$$

Thus we find that

$$S' = \frac{(n - l)(P_{11}P_{22} - P_{12}^2)}{M_{11}P_{22} - 2P_{12}M_{12} + P_{12}^2M_{22}/P_{22}},$$

which is Equation (18).

Now consider the statistic $LR' \equiv (n - l)\hat{\lambda} = (n - l)\lambda(\hat{\beta})$, where $\lambda(\beta) = \kappa(\beta) - 1$, with $\kappa(\beta)$ defined in Equation (11). For our simplified model,

$$\lambda(\beta) = \frac{(\mathbf{y}_1 - \beta\mathbf{y}_2)^\top \mathbf{P}_W (\mathbf{y}_1 - \beta\mathbf{y}_2)}{(\mathbf{y}_1 - \beta\mathbf{y}_2)^\top \mathbf{M}_W (\mathbf{y}_1 - \beta\mathbf{y}_2)} = \frac{P_{11} - 2\beta P_{12} + \beta^2 P_{22}}{M_{11} - 2\beta M_{12} + \beta^2 M_{22}}. \tag{51}$$

The first-order condition for the minimization of $\lambda(\beta)$ leads to the equation

$$(\beta P_{22} - P_{12})(M_{11} - 2\beta M_{12} + \beta^2 M_{22}) = (\beta M_{22} - M_{12})(P_{11} - 2\beta P_{12} + \beta^2 P_{22}).$$

It follows that

$$\hat{\beta}P_{22} - P_{12} = \hat{\lambda}(\hat{\beta}M_{22} - M_{12}).$$

whence

$$\hat{\beta} = \frac{P_{12} - \hat{\lambda}M_{12}}{P_{22} - \hat{\lambda}M_{22}} \quad \text{and} \quad \hat{\lambda} = \frac{P_{12} - \hat{\beta}P_{22}}{M_{12} - \hat{\beta}M_{22}}. \tag{52}$$

The equation above for $\hat{\lambda}$, combined with the definition (51), then shows that

$$\hat{\lambda} = \frac{P_{11} - \hat{\beta}P_{12}}{M_{11} - \hat{\beta}M_{22}} \quad \text{and} \quad \hat{\beta} = \frac{P_{11} - \hat{\lambda}M_{11}}{P_{12} - \hat{\lambda}M_{12}}. \tag{53}$$

By equating the two expressions for $\hat{\beta}$ in (52) and (53), we derive a quadratic equation satisfied by $\hat{\lambda}$, of the form $A\hat{\lambda}^2 - B\hat{\lambda} + C = 0$, where

$$A = M_{11}M_{22} - M_{12}^2, \quad B = P_{11}M_{22} - 2P_{12}M_{12} + P_{22}M_{11}, \quad \text{and} \quad C = P_{11}P_{22} - P_{12}^2 \tag{54}$$

Since we seek to minimize $\hat{\lambda}$, it must be the smaller root of this quadratic equation.

Lemma 1 shows that, although the quadratic forms (15) depend on the value of β , the statistics themselves do not. We are thus at liberty to set $\beta = 0$ in the subsequent analysis, without loss of generality. When $\beta = 0$, we find that

$$\begin{aligned} P_{11} &= Q_{11}, \quad P_{12} = ax_1 + \rho Q_{11} + rQ_{12}, \quad M_{11} = N_{11}, \\ P_{22} &= a^2 + 2a(\rho x_1 + rx_2) + \rho^2 Q_{11} + 2r\rho Q_{12} + r^2 Q_{22}, \\ M_{22} &= \rho^2 N_{11} + 2r\rho N_{12} + r^2 N_{22}, \quad \text{and} \quad M_{12} = \rho N_{11} + rN_{12}, \end{aligned}$$

where $r = \sqrt{1 - \rho^2}$. These are the relations (20).

From the standard formula for the roots of the quadratic equation for $\hat{\lambda}$, we see that

$$LR' = \frac{(n - l)(P_{11}M_{22} - 2P_{12}M_{12} + P_{22}M_{11} - \Delta^{1/2})}{2(M_{11}M_{22} - M_{12}^2)},$$

where the discriminant Δ is given by

$$\Delta = (P_{11}M_{22} - 2P_{12}M_{12} + P_{22}M_{11})^2 - 4(M_{11}M_{22} - M_{12}^2)(P_{11}P_{22} - P_{12}^2).$$

This is Equation (19) □

Proof of Theorem 2. Algebra, tedious but easily handled by computer, shows that the coefficients A , B , and C , defined in (54), are given by

$$\begin{aligned} A &= M_{11}M_{22} - M_{12}^2 = r^2(N_{11}N_{22} - N_{12}^2), \\ C &= P_{11}P_{22} - P_{12}^2 = a^2(Q_{11} - x_1^2) + 2ar(Q_{11}x_2 - Q_{12}x_1) + r^2(Q_{11}Q_{22} - Q_{12}^2), \\ B &= P_{11}M_{22} - 2P_{12}M_{12} + P_{22}M_{11} \\ &= a^2N_{11} + 2ar(N_{11}x_2 - N_{12}x_1) + r^2(Q_{11}N_{22} - 2Q_{12}N_{12} + Q_{22}N_{11}). \end{aligned}$$

It can be seen that these three coefficients are homogeneous of degree two in (a, r) , and it then follows from (19) that LR' is homogeneous of degree zero in these two parameters. Thus if we set $a = 0$ in A , B , and C , these coefficients are all proportional to r^2 , and so we can cancel r^2 from the numerator and denominator of (19). The value of LR' for $a = 0$ is then seen to be given by (21) and (22).

Note that expression (21) no longer depends on r at all. Thus the distribution of LR' in the limit of completely irrelevant instruments is independent of all the model parameters.

If we set $r = 0$, then both the numerator and denominator of (19) are zero. We therefore must divide both by r^2 before setting $r = 0$. This is simple for the denominator, which is explicitly proportional to r^2 . For the numerator, we proceed as follows. Let $B = B_0 + rB_1 + r^2B_2$, with B_0 , B_1 , and B_2 independent of r . Similarly, let $A = r^2A_2$, and let $C = C_0 + rC_1 + r^2C_2$. Then the numerator of (19) can be written as

$$B_0 + rB_1 + r^2B_2 - ((B_0 + rB_1 + r^2B_2)^2 - 4r^2A_2(C_0 + rC_1 + r^2C_2))^{1/2}. \tag{55}$$

The square root in expression (55) is

$$(B_0 + rB_1 + r^2B_2) \left[1 - \frac{4r^2A_2(C_0 + rC_1 + r^2C_2)}{(B_0 + rB_1 + r^2B_2)^2} \right]^{1/2}.$$

A Taylor expansion of this expression for small r shows that the numerator of LR' , expression (55), is

$$\begin{aligned} \frac{2r^2A_2(C_0 + rC_1 + r^2C_2)}{B_0 + rB_1 + r^2B_2} + O(r^3) &= r^2(2A_2C_0/B_0 + O(r)) \\ &= 2r^2[A_2(Q_{11} - x_1^2)/N_{11} + O(r)]. \end{aligned}$$

Thus the limit of LR' when $r \rightarrow 0$ is just

$$(n - l)(Q_{11} - x_1^2)/N_{11}.$$

This is expression (23). □

References

1. Staiger, D.O.; Stock, J.H. Instrumental variables regression with weak instruments. *Econometrica* **1997**, *65*, 557–586.
2. Stock, J.H.; Wright, J.H.; Yogo, M. A survey of weak instruments and weak identification in generalized method of moments. *J. Bus. Econ. Stat.* **2002**, *20*, 518–529.
3. Kleibergen, F. Pivotal statistics for testing structural parameters in instrumental variables regression. *Econometrica* **2002**, *70*, 1781–1803.
4. Moreira, M.J. A conditional likelihood ratio test for structural models. *Econometrica* **2003**, *71*, 1027–1048.
5. Moreira, M.J. Tests with correct size when instruments can be arbitrarily weak. *J. Econom.* **2009**, *152*, 131–140.
6. Andrews, D.W.K.; Moreira, M.J.; Stock, J.H. Optimal two-sided invariant similar tests for instrumental variables regression. *Econometrica* **2006**, *74*, 715–752.
7. Davidson, R.; MacKinnon, J.G. Bootstrap inference in a linear equation estimated by instrumental variables. *Econom. J.* **2008**, *11*, 443–477.
8. Davidson, R.; MacKinnon, J.G. Wild bootstrap tests for IV regression. *J. Bus. Econ. Stat.* **2010**, *28*, 128–144.
9. Guggenberger, P.; Kleibergen, F.; Mavroeidis, S.; Chen, L. On the asymptotic sizes of subset Anderson-Rubin and Lagrange Multiplier tests in linear instrumental variables regression. *Econometrica* **2012**, *80*, 2649–2666.
10. Sargan, J.D. The estimation of economic relationships using instrumental variables. *Econometrica* **1958**, *26*, 393–415.
11. Anderson, T.W.; Rubin, H. Estimation of the parameters of a single equation in a complete set of stochastic equations. *Ann. Math. Stat.* **1949**, *20*, 46–63.
12. Basmann, R.L. On the finite sample distributions of generalized classical identifiability test statistics. *J. Am. Stat. Assoc.* **1960**, *55*, 650–659.
13. Davidson, R.; MacKinnon, J.G. Bootstrap methods in econometrics. In *Palgrave Handbook of Econometrics*; Mills, T.C., Patterson, K., Eds.; Palgrave-Macmillan: Basingstoke, UK, 2006; Volume 1, pp. 812–838.
14. Mariano, R.S.; Sawa, T. The exact finite-sample distribution of the limited-information maximum likelihood estimator in the case of two included endogenous variables. *J. Am. Stat. Assoc.* **1972**, *67*, 159–163.
15. Phillips, P.C.B. Exact small sample theory in the simultaneous equations model. In *Handbook of Econometrics*; Griliches, Z., Intriligator, M.D., Eds.; North Holland: Amsterdam, The Netherlands, 1983; Volume 1, pp. 449–516.
16. Fuller, W.A. Some properties of a modification of the limited information estimator. *Econometrica* **1977**, *45*, 939–953.
17. Davidson, R.; MacKinnon, J.G. The size distortion of bootstrap tests. *Econom. Theory* **1999**, *15*, 361–376.

18. Davidson, R.; MacKinnon, J.G. The power of bootstrap and asymptotic tests. *J. Econom.* **2006**, *133*, 421–441.
19. Horowitz, J.L.; Savin, N.E. Empirically relevant critical values for hypothesis tests. *J. Econom.* **1977**, *95*, 375–389.
20. Davidson, R.; MacKinnon, J.G. Graphical methods for investigating the size and power of hypothesis tests. *Manch. Sch.* **1998**, *66*, 1–26.
21. Andrews, D.W.K.; Moreira, M.J.; Stock, J.H. *Optimal Invariant Similar Tests for Instrumental Variables Regression*; Technical Working Paper 299; NBER: Cambridge, MA, USA, 2004.
22. Davidson, R.; Flachaire, E. The wild bootstrap, tamed at last. *J. Econom.* **2008**, *146*, 162–169.
23. Cameron, A.C.; Gelbach, J.B.; Miller, D.L. Bootstrap-based improvements for inference with clustered errors. *Rev. Econ. Stat.* **2008**, *80*, 414–427.

© 2015 by the authors; licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution license (<http://creativecommons.org/licenses/by/4.0/>).