Econometrica, Vol. 74, No. 2 (March, 2006), 521–538

TESTING A PARAMETRIC MODEL AGAINST A NONPARAMETRIC ALTERNATIVE WITH IDENTIFICATION THROUGH INSTRUMENTAL VARIABLES

BY JOEL L. HOROWITZ[1]

This paper is concerned with inference about a function g that is identified by a conditional moment restriction involving instrumental variables. The paper presents a test of the hypothesis that g belongs to a finite-dimensional parametric family against a nonparametric alternative. The test does not require nonparametric estimation of g and is not subject to the ill-posed inverse problem of nonparametric instrumental variables estimation. Under mild conditions, the test is consistent against any alternative model. In large samples, its power is arbitrarily close to 1 uniformly over a class of alternatives whose distance from the null hypothesis is $O(n^{-1/2})$, where n is the sample size. In Monte Carlo simulations, the finite-sample power of the new test exceeds that of existing tests.

KEYWORDS: Hypothesis test, instrumental variables, specification testing, consistent testing.

1. INTRODUCTION

LET Y BE A SCALAR RANDOM VARIABLE, let X and W be continuously distributed random scalars or vectors, and let g be a function that is identified by the relationship

$$\mathbf{E}[Y - g(X)|W] = 0. \tag{1.1}$$

In (1.1), Y is the dependent variable, X is a possibly endogenous explanatory variable, and W is an instrument for X. This paper presents a test of the null hypothesis that g in (1.1) belongs to a finite-dimensional parametric family against a nonparametric alternative hypothesis. Specifically, let Θ be a compact subset of $\mathbb{R}^d$ for some finite integer $d > 0$. The null hypothesis, H_0, is that

$$g(x) = G(x, \theta) \tag{1.2}$$

for some $\theta \in \Theta$ and almost every x, where G is a known function. The alternative hypothesis, H_1, is that there is no $\theta \in \Theta$ such that (1.2) holds for almost every x. Under mild conditions, the test presented here is consistent against any alternative model. In large samples, its power is arbitrarily close to 1 uniformly over a class of alternative models whose "distance" from H_0 is $O(n^{-1/2})$, where n is the sample size. In Monte Carlo simulations, the finite-sample power of the new test exceeds that of existing tests.

[1]Part of this research was carried out while I was a visitor at the Centre for Microdata Methods and Practice, University College London. I thank Donald Andrews and Richard Blundell for helpful discussions and comments. Research supported in part by NSF Grants SES 991092 and SES 0352675.

There has been much recent interest in nonparametric estimation of g in (1.1). See, for example, Newey, Powell, and Vella (1999), Newey and Powell (2003), Darolles, Florens, and Renault (2002), Blundell, Chen, and Kristensen (2003), and Hall and Horowitz (2005). Methods for testing (1.2) against a nonparametric alternative have been developed by Donald, Imbens, and Newey (2003) and Tripathi and Kitamura (2003). In addition, the test of a conditional mean function developed by Bierens (1990) and Bierens and Ploberger (1997) can be modified to provide a test of (1.2). Horowitz and Spokoiny (2001, 2002) provide extensive references to other tests for conditional mean and quantile functions. The test presented here has power advantages over existing tests that permit X to have endogenous components. In addition, among existing tests of (1.2) against a nonparametric alternative, only the test presented here is uniformly consistent at a known rate over a known set of alternative hypotheses. Uniform consistency is important because it provides some assurance that there are not alternatives against which a test has low power even with large samples. If a test is not uniformly consistent over a specified set, then that set contains alternatives against which the test has low power. Implementation of the test described here is not difficult.

The test developed here is not affected by the ill-posed inverse problem of nonparametric instrumental variables estimation. Consequently, the test's "precision" exceeds that of any nonparametric estimator of g. The rate of convergence of a nonparametric estimator of g is always slower than $O_p(n^{-1/2})$ and, depending on the details of the distribution of (Y, X, W), may be slower than $O_p(n^{-\varepsilon})$ for any $\varepsilon > 0$ (Hall and Horowitz (2005)). In contrast, the test described here can detect a large class of nonparametric alternative models whose distance from the null-hypothesis model is $O(n^{-1/2})$. Nonparametric estimation and testing of conditional mean and median functions is another setting in which the rate of testing is faster than the rate of estimation. See Guerre and Lavergne (2002) and Horowitz and Spokoiny (2001, 2002).

Section 2 describes the test statistic and its properties. Section 3 presents the results of a Monte Carlo investigation of the finite-sample performance of the test, and Section 4 presents an illustrative application to real data. The proofs of theorems are in Horowitz (2005), which is available electronically at *Econometrica*'s web site.

2. THE TEST STATISTIC AND ITS PROPERTIES

Rewrite (1.1) as

$$Y = g(X, Z) + U, \quad \mathbf{E}(U|Z, W) = 0, \tag{2.1}$$

where Y and U are scalar random variables, X and W are random variables whose supports are contained in $[0, 1]^p$ $(p \geq 1)$, and Z is a random variable whose support is contained in $[0, 1]^r$ $(r \geq 0)$. If $r = 0$, then Z is not included in (2.1). Variables X and Z, respectively, are endogenous and exogenous explanatory variables. Variable W is an instrument for X. The assumption that $\mathrm{supp}(X, Z, W) \subset [0, 1]^{2p+r}$ can always be satisfied by carrying out a monotone transformation of (X, Z, W). The inferential problem is to test the null hypothesis, H_0, that

$$g(x, z) = G(x, z, \theta) \tag{2.2}$$

for some unknown $\theta \in \Theta$, known function G, and almost every $(x, z) \in [0, 1]^{p+r}$. The alternative hypothesis, H_1, is that there is no $\theta \in \Theta$ such that (2.2) holds for almost every $(x, z) \in [0, 1]^{p+r}$. The data, $\{Y_i, X_i, Z_i, W_i : i = 1, \ldots, n\}$, are a simple random sample of (Y, X, Z, W).

2.1. *The Test Statistic*

To form the test statistic, let f_{XZW} denote the probability density function of (X, Z, W) and let f_Z denote the probability density function of Z. Let ν be any function in $L_2[0, 1]^{p+r}$. For each $z \in [0, 1]^r$ define the operator T_z on $L_2[0, 1]^p$ by

$$T_z\nu(x, z) = \int t_z(\xi, x)\nu(\xi, z)\, d\xi,$$

where for each $(x_1, x_2) \in [0, 1]^{2p}$,

$$t_z(x_1, x_2) = \int f_{XZW}(x_1, z, w) f_{XZW}(x_2, z, w)\, dw.$$

Assume that T_z is nonsingular for each $z \in [0, 1]^r$. Then H_0 is equivalent to

$$\tilde{S}(x, z) \equiv T_z[g(\cdot, \cdot) - G(\cdot, \cdot, \theta)](x, z) = 0 \tag{2.3}$$

for some $\theta \in \Theta$ and almost every $(x, z) \in [0, 1]^{p+r}$. Hypothesis H_1 is equivalent to the statement that there is no $\theta \in \Theta$ such that (2.3) holds. A test statistic can be based on a sample analog of

$$\int \tilde{S}(x, z)^2\, dx\, dz,$$

but the resulting rate of testing is slower than $n^{-1/2}$ if $r > 0$. A rate of $n^{-1/2}$ can be achieved by carrying out an additional smoothing step. To this end, let $\ell(z_1, z_2)$ denote the kernel of a nonsingular integral operator, L, on $L_2[0, 1]^r$ if $r > 0$. That is, the operator L defined by

$$L\nu(z) = \int \ell(\zeta, z)\nu(\zeta)\, d\zeta$$

is nonsingular. Let L be the identity operator if $r = 0$. Define the operator T on $L_2[0,1]^{p+r}$ by $T\nu(x,z) = LT_z\nu(x,z)$. Then T is nonsingular. Hypothesis H_0 is equivalent to

$$S(x,z) \equiv T[g(\cdot,\cdot) - G(\cdot,\cdot,\theta)](x,z) = 0 \tag{2.4}$$

for some $\theta \in \Theta$ and almost every $(x,z) \in [0,1]^{p+r}$. Hypothesis H_1 is equivalent to the statement that there is no $\theta \in \Theta$ such that (2.4) holds. The test statistic is based on a sample analog of

$$\int S(x,z)^2\,dx\,dz.$$

To form the analog, observe that under H_0,

$$T[g - G(\cdot,\cdot,\theta)](x,z) = \mathbf{E}\big\{[Y - G(X,Z,\theta)]f_{XW}(x,z,W)\ell(Z,z)\big\}.$$

Therefore, it suffices to form a sample analog of $\mathbf{E}\{[Y - G(X,Z,\theta)]f_{XW}(x,z,W)\ell(Z,z)\}$. To do this, let $\hat{f}_{XZW}^{(-i)}$ denote a leave-observation-i-out kernel estimator of f_{XZW}. That is, for $V_i \equiv (X_i, Z_i, W_i)$ and κ a kernel function of a $(2p+r)$-dimensional argument,

$$\hat{f}_{XZW}^{(-i)}(v) = \frac{1}{nh^{2p+r}} \sum_{\substack{j=1 \\ j\neq i}}^{n} \kappa\left(\frac{v - V_i}{h}\right),$$

where h is the bandwidth. Let $\hat{\theta}_n$ be an estimator of θ. The sample analog of $S(x,z)$ is

$$S_n(x,z) = n^{-1/2} \sum_{i=1}^{n} [Y_i - G(X_i, Z_i, \hat{\theta}_n)]\hat{f}_{XZW}^{(-i)}(x, Z_i, W_i)\ell(Z_i, z).$$

The test statistic is

$$\tau_n = \int S_n^2(x,z)\,dx\,dz. \tag{2.5}$$

Hypothesis H_0 is rejected if τ_n is large.

2.2. *Regularity Conditions*

This section states the assumptions that are used to obtain the asymptotic properties of τ_n under the null and alternative hypotheses. Let $\|(x_1, z_1, w_1) - (x_2, z_2, w_2)\|$ denote the Euclidean distance between (x_1, z_1, w_1) and (x_2, z_2, w_2). Let $D_j f_{XZW}$ denote any jth partial or mixed partial derivative of f_{XZW}. Let $D_0 f_{XZW}(x,z,w) = f_{XZW}(x,z,w)$.

ASSUMPTION 1: (i) *The support of* (X, Z, W) *is contained in* $[0,1]^{2p+r}$. (ii) (X, Z, W) *has a probability density function* f_{XZW} *with respect to Lebesgue measure*. (iii) *There is a constant* $C_f < \infty$ *such that* $|D_j f_{XZW}(x, z, w)| \leq C_f$ *for all* $(x, z, w) \in [0,1]^{2p+r}$ *and* $j = 0, 1, 2$. (iv) $|D_2 f_{XZW}(x_1, z_1, w_1) - D_2 f_{XZW}(x_2, z_2, w_2)| \leq C_f \|(x_1, z_1, w_1) - (x_2, z_2, w_2)\|$ *for any second derivative and any* (x_1, z_1, w_1) *and* (x_2, z_2, w_2) *in* $[0,1]^{2p+r}$. (v) *The operator* T_z *is nonsingular for almost every* $z \in [0,1]^r$.

ASSUMPTION 2: (i) $\mathbf{E}(U|Z = z, W = w) = 0$ *and* $\mathbf{E}(U^2|Z = z, W = w) \leq C_U$ *for each* $(z, w) \in [0,1]^{p+r}$ *and some constant* $C_U < \infty$. (ii) $|g(x, z)| \leq C_g$ *for some constant* $C_g < \infty$ *and all* $(x, z) \in [0,1]^{p+r}$.

ASSUMPTION 3: (i) *As* $n \to \infty$, $\theta_n \to^p \theta_0$ *for some* $\theta_0 \in \Theta$, *a compact subset of* $\mathbb{R}^d$. *If* H_0 *is true, then* $g(x, z) = G(x, z, \theta_0)$, $\theta_0 \in \mathrm{int}(\Theta)$, *and*

$$n^{-1/2}(\hat{\theta}_n - \theta_0) = n^{1/2} \sum_{i=1}^{n} \gamma(U_i, X_i, Z_i, W_i, \theta_0) + o_p(1) \tag{2.6}$$

for some function γ *taking values in* $\mathbb{R}^d$ *such that* $\mathbf{E}\gamma(U, X, Z, W, \theta_0) = 0$ *and* $\mathrm{var}[\gamma(U, X, Z, W, \theta_0)]$ *is a finite, nonsingular matrix*.

ASSUMPTION 4: (i) $|G(x, z, \theta)| \leq C_G$ *for all* $(x, z) \in [0,1]^{p+r}$, *all* $\theta \in \Theta$, *and some constant* $C_G < \infty$. (ii) *The first and second derivatives of* $G(x, z, \theta)$ *with respect to* θ *are bounded by* C_G *uniformly over* $(x, z) \in [0,1]^{p+r}$ *and* $\theta \in \Theta$.

ASSUMPTION 5: (i) *The kernel function used to estimate* f_{XZW} *has the form* $\kappa(v) = \prod_{j=1}^{2p+r} K(v_j)$, *where* v_j *is the jth component of* v *and* K *is a symmetrical, twice continuously differentiable probability density function on* $[-1, 1]$. (ii) *The bandwidth*, h, *satisfies* $h = c_h n^{-1/(2p+r+4)}$, *where* c_h *is a constant and* $0 < c_h < \infty$. (iii) *The operator* L *is nonsingular*.

The representation (2.6) of $n^{1/2}(\hat{\theta}_n - \theta_0)$ holds, for example, if $\hat{\theta}_n$ is a generalized method of moments (GMM) estimator.

2.3. *The Asymptotic Distribution of the Test Statistic under the Null Hypothesis*

To obtain the asymptotic distribution of τ_n under H_0, define

$$G_\theta(x, z, \theta) = \frac{\partial G(x, z, \theta)}{\partial \theta},$$

$$\Gamma(x, z) = \mathbf{E}[G_\theta(X, Z, \theta_0) f_{XZW}(x, Z, W)\ell(Z, z)],$$

$$B_n(x,z) = n^{-1/2}\sum_{i=1}^{n}\big[U_i f_{XZW}(x, Z_i, W_i)\ell(Z_i, z) - \Gamma(x,z)'\gamma(U_i, X_i, Z_i, W_i, \theta_0)\big],$$

and

$$V(x_1, z_1; x_2, z_2) = \mathbf{E}[B_n(x_1, z_1)B_n(x_2, z_2)].$$

Define the operator Ω on $L_2[0,1]^{q+r}$ by

$$(\Omega\nu)(x,z) = \int_0^1 V(x, z; \xi, \zeta)\nu(\xi, \zeta)\, d\xi\, d\zeta. \tag{2.7}$$

Let $\{\omega_j : j = 1, 2, \ldots\}$ denote the eigenvalues of Ω sorted so that $\omega_1 \geq \omega_2 \geq \cdots \geq 0$. Let $\{\chi^2_{1j} : j = 1, 2, \ldots\}$ denote independent random variables that are distributed as chi-squared with 1 degree of freedom. The following theorem gives the asymptotic distribution of τ_n under H_0.

THEOREM 1: *If* H_0 *is true and Assumptions* 1–5 *hold*, *then*

$$\tau_n \xrightarrow{d} \sum_{j=1}^{\infty} \omega_j \chi^2_{1j}.$$

2.4. *Obtaining the Critical Value*

The statistic τ_n is not asymptotically pivotal, so its asymptotic distribution cannot be tabulated. This section presents a method for obtaining an approximate asymptotic critical value for the τ_n test. The method replaces the asymptotic distribution of τ_n with an approximate distribution. The difference between the true and approximate distributions can be made arbitrarily small under both the null hypothesis and alternatives. Moreover, the quantiles of the approximate distribution can be estimated consistently as $n \to \infty$. The approximate $1-\alpha$ critical value of the τ_n test is a consistent estimator of the $1-\alpha$ quantile of the approximate distribution.

The approximate critical value is obtained under sampling from a pseudo-true model that coincides with (2.1) if H_0 is true and satisfies a version of $\mathbf{E}[Y - G(X, \theta_0)|Z, W] = 0$ if H_0 is false. The critical value for the case of a false H_0 is used later to establish the properties of τ_n under H_1. The pseudo-true model is defined by

$$\tilde{Y} = G(X, Z, \theta) + \tilde{U}, \tag{2.8}$$

where $\tilde{Y} = Y - \mathbf{E}[Y - G(X, Z, \theta_0)|Z, W]$, $\tilde{U} = \tilde{Y} - G(X, Z, \theta_0)$, and θ_0 is the probability limit of $\hat{\theta}_n$. This model coincides with (2.1) when H_0 is true. Moreover, H_0 holds for the pseudo-true model in the sense that $\mathbf{E}[\tilde{Y} - G(X, Z, \theta_0)|Z, W] = 0$, regardless of whether H_0 holds for (2.1).

To describe the approximation to the asymptotic distribution of τ_n, let $\{\tilde{\omega}_j : j = 1, 2, \dots\}$ be the eigenvalues of the version of Ω (denoted $\tilde{\Omega}$) that is obtained by replacing model (2.1) with model (2.8). Order the $\tilde{\omega}_j$'s such that $\tilde{\omega}_1 \geq \tilde{\omega}_2 \geq \cdots \geq 0$. Then under sampling from (2.8), τ_n is asymptotically distributed as

$$\tilde{\tau} \equiv \sum_{j=1}^{\infty} \tilde{\omega}_j \chi^2_{1j}.$$

Given any $\varepsilon > 0$, there is an integer $K_\varepsilon < \infty$ such that

$$0 < \mathbf{P}\left(\sum_{j=1}^{K_\varepsilon} \tilde{\omega}_j \chi^2_{1j} \leq t\right) - \mathbf{P}(\tilde{\tau} \leq t) < \varepsilon$$

uniformly over t. Define

$$\tilde{\tau}_\varepsilon = \sum_{j=1}^{K_\varepsilon} \tilde{\omega}_j \chi^2_{1j}.$$

Let $z_{\varepsilon\alpha}$ denote the $1 - \alpha$ quantile of the distribution of $\tilde{\tau}_\varepsilon$. Then $0 < \mathbf{P}(\tilde{\tau} > z_{\varepsilon\alpha}) - \alpha < \varepsilon$. Thus, using $z_{\varepsilon\alpha}$ to approximate the asymptotic $1 - \alpha$ critical value of τ_n creates an arbitrarily small error in the probability that a correct H_0 is rejected. Similarly, use of the approximation creates an arbitrarily small change in the power of the τ_n test when H_0 is false. However, the eigenvalues $\tilde{\omega}_j$ are unknown. Accordingly, the approximate $1 - \alpha$ critical value for the τ_n test is a consistent estimator of the $1 - \alpha$ quantile of the distribution of $\tilde{\tau}_\varepsilon$. Specifically, let $\hat{\omega}_j$ $(j = 1, 2, \dots, K_\varepsilon)$ be a consistent estimator of $\tilde{\omega}_j$ under sampling from (2.8). Then the approximate critical value of τ_n is the $1 - \alpha$ quantile of the distribution of

$$\hat{\tau}_n = \sum_{j=1}^{K_\varepsilon} \hat{\omega}_j \chi^2_{1j}.$$

This quantile, which will be denoted $\hat{z}_{\varepsilon\alpha}$, can be estimated with arbitrary accuracy by simulation.

The remainder of this section describes how to obtain the estimated eigenvalues $\{\hat{\omega}_j\}$. Define $\tilde{W}_i = [H(W_i)', Z_i']'$, where H is a known, vector-valued function whose components are linearly independent, and $c_\theta \equiv \dim H + r \geq d$. Assume that $\hat{\theta}_n$ is the GMM estimator

$$\hat{\theta}_n = \arg\min_{\theta \in \Theta} \left\{ \sum_{i=1}^{n} [Y_i - G(X_i, Z_i, \theta)] \tilde{W}_i' \right\} \times A_n \left\{ \sum_{i=1}^{n} [Y_i - G(X_i, Z_i, \theta)] \tilde{W}_i \right\}, \tag{2.9}$$

where $\{A_n\}$ is a sequence of possibly stochastic $c_\theta \times c_\theta$ weight matrices converging in probability to a finite, nonstochastic matrix A. Define the $c_\theta \times d$ matrix $D = \mathbf{E}[\tilde{W} G_\theta(X, Z, \theta_0)']$ and the $d \times c_\theta$ matrix $\tilde{\gamma} = (D'AD)^{-1} D'A$. Then standard calculations for GMM estimators show that

$$\gamma(U_i, X_i, Z_i, W_i, \theta_0) = \tilde{\gamma} \tilde{W}_i \tilde{U}_i.$$

Therefore,

$$\begin{aligned} V(x_1, z_1; x_2, z_2) &= \mathbf{E}\left\{ n^{-1} \sum_{i=1}^{n} [f_{XZW}(x_1, Z_i, W_i)\ell(Z_i, z_1) - \Gamma(x_1, z_1)' \tilde{\gamma} \tilde{W}_i] \tilde{U}_i^2 \right. \\ &\qquad \left. \times [f_{XZW}(x_2, Z_i, W_i)\ell(Z_i, z_2) - \Gamma(x_2, z_2)' \tilde{\gamma} \tilde{W}_i] \right\}. \end{aligned} \tag{2.10}$$

A consistent estimator of V can be obtained by replacing unknown quantities on the right-hand side of (2.10) with estimators. To this end, define

$$\hat{D} = n^{-1} \sum_{i=1}^{n} \tilde{W}_i G_\theta(X, Z, \hat{\theta}_n)',$$

$$\hat{\gamma} = (\hat{D}' A_n \hat{D})^{-1} \hat{D}' A_n,$$

and

$$\hat{\Gamma}(x, z) = n^{-1} \sum_{i=1}^{n} G_\theta(X_i, Z_i, \hat{\theta}_n) \hat{f}_{XZW}(x, Z_i, W_i) \ell(Z_i, z),$$

where $\hat{f}_{XZW}$ is a kernel estimator of f_{XZW}. Also define $\hat{U}_i = Y_i - G(X_i, Z_i, \hat{\theta}_n) - \hat{q}^{(-i)}(Z_i, W_i)$, where $\hat{q}^{(-i)}(z, w)$ is the leave-observation-i-out kernel regression estimator of $Y - G(X, Z, \hat{\theta}_n)$ on (Z, W). Then $V(x_1, z_1; x_2, z_2)$ is estimated consistently by

$$\hat{V}(x_1, z_1; x_2, z_2) = \left\{ n^{-1} \sum_{i=1}^{n} [\hat{f}_{XZW}(x_1, Z_i, W_i)\ell(Z_i, z_1) - \hat{\Gamma}(x_1, z_1)'\hat{\gamma}\tilde{W}_i]\hat{U}_i^2 \times [\hat{f}_{XZW}(x_2, Z_i, W_i)\ell(Z_i, z_2) - \hat{\Gamma}(x_2, z_2)'\hat{\gamma}\tilde{W}_i] \right\}.$$

Let $\hat{\Omega}$ be the integral operator whose kernel is $\hat{V}(x_1, z_1; x_2, z_2)$. The $\hat{\omega}_j$'s are the eigenvalues of $\hat{\Omega}$.

THEOREM 2: *Let Assumptions* 1–5 *hold. Then as* $n \to \infty$, (i) $\sup_{1\le j\le K_\varepsilon} |\hat{\omega}_j - \tilde{\omega}_j| = O[(\log n)/(nh^{2p+r})^{1/2}]$ *almost surely and* (ii) $\hat{z}_{\varepsilon\alpha} \to^p z_{\varepsilon\alpha}$.

To obtain an accurate numerical approximation to the $\hat{\omega}_j$'s, let $\hat{F}(x, z)$ denote the $n \times 1$ vector whose ith component is $\hat{f}_{XW}(x, Z_i, W_i)\ell(Z_i, z)$, let $\hat{G}_\theta$ denote the $n \times d$ matrix whose (i, j) element is $G_\theta(X_i, Z_i, \hat{\theta}_n)$, let Y denote the $n \times n$ diagonal matrix whose (i, i) element is $\hat{U}_i^2$, and let $\tilde{W}$ denote the $n \times d$ matrix $(\tilde{W}_1', \ldots, \tilde{W}_n')'$. Finally, define the matrix $M = I_n - n^{-1}\hat{G}_\theta\tilde{\gamma}\tilde{W}'$, where I_n is the $n \times n$ identity matrix. Then

$$\hat{V}(z_1, z_2) = n^{-1}\hat{F}(x_1, z_1)'MYM'\hat{F}(x_2, z_2).$$

The computation of the $\hat{\omega}_j$'s can now be reduced to finding the eigenvalues of a finite-dimensional matrix. Let $\{\phi_j : j = 1, 2, \ldots\}$ be an orthonormal basis for $L_2[0, 1]^{p+r}$. Then

$$\hat{f}_{XZW}(x, z, W)\ell(Z, z) = \sum_{j=1}^{\infty}\sum_{k=1}^{\infty} \hat{d}_{jk}\phi_j(x, z)\phi_k(Z, W),$$

where

$$\hat{d}_{jk} = \int_0^1 dx \int_0^1 dz_1 \int_0^1 dz_2 \int_0^1 dw\, \hat{f}_{XZW}(x, z_1, w)\ell(z_1, z_2) \times \phi_j(x, z_1)\phi_k(z_2, w).$$

Approximate $\hat{f}_{XZW}(x, z, W)\ell(Z, z)$ by the finite sum

$$\Pi(x, z, W, Z) = \sum_{j=1}^{L}\sum_{k=1}^{L} \hat{d}_{jk}\phi_j(x, z)\phi_k(Z, W)$$

for some integer $L < \infty$. Since $\hat{f}_{XZW}\ell$ is a known function, L can be chosen to make Π approximate $\hat{f}_{XZW}\ell$ with any desired accuracy. Let $\phi(x, z)$ denote the $L \times 1$ vector whose jth component is $\phi_j(x, z)$. Let Φ be the $L \times n$ matrix whose (j, k)th component is $\phi_j(Z_k, W_k)$. Let D be the $L \times L$ matrix $\{d_{jk}\}$. Then $\hat{V}(x_1, z_1; x_2, z_2)$ is approximated by

$$\hat{\hat{V}}(x_1, z_1; x_2, z_2) = n^{-1}\phi(x_1, z_1)' D\Phi M Y M' \Phi' D' \phi(x_2, z_2).$$

The eigenvalues of $\hat{\Omega}$ are approximated by those of the $L \times L$ matrix $D\Phi M Y M' \Phi' D'$.

2.5. *Consistency of the Test Against a Fixed Alternative Model*

In this section, it is assumed that H_0 is false; that is, there is no $\theta \in \Theta$ such that $g(x, z) = G(x, z, \theta)$ for almost every (x, z). Let θ_0 denote the probability limit of $\hat{\theta}_n$. Define $q(x, z) = g(x, z) - G(x, z, \theta_0)$. Let $\tilde{z}_\alpha$ denote the $1 - \alpha$ quantile of the distribution of τ_n under sampling from the pseudo-true model (2.8). Let $\hat{z}_{\varepsilon\alpha}$ denote the $1 - \alpha$ quantile of $\hat{\tau}_n$. The following theorem establishes consistency of the τ_n test against a fixed alternative hypothesis.

THEOREM 3: *Let Assumptions* 1–5 *hold*. *Suppose that* H_0 *is false and that* $\int_0^1 [(Tq)(x, z)]^2\, dx\, dz > 0$. *Then for any* α *such that* $0 < \alpha < 1$,

$$\lim_{n\to\infty} \mathbf{P}(\tau_n > \tilde{z}_\alpha) = 1$$

and

$$\lim_{n\to\infty} \mathbf{P}(\tau_n > \hat{z}_\alpha) = 1.$$

Because T is nonsingular, the τ_n test is consistent against any alternative that differs from $G(x, z, \theta_0)$ on a set of (x, z) values whose Lebesgue measure exceeds zero.

2.6. *Asymptotic Distribution under Local Alternatives*

This section obtains the asymptotic distribution of τ_n under the sequence of local alternative hypotheses

$$Y = G(X, Z, \theta_0) + n^{-1/2}\Delta(X, Z) + U, \tag{2.11}$$

where Δ is a bounded function on $[0, 1]^{p+r}$ and $\theta_0 \in \text{int}(\Theta)$. The following additional notation is used. Let $\hat{\theta}_n$ be the GMM estimator (2.9). Let $\{\omega_j, \psi_j:$

$j = 1, 2, \dots\}$ denote the eigenvalues and orthonormal eigenvectors of the operator Ω that is defined in (2.7), ordered so that $\omega_1 \geq \omega_2 \geq \cdots \geq 0$. Define $\mu(x, z) = T\{\Delta - \mathbf{E}[W\Delta(X, Z)]'\tilde{\gamma}' G_\theta\}(x, z)$ and

$$\mu_j = \int_0^1 \mu(x, z)\psi_j(x, z)\, dx\, dz.$$

Let $\{\chi^2_{1j}(\mu_j^2/\omega_j) : j = 1, 2, \dots\}$ denote independent random variables that are distributed as noncentral chi-squared with 1 degree of freedom and noncentral parameters $\{\mu_j^2/\omega_j\}$. Let $\hat{\theta}_n$ be the GMM estimator (2.9). The following theorem states the result.

THEOREM 4: *Let Assumptions* 1–5 *hold*. *Under the sequence of local alternatives* (2.11),

$$\tau_n \xrightarrow{d} \sum_{j=1}^{\infty} \omega_j \chi^2_{1j}(\mu_j^2/\omega_j).$$

Let z_α denote the $1 - \alpha$ quantile of the distribution of $\sum_{j=1}^{\infty} \omega_j \chi^2_{1j}(\mu_j^2/\omega_j)$. Let $\hat{z}_{\varepsilon\alpha}$ denote the estimated approximate α-level critical value defined in Section 2.2. Then it follows from Theorems 2 and 4 that for any $\varepsilon > 0$,

$$\limsup_{n\to\infty} |\mathbf{P}(\tau_n > \hat{z}_{\varepsilon\alpha}) - \mathbf{P}(\tau_n > z_\alpha)| \leq \varepsilon.$$

It also follows from Theorem 4 that the τ_n test has power against local alternatives whose distance from the null-hypothesis model is $O(n^{-1/2})$. In contrast, the test of Tripathi and Kitamura (2003) has power only against local alternatives whose distance from the null-hypothesis model decreases more slowly than $n^{-1/2}$. If $\mu(x, z) = 0$ for all (x, z), then there is a nonstochastic sequence $\{\theta_n\}$ such that $G(x, z, \theta_n) = G(x, z, \theta_0) + n^{-1/2}\Delta(x, z) + o(n^{-1/2})$. Therefore, the distance between the null and alternative hypotheses is $o(n^{-1/2})$.

2.7. *Uniform Consistency*

This section shows that for any $\varepsilon > 0$, the τ_n test rejects H_0 with probability exceeding $1 - \varepsilon$ uniformly over a class of alternative models whose distance from the null hypothesis is $O(n^{-1/2})$. The following additional notation is used. Let θ_g be the probability limit of $\hat{\theta}_n$ under the hypothesis (not necessarily true) that $g(x, z) = G(x, z, \theta)$ for some $\theta \in \Theta$ and a given function G. Let $\tilde{\Theta}$ be a compact subset of $\mathrm{int}(\Theta)$. Define $q_g(x, z) = g(x, z) - G(x, z, \theta_g)$. Let h denote the bandwidth in $f_{XZW}^{(-i)}$. For each $n = 1, 2, \dots$ and $C > 0$ define $\mathcal{F}_{nc}$ as a set of functions g such that (i) $|g(x, z)| \leq C_g$ for all $(x, z) \in [0, 1]^{p+r}$ and some constant $C_g < \infty$, (ii) $\theta_g \in \tilde{\Theta}$, (iii) (2.6) holds uniformly over $g \in \mathcal{F}_{nc}$, (iv) $\|Tq_g\| \geq n^{-1/2}C$, where $\|\cdot\|$ denotes the L_2 norm, and (v) $h^2\|q_g\|/\|Tq_g\| = o(1)$ as $n \to \infty$. The set $\mathcal{F}_{nc}$ is a set of functions whose distance from H_0 shrinks to zero at the rate $n^{-1/2}$; that is, $\mathcal{F}_{nc}$ includes functions such that $\|q_g\| = O(n^{-1/2})$. Condition (ii) ensures the existence of the critical value defined in Section 2.4. The requirement $\theta_g \in \tilde{\Theta}$ is not restrictive in applications because Θ and $\tilde{\Theta}$ can usually be made large enough to include any reasonable θ_g. Condition (v) rules out alternatives that depend on x only through sequences of eigenvectors of T whose eigenvalues converge to 0 too rapidly. For example, let $p = 1$, $r = 0$, and $\{\lambda_j, \phi_j : j = 1, 2, \ldots\}$ denote the eigenvalues and eigenvectors of T ordered so that $\lambda_1 \geq \lambda_2 \geq \cdots > 0$. Let $G(x, \theta) = \theta\phi_1(x)$, $g(x) = \phi_1(x) + \phi_n(x)$, and $\tilde{W} = \phi_1(W)$. Then $h^2\|q_g\|/\|Tq_g\| = h^2/\lambda_n$. Because $h \propto n^{-1/6}$, condition (v) is violated if $\lambda_n = o(n^{-1/3})$. The practical significance of condition (v) is that the τ_n test has relatively low power against alternatives that differ from H_0 only through eigenvectors of T with very small eigenvalues.

The following theorem states the result of this section.

THEOREM 5: *Let Assumptions* 1, 2, 4, *and* 5 *hold. Assume that* $\hat{\theta}_n$ *satisfies* (2.9). *Then given any* $\delta > 0$, α *such that* $0 < \alpha < 1$, *and any sufficiently large but finite constant* C,

$$\lim_{n\to\infty} \inf_{\mathcal{F}_{nc}} \mathbf{P}(\tau_n > z_\alpha) \geq 1 - \delta \tag{2.12}$$

and

$$\lim_{n\to\infty} \inf_{\mathcal{F}_{nc}} \mathbf{P}(\tau_n > \hat{z}_{\varepsilon\alpha}) \geq 1 - 2\delta. \tag{2.13}$$

2.8. *Alternative Weights*

This section compares τ_n with a generalization of the test of Bierens (1990) and Bierens and Ploberger (1997). To minimize the complexity of the discussion, assume that $p = 1$ and $r = 0$, so Z is not in the model. Let H be a bounded, real-valued function on $[0, 1]^2$ such that

$$\left\| \int_0^1 H(x, w)s(w)\,dw \right\|^2 = 0$$

only if $s(w) = 0$ for almost every $w \in [0, 1]$. Then a test of H_0 can be based on the statistic

$$\tau_{nH} = \int_0^1 S_{nH}^2(x)\,dx,$$

where

$$S_{nH}(x) = n^{-1/2}\sum_{i=1}^{n}[Y_i - G(X_i, \hat{\theta}_n)]H(x, W_i).$$

If $H(x, w) = \tilde{H}(xw)$ for a suitably chosen function $\tilde{H}$, then τ_{nH} is a modification of the statistic of Bierens (1990) and Bierens and Ploberger (1997) for testing the hypothesis that a conditional mean function belongs to a specified, finite-dimensional parametric family. In this section, it is shown that the power of the τ_{nH} test can be low relative to that of the τ_n test. Specifically, there are combinations of density functions of (X, W), f_{XW}, and local alternative models (2.11) such that an α-level τ_{nH} test based on a fixed H has asymptotic local power arbitrarily close to α, whereas the α-level τ_n test has asymptotic local power that is bounded away from α. The opposite situation cannot occur under the assumptions of this paper; that is, it is not possible for the asymptotic power of the α-level τ_n test to approach α while the power of the α-level τ_{nH} test remains bounded away from α.

The conclusion that the power of τ_{nH} can be low relative to that of τ_n is reached by constructing an example in which the α-level τ_n test has asymptotic power that is bounded away from α, but the τ_{nH} test has asymptotic power that is arbitrarily close to α. To minimize the complexity of the example, assume that θ is known and does not have to be estimated. Define

$$\bar{B}_n(x) = n^{-1/2}\sum_{i=1}^{n} U_i f_{XW}(x, W_i),$$

$$\bar{B}_{nH}(x) = n^{-1/2}\sum_{i=1}^{n} U_i H(x, W_i),$$

$$\bar{R}(x_1, x_2) = \mathbf{E}[\bar{B}_n(x_1)\bar{B}_n(x_2)],$$

and

$$\bar{R}_H(x_1, x_2) = \mathbf{E}[\bar{B}_{nH}(x_1)\bar{B}_{nH}(x_2)].$$

Also, define the operators $\bar{\Omega}$ and $\bar{\Omega}_H$ on $L_2[0, 1]$ by

$$(\bar{\Omega}\psi)(x) = \int_0^1 \bar{R}(x, \xi)\psi(\xi)\, d\xi$$

and

$$(\bar{\Omega}_H\psi)(x) = \int_0^1 \bar{R}_H(x, \xi)\psi(\xi)\, d\xi.$$

Let $\{\bar{\omega}_j, \bar{\psi}_j : j = 1, 2, \dots\}$ and $\{\bar{\omega}_{jH}, \bar{\psi}_{jH} : j = 1, 2, \dots\}$ denote the eigenvalues and eigenvectors of $\bar{\Omega}$ and $\bar{\Omega}_H$, respectively, with the eigenvalues sorted in decreasing order. For Δ defined as in (2.11), define

$$\bar{\mu}(x) = (T\Delta)(x),$$

$$\bar{\mu}_H(x) = \int_0^1 \int_0^1 \Delta(\xi) H(x, w) f_{XW}(\xi, w)\, dx\, dw,$$

$$\bar{\mu}_j = \int_0^1 \bar{\mu}(x) \bar{\psi}_j(x)\, dx,$$

and

$$\bar{\mu}_{jH} = \int_0^1 \bar{\mu}_H(x) \bar{\psi}_{jH}(x)\, dx.$$

Then arguments like those used to prove Theorem 4 show that under the sequence of local alternatives (2.11) with a known θ,

$$\tau_n \xrightarrow{d} \sum_{j=1}^{\infty} \bar{\omega}_j \chi^2_{1j}(\bar{\mu}_j^2 / \bar{\omega}_j)$$

and

$$\tau_{nH} \xrightarrow{d} \sum_{j=1}^{\infty} \bar{\omega}_{jH} \chi^2_{1j}(\bar{\mu}_{jH}^2 / \bar{\omega}_{jH})$$

as $n \to \infty$. Therefore, to establish the first conclusion of this section, it suffices to show that for a fixed function H, f_{XW} and Δ can be chosen so that $\|\bar{\mu}\|^2 / \sum_{j=1}^{\infty} \bar{\omega}_j$ is bounded away from 0 and $\|\bar{\mu}_H\|^2 / \sum_{j=1}^{\infty} \bar{\omega}_{jH}$ is arbitrarily close to 0.

To this end, let $\phi_1(x) = 1$ and $\phi_{j+1}(x) = 2^{-1/2} \cos(j\pi x)$ for $j \ge 1$. Let $m > 1$ be a finite integer. Define

$$\lambda_j = \begin{cases} 1, & \text{if } j = 1 \text{ or } m, \\ e^{-2j}, & \text{otherwise.} \end{cases}$$

Let

$$f_{XW}(x, w) = 1 + \sum_{j=1}^{\infty} \lambda_{j+1}^{1/2} \phi_{j+1}(x) \phi_{j+1}(w).$$

Let $\mathbf{E}(U^2 | W = w) = 1$ for all $w \in [0, 1]$. Then $\bar{R}(x_1, x_2) = t(x_1, x_2)$, $\bar{\omega}_j = \lambda_j$, and $\sum_{j=1}^{\infty} \bar{\omega}_j$ is nonzero and finite. Set $\Delta(x) = D\phi_m(x)$ for some finite $D > 0$.

Then $\|\bar{\mu}\|^2 = D^2\lambda_m^2 = D^2$. It suffices to show that m can be chosen so that $\|\bar{\mu}_H\|^2$ is arbitrarily close to 0. To do this, observe that $H(z, w)$ has the Fourier representation

$$H(x, w) = \sum_{j,k=1}^{\infty} h_{jk}\phi_j(x)\phi_k(w),$$

where $\{h_{jk} : j, k = 1, 2, \dots\}$ are constants. Moreover, $\|\bar{\mu}_H\|^2 = D^2\sum_{j=1}^{\infty} h_{jm}^2$. Since H is bounded, m can be chosen so that $\sum_{j=1}^{\infty} h_{jm}^2 < \varepsilon/D^2$ for any $\varepsilon > 0$. With this m, $\|\bar{\mu}_H\|^2 < \varepsilon$, which establishes the first conclusion.

The opposite situation (a sequence of local alternatives for which $\|\bar{\mu}\|^2$ approaches 0 while $\|\bar{\mu}_H\|^2$ remains bounded away from 0) cannot occur. To show this, assume without loss of generality that the marginal distributions of X and W are $U[0, 1]$, $\mathbf{E}(U^2|W = w) = 1$ for all $w \in [0, 1]$, and $\sum_{j=1}^{\infty} \bar{\omega}_{jH} = 1$. Also assume that $\|\Delta\|^2 < C_\Delta$ for some constant $C_\Delta < \infty$. Then

$$\int_0^1 \int_0^1 H(x, w)^2\, dx\, dw = \sum_{j=1}^{\infty} \bar{\omega}_{jH}.$$

It follows from the Cauchy–Schwarz inequality that

$$\begin{aligned}
\|\bar{\mu}_H\|^2 &\le \left[\int_0^1 \int_0^1 H(x, w)^2\, dz\, dw\right] \int_0^1 \left[\int_0 f_{XW}(x, w)\Delta(x)\, dx\right]^2 dw \\
&= \int_0^1 \left[\int_0 f_{XW}(x, w)\Delta(x)\, dx\right]^2 dw \\
&\le \|\Delta\|^2 \|T\Delta\|^2 \\
&\le C_\Delta \|\bar{\mu}\|^2.
\end{aligned}$$

Therefore, $\|\bar{\mu}\|^2$ can approach 0 only if $\|\bar{\mu}_H\|^2$ also approaches 0.

3. MONTE CARLO EXPERIMENTS

This section reports the results of a Monte Carlo investigation of the finite-sample performance of the τ_n test. The experiments consist of testing the null hypotheses, H_0, that

(3.1) $$g(x) = \theta_0 + \theta_1 x$$

and

(3.2) $$g(x) = \theta_0 + \theta_1 x + \theta_2 x^2.$$

The alternative hypotheses are (3.2) if (3.1) is H_0 and

$$(3.3) \qquad g(x) = \theta_0 + \theta_1 x + \theta_2 x^2 + \theta_3 x^3$$

if either (3.1) or (3.2) is H_0.

To provide a basis for judging whether the power of the τ_n test is high or low, this section also reports the results of several other tests. One is an asymptotic t test of the hypothesis $\theta_2 = 0$ if (3.1) is H_0 and of $\theta_3 = 0$ if (3.2) is H_0. The t test is an example of an ad hoc test that might be used in applied research. Another test is the modified test of Bierens (1990) and Bierens and Ploberger (1997) that is described in Section 2.8. The weight function is $H(x, w) = \exp(xw)$. The critical value was computed using the methods described in Section 2.4. We also present Monte Carlo results for the instrumental variables (IV), generalized methods of moments (GMM), and exponential tilting (ET) tests of Donald, Imbens, and Newey (2003) (hereinafter DIN). Splines with 4 knots were used to form instruments for the DIN tests (DIN, equation (2.3)). The results are similar for splines with 2–5 knots. The power of the DIN tests is lower and the errors in the probability of rejecting a correct H_0 are higher with 1 knot or more than 5 knots.

We also carried out experiments with the smoothed empirical likelihood ratio (SELR) test of Tripathi and Kitamura (2003). Its power at the nominal 0.05 level was below 0.14, and well below the powers of the other tests in all experiments and over a wide range of values of the SELR bandwidth parameter. Consequently, the SELR test is not discussed further here.

In all experiments, $\theta_0 = 0$ and $\theta_1 = 0.5$. When (3.2) is the correct model, $\theta_2 = -0.5$. When (3.3) is the correct model, $\theta_2 = -1$, $\theta_3 = 1$ if (3.1) is H_0, and $\theta_3 = 2$ if (3.2) is H_0. Realizations of (X, W) were generated by $X = \Phi(\xi)$, $W = \Phi(\zeta)$, where Φ is the cumulative normal distribution function, $\zeta \sim N(0, 1)$, $\xi = \rho\zeta + (1 - \rho^2)^{1/2}\varepsilon$, $\varepsilon \sim N(0, 1)$, and ρ is a constant parameter whose value varies among experiments. Realizations of Y were generated from $Y = g(x) + \sigma_U U$, where $U = \eta\varepsilon + (1 - \eta^2)^{1/2}\nu$, $\nu \sim N(0, 1)$, $\sigma_U = 0.2$, and η is a constant parameter whose value varies among experiments. The instruments used to estimate (3.1), (3.2), and (3.3), respectively, are $(1, W)$, $(1, W, W^2)$, and $(1, W, W^2, W^3)$. The bandwidth h used to estimate f_{XW} was selected by cross-validation. The kernel is $K(v) = (15/16)(1 - v^2)^2 I(|v| \leq 1)$, where I is the indicator function. The asymptotic critical value was estimated by setting $K_\varepsilon = 25$. The results are not sensitive to the choice of K_ε and the estimated eigenvalues $\hat{\omega}_j$ are very close to 0 when $j > 25$. The sample size is $n = 500$ and the nominal level is 0.05. There are 1,000 Monte Carlo replications in each experiment. Computation of the critical value took approximately 4 seconds on a 900 MHz PC.

The results are shown in Table I. The differences between the nominal and empirical rejection probabilities are small when H_0 is true. When H_0 is false, the powers of the τ_n and t tests are similar. Not surprisingly, the t tests of (3.1)

TABLE I
RESULTS OF MONTE CARLO EXPERIMENTS

Null Model	Alt. Model	ρ	η	Empirical Probability that H_0 Is Rejected Using: τ_n Test	t Test	Bierens' Test	IV Test	GMM Test	ET Test
					H_0 *is true*				
(3.1)		0.8	0.1	0.051	0.052	0.053	0.048	0.041	0.060
		0.8	0.5	0.030	0.034	0.029	0.043	0.041	0.054
		0.7	0.1	0.049	0.052	0.053	0.045	0.043	0.060
(3.2)		0.8	0.1	0.053	0.040	0.054	0.048	0.046	0.054
		0.8	0.5	0.046	0.077	0.043	0.050	0.050	0.048
		0.7	0.1	0.056	0.036	0.036	0.043	0.043	0.052
					H_0 *is false*				
(3.1)	(3.2)	0.8	0.1	0.658	0.714	0.470	0.447	0.426	0.441
		0.8	0.5	0.721	0.827	0.466	0.459	0.440	0.446
		0.7	0.1	0.421	0.444	0.280	0.259	0.244	0.254
(3.1)	(3.3)	0.8	0.1	0.684	0.671	0.479	0.498	0.455	0.468
		0.8	0.5	0.663	0.580	0.464	0.480	0.421	0.437
		0.7	0.1	0.424	0.412	0.274	0.262	0.231	0.270
(3.2)	(3.3)	0.8	0.1	0.890	0.900	0.038	0.722	0.681	0.610
		0.8	0.5	0.972	0.987	0.030	0.685	0.594	0.555
		0.7	0.1	0.527	0.590	0.059	0.298	0.275	0.276

against (3.2) and (3.2) against (3.3) are somewhat more powerful than the τ_n tests. The τ_n test is slightly more powerful for testing (3.1) against (3.3). The Bierens-type and DIN tests are much less powerful than the τ_n test.

4. AN EMPIRICAL EXAMPLE

This section presents an empirical example in which τ_n is used to test two hypotheses about the shape of an Engle curve: one is that the curve is linear; the other is that it is quadratic. The curve is given by (2.1) with $r = 0$, so Z is not in the model. Variable Y denotes the logarithm of the expenditure share of food consumed off the premises where it was purchased, X denotes the logarithm of total expenditures, and W denotes annual income from wages and salaries. The data consist of 785 household-level observations from the 1996 U.S. Consumer Expenditure Survey. The bandwidth for estimating f_{XW} was selected by cross-validation. The kernel is the same as the one used in the Monte Carlo experiments. As in the experiments, $K_\varepsilon = 25$.

The τ_n test of the hypothesis that g is linear (quadratic) gives $\tau_n = 13.4$ (0.32) with a 0.05-level critical value of 3.07 (5.22). Thus, the test rejects the hypothesis that g is linear but not that g is quadratic. The hypotheses were also tested using the t test described in Section 3. This test gives $t = 2.60$ for the hypothesis that g is linear ($\theta_2 = 0$ in (3.2)) and $t = 0.34$ for the hypothesis that g is quadratic ($\theta_3 = 0$ in (3.3)). The 0.05-level critical value is 1.96. Thus, the t test also rejects the hypothesis that g is linear, but not the hypothesis that it is quadratic.

Dept. of Economics, Northwestern University, 2001 Sheridan Road, Evanston, IL 60208, U.S.A.; joel-horowitz@northwestern.edu.

Manuscript received March, 2004; final revision received May, 2005.

REFERENCES

BIERENS, H. J. (1990): "A Consistent Conditional Moment Test of Functional Form," *Econometrica*, 58, 1443–1458.

BIERENS, H. J., AND W. PLOBERGER (1997): "Asymptotic Theory of Integrated Conditional Moment Tests," *Econometrica*, 65, 1129–1151.

BLUNDELL, R., X. CHEN, AND D. KRISTENSEN (2003): "Semi-Nonparametric IV Estimation of Shape Invariant Engle Curves," Working Paper CWP 15/03, Centre for Microdata Methods and Practice, University College London.

DAROLLES, S., J.-P. FLORENS, AND E. RENAULT (2002): "Nonparametric Instrumental Regression," Working Paper, GREMAQ, University of Social Science, Toulouse.

DONALD, S. G., G. W. IMBENS, AND W. K. NEWEY (2003): "Empirical Likelihood Estimation and Consistent Tests with Conditional Moment Restrictions," *Journal of Econometrics*, 117, 55–93.

GUERRE, E., AND P. LAVERGNE (2002): "Optimal Minimax Rates for Nonparametric Specification Testing in Regression Models," *Econometric Theory*, 18, 1139–1171.

HALL, P., AND J. L. HOROWITZ (2005): "Nonparametric Methods for Inference in the Presence of Instrumental Variables," *The Annals of Statistics*, forthcoming.

HOROWITZ, J. L. (2005): "Supplement to 'Testing a Parametric Model Against a Nonparametric Alternative with Identification Through Instrumental Variables'," *Econometrica*, 74, http://www.econometricsociety.org/ecta/supmat/5140proofs.pdf.

HOROWITZ, J. L., AND V. G. SPOKOINY (2001): "An Adaptive, Rate-Optimal Test of a Parametric Mean Regression Model Against a Nonparametric Alternative," *Econometrica*, 69, 599–631.

——— (2002): "An Adaptive, Rate-Optimal Test of Linearity for Median Regression Models," *Journal of the American Statistical Association*, 97, 822–835.

NEWEY, W. K., AND J. L. POWELL (2003): "Instrumental Variable Estimation of Nonparametric Models," *Econometrica*, 71, 1565–1578.

NEWEY, W. K., J. L. POWELL, AND F. VELLA (1999): "Nonparametric Estimation of Triangular Simultaneous Equations Models," *Econometrica*, 67, 565–603.

TRIPATHI, G., AND Y. KITAMURA (2003): "Testing Conditional Moment Restrictions," *The Annals of Statistics*, 31, 2059–2095.