

Nonparametric estimation and inference under shape restrictions

Joel L. Horowitz
Soakbae Lee

The Institute for Fiscal Studies
Department of Economics, UCL

cemmap working paper CWP29/16

NONPARAMETRIC ESTIMATION AND INFERENCE UNDER SHAPE RESTRICTIONS

by

Joel L. Horowitz
Northwestern University
Evanston, IL 60208
U.S.A.

Sokbae Lee
Institute for Fiscal Studies
London WC1E 7AE U.K.

and

Department of Economics
Columbia, University
New York, NY 10027
U.S.A.

July 2016

Abstract

Economic theory often provides shape restrictions on functions of interest in applications, such as monotonicity, convexity, non-increasing (non-decreasing) returns to scale, or the Slutsky inequality of consumer theory; but economic theory does not provide finite-dimensional parametric models. This motivates nonparametric estimation under shape restrictions. Nonparametric estimates are often very noisy. Shape restrictions stabilize nonparametric estimates without imposing arbitrary restrictions, such as additivity or a single-index structure, that may be inconsistent with economic theory and the data. This paper explains how to estimate and obtain an asymptotic uniform confidence band for a conditional mean function under possibly nonlinear shape restrictions, such as the Slutsky inequality. The results of Monte Carlo experiments illustrate the finite-sample performance of the method, and an empirical example illustrates its use in an application.

Research supported in part by European Research Council Grant ERC-2014-CoG-646917-ROMIA. We thank David Jacho-Chávez for providing the data used in this paper. Part of this research was carried out while Joel L. Horowitz was a visitor at the Department of Economics, University College London, and the Centre for Microdata Methods and Practice.

NONPARAMETRIC ESTIMATION AND INFERENCE UNDER SHAPE RESTRICTIONS

1. INTRODUCTION

Let Y be a scalar random variable and X be a scalar random variable or vector. This paper presents a method for nonparametrically estimating and carrying out inference about the conditional mean function

$$g(x) \equiv E(Y | X = x)$$

under a shape restriction on g such as monotonicity, convexity, non-increasing (non-decreasing) returns to scale, or the Slutsky inequality of consumer theory. Economic theory often provides shape restrictions but does not provide finite-dimensional parametric models. For example, cost functions are monotone increasing, concave in input prices, and may exhibit non-increasing or non-decreasing returns to scale. Demand functions satisfy the Slutsky inequality, which is nonlinear. This motivates nonparametric estimation under shape restrictions. This paper explains how to estimate and form a uniform confidence band for g under shape restrictions that are more complicated than monotonicity or convexity and may be nonlinear.

It is well known that g can be estimated consistently and with the optimal rate of convergence without imposing shape restrictions. Fan and Gijbels (1996) and Härdle (1990), among many others, describe nonparametric estimation and rates of convergence without shape restrictions. Mammen (1991a, 1991b), Mammen and Thomas-Agnan (1999), and Wang and Shen (2013) discuss rates of convergence with shape restrictions. However, fully nonparametric estimates can be noisy and inconsistent with economic theory due to random sampling errors. For example, Blundell, Horowitz, and Parey (2012, 2016) found fully nonparametric estimates of demand functions to be wiggly and non-monotonic. Blundell, Horowitz, and Parey (2012, 2016) also found that imposing the Slutsky restriction reduced random noise and led to well-behaved nonparametric estimates without the need for arbitrary and possibly incorrect parametric or semiparametric assumptions.

Many methods are available for carrying out consistent nonparametric estimation under shape restrictions. See, for example, Hall and Huang (2001, 2002); Hall, Huang, Gifford and Gijbels (2001); Hall and Presnell (1999); Matzkin (1994); and the references cited in the foregoing paragraph. Asymptotic inference is not difficult if the values of x at which the shape restriction binds or does not bind in the sampled population are known. Liew (1976) illustrates this in the context of inequality constrained estimation of a linear model. Du, Parmeter, and Racine (2013) carry out kernel nonparametric estimation. In applications, however, it is not known where in the sampled population the shape restriction does or does not bind. This greatly complicates inference, because random sampling errors can cause the shape

restriction to bind or not bind the estimated and true g at different values of x . A similar problem arises in inference about a finite-dimensional parameter that may be on the boundary of the parameter set (Andrews 1999). Existing results on inference about a shape-restricted, nonparametrically estimated conditional mean function are limited to functions that are assumed to be monotonic or convex. The literature on inference under monotonicity or convexity restrictions is vast. See, among many others, Banerjee and Wellner (2001); Birke and Dette (2006); Chernozhukov, Fernandez-Val, and Galichon (2009); Dette, Neumeyer, and Pilz (2006); Dumbgen (2003); Groeneboom and Jongbloed (2015); Groeneboom, Jongbloed, and Wellner (2001); Pal and Woodroffe (2007); and the references therein. Existing results do not treat shape restrictions such as increasing or decreasing returns to scale and the Slutsky inequality that are of particular importance in economics. There is also a large literature on testing the hypothesis that a shape restriction holds. See, for example, Andrews and Shi (2013); Chernozhukov, Lee, and Rosen (2013); Hall and Yatchew (2005); Lee, Song, and Whang (2013); Romano, Shaikh, and Wolf (2014); and the references therein.

This paper is concerned with inference under shape restrictions, such as the Slutsky restriction, that may be nonlinear in a sense that is defined in Section 4. We formulate the estimation problem as minimization of a local quadratic objective function subject to constraints that implement the shape restriction. In general, the shape restriction generates a continuum of constraints. We reduce the number of constraints to a finite value by imposing the shape restriction and estimating g only on a discrete grid of points x in the support of X . We obtain a confidence band that is uniform over points in the grid. The grid becomes finer as the sample size, n , increases, thereby ensuring that, asymptotically, the shape restriction holds everywhere in the support of X . This enables us to obtain a confidence band for g that, asymptotically, is uniform over the support of X and satisfies the shape restriction. In practice, a confidence band can be computed only on a grid, so there is little practical difference between a band that is uniform over grid points and one that is uniform over a continuum.

The use of a discrete grid of points x enables us to overcome the problem of not knowing which constraints are binding in the sampled population. Let $\bar{\mathcal{S}}_n$ be the set of constraints that bind in the population or nearly bind in a sense that is defined in Section 4. This set is unknown. We find a data-based set $\hat{\mathcal{S}}_n$ of “possibly binding” constraints and carry out estimation under the (possibly false) assumption that $\hat{\mathcal{S}}_n = \bar{\mathcal{S}}_n$. We show that $\hat{\mathcal{S}}_n = \bar{\mathcal{S}}_n$ with probability approaching 1 as $n \rightarrow \infty$. Consequently $\bar{\mathcal{S}}_n$ can be treated as known asymptotically, and asymptotic inference can be carried out as if $\bar{\mathcal{S}}_n$ were known and $\hat{\mathcal{S}}_n = \bar{\mathcal{S}}_n$.

Let $g_0(x)$ and $\hat{g}(x)$, respectively, denote the true conditional mean function and the shape-restricted nonparametric estimator. We show that with suitable scaling, $\hat{g}(x) - g_0(x)$ is asymptotically jointly normally distributed with mean 0 over grid points. Asymptotic normality makes it possible to obtain a confidence band for g_0 that is uniform over grid points. As $n \rightarrow \infty$ and the distance between grid points approaches 0, the uniform confidence band over grid points converges to a uniform confidence band over all values of x .

Estimation of $g(x)$ at points x that are not in the grid is unnecessary for forming an asymptotic uniform confidence band for g but may be of interest for other reasons. Estimation of $g(x_{new})$ at a point x_{new} that is not in the grid can be carried out using the methods of this paper by shifting the location of the grid so that x_{new} is a point of the shifted grid. Alternatively, $g(x_{new})$ can be estimated using any of a variety of methods for interpolating $g(x)$ between grid points subject to the shape restrictions. The choice among interpolation methods is arbitrary and, except in special cases, does not yield an estimator that converges in probability as rapidly as an estimator based on the shifted grid.

Section 2 outlines the main steps involved in implementing our method. Section 3 presents the unconstrained and constrained nonparametric estimators of g and defines the grid. Section 4 describes the method for finding the set $\hat{\mathcal{S}}_n$ of possibly binding constraints. Section 5 explains how to carry out inference about g and form a uniform confidence band for g under shape restrictions. The confidence band obtained in Section 5 is uniform over the support of X and also over a class of functions g that includes nearly binding constraints. To minimize notational complexity, the discussion in Sections 2-5 assumes that X is a scalar random variable. The extension to higher dimensions is outlined in Section 6. Section 7 presents the results of Monte Carlo experiments and an empirical example that illustrate the numerical performance of our methods. Section 8 presents concluding comments. The proofs of theorems are in the appendix, which is Section 9.

2. A GUIDE TO IMPLEMENTATION

This section outlines the main steps of our method for estimating and obtaining a uniform confidence band for g . We assume here that X is a scalar random variable whose support is $[0,1]$. The extension to a multidimensional X is presented in Section 6.

1. Define a grid $0 < x_1 < x_2 < \dots < x_J < 1$ of J equally spaced points on $(0,1)$. A data-based method for choosing J in applications is presented in Section 7.

2. Estimate $g(x_j)$ ($j=1,\dots,J$) nonparametrically by using local quadratic estimation with bandwidth h . Let $\tilde{g}(x_j)$ denote the resulting estimate. A method for choosing h in applications is presented in Section 3.1.

3. Use the estimates $\tilde{g}(x_j)$ to find the set $\hat{\mathcal{S}}_n$ of possibly binding shape constraints. $\hat{\mathcal{S}}_n$ is given by equation (4.6).

4. Re-estimate $g(x_j)$ ($j=1,\dots,J$) nonparametrically using constrained local quadratic estimation under the restriction that the shape constraints in $\hat{\mathcal{S}}_n$ are binding (that is, they are equalities) and ignoring all other shape constraints.

5. Form a uniform confidence band for g using either the method of equations (5.8) and (5.9) or the method of Section 5.3.

3. THE ESTIMATORS OF g

This section describes our methods for estimating g with and without shape restrictions. The unrestricted estimator is used to estimate the set of possibly binding constraints. The shape-restricted estimator is an extension of the unrestricted estimator. Section 3.1 presents the unrestricted estimator. Section 3.2 presents grid and the shape-restricted estimator.

3.1 *The Unrestricted Estimator*

This section presents the unrestricted nonparametric estimator of g that is used throughout the remainder of this paper. Let $\{Y_i, X_i : i=1,\dots,n\}$ denote an independent random sample from the distribution of (Y, X) . Assume for now that X is a scalar random variable. The extension to a multidimensional X is presented in Section 6. Also assume that the support of X is a compact interval. Without further loss of generality, let this interval be $[0,1]$.

We use local quadratic estimation with bandwidth $h \propto n^{-1/5}$ to obtain the unrestricted nonparametric estimator of g . In applications, the bandwidth can be chosen by using cross-validation or plug-in methods for local constant or local linear estimation. Under our assumptions, local quadratic estimation with $h \propto n^{-1/5}$ provides an estimator of g that is free of asymptotic bias, and the bandwidth can be selected by standard methods. Local constant, local linear, and series estimation methods with a bandwidth selected by cross-validation or plug-in methods do not have this property. They require undersmoothing or explicit bias correction to prevent asymptotic bias, and this requires choice of an auxiliary bandwidth (or series length in the case of series estimation). There are no satisfactory data-based

methods for choosing the auxiliary bandwidth or series length. Hall and Horowitz (2013) provide numerical illustrations of this problem. Calonico, Cattaneo, and Farrell (2014) present an alternative form of undersmoothing that does not require an auxiliary bandwidth. This method has some desirable theoretical properties but is more complex than the one used here.

The following notation is used to define the unrestricted estimator of $g(x)$ and in the remainder of this paper. Let K denote a probability density function that is supported on $[-1,1]$ and symmetrical about 0. For any $v \in [-1,1]$, let $K_h(v) = K(v/h)$. Let $U = Y - g(X)$ and $\sigma_U^2 = \text{Var}(U)$. For any $x \in [0,1]$, let N_x denote the interval $[x-h, x+h]$. Define

$$n_x = \sum_{i=1}^n I(|X_i - x| \leq h).$$

Let i_1, \dots, i_{n_x} index observations $X_i \in N_x$. For $X_{i_1}, \dots, X_{i_{n_x}} \in N_x$, define the $n_x \times 3$ matrix

$$\mathbf{X}^{(x)} = \begin{bmatrix} 1 & (X_{i_1} - x) & (X_{i_1} - x)^2 \\ & \dots & \\ 1 & (X_{i_{n_x}} - x) & (X_{i_{n_x}} - x)^2 \end{bmatrix},$$

the $n_x \times n_x$ diagonal matrix

$$\mathbf{W}^{(x)} = \text{diag}[K_h(X_i - x) : X_i \in N_x],$$

and the 3×3 matrix

$$\mathbf{S}_n^{(x)} = n_x^{-1} \mathbf{X}^{(x)'} \mathbf{W}^{(x)} \mathbf{X}^{(x)}.$$

Also, for $X_{i_1}, \dots, X_{i_{n_x}} \in N_x$, define the $n_x \times 1$ vectors $\mathbf{Y}^{(x)} = (Y_{i_1}, \dots, Y_{i_{n_x}})'$, $\mathbf{U}^{(x)} = (U_{i_1}, \dots, U_{i_{n_x}})'$, and $\mathbf{g}^{(x)} = [g(X_{i_1}), \dots, g(X_{i_{n_x}})]'$.

Now let $\mathbf{b} = (b_1, b_2, b_3)'$ be a 3×1 vector, and let

$$\begin{aligned} \tilde{\mathbf{b}}(x) &\equiv [\tilde{b}_1(x), \tilde{b}_2(x), \tilde{b}_3(x)]' = \arg \min_{\mathbf{b}} (\mathbf{Y}^{(x)} - \mathbf{X}^{(x)} \mathbf{b})' \mathbf{W}^{(x)} (\mathbf{Y}^{(x)} - \mathbf{X}^{(x)} \mathbf{b}) \\ &= \arg \min_{\mathbf{b}} (\mathbf{b}' \mathbf{X}^{(x)'} \mathbf{W}^{(x)} \mathbf{X}^{(x)} \mathbf{b} - 2 \mathbf{Y}^{(x)'} \mathbf{W}^{(x)} \mathbf{X}^{(x)} \mathbf{b}). \end{aligned}$$

The unrestricted estimator of $g(x)$ is $\tilde{g}(x) = \tilde{b}_1(x)$. Standard algebra of least squares estimation shows that

$$\tilde{g}(x) - g(x) = \mathbf{e}_1' (\mathbf{S}_n^{(x)})^{-1} n_x^{-1} \mathbf{X}^{(x)'} \mathbf{W}^{(x)} \mathbf{U}^{(x)} + \mathbf{e}_1' (\mathbf{S}_n^{(x)})^{-1} n_x^{-1} \mathbf{X}^{(x)'} \mathbf{W}^{(x)} \mathbf{g}^{(x)} - g(x),$$

where $\mathbf{e}_1 = (1, 0, 0)'$.

Now make the following assumptions:

Assumption 1: $\{Y_i, X_i : i = 1, \dots, n\}$ is an independent random sample from the distribution of (Y, X) , where (i) $\text{supp}(X) = [0, 1]$. (ii) $Y = g(X) + U$, $E(U | X = x) = 0$ and $E(U^2 | X = x) = \sigma_U^2$ (a finite constant) for every $x \in [0, 1]$. (iii) There is an operator A such that $(Ag)(x) \leq 0$ for every $x \in [0, 1]$. (iv) $E(|U|^3) < \infty$.

Assumption 2: (i) $g(x)$ is four times continuously differentiable at each $x \in [0, 1]$. (ii) The distribution of X has a probability density function with respect to Lebesgue measure, f_X , that is continuously differentiable everywhere in $[0, 1]$. (iii) $f_X(x) \geq \delta$ for some $\delta > 0$ and every $x \in [0, 1]$.

Assumption 3: (i) K is a bounded probability density function that is supported on $[-1, 1]$ and symmetrical about 0; (ii) $h = cn^{-1/5}$ for some finite $c > 0$.

Except for Assumptions 1(iii) and 1(iv), these are standard assumptions in local polynomial nonparametric estimation. Assumption 1(iii) ensures that the shape restricted model is not misspecified. Section 3.2 provides further information about the operator A . Assumption 1(iv) used in Section 3 to ensure that $P(\hat{\mathcal{S}}_n = \bar{\mathcal{S}}_n) \rightarrow 1$, where $\bar{\mathcal{S}}_n$ is the unknown set of constraints that bind or nearly bind in the population and $\hat{\mathcal{S}}_n$ is the data-based set of possibly binding or nearly binding constraints. Assumption 1(iv) is also used in Section 4 to obtain the asymptotic distribution of the constrained estimator of g . Assumption 1(ii) requires U to be homoscedastic. This assumption can be removed at the cost of a more complex estimation procedure than the one presented here. The extension to a heteroscedastic U is outlined at the end of this section. Assumptions 2 and 3 make the local quadratic estimator undersmoothed, as is necessary to avoid asymptotic bias in the estimator of g . The assumption that g has four continuous derivatives is stronger than needed to obtain the asymptotic distributional results presented in this paper. The results can be obtained under the assumption that g is twice continuously differentiable. However, this requires choosing an undersmoothing bandwidth or an auxiliary bandwidth for explicit bias correction. There are no satisfactory empirical methods for making these choices in applications. The method of Calonico, Cattaneo, and Farrell (2014) permits g to have three derivatives at the cost of greater complexity than the method used here.

The following proposition states the properties of $\tilde{g}(x)$ that are used in this paper.

Proposition 3.1: Let Assumptions 1(i), 1(ii), 2, and 3 hold. For each $x \in (0, 1)$

$$(3.1) \quad (nh)^{1/2}[\tilde{g}(x) - g(x)] = (nh)^{1/2} e_1' (\mathbf{S}_n^{(x)})^{-1} n_x^{-1} \mathbf{X}^{(x)'} \mathbf{W}^{(x)} \mathbf{U}^{(x)} + r_n$$

and

$$(3.2) \quad (nh)^{1/2} [\tilde{g}(x) - g(x)] \rightarrow^d N(0, \sigma_{\tilde{g}(x)}^2),$$

where $r_n \leq ch^2$ for all x and some constant $c < \infty$, and

$$\sigma_{\tilde{g}(x)}^2 = C_K \frac{\sigma_U^2}{f_X(x)},$$

where C_K is a constant that depends on K but not on x , g , or h . ■

Result (3.1) follows from Theorem 3.1 of Fan and Gijbels (1996). See, also, Ruppert and Wand (1994). Result (3.2) is obtained by applying the Lindeberg-Levy central limit theorem to the first term on the right-hand side of (3.1). Fan and Gijbels (1996, p. 62) give the formula for C_K . In applications, $f_X(x)$ and σ_U^2 can be replaced with consistent estimators to give the consistent estimator of $\sigma_{\tilde{g}(x)}^2$,

$$\hat{\sigma}_{\tilde{g}(x)}^2 = C_K \frac{\hat{\sigma}_U^2}{\hat{f}_X(x)},$$

where $\hat{f}_X(x)$ is a consistent estimator of $f_X(x)$ (e.g., a kernel nonparametric density estimator) and $\hat{\sigma}_U^2$ is a consistent estimator of σ_U^2 . When X is a scalar, σ_U^2 can be estimated by the method of Rice (1984); Gasser, Sroka, and Jennen-Steinmetz (1986); and Buckley, Eagleson, and Silverman (1988). To construct this estimator, let $X_{(1)} < X_{(2)} < \dots < X_{(n)}$ be the ordered sequence of X_i 's. and let $\{Y_{(i)}\}$ be the similarly ordered values of the Y_i 's. The estimator of σ_U^2 is

$$\hat{\sigma}_U^2 = \frac{1}{2(n-1)} \sum_{i=1}^{n-1} [Y_{(i+1)} - Y_{(i)}]^2.$$

This estimator is $n^{-1/2}$ -consistent under Assumptions 1 and 2.

The foregoing discussion assumes that U in Assumption 1(ii) is homoscedastic. If U is heteroskedastic, then the constant σ_U^2 in Proposition 1 must be replaced with the function $\sigma_U^2(x) = E(U^2 | X = x)$ (Fan and Gijbels 1996, Theorem 3.1). The estimator $\hat{\sigma}_U^2$ must be replaced by an estimator of $\sigma_U^2(x)$ give the required expression for $\sigma_{\tilde{g}(x)}^2$. Horowitz and Spokoiny (2001) describe such an estimator. The results in the remainder of this paper hold under heteroscedasticity if σ_U^2 and $\hat{\sigma}_U^2$, respectively, are replaced by $\sigma_U^2(x)$ and an estimator of $\sigma_U^2(x)$.

3.2 Shape-Restricted Estimation and the Grid

The shape restriction on g is written $(Ag)(x) \leq 0$ for every $x \in [0,1]$, where A is an operator that maps functions satisfying Assumption 2(i) into a possibly different class of functions on $[0,1]$. For

example, if g is monotone non-increasing, then $Ag = dg/dx$. The shape restriction constitutes infinitely many constraints on g . We represent the shape restriction as a finite number of constraints by imposing it only at a grid of J equally spaced points $0 < x_1 < x_2 < \dots < x_J < 1$. J increases as n increases. Because the grid points are equally spaced, the distance between two consecutive grid points is $\Delta = 1/(J+1)$. We assume that:

Assumption 4: (i) $J \propto n^a (\log n)^{-1/4-b}$ for some $0 < a \leq 1/5$ and $b > 0$. (ii) $h < 1/[2(J+1)]$. (iii) $J \rightarrow \infty$ as $n \rightarrow \infty$.

Assumptions 4(i) and 3(ii) ensure that Assumption 4(ii) holds for all sufficiently large n . Assumption 4(ii) requires $h < 1/[2(J+1)]$ for any n . Under this assumption, for each $i = 1, \dots, n$, there is only one j such that $X_i \in N_{x_j}$. Therefore, $\tilde{g}(x_j)$ and $\tilde{g}(x_k)$ are statistically independent if $j \neq k$. Assumption 4(iii) ensures that the distance between grid points decreases as n increases.

Let $g_j = g(x_j)$, and let \mathbf{g} be the $J \times 1$ vector $\mathbf{g} = (g_1, \dots, g_J)'$. With this notation, a shape restriction that is imposed only at grid points can be written

$$A_k(\mathbf{g}) \leq 0; \quad k = 1, \dots, \kappa \leq J,$$

where the A_k 's are functions. For example, the restriction that g is non-increasing, can be represented as $g_{j+1} - g_j \leq 0$ for every $j = 1, \dots, J-1$. Thus, $A_k(\mathbf{g}) = g_{k+1} - g_k$, and $\kappa = J-1$. J and κ both increase as n increases. This dependence on n is not represented in the notation but is understood throughout this paper.

We impose shape restrictions on the grid by constraining differences between values of $g(x)$ at different values of x , not by constraining derivatives of g . This is because estimators of derivatives of g converge more slowly than the estimator of g . Consequently, the random sampling error of the constrained estimator of g is larger and the uniform confidence band for g wider if shape restrictions are imposed by constraining derivatives than if they are imposed by constraining differences.

Let $\mathcal{S} = \{k : A_k(\mathbf{g}) = 0\}$ denote the set of constraints that bind in the sampled population, and let $|\mathcal{S}|$ denote the number of elements in \mathcal{S} . Estimation of \mathbf{g} subject to $A_k(\mathbf{g}) \leq 0$ ($k = 1, \dots, \kappa$) is asymptotically equivalent to estimating \mathbf{g} subject to $A_k(\mathbf{g}) = 0$ for $k \in \mathcal{S}$. In other words, constraints that do not bind in the population can be dropped, and constraints that do bind can be replaced by equalities. In typical applications, the function $A_k(\mathbf{g})$ depends only on a few components of \mathbf{g} . For example, $A_k(\mathbf{g}) = g_{k+1} - g_k \leq 0$ represents the restriction that g is non-increasing, and $A_k(\mathbf{g})$ depends on only two components of \mathbf{g} . The restriction that g is convex can be represented as

$A_k(\mathbf{g}) = -g(x_k) + 2g(x_{k+1}) - g(x_{k+2}) \leq 0$, and $A_k(\mathbf{g})$ depends on only three components of \mathbf{g} . If $|\mathcal{S}| < \kappa$ (not all constraints are binding), then there may be some components of \mathbf{g} that do not affect the value of $A_k(\mathbf{g})$ for any $k \in \mathcal{S}$. That is, there may be some g_j 's for which $\partial A_k(\mathbf{g}) / \partial g_j = 0$ for every $k \in \mathcal{S}$. These g_j 's are unconstrained. They can be estimated and inference about them carried out using the unrestricted nonparametric method of Section 3.1. It is necessary to carry out constrained estimation and inference only for g_j 's satisfying $\partial A_k(\mathbf{g}) / \partial g_j \neq 0$ for some $k \in \mathcal{S}$. These components affect the value of $A_k(\mathbf{g})$ for some $k \in \mathcal{S}$ and, therefore, are constrained. Call these g_j 's "active components" of \mathbf{g} . Call the remaining components "inactive." Define

$$\mathcal{C} = \{j: g_j \text{ is an active component of } \mathbf{g}\},$$

and let $|\mathcal{C}|$ denote the number of elements of \mathcal{C} . Let $\mathbf{g}^{(a)}$ denote the $|\mathcal{C}| \times 1$ vector of active components of \mathbf{g} .

We use the following notation to define the shape-restricted local quadratic estimator of $\mathbf{g}^{(a)}$. Index the components of $\mathbf{g}^{(a)}$ and the grid points corresponding to them by $\ell = 1, \dots, |\mathcal{C}|$. Define the $n \times 3$ matrix

$$\mathbf{X}^{(\ell)} = \begin{bmatrix} 1 & (X_1 - x_\ell) & (X_1 - x_\ell)^2 \\ \dots & \dots & \dots \\ 1 & (X_n - x_\ell) & (X_n - x_\ell)^2 \end{bmatrix}; \ell = 1, \dots, |\mathcal{C}|.$$

Let $\mathbf{W}^{(\ell)}$ be the $n \times n$ diagonal matrix whose (i, i) component is $K_h(X_i - x_\ell)$, and let \mathbf{Q} be the $3|\mathcal{C}| \times 3|\mathcal{C}|$ block diagonal matrix

$$\mathbf{Q} = \begin{bmatrix} \mathbf{X}^{(1)'} \mathbf{W}^{(1)} \mathbf{X}^{(1)} & 0 & \dots & 0 \\ 0 & \mathbf{X}^{(2)'} \mathbf{W}^{(2)} \mathbf{X}^{(2)} & \dots & 0 \\ \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & \mathbf{X}^{(|\mathcal{C}|)'} \mathbf{W}^{(|\mathcal{C}|)} \mathbf{X}^{(|\mathcal{C}|)} \end{bmatrix}.$$

Finally, define the $3|\mathcal{C}| \times 1$ vector

$$\mathbf{d} = [(\mathbf{X}^{(1)'} \mathbf{W}^{(1)} \mathbf{Y})' \dots (\mathbf{X}^{(|\mathcal{C}|)'} \mathbf{W}^{(|\mathcal{C}|)} \mathbf{Y})']',$$

and let \mathbf{b} the $3|\mathcal{C}| \times 1$ vector $(b_{11}, b_{21}, b_{31}, \dots, b_{1|\mathcal{C}|}, b_{2|\mathcal{C}|}, b_{3|\mathcal{C}|})'$.

If \mathcal{S} were known, the shape restricted local quadratic estimator of g_ℓ ($\ell \in \mathcal{C}$) would be $\tilde{b}_{1\ell}$, where $\tilde{b}_{1\ell}$ is the $(3\ell - 2)$ component of the $3|\mathcal{C}| \times 1$ vector

$$\tilde{\mathbf{b}} = \arg \min_{\mathbf{b}} (0.5\mathbf{b}'\mathbf{Q}\mathbf{b} - \mathbf{d}'\mathbf{b})$$

subject to:

$$A_k(\mathbf{g}) = 0; \quad k \in \mathcal{S}.$$

However, \mathcal{S} is unknown in applications. Therefore, we replace it with the estimate $\hat{\mathcal{S}}_n$ that is described in Section 4. Redefine the active components of \mathbf{g} as the g_j 's satisfying $\partial A_k(\mathbf{g}) / \partial g_j \neq 0$ for some $k \in \hat{\mathcal{S}}_n$.

In the definitions of \mathbf{Q} and \mathbf{d} , replace \mathcal{C} with

$$\hat{\mathcal{C}} = \{j: g_j \text{ is a redefined active component of } \mathbf{g}\}$$

and replace $|\mathcal{C}|$ with $|\hat{\mathcal{C}}|$, which is the number of elements of $\hat{\mathcal{C}}$. We estimate an active component g_ℓ by $\hat{g}_\ell = \hat{b}_{1\ell}$, which is the $(3\ell - 2)$ component of the vector

$$(3.3) \quad \hat{\mathbf{b}} = \arg \min_{\mathbf{b}} (0.5\mathbf{b}'\mathbf{Q}\mathbf{b} - \mathbf{d}'\mathbf{b})$$

subject to:

$$A_k(\mathbf{g}) = 0; \quad k \in \hat{\mathcal{S}}_n.$$

The estimator of the vector of redefined active components $\mathbf{g}^{(a)}$ is $\hat{\mathbf{g}}^{(a)} = (\hat{b}_{11}, \hat{b}_{14}, \dots, \hat{b}_{3|\mathcal{C}|-2})'$.

In summary, we estimate the active components of \mathbf{g} by solving (3.3) and the remaining components by the unrestricted method of Section (3.1). Denote the resulting estimator of \mathbf{g} by $\hat{\mathbf{g}}$. Section 4 obtains the asymptotic distribution of $(nh)^{1/2}(\hat{\mathbf{g}} - \mathbf{g})$ and a uniform confidence band for \mathbf{g} .

4. THE SET OF POSSIBLY BINDING CONSTRAINTS

This section explains how to find the set $\hat{\mathcal{S}}_n$ of possibly binding constraints. Define the $J \times 1$ vectors $\mathbf{g}^{(J)} = [g(x_1), \dots, g(x_J)]'$ and $\tilde{\mathbf{g}}^{(J)} = [\tilde{g}(x_1), \dots, \tilde{g}(x_J)]'$. For each $k = 1, \dots, \kappa$, define the set

$$\mathcal{C}(k) = \{j: \partial A_k(\mathbf{g}) / \partial g_j = 0 \text{ for all } \mathbf{g}\}.$$

Let $|\mathcal{C}(k)|$ denote the number of components of $\mathcal{C}(k)$, and define $J_0 = \max_{k=1, \dots, \kappa} |\mathcal{C}(k)|$. Recall that

$A_k(\mathbf{g}^{(J)}) \leq 0$ for every k . For any finite constant $C > 0$ let \mathcal{M}_{Cn} be the set of constraints k for which $C^{-1}(nh)^{-1/2} \leq -A_k(\mathbf{g}^{(J)}) \leq C[(\log n) / (nh)]^{1/2}$. That is

$$\mathcal{M}_{Cn} = \{k: C^{-1}(nh)^{-1/2} \leq -A_k(\mathbf{g}^{(J)}) \leq C[(\log n) / (nh)]^{1/2}\}.$$

Let $|\mathcal{M}_{Cn}|$ be the number of constraints in \mathcal{M}_{Cn} . Define $\bar{\mathcal{M}}_n$ as the set of constraints that do not belong to \mathcal{M}_{Cn} for any finite $C > 0$. That is

$$\bar{\mathcal{M}}_n = \{k : -A_k(\mathbf{g}^{(J)}) < C^{-1}(nh)^{-1/2} \text{ or } -A_k(\mathbf{g}^{(J)}) > C[(\log n)/(nh)]^{1/2} \quad \forall C < \infty\}$$

Make the following assumption:

Assumption 5: (i) There is a finite constant q not depending on n such that $|\mathcal{C}(k)| \leq q$ for all $k = 1, \dots, \kappa$. (ii) The vector $\mathbf{g}^{(J)}$ is contained in a compact subset \mathcal{G} of \mathbb{R}^J . (iii) A_k is a twice continuously differentiable function of its arguments for each $k = 1, \dots, \kappa$. There is a constant $\mathcal{A} < \infty$ such that $|\partial A_k^2(\mathbf{g}^{(J)}) / \partial g_{j\ell}| \leq \mathcal{A}$ for all $\mathbf{g}^{(J)} \in \mathcal{G}$; $j, \ell = 1, \dots, J$; and $k = 1, \dots, \kappa$. (iv) $|\mathcal{M}_{C_n}| \rightarrow 0$ as $n \rightarrow \infty$ for every finite $C > 0$.

Assumption 5(i) is motivated by the observation that with typical shape restrictions, such as monotonicity, convexity, or the Slutsky inequality, $A_k(\mathbf{g})$ depends on only a few components of \mathbf{g} . Assumption 5(iv) holds for typical shape restrictions if Assumption 4 holds. Examples illustrating this are given in the appendix. Assumption 5(iv) permits some non-binding constraints to be “nearly binding” in the sense that $|A_k(\mathbf{g}^{(J)})| = o[(nh)^{-1/2}]$ but $|A_k(\mathbf{g}^{(J)})| > 0$ for one or more values of k .

The following theorems state properties of $\tilde{\mathbf{g}}^{(J)} - \mathbf{g}^{(J)}$ that are used to construct the set $\hat{\mathcal{S}}$ of possibly binding constraints.

Theorem 4.1: Let Assumptions 1-4 and 5(i)-5(iii) hold. Define the $J \times J$ diagonal matrix $\omega = \text{diag}[\sigma_{\tilde{g}(x_j)}^2 : j = 1, \dots, J]$. Let Υ be the $\kappa \times \kappa$ matrix whose (k, ℓ) component is

$$\Upsilon_{k\ell} = \left[\frac{\partial A_k(\mathbf{g}^{(J)})}{\partial \mathbf{g}} \right]' \omega \left[\frac{\partial A_\ell(\mathbf{g}^{(J)})}{\partial \mathbf{g}} \right].$$

Define the random variables $\mathbf{z} \sim N(0, \omega)$ and $\tilde{\mathbf{z}} \sim N(0, \Upsilon)$. Then

$$(4.2) \quad \limsup_{n \rightarrow \infty} \sup_{\mathbf{t}} |P[(nh)^{1/2}(\tilde{\mathbf{g}}^{(J)} - \mathbf{g}^{(J)}) \leq \mathbf{t}] - P(\mathbf{z} \leq \mathbf{t})| = 0$$

and

$$(4.3) \quad \limsup_{n \rightarrow \infty} \sup_{\mathbf{t}} |P\{(nh)^{1/2}[A_k(\tilde{\mathbf{g}}^{(J)}) - A_k(\mathbf{g}^{(J)})] \leq \mathbf{t}\} - P(\tilde{\mathbf{z}} \leq \mathbf{t})| = 0. \quad \blacksquare$$

Theorem 4.2: Let Assumptions 1-5 hold. Define $c_n = (\log n)^{1/2}$,

$$\bar{\mathcal{S}}_n = \left\{ k : -\Upsilon_{kk}^{-1/2} (nh)^{1/2} A_k(\mathbf{g}^{(J)}) \leq c_n \right\},$$

and

$$\tilde{\mathcal{S}}_n = \left\{ k : -\Upsilon_{kk}^{-1/2} (nh)^{1/2} A_k(\tilde{\mathbf{g}}^{(J)}) \leq c_n \right\}.$$

Then

$$(4.4) \quad \lim_{n \rightarrow \infty} P(k \in \tilde{\mathcal{S}}_n \text{ for any } k \in \mathcal{M}_{Cn} \text{ and finite } C > 0) = 0$$

and

$$(4.5) \quad \lim_{n \rightarrow \infty} P(\tilde{\mathcal{S}}_n = \bar{\mathcal{S}}_n) = 1. \quad \blacksquare$$

These results continue to hold if Υ_{kk} is replaced by the consistent estimator $\hat{\Upsilon}_{kk}$ obtained by replacing $\sigma_{\tilde{g}(x_j)}^2$ with $\hat{\sigma}_{\tilde{g}(x_j)}^2$ in ω and $\mathbf{g}^{(J)}$ with $\tilde{\mathbf{g}}^{(J)}$ in $\Upsilon_{k\ell}$. Accordingly, define

$$(4.6) \quad \hat{\mathcal{S}}_n = \left\{ k : -\hat{\Upsilon}_{kk}^{-1/2} (nh)^{1/2} A_k(\tilde{\mathbf{g}}^{(J)}) \leq c_n \right\}.$$

Then Theorem 4.2 implies

Corollary 4.3:

$$\lim_{n \rightarrow \infty} P(\hat{\mathcal{S}}_n = \bar{\mathcal{S}}_n) = 1. \quad \blacksquare$$

$\hat{\mathcal{S}}_n$ can be calculated from the data and is the desired set of possibly binding constraints. Asymptotically, $\hat{\mathcal{S}}_n$ contains constraints that are nearly binding in the sense defined in the paragraph following Assumption 5 as well as constraints that are truly binding. This property of $\hat{\mathcal{S}}_n$ enables the asymptotic distributional results presented in Section 5 to hold uniformly over a class of functions for which some constraints may be nearly binding.

5. ASYMPTOTIC DISTRIBUTION OF THE CONSTRAINED ESTIMATOR AND UNIFORM CONFIDENCE BAND

This section shows that $(nh)^{1/2}(\hat{\mathbf{g}} - \mathbf{g})$ is asymptotically multivariate normally distributed with mean 0. The distributional result is used to obtain an asymptotic uniform confidence band for $\mathbf{g}^{(J)}$ on the grid and, therefore, for $g(x)$ because the grid becomes dense as $n \rightarrow \infty$. Because $P(\hat{\mathcal{S}}_n = \bar{\mathcal{S}}_n) \rightarrow 1$ as $n \rightarrow \infty$, the asymptotic distribution of $(nh)^{1/2}(\hat{\mathbf{g}} - \mathbf{g})$ is unchanged by replacing the random set $\hat{\mathcal{S}}_n$ in (3.3) with the non-stochastic set $\bar{\mathcal{S}}_n$. Therefore, it suffices to derive the asymptotic distribution of $(nh)^{1/2}(\hat{\mathbf{g}} - \mathbf{g})$ and confidence band for $g(x)$ under the assumption that $\bar{\mathcal{S}}_n$ is known. Accordingly, it is assumed throughout this section that $\bar{\mathcal{S}}_n$ is known. The estimator $\hat{\mathbf{g}}$ is obtained by solving problem (3.3) with $\bar{\mathcal{S}}_n$ in place of $\hat{\mathcal{S}}_n$. Solving (3.3) with the constraint set $\bar{\mathcal{S}}_n$ causes any nearly binding constraints to be treated as binding. The results of this section show that the consequent biases are asymptotically negligible.

Section 5.1 treats the case of linear constraint functions A_k . Section 5.2 extends the results of section 5.1 to nonlinear constraints.

5.1 Linear Constraints

If the functions A_k are linear, the inequalities $A_k(\mathbf{g}) \leq 0$ ($k = 1, \dots, \kappa$) can be written as $\mathbf{A}\mathbf{g} \leq \mathbf{r}$, where \mathbf{A} is a $\kappa \times J$ matrix and \mathbf{r} is a $\kappa \times 1$ vector. For example, if g is monotone non-increasing, then $A_k(\mathbf{g}) = g(x_{k+1}) - g(x_k)$, $\mathbf{r} = \mathbf{0}$, and \mathbf{A} is the $(J-1) \times J$ matrix

$$\mathbf{A} = \begin{pmatrix} -1 & 1 & 0 & 0 & \dots & 0 & 0 \\ 0 & -1 & 1 & 0 & \dots & 0 & 0 \\ & & & \dots & & & \\ 0 & 0 & 0 & 0 & \dots & -1 & 1 \end{pmatrix}.$$

It follows from (3.2) and Assumption 4 that the estimators of the inactive components of $(nh)^{1/2}(\hat{\mathbf{g}} - \mathbf{g})$ are asymptotically normally distributed with means of 0 independently of each other and of the active components. Therefore, asymptotic distribution of $(nh)^{1/2}(\hat{\mathbf{g}} - \mathbf{g})$ will be known after the asymptotic distribution of its active components has been obtained. It follows from Corollary 4.3 and the discussion in the first paragraph of Section 5 that asymptotically, the active components include nearly binding constraints that are treated as if they were binding as well as constraints that are truly binding.

Let $\tilde{\kappa} \leq \kappa$ the number of rows of \mathbf{A} corresponding to constraints affecting active components of \mathbf{g} . Let $\mathbf{A}^{(a)}$ be the $\tilde{\kappa} \times |\mathcal{C}|$ matrix consisting of the columns of \mathbf{A} corresponding to active components of \mathbf{g} . The constraints on these components can be written as $\mathbf{A}^{(a)}\mathbf{g}^{(a)} = \mathbf{r}^{(a)}$, where $\mathbf{r}^{(a)}$ is a $\tilde{\kappa} \times 1$ vector. Let $\tilde{\mathbf{A}}$ denote the $\tilde{\kappa} \times 3|\mathcal{C}|$ matrix in which column $3\ell - 2$ ($\ell = 1, \dots, |\mathcal{C}|$) is column ℓ of $\mathbf{A}^{(a)}$ and the remaining columns are all zeros. Because it can be assumed as $n \rightarrow \infty$ that \mathcal{S} is known, problem (3.3) with linear constraints can be rewritten as

$$(5.1) \quad \hat{\mathbf{b}} = \arg \min_{\mathbf{b}} (0.5\mathbf{b}'\mathbf{Q}\mathbf{b} - \mathbf{d}'\mathbf{b})$$

subject to:

$$\tilde{\mathbf{A}}\mathbf{b} = \mathbf{r}^{(a)}.$$

Note that $\tilde{\mathbf{A}}$ and \mathbf{r} are non-stochastic. Problem (5.1) can be solved analytically using the method of Lagrangian multipliers. The solution is the well-known constrained least squares estimator

$$\hat{\mathbf{b}} = \tilde{\mathbf{b}}^{(a)} - \mathbf{Q}^{-1}\tilde{\mathbf{A}}'(\tilde{\mathbf{A}}\mathbf{Q}^{-1}\tilde{\mathbf{A}}')^+(\tilde{\mathbf{A}}\tilde{\mathbf{b}} - \mathbf{r}^{(a)}),$$

where $\tilde{\mathbf{b}}^{(a)}$ is the subvector of the unconstrained local quadratic estimator of Section (3.1) corresponding to active components of \mathbf{g} and the superscript $+$ denotes the Moore-Penrose generalized inverse. Only columns 1, 4, ..., $3|\mathcal{C}| - 2$ of the $\kappa \times 3|\mathcal{C}|$ matrix $\tilde{\mathbf{A}}$ are non-zero, and components 1, 4, ..., $3|\mathcal{C}| - 2$ of the

$3|\mathcal{C}| \times 1$ vector $\tilde{\mathbf{b}}$ correspond to components of $\tilde{\mathbf{g}}^{(a)}$, which is the local quadratic estimator of $\mathbf{g}^{(a)}$. The submatrix of $\tilde{\mathbf{A}}$ consisting of columns 1, 4, ..., $3|\mathcal{C}|-2$ is $\mathbf{A}^{(a)}$. Therefore, $\tilde{\mathbf{A}}\tilde{\mathbf{b}} = \mathbf{A}^{(a)}\tilde{\mathbf{g}}^{(a)}$,

$$\begin{aligned}\tilde{\mathbf{A}}\tilde{\mathbf{b}} - \mathbf{r}^{(a)} &= \mathbf{A}^{(a)}\tilde{\mathbf{g}}^{(a)} - \mathbf{r}^{(a)} \\ &= \mathbf{A}^{(a)}(\tilde{\mathbf{g}}^{(a)} - E\tilde{\mathbf{g}}^{(a)}) + \mathbf{A}^{(a)}(E\tilde{\mathbf{g}}^{(a)} - \mathbf{g}^{(a)}) + (\mathbf{A}^{(a)}\mathbf{g}^{(a)} - \mathbf{r}^{(a)}),\end{aligned}$$

and

$$\hat{\mathbf{b}} = \tilde{\mathbf{b}}^{(a)} - \mathbf{Q}^{-1}\tilde{\mathbf{A}}'(\tilde{\mathbf{A}}\mathbf{Q}^{-1}\tilde{\mathbf{A}}')^{-1}[\mathbf{A}^{(a)}(\tilde{\mathbf{g}}^{(a)} - E\tilde{\mathbf{g}}^{(a)}) + \mathbf{A}^{(a)}(E\tilde{\mathbf{g}}^{(a)} - \mathbf{g}^{(a)}) + (\mathbf{A}^{(a)}\mathbf{g}^{(a)} - \mathbf{r}^{(a)})].$$

Under Assumptions 1(iii) and 5(iv),

$$\mathbf{A}^{(a)}\mathbf{g}^{(a)} - \mathbf{r}^{(a)} = 0$$

for truly binding constraints,

$$\mathbf{A}^{(a)}\mathbf{g}^{(a)} - \mathbf{r}^{(a)} = o[(nh)^{-1/2}]$$

for nearly binding constraints, and

$$\mathbf{A}^{(a)}(E\tilde{\mathbf{g}}^{(a)} - \mathbf{g}^{(a)}) = O(h^4) = o[(nh)^{-1/2}]$$

for all J . Therefore,

$$\begin{aligned}\hat{\mathbf{g}}^{(a)} - \mathbf{g}^{(a)} &= (\tilde{\mathbf{g}}^{(a)} - \mathbf{g}^{(a)}) - [\mathbf{Q}^{-1}\tilde{\mathbf{A}}'(\tilde{\mathbf{A}}\mathbf{Q}^{-1}\tilde{\mathbf{A}}')^{-1}\mathbf{A}^{(a)}]^{(a)}(\tilde{\mathbf{g}}^{(a)} - E\tilde{\mathbf{g}}^{(a)}) + o[(nh)^{-1/2}] \\ &= \left\{ I_{|\mathcal{C}|\times|\mathcal{C}|} - [\mathbf{Q}^{-1}\tilde{\mathbf{A}}'(\tilde{\mathbf{A}}\mathbf{Q}^{-1}\tilde{\mathbf{A}}')^{-1}\mathbf{A}^{(a)}]^{(a)} \right\} (\tilde{\mathbf{g}}^{(a)} - E\tilde{\mathbf{g}}^{(a)}) + o[(nh)^{-1/2}]\end{aligned}$$

for all J , where $[\cdot]^{(a)}$ denotes the $|\mathcal{C}| \times |\mathcal{C}|$ submatrix of the $|\mathcal{C}| \times |\mathcal{C}|$ matrix $[\cdot]$ consisting of rows and columns 1, 4, ..., $3|\mathcal{C}|-2$. Define

$$\Xi^{(a)} = \left\{ I_{|\mathcal{C}|\times|\mathcal{C}|} - [\mathbf{Q}^{-1}\tilde{\mathbf{A}}'(\tilde{\mathbf{A}}\mathbf{Q}^{-1}\tilde{\mathbf{A}}')^{-1}\mathbf{A}^{(a)}]^{(a)} \right\}.$$

Then

$$(5.2) \quad \hat{\mathbf{g}}^{(a)} - \mathbf{g}^{(a)} = \Xi^{(a)}(\tilde{\mathbf{g}}^{(a)} - E\tilde{\mathbf{g}}^{(a)}) + o[(nh)^{-1/2}]$$

for all J .

Now rearrange the components of \mathbf{g} and $\hat{\mathbf{g}}$ into the vectors $[\mathbf{g}^{(-a)}, \mathbf{g}^{(a)}]'$ and $[\hat{\mathbf{g}}^{(-a)}, \hat{\mathbf{g}}^{(a)}]'$, respectively, where the superscript $(-a)$ denotes inactive components of \mathbf{g} and $\hat{\mathbf{g}}$. It follows from Theorem 4.1 that $(nh)^{1/2}(\tilde{\mathbf{g}}^{(a)} - E\tilde{\mathbf{g}}^{(a)})$ is asymptotically multivariate normally distributed with a diagonal covariance matrix. Therefore, (5.2) implies that $(nh)^{1/2}(\hat{\mathbf{g}}^{(a)} - \mathbf{g}^{(a)})$ and $(nh)^{1/2}(\hat{\mathbf{g}} - \mathbf{g})$ are linear combinations of asymptotic multivariate normals and are, themselves, asymptotically multivariate normal. Consequently, we have the following theorem.

Theorem 5.1: Let Assumptions 1-5 hold. Let Σ_n be the $J \times J$ matrix

$$\Sigma_n = \begin{bmatrix} \Sigma_n^{(-a)} & 0 \\ 0 & \Sigma_n^{(a)} \end{bmatrix},$$

$$\Sigma_n^{(-a)} = \text{diag}[\sigma_{\hat{g}(x_j)}^2 : g(x_j) \text{ is inactive}],$$

$$\Sigma_n^{(a)} = \Xi^{(a)} \Omega \Xi^{(a)'},$$

and

$$\Omega = \text{diag}[\sigma_{\hat{g}(x_j)}^2 : g(x_j) \text{ is active}].$$

Let \mathbf{z} be a random vector that is distributed as $N(0, \Sigma_n)$. If the constraints on g are linear, then

$$(5.3) \quad \lim_{n \rightarrow \infty} \sup_{g \in \bar{\mathcal{M}}_n} \sup_t |P[(nh)^{1/2}(\hat{\mathbf{g}} - \mathbf{g}) \leq \mathbf{t}] - P(\mathbf{z} \leq \mathbf{t})| = 0 \quad . \quad \blacksquare$$

The asymptotic normality property stated in (5.3) holds uniformly over functions $g \in \bar{\mathcal{M}}_n$. That is, it holds uniformly over a class of functions in which some constraints may be nearly binding. This class does not include functions $g \in \mathcal{M}_{C_n}$ for any finite C . Thus, there is a ‘‘gap’’ in the class of functions for which uniformity holds. A similar gap arises in penalized least squares or maximum likelihood estimation of high-dimensional models. See for example, Bühlmann and van de Geer (2011) and Horowitz and Huang (2013). Results of Leeb and Pötscher (2005, 2006) indicate that removing the gap for the large class of shape restrictions treated in this paper would be difficult or impossible. The methods for inference under monotonicity or convexity cited in Section 1 do not establish uniformity. It is not known if a computationally tractable method can be developed for achieving uniformity without a gap like $g \in \mathcal{M}_{C_n}$ for the general class of shape restrictions treated in this paper.

In the econometrics literature on moment inequalities, interest often centers on inference about a partially identified finite-dimensional parameter. As in this paper, it is not known which inequality constraints are or are not equalities. In moment inequalities, a confidence region for the parameter that has the correct asymptotic coverage probability uniformly over a class of distributions that includes nearly binding inequality constraints can be achieved without a gap by inverting a statistic for testing the hypothesis that the inequalities hold for a given parameter value. See, for example, Andrews and Soares (2010). In principle, a confidence region for a conditional mean function might be found by inverting a statistic such as that of Lee, Song, and Whang (2013, 2014) for testing the hypothesis that a shape restriction holds. However, a computationally tractable method for finding the resulting confidence region even for a finite-dimensional parameter has not been found.

Another possibility is to relax the shape restrictions slightly so that no shape constraint is binding in the sampled population. We do not adopt this approach here because it would allow the shape restrictions to be violated in finite samples. Depending on the application, such violations of the shape restrictions could have undesirable consequences. For example, violation of the Slutsky restriction could cause estimates of deadweight losses and other welfare measures to have incorrect signs.

Result (5.3) continues to hold if $\sigma_{\hat{g}(x_j)}^2$ is replaced in Σ_n with the estimator $\hat{\sigma}_{\hat{g}(x_j)}^2$ described in Section 3.1. Denote the resulting matrix by $\hat{\Sigma}_n$. Then it follows from (5.3) that a two-sided, asymptotic $1-\alpha$ confidence interval for $g(x_j)$ is

$$(5.4) \quad \hat{g}_j - (nh)^{-1/2} \hat{\Sigma}_{n,jj}^{1/2} z_{1-\alpha/2} \leq g_j \leq \hat{g}_j + (nh)^{-1/2} \hat{\Sigma}_{n,jj}^{1/2} z_{1-\alpha/2},$$

where \hat{g}_j is the j 'th component of $\hat{\mathbf{g}}$ and $z_{1-\alpha/2}$ is the $1-\alpha/2$ quantile of the standard normal distribution. It follows from the properties of constrained least squares estimators that $\Sigma_n - \text{cov}(\tilde{\mathbf{g}}^{(J)})$ is positive semidefinite. Therefore, the confidence interval (5.4) is no wider and may be narrower than a confidence interval for g_j based on the unconstrained estimator \tilde{g}_j .

An asymptotic $1-\alpha$ uniform confidence band for $g(x_j)$ ($j=1, \dots, J$) is

$$\mathcal{I}_{1-\alpha} = \{g_1, \dots, g_J : -\gamma_{j2} \leq \hat{g}_j - g_j \leq \gamma_{j1}; j=1, \dots, J\},$$

where γ_{j1} and γ_{j2} are critical values. The coverage probability is

$$P \left\{ \bigcap_{j=1}^J [-\gamma_{j2} \leq \hat{g}_j - g_j \leq \gamma_{j1}] \right\}$$

or, equivalently,

$$(5.5) \quad P \left\{ \bigcap_{j=1}^J [\hat{g}_j - \gamma_{j1} \leq g_j \leq \hat{g}_j + \gamma_{j2}] \right\}.$$

An asymptotic coverage probability of $1-\alpha$ can be obtained by choosing the γ_{j1} 's and γ_{j2} 's so that

$$(5.6) \quad P \left\{ \bigcap_{j=1}^J \{-\gamma_{j2} \leq Z_j \leq \gamma_{j1}\} \right\} = 1-\alpha,$$

where $Z = (Z_1, \dots, Z_J)'$ is a random variable with the $N[0, (nh)^{-1} \hat{\Sigma}_n]$ distribution. For a symmetrical confidence band, $\gamma_{j1} = \gamma_{j2} = \gamma_j$, where $\gamma_j > 0$, and

$$(5.7) \quad P\left\{\bigcap_{j=1}^J \{|Z_j| \leq \gamma_j\}\right\} = 1 - \alpha.$$

The probability in (5.7) can be estimated by Monte Carlo for any γ_j 's by drawing random samples from the $N[0, (nh)^{-1}\hat{\Sigma}_n]$ distribution.

A symmetrical confidence band with minimum average width and asymptotic coverage probability $1 - \alpha$ is

$$\mathcal{I}_{1-\alpha} = \{g_1, \dots, g_J : |\hat{g}_j - g_j| \leq \gamma_j; j = 1, \dots, J\},$$

where the γ_j 's satisfy

$$(5.8) \quad \text{minimize: } \sum_{j=1}^J \gamma_j$$

$\gamma_1, \dots, \gamma_J$

subject to

$$P\left\{\bigcap_{j=1}^J \{-\gamma_j \leq Z_j \leq \gamma_j\}\right\} = 1 - \alpha$$

$$\gamma_j \geq 0; \quad j = 1, \dots, J,$$

and $(Z_1, \dots, Z_J)'$ is a random variable with the $N[0, (nh)^{-1}\hat{\Sigma}_n]$ distribution. Denote the optimal solution to this problem by $(\gamma_1, \dots, \gamma_J)' \equiv \boldsymbol{\gamma}_{opt}$. The average width of the confidence band is $2\mathbf{e}'\boldsymbol{\gamma}_{opt} / J$, where \mathbf{e} is a column vector of ones.

An equivalent result can be obtained by Studentizing $(nh)^{1/2}(\hat{\mathbf{g}} - \mathbf{g})$. Let \hat{D}_n be the diagonal matrix whose (j, j) component is $\hat{\Sigma}_{jj}$. Define

$$V_n = \hat{D}_n^{-1/2} \hat{\Sigma}_n \hat{D}_n^{-1/2}.$$

Then the minimum-average-width symmetrical confidence band is

$$\mathcal{I}_{1-\alpha} = \{g_1, \dots, g_J : |\hat{g}_j - g_j| \leq (nh)^{-1/2} \hat{\Sigma}_{jj}^{1/2} \kappa_j; j = 1, \dots, J\},$$

where

$$\kappa_j = (nh)^{1/2} \hat{\Sigma}_{jj}^{-1/2} \gamma_j$$

and the γ_j 's solve (5.8).

An upper bound on the width of the minimum-average-width confidence band can be obtained by adding to (5.8) the constraint

$$\kappa_1 = \kappa_2 = \dots = \kappa_J.$$

Let κ denote the common value of $\kappa_1, \dots, \kappa_J$, and let $\tilde{Z} \equiv (\tilde{Z}_1, \dots, \tilde{Z}_J)'$ be a random vector with the $N(0, V_n)$ distribution. Then (5.8) becomes

$$(5.9) \quad \text{minimize: } \kappa$$

subject to

$$P \left\{ \bigcap_{j=1}^J \{-\kappa \leq \tilde{Z}_j \leq \kappa\} \right\} = 1 - \alpha$$

$$\kappa \geq 0.$$

Denote the optimal value of κ by κ_{opt} . The width of the resulting uniform confidence band is $w = 2(nh)^{-1/2} \Sigma_{jj}^{1/2} \kappa_{opt}$. Let the notation $a_n \asymp b_n$ for sequences of positive constants $\{a_n\}$ and $\{b_n\}$ mean that a_n / b_n is bounded away from zero and infinity. Then we have

Theorem 5.2: Let assumptions 1-5 hold. Then there is a sequence $\{d_n\}$ such that $w \leq d_n$ and $d_n \asymp [\log n / (nh)]^{1/2}$. ■

The width of a uniform confidence band for a conditional mean function without shape restrictions is asymptotic to $[\log n / (nh)]^{1/2}$ (Liero 1982). The shape restrictions treated in this paper do not increase the rate of convergence of an estimator. A uniform confidence band based on the shape restricted estimator is narrower than a band based on an unrestricted estimator, but the rate of convergence to zero of the band's width is the same with or without shape restrictions. Chernozhukov, Fernandez-Val, and Galichon (2009) present a related finite-sample result.

Except in special cases, the boundaries of the confidence band (5.8) do not satisfy the shape restrictions that are assumed to hold for \mathbf{g} . To obtain a confidence band of minimum average width whose boundaries satisfy the shape restrictions, define \mathbf{e} to be a $J \times 1$ vector of ones, $\boldsymbol{\gamma}_1 = (\gamma_{11}, \dots, \gamma_{J1})'$, and $\boldsymbol{\gamma}_2 = (\gamma_{12}, \dots, \gamma_{J2})'$. The uniform confidence band that has the minimum average width among bands satisfying the shape restrictions can be obtained by solving the nonlinear programming problem

$$(5.10) \quad (\boldsymbol{\gamma}_1, \boldsymbol{\gamma}_2)' = \arg \min_{\boldsymbol{\zeta}_1, \boldsymbol{\zeta}_2} \mathbf{e}'(\boldsymbol{\zeta}_2 - \boldsymbol{\zeta}_1)$$

subject to

$$P\left\{\bigcap_{j=1}^J\{-\zeta_{j2} \leq Z_j \leq \zeta_{j1}\}\right\} = 1 - \alpha$$

$$A(\hat{\mathbf{g}} - \zeta_1) \leq 0$$

$$A(\hat{\mathbf{g}} + \zeta_2) \leq 0.$$

The resulting confidence band is $\hat{\mathbf{g}} - \gamma_1 \leq \mathbf{g} \leq \hat{\mathbf{g}} + \gamma_2$. The minimum width symmetrical confidence band that satisfies the shape restrictions can be obtained by setting $\gamma_2 = -\gamma_1$ in problem (5.10).

The confidence bands described in the foregoing paragraphs are uniform only over grid points. However, the grid points $\{x_j\}$ become dense in $[0,1]$ as $n \rightarrow \infty$. Therefore, $\{\gamma_j : j=1, \dots, J\}$ and (5.6)-(5.10) provide an asymptotic uniform confidence region for $g(x)$. As is explained in Section 1, there is little practical difference between a band that is uniform over grid points and one that is uniform over a continuum.

5.2 Nonlinear Constraints

This section explains how to obtain an asymptotic uniform confidence band for g when one or more of the functions $A_k(\mathbf{g})$ specifying the shape constraints is nonlinear. As was explained in the introduction to Section 4, we assume that $\hat{\mathcal{S}}_n = \bar{\mathcal{S}}_n$ in deriving asymptotic uniform confidence band.

As in Section 3.1, let $\tilde{\mathbf{g}}$ denote the unconstrained local quadratic estimator of \mathbf{g} . For $k \in \mathcal{S}$, define the scalar η_k by $A_k(\tilde{\mathbf{g}}) = \eta_k$. Let $\mathbf{b}^{(1)}$ denote the intercept components of \mathbf{b} (that is, components 1, 4, ..., $3|\mathcal{C}|-2$). Then unrestricted estimates of the active components of \mathbf{g} , denoted by $\tilde{\mathbf{b}}^{(a)}$, can be obtained by solving

$$(5.11) \quad \tilde{\mathbf{b}}^{(a)} = \arg \min_{\mathbf{b}} (0.5\mathbf{b}'\mathbf{Q}\mathbf{b} - \mathbf{d}'\mathbf{b})$$

subject to:

$$A_k(\mathbf{b}^{(1)}) = \eta_k \quad (k \in \bar{\mathcal{S}}_n).$$

The constrained estimates are obtained by solving

$$(5.12) \quad \hat{\mathbf{b}} = \arg \min_{\mathbf{b}} (0.5\mathbf{b}'\mathbf{Q}\mathbf{b} - \mathbf{d}'\mathbf{b})$$

subject to:

$$A_k(\mathbf{b}^{(1)}) = 0 \quad (k \in \bar{\mathcal{S}}_n).$$

But $k \in \bar{\mathcal{S}}_n$ implies that $\eta_k \rightarrow 0$ as $n \rightarrow \infty$. Therefore, finding $\hat{\mathbf{g}}$ is equivalent to finding the effect of a small change in the η_k 's on the optimal solution to (5.11). This can be done by using the theory of sensitivity analysis in nonlinear programming (Fiacco 1983).

To state the result of the sensitivity analysis, modify the definition of $A_k(\cdot)$ to include all components of \mathbf{b} in its arguments. Let $|\bar{\mathcal{S}}_n|$ denote the number of constraints in $\bar{\mathcal{S}}_n$. Let $\mathbf{A}(\mathbf{b})$ be the $|\bar{\mathcal{S}}_n| \times 1$ vector whose k 'th component is $A_k(\mathbf{b})$. Let \mathbf{A}_b be the $|\bar{\mathcal{S}}_n| \times 3|\mathcal{C}|$ matrix whose (k, j) component is $\partial A_k(\mathbf{b}_0) / \partial b_j$, where \mathbf{b}_0 is the $3|\mathcal{C}| \times 1$ vector

$$\mathbf{b}_0 = \{[g(x_j), g'(x_j), g''(x_j)]' : j \in \mathcal{C}\}.$$

Let $[\cdot]^{(a)}$ denote the $|\mathcal{C}| \times |\mathcal{C}|$ submatrix of the $|3\mathcal{C}| \times |3\mathcal{C}|$ matrix $[\cdot]$ consisting of rows and columns 1, 4, ..., $3|\mathcal{C}|-2$. Define

$$\bar{\Xi}^{(a)} = \left\{ I_{|\mathcal{C}| \times |\mathcal{C}|} - [\mathbf{Q}^{-1} \mathbf{A}_b' (\mathbf{A}_b \mathbf{Q}^{-1} \mathbf{A}_b')^+ \mathbf{A}_b]^{(a)} \right\}.$$

The result is given by the following theorem.

Theorem 5.3: Let Assumptions 1-5 hold. Define $\bar{\Sigma}_n$ as the $J \times J$ matrix whose (j, k) element is

$$\bar{\Sigma}_{n, jk} = \begin{bmatrix} \bar{\Sigma}_n^{(-a)} & 0 \\ 0 & \bar{\Sigma}_n^{(a)} \end{bmatrix},$$

$$\bar{\Sigma}_n^{(-a)} = \text{diag}[\sigma_{\hat{g}(x_j)}^2 : g(x_j) \text{ is inactive}],$$

$$\bar{\Sigma}_n^{(a)} = \bar{\Xi}^{(a)} \Omega \bar{\Xi}^{(a)'},$$

and

$$\Omega = \text{diag}[\sigma_{\hat{g}(x_j)}^2 : g(x_j) \text{ is active}] .$$

Let \mathbf{z}_{NL} be a random vector that is distributed as $N(0, \bar{\Sigma}_n)$. Then

$$(5.13) \quad \lim_{n \rightarrow \infty} \sup_{g \in \mathcal{M}_n} \sup_t |P[(nh)^{1/2}(\hat{\mathbf{g}} - \mathbf{g}) \leq \mathbf{t}] - P(\mathbf{z}_{NL} \leq \mathbf{t})| = 0. \quad \blacksquare$$

An asymptotic uniform confidence region for \mathbf{g} and $g(x)$ can now be constructed as in Section 5.1 by replacing (5.3) with (5.13).

5.3 Finite-Sample Issues

Corollary 4.3 shows that $\hat{\mathcal{S}}_n = \bar{\mathcal{S}}_n$ with probability approaching 1 as $n \rightarrow \infty$. In a finite sample, however, it is possible that constraints k for which $A_k(\mathbf{g}) < 0$ are included in $\hat{\mathcal{S}}_n$ and, therefore,

erroneously treated as binding in forming the constrained estimator $\hat{\mathbf{g}}$. The resulting bias of $\hat{\mathbf{g}}$ can be large enough to cause the coverage probability of a confidence band for \mathbf{g} to be well below the nominal probability obtained from the asymptotic distribution of $(nh)^{1/2}(\hat{\mathbf{g}} - \mathbf{g})$.

The possibility of a large error in coverage probability can be reduced at the cost of a wider confidence band by reducing the size of $\hat{\mathcal{S}}_n$ and including a bias correction term in the asymptotic distribution of $(nh)^{1/2}(\hat{\mathbf{g}} - \mathbf{g})$. Specifically, redefine the set of possibly binding constraints as

$$\hat{\mathcal{S}}_n(c) = \left\{ k : -c(\log n)^{1/2} \leq \hat{\Upsilon}_{kk}^{-1/2} (nh)^{1/2} A_k(\tilde{\mathbf{g}}^{(J)}) \right\}$$

for some constant $c > 0$ and replace (5.3) and (5.13) by

$$\lim_{n \rightarrow \infty} \sup_{\mathbf{g} \in \mathcal{M}_n} \sup_t |P[(nh)^{1/2}(\hat{\mathbf{g}} - \mathbf{g}) \leq \mathbf{t}] - P(\mathbf{z} - \boldsymbol{\mu} \leq \mathbf{t})| = 0,$$

where $\boldsymbol{\mu}$ is the bias correction term described in the next two paragraphs. If c is small, the bias caused by erroneously treating a constraint as binding and the bias correction $\boldsymbol{\mu}$ are small. In addition, the estimator $\hat{\mathbf{g}}$ obtained with $\hat{\mathcal{S}}_n$ in place of $\hat{\mathcal{S}}_n$ satisfies the shape restrictions $A_k(\hat{\mathbf{g}}) \leq 0$ for all $k = 1, \dots, \kappa$. However, if c and $\boldsymbol{\mu}$ are small, the probability that $A_k(\tilde{\mathbf{g}}) \notin \hat{\mathcal{S}}_n$ may be high in a finite sample even if $A_k(\mathbf{g}) = 0$. If $A_k(\mathbf{g}) = 0$ but $A_k(\tilde{\mathbf{g}}) \notin \hat{\mathcal{S}}_n$, the confidence band for \mathbf{g} is wider than it would be if the constraint $A_k(\hat{\mathbf{g}}) = 0$ were imposed. The size of $\boldsymbol{\mu}$ can be chosen by the user, so there is a tradeoff between the size of $\boldsymbol{\mu}$ and the width of a confidence band.

We now describe a method for choosing $\boldsymbol{\mu}$. We focus on linear constraints with $\mathbf{r} = 0$. As is shown in Section 5.2, nonlinear constraints are asymptotically equivalent to properly constructed linear ones. Therefore, the following discussion also applies to nonlinear constraints. Define

$$\boldsymbol{\mu} = \begin{pmatrix} \boldsymbol{\theta}^{(-a)} \\ \boldsymbol{\mu}^{(a)} \end{pmatrix},$$

where $\boldsymbol{\theta}^{(-a)}$ is a column vector of 0's corresponding to inactive components of \mathbf{g} and

$$(5.14) \quad \boldsymbol{\mu}^{(a)} = [\mathbf{Q}^{-1} \tilde{\mathbf{A}}' (\tilde{\mathbf{A}} \mathbf{Q}^{-1} \tilde{\mathbf{A}}')^+]^{(a)} (\mathbf{A}^{(a)} \mathbf{g}^{(a)} - \mathbf{r}^a).$$

Let μ_j denote the j 'th component of $\boldsymbol{\mu}$. Then equation (5.7) becomes

$$P \left\{ \bigcap_{j=1}^J \{-\gamma_j \leq Z_j - \mu_j \leq \gamma_j\} \right\} = 1 - \alpha$$

or

$$(5.15) \quad P \left\{ \bigcap_{j=1}^J \{-\gamma_j + \mu_j \leq Z_j \leq \gamma_j + \mu_j\} \right\} = 1 - \alpha .$$

Because $\gamma_j > 0$ for all j , the probability on the left-hand side of (5.14) is a decreasing function of $|\mu_j|$ for any j . Define

$$m_j = \begin{cases} 0 & \text{if } j \text{ is an inactive component of } \mathbf{g} \\ m & \text{if } j \text{ is an active component of } \mathbf{g} \end{cases}$$

where $m \geq |\mu_j|$ for all active components j . Choose γ_{mj} ($j = 1, \dots, J$) so that

$$(5.16) \quad P \left\{ \bigcap_{j=1}^J \{-\gamma_{mj} + m_j \leq Z_j \leq \gamma_{mj} + m_j\} \right\} = 1 - \alpha .$$

Then

$$P \left\{ \bigcap_{j=1}^J \{-\gamma_{mj} + \mu_j \leq Z_j \leq \gamma_{mj} + \mu_j\} \right\} \geq 1 - \alpha .$$

Redefine the estimated set of active components of \mathbf{g} as

$$\hat{C}(c) = \{j : \partial A_k(\mathbf{g}) / \partial g_j \neq 0 \text{ for some } k \in \hat{S}_n(c)\} .$$

Note that $\hat{C}(c)$ depends on c .

The distribution of Z_j depends on which components of \mathbf{g} are active and, therefore, on the possibly identified set. Therefore, it is necessary to find the possibly identified set for which $|\mu_j| \leq m$.

This set is

$$\begin{aligned} \hat{S}_n(c) &= \left\{ k : -c(\log n)^{1/2} \leq \hat{Y}_{kk}^{-1/2} (nh)^{1/2} A_k \tilde{\mathbf{g}}^{(J)} \right\} \\ &= \left\{ k : -c \hat{Y}_{kk}^{1/2} [\log n / (nh)]^{1/2} \leq A_k \tilde{\mathbf{g}}^{(J)} \right\} \end{aligned}$$

for some $c > 0$. Because $\mathbf{r}^{(a)} = 0$ in (5.14), $\boldsymbol{\mu}^{(a)} = [\mathbf{Q}^{-1} \tilde{\mathbf{A}}' (\tilde{\mathbf{A}} \mathbf{Q}^{-1} \tilde{\mathbf{A}}')^+]^{(a)} \mathbf{A}^{(a)} \mathbf{g}^{(a)}$. Therefore, if j is active,

$$\begin{aligned} \mu_j &= \sum_{k \in \hat{S}_n(c)} \{ [\mathbf{Q}^{-1} \tilde{\mathbf{A}}' (\tilde{\mathbf{A}} \mathbf{Q}^{-1} \tilde{\mathbf{A}}')^+]^{(a)} \}_{jk} (\mathbf{A}^{(a)} \mathbf{g}^{(a)})_k \\ &= \sum_{k \in \hat{S}_n(c)} \{ [\mathbf{Q}^{-1} \tilde{\mathbf{A}}' (\tilde{\mathbf{A}} \mathbf{Q}^{-1} \tilde{\mathbf{A}}')^+]^{(a)} \}_{jk} A_k \mathbf{g}^{(J)} \end{aligned}$$

and

$$|\mu_j| \leq \sum_{k \in \hat{\mathcal{S}}_n(c)} | \{ [\mathbf{Q}^{-1} \tilde{\mathbf{A}}' (\tilde{\mathbf{A}} \mathbf{Q}^{-1} \tilde{\mathbf{A}}')^+]^{(a)} \}_{jk} \| A_k \mathbf{g}^{(J)} |.$$

$A_k \mathbf{g}^{(J)} \leq 0$ for all k . Therefore, with probability approaching 1 as $n \rightarrow \infty$,

$$-c \hat{\Upsilon}_{kk}^{1/2} [\log n / (nh)]^{1/2} \leq A_k \tilde{\mathbf{g}}^{(J)} - A_k \mathbf{g}^{(J)} \leq c \hat{\Upsilon}_{kk}^{1/2} [\log n / (nh)]^{1/2},$$

$$-A_k \mathbf{g}^{(J)} \leq -A_k \tilde{\mathbf{g}}^{(J)} + c \hat{\Upsilon}_{kk}^{1/2} [\log n / (nh)]^{1/2},$$

$$|A_k \mathbf{g}^{(J)}| \leq c \hat{\Upsilon}_{kk}^{1/2} [\log n / (nh)]^{1/2} - A_k \tilde{\mathbf{g}}^{(J)} \leq 2c \hat{\Upsilon}_{kk}^{1/2} [\log n / (nh)]^{1/2},$$

and

$$|\mu_j| \leq 2c [\log n / (nh)]^{1/2} \sum_{k \in \hat{\mathcal{S}}_n(c)} | \{ [\mathbf{Q}^{-1} \tilde{\mathbf{A}}' (\tilde{\mathbf{A}} \mathbf{Q}^{-1} \tilde{\mathbf{A}}')^+]^{(a)} \}_{jk} | \hat{\Upsilon}_{kk}^{1/2}.$$

Choose c so that

$$(5.17) \quad \max_{j \in \hat{\mathcal{C}}(c)} \left\{ 2c [\log n / (nh)]^{1/2} \sum_{k \in \hat{\mathcal{S}}_n(c)} | \{ [\mathbf{Q}^{-1} \tilde{\mathbf{A}}' (\tilde{\mathbf{A}} \mathbf{Q}^{-1} \tilde{\mathbf{A}}')^+]^{(a)} \}_{jk} | \hat{\Upsilon}_{kk}^{1/2} \right\} = m.$$

The following sequence of steps implement the bias-corrected confidence band.

1. Choose m .
2. Use m to choose $c > 0$ by solving (5.17).
3. Use c to specify $\hat{\mathcal{S}}_n$.
4. Use $\hat{\mathcal{S}}_n$ to obtain the asymptotic distribution of $(nh)^{1/2}(\hat{\mathbf{g}} - \mathbf{g})$ from Theorem 5.1 or 5.3.
5. Find confidence band by simulation from (5.16).

6. MULTIVARIATE EXTENSION

This section outlines the extension of the results of Sections 3-5 to a two-dimensional explanatory vector X , such as price and income in a demand function. Extensions to higher dimensions are possible but are less useful for economics and likely to yield low estimation precision because of the curse of dimensionality. The extension to a two-dimensional X involves mainly notational adjustments to the results of Sections 3-5. Therefore, the results of the extension are presented without proofs.

Denote the data by $\{Y_i, X_{i1}, X_{i2} : i = 1, \dots, n\}$. Assume that $\text{supp}(X_{i1}, X_{i2}) = [0, 1]^2$ for all $i = 1, \dots, n$. The following notation is used to state the unrestricted estimator of $g(x_1, x_2)$. Let K denote a probability density function that is supported on $[-1, 1]^2$ and whose odd moments are all zero. For any $\nu \equiv (\nu_1, \nu_2) \in [-1, 1]^2$ and bandwidths h_1 and h_2 , let

$$(6.1) \quad K_h(\nu) = K(\nu_1 / h_1, \nu_2 / h_2).$$

For any $x_1, x_2 \in [0, 1]^2$, let N_x denote the rectangle $\{\xi, \zeta : |\xi - x_1| \leq h_1, |\zeta - x_2| \leq h_2\}$. Define

$$n_x = \sum_{i=1}^n I(|X_{i1} - x_1| \leq h_1) I(|X_{i2} - x_2| \leq h_2).$$

For $i = i_1, i_2, \dots, i_{n_x}$ and $(X_{i_1}, X_{i_2}), \dots, (X_{i_{n_x}}, X_{i_{n_x}}) \in N_x$, define the 1×6 vector

$$\mathbf{X}_i^{(x)} = \left[1 \quad (X_{i1} - x_1) \quad (X_{i2} - x_2) \quad (X_{i1} - x_1)^2 \quad (X_{i2} - x_2)^2 \quad (X_{i1} - x_1)(X_{i2} - x_2) \right]$$

and the $n_x \times 6$ matrix

$$\mathbf{X}^{(x)} = \begin{pmatrix} \mathbf{X}_{i_1}^{(x)} \\ \dots \\ \mathbf{X}_{i_{n_x}}^{(x)} \end{pmatrix}.$$

For $X_i = (X_{i1}, X_{i2})$ and $x = (x_1, x_2)$, define the $n_x \times n_x$ diagonal matrix

$$\mathbf{W}^{(x)} = \text{diag}[K_h(X_i - x)]: X_i \in N_x]$$

and the 6×6 matrix

$$\mathbf{S}_n^{(x)} = n_x^{-1} \mathbf{X}^{(x)'} \mathbf{W}^{(x)} \mathbf{X}^{(x)}.$$

For $X_{i_1}, \dots, X_{i_{n_x}} \in N_x$, define the $n_x \times 1$ vector $\mathbf{Y}^{(x)} = (Y_{i_1}, \dots, Y_{i_{n_x}})'$ and $\mathbf{g}^{(x)} = [g(X_{i_1}), \dots, g(X_{i_{n_x}})]'$.

Now let $\mathbf{b} = (b_1, \dots, b_6)'$ be a 6×1 vector, and let

$$\begin{aligned} \tilde{\mathbf{b}}(x) &\equiv [\tilde{b}_1(x), \dots, \tilde{b}_6(x)]' = \arg \min_{\mathbf{b}} (\mathbf{Y}^{(x)} - \mathbf{X}^{(x)} \mathbf{b})' \mathbf{W}^{(x)} (\mathbf{Y}^{(x)} - \mathbf{X}^{(x)} \mathbf{b}) \\ &= \arg \min_{\mathbf{b}} (\mathbf{b}' \mathbf{X}^{(x)'} \mathbf{W}^{(x)} \mathbf{X}^{(x)} \mathbf{b} - 2 \mathbf{Y}^{(x)'} \mathbf{W}^{(x)} \mathbf{X}^{(x)} \mathbf{b}). \end{aligned}$$

The unrestricted estimator of $g(x)$ is $\tilde{g}(x) = \tilde{b}_1(x)$. Standard algebra of least squares estimation shows that

$$\tilde{\mathbf{b}}(x) = (\mathbf{S}_n^{(x)})^{-1} n_x^{-1} \mathbf{X}^{(x)'} \mathbf{W}^{(x)} \mathbf{Y}^{(x)}$$

Therefore,

$$\tilde{g}(x) - g(x) = e_1' (\mathbf{S}_n^{(x)})^{-1} n_x^{-1} \mathbf{X}^{(x)'} \mathbf{W}^{(x)} \mathbf{Y}^{(x)} - g(x),$$

where $e_1 = (1, 0, 0, 0, 0, 0)'$.

Now make the following assumptions, which are modifications of Assumptions 1-3:

Assumption 1': $\{Y_i, X_i : i = 1, \dots, n\}$ is an independent random sample from the distribution of (Y, X) , where (i) $\text{supp}(X) = [0, 1]^2$. (ii) $Y = g(X) + U$, $E(U | X = x) = 0$ and $E(U^2 | X = x) = \sigma_U^2$ (a finite

constant) for every $x \in [0,1]^2$. (iii) There is an operator A such that $(Ag)(x) \leq 0$ for every $x \in [0,1]^2$. (iv) $E(|U|^3) < \infty$.

Assumption 2': (i) $g(x)$ is four times continuously differentiable with respect to any combination of the components of $x \in [0,1]^2$. (ii) The distribution of X has a probability density function with respect to Lebesgue measure, f_X , that is continuously differentiable everywhere in $[0,1]^2$. (iii) $f_X(x) \geq \delta$ for some $\delta > 0$ and every $x \in [0,1]^2$.

Assumption 3': (i) K is a bounded probability density function that is supported on $[-1,1]^2$, and all odd moments of K are zero; (ii) $h_j = c_j n^{-1/6}$ ($j=1,2$) for some finite $c_j > 0$.

The bandwidths undersmooth \tilde{g} , so there is no asymptotic bias, and can be selected by applying cross-validation or plug-in methods to the local linear estimator of g . The following proposition generalizes Proposition 1 to the case of a bivariate X .

Proposition 1': Let Assumptions 1'(i), 1'(ii), 2', and 3' hold. For each $x \in (0,1)^2$,

$$(nh_1h_2)^{1/2}[\tilde{g}(x) - g(x)] \rightarrow^d N(0, \sigma_{\tilde{g}(x)}^2),$$

where $\sigma_{\tilde{g}(x)}^2 > 0$ is finite. ■

The expression for $\sigma_{\tilde{g}(x)}^2$ is lengthy and is given by Ruppert and Wand (1994, equation (4.7)). It requires a consistent estimator of σ_U^2 for a bivariate X . To do this, let $j(i)$ ($i = 1, \dots, n$) be a set of indices that is defined through the following recursion:

$$j(1) = \arg \min_{j=2, \dots, n} \|X_j - X_1\|$$

and

$$j(i) = \arg \min_{j \neq i, j(1), \dots, j(i-1)} \|X_j - X_i\|; \quad i = 2, \dots, n.$$

The number $j(i)$ is the index of the design point that is nearest to X_i among those whose indices are not $j(1), \dots, j(i-1)$. Then σ_U^2 can be estimated by

$$\hat{\sigma}_U^2 = \frac{1}{2n} \sum_{i=1}^n (Y_i - Y_{j(i)})^2.$$

Under Assumption 1', $\hat{\sigma}_U^2$ is a $n^{1/2}$ -consistent estimator of σ_U^2 (Horowitz and Spokoiny 2001).

The grid consists of J^2 equally spaced points $\{x_{jk} = (x_{1j}, x_{2k}) : j, k = 1, \dots, J\}$. The shape restriction is

$$A_k(\mathbf{g}) \leq 0; \quad k = 1, \dots, \kappa,$$

where \mathbf{g} is the $J^2 \times 1$ vector whose components are $g(x_{1j}, x_{2k})$ ($j, k = 1, \dots, J$). To obtain the bivariate extension of Theorem 5.1 replace Assumption 4 by

Assumption 4': (i) $J \propto n^a (\log n)^{-1/4-b}$ for some $0 < a \leq 1/6$ and $b > 0$. (ii) $h_1, h_2 < 1/[2(J+1)]$.
(iii) $J \rightarrow \infty$ as $n \rightarrow \infty$.

The bivariate extension of Theorem 5.1 is obtained from the scalar version by replacing Assumptions 1-4 with Assumptions 1'-4', J with J^2 , and $\sigma_{\tilde{g}(x_j)}^2$ with its bivariate extension, $\sigma_{\tilde{g}(x_{jk})}^2$.

To define the constrained estimator, index the grid points by $\ell \in \mathcal{C}$. For $\ell \in \mathcal{C}$ let

$$(6.2) \quad \mathbf{X}_i^{(\ell)} = \left[1 \quad (X_{i1} - x_{\ell 1}) \quad (X_{i2} - x_{\ell 2}) \quad (X_{i1} - x_{\ell 1})^2 \quad (X_{i2} - x_{\ell 2})^2 \quad (X_{i1} - x_{\ell 1})(X_{i2} - x_{\ell 2}) \right].$$

Define the matrix \mathbf{Q} and vector \mathbf{d} as in Section 2 but with K_h and $\mathbf{X}^{(\ell)}$ as in (6.1) and (6.2). Let \mathbf{b} be the $6|\mathcal{C}| \times 1$ vector $(b_{11}, b_{21}, \dots, b_{61}, \dots, b_{1|\mathcal{C}|}, b_{2|\mathcal{C}|}, \dots, b_{6|\mathcal{C}|})'$. Define

$$(6.3) \quad \hat{\mathbf{b}} = \arg \min_{\mathbf{b}} (0.5\mathbf{b}'\mathbf{Q}\mathbf{b} - \mathbf{d}'\mathbf{b})$$

subject to:

$$A_k(\mathbf{g}) = 0; \quad k \in \hat{\mathcal{S}}.$$

Problem (6.3) is the same as (3.3) but with K_h and $\mathbf{X}^{(\ell)}$ as in (6.1) and (6.2). Define active components of \mathbf{g} as in Section 3. The constrained estimator of the active components of \mathbf{g} is

$$\hat{\mathbf{g}}^{(a)} = (\hat{b}_{11}, \hat{b}_{71}, \dots, \hat{b}_{6|\mathcal{C}|-5})'.$$

To obtain the bivariate extensions of Theorems 5.1 and 5.3, redefine $|\mathcal{M}_{C_n}|$ as the number of constraints for which $C^{-1}(nh_1h_2)^{-1/2} \leq -A_k(\mathbf{g}^{(J)}) \leq C[(\log n)/(nh_1h_2)]^{1/2}$, where $\tilde{\mathbf{g}}^{(J)}$ is the $J^2 \times 1$ vector $[\tilde{g}(x_{11}), \dots, \tilde{g}(x_{JJ})]'$. The bivariate extensions of Theorems 5.1 and 5.2 are obtained from the scalar versions of these theorems by replacing Assumptions 1-4 with Assumptions 1'-4', using the redefined $|\mathcal{M}_{C_n}|$ in Assumption 5(iv), replacing J with J^2 , $(nh)^{1/2}$ with $(nh_1h_2)^{1/2}$, and the scalar version of $\sigma_{\tilde{g}(x_j)}^2$ with its bivariate extension. The appendix gives an example in which the redefined $|\mathcal{M}_{C_n}| \rightarrow 0$ in the bivariate case.

7. MONTE CARLO EXPERIMENTS AND AN EMPIRICAL EXAMPLE

This section presents the results of Monte Carlo experiments and an empirical example that illustrate the usefulness of the shape-restricted estimator described in Sections 3-6. The empirical example consists of estimating a production function under a shape constraint. The Monte Carlo experiments are designed to mimic the empirical example and illustrate the finite-sample performance of the uniform confidence band based on the shape-restricted estimator.

To describe the model used in the experiments and example, let Y , K , and L , respectively, denote value-added output, capital, and labor. Suppose that

$$(7.1) \quad \log Y = f(K, L) + U,$$

where U is an unobserved random variable that is independent of K and L and satisfies $E(U) = 0$.

Suppose that the function $\exp[f(K, L)]$ satisfies constant or decreasing returns to scale in levels. That is

$$(7.2) \quad \exp[f(\lambda K, \lambda L)] \leq \lambda \exp[f(K, L)]$$

for all $\lambda > 0$. It is customary to use the log transformation in empirical economics, and we follow that convention here. Taking logarithms on both sides of (7.2) yields

$$(7.3) \quad f(\lambda K, \lambda L) \leq \log \lambda + f(K, L)$$

for all $\lambda > 0$. Define $y = \log Y$, $k = \log K$, $\ell = \log L$, $g(k, \ell) = f(e^k, e^\ell)$, and $\tilde{\lambda} = \log \lambda$. Then (7.3) is equivalent to

$$(7.4) \quad g(\tilde{\lambda} + k, \tilde{\lambda} + \ell) \leq \tilde{\lambda} + g(k, \ell).$$

Model (7.1) is equivalent to

$$(7.5) \quad y = g(k, \ell) + U.$$

The Monte Carlo experiments and empirical example are based on (7.4) and (7.5).

7.1 Monte Carlo Experiments

This section presents the results of a small set of Monte Carlo experiments that illustrate the finite-sample performance of the uniform confidence band for $g(k, \ell)$ in (7.5) using the shape restricted estimator. In the experiments, samples of size $n = 1000$ and $n = 2000$ were generated from the production function

$$g(k, \ell) = \log \left(\left\{ [\exp(k)]^{1/2} + [\exp(\ell)]^{1/2} \right\}^{2\tau} \right)$$

for some constant $\tau > 0$. The resulting model is

$$(7.6) \quad y = \log \left(\left\{ [\exp(k)]^{1/2} + [\exp(\ell)]^{1/2} \right\}^{2\tau} \right) + U.$$

Model (7.6) is equivalent to the following production function model in levels:

$$Y = [(K^{1/2} + L^{1/2})^{2\tau}] e^U.$$

Values of k and ℓ in (7.6) were generated randomly and independently of each other from the $U[0,1]$ distribution. U was sampled independently of (k, ℓ) from the $N(0,0.01)$ distribution.

We report results for $\tau = 1$ (constant returns to scale), $\tau = 0.9$ (slightly decreasing returns to scale), and $\tau = 0.5$ (strongly decreasing returns to scale). In each experiment, the shape restriction is that g satisfies non-increasing returns to scale. Thus, $\tau = 1$ when the shape constraint is binding.

We used the grid

$$(7.7) \quad \left\{ (k_i, \ell_j) = \left(\frac{i}{J+1}, \frac{j}{J+1} \right) : i, j = 1, \dots, J \right\},$$

where J is chosen using the method described in the next paragraph. Using this grid, the discrete version of (7.4) is

$$g \left(\frac{i+1}{J+1}, \frac{j+1}{J+1} \right) \leq \frac{1}{J+1} + g \left(\frac{i}{J+1}, \frac{j}{J+1} \right)$$

for every $i, j = 1, \dots, J-1$.

We used the local quadratic estimator of g described in Section 3 with the uniform kernel function. A baseline bandwidth $h_0 = 0.15$ was determined by auxiliary simulations. Then we set $h = C_h h_0 (n/1000)^{-1/6}$, where $C_h = 0.95, 1$, or 1.05 , depending on the experiment. In practice, as is explained in Section 3, h can be chosen by cross-validation for local linear estimation. We did not use cross-validation in the Monte Carlo experiments because of its computational cost. The size of the grid, J , was $J = \lfloor J(h) \rfloor$, the largest integer not greater than $J(h)$, where

$$(7.8) \quad J(h) = \min \left[\frac{1}{h} (\log n)^{-1/2}, \frac{1}{2h} - 1.1 \right].$$

This choice of J satisfies Assumptions 4'(i) and 4'(ii).

There were 1,000 Monte Carlo replications in each experiment with 10,000 draws used to estimate the distribution of Z . When the estimated set of possibly binding constraints is nonempty, the limiting distribution of Z is degenerate. We used the singular value decomposition (SVD) to deal with singularity.

We present results for symmetrical nominal 95% uniform confidence bands for g obtained from the shape-restricted estimator. We also present results on uniform confidence bands using the unconstrained estimator (that is, the returns to scale constraint was not imposed), the infeasible constrained

(or oracle) estimator in which the true set of binding constraints \mathcal{S} is used in place of the estimated set $\hat{\mathcal{S}}$, and the bias corrected estimator of Section 5.3 with $m = 0.005$.

The results of the experiments are shown in Tables 1-3. Column 4 of the tables shows the empirical coverage probabilities of the uniform confidence bands. Column 5 shows the relative average widths of the various confidence bands. These are the ratios of the average widths of the bands to the average width of the band based on the unconstrained estimator. The relative width of the latter band is 1.0 by definition. The relative widths of the bands based on the constrained and oracle estimators are smaller (larger) than 1.0 according to whether the average widths of these bands are smaller (larger) than the average widths of the bands obtained from the unconstrained estimator. The results show that when $\tau = 1.0$ or $\tau = 0.5$, the empirical coverage probabilities of the bands obtained with all estimators are close to the nominal probability of 0.95. The confidence bands based on the constrained estimator are narrower than the bands based on the unconstrained estimator when the constraints are binding ($\tau = 1$) and have almost the same width when the constraints are not binding ($\tau = 0.5$). In this set of experiments, the constrained estimator performs as well as the oracle estimator. When $\tau = 1$, the relative width of the confidence band obtained from the bias-corrected estimator is larger than the relative widths of the bands obtained from the constrained and oracle estimators. As is explained in Section 5.3, this is because the bias-corrected estimator uses a set of possibly binding constraints $\hat{\mathcal{S}}_n$ that is smaller than the set $\hat{\mathcal{S}}_n$ used by the constrained estimator. Consequently, bias correction results in a higher probability that the constraints are not in the possibly binding set and not imposed in the second estimation step. Although the width of the confidence band obtained from the bias corrected estimator is closer to the width of the band obtained from the unconstrained estimator, the bias-corrected estimator satisfies the shape restrictions. The unconstrained estimator does not.

When $\tau = 0.9$, the constraints are nearly binding. Consequently, they have a relatively high probability of being in $\hat{\mathcal{S}}_n$ and erroneously treated as binding in the second estimation step. As is explained in Section 5.3, this causes the constrained estimator to be biased and the coverage probability of the resulting confidence band to be too low. The bias-corrected estimator largely overcomes this problem and yields coverage probabilities that are close to the nominal probability. The width of the bias-corrected confidence band is less than that of the oracle band because the constraints can be in $\hat{\mathcal{S}}_n$ when bias correction is used, in which case the bias-corrected estimator is constrained.

7.2 Empirical Example

This section reports the results of estimating a production function for the Chinese chemical industry using the firm-level data of Jacho-Chávez, Lewbel, and Linton (2010). We estimated the

production function using data for 1995 and 2001. The dependent variable, y , is the logarithm of value-added real output. The explanatory variables are the logarithm of the net value of real fixed assets, k , and the logarithm of the number of employees, ℓ . As in Jacho-Chávez, Lewbel, and Linton (2010), observations with outliers are removed and both regressors are normalized by their respective medians. As in the Monte Carlo experiments, we used the local quadratic estimator with a uniform kernel function. For each year, the bandwidth, h , was chosen by cross-validation. The grid points were chosen to be within the support of (k, ℓ) in the data. The number of points, J , was determined by (7.8). The sample sizes were $n = 1560$ for 1995 and $n = 1638$ for 2001. Increasing returns to scale are unlikely in the chemical industry. Accordingly, we carried out unconstrained estimation of g and estimation under the restriction of non-increasing returns to scale.

Table 4 and Figures 1-2 present the estimation results at several points (k, ℓ) for which the normalized values of k and ℓ are equal. The constrained, unconstrained, and bias-corrected point estimates are similar, as is to be expected in an industry that has non-increasing returns. However, the constrained estimates are more precise than the unconstrained ones. For example, in 1995 the constrained and unconstrained point estimates of $g(2.524, 2.524)$ are nearly the same, but the standard error of the constrained estimate is much less than that of the unconstrained estimate. It can be seen from Figure 1 that the constrained estimates are slightly more precise than the unconstrained ones in the middle of the distribution of (k, ℓ) and much more precise near the boundaries of the support of (k, ℓ) . As expected, the bias-corrected estimates are less precise than the constrained ones are. It can be seen from Figure 2 that the bias-corrected confidence band is similar to that of the unconstrained band. However, the bias-corrected estimate of g is constrained to satisfy the shape restriction, whereas the unconstrained estimate is not.

8. CONCLUSIONS

Economic theory often provides shape restrictions on functions of interest in applications, but it does not provide finite-dimensional parametric models. This motivates nonparametric estimation under shape restrictions. Shape restrictions can stabilize noisy nonparametric estimates without imposing arbitrary restrictions, such as additivity or a single-index structure, that may be inconsistent with economic theory and the data. This paper has explained how to estimate and obtain an asymptotic uniform confidence band for a conditional mean function under a possibly nonlinear shape restriction. There is a large literature in statistics and econometrics on estimating a conditional mean function under linear shape restrictions, such as monotonicity or convexity. To our knowledge, this paper is the first to construct a uniform confidence band under shape restrictions such as non-increasing or non-decreasing returns to scale and the Slutsky inequality of consumer theory. The results of Monte Carlo experiments and an empirical

application have illustrated the finite-sample performance and usefulness of our method. The methods of this paper can be extended to conditional quantile functions with shape restrictions, though doing so is complicated technically because of the non-differentiability of the objective function of quantile estimation. Estimation of conditional quantile functions under shape restrictions is a topic for further research.

9. PROOFS OF THEOREMS

9.1 *Examples in which $|\mathcal{M}_{Cn}| = o(1)$ as $n \rightarrow \infty$ for every $C < \infty$*

Scalar case: Let X be a scalar and

$$g(x) = \begin{cases} 0 & \text{if } 0 \leq x \leq 0.5 \\ (x-0.5)^2 & \text{if } 0.5 < x \leq 1. \end{cases}$$

Let there be J equally spaced grid points in $[0,1]$. Assume that J is odd so that $x_{(J+1)/2} = 0.5$ and $x_j > 0.5$ implies that $j > (J+1)/2$. The shape restriction is that g is non-decreasing, so $g(x_{j-1}) - g(x_j) \leq 0$. Let $x_j > x_{j-1} \geq 0.5$ be grid points. Then $j \geq (J+3)/2$, and

$$\begin{aligned} |g(x_j) - g(x_{j-1})| &= \frac{1}{(J+1)^2} [2j-1 - (J+1)] \\ &< \frac{2j}{(J+1)^2}. \end{aligned}$$

Therefore,

$$(9.1) \quad 0 < |g(x_j) - g(x_{j-1})| \leq \left(\frac{\log n}{nh} \right)^{1/2}$$

implies that

$$j \leq \frac{(J+1)^2}{2} \left(\frac{\log n}{nh} \right)^{1/2}$$

Under Assumptions 3(ii) and 4(i),

$$\frac{(J+1)^2}{2} \left(\frac{\log n}{nh} \right)^{1/2} \rightarrow 0$$

as $n \rightarrow \infty$. Therefore, there can be no grid points x_j and x_{j-1} satisfying (9.1) if n is sufficiently large, which implies that $|\mathcal{M}_{Cn}| = 0$ if n is sufficiently large.

Bivariate case: Let $X = (X_1, X_2)'$ be two dimensional. Let

$$g(x_1, x_2) = \begin{cases} 0 & \text{if } x_1^2 + x_2^2 \leq 0.25 \\ x_1^2 + x_2^2 - 0.25 & \text{if } x_1^2 + x_2^2 > 0.25 \end{cases}$$

Let there be J equally spaced grid points in each of the two dimensions. Assume that J is odd so that $x_{1,(J+1)/2} = x_{2,(J+1)/2} = 0.5$. Therefore, $x_{1,j} > 0.5$ or $x_{2,k} > 0.5$ implies that $j+k > J+1$. The shape restriction is

$$g(x_{1,j-1}, x_{2k}) - g(x_{1j}, x_{2k}) \leq 0$$

$$g(x_{1j}, x_{2,k-1}) - g(x_{1j}, x_{2k}) \leq 0.$$

This is the finite-difference analog of the restriction $\partial g(x_1, x_2) / \partial x_1 \geq 0$ and $\partial g(x_1, x_2) / \partial x_2 \geq 0$.

If $g(x_{1j}, x_{2k}) > 0$ and $j \geq 2$, then

$$g(x_{1j}, x_{2k}) - g(x_{1,j-1}, x_{2k}) = \frac{1}{(J+1)^2} (2j-1) \text{ if } g(x_{1,j-1}, x_{2k}) > 0$$

$$g(x_{1j}, x_{2k}) - g(x_{1,j-1}, x_{2k}) \leq \frac{1}{(J+1)^2} (2j-1) \text{ if } g(x_{1,j-1}, x_{2k}) = 0.$$

Therefore,

$$g(x_{1,j}, x_{2k}) - g(x_{1,j-1}, x_{2k}) \leq \frac{2j}{(J+1)^2}.$$

Similarly, if $g(x_{1j}, x_{2k}) > 0$ and $k \geq 2$, then

$$g(x_{1j}, x_{2k}) - g(x_{1j}, x_{2,k-1}) \leq \frac{2k}{(J+1)^2}.$$

Now proceed as in the example for a scalar X .

9.2 Proofs of Theorems 4.1, 4.2, and 5.1-5.3

Proof of Theorem 4.1: By (3.1),

$$(nh)^{1/2} [\tilde{g}(x_j) - g(x_j)] \rightarrow^d e_1'(nh/n_{x_j})(S_n^{(x_j)})^{-1}(nh)^{-1/2} \mathbf{X}^{(x_j)'} \mathbf{W}^{(x_j)} \mathbf{U}^{(x_j)}.$$

Define

$$\tilde{\mathbf{X}}^{(x_j)} = \begin{bmatrix} 1 & h^{-1}(X_{i_1} - x_j) & h^{-2}(X_{i_1} - x_j)^2 \\ & \dots & \\ 1 & h^{-1}(X_{i_{n_x}} - x_j) & h^{-2}(X_{i_{n_x}} - x_j)^2 \end{bmatrix}$$

and

$$\mathbf{e}_h = \begin{pmatrix} 1 & 0 & 0 \\ 0 & h & 0 \\ 0 & 0 & h^2 \end{pmatrix}.$$

Then,

$$(nh)^{1/2}[\tilde{g}(x_j) - g(x_j)] \rightarrow^d \mathbf{e}'_1(nh/n_{x_j})(\mathbf{S}_n^{(x_j)})^{-1} \mathbf{e}_h(nh)^{-1/2} \tilde{\mathbf{X}}^{(x_j)'} \mathbf{W}^{(x_j)} \mathbf{U}^{(x_j)}.$$

The 3×1 vector $(nh)^{-1/2} \tilde{\mathbf{X}}^{(x_j)'} \mathbf{W}^{(x_j)} \mathbf{U}^{(x_j)}$ is asymptotically trivariate normally distributed with mean 0 by the multivariate extension of the Lindeberg-Levy central limit theorem. Let $\tilde{\Sigma}_j$ denote the covariance matrix of the limiting distribution. Let \mathcal{B}^J denote the set of all convex sets in \mathbb{R}^J and Z_j be a random vector with the $N(0, \tilde{\Sigma}_j)$ distribution. It follows from Theorem 1.1 of Bentkus (2003) (see, also, Corollary 11.1 of DasGupta (2008)) that for some constant $c_1 < \infty$

$$(9.2) \quad \sup_{B \in \mathcal{B}^3} \left| P[(nh)^{-1/2} \tilde{\mathbf{X}}^{(x_j)'} \mathbf{W}^{(x_j)} \mathbf{U}^{(x_j)} \in B] - P(Z_j \in B) \right| \leq c_1(nh)^{-1/2}.$$

Let ξ denote the $3J \times 1$ vector whose $3j-2, \dots, 3j$ components are $(nh)^{-1/2} \tilde{\mathbf{X}}^{(x_j)'} \mathbf{W}^{(x_j)} \mathbf{U}^{(x_j)}$. Let \tilde{Z} denote a random vector with the $N(0, \tilde{\Sigma})$ distribution, where

$$\tilde{\Sigma} = \begin{pmatrix} \tilde{\Sigma}_1 & \dots & 0 \\ 0 & \dots & 0 \\ 0 & \dots & \tilde{\Sigma}_J \end{pmatrix}.$$

Then (9.2) implies that for some $c_2 < \infty$

$$(9.3) \quad \sup_{B \in \mathcal{B}^{3J}} \left| P(\xi \in B) - P(\tilde{Z} \in B) \right| \leq c_2 J(nh)^{-1/2}.$$

Let L_j denote the probability limit of $\mathbf{e}'_1(nh/n_{x_j})(\mathbf{S}_n^{(x_j)})^{-1}$ as $n \rightarrow \infty$. Standard calculations for kernel estimators show that

$$(9.4) \quad \mathbf{e}'_1(nh/n_{x_j})(\mathbf{S}_n^{(x_j)})^{-1} = L_j + O_p[(nh)^{-1/2}].$$

Result (4.2) follows by combining (9.3), (9.4), and Assumption 4.

To obtain (4.3), observe that under Assumption 5(iii), a Taylor series expansion gives

$$(9.5) \quad A_k(\tilde{\mathbf{g}}^{(J)}) - A_k(\mathbf{g}^{(J)}) = \left[\frac{\partial A_k(\mathbf{g}^{(J)})}{\partial \mathbf{g}} \right]' (\tilde{\mathbf{g}}^{(J)} - \mathbf{g}^{(J)}) + O_p\left(\|\tilde{\mathbf{g}}^{(J)} - \mathbf{g}^{(J)}\|^2\right)$$

$$= \left[\frac{\partial A_k(\mathbf{g}^{(J)})}{\partial \mathbf{g}} \right]' (\tilde{\mathbf{g}}^{(J)} - \mathbf{g}^{(J)}) + O_p[J / (nh)].$$

Result (4.3) follows by combining (9.3) and (9.5). Q.E.D.

Proof of Theorem 4.2: By the definition of c_n , $1 - \Phi(c_n) = O[(n \log n)^{-1/2}]$, where Φ is the standard normal distribution function. It follows from (4.3) that asymptotically,

$$P\left[\Upsilon_{kk}^{-1/2}(nh)^{1/2} |A_k(\tilde{\mathbf{g}}^{(J)}) - A_k(\mathbf{g}^{(J)})| > c_n\right] = O\left[\left(\sqrt{n \log n}\right)^{-1}\right]$$

and

$$(9.6) \quad P\left[\Upsilon_{kk}^{-1/2}(nh)^{1/2} |A_k(\tilde{\mathbf{g}}^{(J)}) - A_k(\mathbf{g}^{(J)})| > c_n \text{ for any } k = 1, \dots, \kappa\right] = O\left[\kappa \left(\sqrt{n \log n}\right)^{-1}\right] = o(1).$$

If $k \in \bar{\mathcal{M}}_n$, then either $A_k(\mathbf{g}^{(J)}) = o[(nh)^{-1/2}]$ or $|A_k(\mathbf{g}^{(J)})| \gg [\log n / (nh)]^{1/2}$. By (9.6)

$$(9.7) \quad P\{|A_k(\tilde{\mathbf{g}}^{(J)})| \leq \Upsilon_{kk}^{1/2}(nh)^{-1/2} c_n \text{ for every } k \text{ such that } |A_k(\mathbf{g}^{(J)})| = o[(nh)^{-1/2}]\}$$

$$= 1 - O\left[\kappa \left(\sqrt{n \log n}\right)^{-1}\right] \rightarrow 1.$$

In addition, by the triangle inequality,

$$\Upsilon_{kk}^{1/2} |A_k(\tilde{\mathbf{g}}^{(J)})| \geq \Upsilon_{kk}^{1/2} |A_k(\mathbf{g}^{(J)})| - \Upsilon_{kk}^{1/2} |A_k(\tilde{\mathbf{g}}^{(J)}) - A_k(\mathbf{g}^{(J)})|.$$

Therefore, if $|A_k(\mathbf{g}^{(J)})| \gg [(\log n) / (nh)]^{1/2}$ and $\Upsilon_{kk}^{-1/2}(nh)^{1/2} |A_k(\tilde{\mathbf{g}}^{(J)}) - A_k(\mathbf{g}^{(J)})| \leq c_n$, then for all sufficiently large n ,

$$\Upsilon_{kk}^{1/2} |A_k(\tilde{\mathbf{g}}^{(J)})| \geq \Upsilon_{kk}^{1/2} |A_k(\mathbf{g}^{(J)})| - (nh)^{-1/2} c_n > (nh)^{-1/2} c_n.$$

Consequently,

$$(9.8) \quad P\left[\Upsilon_{kk}^{-1/2}(nh)^{1/2} |A_k(\tilde{\mathbf{g}}^{(J)})| \leq c_n \text{ for any } k \text{ for which } |A_k(\mathbf{g}^{(J)})| \gg [(\log n) / (nh)]^{1/2}\right] = o(1).$$

$|\mathcal{M}_{C_n}|$ is an integer, and $|\mathcal{M}_{C_n}| \rightarrow 0$ as $n \rightarrow \infty$ for every $C < \infty$ by Assumption 5(iv). Therefore, $\Upsilon_{kk}^{1/2}(nh)^{1/2} |A_k(\mathbf{g}^{(J)})| > c_n$ or $\Upsilon_{kk}^{1/2}(nh)^{1/2} |A_k(\mathbf{g}^{(J)})| = o[(nh)^{-1/2}]$ for every k and all sufficiently large n . But $P[\Upsilon_{kk}^{1/2}(nh)^{1/2} |A_k(\tilde{\mathbf{g}}^{(J)})| \geq c_n] \rightarrow 0$ as $n \rightarrow \infty$. Combining these facts with (9.7) and (9.8) gives (4.4). Combining (4.4) with (9.7) and (9.8) gives (4.5). Q.E.D.

Proof of Theorem 5.1: By Theorem 4.1, $(nh)^{1/2}(\tilde{\mathbf{g}}^{(a)} - \hat{\mathbf{g}}^{(a)})$ and $(nh)^{1/2}(\hat{\mathbf{g}}^{(-a)} - \mathbf{g}^{(-a)})$ are asymptotically multivariate normal with means of zero. Therefore, $(nh)^{1/2}(\hat{\mathbf{g}} - \mathbf{g})$ is a linear combination of asymptotic multivariate normals with means of zero and is asymptotically multivariate normally distributed with a mean of 0. The covariance matrix Σ_n follows from (3.2), (5.2), and independence of the components of $\hat{\mathbf{g}} - \mathbf{g}$ from one another. Q.E.D.

Proof of Theorem 5.2: The optimal κ in (5.9) satisfies

$$P\left\{\bigcap_{j=1}^J\{-\kappa_{opt} \leq \tilde{Z}_j \leq \kappa_{opt}\}\right\} = 1 - \alpha.$$

Equivalently,

$$P\left\{\bigcup_{j=1}^J(|\tilde{Z}_j| \geq \kappa_{opt})\right\} = \alpha.$$

But $\tilde{Z}_j \sim N(0,1)$, so

$$\begin{aligned} P\left\{\bigcup_{j=1}^J(|\tilde{Z}_j| \geq \kappa)\right\} &\leq \sum_{j=1}^J P(|\tilde{Z}_j| \geq \kappa) \\ &= 2J[1 - \Phi(\kappa)] \end{aligned}$$

for any κ . Let $\tilde{\kappa}$ satisfy $2J[1 - \Phi(\tilde{\kappa})] = \alpha$. Then

$$P\left\{\bigcup_{j=1}^J(|\tilde{Z}_j| \geq \tilde{\kappa})\right\} \leq \alpha$$

and $\kappa_{opt} \leq \tilde{\kappa}$. But

$$\frac{2}{(2\pi)^{1/2} c\tilde{\kappa}} e^{-\tilde{\kappa}^2/2} \leq 2[1 - \Phi(\tilde{\kappa})] \leq \frac{2}{(2\pi)^{1/2} \tilde{\kappa}} e^{-\tilde{\kappa}^2/2}$$

and

$$\frac{2J}{(2\pi)^{1/2} c\tilde{\kappa}} e^{-\tilde{\kappa}^2/2} \leq 2J[1 - \Phi(\tilde{\kappa})] \leq \frac{2J}{(2\pi)^{1/2} \tilde{\kappa}} e^{-\tilde{\kappa}^2/2}$$

for any constant $c > 1$ when $\tilde{\kappa}$ is sufficiently large. Set

$$\frac{1}{2}\tilde{\kappa}^2 = \log J - \log[(2\log J)^{1/2}] + \log\left[\left(\frac{2}{\pi\alpha^2}\right)^{1/2}\right].$$

Then $\tilde{\kappa} \asymp (\log n)^{1/2}$,

$$\frac{2J}{(2\pi)^{1/2} \tilde{\kappa}} e^{-\tilde{\kappa}^2/2} = \frac{\alpha}{1+o(1)}$$

as $n \rightarrow \infty$, and

$$\frac{\alpha}{c[1+o(1)]} \leq 2J[1 - \Phi(\tilde{\kappa})] \leq \frac{\alpha}{1+o(1)}.$$

for any $c > 1$ as $n \rightarrow \infty$. The width of the confidence band for the Studentized $(\hat{\mathbf{g}} - \mathbf{g})$ is $2\kappa_{opt} \leq 2\tilde{\kappa}$.

Therefore, the width, w , of the minimum-average-width confidence band for \mathbf{g} without Studentization satisfies $w \leq 2(nh)^{-1/2} \sum_{jj}^{1/2} \tilde{\kappa} \asymp [\log n / (nh)]^{1/2}$. Q.E.D.

The following notation is used in the proof of Theorem 5.3. Define the $(3|\mathcal{C}| + |\mathcal{S}|) \times (3|\mathcal{C}| + |\mathcal{S}|)$ matrix

$$H = \begin{pmatrix} \mathbf{Q} & \mathbf{A}'_b \\ \mathbf{A}_b & 0_{|\mathcal{S}| \times |\mathcal{S}|} \end{pmatrix}$$

and the $(3|\mathcal{C}| + |\mathcal{S}|) \times |\mathcal{S}|$ matrix

$$V = \begin{pmatrix} \mathbf{0}_{3|\mathcal{C}| \times |\mathcal{S}|} \\ \mathbf{I}_{3|\mathcal{S}| \times |\mathcal{S}|} \end{pmatrix}.$$

Then

$$H^+ = \begin{pmatrix} R_{11} & R_{12} \\ R'_{12} & R_{22} \end{pmatrix},$$

where

$$R_{11} = \mathbf{Q}^{-1} [I_{3|\mathcal{C}| \times 3|\mathcal{C}|} - \mathbf{A}'_b (\mathbf{A}_b \mathbf{Q}^{-1} \mathbf{A}'_b)^+ \mathbf{A}_b \mathbf{Q}^{-1}],$$

$$R_{12} = \mathbf{Q}^{-1} \mathbf{A}'_b (\mathbf{A}_b \mathbf{Q}^{-1} \mathbf{A}'_b)^+,$$

and

$$R_{22} = -(\mathbf{A}_b \mathbf{Q}^{-1} \mathbf{A}'_b)^+.$$

In addition,

$$H^+ V = \begin{pmatrix} R_{12} \\ R_{22} \end{pmatrix}.$$

Proof of Theorem 5.3: Let $\hat{\mathbf{b}}$ denote the solution to (5.10), and note that only components of \mathbf{g} in \mathcal{C} are affected by the constraints. It follows by corollaries 3.2.4 and 3.2.5 of Fiacco (1983) that

$$\hat{\mathbf{b}} - \tilde{\mathbf{b}}^{(a)} = -R_{12} \mathbf{A}(\tilde{\mathbf{b}}^{(a)}) + O\left(\left\| \mathbf{A}(\tilde{\mathbf{b}}^{(a)}) \right\|^2\right),$$

where $A(\cdot) = [A_1(\cdot), \dots, A_{|\mathcal{S}|}(\cdot)]'$ and $\|\cdot\|$ is the Euclidean norm in $\mathbb{R}^{|\mathcal{S}|}$. Therefore,

$$\hat{\mathbf{b}} = \tilde{\mathbf{b}}^{(a)} - \mathbf{Q}^{-1} \mathbf{A}'_b (\mathbf{A}_b \mathbf{Q}^{-1} \mathbf{A}'_b)^+ \mathbf{A}(\tilde{\mathbf{b}}^{(a)}) + O\left(\left\| \mathbf{A}(\tilde{\mathbf{b}}^{(a)}) \right\|^2\right),$$

Recall that \mathbf{b}_0 denotes the $3|\mathcal{C}| \times 1$ vector

$$\mathbf{b}_0 = \{[g(x_j), g'(x_j), g''(x_j)]' : j \in \mathcal{C}\}$$

and that the modified definition of $A_k(\cdot)$ includes all components of \mathbf{b} in its arguments. Then Taylor series expansions yield

$$\begin{aligned} A_k(\tilde{\mathbf{b}}^{(a)}) &= A_k(\mathbf{b}_0) + \frac{\partial A_k(\tilde{\mathbf{b}}_k)}{\partial \mathbf{b}'} (\tilde{\mathbf{b}}^{(a)} - \mathbf{b}_0) \\ &= A_k(\mathbf{b}_0) + \frac{\partial A_k(\mathbf{b}_0)}{\partial \mathbf{b}'} (\tilde{\mathbf{b}}^{(a)} - \mathbf{b}_0) + \left[\frac{\partial A_k(\mathbf{b}_0)}{\partial \mathbf{b}'} - \frac{\partial A_k(\tilde{\mathbf{b}}_k)}{\partial \mathbf{b}'} \right] (\tilde{\mathbf{b}}^{(a)} - \mathbf{b}_0) \end{aligned}$$

for each $k = 1, \dots, |\mathcal{S}|$, where $\tilde{\mathbf{b}}$ is between $\tilde{\mathbf{b}}^{(a)}$ and \mathbf{b}_0 . Under Assumption 5, $\partial A_k(\mathbf{b}) / \partial \mathbf{b}$ is a continuous function of \mathbf{b} . Moreover columns of $\partial A_k(\mathbf{b}) / \partial \mathbf{b}$ corresponding to derivatives of g are zero, and $\mathbf{A}(\mathbf{b}_0) = 0$ because \mathbf{A} is defined as the vector giving the binding constraints. Therefore,

$$\mathbf{A}(\tilde{\mathbf{b}}^{(a)}) = \mathbf{A}_b(\tilde{\mathbf{g}}^{(a)} - \mathbf{g}^{(a)}) + O\left(\left\| \tilde{\mathbf{g}}^{(a)} - \mathbf{g}^{(a)} \right\|^2\right)$$

and

$$\hat{\mathbf{b}} - \mathbf{b}_0 = [I_{3|\mathcal{C}| \times 3|\mathcal{C}|} - \mathbf{Q}^{-1} \mathbf{A}'_b (\mathbf{A}_b \mathbf{Q}^{-1} \mathbf{A}'_b)^+ \mathbf{A}_b] (\tilde{\mathbf{b}} - \mathbf{b}_0) + O\left(\left\| \tilde{\mathbf{g}}^{(a)} - \mathbf{g}^{(a)} \right\|^2\right).$$

But $\mathbf{A}_b(\tilde{\mathbf{b}} - \mathbf{b}_0) = \mathbf{A}_b(\tilde{\mathbf{b}} - E\tilde{\mathbf{b}})$, because the estimator of \mathbf{g} is undersmoothed and columns of \mathbf{A}_b corresponding to derivatives of \mathbf{g} are zero. Therefore,

$$\begin{aligned} \hat{\mathbf{g}} - \mathbf{g}^{(a)} &= [I_{3|\mathcal{C}| \times 3|\mathcal{C}|} - \mathbf{Q}^{-1} \mathbf{A}'_b (\mathbf{A}_b \mathbf{Q}^{-1} \mathbf{A}'_b)^+ \mathbf{A}_b]^{(a)} (\tilde{\mathbf{g}}^{(a)} - \mathbf{g}^{(a)}) + O\left(\left\| \tilde{\mathbf{g}}^{(a)} - \mathbf{g}^{(a)} \right\|^2\right) \\ &= \{I_{|\mathcal{C}| \times |\mathcal{C}|} - [\mathbf{Q}^{-1} \mathbf{A}'_b (\mathbf{A}_b \mathbf{Q}^{-1} \mathbf{A}'_b)^+ \mathbf{A}_b]^{(a)}\} (\tilde{\mathbf{g}}^{(a)} - \mathbf{g}^{(a)}) + O\left(\left\| \tilde{\mathbf{g}}^{(a)} - \mathbf{g}^{(a)} \right\|^2\right), \end{aligned}$$

and

$$(nh)^{1/2} (\hat{\mathbf{g}} - \mathbf{g}^{(a)}) = \{I_{|\mathcal{C}| \times |\mathcal{C}|} - [\mathbf{Q}^{-1} \mathbf{A}'_b (\mathbf{A}_b \mathbf{Q}^{-1} \mathbf{A}'_b)^+ \mathbf{A}_b]^{(a)}\} (nh)^{1/2} (\tilde{\mathbf{g}}^{(a)} - \mathbf{g}^{(a)}) + o_p(1).$$

Now proceed as in the proof of Theorem 5.1. Q.E.D.

REFERENCES

- Andrews, D.W.K. (1999). Estimation when a parameter is on a boundary. *Econometrica*, 67, 1341-1383.
- Andrews, D.W.K. and X. Shi (2013). Inference based on conditional moment inequalities. *Econometrica*, 81, 609-666.
- Andrews, D.W.K. and G. Soares (2010). Inference for parameters defined by moment inequalities using generalized moment selection. *Econometrica*, 78, 119-157.
- Banerjee, M. and J.A. Wellner (2001). Likelihood ratio tests for monotone functions. *Annals of Statistics*, 29, 1699-1731.
- Bentkus, V. (2003). On the dependence of the Berry-Esséen bound on dimension. *Journal of Statistical Planning and Inference*, 113, 385-402.
- Birke, M. and H. Dette (2006). Estimating a convex function in nonparametric regression. *Scandinavian Journal of Statistics*, 34, 384-404.
- Blundell, R., Horowitz, J.L., and Parey, M. (2012). Measuring the price responsiveness of gasoline demand: economic shape restrictions and nonparametric demand estimation. *Quantitative Economics*, 3, 29-51.
- Blundell, R., Horowitz, J.L., and Parey, M. (2016). Nonparametric estimation of a heterogeneous demand function under the Slutsky inequality restriction. *Review of Economics and Statistics*, forthcoming.
- Bühlmann, P. and S. van de Geer (2011). *Statistics for High-Dimensional Data: Theory and Applications*. Heidelberg, Springer-Verlag.
- Buckley, M.J., G.K. Eagleson, and B.W. Silverman (1988). The Estimation of Residual Variance in Nonparametric Regression, *Biometrika*, 75, 189-199.
- Calonico, S., M.D. Cattaneo, and M.H. Farrell (2014). On the effect of bias estimation on coverage accuracy in nonparametric inference. Working paper, University of Michigan.
- Chernozhukov, V., I. Fernández-Val, and A. Galichon (2009). Improving point and interval estimators of monotone functions by rearrangement. *Biometrika*, 96, 559-575.
- Chernozhukov, V., S. Lee, and A. Rosen (2013). Intersection bounds: estimation and inference. *Econometrica*, 81, 667-737.
- Das Gupta, A. (2008). *Asymptotic theory of statistics and probability*. New York: Springer.
- Dette, H., N. Neumeyer, and K.F. Pilz (2006). A simple nonparametric estimator of a strictly monotone regression function. *Bernoulli*, 12, 469-490.
- Du, P. C. Parmeter, and J.C. Racine (2013). Nonparametric kernel regression with multiple predictors and multiple shape constraints. *Statistica Sinica*, 23, 1343-1372.
- Dumbgen, L. (2003). Optimal confidence bands for shape-restricted curves. *Bernoulli*, 9, 423-449.

- Fan, J. and I. Gijbels (1996). *Local Polynomial Modelling and Its Applications*. London: Chapman & Hall.
- Fiacco, A.V. (1983). *Introduction to Sensitivity and Stability Analysis in Nonlinear Programming*. New York: Academic Press.
- Gasser, T., L. Sroka, and C. Jennen-Steinmetz (1986). Residual variance and residual pattern in nonlinear regression. *Biometrika*, 73, 625-633.
- Groeneboom, P. and G. Jongbloed (2015). Nonparametric confidence intervals for monotone functions. *Annals of Statistics*, 43, 2019-2054.
- Groeneboom, P., G. Jongbloed, and J.A. Wellner (2001). Estimation of a convex function: characterizations and asymptotic theory. *Annals of Statistics*, 29, 1653-1698.
- Hall, P. and J.L. Horowitz (2013). A simple bootstrap method for constructing nonparametric confidence bands for functions. *Annals of Statistics*, 41, 1892-1921.
- Hall, P. and L.-S. Huang (2001). Nonparametric kernel regression subject to monotonicity constraints. *Annals of Statistics*, 29, 624-647.
- Hall, P. and L.-S. Huang (2002). Unimodal density estimation using kernel methods. *Statistica Sinica*, 12, 965-990.
- Hall, P., L.-S. Huang, J.A. Gifford, and I. Gijbels (2001). Nonparametric estimation of hazard rate under the constraint of monotonicity. *Journal of Computational and Graphical Statistics*, 10, 592-614.
- Hall, P. and B. Presnell (1999). Density estimation under constraints. *Journal of Computational and Graphical Statistics*, 8, 259-277.
- Hall, P. and A. Yatchew (2005). Unified approach to testing functional hypotheses in semiparametric contexts. *Journal of Econometrics*, 127, 225-252.
- Härdle, W. (1990). *Applied Nonparametric Regression*. Cambridge: Cambridge University Press.
- Horowitz, J.L. and J. Huang (2013). Penalized estimation of high-dimensional models under a generalized sparsity condition. *Statistica Sinica*, 23, 725-748.
- 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.
- Jacho-Chávez, D., A. Lewbel, and O. Linton (2010). Identification and nonparametric estimation of a transformed additively separable model. *Journal of Econometrics*, 156, 392-407.
- Lee, S., K. Song, and Y.-J. Whang (2013). Testing functional inequalities. *Journal of Econometrics*, 172, 14-32.
- Lee, S., K. Song, and Y.-J. Whang (2014). Testing for a general class of functional inequalities. Cemmap working paper cwp 9/14, University College London.

- Leeb, H. and B.M. Pötscher (2005). Model selection and inference: facts and fiction. *Econometric Theory*, 21, 21-59.
- Leeb, H. and B. Pötscher (2006). Can one estimate the conditional distribution of post-model-selection estimators? *Annals of Statistics*, 34, 2554-2591.
- Liero, H. (1982). On the maximal deviation of the kernel regression function estimate. *Mathematische Operationsforschung, Serie Statistics*, 13, 171-182.
- Liew, C.K. (1976). Inequality constrained least-squares estimation. *Journal of the American Statistical Association*, 71, 746-751.
- Mammen, E. (1991a). Estimating a smooth monotone regression function. *Annals of Statistics*, 19, 724-740.
- Mammen, E. (1991b). Nonparametric regression under qualitative smoothness assumptions. *Annals of Statistics*, 19, 741-759.
- Mammen, E. and C. Thomas-Agnan (1999). Smoothing splines and shape restrictions. *Scandinavian Journal of Statistics*, 26, 239-252.
- Matzkin, R.L. (1994). Restrictions of economic theory in nonparametric methods. In Engle, R.F. and D.L. McFadden, eds., *Handbook of Econometrics*, vol. 4. pp. 2523-2558. Amsterdam: Elsevier.
- Pal, J.K. and Woodroffe M. (2007). Large sample properties of shape restricted regression estimators with smoothness adjustments. *Statistica Sinica*, 17, 1601-1616.
- Rice, J. (1984). Bandwidth Choice for Nonparametric Regression, *Annals of Statistics*, 12, 1215-1230.
- Romano, J., A. Shaikh, and M. Wolf (2014). A practical two-step method for testing moment inequality models. *Econometrica*, 82, 1979-2002.
- Ruppert, D. and M.P. Wand (1994). Multivariate locally weighted least squares regression. *Annals of Statistics*, 22, 1346-1370.
- Wang, X. and J. Shen (2013). Uniform convergence and rate adaptive estimation of convex functions via constrained optimization. *SIAM Journal on Control and Optimization*, 51, 2753-2787.

Table 1: Results of Monte Carlo Experiments when the Constraints Are Binding

Model: Constant returns to scale; $\tau = 1$
 Nominal coverage probability: 0.95

n	C_h	Estimation Method	Empirical Cov. Prob.	Relative Width
1000	0.95	Unconstrained	0.946	1.0
	1.00		0.942	1.0
	1.05		0.945	1.0
2000	0.95		0.954	1.0
	1.00		0.952	1.0
	1.05		0.951	1.0
1000	0.95	Constrained	0.947	0.818
	1.00		0.933	0.818
	1.05		0.944	0.818
2000	0.95		0.953	0.819
	1.00		0.954	0.819
	1.05		0.960	0.818
1000	0.95	Oracle	0.949	0.816
	1.00		0.935	0.817
	1.05		0.944	0.817
2000	0.95		0.958	0.817
	1.00		0.959	0.817
	1.05		0.962	0.817
1000	0.95	Bias Corrected	0.936	0.906
	1.00		0.956	0.901
	1.05		0.955	0.907
2000	0.95		0.946	0.908
	1.00		0.938	0.911
	1.05		0.949	0.911

Table 2: Results of Monte Carlo Experiments when the Constraints Are Nearly BindingModel: Decreasing returns to scale; $\tau = 0.9$

Nominal coverage probability: 0.95

n	C_h	Estimation Method	Empirical Cov. Prob.	Relative Width
1000	0.95	Unconstrained	0.946	1.0
	1.00		0.942	1.0
	1.05		0.945	1.0
2000	0.95		0.954	1.0
	1.00		0.952	1.0
	1.05		0.951	1.0
1000	0.95	Constrained	0.771	0.828
	1.00		0.742	0.829
	1.05		0.729	0.831
2000	0.95		0.680	0.832
	1.00		0.660	0.836
	1.05		0.615	0.837
1000	0.95	Oracle	0.948	1.0
	1.00		0.943	1.0
	1.05		0.944	1.0
2000	0.95		0.953	1.0
	1.00		0.950	1.0
	1.05		0.952	1.0
1000	0.95	Bias Corrected	0.934	0.972
	1.00		0.909	0.970
	1.05		0.931	0.977
2000	0.95		0.931	0.983
	1.00		0.930	0.987
	1.05		0.923	0.986

Table 3: Results of Monte Carlo Experiments when the Constraints Are Far from BindingModel: Decreasing returns to scale; $\tau = 0.5$


Nominal coverage probability: 0.95

n	C_h	Estimation Method	Empirical Cov. Prob.	Relative Width
1000	0.95	Unconstrained	0.946	1.0
	1.00		0.941	1.0
	1.05		0.943	1.0
2000	0.95		0.954	1.0
	1.00		0.951	1.0
	1.05		0.951	1.0
1000	0.95	Constrained	0.946	0.999
	1.00		0.942	1.0
	1.05		0.944	1.0
2000	0.95		0.953	1.0
	1.00		0.949	1.0
	1.05		0.952	1.0
1000	0.95	Oracle	0.948	1.0
	1.00		0.942	1.0
	1.05		0.944	1.0
2000	0.95		0.953	1.0
	1.00		0.949	1.0
	1.05		0.952	1.0
1000	0.95	Bias Corrected	0.941	1.0
	1.00		0.940	1.0
	1.05		0.942	1.0
2000	0.95		0.962	1.0
	1.00		0.958	1.0
	1.05		0.955	1.0

Table 4: Results of Estimating a Production Function

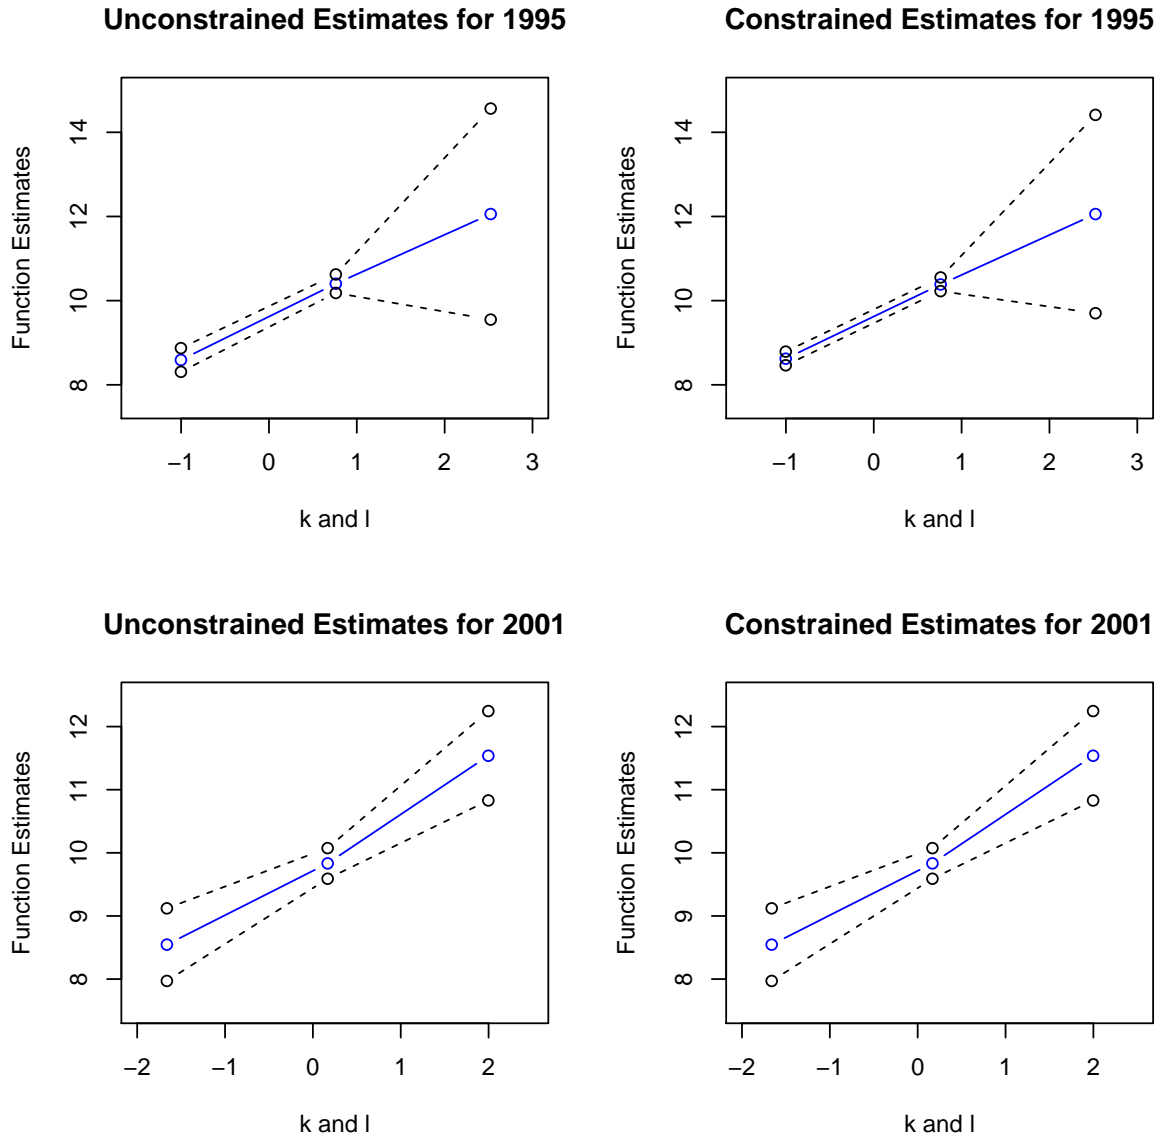
Year	Estimation Method	k, ℓ	Function Estimate	Standard Error	Lower 95% Uniform Bound	Upper 95% Uniform Bound
1995	Unconstrained	-1.002	8.591	0.117	8.309	8.872
		0.761	10.403	0.092	10.183	10.623
		2.524	12.058	1.044	9.550	14.565
	Constrained	-1.002	8.621	0.072	8.481	8.761
		0.761	10.384	0.072	10.244	10.524
		2.524	12.147	0.072	12.007	12.287
	Bias Corr.	-1.002	8.621	0.072	8.464	8.789
		0.761	10.384	0.072	10.226	10.552
		2.524	12.058	1.044	9.701	14.414
2001	Unconstrained	-1.662	8.546	0.239	7.971	9.121
		0.168	9.833	0.101	9.590	10.075
		1.997	11.538	0.295	10.830	12.246
	Constrained	-1.662	8.066	0.089	7.894	8.239
		0.168	9.896	0.089	9.723	10.068
		1.997	11.725	0.089	11.553	11.898
	Bias Corr.	-1.662	8.546	0.239	7.971	9.121
		0.168	9.833	0.101	9.590	10.075
		1.997	11.538	0.295	10.830	12.246

FIGURE 1. Nonparametric Estimates and their Uniform Confidence Bands (No Bias Correction)



Note: The solid lines represent nonparametric estimates, whereas the dashed lines show 95% uniform confidence bands. The circles correspond to the grid. On one hand, the top and bottom panels show estimates for 1995 and 2001, respectively. On the other hand, the left and right panels show unconstrained and constrained estimates, respectively.

FIGURE 2. Nonparametric Estimates and their Uniform Confidence Bands
(Bias Correction Implemented)



Note: The solid lines represent nonparametric estimates, whereas the dashed lines show 95% uniform confidence bands. The circles correspond to the grid. On one hand, the top and bottom panels show estimates for 1995 and 2001, respectively. On the other hand, the left and right panels show unconstrained and constrained estimates, respectively. The bias correction is implemented with $m = 0.005$.