

Inference for Identifiable Parameters in Partially Identified Econometric Models

Joseph P. Romano

Departments of Economics and Statistics
Stanford University
romano@stanford.edu

Azeem M. Shaikh

Department of Economics
University of Chicago
amshaikh@uchicago.edu

February 2008

Abstract

This paper considers the problem of inference for partially identified econometric models. The class of models studied are defined by a population objective function $Q(\theta, P)$ for $\theta \in \Theta$. The second argument indicates the dependence of the objective function on P , the distribution of the observed data. Unlike the classical extremum estimation framework, it is not assumed that $Q(\theta, P)$ has a unique minimizer in the parameter space Θ . The goal may be either to draw inferences about some unknown point in the set of minimizers of the population objective function or to draw inferences about the set of minimizers itself. In this paper, the object of interest is some unknown point $\theta \in \Theta_0(P)$, where $\Theta_0(P) = \arg \min_{\theta \in \Theta} Q(\theta, P)$, and so we seek random sets that contain each $\theta \in \Theta_0(P)$ with at least some prespecified probability asymptotically. We also consider situations where the object of interest is the image of some point $\theta \in \Theta_0(P)$ under a known function. Computationally intensive, yet feasible procedures for constructing random sets satisfying the desired coverage property under weak assumptions are provided. We also provide conditions under which the confidence regions are uniformly consistent in level.

KEYWORDS: Extremum Estimation, Partially Identified Model, Incomplete Model, Identified Set, Identifiable Parameter, Subsampling, Uniform Coverage, Confidence Region, Interval Regression, Moment Inequalities

1 Introduction

This paper provides computationally intensive, yet feasible methods for inference for a large class of partially identified econometric models. A partially identified model is a model in which the parameter of interest is not uniquely defined by the distribution of the observed data. Such models arise naturally in many parts of empirical work in economics. The class of models considered is defined by a population objective function $Q(\theta, P)$ for $\theta \in \Theta$. The second argument indicates the dependence of the objective function on P , the distribution of the observed data. Unlike the classical extremum estimation framework, it is not assumed that $Q(\theta, P)$ has a unique minimizer in the parameter space Θ . The goal may be either to draw inferences about some unknown point in the set of minimizers of the population objective function or to draw inferences about the set of minimizers itself. In this paper, we consider the first of these two goals. The object of interest is some unknown point $\theta \in \Theta_0(P)$, where

$$\Theta_0(P) = \arg \min_{\theta \in \Theta} Q(\theta, P) . \quad (1)$$

We henceforth refer to any $\theta \in \Theta_0(P)$ as an *identifiable parameter* and $\Theta_0(P)$ as the *identified set*. In this instance, given i.i.d. data $X_i, i = 1, \dots, n$, generated from P , we seek random sets $\mathcal{C}_n = \mathcal{C}_n(X_1, \dots, X_n)$ that contain each identifiable parameter with at least some prespecified probability asymptotically. That is, for all $\theta \in \Theta_0(P)$ we have

$$\liminf_{n \rightarrow \infty} P\{\theta \in \mathcal{C}_n\} \geq 1 - \alpha . \quad (2)$$

We term such sets as *confidence regions for identifiable parameters that are pointwise consistent in level*. This terminology reflects the fact that the confidence regions are valid only for a *fixed* probability distribution P and helps distinguish this coverage requirement from others discussed later in which we will demand that the confidence regions are valid uniformly in P . We construct random sets \mathcal{C}_n satisfying (2) under weak assumptions on P . These assumptions are formulated in terms of restrictions on the asymptotic behavior of the estimate of the population objective function, $\hat{Q}_n(\theta)$. Most often $\hat{Q}_n(\theta) = Q(\theta, \hat{P}_n)$ for some estimate \hat{P}_n of P . Our construction is based on the usual duality between confidence regions and hypothesis testing: we invert tests of each of the null hypotheses $H_\theta : \theta \in \Theta_0(P)$ for $\theta \in \Theta$ that control the usual probability of a Type 1 error at level α . Typically, Θ is a subset of Euclidean space and so the number of null hypotheses may be uncountably infinite. The idea of inverting tests using computer-intensive methods can be traced to DiCiccio and Romano (1990).

In the second goal, the object of interest is the identified set, $\Theta_0(P)$, itself. In this case, given i.i.d. data $X_i, i = 1, \dots, n$, generated from P , we seek instead random sets $\mathcal{C}_n = \mathcal{C}_n(X_1, \dots, X_n)$

that contain the identified set with at least some prespecified probability asymptotically. The problem of constructing such sets is treated in a companion paper, Romano and Shaikh (2006b).

Our results on confidence regions for identifiable parameters build upon the earlier work of Chernozhukov et al. (2007), who were the first to consider inference for the same class of partially identified models. Our analysis also extends the work of Imbens and Manski (2004), who analyze the special case of the above class of partially identified models in which the identified set is an interval whose upper and lower endpoints are means or at least behave like means asymptotically. The authors discuss different definitions for confidence regions in the context of this setting and argue forcefully that it may be desirable to demand validity of confidence regions not just for each fixed probability distribution P satisfying weak assumptions, but rather uniformly over all P in some large class of distributions \mathbf{P} . Confidence regions that fail to satisfy this requirement have the feature that for every sample size n , however large, there is some probability distribution $P \in \mathbf{P}$ for which the coverage probability of the confidence region under P is not close to the prescribed level. Researchers may therefore feel that inferences made on the basis of asymptotic approximations are more reliable if the confidence regions exhibit good uniform behavior. Of course, such a requirement will typically require restrictions on P beyond those required for pointwise consistency in level. Bahadur and Savage (1956), for example, show that if \mathbf{P} is suitably large, then there exists no confidence interval for the mean with finite length and good uniform behavior. Romano (2004) extends this non-existence result to a number of other problems. We provide restrictions on \mathbf{P} under which both types of confidence regions suggested in this paper have good uniform behavior. Concretely, we provide conditions under which \mathcal{C}_n satisfies

$$\liminf_{n \rightarrow \infty} \inf_{\theta \in \Theta} \inf_{P \in \mathbf{P}: \theta \in \Theta_0(P)} P\{\theta \in \mathcal{C}_n\} \geq 1 - \alpha . \quad (3)$$

By analogy with our earlier terminology, sets satisfying (3) are referred to as *confidence regions for identifiable parameters that are uniformly consistent in level*. Note that if the identified set $\Theta_0(P)$ consists of a single point $\theta_0(P)$, then this definition reduces to the usual definition of confidence regions that are uniformly consistent in level; that is,

$$\liminf_{n \rightarrow \infty} \inf_{P \in \mathbf{P}} P\{\theta_0(P) \in \mathcal{C}_n\} \geq 1 - \alpha .$$

In order to provide conditions under which these coverage properties hold, we use recently developed results that provide general conditions under which confidence regions for a parameter constructed using subsampling are uniformly consistent in level.

We have so far assumed that the object of interest is an identifiable parameter, $\theta \in \Theta_0(P)$. More generally, the object of interest may be the image of an identifiable parameter under a

known function. A typical example of such a function is the projection of \mathbf{R}^k onto one of the axes. We extend the above definitions of confidence regions to this setting as follows. Consider a function $f : \Theta \rightarrow \Lambda$. Denote by $\Lambda_0(P)$ the image of $\Theta_0(P)$ under f ; that is,

$$\Lambda_0(P) = \{f(\theta) : \theta \in \Theta_0(P)\} . \quad (4)$$

We refer to a set \mathcal{C}_n^f as a *confidence region for a function of identifiable parameters that is pointwise consistent in level* if for any $\lambda \in \Lambda_0(P)$ we have that

$$\liminf_{n \rightarrow \infty} P\{\lambda \in \mathcal{C}_n^f\} \geq 1 - \alpha . \quad (5)$$

As before, we may also demand uniformly good behavior over a class of probability distributions \mathbf{P} by requiring that \mathcal{C}_n^f satisfy

$$\liminf_{n \rightarrow \infty} \inf_{\lambda \in \Lambda} \inf_{P \in \mathbf{P} : \lambda \in \Lambda_0(P)} P\{\lambda \in \mathcal{C}_n^f\} \geq 1 - \alpha . \quad (6)$$

By analogy with our earlier terminology, sets satisfying (6) are referred to as *confidence regions for a function of identifiable parameters that are uniformly consistent in level*. We adapt our constructions of confidence regions for identifiable parameters to provide constructions of confidence sets satisfying these alternative coverage requirements.

Methods for drawing inferences in partially identified models have allowed economists to solve many empirical problems that previously either were intractable or relied on untenable assumptions to achieve identification. Manski (2003) has argued forcefully against invoking such assumptions to make inferences in the context of missing outcome data and treatment response, as it degrades the credibility of inferences made under those assumptions. Another class of models that may only be partially identified is given by game-theoretic models with multiple equilibria. When confronted with such models, researchers often impose identification by assuming some sort of an ad hoc equilibrium selection mechanism. More recently, the empirical literature has explored inferences based on exploiting only those restrictions implied by equilibrium behavior that do not depend on the particular equilibrium being played by the agents. These restrictions are often defined by a system of moment inequalities, which is a special case of the class of models considered in this paper. Ciliberto and Tamer (2004) and Borzekowski and Cohen (2005), for example, use this idea to analyze an entry model and a coordination game, respectively. Benkard et al. (2007) consider the problem of making inferences in dynamic models of imperfect competition, where, in the absence of additional assumptions, the unknown parameters of the model are naturally restricted to a set defined by moment inequalities. Andrews et al. (2004), Pakes et al. (2005), and Rosen (2007) consider

other economic applications of systems of moment inequalities, but develop independent methods of inference. These methods all assume some conservative upper bound on the number of moments where the inequality holds with equality and are therefore more conservative than the methods developed in this paper. Andrews and Guggenberger (2007) avoid this problem by considering the use of subsampling for models defined by moment inequalities. Andrews and Soares (2007) consider the use of “moment selection” techniques to select the moments for which the inequality holds with equality in a first stage and then use multivariate normal approximations to construct critical values. Similar ideas have been proposed by Bugni (2007) and Canay (2007), who use the bootstrap instead of multivariate normal approximations to construct critical values. Such methods, however, depend on the unique structure of the moment inequality problem and therefore can not be easily generalized to more general partially identified models.

Importantly, the results in this paper apply not only to systems of moment inequalities, but also to other examples of partially identified models. Bajari et al. (2007) consider inference in a semiparametric discrete choice model, which gives rise to a maximum score-type of objective function with possibly more than one maximizer. Santos (2007) considers the problem of inference in a nonparametric regression model with identification through instrumental variables. In this case, if one is not willing to assume that the instrumental variables are a complete and sufficient statistic for the regressors, then the model is only partially identified. This example is also noteworthy because the object of interest in this case, the unknown regression function, is infinite dimensional.

Beresteanu and Molinari (2006) propose an alternative approach to inference in partially identified models that is based on the theory of random sets. Their approach applies to a more restrictive class of models than considered here. Moreover, their construction of confidence regions for identifiable parameters is more conservative than ours. Galichon and Henry (2006) also provide an approach to inference for partially identified models, but it too is more conservative than ours.

The remainder of the paper is organized as follows. We begin in Section 2 by providing a number of concrete examples that fall within the class of models we are considering. These examples not only serve as motivation, but will also be useful later for illustrating our methodology. In Section 3, we analyze the problem of constructing confidence sets that satisfy the coverage requirements (2) and (3). We then extend this methodology to construct confidence regions satisfying (5) and (6). Finally, we conclude and provide directions for future research in Section 4.

2 Some Motivating Examples

In this section, we provide motivation for our study of partially identified models by describing several specific examples of such models. We will return to these examples repeatedly in the sequel as illustrations of our methodology.

Example 2.1 (*One-Sided Mean*) Perhaps the simplest example of a partially identified model is given by the following setup. Let $X_i, i = 1, \dots, n$, be an i.i.d. sequence of random variables with distribution P on \mathbf{R} . Denote by $\mu(P)$ the mean of the distribution P . The parameter of interest, θ_0 , is known to satisfy $\theta_0 \geq \mu(P)$. For example, θ_0 might be the mean of an unobservable random variable Z_i with distribution Q on \mathbf{R} that is known to satisfy $Z_i \geq X_i$, and so $\mu(Q) \geq \mu(P)$. The identified set is therefore given by $\Theta_0(P) = \{\theta \in \mathbf{R} : \theta \geq \mu(P)\}$. We may characterize this set as the set of minimizers of the population objective function

$$Q(\theta, P) = (\mu(P) - \theta)_+^2,$$

where the notation $(a)_+$ is used as shorthand for $\max\{a, 0\}$. The sample analog of $Q(\theta, P)$ is given by $\hat{Q}_n(\theta) = (\bar{X}_n - \theta)_+^2$. ■

Example 2.2 (*Two-Sided Mean*) A natural generalization of Example 2.1 is to consider bivariate random variables. To this end, let $(X_i, Y_i), i = 1, \dots, n$, be an i.i.d. sequence of random variables with distribution P on \mathbf{R}^2 . Let $\mu_X(P)$ denote the mean of the first component of the distribution P and $\mu_Y(P)$ the mean of the second component of the distribution P . The parameter of interest, θ_0 , is known to satisfy $\mu_X(P) \leq \theta_0 \leq \mu_Y(P)$. For example, θ_0 might be the mean of another distribution Q on \mathbf{R} that is known to satisfy $\mu_X(P) \leq \mu(Q) \leq \mu_Y(P)$. The identified set is therefore given by $\Theta_0(P) = \{\theta \in \mathbf{R} : \mu_X(P) \leq \theta \leq \mu_Y(P)\}$. This set may be characterized as the set of minimizers of

$$Q(\theta, P) = (\mu_X(P) - \theta)_+^2 + (\theta - \mu_Y(P))_+^2.$$

The sample analog of $Q(\theta, P)$ is given by $\hat{Q}_n(\theta) = (\bar{X}_n - \theta)_+^2 + (\theta - \bar{Y}_n)_+^2$. ■

Remark 2.1 Let $W_i \in [0, 1]$ and $D_i \in \{0, 1\}$. For example, W_i may be the answer to the question, “Do you vote Republican or Democrat?”, and D_i may be the indicator variable for whether the person asked chooses to answer the question. Suppose the researcher observes an i.i.d. sequence $(W_i D_i, D_i), i = 1, \dots, n$, with distribution P ; i.e., W_i is observed if and only if $D_i = 1$. The parameter of interest, $\theta_0 = E\{W_i\}$, is not determined by the distribution of the

observed data, but the researcher can say with certainty that θ_0 satisfies $E_P\{W_i D_i\} \leq \theta_0 \leq E_P\{W_i D_i + 1 - D_i\}$. By identifying $X_i = W_i D_i$ and $Y_i = W_i D_i + 1 - D_i$, this example can be seen to be a special case of Example 2.2. This example is analyzed in detail by Imbens and Manski (2004). See Stoye (2007) for further analysis. ■

Example 2.3 (*Regression with Interval Outcomes*) The following example allows for inference in a linear regression model in which the dependent variable is interval-censored. Let $(X_i, Y_{1,i}, Y_{2,i}, Y_i^*), i = 1, \dots, n$, be an i.i.d. sequence of random variables with distribution Q on $\mathbf{R}^k \times \mathbf{R} \times \mathbf{R} \times \mathbf{R}$. The parameter of interest, θ_0 , is known to satisfy $E_Q\{Y_i^* | X_i\} = X_i' \theta_0$, but Y_i^* is unobserved, which precludes conventional estimation of θ_0 . Let P denote the distribution of the observed random variables $(X_i, Y_{1,i}, Y_{2,i})$. The random variables $(Y_{1,i}, Y_{2,i})$ are known to satisfy $Y_{1,i} \leq Y_i^* \leq Y_{2,i}$ with probability 1 under Q . Thus, $\theta_0 \in \Theta_0(P) = \{\theta \in \mathbf{R}^k : E_P\{Y_{1,i} | X_i\} \leq X_i' \theta \leq E_P\{Y_{2,i} | X_i\} \text{ } P\text{-a.s.}\}$. This set may be characterized as the set of minimizers of

$$Q(\theta, P) = E_P\{(E_P\{Y_{1,i} | X_i\} - X_i' \theta)_+^2 + (X_i' \theta - E_P\{Y_{2,i} | X_i\})_+^2\}.$$

The sample analog $\hat{Q}_n(\theta)$ of $Q(\theta, P)$ is given by replacing expectations with appropriately defined estimators, the nature of which may depend on further assumptions about P . Manski and Tamer (2002) characterize the identified set in this setting and also consider the case where Y_i^* is observed, but X_i is interval-censored. ■

Remark 2.2 A Tobit-like model is a special case of the previous example if we suppose further that $Y_{2,i} = Y_i^*$ and $Y_{1,i} = Y_i^*$ if $Y_i^* > 0$, and $Y_{2,i} = 0$ and $Y_{1,i} = -\infty$ (or some large negative number if there is a plausible lower bound on Y_i^*) if $Y_i^* \leq 0$. It is worthwhile to note that conventional approaches to inference for such a model, while enforcing identification, rely on the much stronger assumption that $\epsilon_i = Y_i^* - X_i' \theta_0$ is independent of X_i . By allowing for only partial identification, we will be able to draw inferences under much weaker assumptions. ■

Example 2.4 (*Moment Inequalities*) Consider the following generalization of Examples 2.1 and 2.2. Let $X_i, i = 1, \dots, n$, be an i.i.d. sequence of random variables with distribution P on \mathbf{R}^k . For $j = 1, \dots, m$, let $g_j(x, \theta)$ be a real-valued function on $\mathbf{R}^k \times \mathbf{R}^l$. The identified set is assumed to be $\Theta_0(P) = \{\theta \in \mathbf{R}^l : E_P\{g_j(X_i, \theta)\} \leq 0 \forall j \text{ s.t. } 1 \leq j \leq m\}$. This set may be characterized as the set of minimizers of

$$Q(\theta, P) = \sum_{1 \leq j \leq m} (E_P\{g_j(X_i, \theta)\})_+^2,$$

or equivalently as $\{\theta \in \Theta : Q(\theta, P) = 0\}$. The sample analog of $Q(\theta, P)$ is given by $\hat{Q}_n(\theta) = \sum_{1 \leq j \leq m} (\frac{1}{n} \sum_{1 \leq i \leq n} g_j(X_i, \theta))_+^2$. This choice of $Q(\theta, P)$ is especially noteworthy because it is used by Cilberto and Tamer (2004), Benkard et al. (2007) and Borzekowski and Cohen (2005) in their empirical applications. Since any equality restriction may be thought of as two inequality restrictions, this example may also be viewed as a generalization of the method of moments to allow for equality and inequality restrictions on the moments, rather than just equality restrictions. ■

Remark 2.3 A prominent example of an econometric model which gives rise to moment inequalities is an entry model. See, for example, Andrews et al. (2004) or Cilberto and Tamer (2004) for a detailed description of such models and a derivation of the inequalities. Briefly, consider an entry model with two firms and let X_i be the indicator for the event “firm 1 enters”. Because of the multiplicity of Nash equilibria, the model only gives upper and lower *bounds*, $L(\theta)$ and $U(\theta)$, on the probability of this event as a function of the unknown parameter, θ , of the econometric model. It is therefore natural to use the functions

$$\begin{aligned} g_1(X_i, \theta) &= L(\theta) - X_i \\ g_2(X_i, \theta) &= X_i - U(\theta) \end{aligned}$$

as a basis for inference in such a model. ■

3 Confidence Regions for Identifiable Parameters

In this section, we consider the problem of constructing confidence regions for identifiable parameters. We begin in Section 3.1 by recalling a general result about subsampling that will be used frequently throughout the remainder of the paper. The proof of the result, which is lengthy, can be found in Romano and Shaikh (2006a). In Section 3.2, we then treat the construction of sets satisfying (2), before turning our attention in Section 3.3 to the problem of constructing sets satisfying (3). In Section 3.4, we generalize these constructions to confidence regions for functions of identifiable parameters.

3.1 A Useful Result about Subsampling

The following theorem describes conditions under which confidence sets constructed using subsampling are asymptotically valid for a fixed P as well as uniformly over \mathbf{P} . The conditions

for a fixed P are distinct from those described in Politis et al. (1999) and can accomodate certain degenerate situations not covered by the results presented there. As we will see below, this feature of the conditions will be essential for the types of problems we are considering.

Theorem 3.1 *Let $X_i, i = 1, \dots, n$, be an i.i.d. sequence of random variables with distribution P . Let $\vartheta(P)$ be a real-valued parameter and let $\hat{\vartheta}_n = \hat{\vartheta}_n(X_1, \dots, X_n)$ be some estimator of $\vartheta(P)$ that is symmetric in its arguments. Denote by $J_n(\cdot, P)$ the distribution of the root $\tau_n(\hat{\vartheta}_n - \vartheta(P))$, where τ_n is a sequence of known constants. Let $b = b_n < n$ be a sequence of positive integers tending to infinity, but satisfying $b/n \rightarrow 0$. Let $N_n = \binom{n}{b}$ and*

$$L_n(x) = \frac{1}{N_n} \sum_{1 \leq i \leq N_n} I\{\tau_b(\hat{\vartheta}_{n,b,i} - \vartheta(P)) \leq x\}, \quad (7)$$

where $\hat{\vartheta}_{n,b,i}$ denotes the estimate $\hat{\vartheta}_n$ evaluated at the i th subset of data of size b . Finally, let $L_n^{-1}(1 - \alpha) = \inf\{x \in \mathbf{R} : L_n(x) \geq 1 - \alpha\}$. Then, for $\alpha \in (0, 1)$, the following statements are true:

(i) *If $\limsup_{n \rightarrow \infty} \sup_{x \in \mathbf{R}} \{J_b(x, P) - J_n(x, P)\} \leq 0$, then*

$$\liminf_{n \rightarrow \infty} P\{\tau_n(\hat{\vartheta}_n - \vartheta(P)) \leq L_n^{-1}(1 - \alpha)\} \geq 1 - \alpha.$$

(ii) *If $\limsup_{n \rightarrow \infty} \sup_{x \in \mathbf{R}} \{J_n(x, P) - J_b(x, P)\} \leq 0$, then*

$$\liminf_{n \rightarrow \infty} P\{\tau_n(\hat{\vartheta}_n - \vartheta(P)) \geq L_n^{-1}(\alpha)\} \geq 1 - \alpha.$$

(iii) *If $\lim_{n \rightarrow \infty} \sup_{x \in \mathbf{R}} |J_b(x, P) - J_n(x, P)| = 0$, then*

$$\liminf_{n \rightarrow \infty} P\{L_n^{-1}(\frac{\alpha}{2}) \leq \tau_n(\hat{\vartheta}_n - \vartheta(P)) \leq L_n^{-1}(1 - \frac{\alpha}{2})\} \geq 1 - \alpha.$$

Moreover, we have that:

(iv) *If $\limsup_{n \rightarrow \infty} \sup_{P \in \mathbf{P}} \sup_{x \in \mathbf{R}} \{J_b(x, P) - J_n(x, P)\} \leq 0$, then*

$$\liminf_{n \rightarrow \infty} \inf_{P \in \mathbf{P}} P\{\tau_n(\hat{\vartheta}_n - \vartheta(P)) \leq L_n^{-1}(1 - \alpha)\} \geq 1 - \alpha.$$

(v) *If $\limsup_{n \rightarrow \infty} \sup_{P \in \mathbf{P}} \sup_{x \in \mathbf{R}} \{J_n(x, P) - J_b(x, P)\} \leq 0$, then*

$$\liminf_{n \rightarrow \infty} \inf_{P \in \mathbf{P}} P\{\tau_n(\hat{\vartheta}_n - \vartheta(P)) \geq L_n^{-1}(\alpha)\} \geq 1 - \alpha.$$

(vi) If $\lim_{n \rightarrow \infty} \sup_{P \in \mathbf{P}} \sup_{x \in \mathbf{R}} |J_b(x, P) - J_n(x, P)| = 0$, then

$$\liminf_{n \rightarrow \infty} \inf_{P \in \mathbf{P}} P\{L_n^{-1}(\frac{\alpha}{2}) \leq \tau_n(\hat{\vartheta}_n - \vartheta(P)) \leq L_n^{-1}(1 - \frac{\alpha}{2})\} \geq 1 - \alpha .$$

Remark 3.1 The subsampling distribution $L_n(x)$ defined by (7) differs from the typical definition given in Politis et al. (1999) in that it centers about the true value of the parameter $\vartheta(P)$ instead of its estimate $\hat{\vartheta}_n$. Strictly speaking, it is therefore not possible to compute $L_n(x)$ as it is defined above. We may, however, compute $L_n(x)$ under the null hypothesis that $\vartheta(P)$ equals a particular value. For our purposes below, this will suffice. ■

3.2 Pointwise Consistency in Level

We assume in the remainder of the paper that $\inf_{\theta \in \Theta} Q(\theta, P) = 0$. Note that this restriction is without loss of generality, as it can always be enforced simply by subtracting the quantity $\inf_{\theta \in \Theta} Q(\theta, P)$ from the objective function whenever $\inf_{\theta \in \Theta} Q(\theta, P)$ exists.

Our construction exploits the well-known duality between $1 - \alpha$ confidence regions and tests of the individual null hypotheses

$$H_\theta : Q(\theta, P) = 0 \text{ for } \theta \in \Theta \tag{8}$$

that control the usual probability of a Type 1 error at level α . The hypotheses H_θ are one-sided in the sense that the alternative hypotheses are understood to be

$$K_\theta : Q(\theta, P) > 0 \text{ for } \theta \in \Theta .$$

Typically, Θ will be a subset of Euclidean space, in which case there may be an uncountably infinite number of hypotheses in this family. For each null hypothesis H_θ , a test statistic is required such that large values of the test statistic provide evidence against the null hypothesis. We will use the statistic $a_n \hat{Q}_n(\theta)$ for some sequence $a_n \rightarrow \infty$ for this purpose.

The construction of critical values for these tests will be based on subsampling. In order to define the critical values precisely, some further notation is required. Let $b = b_n < n$ be a sequence of positive integers tending to infinity, but satisfying $b/n \rightarrow 0$. Let $N_n = \binom{n}{b}$ and let $\hat{Q}_{n,b,i}(\theta)$ denote the statistic $\hat{Q}_n(\theta)$ evaluated at the i th subset of data of size b from the n observations. For $\alpha \in (0, 1)$, define

$$\hat{d}_n(\theta, 1 - \alpha) = \inf\{x : \frac{1}{N_n} \sum_{1 \leq i \leq N_n} I\{a_b \hat{Q}_{n,b,i}(\theta) \leq x\} \geq 1 - \alpha\} . \tag{9}$$

We will test each null hypothesis H_θ by comparing the test statistic $a_n\hat{Q}_n(\theta)$ with the critical value $\hat{d}_n(\theta, 1 - \alpha)$. The set of θ values corresponding to the accepted hypotheses from this procedure will form \mathcal{C}_n , i.e.,

$$\mathcal{C}_n = \{\theta \in \Theta : a_n\hat{Q}_n(\theta) \leq \hat{d}_n(\theta, 1 - \alpha)\} . \quad (10)$$

We now provide conditions under which \mathcal{C}_n defined by (10) satisfies the desired coverage property (2).

Theorem 3.2 *Let $X_i, i = 1, \dots, n$, be an i.i.d. sequence of random variables with distribution P and let $b = b_n < n$ be a sequence of positive integers tending to infinity, but satisfying $b/n \rightarrow 0$. Denote by $J_n(\cdot, \theta, P)$ the distribution of $a_n(\hat{Q}_n(\theta) - Q(\theta, P))$. Suppose that for every $\theta \in \Theta_0(P)$*

$$\limsup_{n \rightarrow \infty} \sup_{x \in \mathbf{R}} \{J_b(x, \theta, P) - J_n(x, \theta, P)\} \leq 0 . \quad (11)$$

Then, \mathcal{C}_n defined by (10) satisfies (2).

PROOF: Fix $\theta \in \Theta_0(P)$. By Theorem 3.1(i), we have immediately that

$$\liminf_{n \rightarrow \infty} P\{a_n\hat{Q}_n(\theta) \leq \hat{d}_n(\theta, 1 - \alpha)\} \geq 1 - \alpha .$$

Thus, by the definition (10), it follows that (2) holds. ■

Remark 3.2 It may be useful to consider confidence regions based on critical values given by

$$\tilde{d}_n(\theta, 1 - \alpha) = \inf\{x : \frac{1}{N_n} \sum_{1 \leq i \leq N_n} I\{a_b(\hat{Q}_{n,b,i}(\theta) - \hat{Q}_n(\theta)) \leq x\} \geq 1 - \alpha\} \quad (12)$$

instead of (9). Such confidence regions will satisfy the coverage requirement (2) under the assumptions of Theorem 3.2 and the additional assumption that $a_b/a_n \rightarrow 0$. The critical values (12) will typically be smaller than the ones given by (9), so the resulting confidence regions will typically be smaller as well. ■

Remark 3.3 Because $\binom{n}{b}$ may be large, it is often more practical to use the following approximation to (9). Let B_n be a positive sequence of numbers tending to ∞ as $n \rightarrow \infty$ and let I_1, \dots, I_{B_n} be chosen randomly with or without replacement from the numbers $1, \dots, N_n$. Then, it follows from Corollary 2.4.1 of Politis et al. (1999) that one may approximate (9) by

$$\inf\{x : \frac{1}{B_n} \sum_{1 \leq i \leq B_n} I\{a_b\hat{Q}_{n,b,I_i}(\theta) \leq x\} \geq 1 - \alpha\}$$

without affecting the conclusions of Theorem 3.2. ■

We now revisit each of the examples described in Section 2 and provide conditions under which the hypothesis (11) of Theorem 3.2 holds. It follows that for each of these examples \mathcal{C}_n defined by (10) satisfies (2).

Example 3.1 (*One-Sided Mean*) Recall the setup of Example 2.1. Let $a_n = n$ and suppose P is such that the variance $\sigma^2(P)$ exists. We will verify that the required condition (11) holds for any such P . Consider $\theta \in \Theta_0(P)$. We may assume without loss of generality that $\mu(P) = 0$. First note that because $\hat{Q}_n(\theta) \geq 0$, we may restrict attention to $x \geq 0$ in (11). There are two cases to consider: $\theta > 0$ and $\theta = 0$. In the former case, we have for any $\epsilon > 0$ that $\bar{X}_n - \theta < \epsilon - \theta$ with probability approaching one. By choosing ϵ sufficiently small, we therefore have that $a_n(\hat{Q}_n(\theta) - Q(\theta, P)) = n(\bar{X}_n - \theta)_+^2 = 0$ with probability approaching one. Hence, (11) holds trivially. In the latter case, $a_n(\hat{Q}_n(\theta) - Q(\theta, P)) = n(\bar{X}_n)_+^2$. Thus, for $x \geq 0$,

$$J_n(x, \theta, P) = P\{n(\bar{X}_n)_+^2 \leq x\} = P\{\sqrt{n}\bar{X}_n \leq \sqrt{x}\}.$$

Since $\sqrt{n}\bar{X}_n \xrightarrow{\mathcal{L}} \Phi_{\sigma(P)}$, where $\Phi_{\sigma(P)}$ is the distribution of a mean-zero normal random variable with variance $\sigma^2(P)$, Polya's Theorem implies that

$$\sup_{x \in \mathbf{R}} |P\{\sqrt{n}\bar{X}_n \leq x\} - \Phi_{\sigma(P)}(x)| \rightarrow 0.$$

The desired condition (11) now follows from an appeal to the triangle inequality. ■

Example 3.2 (*Two-Sided Mean*) Recall the setup of Example 2.2. Let $a_n = n$ and suppose P is such that the variances $\sigma_X^2(P)$ and $\sigma_Y^2(P)$ exist and $\mu_X(P) \leq \mu_Y(P)$. Denote by $\sigma_{X,Y}(P)$ the covariance of X_i and Y_i under P . We will verify that the required condition (11) holds for any such P . Consider $\theta \in \Theta_0(P)$. First note, as before, that because $\hat{Q}_n(\theta) \geq 0$, we may restrict attention to $x \geq 0$ in (11). There are two cases to consider: either $\mu_X(P) < \mu_Y(P)$ or $\mu_X(P) = \mu_Y(P)$.

If $\mu_X(P) < \mu_Y(P)$, then we must consider the case in which $\theta \in (\mu_X(P), \mu_Y(P))$ and the case in which $\theta \in \{\mu_X(P), \mu_Y(P)\}$ separately. In the former case, it is easy to see that $a_n(\hat{Q}_n(\theta) - Q(\theta, P)) = n(\bar{X}_n - \theta)_+^2 + n(\theta - \bar{Y}_n)_+^2 = 0$ with probability approaching one. Hence, (11) holds trivially. In the latter case, we may assume without loss of generality that $\theta = \mu_X(P)$ and $\mu_X(P) = 0$; the case in which $\theta = \mu_Y(P)$ is symmetric. In this case, $a_n(\hat{Q}_n(\theta) - Q(\theta, P)) = n(\bar{X}_n)_+^2 + n(-\bar{Y}_n)_+^2$. Because $\mu_Y(P) > 0$, we have that with probability tending to 1, $a_n(\hat{Q}_n(\theta) - Q(\theta, P)) = n(\bar{X}_n)_+^2$. The desired condition (11) now follows from the reasoning given in Example 3.1.

If $\mu_X(P) = \mu_Y(P)$, then we may assume without loss of generality that this common value is 0. Then, $a_n(\hat{Q}_n(\theta) - Q(\theta, P)) = n(\bar{X}_n)_+^2 + n(-\bar{Y}_n)_+^2$. Thus, for $x \geq 0$,

$$J_n(x, \theta, P) = P\{n(\bar{X}_n)_+^2 + n(-\bar{Y}_n)_+^2 \leq x\} = P\{(\sqrt{n}\bar{X}_n, -\sqrt{n}\bar{Y}_n)' \in S_x\},$$

where S_x is the appropriate convex set. Note that

$$(\sqrt{n}\bar{X}_n, -\sqrt{n}\bar{Y}_n)' \xrightarrow{\mathcal{L}} \Phi_{\sigma_X(P), \sigma_Y(P), -\sigma_{X,Y}(P)},$$

where $\Phi_{\sigma_X(P), \sigma_Y(P), -\sigma_{X,Y}(P)}$ is a mean-zero bivariate normal distribution with variances $\sigma_X^2(P)$ and $\sigma_Y^2(P)$ and covariance $-\sigma_{X,Y}(P)$. It follows from Theorem 2.11 of Bhattacharya and Rao (1976) that

$$\sup_{S \in \mathcal{S}} |P\{(\sqrt{n}\bar{X}_n, -\sqrt{n}\bar{Y}_n)' \in S\} - \Phi_{\sigma_X(P), \sigma_Y(P), -\sigma_{X,Y}(P)}(S)| \rightarrow 0,$$

where \mathcal{S} is the set of all convex subsets of \mathbf{R}^2 . Hence, the desired condition (11) again follows from an appeal to the triangle inequality. ■

Example 3.3 (*Regression with Interval Outcomes*) Recall the setup of Example 2.3. Let $a_n = n$ and let $\{x_1, \dots, x_J\}$ be a set of vectors in \mathbf{R}^k . For $l \in \{1, 2\}$ and $j \in \{1, \dots, J\}$, let

$$\tau_l(x_j, P) = E_P\{Y_{l,i} | X_i = x_j\} \tag{13}$$

$$\hat{\tau}_l(x_j) = \frac{1}{n(x_j)} \sum_{1 \leq i \leq n: X_i = x_j} Y_{l,i}, \tag{14}$$

where $n(x_j) = |\{1 \leq i \leq n : X_i = x_j\}|$. Let

$$\hat{Q}_n(\theta) = \sum_{1 \leq j \leq J} \frac{n(x_j)}{n} \{(\hat{\tau}_1(x_j) - x_j'\theta)_+^2 + (x_j'\theta - \hat{\tau}_2(x_j))_+^2\}. \tag{15}$$

Suppose P is such that (i) $\text{supp}_P(X_i) = \{x_1, \dots, x_J\}$ and (ii) the variances of Y_1 and Y_2 , $\sigma_1^2(P)$ and $\sigma_2^2(P)$, exist.

Note that it follows from the above assumptions on P that the vector whose components are given by $\sqrt{n(x_j)}(\hat{\tau}_l(x_j) - \tau_l(x_j, P))$ for $j = 1, \dots, J$ converges in distribution to a multivariate normal distribution.

We will now verify that the required condition (11) holds for this setup. Consider $\theta \in \Theta_0(P)$. First note, as before, that because $\hat{Q}_n(\theta) \geq 0$, we may restrict attention to $x \geq 0$ in (11). There are two cases to consider: either $\theta \in \text{int}(\Theta_0(P))$ or $\theta \in \partial\Theta_0(P)$. In the former case, $\tau_1(x_j, P) < x_j'\theta < \tau_2(x_j, P)$ for all j . Thus, it is easy to see that $a_n\hat{Q}_n(\theta) = 0$ with

probability tending to 1. Hence, (11) holds trivially. In the latter case, let I_l denote the set of j indices such that $\pi_l(x_j, P) = x'_j \theta$. Note that at most one of I_1 or I_2 may be empty. Then, with probability approaching one, we have that

$$a_n \hat{Q}_n(\theta) = \sum_{j \in I_1} (\sqrt{n(x_j)}(\hat{\tau}_1(x_j) - x'_j \theta))_+^2 + \sum_{j \in I_2} (\sqrt{n(x_j)}(x'_j \theta - \hat{\tau}_2(x_j)))_+^2 .$$

Let $\hat{W}_n(\theta)$ denote the vector whose components are given by $\sqrt{n(x_j)}(\hat{\tau}_1(x_j) - x'_j \theta)$ for $j \in I_1$ and $\sqrt{n(x_j)}(x'_j \theta - \hat{\tau}_2(x_j))$ for $j \in I_2$ and let $W(\theta)$ denote the limiting multivariate normal distribution of $\hat{W}_n(\theta)$. For $x \geq 0$, we have that

$$J_n(x, \theta, P) = P\{\hat{W}_n(\theta) \in S_x\} ,$$

where S_x is the appropriate set in \mathcal{S} , defined by (24). It follows from Theorem 2.11 of Bhattacharya and Rao (1976) that

$$\sup_{S \in \mathcal{S}} |P\{\hat{W}_n(\theta) \in S\} - P\{W(\theta) \in S\}| \rightarrow 0 ,$$

where \mathcal{S} is the set of all convex subsets of $\mathbf{R}^{|I_1|+|I_2|}$. Hence, the desired condition (11) again follows from an appeal to the triangle inequality. ■

Example 3.4 (*Moment Inequalities*) Recall the setup of Example 2.4. Let $a_n = n$ and suppose P is such that the variance of the vector whose components are given by $g_j(X_i, \theta)$ for $j = 1, \dots, m$ exists for all $\theta \in \Theta$. It follows from this assumption that the vector whose components are given by

$$\frac{1}{\sqrt{n}} \sum_{1 \leq i \leq n} (g_j(X_i, \theta) - E_P\{g_j(X_i, \theta)\})$$

for $j = 1, \dots, m$ converges in distribution to a multivariate normal random variable for each $\theta \in \Theta$. We will verify that the required condition (11) holds for any such P .

Consider $\theta \in \Theta_0(P)$. There are two cases to consider: either $\theta \in \text{int}(\Theta_0(P))$ or $\theta \in \partial\Theta_0(P)$. In the former case, $E_P\{g_j(X_i, \theta)\} < 0$ for all j , so $a_n \hat{Q}_n(\theta) = 0$ with probability tending to 1. Hence, (11) holds trivially. In the latter case, let I denote the set of j indices such that $E_P\{g_j(X_i, \theta)\} = 0$. Note that I must be non-empty. Then, with probability approaching one, we have that

$$a_n \hat{Q}_n(\theta) = \sum_{j \in I} \left(\frac{1}{\sqrt{n}} \sum_{1 \leq i \leq n} g_j(X_i, \theta) \right)_+^2 .$$

Let $\hat{W}_n(\theta)$ denote the vector whose components are given by

$$\frac{1}{\sqrt{n}} \sum_{1 \leq i \leq n} g_j(X_i, \theta)$$

for $j \in I$ and let $W(\theta)$ denote the limiting multivariate normal distribution. For $x \geq 0$, we have that

$$J_n(x, \theta, P) = P\{\hat{W}_n(\theta) \in S_x\},$$

where S_x is the appropriate convex set. It follows from Theorem 2.11 of Bhattacharya and Rao (1976) that

$$\sup_{S \in \mathcal{S}} |P\{\hat{W}_n(\theta) \in S\} - P\{W(\theta) \in S\}| \rightarrow 0,$$

where \mathcal{S} is the set of all convex subsets of $\mathbf{R}^{|I|}$. Hence, the desired condition (11) again follows from an appeal to the triangle inequality. ■

Remark 3.4 In each of the examples from Section 2, there are many other choices of $Q(\theta, P)$ sharing the feature that

$$\Theta_0(P) = \arg \min_{\theta \in \Theta} Q(\theta, P). \quad (16)$$

Consider Example 2.4. Let $g(x, \theta)$ denote the vector whose components are given by $g_j(x, \theta)$, $j = 1, \dots, m$ and extend the definition of $(\cdot)_+$ on the real line to vectors in the natural way. It is easy to see that

$$Q(\theta, P) = (E_P\{g(X, \theta)\})'_+ W(\theta) (E_P\{g(X, \theta)\})_+,$$

where $W(\theta)$ is a positive definite matrix for each $\theta \in \Theta$ also satisfies (16) for this example. The objective function given by

$$Q(\theta, P) = \max_{1 \leq j \leq m} w_j(\theta) (E_P\{g_j(X, \theta)\})_+,$$

where $w_j(\theta)$, $j = 1, \dots, m$ are positive for each $\theta \in \Theta$, also satisfies (16) for this example. Each of these choices of $Q(\theta, P)$ will lead to different test statistics $a_n \hat{Q}_n(\theta)$, which in turn will lead to different tests of the individual null hypotheses in (8). Of course, more powerful tests will reject more false hypotheses and thereby lead to smaller confidence regions, so the choice of $Q(\theta, P)$ is an important direction for future research. ■

Remark 3.5 If it was the case that under the null hypothesis H_θ the statistic $a_n \hat{Q}_n(\theta)$ converged in distribution to a random variable that was continuous at its $1 - \alpha$ quantile, then it would follow immediately from Theorem 2.2.1 of Politis et al. (1999) that \mathcal{C}_n defined by (10) satisfies (2). Unfortunately, in the examples from Section 2, this assertion does not always hold. Indeed, in many instances, for $\theta \in \text{int}(\Theta_0(P))$, it is true that $a_n \hat{Q}_n(\theta) = 0$ with probability tending to 1. For this reason, it is necessary to appeal to Theorem 3.1 to establish the desired conclusions. ■

Remark 3.6 Chernozhukov et al. (2007) also consider the problem of constructing confidence regions \mathcal{C}_n satisfying (2). In an earlier version of their paper (Chernozhukov et al., 2004), they propose a construction that relies upon an initial estimate of $\Theta_0(P)$. Their estimate of $\Theta_0(P)$ is given by

$$\hat{\Theta}_{0,n} = \{\theta \in \Theta : \hat{Q}_n(\theta) < \epsilon_n\} ,$$

where ϵ_n is a positive sequence of constants tending to zero sufficiently slowly. Because of this rate restriction, they are able to show that $\hat{\Theta}_{0,n}$ satisfies

$$P\{\Theta_0(P) \subseteq \hat{\Theta}_{0,n}\} \rightarrow 1 .$$

They then propose the confidence region

$$\mathcal{C}_n = \{\theta \in \Theta : a_n \hat{Q}_n(\theta) \leq \sup_{\theta' \in \hat{\Theta}_{0,n}} \hat{d}_n(\theta', 1 - \alpha)\} ,$$

where $\hat{d}_n(\theta, 1 - \alpha)$ is given by (9). Of course, such a construction is clearly more conservative than the one given by (10). Moreover, the level of ϵ_n is completely arbitrary and the confidence region resulting from application of their method may thus be very large or very small depending on the choice of ϵ_n . For this reason, after learning of the results in an earlier version of our paper, Chernozhukov et al. (2007) suggested the confidence region

$$\mathcal{C}_n = \{\theta \in \Theta : a_n \hat{Q}_n(\theta) \leq \min\{\sup_{\theta' \in \hat{\Theta}_{0,n}} \hat{d}_n(\theta', 1 - \alpha), \hat{d}_n(\theta, 1 - \alpha)\} ,$$

where $\hat{d}_n(\theta, 1 - \alpha)$ is again given by (9). In a first-order asymptotic sense, this confidence region is equivalent to our confidence region (10). Moreover, the extent to which this construction is smaller than ours in finite-samples depends again on the level of ϵ_n , which is left unspecified. For this reason, we prefer our original construction, which does not suffer from such arbitrariness. ■

Remark 3.7 In certain applications, it may be of interest to test the null hypothesis that there exists some point $\theta \in \Theta$ for which $Q(\theta, P) = 0$. Such a test may be viewed as a specification test of the model. One may test such a hypothesis by first constructing a confidence region for identifiable parameters and then rejecting the hypothesis if the confidence region is empty. To see this, note that under the null hypothesis, $\Theta_0(P) \neq \emptyset$. Thus, there is some $\theta' \in \Theta_0(P)$. Under the assumptions of Theorem 3.2, it must be the case that

$$\liminf_{n \rightarrow \infty} P\{\theta' \in \mathcal{C}_n\} \geq 1 - \alpha .$$

Thus, under the null hypothesis \mathcal{C}_n will be nonempty with probability at least $1 - \alpha$ asymptotically. ■

Remark 3.8 Our construction of critical values for testing $H_\theta : Q(\theta, P) = 0$ versus $K_\theta : Q(\theta, P) > 0$ has used subsampling. We now show by example that the bootstrap may fail to approximate the distribution of the test statistic, $a_n \hat{Q}_n(\theta)$, under the null hypothesis.

To this end, recall the setup of Example 2.1. Let $a_n = n$ and suppose that $P = N(0, 1)$. Note that $Q(\theta, P) = 0$ at $\theta = 0$, so the null hypothesis is true when $\theta = 0$. Our argument follows the one given by Andrews (2000) to show that the bootstrap is inconsistent when the parameter is on the boundary of the parameter space. Note that $a_n \hat{Q}_n(0) = (\sqrt{n} \bar{X}_n)_+^2 \sim (Z)_+^2$, where Z is a standard normal random variable. Denote by X_i^* , $i = 1, \dots, n$ an i.i.d. sequence of random variables with distribution \hat{P}_n given by the empirical distribution of the original observations X_i , $i = 1, \dots, n$. Define

$$\hat{Q}_n^*(\theta) = (\bar{X}_n^* - \theta)_+^2$$

and consider the event

$$A_c = \{\omega : \liminf_{n \rightarrow \infty} \sqrt{n} \bar{X}_n < -c\} \quad (17)$$

for $c \in (0, \infty)$. The Law of the Iterated Logarithm asserts that

$$\limsup_{n \rightarrow \infty} \frac{\sqrt{n} \bar{X}_n}{\sqrt{2 \log \log n}} = 1 \quad P\text{-a.s. .}$$

This in turn implies that

$$\{\sqrt{n} \bar{X}_n > \sqrt{2 \log \log n} \text{ i.o.}\} \quad P\text{-a.s.}$$

and by symmetry that

$$\{\sqrt{n} \bar{X}_n < -\sqrt{2 \log \log n} \text{ i.o.}\} \quad P\text{-a.s.}$$

It follows that $P\{A_c\} = 1$.

Now consider the bootstrap approximation to the distribution of $a_n \hat{Q}_n(\theta)$ at $\theta = 0$, which is given by

$$\mathcal{L}(a_n(\hat{Q}_n^*(0) - \hat{Q}_n(0)) | X_i, i = 1, \dots, n) . \quad (18)$$

This approximation mimics the hypothesis that $Q(0, P) = 0$ by centering $\hat{Q}_n^*(0)$ about $\hat{Q}_n(0)$. For $\omega \in A_c$, consider a subsequence n_k of $n \geq 1$ for which $\sqrt{n_k} \bar{X}_{n_k}(\omega) < -c$ for all k . For such a subsequence, we have, conditionally on $X_i(\omega), i = 1, \dots, n$, that

$$\begin{aligned} n_k(\hat{Q}_{n_k}^*(0) - \hat{Q}_{n_k}(0)(\omega)) &= (\sqrt{n_k} \bar{X}_{n_k}^*)_+^2 \\ &= (\sqrt{n_k}(\bar{X}_{n_k}^* - \bar{X}_{n_k}(\omega)) + \sqrt{n_k} \bar{X}_{n_k}(\omega))_+^2 \\ &\leq (\sqrt{n_k}(\bar{X}_{n_k}^* - \bar{X}_{n_k}(\omega)) - c)_+^2 \\ &\xrightarrow{\mathcal{L}} (Z - c)_+^2 , \end{aligned}$$

where Z is again a standard normal random variable. The first equality and the inequality follow from the fact that $\sqrt{n_k}\bar{X}_{n_k}(\omega) < -c < 0$. It follows that the bootstrap fails to approximate the distribution of $a_n\hat{Q}_n(\theta)$ at $\theta = 0$.

It is worthwhile to consider the actual probability of a Type 1 error if the bootstrap approximation above is used instead of subsampling. To this end, first note that the probability of such an error is given by

$$P\{a_n\hat{Q}_n(0) > \hat{c}_{\text{boot},n}(1 - \alpha)\}, \quad (19)$$

where $\hat{c}_{\text{boot},n}(1 - \alpha)$ is the $1 - \alpha$ quantile of (18). Recall that $a_n\hat{Q}_n(0) = (\sqrt{n}\bar{X}_n)_+^2 \sim (Z)_+^2$, where Z is a standard normal random variable. Note that

$$a_n(\hat{Q}_n^*(0) - \hat{Q}_n(0)) = (\sqrt{n}(\bar{X}_n^* - \bar{X}_n) + \sqrt{n}\bar{X}_n)_+^2 - (\sqrt{n}\bar{X}_n)_+^2.$$

Therefore, (18) converges in distribution to

$$\mathcal{L}((Z^* + Z)_+^2 - (Z)_+^2 | Z), \quad (20)$$

where Z^* is a standard normal random variable distributed independently of Z . The probability (19) is therefore asymptotically equal to

$$P\{(Z)_+^2 > c(1 - \alpha|Z)\},$$

where $c(1 - \alpha|Z)$ is the $1 - \alpha$ quantile of (20). Using this representation, we simulate the asymptotic value of (19) and find that the probability of Type 1 error is controlled, but too conservatively for practical purposes. In fact, with 2000 simulations, we estimate that (19) is asymptotically equal to 0 to three significant digits for both $\alpha = .1$ and $\alpha = .05$. ■

3.3 Uniform Consistency in Level

We now turn our attention to providing conditions under which \mathcal{C}_n defined by (10) satisfies (3). Recall from the discussion in the introduction that inferences made on the basis of asymptotic arguments with confidence regions that are not uniformly consistent in level may be very misleading in the sense that asymptotic approximations may be poor for arbitrarily large sample sizes.

Theorem 3.3 *Let $X_i, i = 1, \dots, n$, be an i.i.d. sequence of random variables with distribution P and let $b = b_n < n$ be a sequence of positive integers tending to infinity, but satisfying*

$b/n \rightarrow 0$. Denote by $J_n(\cdot, \theta, P)$ the distribution of $a_n(\hat{Q}_n(\theta) - Q(\theta, P))$. Suppose that

$$\limsup_{n \rightarrow \infty} \sup_{\theta \in \Theta} \sup_{P \in \mathbf{P}: \theta \in \Theta_0(P)} \sup_{x \in \mathbf{R}} \{J_b(x, \theta, P) - J_n(x, \theta, P)\} \leq 0 . \quad (21)$$

Then, \mathcal{C}_n defined by (10) satisfies (3).

PROOF: By Theorem 3.1(iv), we have immediately that

$$\liminf_{n \rightarrow \infty} \inf_{\theta \in \Theta} \inf_{P \in \mathbf{P}: \theta \in \Theta_0(P)} P\{a_n \hat{Q}_n(\theta) \leq \hat{d}_n(\theta, 1 - \alpha)\} \geq 1 - \alpha .$$

Thus, by the definition (10), it follows that (3) holds. ■

We now apply Theorem 3.3 to construct confidence regions satisfying the coverage requirement (3) for each of the examples described in Section 2.

Example 3.5 (One-Sided Mean) Recall the setup of Example 2.1. We will now use Theorem 3.3 to show that for this example \mathcal{C}_n defined by (10) satisfies (3) for a large class of distributions \mathbf{P} . To this end, let $a_n = n$ and let \mathbf{P} be a set of distributions satisfying

$$\lim_{\lambda \rightarrow \infty} \sup_{P \in \mathbf{P}} E_P \left[\frac{|X - \mu(P)|^2}{\sigma^2(P)} I \left\{ \frac{|X - \mu(P)|}{\sigma(P)} > \lambda \right\} \right] = 0 . \quad (22)$$

We will argue by contradiction that the required condition (21) holds. If the result were false, then there would exist a subsequence n_j and a corresponding sequence $(\theta_{n_j}, P_{n_j}) \in \Theta \times \mathbf{P}$ such that $\mu(P_{n_j}) \leq \theta_{n_j}$ and

$$\sup_{x \in \mathbf{R}} \{J_{b_{n_j}}(x, \theta_{n_j}, P_{n_j}) - J_{n_j}(x, \theta_{n_j}, P_{n_j})\} \rightarrow \delta \quad (23)$$

for some $\delta > 0$. Since $a_n \hat{Q}_n(\theta_{n_j}) \geq 0$, we may restrict attention to $x \geq 0$. We may thus rewrite (23) as

$$\sup_{x \geq 0} \{P_{n_j} \{(Z_{b_{n_j}}(P_{n_j}) + T_{b_{n_j}}(\theta_{n_j}, P_{n_j}))_+^2 \leq x\} - P_{n_j} \{(Z_{n_j}(P_{n_j}) + T_{n_j}(\theta_{n_j}, P_{n_j}))_+^2 \leq x\}\} ,$$

where

$$\begin{aligned} Z_n(P) &= \sqrt{n}(\bar{X}_n - \mu(P)) \\ T_n(\theta, P) &= \sqrt{n}(\mu(P) - \theta) . \end{aligned}$$

Since $b_{n_j} \leq n_j$ and $\mu(P_{n_j}) \leq \theta_{n_j}$, we have that $T_{b_{n_j}}(\theta_{n_j}, P_{n_j}) \geq T_{n_j}(\theta_{n_j}, P_{n_j})$. Thus, (23) is bounded above by

$$\sup_{x \geq 0} \{P_{n_j} \{(Z_{b_{n_j}}(P_{n_j}) + T_{n_j}(\theta_{n_j}, P_{n_j}))_+^2 \leq x\} - P_{n_j} \{(Z_{n_j}(P_{n_j}) + T_{n_j}(\theta_{n_j}, P_{n_j}))_+^2 \leq x\}\} .$$

By redefining x appropriately and absorbing the term $T_{n_j}(\theta_{n_j}, P_{n_j})$ into x , it follows that

$$\sup_{x \in \mathbf{R}} \{P_{n_j}\{Z_{b_{n_j}}(P_{n_j}) \leq x\} - P_{n_j}\{Z_{n_j}(P_{n_j}) \leq x\}\} \not\rightarrow 0 .$$

By Lemma 11.4.1 of Lehmann and Romano (2005), this yields the desired contradiction. ■

Example 3.6 (*Two-Sided Mean*) Recall the setup of Example 2.2. As in Example 3.5, we may use Theorem 3.3 to show that for this example \mathcal{C}_n defined by (10) satisfies (3) for a large class of distributions \mathbf{P} . To this end, let $a_n = n$, let \mathbf{P} be a set of bivariate distributions such that the marginal distributions satisfy (22) and $\mu_X(P) \leq \mu_Y(P)$. As before, we will argue by contradiction that the required condition (21) holds, but we will require the following result, which generalizes Lemma 11.4.1 of Lehmann and Romano (2005), in order to do so. For the sake of continuity, the proof of the result is found in the appendix.

Lemma 3.1 *Let \mathbf{P} be a set of distributions on \mathbf{R}^k such that the marginal distributions satisfy (22). Let $X_i = (X_{1,i}, \dots, X_{k,i}), i = 1, \dots, n$ be an i.i.d. sequence of random variables with distribution $P \in \mathbf{P}$. Denote by $\Phi_V(\cdot)$ the probability distribution of a multivariate normal random variable with mean 0 and variance V and by $\Sigma(P)$ the variance of P . Then,*

$$\sup_{P \in \mathbf{P}} \sup_{S \in \mathcal{S}} |P\{(\sqrt{n}(\bar{X}_{1,n} - \mu_1(P)), \dots, \sqrt{n}(\bar{X}_{k,n} - \mu_k(P))) \in S\} - \Phi_{\Sigma(P)}(S)| \rightarrow 0 ,$$

where

$$\mathcal{S} = \{S \subseteq \mathbf{R}^k : S \text{ convex and } \Phi_V(\partial S) = 0 \text{ for all p.s.d. } V \text{ s.t. } V_{i,i} = 1 \text{ for } 1 \leq i \leq k\} . \quad (24)$$

PROOF: See appendix. ■

We now return to verifying (21). If the result were false, then there would exist a subsequence n_j and a corresponding sequence $(\theta_{n_j}, P_{n_j}) \in \Theta \times \mathbf{P}$ such that $\mu_X(P_{n_j}) \leq \theta_{n_j} \leq \mu_Y(P_{n_j})$ and

$$\sup_{x \in \mathbf{R}} \{J_{b_{n_j}}(x, \theta_{n_j}, P_{n_j}) - J_{n_j}(x, \theta_{n_j}, P_{n_j})\} \rightarrow \delta \quad (25)$$

for some $\delta > 0$. Note that $a_{b_{n_j}} \hat{Q}_{b_{n_j}}(\theta_{n_j})$

$$(Z_{1,b_{n_j}}(P_{n_j}) + T_{1,b_{n_j}}(\theta_{n_j}, P_{n_j}))_+^2 + (Z_{2,b_{n_j}}(P_{n_j}) + T_{2,b_{n_j}}(\theta_{n_j}, P_{n_j}))_+^2 ,$$

where

$$\begin{aligned}
Z_{1,n}(P) &= \sqrt{n}(\bar{X}_n - \mu_X(P)) \\
Z_{2,n}(P) &= \sqrt{n}(\mu_Y(P) - \bar{Y}_n) \\
T_{1,n}(\theta, P) &= \sqrt{n}(\mu_X(P) - \theta) \\
T_{2,n}(\theta, P) &= \sqrt{n}(\theta - \mu_Y(P)) .
\end{aligned}$$

Let $\tilde{J}_{b_{n_j}}(x, \theta_{n_j}, P_{n_j})$ be the distribution of the statistic

$$(Z_{1,b_{n_j}}(P_{n_j}) + T_{1,n_j}(\theta_{n_j}, P_{n_j}))_+^2 + (Z_{2,b_{n_j}}(P_{n_j}) + T_{2,n_j}(\theta_{n_j}, P_{n_j}))_+^2 .$$

Since $T_{k,b_{n_j}}(\theta_{n_j}, P_{n_j}) \geq T_{k,n_j}(\theta_{n_j}, P_{n_j})$, we have that

$$\tilde{J}_{b_{n_j}}(x, \theta_{n_j}, P_{n_j}) \geq J_{b_{n_j}}(x, \theta_{n_j}, P_{n_j}) .$$

Thus, (25) implies that

$$\sup_{x \in \mathbf{R}} \{ \tilde{J}_{b_{n_j}}(x, \theta_{n_j}, P_{n_j}) - J_{n_j}(x, \theta_{n_j}, P_{n_j}) \} \not\rightarrow 0 . \quad (26)$$

Since $a_n \hat{Q}_n(\theta) \geq 0$, we may restrict attention to $x \geq 0$. For such x , $\tilde{J}_{b_{n_j}}(x, \theta_{n_j}, P_{n_j})$ is simply the probability under P_{n_j} that the vector $(Z_{1,b_{n_j}}(P_{n_j}), Z_{2,b_{n_j}}(P_{n_j}))'$ lies in a set $S_x \in \mathcal{S}$, where \mathcal{S} is defined by (24). Importantly, $J_{n_j}(x, \theta_{n_j}, P_{n_j})$ is simply the probability under P_{n_j} that the vector $(Z_{1,n_j}(P_{n_j}), Z_{2,n_j}(P_{n_j}))'$ lies in the same set. But, by Lemma 3.1, however, we know that

$$\sup_{S \in \mathcal{S}} |P_{n_j} \{ (Z_{1,n_j}(P_{n_j}), Z_{2,n_j}(P_{n_j}))' \in S \} - \Phi_{\sigma_X(P_{n_j}), \sigma_Y(P_{n_j}), \rho(P_{n_j})}(S)| \rightarrow 0 ,$$

and

$$\sup_{S \in \mathcal{S}} |P_{n_j} \{ (Z_{1,b_{n_j}}(P_{n_j}), Z_{2,b_{n_j}}(P_{n_j}))' \in S \} - \Phi_{\sigma_X(P_{n_j}), \sigma_Y(P_{n_j}), \rho(P_{n_j})}(S)| \rightarrow 0 .$$

An appeal to the triangle inequality yields the desired contradiction to (26). ■

Remark 3.9 Recall that Imbens and Manski (2004) analyze a special case of Example 2.2 in which $X_i = W_i D_i$ and $Y_i = W_i D_i + 1 - D_i$ where $W_i \in [0, 1]$ and $D_i \in \{0, 1\}$. The motivation for their study of this problem stems from considering

$$\mathcal{C}_n = \left[\bar{X}_n - \frac{z_{1-\alpha} \hat{\sigma}_{X,n}}{\sqrt{n}}, \bar{Y}_n + \frac{z_{1-\alpha} \hat{\sigma}_{Y,n}}{\sqrt{n}} \right] , \quad (27)$$

where $\hat{\sigma}_{X,n}$ and $\hat{\sigma}_{Y,n}$ are the usual consistent estimators of $\sigma_X(P)$ and $\sigma_Y(P)$ respectively, and $z_{1-\alpha}$ is the $1 - \alpha$ quantile of the standard normal distribution. For any P such that $P\{D_i =$

$1\} < 1$, \mathcal{C}_n defined by (27) satisfies (2). To see this, consider the two cases $\theta \in \text{int}(\Theta_0(P))$ and $\theta \in \partial\Theta_0(P)$ separately. In the former case, $\bar{X}_n < \theta < \bar{Y}_n$ with probability approaching 1, which implies that $\theta \in \mathcal{C}_n$ with probability approaching 1 as well. In the latter case, $\theta = \mu_X(P)$ or $\theta = \mu_Y(P)$. Suppose $\theta = \mu_X(P)$; the case in which $\theta = \mu_Y(P)$ is completely symmetric. Then, $\theta < \bar{Y}_n$ with probability approaching 1 and $\theta > \bar{X}_n - \frac{z_{1-\alpha}\hat{\sigma}_{X,n}}{\sqrt{n}}$ with probability at least $1 - \alpha$ asymptotically. We therefore have that $\theta \in \mathcal{C}_n$ with probability at least $1 - \alpha$ asymptotically, as desired. Imbens and Manski (2004) show, however, that this convergence is not uniform over all P satisfying $P\{D_i = 1\} < 1$. Specifically, they show that for every sample size n there is a P with $P\{D_i = 1\}$ sufficiently close, but not equal, to 1 and a $\theta \in \Theta_0(P)$ for which

$$P\{\theta \in \mathcal{C}_n\} \approx 1 - 2\alpha .$$

To rectify this shortcoming, Imbens and Manski (2004) propose an alternative to \mathcal{C}_n defined by (27) which satisfies the uniform coverage requirement (3) for $\mathbf{P} = \mathbf{P}'$ such that

$$\begin{aligned} \inf_{P \in \mathbf{P}'} P\{D_i = 1\} &> 0 \\ \inf_{P \in \mathbf{P}'} \sigma_{W_i|D_i=1}^2(P) &> 0 . \end{aligned}$$

We may apply the analysis of Example 3.6 and conclude immediately that \mathcal{C}_n defined by (10) also satisfies (3) for \mathbf{P} such that the distributions of X_i and Y_i defined above satisfy (22). This class of distributions is larger than the one considered by Imbens and Manski (2004). To see this, first note that

$$\begin{aligned} \inf_{P \in \mathbf{P}'} \sigma_{W_i D_i}^2(P) &\geq \inf_{P \in \mathbf{P}'} P\{D_i = 1\} \sigma_{W_i|D_i=1}^2(P) \\ &\geq \inf_{P \in \mathbf{P}'} P\{D_i = 1\} \inf_{P \in \mathbf{P}'} \sigma_{W_i|D_i=1}^2(P) > 0 . \end{aligned}$$

Since $W_i D_i$ is supported on a compact set, it follows that for any $\delta > 0$

$$\sup_{P \in \mathbf{P}'} E_P \left\{ \left(\frac{|W_i D_i - \mu_{W_i D_i}(P)|}{\sigma_{W_i D_i}(P)} \right)^{2+\delta} \right\} < \infty .$$

From the pointwise inequality

$$x^2 I\{x > \lambda\} \leq \frac{|x|^{2+\delta}}{\lambda^\delta}$$

for $\lambda > 0$ and $\delta > 0$, we therefore have that (22) holds for $W_i D_i$ and $\mathbf{P} = \mathbf{P}'$. Next note that to show that

$$\inf_{P \in \mathbf{P}'} \sigma_{W_i D_i + 1 - D_i}^2(P) > 0$$

it suffices to show that

$$\inf_{P \in \mathbf{P}'} \rho_{W_i D_i, 1 - D_i}(P) > -1 . \tag{28}$$

But, by direct calculation, we have that

$$\rho_{W_i D_i, 1-D_i}(P) = \frac{-1}{\sqrt{1 + \frac{\sigma_{W_i|D_i=1}^2(P)}{E_P\{W_i|D_i=1\}^2 P\{D_i=0\}}}} .$$

It now follows immediately that (28) holds. We may therefore argue as above for $W_i D_i$ to establish that (22) holds for $W_i D_i + 1 - D_i$ and $\mathbf{P} = \mathbf{P}'$. Hence, \mathbf{P}' is a subset of the set of distributions obtained by applying the analysis of Example 3.6.

It is worthwhile to note that the shortcoming of the construction given by (27) pointed out by Imbens and Manski (2004) disappears if we require further that the confidence region $\mathcal{C}_n = [\hat{L}_n, \hat{U}_n]$ is *equitailed* in the sense that for any $\theta \in \Theta_0(P)$

$$\begin{aligned} \limsup_{n \rightarrow \infty} P\{\theta < \hat{L}_n\} &\leq \frac{\alpha}{2} \\ \limsup_{n \rightarrow \infty} P\{\theta > \hat{U}_n\} &\leq \frac{\alpha}{2} . \end{aligned}$$

One such confidence region is given by (27) in which $z_{1-\alpha}$ is replaced by $z_{1-\frac{\alpha}{2}}$. Note that even if $P\{D_i = 1\} = 1$, so $\Theta_0(P)$ is a singleton, the confidence region \mathcal{C}_n defined in this way satisfies (2). ■

Example 3.7 (*Regression with Interval Outcomes*) Recall the setup of Example 2.3. Let $a_n = n$ and let $\{x_1, \dots, x_J\}$ be a set of vectors in \mathbf{R}^k . For $\ell \in \{1, 2\}$ and $j \in \{1, \dots, J\}$, let $\tau_\ell(x_j, P)$ and $\hat{\tau}_\ell(x_j)$ be given by (13) and (14), respectively. Let $\hat{Q}_n(\theta)$ be given by (15). Suppose \mathbf{P} is a set of distributions such that (i) $\text{supp}_P(X_i) = \{x_1, \dots, x_J\}$ for all $P \in \mathbf{P}$, (ii) for $\ell \in \{1, 2\}$ and $j \in \{1, \dots, J\}$

$$(Y_{\ell,i} - E_P\{Y_{\ell,i}|X_i = x_j\})I\{X_i = x_j\} \tag{29}$$

satisfies (22), and (iii)

$$\inf_{P \in \mathbf{P}} P\{X_i = x_j\} > 0 . \tag{30}$$

We now argue that the required condition (21) holds under these assumptions. Note that we may rewrite $a_n \hat{Q}_n(\theta)$ as

$$\sum_{1 \leq j \leq m} \{(\sqrt{n(x_j)}(\hat{\tau}_1(x_j) - x'_j \theta))_+^2 + (\sqrt{n(x_j)}(x'_j \theta - \hat{\tau}_2(x_j)))_+^2\} .$$

The analysis for this example will therefore be identical to the analysis for Example 3.8 presented below provided that we can show that a central limit theorem holds for the vector whose

components are given by $\sqrt{n(x_j)}(\hat{\tau}_\ell(x_j) - x'_j\theta)$ uniformly in $P \in \mathbf{P}$. To see that this holds under the above assumptions, write

$$\sqrt{n(x_j)}(\hat{\tau}_\ell(x_j) - \tau_\ell(x_j, P)) = \sqrt{\frac{n}{n(x_j)}} \frac{1}{\sqrt{n}} \sum_{1 \leq i \leq n} (Y_{\ell,i} - E_P\{Y_{\ell,i}|X_i = x_j\})I\{X_i = x_j\}.$$

From Lemma 3.1, we have immediately that such a central limit theorem holds for the vector whose components are given by (29). From Chebychev's inequality and the fact that $I\{X_i = x_j\}$ is bounded, we also have that $n(x_j)/n$ converges in probability to $P\{X_i = x_j\}$ uniformly in $P \in \mathbf{P}$. It follows from (30) that $\sqrt{n/n(x_j)}$ converges in probability to $1/P\{X_i = x_j\}$ uniformly in $P \in \mathbf{P}$. The desired conclusion now follows from Slutsky's Theorem. ■

Example 3.8 (*Moment Inequalities*) Recall the setup of Example 2.4. Let $a_n = n$. Denote by $\Sigma(\theta, P)$ the variance of $(g_1(X_i, \theta), \dots, g_m(X_i, \theta))$ under P and by $\sigma_j^2(\theta, P)$ the j th diagonal element of $\Sigma(\theta, P)$. Suppose \mathbf{P} is a set of distributions such that for each $1 \leq j \leq m$

$$\limsup_{\lambda \rightarrow 0} \sup_{P \in \mathbf{P}} \sup_{\theta \in \Theta_0(P)} E_P \left[\frac{(g_j(X_i, \theta) - E_P\{g_j(X_i, \theta)\})^2}{\sigma_j^2(\theta, P)} I\left\{ \frac{|g_j(X_i, \theta) - E_P\{g_j(X_i, \theta)\}|}{\sigma_j(\theta, P)} > \lambda \right\} \right] = 0.$$

We now argue by contradiction that the required condition (21) holds under this assumption. If the result were false, then there would exist a subsequence n_k and a corresponding sequence $(\theta_{n_k}, P_{n_k}) \in \Theta \times \mathbf{P}$ such that $\theta_{n_k} \in \Theta_0(P_{n_k})$ and

$$\sup_{x \in \mathbf{R}} \{J_{b_{n_k}}(x, \theta_{n_k}, P_{n_k}) - J_{n_k}(x, \theta_{n_k}, P_{n_k})\} \rightarrow \delta \quad (31)$$

for some $\delta > 0$. Note that $a_{b_{n_k}} \hat{Q}_{b_{n_k}}(\theta_{n_k})$ equals

$$\sum_{1 \leq j \leq m} (Z_{j, b_{n_k}}(\theta_{n_k}, P_{n_k}) + T_{j, b_{n_k}}(\theta_{n_k}, P_{n_k}))_+^2,$$

where

$$\begin{aligned} Z_{j,n}(\theta, P) &= \frac{1}{\sqrt{n}} \sum_{1 \leq i \leq n} (g_j(X_i, \theta) - E_P\{g_j(X_i, \theta)\}) \\ T_{j,n}(\theta, P) &= \sqrt{n} E_P\{g_j(X_i, \theta)\}. \end{aligned}$$

Let $\tilde{J}_{b_{n_k}}(x, \theta_{n_k}, P_{n_k})$ be the distribution of the statistic

$$\sum_{1 \leq j \leq m} (Z_{j, b_{n_k}}(\theta_{n_k}, P_{n_k}) + T_{j, b_{n_k}}(\theta_{n_k}, P_{n_k}))_+^2.$$

Since $b_{n_k} \leq n_k$ and $\theta_{n_k} \in \Theta_0(P_{n_k})$, we have that $T_{j, b_{n_k}}(\theta_{n_k}, P_{n_k}) \geq T_{j, n_k}(\theta_{n_k}, P_{n_k})$, which implies, in turn, that

$$\tilde{J}_{b_{n_k}}(x, \theta_{n_k}, P_{n_k}) \geq J_{n_k}(x, \theta_{n_k}, P_{n_k}).$$

Thus, (31) implies that

$$\sup_{x \in \mathbf{R}} \{ \tilde{J}_{b_{n_k}}(x, \theta_{n_k}, P_{n_k}) - J_{n_k}(x, \theta_{n_k}, P_{n_k}) \} \neq 0 . \quad (32)$$

Since $a_n \hat{Q}_n(\theta) \geq 0$, we may restrict attention to $x \geq 0$. For such x , $\tilde{J}_{b_{n_k}}(x, \theta_{n_k}, P_{n_k})$ is the probability under P_{n_j} that the vector $(Z_{1,b_{n_k}}(\theta_{n_k}, P_{n_k}), \dots, Z_{m,b_{n_k}}(\theta_{n_k}, P_{n_k}))'$ lies in a set $S_x \in \mathcal{S}$, where \mathcal{S} is defined by (24). Importantly, $J_{n_k}(x, \theta_{n_k}, P_{n_k})$ is simply the probability under P_{n_k} that the vector $(Z_{1,n_k}(\theta_{n_k}, P_{n_k}), \dots, Z_{m,n_k}(\theta_{n_k}, P_{n_k}))'$ lies in the same set. But, by Lemma 3.1, however, we know that

$$\sup_{S \in \mathcal{S}} |P_{n_k} \{ (Z_{1,n_k}(\theta_{n_k}, P_{n_k}), \dots, Z_{m,n_k}(\theta_{n_k}, P_{n_k}))' \in S \} - \Phi_{\Sigma(\theta_{n_k}, P_{n_k})}(S)| \rightarrow 0 ,$$

and

$$\sup_{S \in \mathcal{S}} |P_{n_k} \{ (Z_{1,b_{n_k}}(\theta_{n_k}, P_{n_k}), \dots, Z_{m,b_{n_k}}(\theta_{n_k}, P_{n_k}))' \in S \} - \Phi_{\Sigma(\theta_{n_k}, P_{n_k})}(S)| \rightarrow 0 .$$

An appeal to the triangle inequality yields the desired contradiction to (32). ■

Remark 3.10 Independently, Andrews and Guggenberger (2007) show using a different argument that \mathcal{C}_n defined by (10) satisfies (3) whenever \mathbf{P} is such that for some $\delta > 0$

$$\sup_{P \in \mathbf{P}} \sup_{\theta \in \Theta} E_P \left[\frac{(g_j(X_i, \theta) - E_P\{g_j(X_i, \theta)\})^{2+\delta}}{\sigma_j^{2+\delta}(\theta, P)} \right] < \infty .$$

This condition is mildly stronger than the condition imposed in Example 3.8. ■

3.4 Confidence Regions for Functions of Identifiable Parameters

In this section, we consider the problem of constructing sets satisfying (5) and (6). Let $f : \Theta \rightarrow \Lambda$ be given. Our construction again relies upon test inversion, but in this case the individual null hypotheses are given by

$$H_\lambda : \lambda \in \Lambda_0(P) \text{ for } \lambda \in \Lambda , \quad (33)$$

where $\Lambda_0(P)$ is defined by (4). The alternative hypotheses are understood to be

$$K_\lambda : \lambda \notin \Lambda_0(P) \text{ for } \lambda \in \Lambda .$$

For $\lambda \in \Lambda$, let $f^{-1}(\lambda) = \{\theta \in \Theta : f(\theta) = \lambda\}$. Note that

$$\begin{aligned} \lambda \in \Lambda_0(P) &\iff \exists \theta \in f^{-1}(\lambda) \text{ s.t. } Q(\theta, P) = 0 \\ &\implies \inf_{\theta \in f^{-1}(\lambda)} Q(\theta, P) = 0 . \end{aligned} \quad (34)$$

This suggests a natural test statistic for each of these null hypotheses H_λ :

$$\inf_{\theta \in f^{-1}(\lambda)} a_n \hat{Q}_n(\theta) , \quad (35)$$

where $a_n \hat{Q}_n(\theta)$ is the test statistic used earlier to test the null hypothesis that $Q(\theta, P) = 0$.

Remark 3.11 Under further assumptions, the implication (34) will hold in reverse as well. In particular, if Θ is a compact metric space, then it suffices to assume that $f(\theta)$ and $Q(\theta, P)$ are both continuous. But in many instances, such as Example 3.9 below, it is easy to verify that the reverse implication holds directly without relying upon such assumptions. ■

We may now proceed as before, but with the test statistic (35) in place of our earlier test statistic $a_n \hat{Q}_n(\theta)$. For $\alpha \in (0, 1)$, define

$$\hat{d}_n^f(\lambda, 1 - \alpha) = \inf \left\{ x : \frac{1}{N_n} \sum_{1 \leq i \leq N_n} I \left\{ \inf_{\theta \in f^{-1}(\lambda)} a_b \hat{Q}_{n,b,i}(\theta) \leq x \right\} \geq 1 - \alpha \right\} ,$$

and let

$$\mathcal{C}_n^f = \left\{ \lambda \in \Lambda : \inf_{\theta \in f^{-1}(\lambda)} a_n \hat{Q}_n(\theta) \leq \hat{d}_n^f(\lambda, 1 - \alpha) \right\} . \quad (36)$$

We now have the following theorem, which generalizes Theorems 3.2 and 3.3.

Theorem 3.4 *Let $X_i, i = 1, \dots, n$, be an i.i.d. sequence of random variables with distribution P and let $b = b_n < n$ be a sequence of positive integers tending to infinity, but satisfying $b/n \rightarrow 0$. Denote by $J_n(\cdot, \lambda, P)$ the distribution of*

$$a_n \left(\inf_{\theta \in f^{-1}(\lambda)} \hat{Q}_n(\theta) - \inf_{\theta \in f^{-1}(\lambda)} Q(\theta, P) \right) .$$

(i) *Suppose that for every $\lambda \in \Lambda_0(P)$*

$$\limsup_{n \rightarrow \infty} \sup_{x \in \mathbf{R}} \{ J_b(x, \lambda, P) - J_n(x, \lambda, P) \} \leq 0 . \quad (37)$$

Then, \mathcal{C}_n^f defined by (36) satisfies (5).

(ii) *Suppose that*

$$\limsup_{n \rightarrow \infty} \sup_{\lambda \in \Lambda} \sup_{P \in \mathbf{P} : \lambda \in \Lambda_0(P)} \sup_{x \in \mathbf{R}} \{ J_b(x, \lambda, P) - J_n(x, \lambda, P) \} \leq 0 . \quad (38)$$

Then, \mathcal{C}_n^f defined by (36) satisfies (6).

PROOF: The proof is identical to the proofs of Theorems 3.2 and 3.3. ■

We now provide a simple illustration of the use of Theorem 3.4 to construct sets satisfying (5) and (6).

Example 3.9 Consider the following straightforward generalization of Example 2.1. Let $(X_i, Y_i), i = 1, \dots, n$ be an i.i.d. sequence of random variables with distribution P on \mathbf{R}^2 . The parameter of interest, θ_0 , is known to satisfy $\theta_{0,1} \geq \mu_X(P)$ and $\theta_{0,2} \geq \mu_Y(P)$. The identified set is therefore given by $\Theta_0(P) = \{\theta \in \mathbf{R}^2 : \theta_1 \geq \mu_X(P) \text{ and } \theta_2 \geq \mu_Y(P)\}$. This set may be characterized as the set of minimizers of

$$Q(\theta, P) = (\mu_X(P) - \theta_1)_+^2 + (\mu_Y(P) - \theta_2)_+^2 .$$

The sample analog of $Q(\theta, P)$ is given by $\hat{Q}_n(\theta) = (\bar{X}_n - \theta_1)_+^2 + (\bar{Y}_n - \theta_2)_+^2$. Suppose interest focuses on $\theta_{0,1}$ rather than the entire vector θ_0 ; that is, the object of interest is $f(\theta_0)$, where $f : \mathbf{R}^2 \rightarrow \mathbf{R}$ is defined by $f(\theta) = \theta_1$, instead of θ_0 . Note that in this instance $\Lambda_0(P)$ is simply $\{\theta_1 \in \mathbf{R} : \theta_1 \geq \mu_X(P)\}$.

First, consider the problem of constructing sets satisfying (5). To this end, let $a_n = n$ and suppose P is such that $\sigma_X^2(P)$ exists. Consider $\theta_1 \in \Lambda_0(P)$. Assume without loss of generality that $\mu_X(P) = 0$. Then,

$$\begin{aligned} \inf_{\theta_2 \in f^{-1}(\theta_1)} a_n \hat{Q}_n(\theta) &= \inf_{\theta_2 \in \mathbf{R}} n(\bar{X}_n - \theta_1)_+^2 + n(\bar{Y}_n - \theta_2)_+^2 \\ &= n(\bar{X}_n - \theta_1)_+^2 . \end{aligned}$$

It now follows immediately from the analysis of Example 3.1 that (37) holds. Therefore, \mathcal{C}_n^f defined by (36) satisfies (5).

Now consider the problem of constructing sets satisfying (6). As before, let $a_n = n$ and let \mathbf{P} be a set of distributions for which the marginal distribution of X satisfies (22). Consider $\theta_1 \in \Lambda_0(P)$. Since $\inf_{\theta_2 \in f^{-1}(\theta_1)} a_n \hat{Q}_n(\theta)$ is simply $n(\bar{X}_n - \theta_1)_+^2$, it follows immediately from the analysis of Example (3.5) that (38) holds. Therefore, \mathcal{C}_n^f defined by (36) satisfies (6). ■

Remark 3.12 Of course, given a confidence region for identifiable parameters \mathcal{C}_n , one crude construction of a confidence region for a function of identifiable parameters is available as the image of \mathcal{C}_n under the function of interest. Unfortunately, such a construction will typically be very conservative. ■

4 Conclusion

This paper has provided computationally intensive, yet feasible methods for inference for a large class of partially identified models. The class of models we have considered are defined by a population objective function $Q(\theta, P)$. The main problem we have considered is the construction of random sets that contain each identifiable parameter with at least some pre-specified probability asymptotically. We have also extended these constructions to situations in which the object of interest is the image of an identifiable parameter under a known function. We have verified that these constructions can be applied in several examples of interest.

The results developed in this paper build upon earlier work by Chernozhukov et al. (2007), who also consider the problem of inference for the same class of partially identified models. In contrast to their analysis, we have also provided conditions under which our confidence regions are uniformly consistent in level. Imbens and Manski (2004) consider the problem of constructing confidence regions for identifiable parameters that are uniformly consistent in level, but their analysis considers only the special case of our class of partially identified models in which the identified set is an interval whose upper and lower endpoints are means or at least behave like means asymptotically. Our results therefore provide a generalization of their results to this broader class of models. In order to prove these results, we have relied upon recently developed results that provide general conditions under which confidence regions for a parameter constructed using subsampling are uniformly consistent in level.

In a companion paper, Romano and Shaikh (2006b), we consider the problem of constructing random sets that contain the entire identified set with at least some prespecified probability asymptotically. There we also provide an empirical illustration of the methodologies in both papers.

5 Appendix: Proof of Lemma 3.3.1

First note that we may assume without loss of generality that $\mu_1(P) = \dots = \mu_k(P) = 0$ and that the diagonal elements of $\Sigma(P)$ are all equal to 1. We will now argue by induction. Induction is anchored at $k = 1$ by Lemma 11.4.1 of Lehmann and Romano (2005) and Theorem 2.11 of Bhattacharya and Rao (1976), so assume that the result is true for $1, \dots, k - 1$. To show that the result holds for k , we will argue by contradiction. If the result were false, then there would exist a subsequence n_j and a corresponding sequence $P_{n_j} \in \mathbf{P}$ such that $\Sigma(P_{n_j}) \rightarrow \Sigma^*$

and

$$\sup_{S \in \mathcal{S}} |P_{n_j} \{ \sqrt{n_j}(\bar{X}_{1,n_j}, \dots, \bar{X}_{k,n_j}) \in S \} - \Phi_{\Sigma(P_{n_j})}(S)| \not\rightarrow 0 .$$

By Theorem 2.11 of Bhattacharya and Rao (1976), we have that

$$\sup_{S \in \mathcal{S}} |\Phi_{\Sigma(P_{n_j})}(S) - \Phi_{\Sigma^*}(S)| \rightarrow 0 ,$$

so we have that

$$\sup_{S \in \mathcal{S}} |P_{n_j} \{ \sqrt{n_j}(\bar{X}_{1,n_j}, \dots, \bar{X}_{k,n_j}) \in S \} - \Phi_{\Sigma^*}(S)| \not\rightarrow 0 .$$

There are two cases to consider.

First consider the case in which there exists $(0, \dots, 0) \neq \eta \in \mathbf{R}^k$ such that $\eta' \Sigma^* \eta = 0$. We may assume without loss of generality that $\eta_1 = -1$. It follows that $\eta' \Sigma(P_{n_j}) \eta$, the variance of $\sqrt{n_j}(\bar{X}_{1,n_j}, \dots, \bar{X}_{k,n_j})' \eta$, tends to zero. Hence,

$$\sqrt{n_j}(\bar{X}_{1,n_j}, \dots, \bar{X}_{k,n_j})' \eta \xrightarrow{P_{n_j}} 0 .$$

Therefore,

$$\begin{aligned} & P_{n_j} \{ \sqrt{n_j} \bar{X}_{1,n_j} \leq x_1, \dots, \sqrt{n_j} \bar{X}_{k,n_j} \leq x_k \} \\ &= P_{n_j} \{ \sqrt{n_j}(\bar{X}_{2,n_j} \eta_2 + \dots + \bar{X}_{k,n_j} \eta_k) + o_{P_{n_j}}(1) \leq x_1, \dots, \sqrt{n_j} \bar{X}_{k,n_j} \leq x_k \} . \end{aligned}$$

By the inductive hypothesis and Slutsky's Theorem, we have that this last expression converges to $\Phi_{\tilde{\Sigma}}(S_{x,\eta})$, where $\tilde{\Sigma}$ is the lower right $k-1 \times k-1$ submatrix of Σ and $S_{x,\eta}$ is the appropriate convex subset of \mathbf{R}^{k-1} . Since $\Phi_{\tilde{\Sigma}}(S_{x,\eta}) = \Phi_{\Sigma^*}((-\infty, x_1] \times \dots \times (-\infty, x_k])$, we have that

$$P_{n_j} \{ \sqrt{n_j} \bar{X}_{1,n_j} \leq x_1, \dots, \sqrt{n_j} \bar{X}_{k,n_j} \leq x_k \} \rightarrow \Phi_{\Sigma^*}((-\infty, x_1] \times \dots \times (-\infty, x_k]) .$$

By Theorem 2.11 of Bhattacharya and Rao (1976), we therefore have that

$$\sup_{S \in \mathcal{S}} |P_{n_j} \{ \sqrt{n_j}(\bar{X}_{1,n_j}, \dots, \bar{X}_{k,n_j}) \in S \} - \Phi_{\Sigma^*}(S)| \rightarrow 0 ,$$

which yields the desired contradiction in the first case.

Now consider the case in which for all $(0, \dots, 0) \neq \eta \in \mathbf{R}^k$, we have that $\eta' \Sigma^* \eta > 0$. We will use the Cramer-Wold device to establish that

$$\sqrt{n_j}(\bar{X}_{1,n_j}, \dots, \bar{X}_{k,n_j}) \xrightarrow{\mathcal{L}} \Phi_{\Sigma^*}(\cdot)$$

under P_{n_j} . Let $a \in \mathbf{R}^k$ and consider $\sqrt{n_j}(\bar{X}_{1,n_j}, \dots, \bar{X}_{k,n_j})' a$. To show that this expression converges in distribution to the appropriate normal distribution, it suffices by Lemma 11.4.1

of Lehmann and Romano (2005) to show that $a'X$ satisfies (22) when $X \sim P \in \mathbf{P}$ and the diagonal elements of $\Sigma(P)$ are all equal to 1. We may assume without loss of generality that all elements of a are nonzero, since otherwise the result would follow from the inductive hypothesis. Moreover, since (a_1X_1, \dots, a_kX_k) has a distribution in \mathbf{P} whenever X does, we may assume further that $a_1 = \dots = a_k = 1$. Since $\sigma_{X_1+\dots+X_k}^2(P) = (1, \dots, 1)' \Sigma(P) (1, \dots, 1)$ and

$$(1, \dots, 1)' \Sigma(P_{n_j}) (1, \dots, 1) \rightarrow (1, \dots, 1)' \Sigma^* (1, \dots, 1) > 0 ,$$

we may also assume that

$$\sigma_{X_1+\dots+X_k}^2(P) > \delta^2 > 0 .$$

For any $\lambda > 0$ and any such P , note that

$$\begin{aligned} & E_P \left[\frac{(X_1 + \dots + X_k)^2}{\sigma_{X_1+\dots+X_k}^2(P)} I_{\left\{ \frac{|X_1 + \dots + X_k|}{\sigma_{X_1+\dots+X_k}(P)} > \lambda \right\}} \right] \\ & \leq E_P \left[\frac{(X_1 + \dots + X_k)^2}{\delta^2} I_{\left\{ \frac{|X_1 + \dots + X_k|}{\delta} > \lambda \right\}} \right] . \end{aligned} \quad (39)$$

Using the inequality

$$(A_1 + \dots + A_k)^2 I_{\{|A_1 + \dots + A_k| > \lambda\}} \leq k(A_1^2 I_{\{|A_1| > \lambda/k\}} + \dots + A_k^2 I_{\{|A_k| > \lambda/k\}}) ,$$

we may further bound (39) by

$$k \sum_{1 \leq \ell \leq k} E_P \left[\frac{X_\ell^2}{\delta^2} I_{\left\{ \frac{|X_\ell|}{\delta} > \frac{\lambda}{k} \right\}} \right] .$$

Thus, by the assumption on the marginal distributions of $P \in \mathbf{P}$, we have that

$$\lim_{\lambda \rightarrow 0} \sup_{P \in \mathbf{P}} E_P \left[\frac{(X_1 + \dots + X_k)^2}{\sigma_{X_1+\dots+X_k}^2(P)} I_{\left\{ \frac{|X_1 + \dots + X_k|}{\sigma_{X_1+\dots+X_k}(P)} > \lambda \right\}} \right] = 0 ,$$

which establishes the desired convergence. By Theorem 2.11 of Bhattacharya and Rao (1976), we have that

$$\sup_{S \in \mathcal{S}} |P_{n_j} \{\sqrt{n_j}(\bar{X}_{1,n_j}, \dots, \bar{X}_{k,n_j}) \in S\} - \Phi_{\Sigma^*}(S)| \rightarrow 0 ,$$

which yields the desired contradiction in the second case. ■

References

ANDREWS, D. W. K. (2000). Inconsistency of the bootstrap when a parameter is on the boundary of the parameter space. *Econometrica*, **68** 399–405.

- ANDREWS, D. W. K., BERRY, S. and JIA, P. (2004). Confidence regions for parameters in discrete games with multiple equilibria with an application to discount chain store location. Working paper.
- ANDREWS, D. W. K. and GUGGENBERGER, P. (2007). Validity of subsampling and 'plug-in asymptotic' inference for parameters defined by moment inequalities. Working paper.
- ANDREWS, D. W. K. and SOARES, G. (2007). Inference for parameters defined by moment inequalities using generalized moment selection. Working paper.
- BAHADUR, R. R. and SAVAGE, L. J. (1956). The nonexistence of certain statistical procedures in nonparametric problems. *Annals of Mathematical Statistics*, **27** 1115–1122.
- BAJARI, P., FOX, J. T. and RYAN, S. P. (2007). Evaluating wireless carrier consolidation using semiparametric demand estimation. Working paper.
- BENKARD, C. L., BAJARI, P. and LEVIN, J. (2007). Estimating dynamic models of imperfect competition. *Econometrica*, **75** 1331–1370.
- BERESTEANU, A. and MOLINARI, F. (2006). Asymptotic properties for a class of partially identified models. Working paper.
- BHATTACHARYA, R. N. and RAO, R. (1976). *Normal Approximation and Asymptotic Expansions*. Wiley, New York.
- BORZEKOWSKI, R. and COHEN, A. (2005). Estimating strategic complementarities in credit unions' outsourcing decisions. Working paper.
- BUGNI, F. (2007). Bootstrap inference in partially identified models. Working paper.
- CANAY, I. (2007). El inference for partially identified models: Large deviations optimality and bootstrap validity. Working paper.
- CHERNOZHUKOV, V., HONG, H. and TAMER, E. (2004). Inference for identified parameter sets in econometric models. Working paper, Department of Economics, M.I.T.
- CHERNOZHUKOV, V., HONG, H. and TAMER, E. (2007). Estimation and confidence regions for parameter sets in econometric models. *Econometrica*, **75** 1243–1284.
- CILBERTO, F. and TAMER, E. (2004). Market structure and multiple equilibria in airline markets. Working paper.

- DiCiccio, T. and Romano, J. P. (1990). Nonparametric confidence limits by resampling and least favorable distributions. *International Statistical Review*, **58** 59–76.
- GALICHON, A. and HENRY, M. (2006). Dilation bootstrap: A methodology for constructing confidence regions with partially identified models. Working paper.
- IMBENS, G. and MANSKI, C. F. (2004). Confidence intervals for partially identified parameters. *Econometrica*, **72** 1845–1857.
- LEHMANN, E. L. and ROMANO, J. P. (2005). *Testing Statistical Hypotheses*. 3rd ed. Springer, New York.
- MANSKI, C. F. (2003). *Partial Identification of Probability Distributions*. Springer-Verlag, New York.
- MANSKI, C. F. and TAMER, E. (2002). Inference on regressions with interval data on a regressor or outcome. *Econometrica*, **70** 519–546.
- PAKES, A., PORTER, J., HO, K. and ISHII, J. (2005). Moment inequalities and their application. Working paper.
- POLITIS, D. N., ROMANO, J. P. and WOLF, M. (1999). *Subsampling*. Springer, New York.
- ROMANO, J. P. (2004). On non-parametric testing, the uniform behavior of the t -test, and related problems. *Scandinavian Journal of Statistics*, **31** 567–584.
- ROMANO, J. P. and SHAIKH, A. M. (2006a). Inference for identifiable parameters in partially identified econometric models. Technical Report 2006–9, Department of Statistics, Stanford University.
- ROMANO, J. P. and SHAIKH, A. M. (2006b). Inference for the identified set in partially identified econometric models. Technical Report 2006–10, Department of Statistics, Stanford University.
- ROSEN, A. (2007). Confidence sets for partially identified parameters that satisfy a finite number of moment inequalities. Working paper.
- SANTOS, A. (2007). Inference in nonparametric instrumental variables with partial identification. Working paper.
- STOYE, J. (2007). More on confidence intervals for partially identified models. Working paper.