ARTICLE IN PRESS

Journal of Econometrics ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

Journal of Econometrics

journal homepage: www.elsevier.com/locate/jeconom

Specification analysis of linear quantile models

J.C. Escanciano [a], S.C. Goh [b,*]

[a] Department of Economics, Indiana University, Wylie Hall, 100 South Woodlawn Avenue, Bloomington, IN 47405, United States

[b] Department of Economics, The University of Wisconsin-Milwaukee, P.O. Box 413, Milwaukee, WI 53201, United States

ARTICLE INFO

Article history:
Received 22 September 2011
Received in revised form
6 December 2012
Accepted 27 July 2013
Available online xxxx

JEL classification:
C12
C31
C52

Keywords:
Quantile regression
Specification tests
Empirical processes
Wild bootstrap

ABSTRACT

This paper introduces a nonparametric test for the correct specification of a linear conditional quantile function over a continuum of quantile levels. These tests may be applied to assess the validity of post-estimation inferences regarding the effect of conditioning variables on the distribution of outcomes. We show that the use of an orthogonal projection on the tangent space of nuisance parameters at each quantile index both improves power and facilitates the simulation of critical values via the application of a simple multiplier bootstrap procedure. Monte Carlo evidence and an application to the empirical analysis of age–earnings curves are included.

1. Introduction

Let Y be a random variable, and X a d-dimensional random vector. Consider the continuum of conditional probability restrictions given by

$$P\left[Y \leq X^{\top}\theta_0(\alpha)\middle|X\right] = \alpha, \quad \alpha \in \mathcal{A}, \tag{1.1}$$

where $\mathcal{A}$ is a compact subset of $(0, 1)$, $X^{\top}$ denotes the transpose of X, and $\theta_0(\cdot)$ is a measurable unknown function from $\mathcal{A}$ to a compact subset Θ of $\mathbb{R}^d$. As such, (1.1) involves a continuum of conditional quantile models that are each linear in a vector of parameters $\theta_0(\alpha)$. These models generalize conventional linear regression models identified by a median restriction. Estimators of $\theta_0(\alpha)$ in the context of (1.1) were pioneered in econometrics by Koenker and Bassett (1978). Their methodology has proven to be quite popular in recent years (e.g., Koenker and Hallock, 2001; Koenker, 2005, and references cited therein).

This paper develops omnibus tests for the correct specification of the conditional quantile function $X^{\top}\theta_0(\cdot)$ in (1.1) over $\mathcal{A}$. The hypothesis is that (1.1) holds with probability one for some $\theta_0(\cdot)$ in the corresponding parameter space, while the alternative is simply the negation of the null. We aim to construct consistent tests, i.e., tests that reject with probability tending to one in large samples if the model is misspecified. Such tests are important in applications because the conclusions of any post-estimation inferences based on the fitted quantile model will be sensitive to the implicit assumption that the conditional quantile function $X^{\top}\theta_0(\cdot)$ is correctly specified for all quantiles $\alpha \in \mathcal{A}$. In particular, if the conditional α'-quantile of Y for some $\alpha' \in \mathcal{A}$ is incorrectly assumed to have the form $X^{\top}\theta_0(\alpha')$, then estimators of $\theta_0(\alpha')$ will result in misleading inferences of the marginal effect of X on the α'-quantile of Y; see Angrist et al. (2006).

While omnibus tests of the validity of a linear-in-parameters conditional quantile function against unspecified alternatives have already been developed in a number of different papers, the analysis has to date been mostly limited to a single quantile, generally taken without loss of generality to be the median. See e.g., the papers of Zheng (1998), Bierens and Ginther (2001), Horowitz and Spokoiny (2002) and Whang (2006a,b)). Horowitz and Lee (2009) develop a specification test for the more general case where X is possibly endogenous and the single-quantile restriction holds conditionally on a vector of instruments. It is straightforward in theory to extend the testing approach that we propose to the case where X is endogenous and the "structural" quantile function $X^{\top}\theta_0(\alpha)$ is identified by a vector of instruments satisfying the conditional rank

* Corresponding author. Tel.: +1 414 229 3395; fax: +1 414 229 3860.
E-mail addresses: jescanci@indiana.edu (J.C. Escanciano), goh@uwm.edu (S.C. Goh).

0304-4076/$ – see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.jeconom.2013.07.006

invariance condition of Chernozhukov and Hansen (2005). An extension of this sort, however, is complicated in practice by the apparent unavailability of a computationally convenient estimator of the null nuisance parameter $\theta_0(\alpha)$ in the case of a structural quantile model. In particular, an extension to the case of endogenous X would require in practice estimates of the null nuisance parameters $\theta_0(\alpha_j)$ over a grid of quantiles $\{\alpha_j\}$ that becomes progressively dense in the sample size.

To the best of our knowledge, essentially the only proposal to date for omnibus specification tests of the functional form of a conditional quantile model over a continuum of quantiles is given in Escanciano and Velasco (2010). These authors considered tests of possibly nonlinear dynamic quantile models implemented using subsampling. A specialization of an approach recently proposed by Rothe and Wied (2013) and based on estimates of conditional *distribution* functions might also be considered in this context.

In this paper, we consider an approach in the specific framework of linear quantile regression models for independently and identically distributed (iid) data. What appears distinctive to our approach is the specific treatment of the nuisance parameters in the testing problem. In contrast to Escanciano and Velasco (2010), we acknowledge our lack of knowledge of $\theta_0(\cdot)$ in testing for (1.1), which can be accounted for by an orthogonal projection of a certain weight function into the so-called tangent space of nuisance parameters at each fixed quantile $\alpha \in \mathcal{A}$. The result of this projection is a test with improved power properties. The improvement in power of such projections has been noticed before in different contexts. Neyman (1959) first applied this idea in the context of fully parametric models. A more recent extension of this idea to the semiparametric context has been proposed by Bickel et al. (2006).

To illustrate the main ideas in the simplest possible terms, consider an example in a fully parametric case where one observes a random sample $\{W_i : i = 1, \ldots, n\}$ from the population of W with density f_{θ_0} satisfying the moment restriction $E[m(W, \theta_0)] = 0$, where θ_0 is an unknown nuisance parameter with values in $\Theta \subset \mathbb{R}^d$. Suppose $m(W, \cdot)$ is continuously differentiable at θ_0, with bounded derivative, and that a $\sqrt{n}$-consistent estimator of θ_0 is available, say θ_n. Furthermore, assume that $E[m^2(W, \theta_0)] < \infty$ and that $E[|s_\theta(W, \theta_0)|^2] < \infty$, where s_θ denotes the derivative of the log density with respect to θ. To test the moment restriction, it is natural to base a test on the sample analog of moments, i.e., on the statistic $\hat{R}_n = n^{-1/2} \sum_{i=1}^n m(W_i, \theta_n)$.

Using a standard Taylor expansion, a uniform law of large numbers and the generalized information equality (i.e., $E[\partial m(W, \theta_0)/\partial\theta] = -E[m(W, \theta_0)s_{\theta_0}(W)]$), we obtain the expansion

$$\hat{R}_n = \frac{1}{\sqrt{n}} \sum_{i=1}^n m(W_i, \theta_0) - \sqrt{n}\,(\theta_n - \theta_0)' \times E[m(W, \theta_0)s_{\theta_0}(W)] + o_P(1). \quad (1.2)$$

It is clear that the asymptotic distribution of $\hat{R}_n$ generally depends on that of the estimator θ_n, which complicates its approximation by bootstrap methods. But if the moment function $m(W, \theta_0)$ is orthogonal to the score s_{θ_0}, i.e., if $E[m(W, \theta_0)s_{\theta_0}(W)] = 0$, then estimation of θ_0 has no asymptotic impact on $\hat{R}_n$. Moreover, similar arguments to those used by Neyman (1959) show that a test satisfying such an orthogonality condition is optimal. See also Bickel et al. (2006).

The major innovation in our paper involves a by-product of an orthogonality condition of the form $E\left[m(W, \theta_0)s_{\theta_0}(W)\right] = 0$ that appears not to have been noticed in the existing literature. We exploit the fact that orthogonality enables a simple multiplier bootstrap approximation of the resulting test. That is, we exploit the fact that if $\{V_i\}_{i=1}^n$ is a sequence of iid random variables with zero mean, unit variance and independent of the sequence $\{W_i\}_{i=1}^n$, then by the same arguments leading to (1.2) above we have

$$\begin{aligned}\frac{1}{\sqrt{n}} \sum_{i=1}^n m(W, \theta_n)V_i &= \frac{1}{\sqrt{n}} \sum_{i=1}^n m(W, \theta_0)V_i - \sqrt{n}\,(\theta_n - \theta_0)' \\ &\quad \times E[m(W, \theta_0)s_{\theta_0}(W)V_i] + o_P(1) \\ &= \frac{1}{\sqrt{n}} \sum_{i=1}^n m(W, \theta_0)V_i + o_P(1) \\ &\to_d N(0, \Omega) \end{aligned} \quad (1.3)$$

for $\Omega = E[m^2(W, \theta_0)]$. It follows that the orthogonality condition $E[m(W, \theta_0)s_{\theta_0}(W)] = 0$ is critical for consistency of this resampling scheme. In particular, the limiting distribution of $\hat{R}_n$ will depend on that of θ_n if $E[m(W, \theta_0)s_{\theta_0}(W)] \neq 0$.

In the present context of quantile regression, these ideas turn out to be quite important in a practical sense, given the popularity of subsampling schemes in this context. Subsampling-based methods tend both to be more computationally intensive than the multiplier bootstrap scheme described above and to be particularly sensitive to the subjective choice of the subsample size. In sum, we believe that the methodology proposed in our paper is a useful extension of existing methods for testing the linearity of conditional quantiles, both in terms of the improvement in power it offers, and also in terms of the relative simplicity of implementation it offers via a multiplier bootstrap.

A natural alternative to the omnibus testing approach that we advocate is a "directional" one in which the researcher specifies a direction of departure from the hypothesized model. For instance, consider augmenting an existing vector of regressors X with polynomial or other nonlinear transformations of certain components of X. The significance of the additional coefficients in the resulting augmented quantile regression over a relevant continuum of quantiles is interpreted as evidence against correct specification of the original linear quantile regression model in X; see Otsu (2009). The implementation of such a test can be carried out using any of the three approaches outlined in Koenker and Machado (1999). This approach is simply an extension of the classical RESET approach of Ramsey (1969) to the setting of quantile regression and shares the same limitation of inconsistency that it possesses in the original setting of mean regression (Bierens, 1982). That is, there are uncountably many misspecifications that cannot be detected with such an approach; we provide a formal discussion of this point below in Section 3.2. On the other hand, directional tests are relatively simple to apply and often have good power properties if the specified direction of departure turns out to imply a model that is "close" in an appropriate sense to the true model. As such, we believe that the omnibus and directional approaches should be viewed as complements, rather than as substitutes. We show that typical testing strategies based on the significance of coefficients corresponding to nonlinear transformations of certain regressors are not optimal against local alternatives in the sense discussed below in Section 3.2. We also derive optimal directional tests based on the same empirical process on which our proposed omnibus test is based.

The remainder of this paper is organized as follows. In Section 2 we introduce the weighted empirical processes that constitute the basis upon which the new testing procedure is developed. In Section 3.2 we study the asymptotic distribution of the proposed tests under the null as well as under fixed and local alternatives. Section 3.2 also investigates optimal directional tests and their connection with our omnibus test. Section 3.3 discusses the use of a multiplier-bootstrap technique to approximate the asymptotic distributions of test statistics under the null as well as associated issues of implementation. Section 4 summarizes the results of Monte Carlo experiments designed to assess the finite-sample performance of our proposed testing procedures. Section 5 illustrates the applicability of the tests proposed here in the context of an empirical analysis of individual age–earnings profiles using US labor-market data. Mathematical proofs appear in the Appendix.

Please cite this article in press as: Escanciano, J.C., Goh, S.C., Specification analysis of linear quantile models. Journal of Econometrics (2013), http://dx.doi.org/10.1016/j.jeconom.2013.07.006