

Victor Chernozhukov, Christian Hansen, Kaspar Wüthrich

Krantz Template



Contents

List of Figures	ix
List of Tables	xi
I Title	1
1 Instrumental Variable Quantile Regression	3
1.1 Introduction	4
1.2 Model Overview	5
1.2.1 The Instrumental Variable Quantile Regression Model	5
1.2.2 Conditions for Point Identification	9
1.2.3 Discussion of the IVQR Model	10
1.2.4 Examples	11
1.2.5 Comparison to Other Approaches	14
1.3 Basic Estimation and Inference Approaches	16
1.3.1 Generalized Methods of Moments and Related Ap- proaches	17
1.3.2 Inverse Quantile Regression	19
1.3.2.1 A Useful Interpretation of IQR as a GMM es- timator.	21
1.3.3 Weak Identification Robust Inference	22
1.3.4 Finite Sample Inference	24
1.4 Advanced Inference with High-Dimensional X	25
1.4.1 Neyman-Orthogonal Scores	25
1.4.2 Estimation and Inference Using Orthogonal Scores . .	26
1.5 Conclusion	28
Bibliography	31



List of Figures



List of Tables



Part I

Title



1

Instrumental Variable Quantile Regression

CONTENTS

1.1	Introduction	3
1.2	Model Overview	5
1.2.1	The Instrumental Variable Quantile Regression Model ..	5
1.2.2	Conditions for Point Identification	8
1.2.3	Discussion of the IVQR Model	9
1.2.4	Examples	11
1.2.5	Comparison to Other Approaches	14
1.3	Basic Estimation and Inference Approaches	16
1.3.1	Generalized Methods of Moments and Related Approaches	17
1.3.2	Inverse Quantile Regression	19
1.3.2.1	A Useful Interpretation of IQR as a GMM estimator.	20
1.3.3	Weak Identification Robust Inference	22
1.3.4	Finite Sample Inference	24
1.4	Advanced Inference with High-Dimensional X	25
1.4.1	Neyman-Orthogonal Scores	25
1.4.2	Estimation and Inference Using Orthogonal Scores	26
1.5	Conclusion	28

This chapter reviews the instrumental variable quantile regression model of Chernozhukov and Hansen (2005). We discuss the key conditions used for identification of structural quantile effects within this model which include the availability of instruments and a restriction on the ranks of structural disturbances. We outline several approaches to obtaining point estimates and performing statistical inference for model parameters. Finally, we point to possible directions for future research.

Keywords: instrumental variables, ranks, $C(\alpha)$ -statistic, treatment effects, causal effects

JEL classification: C21, C26

1.1 Introduction

Empirical analyses often focus on understanding the structural (causal) relationship between an outcome, Y , and variables of interest, D . In many cases, interest is not just on how D affects measures of the center of the distribution Y but also on other features of the distribution. For example, in understanding the effect of a government subsidized saving program, one might be more interested in the effect of the program on the lower tail of the savings distribution conditional on individual characteristics than on the effect of the program on the mean of the savings distribution. Quantile regression, as introduced by Koenker and Bassett (1978), offers one useful way to estimate such effects and to summarize the impact of changes in D on the conditional distribution of Y .

Of course, variables of interest are often endogenous or self-selected in observational data. For example, individuals choose whether to participate in government subsidized savings plans. Similarly, in trying to understand the demand relationship between quantity and price, one must face that prices and quantities are jointly determined. Endogeneity of covariates renders conventional quantile regression inconsistent for estimating the causal effects of variables on the quantiles of outcomes of interest. Instrumental variables (IV) provide a powerful tool for learning about structural effects in the presence of endogenous right-hand-side variables, and we focus this review on a generalization of the classical linear instrumental variables model to accommodate estimating structural quantile treatment effects (QTE) in the presence of endogenous covariates.

We specifically focus on the instrumental variable quantile regression model developed in Chernozhukov and Hansen (2005). Chernozhukov and Hansen (2005) provide conditions under which structural QTE are nonparametrically identified through the use of instrumental variables. The key identifying assumption is a condition that restricts how structural errors, which we will refer to as rank variables, vary across different potential states of the endogenous variables. The simplest, though strongest, version of this condition is rank invariance which requires that individual ranks are invariant to the potential states of the endogenous variable. Rank invariance is implied by many classical structural models which posit a single source of unobserved heterogeneity, and the framework developed in Chernozhukov and Hansen (2005) is indeed a natural generalization of the classical structural simultaneous equation model, corresponding to a structural simultaneous equation model with non-additive errors.

There are alternative sets of modeling assumptions that one could employ to build a quantile model with endogeneity. Abadie et al. (2002) offer an approach within the local average treatment effect framework of Imbens and Angrist (1994). This approach differs from the framework discussed in

this review in a few key respects. The Abadie et al. (2002) framework does not restrict the behavior of rank variables across potential treatment states and thus allows for essentially unrestricted heterogeneity of effects at the cost of only identifying QTE for the subpopulation of compliers. To achieve identification without restricting structural errors, Abadie et al. (2002) restrict attention to a setting with a binary endogenous treatment variable and impose a monotonicity restriction on the relation between the instrument and treatment. When the endogenous variable of interest is continuous, triangular models as in Imbens and Newey (2009) provide another alternative framework for identifying and estimating QTE. As in Abadie et al. (2002), the Imbens and Newey (2009) framework does not restrict the evolution of ranks across treatment states. Instead, it relies on monotonicity of the selection mechanism in a scalar disturbance. We refer interested readers to Melly and Wüthrich (2016) in this handbook, which discusses the Abadie et al. (2002) approach in detail and contains further comparative discussion of the two modeling frameworks, and to Wüthrich (2014), which establishes a connection between the estimands of both models under the assumptions of the Abadie et al. (2002) framework. Section 1.2.5 discusses the approach by Imbens and Newey (2009) and compares it to the framework discussed in this chapter.

We devote the remainder of this review to providing an overview of the model of Chernozhukov and Hansen (2005) along with outlining approaches to estimating parameters and performing inference within this model.

1.2 Model Overview

1.2.1 The Instrumental Variable Quantile Regression Model

The instrumental variables quantile regression (IVQR) model is developed within the conventional potential outcome framework. Potential real-valued outcomes, which vary among observational units, are indexed against potential treatment states $d \in \mathcal{D}$ and denoted Y_d . The potential outcomes $\{Y_d\}$ are latent because, given the observed treatment D , the observed outcome for each observational unit is only one component

$$Y := Y_D$$

of the potential outcomes vector $\{Y_d\}$. Note that we use capital letters to denote random variables and lower case letters to denote the potential values the random variables may take throughout this review. We also do not explicitly state technical measurability assumptions as these can be deduced from the context.

The objective of causal or structural analysis is to learn about features of the distributions of potential outcomes Y_d . Of primary interest to us are the

τ^{th} quantiles of potential outcomes under various potential treatment states d , conditional on observed characteristics $X = x$, denoted as

$$Q_{Y_d}(\tau|x) = q(\tau, d, x).$$

We note that, after conditioning on observed characteristics $X = x$, each potential outcome Y_d can be related to its quantile function $q(\tau, d, x)$ as

$$Y_d = q(U_d, d, x), \text{ where } U_d \sim U(0, 1) \quad (1.1)$$

is the structural error term and (1.1) follows from the Fisher-Skorohod representation of random variables.

Given the conditional quantiles of the potential outcomes, we are then interested in QTE which are given by the difference in τ^{th} quantiles of two different conditional potential outcomes Y_{d_1} and Y_{d_0} :

$$q(\tau, d_1, x) - q(\tau, d_0, x).$$

These QTE may then be used to summarize the impact of variables of interest D on the quantiles of potential outcomes as suggested in Doksum (1974) and Lehmann (1975).

It is important to note that the structural error U_d in (1.1) is responsible for heterogeneity of potential outcomes among individuals with the same observed characteristics x . This error term determines the relative ranking of observationally equivalent individuals in the distribution of potential outcomes given the individuals' observed characteristics, and thus we refer to U_d as the rank variable. Because U_d drives differences between observationally equivalent individuals, one may think of U_d as representing some unobserved characteristic, e.g. ability or "proneness," where we adopt the term proneness from Doksum (1974) who uses the term as in "prone to learn fast" or "prone to grow taller". This interpretation of the structural error makes quantile analysis an interesting tool for describing and learning the structure of heterogeneous treatment effects while accounting for unobserved heterogeneity; see Doksum (1974), Heckman et al. (1997), and Koenker (2005). For example, consider a returns-to-training model, where Y_d 's are potential earnings under different training levels d , and $q(\tau, d, x)$ is the conditional earnings function which describes how an individual with training d , characteristics x , and latent "ability" τ is rewarded by the labor market. The earnings function may differ for different levels of τ , implying heterogeneous effects of training on earnings of people that have different levels of "ability". For example, it may be that the largest returns to training accrue to those in the upper tail of the conditional distribution, that is, to the "high-ability" workers.

In observational data, the realized treatment D is often selected in relation to potential outcomes, inducing endogeneity. This endogeneity makes the conventional quantile regression of Y on D and X , which relies upon the restriction

$$P[Y \leq Q_Y(\tau|D, X)|D, X] = \tau \text{ a.s.},$$

inappropriate for measuring the structural quantile function $q(\tau, d, x)$ and thus for learning about QTE. Indeed, the conditional quantile function, $Q_Y(\tau|d, x)$, solving these equations will generally differ from the structural quantile function of latent potential outcomes, $q(\tau, d, x)$, under endogeneity. The IVQR model presented below provides conditions under which we can identify and estimate the quantiles of the latent potential outcomes through the use of instruments Z that affect D but are independent of potential outcomes by making use of the nonlinear quantile-type conditional moment restrictions

$$P[Y \leq q(\tau, D, X)|X, Z] = \tau \text{ a.s.}$$

Formally, the IVQR model consists of five key conditions (some are representations).

Assumption 1 (IVQR Model) *Consider a common probability space (Ω, \mathcal{F}, P) and the set of potential outcome variables $(Y_d, d \in \mathcal{D})$, endogenous variables D , exogenous covariates X , and instrumental variables Z . The following conditions hold jointly with probability one:*

- A1** POTENTIAL OUTCOMES. *Conditional on X and for each d , $Y_d = q(U_d, d, X)$, where $\tau \mapsto q(\tau, d, X)$ is non-decreasing on $[0, 1]$ and left-continuous and $U_d \sim U(0, 1)$.*
- A2** INDEPENDENCE. *Conditional on X and for each d , U_d is independent of instrumental variables Z .*
- A3** SELECTION. *$D := \delta(Z, X, \nu)$ for some unknown function δ and random vector ν .*
- A4** RANK SIMILARITY. *Conditional on (X, Z, ν) , $\{U_d\}$ are identically distributed.*
- A5** OBSERVABLES. *The observed random vector consists of $Y := Y_D$, D , X and Z .*

The following theorem summarizes the main econometric implications of the model.

Theorem 1 (Main Implications of the IVQR Model) *Suppose conditions A1-A5 hold. (i) Then we have for $U := U_D$, with probability one,*

$$Y = q(U, D, X), \quad U \sim U(0, 1)|X, Z. \quad (1.2)$$

(ii) If (1.2) holds and $\tau \mapsto q(\tau, d, X)$ is strictly increasing for each d , then for each $\tau \in (0, 1)$, a.s

$$P[Y \leq q(\tau, D, X)|X, Z] = \tau. \quad (1.3)$$

(iii) If (1.2) holds, then for any closed subset I of $[0, 1]$, a.s.

$$P(U \in I) \leq P[Y \in q(I, D, X)|X, Z], \quad (1.4)$$

where $q(I, d, x)$ is the image of I under the mapping $\tau \mapsto q(\tau, d, x)$.

The first result states that the main consequence of A1-A5 is a simultaneous equation model (1.2) with non-separable error U that is independent of Z, X , and normalized so that $U \sim U(0, 1)$. The second result considers econometric implications when $\tau \mapsto q(\tau, D, X)$ is strictly increasing, which requires that Y is non-atomic conditional on X and Z . In this case, we obtain the conditional moment restriction (1.3). This implication follows from the first result and the fact that

$$\{Y \leq q(\tau, D, X)\} \text{ is equivalent to } \{U \leq \tau\}$$

when $q(\tau, D, X)$ is strictly increasing in τ . The final result deals with the case where Y may have atoms conditional on X and Z , e.g. when Y is a count or discrete response variable. The first two results were obtained in Chernozhukov and Hansen (2005), and the third result is in the spirit of results given in Chesher et al. (2013), Chesher (2005), and Chesher and Smolinski (2010).

The model and the results of Theorem 1 are useful for two reasons. First, Theorem 1 serves as a means of identifying QTE in a reasonably general heterogeneous effects model. Second, by demonstrating that the IVQR model leads to the conditional moment restrictions (1.3) and (1.4), Theorem 1 provides an economic and causal foundation for estimation based on these restrictions.

Equations (1.3) and (1.4) implicitly define the identification region for the structural quantile function $(\tau, d, x) \mapsto q(\tau, d, x)$. The identification region for the case of strictly increasing $\tau \mapsto q(\tau, d, x)$ can be stated as the set \mathcal{M} of functions $(\tau, d, x) \mapsto m(\tau, d, x)$ that satisfy the following relations, for all $\tau \in (0, 1]$

$$P[Y < m(\tau, D, X)|X, Z] = \tau \text{ a.s.} \quad (1.5)$$

This representation of the identification region \mathcal{M} is implicit. Without imposing additional conditions, statistical inference about $q \in \mathcal{M}$ from (1.5) can be performed using weak-identification robust inference as described in Chernozhukov and Hansen (2008), Jun (2008), Santos (2012), or Chernozhukov et al. (2009). Section 1.2.2 discusses conditions under which point identification is obtained; and we mainly focus on the point-identified case in discussing estimation and inference in this review.

The identification region for the case of weakly increasing $\tau \mapsto q(\tau, d, x)$ can be stated as the set \mathcal{M} of functions $(\tau, d, x) \mapsto m(\tau, d, x)$ that satisfy the following relations: For any closed subset I of $(0, 1]$,

$$P(U \in I) \leq P[Y \in m(I, D, X)|X, Z] \text{ a.s.},$$

where $m(I, D, X)$ is the image of I under the mapping $\tau \mapsto m(\tau, D, X)$. The inference problem here falls in the class of conditional moment inequalities and approaches such as those described in Andrews and Shi (2013) or Chernozhukov et al. (2013b) can be used.

1.2.2 Conditions for Point Identification

Here we briefly discuss the key conditions under which the moment equations (1.3) point identify the structural quantile function $q(\tau, d, x)$. We focus on the simplest case where $D \in \{0, 1\}$ and $Z \in \{0, 1\}$ and refer to Chernozhukov and Hansen (2013) for more details and extensions to multivalued and continuous D and to Chernozhukov and Hansen (2006) for a discussion of identification in linear-in-parameters models. The following analysis is conditional on $X = x$, but we suppress this dependence for the ease of notation.

It follows from Theorem 1 that there is at least one function $q(\tau, d)$ that solves $P[Y \leq q(\tau, D)|Z] = \tau$ a.s. The function $q(\tau, d)$ can be equivalently represented by a vector of its values $q = (q(\tau, 0), q(\tau, 1))'$. Therefore, for vectors of the form $y = (y_0, y_1)'$, we have a vector of moment equations

$$\Pi(y) := (P[Y \leq y_D|Z = 0] - \tau, P[Y \leq y_D|Z = 1] - \tau)',$$

where $y_D := (1 - D) \cdot y_0 + D \cdot y_1$. We say that $q(\tau, d)$ is identified in some parameter space, \mathcal{L} , if $y = q$ is the only solution to $\Pi(y) = 0$ among all $y \in \mathcal{L}$. Define the Jacobian $\partial\Pi(y)$ of $\Pi(y)$ with respect to $y = (y_0, y_1)'$ as

$$\begin{aligned} \partial\Pi(y) &:= \begin{bmatrix} f_Y(y_0|D = 0, Z = 0) P[D = 0|Z = 0] \\ f_Y(y_0|D = 0, Z = 1) P[D = 0|Z = 1] \\ f_Y(y_1|D = 1, Z = 0) P[D = 1|Z = 0] \\ f_Y(y_1|D = 1, Z = 1) P[D = 1|Z = 1] \end{bmatrix} \\ &:= \begin{bmatrix} f_{Y,D}(y_0, 0|Z = 0) & f_{Y,D}(y_1, 1|Z = 0) \\ f_{Y,D}(y_0, 0|Z = 1) & f_{Y,D}(y_1, 1|Z = 1) \end{bmatrix} \end{aligned} \quad (1.6)$$

The key condition for point identification is full rank of $\partial\Pi(y)$ at $y = q$. This local identification condition can be extended to a global condition; see Chernozhukov and Hansen (2005, 2013).

Full rank of $\partial\Pi(y)$ requires the impact of Z on the joint distribution of (Y, D) to be rich enough. To illustrate, note that full rank of $\partial\Pi(y)$ is equivalent to $\det(\partial\Pi(y)) \neq 0$, which implies that

$$\frac{f_{Y,D}(y_1, 1|Z = 1)}{f_{Y,D}(y_0, 0|Z = 1)} > \frac{f_{Y,D}(y_1, 1|Z = 0)}{f_{Y,D}(y_0, 0|Z = 0)} \quad (1.7)$$

(or the same condition with $>$ replaced by $<$). Inequality (1.7) may be interpreted as a *monotone likelihood ratio condition*. That is, the instrument Z should have a monotonic impact on the likelihood ratio in (1.7), which is generally stronger than the usual condition that D is correlated with Z . Nevertheless, the full rank condition will be trivially satisfied in many useful contexts. For instance, if the instrument satisfies one-sided non-compliance (e.g., those not offered the treatment cannot receive that treatment), $P[D = 1|Z = 0] = 0$, so that the right-hand side of (1.7) equals 0, which makes (1.7) hold trivially.

1.2.3 Discussion of the IVQR Model

Condition A1 imposes monotonicity on the structural function of interest which makes its relation to $q(\tau, d, x)$ apparent. Condition A2 states that potential outcomes are independent of Z , given X , which is a conventional independence restriction employed in nonlinear IV models. Condition A3 provides a convenient representation of a treatment selection mechanism, stated for the purpose of discussion. In A3, the unobserved random vector ν is responsible for the difference in treatment choices D across observationally identical individuals. Dependence between ν and $\{U_d\}$ is the source of endogeneity that makes the conventional exogeneity assumption $U \sim U(0, 1)|X, D$ break down. This failure leads to inconsistency of exogenous quantile methods for estimating the structural quantile function. Within the model outlined above, this breakdown is resolved through the use of instrumental variables.

The independence imposed in A2 and A3 is weaker than the assumption that both the disturbances $\{U_d\}$ in the outcome equation and the disturbances ν in the selection equation are *jointly* independent of the instrument Z which is maintained, for example, in Abadie et al. (2002). The assumption that structural errors $\{U_d\}$ and first-stage unobservables ν are jointly independent of instruments may be violated in practical examples. For example, this condition would not hold when the instrument is measured with error as discussed in Hausman (1977) or when the instrument is not assigned exogenously relative to the selection equation as in Example 2 in Imbens and Angrist (1994).

Condition A4 is the key restriction of the IVQR model. This assumption restricts the variation in ranks across potential outcomes and is key for identifying the structural quantile function $q(\tau, d, x)$ and the associated QTE. The simplest, though strongest, version of this condition is rank invariance which imposes that ranks U_d do not vary with potential treatment states d :

$$U_d = U \text{ for each } d \in \mathcal{D}. \quad (1.8)$$

Rank invariance is a strong condition that has been used in many interesting models without endogeneity such as Doksum (1974), Heckman et al. (1997), and Koenker and Geling (2001). Rank invariance implies that a common unobserved factor U , such as innate ability, determines the ranking of a given person across treatment states. For example, under rank invariance, people who are strong (highly ranked) earners without a training program ($d = 0$) remain strong earners having done the training ($d = 1$). Indeed, the earnings of a person with characteristics x and rank $U = \tau$ in the training state “0” is $Y_0 = q(\tau, 0, x)$ and in the state “1” is $Y_1 = q(\tau, 1, x)$; that is, the individual’s rank, τ , in the earnings distribution is exactly the same whether or not the person receives training. Finally, note that Condition A3 is a pure representation under rank invariance as nothing restricts the unobserved component ν in this case.

While convenient, rank invariance seems too strong a condition for many

applications as discussed, for example, in Heckman et al. (1997). Rank invariance maintains that an individual's rank in the outcome distribution under every possible state of the endogenous variables is exactly the same. Thus, the potential outcomes $\{Y_d\}$ are jointly degenerate which allows identification of individual treatment effects even though no individual is ever observed in more than one state of the endogenous variable. Rank invariance also rules out the possibility that there may be many unobserved factors that determine individual ranks which may be differentially relevant under different states of the endogenous variables.

Rank similarity A4 relaxes these undesirable features of rank invariance by allowing the rank variables $\{U_d\}$ to change across d in a way that reflects unobserved, asystematic variation in ranks across states of the endogenous variables while also providing sufficient structure to allow identification of QTE via the moment restrictions in Theorem 1. More specifically, rank similarity A4 relaxes exact rank invariance by allowing “slippages”, in the terminology of Heckman et al. (1997), in an individual's rank away from some common level U . Conditional on U , which may enter disturbance ν in the selection equation, and any other components of ν from the selection equation A3, rank similarity yields that the slippages of ranks away from common level U under different potential states of the endogenous variable, $U_d - U$, are identically distributed across $d \in \mathcal{D}$. In this formulation, we implicitly assume that any selection of the state of the endogenous variables occurs without knowing the exact potential outcomes. That is, selection may depend on U and even the distribution of slippages, but does not depend on the exact slippage $U_d - U$. This assumption is consistent with many empirical situations where the exact latent outcomes are not known before receipt of treatment. We also note that conditioning on appropriate covariates X may be important to achieve rank similarity. Finally, we note that rank similarity has testable implications. Dong and Shen (2015) and Frandsen and Lefgren (2015) exploit these conditions to develop tests of unconditional rank similarity, and their approaches could be extended to test some forms of conditional rank similarity.

1.2.4 Examples

We present two examples that highlight the nature of the model, its strengths, and its limitations.

Example 1 (Demand with Non-Separable Error). The following is a generalization of the classic supply-demand example taken from Chernozhukov and Hansen (2006). Consider the model

$$\begin{aligned} Y_p &= q(U, p), \\ \tilde{Y}_p &= \rho(\mathcal{U}, p, z), \\ P &\in \{p : \rho(U, p, Z) = q(\mathcal{U}, p)\}, \end{aligned} \tag{1.9}$$

where functions q and ρ are increasing in their first argument. The function

$p \mapsto Y_p$ is the random demand function, and $p \mapsto \tilde{Y}_p$ is the random supply function. Additionally, functions q and ρ may depend on covariates X , but this dependence is suppressed.

Random variable U is the level of demand and describes the demand curve at different states of the world. Demand is maximal when $U = 1$ and minimal when $U = 0$, holding p fixed. Note that we imposed rank invariance (1.8), as is typical in classic supply-demand models, by making U invariant to p .

Model (1.9) incorporates traditional additive error models for demand which have $Y_p = q(p) + \epsilon$ where $\epsilon = Q_\epsilon(U)$. The model is much more general in that the price can affect the entire distribution of the demand curve, while in traditional models it only affects the location of the distribution of the demand curve.

The τ -quantile of the demand curve $p \mapsto Y_p$ is given by $p \mapsto q(\tau, p)$. Thus, the curve $p \mapsto Y_p$ lies below the curve $p \mapsto q(\tau, p)$ with probability τ . Therefore, the various quantiles of the potential outcomes play an important role in describing the distribution and heterogeneity of the stochastic demand curve. The QTE may be characterized by $\partial q(\tau, p)/\partial p$ or by an elasticity $\partial \ln q(\tau, p)/\partial \ln p$. For example, consider the model $q(\tau, p) = \exp(\beta(\tau) + \alpha(\tau) \ln p)$ which corresponds to a Cobb-Douglas model for demand with non-separable error $Y_p = \exp(\beta(U) + \alpha(U) \ln p)$. The log transformation gives $\ln Y_p = \beta(U) + \alpha(U) \ln p$, and the QTE for the log-demand equation is given by the elasticity of the original τ -demand curve $\alpha(\tau) = \partial Q_{\ln Y_p}(\tau)/\partial \ln p = \partial \ln q(\tau, p)/\partial \ln p$.

The elasticity $\alpha(U)$ is random and depends on the state of the demand U and may vary considerably with U . For example, this variation could arise when the number of buyers varies and aggregation induces a non-constant elasticity across the demand levels. Chernozhukov and Hansen (2008) estimate a simple demand model based on data from a New York fish market that was first collected and used by Graddy (1995). They find point estimates of the demand elasticity, $\alpha(\tau)$, that vary quite substantially from -2 for low quantiles to -0.5 for high quantiles of the demand curve.

The third condition in (1.9), $P \in \{p : \rho(U, p, Z) = q(U, p)\}$, is the equilibrium condition that generates endogeneity; the selection of the clearing price P by the market depends on the potential demand and supply outcomes. As a result, we have a representation that is consistent with A3, $P = \delta(Z, \nu)$, where ν consists of U and \mathcal{U} and may include “sunspot” variables if the equilibrium price is not unique. Thus what we observe can be written as

$$Y := q(U, P), \quad P := \delta(Z, \nu), \quad U \text{ is independent of } Z. \quad (1.10)$$

Identification of the τ^{th} quantile of the demand function, $p \mapsto q(p, \tau)$ is obtained through the use of instrumental variables Z , like weather conditions or factor prices, that shift the supply curve and do not affect the level of the demand curve, U , so that independence assumption A2 is met. Furthermore, the IVQR model allows arbitrary correlation between Z and ν . This property is important as it allows, for example, Z to be measured with error or to be

exogenous relative to the demand equation but endogeneous relative to the supply equation.

Example 2 (Savings). Chernozhukov and Hansen (2004) use the framework of the IVQR model to examine the effects of participating in a 401(k) plan on an individual's accumulated wealth. Since wealth is continuous, wealth, Y_d , in the participation state $d \in \{0, 1\}$ can be represented as

$$Y_d = q(U_d, d, X), \quad U_d \sim U(0, 1)$$

where $\tau \mapsto q(\tau, d, X)$ is the conditional quantile function of Y_d and U_d is an unobserved random variable. U_d is an unobservable that drives differences in accumulated wealth conditional on X under participation state d . Thus, one might think of U_d as the preference for saving and interpret the quantile index τ as indexing rank in the preference for saving distribution. One could also model the individual as selecting the 401(k) participation state to maximize expected utility:

$$\begin{aligned} D &= \arg \max_{d \in \mathcal{D}} E \left[W\{Y_d, d\} \middle| X, Z, \nu \right] \\ &= \arg \max_{d \in \mathcal{D}} E \left[W\{q(U_d, d, x), d\} \middle| X, Z, \nu \right], \end{aligned} \tag{1.11}$$

where $W\{Y_d, d\}$ is the random indirect utility derived under participation state d . Of course, utility may depend on both observables in X as well as realized and unrealized unobservables. Only dependence on Y_d and d is highlighted. As a result, the participation decision is represented by

$$D = \delta(Z, X, \nu),$$

where Z and X are observed, ν is an unobserved information component that may be related to ranks U_d and includes other unobserved variables that affect the participation state, and function δ is unknown. This model fits into the IVQR model with the independence condition A2 requiring that U_d is independent of Z , conditional on X .

Under rank invariance (1.8) the preference for saving vector U_d may be collapsed to a single random variable $U = U_0 = U_1$. In this case, a single preference for saving is responsible for an individual's ranking across both treatment states. The more general rank similarity condition A4 relaxes the exact invariance of ranks U_d across d by allowing noisy, asystematic variations of U_d across d , conditional on (ν, X, Z) . This relaxation allows for variation in rank across the treatment states, requiring only an "expectational rank invariance." Similarity implies that given the information in (ν, X, Z) employed to make the selection of treatment D , the expectation of any function of rank U_d does not vary across the treatment states. That is, ex-ante, conditional on (ν, X, Z) , the ranks may be considered to be the same across potential treatments, but the realized, ex-post, rank may be different across treatment states.

From an econometric perspective, the similarity assumption is nothing but a restriction on the unobserved heterogeneity component which precludes systematic variation of U_d across the treatment states. To be more concrete, consider the following simple example where

$$U_d = F_{\nu+\eta_d}(\nu + \eta_d),$$

where $F_{\nu+\eta_d}(\cdot)$ is the distribution function of $\nu + \eta_d$ and $\{\eta_d\}$ are mutually i.i.d. conditional on ν , X , and Z . The variable ν represents an individual's "mean" saving preference, while η_d is a noisy adjustment. Clearly similarity holds in this case, $U_d \stackrel{d}{=} U_d$ given ν , X , and Z . This more general assumption leaves the individual optimization problem (1.11) unaffected, while allowing variation in an individual's rank across different potential outcomes.

While we feel that rank similarity may be a reasonable assumption in many contexts, imposing rank similarity is not innocuous. In the context of 401(k) participation, matching practices of employers could jeopardize the validity of the similarity assumption. To be more concrete, let $U_d = F_{\nu+\eta_d}(\nu + \eta_d)$ as before but let $\eta_d = dM$ for random variable M that depends on the match rate and is independent of ν , X , and Z . Then conditional on $\nu = v$, X , and Z , $U_0 = F_\nu(v)$ is degenerate but $U_1 = F_{\nu+M}(v + M)$ is not. Therefore, U_1 is not equal to U_0 in distribution. Similarity may still hold in the presence of the employer match if the rank, U_d , in the asset distribution is insensitive to the match rate. The rank may be insensitive if, for example, individuals follow simple rules of thumb such as target saving when they make their savings decisions. Also, if the variation of match rates is small relative to the variation of individual heterogeneity or if the covariates capture most of the variation in match rates, then similarity may be satisfied approximately.

1.2.5 Comparison to Other Approaches

There are, of course, other assumptions that one could employ to build a quantile model with endogeneity. In this section, we briefly compare the IVQR framework to triangular models as in Imbens and Newey (2009); see Chesher (2003), Koenker and Ma (2006), Lee (2007) and Chernozhukov et al. (2015a) for related models and results. We also note that triangular models are related to the Rosenblatt transform; see for example the chapter by Hallin and Šiman (2016) in this handbook. A comparison between the IVQR model and the popular Abadie et al. (2002) approach is provided in Melly and Wüthrich (2016) in this handbook.

The triangular model takes the form of a triangular system of equations

$$\begin{aligned} Y &= g(D, \epsilon), \\ D &= h(Z, \eta), \end{aligned}$$

where Y is the outcome, D is a continuous scalar endogenous variable, ϵ is a

vector of disturbances, Z is a vector of instruments with a continuous component, η is a scalar reduced form error, and we ignore other covariates X for simplicity. It is important to note that the triangular system generally rules out simultaneous equations which typically have that the reduced form relating D to Z depends on a vector of disturbances. For example, in a supply and demand system, the reduced form for both price and quantity will generally depend on the unobservables from both the supply equation and the demand equation; see Example 1 in Section 1.2.4.

Outside of η being a scalar, the key conditions that allow identification of quantile effects in the triangular system are (a) the function $\eta \mapsto h(Z, \eta)$ is strictly increasing in η and (b) D and ϵ are independent conditional on V for some observable or estimable V . The variable V is thus the “control function” conditional on which changes in D may be taken as causal. Imbens and Newey (2009) use $V = F_{D|Z}(d, z) = F_\eta(\eta)$ as a control variable and show that this variable satisfies condition (b) under the additional condition that (ϵ, η) is independent of Z . Identification then proceeds as follows. Under the assumed monotonicity of $h(Z, \eta)$ in η , $D = h(Z, \eta)$ can be used to identify V . Using V obtained in this first step, one may then construct the distribution of $Y|D, V$. Integrating over the distribution of V and using iterated expectations, one has

$$\begin{aligned} \int F_{Y|D,V}(y|d, v)F_V(dv) &= \int \mathbf{1}(g(d, \epsilon) \leq y)F_\epsilon(d\epsilon) \\ &= \Pr(g(d, \epsilon) \leq y) := G(y, d) \end{aligned}$$

and the structural quantile function Y_d can be obtained as $G^{-1}(\tau, d)$.

It should be emphasized that the triangular model is neither more nor less general than the IVQR model reviewed here. The key difference between the approaches is that the IVQR model uses an essentially unrestricted selection equation (ν may be vector valued) but requires monotonicity and a scalar disturbance (U) in the structural equation. The triangular system on the other hand relies on monotonicity of the selection mechanism in a scalar disturbance (η) but does not restrict the unobserved heterogeneity in the outcome equation (ϵ may be a vector of disturbances). In addition, the triangular system, as developed in Imbens and Newey (2009), requires a more stringent independence condition in that the instruments Z needs to be independent of both the structural disturbances, ϵ , and the reduced form disturbance, η . That the approaches impose structure on different parts of the model makes them complementary with a researcher’s choice between the two being dictated by whether it is more natural to impose restrictions on the structural function or the reduced form in a given application.

Finally, we note that the triangular model and the IVQR model can be made compatible by imposing the conditions from the triangular model on the selection equation and the conditions from the IVQR model on the structural model. Torgovitsky (2015) studies identification when both sets of conditions are imposed and shows that the requirements on the instruments may be

substantially relaxed relative to the IVQR model or Imbens and Newey (2009) in this case.

1.3 Basic Estimation and Inference Approaches

In this section, we present various approaches to estimating and doing inference for the parameters of the IVQR model under the leading case where $\tau \mapsto q(\tau, d, X)$ is strictly increasing. We focus on linear-in-parameters structural quantile models at a single quantile of interest τ :

$$q(\tau, d, x) = d' \alpha_0(\tau) + x' \beta_0(\tau). \quad (1.12)$$

In (1.12), $\alpha_0(\tau)$ captures the causal effect of the endogenous variables D on the τ^{th} quantile of the conditional distribution of potential outcomes Y_d given $X = x$. Similarly, $\beta_0(\tau)$ provides the causal effect of controls X on the τ^{th} quantile of the conditional potential outcome distributions. We note that D may also contain interactions of endogenous variables and covariates. Because $\alpha_0(\tau)$ is the chief object of interest in many studies, we focus most of our discussion on estimating and doing inference for $\alpha_0(\tau)$ treating $\beta_0(\tau)$ as a nuisance parameter. Note that in what follows we will often suppress the dependence of $\alpha_0(\tau)$ and $\beta_0(\tau)$ on the quantile level τ .

In interpreting the parameters in (1.12), it is important to note that the quantile index, τ , refers to the quantile of potential outcome Y_d given that exogenous variables are set to $X = x$ and not to the unconditional quantile of Y_d . For example, suppose that one of the control variables in the savings example in Section 1.2.4 is income. An individual at the 10th percentile of the distribution of Y_d given an income of \$200,000, which is far above the median income, may not necessarily be at the low tail of the unconditional distribution of Y_d as even a relatively low saver with a high level of income may still save substantially more than the median saver in the overall population, i.e., without conditioning on income; see Frölich and Melly (2013) for a further discussion of this point. In some applications, features of the conditional distribution are not the chief objects of interest and researchers are interested in effects of treatments on unconditional quantiles. Unconditional QTE can be obtained from the conditional quantile functions in three steps. First, obtain the conditional potential outcome distribution functions, $F_{Y_d}(y|x)$, as

$$F_{Y_d}(y|x) = \int_0^1 \mathbb{1}(d' \alpha_0(\tau) + x' \beta_0(\tau) \leq y) d\tau,$$

where $\mathbb{1}(\cdot)$ is the indicator function that returns one when the expression inside the parentheses is true and zero otherwise. Second, the unconditional potential outcome distributions, $F_{Y_d}(y)$, are obtained by integrating $F_{Y_d}(y|x)$

with respect to the marginal distribution of covariates, $F_X(x)$:

$$F_{Y_d}(y) = \int F_{Y_d}(y|x) dF_X(x).$$

Finally, the unconditional τ -QTE is given by $F_{Y_{d_1}}^{-1}(\tau) - F_{Y_{d_0}}^{-1}(\tau)$. This discussion suggests that given estimators of the parameters $\alpha_0(\tau)$ and $\beta_0(\tau)$ and the distribution of covariates $F_X(x)$, unconditional QTE can be estimated based on the plug-in principle; see for instance Machado and Mata (2005), Melly (2005) or Chernozhukov et al. (2013a).

Model (1.12) provides a simple and widely used baseline for discussion of estimation and inference. Extending the discussion to allow for nonlinear parametric specifications of the potential outcome quantile functions or to estimation at a small number of quantile indices that are widely spaced is straightforward. In some applications, we may be interested in understanding QTE across a range of quantile indices, say $\tau \in [\delta, 1 - \delta]$ for some $\delta > 0$. Chernozhukov and Hansen (2006) explicitly consider this case and provide uniform convergence results which allow for inference about a variety of hypotheses surrounding the behavior of QTE viewed as a function of τ such as tests of monotonicity of treatment effects or tests that treatment effects are uniformly 0 across a range of τ . Finally, we note that Chernozhukov et al. (2007), Horowitz and Lee (2007), Chen and Pouzo (2009), Chen and Pouzo (2012), and Gagliardini and Scaillet (2012) consider fully nonparametric approaches to estimating structural quantile models.

1.3.1 Generalized Methods of Moments and Related Approaches

The most direct way to estimate the parameters of the linear IVQR model is to note that the main implication of the model, equation (1.3), implies unconditional moment conditions

$$E[(\tau - \mathbb{1}(Y - D'\alpha_0 - X'\beta_0 \leq 0))\Psi] = 0 \quad (1.13)$$

where $\Psi := \Psi(X, Z)$ is a vector of functions of the instruments and endogenous variables.¹ Supposing that α_0 is an $s \times 1$ vector and β_0 is a $k \times 1$ vector, a minimal necessary condition for identifying the model parameters will be $\dim(\Psi) = r \geq k + s$.

Let, for $\theta := (\alpha, \beta)$ and $V := (Y, D, X, Z)$,

$$g_\tau(V, \theta) = (\tau - \mathbb{1}(Y - D'\alpha - X'\beta \leq 0))\Psi.$$

With a given set of instruments, Ψ , and observables $\{V_i\}_{i=1}^N =$

¹A natural choice of instruments would be $\Psi = (Z', X)'$ though the instruments and GMM weighting matrix could be chosen to produce a pointwise efficient procedure following Chamberlain (1987).

$\{Y_i, D_i, X_i, Z_i\}_{i=1}^N$, one may then form the sample analog of the right-hand-side of the equation (1.13),

$$\widehat{g}_N(\theta) = \frac{1}{N} \sum_{i=1}^N g_\tau(V_i, \theta), \quad (1.14)$$

and estimate $\theta_0 = (\alpha'_0, \beta'_0)'$ by generalized method of moments (GMM) as

$$\widehat{\theta} = (\widehat{\alpha}', \widehat{\beta}')' = \arg \min_{\theta \in \Theta} m_N(\theta) \quad (1.15)$$

for

$$m_N(\theta) := N \widehat{g}_N(\theta)' \Omega_N \widehat{g}_N(\theta)$$

where Ω_N is the GMM weighting matrix that will typically be set as

$$\Omega_N = \left(\tau(1 - \tau) \frac{1}{N} \sum_{i=1}^N \Psi_i \Psi_i' \right)^{-1}.$$

Maintaining sufficient conditions for point identification as in Chernozhukov and Hansen (2005, 2006, 2013) and assuming that a suitable solution to the GMM optimization problem (1.15) can be found, asymptotic properties of $\widehat{\theta}(\tau)$ would then follow from standard results for GMM with non-smooth moment conditions as in Newey and McFadden (1994); see Abadie (1995) and Chernozhukov and Hong (2003). We note that if the GMM problem (1.13) is overidentified, overidentification-type tests can be used to assess the joint validity of the underlying assumptions.

The chief difficulty in implementing estimation based on (1.15) is that the function being minimized is both non-smooth and non-convex in general. We also note that in many applications, s will be small, often one, but k may be quite large. Solving (1.15) then involves optimizing a non-smooth, non-convex function over $s + k$ arguments where $s + k$ may be quite large. Directly solving this problem thus poses a substantial computational challenge and has led to the adoption of different approaches to estimating the parameters of the IVQR model.

Within the conventional GMM framework, one option is to take the quasi-Bayesian approach of Chernozhukov and Hong (2003); see also Wang and Yang (2016) in this handbook for a review of subsequent work on related methods. The Chernozhukov and Hong (2003) approach uses the GMM criterion function to form a “quasi-likelihood”,

$$L_N(\theta) = \exp \left(-\frac{1}{2} N \widehat{g}_N(\theta)' \Omega_N \widehat{g}_N(\theta) \right),$$

which when coupled with a prior density $\pi(\theta)$ over model parameters θ , defines a “quasi-posterior” density for θ :

$$\pi_N(\theta) = L_N(\theta)\pi(\theta) / \int L_N(\tilde{\theta})d\pi(\tilde{\theta}) \propto L_N(\theta)\pi(\theta).$$

Rather than try to solve the optimization problem (1.15), one can then use MCMC sampling to attempt to explore the implied quasi-posterior distribution. Chernozhukov and Hong (2003) show that measures of central tendency from the quasi-posterior, such as the quasi-posterior mean,

$$\hat{\theta} = (\hat{\alpha}', \hat{\beta}')' = \int \theta d\pi_N(\theta)$$

and quasi-posterior median are consistent for model parameters with the same asymptotic distribution as the solution to (1.15). Chernozhukov and Hong (2003) also demonstrate that valid frequentist confidence intervals may be obtained by taking quasi-posterior quantiles. For example, a frequentist 95% confidence interval may be constructed as by taking the 2.5 and 97.5 quantiles of the quasi-posterior distribution. This approach bypasses the need to optimize a non-convex and non-smooth criterion at the cost of needing to design a sampler that adequately explores the quasi-posterior in a reasonable amount of computation time.

A second option is to directly smooth the GMM-criterion function as in Kaplan and Sun (2016), building upon ideas in Amemiya (1982) and Horowitz (1998). Specifically, one modifies the moment condition (1.14) to

$$\hat{g}_N^{h_N}(\theta) = \frac{1}{N} \sum_{i=1}^N (\tau - G_{h_N}(Y_i - D_i'\alpha - X_i'\beta)) \Psi_i, \quad (1.16)$$

by smoothing the indicator function, where $G_h(\cdot)$ denotes a smoothing function with smoothing parameter h . $G_h(\cdot)$ can be defined as the survival function associated with any kernel function $K_h(\cdot)$, i.e. $G_h(u) = \int_u^\infty K_h(v)dv$, that satisfies regularity conditions provided in Kaplan and Sun (2016). One can then proceed to estimate model parameters by replacing $\hat{g}_N(\theta)$ in (1.15) with $\hat{g}_N^{h_N}(\theta)$ and applying any optimizer which is appropriate for smooth, non-convex optimization problems or the quasi-Bayesian approach described above. Solving the smoothed problem can offer some computational gains relative to attempting to solve the original problem, though non-convexities remain after smoothing. The resulting estimator is first-order-equivalent to the GMM estimator for the original problem. The estimator can, however, enjoy higher-order improved performance. Kaplan and Sun (2016) provide a plug-in approach to choosing the smoothing parameter h_N and also demonstrate that the estimated parameters obtained from solving the smoothed problem may perform better in small samples than those from solving the unsmoothed problem or the inverse quantile regression discussed in Section 1.3.2.

1.3.2 Inverse Quantile Regression

Rather than work directly with moment condition (1.13), Chernozhukov and Hansen (2006) and Chernozhukov and Hansen (2008) take a different approach which they label the inverse quantile regression (IQR). The IQR is based on

the observation that (1.3) coupled with the linear quantile model (1.12) implies that the τ^{th} quantile of $Y - D'\alpha_0$ conditional on covariates X and instruments Z is equal to $X'\beta_0(\tau)$:

$$Q_{Y-D'\alpha_0}(\tau|X, Z) = X'\beta_0 + Z'\gamma_0 \text{ with } \gamma_0 \equiv 0. \quad (1.17)$$

That is, at the true value of the coefficient vector on the endogenous variables α_0 , the conventional linear τ -quantile regression of $Y - D'\alpha_0$ onto X and Z would yield coefficients on the instruments of exactly 0 in the population. This observation then suggests an estimation approach based on concentrating X out of the problem using conventional quantile regression, which is convex and can be solved very quickly, and then solving a lower dimensional non-convex optimization problem over only the dimension of D to find $\hat{\alpha}$.

Specifically, the IQR procedure works as follows. Let a denote an arbitrary hypothesized value for α_0 . Using the hypothesized value a , estimate coefficients $\beta(a)$ and $\gamma(a)$ from the model $Q_{Y-D'a}(\tau|X, Z) = X'\beta(a) + Z'\gamma(a)$ by running the ordinary linear τ -quantile regression of $Y - D'a$ onto X and Z . Let $\hat{\beta}(a)$ and $\hat{\gamma}(a)$ denote the resulting estimators of $\beta(a)$ and $\gamma(a)$. Also, let $\hat{\Omega}_N(a)$ denote the estimated covariance matrix of $\sqrt{N}(\hat{\gamma}(a) - \gamma(a))$, and note that this covariance matrix is available in any common implementation of the ordinary quantile regression. We can then define the IQR estimator of α_0 as

$$\hat{\alpha} = \arg \min_{a \in \mathcal{A}} W_N(a), \quad (1.18)$$

where

$$W_N(a) := N\hat{\gamma}(a)'\hat{\Omega}_N(a)^{-1}\hat{\gamma}(a). \quad (1.19)$$

Given $\hat{\alpha}$, we can then estimate β_0 as $\hat{\beta}(\hat{\alpha})$.

In terms of point estimation, the main virtue of the IQR is that, by concentrating out the coefficients on exogenous variables X , it produces a non-convex optimization problem over only the parameters α . In many applications, the dimension of D is small, so one can approach the non-convex optimization problem using highly robust optimization procedures that deal effectively with objectives with many local optima. Chernozhukov and Hansen (2006) recommend using a grid-search to solve (1.18) though other approaches are certainly available. Using a grid-search is particularly appealing when coupled with weak-identification robust inference as discussed in Section 1.3.3.

Chernozhukov and Hansen (2006) analyze the properties of $(\hat{\alpha}(\tau)', \hat{\beta}(\tau)')'$ under assumptions that guarantee strong identification. They verify asymptotic normality of the estimator, provide a consistent estimator of the asymptotic variance, and show how instruments and observation weights can be chosen to produce an efficient estimator of the coefficients for a single quantile following Chamberlain (1987). Chernozhukov and Hansen (2006) also analyze the behavior of the process $(\hat{\alpha}(\tau)', \hat{\beta}(\tau)')'$ not just at a point but viewed as a function of τ , providing uniform convergence results and discussing in detail applications of these convergence results to testing hypotheses about the behavior of $(\alpha_0(\tau)', \beta_0(\tau)')'$ across the index τ .

1.3.2.1 A Useful Interpretation of IQR as a GMM estimator.

It is useful to interpret IQR as first-order-equivalent to a particular GMM estimator, where we first profile out the coefficients on exogenous variables.

To this end, let us define

$$g_\tau(V, \alpha; \beta, \delta) = (\tau - \mathbb{1}(Y \leq D'\alpha + X'\beta)) \Psi(\alpha, \delta(\alpha)), \quad (1.20)$$

with “instrument”

$$\Psi(\alpha, \delta(\alpha)) := (Z - \delta(\alpha)X). \quad (1.21)$$

In (1.21),

$$\delta(\alpha) = M(\alpha)J^{-1}(\alpha)$$

where δ is a matrix parameter,

$$M(\alpha) = \mathbb{E}[ZX'f_\varepsilon(0 | X, Z)], \quad J(\alpha) = \mathbb{E}[XX'f_\varepsilon(0 | X, Z)],$$

and $f_\varepsilon(0|X, Z)$ is the conditional density of $\varepsilon = Y - D'\alpha - X'\beta(\alpha)$ where $\beta(\alpha)$ is defined by

$$\mathbb{E}[(\tau - \mathbb{1}(Y \leq D'\alpha + X'\beta(\alpha))X] = 0.$$

To proceed with estimation, for a hypothesized value a , we first profile out the coefficients on the exogenous variables as in IQR,

$$\hat{\beta}(a) = \arg \min_{b \in \mathcal{B}} \frac{1}{N} \sum_{i=1}^N \rho_\tau(Y_i - D_i'a - X_i'b). \quad (1.22)$$

We may then plug the solution of (1.22) into (1.20) to form

$$\hat{g}_N(a) = \frac{1}{N} \sum_{i=1}^N g(V_i, a, \hat{\beta}(a), \hat{\delta}(a)), \quad (1.23)$$

where

$$\hat{\delta}(a) = \widehat{M}(a)\widehat{J}^{-1}(a),$$

for

$$\begin{aligned} \widehat{M}(a) &= \frac{1}{Nh_N} \sum_{i=1}^N Z_i X_i' K_{h_N}(Y_i - D_i'a - X_i'\hat{\beta}(a)), \\ \widehat{J}(a) &= \frac{1}{Nh_N} \sum_{i=1}^N X_i X_i' K_{h_N}(Y_i - D_i'a - X_i'\hat{\beta}(a)), \end{aligned}$$

and $K_{h_N}(\cdot)$ a kernel function with bandwidth h_N . Then, we consider the GMM estimator based on the concentrated moments (1.23):

$$\hat{\alpha}(\tau) = \arg \min_{a \in \mathcal{A}} m_N(a),$$

for

$$m_N(a) := N\widehat{g}_N(a)' \widehat{\Sigma}(a, a)^{-1} \widehat{g}_N(a). \quad (1.24)$$

$\widehat{\Sigma}(a, a)$ in $m_N(a)$ is an estimator of the covariance function of the sample concentrated moment functions (1.23) such as

$$\widehat{\Sigma}(a_1, a_2) = \frac{1}{N} \sum_{i=1}^N g(V_i, a_1, \widehat{\beta}(a_1)) g(V_i, a_2, \widehat{\beta}(a_2))'. \quad (1.25)$$

The estimator $\widehat{\alpha}$ is first-order equivalent to the estimator $\tilde{\alpha}$ which employs the moment function:

$$g_\tau^*(V, \alpha) = (\tau - \mathbb{1}(Y \leq D'\alpha + X'\beta_0)) \Psi(\alpha_0, \delta(\alpha_0)).$$

That is, the sample objective function for $\tilde{\alpha}$ uses

$$\widehat{g}_N(\alpha) = \frac{1}{N} \sum_{i=1}^N g_\tau^*(V_i, \alpha) \quad (1.26)$$

where

$$\widehat{\Sigma}(\alpha, \alpha) = \text{E}[g_\tau^*(V_i, \alpha_0) g_\tau^*(V_i, \alpha_0)'].$$

This equivalence holds because the moments possess the Neyman orthogonality property that we discuss later. Moreover, by examining the first-order properties of the IQR estimator we can conclude that $\tilde{\alpha}$ and IQR are first-order equivalent.

1.3.3 Weak Identification Robust Inference

The good behavior of asymptotic approximation results for the point estimators provided in Sections 1.3.1-1.3.2 rely on strong identification of the model parameters as discussed in Section 1.2.2. Because checking these conditions may be difficult, it is useful to have inference procedures that are robust to weak- or non-identification.

Chernozhukov and Hansen (2008) present a simple weak-identification robust inference procedure that results naturally from the IQR estimator. The basic idea underlying this procedure is exactly the relation (1.17) which states that the instruments Z should have no explanatory power in the conventional τ -quantile regression of $Y - D'\alpha_0$ on X and Z at the true value of the structural parameter α_0 . Thus, a valid test of the hypothesis that $\alpha_0 = a$ for some hypothesized a can be obtained by considering a test of the hypothesis that $\gamma(a) = 0$ for $\gamma(a)$ denoting the population value of the τ -quantile regression coefficients defined in Section 1.3.2. Also, note that $W_N(a)$ in (1.19) is simply the standard Wald statistic for testing $\gamma(a) = 0$ and that $W_N(\alpha_0)$ converges in distribution to a $\chi_{\dim(Z)}^2$ regardless of the strength of identification of α_0 ;

see Chernozhukov and Hansen (2008) for details.² It then follows that a valid $(1 - p)\%$ confidence region for α_0 may be constructed as the set

$$\{a \in \mathcal{A} : W_N(a) \leq c_{1-p}\} \quad (1.27)$$

where c_{1-p} is such that $P\left[\chi_{\dim(Z)}^2 > c_{1-p}\right] = p$, and the set may be approximated numerically by considering a 's in the grid $\{a_j, j = 1, \dots, J\}$. Thus, a natural byproduct of solving (1.18) through a grid search is a confidence set for the structural parameter α_0 that is valid regardless of the strength of identification of the parameter. We note that this procedure could also be adapted to be used with the orthogonal scores defined in Section 1.4.1 to provide weak-identification robust inference in settings with high-dimensional X or other settings where robustness to estimation of the nuisance parameter β_0 is a major concern.

The approach of Chernozhukov and Hansen (2008) outlined above is in the spirit of the weak identification robust procedure of Anderson and Rubin (1949). The procedure is relatively simple to implement, but suffers from the same well-known lack of power as other Anderson-Rubin-type statistics in overidentified models under strong identification. To overcome this potential inefficiency, Jun (2008) proposes a different statistic analogous to the proposal of Kleibergen (2005) which is locally efficient under strong identification but may suffer from substantial declines in power against alternatives that are distant from the true parameter value. In the following, we discuss the related approach of Andrews and Mikusheva (2016) which extends the conditional likelihood ratio approach of Moreira (2003) to general nonlinear settings. This approach retains efficiency under strong identification but also maintains good power against distant alternatives.

The Andrews and Mikusheva (2016) approach employs a quasi-likelihood ratio (QLR) statistic as

$$QLR_N(a) = m_N(a) - \inf_{a \in \mathcal{A}} m_N(a) \quad (1.28)$$

where $m_N(a)$ is the GMM objective function (1.24).

Under weak identification, the distribution of $QLR_N(a)$ is non-standard and depends on a nuisance function that is not consistently estimable. Andrews and Mikusheva (2016) provide a sufficient statistic (in LeCam's Gaussian limit experiment)

$$S(a) = \sqrt{N} \left(\hat{g}_N(a) - \hat{\Sigma}(a, \alpha_0) \hat{\Sigma}(\alpha_0, \alpha_0)^{-1} \hat{g}_N(\alpha_0) \right)$$

for this functional nuisance parameter, where $g_N(a)$ and $\hat{\Sigma}(a_1, a_2)$ are defined in (1.23) and (1.25). Andrews and Mikusheva (2016) also outline a procedure

²The same statement would also hold for the GMM objective function based on (1.23) discussed in Section 1.3.2.1.

to simulate the distribution of $QLR(a)$ conditional on $S(a)$ that proceeds as follows. First, draw $\zeta_b^* \sim N\left(0, \widehat{\Sigma}(\alpha_0, \alpha_0)\right)$ for $b = 1, \dots, B$ for a large number B . For each ζ_b^* , the QLR statistic for that draw is then calculated as

$$QLR_{N,b}^*(a) = m_{N,b}^*(a) - \inf_{a \in \mathcal{A}} m_{N,b}^*(a)$$

where

$$m_{N,b}^*(a) = N \widehat{g}_{N,b}^*(a)' \widehat{\Sigma}(a, a)^{-1} \widehat{g}_{N,b}^*(a)$$

for

$$\widehat{g}_{N,b}^*(a) = S(a) + \widehat{\Sigma}(a, \alpha_0) \widehat{\Sigma}(\alpha_0, \alpha_0)^{-1} \zeta_b^*.$$

The simulated distribution then provides an appropriate critical value, $c_{1-p}(S(a))$, for performing a valid p -level test of the null hypothesis that $\alpha_0 = a$ by rejecting when $QLR_N(a) > c_{1-p}(S(a))$. It then follows that a valid $(1-p)\%$ confidence region for α_0 is given by

$$\{a \in \mathcal{A} : QLR_N(a) \leq c_{1-p}(S(a))\}.$$

1.3.4 Finite Sample Inference

The inference procedures reviewed in the previous sections all rely on asymptotic approximations. Chernozhukov et al. (2009) provide a finite sample inference approach which can also be used if the validity of the assumptions necessary to justify these approximations is questionable and is valid in setups with weak or set identification.

Their approach makes use of the fact under the assumptions of the IVQR model, the event $\{Y \leq q(\tau, D, X)\}$ conditional on (Z, X) is distributed exactly as a Bernoulli(τ) random variable regardless of the sample size. This random variable depends only on τ , which is known, and so is pivotal in finite samples. For the GMM objective function $m_N(\theta_0)$ defined in (1.15), this implies that $m_N(\theta_0) \stackrel{d}{=} \tilde{m}_N$ conditional on $\{X_i, Z_i\}_{i=1}^N$, where

$$\tilde{m}_N := \left(\frac{1}{\sqrt{N}} \sum_{i=1}^N (\tau - B_i) \cdot \Psi_i \right)' \Omega_N \left(\frac{1}{\sqrt{N}} \sum_{i=1}^N (\tau - B_i) \cdot \Psi_i \right)$$

and $\{B_i\}_{i=1}^N$ are i.i.d. Bernoulli random variables that are independent of $\{X_i, Z_i\}_{i=1}^N$ and have $E[B_i] = \tau$. This result provides the finite sample distribution of the GMM function $m_N(\theta)$ at $\theta = \theta_0$, which does not depend on any unknown parameters. Given the finite sample distribution of $m_N(\theta_0)$, a p -level test of the null hypothesis that $\theta = \theta_0$ is given by the rule that rejects the null if $m_N(\theta) > c_{1-p}$, where the critical value c_{1-p} is the $(1-p)^{\text{th}}$ quantile of $\tilde{m}_N(\tau)$. It then follows that a valid $(1-p)\%$ joint confidence set for θ is given by

$$\{\theta \in \Theta : m_N(\theta) \leq c_{1-p}\}.$$

We note that inference is simultaneous on all components of θ and that for joint inference the approach is not conservative. Inference about subcomponents of θ such as α may be made by projections and may be conservative.

The chief difficulty with the finite sample approach is computational. Implementing the approach requires inversion of the function $m_N(\theta)$, which may be quite difficult if the number of parameters is large. To alleviate this problem, Chernozhukov et al. (2009) develop suitable MCMC algorithms.

1.4 Advanced Inference with High-Dimensional X

1.4.1 Neyman-Orthogonal Scores

Here we deal with the case where we have high-dimensional covariates. Such cases are common in current high-dimensional data sets where one may see very many potential control variables. High-dimensional covariates also arises in semiparametric problems; for example, we may be interested in a partially linear structural quantile model

$$q(\tau, d, w) = \alpha_0(\tau)d + g(\tau, w)$$

where W is a low-dimensional set of variables and we approximate $g(\tau, w) \approx x' \beta_0(\tau)$ using a collection of approximating functions $x = h(w)$. In settings with high-dimensional X , estimation of $\beta_0(\tau)$ may contaminate estimation of the parameters of interest, $\alpha_0(\tau)$, leading to a breakdown of estimation and inference based directly on (1.13). The potential for contamination is especially acute in high-dimensional settings where some form of regularization will be used to make informative estimation feasible but may arise more generally.

Due to the potentially poor finite sample performance of estimators based directly on (1.13), one might prefer to base estimation and inference on “orthogonal” moment conditions that are relatively insensitive to estimation of the nuisance parameters β_0 . Specifically, we may prefer to base estimation and inference for α_0 on moment functions

$$g(V, \alpha; \eta), \text{ where } V = (Y, D, Z, X)$$

and η denotes nuisance parameters with true values η_0 that include β_0 as a sub-component, that identify α_0 via

$$E[g(V, \alpha_0; \eta_0)] = 0 \tag{1.29}$$

and obey the Neyman orthogonality condition:

$$\partial_\eta E[g(V, \alpha_0; \eta)] \Big|_{\eta=\eta_0} = 0 \tag{1.30}$$

where ∂_η denotes a functional derivative operator. (1.30) is the key orthogonality condition that ensures that the moment conditions defining α_0 are locally insensitive to perturbations in the nuisance parameters. This property results in the first-order properties of estimation and inference of α_0 based on sample analogs to (1.29) being insensitive to estimation of nuisance functions as long as sufficiently high-quality estimators of the nuisance functions are available. The idea of using orthogonal estimating equations goes back at least to Neyman (1959) and Neyman (1979) where they were used in construction of Neyman's celebrated $C(\alpha)$ -statistic. The use of moment conditions satisfying the orthogonality condition (1.30) is crucial for establishing good properties of semi-parametric estimators in modern, high-dimensional estimation settings when regularized estimation or other machine learning tools are used in estimation of nuisance functions; see, e.g. Belloni et al. (2016), Chernozhukov et al. (2015b), and Chernozhukov et al. (2016).

The orthogonal moment functions for the IVQR setting are given by

$$g_\tau(V, \alpha, \eta) = (\tau - \mathbb{1}(Y \leq D'\alpha + X'\beta))\Psi(\alpha, \delta(\alpha)),$$

where $\Psi(\alpha, \delta(\alpha))$ and $\delta(\alpha)$ are defined in Section 1.3.2.1. The nuisance parameter and its true value are then given by

$$\eta := (\beta, \delta(\alpha)), \quad \text{and} \quad \eta_0 := (\beta_0, \delta(\alpha_0)).$$

Observe that the Neyman orthogonality condition holds for these moment conditions because, under appropriate smoothness conditions,

$$\begin{aligned} \partial_\beta \mathbb{E}[g(V, \alpha_0; \eta)] \Big|_{\eta=\eta_0} &= M(\alpha_0) - M(\alpha_0)J^{-1}(\alpha_0)J(\alpha_0) = 0, \\ \partial_\delta \mathbb{E}[g(V, \alpha_0; \eta)] \Big|_{\eta=\eta_0} &= \mathbb{E}[(\tau - \mathbb{1}(Y \leq D'\alpha_0 + X'\beta_0))X] = 0. \end{aligned}$$

1.4.2 Estimation and Inference Using Orthogonal Scores

We start similarly to the IQR estimator by first profiling out the coefficients on exogenous variables using an ℓ_1 -penalized quantile regression estimator to define

$$\widehat{\beta}(a) = \arg \min_{b \in \mathcal{B}} \frac{1}{n} \sum_{i=1}^N \rho_\tau(Y_i - D_i'a - X_i'b) + \lambda \sum_{j=1}^{\dim(b)} \psi_j |b_j|. \quad (1.31)$$

for a hypothesized value a . We then estimate

$$\begin{aligned} \widehat{M}(a) &= \frac{1}{Nh_N} \sum_{i=1}^N Z_i X_i' K_{h_N} \left(Y_i - D_i'a - X_i'\widehat{\beta}(a) \right), \\ \widehat{J}(a) &= \frac{1}{Nh_N} \sum_{i=1}^N X_i X_i' K_{h_N} \left(Y_i - D_i'a - X_i'\widehat{\beta}(a) \right), \end{aligned}$$

for $K_{h_N}(\cdot)$ a kernel function with bandwidth h_N as before. Since $\hat{J}(a)$ is high-dimensional and is not invertible, we may estimate row-components $\delta_j(a)$ of matrix $\delta(a)$ by solving the ℓ_1 -regularized problem

$$\hat{\delta}_j(a) = \arg \min_{\delta} \frac{1}{2} \delta' \hat{J}(a) \delta - \hat{M}_j(a) \delta + \vartheta \|\delta\|_1,$$

where $\hat{M}_j(a)$ is the j -th row of $\hat{M}(a)$, interpreted as a row vector itself, and ϑ is a penalty level. The solution $\hat{\delta}_j(a)$ obeys the Karush-Kuhn-Tucker condition

$$\|\hat{\delta}_j(a)' \hat{J}(a) - \hat{M}_j(a)\|_{\infty} \leq \vartheta, \forall j, \quad (1.32)$$

so we may think of $\hat{\delta}_j(a)$ as a regularized estimator of $M_j(a)J^{-1}(a)$.

Alternatively we can the regularized estimator via Dantzig form of Lasso by minimizing a norm of $\hat{\delta}(a)$ subject to the above constraints (1.32).

We may then plug in the solution of (1.31) to form a concentrated sample moment function analogous to (1.14) as

$$\hat{g}_N(a) = \frac{1}{N} \sum_{i=1}^N \left(\tau - \mathbb{1} \left(Y_i - D_i' a - X_i' \hat{\beta}(a) \leq 0 \right) \right) \Psi(a, \hat{\delta}(a)). \quad (1.33)$$

These concentrated moments can be used to set-up the continuously-updated GMM estimator:

$$\hat{\alpha} = \arg \min_{a \in \mathcal{A}} N \hat{g}_N(a)' \hat{\Sigma}(a, a)^{-1} \hat{g}_N(a),$$

where again $\hat{\Sigma}(a, a)$ is an estimator of the covariance function of the sample concentrated moment functions (1.33). The estimator $\hat{\alpha}$ would then follow standard properties of the infeasible GMM estimator that replaced the estimators $\hat{\beta}(a)$ and $\hat{\delta}(a)$ with their true values β_0 and $\delta(\alpha_0)$ as long as instruments are low dimensional and identification is strong. If the set of instruments was also high-dimensional, further regularization would be called for to make reliable estimation and inference feasible.

We can also directly use the concentrated moments to set-up standard Anderson-Rubin-type inference for α_0 under weak or partial identification as in Section 1.3.3. Similarly, we could base inference from more refined approaches, such as Andrews and Mikusheva (2016), on the concentrated moments. Indeed, we can use these concentrated moments to form a quasi-likelihood ratio (QLR) statistic as

$$QLR_N(a) = N \hat{g}_N(a)' \hat{\Sigma}(a, a)^{-1} \hat{g}_N(a) - \inf_{a \in \mathcal{A}} N \hat{g}_N(a)' \hat{\Sigma}(a, a)^{-1} \hat{g}_N(a). \quad (1.34)$$

Because of the orthogonality property, estimation of the nuisance parameters does not affect the first-order behavior of the empirical moments, so inference based on (1.34) falls back exactly in the setting of Andrews and Mikusheva (2016). One could then employ their approach to compute the critical values

for $QLR_N(a)$ conditional on a sufficient statistic, $c_{1-p}(QLR_N(a))$. It then follows that a valid $(1 - p)\%$ confidence region for α_0 may be constructed by considering a 's in the grid $\{a_j, j = 1, \dots, J\}$ exactly as in approximating (1.27).

1.5 Conclusion

In this chapter, we have reviewed the structural IVQR model developed in Chernozhukov and Hansen (2005) which can be used to estimate causal quantile effects in the presence of endogeneity. The model makes use of instrumental variables that satisfy conventional independence and relevance conditions from the nonlinear instrumental variables literature. Specifically, instruments are assumed to be independent of unobservables associated to potential outcomes but related to endogenous right-hand-side variables in the model. The presence of instruments alone is insufficient to identify QTE, and the IVQR models imposes an additional condition on structural unobservables, termed rank similarity, that restricts the distribution of unobservables in potential outcomes across different potential states of the endogenous variables. Under these conditions, an IV-style moment condition can be derived which then provides a basis for identification and estimation of QTE. We provided two concrete examples of economic models that fall within the IVQR framework.

We then reviewed leading approaches to estimating model parameters and performing inference for QTE within the IVQR model based on the moment conditions implied by the model. Estimation and inference is complicated by the non-smooth and non-convex nature of the IVQR moment conditions. We discuss estimation and inference approaches that attempt to alleviate this issue. We also review approaches to inference which remain valid under weak or even non-identification.

There are, of course, many open areas for research in quantile models with endogeneity. As discussed in Section 1.2.5, Abadie et al. (2002) and Imbens and Newey (2009) offer alternative approaches to identifying QTE by imposing alternate sets of assumptions to those used in the IVQR model. These approaches and the IVQR model are non-nested and further understanding their connections may be interesting. Wüthrich (2014) provides a contribution in this direction by showing the connection between the estimands of both models within the structure of the Abadie et al. (2002) framework. It would also be interesting to analyze the properties of the IVQR estimands when some of the underlying assumptions are violated. Towards this end, Wüthrich (2014) provides a characterization of QTE estimands based on the IVQR model with binary treatments in the absence of rank similarity. Another topic that may deserve further consideration is the systematic analysis of estimation and inference based on the orthogonal moment equations sketched

in Section 1.4.1, especially in high-dimensional settings. We also note that the IVQR model may be useful for uncovering structural objects even if quantile effects are not the chief objects of interest; see, for example, Berry and Haile (2014). It may be interesting to further explore application of the IVQR model and related estimation methods in structural economic applications. Finally, a potentially interesting but more unexplored area may be to think about quantile-like quantities for multivariate outcomes with endogenous covariates.



Bibliography

- Abadie, A., October 1995. Changes in spanish labor income structure during the 1980s: A quantile regression approach, CEMFI Working Paper No. 9521.
- Abadie, A., Angrist, J., Imbens, G., 2002. Instrumental variables estimates of the effect of subsidized training on the quantiles of trainee earnings. *Econometrica* 70 (1), 91–117.
- Amemiya, T., 1982. Two stage least absolute deviations estimators. *Econometrica* 50, 689–711.
- Anderson, T. W., Rubin, H., 1949. Estimation of the parameters of single equation in a complete system of stochastic equations. *Annals of Mathematical Statistics* 20, 46–63.
- Andrews, D. W. K., Shi, X., 2013. Inference based on conditional moment inequalities. *Econometrica* 81 (2), 609–666.
- Andrews, I., Mikusheva, A., 2016. Conditional inference with a functional nuisance parameter. *Econometrica* 84 (4), 1571–1612.
- Belloni, A., Chernozhukov, V., Fernández-Val, I., Hansen, C., 2016. Program evaluation with high-dimensional data, forthcoming *Econometrica*.
URL <https://arxiv.org/abs/1311.2645>
- Berry, S. T., Haile, P. A., 2014. Identification in differentiated products markets using market level data. *Econometrica* 82 (5), 1749–1797.
- Chamberlain, G., 1987. Asymptotic efficiency in estimation with conditional moment restrictions. *Journal of Econometrics* 34 (3), 305–334.
- Chen, X., Pouzo, D., 2009. Efficient estimation of semiparametric conditional moment models with possibly nonsmooth residuals. *Journal of Econometrics* 152 (1), 46–60.
- Chen, X., Pouzo, D., 2012. Estimation of nonparametric conditional moment models with possibly nonsmooth moments. *Econometrica* 80 (1), 277–322.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., 2016. Double machine learning for treatment and causal parameters. arXiv:1608.00060.
URL <https://arxiv.org/abs/1608.00060>

- Chernozhukov, V., Fernandez-Val, I., Melly, B., 2013a. Inference on counterfactual distributions. *Econometrica* 81 (6), pp. 2205–2268.
- Chernozhukov, V., Fernandez-Val, I., Kowalski, A. E., 2015a. Quantile regression with censoring and endogeneity. *Journal of Econometrics* 186 (1), 201–221.
- Chernozhukov, V., Hansen, C., 2004. The effects of 401(k) participation on the wealth distribution: An instrumental quantile regression analysis. *Review of Economics and Statistics* 86(3), 735–751.
- Chernozhukov, V., Hansen, C., 2005. An IV model of quantile treatment effects. *Econometrica* 73 (1), 245–262.
- Chernozhukov, V., Hansen, C., 2006. Instrumental quantile regression inference for structural and treatment effect models. *Journal of Econometrics* 132 (2), 491–525.
- Chernozhukov, V., Hansen, C., 2008. Instrumental variable quantile regression: A robust inference approach. *Journal of Econometrics* 142 (1), 379–398.
- Chernozhukov, V., Hansen, C., 2013. Quantile models with endogeneity. *Annual Review of Economics* 5, 57–81.
- Chernozhukov, V., Hansen, C., Jansson, M., 2009. Finite sample inference for quantile regression models. *Journal of Econometrics* 152 (2), 93–103.
- Chernozhukov, V., Hansen, C., Spindler, M., 2015b. Valid post-selection and post-regularization inference: An elementary, general approach. *Annual Review of Economics* 7, 649–688.
- Chernozhukov, V., Hong, H., 2003. An mcmc approach to classical estimation. *Journal of Econometrics* 115 (2), 293–346.
- Chernozhukov, V., Imbens, G. W., Newey, W. K., 2007. Instrumental variable estimation of nonseparable models. *Journal of Econometrics* 139 (1), 4–14.
- Chernozhukov, V., Lee, S., Rosen, A., 2013b. Intersection bounds: Estimation and inference. *Econometrica* 81 (2), 667–737.
- Chesher, A., 2003. Identification in nonseparable models. *Econometrica* 71 (5), 1405–1441.
- Chesher, A., 2005. Nonparametric identification under discrete variation. *Econometrica* 73 (5), 1525–1550.
- Chesher, A., Rosen, A., Smolinski, K., 2013. An instrumental variable model of multiple discrete choice. *Quantitative Economics* 4 (2), 157–196.

- Chesher, A., Smolinski, K., 2010. Sharp identified sets for discrete variable IV models, ceMMAP Working Paper CWP11/10.
URL <http://www.cemmap.ac.uk/wps/cwp1110.pdf>
- Doksum, K., 1974. Empirical probability plots and statistical inference for nonlinear models in the two-sample case. *Annals of Statistics* 2, 267–277.
- Dong, Y., Shen, S., 2015. Testing for rank invariance or similarity in program evaluation, working paper.
URL http://www.yingyingdong.com/Research/Rank_Invariance_Test_R1_Maintext.pdf
- Frandsen, B. R., Lefgren, L. J., 2015. Testing rank similarity, working paper.
URL <https://economics.byu.edu/frandsen/Documents/testingranksimilarity20151111.pdf>
- Frölich, M., Melly, B., 2013. Unconditional quantile treatment effects under endogeneity. *Journal of Business and Economic Statistics* 31 (3), 346–357.
- Gagliardini, P., Scaillet, O., 2012. Nonparametric instrumental variable estimation of structural quantile effects. *Econometrica* 80 (4), 1533–1562.
- Graddy, K., 1995. Testing for imperfect competition at the Fulton fish market. *Rand Journal of Economics* 26(1), 75–92.
- Hallin, M., Šíman, M., 2016. Multiple-output quantile regression. In: Chernozhukov, V., He, X., Koenker, R., Peng, L. (Eds.), *Handbook of Quantile Regression*. CRC Chapman-Hall, forthcoming.
- Hausman, J. A., 1977. Errors in variables in simultaneous equation models. *Journal of Econometrics* 5 (3), 389–401.
- Heckman, J. J., Smith, J., Clements, N., 1997. Making the most out of programme evaluations and social experiments: Accounting for heterogeneity in programme impacts. *The Review of Economic Studies* 64 (4), 487–535.
- Horowitz, J. L., 1998. Bootstrap methods for median regression models. *Econometrica* 66 (6), 1327–1351.
- Horowitz, J. L., Lee, S., 2007. Nonparametric instrumental variables estimation of a quantile regression model. *Econometrica* 75 (4), 1191–1208.
- Imbens, G. W., Angrist, J. D., 1994. Identification and estimation of local average treatment effects. *Econometrica* 62 (2), 467–475.
- Imbens, G. W., Newey, W. K., 2009. Identification and estimation of triangular simultaneous equations models without additivity. *Econometrica* 77 (5), 1481–1512.
- Jun, S. J., 2008. Weak identification robust tests in an instrumental quantile model. *Journal of Econometrics* 144, 118–138.

- Kaplan, D. M., Sun, Y., 2016. Smoothed estimating equations for instrumental variables quantile regression, forthcoming *Econometric Theory*.
- Kleibergen, F., 2005. Testing parameters in gmm without assuming that they are identified. *Econometrica* 73 (4), 1103–1124.
- Koenker, R., 2005. *Quantile Regression*. Cambridge University Press.
- Koenker, R., Bassett, G. S., 1978. Regression quantiles. *Econometrica* 46, 33–50.
- Koenker, R., Geling, O., 2001. Reappraising medfly longevity: A quantile regression survival analysis. *Journal of the American Statistical Association* 96, 458–468.
- Koenker, R., Ma, L., 2006. Quantile regression methods for recursive structural equation models. *Journal of Econometrics* 134 (2), 471–506.
- Lee, S., 2007. Endogeneity in quantile regression models: A control function approach. *Journal of Econometrics* 141 (2), 1131–1158.
- Lehmann, E. L., 1975. *Nonparametrics: statistical methods based on ranks*. Holden-Day Inc., San Francisco, Calif.
- Machado, J. A. F., Mata, J., 2005. Counterfactual decomposition of changes in wage distributions using quantile regression. *Journal of Applied Econometrics* 20 (4), 445–465.
- Melly, B., 2005. Decomposition of differences in distribution using quantile regression. *Labour Economics* 12 (4), 577 – 590, european Association of Labour Economists 16th Annual Conference, Universidade Nova de Lisboa, Lisbon, 9th – 11th September, 2004.
- Melly, B., Wüthrich, K., 2016. Local quantile treatment effects. In: Chernozhukov, V., He, X., Koenker, R., Peng, L. (Eds.), *Handbook of Quantile Regression*. CRC Chapman-Hall, forthcoming.
- Moreira, M. J., 2003. A conditional likelihood ratio test for structural models. *Econometrica* 71, 1027–1048.
- Newey, W., McFadden, D., 1994. Large sample estimation and hypothesis testing. In: Engle, R. F., McFadden, D. (Eds.), *Handbook of Econometrics*, Vol. IV. Elsevier B. V., pp. 2111–2245.
- Neyman, J., 1959. Optimal asymptotic tests of composite statistical hypotheses. In: Grenander, U. (Ed.), *Probability and Statistics, the Harald Cramer Volume*. New York, Wiley.
- Neyman, J., 1979. $c(\alpha)$ tests and their use. *Sankhya* 41, 1–21.

- Santos, A., 2012. Inference in nonparametric instrumental variables with partial identification. *Econometrica* 80 (1), 213–275.
- Torgovitsky, A., 2015. Identification of nonseparable models using instruments with small support. *Econometrica* 83 (3), 1185–1197.
- Wang, H. J., Yang, Y., 2016. Bayesian quantile regression. In: Chernozhukov, V., He, X., Koenker, R., Peng, L. (Eds.), *Handbook of Quantile Regression*. CRC Chapman-Hall, forthcoming.
- Wüthrich, K., 2014. A comparison of two quantile models with endogeneity, working Paper, Universität Bern, Department of Economics.
URL <http://www.vwl.unibe.ch/papers/dp/dp1408.pdf>