The Econometrics Journal

Econometrics Journal (2016), volume **19**, pp. C61–C94.
doi: 10.1111/ectj.12062

Nonlinear panel data estimation via quantile regressions

MANUEL ARELLANO[†] AND STÉPHANE BONHOMME[‡]

[†]*CEMFI, Calle Casado del Alisal, 5, 28014 Madrid, Spain.*
E-mail: arellano@cemfi.es

[‡]*Department of Economics, University of Chicago, 1126 East 59th Street, Chicago, IL 60637, USA.*
E-mail: sbonhomme@uchicago.edu

First version received: July 2015; final version accepted: March 2016

Summary We introduce a class of quantile regression estimators for short panels. Our framework covers static and dynamic autoregressive models, models with general predetermined regressors and models with multiple individual effects. We use quantile regression as a flexible tool to model the relationships between outcomes, covariates and heterogeneity. We develop an iterative simulation-based approach for estimation, which exploits the computational simplicity of ordinary quantile regression in each iteration step. Finally, an application to measure the effect of smoking during pregnancy on birthweight completes the paper.

Keywords: *Dynamic models, Expectation-maximization, Non-separable heterogeneity, Panel data, Quantile regression.*

1. INTRODUCTION

Nonlinear panel data models are central to applied research. However, despite some recent progress, the literature is still short of answers for panel versions of many models commonly used in empirical work (Arellano and Bonhomme, 2011). More broadly, to date no approach is yet available to specify and estimate general panel data relationships in static or dynamic settings.

In this paper, we rely on quantile regression as a flexible estimation tool for nonlinear panel models. Since Koenker and Bassett (1978), quantile regression techniques have proven useful tools to document distributional effects in cross-sectional settings. Koenker (2005) provides a comprehensive account of these methods. Quantile-based specifications have the ability to deal with complex interactions between covariates and latent heterogeneity, and to provide a rich description of heterogeneous responses of outcomes to variations in covariates. In panel data, quantile methods are particularly well suited as they allow us to build flexible models for the dependence of unobserved heterogeneity on exogenous covariates or initial conditions, and for the feedback processes of covariates in dynamic models with general predetermined regressors.

We consider classes of panel data models with continuous outcomes that satisfy conditional independence restrictions. In static settings, these conditions restrict the time-series dependence of the time-varying disturbances. Imposing some form of dynamic restrictions is necessary in order to separate out what part of the overall time variation is due to unobserved heterogeneity; see Evdokimov (2010) and Arellano and Bonhomme (2012). In dynamic settings, finite-order Markovian set-ups naturally imply conditional independence restrictions. In both static and dynamic settings, results from the literature on nonlinear measurement error models – see Hu and Schennach (2008) and Hu and Shum (2012) – can then be used to provide sufficient conditions for nonparametric identification for a fixed number of time periods.

The main goal of the paper is to develop a tractable estimation strategy for nonlinear panel models. For this purpose, we specify outcomes Y_{it} as a function of covariates X_{it} and latent heterogeneity η_i as

$$Y_{it} = \sum_{k=1}^{K_1} \theta_k(U_{it}) g_k(X_{it}, \eta_i). \tag{1.1}$$

Similarly, we specify the dependence of η_i on covariates $X_i = (X_{i1}', \dots, X_{iT}')'$ as

$$\eta_i = \sum_{k=1}^{K_2} \delta_k(V_i) h_k(X_i). \tag{1.2}$$

Here, $U_{i1}, \dots, U_{iT}, V_i$ are independent uniform random variables, and g and h belong to some families of functions. Outcomes Y_{it} and heterogeneity η_i are monotonic in U_{it} and V_i, respectively, so (1.1) and (1.2) are models of conditional quantile functions.

The g and h are anonymous functions without an economic interpretation. They are just building blocks of flexible models. Objects of interest will be summary measures of derivative effects constructed from the models.

The linear quantile specifications (1.1) and (1.2) allow us to document interactions between covariates and heterogeneity at various quantiles. In particular, (1.2) is a correlated random-effects model that can become arbitrarily flexible as K_2 increases. Linearity in the quantile parameters, though not essential to our approach, is helpful for computational purposes. Moreover, while (1.1) and (1.2) are stated for the static case and a scalar unobserved effect, we show how to extend the framework to allow for dynamics and multidimensional latent components.

The main econometric challenge is that the researcher has no data on heterogeneity η_i. If η_i were observed, then one would simply run an ordinary quantile regression of Y_{it} on the $g_k(X_{it}, \eta_i)$ variables in (1.1). As η_i is not observed, we need to construct some imputations, say M imputed values $\eta_i^{(m)}$, $m = 1, \dots, M$, for each individual in the panel. Having obtained these, we can get estimates by running a quantile regression averaged over imputed values.

For the imputed values to be valid, they have to be draws from the posterior distribution of η_i conditioned on the data, which depends on the parameters to be estimated (θ and δ). Our approach is thus iterative. We start by selecting initial values for conditional quantiles of Y_{it} and η_i, which then allows us to generate imputes of η_i, which we can use to update the quantile parameter estimates, and so on.

A difficulty in applying this idea is that the unknown parameters θ and δ are functions, and hence infinite-dimensional. This is because we need to model the full conditional distribution of outcomes and latent individual effects, as opposed to a single quantile, as is typically the case in applications of ordinary quantile regression. To deal with this issue, we follow Wei and Carroll (2009), and we use a finite-dimensional approximation to θ and δ based on interpolating splines.

In the case of model (1.1) and (1.2), the estimation method works as follows, starting with initial parameter values for $\theta_k(\tau)$ and $\delta_k(\tau)$ and iterating the two steps below until convergence to a stationary distribution.

STEP 1. Given values for $\theta_k(\tau)$ and $\delta_k(\tau)$ on a grid of τ, we compute the implied posterior distribution of the individual effects and draw, for each individual unit i in the sample, a sequence $\eta_i^{(1)}, \ldots, \eta_i^{(M)}$ from that distribution.

STEP 2. With draws of η at hand, we update the parameters $\theta_k(\tau)$ and $\delta_k(\tau)$ by means of two sets of quantile regressions, regressing outcomes Y_{it} on the $g_k(X_{it}, \eta_i^{(m)})$ to update $\theta_k(\tau)$, and regressing the individual draws $\eta_i^{(m)}$ on the $h_k(X_i)$ to update $\delta_k(\tau)$.

The resulting algorithm is a variant of the expectation-maximization (EM) algorithm of Dempster et al. (1977), sometimes referred to as stochastic EM. The sequence of parameter estimates converges to an ergodic Markov chain in the limit. Following Nielsen (2000a, 2000b) we characterize the asymptotic distribution of our sequential simulated method-of-moments estimator based on M imputations. A difference with most applications of EM-type algorithms is that we do not update parameters in each iteration using maximum likelihood, but using quantile regressions.[1] This is an important feature of our approach, as the fact that quantile regression estimates can be computed in a quantile-by-quantile fashion, and the convexity of the quantile regression objective function, make each parameter update in Step 2 in the above algorithm fast and reliable.

We apply our estimator to assess the effect of smoking during pregnancy on a child's birthweight. Following Abrevaya (2006), we allow for mother-specific fixed effects in estimation. Both nonlinearities and unobserved heterogeneity are thus allowed for, using our panel data quantile regression estimator. We find that, while allowing for time-invariant mother-specific effects decreases the magnitude of the negative coefficient of smoking, the latter remains sizable, especially at low birthweights, and exhibits substantial heterogeneity across mothers.

Literature review and outline

Starting with Koenker (2004), most panel data approaches to date proceed in a quantile-by-quantile fashion, and include individual indicators as additional covariates in the quantile regression. As shown by some recent work, however, this fixed-effects approach faces special challenges when applied to quantile regression. Galvao et al. (2012) and Arellano and Weidner (2015) study the large N, T properties of the fixed-effects quantile regression estimator, and show that it may suffer from large biases in short panels. Rosen (2012) shows that a fixed-effects model for a single quantile may not be point-identified. Recent related contributions are Lamarche (2010), Galvao (2011) and Canay (2011). In contrast, our approach relies on specifying a semiparametric model for individual effects given covariates and initial conditions, as in (1.2). As a result, in this paper, the analysis is conducted for fixed T, as N tends to infinity.

Our approach is closer in spirit to other random-effects approaches in the literature. For example, Abrevaya and Dahl (2008) consider a correlated random-effects model to study the effects of smoking and pre-natal care on birthweight. Their approach mimics control function

[1] Related sequential method-of-moments estimators are considered in Arcidiacono and Jones (2003), Arcidiacono and Miller (2011) and Bonhomme and Robin (2009), among others. Elashoff and Ryan (2004) present an algorithm for accommodating missing data in situations where a natural set of estimating equations exists for the complete data setting.

approaches used in linear panel models. Geraci and Bottai (2007) consider a random-effects approach for a single quantile, assuming that the outcome variable is distributed as an asymmetric Laplace distribution conditional on covariates and individual effects. Recent related approaches to quantile panel data models include Chernozhukov et al. (2013, 2015) and Graham et al. (2015). These approaches are non-nested with ours. In particular, they will generally not recover the quantile effects we focus on in this paper. More broadly, compared to existing work, our aim is to build a framework that can deal with general nonlinear and dynamic relationships, thus providing an extension of standard linear panel data methods to nonlinear settings.

The analysis also relates to method-of-moments estimators for models with latent variables. Compared to Schennach (2014), here we rely on conditional moment restrictions and focus on cases where the entire model specification is point-identified. Finally, our analysis is most closely related to Wei and Carroll (2009), who proposed a consistent estimation method for cross-sectional linear quantile regression subject to covariate measurement error. A key difference with Wei and Carroll is that, in our set-up, the conditional distribution of individual effects is unknown, and needs to be estimated along with the other parameters of the model.

The outline of the paper is as follows. In Section 2, we present static models and discuss identification. In Section 3, we present our estimation method and study some of its properties. In Section 4, we extend the approach to dynamic settings. In Section 5, we show how our method can be used to estimate average marginal effects, which are of interest in a number of applications. In Section 6, we present the empirical illustration. Lastly, we conclude in Section 7. Proofs and further discussion are contained in the Appendices. Computer codes implementing the method are available as Supporting Information.

2. QUANTILE MODELS FOR PANEL DATA

In this section, we start by introducing a class of static panel data models. At the end of the section, we provide conditions for nonparametric identification.

2.1. Model and assumptions

Outcome variables. Let $Y_i = (Y_{i1}, \ldots, Y_{iT})'$ denote a sequence of T scalar continuous outcomes for individual i, and let $X_i = (X'_{i1}, \ldots, X'_{iT})'$ denote a sequence of strictly exogenous regressors, which may contain a constant. Let η_i denote a q-dimensional vector of individual-specific effects, and let U_{it} denote a scalar error term. We specify the conditional quantile response function of Y_{it} given X_{it} and η_i as follows:

$$Y_{it} = Q_Y(X_{it}, \eta_i, U_{it}), \qquad i = 1, \ldots, N, \quad t = 1, \ldots, T. \tag{2.1}$$

Model (2.1) can be used to empirically document nonlinear and heterogeneous effects of covariates. In our illustration to smoking and birthweight, the model allows smoking effects to differ across mothers (through the dependence on η_i) and along the distribution of birthweights (through the dependence on U_{it}). In Section 5, we describe a set of treatment effect parameters that our method allows us to estimate.

We make the following assumption.

ASSUMPTION 2.1. (OUTCOMES) *(a) U_{it} follows a standard uniform distribution, independent of (X_i, η_i); (b) $\tau \mapsto Q_Y(x, \eta, \tau)$ is strictly increasing on* $(0, 1)$*, for almost all* (x, η) *in the support of (X_{it}, η_i); (c) for all $t \neq s$, U_{it} is independent of U_{is}.*

Assumption 2.1(a) contains two parts. First, U_{it} is assumed independent of the full sequence $X_{i1}, \ldots, X_{iT}$, and independent of individual effects. Strict exogeneity of X can be relaxed to allow for predetermined covariates (see Section 4). Second, the marginal distribution of U_{it} is normalized to be uniform on the unit interval. Assumption 2.1(b) guarantees that outcomes have absolutely continuous distributions. Together, Assumption 2.1(a) and (b) imply that, for all $\tau \in (0, 1)$, $Q_Y(X_{it}, \eta_i, \tau)$ is the τ-conditional quantile of Y_{it} given (X_i, η_i).[2]

Assumption 2.1(c) imposes independence restrictions on the process $U_{i1}, \ldots, U_{iT}$. Restricting the dynamics of error variables U_{it} is needed when aiming at separating the time-varying unobserved errors U_{it} from the time-invariant unobserved individual effects η_i. In Assumption 2.1(c), U_{it} are assumed to be independent over time. In Section 4, we develop various extensions of the model that allow for dynamic effects. Finally, although we have assumed in (2.1) that Q_Y does not depend on time, one could easily allow $Q_Y = Q_Y^t$ to depend on t, reflecting, for example, age or calendar time effects depending on the application.

Unobserved heterogeneity. Next, we specify the conditional quantile response function of η_i given X_i as

$$\eta_i = Q_\eta(X_i, V_i), \qquad i = 1, \ldots, N. \tag{2.2}$$

Provided η_i is continuously distributed given X_i and Assumption 2.2 below holds, (2.2) is a representation that comes without loss of generality, corresponding to a fully unrestricted correlated random-effects specification.

ASSUMPTION 2.2. (INDIVIDUAL EFFECTS) *(a) V_i follows a standard uniform distribution, independent of X_i; (b) $\tau \mapsto Q_\eta(x, \tau)$ is strictly increasing on* $(0, 1)$*, for almost all x in the support of X_i.*

2.2. *Examples*

We next describe several examples to illustrate the static set-up introduced above.

EXAMPLE 2.1. (LOCATION SCALE) As a first special case of model (2.1), consider the following panel generalization of the location-scale model (He, 1997)

$$Y_{it} = X_{it}'\beta + \eta_i + (X_{it}'\gamma + \mu\eta_i)\varepsilon_{it}, \tag{2.3}$$

[2] Indeed, using Assumption 2.1(a) and (b), we have

$$\begin{aligned}\Pr(Y_{it} \leq Q_Y(X_{it}, \eta_i, \tau)|X_i, \eta_i) &= \Pr(Q_Y(X_{it}, \eta_i, U_{it}) \leq Q_Y(X_{it}, \eta_i, \tau)|X_i, \eta_i)\\ &= \Pr(U_{it} \leq \tau|X_i, \eta_i) = \tau.\end{aligned}$$

where ε_{it} are independent and identically distributed (i.i.d.) across periods, and independent of all regressors and individual effects.[3] Denoting $U_{it} = F(\varepsilon_{it})$, where F is the cumulative distribution function (CDF) of ε_{it}, the conditional quantiles of Y_{it} are given by

$$Q_Y(X_{it}, \eta_i, \tau) = X_{it}'\beta + \eta_i + (X_{it}'\gamma + \mu\eta_i)F^{-1}(\tau), \qquad \tau \in (0, 1).$$

EXAMPLE 2.2. (PANEL QUANTILE REGRESSION) Consider, next, the following linear quantile specification with scalar η_i, which generalizes (2.3):

$$Y_{it} = X_{it}'\beta(U_{it}) + \eta_i\gamma(U_{it}). \tag{2.4}$$

Given Assumption 2.1(a) and (b), the conditional quantiles of Y_{it} are given by

$$Q_Y(X_{it}, \eta_i, \tau) = X_{it}'\beta(\tau) + \eta_i\gamma(\tau).$$

Model (2.4) is a panel data generalization of the classical linear quantile model of Koenker and Bassett (1978). Were we to observe the individual effects η_i along with the covariates X_{it}, it would be reasonable to postulate a model of this form. It is instructive to compare model (2.4) with the following more general but different type of model

$$Y_{it} = X_{it}'\beta(U_{it}) + \eta_i(U_{it}), \tag{2.5}$$

where $\eta_i(\tau)$ is an individual-specific nonparametric function of τ. Koenker (2004) and subsequent fixed-effects approaches considered this more general model. Unlike (2.4), the presence of the process $\eta_i(\tau)$ in (2.5) introduces an element of nonparametric functional heterogeneity in the conditional distribution of Y_{it}. In contrast, a key aspect of our approach is that we view η as missing data, and introduce them as additional (latent) covariates in the quantile regression model.

The term $\eta_i(U_{it})$ in model (2.5) can be regarded as a function of U_{it} and a vector of unobserved individual effects of unspecified dimension. In this way, model (2.5) allows for multiple individual characteristics that affect differently individuals with different error rank U_{it}. However, while being agnostic about the number of unobserved individual factors affecting outcomes is attractive, sometimes substantive reasons suggest that only a small number of underlying factors play a role. Additionally, as our analysis makes clear, whether one uses a quantile model with a different individual effect at each quantile or a model with a small number of unobserved effects has implications for identification.[4]

In order to complete model (2.4), one can use another linear quantile specification for the conditional distribution of individual effects:

$$\eta_i = X_i'\delta(V_i). \tag{2.6}$$

Given Assumption 2.2, the conditional quantiles of η_i are then given by

$$Q_\eta(X_i, \tau) = X_i'\delta(\tau).$$

[3] A generalization of (2.3) that allows for two-dimensional individual effects (as in Example 2.3) is

$$Y_{it} = X_{it}'\beta + \eta_{i1} + (X_{it}'\gamma + \eta_{i2})\varepsilon_{it}.$$

[4] As mentioned in the introduction, Rosen (2012) shows that a fixed-effects model for a single quantile may not be point-identified.

Model (2.6) corresponds to a correlated random-effects approach. However, it is more flexible than alternative specifications in the literature. A commonly used specification is (Chamberlain, 1984)

$$\eta_i = X_i'\mu + \sigma\varepsilon_i, \quad \varepsilon_i | X_i \sim \mathcal{N}(0, 1). \tag{2.7}$$

For example, in contrast with (2.7), model (2.6) is fully nonparametric in the absence of covariates (i.e., when an independent random-effects specification is assumed). Model (2.6) and its extensions based on series specifications may also be of interest in other nonlinear panel data models, where the outcome equation does not follow a quantile model. We return to this point in the conclusion.

EXAMPLE 2.3. (MULTIDIMENSIONAL HETEROGENEITY) Model (2.4) can easily be modified to allow for more general interactions between observables and unobservables, thus permitting the effects of covariates to be heterogeneous at different quantiles. A random coefficients generalization that allows for heterogeneous effects is

$$Q_Y(X_{it}, \eta_i, \tau) = X_{it}'\beta(\tau) + \gamma_1(\tau)\eta_{i1} + X_{it}'\gamma_2(\tau)\eta_{i2}, \tag{2.8}$$

where $\eta_i = (\eta_{i1}, \eta_{i2})'$ is bivariate.

In order to extend (2.6) to the case with bivariate unobserved heterogeneity, it is convenient to assume a triangular structure such as

$$\begin{aligned} \eta_{i1} &= X_i'\delta_{11}(V_{i1}), \\ \eta_{i2} &= \eta_{i1}\delta_{21}(V_{i2}) + X_i'\delta_{22}(V_{i2}), \end{aligned} \tag{2.9}$$

where V_{i1} and V_{i2} follow independent standard uniform distributions. Though not invariant to permutation of (η_{i1}, η_{i2}), except if fully nonparametric, model (2.9) provides a flexible specification for the bivariate conditional distribution of (η_{i1}, η_{i2}) given X_i.[5]

2.3. Nonparametric identification

The class of panel data models introduced above satisfies conditional independence restrictions, as period-specific outcomes $Y_{i1}, \ldots, Y_{iT}$ are mutually independent conditional on exogenous covariates and individual heterogeneity X_i, η_i. A body of work, initially developed in the context of nonlinear measurement error models, has established nonparametric identification results in related models under conditional independence restrictions; see Hu (2015) for a recent survey. Here we show how the result in Hu and Schennach (2008) can be used to show nonparametric identification. In Section 4, we build on Hu and Shum (2012) to provide conditions for identification in dynamic models, under Markovian restrictions.

Consider model (2.1) and (2.2), with a scalar unobserved effect η_i. At least three periods are needed for identification, and we set $T = 3$. In the case where η_i is multivariate, identification requires using additional time periods (see below). Throughout we use f_Z and $f_{Z|W}$ as generic notation for the distribution function of a random vector Z and for the conditional distribution of Z given W, respectively.

[5] It is worth pointing out that quantiles appear not to generalize easily to the multivariate case.

Under conditional independence over time (Assumption 2.1(c)), we have, for all y_1, y_2, y_3, $x = (x_1', x_2', x_3')'$ and η:

$$f_{Y_1,Y_2,Y_3|\eta,X}(y_1, y_2, y_3 \mid \eta, x) = f_{Y_1|\eta,X}(y_1 \mid \eta, x) f_{Y_2|\eta,X}(y_2 \mid \eta, x) f_{Y_3|\eta,X}(y_3 \mid \eta, x). \quad (2.10)$$

Hence the data distribution function relates to the densities of interest as follows:

$$f_{Y_1,Y_2,Y_3|X}(y_1, y_2, y_3 \mid x) = \int f_{Y_1|\eta,X}(y_1 \mid \eta, x) f_{Y_2|\eta,X}(y_2 \mid \eta, x) f_{Y_3|\eta,X}(y_3 \mid \eta, x) f_{\eta|X}(\eta \mid x) d\eta. \quad (2.11)$$

The goal is the identification of $f_{Y_1|\eta,X}$, $f_{Y_2|\eta,X}$, $f_{Y_3|\eta,X}$ and $f_{\eta|X}$ given knowledge of $f_{Y_1,Y_2,Y_3|X}$.

The setting of (2.11) is formally equivalent (conditional on x) to the instrumental variables set-up of Hu and Schennach (2008) for nonclassical nonlinear errors-in-variables models. Specifically, according to the terminology of Hu and Schennach, Y_{i3} would be the outcome variable, Y_{i2} would be the mismeasured regressor, Y_{i1} would be the instrumental variable and η_i would be the latent, error-free regressor. We closely rely on their analysis and make the following assumption.

ASSUMPTION 2.3. (IDENTIFICATION) *Almost surely in covariate values x, (a) the joint density $f_{Y_1,Y_2,Y_3,\eta|X=x}$ is bounded, as well as all its joint and marginal densities; (b) for all $\eta_1 \neq \eta_2$,* $\Pr[f_{Y_3|\eta,X}(Y_{i3}|\eta_1, x) \neq f_{Y_3|\eta,X}(Y_{i3}|\eta_2, x) \mid X_i = x] > 0$*; (c) there exists a known functional Γ_x such that $\Gamma_x(f_{Y_2|\eta,X}(\cdot|\eta, x)) = \eta$; (d) the linear operators $L_{Y_2|\eta,x}$ and $L_{Y_1|Y_2,x}$, associated with the conditional densities $f_{Y_2|\eta,X=x}$ and $f_{Y_1|Y_2,X=x}$, respectively, are injective.*

Assumption 2.3(a) requires bounded densities. Assumption 2.3(b) requires that $f_{Y_3|\eta,X}$ be non-identical at different values of η. Assumption 2.3(c) imposes a centred measure of location on $f_{Y_2|\eta,X=x}$. In Example 2.2, the following normalization implies Assumption 2.3(c),

$$\int_0^1 \beta_0(\tau) d\tau = 0, \quad \text{and} \quad \int_0^1 \gamma(\tau) d\tau = 1, \quad (2.12)$$

where $\beta_0(\tau)$ corresponds to the coefficient of the constant in X_{it}. We use (2.12) in our empirical implementation.[6] Lastly, Assumption 2.3(d) is an injectivity condition. As pointed out by Hu and Schennach (2008), injectivity is closely related to completeness conditions commonly assumed in the literature on nonparametric instrumental variables. Similarly to completeness, injectivity is a high-level condition.[7] In Appendix A, we further discuss the different parts of Assumption 2.3.

We then have the following result, which is a direct application of the identification theorem in Hu and Schennach (2008).

PROPOSITION 2.1. (HU AND SCHENNACH (2008)) *Let Assumptions 2.1, 2.2 and 2.3 hold. Then all conditional densities $f_{Y_1|\eta,X=x}$, $f_{Y_2|\eta,X=x}$, $f_{Y_3|\eta,X=x}$ and $f_{\eta|X=x}$, are nonparametrically identified for almost all x.*

This result places no restrictions on the form of $f_{Y_t|\eta,X=x}$, thus allowing for general distributional time effects.

[6] In fact, Assumption 2.3(c) is also implied by (2.12) in the following model with first-order interactions, a version of which we estimate in the empirical application: $Y_{it} = X_{it}'\beta(U_{it}) + \eta_i X_{it}'\gamma(U_{it})$.

[7] See, e.g., Canay et al. (2013) for results on the testability of completeness assumptions, and D'Haultfoeuille (2011), Andrews (2011) and Hu and Shiu (2012) for primitive conditions in several settings.

Lastly, the identification result extends to models with multiple, q-dimensional individual effects η_i, by taking a larger $T > 3$. For example, with $T = 5$, it is possible to apply the identification theorem of Hu and Schennach (2008) to a bivariate η_i using (Y_{i1}, Y_{i2}) instead of Y_{i1}, (Y_{i3}, Y_{i4}) instead of Y_{i2}, and Y_{i5} instead of Y_{i3}.

3. QUANTILE REGRESSION ESTIMATORS

In this section, we introduce our estimation strategy and discuss several of its statistical properties.

3.1. Model specification and moment restrictions

We specify the conditional quantile function of Y_{it} in (2.1), for scalar η_i, as

$$Q_Y(X_{it}, \eta_i, \tau) = W_{it}(\eta_i)'\theta(\tau). \tag{3.1}$$

In (3.1), the vector $W_{it}(\eta_i)$ contains a finite number of functions of X_{it} and η_i. One possibility is to adopt a simple linear quantile specification as in Example 2.2, in which case $W_{it}(\eta_i) = (X_{it}', \eta_i)'$. A more flexible approach is to use a series specification of the quantile function as in (1.1), and to set $W_{it}(\eta_i) = (g_1(X_{it}, \eta_i), \dots, g_{K_1}(X_{it}, \eta_i))'$ for a set of K_1 functions $g_1, \dots, g_{K_1}$. In practice, one can use orthogonal polynomials, wavelets or splines, for example; see Chen (2007) for a comprehensive survey of sieve methods.

Similarly, we specify the conditional quantile function of η_i in (2.2) as

$$Q_\eta(X_i, \tau) = Z_i'\delta(\tau). \tag{3.2}$$

In (3.2), the vector Z_i contains a finite number of functions of covariates X_i, such as $Z_i = (h_1(X_i), \dots, h_{K_2}(X_i))$ for a set of K_2 functions $h_1, \dots, h_{K_2}$. Extensions to vector-valued η_i can be done along the lines of Example 2.3.

The posterior density of the individual effects $f_{\eta|Y,X}$ plays an important role in the analysis. It is given by

$$f_{\eta|Y,X}(\eta \mid y, x; \theta(\cdot), \delta(\cdot)) = \frac{\prod_{t=1}^T f_{Y_t|X_t,\eta}(y_t \mid x_t, \eta; \theta(\cdot)) f_{\eta|X}(\eta \mid x; \delta(\cdot))}{\int \prod_{t=1}^T f_{Y_t|X_t,\eta}(y_t \mid x_t, \widetilde{\eta}; \theta(\cdot)) f_{\eta|X}(\widetilde{\eta} \mid x; \delta(\cdot)) d\widetilde{\eta}}, \tag{3.3}$$

where we have used conditional independence in Assumption 2.1(c), and we have explicitly indicated the dependence of the various densities on model parameters.

Let $\psi_\tau(u) = \tau - \mathbf{1}\{u < 0\}$. The function ψ_τ is the first derivative (outside the origin) of the 'check' function ρ_τ, which is familiar from the quantile regression literature (Koenker and Basset, 1978):

$$\rho_\tau(u) = (\tau - \mathbf{1}\{u < 0\})u, \qquad \psi_\tau(u) = \frac{d\rho_\tau(u)}{du}.$$

In order to derive the main moment restrictions, we start by noting that, for all $\tau \in (0, 1)$, the following infeasible moment restrictions hold, as a direct implication of Assumptions 2.1 and 2.2:

$$E\Big[\sum_{t=1}^{T} W_{it}(\eta_i)\psi_\tau(Y_{it} - W_{it}(\eta_i)'\theta(\tau))\Big] = 0, \tag{3.4}$$

and

$$E[Z_i \psi_\tau(\eta_i - Z_i'\delta(\tau))] = 0. \tag{3.5}$$

Indeed, (3.4) is the first-order condition associated with the infeasible population quantile regression of Y_{it} on $W_{it}(\eta_i)$. Similarly, (3.5) corresponds to the infeasible quantile regression of η_i on Z_i.

Applying the law of iterated expectations to (3.4) and (3.5), respectively, we obtain the following integrated moment restrictions, for all $\tau \in (0, 1)$:

$$E\Big[\int \Big(\sum_{t=1}^{T} W_{it}(\eta)\psi_\tau(Y_{it} - W_{it}(\eta)'\theta(\tau))\Big) f(\eta \mid Y_i, X_i; \theta(\cdot), \delta(\cdot)) d\eta\Big] = 0, \tag{3.6}$$

and

$$E\Big[\int \big(Z_i \psi_\tau(\eta - Z_i'\delta(\tau))\big) f(\eta \mid Y_i, X_i; \theta(\cdot), \delta(\cdot)) d\eta\Big] = 0. \tag{3.7}$$

Here, and in the rest of the analysis, we use f as a shorthand for the posterior density $f_{\eta|Y,X}$.

It follows from (3.6) and (3.7) that, if the posterior density of the individual effects were known, then estimating the model's parameters could be done using two sets of linear quantile regressions, weighted by the posterior density. However, as the notation makes clear, the posterior density in (3.3) depends on the entire processes $\theta(\cdot)$ and $\delta(\cdot)$. Specifically, for absolutely continuous conditional densities of outcomes and individual effects, we have

$$f_{Y_t|X_t,\eta}(y_t \mid x_t, \eta; \theta(\cdot)) = \lim_{\epsilon \to 0} \frac{\epsilon}{w_t(\eta)'(\theta(u_t + \epsilon) - \theta(u_t))}, \tag{3.8}$$

and

$$f_{\eta|X}(\eta \mid x; \delta(\cdot)) = \lim_{\epsilon \to 0} \frac{\epsilon}{z'(\delta(v + \epsilon) - \delta(v))}, \tag{3.9}$$

where u_t and v are defined by $w_t(\eta)'\theta(u_t) = y_t$ and $z'\delta(v) = \eta$, respectively. Equations (3.8) and (3.9) come from the fact that the density of a random variable and the derivative of its quantile function are the inverse of each other.

The dependence of the posterior density on the entire set of model parameters makes it impossible to directly recover $\theta(\tau)$ and $\delta(\tau)$ in (3.6) and (3.7) in a τ-by-τ fashion. The main idea of the algorithm that we present in the next subsection is to circumvent this difficulty by iterating back-and-forth between computation of the posterior density, and computation of the model's parameters given the posterior density. The latter is easy to do, as it is based on weighted quantile regressions. Similar ideas have been used in the literature; see, e.g. Arcidiacono and Jones (2003). However, an additional difficulty in our case is that the posterior density depends on a continuum of parameters. In order to develop a practical approach, we now introduce a finite-dimensional, tractable approximating model.

Parametric specification. Building on Wei and Carroll (2009), we approximate $\theta(\cdot)$ and $\delta(\cdot)$ using splines, with L knots $0 < \tau_1 < \tau_2 < \ldots < \tau_L < 1$. A practical possibility is to use piecewise-linear splines as in Wei and Carroll, but other choices are possible, such as cubic splines or shape-preserving B-splines. When using interpolating splines, the approximation argument requires suitable smoothness assumptions on $\theta(\tau)$ and $\delta(\tau)$ as functions of $\tau \in (0, 1)$. For fixed L, the spline specification can be seen as an approximation to the underlying quantile functions.

Let us define $\xi = (\xi_A', \xi_B')'$, where

$$\xi_A = (\theta(\tau_1)', \theta(\tau_2)', \ldots, \theta(\tau_L)')' \quad \text{and} \quad \xi_B = (\delta(\tau_1)', \delta(\tau_2)', \ldots, \delta(\tau_L)')'.$$

The approximating model depends on the finite-dimensional parameter vector ξ that is used to construct interpolating splines. The associated likelihood function and density of individual effects are then denoted as $f_{Y_t|X_t,\eta}(y_t \mid x_t, \eta; \xi_A)$ and $f_{\eta|X}(\eta \mid x; \xi_B)$, respectively, and the implied posterior density is

$$f(\eta \mid y, x; \xi) = \frac{\prod_{t=1}^T f_{Y_t|X_t,\eta}(y_t \mid x_t, \eta; \xi_A) f_{\eta|X}(\eta \mid x; \xi_B)}{\int \prod_{t=1}^T f_{Y_t|X_t,\eta}(y_t \mid x_t, \widetilde{\eta}; \xi_A) f_{\eta|X}(\widetilde{\eta} \mid x; \xi_B) d\widetilde{\eta}}. \tag{3.10}$$

The approximating densities take closed-form expressions when using piecewise-linear splines. Moreover, when implementing the algorithm in practice we augment the specification with parametric models in the tail intervals of the intercepts of $\theta(\tau)$ and $\delta(\tau)$. In this case, the estimation algorithm needs to be modified slightly. See Section 6.1 for a discussion of implementation.

Finally, the integrated moment restrictions of the approximating model are, for all $\ell = 1, \ldots, L$

$$E\Big[\int \Big(\sum_{t=1}^{T} W_{it}(\eta)\psi_{\tau_\ell}(Y_{it} - W_{it}(\eta)'\theta(\tau_\ell))\Big) f(\eta \mid Y_i, X_i; \xi) d\eta\Big] = 0, \tag{3.11}$$

and

$$E\Big[\int \big(Z_i \psi_{\tau_\ell}(\eta - Z_i'\delta(\tau_\ell))\big) f(\eta \mid Y_i, X_i; \xi) d\eta\Big] = 0. \tag{3.12}$$

3.2. Estimation algorithm

Let (Y_i, X_i'), $i = 1, \ldots, N$, be an i.i.d. sample. Motivated by the integrated moment restrictions (3.11) and (3.12), we propose to estimate the model's parameters by using an iterative method. In practice, we use a simulation-based approach to replace the integrals in (3.11) and (3.12) by sums. Starting with initial parameter values $\widehat{\xi}^{(0)}$, we iterate the following two steps in a stochastic EM algorithm until convergence to a stationary distribution.

STEP 1. For all $i = 1, \ldots, N$, compute the posterior density

$$\widehat{f}_i^{(s)}(\eta) = f(\eta \mid Y_i, X_i; \widehat{\xi}^{(s)}), \tag{3.13}$$

and draw M values $\eta_i^{(1)}, \ldots, \eta_i^{(M)}$ from $\widehat{f}_i^{(s)}$.

STEP 2. Solve, for $\ell = 1, \ldots, L$,

$$\widehat{\theta}(\tau_\ell)^{(s+1)} = \underset{\theta}{\operatorname{argmin}} \sum_{i=1}^{N} \sum_{m=1}^{M} \sum_{t=1}^{T} \rho_{\tau_\ell}(Y_{it} - W_{it}(\eta_i^{(m)})'\theta),$$

$$\widehat{\delta}(\tau_\ell)^{(s+1)} = \underset{\delta}{\operatorname{argmin}} \sum_{i=1}^{N} \sum_{m=1}^{M} \rho_{\tau_\ell}(\eta_i^{(m)} - Z_i'\delta).$$

This sequential simulated method-of-moment method is related to, but different from, the standard EM algorithm (Dempster et al., 1977). As in EM, the algorithm iterates back-and-forth between computation of the posterior density of the individual effects ('E'-step) and computation of the parameters given the posterior density ('M'-step). Unlike in EM, however, in the second step of the algorithm (the M-step), estimation is not based on a likelihood function, but on the check function of quantile regression.

Proceeding in this way has two major computational advantages compared to maximizing the full likelihood of the approximating model. First, as opposed to the likelihood function, which is a complicated function of all quantile regression coefficients, the M-step problem nicely decomposes into L different τ_ℓ-specific subproblems. Secondly, using the check function yields a globally convex objective function in each step. In fact, the M-step simply consists of $2L$ ordinary quantile regressions, where the simulated values of the individual effects are treated, in turn, as covariates and dependent variables.

At the same time, two features of the standard EM algorithm differ in our sequential method-of-moment method. First, as our algorithm is not likelihood-based, the resulting estimator will not be efficient in general, even as the number of draws M tends to infinity.[8]

Second, unlike in deterministic versions of EM, in the E-step we draw M values for the individual effects according to their posterior density $\widehat{f}_i^{(s)}(\eta) = f(\eta \mid Y_i, X_i; \widehat{\xi}^{(s)})$. We use a random-walk Metropolis–Hastings sampler for this purpose, but other choices are possible (such as particle filter methods).[9] An advantage of Metropolis–Hastings over grid approximations and importance sampling weights is that the integral in the denominator of the posterior density of η is not needed. The output of this algorithm is a Markov chain. In practice, we stop the chain after a large number of iterations and we report an average across the last $\widetilde{S}$ values: $\widehat{\xi} = (1/\widetilde{S}) \sum_{s=S-\widetilde{S}+1}^{S} \widehat{\xi}^{(s)}$.

In each iteration of the algorithm, the draws $\eta_i^{(1)}, \ldots, \eta_i^{(M)}$ are randomly redrawn. This approach, sometimes referred to as stochastic EM, thus differs from the simulated EM algorithm of McFadden and Ruud (1994) where the same underlying uniform draws are used in each iteration. Nielsen (2000a, 2000b) studies and compares various statistical properties of simulated EM and stochastic EM in a likelihood context. In particular, he provides conditions under which the Markov chain output of stochastic EM is ergodic. As M tends to infinity, the sum converges to the true integral. The problem is then smooth (because of the integral with respect to η). Building on Nielsen's work, we next analyse the statistical properties of estimators based on fixed-M and large-M versions of the algorithm.

[8] This loss of efficiency relative to maximum likelihood is similar to the one documented in Arcidiacono and Jones (2003), for example.

[9] Note that the posterior density is non-negative by construction. In particular, drawing from $\widehat{f}_i^{(s)}(\eta)$ automatically produces rearrangement of the various quantile curves, as in Chernozhukov et al. (2010).

3.3. *Asymptotic properties*

We now discuss the asymptotic properties of the estimation algorithm. Throughout, T is fixed while N tends to infinity.

Parametric inference. We start by discussing the asymptotic properties of the estimator based on the stochastic EM algorithm, for fixed number of draws M, in the case where the parametric model is assumed to be correctly specified. That is, K_1, K_2 (the number of series terms) and L (the size of the grid on the unit interval) are held fixed as N tends to infinity. In the following subsection, we study consistency as K_1, K_2 and L tend to infinity with N, in the large-M limit.

Nielsen (2000a) studies the statistical properties of the stochastic EM algorithm in a likelihood case. He provides conditions under which the Markov chain $\widehat{\xi}^{(s)}$ is ergodic, for a fixed sample size. In addition, he also characterizes the asymptotic distribution of $\sqrt{N}(\widehat{\xi}^{(s)} - \overline{\xi})$ as N increases, where $\overline{\xi}$ denotes the population parameter vector.

In Appendix B, we rely on Nielsen's work to characterize the asymptotic distribution of $\widehat{\xi}^{(s)} = ((\widehat{\theta}^{(s)})', (\widehat{\delta}^{(s)})')'$ in our model, where the optimization step is not likelihood-based but relies on quantile-based estimating equations. Specifically, if s corresponds to a draw from the ergodic distribution of the Markov chain, and M is the number of draws per iteration, then

$$\sqrt{N}(\widehat{\xi}^{(s)} - \overline{\xi}) \xrightarrow{d} \mathcal{N}(0, \mathcal{V} + \mathcal{V}_M),$$

where the expressions of $\mathcal{V}$ and $\mathcal{V}_M$ are given in Appendix B.

In addition, if $\widehat{\xi}$ is a parameter draw and M tends to infinity, or alternatively if $\widehat{\xi}$ is computed as the average of $\widehat{\xi}^{(s)}$ over $\widetilde{S}$ iterations with $\widetilde{S}$ tending to infinity (as in our implementation), then

$$\sqrt{N}(\widehat{\xi} - \overline{\xi}) \xrightarrow{d} \mathcal{N}(0, \mathcal{V}),$$

where $\mathcal{V}$ is the asymptotic variance of the method-of-moments estimator based on the integrated moment restrictions (3.11) and (3.12).

Nonparametric consistency. In the asymptotic theory of the previous subsection, K_1, K_2 and L are held fixed as N tends to infinity. It might be more appealing to see the parametric specification based on series and splines as an approximation to the quantile functions, which becomes more accurate as the dimensions K_1, K_2 and L increase. Here, our aim is to provide conditions under which the estimator is consistent as N, K_1, K_2 and L tend to infinity.

To proceed, we consider the following assumption on the data-generating process, as in Belloni et al. (2011),

$$Y_{it} = W_{it}(\eta_i)'\overline{\theta}(U_{it}) + R_Y(X_{it}, \eta_i, U_{it}),$$

and, similarly,

$$\eta_i = Z_i'\overline{\delta}(V_i) + R_\eta(X_i, V_i),$$

where $\sup_{(x,e,u)} |R_Y(x, e, u)| = o(1)$ as K_1 tends to infinity, and $\sup_{(x,v)} |R_\eta(x, v)| = o(1)$ as K_2 tends to infinity.

Let $\xi(\tau) = (\theta(\tau)', \delta(\tau)')'$ be a $(K_1 + K_2) \times 1$ vector for all $\tau \in (0, 1)$, and let $\xi : (0, 1) \to \mathbb{R}^{K_1+K_2}$ be the associated function. Let us consider the estimator $\widehat{\xi} = (\widehat{\theta}', \widehat{\delta}')'$ based on the integrated moment restrictions (3.11) and (3.12). This analysis as $M \to \infty$ thus ignores the impact of small-M simulation error. Note that $\widehat{\xi}$ is a function defined on the unit interval.

In Appendix B, we provide and discuss conditions that guarantee that $\widehat{\xi}$ is uniformly consistent for $\overline{\xi} = (\overline{\theta}', \overline{\delta}')'$; that is,

$$\sup_{\tau \in (0,1)} \|\widehat{\xi}(\tau) - \overline{\xi}(\tau)\| = o_p(1), \tag{3.14}$$

where $\| \cdot \|$ denotes the Euclidean norm on $\mathbb{R}^{K_1+K_2}$.

Some of the conditions for consistency given in Appendix B are non-primitive. In particular, an identification condition is required, which is related to Assumption 2.3, though it differs from it because our estimator is based on a set of moment conditions rather than the likelihood. More generally, models with latent distributions, such as the nonlinear panel data models we analyse in this paper, are subject to ill-posedness, making a complete characterization of asymptotic distributions challenging.[10] A practical possibility, for which we do not yet have a formal justification, is to use empirical counterparts of the fixed-(K_1, K_2, L) asymptotic formulae derived in the previous subsection or, alternatively the bootstrap, to conduct inference. A related question is that of the practical choice of K_1, K_2 and L. In this paper, we do not characterize the asymptotic distribution of our estimator as N, K_1, K_2 and L tend to infinity, and we leave these important questions to future work.

4. DYNAMIC MODELS

In this section, we extend the method to dynamic models with dependence on lagged outcomes or predetermined covariates.

4.1. Models, examples and identification

In a dynamic extension of the static model (2.1), we specify the conditional quantile function of Y_{it} given $Y_{i,t-1}$, X_{it} and η_i as

$$Y_{it} = Q_Y(Y_{i,t-1}, X_{it}, \eta_i, U_{it}), \quad i = 1, \ldots, N, \quad t = 2, \ldots, T. \tag{4.1}$$

A simple extension is obtained by replacing $Y_{i,t-1}$ by a vector containing various lags of the outcome variable. As in the static case, Q_Y could depend on t.

Linear versions of (4.1) are widely used in applications, including in the study of individual earnings, firm-level investment and cross-country growth, or in the numerous applications of panel VAR models. In these applications, interactions between heterogeneity and dynamics are often of great interest. A recent example is the analysis of institutions and economic growth in Acemoglu et al. (2015).

The assumptions we impose in model (4.1), and the modelling of unobserved heterogeneity, both depend on the nature of the covariates process. We consider two cases in turn: strictly exogenous and predetermined covariates.

[10] In particular, the class of models we consider nests nonparametric deconvolution models with repeated measurements; see, e.g. Kotlarski (1967), Horowitz and Markatou (1996), Delaigle et al. (2008), Bonhomme and Robin (2010). In such settings, quantiles are generally not root-N estimable (Hall and Lahiri, 2008).

Autoregressive models. In the case where covariates are strictly exogenous, with some abuse of notation we suppose that Assumption 2.1 holds with $(Y_{i,t-1}, X'_{iT})'$ instead of X_{it} and $(Y_{i1}, X'_{i1}, \ldots, X'_{iT})'$ instead of X_i. Note that the latter contains both strictly exogenous covariates and first-period outcomes. Individual effects can be written without loss of generality as

$$\eta_i = Q_\eta(Y_{i1}, X_i, V_i), \qquad i = 1, \ldots, N, \tag{4.2}$$

and we suppose that Assumption 2.2 holds with $(Y_{i1}, X'_i)'$ instead of X_i.

Predetermined covariates. In dynamic models with predetermined regressors, current values of U_{it} can affect future values of covariates X_{is}, $s > t$. Given the presence of latent variables in our nonlinear set-up, a model for the feedback process is needed. That is, we need to specify the conditional distribution of X_{it} given $(Y_i^{t-1}, X_i^{t-1}, \eta_i)$, where $Y_i^{t-1} = (Y_{i,t-1}, \ldots, Y_{i1})'$ and $X_i^{t-1} = (X'_{i,t-1}, \ldots, X'_{i1})'$. We use additional quantile specifications for this purpose. In the case where X_{it} is scalar, and under a conditional first-order Markov assumption for (Y_{it}, X_{it}), $t = 1, \ldots, T$, given η_i, we specify, without further loss of generality:

$$X_{it} = Q_X(Y_{i,t-1}, X_{i,t-1}, \eta_i, A_{it}), \quad i = 1, \ldots, N, \quad t = 2, \ldots, T. \tag{4.3}$$

We suppose that Assumptions 2.1 and 2.2 hold, with $(Y_{i,t-1}, X'_{it})'$ instead of X_{it} and $(Y_{i1}, X'_{i1})'$ instead of X_i, and

$$\eta_i = Q_\eta(Y_{i1}, X_{i1}, V_i), \quad i = 1, \ldots, N. \tag{4.4}$$

We then complete the model with the following assumption on the feedback process.

ASSUMPTION 4.1. (PREDETERMINED COVARIATES) *(a) A_{it} follows a standard uniform distribution, independent of $(Y_{i,t-1}, X_{i,t-1}, \eta_i)$; (b) $\tau \mapsto Q_X(y, x, \eta, \tau)$ is strictly increasing on $(0, 1)$, for almost all (y, x, η) in the support of $(Y_{i,t-1}, X_{i,t-1}, \eta_i)$; (c) for all $t \neq s$, A_{it} is independent of A_{is}.*

Model (4.3) can be extended to multidimensional predetermined covariates using a triangular approach in the spirit of the one introduced in Example 2.3. For example, with two-dimensional $X_{it} = (X_{1it}, X_{2it})'$,

$$\begin{aligned} X_{1it} &= Q_{X_1}(Y_{i,t-1}, X_{1i,t-1}, X_{2i,t-1}, \eta_i, A_{1it}), \\ X_{2it} &= Q_{X_2}(Y_{i,t-1}, X_{1it}, X_{1i,t-1}, X_{2i,t-1}, \eta_i, A_{2it}), \end{aligned} \tag{4.5}$$

where η_i can be scalar or multidimensional as in Example 2.3.

EXAMPLE 4.1. (PANEL QUANTILE AUTOREGRESSION) A dynamic counterpart to Example 2.2 is the following linear quantile regression model:

$$Y_{it} = \rho(U_{it})Y_{i,t-1} + X'_{it}\beta(U_{it}) + \eta_i\gamma(U_{it}). \tag{4.6}$$

Model (4.6) differs from the more general model studied in Galvao (2011):

$$Y_{it} = \rho(U_{it})Y_{i,t-1} + X'_{it}\beta(U_{it}) + \eta_i(U_{it}). \tag{4.7}$$

Similarly as in (2.5), and in contrast with the models introduced in this paper, the presence of the functional heterogeneity term $\eta_i(\tau)$ makes fixed-T consistent estimation problematic in (4.7).

An extension of (4.6) is

$$Y_{it} = h(Y_{i,t-1})'\rho(U_{it}) + X'_{it}\beta(U_{it}) + \eta_i\gamma(U_{it}), \quad t = 2, \ldots, T, \tag{4.8}$$

where h is a univariate function. For example, when $h(y) = |y|$ model (4.8) is a panel data version of the CAViaR model of Engle and Manganelli (2004). Other choices will lead to panel counterparts of various dynamic quantile models; see, e.g. Gouriéroux and Jasiak (2008). The approach developed in this paper allows for more general, nonlinear series specifications of dynamic quantile functions in a panel data context.

EXAMPLE 4.2. (QUANTILE AUTOREGRESSION WITH PREDETERMINED COVARIATES) Extending Example 4.1 to allow for a scalar predetermined covariate X_{it}, we may augment (4.6) with the following linear quantile specification for X_{it}:

$$X_{it} = \mu(A_{it})Y_{i,t-1} + \xi_1(A_{it})X_{i,t-1} + \xi_0(A_{it}) + \zeta(A_{it})\eta_i.$$

This specification can be extended to allow for multidimensional predetermined regressors, as in (4.5).

Identification. In dynamic models, nonparametric identification requires $T \geq 4$. Under Assumption 2.1, U_{it} is independent of X_{is} for all s and uniformly distributed, and independent of U_{is} for all $s \neq t$. So, taking $T = 4$, we have

$$f_{Y_1,Y_2,Y_3,Y_4|X}(y_1, y_2, y_3, y_4 \mid x) = \int f_{Y_2|Y_1,\eta,X}(y_2 \mid y_1, \eta, x) f_{Y_3|Y_2,\eta,X}(y_3 \mid y_2, \eta, x) \times f_{Y_4|Y_3,\eta,X}(y_4 \mid y_3, \eta, x) f_{\eta,Y_1|X}(\eta, y_1 \mid x) d\eta, \tag{4.9}$$

where we have used that Y_{i4} is conditionally independent of (Y_{i2}, Y_{i1}) given (Y_{i3}, X_i, η_i), and that Y_{i3} is conditionally independent of Y_{i1} given (Y_{i2}, X_i, η_i).

An extension of the theorem of Hu and Schennach (2008), along the lines of Hu and Shum (2012), then shows nonparametric identification of all conditional densities $f_{Y_2|Y_1,\eta,X}$, $f_{Y_3|Y_2,\eta,X}$, $f_{Y_4|Y_3,\eta,X}$ and $f_{\eta,Y_1|X}$, in the autoregressive model, under suitable assumptions.[11]

Lastly, autoregressive models with predetermined covariates can be shown to be nonparametrically identified using similar arguments, provided the feedback process is first-order Markov.

4.2. *Estimation in dynamic models*

The estimation algorithm of Section 3 can be directly modified to deal with autoregressive models with strictly exogenous covariates. Consider a linear specification of the quantile functions (4.1) and (4.2), possibly based on series. Then, the stochastic EM algorithm essentially takes the same form as in the static case, except for the posterior density of the individual effects, which is now computed as

$$f(\eta \mid y, x; \xi) = \frac{\prod_{t=2}^{T} f_{Y_t|Y_{t-1},X_t,\eta}(y_t \mid y_{t-1}, x_t, \eta; \xi_A) f_{\eta|Y_1,X}(\eta \mid y_1, x; \xi_B)}{\int \prod_{t=2}^{T} f_{Y_t|Y_{t-1},X_t,\eta}(y_t \mid y_{t-1}, x_t, \widetilde{\eta}; \xi_A) f_{\eta|Y_1,X}(\widetilde{\eta} \mid y_1, x; \xi_B) d\widetilde{\eta}}. \tag{4.10}$$

[11] In the dynamic model (4.8), it follows from the analysis of Hu and Shum (2012) that one can rely on (2.12) as in the static case, provided the averages across τ values of the coefficients of exogenous regressors and lagged outcome are identified based on $E[Y_{it} - Y_{i,t-1} \mid Y_i^{t-2}, X_i] = E[h(Y_{i,t-1}) - h(Y_{i,t-2}) \mid Y_i^{t-2}, X_i]' \int_0^1 \rho(\tau)d\tau + (X_{it} - X_{i,t-1})' \int_0^1 \beta(\tau)d\tau$.

General predetermined regressors. In models with predetermined covariates, the critical difference is in the nature of the posterior density of the individual effects. Letting $W_{it} = (Y_{it}, X_{it}')'$ and $W_i^t = (W_{i1}', \ldots, W_{it}')'$, we have

$$\begin{aligned} f(\eta \mid y, x; \xi) &= \frac{f_{W_2,\ldots,W_T}(w_2, \ldots, w_T \mid w_1, \eta) f_{\eta|W_1}(\eta \mid w_1)}{\int f_{W_2,\ldots,W_T}(w_2, \ldots, w_T \mid w_1, \eta) f_{\eta|W_1}(\eta \mid w_1) d\eta} \\ &= \frac{f_{\eta|W_1}(\eta \mid w_1; \xi_B) \prod_{t=2}^T f_{Y_t|Y_{t-1},X_t,\eta}(y_t \mid y_{t-1}, x_t, \eta; \xi_A) f_{X_t|W^{t-1},\eta}(x_t \mid w^{t-1}, \eta; \xi_C)}{\int f_{\eta|W_1}(\widetilde{\eta} \mid w_1; \xi_B) \prod_{t=2}^T f_{Y_t|Y_{t-1},X_t,\eta}(y_t \mid y_{t-1}, x_t, \widetilde{\eta}; \xi_A) f_{X_t|W^{t-1},\eta}(x_t \mid w^{t-1}, \widetilde{\eta}; \xi_C) d\widetilde{\eta}}, \end{aligned}$$

where now $\xi = (\xi_A', \xi_B', \xi_C')'$ includes additional parameters that correspond to the model of the feedback process from past values of Y_{it} and X_{it} to future values of X_{is}, for $s > t$.

Under predeterminedness, the quantile model only specifies the partial likelihood:

$$\prod_{t=2}^T f_{Y_t|Y_{t-1},X_t,\eta}(y_t \mid y_{t-1}, x_t, \eta; \xi_A).$$

However, the posterior density of the individual effects also depends on the feedback process,

$$f_{X_t|W^{t-1},\eta}(x_t \mid w^{t-1}, \eta; \xi_C),$$

in addition to the density of individual effects. Note that the feedback process could depend on an additional vector of individual effects different from η_i.

In line with our approach, we also specify the quantile function of covariates in (4.3) using linear (series) quantile regression models. Specifically, letting X_{pit}, $p = 1, \ldots, P$, denote the various components of X_{it}, we specify the following triangular, recursive system that extends Example 4.2 to multidimensional predetermined covariates:

$$\begin{aligned} X_{1it} &= W_{1it}(\eta_i)\mu_1(A_{1it}), \\ \ldots & \quad \ldots \quad \ldots \\ X_{Pit} &= W_{Pit}(\eta_i)\mu_P(A_{Pit}). \end{aligned} \tag{4.11}$$

Here, $A_{1it}, \ldots, A_{Pit}$ follow independent standard uniform distributions, independent of all other random variables in the model, $W_{1it}(\eta_i)$ contains functions of $(Y_{i,t-1}, X_{i,t-1}, \eta_i)$ and $W_{pit}(\eta_i)$ contains functions of $(X_{1it}, \ldots X_{p-1,it}, Y_{i,t-1}, X_{i,t-1}, \eta_i)$ for $p > 1$. The parameter vector ξ_C includes all $\mu_p(\tau_\ell)$, for $p = 1, \ldots, P$ and $\ell = 1, \ldots, L$.

Thus, the model with predetermined regressors has three layers of quantile regressions: the outcome model (4.1) specified as a linear quantile regression, the model of the feedback process (4.11) and the model of individual effects (4.4), which here depends on first-period outcomes and covariates. The estimation algorithm is similar to the one for static models, with minor differences in both steps.[12]

[12] In addition, in Appendix C, we describe how to allow for autocorrelated errors in model (2.1) and (2.2).

5. QUANTILE MARGINAL EFFECTS

In nonlinear panel data models, it is often of interest to compute the effect of marginal changes in covariates on the entire distribution of outcome variables. As an example, let us consider the following average quantile marginal effect (QME) for continuous X_{it}:

$$M(\tau) = E\left[\frac{\partial Q_Y(X_{it}, \eta_i, \tau)}{\partial x}\right].$$

Here, $\partial Q_Y/\partial x$ denotes the vector of partial derivatives of Q_Y with respect to its first $\dim(X_{it})$ arguments.

In the quantile regression model of Example 2.2, individual QMEs are equal to $\partial Q_Y (X_{it}, \eta_i, \tau)/\partial x = \beta(\tau)$ and $M(\tau) = \beta(\tau)$. In Example 2.3, individual QME are heterogeneous, equal to $\beta(\tau) + \gamma_2(\tau)\eta_{i2}$ and $M(\tau) = \beta(\tau) + \gamma_2(\tau)E[\eta_{i2}]$. Series specifications of the quantile function as in (1.1) can allow for rich heterogeneity in individual QMEs.

Dynamic models. QMEs are also of interest in dynamic models. One can define a short-run average QME as

$$M_t(\tau) = E\left[\frac{\partial Q_Y(Y_{i,t-1}, X_{it}, \eta_i, \tau)}{\partial x}\right].$$

Moreover, when considering marginal changes in the lagged outcome $Y_{i,t-1}$, the average QME, $E[\partial Q_Y(Y_{i,t-1}, X_{it}, \eta_i, \tau)/\partial y]$, can be interpreted as a nonlinear measure of state dependence. In that case, $\partial Q_Y/\partial y$ denotes the derivative of Q_Y with respect to its first argument.

Dynamic models also provide the opportunity to document dynamic QMEs, such as the following one-period-ahead average QME:

$$M_{t+1/t}(\tau_1, \tau_2) = E\left[\frac{\partial Q_Y(Q_Y(Y_{i,t-1}, X_{it}, \eta_i, \tau_1), X_{i,t+1}, \eta_i, \tau_2)}{\partial y} \times \frac{\partial Q_Y(Y_{i,t-1}, X_{it}, \eta_i, \tau_1)}{\partial x}\right].$$

Here, $M_{t+1/t}(\tau_1, \tau_2)$ measures the average effect of a marginal change in X_{it} when η_i is kept fixed, and the innovations in periods t and $t + 1$ have rank τ_1 and τ_2, respectively.

Panel quantile treatment effects. When the covariate of interest is binary, as in our empirical application in Section 6, one can define panel data versions of quantile treatment effects. To see this, let D_{it} be the binary covariate of interest, and let X_{it} include all other time-varying covariates. Consider the static model (2.1), the argument extending directly to dynamic models. Potential outcomes are defined as

$$Y_{it}(d) = Q_Y(d, X_{it}, \eta_i, U_{it}), \qquad d \in \{0, 1\}.$$

Under Assumption 2.1, $(Y_{it}(0), Y_{it}(1))$ is conditionally independent of D_{it} given (X_i, η_i). This amounts to assuming selection on observables and unobservables, when unobserved effects η_i are identified from the panel dimension.

The average conditional quantile treatment effect is then defined as

$$E[Q_Y(1, X_{it}, \eta_i, \tau) - Q_Y(0, X_{it}, \eta_i, \tau)].$$

In the linear quantile regression model of Example 2.2, this is simply the coefficient of the vector $\beta(\tau)$ corresponding to D_{it}. In fact, the distribution of treatment effects is identified for this model, under the conditions spelled out in Section 2. The key assumption is rank invariance of U_{it} given X_i and η_i.

It is also possible to define unconditional quantile treatment effects, as

$$F^{-1}_{Y_{it}(1)}(\tau) - F^{-1}_{Y_{it}(0)}(\tau), \tag{5.1}$$

where the CDFs $F_{Y_{it}(0)}$ and $F_{Y_{it}(1)}$ are given by[13]

$$F_{Y_{it}(d)}(y) = E\Big[\int_0^1 \mathbf{1}\{Q_Y(d, X_{it}, \eta_i, \tau) \le y\} d\tau\Big], \qquad d \in \{0, 1\}. \tag{5.2}$$

All these quantities can readily be estimated using our panel quantile estimator.

6. EMPIRICAL APPLICATION

In this section, we present an empirical illustration to the link between mothers' smoking during pregnancy and birthweight. We start by discussing how we implement the estimation algorithm in practice.

6.1. Implementation

Piecewise-linear splines. We use piecewise-linear splines as an approximating model. Although other spline families could be used instead, computing the implied likelihood functions would then require inverting quantile functions numerically. In contrast, for linear splines, for all $\ell = 1, \ldots, L-1$, we have

$$\theta(\tau) = \theta(\tau_\ell) + \frac{\tau - \tau_\ell}{\tau_{\ell+1} - \tau_\ell}\big(\theta(\tau_{\ell+1}) - \theta(\tau_\ell)\big), \qquad \tau_\ell < \tau \le \tau_{\ell+1},$$

$$\delta(\tau) = \delta(\tau_\ell) + \frac{\tau - \tau_\ell}{\tau_{\ell+1} - \tau_\ell}\big(\delta(\tau_{\ell+1}) - \delta(\tau_\ell)\big), \qquad \tau_\ell < \tau \le \tau_{\ell+1},$$

and the implied approximating period-t density of outcomes and the implied approximating density of individual effects take the following simple closed-form expressions,

$$f_{Y_t|X_t,\eta}(y_t \mid x_t, \eta; \xi_A) = \frac{\tau_{\ell+1} - \tau_\ell}{w_t(\eta)'(\theta(\tau_{\ell+1}) - \theta(\tau_\ell))} \quad \text{if } w_t(\eta)'\theta(\tau_\ell) < y_t \le w_t(\eta)'\theta(\tau_{\ell+1}), \tag{6.1}$$

$$f_{\eta|X}(\eta \mid x; \xi_B) = \frac{\tau_{\ell+1} - \tau_\ell}{z'(\delta(\tau_{\ell+1}) - \delta(\tau_\ell))} \quad \text{if } z'\delta(\tau_\ell) < \eta \le z'\delta(\tau_{\ell+1}), \tag{6.2}$$

augmented with a specification in the tail intervals $(0, \tau_1)$ and $(\tau_L, 1)$.

Tail intervals. In order to model quantile functions in the intervals $(0, \tau_1)$ and $(\tau_L, 1)$ one could assume, following Wei and Carroll (2009), that $\theta(\cdot)$ and $\delta(\cdot)$ are constant on these intervals, so the implied distribution functions have mass points at the two ends of the support. In

[13] Note that unconditional quantile treatment effects cannot be directly estimated as in Firpo (2007) in this context, due to the presence of the unobserved η_i and the lack of fixed-T identification for fixed-effects binary choice models.

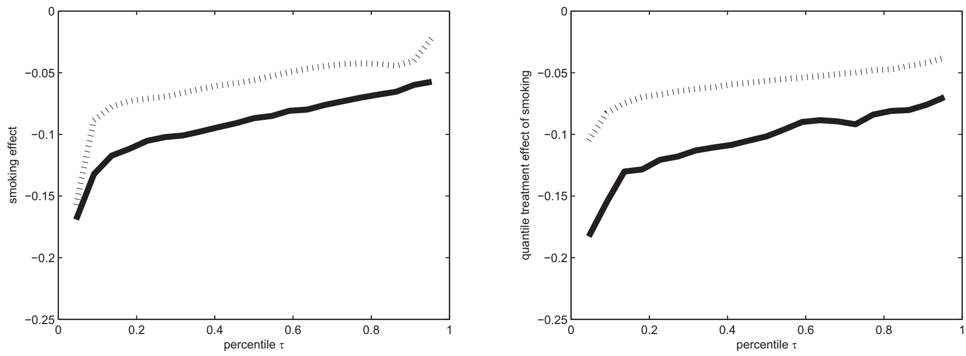

Figure 1. Quantile effects of smoking during pregnancy on log-birthweight.

Appendix D, we outline a different, exponential-based modelling of the extreme intervals, motivated by the desire to avoid the fact that the support of the likelihood function depends on the parameter value. We use this method in the empirical application.

6.2. Application: smoking and birthweight

Here, we revisit the effect of maternal inputs of children's birth outcomes. Specifically, we study the effect of smoking during pregnancy on children's birthweights. Abrevaya (2006) uses a mother fixed-effects approach to address endogeneity of smoking. In this paper, we use quantile regression with mother-specific effects to allow for both unobserved heterogeneity and nonlinearities in the relationship between smoking and weight at birth. As a complement, in Appendix E, we report the results of a Monte Carlo simulation broadly calibrated to this application, in order to assess the performance of our estimator in finite samples.

We focus on a balanced subsample from the US natality data used in Abrevaya (2006), which comprises 12,360 women with three children each. Our outcome is the log-birthweight. The main covariate is a binary smoking indicator. Age of the mother and gender of the child are used as additional controls.

An ordinary least-squares (OLS) regression yields a significantly negative point estimate of the smoking coefficient: -0.095. The fixed-effects estimate is also negative, but it is twice as small: -0.050, significant. This suggests a negative endogeneity bias in OLS, and is consistent with the results in Abrevaya (2006).

In the left panel of Figure 1 (data from Abrevaya, 2006), the solid line shows the smoking coefficient estimated from pooled quantile regressions, on a fine grid of τ values. According to these estimates, the effect of smoking is more negative at lower quantiles of birthweights. The dashed line in the left panel of Figure 1 shows the quantile estimate of the smoking effect. We use a linear quantile regression specification as in Example 2.2, augmented with a parametric exponential model in the tail intervals. The covariates are smoking status, age and gender, with an intercept. We use individual-specific averages of these variables as covariates in the specification for η_i. Estimates are computed using $L = 21$ knots. The stochastic EM algorithm is run for 100

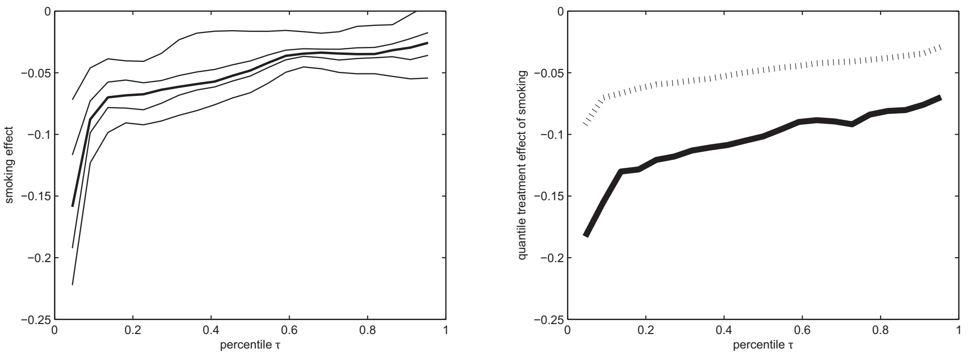

Figure 2. Quantile effects of smoking during pregnancy on log-birthweight.

iterations, with 500 random-walk Metropolis–Hastings draws within each iteration.[14] Parameter estimates are computed as averages of the 50 last iterations of the algorithm.[15]

In the left panel of Figure 1, we see that the smoking effect becomes less negative when correcting for time-invariant endogeneity through the introduction of mother-specific fixed effects. At the same time, the effect is still sizable, and it remains increasing along the distribution.

As another exercise, in the right panel of Figure 1 we compute the unconditional quantile treatment effect of smoking as the difference in log-birthweights between a sample of smoking women, and a sample of non-smoking women, keeping all other characteristics (i.e. observed X_i and unobserved η_i) constant; see (5.1) and (5.2). We report differences in quantiles of simulated potential outcomes obtained using the method of Machado and Mata (2005). This exercise illustrates the usefulness of specifying and estimating a complete semiparametric model of the joint distribution of outcomes and unobservables, in order to compute counterfactual distributions that take into account the presence of unobserved heterogeneity. In this panel, the solid line shows the empirical difference between unconditional quantiles, while the dashed line shows the quantile treatment effect that accounts for both observables and unobservables.

The results in the right panel of Figure 1 are broadly similar to the results reported in the left panel. An interesting finding is that in this case the endogeneity bias (i.e. the difference between the dashed and solid lines) is slightly larger, and that it tends to decrease as one moves from lower to higher quantiles of birthweight.

Finally, in Figure 2 (data from Abrevaya, 2006), we report the results of an interacted quantile model, as in (1.1) and (1.2), where the specification allows for all first-order interactions between covariates (i.e. smoking status, age and gender) and the unobserved mother-specific effect. In this model, the quantile effect of smoking is mother-specific. In the left panel, lines represent the percentiles 0.05, 0.25, 0.50, 0.75 and 0.95 of the heterogeneous smoking effect

[14] The variance of the random-walk proposal is set to achieve an acceptance rate of $\approx 30\%$.

[15] For θ, starting parameter values are taken based on ordinary quantile regressions of log-birthweight on smoking status, age and gender, with an intercept, setting the coefficient of η_i in the outcome equation to one. For δ, we set all initial quantile parameters to $\{0.1, 0.2, \dots, 2.1\}$. The initial values for the exponential parameters in the tails are all set to 20. We experimented with other starting values for the model's parameters (e.g. we initialized δ based on quantile regressions of individual-specific means $\overline{Y}_i$ on $\overline{X}_i$) and found no qualitative differences compared to the results we report.

across mothers, at various percentiles τ. In the right panel, the solid line is the raw quantile treatment effect of smoking, and the dashed line is the quantile treatment effect estimate based on panel quantile regression with interactions. The results in the right panel show the unconditional quantile treatment effect of smoking, and are similar to the results obtained for a simple linear specification (see the right panel of Figure 1). However, in the left panel of Figure 2, we see substantial mother-specific heterogeneity in the conditional quantile treatment effect of smoking, as for some mothers, smoking appears particularly detrimental to children's birthweight, whereas for other mothers, the smoking effect, while consistently negative, is much smaller. This evidence is in line with the results of a linear random coefficients model reported in Arellano and Bonhomme (2012).

7. CONCLUSION

Quantile methods are flexible tools to model nonlinear panel data relationships. In this paper, quantile regression is used to model the dependence between outcomes, covariates and individual heterogeneity, and between individual effects and exogenous regressors or initial conditions. Quantile specifications also allow modelling feedback processes in models with predetermined covariates. The empirical application illustrates the benefits of having a flexible approach to allow for heterogeneity and nonlinearity within the same model in a panel data context.

Our approach leads to fixed-T identification of complete models. The estimation algorithm exploits the computational advantages of linear quantile regression, within an iterative scheme that allows us to deal with the presence of unobserved individual effects. Beyond static or dynamic quantile regression models with single or multiple individual effects, our approach naturally extends to series specifications, thus allowing for rich interactions between covariates and heterogeneity at various points of the distribution.

Our quantile-based modelling of the distribution of individual effects could also be of interest in other models. For example, one could consider semiparametric likelihood panel data models, where the conditional likelihood of the outcome Y_i given X_i and η_i depends on a finite-dimensional parameter vector α, and the conditional distribution of η_i given X_i is left unrestricted. The approach of this paper is easily adapted to this case, and delivers a semiparametric likelihood of the form,

$$f_{Y|X}(y|x;\alpha,\delta(\cdot)) = \int f_{Y|X,\eta}(y|x,\eta;\alpha) f_{\eta|X}(\eta|x;\delta(\cdot))d\eta,$$

where $\delta(\cdot)$ is a process of quantile coefficients.

Our framework also naturally extends to models with time-varying unobservables, such as

$$Y_{it} = Q_Y(X_{it}, \eta_{it}, U_{it}),$$

$$\eta_{it} = Q_\eta(\eta_{i,t-1}, V_{it}),$$

where U_{it} and V_{it} are i.i.d. and uniformly distributed. Arellano et al. (2015) use a quantile-based approach to document nonlinear relationships between earnings shocks to households and their lifetime profiles of earnings and consumption. This application illustrates the potential of our estimation approach in dynamic settings.

A relevant issue for empirical practice is measurement error. Our approach can be extended to allow covariates to be measured with error, as the analysis in Wei and Carroll (2009) illustrates. When a validation sample is available, our algorithm can also be modified to allow for measurement error in outcome variables. In both cases, true variables are treated similarly as latent individual effects in the above analysis, and they are repeatedly drawn from their posterior densities in each iteration of the algorithm.

Lastly, this paper leaves a number of important questions unanswered. Statistical inference in the nonparametric problem, where the complexity of the approximating model increases together with the sample size, is one of them. Providing primitive conditions for identification, and devising efficient computational routines, are other important questions for future work.

ACKNOWLEDGEMENTS

A presentation based on this paper was given at the *Econometrics Journal* RES Conference 2013. A previous version of the paper was given the title ‘Random-effects quantile regression’. We thank the editor and two anonymous referees for comments and suggestions. We also thank Richard Blundell, Xiaohong Chen, Andrew Chesher, Bryan Graham, Roger Koenker, Arthur Lewbel, Konrad Menzel, Lars Nesheim, Yin Wei, and seminar participants at various places for comments. Raffaele Saggio provided excellent research assistance. Support from the European Research Council (ERC) grant agreement no. 263107 is gratefully acknowledged. All errors are our own.

REFERENCES

Abrevaya, J. (2006). Estimating the effect of smoking on birth outcomes using a matched panel data approach. *Journal of Applied Econometrics 21*, 489–519.

Abrevaya, J. and C. M. Dahl (2008). The effects of birth inputs on birthweight. *Journal of Business and Economic Statistics 26*, 379–97.

Acemoglu, D., S. Naidu, P. Restrepo and J. Robinson (2015). Democracy does cause growth. Working paper, MIT.

Andrews, D. (2011). Examples of L^2-complete and boundedly-complete distributions. Cowles Foundation Discussion Paper 1801, Yale University.

Arcidiacono, P. and J. B. Jones (2003). Finite mixture distributions, sequential likelihood and the EM algorithm. *Econometrica 71*, 933–46.

Arcidiacono, P. and R. Miller (2011). Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica 79*, 1823–67.

Arellano, M. and S. Bonhomme (2011). Nonlinear panel data analysis. *Annual Review of Economics 3*, 395–424.

Arellano, M. and S. Bonhomme (2012). Identifying distributional characteristics in random coefficients panel data models. *Review of Economic Studies 79*, 987–1020.

Arellano, M. and M. Weidner (2015). Instrumental variable quantile regressions in large panels with fixed effects. Working paper, CEMFI.

Arellano, M., R. Blundell and S. Bonhomme (2015). Earnings and consumption dynamics: a nonlinear panel data framework. CWP 53/15, Centre for Microdata Methods and Practice, Institute for Fiscal Studies and University College London.

Belloni, A., V. Chernozhukov and I. Fernández-Val (2011). Conditional quantile processes based on series or many regressors. CWP 19/11, Centre for Microdata Methods and Practice, Institute for Fiscal Studies and University College London.

Bonhomme, S. and J.-M. Robin (2009). Assessing the equalizing force of mobility using short panels: France, 1990–2000. *Review of Economic Studies 76*, 63–92.

Bonhomme, S. and J.-M. Robin (2010). Generalized nonparametric deconvolution with an application to earnings dynamics. *Review of Economic Studies 77*, 491–533.

Canay, I. A. (2011). A simple approach to quantile regression for panel data. *Econometrics Journal 14*, 368–86.

Canay, I. A., A. Santos and A. Shaikh (2013). On the testability of identification in some nonparametric models with endogeneity. *Econometrica 81*, 2535–59.

Chamberlain, G. (1984). Panel data. In Z. Griliches and M. D. Intriligator (Eds.), *Handbook of Econometrics*, *Volume 2*, 1248–318. Amsterdam: Elsevier.

Chen, X. (2007). Large sample sieve estimation of semi-nonparametric models. In J. J. Heckman and E. E. Leamer (Eds.), *Handbook of Econometrics*, *Volume 6*, *Part B*, 5559–632. Amsterdam: Elsevier.

Chernozhukov, V., I. Fernández-Val and A. Galichon (2010). Quantile and probability curves without crossing. *Econometrica 78*, 1093–125.

Chernozhukov, V., I. Fernández-Val, J. Hahn and W. K. Newey (2013). Average and quantile effects in nonseparable panel models. *Econometrica 81*, 535–80.

Chernozhukov, V., I. Fernández-Val, S. Hoderlein, H. Holzmann and W. K. Newey (2015). Nonparametric identification in panels using quantiles. *Journal of Econometrics 188*, 378–92.

Delaigle, A., P. Hall and A. Meister (2008). On deconvolution with repeated measurements. *Annals of Statistics 36*, 665–85.

Dempster, A. P., N. M. Laird and D. B. Rubin (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society*, Series B, *39*, 1–38.

D'Haultfoeuille, X. (2011). On the completeness condition for nonparametric instrumental problems. *Econometric Theory 27*, 460–71.

Elashoff, M. and L. Ryan (2004). An EM algorithm for estimating equations. *Journal of Computational and Graphical Statistics 13*, 48–65.

Engle, R. F. and S. Manganelli (2004). CAViaR: conditional autoregressive value at risk by regression quantiles. *Journal of Business and Economic Statistics 22*, 367–81.

Evdokimov, K. (2010). Identification and estimation of a nonparametric panel data model with unobserved heterogeneity. Working paper, Princeton University.

Firpo, S. (2007). Efficient semiparametric estimation of quantile treatment effects. *Econometrica 75*, 259–76.

Galvao, A. F. (2011). Quantile regression for dynamic panel data with fixed effects. *Journal of Econometrics 164*, 142–57.

Galvao, A. F., K. Kato and G. Montes-Rojas (2012). Asymptotics for quantile regression models with individual effects. *Journal of Econometrics 170*, 76–91.

Geraci, M. and M. Bottai (2007). Quantile regression for longitudinal data using the asymmetric Laplace distribution. *Biostatistics 8*, 140–54.

Gouriéroux, C. and J. Jasiak (2008). Dynamic quantile models. *Journal of Econometrics 147*, 198–205.

Graham, B., J. Hahn, A. Poirier and J. Powell (2015). Quantile regression with panel data. NBER Working Paper 21034.

Hall, P. and S. N. Lahiri (2008). Estimation of distributions, moments and quantiles in deconvolution problems. *Annals of Statistics 36*, 2110–34.

He, X. (1997). Quantile curves without crossing. *The American Statistician 51*, 186–92.

Horowitz, J. L. and M. Markatou (1996). Semiparametric Estimation of regression models for panel data. *Review of Economic Studies 63*, 145–68.

Hu, Y. (2015). Microeconomic models with latent variables: applications of measurement error models in empirical industrial organization and labor economics. CWP 03/15, Centre for Microdata Methods and Practice, Institute for Fiscal Studies and University College London.

Hu, Y. and S. M. Schennach (2008). Instrumental variable treatment of nonclassical measurement error models. *Econometrica 76*, 195–216.

Hu, Y. and J-L. Shiu (2012). Nonparametric identification using instrumental variables: sufficient conditions for completeness. CWP 25/11, Centre for Microdata Methods and Practice, Institute for Fiscal Studies and University College London.

Hu, Y. and M. Shum (2012). Nonparametric identification of dynamic models with unobserved state variables. *Journal of Econometrics 171*, 32–44.

Koenker, R. and G. J. Bassett (1978). Regression quantiles. *Econometrica 46*, 33–50.

Koenker, R. (2004). Quantile regression for longitudinal data. *Journal of Multivariate Analysis 91*, 74–89.

Koenker, R. (2005). *Quantile Regression*. Econometric Society Monograph Series. Cambridge: Cambridge University Press.

Kotlarski, I. (1967). On characterizing the gamma and normal distribution. *Pacific Journal of Mathematics 20*, 69–76.

Lamarche, C. (2010). Robust penalized quantile regression for panel data. *Journal of Econometrics 157*, 396–408.

Machado, J. A. F. and J. Mata (2005). Counterfactual decomposition of changes in wage distributions using quantile regression. *Journal of Applied Econometrics 20*, 445–65.

McFadden, D. L. and P. A. Ruud (1994). Estimation by simulation. *Review of Economics and Statistics 76*, 591–608.

Nielsen, S. F. (2000a). The stochastic EM algorithm: estimation and asymptotic results. *Bernouilli 6*, 457–89.

Nielsen, S. F. (2000b). On simulated EM algorithms. *Journal of Econometrics 96*, 267–92.

Rosen, A. (2012). Set identification via quantile restrictions in short panels. *Journal of Econometrics 166*, 127–37.

Schennach, S. (2014). Entropic latent variable integration via simulation. *Econometrica 82*, 345–85.

Wei, Y. and R. J. Carroll (2009). Quantile regression with measurement error. *Journal of the American Statistical Association 104*, 1129–43.

APPENDIX A: IDENTIFICATION – DISCUSSION OF ASSUMPTION 2.3

Assumption 2.3(a) requires that all densities under consideration be bounded. This imposes mild restrictions on the model's parameters. Assumption 2.3(b) requires that $f_{Y_3|\eta,X}$ be non-identical at different values of η. This assumption will be satisfied if, for some τ in small open neighbourhood $Q_{Y_3}(x,\eta_1,\tau) \neq Q_{Y_3}(x,\eta_2,\tau)$. In Example 2.2, Assumption 2.3(a) requires strict monotonicity of quantile functions – that is, $x'\nabla\beta(\tau) + \eta\nabla\gamma(\tau) \geq \underline{c} > 0$, where $\nabla\xi(\tau)$ denotes the first derivative of $\xi(\cdot)$ evaluated at τ – while Assumption 2.3(b) holds if $\gamma(\tau) \neq 0$ for τ in some open neighbourhood.

Assumption 2.3(c) imposes a centred measure of location on $f_{Y_2|\eta,X=x}$. In order to apply the identification theorem in Hu and Schennach (2008), it is not necessary that Γ_x be known. If, instead, Γ_x is a known function of the data distribution, their argument goes through. For example, in Example 2.2, one convenient normalization is obtained by noting that

$$E[Y_{it} \mid \eta_i, X_{it}] = X_{it}'\left[\int_0^1 \beta(\tau)d\tau\right] + \eta_i\left[\int_0^1 \gamma(\tau)d\tau\right] \equiv \widetilde{X}_{it}'\overline{\beta}_1 + \overline{\beta}_0 + \eta_i\overline{\gamma},$$

where $\overline{\beta}_0 = \int_0^1 \beta_0(\tau)d\tau$ corresponds to the coefficient of the constant in $X_{it} = (\widetilde{X}_{it}', 1)'$. Now, if $\widetilde{X}_{it}$ varies over time and a rank condition is satisfied, $\overline{\beta}_1$ is a known function of the data distribution, simply given by the within-group estimand. Thus, in this case, we can take

$$\Gamma_x(g) = \int yg(y)dy - \widetilde{x}_2'\overline{\beta}_1,$$

and note that the following normalization implies Assumption 2.3(c):

$$\overline{\beta}_0 = \int_0^1 \beta_0(\tau)d\tau = 0 \quad \text{and} \quad \overline{\gamma} = \int_0^1 \gamma(\tau)d\tau = 1.$$

In a fully nonparametric setting and arbitrary t, to ensure that Assumption 2.3(c) holds for some period (i.e. $t = 1$) we can proceed as follows. First, let us define

$$\widetilde{\eta}_i \equiv E(Y_{i1} \mid \eta_i, X_{i1}).$$

Then, in every period t, provided $\eta \mapsto E[Y_{i1}|\eta_i = \eta, X_{i1} = x_1]$ is invertible for almost all x_1, we have

$$Y_{it} = Q_Y(X_{it}, \eta_i, U_{it}) \equiv \widetilde{Q}_Y(X_{it}, X_{i1}, \widetilde{\eta}_i, U_{it}).$$

Estimating specifications of this form will deliver estimates of $\widetilde{Q}_Y$, from which the average marginal effects defined in Section 5 can be recovered as estimates of

$$M_t(\tau) = E\left[\frac{\partial Q_Y(X_{it}, \eta_i, \tau)}{\partial x_t}\right] = E\left[\frac{\partial \widetilde{Q}_Y(X_{it}, X_{i1}, \widetilde{\eta}_i, \tau)}{\partial x_t}\right],$$

where $\partial \widetilde{Q}_Y/\partial x_t$ denotes the vector of partial derivatives of $\widetilde{Q}_Y$ with respect to its first $\dim(X_{it})$ arguments.

Assumption 2.3(d) is an injectivity condition. The operator $L_{Y_2|\eta,x}$ is defined as $[L_{Y_2|\eta,x}\ h](y_2) = \int f_{Y_2|\eta,X}(y_2|\eta, x)h(\eta)d\eta$, for all bounded functions h. Here, $L_{Y_2|\eta,x}$ is injective if the only solution to $L_{Y_2|\eta,x}h = 0$ is $h = 0$. As pointed out by Hu and Schennach (2008), injectivity is closely related to completeness conditions commonly assumed in the literature on nonparametric instrumental variable estimation. Similarly as completeness, injectivity is a high-level condition; see, e.g. Canay et al. (2013) for results on the testability of completeness assumptions.

Several recent papers provide explicit conditions for completeness or injectivity in specific models. Andrews (2011) constructs classes of distributions that are L^2-complete and boundedly complete. D'Haultfoeuille (2011) provides primitive conditions for completeness in a linear model with homoscedastic errors. The results of Hu and Shiu (2012) apply to the location-scale quantile model of Example 2.1. In this case, conditions that guarantee that $L_{Y_2|\eta,x}$ is injective involve the tail properties of the conditional density of Y_{i2} given η_i (and X_i) and its characteristic function.[16] Providing primitive conditions for injectivity/completeness in more general models, such as the linear quantile regression model of Example 2.2, is an interesting question but exceeds the scope of this paper.

[16] See Lemma 4 in Hu and Shiu (2012).

APPENDIX B: ASYMPTOTIC PROPERTIES

B.1. Parametric inference

Here, we rely on Nielsen's work to characterize the asymptotic distribution of $\widehat{\xi}^{(s)}$ in our model, where the optimization step is not likelihood-based. To do so, let us rewrite the moment restrictions in a compact notation,

$$E[\Psi_i(\eta_i;\overline{\xi})] = 0,$$

where ξ (with true value $\overline{\xi}$) is a finite-dimensional parameter vector of the same dimension as Ψ. Equivalently, we have

$$E\left[\int \Psi_i(\eta;\overline{\xi}) f(\eta|W_i;\overline{\xi}) d\eta\right] = 0,$$

where $W_i = (Y_i, X_i')'$.

The stochastic EM algorithm for this problem works as follows, based on an i.i.d. sample $(W_1, \ldots, W_N)$. Iteratively, one draws $\widehat{\xi}^{(s+1)}$ given $\widehat{\xi}^{(s)}$ in two steps.

STEP 1. For $i = 1, \ldots, N$, draw $\eta_i^{(1,s)}, \ldots, \eta_i^{(M,s)}$ from the posterior distribution $f(\eta_i|W_i;\widehat{\xi}^{(s)})$.[17]
STEP 2. Solve for $\widehat{\xi}^{(s+1)}$ in

$$\sum_{i=1}^{N}\sum_{m=1}^{M} \Psi_i(\eta_i^{(m,s)};\widehat{\xi}^{(s+1)}) = 0.$$

This results in a Markov chain $(\widehat{\xi}^{(0)}, \widehat{\xi}^{(1)}, \ldots)$, which is ergodic under suitable conditions. Moreover, under conditions given in Nielsen (2000a), asymptotically as N tends to infinity, and for almost every W-sequence and conditional on W (hereafter, simply 'conditional on W') the process $\sqrt{N}(\widehat{\xi}^{(s)} - \widehat{\xi})$ converges to a Gaussian AR(1) process, where $\widehat{\xi}$ solves the integrated moment restrictions:

$$\sum_{i=1}^{N} \int \Psi_i(\eta;\widehat{\xi}) f(\eta|W_i;\widehat{\xi}) d\eta = 0. \tag{B.1}$$

In the rest of this section, we characterize the unconditional asymptotic distribution of $\sqrt{N}(\widehat{\xi}^{(s)} - \overline{\xi})$. The derivations in this section are heuristic, and throughout we assume sufficient regularity conditions to justify all the steps.[18]

Using a conditional quantile representation, we have

$$\eta_i^{(m,s)} = Q_{\eta|W}(W_i, V_i^{(m,s)};\widehat{\xi}^{(s)}),$$

where $V_i^{(m,s)}$ are standard uniform draws, independent of each other and independent of W_i.

Thus, we have

$$\sum_{i=1}^{N}\sum_{m=1}^{M} \Psi_i(Q_{\eta|W}(W_i, V_i^{(m,s)};\widehat{\xi}^{(s)});\widehat{\xi}^{(s+1)}) = 0.$$

[17] For simplicity, we consider the case where $\eta_i^{(1,s)}, \ldots, \eta_i^{(M,s)}$ are independent draws.

[18] Note that in our quantile model some of the moment restrictions involve derivatives of 'check' functions, which are not smooth. However, this is not central to the discussion that follows, as it does not affect the form of the asymptotic variance.

Expanding around $\widehat{\xi}$ conditional on W, and using the fact that $\widehat{\xi}$ tends to $\overline{\xi}$ as N tends to infinity, we obtain

$$A(\widehat{\xi}^{(s+1)} - \widehat{\xi}) + B(\widehat{\xi}^{(s)} - \widehat{\xi}) + \varepsilon^{(s)} = o_p(N^{-(1/2)}), \tag{B.2}$$

where

$$A \equiv \left.\frac{\partial}{\partial \xi'}\right|_{\overline{\xi}} E[\Psi_i(Q_{\eta|W}(W_i, V_i; \overline{\xi}); \xi)] = \left.\frac{\partial}{\partial \xi'}\right|_{\overline{\xi}} E[\Psi_i(\eta_i; \xi)],$$

$$B \equiv \left.\frac{\partial}{\partial \xi'}\right|_{\overline{\xi}} E[\Psi_i(Q_{\eta|W}(W_i, V_i; \xi); \overline{\xi})] = \left.\frac{\partial}{\partial \xi'}\right|_{\overline{\xi}} E[\int \Psi_i(\eta; \overline{\xi}) f(\eta|W_i; \xi) d\eta],$$

$$\varepsilon^{(s)} \equiv \frac{1}{NM} \sum_{i=1}^{N} \sum_{m=1}^{M} \Psi_i(Q_{\eta|W}(W_i, V_i^{(m,s)}; \overline{\xi}); \overline{\xi}).$$

Note that

$$A + B = \left.\frac{\partial}{\partial \xi'}\right|_{\overline{\xi}} E\left[\int \Psi_i(\eta; \xi) f(\eta|W_i; \xi) d\eta\right].$$

The identification condition for the method-of-moments problem thus requires $A + B < 0$, so $(-A)^{-1}B < I$. This implies that the Gaussian AR(1) limit of $\sqrt{N}(\widehat{\xi}^{(s)} - \widehat{\xi})$ conditional on W is stable. Thus, we have

$$\sqrt{N}(\widehat{\xi}^{(s)} - \widehat{\xi}) = \sum_{k=0}^{\infty} (-A^{-1}B)^k (-A^{-1}) \sqrt{N} \varepsilon^{(s-1-k)} + o_p(1). \tag{B.3}$$

Moreover, $\sqrt{N}\varepsilon^{(s)}$ are asymptotically i.i.d. normal with zero mean and variance Σ/M, where

$$\Sigma = E[\Psi_i(\eta_i; \overline{\xi}) \Psi_i(\eta_i; \overline{\xi})'].$$

Hence, conditional on W,

$$\sqrt{N}(\widehat{\xi}^{(s)} - \widehat{\xi}) \xrightarrow{d} \mathcal{N}(0, \mathcal{V}_M),$$

where

$$\mathcal{V}_M = \sum_{k=0}^{\infty} (-A^{-1}B)^k (-A^{-1}) \frac{\Sigma}{M} (-A^{-1})' ((-A^{-1}B)^k)'.$$

Note that $\mathcal{V}_M$ can be recovered from the following matrix equation

$$A^{-1} B \mathcal{V}_M B' (A^{-1})' = \mathcal{V}_M - A^{-1} \frac{\Sigma}{M} (A^{-1})',$$

which can be easily solved in vector form.

Finally, unconditionally, we have by asymptotic independence,

$$\sqrt{N}(\widehat{\xi}^{(s)} - \overline{\xi}) = \sqrt{N}(\widehat{\xi}^{(s)} - \widehat{\xi}) + \sqrt{N}(\widehat{\xi} - \overline{\xi}) \xrightarrow{d} \mathcal{N}(0, \mathcal{V} + \mathcal{V}_M),$$

where $\mathcal{V}$ is the asymptotic variance of $\sqrt{N}(\widehat{\xi} - \overline{\xi})$; that is,

$$\mathcal{V} = (A + B)^{-1} \Omega ((A + B)^{-1})',$$

where $\Omega = E[(\int \Psi_i(\eta; \overline{\xi}) f(\eta|W_i; \overline{\xi}) d\eta)(\int \Psi_i(\eta; \overline{\xi}) f(\eta|W_i; \overline{\xi}) d\eta)']$.

B.2. Nonparametric consistency

Let $\overline{\xi}(\tau) = (\overline{\theta}(\tau)', \overline{\delta}(\tau)')'$, and let $\varphi_i(\xi(\cdot), \tau)$ be the $(K_1 + K_2) \times 1$ moment vector that corresponds to the integrated moment restrictions (3.6) and (3.7). Let $\|\cdot\|$ denote the Euclidean norm on $\mathbb{R}^{K_1+K_2}$, and let $\|\xi(\cdot)\|_\infty = \sup_{\tau\in(0,1)} \|\xi(\tau)\|$ denote the associated uniform norm.

Let $K = K_1 + K_2$. Consider a space $\mathcal{H}_K$ of functions $\xi(\cdot)$, which contains differentiable functions whose first derivatives (component-wise) are bounded and Lipschitz continuous on (0, 1). Moreover, suppose there exists $\underline{c}$ such that, for all $\tau_1 < \tau_2$ and with probability one, $W_{it}(\eta_i)'(\theta(\tau_2) - \theta(\tau_1)) \geq \underline{c}(\tau_2 - \tau_1)$ and $Z_i'(\delta(\tau_2) - \delta(\tau_1)) \geq \underline{c}(\tau_2 - \tau_1)$. This last requirement imposes strict monotonicity of the conditional quantile functions. These assumptions guarantee that the implied likelihood functions and posterior density of the individual effects are bounded from above and away from zero. Finally, all functions $\xi(\cdot) \in \mathcal{H}_K$ are assumed to satisfy a location restriction as in Assumption 2.3(c).

To every function $\xi(\cdot) \in \mathcal{H}_K$, we associate an interpolating spline $\pi_L\xi(\cdot)$ in a space $\mathcal{H}_{KL}$. We use piecewise-linear splines on $(\tau_1, \ldots, \tau_L)$, as in Section 6.1. For simplicity, we consider the case where quantile functions are constant on the tail intervals, so $\pi_L\xi(\tau) = \xi(\tau_1)$ for $\tau \in (0, \tau_1)$, and $\pi_L\xi(\tau) = \xi(\tau_L)$ for $\tau \in (\tau_L, 1)$. Moreover, the minimum and maximum of $L|\tau_{\ell+1} - \tau_\ell|$ are assumed to be asymptotically bounded away from zero and infinity. We also assume that L tends to infinity sufficiently fast relative to K so that $\|\xi(\cdot) - \pi_L\xi(\cdot)\|_\infty = o(1)$ for all $\xi(\cdot) \in \mathcal{H}_K$.

Let us define

$$Q_K(\xi(\cdot)) = \int_0^1 \|E[\varphi_i(\xi(\cdot), \tau)]\|^2 d\tau,$$

and

$$\widehat{Q}_{KL}(\xi(\cdot)) = \frac{1}{L}\sum_{\ell=1}^{L} \left\| \frac{1}{N}\sum_{i=1}^{N} \varphi_i(\pi_L\xi(\cdot), \tau_\ell) \right\|^2.$$

The estimator $\widehat{\xi}(\cdot)$ minimizes $\widehat{Q}_{KL}$ on $\mathcal{H}_{KL}$.

Consistency follows from the following high-level assumptions, which we briefly and informally discuss below.

ASSUMPTION B.1. (IDENTIFICATION; UNIFORM CONVERGENCE) *(a) For all $\epsilon > 0$ there is a $c > 0$ such that, for all K_1, K_2, L,*

$$\inf_{\xi(\cdot)\in\mathcal{H}_K,\, \|\xi(\cdot)-\overline{\xi}(\cdot)\|_\infty > \epsilon} Q_K(\xi(\cdot)) > Q_K(\overline{\xi}(\cdot)) + c;$$

(b) as N, K_1, K_2, L tend to infinity,

$$\sup_{\xi(\cdot)\in\mathcal{H}_K} |\widehat{Q}_{KL}(\xi(\cdot)) - Q_K(\xi(\cdot))| = o_p(1).$$

PROPOSITION B.1. (NONPARAMETRIC CONSISTENCY) *Under Assumption B.1, $\widehat{\xi}(\cdot)$ is uniformly consistent for $\overline{\xi}(\cdot)$ in the sense that (3.14) holds.*

Proof: Let $\widetilde{\xi}(\cdot) \in \mathcal{H}_K$ such that $\widehat{\xi}(\cdot) = \pi_L\widetilde{\xi}(\cdot)$. We have $\|\widetilde{\xi}(\cdot) - \widehat{\xi}(\cdot)\|_\infty = \|\widetilde{\xi}(\cdot) - \pi_L\widetilde{\xi}(\cdot)\|_\infty = o_p(1)$.

By definition of $\widehat{\xi}$, we have $\widehat{Q}_{KL}(\widehat{\xi}(\cdot)) \leq \widehat{Q}_{KL}(\overline{\xi}(\cdot))$. Hence, by Assumption B.1(b), and as $\widehat{Q}_{KL}(\widehat{\xi}(\cdot)) = \widehat{Q}_{KL}(\widetilde{\xi}(\cdot))$:

$$Q_K(\widetilde{\xi}(\cdot)) \leq Q_K(\overline{\xi}(\cdot)) + o_p(1).$$

Let $\epsilon > 0$. By Assumption B.1(a), it thus follows that $\|\widetilde{\xi}(\cdot) - \overline{\xi}(\cdot)\|_\infty \leq \epsilon$ with probability approaching one.

Hence, $\|\widehat{\xi}(\cdot) - \overline{\xi}(\cdot)\|_\infty \leq \|\widehat{\xi}(\cdot) - \widetilde{\xi}(\cdot)\|_\infty + \|\widetilde{\xi}(\cdot) - \overline{\xi}(\cdot)\|_\infty = o_p(1)$. This shows (3.14). □

Discussion of Assumption B.1(a). To provide intuition on the identification condition in Assumption B.1(a), consider the case where the posterior density $f(\eta|Y_i, X_i)$ is known. Consider the last K_2 elements of φ_i, the argument for the first K_1 elements being similar. Showing Assumption B.1(a) requires bounding the following quantity from below:

$$\Delta \equiv \int_0^1 \|E[Z_i(\tau - F(Z_i'\delta(\tau)|Y_i, X_i))]\|^2 - \|E[Z_i(\tau - F(Z_i'\overline{\delta}(\tau)|Y_i, X_i))]\|^2 d\tau.$$

Expanding around $\overline{\delta}(\tau)$ yields

$$\begin{aligned} E[Z_i(\tau - F(Z_i'\delta(\tau)|Y_i, X_i))] = {} & E[Z_i(\tau - F(Z_i'\overline{\delta}(\tau)|Y_i, X_i))] \\ & - E[Z_i Z_i' f(A_i(\tau;\delta)|Y_i, X_i)](\delta(\tau) - \overline{\delta}(\tau)), \end{aligned}$$

where $A_i(\tau;\delta)$ lies between $Z_i'\delta(\tau)$ and $Z_i'\overline{\delta}(\tau)$. Now, $E[Z_i(\tau - F(Z_i'\overline{\delta}(\tau)|Y_i, X_i))] = o(1)$, provided the remainder R_η tends to zero sufficiently fast as K_2 increases. Moreover, if $f(\eta|Y_i, X_i)$ is bounded away from zero as well as from above, and if the eigenvalues of the Gram matrix $E[Z_i Z_i']$ are bounded away from zero as well as from above, then there exists a constant $\mu > 0$ such that, for all τ:

$$\|E[Z_i Z_i' f(A_i(\tau;\delta)|Y_i, X_i)](\delta(\tau) - \overline{\delta}(\tau))\|^2 \geq \mu \|\delta(\tau) - \overline{\delta}(\tau)\|^2.$$

Finally, suppose $\|\delta(\cdot) - \overline{\delta}(\cdot)\|_\infty > \epsilon$. Then, by continuity of $\delta(\cdot) - \overline{\delta}(\cdot)$, there exists a non-empty interval (τ_1, τ_2) such that $\|\delta(\tau) - \overline{\delta}(\tau)\| > \epsilon$ for $\tau \in (\tau_1, \tau_2)$. Hence, $\Delta > \mu\epsilon^2|\tau_2 - \tau_1| + o(1)$.

In the panel quantile models considered in this paper, $f(\eta|Y_i, X_i; \xi(\cdot))$ depends on the unknown function $\xi(\cdot) = (\theta(\cdot)', \delta(\cdot)')'$. As we pointed out in Section 2.3, identification then depends on high-level conditions such as operator injectivity. Here, we do not provide primitive conditions for Assumption B.1(a) to hold in this case.

Discussion of Assumption B.1(b). The uniform convergence condition in Assumption B.1(b) will hold if the following conditions are satisfied:

$$A \equiv \sup_{\xi(\cdot)\in\mathcal{H}_K} \left| \frac{1}{L}\sum_{\ell=1}^{L} \left\| \frac{1}{N}\sum_{i=1}^{N} \varphi_i(\pi_L\xi(\cdot), \tau_\ell) \right\|^2 - \frac{1}{L}\sum_{\ell=1}^{L} \|E[\varphi_i(\pi_L\xi(\cdot), \tau_\ell)]\|^2 \right| = o_p(1),$$

$$B \equiv \sup_{\xi(\cdot)\in\mathcal{H}_K} \left| \frac{1}{L}\sum_{\ell=1}^{L} \|E[\varphi_i(\pi_L\xi(\cdot), \tau_\ell)]\|^2 - \frac{1}{L}\sum_{\ell=1}^{L} \|E[\varphi_i(\xi(\cdot), \tau_\ell)]\|^2 \right| = o(1),$$

$$C \equiv \sup_{\xi(\cdot)\in\mathcal{H}_K} \left| \frac{1}{L}\sum_{\ell=1}^{L} \|E[\varphi_i(\xi(\cdot), \tau_\ell)]\|^2 - \int_0^1 \|E[\varphi_i(\xi(\cdot), \tau)]\|^2 d\tau \right| = o(1).$$

The A quantity involves the difference between the empirical and population objective functions of the approximating parametric model. In the second term in B, the posterior density of individual effects depends on the entire function $\xi(\cdot)$, as opposed to its spline approximation $\pi_L\xi(\cdot)$. Lastly, the second term in C involves an integral on the unit interval, which needs to be compared to an average on the grid of τ_ℓ.

A, B, C can be bounded if it can first be established that there exist constants $C_1 > 0, C_2 > 0, \nu > 0$ such that, for all $\xi_1(\cdot), \xi_2(\cdot)$ in $\mathcal{H}_{KL}$ and τ_1, τ_2 in $(0, 1)$:[19]

$$\|\varphi_i(\xi_2(\cdot), \tau_2) - \varphi_i(\xi_1(\cdot), \tau_1)\| \leq C_1\sqrt{K}\|\xi_2(\cdot) - \xi_1(\cdot)\|_\infty^\nu + C_2\sqrt{K}|\tau_2 - \tau_1|. \tag{B.4}$$

The $\pi_L\xi(\cdot)$ belong to a compact KL-dimensional space. Given (B.4), it can be shown that $A = o_p(1)$, provided L tends to infinity sufficiently fast relative to K and KL/N tends to zero. The latter condition arises as $\pi_L\xi(\cdot)$ is finite-dimensional, with dimension KL. Wei and Carroll (2009) establish this result formally for a related model, in a case where K does not increase with the sample size.

Next, provided (B.4) can be extended to hold for any $\xi_1(\cdot)$ and $\xi_2(\cdot)$ in $\mathcal{H}_K$, and using that $\|\xi(\cdot) - \pi_L\xi(\cdot)\|_\infty = o(1)$, we find that $B = o(1)$ as long as L tends to infinity sufficiently fast relative to K.

Lastly, again using (B.4) but now for $\xi_1(\cdot) = \xi_2(\cdot)$, we obtain $C = o(1)$, again provided L tends to infinity sufficiently fast relative to K.

APPENDIX C: EXTENSION – AUTOCORRELATED DISTURBANCES

To allow for autocorrelated errors in model (2.1) and (2.2), we replace Assumption 2.1(c) by the following.

ASSUMPTION C.1. (AUTOCORRELATED ERRORS) *$(U_{i1}, \ldots, U_{iT})$ is distributed as a copula $C(u_1, \ldots, u_T)$, independent of (X_i, η_i).*

Nonparametric identification of the model (including the copula) can be shown under Markovian assumptions, as in the autoregressive model of Section 4. For estimation, we let the copula depend on a finite-dimensional parameter ϕ, which we estimate along with all quantile parameters. The iterative estimation algorithm is then easily modified by adding an update in Step 2 (the M-step):

$$\widehat{\phi}^{(s+1)} = \underset{\phi}{\operatorname{argmax}} \sum_{i=1}^{N}\sum_{m=1}^{M} \ln\Big[c(F(Y_{i1}|X_{i1}, \eta_i^{(m)}; \widehat{\xi}_A^{(s+1)}), \ldots, F(Y_{it}|X_{it}, \eta_i^{(m)}; \widehat{\xi}_A^{(s+1)}); \phi)\Big]. \tag{C.1}$$

Here, $c(u_1, \ldots, u_T) \equiv \partial^T C(u_1, \ldots, u_T)/\partial u_1 \ldots \partial u_T$ is the copula density and, for any y_t such that $w_t(\eta)'\theta(\tau_\ell) < y_t \leq w_t(\eta)'\theta(\tau_{\ell+1})$,

$$F(y_t|x_t, \eta; \xi_A) = \tau_\ell + (\tau_{\ell+1} - \tau_\ell)\frac{y_t - w_t(\eta)'\theta(\tau_\ell)}{w_t(\eta)'(\theta(\tau_{\ell+1}) - \theta(\tau_\ell))},$$

augmented with a specification outside the interval $(w_t(\eta)'\theta(\tau_1), w_t(\eta)'\theta(\tau_L))$. Here, F is a shorthand for $F_{Y_t|X_t,\eta}$.

The posterior density is then given by

$$f(\eta|y, x; \xi, \phi) = \frac{\prod_{t=1}^T f_{Y_t|X_t,\eta}(y_t \mid x_t, \eta; \xi_A) c[F(y_1|x_1, \eta; \xi_A), \ldots, F(y_T|x_T, \eta; \xi_A); \phi] f(\eta \mid x; \xi_B)}{\int \prod_{t=1}^T f_{Y_t|X_t,\eta}(y_t \mid x_t, \widetilde{\eta}; \xi_A) c[F(y_1|x_1, \widetilde{\eta}; \xi_A), \ldots, F(y_T|x_T, \widetilde{\eta}; \xi_A); \phi] f(\widetilde{\eta} \mid x; \xi_B) d\widetilde{\eta}}.$$

Lastly, note that the approach outlined here does not seem to easily generalize to allow for autocorrelated disturbances in autoregressive models (i.e. for ARMA-type quantile regression models).

[19] Consider the first K_1 elements of φ_i (the last K_2 elements having a similar structure):

$$\int \sum_{t=1}^{T} W_{it}(\eta)\psi_\tau(Y_{it} - W_{it}(\eta)'\theta(\tau)) f(\eta|Y_i, X_i; \pi_L\xi(\cdot)) d\eta.$$

A possibility to establish (B.4) could be to assume that $\eta \mapsto W_{it}(\eta)'\theta(\tau)$ is invertible almost surely (such a condition requires that the conditional quantile function of outcomes be monotonic in η_i), and that its inverse is Lipschitz continuous in $\theta(\tau)$, and then to use the expression of $f(\eta|Y_i, X_i; \pi_L\xi(\cdot))$, which involves the piecewise-linear expressions (6.1) and (6.2).

APPENDIX D: EXPONENTIAL MODELLING OF THE TAILS

For implementation, we use the following modelling for the splines in the extreme intervals indexed by $\lambda_1 > 0$ and $\lambda_L > 0$,

$$\theta(\tau) = \theta(\tau_1) + \frac{\ln(\tau/\tau_1)}{\lambda_1}\iota_c, \quad \tau \leq \tau_1,$$

$$\theta(\tau) = \theta(\tau_L) - \frac{\ln((1-\tau)/(1-\tau_L))}{\lambda_L}\iota_c, \quad \tau > \tau_L,$$

where ι_c is a vector of zeros, with a one at the position of the constant term in $\theta(\tau)$. We adopt a similar specification for $\delta(\tau)$, with parameters $\lambda_1^\eta > 0$ and $\lambda_L^\eta > 0$. Modelling the constant terms in $\theta(\tau)$ and $\delta(\tau)$, as we do, avoids the inconvenient that the support of the likelihood function depends on the parameter value. Moreover, our specification boils down to the Laplace model of Geraci and Bottai (2007) when $L = 1$, $\lambda_1 = 1 - \tau_1$ and $\lambda_L = \tau_L$.

The implied approximating period-t outcome density is then

$$\begin{aligned} f_{Y_t|X_t,\eta}(y_t \mid x_t, \eta; \xi_A) = &\sum_{\ell=1}^{L-1} \frac{\tau_{\ell+1} - \tau_\ell}{w_t(\eta)'(\theta(\tau_{\ell+1}) - \theta(\tau_\ell))} \mathbf{1}\{w_t(\eta)'\theta(\tau_\ell) < y_t \leq w_t(\eta)'\theta(\tau_{\ell+1})\} \\ &+ \tau_1\lambda_1 e^{\lambda_1(y_t - w_t(\eta)'\theta(\tau_1))}\mathbf{1}\{y_t \leq w_t(\eta)'\theta(\tau_1)\} \\ &+ (1-\tau_L)\lambda_L e^{-\lambda_L(y_t - w_t(\eta)'\theta(\tau_L))}\mathbf{1}\{y_t > w_t(\eta)'\theta(\tau_L)\}. \end{aligned}$$

Similarly, the approximating density of individual effects is

$$\begin{aligned} f_{\eta|X}(\eta \mid x; \xi_B) = &\sum_{\ell=1}^{L-1} \frac{\tau_{\ell+1} - \tau_\ell}{z'(\delta(\tau_{\ell+1}) - \delta(\tau_\ell))} \mathbf{1}\{z'\delta(\tau_\ell) < \eta \leq z'\delta(\tau_{\ell+1})\} \\ &+ \tau_1\lambda_1^\eta e^{\lambda_1^\eta(\eta - z'\delta(\tau_1))}\mathbf{1}\{\eta \leq z'\delta(\tau_1)\} \\ &+ (1-\tau_L)\lambda_L^\eta e^{-\lambda_L^\eta(\eta - z'\delta(\tau_L))}\mathbf{1}\{\eta > z'\delta(\tau_L)\}. \end{aligned}$$

Update rules for exponential parameters. We adopt a likelihood approach to update the parameters $\lambda_1, \lambda_L, \lambda_1^\eta, \lambda_L^\eta$. This yields the following moment restrictions:

$$\overline{\lambda}_1^\eta = \frac{-E[\int \mathbf{1}\{\eta \leq Z_i'\overline{\delta}(\tau_1)\} f(\eta|Y_i, X_i; \overline{\xi}) d\eta]}{E[\int (\eta - Z_i'\overline{\delta}(\tau_1))\mathbf{1}\{\eta \leq Z_i'\overline{\delta}(\tau_1)\} f(\eta|Y_i, X_i; \overline{\xi}) d\eta]},$$

and

$$\overline{\lambda}_L^\eta = \frac{E[\int \mathbf{1}\{\eta > Z_i'\overline{\delta}(\tau_L)\} f(\eta|Y_i, X_i; \overline{\xi}) d\eta]}{E[\int (\eta - Z_i'\overline{\delta}(\tau_L))\mathbf{1}\{\eta > Z_i'\overline{\delta}(\tau_L)\} f(\eta|Y_i, X_i; \overline{\xi}) d\eta]},$$

with similar equations for λ_1, λ_L.

Hence, the update rules in Step 2 of the algorithm (the M-step) are

$$\widehat{\lambda}_1^{\eta,(s+1)} = \frac{-\sum_{i=1}^N \sum_{m=1}^M \mathbf{1}\{\eta_i^{(m)} \leq Z_i'\widehat{\delta}(\tau_1)^{(s)}\}}{\sum_{i=1}^N \sum_{m=1}^M (\eta_i^{(m)} - Z_i'\widehat{\delta}(\tau_1)^{(s)})\mathbf{1}\{\eta_i^{(m)} \leq Z_i'\widehat{\delta}(\tau_1)^{(s)}\}},$$

and

$$\widehat{\lambda}_L^{\eta,(s+1)} = \frac{\sum_{i=1}^N \sum_{m=1}^M \mathbf{1}\{\eta_i^{(m)} > Z_i'\widehat{\delta}(\tau_L)^{(s)}\}}{\sum_{i=1}^N \sum_{m=1}^M (\eta_i^{(m)} - Z_i'\widehat{\delta}(\tau_L)^{(s)})\mathbf{1}\{\eta_i^{(m)} > Z_i'\widehat{\delta}(\tau_L)^{(s)}\}}.$$

APPENDIX E: MONTE CARLO ILLUSTRATION

The data-generating process is

$$Y_{it} = \beta_0(U_{it}) + \beta_1(U_{it})X_{1it} + \beta_2(U_{it})X_{2it} + \beta_3(U_{it})X_{3it} + \gamma(U_{it})\eta_i,$$

and

$$\eta_i = \delta_0(V_i) + \delta_1(V_i)\overline{X}_{1i} + \delta_2(V_i)\overline{X}_{2i} + \delta_3(V_i)\overline{X}_{3i}.$$

The covariates X_{i1} (smoking status), X_{i2} (age) and X_{i3} (gender) are taken from the data set of the empirical illustration. $T = 3$, and we extract a random subsample of 1000 mothers from the original data set. The true parameter values correspond to estimates on the full sample. Parameters β, γ and δ are taken to be piecewise-linear on an equidistant grid with $L = 11$ knots, with exponential specifications in the tails of intercept coefficients. For computation, we use the same method as in the application to select starting values, and we let the EM algorithm run for 100 iterations, with 100 random-walk Metropolis–Hastings draws within each iteration, reporting averages over the last 50 iterations. We report the results of 500 simulations in Figure E.1, which shows the data-generating process with $L = 11$ knots, $N = 1000$ and $T = 3$. The x-axis shows τ percentiles. True parameter values are shown by solid lines, Monte Carlo means

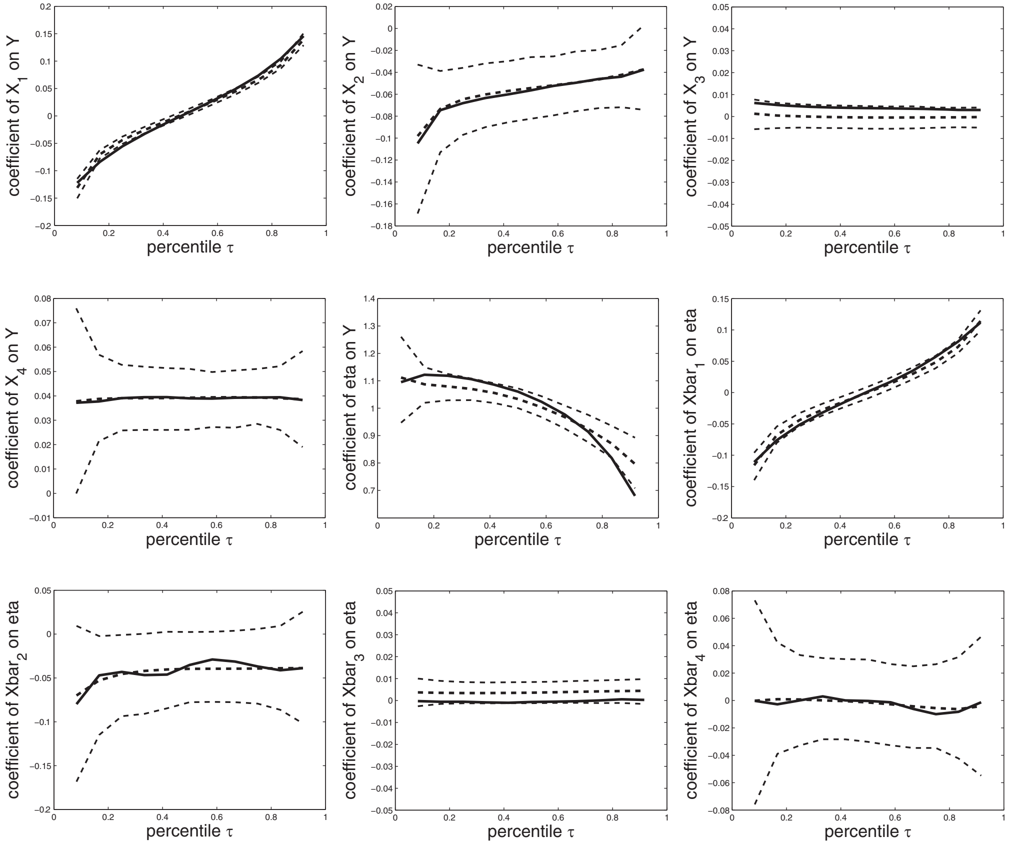

Figure E.1. Monte Carlo results.

are shown by thick dashed lines and 95% pointwise confidence intervals are shown by thin dashed lines. For example, the confidence intervals of the quantile parameters $\beta_1(\tau)$ corresponding to the effect of smoking are quite tight, even though the sample size is about 12 times smaller than the one of the application. Overall, the results provide encouraging evidence on the finite sample performance of the estimator.

SUPPORTING INFORMATION

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Replication Files