QUANTILE REGRESSION FOR PANEL DATA AND FACTOR MODELS*

CARLOS LAMARCHE†

Summary: For nearly 25 years, advances in panel data and quantile regression were developed almost completely in parallel, with no intersection until the work by Koenker (2004). The early theoretical work in statistics and economics raised more questions than answers, but it encouraged the development of several promising new approaches and research that offered a better understanding of the challenges and possibilities in the intersection of the literatures. Panel data quantile regression allows the estimation of effects that are heterogeneous throughout the conditional distribution of the response variable, while controlling for individual and time specific confounders. This type of heterogeneous effect is not be well summarized by the average effect. For instance, the relationship between number of students in a class and average educational achievement has been extensively investigated, but research also shows that class size affects differently low-achieving and high-achieving students. The recent advances in panel data include several methods and algorithms that created opportunities for more informative and robust empirical analysis in models with subject heterogeneity and factor structure.

Keywords: Quantile Regression; Panel Data; Subject Heterogeneity; Factor Models; Fixed Effects; Shrinkage; Endogenous Variables; Inference; Microeconometrics; Labor Economics
JEL codes: C21, C23, C31, C33, C38

Introduction

In 1978, Roger Koenker and Gilbert Bassett published a seminal paper titled "Regression Quantiles", creating a literature that goes beyond mean regression. Its predecessors, the median and sample quantiles, have a long history and were considered by Boscovich, Laplace, and Galton (Koenker, 2017). Quantile regression is an estimation technique to estimate covariate effects at different quantiles (or percentiles, such as the median or 50th percentile)

*March 2, 2021. I would like to thank Matt Harding and Rob Hartley for comments on a previous version of this paper.

†Carlos Lamarche: Department of Economics, University of Kentucky, 223G Gatton College of Business & Economics, Lexington, KY 40506. Email: clamarche@uky.edu.

of the conditional distribution of the dependent variable. Quantile regression can be used to estimate these effects at any quantile of the distribution, offering a more informative regression picture than familiar least squares methods. After the introduction of quantile regression, Steve Portnoy, James Powell, Roger Koenker, Xuming He, Victor Chernozhukov, and others generalized quantile regression, facilitating the application of the techniques in several scientific areas ranging from medicine and ecology to finance and economics. Part of the modern quantile regression literature includes models and methods for panel data.

Panel data quantile regression, broadly defined, can be viewed as a collection of statistical methods designed to estimate a conditional quantile regression model with latent individual heterogeneity. Panel data consists of multiple observations on the response variable, y_{it}, and a p-dimensional vector of regressors, $\boldsymbol{x}_{it}$, for subject $i = 1, 2, \dots, n$ over time $t = 1, 2, \dots, T$. With the exception of Lipsitz et al. (1997), the literature on panel data quantile regression was scarce until the work by Koenker (2004). Motivated by the lack of suitable transformations to deal with latent individual heterogeneity, the first class of models considered joint estimation of the slope parameter and individual effects. The simplest version of a panel quantile regression model involved finding

$$\underset{\boldsymbol{\alpha},\boldsymbol{\beta} \in \mathcal{A} \times \mathcal{B}}{\operatorname{argmin}} \mathbb{E}(\rho_\tau(y - \boldsymbol{x}'\boldsymbol{\beta} - \boldsymbol{\alpha})), \tag{1}$$

where the quantile regression function $\rho_\tau = u(\tau - I(u < 0))$, and $\tau \in (0, 1)$. The parameter of interest is $\boldsymbol{\beta} \in \mathcal{B} \subseteq R^p$, and $\boldsymbol{\alpha} \in \mathcal{A} \subseteq R^n$ is a vector of individual parameters. The formulation of the problem is analogous to the classical fixed effects estimator, which is widely adopted in practice. Similar problems were investigated by Abrevaya and Dahl (2008), Lamarche (2010), Canay (2011), Galvao (2011), and Rosen (2012).

Panel quantile regression with fixed effects is highly parametrized and its asymptotic properties were not well established until the work by Kato et al. (2012). For asymptotic results, they found that T has to grow at most polynomially in n in order to remove an additional term found in the Bahadur representation of the estimator for a model with individual parameters. Importantly, this condition is interpreted as a warning device to practitioners when asymptotic results are likely to provide a good approximation in practice. More recently, Galvao et al. (2020) establish consistency and asymptotic normality when T grows faster than n at rates that are, at best, similar to standard non-linear panel data models. Given these important caveats, it is not surprising that the most recent literature offers a variety of alternative approaches to estimate a panel quantile model, including Chernozhukov et al.

(2013), Chernozhukov et al. (2015), Arellano and Bonhomme (2016), Harding and Lamarche (2017), Graham et al. (2018), and Machado and Santos Silva (2019).

Recently, the emphasis has been on introducing more flexible forms of latent heterogeneity. This follows the literature on linear panel data that has made significant progress on the estimation of models with interactive effects and factor structure (Pesaran, 2006; Bai, 2009; Moon and Weidner, 2015, 2017). In large T panels, it is possible to accommodate (1) as

$$\underset{\boldsymbol{\gamma},\boldsymbol{f},\boldsymbol{\beta}\in\mathcal{G}\times\mathcal{F}\times\mathcal{B}}{\operatorname{argmin}} \mathbb{E}(\rho_\tau(y-\boldsymbol{x}'\boldsymbol{\beta}-\boldsymbol{\gamma}'\boldsymbol{f})), \tag{2}$$

where $\boldsymbol{\gamma}$ is a r-dimensional vector of individual factor loadings, $\boldsymbol{f}$ is a vector of unobserved time-varying factors, and r is the number of latent factors. It is straightforward to adapt this problem to include individual effects by noting that $\boldsymbol{\gamma}'\boldsymbol{f} = \gamma_1 + \sum_{k=2}^{r} \gamma_k f_k$ if f_1 is normalized to one. This generalization is important in empirical applications as it moves away from the classical additive separable structure on unobserved heterogeneity that allows one to control for individual and time heterogeneity, $\gamma + f$, for the possibility that $\boldsymbol{f}$ affects the response variable depending on subject-specific heterogeneity $\boldsymbol{\gamma}$. Problems similar to (2) were investigated in Harding and Lamarche (2014), Ando and Bai (2020), and Harding et al. (2020). Generalizations of standard factor models to allow for quantile specific loadings and factors are investigated in Ma et al. (2020), Chen et al. (2020), and Battistin et al. (2020).

Applications of panel quantile regression are now found in several fields of study within Economics and Finance. Specifically, there has been considerable work in labor economics. Using a panel of private sector workers, Dahl et al. (2013) investigate the impact of different bargaining systems on the distribution of wages. Hartley and Lamarche (2018) study behavioral-induced welfare participation at the highest quantiles of earnings using data on a reform experiment. Powell (2020) proposes an estimation approach and studies labor supply changes to transitory income shocks using data from economic stimulus payments. There is also a growing empirical literature in other applied micro areas such as environmental economics (see, e.g., Fell and Kaffine, 2018), empirical auctions (De Silva et al., 2009, 2017), and empirical finance (Fuchs and Gehring, 2017; Corradin and Maddaloni, 2020). As discussed in the section titled An Empirical Illustration, the use of quantile regression in panel data has been also noted in the evaluation of policies. We briefly illustrate the application of panel methods by estimating the effect of the Civil Rights Act of 1964 on earnings using data from Chay and Powell (2001). Lastly, considering a factor model, we illustrate the application of new methods to the evaluation of an energy savings program in Ireland using data from Harding and Lamarche (2019a).

The simplest formulation of a panel quantile model follows the tradition in panel econometrics of modeling individual heterogeneity as a time-invariant latent variable (Hsiao, 2014; Baltagi, 2013). The basic model is,

$$y_{it} = \alpha_i + \boldsymbol{x}_{it}'\boldsymbol{\beta} + u_{it}, \tag{3}$$

where $y_{it} \in \mathbb{R}$ is the response variable for cross-sectional unit i at time t, $\boldsymbol{x}_{it}$ is a vector of regressors specific to cross-sectional unit i at time t, and u_{it} is an error term. The variable α_i is considered to be latent and possibly correlated with the independent variables. The endogeneity of the vector $\boldsymbol{x}_{it}$ motivates the use of within or first difference transformations to consistently estimate the parameter of interest, $\boldsymbol{\beta}$. These transformations are easy to implement and allow identification of $\boldsymbol{\beta}$. For instance, the moment condition for the first difference estimator is $\mathbb{E}((\boldsymbol{x}_{it} - \boldsymbol{x}_{it-1})(u_{it} - u_{it-1})) = \mathbf{0}$, under independence between the error term and the regressors. This naturally implies that $\boldsymbol{\beta}$ can be consistently estimated by a linear regression model using variables in differences.

On the other hand, these classical transformations that deal with individual heterogeneity fail when the model is nonlinear, and therefore, the early work in panel quantiles adopted a different strategy (Koenker, 2004; Lamarche, 2006, 2010; Galvao, 2011). The initial approach was to estimate a model with a large number of individual intercepts, directly implied by equation (3):

$$Q_{y_{it}}(\tau|\boldsymbol{x}_{it}) = \alpha_i(\tau) + \boldsymbol{x}_{it}'\boldsymbol{\beta}(\tau), \tag{4}$$

where the conditional quantile function is defined as $Q_{y_{it}}(\tau|\boldsymbol{x}_{it}) := \inf\{y : P(y_{it} \leq y|\boldsymbol{x}_{it}) \geq \tau\}$. In this model, α_i is treated as a parameter. In the case that α_i is treated as a random variable, one can write $Q_{y_{it}}(\tau|\boldsymbol{x}_{it}, \alpha_i(\tau))$, but the fixed effects estimator and its large sample results are identical whether the model conditions on individual effects or not.

A direct implication of equation (4) is that $y_{it} = Q_{y_{it}}(\tau|\boldsymbol{x}_{it}) + u_{it}(\tau)$, where $u_{it}(\tau)$ is a quantile-specific random variable with zero τ-th conditional quantile: $P(u_{it}(\tau) \leq 0|\boldsymbol{x}_{it}) = \tau$. It then follows that,

$$\mathbb{E}\{\tau - I(y_{it} \leq Q_{y_{it}}(\tau|\boldsymbol{x}_{it}))|\boldsymbol{x}_{it}\} = 0, \tag{5}$$

which is a non-linear moment condition and is different than the moment condition corresponding to a response variable of a model in differences. For instance, consider a Gaussian model without an overall intercept and with covariate effects that do not vary by quantiles, $Q_{y_{it}}(\tau|\boldsymbol{x}_{it}) = \alpha_i(\tau) + \boldsymbol{x}_{it}'\boldsymbol{\beta}$. It is clear that $u_{it}(\tau) = u_{it} + (\alpha_i - \alpha_i(\tau)) = u_{it} - \Phi(\tau)^{-1}$, while the quantile error term in a model in differences is $\Delta u_{it}(\tau) = u_{it} - u_{it-1}$. Thus, the moment conditions corresponding to the model in levels and the model in differences are identical only at the median quantile.

To sum up, the lack of feasible transformations to deal with individual heterogeneity created a series of theoretical and computational challenges. Importantly, estimating a large number of parameters in a non-linear model can lead to estimates that are severely biased, as originally pointed out by Neyman and Scott (1948), and Hahn and Newey (2004) in the case of panel models. At the same time, estimating a large number of parameters can be computationally demanding in situations when n is large. However, fixed effects approaches (Koenker, 2004; Lamarche, 2010; Galvao, 2011) take advantage of modern developments on Sparse matrix algebra and solve a relatively simple linear programming problem that performs well in large panel applications.

Panel Data Quantile Regression

Fixed Effects. Koenker (2004) proposes to jointly estimate $\boldsymbol{\beta}$ and individual effects α_i for $1 \leq i \leq n$, which are assumed to be location-shift effects. This idea is implemented by simultaneously estimating m quantiles as follows:

$$\underset{\boldsymbol{\beta},\boldsymbol{\alpha}\in\mathcal{A}\times\mathcal{B}}{\operatorname{argmin}} \sum_{j=1}^{m}\sum_{i=1}^{n}\sum_{t=1}^{T} \omega_j \rho_{\tau_j}(y_{it} - \boldsymbol{x}_{it}'\boldsymbol{\beta}(\tau_j) - \alpha_i), \tag{6}$$

where ω_j is a relative weight given to the j-th quantile. The weight controls the influence of the m quantiles on the estimation of the parameters of the model. As Koenker (2004) explains, selecting the vector of weights, $\boldsymbol{\omega} = (\omega_1, \ldots, \omega_m)'$, is similar to the choice of discretely weighted L-statistics. To the best of our knowledge, selecting an optimal vector of weights has not been studied in the literature, and standard ad-hoc choices are 0.25, 0.5, and 0.25 for the quartiles, or alternatively, $\omega_j = 1/m$ for $1 \leq j \leq m$.

The fixed effects estimator (6) estimates the following conditional quantile model:

$$Q_{y_{it}}(\tau_j|\boldsymbol{x}_{it}) = \boldsymbol{x}_{it}'\boldsymbol{\beta}(\tau_j) + \alpha_i. \tag{7}$$

The parameter $\boldsymbol{\beta}(\tau_j)$ models how covariates influence the location, scale, and shape of the conditional distribution of the response. For instance, if the model has one regressor and the error term is independently and identically distributed (i.i.d.) as F, then the quantile functions $Q_{y_{it}}(\tau_j|x_{it})$ are parallel lines with parameter $\boldsymbol{\beta}(\tau_j) = (\beta_0(\tau_j), \beta_1)' = (\beta_0 + F_u(\tau_j)^{-1}, \beta_1)'$.

Moreover, it is important to note that the individual effect does not represent a distributional shift (Koenker, 2004). The individual effect is a location shift parameter on the conditional quantiles of the response, implying that the conditional distribution for each subject have the same shape but different locations as long as these parameters are different. From (7), one can identify $n-1$ individual intercepts because the model includes an overall intercept. If β_0 is dropped from the model, the individual intercepts depend on the estimated quantile τ.

Under the assumption that n nuisance parameters in (7) are location shifters, a number of approaches are based on simple transformations of the data to remove the fixed effects from the model. One of the most popular approaches is the two-step estimator proposed by Canay (2011). The method is relatively simple to implement in practice and has been frequently adopted in the empirical literature. In the first step, he obtains $\hat{\alpha}_i = T^{-1}\sum_{t=1}^{T}(y_{it} - \boldsymbol{x}_{it}'\tilde{\boldsymbol{\beta}})$, where $\tilde{\boldsymbol{\beta}}$ is a consistent estimator of $\boldsymbol{\beta}$. In the second step, he employs quantile regression to estimate $\boldsymbol{\beta}(\tau)$ using $y_{it} - \hat{\alpha}_i$ as a response variable. Despite the computational appeal of the method, Besstremyannaya and Golovan (2019) argue that the estimator suffers from several issues and show that the approach can lead to incorrect inference in panels typically considered in applied microeconomics.

Although the literature offered evidence that the estimator performed well in small and large samples, the asymptotic behavior of the fixed effects estimator was not rigorously established until the work by Kato et al. (2012). Their contribution is important because they derive results under *general conditions*, providing a better understanding of large sample results in relation to other fixed effects estimators for non-linear models. Kato et al. (2012) obtain consistency and asymptotic normality results for the fixed effects estimator under n and T jointly tending to infinity. They consider $m = 1$, and they estimate a model without an overall intercept, so that α_i cannot be a location shift as in Koenker (2004). Under regularity conditions, they obtain the Bahadur representation of $\hat{\boldsymbol{\beta}}(\tau)$ as,

$$\hat{\boldsymbol{\beta}}(\tau) - \boldsymbol{\beta}(\tau) = \boldsymbol{\Gamma}_n^{-1}\left[\frac{1}{nT}\sum_{i=1}^{n}\sum_{t=1}^{T}(\boldsymbol{x}_{it} - \boldsymbol{\gamma}_i)\psi_\tau(u_{it})\right] + O_p((T/\log(n))^{-3/4}) + o_p(\|\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|), \quad (8)$$

where the quantile influence function $\psi_\tau(u) = \tau - I(u < 0)$, $\boldsymbol{\gamma}_i = \mathbb{E}(f_i(0|\boldsymbol{x}_{i1})\boldsymbol{x}_{i1})/f_i(0)$, and f_i is the density of $u_{it}(\tau) = y_{it} - \boldsymbol{x}_{it}'\hat{\boldsymbol{\beta}}(\tau) - \hat{\alpha}_i(\tau)$. The matrix $\boldsymbol{\Gamma}_n^{-1} = n^{-1}\sum_{i=1}^{n}\mathbb{E}[f_i(0|\boldsymbol{x}_{i1})\boldsymbol{x}_{i1}(\boldsymbol{x}_{i1} - \boldsymbol{\gamma}_i)']$.

If $n^2(\log(n))^3/T \to 0$, then

$$\sqrt{nT}(\hat{\boldsymbol{\beta}}(\tau) - \boldsymbol{\beta}(\tau)) \xrightarrow{d} \mathcal{N}(\mathbf{0}, \tau(1-\tau)\boldsymbol{\Gamma}^{-1}\boldsymbol{V}\boldsymbol{\Gamma}^{-1}), \quad (9)$$

where $\boldsymbol{V}$ is the limiting covariance matrix of $V_{ni} = T^{-1/2} \sum_{t=1}^{T} (\boldsymbol{x}_{it} - \boldsymbol{\gamma}_i)\psi_\tau(u_{it})$. They also show that $\hat{\boldsymbol{\beta}}(\tau)$ is consistent if $\log(n)^2/T \to 0$, which is more stringent that the standard rate for non-linear panel data models, $n/T \to 0$. More recent results obtained by a different approach indicate that standard rates can indeed be obtained in models with fixed effects (Galvao et al., 2020).

The early work on fixed effects models allowed researchers to investigate other important issues that are common in applied work, such as endogenous regressors and dynamic models. Harding and Lamarche (2009) address demand-supply type endogeneity by considering an instrumental variable approach in a model with fixed effects. Galvao (2011) studies estimation of slope coefficients in models with fixed effects and uses instrumental variables to address the inconsistency of the fixed effects approach in a dynamic panel quantile model. Zhang et al. (2019) proposes an iterative estimation procedure for a model with individual effects and heterogeneous parameters for subgroups of subjects with identical slope effects.

Penalized Effects. Motivated by an interpretation of the Gaussian random effects estimator as a penalized least squares estimator, Koenker (2004) develops a variation of the fixed effects estimator by augmenting the fidelity term in (6) with a LASSO-type penalty. The penalized quantile regression estimator is defined as:

$$\underset{\boldsymbol{\beta},\boldsymbol{\alpha}\in\mathcal{A}\times\mathcal{B}}{\operatorname{argmin}} \sum_{j=1}^{m}\sum_{i=1}^{n}\sum_{t=1}^{T} \omega_j \rho_{\tau_j}(y_{it} - \boldsymbol{x}_{it}'\boldsymbol{\beta}(\tau_j) - \alpha_i) + \lambda \sum_{i=1}^{n} |\alpha_i|, \tag{10}$$

where ω_j controls the influence of the quantiles on the estimation of the individual effects, and the tuning parameter λ controls the degree of shrinkage towards zero. The parameter $\lambda \geq 0$ depends on T and it may also depend on data. Note that $\lambda = 0$ implies (6).

The estimation of n individual parameters increases the variability of the estimates of the covariate effects. However, shrinkage of these effects toward a common value reduces the variability of the fixed effects estimator of the slope parameter $\boldsymbol{\beta}$. Thus, as in all regularization problems, the selection of the tuning parameter λ is important. As shown in Lamarche (2010), under independence between the individual effects and the independent variables, the estimator $\hat{\boldsymbol{\beta}}(\tau, \lambda)$ is unbiased for all $\lambda > 0$. Thus, it is reasonable to consider choosing λ to minimize variance. The optimal choice of λ can be viewed as selecting the minimum variance estimator in the class of penalized estimators, the analog of the generalized least squares estimator for panel data. However, the selection of the tuning parameter remains a topic of investigation in the case of endogenous regressors.

More recently, Gu and Volgushev (2019) propose a method to estimate a model with group effects where individual parameters are allowed to have a group membership. The particular group structure and the slope parameters of the model are estimated using a convex minimization approach:

$$\arg\min_{\boldsymbol{\beta},\boldsymbol{\alpha}} \sum_{i=1}^{n} \sum_{t=1}^{T} \rho_\tau(y_{it} - \boldsymbol{x}_{it}'\boldsymbol{\beta} - \alpha_i) + \sum_{i \neq j} \lambda_{ij} |\alpha_i - \alpha_j|, \tag{11}$$

where the tuning parameter $\lambda_{ij} = \lambda|\check{\alpha}_i - \check{\alpha}_j|^{-2}$, and $\check{\alpha}_i$ is an individual effect estimated by fixed effects quantile regression. The individual fixed effects are grouped into K distinct groups and are allowed to be different by quantiles. The paper offers conditions and theoretical results on consistent group structure estimation and establishes asymptotic properties of the resulting joint and group-specific estimators provided that $(\log T)^3(\log n)^2/T \to 0$.

The penalized estimator can also be adapted to the analysis of randomized control trials and quasi-experimental research designs. For instance, in applications related to the evaluation of welfare experiments, individual effects can represent cost of participation on welfare and, the subject heterogeneity is therefore sparse in relationship to treatment (Hartley and Lamarche, 2018). Thus, $\sum_{i=1}^{n} |\alpha_i| = \sum_{i=1}^{n}(d_i|\alpha_{i,1}| + (1-d_i)|\alpha_{i,0}|) = \sum_{j \in \mathcal{D}_0} |\alpha_j| + \sum_{k \in \mathcal{D}_1} |\alpha_k|$, where $d_i \in \{0, 1\}$ indicates treatment status, $\mathcal{D}_0$ includes participants in the control group, and $\mathcal{D}_1$ includes participants in the treatment group. The quantile treatment effect $\delta(\tau)$ can be estimated by finding:

$$\arg\min \sum_{i=1}^{n} \sum_{t=1}^{T} \hat{w}_i(\boldsymbol{x}_{i0}) \cdot \rho_\tau \left(y_{it} - \beta_0 - \delta d_i - \alpha_i\right) + \lambda_0 \sum_{j \in \mathcal{D}_0} |\alpha_j| + \lambda_1 \sum_{k \in \mathcal{D}_1} |\alpha_k|, \tag{12}$$

where $\hat{w}_i(\boldsymbol{x}_{i0}) = d_i/\hat{p}_i(\boldsymbol{x}_{i0}) + (1-d_i)/(1-\hat{p}_i(\boldsymbol{x}_{i0}))$, and (λ_0, λ_1) is a vector of tuning parameters. The use of weights is common in quantile regression, and in (12), the weights are introduced to address selection on observables. The variable $p_i(\boldsymbol{x}_{i0}) = P(d_i = 1|\boldsymbol{x}_{i0})$ can be estimated using standard propensity score methods using observed characteristics measured before the intervention, $\boldsymbol{x}_{i0}$.

Correlated Random Effects. The correlated random-effects framework provides a compromise between misspecification arising from the omission of controlling for unobserved heterogeneity and incidental parameters bias arising from estimating a large number of nuisance parameters in a nonlinear panel model.

Using data from the states of Arizona and Washington on maternally linked births, Abrevaya and Dahl (2008) estimate a panel quantile model adopting the correlated random-effects framework of Chamberlain (1982, 1984). Consider $\alpha_i = g(\boldsymbol{x}_i) + a_i$, where $g(\cdot)$ is an unknown function, and a_i is a variable independent of $\boldsymbol{x}_i = (\boldsymbol{x}'_{i1}, \ldots, \boldsymbol{x}'_{iT})'$. The panel quantile regression model is now,

$$Q_{y_{it}}(\tau|\boldsymbol{x}_{it}, \alpha_i) = \boldsymbol{x}'_{it}\boldsymbol{\beta} + g(\boldsymbol{x}_i) + Q_\epsilon(\tau|\boldsymbol{x}_i) = \boldsymbol{x}'_{it}\boldsymbol{\beta}(\tau) + g(\boldsymbol{x}_i), \tag{13}$$

where $Q_\epsilon(\tau|\boldsymbol{x}_i)$ is the quantile of the composite error $\epsilon_{it} = a_i + u_{it}$. Importantly, T is assumed fixed and the number of parameters do not increase with the sample size. Abrevaya and Dahl (2008) estimate model (13) using linear quantile regression and they propose a series of tests for the correlated effects model.

Harding and Lamarche (2017) adopt a similar framework to Abrevaya and Dahl (2008). The function $g(\boldsymbol{x}_i)$, however, is non-parametrically estimated and a_i is treated as a parameter. This flexible correlated random effects procedure involves two steps. First, they model $g(\boldsymbol{x}_i)$ as a linear expansion of B-splines:

$$g(\boldsymbol{x}_i)'\boldsymbol{\phi}(\tau) \approx \boldsymbol{b}(\boldsymbol{x}_{i1})'\boldsymbol{\phi}_1(\tau) + \boldsymbol{b}(\boldsymbol{x}_{i2})'\boldsymbol{\phi}_2(\tau) + ... + \boldsymbol{b}(\boldsymbol{x}_{iT})'\boldsymbol{\phi}_T(\tau), \tag{14}$$

where $\boldsymbol{b}(\boldsymbol{x}_{ij}) = (b_1(\boldsymbol{x}_{ij}), \ldots, b_{k_n+h+1}(\boldsymbol{x}_{ij}))'$ is a B-spline basis function, k_n is the number of knots, h is the degree of the B-spline basis, and $\boldsymbol{\phi}$ is the spline coefficient vector. The model becomes a linear quantile regression model in all coefficients and can be estimated using:

$$\underset{\boldsymbol{\beta},\boldsymbol{\phi},\boldsymbol{a} \in \mathcal{B}\times\mathcal{G}\times\mathcal{A}}{\operatorname{argmin}} \sum_{j=1}^{m}\sum_{i=1}^{n}\sum_{t=1}^{T} \omega_j \rho_{\tau_j}(y_{it} - \boldsymbol{x}'_{it}\boldsymbol{\beta}(\tau_j) - g(\boldsymbol{x}_i)'\boldsymbol{\phi}(\tau_j) - a_i) + \lambda \sum_{i=1}^{n} |a_i|. \tag{15}$$

Although flexibility in specification is important, an alternative practical specification for $g(\cdot)$ is to use a known parametric function of time-series averages or, alternatively, a vector of covariates for each of the n subjects.

As in the case of (10), the shrinkage of individual effects can improve the performance of the slope estimator. Moreover, the inclusion of the penalty term offers the possibility of investigating if the model specification and assumptions are supported by data. If $\alpha_i = g(\boldsymbol{x}_i) + a_i$, and $g(\cdot)$ is correctly specified, then a_i cannot be correlated with $\boldsymbol{x}_i$. Thus, shrinking a_i towards zero should not affect the bias of the slope parameter estimator at any value of λ. On the other hand, the limiting case of the penalized estimator when $\lambda \to 0$ is not affected by the possible correlation between a_i and $\boldsymbol{x}_i$. Harding and Lamarche (2017) propose a Hausman-type test to evaluate whether the fixed effects estimator, $\lim_{\lambda\to 0} \hat{\beta}(\tau, \lambda)$, and the penalized estimator, $\hat{\beta}(\tau, \lambda)$, offer significantly different results. If a_i and $\boldsymbol{x}_i$ are independent, $\lim_{\lambda\to 0} \hat{\beta}(\tau, \lambda)$ and $\hat{\beta}(\tau, \lambda)$ should be relatively similar for any value of λ.

Considering a correlated effects framework, Arellano and Bonhomme (2016) develop a flexible approach that offers practical important advantages over competing techniques. They assume an outcome $y_{it} = \sum_{k=1}^{K_1} \theta_k(U_{it}) h_k(\boldsymbol{x}_{it}, \alpha_i)$, and employ an equation that models dependence between latent-specific heterogeneity and the independent variables, $\alpha_i = \sum_{k=1}^{K_2} \delta_k(V_i) g_k(\boldsymbol{x}_i)$. The variables U_{it} and V_i are uniform random variables, and $h_k(\cdot)$ and $g_k(\cdot)$ are unknown functions. They consider two moment conditions of the form,

$$\begin{aligned} \sum_{t=1}^{T} E\left(\boldsymbol{W}_{it}(\alpha_i)\psi_\tau(y_{it} - \boldsymbol{W}_{it}(\alpha_i)'\boldsymbol{\theta}(\tau))\right) &= \mathbf{0} \\ \sum_{t=1}^{T} E\left(\boldsymbol{Z}_i \psi_\tau(\alpha_i - \boldsymbol{Z}_i'\boldsymbol{\delta}(\tau))\right) &= \mathbf{0}, \end{aligned}$$

where $\boldsymbol{W}_{it}(\alpha_i) = (h_1(\boldsymbol{x}_{it}, \alpha_i), \ldots, h_{K_1}(\boldsymbol{x}_{it}, \alpha_i))'$ and $\boldsymbol{Z}_i = (g_1(\boldsymbol{x}_i), \ldots, g_{K_2}(\boldsymbol{x}_i))'$ are vectors of $\boldsymbol{x}_{it}$, $\boldsymbol{x}_i$ and α_i. These moment conditions are unfeasible since α_i is a latent variable, but Arellano and Bonhomme (2016) demonstrate that identification and estimation can be achieved by integrating out the latent variable of the model. They propose a sequential EM algorithm, which is similar to the standard EM algorithm, with the exception that the M-step is performed using quantile regression check functions rather than likelihood functions.

Statistical Inference. In the last years, a number of papers studied inference for fixed and penalized effects methods in models with i.i.d. and non-i.i.d errors. Kato et al. (2012) show that the asymptotic covariance matrix of the slope estimator $\hat{\boldsymbol{\beta}}(\tau)$ in (9) can be consistently estimated by $\tau(1-\tau)\hat{\boldsymbol{\Gamma}}^{-1}\hat{\boldsymbol{V}}\hat{\boldsymbol{\Gamma}}^{-1}$, where $\hat{\boldsymbol{\Gamma}}$ and $\hat{\boldsymbol{V}}$ are matrices estimated by kernel methods. Considering a similar estimator, Yoon and Galvao (2020) allow for temporal correlation within each individual. Progress on statistical inference for the penalized quantile regression estimator has been moving at a slower pace. Lamarche and Parker (2020) recently propose a wild residual bootstrap procedure and show that it is asymptotically valid for approximating the distribution of the penalized estimator. They also show that the pairs bootstrap that samples cross-sectional units with replacement does not approximate well the limiting distribution of the penalized estimator.

The recent literature also includes models with clustered standard errors (Hagemann, 2017), although addressing individual heterogeneity and clustered observations in high-dimensional panel data models has been an elusive topic of investigation.

Computational Aspects and Sparsity-based Estimation. Feasible estimation of fixed effects models with large N has been possible due to developments for sparse matrices in

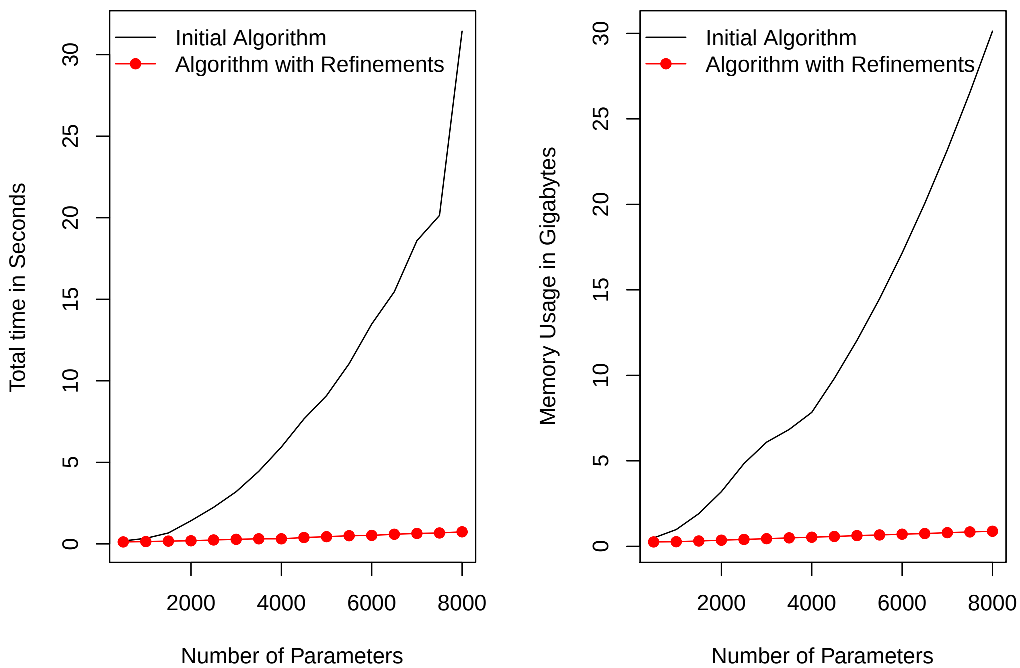

FIGURE 1. *Estimation of Fixed Effects Models using Sparse Matrices.*

quantile regression, which are available in the `R` libraries `SparseM` and `quantreg` by Koenker (2021). The sparsity-based estimation was a key development employed first in Koenker (2004). More recently, a simple refinement of the initial algorithm for panel quantiles led to reductions of the computational cost of estimating fixed effects models, as shown in Figure 1.[1] For instance, to estimate over 25,000 parameters, Harding and Lamarche (2019b) first create a sparse matrix, using the `Matrix` library developed by Bates and Maechler (2021), and then, they accommodate the algorithm to employ other routines in `SparseM` and `quantreg`.

An Empirical Application. We now illustrate the importance of addressing unobserved heterogeneity in quantile regression models using data from Chay and Powell (2001). We are interested in estimating the black-white earnings gap after the introduction of Title VII of

[1]Figure 1 presents results based on estimating a panel quantile regression model with $T = 20$ and N ranging from 500 to 8,000. The model also includes a slope parameter. The simulations using `R` version 4.0.4 were conducted on a Dell Optiplex 7070 workstation which has a i7-9700 Intel Core 4.7GHz processor and 32GB RAM.

the Civil Rights Act of 1964, which prohibited discrimination by employers on the basis of race and gender. The data is from the Current Population Survey and it covers the period between 1957 and 1971. The original data is heavily censored by the maximum taxable earnings level for Social Security (Galvao et al., 2013), but in this illustration we use a sub-sample of workers whose earnings are below the maximum during the entire observation window. We estimate the effect of relative earnings of black workers in the southern states of the U.S. after the introduction of the Civil Rights Act, controlling for education and age of the worker. The response variable is the natural logarithm of earnings. The black-white gap estimates range from -0.25 at the 0.1 quantile to -0.08 at the 0.9 quantile. With that in mind, we concentrate on the earnings gap after the introduction of the Civil Rights Act of 1964.

Figure 2 shows results obtained from standard quantile regression (QR), fixed effects quantile regression (FEQR), and penalized quantile regression (PQR). It also presents point-wise confidence intervals obtained by a wild bootstrap procedure (Lamarche and Parker, 2020). First, we note the significant differences between QR estimates and the other panel quantile estimates. Moreover, the QR estimate is 0.26 at the 0.1 quantiles, and the estimates tend to decrease as we go across quantiles. The results in the lower tail mistakenly suggest that the Civil Rights Act of 1964 eliminated the earnings gap among conditionally low-earners. These results show that not controlling for individual heterogeneity can lead to biased results, as in linear panel data models. The other panels present results for the FEQR and PQR estimators. They suggest that the relative earnings of black workers improved after the introduction of the policy, and the effect seems larger at the upper conditional quantiles of the earnings distribution.

Panel Data Quantile Regression 2.0: Beyond Individual Heterogeneity

The initial class of panel data models focused on latent time-invariant heterogeneity. A recent number of papers propose to estimate variations of the original panel quantile regression model with individual effects. Graham et al. (2018) study identification and estimation of a generalization of the random coefficient representation of the quantile regression model (Koenker, 2005), introducing dependence between regressors and random coefficients. They consider $Q_{Y|\boldsymbol{X}}(\tau|\boldsymbol{x}) = \boldsymbol{x}'\boldsymbol{\beta}(\tau,\boldsymbol{x}) + \boldsymbol{w}'\boldsymbol{\delta}(\tau)$, where $\boldsymbol{w} = w(\boldsymbol{x})$ is a matrix of functions of the regressors $\boldsymbol{x}$, and $\boldsymbol{\beta}(\tau,\boldsymbol{x})$ is a nonparametric function. Moreover, Machado and Santos Silva (2019) propose a method-of-moments estimator for a linear location-scale shift model,

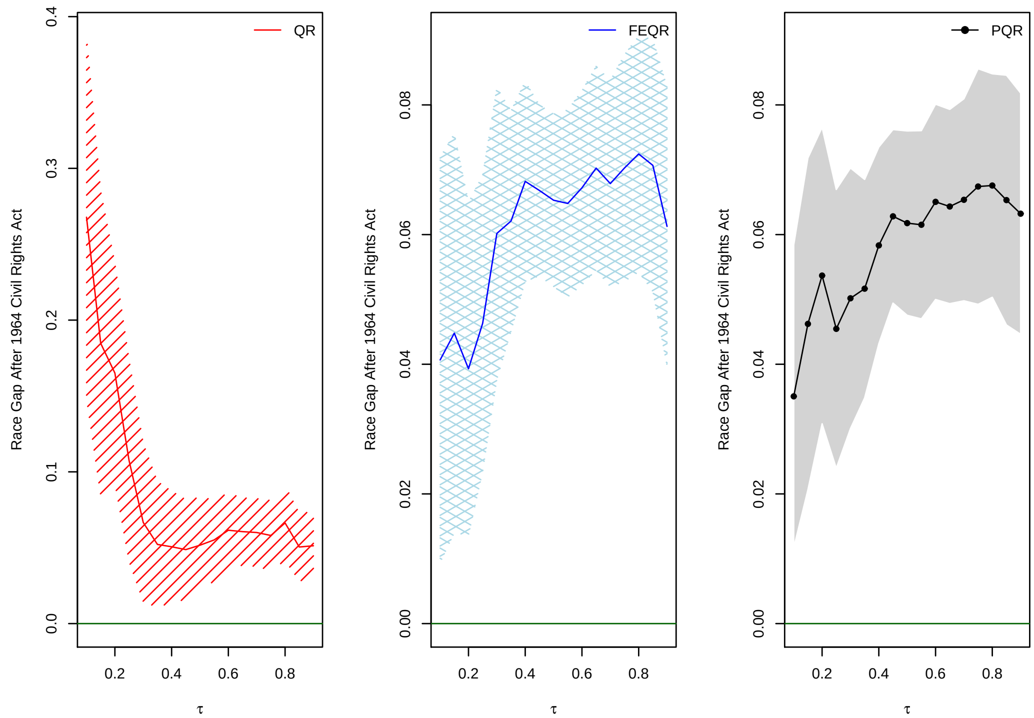

FIGURE 2. *The black-white race gap after the Civil Rights Act of 1964.*

$y_{it} = \alpha_i + \boldsymbol{x}'_{it}\boldsymbol{\beta} + (\delta_i + \boldsymbol{z}'_{it}\boldsymbol{\gamma})u_{it}$, where $\boldsymbol{z}_{it}$ is a vector of known transformations of the independent variables. The approach is simple to implement in practice and allows estimation of $(\alpha_i, \delta_i, \boldsymbol{\beta}', \boldsymbol{\gamma}')$ and the quantile function of the error term.

In what follows, we review new models and methods, including a number of generalizations to address other sources of misspecification.

Instrumental Variables. Harding and Lamarche (2009) consider the estimation of a panel data model with endogenous independent variables, allowing correlation between the error term of the model and the regressors. The estimation procedure is similar to the framework developed by Chernozhukov and Hansen (2005, 2006) for the use of instrumental variables in quantile regression models. Consider the objective function:

$$R(\tau, \boldsymbol{\delta}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\alpha}) = \sum_{i=1}^{n} \sum_{t=1}^{T} \rho_\tau \left(y_{it} - \boldsymbol{d}'_{it}\boldsymbol{\delta} - \boldsymbol{x}'_{it}\boldsymbol{\beta} - \boldsymbol{z}'_{it}\boldsymbol{\alpha} - \hat{\boldsymbol{w}}'_{it}\boldsymbol{\gamma} \right), \tag{16}$$

where $\rho_\tau = u(\tau - I(u \leq 0))$ is the quantile regression loss function, $\hat{\boldsymbol{w}}$ is the least squares projection of the endogenous variables $\boldsymbol{d}$ on the instruments $\boldsymbol{w}$, $\boldsymbol{x}$ is a vector of the exogenous variables, and $\boldsymbol{z}$ is an incidence vector for the individual effects. The estimation proceeds in two steps. First, they minimize $R(\tau, \boldsymbol{\delta}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\alpha})$ for $\boldsymbol{\beta}$, $\boldsymbol{\gamma}$, and $\boldsymbol{\alpha}$ as functions of τ and $\boldsymbol{\delta}$. Then, they estimate the coefficient on the endogenous variable by finding the value of $\boldsymbol{\delta}$, which minimizes a weighted distance function defined on $\boldsymbol{\gamma}$: $\hat{\boldsymbol{\delta}}(\tau) = \arg\min_{\boldsymbol{\delta}} \ \hat{\boldsymbol{\gamma}}(\tau, \boldsymbol{\delta})' \boldsymbol{A} \hat{\boldsymbol{\gamma}}(\tau, \boldsymbol{\delta})$, for a given positive definite matrix $\boldsymbol{A}$. The two-step IV approach was also adopted by Galvao (2011) to estimate a dynamic panel data model with fixed effects.

Chetverikov et al. (2016) extend the ideas of Hausman and Taylor (1981) to a quantile model. The model for individual i in group g is $Q_{y_{ig}}(\tau | \boldsymbol{z}_{ig}, \boldsymbol{x}_g, \boldsymbol{\alpha}_g) = \boldsymbol{z}_{ig}' \boldsymbol{\alpha}_g(\tau)$, where $\boldsymbol{\alpha}_{g,1}(\tau) = \boldsymbol{x}_g' \boldsymbol{\beta}(\tau) + \epsilon_g(\tau)$. The variable $\boldsymbol{\alpha}_g(\tau)$ is a group-specific vector of effects and $\boldsymbol{\alpha}_{g,1}(\tau)$ is the first element of the vector $\boldsymbol{\alpha}_g(\tau)$. The parameter of interest $\boldsymbol{\beta}(\tau)$ is estimated in two steps. First, $\hat{\boldsymbol{\alpha}}(\tau)$ is obtained by employing quantile regression of y_{ig} on $\boldsymbol{z}_{ig}$. Second, $\hat{\boldsymbol{\beta}}(\tau)$ is obtained by employing a linear regression of $\hat{\boldsymbol{\alpha}}_{g,1}(\tau)$ on $\boldsymbol{x}_g$ using internally generated instruments. The paper shows that the estimator is consistent and asymptotically normal under $G^{2/3}(\log(N_G))/N_G \to 0$ as $G \to \infty$, where G is the number of groups and N_G is the number of individuals within each group $g = 1, 2, \ldots, N_G$. Therefore, the approach is most useful in applications when both N_G and G are large. The advantage, however, is that the method allows the use of internal instruments, which are transformations of the exogenous component of the vector $\boldsymbol{z}_{ig}$.

Censored Models and Attrition. The literature on quantile regression for longitudinal data with censored observations or drop-outs is very limited. Wang and Fygenson (2009) study inference in a quantile regression for longitudinal data where some of the responses are left censored by fixed values. They propose a rank score test for statistical inference considering censoring and intra-subject correlation. Galvao et al. (2013) investigate the estimation of a panel quantile regression model with fixed effects when the response variable is censored. The paper proposes two-step estimators that are obtained by applying fixed effects quantile regression to subsets of observations. This approach overcomes practical challenges in situations where the number of estimated parameters is large and the degree of censoring is high.

Let y_{it}^* denote the potentially left-censored t-th response of the i-th individual and let $y_{it} = \max(C_{it}, y_{it}^*)$ be its corresponding observed value, where C_{it} is a known censoring point. Let $\delta_{it} = 1(y_{it}^* > C_{it})$ indicate uncensored observations and $\pi_0(\boldsymbol{x}_{it}, C_{it}) = P(\delta_{it} = 1 | \boldsymbol{x}_{it}, C_{it})$ be the propensity score. The model is $Q_{y^*_{it}}(\tau|\boldsymbol{x}_{it}) = \alpha_i(\tau) + \boldsymbol{x}'_{it}\boldsymbol{\beta}(\tau)$, and based on equivariance to monotone transformations, $Q_{y_{it}}(\tau|\boldsymbol{x}_{it}) = \max(C_{it}, Q_{y^*_{it}}(\tau|\boldsymbol{x}_{it}))$. This leads to a version of the Powell (1986) estimator with fixed effects based on $\rho_\tau(y_{it} - \max(C_{it}, \alpha_i + \boldsymbol{x}'_{it}\boldsymbol{\beta}))$. An asymptotically equivalent estimator can be obtained by minimizing $\sum_i \sum_t \rho_\tau(y_{it} - \alpha_i - \boldsymbol{x}'_{it}\boldsymbol{\beta})$ restricted to a subset of observations i and t where the conditional quantile function is above C_{it}. Alternatively, as Galvao et al. (2013) show, one can simply apply fixed effects quantile regression to observations that satisfy $\pi_0(\boldsymbol{x}_{it}, C_{it}) > 1 - \tau$.

Naturally, the true propensity score function is unknown and needs to be estimated. They propose a feasible two-step estimator. In step 1, $\pi_0(\boldsymbol{x}_{it}, C_{it})$ is estimated by using either a parametric or nonparametric regression method for binary data. Then, in step 2, $(\boldsymbol{\alpha}', \boldsymbol{\beta}')$ can be estimated by applying fixed effects quantile regression to a subset of observations:

$$\sum_{i=1}^{n}\sum_{t=1}^{T} \rho_\tau(y_{it} - \alpha_i - \boldsymbol{x}'_{it}\boldsymbol{\beta}) 1\left(\hat{\pi}(\boldsymbol{x}_{it}, C_{it}) > 1 - \tau + c_n\right), \tag{17}$$

where $\hat{\pi}(\boldsymbol{x}_{it}, C_{it})$ is the estimated propensity score and c_n is a small constant that tends to zero as n and T tend to infinity. These estimators are simple to compute and easy to implement in panel data applications with a large number of subjects.

Recent papers also propose quantile regression estimators for panel data when units drop out of the sample after an initial random sample of subjects is considered. Lipsitz et al. (1997) and Maitra and Vahid (2006) propose a weighting scheme for longitudinal data, but their estimating equations would lead to inconsistent and inefficient results for general models of attrition. Harding and Lamarche (2019a) study identification of slope parameters in unbalanced panels and focus the investigation on attrition arising from observable and unobservable factors. They propose a two-step estimator that accommodates penalized estimation, which is shown to reduce the bias of the quantile fixed effects estimator. The estimator is shown to be consistent under selection on unobservables, and a key condition is that the researcher has access to streaming samples, which are similar to the refreshment samples used in additive-non-ignorable models (Hirano et al., 2001).

Quantile Treatment Effects. The estimation of quantile treatment effects (QTE) has a long tradition in quantile regression, but, until recently, the literature abstracted away from panel data (Koenker, 2005). There are now several important contributions that expand our ability to investigate the effect of interventions beyond average treatment effects, which can be misleading in practice (Bitler et al., 2006).

Chernozhukov et al. (2013) and Chernozhukov et al. (2015) study identification and estimation of QTE and quantile structural functions using panel data. Chernozhukov et al. (2013) develop an elegant identification approach of quantile effects in nonseparable models. The paper provides an approach for the important problem of nonparametric quantile regression in panel data, allowing for multidimensional heterogeneity.

Using panel data but assuming a difference-in-differences model, Callaway et al. (2018) and Callaway and Li (2019) investigate identification and estimation of QTE and quantile treatment effects on the treated units (QTT). Callaway and Li (2019) show identification under a Copula Stability Condition, which implies that dependence between potential outcomes does not change over time. They estimate QTT by $\hat{\delta}(\tau) = \hat{Q}_{y_{1t}}(\tau|D = 1) - \hat{Q}_{y_{0t}}(\tau|D = 1)$, where the quantile functions for the treated group is estimated by inverting the empirical probability function, and the counterfactual quantile is,

$$\hat{Q}_{y_{0t}}(y) = \hat{F}^{-1}_{0t|D=1}(y) = \frac{1}{n_D} \sum_{i \in \mathcal{D}} 1\{\hat{F}_{\Delta y_{1|D=0}}(y_0) \leq y - \hat{F}_{y_{t-1|D=1}}(y_1)\},$$

where $y_0 = \hat{F}_{\Delta y_{t-1}|D=1}(\Delta y_{it-1})$, $y_1 = \hat{F}_{\Delta y_{it-2}|D=1}(\Delta y_{it-2})$, n_D is the number of observations in the treated group, and $\mathcal{D}$ is the set of treated individuals. Uniform confidence bands for the QTT are obtained using the bootstrap, and the use of the approach is illustrated by estimating the effect of a minimum wage increase on county level unemployment.

Factor Models

A number of recent approaches generalize factor models to allow for quantile-specific parameters. Ma et al. (2020) consider a semiparametric quantile model with a factor structure. Chen et al. (2020), and Battistin et al. (2020) consider generalizations of the classical factor model, allowing factors to have a location-scale shift effect on the distribution of the response variable. Chen et al. (2020) study estimation of $Q_{X_{it}}(\tau|\boldsymbol{f}_t(\tau)) = \boldsymbol{\gamma}_i'(\tau)\boldsymbol{f}_t(\tau)$, for some observed variable X_{it}. Note that this representation allows both loadings and factors to have a distributional effect, and it implicitly assumes that the loadings are quantile-specific parameters to be estimated. They propose an interactive approach and show that the estimator achieves the same rate of convergence as other popular estimators for an approximate factor model. Battistin et al. (2020) study identification and estimation of quantiles of the distribution of treatment effects considering a generalized factor model. Their method differs from competing approaches as they consider an approximation of quantile functions of counterfactual outcomes based on a small-variance measurement error strategy.

There is also progress on factor-augmented panel quantile regression models. The assumption of time-invariant heterogeneity associated with the traditional fixed effects framework can be restrictive in applications, in particular when applied researchers have a large number of repeated observations on each subject. In recent years, a number of papers consider a generalization of the familiar panel data models (3) by augmenting the model with $\eta_{it} = \boldsymbol{\gamma}_i' \boldsymbol{f}_t$, where $\boldsymbol{\gamma}_i$ is a vector of factor loadings and $\boldsymbol{f}_t$ corresponds to common time-varying factors (Pesaran, 2006; Bai, 2009; Moon and Weidner, 2015, 2017). This extension substantially increases the flexibility of controlling for unobserved heterogeneity, but until recently, existing approaches did not offer the possibility of estimating a quantile regression model.

The recent literature on models with interactive fixed effects and factor structure includes papers by Harding and Lamarche (2014), Ando and Bai (2020), and Harding et al. (2020). This research considers a class of models similar to the following panel data model:

$$y_{it} = \alpha_i + \boldsymbol{\theta}_i' \boldsymbol{d}_{it} + \boldsymbol{\beta}_i' \boldsymbol{x}_{it} + \boldsymbol{\gamma}_i' \boldsymbol{f}_t + u_{it}, \tag{18}$$

where $y_{it} \in \mathbb{R}$ is the response variable for cross-sectional unit i at time t, $\boldsymbol{d}_{it}$ is a vector of endogenous variables that can include a lagged dependent variable y_{it-1}, $\boldsymbol{x}_{it}$ is a vector of exogenous regressors, and u_{it} is the error term. As before, $\boldsymbol{f}_t$ is a vector of unobserved factors, $\boldsymbol{\gamma}_i$ is a vector of latent factor loadings, and α_i is an individual specific effect potentially correlated with the regressor variables.

Models with Homogeneous Slopes. Harding and Lamarche (2014) propose an estimator for a panel quantile model with homogeneous slope coefficients and interactive effects. They allow for dependence between the two static conditional quantile functions of the response variable y_{it} and the endogeneous variable $\boldsymbol{d}_{it}$:

$$\begin{aligned} Q_{y_{it}}(\tau | \boldsymbol{d}_{it}, \boldsymbol{x}_{it}, \boldsymbol{\gamma}_i, \boldsymbol{f}_t) &= \boldsymbol{d}_{it}' \boldsymbol{\theta}(\tau) + \boldsymbol{x}_{it}' \boldsymbol{\beta}(\tau) + \boldsymbol{f}_t' \boldsymbol{\gamma}_i(\tau), && (19) \\ Q_{d_{it}}(\tau | \boldsymbol{w}_{it}, \boldsymbol{x}_{it}, \boldsymbol{\gamma}_i, \boldsymbol{f}_t) &= g(\tau, \boldsymbol{w}_{it}, \boldsymbol{x}_{it}, \boldsymbol{\gamma}_i, \boldsymbol{f}_t) && (20) \end{aligned}$$

where $\boldsymbol{w}_{it}$ is a vector of instrumental variables and $g(\cdot)$ is a known linear conditional quantile function. Using the convention that the conditional quantile function $Q_{y_{it}}(\tau | \boldsymbol{d}_{it}, \boldsymbol{x}_{it}, \boldsymbol{\gamma}_i, \boldsymbol{f}_t)$ is evaluated at $\boldsymbol{d}_{it} = Q_{d_{it}}(\tau | \boldsymbol{w}_{it}, \boldsymbol{x}_{it}, \boldsymbol{\gamma}_i, \boldsymbol{f}_t)$, they substitute (20) into (19). After basic algebraic manipulations, they obtain an expression for the unknown factors, $\boldsymbol{f}_t(\tau) \approx \boldsymbol{\Psi}(\tau; \bar{\boldsymbol{z}}_t, \bar{\boldsymbol{w}}_t, \bar{\boldsymbol{x}}_t)$, where $\bar{z}_t = n^{-1} \sum_{i=1}^n z_{it}$, $\boldsymbol{z}_{it} = (y_{it}, \boldsymbol{d}_{it}')'$ and $\boldsymbol{\Psi}$ is a known parametric function of cross-sectional averages of the endogenous and exogenous variables.

The model (19) can be estimated as follows. Consider,

$$C_{it}(\tau, \boldsymbol{\theta}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\phi}) = \rho_\tau \left(y_{it} - \boldsymbol{d}_{it}'\boldsymbol{\theta} - \boldsymbol{x}_{it}'\boldsymbol{\beta} - \hat{\boldsymbol{f}}_t'\boldsymbol{\gamma} - \hat{\boldsymbol{\Phi}}_{it}'\boldsymbol{\phi} \right). \tag{21}$$

where $\boldsymbol{\Phi}_{it}(\tau) = \boldsymbol{\Phi}(\tau; \boldsymbol{w}_{it}, \boldsymbol{x}_{it}, \boldsymbol{f}_t, \boldsymbol{\gamma}_i)$ is a vector of transformations of instruments. First, obtain $\boldsymbol{\beta}(\tau, \boldsymbol{\theta})$, $\hat{\boldsymbol{\gamma}}(\tau, \boldsymbol{\theta})$, and $\hat{\boldsymbol{\phi}}(\tau, \boldsymbol{\theta})$ as the solution of

$$\min_{\beta, \phi, \gamma \in \mathcal{B} \times \mathcal{G} \times \mathcal{F}} \sum_{i=1}^{n} \sum_{t=1}^{T} C_{it}(\tau, \boldsymbol{\theta}; \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\phi}). \tag{22}$$

Second, find $\boldsymbol{\theta}$ that minimizes a weighted distance function defined on $\boldsymbol{\phi}$:

$$\hat{\boldsymbol{\theta}}(\tau) = \underset{\boldsymbol{\theta} \in \boldsymbol{\Theta}}{\operatorname{argmin}} \left\{ \hat{\boldsymbol{\phi}}(\tau, \boldsymbol{\theta})' \hat{\boldsymbol{A}}(\tau) \hat{\boldsymbol{\phi}}(\tau, \boldsymbol{\theta}) \right\}, \tag{23}$$

for a positive definite matrix $\boldsymbol{A}$. The quantile regression estimator for a model with interactive effects is then defined as $(\hat{\boldsymbol{\theta}}(\tau)', \hat{\boldsymbol{\beta}}(\hat{\boldsymbol{\theta}}(\tau), \tau)', \hat{\boldsymbol{\gamma}}(\hat{\boldsymbol{\theta}}(\tau), \tau)')$. As shown in Harding and Lamarche (2014), the method can be accommodated to include individual effects.

Models with Heterogeneous Slopes. Ando and Bai (2020) investigate the estimation of a model with interactive effects, but contrary to Harding and Lamarche (2014), they allow for heterogeneous slope coefficients and a factor structure that vary across quantiles. They consider the following conditional quantile function:

$$Q_{y_{it}}(\tau | \boldsymbol{x}_{it}, \boldsymbol{\gamma}_i, \boldsymbol{f}_t) = \boldsymbol{x}_{it}'\boldsymbol{\beta}_i(\tau) + \boldsymbol{\gamma}_i'(\tau) \boldsymbol{f}_t(\tau), \tag{24}$$

which is estimated by simultaneously minimizing

$$\ell(y | \boldsymbol{X}, \boldsymbol{B}, \boldsymbol{\Lambda}, \boldsymbol{F}) = \sum_{i=1}^{n} \sum_{t=1}^{T} \rho_\tau (y_{it} - \boldsymbol{x}_{it}'\boldsymbol{\beta}_i(\tau) - \boldsymbol{\gamma}_i'(\tau) \boldsymbol{f}_t(\tau)), \tag{25}$$

where $\boldsymbol{B} = (\boldsymbol{\beta}_1'(\tau), \ldots, \boldsymbol{\beta}_n'(\tau))'$, $\boldsymbol{\Lambda} = (\boldsymbol{\gamma}_1'(\tau), \ldots, \boldsymbol{\gamma}_n'(\tau))'$, and $\boldsymbol{F} = (\boldsymbol{f}_1'(\tau), \ldots, \boldsymbol{f}_T'(\tau))'$. They propose a frequentist approach and a Bayesian algorithm that is based on data augmentation, and they apply the method to investigate quantile co-movement of the global financial market.

The asymptotic theory of these estimators is different than the estimators for linear models with interactive effects. Under regularity conditions, Ando and Bai (2020) demonstrate that the proposed estimator $\hat{\boldsymbol{\beta}}_i(\tau)$ for $1 \leq i \leq n$ is uniformly consistent. Moreover, they show the estimator for the loadings and factors are uniformly consistent, i.e. $\max_{1 \leq i \leq n} \|\hat{\boldsymbol{\gamma}}_i(\tau) - \boldsymbol{\gamma}_i(\tau)\| = o_p(1)$, and $\max_{1 \leq t \leq T} \|\hat{\boldsymbol{f}}_t(\tau) - \boldsymbol{f}_t(\tau)\| = o_p(1)$, provided that $\log(T)/\sqrt{n} \to 0$ and $\log(n)/\sqrt{T} \to 0$. They also show that the estimator is asymptotically normal if $T^{1/2}/n^{1-c} \to 0$ and $n^{1/2}/T^{1-c} \to 0$ for $1/16 < c < 1/2$.

As in Ando and Bai (2020), Harding et al. (2020) consider the estimation of a panel quantile function with heterogeneous coefficients. However, they consider a dynamic quantile regression model:

$$Q_{y_{it}}(\tau|y_{it-1}, \boldsymbol{x}_{it}, \boldsymbol{f}_t) = \alpha_i(\tau) + \theta_i(\tau)y_{it-1} + \boldsymbol{x}_{it}'\boldsymbol{\beta}_i(\tau) + \boldsymbol{f}_t'\boldsymbol{\gamma}_i(\tau). \tag{26}$$

The $p_x \times 1$ vector of regressors is assumed to follow the general linear process $\boldsymbol{x}_{it} = \boldsymbol{\alpha}_{ix} + \boldsymbol{\Gamma}_i'\boldsymbol{f}_t + \boldsymbol{v}_{it}$, where $\boldsymbol{\alpha}_{ix}$ is an individual effect, $\boldsymbol{\Gamma}_i$ is a matrix of factor loadings, and $\boldsymbol{v}_{it}$ follows a stationary process independently distributed of other variables in the model. Due to the combination of cross-sectional error dependence and dynamics, existing panel quantile regression approaches are inconsistent for the estimation of $(\theta_i, \boldsymbol{\beta}_i')'$ for $i = 1, \ldots, n$.

Harding et al. (2020) derive a large n representation for a linear combination of the latent factors. Let $\boldsymbol{\delta}_i(L) = \sum_{l=0}^{\infty} \boldsymbol{\delta}_{il}L^l$, $\boldsymbol{\delta}_{il} = (\boldsymbol{\delta}_{iy,l}', \boldsymbol{\delta}_{ix,l}')'$, $\boldsymbol{\delta}_{iy,l}$ is a reduced form coefficient for the cross-sectional average of y_{it-l}, $\boldsymbol{\delta}_{ix,l}$ is a reduced form coefficient for the cross-sectional average of $\mathbf{x}_{it-l}$, and $\bar{\mathbf{z}}_{t-l} = (\bar{y}_{t-l}, \bar{\mathbf{x}}_{t-l}')'$. The lag operator is denoted by L. Substituting the representation of the factors in the model, they obtain

$$y_{it} = \beta_{i0} + \theta_i y_{it-1} + \mathbf{x}_{it}'\boldsymbol{\beta}_i + \sum_{l=0}^{p_T} \bar{\mathbf{z}}_{t-l}'\boldsymbol{\delta}_{il} + u_{it} + h_{it,n}, \tag{27}$$

where $\beta_{i0} = \alpha_i + \boldsymbol{\gamma}_i'\mathbf{f}_0$ and $h_{it,n} = \sum_{l=p_T+1}^{\infty} \bar{\mathbf{z}}_{t-l}'\boldsymbol{\delta}_{il} + O_p(n^{-1/2})$. It is important to note that (27) includes the original error term, u_{it}, and an approximation error, $h_{it,n}$. Under regularity conditions, they obtain an expression for the difference between the unfeasible quantile regression problem and the quantile regression problem that corresponds to (27):

$$\left| \frac{1}{T} \sum_{t=1}^{T} (\rho_\tau(u_{it} + h_{it,n}) - \rho_\tau(u_{it})) \right| \leq \frac{\rho^{p_T+1}}{1-\rho} B_z + O_p\left(\frac{1}{\sqrt{n}}\right), \tag{28}$$

where B_z is a variable that depends on the vector $\bar{\mathbf{z}}$ and is asymptotically bounded. They show that this approximation error is asymptotically negligible because the right hand side of (28) tends to zero as n, T, and $p_T \to \infty$, since $0 < \rho < 1$. Therefore, $\boldsymbol{f}_t$ can be replaced by the current and lagged cross-section averages $\bar{\mathbf{z}}_t$, leading to

$$Q_{y_{it}}(\tau|\bar{\mathfrak{F}}_{it}) = \alpha_i(\tau) + \theta_i(\tau)y_{it-1} + \mathbf{x}_{it}'\boldsymbol{\beta}_i(\tau) + \sum_{l=0}^{p_T} \bar{\mathbf{z}}_{t-l}'\boldsymbol{\delta}_{il}(\tau), \tag{29}$$

where the feasible set $\bar{\mathfrak{F}}_{it}$ includes $\mathbf{x}_{it}$, y_{it-1}, and $\bar{\mathbf{z}}_{t-s}$ for all $s = 0, 1, 2, \ldots, p_T$.

This quantile function can now be used to estimate the parameters of interest $\boldsymbol{\vartheta}_i(\tau) := (\theta_i(\tau), \boldsymbol{\beta}_i(\tau)')'$ for each i and $0 < \tau < 1$. In large T panels, one can obtain

$$\hat{\boldsymbol{\pi}}_i(\tau) = \arg \min_{\boldsymbol{\pi}_i \in \boldsymbol{\Pi}_i} \frac{1}{T} \sum_{t=1}^{T} \rho_\tau(y_{it} - \mathbf{X}_{it}' \boldsymbol{\pi}_i), \tag{30}$$

where $\boldsymbol{\pi}_i(\tau) := (\theta_i(\tau), \boldsymbol{\beta}_i(\tau)', \alpha_i(\tau), \boldsymbol{\delta}_i(\tau)')'$ and $\mathbf{X}_{it} = (y_{it-1}, \mathbf{x}_{it}', 1, \bar{\mathbf{z}}_t', \bar{\mathbf{z}}_{t-1}', ..., \bar{\mathbf{z}}_{t-p_T}')'$ is a vector of regressors. They also propose a quantile mean group estimator for $\boldsymbol{\vartheta}(\tau) := E((\theta_i(\tau), \boldsymbol{\beta}_i(\tau)')')$. The estimator is,

$$\hat{\boldsymbol{\vartheta}}(\tau) = \frac{1}{n} \sum_{i=1}^{n} \hat{\boldsymbol{\vartheta}}_i(\tau) = \frac{1}{n} \sum_{i=1}^{n} \left(\boldsymbol{\Xi}_i \circ \hat{\boldsymbol{\pi}}_i(\tau) \right), \tag{31}$$

where $\circ$ denotes Hadamard product and $\boldsymbol{\Xi}_i = (\boldsymbol{\iota}_i', \mathbf{0}_i')'$ with $\boldsymbol{\iota}_i$ denoting a $p_x + 1$ dimensional vector of ones and $\mathbf{0}_i$ a $(p_x + 1)(p_T + 1)$ dimensional vector of zeros.

The increasing availability of large T microeconometric panels creates opportunities to estimate individual-specific treatment effects that vary across the quantiles of the response distribution. The next section provides an illustration of the previous approaches in models with a factor structure.

An Empirical Illustration. We employ residential electricity consumption data from a randomized control trial in Ireland. Smart meter data is obtained from the Irish Social Science Data Archive (ISSDA) and CER Smart Metering Project, and electricity consumption is measured over 30-minute intervals for $n = 670$ households. The participants of the program were assigned to different treatment types but we only consider one type here for simplicity. Treated customers were charged 13.5 cents per kilowatt hour (kwh) from 8 am to 11 pm, with the exception of 5 pm to 7 pm when they were charged 26 cents. The households in the treatment group were provided with an in-home display device which shows electricity usage and price in real time as well as a bimonthly bill combined with an energy usage statement. The control group has a time invariant rate of 14.1 cents per kwh and receives bimonthly electricity bills. We do not use data from 11 pm to 8 am, and we consider usage on Wednesdays in the first six months of the experiment. The total number of observations per household is $T = 780$.

The dependent variable is the natural logarithm of electricity consumption, measured in kwh at the residential level. Following Harding and Lamarche (2019a), we introduce a large number of control variables including average temperature in Ireland, average relative humidity, an indicator for household size, an indicator variable for whether the head of the household is employed, indicators for the number of rooms in the house, and indicators for

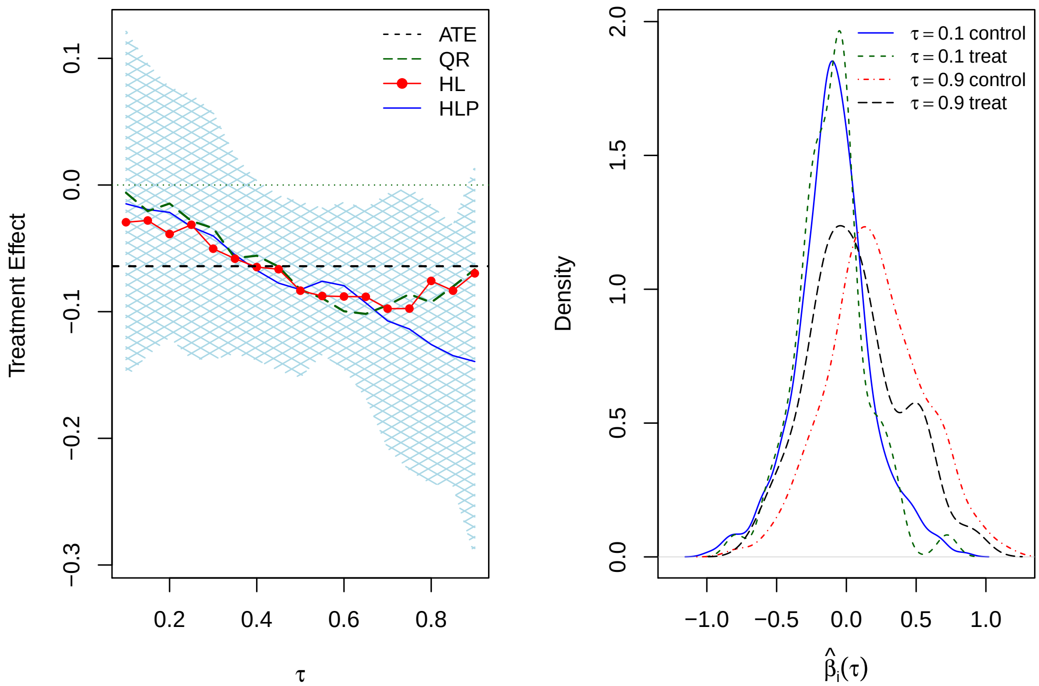

FIGURE 3. *Panel quantile regression results for a factor model of residential electricity consumption.*

other characteristics of the house. We present results for the treatment effect using Figure 3.

The left panel of Figure 3 shows average treatment effects (ATE) and quantile treatment effects, which are estimated by quantile regression (QR), quantile regression for a model with homogeneous treatments considered by Harding and Lamarche (2014) (HL), and quantile regression for a model with heterogeneous treatments as in Harding et al. (2020) (HLP). The figure also shows point-wise confidence intervals for the HLP estimator. We see that the QTEs are negative, and savings (reductions in electricity usage) tend to be bigger in the upper tail. Importantly, ignoring slope heterogeneity seems to produce different estimated QTEs among households with conditional higher use. The right panel presents the empirical distribution of the estimated coefficients by treatment status, and shows how the largest reductions in electricity usage are estimated at the 0.9 quantile of the conditional distribution.

From Theory to Practice: An Incomplete Journey

The initial interest in panel data quantile regression was primarily motivated by the limitations of classical panel data methods to uncover heterogeneous effects. The first applications of quantile regression using longitudinal data illustrated the desire for more flexible approaches "capable of revealing departures from underlying assumptions of parametric models" (Wei et al. (2006); see also Wei and He (2006)). Moreover, as pointed out by Arias et al. (2002), quantile regression applied to longitudinal data did not provide a useful tool for empirical researchers when potentially different sources of latent heterogeneity determined a reduced-form parameter.

In recent years, however, advances in panel data quantile regression allowed practitioners to estimate flexible models under a variety of different assumptions. These advances included a number of methods and algorithms that created opportunities for more informative and robust empirical analysis. The trend is reflected in the number of applications published in top-general and top-field journals in Economics and Finance such as the *American Economic Journal: Policy*, *Journal of the European Economic Association*, *Journal of Development Economics*, *Journal of Financial Economics*, *Journal of Labor Economics*, and *Journal of Public Economics*, to name a few.

Despite the relative success, panel quantile regression is still in its infant stage and many theoretical and practical issues are not fully addressed. The increased availability of large T micro-econometric panels, administrative data, experimental data, and digital records will open up new challenges and possibilities. As we encounter vast amounts of data, new models and innovative computational frameworks will become necessary. These future developments are likely to solidify the growing consensus that heterogeneity of effects are consistent with both static and dynamic economic models (Bitler et al., 2006; de Castro and Galvao, 2019).

Further Reading

Abadie A., Angrist J. and Imbens G. (2002). Instrumental variables estimates of the effect of subsidized raining on the quantiles of trainee earnings. *Econometrica*, 70(1):91–117.

Arellano, M. and Bonhomme, S. (2017). Quantile selection models with an application to understanding changes in wage inequality. *Econometrica*, 85(1):1–28.

Athey, S. and G. W. Imbens (2006). Identification and inference in nonlinear difference-in-differences models. *Econometrica*, 74(2), 431–497.

Chernozhukov, V. and Hong, H. (2002). Three-step censored quantile regression and extramarital affairs, *Journal of the American Statistical Association*, 97: 872–882.

Chernozhukov, V., Fernández-Val, I. and Weidner, M. (2021). Network and panel quantile effects via distribution regression, *Journal of Econometrics*, forthcoming.

Hausman, J. A., H. Liu, Y. Luo and C. Palmer (2021). Errors in the dependent variable of quantile regression models, *Econometrica*, forthcoming.

Firpo, S. (2007). Efficient semiparametric estimation of quantile treatment effects. *Econometrica*, 75 (1):259–276.

Frölich, M. and Melly, B. (2013). Unconditional quantile treatment effects under endogeneity, *Journal of Business and Economic Statistics*, 31: 346–357.

Galvao, A. F. and Wang, L. (2015). Efficient minimum distance estimator for quantile regression fixed effects panel data. *Journal of Multivariate Analysis*, 133(C): 1–26.

Galvao, A. F. and Kato, K. (2016). Smoothed quantile regression for panel data, *Journal of Econometrics*, 193(1): 92-112.

Parente, P.M.D.C. and Santos Silva, J.M.C. (2016). Quantile regression with clustered data, *Journal of Econometric Methods*, 5(1): 1-15.

Powell, D. (2020). Quantile treatment effects in the presence of covariates, *Review of Economics and Statistics*, 102 (5): 994-1005.

References

Abrevaya, J. and Dahl, C. (2008). The effects of smoking and prenatal care on birth outcomes: Evidence from quantile regression estimation on panel data. *Journal of Business and Economics Statistics*, 26(4):379–397.

Ando, T. and Bai, J. (2020). Quantile co-movement in financial markets: A panel quantile model with unobserved heterogeneity. *Journal of the American Statistical Association*, 115(529):266–279.

Arellano, M. and Bonhomme, S. (2016). Nonlinear panel data estimation via quantile regressions. *The Econometrics Journal*, 19(3):C61–C94.

Arias, O., Hallock, K. F., and Sosa-Escudero, W. (2002). Individual heterogeneity in the returns to schooling: instrumental variables quantile regression using twins data. In Fitzenberger, B., Koenker, R., and Machado, J. A. F., editors, *Economic Applications of Quantile Regression*, pages 7–40. Physica-Verlag HD.

Bai, J. (2009). Panel data models with interactive fixed effects. *Econometrica*, 77(4):1229–1279.

Baltagi, B. (2013). *Econometric Analysis of Panel Data*. Wiley, New York, 5th edition.

Bates, D. and Maechler, M. (2021). Matrix: Sparse and dense matrix classes and methods. R package version 1.3-2, `www.r-project.org`.

Battistin, E., Lamarche, C., and Rettore, E. (2020). Quantiles of the gain distribution of an early child intervention. CEPR Discussion Paper No. DP14721. Available at SSRN: https://ssrn.com/abstract=3603950.

Besstremyannaya, G. and Golovan, S. (2019). Reconsideration of a simple approach to quantile regression for panel data. *Econometrics Journal*, 22(3):292–308.

Bitler, M. P., Gelbach, J. B., and Hoynes, H. W. (2006). What mean impacts miss: Distributional effects of welfare reform experiments. *American Economic Review*, 96(4):988–1012.

Callaway, B. and Li, T. (2019). Quantile treatment effects in difference in differences models with panel data. *Quantitative Economics*, 10(4):1579–1618.

Callaway, B., Li, T., and Oka, T. (2018). Quantile treatment effects in difference in differences models under dependence restrictions and with only two time periods. *Journal of Econometrics*, 206(2):395 – 413.

Canay, I. A. (2011). A simple approach to quantile regression for panel data. *The Econometrics Journal*, 14(3):368–386.

Chamberlain, G. (1982). Multivariate regression models for panel data. *Journal of Econometrics*, 18:5–46.

Chamberlain, G. (1984). Panel data. *Handbook of Econometrics, Edited by Z. Griliches and M.D. Intriligator*, 2, Amsterdam: North-Holland.

Chay, K. and Powell, J. L. (2001). Semiparametric censored regression models. *Journal of Economic Perspectives*, 15(4):29–42.

Chen, L., Dolado, J. J., and Gonzalo, J. (2020). Quantile factor models. arXiv:1911.02173 [econ.EM].

Chernozhukov, V., Fernández-Val, I., Hahn, J., and Newey, W. (2013). Average and quantile effects in nonseparable panel models. *Econometrica*, 81:535–580.

Chernozhukov, V., Fernández-Val, I., Hoderlein, S., Holzmann, H., and Newey, W. (2015). Nonparametric identification in panels using quantiles. *Journal of Econometrics*, 188(2):378 – 392.

Chernozhukov, V. and Hansen, C. (2005). An iv model of quantile treatment effects. *Econometrica*, 73(1):245–262.

Chernozhukov, V. and Hansen, C. (2006). Instrumental quantile regression inference for structural and treatment effect models. *Journal of Econometrics*, 132(2):491 – 525.

Chetverikov, D., Larsen, B., and Palmer, C. (2016). IV quantile regression for group-level treatments, with an application to the distributional effects of trade. *Econometrica*, 84(2):809–833.

Corradin, S. and Maddaloni, A. (2020). The importance of being special: Repo markets during the crisis. *Journal of Financial Economics*, 137(2):392–429.

Dahl, C. M., le Maire, D., and Munch, J. R. (2013). Wage dispersion and decentralization of wage bargaining. *Journal of Labor Economics*, 31(3):501–533.

de Castro, L. and Galvao, A. F. (2019). Dynamic quantile models of rational behavior. *Econometrica*, 87(6):1893–1939.

De Silva, D. G., Kosmopoulou, G., and Lamarche, C. (2009). The effect of information on the bidding and survival of entrants in procurement auctions. *Journal of Public Economics*, 93(1):56–72.

De Silva, D. G., Kosmopoulou, G., and Lamarche, C. (2017). Subcontracting and the survival of plants in the road construction industry: A panel quantile regression analysis. *Journal of Economic Behavior Organization*, 137:113–131.

Fell, H. and Kaffine, D. T. (2018). The fall of coal: Joint impacts of fuel prices and renewables on generation and emissions. *American Economic Journal: Economic Policy*, 10(2):90–116.

Fuchs, A. and Gehring, K. (2017). The Home Bias in Sovereign Ratings. *Journal of the European Economic Association*, 15(6):1386–1423.

Galvao, A., Gu, J., and Volgushev, S. (2020). On the unbiased asymptotic normality of quantile regression with fixed effects. *Journal of Econometrics*, 218:178–215.

Galvao, A. F. (2011). Quantile regression for dynamic panel data with fixed effects. *Journal of Econometrics*, 164(1):142 – 157.

Galvao, A. F., Lamarche, C., and Lima, L. R. (2013). Estimation of censored quantile regression for panel data with fixed effects. *Journal of the American Statistical Association*, 108(503):1075–1089.

Graham, B. S., Hahn, J., Poirier, A., and Powell, J. L. (2018). A quantile correlated random coefficients panel data model. *Journal of Econometrics*, 206(2):305 – 335.

Gu, J. and Volgushev, S. (2019). Panel data quantile regression with grouped fixed effects. *Journal of Econometrics*, 213(1):68 – 91.

Hagemann, A. (2017). Cluster-robust bootstrap inference in quantile regression models. *Journal of the American Statistical Association*, 112(517):446–456.

Hahn, J. and Newey, W. (2004). Jackknife and analytical bias reduction for nonlinear panel models. *Econometrica*, 72(4):1295–1319.

Harding, M. and Lamarche, C. (2009). A quantile regression approach for estimating panel data models using instrumental variables. *Economics Letters*, 104:133–135.

Harding, M. and Lamarche, C. (2014). Estimating and testing a quantile regression model with interactive effects. *Journal of Econometrics*, pages 101 – 113.

Harding, M. and Lamarche, C. (2017). Penalized quantile regression with semiparametric correlated effects: An application with heterogeneous preferences. *Journal of Applied Econometrics*, 32(2):342–358.

Harding, M. and Lamarche, C. (2019a). A panel quantile approach to attrition bias in big data: Evidence from a randomized experiment. *Journal of Econometrics*, 211(1):61 – 82.

Harding, M. and Lamarche, C. (2019b). Penalized Estimation of a Quantile Count Model for Panel Data. *Annals of Economics and Statistics*, (134):177–206.

Harding, M., Lamarche, C., and Pesaran, M. H. (2020). Common correlated effects estimation of heterogeneous dynamic panel quantile regression models. *Journal of Applied Econometrics*, 35(3):294–314.

Hartley, R. P. and Lamarche, C. (2018). Behavioral responses and welfare reform: Evidence from a randomized experiment. *Labour Economics*, 54(C):135–151.

Hausman, J. A. and Taylor, W. E. (1981). Panel data and unobservable individual effects. *Econometrica*, 49(6):1377–1398.

Hirano, K., Imbens, G. W., Ridder, G., and Rubin, D. B. (2001). Combining panel data sets with attrition and refreshment samples. *Econometrica*, 69(6):pp. 1645–1659.

Hsiao, C. (2014). *Analysis of Panel Data*. Cambridge University Press, New York, 3rd edition.

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

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

Koenker, R. (2005). *Quantile Regression*. Cambridge University Press.

Koenker, R. (2017). Quantile regression: 40 years on. *Annual Review of Economics*, 9(1):155–176.

Koenker, R. (2021). Quantreg: Quantile regression. R package version 5.85, `www.r-project.org`.

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

Lamarche, C. (2006). *Panel Data Quantile Regression*. PhD thesis, University of Illinois at Urbana-Champaign.

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

Lamarche, C. and Parker, T. (2020). Wild bootstrap inference for penalized quantile regression for longitudinal data. arXiv:2004.05127 [econ.EM].

Lipsitz, S. R., Fitzmaurice, G. M., Molenberghs, G., and Zhao, L. P. (1997). Quantile regression methods for longitudinal data with drop-outs: Application to cd4 cell counts of patients infected with the human immunodeficiency virus. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 46(4):463–476.

Ma, S., Linton, O., and Gao, J. (2020). Estimation and inference in semiparametric quantile factor models. *Journal of Econometrics*.

Machado, J. A. and Santos Silva, J. (2019). Quantiles via moments. *Journal of Econometrics*, 213(1):145 – 173.

Maitra, P. and Vahid, F. (2006). The effect of household characteristics on living standards in south africa 1993-1998: a quantile regression analysis with sample attrition. *Journal of Applied Econometrics*, 21(7):999–1018.

Moon, H. R. and Weidner, M. (2015). Linear regression for panel with unknown number of factors as interactive fixed effects. *Econometrica*, 83(4):1543–1579.

Moon, H. R. and Weidner, M. (2017). Dynamic linear panel regression models with interactive effects. *Econometric Theory*, pages 158–195.

Neyman, J. and Scott, E. L. (1948). Consistent estimates based on partially consistent observations. *Econometrica*, 16:1–32.

Pesaran, M. H. (2006). Estimation and inference in large heterogeneous panels with a multifactor error structure. *Econometrica*, 74(4):967–1012.

Powell, D. (2020). Does labor supply respond to transitory income? evidence from the economic stimulus payments of 2008. *Journal of Labor Economics*, 38(1):1–38.

Powell, J. L. (1986). Censored regression quantiles. *Journal of Econometrics*, 32:143–155.

Rosen, A. M. (2012). Set identification via quantile restrictions in short panels. *Journal of Econometrics*, 166(1):127 – 137.

Wang, H. J. and Fygenson, M. (2009). Inference for censored quantile regression in longitudinal studies. 37:756–781.

Wei, Y. and He, X. (2006). Conditional growth charts. *Annals of Statistics*, 34(5):2069–2097.

Wei, Y., Pere, A., Koenker, R., and He, X. (2006). Quantile regression methods for reference growth charts. *Statistics in Medicine*, 25(8):1369–1382.

Yoon, J. and Galvao, A. F. (2020). Cluster robust covariance matrix estimation in panel quantile regression with individual fixed effects. *Quantitative Economics*, 11(2):579–608.

Zhang, Y., Wang, H. J., and Zhu, Z. (2019). Quantile-regression-based clustering for panel data. *Journal of Econometrics*, 213(1):54 – 67.