

PARAMETRIC MODELING OF QUANTILE REGRESSION COEFFICIENT FUNCTIONS WITH LONGITUDINAL DATA

Paolo Frumento,

Matteo Bottai

and

Iván Fernández-Val

University of Pisa,
Karolinska Institutet
and Boston University

June 2, 2020

Author's Footnote

Paolo Frumento (paolo.frumento@unipi.it) is Associate Professor at the Department of Political Sciences, University of Pisa, Italy.

Matteo Bottai (matteo.bottai@ki.se) is Professor at the Unit of Biostatistics at Karolinska Institutet, Institute of Environmental Medicine, Stockholm, Sweden.

Iván Fernández-Val (ivanf@bu.edu) is Professor at the Department of Economics, Boston University.

Address for correspondence: Department of Political Sciences, University of Pisa, Via F. Serafini, 3, 56126 Pisa, Italy.

Abstract

In ordinary quantile regression, quantiles of different order are estimated one at a time. An alternative approach, which is referred to as *quantile regression coefficients modeling* (QRCM), is to model quantile regression coefficients as parametric functions of the order of the quantile. In this paper, we describe how the QRCM paradigm can be applied to longitudinal data. We introduce a two-level quantile function, in which two different quantile regression models are used to describe the (conditional) distribution of the within-subject response and that of the individual effects. We propose a novel type of penalized fixed-effects estimator, and discuss its advantages over standard methods based on ℓ_1 and ℓ_2 penalization. We provide model identifiability conditions, derive asymptotic properties, describe goodness-of-fit measures and model selection criteria, present simulation results, and discuss an application. The proposed method has been implemented in the R package `qrcm`.

Keywords: Longitudinal quantile regression, two-level quantile function, parametric quantile function, penalized fixed-effects, R package `qrcm`.

1 INTRODUCTION

Quantile regression (e.g., Koenker and Bassett, 1978; Koenker, 2005) has become a standard method in many fields, including medicine, epidemiology, economics, and social sciences. Different solutions have been proposed to extend quantile regression to longitudinal data, in which the same individuals or clusters are observed repeatedly.

In conditional models, that include fixed- and random-effects models, the dependence between observations is accounted for by introducing individual-specific parameters, or “individual effects”. In fixed-effects models, the individual effects are treated as parameters, avoiding distributional assumptions and allowing for a simple computation. A penalized fixed-effects estimator for longitudinal quantile regression has been proposed by Koenker (2004), and similar approaches have been used in Lamarche (2010), Canay (2011), and Kato et al. (2012).¹ In random-effects models, the individual effects are described by a parametric distribution. Different methods have been proposed to combine the parametric likelihood of the random effects with the estimating equation of ordinary quantile regression. Geraci and Bottai (2007, 2014) used the log-likelihood of an asymmetric Laplace distribution, and Kim and Yang (2011) described an empirical likelihood method. Abrevaya and Dahl (2008) adapted the correlated random effect approach of Chamberlain (1984) to quantile regression, and Arellano and Bonhomme (2016) marginalized the loss function of quantile regression with respect to the posterior distribution of the individual effects. Farcomeni (2012), Marino, Tzavidis, and Alfó (2016), and Alfó, Salvati, and Ranalli (2017) used finite mixtures to approximate the probability density function of the individual effects through a discrete distribution.

Marginal models have also been described in the literature. Leng and Zhang (2014)

¹Chernozhukov et al. (2018) considered an alternative to quantile regression for estimation of quantile effects in longitudinal data based on distribution regression.

defined a set of unbiased estimating equations carrying information on the correlation structure. A similar approach was used by Zhao, Lian, and Liang (2017) to implement longitudinal single-index quantile regression.

In this paper, we adopt the conditional paradigm and introduce a two-level quantile function, in which both the distribution of the within-subject response (level 1) and that of the individual effects (level 2) are described by quantile regression models. With this approach, the distribution of the individual effects is not subject to strong parametric assumptions and is allowed to depend on level-2 covariates. Following Frumento and Bottai (2016, 2017) and Yang, Chen, and Chang (2017), we describe the level-1 and level-2 quantile regression coefficients by (flexible) parametric functions of the order of the quantile. Compared with standard quantile regression, in which quantiles are estimated one at a time, this modeling approach presents numerous advantages, that include a simpler computation and inference (owing to a smooth objective function), increased statistical efficiency, and easy interpretability of the results.

To fit the model, we introduce a new form of penalized fixed-effects estimator in which the penalty term carries information on level-2 parameters. This method presents important advantages over standard ℓ_1 and ℓ_2 penalization. In particular, it avoids the problem of selecting a tuning constant, and allows to estimate the level-2 coefficients of the model using fixed-effects techniques.

The paper is structured as follows. We describe a general model in Section 2, and discuss model building in Section 3. We introduce an estimator in Section 4, and in Section 5 we derive its asymptotic properties. In Section 6 we present goodness-of-fit measures and tools for model selection, and in Section 7 we report simulation results. Section 8 concludes the paper with the analysis of a dataset relating plasma neutrophil gelatinase-associated lipocalin (NGAL) to sepsis status. Appendix A provides a general

asymptotic expansion for fixed effects estimators with mixed-rates asymptotics and applies it to derive the asymptotic distribution of the proposed estimator. We discuss computation in Appendix B, and present extended simulation results in Appendix C. The R package `qrqm` implements the described estimator and provides a variety of auxiliary functions for model building, summary, plotting, prediction, and goodness-of-fit assessment.

2 THE MODEL

2.1 A two-level quantile function

We consider a cluster data structure, in which N individuals or clusters are observed repeatedly. We denote by $i = 1, \dots, N$ the index of the subject, and by $t = 1, \dots, T$ the within-subject index, such that the total sample size is NT . Designs in which T varies across clusters are also possible, at the cost of a slightly more complicated notation.

We denote by Y_{it} a response variable of interest, and assume that

$$Y_{it} = \mathbf{x}_{it}^T \boldsymbol{\beta}(U_{it}) + \mathbf{z}_i^T \boldsymbol{\gamma}(V_i) \quad (1)$$

where \mathbf{x}_{it} is a d_x -dimensional vector of level-1 covariates, with associated parameter $\boldsymbol{\beta}(\cdot)$; and \mathbf{z}_i is a d_z -dimensional vector of level-2 covariates, with associated parameter $\boldsymbol{\gamma}(\cdot)$.

We assume that (i) $\mathbf{x}_{it}^T \boldsymbol{\beta}(\cdot)$ and $\mathbf{z}_i^T \boldsymbol{\gamma}(\cdot)$ are a.s. non-decreasing functions of their arguments, and (ii) U_{it} and V_i are $U(0, 1)$ variables, independent of each other and of the covariates. Based on model (1), $\alpha_i = \mathbf{z}_i^T \boldsymbol{\gamma}(V_i)$ is an individual effect with conditional quantile function $\mathbf{z}_i^T \boldsymbol{\gamma}(\cdot)$, while $\mathbf{x}_{it}^T \boldsymbol{\beta}(\cdot)$ is the conditional quantile function of $Y_{it} - \alpha_i$.

The level-1 quantile regression model, $\mathbf{x}_{it}^T \boldsymbol{\beta}(\cdot)$, has the standard interpretation (e.g., Koenker, 2004): it characterizes the “within” part of the distribution, purged of the individual effects. The level-2 regression model, $\mathbf{z}_i^T \boldsymbol{\gamma}(\cdot)$, describes the distribution of the

between-subject differences with respect to a reference value which typically corresponds to a “mean” or “median” individual.

Consider, for example, a clinical study in which patients are repeatedly measured their body mass index (BMI) during their lifetime. The level-1 part of the model describes the conditional quantiles of BMI in a “typical” patient, i.e., someone with an individual effect equal to 0. Level-1 predictors include time-varying characteristics, such as the age of the patient at each observation, as well as constant traits, such as the gender of the patient. The level-2 model accounts for the between-patient heterogeneity, and describes the conditional quantiles of the individual effects. Level-2 covariates can only include time-invariant traits, such as the gender, and summary statistics of level-1 covariates, e.g., the age at the first examination. Note that the dimension of the level-1 covariates, \mathbf{x}_{it} , is NT , while that of the level-2 covariates, \mathbf{z}_i , is N .

Unlike the “standard” approaches, that do not consider the effect of level-2 covariates, our modeling framework allows to investigate the determinants of the between-subject variability. For example, in the linear random-intercept model, the level-2 response is described by a $N(0, \phi^2)$ distribution, in which $\phi^2 = \text{var}(\alpha_i)$ is interpreted as the “between” variance and is assumed to be unaffected by predictors. This model may fail to capture important features of the data, such as the fact that the variance of the individual effects is different in males and females. Model (1), instead, allows including gender as level-2 predictor.

Using a quantile regression approach permits avoiding strong parametric assumptions such as normality and homoskedasticity, that are often used in likelihood-based modeling. In the existing literature on longitudinal quantile regression, however, a quantile regression model is usually applied to the level-1 response, but not to the individual effects, that are treated as nuisance parameters. In our paradigm, instead, the two parts of the distribution

are considered “equally important”, in the sense that the same modeling structure is used to describe the quantiles of the within-subject response, and those of the between-subject differences. As shown later in the paper, working with model (1) permits using the same techniques to estimate both the level-1 and the level-2 parameters, and avoids combining level-1 quantile regression methods with likelihood-based level-2 estimators as for example in Kim and Yang (2011). This leads to rather simple procedures for estimation and inference, in which a fundamental role is played by the two independent uniform random variables (U_{it}, V_i) that generate the data.

2.2 Parametric coefficient functions

Through the paper, we assume that the quantile regression coefficient functions, $\beta(\cdot)$ and $\gamma(\cdot)$, can be modeled parametrically:

$$\beta(u) = \beta(u \mid \boldsymbol{\theta}), \quad \gamma(v) = \gamma(v \mid \boldsymbol{\phi}), \quad (2)$$

where $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ are unknown model parameters. This modeling approach was used by Frumento and Bottai (2016, 2017), and is exemplified in Figure 1. The broken line in figure represents standard regression coefficients at quantiles $u = (0.01, 0.02, \dots, 0.99)$. The estimated coefficients show a non-smooth, volatile trend and, although consistently positive, are almost never significant. A parametric model can be used to characterize the coefficient function with few parameters and describe it by a simple, closed-form mathematical expression. In Figure 1 we propose a linear fit, $\beta(u \mid \boldsymbol{\theta}) = \theta_0 + \theta_1 u$, that is represented by a dashed line. This simple model reveals the underlying trend and permits achieving statistical significance.

Compared with standard quantile regression, which works in a quantile-by-quantile fashion, modeling quantile functions parametrically simplifies estimation and inference and

yields important advantages in terms of parsimony, efficiency, and ease of interpretation. Moreover, it allows for model identification in the presence of latent structures or missing information, making it simple to apply quantile regression to censored and truncated data (Frumento and Bottai, 2017).

On the other hand, this approach requires formulating a parametric model for the coefficient functions, $\beta(u | \theta)$ and $\gamma(v | \phi)$. This task is not straightforward and the existing literature on the subject is lacking. In Section 3 we describe in details model building, provide guidelines, and suggest a variety of possible parametrizations.




Figure 1: The broken line exemplifies a typical behavior of quantile regression estimators, while the dashed line suggests a hypothetical linear trend, $\beta(u | \theta) = \theta_0 + \theta_1 u$. Pointwise confidence intervals are represented by the light and dark shaded areas, respectively. The dotted line indicates the zero.

3 TWO-LEVEL MODELING OF QUANTILE REGRESSION COEFFICIENT FUNCTIONS

We assume model (1) to hold, and parametrize the quantile regression coefficient functions as follows:

$$\boldsymbol{\beta}(u \mid \boldsymbol{\theta}) = \boldsymbol{\theta}\mathbf{b}(u), \quad \boldsymbol{\gamma}(v \mid \boldsymbol{\phi}) = \boldsymbol{\phi}\mathbf{c}(v), \quad (3)$$

where $\mathbf{b}(u) = [b_1(u), \dots, b_{d_b}(u)]^\top$ and $\mathbf{c}(v) = [c_1(v), \dots, c_{d_c}(v)]^\top$ are d_b - and d_c -dimensional sets of known functions. With this notation, $\boldsymbol{\theta}$ is a $d_x \times d_b$ matrix, and $\boldsymbol{\phi}$ is a $d_z \times d_c$ matrix. The data-generating process can be written as

$$Y_{it} = \mathbf{x}_{it}^\top \boldsymbol{\theta} \mathbf{b}(U_{it}) + \mathbf{z}_i^\top \boldsymbol{\phi} \mathbf{c}(V_i). \quad (4)$$

Although other parametrizations are possible (e.g., $\boldsymbol{\beta}(u \mid \boldsymbol{\theta})$ and $\boldsymbol{\gamma}(v \mid \boldsymbol{\phi})$ may be allowed to be nonlinear functions of $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$), model (3) is very flexible and computationally convenient. We illustrate the potentials of this modeling approach with a number of examples, and provide general guidelines for model building.

3.1 A simple model

Consider the following model with a single level-1 covariate x , and no level-2 predictors:

$$Y_{it} = \beta_0(U_{it}) + \beta_1(U_{it})x_{it} + \gamma_0(V_i).$$

Denote by $\zeta(\cdot)$ the quantile function of a standard normal distribution, and assume that

$$\beta_0(u \mid \boldsymbol{\theta}) = \theta_{00} + \theta_{01}\zeta(u),$$

$$\beta_1(u \mid \boldsymbol{\theta}) = \theta_{10},$$

$$\gamma(v | \boldsymbol{\phi}) = \phi \zeta(v).$$

This is just a reformulation of the standard linear random-intercept model, in which $Y_{it} = \theta_{00} + \theta_{10}x_{it} + \alpha_i + \epsilon_{it}$ with $\alpha_i \sim N(0, \phi^2)$ and $\epsilon_{it} \sim N(0, \theta_{01}^2)$. In this model, θ_{00} corresponds to the intercept of the “fixed” part, while θ_{01}^2 and ϕ^2 are interpreted as the “within” and “between” variance components. In the equivalent quantile regression model, θ_{00} is the “intercept” of $\beta_0(u | \boldsymbol{\theta})$ and corresponds to $\beta_0(0.5 | \boldsymbol{\theta})$, while θ_{01} and ϕ are “slopes” associated with $\zeta(\cdot)$ in the level-1 and level-2 part of the quantile function, respectively. The regression coefficient of x , $\beta_1(u | \boldsymbol{\theta})$, is assumed to be constant across quantiles, forcing homoskedasticity.

3.2 A more flexible model

The standard linear random-intercept model is rather restrictive and, within the described framework, can be easily generalized by choosing a different specification of $\beta(\cdot | \boldsymbol{\theta})$ and $\gamma(\cdot | \boldsymbol{\phi})$. For example, one may define

$$\beta_0(u | \boldsymbol{\theta}) = \theta_{00} + \theta_{01}u + \theta_{02}u^2 + \theta_{03}u^3 + \theta_{04}\zeta(u),$$

$$\beta_1(u | \boldsymbol{\theta}) = \theta_{10} + \theta_{11}u,$$

$$\gamma(v | \boldsymbol{\phi}) = \phi_1 \log(2v) + \phi_2 \log(2(1-v)).$$

The intercept, $\beta_0(u | \boldsymbol{\theta})$, is modeled by a linear combination of $\zeta(u)$, the quantile function of a standard normal distribution, and three additional components, u , u^2 and u^3 , that allow for a deviation from the normal model. The resulting quantile function can be asymmetric or multimodal and does not correspond to any “standard” family of random variables. The coefficient associated with x , $\beta_1(u | \boldsymbol{\theta})$, is now assumed to be a linear function of u , allowing for data heteroskedasticity. In particular, the variance of the level-1 response is

an increasing function of x , if $\theta_{11} > 0$, and a decreasing function of it, if $\theta_{11} < 0$. Finally, the individual effects are assumed to follow a zero-median asymmetric logistic distribution, which is much more flexible than the commonly used normal model.

As shown in this example, $\beta(\cdot)$ and $\gamma(\cdot)$ can be constructed as linear combinations of relatively simple functions, $\mathbf{b}(\cdot)$ and $\mathbf{c}(\cdot)$, such that $\beta(u | \boldsymbol{\theta}) = \boldsymbol{\theta}\mathbf{b}(u)$, and $\gamma(v | \boldsymbol{\phi}) = \boldsymbol{\phi}\mathbf{c}(v)$. In this framework, the model is entirely determined by the choice of $\mathbf{b}(\cdot)$ and $\mathbf{c}(\cdot)$. Useful guidelines for model building are provided in the rest of this section. Various modeling approaches are illustrated in Sections 7 and 8 of this paper, while a general discussion on quantile modeling can be found in the book by Gilchrist (2000). Finally, the documentation of the `qrqm` package (in particular the functions `iqr` and `iqrL`) includes an extensive tutorial for the practitioners.

3.3 Model building: level 1

Modeling $\beta_0(u | \boldsymbol{\theta})$. Assuming that the support of \mathbf{x} includes the zero (which can be obtained by centering the covariates), $\beta_0(\cdot | \boldsymbol{\theta})$ must be a monotonically increasing function. Prior belief or knowledge can be used to identify a meaningful parametric model. For instance, one may use the quantile function of a known distribution. Possible parametrizations of $\beta_0(u | \boldsymbol{\theta})$ include: $\theta_{00} + \theta_{01}\zeta(u)$, the normal distribution, $N(\theta_{00}, \theta_{01}^2)$; $-\theta_{01} \log(1 - u)$, the exponential distribution, $\text{Exp}(\theta_{01})$; $\theta_{00} + \theta_{01} \log(u/(1 - u))$, the logistic distribution, $\text{Logis}(\theta_{00}, \theta_{01})$; $\theta_{00} + \theta_{01} \log(u) + \theta_{02} \log(1 - u)$, the asymmetric logistic, $\text{aLogis}(\theta_{00}, \theta_{01}, \theta_{02})$; $\theta_{00} + \theta_{01}u$, the uniform distribution, $U(\theta_{00}, \theta_{00} + \theta_{01})$. Note that, in this framework, the parameters of well-known distributions may have an unusual interpretation. For example, the value of θ_{01} in a $U(\theta_{00}, \theta_{00} + \theta_{01})$ distribution corresponds to its range, but can also be seen as the slope of a linear quantile function, $\theta_{00} + \theta_{01}u$.

Modeling $\beta_1(u \mid \boldsymbol{\theta}), \beta_2(u \mid \boldsymbol{\theta}), \dots$. There are no general constraints to the parametric form of the regression coefficients associated with the covariates. However, the coefficient functions are usually bounded and exhibit a rather simple behavior. Sometimes, it is possible to assume that covariates only affect the location of the level-1 response, and force homoskedasticity by choosing a constant-slope model in which $\beta_j(u \mid \boldsymbol{\theta}) = \theta_{j0}, j = 1, 2, \dots$. In a more general scenario, a useful approximation is often given by a linear-slope model, $\beta_j(u \mid \boldsymbol{\theta}) = \theta_{j0} + \theta_{j1}u$, or a quadratic-slope model, $\beta_j(u \mid \boldsymbol{\theta}) = \theta_{j0} + \theta_{j1}u + \theta_{j2}u^2$, which does not impose monotone effect with respect to u .

3.4 Model building: level 2

A similar model strategy can be applied to the level-2 quantile function. There are, however, some important differences.

Modeling $\gamma_0(v \mid \boldsymbol{\phi})$. The distribution of the individual effects is typically assumed to have zero mean or median, and, for identifiability, $\gamma_0(v \mid \boldsymbol{\phi})$ does not usually include a constant term. Meaningful definitions of $\gamma_0(v \mid \boldsymbol{\phi})$ include: $\phi_{01}\zeta(v)$, the normal distribution, $N(0, \phi_{01}^2)$; $-\phi_{01} \log(1 - v)$, the exponential distribution, $\text{Exp}(\phi_{01})$; $\phi_{01} \log(v/(1 - v))$, the logistic distribution, $\text{Logis}(0, \phi_{01})$; $\phi_{01} \log(2v) + \phi_{02} \log(2(1 - v))$, a zero-median asymmetric logistic; $\phi_{01}[\log(v) + 1] + \phi_{02}[\log(1 - v) + 1]$, a zero-mean asymmetric logistic; $\phi_{01}2(v - 0.5)$, a centered uniform distribution, $U(-\phi_{01}, \phi_{01})$. In most cases, the coefficients can be interpreted as scale parameters, while the centrality parameter is fixed and equal to zero. In the exponential case, the value 0 is the minimum of the support of the individual effects, and not a measure of central tendency, while both the mean and the standard deviation of the individual effects correspond to ϕ_{01} .

Modeling $\gamma_1(v | \boldsymbol{\phi}), \gamma_2(v | \boldsymbol{\phi}), \dots$. Importantly, the described framework permits investigating how the conditional quantile function of the individual effects depends on level-2 covariates \mathbf{z}_i , which typically include cluster-invariant characteristics (e.g., gender) or summary measures of the level-1 covariates, e.g., the cluster means or medians. The variance of the individual effects is likely to differ across subgroups of the population. Also, as suggested by some authors (e.g., Lancaster, 2000), agents may select their covariates' values based on prior knowledge about their own individual effect, which induces a correlation between α_i and \mathbf{z}_i .

Modeling the effect of level-2 covariates is not trivial. To make an example, suppose that $\alpha_i = \gamma_0(V_i) + \gamma_1(V_i)z_i$, and consider the following alternative parametrizations:

$$\gamma_0(v | \boldsymbol{\phi}) = \phi_{01}\zeta(v), \quad \gamma_1(v | \boldsymbol{\phi}) = \phi_{11}\zeta(v) \quad (\text{i})$$

$$\gamma_0(v | \boldsymbol{\phi}) = \phi_{01}\zeta(v), \quad \gamma_1(v | \boldsymbol{\phi}) = \phi_{10} + \phi_{11}v \quad (\text{ii})$$

$$\gamma_0(v | \boldsymbol{\phi}) = \phi_{01}v, \quad \gamma_1(v | \boldsymbol{\phi}) = \phi_{10} + \phi_{11}v \quad (\text{iii})$$

In model (i), where $\gamma_0(v | \boldsymbol{\phi})$ and $\gamma_1(v | \boldsymbol{\phi})$ are symmetric around the zero, the conditional distribution of α_i has zero mean and median at all values of z_i . The covariate only affects the scale of the individual effects by introducing heteroskedasticity, while no linear correlation between z_i and α_i is present. Model (i) assumes normality, but allows the variance of the individual effects to be a function of the level-2 covariates, i.e. $\alpha_i | z_i \sim N(0, \phi_{01}^2 + \phi_{11}^2 z_i^2)$. For example, if z_i is binary, the “between” variance is ϕ_{01}^2 when $z_i = 0$, and $\phi_{01}^2 + \phi_{11}^2$ when $z_i = 1$.

In models (ii) and (iii), z_i and α_i have a non-zero correlation unless $\phi_{10} = 0$. In model (ii), where $\int_0^1 \gamma_0(v | \boldsymbol{\phi}) dv = 0$, the *marginal* distribution of the individual effects has zero mean if z_i is centered around its mean or $\phi_{10} + \phi_{11}/2 = 0$. In model (iii) the mean and the median of the individual effects are functions of the parameters and cannot be determined

in advance. However, if $z_i \geq 0$, model (iii) generates $\alpha_i \geq 0$ for any positive value of the parameters, implying that the “reference” individual ($\alpha_i = 0$) corresponds to someone with the smallest possible individual effect.

3.5 Additional remarks

The problem of formulating a parametric quantile function is equivalent, at least in principle, to that of choosing a parametric form for a probability density function, a hazard function, or a survival function. For example, as shown in Section 3.1, standard parametric assumptions such as normality and homoskedasticity can be directly translated into a quantile function with a simple closed-form expression. However, as suggested in Section 3.2, the models that can be used to describe a quantile function are often very different from most of the “conventional” parametric distributions, and frequently much more flexible.

An exploratory semiparametric fit can be obtained by letting $\mathbf{b}(\cdot)$ and $\mathbf{c}(\cdot)$ be the basis of a linear or polynomial spline. A flexible model can be used as a guide to find more parsimonious and efficient parametrizations. Note that standard quantile regression, in which quantiles are estimated one at a time, can be thought of as a model in which $\mathbf{b}(\cdot)$ and $\mathbf{c}(\cdot)$ are allowed to be arbitrarily flexible and the parameters $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ are virtually infinite-dimensional.

In absence of prior knowledge, one may define $\mathbf{b}(\cdot)$ and $\mathbf{c}(\cdot)$ using polynomials [e.g., u, u^2, u^3, \dots], roots [e.g., $u^{1/2}, (1-u)^{1/2}, u^{1/3}, (1-u)^{1/3}, \dots$], trigonometric functions [e.g., $\cos(2\pi u), \sin(2\pi u)$], splines, and combinations of the above. A possible strategy is to consider a “simple” quantile function (e.g., that of a normal or an exponential distribution, depending on the nature of the outcome) and allow for a departure from it, as suggested in Section 3.2.

Importantly, the model specification should reflect assumptions on the shape, support, and boundedness (or unboundedness) of the level-1 and level-2 responses. For example, if

the individual effects are believed to be symmetric, $\gamma_0(v | \boldsymbol{\phi})$ could be formed by the quantile function of a normal or logistic distribution. If the level-1 distribution has a long right tail, $\beta_0(u | \boldsymbol{\theta})$ may have a positive asymptote in $u = 1$, e.g., $\beta_0(u | \boldsymbol{\theta}) = \theta_{00} - \theta_{01} \log(1 - u) + \dots$. On the other hand, if the outcome is strictly positive, building blocks such as $\log(u)$ or $\zeta(u)$, that present a negative asymptote in $u = 0$, may not be appropriate.

Apart from the above important considerations, the choice of $\mathbf{b}(\cdot)$ and $\mathbf{c}(\cdot)$ is not as crucial as it appears. For example, the coefficient function defined by $\beta(u) = (u - 0.3)^3$ is almost identical to $\beta(u) = -1.87 + 6.20u + 1.84 \cos(u) - 5.92 \sin(u)$, the correlation between the two being about 0.99999. The fact that very different model specifications can be indistinguishable in terms of model fit is unsurprising (for example, it is almost impossible to distinguish a Normal distribution, a Student's t distribution with large degrees of freedoms, and a Gamma distribution with large shape parameter), and suggests that meaningful criteria for model selection should include parsimony and interpretability.

Often, a rather restrictive model may provide a reasonable approximation of the true data distribution, and can be preferred to a more correct model because of its simplicity. Also, parsimonious models are very rewarding in terms of precision, although they may introduce some bias. This explains why strong parametric assumptions, such as homoskedasticity and proportionality of hazards or odds, are used routinely in statistical analysis. In quantile regression, very convenient assumptions are represented by the constant-slope model (e.g., $\beta(u | \boldsymbol{\theta}) = \theta_0$), in which a certain predictor has the same effect at all quantiles, and the linear-slope model (e.g., $\beta(u | \boldsymbol{\theta}) = \theta_0 + \theta_1 u$), in which a quantile regression coefficient is assumed to be a linear function.

4 THE ESTIMATOR

Frumento and Bottai (2016) considered cross-sectional data (y_i, \mathbf{x}_i) and defined $\boldsymbol{\beta}(u | \boldsymbol{\theta}) = \boldsymbol{\theta}\mathbf{b}(u)$ as in (3). To estimate $\boldsymbol{\theta}$, they suggested minimizing

$$L(\boldsymbol{\theta}) = \int_0^1 \sum_i \rho_u(y_i - \mathbf{x}_i^T \boldsymbol{\beta}(u | \boldsymbol{\theta})) du, \quad (5)$$

which is the integral, with respect to the order of the quantile, of the loss function of standard quantile regression, $\rho_u(w) = w(u - I(w \leq 0))$ being the “check” function. This estimation method is referred to as *integrated loss minimization* (ILM) and is currently implemented in the `qrqm` R package.

To generalize this idea to longitudinal data, assume model (4) holds,

$$Y_{it} = \mathbf{x}_{it}^T \boldsymbol{\beta}(U_{it} | \boldsymbol{\theta}) + \mathbf{z}_i^T \boldsymbol{\gamma}(V_i | \boldsymbol{\phi}) = \mathbf{x}_{it}^T \boldsymbol{\theta} \mathbf{b}(U_{it}) + \mathbf{z}_i^T \boldsymbol{\phi} \mathbf{c}(V_i),$$

and denote by y_{it} a realization of Y_{it} . If the individual effects $\alpha_i = \mathbf{z}_i^T \boldsymbol{\phi} \mathbf{c}(V_i)$ were known, one could directly apply the ILM estimator to $y_{it} - \alpha_i$, to compute an estimate of $\boldsymbol{\theta}$; and to α_i , to compute an estimate of $\boldsymbol{\phi}$. This would require solving

$$\min_{\boldsymbol{\theta}} L_1(\boldsymbol{\theta}, \boldsymbol{\alpha}_N), \quad \min_{\boldsymbol{\phi}} L_2(\boldsymbol{\phi}, \boldsymbol{\alpha}_N)$$

where $\boldsymbol{\alpha}_N = (\alpha_1, \dots, \alpha_N)$,²

$$\begin{aligned} L_1(\boldsymbol{\theta}, \boldsymbol{\alpha}_N) &= \int_0^1 \sum_{i=1}^N \sum_{t=1}^T \rho_u(y_{it} - \alpha_i - \mathbf{x}_{it}^T \boldsymbol{\beta}(u | \boldsymbol{\theta})) du \\ &= \sum_{i=1}^N \sum_{t=1}^T \left\{ (y_{it} - \alpha_i)(u_{it}(\boldsymbol{\theta}, \alpha_i) - 0.5) + \mathbf{x}_{it}^T \boldsymbol{\theta} [\bar{\mathbf{B}} - \mathbf{B}(u_{it}(\boldsymbol{\theta}, \alpha_i))] \right\}, \end{aligned} \quad (6)$$

²We index $\boldsymbol{\alpha}_N$ by N to emphasize that the dimension grows with the sample size.

$$\begin{aligned}
L_2(\boldsymbol{\phi}, \boldsymbol{\alpha}_N) &= \int_0^1 \sum_{i=1}^N \rho_v(\alpha_i - \mathbf{z}_i^\top \boldsymbol{\gamma}(v | \boldsymbol{\phi})) dv \\
&= \sum_{i=1}^N \left\{ \alpha_i (v_i(\boldsymbol{\phi}, \alpha_i) - 0.5) + \mathbf{z}_i^\top \boldsymbol{\phi} [\bar{\mathbf{C}} - \mathbf{C}(v_i(\boldsymbol{\phi}, \alpha_i))] \right\}.
\end{aligned} \tag{7}$$

To obtain expressions (6)³ and (7), we used equation (9) from Frumento and Bottai (2016), and define

$$\mathbf{B}(u) = \int_0^u \mathbf{b}(s) ds, \quad \bar{\mathbf{B}} = \int_0^1 \mathbf{B}(u) du, \tag{8}$$

$$\mathbf{C}(v) = \int_0^v \mathbf{c}(s) ds, \quad \bar{\mathbf{C}} = \int_0^1 \mathbf{C}(v) dv. \tag{9}$$

In the formulas, $u_{it}(\boldsymbol{\theta}, \alpha_i)$ and $v_i(\boldsymbol{\phi}, \alpha_i)$ are such that $y_{it} - \alpha_i = \mathbf{x}_{it}^\top \boldsymbol{\theta} \mathbf{b}(u_{it}(\boldsymbol{\theta}, \alpha_i))$ and $\alpha_i = \mathbf{z}_i^\top \boldsymbol{\phi} \mathbf{c}(v_i(\boldsymbol{\phi}, \alpha_i))$, respectively. This also implies that

$$u_{it}(\boldsymbol{\theta}, \alpha_i) = F_{y-\alpha}(y_{it} - \alpha_i | \mathbf{x}_{it}, \boldsymbol{\theta}) \tag{10}$$

is the cumulative distribution of $Y_{it} - \alpha_i$, given \mathbf{x}_i , with parameter $\boldsymbol{\theta}$; and

$$v_i(\boldsymbol{\phi}, \alpha_i) = F_\alpha(\alpha_i | \mathbf{z}_i, \boldsymbol{\phi}) \tag{11}$$

is the cumulative distribution of α_i , given \mathbf{z}_i , with parameter $\boldsymbol{\phi}$.

In practice, the vector $\boldsymbol{\alpha}_N$ of individual effects is not known and must be estimated. We propose estimating $(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\alpha}_N)$ by solving

$$\min_{\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\alpha}_N} L_1(\boldsymbol{\theta}, \boldsymbol{\alpha}_N) + L_2(\boldsymbol{\phi}, \boldsymbol{\alpha}_N). \tag{12}$$

The proposed loss function is similar to that of a penalized fixed-effects estimator in which $L_2(\boldsymbol{\phi}, \boldsymbol{\alpha}_N)$ plays the role of a penalty term. Intuitively, $L_2(\boldsymbol{\phi}, \boldsymbol{\alpha}_N)$ shrinks the estimated

³ The expression for $L_1(\boldsymbol{\theta}, \boldsymbol{\alpha}_N)$ bears some similarity to Koenker's (2004) loss function for unpenalized fixed-effects quantile regression, which is defined by $L(\boldsymbol{\beta}, \boldsymbol{\alpha}_N) = \sum_j \sum_i \sum_t w_j \rho_{u_j}(y_{it} - \alpha_i - \mathbf{x}_{it}^\top \boldsymbol{\beta}(u_j))$ and can be seen as a discretized, non-parametrized, and weighted version of $L_1(\boldsymbol{\theta}, \boldsymbol{\alpha}_N)$.

fixed effects towards their assumed conditional distribution, introducing some degree of smoothing, improving model identification and efficiency, and avoiding overfitting. At the same time, $L_2(\boldsymbol{\phi}, \boldsymbol{\alpha}_N)$ carries information on the parameter $\boldsymbol{\phi}$ that describes the quantile function of $\boldsymbol{\alpha}_N$.

Since both $\boldsymbol{\alpha}_N$ and $\boldsymbol{\phi}$ are treated as parameters, this approach combines features of fixed-effects estimators, which only estimate $\boldsymbol{\theta}$ and $\boldsymbol{\alpha}_N$, and random-effects models, which directly estimate $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$. Computation, however, is much simpler than that of purely random-effects methods (e.g., Kim and Yang, 2011; Arellano and Bonhomme, 2016).

The gradient functions of $L(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\alpha}_N) = L_1(\boldsymbol{\theta}, \boldsymbol{\alpha}_N) + L_2(\boldsymbol{\phi}, \boldsymbol{\alpha}_N)$ can be written as

$$\mathbf{G}_{\boldsymbol{\theta}}(\boldsymbol{\theta}, \boldsymbol{\alpha}_N) = \nabla_{\text{vec}(\boldsymbol{\theta})} L(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\alpha}_N) = \sum_{i=1}^N \sum_{t=1}^T [\bar{\mathbf{B}} - \mathbf{B}(u_{it}(\boldsymbol{\theta}, \alpha_i))] \otimes \mathbf{x}_{it}, \quad (13)$$

$$\mathbf{G}_{\boldsymbol{\phi}}(\boldsymbol{\phi}, \boldsymbol{\alpha}_N) = \nabla_{\text{vec}(\boldsymbol{\phi})} L(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\alpha}_N) = \sum_{i=1}^N [\bar{\mathbf{C}} - \mathbf{C}(v_i(\boldsymbol{\phi}, \alpha_i))] \otimes \mathbf{z}_i, \quad (14)$$

$$G_{\alpha_i}(\alpha_i, \boldsymbol{\theta}, \boldsymbol{\phi}) = \nabla_{\alpha_i} L(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\alpha}_N) = \left[\sum_{t=1}^T (0.5 - u_{it}(\boldsymbol{\theta}, \alpha_i)) \right] + (v_i(\boldsymbol{\phi}, \alpha_i) - 0.5), \quad (15)$$

where vec denotes the vectorization operator and \otimes the kronecker product. The model parameters, $(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\alpha}_N)$, only enter equations (13)–(15) through the cumulative distribution functions $u_{it}(\boldsymbol{\theta}, \alpha_i)$ and $v_i(\boldsymbol{\phi}, \alpha_i)$ defined in (10) and (11). Note that $\mathbf{G}_{\boldsymbol{\theta}}(\boldsymbol{\theta}, \boldsymbol{\alpha}_N)$ does not carry information on $\boldsymbol{\phi}$, and $\mathbf{G}_{\boldsymbol{\phi}}(\boldsymbol{\phi}, \boldsymbol{\alpha}_N)$ does not carry information on $\boldsymbol{\theta}$; while $G_{\alpha_i}(\alpha_i, \boldsymbol{\theta}, \boldsymbol{\phi})$ depends on both $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$. As shown by Frumento and Bottai (2016), $\mathbf{G}_{\boldsymbol{\theta}}(\boldsymbol{\theta}, \boldsymbol{\alpha}_N)$ and $\mathbf{G}_{\boldsymbol{\phi}}(\boldsymbol{\phi}, \boldsymbol{\alpha}_N)$ approach zero when the distributions of $u_{it}(\boldsymbol{\theta}, \alpha_i)$ and $v_i(\boldsymbol{\phi}, \alpha_i)$ tend to be uniform. This reflects the data-generating process described in (1), which involves the two independent uniform variables U_{it} and V_i .

Equation (15) clarifies the role of the “penalty” term $L_2(\boldsymbol{\phi}, \boldsymbol{\alpha}_N)$:

- the left-hand side of (15), $\sum_{t=1}^T (0.5 - u_{it}(\boldsymbol{\theta}, \alpha_i)) = \nabla_{\alpha_i} L_1(\boldsymbol{\theta}, \boldsymbol{\alpha}_N)$, is an unpenalized estimating equation for α_i . It approaches zero when $u_{i1}(\boldsymbol{\theta}, \alpha_i), u_{i2}(\boldsymbol{\theta}, \alpha_i), \dots, u_{iT}(\boldsymbol{\theta}, \alpha_i)$ are evenly spaced in $(0, 1)$, imposing a within-cluster uniformity of $u_{it}(\boldsymbol{\theta}, \alpha_i)$ which mirrors the assumption of independence between U_{it} and V_i ;
- the right-hand side, $(v_i(\boldsymbol{\phi}, \alpha_i) - 0.5) = \nabla_{\alpha_i} L_2(\boldsymbol{\phi}, \boldsymbol{\alpha}_N)$, is a penalty term that shrinks the value of α_i towards its conditional median, $\mathbf{z}_i^T \boldsymbol{\gamma}(0.5 \mid \boldsymbol{\phi}) = \mathbf{z}_i^T \boldsymbol{\phi} \mathbf{c}(0.5)$.

A desirable property of the proposed penalization is that it only affects the estimates of $\boldsymbol{\alpha}_N$ when the clusters are relatively small. As $T \rightarrow \infty$, each cluster contains sufficient information to estimate its own individual effect and, consistently, the penalty term $(v_i(\boldsymbol{\phi}, \alpha_i) - 0.5)$ in equation (15) becomes irrelevant.

Estimation can be performed by the following iterative process: (i) given $\boldsymbol{\alpha}_N$, estimate $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ separately by solving $\mathbf{G}_{\boldsymbol{\theta}}(\boldsymbol{\theta}, \boldsymbol{\alpha}_N) = 0$ and $\mathbf{G}_{\boldsymbol{\phi}}(\boldsymbol{\phi}, \boldsymbol{\alpha}_N) = 0$; (ii) given $(\boldsymbol{\theta}, \boldsymbol{\phi})$, compute a new estimate of $\boldsymbol{\alpha}_N$ by solving $G_{\alpha_i}(\alpha_i, \boldsymbol{\theta}, \boldsymbol{\phi}) = 0$, $i = 1, \dots, N$. Step (i) can be implemented with standard routines available in the `qrqm` package, while step (ii) requires finding the zero of N univariate estimating equations. Neither $u_{it}(\boldsymbol{\theta}, \alpha_i)$ nor $v_i(\boldsymbol{\phi}, \alpha_i)$ are generally available in closed form, and can be evaluated by using a bisection algorithm. Note that the objective function defined by (12) is a smooth function of all parameters, unlike the loss function of standard quantile regression.

The fact that the quantile function may be ill-defined at some value of the parameters can be an issue during estimation. In the implementation of the `qrqm` package, we use unconstrained optimization from carefully chosen initial values. The algorithm is described in detail in Appendix B.

4.1 A new family of penalized fixed-effects estimators

A possible interpretation of the proposed loss function,

$$L(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\alpha}_N) = L_1(\boldsymbol{\theta}, \boldsymbol{\alpha}_N) + L_2(\boldsymbol{\phi}, \boldsymbol{\alpha}_N),$$

is to consider $L_2(\boldsymbol{\phi}, \boldsymbol{\alpha}_N)$ as a penalty term that shrinks the estimated individual effects towards their conditional median, $\mathbf{z}_i^T \boldsymbol{\gamma}(0.5 \mid \boldsymbol{\phi})$. Unlike standard penalizations, however, $L_2(\boldsymbol{\phi}, \boldsymbol{\alpha}_N)$ may depend on level-2 covariates and is a function of estimated parameters.

To clarify this idea, consider a more traditional penalized loss function

$$L_\lambda(\boldsymbol{\theta}, \boldsymbol{\alpha}_N) = L_1(\boldsymbol{\theta}, \boldsymbol{\alpha}_N) + \lambda L_2(\boldsymbol{\alpha}_N),$$

where $L_2(\boldsymbol{\alpha}_N)$ is a penalty term which does not contain $\boldsymbol{\phi}$, and λ is a tuning parameter. Common choices of $L_2(\boldsymbol{\alpha}_N)$ are the ℓ_1 -penalization, $L_2(\boldsymbol{\alpha}_N) = \sum_{i=1}^N |\alpha_i|$, which was used by Koenker (2004) to implement longitudinal quantile regression, and the ℓ_2 -penalization, $L_2(\boldsymbol{\alpha}_N) = \sum_{i=1}^N \alpha_i^2$.

Standard ℓ_1 - and ℓ_2 -penalized fixed-effects methods are computationally simple and can substantially improve efficiency of the estimates of the structural parameters. However, besides the fact that they do not allow for estimation of $\boldsymbol{\phi}$, they present some important limitations: (i) they do not use prior knowledge on the distribution of the individual effects; (ii) they can introduce bias; (iii) they apply the same penalization to all clusters; and (iv) they require to specify a tuning parameter.

For instance, ℓ_1 -penalized estimators of quantile regression coefficients are asymptotically biased unless $(\alpha_1, \dots, \alpha_N)$ are independent and identically distributed with zero median (Lamarche, 2010). This is just a consequence of the ℓ_1 -penalty term being a sum of absolute deviations from zero, which does not generally reflect the true distribution of $\boldsymbol{\alpha}_N$ and the effect of level-2 covariates on it. Moreover, the same value of λ is used for

all clusters, ignoring the fact that the variance of the individual effects may differ across subgroups of the population.

The tuning constant λ determines the degree of shrinking and, in the standard random-intercept linear model, its optimal value is $\sigma_\epsilon^2/\sigma_\alpha^2$, i.e., a function of nuisance scale parameters (e.g., Koenker, 2004). Outside the restrictive conditions of linear models, not only the choice of λ becomes problematic, but also the use of a single value of λ for all clusters is questionable.

The novelty of our approach is that, unlike the ℓ_1 - and ℓ_2 -penalizations, the term $L_2(\boldsymbol{\phi}, \boldsymbol{\alpha}_N)$ reflects the true (conditional) distribution of $\boldsymbol{\alpha}_N$ and carries information about its parameters, $\boldsymbol{\phi}$. Our estimator presents the following advantages over standard penalized methods: (i) it enables incorporating parametric assumptions on the distribution of $\boldsymbol{\alpha}_N$; (ii) it permits estimating all parameters consistently; (iii) it applies a different degree of shrinking to each cluster, by modeling the effect of level-2 covariates on the distribution of the individual effects; and (iv) it does not require selecting a tuning constant, as no nuisance parameters are present.

To clarify point (iv), consider the loss function of an ℓ_2 -penalized linear regression model: $L_\lambda(\boldsymbol{\beta}, \boldsymbol{\alpha}_N) = L_1(\boldsymbol{\beta}, \boldsymbol{\alpha}_N) + \lambda L_2(\boldsymbol{\alpha}_N) = \sum_{i=1}^N \sum_{t=1}^T (y_{it} - \mathbf{x}_i^T \boldsymbol{\beta} - \alpha_i)^2 + \lambda \sum_{i=1}^N \alpha_i^2$. Here, $L_1(\boldsymbol{\beta}, \boldsymbol{\alpha}_N)$ and $L_2(\boldsymbol{\alpha}_N)$ lack information on the nuisance scale parameters $\sigma_\epsilon^2 = \text{var}(Y_{it} - \mathbf{x}_i^T \boldsymbol{\beta} - \alpha_i)$ and $\sigma_\alpha^2 = \text{var}(\alpha_i)$. This is adjusted for by the tuning constant $\lambda = \sigma_\epsilon^2/\sigma_\alpha^2$. In our special type of penalized estimator, instead, $L_1(\boldsymbol{\theta}, \boldsymbol{\alpha}_N)$ and $L_2(\boldsymbol{\phi}, \boldsymbol{\alpha}_N)$ carry information on *all* model parameters. Intuitively, this means that $L_1(\boldsymbol{\theta}, \boldsymbol{\alpha}_N)$ and $L_2(\boldsymbol{\phi}, \boldsymbol{\alpha}_N)$ are already “properly scaled”. The tuning constant can be thought of as an implicit parameter, a function of $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$. Although a more general estimator with criterion function $L_1(\boldsymbol{\theta}, \boldsymbol{\alpha}_N) + \lambda L_2(\boldsymbol{\phi}, \boldsymbol{\alpha}_N)$ could in principle be formulated, choosing $\lambda = 1$ appears natural and avoids the problem of selecting the tuning parameter.

5 INFERENCE

The asymptotic properties of fixed-effects estimators are complicated by the fact that, as $N \rightarrow \infty$, the dimension of the parameter $\boldsymbol{\alpha}_N$ tends to infinity. Unless $T \rightarrow \infty$, the individual effects α_i are estimated using a fixed number of observations. This is often referred to as the “incidental parameter” problem (Neyman and Scott, 1948; Lancaster, 2000), which causes widely used estimators, such as maximum likelihood and M-estimators, to be inconsistent.

To develop the asymptotic theory of our estimator, we follow the recent panel data literature in econometrics and deal with the incidental parameter problem by considering asymptotic sequences where both N and T tend to infinity (e.g., Phillips & Moon, 1999, Hahn and Newey, 2004, Koenker, 2004, Fernández-Val, 2005, Arellano & Hahn, 2007, Lamarche, 2010, and Kato et al., 2012). Under this approximation, we show that our estimators are consistent but might have biases in the asymptotic distribution depending on the relative rate of convergence of N and T . We apply the theory of M-estimators (e.g., Newey and McFadden, 1994), and use well-established results to handle the following non-standard features of our problem: (i) the estimators of $\boldsymbol{\theta}$, $\boldsymbol{\phi}$ and $\boldsymbol{\alpha}_N$ converge at different rates (e.g., Radchenko, 2008; Cheng and Shang, 2015; Masuda and Shimizu, 2017); and (ii) additional conditions are required on the relative growth rate of N and T (e.g., Hahn and Newey, 2004; Fernández-Val, 2005; Newey, 2007).

Let $\boldsymbol{x}_{it} = (\boldsymbol{x}_{1i}^T, \boldsymbol{x}_{2it}^T)^T$, where \boldsymbol{x}_{1i} contains the time-invariant components including the constant and \boldsymbol{x}_{2it} contains the time-varying covariates. We use the following sufficient conditions to establish the identification of the parameters and derive the asymptotic properties of the estimators:

Assumption 1 (Longitudinal ILM Estimator) (i) *The data generating process is $Y_{it} =$*

$\mathbf{x}_{it}^\top \boldsymbol{\theta}^0 \mathbf{b}(U_{it}) + \mathbf{z}_i^\top \boldsymbol{\phi}^0 \mathbf{c}(V_i)$, where $\text{plim}_{N \rightarrow \infty} N^{-1} \sum_{i=1}^N \alpha_i^0 = 0$ for $\alpha_i^0 := \mathbf{z}_i^\top \boldsymbol{\phi}^0 \mathbf{c}(V_i)$ and, conditional on $\{(\mathbf{x}_{it}, \mathbf{z}_i) : 1 \leq i \leq N, 1 \leq t \leq T\}$, $U_{it} \sim U(0, 1)$ independently over i and t , $V_i \sim U(0, 1)$ independently over i , and U_{it} and V_j are independent over all i, t, j . (ii) For each i , $\text{plim}_{T \rightarrow \infty} T^{-1} \sum_{t=1}^T \bar{\mathbf{x}}_{2it} \bar{\mathbf{x}}_{2it}^\top$ exists and is positive definite for $\bar{\mathbf{x}}_{2it} = (1, \mathbf{x}_{2it}^\top)^\top$, and $\text{plim}_{N \rightarrow \infty} N^{-1} \sum_{i=1}^N \mathbf{x}_{1i} \mathbf{x}_{1i}^\top$ and $\text{plim}_{N \rightarrow \infty} N^{-1} \sum_{i=1}^N \mathbf{z}_i \mathbf{z}_i^\top$ exist and are positive definite. (iii) The variables $\mathbf{x}_{it}^\top \boldsymbol{\theta}^0 \mathbf{b}(U_{it})$ and $\mathbf{z}_i^\top \boldsymbol{\phi}^0 \mathbf{c}(V_i)$ are finite a.s., and $\mathbf{x}_{it}^\top \boldsymbol{\theta}^0 \mathbf{b}'(U_{it})$ and $\mathbf{z}_i^\top \boldsymbol{\phi}^0 \mathbf{c}'(V_i)$ are positive a.s. (iv) There exist two sets of quantile indexes $\{u_1, \dots, u_{d_b}\}$ and $\{v_1, \dots, v_{d_c}\}$ such that the matrices $\mathbf{b}(u_1, \dots, u_{d_b}) := [\mathbf{b}(u_1), \dots, \mathbf{b}(u_{d_b})]$ and $\mathbf{c}(v_1, \dots, v_{d_c}) := [\mathbf{c}(v_1), \dots, \mathbf{c}(v_{d_c})]$ have full rank. (v) The functions $u \mapsto \mathbf{b}(u)$ and $v \mapsto \mathbf{c}(v)$ are three times continuously differentiable on $(0, 1)$, $\sup_i \text{plim}_{T \rightarrow \infty} T^{-1} \sum_{t=1}^T \|\mathbf{x}_{2it}\|^\xi < \infty$ for some $\xi > 32$, and $\|(\mathbf{x}_{1i}^\top, \mathbf{z}_i^\top)^\top\|$ is bounded a.s. (vi) The following probability limits exist:

$$\begin{aligned} \bar{\mathbf{H}}_\theta &= \text{plim}_{N, T \rightarrow \infty} \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \mathbf{E}(\mathbf{H}_{\theta it}), \\ \bar{\mathbf{H}}_\phi &= \text{plim}_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \mathbf{E} \left(\frac{\mathbf{c}(V_i) \mathbf{c}(V_i)^\top}{\mathbf{z}_i^\top \boldsymbol{\phi}^0 \mathbf{c}'(V_i)} \otimes \mathbf{z}_i \mathbf{z}_i^\top \right), \\ \bar{\mathbf{H}}_{\phi\theta} &= \text{plim}_{N, T \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \mathbf{E} \left(\frac{\mathbf{c}(V_i)}{\mathbf{z}_i^\top \boldsymbol{\phi}^0 \mathbf{c}'(V_i)} \otimes \mathbf{z}_i \right) \frac{\sigma_i^2}{T} \sum_{t=1}^T \mathbf{E} \left(\frac{\mathbf{b}(U_{it})}{\mathbf{x}_{it}^\top \boldsymbol{\theta}^0 \mathbf{b}'(U_{it})} \otimes \mathbf{x}_{it} \right)^\top, \\ \bar{\mathbf{b}}_\theta &= \text{plim}_{N, T \rightarrow \infty} \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \mathbf{E}(\mathbf{b}_{\theta it}), \quad \bar{\mathbf{b}}_\phi = \text{plim}_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \mathbf{E}(\mathbf{b}_{\phi i}), \\ \bar{\mathbf{Q}}_\theta &= \text{plim}_{N, T \rightarrow \infty} \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \mathbf{E}(\boldsymbol{\varphi}_{\theta it} \boldsymbol{\varphi}_{\theta it}^\top), \\ \bar{\mathbf{Q}}_\phi &= \text{plim}_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \mathbf{E} \left\{ [\bar{\mathbf{C}} - \mathbf{C}(V_i)] [\bar{\mathbf{C}} - \mathbf{C}(V_i)]^\top \otimes \mathbf{z}_i \mathbf{z}_i^\top \right\}, \end{aligned}$$

where the expectation \mathbf{E} is taken with respect to the distribution of U_{it} and V_i , \otimes denotes the

Kronecker product, \mathbf{c}' and \mathbf{c}'' denote the vectors of first and second derivatives of $v \mapsto \mathbf{c}(v)$,

$$\begin{aligned}
\mathbf{H}_{\theta it} &= \frac{\mathbf{b}(U_{it})\mathbf{b}(U_{it})^\top}{\mathbf{x}_{it}^\top \boldsymbol{\theta}^0 \mathbf{b}'(U_{it})} \otimes \mathbf{x}_{it} \mathbf{x}_{it}^\top - \left(\frac{\mathbf{b}(U_{it})}{\mathbf{x}_{it}^\top \boldsymbol{\theta}^0 \mathbf{b}'(U_{it})} \otimes \mathbf{x}_{it} \right) \frac{\sigma_i^2}{T} \sum_{t=1}^T \mathbb{E} \left(\frac{\mathbf{b}(U_{it})}{\mathbf{x}_{it}^\top \boldsymbol{\theta}^0 \mathbf{b}'(U_{it})} \otimes \mathbf{x}_{it} \right)^\top, \\
\mathbf{b}_{\phi i} &= \left\{ \frac{[\sigma_i^2(V_i - 0.5) - \beta_i] \mathbf{c}(V_i)}{\mathbf{z}_i^\top \boldsymbol{\phi}^0 \mathbf{c}'(V_i)} + \frac{\sigma_i^4 \mathbf{c}(V_i) \mathbf{z}_i^\top \boldsymbol{\phi}^0 \mathbf{c}''(V_i)}{24[\mathbf{z}_i^\top \boldsymbol{\phi}^0 \mathbf{c}'(V_i)]^3} - \frac{\sigma_i^4 \mathbf{c}'(V_i)}{24[\mathbf{z}_i^\top \boldsymbol{\phi}^0 \mathbf{c}'(V_i)]^2} \right\} \otimes \mathbf{z}_i, \\
\mathbf{b}_{\theta it} &= \left\{ \frac{[\sigma_i^2(U_{it} - 0.5) + \beta_i] \mathbf{b}(U_{it})}{\mathbf{x}_{it}^\top \boldsymbol{\theta}^0 \mathbf{b}'(U_{it})} - \frac{\sigma_i^4 \mathbf{b}(U_{it}) \mathbf{x}_{it}^\top \boldsymbol{\theta}^0 \mathbf{b}''(U_{it})}{24[\mathbf{x}_{it}^\top \boldsymbol{\theta}^0 \mathbf{b}'(U_{it})]^3} - \frac{\sigma_i^4 \mathbf{b}'(U_{it})}{24[\mathbf{x}_{it}^\top \boldsymbol{\theta}^0 \mathbf{b}'(U_{it})]^2} \right\} \otimes \mathbf{x}_{it}, \\
\boldsymbol{\varphi}_{\theta it} &= [\bar{\mathbf{B}} - \mathbf{B}(U_{it})] \otimes \mathbf{x}_{it} + \frac{\sigma_i^2}{T} \sum_{s=1}^T \mathbb{E} \left(\frac{\mathbf{b}(U_{is})}{\mathbf{x}_{is}^\top \boldsymbol{\theta}^0 \mathbf{b}'(U_{is})} \otimes \mathbf{x}_{is} \right) (U_{it} - 0.5), \\
\beta_i &= \sigma_i^4 \text{plim}_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbb{E} \left(\frac{U_{it} - 0.5}{\mathbf{x}_{it}^\top \boldsymbol{\theta}^0 \mathbf{b}'(U_{it})} - \frac{\sigma_i^2 \mathbf{x}_{it}^\top \boldsymbol{\theta}^0 \mathbf{b}''(U_{it})}{24[\mathbf{x}_{it}^\top \boldsymbol{\theta}^0 \mathbf{b}'(U_{it})]^3} \right), \\
\sigma_i^2 &= \left[\text{plim}_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbb{E} \left(\frac{1}{\mathbf{x}_{it}^\top \boldsymbol{\theta}^0 \mathbf{b}'(U_{it})} \right) \right]^{-1},
\end{aligned}$$

and \mathbf{b}' and \mathbf{b}'' denote the vectors of first and second derivatives of $u \mapsto \mathbf{b}(u)$. (vii) The minimum eigenvalues of the matrices $\bar{\mathbf{H}}_{\boldsymbol{\theta}}$ and $\bar{\mathbf{H}}_{\boldsymbol{\phi}}$ are bounded away from zero, and $\sup_i \sigma_i^2 < c$ for some constant $c > 0$.

We use Assumptions 1(i)-(iv) to establish the identification of all the model parameters. Assumption 1(i) imposes that the model is correctly specified. It also requires α_i and $Y_{it} - \alpha_i$ to be conditionally independent across (i, t) , and to be independent of each other. We do not impose any sampling condition on the covariate sequences $\{(\mathbf{x}_{it}, \mathbf{z}_i) : 1 \leq i \leq N, 1 \leq t \leq T\}$, other than the existence of some limits. Assumptions 1(ii)-(iii) apply standard regularity conditions for parameter identification in quantile regression to our longitudinal model (e.g., Angrist et al., 2006). For example, Assumption 1(iii) imposes that the conditional quantile and density functions of $Y_{it} - \alpha_i$ and α_i are bounded. These conditions, together with a location normalization on the fixed effects in Assumption 1(i),

guarantee that $\beta(\cdot)$ and $\gamma(\cdot)$ in the model (1) are identified.⁴ Then, Assumption 1(iv) pins down θ^0 and ϕ^0 from the system of linear equations $\beta(\cdot) = \theta^0 \mathbf{b}(\cdot)$ and $\gamma(\cdot) = \phi^0 \mathbf{c}(\cdot)$. Assumption 1(iv) provides a sufficient condition to guarantee existence and uniqueness of solution to the system from a subset of equations, which is easy to verify in practice. It can be replaced by any other existence and uniqueness condition.

Assumptions 1(v)-(vii) impose regularity conditions to derive the distribution of the estimators in large samples. The derivation relies on a general asymptotic expansion for fixed effects M-estimators given in Appendix A, which extend the results of Hahn and Newey (2004) and Fernández-Val (2005) to estimators with mixed-rates asymptotics. Assumption 1(v) requires sufficient smoothness and bounded moments of the objective functions (6) and (7) and their partial derivatives, which are needed to carry out higher-order expansions of these functions. Assumption 1(vi) guarantees that all the terms of the expansions are well-defined. Finally, Assumption 1(vii) is a standard condition imposing that the limit Hessian matrices of the objective functions are non-singular.

Theorem 1 *Suppose that Assumption 1 holds. Then, (i) θ^0 and ϕ^0 are identified. (ii) If $N = O(T)$,*

$$\sqrt{NT} \left(\text{vec} \left[\hat{\theta} - \theta^0 \right] + \frac{\bar{\mathbf{H}}_{\theta}^{-1} \bar{\mathbf{b}}_{\theta}}{T} \right) \rightarrow_d \bar{\mathbf{H}}_{\theta}^{-1} N(0, \bar{\mathcal{O}}_{\theta}).$$

(iii) If $N = O(T^2)$,

$$\sqrt{N} \left(\text{vec} \left[\hat{\phi} - \phi^0 \right] + \frac{\bar{\mathbf{H}}_{\phi}^{-1} (\bar{\mathbf{b}}_{\phi} + \bar{\mathbf{H}}_{\phi\theta} \bar{\mathbf{H}}_{\theta}^{-1} \bar{\mathbf{b}}_{\theta})}{T} \right) \rightarrow_d \bar{\mathbf{H}}_{\phi}^{-1} N(0, \bar{\mathcal{O}}_{\phi}).$$

The expressions of all the terms are given in Assumption 1.

⁴We normalize the mean of the fixed effects. Alternative normalizations on the median or other quantile of the fixed effects are also possible.

Theorem 1 shows that the parameters $\boldsymbol{\theta}^0$ and $\boldsymbol{\phi}^0$ are identified and their estimators $\hat{\boldsymbol{\theta}}$ and $\hat{\boldsymbol{\phi}}$ have a normal distribution in large samples with different rates of convergence. The large sample distribution of the plugin estimators of $\boldsymbol{\beta}^0(\cdot) = \boldsymbol{\theta}^0 \mathbf{b}(\cdot)$ and $\boldsymbol{\gamma}^0(\cdot) = \boldsymbol{\phi}^0 \mathbf{c}(\cdot)$ can be obtained by the delta method. Let $\hat{\boldsymbol{\beta}}(u) = \hat{\boldsymbol{\theta}} \mathbf{b}(u)$ and $\hat{\boldsymbol{\gamma}}(v) = \hat{\boldsymbol{\phi}} \mathbf{c}(v)$, for any $u, v \in (0, 1)$. Then, if $N = O(T)$,

$$\sqrt{NT} \left(\hat{\boldsymbol{\beta}}(u) - \boldsymbol{\beta}^0(u) + \frac{\text{vec}_{d_x, d_b}^{-1}(\bar{\mathbf{H}}_{\boldsymbol{\theta}}^{-1} \bar{\mathbf{b}}_{\boldsymbol{\theta}}) \mathbf{b}(u)}{T} \right) \rightarrow_d \text{N}(0, \tilde{\mathbf{b}}(u)^\top \bar{\mathbf{H}}_{\boldsymbol{\theta}}^{-1} \bar{\mathcal{O}}_{\boldsymbol{\theta}} \bar{\mathbf{H}}_{\boldsymbol{\theta}}^{-1} \tilde{\mathbf{b}}(u)),$$

where $\text{vec}_{d,k}^{-1}$ is the inverse vectorization operator that maps a dk -vector to a $d \times k$ matrix, i.e. $\text{vec}_{d,k}^{-1}(\mathbf{v}) = \{[\text{vec}(\mathbf{I}_k)]^\top \otimes \mathbf{I}_d\}(\mathbf{I}_k \otimes \mathbf{v})$, \mathbf{I}_n is the identity matrix of size n , and $\tilde{\mathbf{b}}(u) = \mathbf{b}(u) \otimes \mathbf{I}_{d_x}$. Similarly, if $N = O(T^2)$,

$$\sqrt{N} \left(\hat{\boldsymbol{\gamma}}(v) - \boldsymbol{\gamma}^0(v) + \frac{\text{vec}_{d_z, d_c}^{-1}(\bar{\mathbf{H}}_{\boldsymbol{\phi}}^{-1}(\bar{\mathbf{b}}_{\boldsymbol{\phi}} + \bar{\mathbf{H}}_{\boldsymbol{\phi}\boldsymbol{\theta}} \bar{\mathbf{H}}_{\boldsymbol{\theta}}^{-1} \bar{\mathbf{b}}_{\boldsymbol{\theta}})) \mathbf{c}(v)}{T} \right) \rightarrow_d \text{N}(0, \tilde{\mathbf{c}}(v)^\top \bar{\mathbf{H}}_{\boldsymbol{\phi}}^{-1} \bar{\mathcal{O}}_{\boldsymbol{\phi}} \bar{\mathbf{H}}_{\boldsymbol{\phi}}^{-1} \tilde{\mathbf{c}}(v)),$$

where $\tilde{\mathbf{c}}(v) = \mathbf{c}(v) \otimes \mathbf{I}_{d_z}$.

The rates of convergence of all the estimators agree with the square roots of the dimensions of the observations that are informative about the corresponding parameters. Thus, the rate is \sqrt{NT} for $\boldsymbol{\theta}^0$ and $\boldsymbol{\beta}^0(u)$, and \sqrt{N} for $\boldsymbol{\phi}^0$ and $\boldsymbol{\gamma}^0(v)$. All the estimators might suffer from bias in short panels due to the estimation of the fixed effects. The order of this bias is the inverse of the number of observations that are informative about each fixed effect, i.e. T^{-1} . Comparing the rates of convergence with the order of the bias, we can see that the biases of $\hat{\boldsymbol{\theta}}$ and $\hat{\boldsymbol{\beta}}(u)$ are negligible in the asymptotic distribution when $N/T \rightarrow 0$, whereas the biases of $\hat{\boldsymbol{\phi}}$ and $\hat{\boldsymbol{\gamma}}(v)$ are negligible when $N/T^2 \rightarrow 0$. These biases can be reduced by using analytical or jackknife corrections (e.g., Hahn and Newey, 2004, Fernández-Val, 2005, or Dhaene & Jochmans, 2015). We provide consistent analytical estimators of the components of the biases and variances below.

We construct estimators of the components of the asymptotic distribution using sample analogs evaluated at the estimated value of the parameters, e.g., $\hat{u}_{it} = u_{it}(\hat{\boldsymbol{\theta}}, \hat{\alpha}_i)$ and $\hat{v}_i = v_i(\hat{\boldsymbol{\phi}}, \hat{\alpha}_i)$. Then,

$$\begin{aligned}\hat{\mathbf{H}}_{\boldsymbol{\theta}} &= \sum_{i=1}^N \sum_{t=1}^T \hat{\mathbf{H}}_{\boldsymbol{\theta}it}, \quad \hat{\mathbf{H}}_{\boldsymbol{\phi}} = \frac{1}{N} \sum_{i=1}^N \left(\frac{\mathbf{c}(\hat{v}_i) \mathbf{c}(\hat{v}_i)^{\top}}{\mathbf{z}_i^{\top} \hat{\boldsymbol{\phi}}^0 \mathbf{c}'(\hat{v}_i)} \otimes \mathbf{z}_i \mathbf{z}_i^{\top} \right), \\ \hat{\mathbf{H}}_{\boldsymbol{\phi}\boldsymbol{\theta}} &= \frac{1}{N} \sum_{i=1}^N \left(\frac{\mathbf{c}(\hat{v}_i)}{\mathbf{z}_i^{\top} \hat{\boldsymbol{\phi}} \mathbf{c}'(\hat{v}_i)} \otimes \mathbf{z}_i \right) \frac{\hat{\sigma}_i^2}{T} \sum_{t=1}^T \left(\frac{\mathbf{b}(\hat{u}_{it})}{\mathbf{x}_{it}^{\top} \hat{\boldsymbol{\theta}} \mathbf{b}'(\hat{u}_{it})} \otimes \mathbf{x}_{it} \right)^{\top}, \\ \hat{\mathbf{b}}_{\boldsymbol{\theta}} &= \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \hat{\mathbf{b}}_{\boldsymbol{\theta}it}, \quad \hat{\mathbf{b}}_{\boldsymbol{\phi}} = \frac{1}{N} \sum_{i=1}^N \hat{\mathbf{b}}_{\boldsymbol{\phi}i}, \\ \hat{\boldsymbol{\varphi}}_{\boldsymbol{\theta}} &= \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \hat{\boldsymbol{\varphi}}_{\boldsymbol{\theta}it} \hat{\boldsymbol{\varphi}}_{\boldsymbol{\theta}it}^{\top}, \quad \hat{\boldsymbol{\varphi}}_{\boldsymbol{\phi}} = \frac{1}{N} \sum_{i=1}^N \left\{ [\bar{\mathbf{C}} - \mathbf{C}(\hat{v}_i)] [\bar{\mathbf{C}} - \mathbf{C}(\hat{v}_i)]^{\top} \otimes \mathbf{z}_i \mathbf{z}_i^{\top} \right\},\end{aligned}$$

where

$$\begin{aligned}\hat{\mathbf{H}}_{\boldsymbol{\theta}it} &= \frac{\mathbf{b}(\hat{u}_{it}) \mathbf{b}(\hat{u}_{it})^{\top}}{\mathbf{x}_{it}^{\top} \hat{\boldsymbol{\theta}} \mathbf{b}'(\hat{u}_{it})} \otimes \mathbf{x}_{it} \mathbf{x}_{it}^{\top} - \left(\frac{\mathbf{b}(\hat{u}_{it})}{\mathbf{x}_{it}^{\top} \hat{\boldsymbol{\theta}} \mathbf{b}'(\hat{u}_{it})} \otimes \mathbf{x}_{it} \right) \frac{\hat{\sigma}_i^2}{T} \sum_{t=1}^T \left(\frac{\mathbf{b}(\hat{u}_{it})}{\mathbf{x}_{it}^{\top} \hat{\boldsymbol{\theta}} \mathbf{b}'(\hat{u}_{it})} \otimes \mathbf{x}_{it} \right)^{\top}, \\ \hat{\mathbf{H}}_{\boldsymbol{\phi}i} &= \frac{\mathbf{c}(\hat{v}_i) \mathbf{c}(\hat{v}_i)^{\top}}{\mathbf{z}_i^{\top} \hat{\boldsymbol{\phi}}^0 \mathbf{c}'(\hat{v}_i)} \otimes \mathbf{z}_i \mathbf{z}_i^{\top} - \left(\frac{\mathbf{c}(\hat{v}_i)}{\mathbf{z}_i^{\top} \hat{\boldsymbol{\phi}} \mathbf{c}'(\hat{v}_i)} \otimes \mathbf{z}_i \right) \frac{\hat{\sigma}_i^2}{T} \left(\frac{\mathbf{c}(\hat{v}_i)}{\mathbf{z}_i^{\top} \hat{\boldsymbol{\phi}} \mathbf{c}'(\hat{v}_i)} \otimes \mathbf{z}_i \right)^{\top}, \\ \hat{\mathbf{b}}_{\boldsymbol{\theta}it} &= \left\{ \frac{[\hat{\sigma}_i^2(\hat{u}_{it} - 0.5) + \hat{\beta}_i] \mathbf{b}(\hat{u}_{it})}{\mathbf{x}_{it}^{\top} \hat{\boldsymbol{\theta}} \mathbf{b}'(\hat{u}_{it})} - \frac{\hat{\sigma}_i^4 \mathbf{b}(\hat{u}_{it}) \mathbf{x}_{it}^{\top} \hat{\boldsymbol{\theta}}^0 \mathbf{b}''(\hat{u}_{it})}{24[\mathbf{x}_{it}^{\top} \hat{\boldsymbol{\theta}} \mathbf{b}'(\hat{u}_{it})]^3} - \frac{\hat{\sigma}_i^4 \mathbf{b}'(\hat{u}_{it})}{24[\mathbf{x}_{it}^{\top} \hat{\boldsymbol{\theta}} \mathbf{b}'(\hat{u}_{it})]^2} \right\} \otimes \mathbf{x}_{it}, \\ \hat{\mathbf{b}}_{\boldsymbol{\phi}i} &= \left\{ \frac{[\hat{\sigma}_i^2(\hat{v}_i - 0.5) - \hat{\beta}_i] \mathbf{c}(\hat{v}_i)}{\mathbf{z}_i^{\top} \hat{\boldsymbol{\phi}} \mathbf{c}'(\hat{v}_i)} + \frac{\hat{\sigma}_i^4 \mathbf{c}(\hat{v}_i) \mathbf{z}_i^{\top} \hat{\boldsymbol{\phi}} \mathbf{c}''(\hat{v}_i)}{24[\mathbf{z}_i^{\top} \hat{\boldsymbol{\phi}} \mathbf{c}'(\hat{v}_i)]^3} - \frac{\hat{\sigma}_i^4 \mathbf{c}'(\hat{v}_i)}{24[\mathbf{z}_i^{\top} \hat{\boldsymbol{\phi}}^0 \mathbf{c}'(\hat{v}_i)]^2} \right\} \otimes \mathbf{z}_i, \\ \hat{\boldsymbol{\varphi}}_{\boldsymbol{\theta}it} &= [\bar{\mathbf{B}} - \mathbf{B}(\hat{u}_{it})] \otimes \mathbf{x}_{it} + \frac{\hat{\sigma}_i^2}{T} \sum_{s=1}^T \left(\frac{\mathbf{b}(\hat{u}_{is})}{\mathbf{x}_{is}^{\top} \hat{\boldsymbol{\theta}} \mathbf{b}'(\hat{u}_{is})} \otimes \mathbf{x}_{is} \right) (\hat{u}_{it} - 0.5), \\ \hat{\beta}_i &= \hat{\sigma}_i^4 \frac{1}{T} \sum_{t=1}^T \left(\frac{\hat{u}_{it} - 0.5}{\mathbf{x}_{it}^{\top} \hat{\boldsymbol{\theta}} \mathbf{b}'(\hat{u}_{it})} - \frac{\hat{\sigma}_i^2 \mathbf{x}_{it}^{\top} \hat{\boldsymbol{\theta}} \mathbf{b}''(\hat{u}_{it})}{24[\mathbf{x}_{it}^{\top} \hat{\boldsymbol{\theta}} \mathbf{b}'(\hat{u}_{it})]^3} \right), \quad \hat{\sigma}_i^2 = \left[\frac{1}{T} \sum_{t=1}^T \frac{1}{\mathbf{x}_{it}^{\top} \hat{\boldsymbol{\theta}} \mathbf{b}'(\hat{u}_{it})} \right]^{-1}.\end{aligned}$$

The consistency of these estimators follows from the law of large numbers and consistency

of $\hat{\boldsymbol{\theta}}$, $\hat{\boldsymbol{\phi}}$ and $\hat{\alpha}_i$, $1 \leq i \leq N$, together with the continuous mapping theorem, as all the components are continuous functions of the parameters.

6 GOODNESS-OF-FIT MEASURES AND MODEL SELECTION

6.1 Goodness-of-fit

To assess the model fit, we use the fact that, under the true model, \hat{u}_{it} and \hat{v}_i consistently estimate the realizations of the two independent uniform variables U_{it} and V_i that generated the data.

A graphical inspection of the joint and marginal distributions of \hat{u}_{it} and \hat{v}_i is always recommended. For a formal test, we suggest comparing the empirical distribution of \hat{u}_{it} and \hat{v}_i , $\hat{F}_{\hat{u}\hat{v}}$, with the distribution of U_{it} and V_i , given by

$$F_{uv}(u, v) = uv.$$

Following Frumento and Bottai (2016; 2017), we compute a p-value for the null hypothesis $H_0 : \{the\ model\ is\ correct\}$ using a Monte Carlo procedure:

- *step 0*: compute a test statistic D that measures the distance between $\hat{F}_{\hat{u}\hat{v}}(\hat{u}_{it}, \hat{v}_i)$ and $F_{uv}(\hat{u}_{it}, \hat{v}_i) = \hat{u}_{it}\hat{v}_i$;
- *step 1*: simulate new data as $Y_{it}^* = \mathbf{x}_{it}^T \hat{\boldsymbol{\theta}} \mathbf{b}(U_{it}^*) + \mathbf{z}_i^T \hat{\boldsymbol{\phi}} \mathbf{c}(V_i^*)$ by randomly generating (U_{it}^*, V_i^*) from two independent $U(0, 1)$ distributions;
- *step 2*: fit the model on the simulated data and compute the corresponding value D^* of the test statistic.

After repeating steps 1-2 for a sufficient number of times, the p-value is computed as the empirical proportion of cases in which $D^* > D$. In step 1, it is also possible to take a random sample of clusters, and to resample the covariates' value within each cluster. To assess local fit, the test could be repeated within subsets of the original sample identified by specific values of the covariates.

In the implementation of the `qrqm` package, we chose D to be the Kolmogorov-Smirnov statistic, $\sup_{u,v} |\hat{F}_{\hat{u}\hat{v}}(u, v) - uv|$. This testing procedure is usually reliable, as indicated by the simulation results reported in Section 7.

6.2 Model selection

As suggested in Section 3.5 and exemplified in the real-data example presented in Section 8, it is usually possible to identify numerous alternative models that have a similar fit and are not rejected by a goodness-of-fit test. This can be explained by the fact that the same coefficient functions can be well approximated by different parametric functions.

Important criteria for model selection include parsimony, flexibility, and interpretability. Nested models can be compared by standard Wald test. Let $\hat{\alpha}_N = (\hat{\alpha}_1, \dots, \hat{\alpha}_N)$. To compare non-nested models, the value of $L_1(\hat{\theta}, \hat{\alpha}_N)$ and $L_2(\hat{\phi}, \hat{\alpha}_N)$ can be used to construct information criteria such as the AIC (Akaike, 1974) and the BIC (Schwarz, 1978). These criteria were initially designed for likelihood-based estimators, but can be extended to estimators defined by the minimizer of a loss function. For example, a modification of BIC criterion for M-estimators has been described by Machado (1993), while Koenker (2005) used AIC to compare quantile regression models. Consider the probability density function of the asymmetric Laplace distribution,

$$f_p(y | \mu, \sigma) = \frac{p(1-p)}{\sigma(p)} \exp\left\{-\frac{\rho(y - \mu(p))}{2\sigma(p)}\right\},$$

where $\mu(p)$ is a location parameter and corresponds to the p -th quantile of the distribution, while $\sigma(p)$ is a scale parameter. Although this distribution is not generally considered a plausible model, its log-likelihood has been used by numerous authors, including Koenker and Machado (1999) and Lee, Noh, and Park (2014), to obtain measures of goodness-of-fit for quantile regression. Simple algebra permits showing that

$$\begin{aligned}\hat{\boldsymbol{\theta}} &= \arg \min_{\boldsymbol{\theta}} \int_0^1 \sum_{i=1}^N \sum_{t=1}^T \log f_u(y_{it} - \hat{\alpha}_i \mid \mathbf{x}_{it}^T \boldsymbol{\theta} \mathbf{b}(u), \sigma_1) du, \\ \hat{\boldsymbol{\phi}} &= \arg \min_{\boldsymbol{\phi}} \int_0^1 \sum_{i=1}^N \log f_v(\hat{\alpha}_i \mid \mathbf{z}_i^T \boldsymbol{\phi} \mathbf{c}(v), \sigma_2) dv,\end{aligned}$$

i.e., $\hat{\boldsymbol{\theta}}$ and $\hat{\boldsymbol{\phi}}$ minimize an ‘‘average’’ Laplace log-likelihood, in which u and v have been integrated away. After substituting $\hat{\sigma}_1 = L_1(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\alpha}}_N)/(2NT)$ and $\hat{\sigma}_2 = L_2(\hat{\boldsymbol{\phi}}, \hat{\boldsymbol{\alpha}}_N)/(2N)$, we obtain the following AIC and BIC:

$$\begin{aligned}\text{AIC}_1 &= \log L_1(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\alpha}}_N) + \frac{q_1}{NT}, \quad \text{BIC}_1 = \log L_1(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\alpha}}_N) + \frac{q_1 \log(NT)}{2NT}, \\ \text{AIC}_2 &= \log L_2(\hat{\boldsymbol{\phi}}, \hat{\boldsymbol{\alpha}}_N) + \frac{q_2}{N}, \quad \text{BIC}_2 = \log L_2(\hat{\boldsymbol{\phi}}, \hat{\boldsymbol{\alpha}}_N) + \frac{q_2 \log(N)}{2N}\end{aligned}$$

where q_1 and q_2 are the number of non-zero elements of $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$, respectively. Note that BIC_1 can be obtained from equation 2.3 of Lee, Noh, and Park (2014) by replacing the loss of standard quantile regression with $L_1(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\alpha}}_N)$.

The proposed criteria seem to work well in simulation (see Appendix C). However, they often tend to reward parsimony, possibly sacrificing goodness of fit. The testing procedure described in Section 6.1 should always be used to perform a preliminary screening of the candidate models.

7 SIMULATION RESULTS

We analyze the performance of the estimators $\hat{\boldsymbol{\beta}}(u)$ and $\hat{\boldsymbol{\gamma}}(v)$ in finite samples through numerical simulations. In particular, we report the biases and standard errors of these estimators for different values of the dimensions T and N and the orders of the quantiles u and v . We also evaluate the empirical size and power of the goodness of fit test.

We used the following design to generate the data:

$$Y_{it} = \beta_0(U_{it}) + \beta_1(U_{it})x_{it} + \gamma_0(V_i) + \gamma_1(V_i)z_i$$

where $x_{it} \sim \text{Beta}(2, 2)$ and $z_i \sim \text{U}(0, 3)$. In simulation 1, we defined:

$$\begin{aligned} \beta_0(u) &= 1 - 0.5 \log(1 - u), & \beta_1(u) &= 1 + 10(u - 0.5)^3, \\ \gamma_0(v) &= 0.5\zeta(v), & \gamma_1(v) &= 0.5\zeta(v), \end{aligned}$$

where $\zeta(v)$ is the quantile function of a standard normal distribution. In simulation 2, we defined:

$$\begin{aligned} \beta_0(u) &= 2(1 - (1 - u)^{1/4}), & \beta_1(u) &= 3(1 + u), \\ \gamma_0(v) &= \log(1 - \log(1 - v)), & \gamma_1(v) &= 0.5 \log(1 - \log(1 - v)). \end{aligned}$$

To fit the true model, we used $\mathbf{b}(u) = [1, -\log(1 - u), (u - 0.5)^3]^T$ and $\mathbf{c}(v) = [\zeta(v)]$ in simulation 1, and $\mathbf{b}(u) = [1, 1 - (1 - u)^{1/4}, u]^T$ and $\mathbf{c}(v) = [\log(1 - \log(1 - v))]$ in simulation 2. We ran $R = 1000$ Monte Carlo simulations, with $N = \{150, 300\}$ and $T = \{5, 10\}$. In Tables 1 and 2, we report the true value of $\boldsymbol{\beta}(\cdot)$ and $\boldsymbol{\gamma}(\cdot)$ at the quintiles, their average estimates, the empirical standard errors across simulations, and the average estimates of the asymptotic standard errors. Despite the incidental parameters problem, a small bias was found, even with small values of T . Also, as T increased, the observed bias decreased

rapidly as predicted by the asymptotic theory of Section 5. The estimated standard errors were, on average, very close to their true values.

To assess the performance of the goodness-of-fit procedure described in Section 6.1, we selected two nominal significance levels, $\alpha = 0.05$ and $\alpha = 0.10$, and computed the empirical probability of type I error ($\tilde{\alpha}$) and the power ($1 - \tilde{\beta}$) of the Kolmogorov-Smirnov goodness-of-fit test described in Section 6. The power was estimated by the empirical probability to reject a misspecified model in which the quantile function was described by an incorrect basis function. In simulation 1, we incorrectly parametrized $\beta_1(u)$ as a linear function, $\beta_1(u) = \theta_{01} + \theta_{11}u$. In simulation 2, we incorrectly assumed that the individual effects have a logistic distribution, defined by $\mathbf{c}(v) = [\log(v/(1 - v))]$. Results are shown in the bottom rows of Tables 1 and 2. The risk of type I error was very close to its nominal level, and approached it as the value of T increased. With small values of N and T , the risk of type II error was relatively large, and the power was often less than 50%. However, with $N = 300$ and $T = 10$, and a nominal level of 0.10 for rejection, the incorrect models were rejected in more than 90% of cases in both scenarios.

Additional simulation results are reported in Appendix C, where we compare our estimator with Koenker's (2004) penalized fixed-effects quantile regression, and discuss the performance of the model selection criteria presented in Section 6.2.

Table 1: Simulation results with $N = 150$

		Simulation 1				Simulation 2											
$T = 5$	u	β_0	$\hat{\beta}_0$	se	$\hat{s}e$	β_1	$\hat{\beta}_1$	se	$\hat{s}e$	β_0	$\hat{\beta}_0$	se	$\hat{s}e$	β_1	$\hat{\beta}_1$	se	$\hat{s}e$
	0.2	1.11	1.17	.10	.11	0.73	0.72	.10	.11	0.11	0.11	.01	.01	3.60	3.79	.11	.12
	0.4	1.26	1.30	.11	.11	0.99	0.97	.11	.11	0.24	0.24	.03	.03	4.20	4.28	.12	.12
	0.6	1.46	1.48	.11	.11	1.01	0.99	.11	.11	0.41	0.41	.05	.05	4.80	4.77	.15	.15
	0.8	1.80	1.79	.12	.12	1.27	1.24	.12	.13	0.66	0.66	.08	.09	5.40	5.26	.19	.19
	v	γ_0	$\hat{\gamma}_0$	se	$\hat{s}e$	γ_1	$\hat{\gamma}_1$	se	$\hat{s}e$	γ_0	$\hat{\gamma}_0$	se	$\hat{s}e$	γ_1	$\hat{\gamma}_1$	se	$\hat{s}e$
	0.2	0.42	0.37	.09	.10	0.42	0.40	.08	.08	0.20	0.20	.04	.04	0.10	0.10	.02	.03
	0.4	0.13	0.11	.03	.03	0.13	0.12	.02	.02	0.41	0.41	.07	.08	0.21	0.20	.05	.05
$T = 10$	u	β_0	$\hat{\beta}_0$	se	$\hat{s}e$	β_1	$\hat{\beta}_1$	se	$\hat{s}e$	β_0	$\hat{\beta}_0$	se	$\hat{s}e$	β_1	$\hat{\beta}_1$	se	$\hat{s}e$
	0.2	1.11	1.14	.10	.10	0.73	0.73	.06	.07	0.11	0.11	.01	.01	3.60	3.68	.07	.08
	0.4	1.26	1.27	.10	.10	0.99	0.98	.07	.07	0.24	0.23	.02	.02	4.20	4.23	.08	.08
	0.6	1.46	1.47	.10	.10	1.01	1.00	.07	.07	0.41	0.40	.03	.03	4.80	4.78	.10	.10
	0.8	1.80	1.80	.10	.10	1.27	1.25	.08	.08	0.66	0.64	.05	.06	5.40	5.33	.12	.13
	v	γ_0	$\hat{\gamma}_0$	se	$\hat{s}e$	γ_1	$\hat{\gamma}_1$	se	$\hat{s}e$	γ_0	$\hat{\gamma}_0$	se	$\hat{s}e$	γ_1	$\hat{\gamma}_1$	se	$\hat{s}e$
	0.2	0.42	0.40	.09	.09	0.42	0.41	.08	.08	0.20	0.21	.03	.04	0.10	0.10	.02	.03
	0.4	0.13	0.12	.03	.03	0.13	0.12	.02	.02	0.41	0.42	.07	.07	0.21	0.20	.05	.05
0.6	0.13	0.12	.03	.03	0.13	0.12	.02	.02	0.65	0.66	.11	.11	0.33	0.32	.08	.08	
0.8	0.42	0.40	.09	.09	0.42	0.41	.08	.08	0.96	0.98	.16	.17	0.48	0.47	.12	.12	
		$T = 5$		$T = 10$		$T = 5$		$T = 10$									
α		$\tilde{\alpha}$	$1 - \tilde{\beta}$		$\tilde{\alpha}$	$1 - \tilde{\beta}$		$\tilde{\alpha}$	$1 - \tilde{\beta}$		$\tilde{\alpha}$	$1 - \tilde{\beta}$					
0.05		0.05	0.22		0.04	0.50		0.04	0.30		0.04	0.53					
0.10		0.10	0.36		0.09	0.80		0.10	0.45		0.09	0.69					

Summary of simulation results, based on $R = 1000$ Monte Carlo replications, with $N = 150$ and $T = \{5, 10\}$. For each coefficient, we report the true *absolute* value (β, γ) at the quintiles $(0.2, 0.4, 0.6, 0.8)$, the average estimate $(\hat{\beta}, \hat{\gamma})$, the standard error (se), computed as the standard deviation of the estimated model parameters across simulations, and the average estimated asymptotic standard error ($\hat{s}e$). The bottom table reports, for two different nominal levels $\alpha = 0.05, 0.10$, the empirical probability of type I error ($\tilde{\alpha}$) and the power $(1 - \tilde{\beta})$ of the Kolmogorov-Smirnov goodness-of-fit test described in Section 6.

Table 2: Simulation results with $N = 300$

		Simulation 1								Simulation 2							
$T = 5$	u	β_0	$\hat{\beta}_0$	se	\hat{se}	β_1	$\hat{\beta}_1$	se	\hat{se}	β_0	$\hat{\beta}_0$	se	\hat{se}	β_1	$\hat{\beta}_1$	se	\hat{se}
	0.2	1.11	1.17	.08	.08	0.73	0.72	.07	.07	0.11	0.11	.01	.01	3.60	3.79	.08	.08
	0.4	1.26	1.30	.08	.08	0.99	0.97	.08	.08	0.24	0.24	.02	.02	4.20	4.28	.08	.08
	0.6	1.46	1.48	.08	.08	1.01	1.00	.08	.08	0.41	0.41	.03	.04	4.80	4.77	.10	.10
	0.8	1.80	1.79	.08	.08	1.27	1.24	.09	.09	0.66	0.66	.06	.06	5.40	5.27	.13	.13
$T = 5$	v	γ_0	$\hat{\gamma}_0$	se	\hat{se}	γ_1	$\hat{\gamma}_1$	se	\hat{se}	γ_0	$\hat{\gamma}_0$	se	\hat{se}	γ_1	$\hat{\gamma}_1$	se	\hat{se}
	0.2	0.42	0.37	.06	.07	0.42	0.41	.06	.06	0.20	0.20	.03	.03	0.10	0.10	.02	.02
	0.4	0.13	0.11	.02	.02	0.13	0.12	.02	.02	0.41	0.41	.05	.06	0.21	0.20	.04	.04
	0.6	0.13	0.11	.02	.02	0.13	0.12	.02	.02	0.65	0.65	.09	.09	0.33	0.32	.06	.06
	0.8	0.42	0.37	.06	.07	0.42	0.41	.06	.06	0.96	0.96	.13	.13	0.48	0.47	.09	.09
$T = 10$	u	β_0	$\hat{\beta}_0$	se	\hat{se}	β_1	$\hat{\beta}_1$	se	\hat{se}	β_0	$\hat{\beta}_0$	se	\hat{se}	β_1	$\hat{\beta}_1$	se	\hat{se}
	0.2	1.11	1.14	.07	.07	0.73	0.73	.05	.05	0.11	0.10	.01	.01	3.60	3.67	.05	.05
	0.4	1.26	1.27	.07	.07	0.99	0.98	.05	.05	0.24	0.23	.01	.01	4.20	4.23	.06	.06
	0.6	1.46	1.47	.07	.07	1.01	1.00	.05	.05	0.41	0.40	.02	.02	4.80	4.78	.07	.07
	0.8	1.80	1.80	.07	.07	1.27	1.25	.05	.06	0.66	0.64	.04	.04	5.40	5.33	.09	.09
$T = 10$	v	γ_0	$\hat{\gamma}_0$	se	\hat{se}	γ_1	$\hat{\gamma}_1$	se	\hat{se}	γ_0	$\hat{\gamma}_0$	se	\hat{se}	γ_1	$\hat{\gamma}_1$	se	\hat{se}
	0.2	0.42	0.39	.06	.07	0.42	0.41	.05	.06	0.20	0.21	.02	.02	0.10	0.10	.02	.02
	0.4	0.13	0.12	.02	.02	0.13	0.12	.02	.02	0.41	0.42	.05	.05	0.21	0.20	.03	.04
	0.6	0.13	0.12	.02	.02	0.42	0.12	.02	.02	0.65	0.66	.07	.08	0.33	0.32	.05	.06
	0.8	0.42	0.39	.06	.07	0.13	0.41	.05	.06	0.96	0.98	.11	.12	0.48	0.48	.08	.08
		$T = 5$				$T = 10$				$T = 5$				$T = 10$			
α		$\tilde{\alpha}$		$1 - \tilde{\beta}$		$\tilde{\alpha}$		$1 - \tilde{\beta}$		$\tilde{\alpha}$		$1 - \tilde{\beta}$		$\tilde{\alpha}$		$1 - \tilde{\beta}$	
	0.05	0.07		0.43		0.05		0.99		0.04		0.52		0.06		0.81	
	0.10	0.11		0.65		0.10		1.00		0.08		0.67		0.10		0.90	

Summary of simulation results with $N = 300$ and $T = \{5, 10\}$.

8 ANALYSIS OF NGAL DATA

We analyzed data from Mårtensson et al. (2013), aiming to investigate the role of plasma neutrophil gelatinase-associated lipocalin (NGAL) as a marker of sepsis and acute kidney dysfunction. The dataset included 139 patients admitted to the general intensive care unit at Karolinska University Hospital in Solna, Sweden, between August 2007 and November 2010. Baseline information was collected, and patients were classified daily as having sepsis or not. NGAL (mg/mL), procalcitonin (PCT), C-reactive protein (CRP), and creatinine changes relative to baseline (Δcreat) were measured daily before discharge, for a total of 1317 plasma samples. After removing missing data, individuals with only one observation, and one patient with severe complications, the final sample included 135 patients for a total sample size of $\sum_{i=1}^{135} T_i = 1263$. The number of observations per patient varied between 2 and 38, and more than 80% of patients had $T_i \leq 14$.

The goal of our analysis was to estimate conditional quantiles of NGAL, and in particular to measure its association with sepsis. The between-patient variability appeared to be very large, reflecting the presence of important differences in the initial health conditions. We formulated a regression model with the following predictors: a binary indicator of sepsis status, an indicator of $\Delta\text{creat} \geq 50$, age (centered at its median, 52 years, and divided by 10), an indicator of female gender, and time since hospitalization (weeks). Age and gender were cluster-invariant and were also included as level-2 predictors.

The response variable was log-transformed, which made it more plausible to define individual effects on the additive scale as in model (1). The regression function was

$$\begin{aligned} \log(\text{NGAL}_{it}) &= \beta_0(U_{it}) + \beta_1(U_{it})\text{sepsis}_{it} + \beta_2(U_{it})I(\Delta\text{creat}_{it} > 50) \\ &+ \beta_3(U_{it})(\text{age}_i - 52)/10 + \beta_4(U_{it})\text{female}_i + \beta_5(U_{it})\text{time}_{it} \\ &+ \gamma_0(V_i) + \gamma_1(V_i)(\text{age}_i - 52)/10 + \gamma_2(V_i)\text{female}_i. \end{aligned}$$

We formulated a variety of models, in which $\beta_0(\cdot)$ and $\gamma_0(\cdot)$ were unbounded, while the other coefficients were modeled by bounded functions. To facilitate interpretation, we forced $\gamma(0.5) = 0$, assigning the individual effects a zero-median distribution in which level-2 covariates only affect the scale parameter.

Selected modeling options are illustrated in Table 3. Different models appeared to fit the data well, and were not rejected by the goodness-of-fit test described in Section 6. The following model combined simplicity and flexibility, and was selected for illustrative purposes:

$$\begin{aligned}\beta_0(u \mid \boldsymbol{\theta}) &= \theta_{00} + \theta_{01} \log(u) + \theta_{02} \log(1 - u), \\ \beta_j(u \mid \boldsymbol{\theta}) &= \theta_{j0} + \theta_{j3}u + \theta_{j4}u^{1/4} + \theta_{j5}(1 - u)^{1/4}, \quad j = 1, \dots, 5, \\ \gamma_0(v \mid \boldsymbol{\phi}) &= \phi_{01} \log(2v) + \phi_{02} \log(2(1 - v)), \\ \gamma_j(v \mid \boldsymbol{\phi}) &= \phi_{j3}(v - 0.5), \quad j = 1, \dots, 2.\end{aligned}$$

The level-1 and level-2 intercepts were described by different versions of the asymmetric Logistic distribution. The coefficient functions associated with level-1 covariates were a combination of linear and root-4 functions, while those of the level-2 predictors were assumed to be linear.

The p-value of the Kolmogorov-Smirnov test was 0.21. To assess local fit, the test was repeated in subsamples with different values of the covariates (e.g., the females, those with $\Delta_{\text{creat}} > 50$, etc.). No significant evidence of model misspecification was found.

All 27 model parameters are reported in Table 4, while regression coefficients at selected quantiles are summarized in Table 5. We represent graphically the quantile regression coefficient functions in Figure 2, where we also report a “nonparametric” fit obtained by modeling all coefficients as piecewise linear functions with knots at the deciles.

Results showed that the distribution of the individual effects was almost symmetric

(as suggested by the fact that $\hat{\phi}_{01} \simeq -\hat{\phi}_{02}$) and that its variance was not significantly affected by cluster-level predictors. Instead, all predictors apart from gender appeared to be associated with the level-1 response. In particular, the coefficients associated with sepsis, $\Delta\text{creat}_{it} > 50$ and age were consistently positive at all quantiles, while the coefficient of time was always negative. The sepsis status was associated with a percentile difference of about 0.12 at quantiles 0.2, 0.4, 0.6, 0.8. As shown by Figure 2, an even larger percentile difference was found at quantiles above 0.8.

Table 3: Alternative specifications of $\mathbf{b}(u)$ and $\mathbf{c}(v)$

$\beta_0(u)$	$\beta_j(u)$	$\gamma_0(v)$	$\gamma_j(v)$
$\zeta(u)$	u	$\zeta(v)$	$(v - 0.5)$
$\log(\frac{u}{1-u})$	u, u^2, u^3, \dots	$\log(\frac{v}{1-v})$	$(v - 0.5)^3$
$\log(u), \log(1 - u)$	$u, \cos(\pi u), \sin(\pi u)$	$\log(2v), \log(2(1 - v))$	$(v - 0.5), (v - 0.5)^3$
$u, \log(u), \log(1 - u)$	$u^\delta, (1 - u)^\delta$	$v - 0.5, \log(2v), \log(2(1 - v))$	$(v - 0.5)_-^3, (v - 0.5)_+^3$

In the table, $\beta_0(u)$ and $\gamma_0(v)$ denote level-1 and level-2 intercept, respectively, while $\beta_j(u)$ and $\gamma_j(v)$ represent coefficients associated with generic level-1 and level-2 covariates. Different parametric models are represented by the functions that compose $\mathbf{b}(u)$ and $\mathbf{c}(v)$. A constant term $b(u) = 1$ was always included. The notation $\zeta(\cdot)$ is used for the quantile function of a standard normal distribution, and we defined $(v - 0.5)_-^3 = I(v \leq 0.5)(v - 0.5)^3$ and $(v - 0.5)_+^3 = I(v > 0.5)(v - 0.5)^3$.

Table 4: Estimated model parameters

Level 1 ($\hat{\theta}$)	1	$\log(u)$	$\log(1 - u)$	u	$u^{1/4}$	$(1 - u)^{1/4}$
Intercept	4.67 (.07)	0.12 (.03)	-0.13 (.03)	-	-	-
Sepsis	0.39 (.26)	-	-	-0.20 (.23)	0.14 (.29)	-0.35 (.27)
$\Delta_{\text{creat}} \geq 50$	0.39 (.61)	-	-	-0.44 (.40)	0.46 (.46)	-0.54 (.57)
(age - 52)/10	0.00 (.09)	-	-	0.02 (.07)	0.00 (.08)	0.06 (.09)
female gender	0.16 (.31)	-	-	-0.40 (.23)	0.34 (.28)	-0.29 (.30)
time (weeks)	-0.14 (.09)	-	-	0.12 (.08)	-0.12 (.10)	0.12 (.09)

Level 2 ($\hat{\phi}$)	$\log(2v)$	$\log(2(1 - v))$	$(v - 0.5)$
Intercept	0.33 (.08)	-0.29 (.08)	-
(age - 52)/10	-	-	0.11 (.07)
female gender	-	-	-0.20 (.28)

Summary of the selected model (top: $\hat{\theta}$; bottom: $\hat{\phi}$), with estimated standard errors in brackets. The model is represented graphically in Figure 2, and selected quantiles are summarized in Table 5.

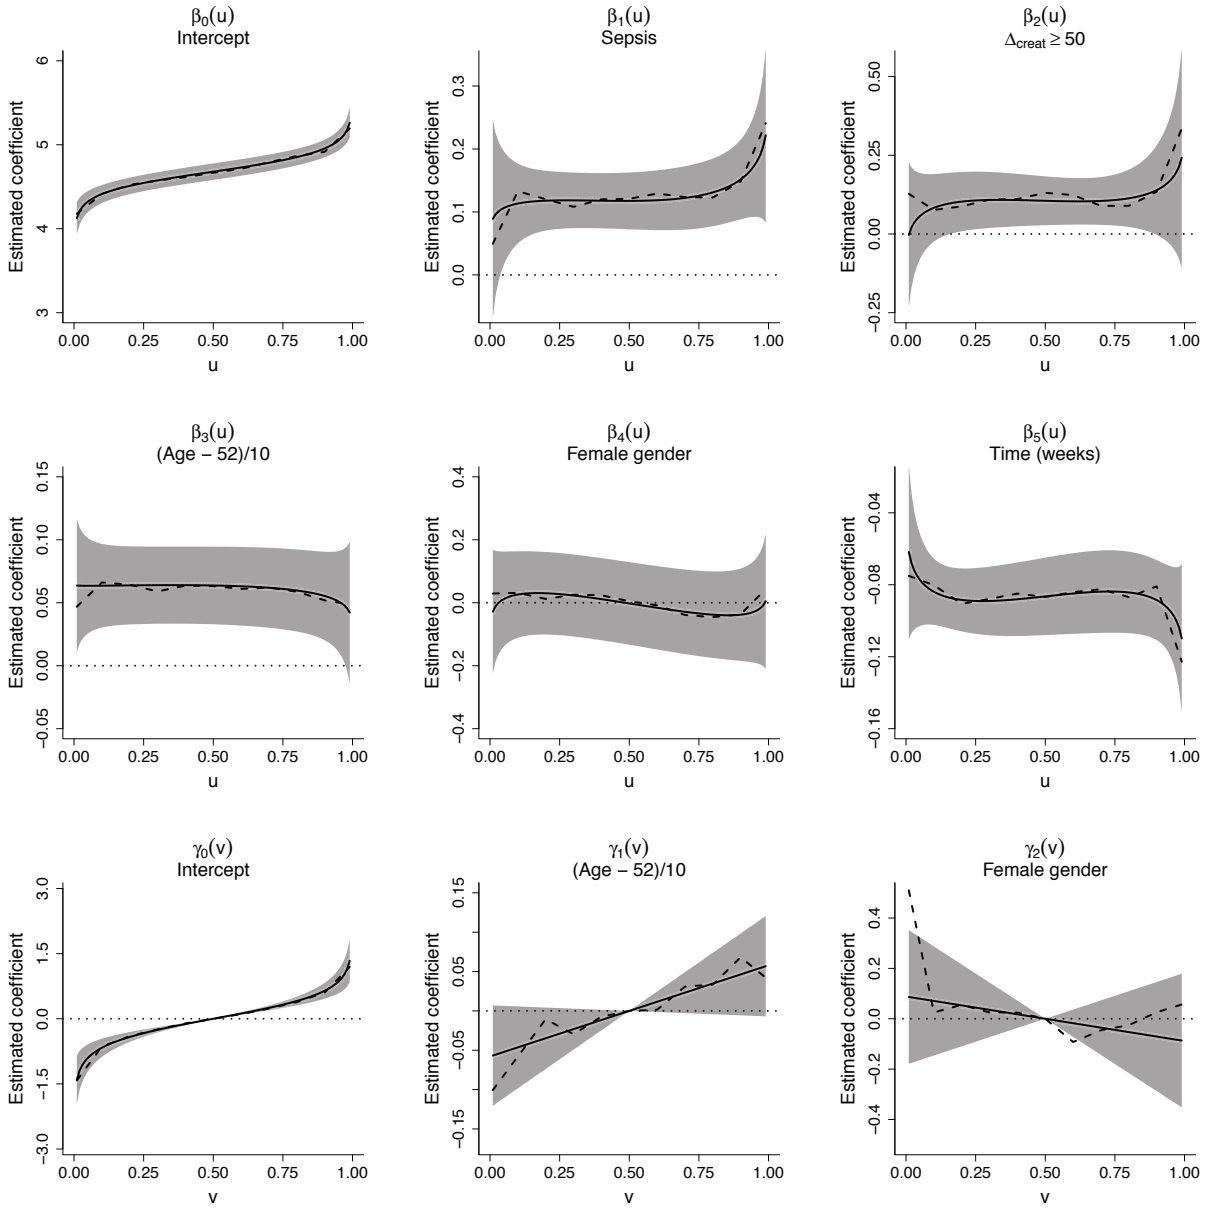


Figure 2: Continuous lines represent the estimated quantile regression coefficient functions, based on the parametric model summarized in Table 4. Shaded areas represents pointwise 95% confidence intervals. The dashed lines are obtained from a “nonparametric” model in which $\beta(u)$ and $\gamma(v)$ were fitted by piecewise linear functions with knots at the deciles. The dotted horizontal lines indicate the zero.

Table 5: Summary of selected quantiles

quantile	0.2	0.4	0.6	0.8
Intercept	4.51 (.05)	4.63 (.05)	4.73 (.05)	4.85 (.05)
Sepsis	0.12 (.02)	0.12 (.02)	0.12 (.02)	0.13 (.03)
Δcreat	0.10 (.05)	0.11 (.04)	0.10 (.04)	0.11 (.05)
(age - 52)/10	0.06 (.02)	0.06 (.02)	0.06 (.02)	0.06 (.02)
female gender	0.03 (.07)	0.01 (.07)	-0.02 (.07)	-0.04 (.07)
time (weeks)	-0.09 (.01)	-0.09 (.01)	-0.08 (.01)	-0.08 (.01)
Intercept	-0.42 (.06)	-0.12 (.01)	0.12 (.01)	0.41 (.05)
(age - 52)/10	-0.03 (.02)	-0.01 (.01)	0.01 (.01)	0.03 (.02)
female gender	0.05 (.08)	0.02 (.03)	-0.02 (.03)	-0.05 (.08)

Estimated regression coefficients at quantiles (0.2, 0.4, 0.6, 0.8). Top table: $\hat{\beta}(u) = \beta(u | \hat{\theta}) = \hat{\theta}b(u)$;
bottom table: $\hat{\gamma}(v) = \gamma(v | \hat{\phi}) = \hat{\phi}c(v)$. Estimated standard errors in brackets.

9 CONCLUSIONS

We introduced a general framework for longitudinal quantile regression, extending the work of Frumento and Bottai (2016; 2017) on quantile regression coefficients modeling. We defined a two-level quantile function in which both the “within” and the “between” part of the distribution are described by a quantile regression model. This allows to investigate how covariates affect not only the level-1 response, but also the distribution of the individual effects, which is generally overlooked in the existing literature on longitudinal quantile regression. Identification is achieved by modeling the coefficient functions parametrically, and estimation is carried out by minimizing a smooth objective function.

The proposed method is computationally simple and can be viewed as a special type of penalized fixed-effects estimator that presents important elements of novelty of its own. The penalty term carries information on the parameters that describe the conditional distribution of the individual effects. This permits estimating both level-1 and level-2 parameters, as in random-effects models, but allows carrying out estimation and inference using fixed-effects techniques. Moreover, it avoids the problem of choosing a tuning constant as in standard ℓ_1 - or ℓ_2 -penalization. The described form of penalized fixed-effects method is not limited to a quantile regression framework and could be applied to different estimation problems.

The proposed modeling framework can be generalized in different directions. An interesting possibility is to include multiple individual effects as in random-slope models. In our framework, individual effects are represented by a pure location shift as in Koenker (2004), Geraci and Bottai (2007, 2014), and Canay (2011). Using the proposed penalized fixed-effects approach, it is relatively simple to incorporate not only an individual intercept, α_i , but also a set of individual slopes, say $\{\delta_{1i}, \delta_{2i}, \dots\}$. This, however, would typically result

in cumbersome computation and, unless N and T are sufficiently large, would probably undermine model identifiability. Using Koenker’s (2004) words: “At best we may be able to estimate an individual specific location-shift effect, and even this may strain credulity”.

Another interesting extension is represented by varying-coefficients models (e.g., Hastie and Tibshirani, 1993; Fan and Zhang, 1999, 2000; Chiang, Rice, and Wu, 2001; Kim, 2007) that could be implemented by allowing the level-1 regression coefficients to be functions of time. A possible approach is to describe the coefficients, say $\beta(u, t)$, using tensor products of splines. Finally, the proposed method could be used to estimate static and dynamic quantile autoregressive models (e.g., Arellano and Bonhomme, 2016).

An important problem that has not been discussed in the paper is represented by quantile crossing, occurring when either $\mathbf{x}_{it}^T \beta'(u | \hat{\theta}) < 0$ or $\mathbf{z}_i^T \gamma'(v | \hat{\phi}) < 0$. One may want to determine in advance which values of the parameters θ and ϕ would ensure that no crossing occurs, i.e., that $\mathbf{x}_{it}^T \beta(\cdot | \hat{\theta})$ and $\mathbf{z}_i^T \gamma(\cdot | \hat{\phi})$ are monotonically increasing functions. This is only possible in very simple models with few covariates, or in presence of restrictive assumptions. However, simulation evidence suggests that parametric models are relatively immune to quantile crossing, compared with the “nonparametric” approaches based on ordinary quantile regression. Additionally, the parametric structure makes it particularly simple to verify crossing, taking advantage of the closed-form analytical expression of the quantile function, and admits the application of monotonization methods such as the rearrangement of Chernozhukov et al. (2010) to produce increasing estimates of conditional quantiles.

This paper is accompanied by an R package `qrqm`, that includes a function named `iqrL` that performs model fitting, and a variety of auxiliary functions for prediction, plotting, and goodness-of-fit assessment. The documentation contains a rich set of examples, and can serve as tutorial for the practitioners. The package is available upon request to the

authors.

References

- Abrevaya, J., and Dahl, C. M., (2008). “The Effects of Birth Inputs on Birthweight”, *Journal of Business & Economic Statistics* **26**(4), 379–397.
- Alfó, M., Salvati, N., and Ranalli M.G., (2017). “Finite mixtures of quantile and M-quantile regression models”, *Statistics and Computing* **27**(2), 547–570.
- Akaike, H. (1974). “A new look at the statistical model identification”, *IEEE Transactions on Automatic Control*, **19**(6), 716–723.
- Angrist, J. D., Chernozhukov, V., and I. Fernandez-Val (2006). “Quantile Regression under Misspecification, with an Application to the U.S. Wage Structure,” *Econometrica* **74**, pp. 539–563.
- Arellano, M., and Bonhomme, S. (2016). “Nonlinear panel data estimation via quantile regression”, *The Econometric Journal*, **19**(3), C61–C94, doi: 10.1111/ectj.12062.
- Arellano M, Hahn J. 2007. Understanding bias in nonlinear panel models: Some recent developments. *Econometric Society Monographs* 43:381
- Barrodale, I. and Roberts, F. D. K. (1973). “An improved algorithm for discrete L_1 linear approximation”, *S.I.A.M. Journal on Numerical Analysis* **10**(5), 839–848.
- Canay, I.A. (2011). “A simple approach to quantile regression for panel data”, *Econometrics Journal*, **14**(3), 368–386.

- Chamberlain G. 1984. Panel Data. *Griliches and M. Intrilligator, eds., Handbook of Econometrics, Chapter 22* :1247–1318
- Cheng, G., and Shang, Z. (2015). “Joint asymptotics for semi-nonparametric regression models with partially linear structure”, *Annals of Statistics*, **43**(3), 1351–1390.
- Chernozhukov, V., Fernández-Val, I., and Galichon, A., (2010). “Quantile and Probability Curves without Crossing”, *Econometrica*, **78**(3), 1093–1125.
- Chernozhukov, V., Fernández-Val, I., and Weidner, M., (2018). “Network and Panel Quantile Effects Via Distribution Regression”, *arXiv eprint 1803.08154*.
- Chiang, C.T., Rice, J. A., and Wu, C.O. (2001). “Smoothing spline estimation for varying coefficient models with repeatedly measured dependent variables”, *Journal of the American Statistical Association*, **96**, 605–619.
- Dhaene G, Jochmans K. 2015. “Split-panel jackknife estimation of fixed-effect models,” *The Review of Economic Studies* **82**, 991–1030
- Farcomeni, A. (2012). “Quantile regression for longitudinal data based on latent Markov subject-specific parameters”, *Statistics and Computing*, **22**(1), 141–152.
- Fan, J. and Zhang, W. (1999). “Statistical estimation in varying coefficient models”, *Annals of Statistics*, **27**, 1491–1518.
- Fan, J. and Zhang, J. T. (2000). “Functional linear models for longitudinal data”, *Journal of the Royal Statistical Society, Series B*, **62**, 303–322.
- Fernández-Val, I. (2005). “Bias correction in panel data models with individual specific parameters”, *mimeo, Department of Economics, Boston University, Boston, MA*.

- Phillips PCB, Moon H. 1999. Linear regression limit theory for nonstationary panel data. *Econometrica* 67:1057–1111
- Frumento, P. (2016). `qrqm`: *Quantile regression coefficients modeling*. R package version 2.0, url: <http://CRAN.R-project.org/package=qrqm>
- Frumento, P., and Bottai, M. (2016). “Parametric modeling of quantile regression coefficient functions”, *Biometrics*, **72**(1), 74–84, doi: 10.1111/biom.12410.
- Frumento, P., and Bottai, M. (2017). “Parametric modeling of quantile regression coefficient functions with censored and truncated data”, *Biometrics*, **73**(4), 1179–1188, doi: 10.1111/biom.12675.
- Geraci, M., and Bottai, M. (2007). “Quantile regression for longitudinal data using the asymmetric Laplace distribution”, *Biostatistics*, **8**(1), 140–154.
- Geraci, M., and Bottai, M. (2014). “Linear quantile mixed models”, *Statistics and Computing*, **24**(3), 461–479.
- Gilchrist, W. (2000). *Statistical modeling with quantile functions*, Chapman & Hall, ISBN 1-58488-174-7.
- Hahn, J., and Newey, W. (2004). “Jackknife and analytical bias reduction for nonlinear panel models”, *Econometrica*, **72**(4), 1295–1319.
- Hastie, T. J. and Tibshirani, R. J. (1993). “Varying-coefficient models”, *Journal of the Royal Statistical Society, Series B*, **55**, 757–796.
- Kato, K., Galvao, A.F., Jr., and Montes-Rojas, G.V. (2012). “Asymptotics for panel quantile regression models with individual effects”, *Journal of Econometrics*, **170**(1), 76–91.

- Kim, M.O., and Yang, Y. (2011). “Semiparametric approach to a random effects quantile regression model”, *Journal of the American Statistical Association*, **106**(496), 1405–1417.
- Kim, M.O. (2007). “Quantile regression with varying coefficients”, *Annals of Statistics*, **35**(1), 92–108.
- Koenker, R. (2004). “Quantile regression for longitudinal data”, *Journal of Multivariate Analysis*, **91**(1), 74–89.
- Koenker, R. (2005). “*Quantile regression*”, Econometric Society Monograph Series, Cambridge: Cambridge University Press.
- Koenker, R. and Bassett, G., Jr. (1978). “Regression Quantiles”, *Econometrica* **46**(1), 33–50.
- Koenker, R. and Machado, J.A.F. (1999). “Goodness of fit and related inference processes for quantile regression”, *Journal of the American Statistical Association* **94**(448), 1296–1310.
- Lamarche, C. (2010). “Robust penalized quantile regression estimation for panel data”, *Journal of Econometrics*, **157**(2), 396–408.
- Lancaster, T. (2000). “The incidental parameter problem since 1948”, *Journal of Econometrics*, **95**(2), 391–413.
- Lee, E.R., Noh, H., and Park, B.U. (2014). “Model selection via Bayesian information criterion for quantile regression models”, *Journal of the American Statistical Association*, **109**(505), 216–229, DOI: 10.1080/01621459.2013.836975

- Leng, C, and Zhang, W. (2014). “Smoothing combined estimating equations in quantile regression for longitudinal data”, *Statistics and Computing*, **24**(1), 123–136.
- Machado, J. A. F. (1993). “Robust model selection and M-estimation”, *Econometric Theory*, **9**(3), 478–493.
- Marino, F., Tzavidis, N., and Alfó, M. (2016). “Mixed hidden Markov quantile regression models for longitudinal data with possibly incomplete sequences”, *Statistical methods in medical research*, 962280216678433.
- Masuda, H., and Shimizu, Y. (2017). “Moment convergence in regularized estimation under multiple and mixed-rates asymptotics”, *Mathematical Methods of Statistics*, **26**(2), 81–110.
- Mårtensson, J., Bell, M., Xu, S., Bottai, M., Ravn, B., Venge, P., Martling, C.R. (2013). “Association of plasma neutrophil gelatinase-associated lipocalin (NGAL) with sepsis and acute kidney dysfunction”, *Biomarkers*, **18**(4), 349–356, DOI: 10.3109/1354750X.2013.787460.
- Newey, W.K. (2007). “Course materials for 14.386 *New Econometric Methods, Spring 2007*”, *MIT OpenCourseWare* (<http://ocw.mit.edu>), Massachusetts Institute of Technology. Downloaded on 15 June 2017.
- Newey, W. K. and McFadden, D. (1994). *Large sample estimation and hypothesis testing*. in R. F. Engle and D. L. McFadden (eds.), *Handbook of Econometrics* **4**, Ch. 36, 2111–2245, Handbooks in Econometrics, 2, North-Holland, Amsterdam.
- Neyman, J., and Scott, E.L. (1948). “Consistent estimation from partially consistent observations”, *Econometrica*, **16**(1), 11–32.

- Radchenko, P. (2008). “Mixed-rates asymptotics”, *The Annals of Statistics*, **36**(1), 287–309, DOI: 10.1214/009053607000000668.
- Rilstone, P., Srivastava, V.K., and Ullah, A. (1996). “The second-order bias and mean squared error of nonlinear estimators”, *Journal of Econometrics*, **75**(2), 369–395.
- Schwarz, G. (1978). “Estimating the dimension of a model”, *The Annals of Statistics*, **6**(2), 461–464.
- Yang, C.C., Chen, Y.H., and Chang, H.Y. (2017). “Composite marginal quantile regression analysis for longitudinal adolescent body mass index data”, *Statistics in Medicine*, **36**(21), 3380–3397, DOI: 10.1002/sim.7355.
- Zhao, W., Lian, H., and Liang, H. (2017). “GEE analysis for longitudinal single-index quantile regression”, *Journal of Statistical Planning and Inference*, **187**, 78–102.

Appendix A - Proof of Theorem 1

Asymptotic Expansion for Estimators with Different Rates of Convergence

We start by extending the asymptotic expansions of Hahn and Newey (2004) and Fernández-Val (2005) for fixed effects estimators in nonlinear panel data models to models with mixed-rates asymptotics. In particular, we consider the M-estimators:

$$(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\phi}}) \in \arg \min_{\{\boldsymbol{\theta} \in \Theta, \boldsymbol{\phi} \in \Phi\}} L(\boldsymbol{\theta}, \boldsymbol{\phi}), \quad \hat{L}(\boldsymbol{\theta}, \boldsymbol{\phi}) = \frac{1}{NT} \left[\sum_{i=1}^N \sum_{t=1}^T \ell_{1it}(\boldsymbol{\theta}, \hat{\alpha}_i(\boldsymbol{\theta}, \boldsymbol{\phi})) + \sum_{i=1}^N \ell_{2i}(\boldsymbol{\phi}, \hat{\alpha}_i(\boldsymbol{\theta}, \boldsymbol{\phi})) \right], \quad (16)$$

where

$$\hat{\alpha}_i(\boldsymbol{\theta}, \boldsymbol{\phi}) \in \arg \min_{\alpha_i \in \mathcal{A}} \hat{L}_i(\boldsymbol{\theta}, \boldsymbol{\phi}, \alpha_i), \quad \hat{L}_i(\boldsymbol{\theta}, \boldsymbol{\phi}, \alpha_i) = \frac{1}{T} \left[\sum_{t=1}^T \ell_{1it}(\boldsymbol{\theta}, \alpha_i) + \ell_{2i}(\boldsymbol{\phi}, \alpha_i) \right], \quad (17)$$

ℓ_{1it} and ℓ_{2i} are random functions that depend on the data, and \mathcal{A} , Θ and Φ denote the parameter spaces for α_i , $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$, respectively.

Let $\mathcal{B}_\varepsilon(\alpha_i^0)$, $1 \leq i \leq n$, $\mathcal{B}_\varepsilon(\boldsymbol{\theta}^0)$ and $\mathcal{B}_\varepsilon(\boldsymbol{\phi}^0)$ be ε -neighborhoods of the true value of the parameters α_i^0 , $1 \leq i \leq n$, $\boldsymbol{\theta}^0$ and $\boldsymbol{\phi}^0$, for some $\varepsilon > 0$. In what follows, we use superscripts for partial derivatives, e.g. $\ell_{1it}^\theta(\boldsymbol{\theta}, \alpha_i) := \partial \ell_{1it}(\boldsymbol{\theta}, \alpha_i) / \partial \boldsymbol{\theta}$ and $\ell_{1it}^{\alpha\alpha}(\boldsymbol{\theta}, \alpha_i) := \partial^2 \ell_{1it}(\boldsymbol{\theta}, \alpha_i) / \partial \alpha_i^2$, and often drop the arguments when the functions are evaluated at the true values, e.g. $\ell_{1it}^\theta := \ell_{1it}^\theta(\boldsymbol{\theta}^0, \alpha_i^0)$. We assume that the parameters $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ are vector-valued. We explain below how to adapt the expansions to the integrated loss minimization estimator where $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ can be matrices.

Assumption 2 (i) *Consistency:* $\hat{\boldsymbol{\theta}} \rightarrow_P \boldsymbol{\theta}^0$, $\hat{\boldsymbol{\phi}} \rightarrow_P \boldsymbol{\phi}^0$ and $\sup_i |\hat{\alpha}_i - \alpha_i^0| \rightarrow_P 0$ for $\hat{\alpha}_i := \hat{\alpha}_i(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\phi}})$. (ii) *For each i , $(\boldsymbol{\theta}, \alpha_i) \mapsto \ell_{1it}(\boldsymbol{\theta}, \alpha_i)$ is four times continuously differentiable a.s. on $\mathcal{B}_\varepsilon(\boldsymbol{\theta}^0) \times \mathcal{B}_\varepsilon(\alpha_i^0)$ and the partial derivatives up to fourth order are bounded in absolute value by random variables $M_{it} > 0$ a.s. such that $\sup_i \text{plim}_{T \rightarrow \infty} \sum_{t=1}^T M_{it}^\xi / T < \infty$ for some $\xi > 8$; and $(\boldsymbol{\phi}, \alpha_i) \mapsto \ell_{2i}(\boldsymbol{\phi}, \alpha_i)$ is four times continuously differentiable a.s. on $\mathcal{B}_\varepsilon(\boldsymbol{\phi}^0) \times \mathcal{B}_\varepsilon(\alpha_i^0)$ and the partial derivatives up to fourth order are bounded in absolute value a.s. (iii) *The true value of the parameters are in the interiors of the parameter spaces, i.e.,**

$\theta^0 \in \Theta^\circ$, $\phi^0 \in \Phi^\circ$, and $\alpha_i^0 \in \mathcal{A}^\circ$ for each i . (iv) The following limits exist:

$$\begin{aligned}\bar{\mathbf{H}}_\theta &= \text{plim}_{N,T \rightarrow \infty} \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \left[\ell_{1it}^{\theta\theta} - \ell_{1it}^{\theta\alpha} \frac{\sum_{t=1}^T \ell_{1it}^{\alpha\theta}}{\sum_{t=1}^T \ell_{1it}^{\alpha\alpha}} \right], \\ \bar{\mathbf{H}}_{\phi\theta} &= -\text{plim}_{N,T \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \ell_{2i}^{\phi\alpha} \frac{\sum_{t=1}^T \ell_{1it}^{\alpha\theta}}{\sum_{t=1}^T \ell_{1it}^{\alpha\alpha}}, \quad \bar{\mathbf{H}}_\phi = \text{plim}_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \ell_{2i}^{\phi\phi}, \\ \bar{b}_\theta &= \text{plim}_{N,T \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N b_{\theta i}, \quad \bar{b}_\phi = \text{plim}_{N,T \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N b_{\phi i}, \\ \bar{\mathcal{O}}_\theta &= \text{plim}_{N,T \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \varphi_{\theta i} \varphi_{\theta i}^\top, \quad \bar{\mathcal{O}}_\phi = \text{plim}_{N,T \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \varphi_{\phi i} \varphi_{\phi i}^\top,\end{aligned}$$

where

$$\begin{aligned}b_{\theta i} &= \frac{1}{T} \sum_{t=1}^T \ell_{1it}^{\theta\alpha} \beta_i - \frac{\sigma_i^2}{T} \sum_{t=1}^T \ell_{1it}^{\theta\alpha} \ell_{1it}^{\alpha\alpha} + \frac{1}{2T} \sum_{t=1}^T \ell_{1it}^{\theta\alpha\alpha} \psi_i^2, \\ b_{\phi i} &= \ell_{2i}^{\phi\alpha} (\sqrt{T} \psi_i + \beta_i) + \frac{1}{2} \ell_{2i}^{\phi\alpha\alpha} \psi_i^2, \\ \varphi_{\theta i} &= \frac{1}{\sqrt{T}} \sum_{t=1}^T \ell_{1it}^\theta + \frac{\psi_i}{T} \sum_{t=1}^T \ell_{1it}^{\theta\alpha}, \quad \varphi_{\phi i} = \ell_{2i}^\phi, \\ \beta_i &= \sigma_i^2 \text{plim}_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \left(\sigma_i^2 \ell_{1it}^{\alpha\alpha} \ell_{1it}^\alpha - \frac{\psi_i^2}{2} \ell_{1it}^{\alpha\alpha\alpha} \right), \\ \psi_i &= -\frac{\sigma_i^2}{\sqrt{T}} \left(\sum_{t=1}^T \ell_{1it}^\alpha + \ell_{2i}^\alpha \right), \quad \sigma_i^2 = \left[\text{plim}_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \ell_{1it}^{\alpha\alpha} \right]^{-1}.\end{aligned}$$

(v) The minimum eigenvalues of the matrices $\bar{\mathbf{H}}_\theta$ and $\bar{\mathbf{H}}_\phi$ are bounded away from zero and $\inf_i \text{plim}_{T \rightarrow \infty} \sum_{t=1}^T \ell_{1it}^{\alpha\alpha}/T$ is bounded away from zero.

Theorem 2 Suppose that Assumption 2 holds. If $N = O(T)$,

$$\sqrt{NT} \left(\hat{\theta} - \theta^0 + \frac{\bar{\mathbf{H}}_\theta^{-1} \bar{b}_\theta}{T} \right) \rightarrow_d \bar{\mathbf{H}}_\theta^{-1} \mathcal{N}(0, \bar{\mathcal{O}}_\theta).$$

If $N = O(T^2)$,

$$\sqrt{N} \left(\hat{\phi} - \phi^0 + \frac{\bar{\mathbf{H}}_{\phi}^{-1}(\bar{b}_{\phi} + \bar{\mathbf{H}}_{\phi\theta} \bar{\mathbf{H}}_{\theta}^{-1} \bar{b}_{\theta})}{T} \right) \rightarrow_d \bar{\mathbf{H}}_{\phi}^{-1} \mathcal{N}(0, \bar{\mathcal{V}}_{\phi}).$$

Proof. We divide the proof in two steps: asymptotic expansions and asymptotic distributions.

Step 1: Asymptotic Expansions. The asymptotic expansions of $\hat{\theta}$ and $\hat{\phi}$ are derived in the following steps:

1. First-order asymptotic expansion of the first order conditions of (16) around the true values θ^0 and ϕ^0 .
2. Second-order expansion of $\hat{\alpha}_i(\theta^0, \phi^0)$ from (17).
3. Second-order expansion of the gradients $\sum_{t=1}^T \ell_{1it}^{\theta}(\theta^0, \hat{\alpha}_i(\theta^0, \phi^0))$ and $\ell_{2i}^{\phi}(\phi^0, \hat{\alpha}_i(\theta^0, \phi^0))$ with respect to $\hat{\alpha}_i(\theta^0, \phi^0)$ around α_i^0 .
4. Asymptotic expansion of $\hat{\theta}$ and $\hat{\phi}$.

Assumption 2 guarantees that in these expansions all the terms are bounded in probability and the remainders are negligible (e.g., Hahn and Newey, 2004, and Fernández-Val, 2005).

1. First-Order Expansion of the First-Order Conditions for $(\hat{\theta}, \hat{\phi})$: the first-order conditions of (16) are

$$\begin{aligned} 0 &= \frac{1}{NT} \sum_{i=1}^N \left\{ \sum_{t=1}^T \ell_{1it}^{\theta}(\hat{\theta}, \hat{\alpha}_i(\hat{\theta}, \hat{\phi})) + \left[\sum_{t=1}^T \ell_{1it}^{\alpha}(\hat{\theta}, \hat{\alpha}_i(\hat{\theta}, \hat{\phi})) + \ell_{2i}^{\phi}(\hat{\phi}, \hat{\alpha}_i(\hat{\theta}, \hat{\phi})) \right] \hat{\alpha}_i^{\theta}(\hat{\theta}, \hat{\phi}) \right\} \\ &= \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \ell_{1it}^{\theta}(\hat{\theta}, \hat{\alpha}_i(\hat{\theta}, \hat{\phi})) \end{aligned}$$

and

$$0 = \frac{1}{NT} \sum_{i=1}^N \left\{ \ell_{2i}^{\phi}(\hat{\phi}, \hat{\alpha}_i(\hat{\theta}, \hat{\phi})) + \left[\sum_{t=1}^T \ell_{1it}^{\alpha}(\hat{\theta}, \hat{\alpha}_i(\hat{\theta}, \hat{\phi})) + \ell_{2i}^{\alpha}(\hat{\phi}, \hat{\alpha}_i(\hat{\theta}, \hat{\phi})) \right] \hat{\alpha}_i^{\phi}(\hat{\theta}, \hat{\phi}) \right\} \\ = \frac{1}{NT} \sum_{i=1}^N \ell_{2i}^{\phi}(\hat{\phi}, \hat{\alpha}_i(\hat{\theta}, \hat{\phi})),$$

where the second equality in both cases follows from

$$\sum_{t=1}^T \ell_{1it}^{\alpha}(\theta, \hat{\alpha}_i(\theta, \phi)) + \ell_{2i}^{\alpha}(\phi, \hat{\alpha}_i(\theta, \phi)) = 0, \quad (18)$$

by the first-order condition of (17).

A first-order expansion in $(\hat{\theta}, \hat{\phi})$ around (θ^0, ϕ^0) yields

$$0 = \frac{1}{NT} \sum_{i=1}^N \left[\begin{array}{c} \sum_{t=1}^T \ell_{1it}^{\theta}(\hat{\theta}, \hat{\alpha}_i(\hat{\theta}, \hat{\phi})) \\ \ell_{2i}^{\phi}(\hat{\phi}, \hat{\alpha}_i(\hat{\theta}, \hat{\phi})) \end{array} \right] = \frac{1}{NT} \sum_{i=1}^N \left[\begin{array}{c} \sum_{t=1}^T \ell_{1it}^{\theta}(\theta^0, \hat{\alpha}_i(\theta^0, \phi^0)) \\ \ell_{2i}^{\phi}(\phi^0, \hat{\alpha}_i(\theta^0, \phi^0)) \end{array} \right] \\ + \frac{1}{NT} \sum_{i=1}^N \left[\begin{array}{cc} \sum_{t=1}^T \frac{d\ell_{1it}^{\theta}(\bar{\theta}, \hat{\alpha}_i(\bar{\theta}, \bar{\phi}))}{d\theta^{\top}} & \sum_{t=1}^T \frac{d\ell_{1it}^{\theta}(\bar{\theta}, \hat{\alpha}_i(\bar{\theta}, \bar{\phi}))}{d\phi^{\top}} \\ \frac{d\ell_{2i}^{\phi}(\bar{\phi}, \hat{\alpha}_i(\bar{\theta}, \bar{\phi}))}{d\theta^{\top}} & \frac{d\ell_{2i}^{\phi}(\bar{\phi}, \hat{\alpha}_i(\bar{\theta}, \bar{\phi}))}{d\phi^{\top}} \end{array} \right] \left[\begin{array}{c} \hat{\theta} - \theta^0 \\ \hat{\phi} - \phi^0 \end{array} \right], \quad (19)$$

where

$$\frac{d\ell_{1it}^{\theta}(\theta, \hat{\alpha}_i(\theta, \phi))}{d\theta^{\top}} = \ell_{1it}^{\theta\theta}(\theta, \hat{\alpha}_i(\theta, \phi)) + \ell_{1it}^{\theta\alpha}(\theta, \hat{\alpha}_i(\theta, \phi)) \hat{\alpha}_i^{\theta}(\theta, \phi)^{\top},$$

$$\frac{d\ell_{1it}^{\theta}(\theta, \hat{\alpha}_i(\theta, \phi))}{d\phi^{\top}} = \ell_{1it}^{\theta\alpha}(\theta, \hat{\alpha}_i(\theta, \phi)) \hat{\alpha}_i^{\phi}(\theta, \phi)^{\top},$$

$$\frac{d\ell_{2i}^{\phi}(\phi, \hat{\alpha}_i(\theta, \phi))}{d\theta^{\top}} = \ell_{2i}^{\phi\alpha}(\phi, \hat{\alpha}_i^{\theta}(\theta, \phi)) \hat{\alpha}_i^{\theta}(\theta, \theta)^{\top},$$

$$\frac{d\ell_{2i}^{\phi}(\phi, \hat{\alpha}_i(\theta, \phi))}{d\phi^{\top}} = \ell_{2i}^{\phi\phi}(\phi, \hat{\alpha}_i(\theta, \phi)) + \ell_{2i}^{\phi\alpha}(\phi, \hat{\alpha}_i^{\theta}(\theta, \phi)) \hat{\alpha}_i^{\phi}(\theta, \phi)^{\top},$$

and $(\bar{\theta}, \bar{\phi})$ lies between $(\hat{\theta}, \hat{\phi})$ and (θ^0, ϕ^0) .

The expressions of $\hat{\alpha}_i^\theta(\boldsymbol{\theta}, \boldsymbol{\phi})$ and $\hat{\alpha}_i^\phi(\boldsymbol{\theta}, \boldsymbol{\phi})$ can be obtained from (18). Thus, differentiation with respect to $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ gives

$$\hat{\alpha}_i^\theta(\boldsymbol{\theta}, \boldsymbol{\phi}) = -\frac{\sum_{t=1}^T \ell_{1it}^{\alpha\theta}(\boldsymbol{\theta}, \hat{\alpha}_i(\boldsymbol{\theta}, \boldsymbol{\phi}))^\top}{\sum_{t=1}^T \ell_{1it}^{\alpha\alpha}(\boldsymbol{\theta}, \hat{\alpha}_i(\boldsymbol{\theta}, \boldsymbol{\phi})) + \ell_{2i}^{\alpha\alpha}(\boldsymbol{\phi}, \hat{\alpha}_i(\boldsymbol{\theta}, \boldsymbol{\phi}))},$$

and

$$\hat{\alpha}_i^\phi(\boldsymbol{\theta}, \boldsymbol{\phi}) = -\frac{\ell_{2i}^{\alpha\phi}(\boldsymbol{\phi}, \hat{\alpha}_i(\boldsymbol{\theta}, \boldsymbol{\phi}))^\top}{\sum_{t=1}^T \ell_{1it}^{\alpha\alpha}(\boldsymbol{\theta}, \hat{\alpha}_i(\boldsymbol{\theta}, \boldsymbol{\phi})) + \ell_{2i}^{\alpha\alpha}(\boldsymbol{\phi}, \hat{\alpha}_i(\boldsymbol{\theta}, \boldsymbol{\phi}))},$$

respectively. Note that $\hat{\alpha}_i^\theta(\boldsymbol{\theta}, \boldsymbol{\phi}) = O_P(1)$ and $\hat{\alpha}_i^\phi(\boldsymbol{\theta}, \boldsymbol{\phi}) = O_P(T^{-1})$, because all the terms are bounded in probability and the denominator is bounded away from zero with probability one.

2. Second-Order Expansion of $\hat{\alpha}_i(\boldsymbol{\theta}^0, \boldsymbol{\phi}^0)$: the first-order condition of (17) at $(\boldsymbol{\theta}, \boldsymbol{\phi}) = (\boldsymbol{\theta}^0, \boldsymbol{\phi}^0)$

is

$$0 = \sum_{t=1}^T \ell_{1it}(\boldsymbol{\theta}^0, \hat{\alpha}_i(\boldsymbol{\theta}^0, \boldsymbol{\phi}^0)) + \ell_{2i}(\boldsymbol{\phi}^0, \hat{\alpha}_i(\boldsymbol{\theta}^0, \boldsymbol{\phi}^0)).$$

A second-order expansion in $\hat{\alpha}_i(\boldsymbol{\theta}^0, \boldsymbol{\phi}^0)$ around α_i^0 yields (e.g., Rilstone, Srivastava and Ullah, 1996)

$$\hat{\alpha}_i(\boldsymbol{\theta}^0, \boldsymbol{\phi}^0) - \alpha_i^0 = \frac{\psi_i}{\sqrt{T}} + \frac{\beta_i}{T} + o_P(T^{-1}) \quad (20)$$

where the expressions of the influence function ψ_i and second-order bias β_i are given in Assumption 2. To bound the remainder term uniformly over i , we use that

$$\sup_i \|\hat{\alpha}_i(\boldsymbol{\theta}^0, \boldsymbol{\phi}^0) - \alpha_i^0\| \leq \sup_i \|\hat{\alpha}_i(\boldsymbol{\theta}^0, \boldsymbol{\phi}^0) - \hat{\alpha}_i(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\phi}})\| + \sup_i \|\hat{\alpha}_i - \alpha_i^0\| \rightarrow_P 0,$$

by the triangle inequality, a.s. continuity of $(\boldsymbol{\theta}, \boldsymbol{\phi}) \mapsto \hat{\alpha}_i(\boldsymbol{\theta}, \boldsymbol{\phi})$ and Assumption 2(i).

3. Second-Order Expansion of the Gradients $\sum_{t=1}^T \ell_{1it}^\theta(\boldsymbol{\theta}^0, \hat{\alpha}_i(\boldsymbol{\theta}^0, \boldsymbol{\phi}^0))$ and $\ell_{2i}^\phi(\boldsymbol{\phi}^0, \hat{\alpha}_i(\boldsymbol{\theta}^0, \boldsymbol{\phi}^0))$:

a second-order expansion of $\sum_{t=1}^T \ell_{1it}^\theta(\boldsymbol{\theta}^0, \hat{\alpha}_i(\boldsymbol{\theta}^0, \boldsymbol{\phi}^0))$ in $\hat{\alpha}_i(\boldsymbol{\theta}^0, \boldsymbol{\phi}^0)$ around α_i^0 yields

$$\frac{1}{T} \sum_{t=1}^T \ell_{1it}^\theta(\boldsymbol{\theta}^0, \hat{\alpha}_i(\boldsymbol{\theta}^0, \boldsymbol{\phi}^0)) = \frac{\varphi_{\theta i}}{\sqrt{T}} + \frac{b_{\theta i}}{T} + o_P(T^{-1}),$$

where the expressions of the influence function φ_{θ_i} , and second-order bias b_{θ_i} are given in Assumption 2, and we use (20). A similar analysis for $\ell_{2i}^\phi(\boldsymbol{\phi}^0, \hat{\alpha}_i(\boldsymbol{\theta}^0, \boldsymbol{\phi}^0))$ gives

$$\ell_{2i}^\phi(\boldsymbol{\phi}^0, \hat{\alpha}_i(\boldsymbol{\theta}^0, \boldsymbol{\phi}^0)) = \varphi_{\phi_i} + \frac{b_{\phi_i}}{T} + o_P(T^{-1}),$$

where the expression of φ_{ϕ_i} and b_{ϕ_i} are given in Assumption 2, and we use again (20).

Step 2: Asymptotic Distributions of $\hat{\boldsymbol{\theta}}$ and $\hat{\boldsymbol{\phi}}$. From the expansions of the gradients in the previous section

$$\sum_{i=1}^N \sum_{t=1}^T \ell_{1it}^\theta(\boldsymbol{\theta}^0, \hat{\alpha}_i(\boldsymbol{\theta}^0, \boldsymbol{\phi}^0)) = O_P(\sqrt{NT} \vee N),$$

and

$$\sum_{i=1}^N \ell_{2i}^\phi(\boldsymbol{\phi}^0, \hat{\alpha}_i(\boldsymbol{\theta}^0, \boldsymbol{\phi}^0)) = O_P(\sqrt{N}),$$

if $N = O(T^2)$. Here we use that $\sum_{i=1}^N \varphi_{\theta_i} = O_P(\sqrt{N})$, $\sum_{i=1}^N b_{\theta_i} = O_P(N)$, $\sum_{i=1}^N \varphi_{\phi_i} = O_P(\sqrt{N})$, and $\sum_{i=1}^N b_{\phi_i} = O_P(N)$ (e.g., Fernández-Val, 2005).

We need to consider two different cases because the asymptotic distribution of $\hat{\boldsymbol{\phi}}$ is non-degenerate under a wider range of sequences for N and T than the distribution of $\hat{\boldsymbol{\theta}}$. These cases are

1. $N = O(T)$: combining the previous results with (19) yields $\sqrt{NT}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^0) = O_P(1)$ and $\sqrt{N}(\hat{\boldsymbol{\phi}} - \boldsymbol{\phi}^0) = O_P(1)$, such that

$$\begin{aligned} \frac{1}{N} \sum_{i=1}^N \begin{bmatrix} \frac{1}{T} \sum_{t=1}^T \frac{d\ell_{1it}^\theta(\bar{\boldsymbol{\theta}}, \hat{\alpha}_i(\bar{\boldsymbol{\theta}}, \bar{\boldsymbol{\phi}}))}{d\boldsymbol{\theta}^\top} & \frac{1}{\sqrt{T}} \sum_{t=1}^T \frac{d\ell_{1it}^\theta(\bar{\boldsymbol{\theta}}, \hat{\alpha}_i(\bar{\boldsymbol{\theta}}, \bar{\boldsymbol{\phi}}))}{d\boldsymbol{\phi}^\top} \\ \frac{1}{\sqrt{T}} \frac{d\ell_{2i}^\phi(\bar{\boldsymbol{\phi}}, \hat{\alpha}_i(\bar{\boldsymbol{\theta}}, \bar{\boldsymbol{\phi}}))}{d\boldsymbol{\theta}^\top} & \frac{d\ell_{2i}^\phi(\bar{\boldsymbol{\phi}}, \hat{\alpha}_i(\bar{\boldsymbol{\theta}}, \bar{\boldsymbol{\phi}}))}{d\boldsymbol{\phi}^\top} \end{bmatrix} \begin{bmatrix} \sqrt{NT}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^0) \\ \sqrt{N}(\hat{\boldsymbol{\phi}} - \boldsymbol{\phi}^0) \end{bmatrix} \\ = \frac{-1}{\sqrt{N}} \sum_{i=1}^N \begin{bmatrix} \frac{1}{\sqrt{T}} \sum_{t=1}^T \ell_{1it}^\theta(\boldsymbol{\theta}^0, \hat{\alpha}_i(\boldsymbol{\theta}^0, \boldsymbol{\phi}^0)) \\ \ell_{2i}^\phi(\boldsymbol{\phi}^0, \hat{\alpha}_i(\boldsymbol{\theta}^0, \boldsymbol{\phi}^0)) \end{bmatrix}. \quad (21) \end{aligned}$$

Note that the off-diagonal elements of the matrix in the LHS are of order $O_P(T^{-1})$, whereas the diagonal elements are of order $O_P(1)$. Hence

$$\begin{aligned} \frac{1}{N} \sum_{i=1}^N \left[\begin{array}{cc} \frac{1}{T} \sum_{t=1}^T \frac{d\ell_{1it}^{\theta}(\bar{\theta}, \hat{\alpha}_i(\bar{\theta}, \bar{\phi}))}{d\theta^T} & \frac{1}{\sqrt{T}} \sum_{t=1}^T \frac{d\ell_{1it}^{\theta}(\bar{\theta}, \hat{\alpha}_i(\bar{\theta}, \bar{\phi}))}{d\phi^T} \\ \frac{1}{\sqrt{T}} \frac{d\ell_{2i}^{\phi}(\bar{\phi}, \hat{\alpha}_i(\bar{\theta}, \bar{\phi}))}{d\theta^T} & \frac{d\ell_{2i}^{\phi}(\bar{\phi}, \hat{\alpha}_i(\bar{\theta}, \bar{\phi}))}{d\phi^T} \end{array} \right] \\ = \frac{1}{N} \sum_{i=1}^N \left[\begin{array}{cc} \frac{1}{T} \sum_{t=1}^T \frac{d\ell_{1it}^{\theta}(\bar{\theta}, \hat{\alpha}_i(\bar{\theta}, \bar{\phi}))}{d\theta^T} & 0 \\ 0 & \frac{d\ell_{2i}^{\phi}(\bar{\phi}, \hat{\alpha}_i(\bar{\theta}, \bar{\phi}))}{d\phi^T} \end{array} \right] + O_P\left(\frac{1}{\sqrt{T}}\right). \end{aligned}$$

2. $T = o(N)$: combining the previous results with (19) yields $T(\hat{\theta} - \theta^0) = O_P(1)$ and $\sqrt{N}(\hat{\phi} - \phi^0) = O_P(1)$, such that

$$\begin{aligned} \frac{1}{N} \sum_{i=1}^N \left[\begin{array}{cc} \frac{1}{T} \sum_{t=1}^T \frac{d\ell_{1it}^{\theta}(\bar{\theta}, \hat{\alpha}_i(\bar{\theta}, \bar{\phi}))}{d\theta^T} & \frac{1}{\sqrt{N}} \sum_{t=1}^T \frac{d\ell_{1it}^{\theta}(\bar{\theta}, \hat{\alpha}_i(\bar{\theta}, \bar{\phi}))}{d\phi^T} \\ \frac{\sqrt{N}}{T} \frac{d\ell_{2i}^{\phi}(\bar{\phi}, \hat{\alpha}_i(\bar{\theta}, \bar{\phi}))}{d\theta^T} & \frac{d\ell_{2i}^{\phi}(\bar{\phi}, \hat{\alpha}_i(\bar{\theta}, \bar{\phi}))}{d\phi^T} \end{array} \right] \begin{bmatrix} T(\hat{\theta} - \theta^0) \\ \sqrt{N}(\hat{\phi} - \phi^0) \end{bmatrix} \\ = \frac{-1}{\sqrt{N}} \sum_{i=1}^N \left[\begin{array}{c} \frac{1}{\sqrt{N}} \sum_{t=1}^T \ell_{1it}^{\theta}(\theta^0, \hat{\alpha}_i(\theta^0, \phi^0)) \\ \ell_{2i}^{\phi}(\phi^0, \hat{\alpha}_i(\theta^0, \phi^0)) \end{array} \right]. \quad (22) \end{aligned}$$

Similar to the other case, some of the off-diagonal elements of the matrix in the LHS are of smaller order than the diagonal elements. Hence, if $N = O(T^2)$,

$$\begin{aligned} \frac{1}{N} \sum_{i=1}^N \left[\begin{array}{cc} \frac{1}{T} \sum_{t=1}^T \frac{d\ell_{1it}^{\theta}(\bar{\theta}, \hat{\alpha}_i(\bar{\theta}, \bar{\phi}))}{d\theta^T} & \frac{1}{\sqrt{N}} \sum_{t=1}^T \frac{d\ell_{1it}^{\theta}(\bar{\theta}, \hat{\alpha}_i(\bar{\theta}, \bar{\phi}))}{d\phi^T} \\ \frac{\sqrt{N}}{T} \frac{d\ell_{2i}^{\phi}(\bar{\phi}, \hat{\alpha}_i(\bar{\theta}, \bar{\phi}))}{d\theta^T} & \frac{d\ell_{2i}^{\phi}(\bar{\phi}, \hat{\alpha}_i(\bar{\theta}, \bar{\phi}))}{d\phi^T} \end{array} \right] \\ = \frac{1}{N} \sum_{i=1}^N \left[\begin{array}{cc} \frac{1}{T} \sum_{t=1}^T \frac{d\ell_{1it}^{\theta}(\bar{\theta}, \hat{\alpha}_i(\bar{\theta}, \bar{\phi}))}{d\theta^T} & 0 \\ \frac{\sqrt{N}}{T} \frac{d\ell_{2i}^{\phi}(\bar{\phi}, \hat{\alpha}_i(\bar{\theta}, \bar{\phi}))}{d\theta^T} & \frac{d\ell_{2i}^{\phi}(\bar{\phi}, \hat{\alpha}_i(\bar{\theta}, \bar{\phi}))}{d\phi^T} \end{array} \right] + o_P(1). \end{aligned}$$

Note that

$$\begin{aligned} \text{plim}_{N,T \rightarrow \infty} \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \frac{d\ell_{1it}^{\boldsymbol{\theta}}(\bar{\boldsymbol{\theta}}, \hat{\alpha}_i(\bar{\boldsymbol{\theta}}, \bar{\boldsymbol{\phi}}))}{d\boldsymbol{\theta}^\top} &= \text{plim}_{N,T \rightarrow \infty} \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \left[\ell_{1it}^{\boldsymbol{\theta}\boldsymbol{\theta}} - \ell_{1it}^{\boldsymbol{\theta}\alpha} \frac{\sum_{t=1}^T \ell_{1it}^{\alpha\boldsymbol{\theta}}}{\sum_{t=1}^T \ell_{1it}^{\alpha\alpha}} \right], \\ \text{plim}_{N,T \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \frac{d\ell_{2i}^{\boldsymbol{\phi}}(\bar{\boldsymbol{\theta}}, \hat{\alpha}_i(\bar{\boldsymbol{\theta}}, \bar{\boldsymbol{\phi}}))}{d\boldsymbol{\theta}^\top} &= -\text{plim}_{N,T \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \ell_{2i}^{\boldsymbol{\phi}\alpha} \frac{\sum_{t=1}^T \ell_{1it}^{\alpha\boldsymbol{\theta}}}{\sum_{t=1}^T \ell_{1it}^{\alpha\alpha}}, \\ \text{plim}_{N \rightarrow \infty} \frac{1}{N} \frac{d\ell_{2i}^{\boldsymbol{\phi}}(\bar{\boldsymbol{\phi}}, \hat{\alpha}_i(\bar{\boldsymbol{\theta}}, \bar{\boldsymbol{\phi}}))}{d\boldsymbol{\phi}^\top} &= \text{plim}_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \ell_{2i}^{\boldsymbol{\phi}\boldsymbol{\phi}}, \end{aligned}$$

because $\sqrt{NT}(\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^0) = O_P(1)$, $\sqrt{N}(\bar{\boldsymbol{\phi}} - \boldsymbol{\phi}^0) = O_P(1)$, and

$$\hat{\alpha}_i(\bar{\boldsymbol{\theta}}, \bar{\boldsymbol{\phi}}) = \hat{\alpha}_i(\boldsymbol{\theta}^0, \boldsymbol{\phi}^0) + \hat{\alpha}_i^{\boldsymbol{\theta}}(\check{\boldsymbol{\theta}}, \check{\boldsymbol{\phi}})^\top (\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^0) + \hat{\alpha}_i^{\boldsymbol{\phi}}(\check{\boldsymbol{\theta}}, \check{\boldsymbol{\phi}})^\top (\bar{\boldsymbol{\phi}} - \boldsymbol{\phi}^0) = \alpha_i^0 + O_P(T^{-1/2}),$$

where $(\check{\boldsymbol{\theta}}, \check{\boldsymbol{\phi}})$ lies between $(\bar{\boldsymbol{\theta}}, \bar{\boldsymbol{\phi}})$ and $(\boldsymbol{\theta}^0, \boldsymbol{\phi}^0)$.

Starting from (21) and combining the results from the previous steps yields

$$0 = \frac{1}{\sqrt{N}} \sum_{i=1}^N \varphi_{\boldsymbol{\theta}i} + \bar{\mathbf{H}}_{\boldsymbol{\theta}} \sqrt{NT} \left(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^0 + \frac{\bar{\mathbf{H}}_{\boldsymbol{\theta}}^{-1} \bar{b}_{\boldsymbol{\theta}}}{T} \right) + o_P \left(\sqrt{\frac{N}{T}} \right) + o_P \left(\sqrt{NT} (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^0) \right)$$

Hence, under asymptotic sequences where $N = O(T)$,

$$\sqrt{NT} \left(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}^0 + \frac{\bar{\mathbf{H}}_{\boldsymbol{\theta}}^{-1} \bar{b}_{\boldsymbol{\theta}}}{T} \right) \rightarrow_d \bar{\mathbf{H}}_{\boldsymbol{\theta}}^{-1} \mathcal{N}(0, \bar{\mathcal{O}}_{\boldsymbol{\theta}}).$$

Similarly, starting from (22), combining the results from the previous steps yields and using block matrix inversion

$$\begin{aligned} 0 &= \frac{1}{\sqrt{N}} \sum_{i=1}^N \varphi_{\boldsymbol{\phi}i} + \bar{\mathbf{H}}_{\boldsymbol{\phi}} \sqrt{N} \left(\hat{\boldsymbol{\phi}} - \boldsymbol{\phi}^0 + \frac{\bar{\mathbf{H}}_{\boldsymbol{\phi}}^{-1} (\bar{b}_{\boldsymbol{\phi}} + \bar{\mathbf{H}}_{\boldsymbol{\phi}\boldsymbol{\theta}} \bar{\mathbf{H}}_{\boldsymbol{\theta}}^{-1} \bar{b}_{\boldsymbol{\theta}})}{T} \right) \\ &\quad + o_P \left(\frac{\sqrt{N}}{T} \right) + o_P \left(\sqrt{N} (\hat{\boldsymbol{\phi}} - \boldsymbol{\phi}^0) \right). \end{aligned}$$

Hence, under asymptotic sequences where $N = O(T^2)$,

$$\sqrt{N} \left(\hat{\boldsymbol{\phi}} - \boldsymbol{\phi}^0 + \frac{\bar{\mathbf{H}}_{\boldsymbol{\phi}}^{-1} (\bar{b}_{\boldsymbol{\phi}} + \bar{\mathbf{H}}_{\boldsymbol{\phi}\boldsymbol{\theta}} \bar{\mathbf{H}}_{\boldsymbol{\theta}}^{-1} \bar{b}_{\boldsymbol{\theta}})}{T} \right) \rightarrow_d \bar{\mathbf{H}}_{\boldsymbol{\phi}}^{-1} \mathcal{N}(0, \bar{\mathcal{O}}_{\boldsymbol{\phi}}).$$

■

Application to ILM Estimator

In the case of the integrated loss minimization (ILM) estimator for longitudinal data

$$\ell_{1it}(\boldsymbol{\theta}, \alpha_i) = (y_{it} - \alpha_i)(u_{it}(\boldsymbol{\theta}, \alpha_i) - 0.5) + \mathbf{x}_{it}^T \boldsymbol{\theta} [\bar{\mathbf{B}} - \mathbf{B}(u_{it}(\boldsymbol{\theta}, \alpha_i))], \quad (23)$$

where $\mathbf{B}(u) = \int_0^u \mathbf{b}(x)dx$, $\bar{\mathbf{B}} = \int_0^1 \mathbf{B}(u)du$, and $u_{it}(\boldsymbol{\theta}, \alpha_i)$ is the solution in u to

$$y_{it} - \alpha_i = \mathbf{x}_{it}^T \boldsymbol{\theta} \mathbf{b}(u);$$

and

$$\ell_{2i}(\boldsymbol{\phi}, \alpha_i) = \alpha_i(v_i(\boldsymbol{\phi}, \alpha_i) - 0.5) + \mathbf{z}_i^T \boldsymbol{\phi} [\bar{\mathbf{C}} - \mathbf{C}(v_i(\boldsymbol{\phi}, \alpha_i))], \quad (24)$$

where $\mathbf{C}(v) = \int_0^v \mathbf{c}(x)dx$, $\bar{\mathbf{C}} = \int_0^1 \mathbf{C}(v)dv$, and $v_i(\boldsymbol{\phi}, \alpha_i)$ is the solution in v to

$$\alpha_i = \mathbf{z}_i^T \boldsymbol{\phi} \mathbf{c}(v).$$

To apply the previous analysis to the ILM estimator, we need to adapt the notation to the case where $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ can be matrix-valued parameters by vectorizing $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ in all the expressions. For example, $\ell_{1it}^{\boldsymbol{\theta}}(\boldsymbol{\theta}, \alpha_i)$ becomes $\partial \ell_{1it}(\boldsymbol{\theta}, \alpha_i) / \partial \text{vec}(\boldsymbol{\theta})$ and $(\hat{\boldsymbol{\phi}} - \boldsymbol{\phi}^0)$ becomes $\text{vec}(\hat{\boldsymbol{\phi}} - \boldsymbol{\phi}^0)$.

Proof of Theorem 1. We divide the proof in three parts: identification, consistency, and asymptotic distribution.

Part 1: Identification. The identification analysis has several steps. First, we show identification of the quantile regression level-1 coefficient functions of the time-varying covariates and an aggregated individual effect that contains the quantile regression level-1 coefficient functions of the time invariant covariates and individual effects, using within-group variation. Second, we separate the quantile regression level-1 coefficient functions of the time invariant covariates from the individual effects using between-group variation

and the location normalization in the distribution of the individual effects. Third, we show identification of the quantile regression level-2 coefficient functions using between-group variation of the individual effects. Fourth, we show that the parametric coefficient functions are identified from the quantile regression coefficient functions as the solutions to linear systems of equations. These solutions exist and are unique by Assumption 1(iv).

In this part, it is convenient to partition $\mathbf{x}_{it} = (\mathbf{x}_{1i}, \mathbf{x}_{2it})$, where \mathbf{x}_{1i} contains the time-invariant components including the constant and \mathbf{x}_{2it} contains the time-varying covariates, $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)$ and $\mathbf{b}(u) = [\mathbf{b}_1(u), \mathbf{b}_2(u)]$, such that $\mathbf{x}_{it}^T \boldsymbol{\theta} \mathbf{b}(u) = \mathbf{x}_{1i}^T \boldsymbol{\theta}_1 \mathbf{b}_1(u) + \mathbf{x}_{2it}^T \boldsymbol{\theta}_2 \mathbf{b}_2(u)$; and use the parametrization $\alpha_i(u) = \alpha_i + \mathbf{x}_{1i}^T \boldsymbol{\theta}_1 \mathbf{b}_1(u)$, $\boldsymbol{\theta}_2(u) := \boldsymbol{\theta}_2 \mathbf{b}(u)$ and $\boldsymbol{\phi}(v) := \boldsymbol{\phi} \mathbf{c}(v)$. Then, we can express the elements of the objective function as

$$\ell_{1it}(\boldsymbol{\theta}, \alpha_i) = \int_0^1 \rho_u(y_{it} - \alpha_i(u) - \mathbf{x}_{2it}^T \boldsymbol{\theta}_2(u)) du,$$

and

$$\ell_{2i}(\boldsymbol{\phi}, \alpha_i) = \int_0^1 \rho_v(\alpha_i - \mathbf{z}_i^T \boldsymbol{\phi}(v)) dv,$$

where $\rho_\tau(t) = t(\tau - I(t \leq 0))$ for any $\tau \in [0, 1]$.

For any u , identification of $\boldsymbol{\theta}_2^0(u) = \boldsymbol{\theta}_2^0 \mathbf{b}_2(u)$ and $\alpha_i^0(u) = \alpha_i^0 + \mathbf{x}_{1i}^T \boldsymbol{\theta}_1^0 \mathbf{b}_1(u)$ follows from Assumption 1(i)–(iii) by standard arguments for quantile regression since

$$f_{Y_{it}}(\alpha_i^0(u) + \mathbf{x}_{2it}^T \boldsymbol{\theta}_2^0(u) \mid \alpha_i, \mathbf{x}_{it}) = \frac{1}{\mathbf{x}_{it}^T \boldsymbol{\theta}^0 \mathbf{b}'(u)} < \infty,$$

where $f_{Y_{it}}(\cdot \mid \alpha_i, \mathbf{x}_{it})$ is the distribution of Y_{it} conditional on α_i and \mathbf{x}_{it} . Then, to identify α_i^0 and $\boldsymbol{\theta}_1^0(u) = \boldsymbol{\theta}_1^0 \mathbf{b}_1(u)$ from $\alpha_i^0(u)$ we use the within expectation

$$E[\alpha_i^0(u) \mid \alpha_i, \mathbf{x}_{1i}] = \alpha_i^0 + \mathbf{x}_{1i}^T E[\boldsymbol{\theta}_1(U_{it})].$$

The between variation of this expression identifies $\text{plim}_{N \rightarrow \infty} N^{-1} \sum_{i=1}^N \alpha_i^0 + E(\boldsymbol{\theta}_1(U_{it}))$ and $E[\boldsymbol{\theta}_{1,-1}(U_{it})]$, where $\theta_1(u)$ and $\boldsymbol{\theta}_{1,-1}(u)$ are the intercept and vector of slopes of $\boldsymbol{\theta}_1(u)$,

respectively. To separate $\text{plim}_{N \rightarrow \infty} N^{-1} \sum_{i=1}^N \alpha_i^0$ from $E(\theta_1(U_{it}))$ we use the normalization in Assumption 1(i). Thus,

$$\text{plim}_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \alpha_i^0(u) = \text{plim}_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \alpha_i^0 + \text{plim}_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \mathbf{x}_{1i}^T E[\theta_1(U_{it})]$$

pins down $E[\theta_1(U_{it})]$ since $\text{plim}_{N \rightarrow \infty} N^{-1} \sum_{i=1}^N \alpha_i^0 = 0$. Then, α_i^0 is identified by

$$\alpha_i^0 = E[\alpha_i^0(u) \mid \alpha_i, \mathbf{x}_{1i}] - \mathbf{x}_{1i}^T E[\theta_1(U_{it})],$$

and $\theta_1(u)$ is identified by

$$\theta_1(u) = \left[\text{plim}_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \mathbf{x}_{1i} \mathbf{x}_{1i}^T \right]^{-1} \text{plim}_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \mathbf{x}_{1i} (\alpha_i^0(u) - \alpha_i^0).$$

Finally, identification of $\phi^0(v) = \phi^0 \mathbf{c}(v)$ for any v follows from identification of α_i^0 by standard arguments for quantile regression since

$$f_{\alpha_i}(\mathbf{z}_i^T \phi^0(v) \mid \mathbf{z}_i) = \frac{1}{\mathbf{z}_i^T \phi^0 \mathbf{c}'(v_i)} < \infty,$$

where $f_{\alpha_i}(\cdot \mid \mathbf{z}_i)$ is the distribution of α_i conditional on \mathbf{z}_i .

To show identification of θ^0 and ϕ^0 from identification of $\theta^0(u) = [\theta_1^0(u), \theta_2^0(u)]$ and $\phi^0(v)$, we use Assumption 1(iv). Let $\theta^0(u_1, \dots, u_{d_b}) = [\theta^0(u_1), \dots, \theta^0(u_{d_b})]$ and $\phi^0(v_1, \dots, v_{d_c}) = [\phi^0(v_1), \dots, \phi^0(v_{d_c})]$. Then, we have two systems of linear equations

$$\theta^0(u_1, \dots, u_{d_b}) = \theta^0 \mathbf{b}(u_1, \dots, u_{d_b}), \quad \phi^0(v_1, \dots, v_{d_c}) = \phi^0 \mathbf{c}(v_1, \dots, v_{d_c}),$$

which have as unique solutions

$$\theta^0 = \theta^0(u_1, \dots, u_{d_b}) \mathbf{b}(u_1, \dots, u_{d_b})^{-1}, \quad \phi^0 = \phi^0(v_1, \dots, v_{d_c}) \mathbf{c}(v_1, \dots, v_{d_c})^{-1}.$$

Part 2: Consistency. The consistency of all the estimators can be established sequentially. Let

$$\hat{L}_i(\boldsymbol{\theta}, \boldsymbol{\phi}, \alpha_i) = \frac{1}{T} \left[\sum_{t=1}^T \ell_{1it}(\boldsymbol{\theta}, \alpha_i) + \ell_{2i}(\boldsymbol{\phi}, \alpha_i) \right],$$

and

$$L_{1i}(\boldsymbol{\theta}, \alpha_i) = \text{plim}_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \ell_{1it}(\boldsymbol{\theta}, \alpha_i).$$

Consistency of $\hat{\boldsymbol{\theta}}$ and uniform consistency of $\hat{\alpha}_i$, i.e., $\hat{\boldsymbol{\theta}} \rightarrow_P \boldsymbol{\theta}^0$ and $\sup_i |\hat{\alpha}_i - \alpha_i^0| \rightarrow_P 0$, follows from standard arguments for quantile regression using the parametrization of step 1, together with the fact that the penalty term of the objective function is asymptotically negligible, i.e.,

$$\sup_i \left[\sup_{\{\boldsymbol{\theta}, \boldsymbol{\phi}, \alpha_i\}} \left| \hat{L}_i(\boldsymbol{\theta}, \boldsymbol{\phi}, \alpha_i) - L_{1i}(\boldsymbol{\theta}, \alpha_i) \right| \right] \rightarrow_P 0,$$

where we use that

$$\sup_i \left[\sup_{\{\boldsymbol{\phi}, \alpha_i\}} \left| \frac{1}{T} \ell_{2i}(\boldsymbol{\phi}, \alpha_i) \right| \right] \rightarrow_P 0,$$

and the triangle inequality.

Then, $\hat{\boldsymbol{\phi}} \rightarrow_P \boldsymbol{\phi}^0$ also follows from standard arguments for quantile regression using that

$$\sup_{\boldsymbol{\phi}} \left| \frac{1}{N} \sum_{i=1}^N \ell_{2i}(\boldsymbol{\phi}, \hat{\alpha}_i) - L_2(\boldsymbol{\phi}) \right| \rightarrow_P 0, \quad L_2(\boldsymbol{\phi}) = \text{plim}_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \ell_{2i}(\boldsymbol{\phi}, \alpha_i^0),$$

where we use that, by the triangle inequality,

$$\begin{aligned} \left| \frac{1}{N} \sum_{i=1}^N \ell_{2i}(\boldsymbol{\phi}, \hat{\alpha}_i) - L_2(\boldsymbol{\phi}) \right| &\leq \left| \frac{1}{N} \sum_{i=1}^N \ell_{2i}(\boldsymbol{\phi}, \hat{\alpha}_i) - \frac{1}{N} \sum_{i=1}^N \ell_{2i}(\boldsymbol{\phi}, \alpha_i^0) \right| \\ &\quad + \left| \frac{1}{N} \sum_{i=1}^N \ell_{2i}(\boldsymbol{\phi}, \alpha_i^0) - L_2(\boldsymbol{\phi}) \right|, \end{aligned}$$

continuity of $\alpha_i \mapsto \ell_{2i}(\boldsymbol{\phi}, \alpha_i)$, and uniform consistency of $\hat{\alpha}_i$.

Part 3: Asymptotic Distribution. The derivation of the asymptotic distribution follows from Theorem 2 after verifying Assumption 2(ii)-(viii) and evaluating the expressions for the ILM estimator.

We start by obtaining the derivatives of $(\boldsymbol{\theta}, \alpha_i) \mapsto \ell_{1it}(\boldsymbol{\theta}, \alpha_i)$ and $\boldsymbol{\phi} \mapsto \ell_{2i}(\boldsymbol{\phi}, \alpha_i)$. We report only the derivatives that show up in the terms of the asymptotic expansions. A similar analysis applies to the additional terms that appear in the remainder terms. Direct calculations from (23) yield

$$\begin{aligned}
\ell_{1it}^\alpha(\boldsymbol{\theta}, \alpha_i) &= -(u_{it}(\boldsymbol{\theta}, \alpha_i) - 0.5), \quad \ell_{1it}^{\alpha\alpha}(\boldsymbol{\theta}, \alpha_i) = -u_{it}^\alpha(\boldsymbol{\theta}, \alpha_i) = [\mathbf{x}_{it}^\top \boldsymbol{\theta} \mathbf{b}'(u_{it}(\boldsymbol{\theta}, \alpha_i))]^{-1}, \\
\ell_{1it}^{\alpha\alpha\alpha}(\boldsymbol{\theta}, \alpha_i) &= -\mathbf{x}_{it}^\top \boldsymbol{\theta} \mathbf{b}''(u_{it}(\boldsymbol{\theta}, \alpha_i)) \ell_{1it}^{\alpha\alpha}(\boldsymbol{\theta}, \alpha_i)^3, \\
\ell_{1it}^\theta(\boldsymbol{\theta}, \alpha_i) &= [\bar{\mathbf{B}} - \mathbf{B}(u_{it}(\boldsymbol{\theta}, \alpha_i))] \otimes \mathbf{x}_{it}, \\
\ell_{1it}^{\theta\theta}(\boldsymbol{\theta}, \alpha_i) &= [\ell_{1it}^{\alpha\alpha}(\boldsymbol{\theta}, \alpha_i) \mathbf{b}(u_{it}(\boldsymbol{\theta}, \alpha_i)) \mathbf{b}(u_{it}(\boldsymbol{\theta}, \alpha_i))^\top] \otimes [\mathbf{x}_{it} \mathbf{x}_{it}^\top], \\
\ell_{1it}^{\theta\alpha}(\boldsymbol{\theta}, \alpha_i) &= [\ell_{1it}^{\alpha\alpha}(\boldsymbol{\theta}, \alpha_i) \mathbf{b}(u_{it}(\boldsymbol{\theta}, \alpha_i))] \otimes \mathbf{x}_{it}, \\
\ell_{1it}^{\theta\alpha\alpha}(\boldsymbol{\theta}, \alpha_i) &= [\ell_{1it}^{\alpha\alpha\alpha}(\boldsymbol{\theta}, \alpha_i) \mathbf{b}(u_{it}(\boldsymbol{\theta}, \alpha_i)) - \ell_{1it}^{\alpha\alpha}(\boldsymbol{\theta}, \alpha_i)^2 \mathbf{b}'(u_{it}(\boldsymbol{\theta}, \alpha_i))] \otimes \mathbf{x}_{it}.
\end{aligned}$$

Analogously, direct calculations from (24) yield

$$\begin{aligned}
\ell_{2i}^\alpha(\boldsymbol{\phi}, \alpha_i) &= v_i(\boldsymbol{\phi}, \alpha_i) - 0.5, \quad \ell_{2i}^{\alpha\alpha}(\boldsymbol{\phi}, \alpha_i) = v_i^\alpha(\boldsymbol{\phi}, \alpha_i) = [\mathbf{z}_i^\top \boldsymbol{\phi} \mathbf{c}'(v_i(\boldsymbol{\phi}, \alpha_i))]^{-1}, \\
\ell_{2i}^{\alpha\alpha\alpha}(\boldsymbol{\phi}, \alpha_i) &= -\mathbf{z}_i^\top \boldsymbol{\phi} \mathbf{c}''(v_i(\boldsymbol{\phi}, \alpha_i)) \ell_{2i}^{\alpha\alpha}(\boldsymbol{\phi}, \alpha_i)^3, \\
\ell_{2i}^\phi(\boldsymbol{\phi}, \alpha_i) &= [\bar{\mathbf{C}} - \mathbf{C}(v_i(\boldsymbol{\phi}, \alpha_i))] \otimes \mathbf{z}_i, \\
\ell_{2i}^{\phi\phi}(\boldsymbol{\phi}, \alpha_i) &= [\ell_{2i}^{\alpha\alpha}(\boldsymbol{\phi}, \alpha_i) \mathbf{c}(v_i(\boldsymbol{\phi}, \alpha_i)) \mathbf{c}(v_i(\boldsymbol{\phi}, \alpha_i))^\top] \otimes [\mathbf{z}_i \mathbf{z}_i^\top], \\
\ell_{2i}^{\phi\alpha}(\boldsymbol{\phi}, \alpha_i) &= -[\mathbf{c}(v_i(\boldsymbol{\phi}, \alpha_i)) \ell_{2i}^{\alpha\alpha}(\boldsymbol{\phi}, \alpha_i)] \otimes \mathbf{z}_i, \\
\ell_{2i}^{\phi\alpha\alpha}(\boldsymbol{\phi}, \alpha_i) &= -[\ell_{2i}^{\alpha\alpha\alpha}(\boldsymbol{\phi}, \alpha_i) \mathbf{c}(v_i(\boldsymbol{\phi}, \alpha_i)) + \ell_{2i}^{\alpha\alpha}(\boldsymbol{\phi}, \alpha_i)^2 \mathbf{c}'(v_i(\boldsymbol{\phi}, \alpha_i))] \otimes \mathbf{z}_i.
\end{aligned}$$

Assumption 2(ii) follows from Assumption 1(v) by inspection of the derivatives. Assumption 2(iii) holds trivially as the parameter spaces are $\mathcal{A} = \mathbb{R}$, $\Theta = \mathbb{R}^{d_x}$ and $\Phi = \mathbb{R}^{d_z}$,

and Assumptions 2(iv)-(v) follow directly from Assumption 1(vi)-(vii). Then, the asymptotic distributions follow from Theorem 2, replacing the derivatives above evaluated at the true parameter values in the expressions of the bias and variance given in Assumption 2, after dropping out some terms that are either asymptotically negligible or zero because $u_{it} \sim U_{it}$ and $v_i \sim V_i$ are independent. ■

Appendix B - Computation

An efficient algorithm has been implemented in the `qrcm` R package, which also includes all the necessary functions for model building, summary, predictions, testing, and plotting. The steps of the algorithm can be summarized as follows:

- step 0. Select starting values $\hat{\boldsymbol{\xi}}^{(0)} = (\hat{\boldsymbol{\theta}}^{(0)}, \hat{\boldsymbol{\phi}}^{(0)}, \hat{\boldsymbol{\alpha}}_N^{(0)})$.
- step 1. Given a current estimate $\hat{\boldsymbol{\alpha}}_N$, minimize $L(\boldsymbol{\theta}, \hat{\boldsymbol{\alpha}}_N)$ (equation 6) with respect to $\boldsymbol{\theta}$, and $L(\boldsymbol{\phi}, \hat{\boldsymbol{\alpha}}_N)$ (equation 7) with respect to $\boldsymbol{\phi}$. Equivalently, find the approximated zeroes of $\mathbf{G}_{\boldsymbol{\theta}}(\boldsymbol{\theta}, \hat{\boldsymbol{\alpha}}_N)$ (equation 13) and $\mathbf{G}_{\boldsymbol{\phi}}(\boldsymbol{\phi}, \hat{\boldsymbol{\alpha}}_N)$ (equation 14).
- step 2. Given a current estimate $(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\phi}})$, find the approximated zeroes of $G_{\alpha_i}(\alpha_i, \hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\phi}})$ (equation 15), for $i = 1, \dots, N$.

Steps 1-2 are repeated until a convergence criterion has been reached. In the `qrcm` package, the algorithm stops when, in two consecutive iterations, either the absolute difference in the estimated parameters, or the absolute change in the loss function defined by $L(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\phi}}, \hat{\boldsymbol{\alpha}}_N) = L(\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\alpha}}_N) + L(\hat{\boldsymbol{\phi}}, \hat{\boldsymbol{\alpha}}_N)$ (equation 12) is below a certain tolerance (default 0.00001).

Note that $\mathbf{G}_{\boldsymbol{\theta}}(\boldsymbol{\theta}, \boldsymbol{\alpha}_N)$, $\mathbf{G}_{\boldsymbol{\phi}}(\boldsymbol{\phi}, \boldsymbol{\alpha}_N)$, and $G_{\alpha_i}(\alpha_i, \boldsymbol{\theta}, \boldsymbol{\phi})$ are smooth functions of their arguments, and can be solved by using a standard Newton-type algorithm. In practice, a bisection algorithm is used to solve $G_{\alpha_i}(\alpha_i, \hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\phi}})$ (see section **B2** of this Appendix).

The table below summarizes the computation times required to estimate the models presented in simulation, using a desktop computer Intel(R) Core(TM) i7-4770 CPU @ 3.40GHz, RAM 8.00 GB, 64-bit Operating System.

	$N = 150$		$N = 300$	
	$T = 5$	$T = 10$	$T = 5$	$T = 10$
Simulation 1	0.7 (0.4–0.8)	1.6 (0.8–2.2)	1.3 (0.7–1.6)	2.7 (1.3–3.8)
Simulation 2	1.7 (1.5–1.8)	2.3 (2.1–2.5)	3.2 (3.0–3.4)	4.7 (4.4–5.0)

Summary statistics of computation times for simulations 1 and 2 described in Section 7. We report the median time (in seconds) and, in brackets, the interquartile range.

B1. Choosing the starting values

Selecting the starting points is fundamental, because not all values of the parameters correspond to a well-defined quantile function. If $\mathbf{x}_{it}^T \boldsymbol{\beta}(\cdot \mid \hat{\boldsymbol{\theta}}^{(0)})$ or $\mathbf{z}_i^T \boldsymbol{\gamma}(\cdot \mid \hat{\boldsymbol{\phi}}^{(0)})$ are not monotonically increasing functions, the algorithm may converge to a nonsense solution.

As shown by equations (13–15), the gradient only depends on the model parameters through $u_{it}(\boldsymbol{\theta}, \alpha_i)$ and $v_i(\boldsymbol{\phi}, \alpha_i)$, that correspond to the values of the cumulative distribution functions of $y_{it} - \alpha_i$ and α_i , respectively (equations 10 and 11). Given an initial estimate of (u_{it}, v_i) , $(\hat{u}_{it}^{(0)}, \hat{v}_i^{(0)})$, a starting value for $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ can be obtained by approximating model (4) by a linear regression model of the form

$$Y_{it} = \mathbf{x}_{it}^T \boldsymbol{\theta} \mathbf{b}(\hat{u}_{it}^{(0)}) + \mathbf{z}_i^T \boldsymbol{\phi} \mathbf{c}(\hat{v}_i^{(0)}) + \epsilon_{it},$$

in which $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ represent regression coefficients associated with the tensor products $\mathbf{x}_i \otimes \mathbf{b}(\hat{u}_{it}^{(0)})$ and $\mathbf{z}_i \otimes \mathbf{c}(\hat{v}_i^{(0)})$.

To obtain the initial values $\hat{u}_{it}^{(0)}$ and $\hat{v}_i^{(0)}$, we proceed as follows: first, we compute a preliminary estimate $\hat{\boldsymbol{\alpha}}_N^{(0)} = \{\hat{\alpha}_1^{(0)}, \dots, \hat{\alpha}_N^{(0)}\}$ of $\boldsymbol{\alpha}_N = \{\alpha_1, \dots, \alpha_N\}$ using the cluster medians. Then, we use the `pch` R package to estimate nonparametrically the cumulative distribution function of $y_{it} - \hat{\alpha}_i^{(0)}$ and, separately, that of $\hat{\boldsymbol{\alpha}}_N^{(0)}$. The fitted values are used to define $\hat{u}_{it}^{(0)}$ and $\hat{v}_i^{(0)}$.

The algorithm appears very stable and, in the simulation and data analysis conducted in this paper, never failed to converge.

B2. Evaluating $\hat{u}_{it} := u_{it}(\hat{\boldsymbol{\theta}}, \hat{\alpha}_i)$ and $\hat{v}_i := v_i(\hat{\boldsymbol{\phi}}, \hat{\alpha}_i)$

At any current estimate $\hat{\boldsymbol{\xi}} = (\hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\phi}}, \hat{\boldsymbol{\alpha}}_N)$, evaluating the loss function and its derivatives requires computing the values $(\hat{u}_{it}, \hat{v}_i)$ such that

$$\mathbf{x}_{it}^T \boldsymbol{\beta}(\hat{u}_{it} \mid \hat{\boldsymbol{\theta}}) = \mathbf{x}_{it}^T \hat{\boldsymbol{\theta}} \mathbf{b}(\hat{u}_{it}) = y_{it} - \hat{\alpha}_i,$$

$$\mathbf{z}_i^T \boldsymbol{\gamma}(\hat{v}_i \mid \hat{\boldsymbol{\phi}}) = \mathbf{z}_i^T \hat{\boldsymbol{\phi}} \mathbf{c}(\hat{v}_i) = \hat{\alpha}_i.$$

These values are not generally available in closed form, and are computed using a bisection algorithm. For example, to compute \hat{u}_{it} , we proceed as follows: (i) start with $\hat{u}_{it}^{(0)} = 0.5$; (ii) for $s = 1, 2, 3, \dots$, define

$$\hat{u}_{it}^{(s)} = \hat{u}_{it}^{(s-1)} + \frac{1}{2^{s+1}} \text{sign} \left(y_{it} - \hat{\alpha}_i - \mathbf{x}_{it}^T \hat{\boldsymbol{\theta}} \mathbf{b}(\hat{u}_{it}^{(s-1)}) \right).$$

By bisecting the unit interval 20 times, it is possible to achieve a precision of about 10^{-6} .

To solve equation (15), which includes both u_{it} and v_i , two nested bisections must be implemented. We start with $\hat{v}_i = 0.5$, and compute $\hat{\alpha}_i = \mathbf{z}_i^T \hat{\boldsymbol{\phi}} \mathbf{c}(\hat{v}_i)$, $i = 1, \dots, N$. We then apply bisection to compute the corresponding value of \hat{u}_{it} . Based on the sign of $G_{\alpha_i}(\hat{\alpha}_i, \hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\phi}})$, we update the current value of \hat{v}_i , and repeat the process until convergence.

Appendix C - Extended simulation results

We present additional simulation results that were not shown in the main text. We refer to Section 7 for details on the simulation scenarios.

C1. Comparison with standard penalized fixed-effects estimators

We compared the performance of the described estimator with that of Koenker’s (2004) penalized fixed-effects method for longitudinal quantile regression, implemented in the `rqpd` R package. To fit the model, the individual effects were centered around the median of their marginal distribution (which is zero in simulation 1, and approximately 0.8744 in simulation 2). The tuning parameter λ was selected as the ratio of the level-1 and level-2 variance components of a linear random-intercept model. Quantiles 0.2, 0.4, 0.6, 0.8 were estimated jointly with equal weight.

A comparison is only possible for level-1 parameters, as no level-2 coefficients are estimated in standard penalized quantile regression. Results are summarized in Tables C1a and C1b. As expected, using a parametric model improved efficiency. Additionally, our estimator appeared to have a smaller bias, as if the imposed parametric structure could alleviate the incidental parameter problem.

C2. Model selection

To assess the performance of the information criteria described in Section 6.2, we estimated three alternative models, of which only one was correctly specified. Then, we used AIC and BIC to select the “best” model. The simulation is described in details in Table C2a. In simulation 1, we compared different specifications of $\beta_1(u \mid \boldsymbol{\theta})$, using level-1 AIC and BIC. In simulation 2, we compared different specifications of $\gamma_1(v \mid \boldsymbol{\phi})$, using level-2 AIC

and BIC. In both scenarios, model II had the same number of parameters as model I, and corresponded to a very good approximation of the true model; model III had less parameters, but was more severely misspecified.

Results are summarized in Table C2b. As expected, BIC tends to reward more parsimonious models, especially when the sample size is relatively small.

Table C1a: Comparison with a penalized fixed-effects estimator ($N = 150$)

Simulation 1											
$T = 5$	u	β_0	$\hat{\beta}_0$	$\hat{\beta}_0^K$	se	se ^K	β_1	$\hat{\beta}_1$	$\hat{\beta}_1^K$	se	se ^K
	0.2	1.11	1.17	1.19	.10	.13	0.73	0.72	0.69	.10	.13
	0.4	1.26	1.30	1.33	.11	.12	0.99	0.97	0.92	.11	.11
	0.6	1.46	1.48	1.48	.11	.13	1.01	0.99	1.02	.11	.12
	0.8	1.80	1.79	1.75	.12	.15	1.27	1.24	1.23	.12	.22
$T = 10$	u	β_0	$\hat{\beta}_0$	$\hat{\beta}_0^K$	se	se ^K	β_1	$\hat{\beta}_1$	$\hat{\beta}_1^K$	se	se ^K
	0.2	1.11	1.14	1.16	.10	.12	0.73	0.73	0.71	.06	.09
	0.4	1.26	1.27	1.30	.10	.12	0.99	0.98	0.95	.07	.07
	0.6	1.46	1.47	1.48	.10	.12	1.01	1.00	1.03	.07	.08
	0.8	1.80	1.80	1.78	.10	.14	1.27	1.25	1.25	.08	.17
Simulation 2											
$T = 5$	u	β_0	$\hat{\beta}_0$	$\hat{\beta}_0^K$	se	se ^K	β_1	$\hat{\beta}_1$	$\hat{\beta}_1^K$	se	se ^K
	0.2	0.11	0.11	0.22	.01	.11	3.60	3.79	3.71	.11	.17
	0.4	0.24	0.24	0.35	.03	.11	4.20	4.28	4.17	.12	.18
	0.6	0.41	0.41	0.46	.05	.11	4.80	4.77	4.72	.15	.19
	0.8	0.66	0.66	0.66	.08	.12	5.40	5.26	5.29	.19	.21
$T = 10$	u	β_0	$\hat{\beta}_0$	$\hat{\beta}_0^K$	se	se ^K	β_1	$\hat{\beta}_1$	$\hat{\beta}_1^K$	se	se ^K
	0.2	0.11	0.11	0.19	.01	.09	3.60	3.68	3.64	.07	.10
	0.4	0.24	0.23	0.31	.02	.09	4.20	4.23	4.18	.08	.12
	0.6	0.41	0.40	0.44	.03	.10	4.80	4.78	4.76	.10	.13
	0.8	0.66	0.64	0.66	.05	.10	5.40	5.33	5.35	.12	.15

Comparison with Koenker's (2004) penalized fixed-effects estimator with $N = 150$ and $T = \{5, 10\}$. We report the true value (β) of the level-1 coefficients at the quintiles (0.2, 0.4, 0.6, 0.8), and the average estimates and empirical standard errors obtained with our method ($\hat{\beta}$, se) and with Koenker's penalized fixed-effects estimator ($\hat{\beta}^K$, se^K), across $B = 1000$ simulated datasets.

Table C1b: Comparison with a penalized fixed-effects estimator ($N = 300$)

Simulation 1											
$T = 5$	u	β_0	$\hat{\beta}_0$	$\hat{\beta}_0^K$	se	se^K	β_1	$\hat{\beta}_1$	$\hat{\beta}_1^K$	se	se^K
	0.2	1.11	1.17	1.20	.08	.09	0.73	0.72	0.69	.07	.09
	0.4	1.26	1.30	1.33	.08	.09	0.99	0.97	0.92	.08	.08
	0.6	1.46	1.48	1.48	.08	.09	1.01	1.00	1.02	.08	.09
	0.8	1.80	1.79	1.74	.08	.11	1.27	1.24	1.24	.09	.16
$T = 10$	u	β_0	$\hat{\beta}_0$	$\hat{\beta}_0^K$	se	se^K	β_1	$\hat{\beta}_1$	$\hat{\beta}_1^K$	se	se^K
	0.2	1.11	1.14	1.15	.07	.08	0.73	0.73	0.71	.05	.06
	0.4	1.26	1.27	1.29	.07	.08	0.99	0.98	0.94	.05	.05
	0.6	1.46	1.47	1.47	.07	.09	1.01	1.00	1.02	.05	.06
	0.8	1.80	1.80	1.78	.07	.10	1.27	1.25	1.25	.05	.12
Simulation 2											
$T = 5$	u	β_0	$\hat{\beta}_0$	$\hat{\beta}_0^K$	se	se^K	β_1	$\hat{\beta}_1$	$\hat{\beta}_1^K$	se	se^K
	0.2	0.11	0.11	0.22	.01	.08	3.60	3.79	3.70	.08	.11
	0.4	0.24	0.24	0.36	.02	.08	4.20	4.28	4.17	.08	.13
	0.6	0.41	0.41	0.47	.03	.08	4.80	4.77	4.72	.10	.14
	0.8	0.66	0.66	0.65	.06	.08	5.40	5.27	5.30	.13	.15
$T = 10$	u	β_0	$\hat{\beta}_0$	$\hat{\beta}_0^K$	se	se^K	β_1	$\hat{\beta}_1$	$\hat{\beta}_1^K$	se	se^K
	0.2	0.11	0.10	0.18	.01	.06	3.60	3.67	3.64	.05	.07
	0.4	0.24	0.23	0.31	.01	.07	4.20	4.23	4.17	.06	.08
	0.6	0.41	0.40	0.44	.02	.07	4.80	4.78	4.76	.07	.09
	0.8	0.66	0.64	0.65	.04	.07	5.40	5.33	5.35	.09	.10

Comparison with Koenker's (2004) penalized fixed-effects estimator with $N = 300$ and $T = \{5, 10\}$.

Table C2a: Three different model specifications

	Simulation 1	Simulation 2
Coefficient	$\beta_1(u \mid \boldsymbol{\theta})$	$\gamma_1(v \mid \boldsymbol{\phi})$
True value	$1 + 10(u - 0.5)^3$	$0.5 \log(1 - \log(1 - v))$
Model I (correct)	$\theta_{10} + \theta_{11}u + \theta_{12}u^2 + \theta_{13}u^3$	$\phi_{10} + \phi_{11} \log(1 - \log(1 - v))$
Model II (misspecified)	$\theta_{10} + \theta_{11}u + \theta_{12} \cos(\pi u) + \theta_{13} \sin(\pi u)$	$\phi_{10} + \phi_{11}(1 - (1 - v)^{0.5})$
Model III (misspecified)	$\theta_{10} + \theta_{11}\sqrt{u} + \theta_{12}\sqrt{1 - u}$	$\phi_{11}(1 - v^2)$

Alternative models to be compared using AIC and BIC. For each scenario, we estimated three different models, of which one (model I) was correctly specified, and the other two (models II and III) were misspecified.

Table C2b: Performance of AIC and BIC

Simulation 1									
	$N = 150$				$N = 300$				
	$T = 5$		$T = 10$		$T = 5$		$T = 10$		
	AIC	BIC	AIC	BIC	AIC	BIC	AIC	BIC	
Model I (correct)	0.77	0.28	0.88	0.84	0.87	0.82	0.94	0.94	
Model II (misspecified)	0.21	0.07	0.12	0.12	0.13	0.12	0.06	0.06	
Model III (misspecified)	0.02	0.65	0.00	0.04	0.00	0.06	0.00	0.00	

Simulation 2									
	$N = 150$				$N = 300$				
	$T = 5$		$T = 10$		$T = 5$		$T = 10$		
	AIC	BIC	AIC	BIC	AIC	BIC	AIC	BIC	
Model I (correct)	0.71	0.55	0.75	0.64	0.86	0.81	0.80	0.79	
Model II (misspecified)	0.20	0.15	0.23	0.20	0.13	0.12	0.20	0.20	
Model III (misspecified)	0.09	0.30	0.02	0.16	0.01	0.07	0.00	0.01	

Model selection based on AIC and BIC. In the table, we report the relative frequency with which each of the three candidate models was selected, across $B = 1000$ simulated datasets. The correct specification is that defined by model I (see Table C2a).