A New Framework for Estimation of Unconditional Quantile Treatment Effects: The Residualized Quantile Regression (RQR) Model*

Nicolai T. Borgen¹, Andreas Haupt², and Øyvind Wiborg³

¹ Department of Special Needs Education, University of Oslo, Norway

² Institute of Sociology, Media and Cultural Studies, Karlsruhe Institute of Technology, Germany

³ Department of Sociology and Human Geography, University of Oslo, Norway

June 2022

Abstract

Identifying unconditional quantile treatment effects (QTE) has become increasingly popular within social sciences. However, current methods to identify unconditional QTEs of continuous treatment variables are incomplete. Contrary to popular belief, the unconditional quantile regression model introduced by Firpo et al. (2009) does not identify QTE, while the propensity score framework of Firpo (2007) allows for only a binary treatment variable, and the generalized quantile regression model of Powell (2020) is unfeasible with high-dimensional fixed effects. This paper introduces a two-step approach to estimate unconditional QTEs. The treatment variable is regressed on the control variables in the first step using, for example, linear regression. Then, in the second step, the residuals are used as the treatment variable in a quantile regression model. The intuition is that the first step decomposes the variance of the treatment variable into a piece explained by the observed control variables and a residual piece independent of the controls. Since the control variables purge the treatment of confounding, they are redundant in the second step. Therefore, the RQR approach circumvents the problem that inclusion of controls together with the treatment variable in CQR changes the interpretation of the treatment coefficients. Unlike much of the literature on quantile regression, this two-step residualized quantile regression framework is easy to understand, computationally fast, and can include high-dimensional fixed effects.

Keywords: Quantile regression, quantile treatment effect, residual regression, residualized quantile regression, fixed effects.

* We thank the anonymous reviewers, which critical reading helped us clarify the assumptions and our contribution. Earlier versions of the paper were presented at the 2021 ISA RC28 Spring Meeting in Turku, the 2022 DAGStat Conference in Hamburg, the Social Inequalities and Population Dynamics (SIPD) group's seminar at the University of Oslo in 2021, and the EQOP seminar at the University of Oslo in 2021. We thank participants for their comments and suggestions. The contribution of Nicolai T. Borgen was financed by a European Research Council grant (#818425). The register data were made available by Statistics Norway.

Introduction

Studying differences between groups has historically been akin to looking at differences in means. However, researchers increasingly turn to quantile regression models to get a complete view of how treatments affect outcomes (Koenker, 2005). One advantage of quantile regression models over standard linear regression models is studying how associations vary across the outcome variable's distribution, allowing researchers to explore new types of research questions. Generally, while linear regression models enable us to examine how the average of the outcome differs between groups, quantile regression models allow for studying how quantile values differ (Firpo, 2007). Thus, quantile regression models can be used to analyze how individuals with a proneness to have high outcomes react differently to treatment than individuals with low proneness, called unconditional quantile treatment effects (QTEs). We can, for example, ask whether pupils with high abilities reacted differently to school closures through the COVID pandemic compared to pupils with low abilities.

Despite significant advances in quantile regression models since the turn of the century, current methods to identify unconditional QTEs are still incomplete. Historically, the conditional quantile regression (CQR) model – which builds upon Roger Koenker and colleagues' work in the mid-1970s – has been used to estimate quantile regression coefficients (Koenker, 2017). CQR coefficients can be interpreted as QTEs whenever we do not need to include any control variables in our model (e.g., randomized treatment). However, unlike in linear regression models, including control variables changes the interpretation of the CQR coefficients, and they can no longer be interpreted as unconditional QTEs without strong assumptions (Borgen et al., 2022; Firpo, 2007; Killewald & Bearak, 2014; Rios-Avila & Maroto, 2022; Wenz, 2018). Therefore, solutions that allow for including control variables in quantile regression models while simultaneously preserving the coefficients' interpretation as QTEs are being developed (Firpo, 2007; Frölich & Melly, 2010; Powell, 2020).

This paper adds to this growing literature by offering a new quantile treatment estimation method, called Residualized Quantile Regression (RQR), which complements existing approaches. In his seminal paper, Firpo (2007) proposed an elegant solution to estimate unconditional QTEs with a

single binary treatment variable using a propensity score matching framework. However, the propensity score framework cannot be used with non-binary treatment variables, and including fixed effects is problematic. Recently, Powell (2020) developed the generalized quantile regression (GQR) model that allows for non-binary treatment variables.¹ However, this method is computationally demanding, with computational issues growing with the model's complexity and the sample size. Thus, including high-dimensional fixed effects in large administrative data sets is challenging or practically impossible using the GQR model. As we will show, our estimation method can easily handle large data sets and complex model specifications.

The RQR model builds upon the simple fact that CQR can estimate QTEs if the treatment is randomized. Thus, we need to find a way to randomize, so to speak, the treatment variable based on the observed control variables, which is a surprisingly uncomplicated task. Our simple yet innovative idea is that unconditional QTEs can be estimated through the following two-step approach. First, regress the treatment variable on control variables using, for example, ordinary least squares and obtain residuals of the treatment variable. Second, regress the outcome variable on the residualized treatment variable using the method of minimum absolute deviation. The intuition is that the first step decomposes the variance of the treatment variable into a piece explained by the observed control variables and a residual piece that is orthogonal to the observed controls. In the case of using OLS in the first step, this decomposition gives us treatment information mean independent of confounders, which can serve as a randomized treatment variable. Furthermore, since the control variables purge the treatment of confounding in the first step, they are redundant in the second step. Therefore, our approach serves as a straightforward solution to estimate QTEs in the presence of covariates.

The RQR model has several advantages over current QTE approaches, making it a valuable addition to the quantile regression toolkit. Most importantly, including high-dimensional fixed effects

¹ We do not discuss Powell (2016)'s non-additive fixed effects panel estimator (QRPD) in-depth in this paper. The motivation behind the RQR model and Powell (2016)' QRPD model is similar – identify unconditional QTEs in the presence of covariates – but the problem it addresses differs; it corrects for biases caused by unit specific trajectories.

is effortless in the RQR model, and the model allows for both binary and non-binary treatment variables. Additionally, the estimation procedure is computationally efficient, and it is straightforward to implement the RQR model in all software that provides a package for CQR or linear programming. Borgen et al. (2021) introduce the `rqr` and `rqrplot` Stata commands that estimate and plot RQR coefficients.

The RQR model belongs to a class of models other than the popular unconditional quantile regression (UQR) model (Firpo et al., 2009). The distinction between these models is discussed in detail by Rios-Avila and Maroto (2022) and Borgen et al. (2022). Here, we briefly clarify their difference. The UQR model partly gained its popularity within fields such as sociology (Budig & Hodges, 2014; Cooke, 2014; England et al., 2016; Glauber, 2018; Killewald & Bearak, 2014), educational science (Porter, 2015), and econometrics (Havnes & Mogstad, 2015; Lindqvist & Vestman, 2011) because it *seemingly* identifies QTEs in the presence of control variables. However, the UQR model was developed to infer how independent variables influence overall quantile values. There is accordingly a mismatch between the quantile regression model used in many studies (the UQR model) and these studies' aim (identify QTEs) (Borgen et al., 2022). These studies' research questions would often be better answered using a fast and easy-to-understand QTE model, such as the RQR model.

In the following, we start by defining the unconditional QTE. Then we describe the RQR model in more detail. Lastly, we demonstrate the RQR model's performance in data simulations and an empirical application on real data, comparing the RQR approach to other quantile regression approaches in both.

Unconditional quantile treatment effects (QTEs)

While ordinary least squares (OLS) and estimation of average treatment effects (ATE) is the main workhorse of quantitative empirical research, an increasing number of scholars are turning to quantile regression models to estimate unconditional quantile treatment effects (QTEs). The main attraction of

unconditional QTEs is that it allows for studying treatment effect heterogeneity by individuals' overall proneness to have high or low outcomes.

Because of their close resemblance, let us briefly define ATE before turning to QTEs. In the potential outcomes framework (Morgan & Winship, 2015), the causal effect of a treatment for a single unit is defined as:

$$[1] \quad \delta_i = Y_{i,1} - Y_{i,0}$$

, where $Y_{i,1}$ is the value of Y for individual i when the treatment is set to 1 and $Y_{i,0}$ is the value of Y for the same individual when the treatment is set to 0. Observing the outcome in both treatment states is impossible in reality, and the unit-level causal effects are accordingly based on hypothetical, what-if states in a thought experiment. Using all observations within their two potential states, we can calculate the commonly used average treatment effect (ATE) as the average difference between the potential outcomes:

$$[2] \quad ATE = E[Y_{i,1}] - E[Y_{i,0}]$$

Likewise, if we know the whole distribution of the potential outcomes, $F_{Y_0}(y)$ and $F_{Y_1}(y)$, then we can define QTEs for the quantile τ in a similar manner:

$$[3] \quad QTE^\tau = Q_{Y_1}^\tau - Q_{Y_0}^\tau$$

, where $Q_{Y_1}^\tau$ and $Q_{Y_0}^\tau$ are the values of the quantile τ under the potential outcomes given the treatment and the non-treatment conditions (Frölich & Melly, 2010). Consequently, QTE differs from ATE primarily by studying differences in quantile values rather than differences in means. Another difference is that QTEs potentially provide information about treatment effects for individuals at

different parts of the outcome distribution. That is, assuming individuals maintain their ranks in the potential outcome distributions (i.e., rank preservation), QTE^τ can be interpreted as treatment effects for individuals at quantile τ (Melly & Wüthrich, 2017).²

Since we obviously cannot observe the counterfactuals, we must compare outcomes (or their distributions) for individuals with different treatment values. For simplicity, let us illustrate with a random treatment assignment. The ATE is calculated by comparing average wages among treated with that of untreated. To calculate the QTE at the median, we start by ranking the treated's outcome from the lowest to the highest and find the value that separates the top 50% from the bottom half. This outcome represents the median outcome among the treated. After repeating this exercise for the untreated, we can compare the treated's median with that of the untreated.³ In the same manner, we can contrast other parts of the outcome distribution, such as the value that separates the bottom 10% and the top 90% (10th percentile) or the bottom 75% and the top 25% (75th percentile).

Whether QTEs differ from ATE depends on how the treatment influences the outcome (Hao & Naiman, 2007). For example, suppose the treatment only shifts the outcome distribution's location left or right, as shown in Panel A in Figure 1. In that case, the difference between the 10th, 50th, 90th, or any other percentile for treated ($T = 1$) and untreated individuals ($T = 0$) is of equal size as the difference in means for treated and untreated individuals. In contrast, when the treatment induces both a location and a scale shift (Panel B), not all QTEs are identical to the difference in means.

² However, if the rank preservation assumption fails to hold up, treatment effects cannot be interpreted as individual-level treatment effects. In that case, QTE coefficients inform on differences between the same quantile between the distributions of the treated and non-treated.

³ Note that treated's (untreated) median outcome are calculated among treated (untreated) only, the same way treated's (untreated) mean outcome are; it is a group-specific measure of the outcome variable's distribution.

Figure 1: Illustrating the difference between OLS and QTE when the treatment variable only induces a location shift (Panel A) and when the treatment induces both a location and a scale shift (Panel B).

Note: The blue and red shaded area represents kernel density estimates for untreated (T=0) and the treated (T=1). The blue and red dotted lines show the same groups' mean (OLS) and 90th percentile (QTE^{90th}). See Online Appendix C for Stata code to replicate the graph.

Estimating QTEs is effortless in the rare case where treatment assignment is random. As shown above, we can directly compare the quantile values of treated and untreated. Alternatively, QTEs can be estimated using linear programming methods to minimize the sum of asymmetrically weighted absolute residuals, as in the CQR model (Koenker, 2005). However, in most cases, we want to add control variables or fixed effects to account for selection bias (Killewald & Bearak, 2014). Furthermore, even in randomized controlled trials, control variables such as baseline levels are often needed to ensure the estimates' internal validity. Thus, the CQR model's ability to estimate QTEs is limited.

Estimating QTEs in the presence of control variables is more complicated than estimating treatment effects in the OLS model; because including control variables in the CQR model changes the coefficients' interpretation (Firpo, 2007; Frölich & Melly, 2010). The implication is that unconditional QTEs cannot be identified using CQR in the presence of covariates (at least not without additional

strong assumptions). Highlighting this dimension of the CQR model is essential, as the CQR model has traditionally been – and still is – misunderstood and erroneously used (Killewald & Bearak, 2014; Wenz, 2018). The UQR model is frequently posed as a solution to this problem; however, as noted above, it is not designed to study QTEs (Borgen et al., 2022; Rios-Avila & Maroto, 2022). The following section discusses approaches to include control variables in QTE models without changing the interpretation of the estimated treatment effect as QTEs.

Estimation of QTE in the presence of control variables

This paper introduces a novel two-step approach to identify QTEs where the treatment variable is purged of confounding in the first step using OLS, followed by a QTE estimation in the second step using a bivariate quantile regression model. The core building block of this two-step approach is that selection into treatment can be accounted for separately from estimating the QTEs. Another building block is the decomposition of the treatment variable into a piece explained by the observed control variables and a piece orthogonal to the controls. The next two subsections elaborate on these points before providing more details about the RQR approach.

A two-step approach to estimating QTEs

The RQR model belongs within the class of two-step estimators. In his seminal paper, Firpo (2007) proposed an elegant two-step solution to estimate unconditional QTE with a single binary treatment variable using a propensity score matching framework, which we call PS-QTE. The unconditional QTEs of the binary treatment variable are estimated using weighted CQR models (Koenker, 2005) with individual weights identical to the inverse probability of individuals being in their treatment group (Rosenbaum & Rubin, 1983).⁴

⁴ This approach was extended to apply to endogenous binary treatment variables by Frölich and Melly (2010).

To get an intuition of Firpo (2007)'s two-step matching approach, let us replace inverse probability weighting (see footnote 4 below for details) with 1-nearest neighbor (1-NN) matching. For each treated individual i , 1-NN matching selects the untreated individual with the smallest distance from individual i , with the distance measure typically being the conditional probability of treatment assignment given the observed control variables (i.e., the propensity score) (Stuart, 2010). The matched set of treated and untreated individuals serve as the counterfactual distributions of the potential outcomes, $F_{Y0}(y)$ and $F_{Y1}(y)$. Thus, the QTE at quantile τ is simply the difference between the treated group's quantile value and the corresponding quantile value in the untreated group, excluding unmatched individuals.

Our RQR framework is similar in spirit to Firpo (2007)'s propensity score approach in that we propose a two-step approach where the treatment variable is purged of confounding in the first step, followed by a QTE estimation in the second step. It is also similar in the sense that quantile regression is only used in the second step. However, instead of relying on a propensity score framework in the first step, we obtain residuals in the first step in order to construct a quasi-randomized treatment assignment. Although different first-step regressions can be used, we will draw upon the Conditional Expectation Function Decomposition property in OLS to de-confound the treatment variable. After regressing the treatment variable on the control variables using OLS, we obtain the residuals of the treatment variable. As explained in the next section, the residualized treatment variable is by construction orthogonal to observed control variables (assuming the treatment selection is modeled perfectly), allowing us to use quantile regression models to regress the outcome variable on the residualized treatment to obtain the QTEs.

Residuals of the treatment variable

In addition to random factors influencing treatment, the variation in treatment variable values may partly arise because of confounding variables (which we here assume are observed). Therefore, we want to decompose the treatment variable into a piece explained by the observed control variables and a residual piece. The sets of control variables may include interactions, polynomials, and fixed

effects. However, to illustrate our approach, we will begin with a correctly specified model in an undemanding confounding scenario before turning to misspecified models and more challenging confounding structures below. Consider the following model where T is the treatment variable, and x is an observed confounder.

$$[4] \quad T_i = \delta_0 + \delta_1 x_i + e_i$$

In this case, the confounder x induces only linear location shifts in the treatment T . Assuming this simple model, we can calculate the residuals $\tilde{T}$

$$[5] \quad \tilde{T}_i = T_i - \hat{T}_i = \delta_0 + \delta_1 x_i + e_i - \delta_0 - \delta_1 x_i = e_i$$

using a simple OLS model.

For our purpose, the main takeaway point is the decomposition of the treatment variable into a piece explained by x_i ($\hat{T}_i$) and a residual piece orthogonal to any function of x_i ($\tilde{T}_i$) (i.e., the so-called Conditional Expectation Function Decomposition property; Cunningham, 2020, p. 56). As proof that the residual piece is mean independent of x_i (assuming the model is correctly specified), consider the conditional expectation of the residuals $\tilde{T}_i$ (from the regression of T_i on x_i) given the observed control variable x_i :

$$[6] \quad E[\tilde{T}_i | x_i] = E[T_i - E[T_i | x_i] | x_i] = E[T_i | x_i] - E[T_i | x_i] = 0$$

Given a correctly specified model, the (treatment) residuals are mean independent of the confounder x_i and any function of x_i ($h(x_i)$): $E(h(x_i)\tilde{T}_i) = 0$. It follows from this that $cov(\tilde{T}_i, x_i) = 0$ as well. The confounder may still influence the variance or other moments of the outcome distribution; however, as discussed below, we only need the treatment to be mean independent of the confounder

to estimate unbiased QTEs. Since the residualized treatment is mean independent of x_i , the second step's regression coefficient of $\tilde{T}_i$ cannot be biased by the observed control variables (from the first step). Thus, this decomposition allows us to de-confound the treatment variable (assuming selection on observables). As detailed in the next section, the de-confounded treatment variable can then be included in a bivariate quantile regression model to identify QTEs.⁵

The Residualized Quantile Regression (RQR) model

Framework

We propose to estimate QTEs using the following two-step approach where the control variables are included in step 1 only:

Step 1: Regress the treatment variable on the control variables using OLS and obtain the residuals.

Step 2: Regress the outcome variable on the residualized treatment variable using the CQR algorithm.

Several assumptions must be met for this two-step estimator to identify QTEs. The most important is the unconfoundedness assumption, which loosely states that all variables affecting the treatment and the outcome (i.e., confounders) are observed and that the treatment selection is correctly specified. More specifically, we assume that the treatment residuals $\tilde{T}_i$ is independent of the potential outcomes under treatment and no-treatment (Assumption 1). This assumption is similar to the conditional independence assumption of Powell (2020) and the unconfoundedness assumption of Firpo (2007).

Assumption 1 (Unconfoundedness Assumption). $(Y_{i1}, Y_{i0}) \perp\!\!\!\perp \tilde{T}_i$

⁵ This decomposition property is also applied in the classical Frisch-Waugh-Lovell (FWL) theorem (Frisch & Waugh, 1933; Michael C Lovell, 1963; Michael C. Lovell, 2008), also known as the regression anatomy (Angrist & Pischke, 2009, pp. 35-36; Filoso, 2013) and double residual regression (Goldberger, 1991, p. 186).

The unconfoundedness assumption, and the fact that the treatment is residualized before estimating the (bivariate) quantile regression model, are the main departures from the CQR model. Thus, other conditions required in the bivariate CQR model, covered in detail by Koenker (2005), also apply to the RQR model, such as the assumption of no quantile crossing (He, 1997) and a continuously distributed outcome variable (Machado & Silva, 2005). Let Y be a continuous random variable with the cumulative distribution function (CDF) of Y defined as $F(y) = P(Y \leq y)$. If the CDF is strictly monotonically increasing, then the quantile function $Q_Y(\tau)$ can be defined in terms of the inverse of $F(Y)$: $F^{-1}(\tau) = \inf\{y: F(y) \geq \tau\}$.

Assumption 2 (CDF is strictly monotonically increasing). $Q_Y(p) = F_Y^{-1}(p)$

Additionally, we invoke the rank invariance assumption for the treatment effects to be interpreted as individual-level QTEs (but see footnote 2 for an alternative interpretation). This assumption is also commonly assumed in the CQR, PS-QTE, and GQR models (Firpo, 2007; Koenker, 2005; Powell, 2020). Consider the binary treatment variable T_i and let Y_{i0} and Y_{i1} be the potential outcomes for $T_i = 0, 1$. Further, let $r_i^0 \sim U[0, 1]$ and $r_i^1 \sim U[0, 1]$ be the potential ranks for $T_i = 0, 1$. We then assume that the ranks in the potential outcomes are the same.

Assumption 3 (Rank Invariance Assumption). $r_i^0 = r_i^1$.

Estimation

Under selection on observables and a correctly specified model, the first step purges the treatment variable of confounding, ensuring treatment ignorability in the second step. Thus, residualization of the treatment allows for including control variables (in the first step) to account for confounding without altering the coefficients' interpretation as QTEs (in the second step). In the second step, linear

programming methods are used to estimate coefficients.⁶ The algorithm is only briefly described here, and details can be found in Koenker (2005) and Hao and Naiman (2007). Known as the method of minimum absolute deviations (MAD), the QR estimator finds the coefficients that minimize the sum of weighted absolute residuals:

$$[7] \quad \sum_{i: y_i \geq x_i' \beta^{(\tau)}}^N \tau |y_i - x_i' \beta^{(\tau)}| + \sum_{i: y_i < x_i' \beta^{(\tau)}}^N (1 - \tau) |y_i - x_i' \beta^{(\tau)}|$$

, where $0 < \tau < 1$ and the superscript (τ) is included to clarify that the betas are allowed to differ by quantile. In the two-step RQR estimator, the coefficient matrix consists of only the residualized treatment variable and the constant. Thus, for the estimate of the p th quantile, we need to minimize the weighted absolute deviation of the coefficients of the residualized treatment variable $\beta_i^{(\tau)}$ and the constant term $\beta_o^{(\tau)}$:

$$[8] \quad \sum_{i: y_i \geq \beta_o^{(\tau)} + \beta_1^{(\tau)} \tilde{T}_i}^N \tau |y_i - \beta_o^{(\tau)} - \beta_1^{(\tau)} \tilde{T}_i| + \sum_{i: y_i < \beta_o^{(\tau)} + \beta_1^{(\tau)} \tilde{T}_i}^N (1 - \tau) |y_i - \beta_o^{(\tau)} - \beta_1^{(\tau)} \tilde{T}_i|$$

⁶ This approach is similar in spirit to the PS-QTE approach described above. The PS-QTE approach uses a logistic regression model to obtain the estimated probability of being treated T_i conditional on the set of control variables X_i , which is the estimated propensity score $\hat{p}(x_i)$. This propensity score is used to generate probability weights that are inversely related to the probability of being treated:

$$\hat{w}_{1,i} = \frac{T_i}{\hat{p}(x_i)} \text{ and } \hat{w}_{0,i} = \frac{1-T_i}{1-\hat{p}(x_i)}$$

where $\hat{w}_{1,i}$ is the weight for treated and $\hat{w}_{0,i}$ is the weight for untreated. Then, we include these inverse probability weights $\hat{w}_{j,i}$ (j index treatment status) in the tradition CQR algorithm along with the binary treatment variable:

$$\hat{w}_{j,i} * \left(\sum_{i: y_i \geq \beta_o^{(\tau)} + \beta_1^{(\tau)} T_i}^N \tau |y_i - \beta_o^{(\tau)} - \beta_1^{(\tau)} T_i| + \sum_{i: y_i < \beta_o^{(\tau)} + \beta_1^{(\tau)} T_i}^N (1 - \tau) |y_i - \beta_o^{(\tau)} - \beta_1^{(\tau)} T_i| \right)$$

Let us provide some further intuition why the RQR model works. There are two main reasons why the β_1 coefficient that minimizes equation [8] is the QTE (given assumptions 1-3 holds). First, as we discussed above, under the unconfoundedness assumption and a correctly specified first-step model, the residualized treatment variable is mean independent of the observed control variables x_i and any functions of x_i . When $\tilde{T}_i$ increases by one unit, this tells us nothing about the average value of the confounder x_i (i.e., $E[\tilde{T}_i|x_i] = 0$, see equation 6 above). Therefore, the tuning of β_1 in the equation above is not influenced by the effect of the (observed) confounder(s) on the outcome.

Second, there is by construction a mechanical association between the residuals $\tilde{T}_i$ and the (non-residualized) treatment variable T_i . Let $T_i = \delta_0 + \delta_1 X_i + e_i$ where we without loss assume that the constant $\delta_0 = 0$. Then regressing T_i on $\tilde{T}_i$ trivially gives

$$\begin{aligned}
\beta_{T_i \tilde{T}_i} &= \frac{\text{cov}(T_i, \tilde{T}_i)}{\text{var}(\tilde{T}_i)} \\
&= \frac{E[\tilde{T}_i T_i] - E[\tilde{T}_i]E[T_i]}{E[\tilde{T}_i^2] - (E[\tilde{T}_i])^2} \\
&= \frac{E\left[\overbrace{(T_i - \hat{T}_i)}^{\tilde{T}_i} * T_i\right]}{E\left[\overbrace{(T_i - \hat{T}_i)}^{\tilde{T}_i}\right]^2} \\
&= \frac{E\left[\left(\overbrace{(\delta_1 X_i + e_i)}^{T_i} - \overbrace{(\delta_1 X_i)}^{\hat{T}_i}\right) * \overbrace{(\delta_1 X_i + e_i)}^{T_i}\right]}{E\left[\left(\overbrace{(\delta_1 X_i + e_i)}^{T_i} - \overbrace{(\delta_1 X_i)}^{\hat{T}_i}\right)^2\right]} \\
[9] \quad &= \frac{E[(\delta_1 X_i + e_i - \delta_1 X_i) * (\delta_1 X_i + e_i)]}{E[(\delta_1 X_i + e_i - \delta_1 X_i)^2]} \\
&= \frac{E[e_i(\delta_1 X_i + e_i)]}{E[e_i^2]} \\
&= \frac{E[e_i^2]}{E[e_i^2]} \\
&= 1
\end{aligned}$$

, where we use the fact that $E[\tilde{T}_i] = 0$ in the third equality and the fact that $E[e_i X_i] = 0$ in the sixth equality. The takeaway is that when the residuals increase by one unit, the (non-residualized) treatment, on average, increases by the same amount (i.e., $\beta_{T_i \tilde{T}_i} = 1$).

In sum, the mechanical association between the residuals and the (non-residualized) treatment and the fact that the residuals are mean independent of observed control variables allows us to use the residuals as an as-if randomized treatment variable. Thus, control variables are not needed when estimating treatment effects, and QTEs can be estimated using a simple CQR model with the residualized treatment as the only independent variable. Our RQR approach hence circumvents the problem that the inclusion of controls together with the treatment variable in CQR changes the interpretation of the coefficient of the treatment variable.

Note that while the coefficients can be interpreted as unconditional QTEs, the constant term in the RQR model rarely has a sensible interpretation. The constant is the predicted quantile value of the outcome when the treatment residuals are zero.

Bias from misspecified first-step model

Adjustment-based regression approaches such as the RQR model provides unbiased coefficients given the unconfoundedness assumption is satisfied. Here, we discuss how having a misspecified first-stage regression would bias the second-step estimation of the QTEs. Importantly, failure to observe confounders or having a misspecified model will bias coefficients in other adjustment-based quantile regression models, including other QTE models, the UQR model, and the CQR model, as well as in linear regression models and propensity score matching approaches. Moreover, an advantage of the RQR model is that we can include fixed effects to adjust for time-invariant unobserved confounders, in which case only failure to measure and correctly specify time-variant confounders is of concern (Allison, 2009). However, we will ignore fixed effects in the following and focus on observed covariates.

Covariates are included in the RQR model to isolate treatment variation unrelated to confounding factors. Unmeasured confounders are correlated with the treatment residuals and would bias the

estimates. However, bias could also occur if observed confounders are modeled incorrectly. Consider a slight change in the example above, where the treatment is a quadratic function of the confounder x_i .

$$[10] \quad T_i = \delta_0 + \delta_1 x + \delta_2 x_i^2 + e_i$$

Further, suppose we estimate an incorrectly specified first-stage model where the confounder is included linearly, which means that the treatment residuals would be equal to $e_i - \delta_2 x_i^2$. In that case, the treatment residuals are mean independent of the linear term of the confounder but not its quadratic term (i.e., $E[\tilde{T}_i | x_i] = 0$ and $E[\tilde{T}_i | x_i^2] \neq 0$), and the unconfoundedness assumption is not satisfied. Since the treatment residuals are correlated with the confounder, plugging this treatment residual into the second-step quantile regression model will yield biased estimates of the QTE. The same logic applies if there are unmodeled interactions between confounders or if confounders are unmeasured. Thus, similar to other adjustment-based regression models, including the propensity score matching model and OLS, selection into treatment must be sufficiently modeled for the RQR model to provide unbiased estimates.

Choosing the correct causal model, including what constitutes potential confounders, must primarily be specified based on the know-how of the field (Morgan & Winship, 2015; Pearl, 2009). For example, not all covariates correlated with the treatment should be included, as highlighted by the discussions on bias amplification of near-instrumental variables (Pearl, 2011) and pre-treatment collider variables (Elwert & Winship, 2014). Nevertheless, building on specification tests such as Ramsey (1969)'s Reset Test and Pregibon (1980)'s Link Test, we can check whether there are unmodeled interactions or higher-order terms among our covariates. After calculating the treatment residuals, we propose to use an F-test to jointly test whether including higher-order terms (for continuous predictors) and interaction terms explain variance in the treatment residuals. Rejecting the null hypothesis suggests an incorrectly specified first-step regression that may result in biased

estimation of the QTE. If the control variables are specified incorrectly, we can include a squared term and reassess our model to see if it passes the test.⁷ (Supplementary Online Appendix C provides an example of this specification test.)

Standard errors

In various quantile regression models, estimation of standard errors using the bootstrap procedure (Hao & Naiman, 2007; Mooney & Duval, 1993) is preferred over the asymptotic procedure, including in the CQR model (Hao & Naiman, 2007; Koenker & Hallock, 2001), the PS-QTE model (Firpo, 2007), and the UQR model (Firpo et al., 2009). Likewise, we propose to estimate standard errors in the RQR model using a bootstrap procedure, where steps 1 and 2 described above are bootstrapped. Bootstrapping means randomly drawing M re-samples of size N with replacement from the original data sample. In each re-sample, a residualizing of the treatment variable is followed by an estimation of the quantile regression coefficients using the residualized treatment variable. The standard deviation of the estimated coefficients are the standard errors of the RQR coefficients.

Comparisons to other methods

There are several benefits of the RQR model compared to other approaches. First, including high-dimensional fixed effects in the RQR model is straightforward. Like other control variables, fixed effects are included exclusively in the first step to account for confounding. After obtaining the treatment variable's residuals (first step), we can estimate the unconditional QTE using the bivariate quantile regression model with the residualized treatment and the outcome variable (second step). In contrast, including high-dimensional fixed effects in the propensity score framework derived in Firpo (2007) is not only computationally burdensome but also potentially problematic because of the logistic

⁷ Moreover, when the treatment variable is binary, we may also consider to replace the first-step regression with a logit or probit model; although OLS provides a convenient approach to calculate the treatment residuals, other regressions can be used. However, in that case, fixed effects cannot be included, and one may consider instead using the propensity score framework of Firpo (2007).

regression model in the first step. In Powell (2020)'s GQR model, fixed effects can be included as information to construct the counterfactual distribution without treatment status (i.e., “prone variables”); however, that computationally-intensive solution is unfeasible with high-dimensional fixed effects.

Second, compared to much of the literature on quantile regression, which is computationally complex (Machado & Silva, 2019), the RQR approach is conceptually easier to grasp. While obviously of minor importance compared to its correctness, a method's transparency and ease of understanding nevertheless matter. The complexity of quantile regression has led to more than a decade of research – published in top-ranked journals across several disciplines – that mistakenly interprets unconditional partial effects from UQR as unconditional QTEs (Borgen et al., 2022). Having a transparent and relatively easy-to-understand method reduces the risk of such errors.

Third, the RQR model is computationally efficient. The advent of big data and complex model specifications have made computational efficiency crucial in all regression models. But even more so in quantile regression models, where regressions are repeated at multiple quantiles and bootstrapping is needed for inference. The computational burden of quantile regressions may discourage researchers from using this approach (Chernozhukov et al., 2020a; Fortin et al., 2011). When unconditional QTEs are estimated across multiple quantiles using the RQR model, one only needs to residualize the treatment variable once, making the estimation procedure more efficient and less computationally demanding. Further, the computational time can be reduced considerably by applying newly developed quantile regression algorithms (Chernozhukov et al., 2020a, 2020b). As an illustration, consider the QTEs in Figure 3, discussed in detail below. Even with this undemanding model specification, the GQR model takes 20 times longer to run than the RQR model (see Online Appendix G for files to replicate).

Fourth, unlike the propensity score framework (Firpo, 2007), the RQR approach extends to continuous treatment variables. Both binary and continuous treatment variables can be included in RQR without altering the coefficients' interpretation or changing the estimation procedure. Although

we can estimate QTEs of non-binary treatment variables using the GQR model, that model cannot include fixed effects, as mentioned above.

Additionally, there are several minor benefits of the RQR approach. A fifth advantage is that the RQR framework can draw on the already extensive literature on CQR, including models to estimate parametric quantile regression (Frumento & Bottai, 2016), methods to deal with count data (Machado & Silva, 2005), and fast algorithms (mentioned above). Finally, it is straightforward to implement the RQR estimator in all software that provides a package for CQR or linear programming, which makes it accessible to most researchers.⁸

Data simulations

Data generating process

This section uses Monte Carlo simulations to compare the RQR model's performance to other quantile regression approaches. We begin with fairly undemanding simulation scenarios before turning to more complex ones, discussed below with results in the supplementary online appendices B-J. The data simulations in the main text consist of running 10,000 draws of $N=2,000$ for two simulation scenarios. In both simulation scenarios, we have a continuous outcome variable y_i , a binary treatment variable t_i , and a binary control variable x_i . The two simulation scenarios differ concerning whether the control variable x_i affects only the outcome (scenario 1) or x_i affects both the outcome and the treatment variable (scenario 2). In simulation scenario 1, the treatment variable t_i is exogenous, and there is no need to include any control variables to identify the unconditional QTE. In contrast, controlling for the observed control variable x_i is necessary (and sufficient) to identify unconditional QTEs in simulation scenario 2.

⁸ Including Stata (qreg), R (quantreg), Python (QuantReg), SAS (quantreg), SPSS (quantile regression), and more. For time being, the GQR model is only available in Stata (genqreg) and the PS-QTE model is available in Stata (ivqte) and R (qte). The RQR model can be estimated in Stata using the rqr command (Borgen et al., 2021).

More specifically, in scenario 1, we begin by randomly assigning 25% of the sample the value 1 on the binary control variable x_i and randomly assigning 10% of the sample the value 1 on the binary treatment variable t_i . Next, we define a random pre-treatment outcome variable y_i^0 as:

$$[11] \quad y_i^0 = x_i * 1 + \epsilon_i, \quad \text{where } \epsilon_i \sim N(0,1)$$

Finally, we allow the strength of the treatment variable (t_i) to depend on the individual i 's percentile rank ($r_i \sim U[0,1]$) in the pre-treatment outcome distribution (y_i^0). The QTE monotonically increases across the outcome distribution.⁹

$$[12] \quad y_i = \overbrace{(r_i - 0.50)}^{QTE} * t_i + y_i^0$$

The setup in scenario 2 is similar, except the conditional probability of being treated depends on x_i : $P(t_i = 1|x_i = 0) = 0.067$ and $P(t_i = 1|x_i = 1) = 0.20$. Thus, the likelihood of being treated is 13.3 percentage points higher for individuals with 1 on the observed control variable than those with the value 0.

We estimate quantile regression coefficients using five different quantile regression models for each of the 10,000 draws and in each of the two simulation scenarios. Three of these estimation strategies identify unconditional QTEs, and their coefficients should, therefore, be nearly identical: the RQR method introduced in this paper, the PS-QTE approach of Firpo (2007), and the GQR approach of Powell (2020). The last two strategies are the CQR model (Koenker, 2005) and the UQR model (Firpo et al., 2009), which are included because they are popular quantile regression methods. The RQR coefficients may differ from the CQR and UQR coefficients, which identify conditional QTEs and unconditional partial effects. All models are run with the treatment variable t_i and the control variable

⁹ For example, the QTE is -0.40 at the 10th percentile, 0 at the 50th percentile, and 0.40 at the 90th percentile.

x_i . The Monte Carlo Error is small and hovers around 0.001-0.002 for all models and quantiles, as shown in Online Appendix Figure A1.

Table 1: Average differences between estimated regression coefficients and the true QTE ($\varphi^{(\tau)}$) and its standard deviation ($\sigma_{\varphi}^{(\tau)}$) from simulation scenarios 1 and 2 for selected quantiles (10,000 draws of $N=2,000$).

	$Q^{.10}$		$Q^{.25}$		$Q^{.50}$		$Q^{.75}$		$Q^{.90}$	
	$\varphi^{(.10)}$	$\sigma_{\varphi}^{(.10)}$	$\varphi^{(.25)}$	$\sigma_{\varphi}^{(.25)}$	$\varphi^{(.50)}$	$\sigma_{\varphi}^{(.50)}$	$\varphi^{(.75)}$	$\sigma_{\varphi}^{(.75)}$	$\varphi^{(.90)}$	$\sigma_{\varphi}^{(.90)}$
Scenario 1:										
RQR	0.003	(0.151)	0.001	(0.133)	-0.002	(0.128)	-0.003	(0.134)	-0.004	(0.154)
PS-QTE	0.003	(0.151)	0.002	(0.133)	-0.002	(0.128)	-0.003	(0.134)	-0.005	(0.155)
GQR	0.005	(0.151)	0.003	(0.133)	-0.001	(0.127)	-0.004	(0.134)	-0.004	(0.154)
CQR	0.031	(0.149)	0.027	(0.128)	0.001	(0.121)	-0.028	(0.129)	-0.038	(0.148)
UQR	-0.003	(0.164)	0.037	(0.117)	-0.002	(0.101)	-0.038	(0.117)	0.002	(0.165)
Scenario 2:										
RQR	0.007	(0.174)	0.005	(0.147)	0.004	(0.132)	0.001	(0.129)	-0.001	(0.142)
PS-QTE	0.006	(0.183)	0.003	(0.159)	0.003	(0.145)	0.000	(0.140)	-0.002	(0.150)
GQR	0.007	(0.176)	0.007	(0.149)	0.007	(0.132)	0.002	(0.128)	-0.001	(0.139)
CQR	0.084	(0.156)	0.100	(0.134)	0.081	(0.122)	0.032	(0.126)	-0.004	(0.145)
UQR	0.080	(0.149)	0.072	(0.108)	0.003	(0.100)	-0.011	(0.125)	0.121	(0.195)

Note: Data simulation is performed in Stata 16.0, and files to replicate the results are available in Online Appendix C. CQR is the conditional quantile regression model (Koenker, 2005) estimated using the qreg command; RQR is the residualized quantile regression model introduced in this paper, PS-QTE is the propensity score framework of Firpo (2007) estimated using the ivqte command (Frölich & Melly, 2010); GQR is the generalized quantile regression (Powell, 2020) estimated using the genqreg command; UQR is the unconditional quantile regression model (Firpo et al., 2009) estimated using the rifreg command.

Main simulation results

The main simulation results are shown in Table 1 and Figure 2. The reported $\varphi^{(\tau)}$ quantities show the average difference between the estimated regression coefficient ($\hat{\beta}_j^{(\tau)}$) and the true QTE ($\beta^{(\tau)}$) at the quantile τ across the 10,000 independent draws j : $\varphi^{(\tau)} = E[\hat{\beta}_j^{(\tau)} - \beta^{(\tau)}]$. The $\varphi^{(\tau)}$ quantity should be interpreted as the estimated bias for the QTE models, where larger values mean more bias. In contrast, since CQR and UQR models do not estimate unconditional QTEs, any differences between their coefficients and the true coefficient do not reflect a bias and should not be interpreted as such. Table 1 also includes the standard deviation of the $\varphi^{(\tau)}$ quantities. Online Appendix Table A1 presents

simulation results using estimated regression coefficients rather than differences between coefficients and true QTE.

As expected, all three QTE models yield coefficients that approximate the true QTE across the outcome distribution, with the bias ranging from 0.001 to 0.005 in most cases. The differences between simulation scenario 1 (where there is no confounding) and simulation scenario 2 (where conditioning on x_i is sufficient to account for confounding) is negligible. There is a slight tendency for the QTE models to overestimate the QTE at the bottom of the outcome distribution and overestimate it at the top. However, this bias amounts to only 1-3% of its standard deviation based on the different draws, which is around 0.14 in the first scenario and slightly higher in the second scenario. Moreover, the bias is small relative to the size of the QTEs. For example, the estimated QTEs at the 10th percentile in scenario 1 are 0.75% off in the RQR and PS-QTE models and 1.25% off in the GQR model. Overall, the similarity of the estimated effects across the three different frameworks provides evidence that running a quantile regression model on a residualized treatment variable is a valuable approach to identifying unconditional QTEs.

Table 2: Average of the absolute values of the $\sigma_{\varphi}^{(\tau)}$ and $\sigma_{\varphi}^{(\tau)}$ quantities across the estimated quantiles.

	$ \varphi^{(\tau)} $	$ \sigma_{\varphi}^{(\tau)} $
Scenario 1:		
RQR	0.003	0.140
PS-QTE	0.003	0.140
GQR	0.003	0.140
Scenario 2:		
RQR	0.003	0.144
PS-QTE	0.002	0.156
GQR	0.005	0.145

Note: In each of the 10,000 draws, the average coefficient ($\sigma_{\varphi}^{(\tau)}$) and its standard deviation ($\sigma_{\varphi}^{(\tau)}$) are calculated (see Table 1 and Figure 2). This table shows the average of these quantities' absolute values at quantiles 5 to 95 at steps of 5 (19 observations for each scenario and method combination).

This paper's data simulation is not a comprehensive comparison of different QTE models. However, a couple of differences between the three QTE models in this specific setup could be noted.

First, the bias is marginally lowest in the PS-QTE approach, followed by the RQR model, and finally, the GQR model, as summarized in Table 2. The average of the absolute value of the bias across the 19 quantiles is indistinguishable at the third decimal in scenario 1. In scenario 2, the PS-QTE, RQR, and GQR models' biases are 0.002, 0.003, and 0.005. Second, the standard deviations across the draws are 7-8% lower in the RQR and GQR model than in the PS-QTE model in scenario 2 (and indistinguishable in scenario 1). All in all, the differences between the QTE models are trivial in this simulation setup.

Figure 2: Average differences between estimated regression coefficients and the true QTE ($\varphi^{(\tau)}$) from simulation scenarios 1 and 2.

Note: See note in Table 1 for details.

The QTE models' coefficients differ from the CQR and the UQR coefficients in both simulation scenarios. Let us start with the CQR results. As noted above, the difference between the CQR coefficients and the true QTE should not be interpreted as a bias, as the CQR model estimates conditional quantile value differences. Clearly shown in Figure 2, the CQR coefficients differ in scenarios 1 and 2, despite the underlying unconditional QTE being the same in both scenarios. In fact,

including the observed control variable x_i changes the interpretation of the conditional QTEs even in scenario 1, where the control variable x_i does *not* influence the treatment variable t_i (no-confounding scenario) (Wenz, 2018). Contrary to this, excluding the control variable x_i does not change the estimated unconditional QTEs using RQR, PS-QTE, and GQR (not shown).

The unconditional partial effects, estimated using the UQR model of Firpo et al. (2009), also differ considerably from the unconditional QTEs. Moreover, the UQR coefficients differ in scenarios 1 and 2, even though the underlying QTEs are the same. The differences between the QTE models (RQR, GQR, and PS-QTE) and the UQR model illustrate that the UQR model does not identify QTEs, although it is widely used in the literature for that purpose (Borgen et al., 2022).

Supplementary simulation scenarios

The simulation scenarios in the main text are supplemented by a broader set of simulation scenarios in Online Appendix B to Online Appendix J. We here discuss these results briefly, with more information included in the online appendices. First, Appendix B varies the distribution of the treatment, the structure of the QTE, and the skewness of the outcome distribution. Specifically, it simulates data with a binary treatment variable (as in the main text), a treatment variable with a uniformly distributed random component, and a treatment variable with a normally distributed random component. Further, we distinguish between three structures on the QTEs: constant, quadratic, and cubic. Finally, we allow the outcome's residual distribution to be either normal or right-skewed. The results in all of these 18 simulation scenarios align with the results discussed above: RQR provides unbiased estimates of the unconditional QTEs given a correctly specified first step.

Second, Appendix C simulates data where confounders interact and has a quadratic effect on the treatment. We show that incorrectly specifying the treatment selection results in biased estimates of the QTE. However, such misspecification can, in some cases, be detected using a specification test, such as in the Appendix example. Moreover, when interactions and quadratic effects are correctly specified, the RQR model provides unbiased estimates of the QTEs.

Third, Appendix D provides an example where the confounder's effect on the treatment is allowed to vary randomly. For some individuals, the confounder strongly influences the treatment values, while it has no effect for others. (As a result, the variance of the treatment residuals increases as a function of the confounder.) The appendix shows that the RQR model provides unbiased QTE estimates in this scenario.

Fourth, Appendix E simulates data where the variance of the treatment is a function of the confounder. For the RQR coefficients to have a causal interpretation, we must assume that the treatment residuals are mean independent of the confounders, not fully independent. The appendix shows that the RQR model provides unbiased estimates in a scenario where the variance of the treatment residuals is a function of the confounder.

Fifth, Appendix F simulates data where the confounder has a curvilinear effect on the outcome and a scenario where the confounder has a quantile treatment effect. Again, the RQR model provides unbiased estimates in these simulation scenarios.

Sixth, Appendix G simulates panel data where there is a time-invariant unobserved confounder. While other QTE approaches fail in such a scenario, we show that we can use the RQR model to identify unbiased QTEs using a fixed effects estimator.

Finally, the three last appendices show simulations where the RQR model (and other QTE models) fails to identify QTEs. Appendix G simulates data where the rank invariance assumption is not met. Rank invariance means that individuals' ranks in the potential outcome distributions remain the same irrespective of their treatment. Without invoking this assumption, QTEs cannot be interpreted as treatment effects for individuals located in different parts of the distribution. The appendix shows that RQR provides biased QTE estimates when rank invariance is not met, as do the GQR and the PS-QTE model. Next, Appendix H illustrates how the RQR model fails to identify QTEs in the presence of quantile crossing, along with the GQR model. Finally, the CQR algorithm requires uniqueness of quantiles (Firpo, 2007), which is not satisfied with a discrete outcome variable (Machado & Silva, 2005). Appendix I shows that the RQR model struggles with discrete outcome variables, and it struggles more

than the GQR model in the specific example included. However, the appendix also shows that the jittering approach of Machado and Silva (2005) can be used to solve the problem.

Empirical applications

After demonstrating the unbiasedness of the RQR in various simulation scenarios where the identifying assumptions are met, let us look at some real data examples. The purpose of these examples is to compare coefficients from different quantile regression models and, unlike the simulation results above, not identify unbiased QTEs. The true effects are unknown in these examples, and the coefficients cannot be interpreted causally for various reasons, including unobserved confounding and issues related to over-control. Thus, we need another baseline to evaluate the RQR model's validity. We indirectly assess the RQR model's validity by comparing it with the established GQR and PS-QTE approaches. Generally, similar coefficients support our approach's validity, although coefficients may differ somewhat because of small-sample "noise" and differences due to identification strategies.

Current Population Survey

This section illustrates the RQR model using the Outgoing Rotation Group supplement of the Current Population Survey (CPS) to study the effects of union status on male wages. Additionally, it again highlights the quantile regression models' differences. Union wage effects served as a motivating example in Firpo et al. (2009)'s UQR paper, which has become popular in studies of QTEs (Borgen et al., 2022). In their paper, Firpo et al. (2009) demonstrate that UQR coefficients differ from CQR coefficients. Here, we use the union wage effects example to showcase that RQR coefficients differ from UQR (and CQR) coefficients but not GQR and PS-QTE coefficients.¹⁰

¹⁰ The supplementary data to the Firpo et al. (2009) paper did not include household identification number necessary to include high-dimensional fixed effects. Appendix Figure A3 replicates the results of Figure 3 using the exact same data as Firpo et al. (2009).

Figure 3: Cross-sectional effects of union status on log wages for full-time working males in the 1983-1986 Outgoing Rotation group supplement of the Current Population Survey (N=251,153).

Note: The sample includes male household heads aged 16-64. The included control variables are the respondents' age, five dummies for educational level, a dummy variable for completed education, a dummy variable for married, and a dummy variable for non-white. PS-QTE is the weighting procedure of Firpo (2007) as implemented by the `ivqte` command in Stata (Frölich & Melly, 2010). The GQR model is the generalized quantile regression model of Powell (2020). Data is obtained from the National Bureau of Economic Research (<https://data.nber.org/cps/>).

Convincingly, differences between estimated union wage effects using QTE models are negligible; RQR, GQR, and PS-QTE coefficients decrease similarly from about 0.40 at the 5th quantile to -0.10 at the 95th quantile (Figure 3). Concerning computational speed, less than a minute is needed to estimate the RQR coefficients on a standard laptop when using the fast quantile regression algorithms provided by Chernozhukov et al. (2020b). In comparison, the PS-QTE model takes more than five times longer to run and the GQR model 20 times longer. These union wage effects are estimated with controls for age, education, marital status, and race. These control variables may not be sufficient to account for confounding; however, as this is similar across models, the model specification allows us to exemplify the RQR model and compare it to the other quantile regression approaches.

Figure 4: Effects of union status on log wages for full-time working males in the 1983-1986 Outgoing Rotation group supplement of the Current Population Survey (N=251,153). Panel A shows cross-sectional estimates, while panel B includes household fixed effects.

Note: See Figure 3 for details about data and variables. The UQR model is estimated using `rifreg` (Panel A) and `xtrifreg` (Panel B) commands (Borgen, 2016).

Figure 4 compares RQR coefficients with UQR and CQR coefficients. Two main results stand out in the cross-sectional analysis in Panel A. First, although the RQR and CQR models' estimated union wage effects are monotonically declining across the outcome distribution, the gradient is steeper in the RQR model. The RQR coefficients are larger at the bottom of the wage distribution before falling and eventually reaching a substantial adverse effect at the top. This finding highlights the difference between conditional and unconditional QTEs, as others have previously discussed (Firpo, 2007; Killewald & Bearak, 2014; Porter, 2015; Wenz, 2018). Second, the UQR coefficients differ from the RQR coefficients, especially in the bottom fifth of the wage distribution. The reason is that the UQR model of Firpo et al. (2009) identifies UQPE rather than unconditional QTE.

In Panel B of Figure 4, union wage effects are estimated with household fixed effects. Including high-dimensional fixed effects in the RQR model is straightforward; fixed effects are treated as any

other control variable and included only in the first step. After obtaining the treatment variable's residuals, unconditional QTEs can be estimated by regressing the outcome on the residualized treatment in a bivariate quantile regression model. Fixed effects can also, with ease, be included in the UQR model (Borgen, 2016; Rios-Avila, 2020). However, as it is challenging in the CQR model (for one approach, see Machado & Silva, 2019), Panel B compares the RQR coefficient solely with those of UQR. In the union example, the estimated within-household treatment effects are generally lower than the cross-sectional estimates, especially at the bottom of the wage distribution. However, the overall trend is similar in the cross-sectional estimate (Panel A) and the fixed effects estimate (Panel B).

Register data

We use Norwegian register data in the following example, which largely alleviates issues caused by imprecise coefficients in small samples. Figure 5 looks at influences on 8th-grade standardized test scores of birth order, birth month (i.e., relative school-starting age), immigrant background, gender, school SES, and parental earnings, where all independent variables are included simultaneously in the model. The three different QTE models provide nearly identical estimates for five of six variables. Importantly, this is the case despite the treatment effects having completely different effects across the outcome distribution. This result mirrors the findings in Appendix B, where we use data simulations to show that the RQR model (and the GQR model) can handle different treatment structures.

The exception is the variable children of immigrants, where the PS-QTE approach differs somewhat from the RQR and GQR models. However, this difference does not warrant a concern; rather, it showcases the well-known fact that regression and matching approaches may differ in some scenarios, despite identifying the same estimand. In this case, the inverse probability weighting matching approach also provides a different estimate of the ATE than the classical OLS model.

Figure 5: Comparing RQR, GQR, and PS-QTE coefficients on 8th-grade standardized test scores in Norwegian register data (N=480,264).

Note: The outcome variable is the average of 8th-grade entry-test scores in reading and math for birth cohorts born in Norway between 1997 and 2005. Birth order is defined based on both mother and father. Born Sept.-Dec. is a dummy variable for born in the last quarter of the year; the school-starting cutoff is January 1st. Children of immigrants are children born to at least one parent born outside of Norway. Girl is a dummy for gender. Parental earnings measure parents' percentile earnings rank. School SES is the average parental earnings rank among 5th-graders at the same school. The RQR and GQR models are used for all variables, while the PS-QTE model is only used for binary predictors. All coefficients are estimated controlling for the other independent variables in the graph.

National Longitudinal Survey of Young Working Women

As a final example, we use a subsample of the National Longitudinal Survey, containing panel data on 4,711 working women aged 14-26 in 1968 from 1968 to 1988. The outcome is log wages, and as in Figure 5, we show coefficients for all independent variables. Although the smaller sample size results in noisier coefficients, we clearly see that the GQR, RQR, and PS-QTE coefficients resemble each other, irrespective of the type of independent variable (binary or metric) and structure of treatment effects across the distribution.

Figure 6: Comparing RQR, GQR, and PS-QTE coefficients on log wages in a subsample of the National Longitudinal Survey (N=3,956).

Note: We keep observations born between 1943 and 1953 and are between ages 25 and 45 during the interview years. Then we keep one observation per woman (the latest recorded one). The RQR and GQR models are used for all variables, while the PS-QTE model is only used for binary predictors. All coefficients are estimated controlling for the other independent variables in the graph.

Discussion and conclusion

Identification of unconditional QTEs has become increasingly popular within social sciences; however, contrary to popular belief, current methods to identify unconditional QTEs of a continuous treatment are incomplete. Firpo (2007) has proposed a solution to estimate unconditional QTE with a single binary treatment variable using a propensity-matching framework. Powell (2020) developed the generalized quantile regression model that extends to non-binary treatment variables. However, none of these approaches easily allow for high dimensional (additive) fixed effects. Although fixed effects are easy to include in the UQR framework (Borgen, 2016; Rios-Avila, 2020), it is important to highlight that the UQR model identifies the treatments' effect on quantile values (UQPE) rather than quantile values differences between treatment groups (QTE).

This paper fills a gap in the literature by introducing a straightforward framework to estimate unconditional QTEs of both continuous and binary treatment variables in the presence of covariates and high dimensional fixed effects, called residualized quantile regression. The main advantage of the RQR model compared to other QTE models is that fixed effects can be included, which is a common workhorse in causal modeling, as exemplified in the literature on mother- and fatherhood wage penalty (Cooke, 2014; England et al., 2016; Killewald & Bearak, 2014).

Furthermore, the RQR model can draw on the already extensive literature on CQR, is easy to implement in all software that provides packages for CQR or linear programming, and is computationally efficient. With better access to large-scale data and more complex model specifications, the computational burden may discourage researchers from using quantile regressions. The RQR model's computational time can be reduced considerably by applying newly developed quantile regression algorithms (Chernozhukov et al., 2020a, 2020b), making it substantially faster than current QTE approaches. See Borgen et al. (2021) for a Stata command that exploits these fast algorithms.

This paper has introduced a new framework for estimating unconditional QTEs and demonstrated that it produces coefficients similar to other QTE models. We have suggested using OLS in the first-step regression to deconfound the treatment variable; however, future research should explore other first-step regressions, especially with binary treatment variables. Furthermore, statistical inference and estimation of RQR standard errors need more work. Bootstrapping the two-step procedure seemingly provides sound standard errors, as the 95% bootstrapped confidence intervals in our data simulations have coverage rates of 94.34% (Normal-Approximation Bootstrap CIs), 95.26% (Percentile Bootstrap CIs), and 94.52% (Bias-Corrected Bootstrap CIs) (Online Appendix Figure A4). Future work should systematically test various bootstrap procedures' performance in different simulation scenarios. Finally, the current paper has only dealt with the case of selection on observables. The RQR framework should be developed to allow for identifying an endogenous treatment variable's QTEs via instrumental variable estimation.

References

- Allison, P. D. (2009). *Fixed Effects Regression Models*. Los Angeles: Sage.
- Angrist, J. D., & Pischke, J.-S. (2009). *Mostly harmless econometrics: an empiricist's companion*. Princeton: Princeton University Press.
- Borgen, N. T. (2016). Fixed effects in unconditional quantile regression. *Stata Journal*, 16, 403-415.
- Borgen, N. T., Haupt, A., & Wiborg, Ø. (2021). Flexible and fast estimation of quantile treatment effects: The `rqr` and `rqrplot` commands. *SocArXiv*, 10.31235/osf.io/4vquh.
- Borgen, N. T., Haupt, A., & Wiborg, Ø. (2022). Quantile Regression Estimands and Models: Revisiting the Motherhood Wage Penalty Debate. *SocArXiv*, doi:10.31235/osf.io/9avrp.
- Budig, M. J., & Hodges, M. J. (2014). Statistical Models and Empirical Evidence for Differences in the Motherhood Penalty across the Earnings Distribution. *American Sociological Review*, 79, 358-364.
- Chernozhukov, V., Fernández-Val, I., & Melly, B. (2020a). Fast algorithms for the quantile regression process. *Empirical Economics*.
- Chernozhukov, V., Fernández-Val, I., & Melly, B. (2020b). Quantile and distribution regression in Stata: algorithms, pointwise and functional inference. *Unpublished paper*.
- Cooke, L. P. (2014). Gendered Parenthood Penalties and Premiums across the Earnings Distribution in Australia, the United Kingdom, and the United States. *European Sociological Review*, 30, 360-372.
- Cunningham, S. (2020). *Causal inference: The mixtape (V. 1.8)*. https://www.scunning.com/causalinference_norap.pdf.
- Elwert, F., & Winship, C. (2014). Endogenous Selection Bias: The Problem of Conditioning on a Collider Variable. *Annual Review of Sociology*, 40, 31-53.
- England, P., Bearak, J., Budig, M. J., & Hodges, M. J. J. A. S. R. (2016). Do highly paid, highly skilled women experience the largest motherhood penalty? , 81, 1161-1189.
- Filoso, V. (2013). Regression Anatomy, Revealed. *The Stata Journal*, 13, 92-106.
- Firpo, S. (2007). Efficient semiparametric estimation of quantile treatment effects. *Econometrica*, 75, 259-276.
- Firpo, S., Fortin, N. M., & Lemieux, T. (2009). Unconditional quantile regressions. *Econometrica*, 77, 953-973.
- Fortin, N., Lemieux, T., & Firpo, S. (2011). Chapter 1 - Decomposition Methods in Economics. In O. Ashenfelter & D. Card (Eds.), *Handbook of Labor Economics* (Vol. 4, pp. 1-102): Elsevier.
- Frandsen, B. R., & Lefgren, L. J. (2018). Testing Rank Similarity. *The Review of Economics and Statistics*, 100, 86-91.
- Frisch, R., & Waugh, F. V. (1933). Partial time regressions as compared with individual trends. *Econometrica: Journal of the Econometric Society*, 387-401.
- Frumento, P., & Bottai, M. (2016). Parametric modeling of quantile regression coefficient functions. *Biometrics*, 72, 74-84.
- Frölich, M., & Melly, B. (2010). Estimation of quantile treatment effects with Stata. *The Stata Journal*, 10, 423-457.
- Glauber, R. (2018). Trends in the Motherhood Wage Penalty and Fatherhood Wage Premium for Low, Middle, and High Earners. *Demography*, 55, 1663-1680.
- Goldberger, A. S. (1991). *A course in econometrics*. Cambridge, Mass.: Harvard University Press.
- Hao, L., & Naiman, D. Q. (2007). *Quantile regression* (Vol. 07-149). Thousand Oaks, California: Sage.
- Havnes, T., & Mogstad, M. (2015). Is universal child care leveling the playing field? *Journal of Public Economics*, 127, 100-114.
- He, X. (1997). Quantile Curves without Crossing. *The American Statistician*, 51, 186-192.
- Heisig, J. P., Schaeffer, M., & Giesecke, J. (2017). The costs of simplicity: Why multilevel models may benefit from accounting for cross-cluster differences in the effects of controls. *American Sociological Review*, 82, 796-827.

- Killewald, A., & Bearak, J. (2014). Is the Motherhood Penalty Larger for Low-Wage Women? A Comment on Quantile Regression. *American Sociological Review*, 79, 350-357.
- Koenker, R. (2005). *Quantile regression* (Vol. no. 38). Cambridge: Cambridge University Press.
- Koenker, R. (2017). Quantile regression: 40 years on. *Annual Review of Economics*, 9, 155-176.
- Koenker, R., & Hallock, K. (2001). Quantile regression: An introduction. *Journal of Economic Perspectives*, 15, 43-56.
- Lindqvist, E., & Vestman, R. (2011). The Labor Market Returns to Cognitive and Noncognitive Ability: Evidence from the Swedish Enlistment. *American Economic Journal: Applied Economics*, 3, 101-128.
- Lovell, M. C. (1963). Seasonal adjustment of economic time series and multiple regression analysis. *Journal of the American Statistical Association*, 58, 993-1010.
- Lovell, M. C. (2008). A simple proof of the FWL theorem. *The Journal of Economic Education*, 39, 88-91.
- Machado, J. A. F., & Silva, J. M. C. S. (2005). Quantiles for counts. *Journal of the American Statistical Association*, 100, 1226-1237.
- Machado, J. A. F., & Silva, J. M. C. S. (2019). Quantiles via moments. *Journal of Econometrics*.
- Melly, B., & Wüthrich, K. (2017). Local quantile treatment effects. In *Handbook of quantile regression* (pp. 145-164): Chapman and Hall/CRC.
- Mooney, C. Z., & Duval, R. D. (1993). *Bootstrapping: A nonparametric approach to statistical inference*: Sage.
- Morgan, S. L., & Winship, C. (2015). *Counterfactuals and causal inference*: Cambridge University Press.
- Pearl, J. (2009). *Causality: models, reasoning, and inference*. Cambridge: Cambridge University Press.
- Pearl, J. (2011). Invited commentary: understanding bias amplification. *American Journal of Epidemiology*, 174, 1223-1227.
- Porter, S. R. (2015). Quantile regression: analyzing changes in distributions instead of means. In *Higher Education: Handbook of Theory and Research* (pp. 335-381): Springer.
- Powell, D. (2016). Quantile regression with nonadditive fixed effects. *Quantile Treatment Effects*, 1-28.
- Powell, D. (2020). Quantile treatment effects in the presence of covariates. *Review of Economics and Statistics*, 102, 994-1005.
- Pregibon, D. (1980). Goodness of Link Tests for Generalized Linear Models. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 29, 15-14.
- Ramsey, J. B. (1969). Tests for specification errors in classical linear least-squares regression analysis. *Journal of the Royal Statistical Society: Series B (Methodological)*, 31, 350-371.
- Rios-Avila, F. (2020). Recentered influence functions (RIFs) in Stata: RIF regression and RIF decomposition. *The Stata Journal*, 20, 51-94.
- Rios-Avila, F., & Maroto, M. L. (2022). Moving Beyond Linear Regression: Implementing and Interpreting Quantile Regression Models With Fixed Effects. *Sociological Methods & Research*, 00491241211036165.
- Rosenbaum, P. R., & Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70, 41-55.
- Stuart, E. A. (2010). Matching methods for causal inference: A review and a look forward. *Statistical science: a review journal of the Institute of Mathematical Statistics*, 25, 1.
- Wenz, S. E. (2018). What Quantile Regression Does and Doesn't Do: A Commentary on Petscher and Logan (2014). *Child development*.
- White, I. R. (2010). Simsum: Analyses of Simulation Studies Including Monte Carlo Error. *The Stata Journal*, 10, 369-385.

Supplementary Online Appendix

A New Framework for Estimation of Unconditional Quantile Treatment Effects: The Residualized Quantile Regression (RQR) Model

Nicolai T. Borgen¹, Andreas Haupt², and Øyvind Wiborg^{3,4}

¹ Department of Special Needs Education, University of Oslo, Norway

² Institute of Sociology, Media and Cultural Studies, Karlsruhe Institute of Technology, Germany

³ Department of Sociology and Human Geography, University of Oslo, Norway

⁴ Center for the Study of Professions, Oslo Metropolitan University, Norway

Content of supplementary online appendix:

Online Appendix A: Supplementary tables and figures

Online Appendix B: Additional data simulations

Online Appendix C: Specification tests and incorrectly specified models

Online Appendix D: Heterogenous effects of confounders on treatment

Online Appendix E: Confounders influence variance

Online Appendix F: Heterogeneous effects of confounders on the outcome

Online Appendix G: Unobserved time-invariant confounding

Online Appendix H: Rank invariance assumption

Online Appendix I: Quantile crossing

Online Appendix J: Discrete outcome variable

Online Appendix K: Replication files

Supplementary Online Appendix A: Supplementary tables and figures

Online Appendix Table A1: Quantile regression coefficients for simulation scenarios 1 and 2 for selected quantiles (10,000 draws of N=2,000).

	$Q^{.10}$		$Q^{.25}$		$Q^{.50}$		$Q^{.75}$		$Q^{.90}$	
	$\varphi^{(.10)}$	$\sigma_{\varphi}^{(.10)}$	$\varphi^{(.25)}$	$\sigma_{\varphi}^{(.25)}$	$\varphi^{(.50)}$	$\sigma_{\varphi}^{(.50)}$	$\varphi^{(.75)}$	$\sigma_{\varphi}^{(.75)}$	$\varphi^{(.90)}$	$\sigma_{\varphi}^{(.90)}$
Scenario 1:										
RQR	-0.397	(0.151)	-0.249	(0.133)	-0.002	(0.128)	0.247	(0.134)	0.396	(0.154)
PS-QTE	-0.397	(0.151)	-0.248	(0.133)	-0.002	(0.128)	0.247	(0.134)	0.395	(0.155)
GQR	-0.395	(0.151)	-0.247	(0.133)	-0.001	(0.127)	0.246	(0.134)	0.396	(0.154)
CQR	-0.369	(0.149)	-0.223	(0.128)	0.001	(0.121)	0.222	(0.129)	0.362	(0.148)
UQR	-0.403	(0.164)	-0.213	(0.117)	-0.002	(0.101)	0.212	(0.117)	0.402	(0.165)
Scenario 2:										
RQR	-0.393	(0.174)	-0.245	(0.147)	0.004	(0.132)	0.251	(0.129)	0.399	(0.142)
PS-QTE	-0.394	(0.183)	-0.247	(0.159)	0.003	(0.145)	0.250	(0.140)	0.398	(0.150)
GQR	-0.393	(0.176)	-0.243	(0.149)	0.007	(0.132)	0.252	(0.128)	0.399	(0.139)
CQR	-0.316	(0.156)	-0.150	(0.134)	0.081	(0.122)	0.282	(0.126)	0.396	(0.145)
UQR	-0.320	(0.149)	-0.178	(0.108)	0.003	(0.100)	0.239	(0.125)	0.521	(0.195)

Note: See note in Table 1 for details.

Appendix Figure A1: The Monte Carlo Error in simulation scenarios 1 (Panel A) and 2 (Panel B) in different quantile regression approaches and at different quantiles.

Note: The Monte Carlo Error is defined as: $\sqrt{\frac{1}{n-1} \sum_{j=1}^n (\hat{\beta}_j^{(\tau)} - E[\hat{\beta}_j^{(\tau)}])^2} / n$, where j index simulated dataset and n is the total number of simulated datasets (White, 2010). Online Appendix C includes Stata Code to replicate the graph

Appendix Figure A2: Quantile regression coefficients from simulation scenarios 1 (Panel A) and 2 (Panel B).

Note: See note in Figure 2 for details. Online Appendix C includes Stata Code to replicate the graph.

Appendix Figure A3: Effects of union status on log wages for full-time working males in the 1983-1985 Outgoing Rotation group supplement of the Current Population Survey (N=266,956).

Note: The included control variables are nine dummies for labor market experience, five dummies for education, a dummy variable for married, and a dummy variable for non-white. The UQR model is implemented using the `rifreg` command. Data is available as supplemental material to Firpo et al. (2009) and can be downloaded from: <https://www.econometricsociety.org/publications/econometrica/2009/05/01/unconditional-quantile-regressions>.

Appendix Figure A4: Coverage rate of 95% confidence intervals based on asymptotic standard errors and various bootstrapped confidence intervals (2000 repetitions) in simulation scenario 2 (1,000 draws of N=2,000).

Note: The standard confidence intervals are based on asymptotic standard errors. These standard errors assume the residuals are independent and identically distributed and do not account for the uncertainty in the first stage (residualization). In each simulation draw, we record whether the 95% confidence intervals include the true value ($C95_j$). The coverage rate calculates the proportion of the confidence intervals that include the true value: $\frac{1}{n} \sum_{j=1}^n (C95_j)$, where j index simulated dataset and n is the total number of simulated datasets (Heisig et al., 2017). The average coverage rate across the five quantiles is 94.86% (standard CIs), 94.34% (Normal-Approximation Bootstrap CIs), 95.26% (Percentile Bootstrap CIs), and 94.52% (Bias-Corrected Bootstrap CIs).

Supplementary Online Appendix B: Additional data simulations

This appendix supplements the data simulations in the main text by providing a broader set of simulation scenarios. To begin, we simulate data with a binary treatment variable (as in the main text) (Appendix Figure B1), a treatment variable with a uniformly distributed random component (Appendix Figure B2), and a treatment variable with a normally distributed random component (Appendix Figure B3). Further, we distinguish between three structures on the QTEs: constant (panels A and D), quadratic (panels B and E), and cubic (panels C and F). Finally, we allow the outcome's residual distribution to be either normal (panels A-C) or right-skewed (panels D-F). See Appendix Table B1 for a description of the data generating process for each of the simulation scenarios.

Appendix Table B1: Data generating process for the graphs in Appendix Figures B1 (binary treatment), B2 (uniform treatment), and B3 (normal treatment).

	Treatment variable	Pre-treatment outcome	Outcome variable
Appendix Figure B1			
Panel A	$P(t_i = 1 x_i = 0) = .20$ $P(t_i = 1 x_i = 1) = .067$	$y_i^0 = x * 1.3 + \varepsilon_i$ where $\varepsilon_i \sim N(0,1)$	$y_i^1 = t * .5 + y_i^0$
Panel B			$y_i^1 = (r_i - r_i^2) * t * 2.15 + y_i^0$
Panel C		$y_i^0 = x * .07 + \varepsilon_i$ where $\varepsilon_i \sim \text{Beta}(1.3,30)$	$y_i^1 = (-3.7 * (r_i - .5)^3 + .5) * t * .8 + y_i^0$
Panel D			$y_i^1 = t * .02 + y_i^0$
Panel E			$y_i^1 = (r_i - r_i^2) * t * .1 + y_i^0$
Panel F			$y_i^1 = (-3.7 * (r_i - .5)^3 + .5) * t * .025 + y_i^0$
Appendix Figure B2			
Panel A	$t_i = x_i * .45 + e_i$ where $e_i \sim U[0,1]$	$y_i^0 = x * 1.3 + \varepsilon_i$ where $\varepsilon_i \sim N(0,1)$	$y_i^1 = t * .5 + y_i^0$
Panel B			$y_i^1 = (r_i - r_i^2) * t * 2.15 + y_i^0$
Panel C			$y_i^1 = (-3.7 * (r_i - .5)^3 + .5) * t * .8 + y_i^0$
Panel D		$y_i^0 = x * .04 + \varepsilon_i$ where $\varepsilon_i \sim \text{Beta}(1.3,30)$	$y_i^1 = t * .02 + y_i^0$
Panel E			$y_i^1 = (r_i - r_i^2) * t * .1 + y_i^0$
Panel F			$y_i^1 = (-3.7 * (r_i - .5)^3 + .5) * t * .025 + y_i^0$
Appendix Figure B3			
Panel A	$t_i = x_i * .90 + e_i * .7$ where $e_i \sim N(0,1)$	$y_i^0 = x * 1.3 + \varepsilon_i$ where $\varepsilon_i \sim N(0,1)$	$y_i^1 = t * .5 + y_i^0$
Panel B			$y_i^1 = (r_i - r_i^2) * t * 2.15 + y_i^0$
Panel C		$y_i^0 = x * .07 + \varepsilon_i$ where $\varepsilon_i \sim \text{Beta}(1.3,30)$	$y_i^1 = (-3.7 * (r_i - .5)^3 + .5) * t * .8 + y_i^0$
Panel D			$y_i^1 = t * .02 + y_i^0$
Panel E			$y_i^1 = (r_i - r_i^2) * t * .1 + y_i^0$
Panel F			$y_i^1 = (-3.7 * (r_i - .5)^3 + .5) * t * .025 + y_i^0$

Note: The r_i is the individual i 's percentile rank ($r_i \sim U[0,1]$) in the pre-treatment outcome distribution (y_i^0). x_i is a binary treatment where 25% = 1.

Appendix Figure B1: Quantile treatment effects of a binary treatment variable in data simulations with different treatment effect structures and outcome variables (1,000 draws of $N=2,000$).

Note: Data simulation is performed in Stata 16.0, and files to replicate the results are available in Online Appendix K. Data generating process is explained in Appendix Table B1. QTEs are constant in panels A and D, quadratic in panels B and E, and cubic in panels C and F. The outcome has a normally distributed error term in panels A-C and a right-skewed error term in panels D-F. The reported coefficients are the regression coefficients at each quantile divided by the outcome variable's standard deviation.

Appendix Figure B2: Quantile treatment effects of a treatment variable with a uniform error term in data simulations with different treatment effect structures and outcome variables (1,000 draws of $N=2,000$).

Note: Data simulation is performed in Stata 16.0, and files to replicate the results are available in Online Appendix K. Data generating process is explained in Appendix Table B1. QTEs are constant in panels A and D, quadratic in panels B and E, and cubic in panels C and F. The outcome has a normally distributed error term in panels A-C and a right-skewed error term in panels D-F. The reported coefficients are the regression coefficients at each quantile divided by the outcome variable's standard deviation.

Appendix Figure B3: Quantile treatment effects of a treatment variable with a normally distributed error term in data simulations with different treatment effect structures and outcome variables (1,000 draws of $N=2,000$).

Note: Data simulation is performed in Stata 16.0, and files to replicate the results are available in Online Appendix K. Data generating process is explained in Appendix Table B1. QTEs are constant in panels A and D, quadratic in panels B and E, and cubic in panels C and F. The outcome has a normally distributed error term in panels A-C and a right-skewed error term in panels D-F. The reported coefficients are the regression coefficients at each quantile divided by the outcome variable's standard deviation.

Supplementary Online Appendix C: Specification tests and incorrectly specified models

Causal modeling based on covariate adjustment strategies requires that confounders are perfectly measured (selection on observables) and that the model is correctly specified. A misspecified model will render the treatment residuals correlated with the observed confounders, which would bias the second-step QTE estimation. Various approaches can be used to test model misspecification. For example, typical specification tests in OLS can be applied (Pregibon, 1980; Ramsey, 1969), or one could add interactions or higher-order polynomials in the first step and check whether they significantly improve model fit. Somewhat related to these approaches, we propose to check whether the treatment residuals are uncorrelated with all interactions and fourth-order polynomials using a model F-test.

This appendix uses data simulations to illustrate the implications of a misspecified first-step model and shows how the specification test can be used to detect misspecification. Consider the following DGP, where the treatment t_i is defined as $t_i = x1_i * .20 + x2_i * .30 - x3_i * .50 + x3_i^2 * .70 + x1_i x2_i * .40 + e_i$, where $e_i, x2_i$, and $x3_i \sim N(0,1)$ and $x1_i$ is a binary variable with 25%=1. The pre-treatment outcome y_i^0 is defined as $y_i^0 = x1_i * 1 + x2_i * 1.2 + x3_i * .90 + \epsilon_i$, where $\epsilon_i \sim N(0,1)$. Finally, we allow the strength of the treatment variable (t_i) to depend on the individual i 's percentile rank ($r_i \sim U[0,1]$) in the pre-treatment outcome distribution (y_i^0). As in the main data simulations, the QTE increases across the outcome distribution monotonically: $y_i = (r_i - .50) * t_i + y_i^0$.

Using this DGP, we calculate treatment residuals based on four different models, where models A-C are misspecified whereas model D is correctly specified:

$$\text{Model A: } t_i = \beta_1 x1_i + \beta_2 x2_i + \beta_3 x3_i + \epsilon_i$$

$$\text{Model B: } t_i = \beta_1 x1_i + \beta_2 x2_i + \beta_3 x3_i + \beta_4 x3_i^2 + \epsilon_i$$

$$\text{Model C: } t_i = \beta_1 x1_i + \beta_2 x2_i + \beta_3 x3_i + \beta_5 x1_i x2_i + \epsilon_i$$

$$\text{Model D: } t_i = \beta_1 x1_i + \beta_2 x2_i + \beta_3 x3_i + \beta_4 x3_i^2 + \beta_5 x1_i x2_i + \epsilon_i$$

Appendix Figure C1: Estimated bias with misspecified (Panel A-C) and correctly specified first-step regression with RQR and GQR (2,0000 draws of $N=2,000$).

Note: Data simulation is performed in Stata 16.0, and files to replicate the results are available in Online Appendix K.

We begin by showing that a misspecified model biases the QTE estimation. Figure C1 shows the bias in the estimated QTE with the different model setups using the RQR and GQR models. As expected, although all relevant confounders are observed, failure to correctly specify the model results in biased estimates (Panel A-C). The bias is similar in magnitude for the RQR (blue lines) and the GQR (red lines) models, but the GQR is more biased at the top while RQR is more biased at the bottom. When the model is correctly specified, there is no bias with any models (Panel D).

The bias caused by a misspecified first-step regression in models A-C could be detected using a specification test. After calculating the treatment residuals based on model specifications A-D above, we regress the treatment residuals on a covariate list that includes fourth-order polynomials of continuous variables and interactions between all variables (including polynomials). Any non-zero

coefficients would indicate that the first-step regression has not sufficiently made the treatment residuals independent of the observed confounders. We then jointly test whether all coefficients are zero using an F-test.

Table C1 shows the F-statistic and the p-value after regressing the treatment residuals on the full list of covariates. In the misspecified models, the F-test rejects the null hypothesis that all coefficients are zero, with p-values below 0.001 (Panels A-C). As a sanity check, we have also examined whether the correctly specified model is rejected using the same F-test. The p-value with the correctly specified model is 0.83, which correctly suggests that there are no unmodeled interactions or higher-order polynomials (Panel D).

Table C1: Joint F-test of covariates.

	F	Prob > F
Panel A: Misspecified interaction and quadratic term	106.5258	0.0000
Panel B: Misspecified interaction	3.9409	0.0002
Panel C: Misspecified quadratic term	103.3983	0.0000
Panel D: Correctly specified	0.7391	0.7372

Note: Average F-values and p-values across the 2,000 simulations.

Supplementary Online Appendix D: Heterogeneous effects of confounders on treatment

The confounder has a constant effect on the treatment in most simulations in this paper. In Appendix C, we relax this assumption by allowing interactions between confounders. In this appendix, we allow for the effect of the confounder on the treatment to vary randomly across individuals. We will show that the RQR model handles such heterogeneity well. We define a confounder $x_i \sim N(10, 2)$ and let the treatment be equal to $t_i = a_i x_i + \epsilon_i$ where $a_i \sim U[0, 1]$ and $\epsilon_i \sim N(0, 1)$. Then, we define the pre-treatment outcome as $y_i^0 = x_i * 1 + \epsilon_i$, where $\epsilon_i \sim N(0, 2)$. Finally, we define the observed outcome as $y_i = ((r_i - 0.50) * .1)t_i + y_i^0$, where $r_i \sim U[0, 1]$ is the rank in the pre-treatment distribution. One implication of this simulation setup is that the variance of the treatment residuals increases as a function of the confounder, as shown in Appendix Figure D1. (See Appendix E for more about confounders influencing the variance of the treatment residuals.) This data simulation shows that the bias is small and close to zero with the RQR and GQR models, as shown in Appendix Figure D2.

Appendix Figure D1: Scatter plot between treatment residuals and the confounder x_i in a single draw (N=10,000).

Appendix Figure D2: Estimated bias with RQR and GQR and a non-constant effect of the confounder (2,000 draws of $N=2,000$).

Supplementary Online Appendix E: Counfounders influence variance

In the RQR model, we must assume that the treatment residuals are mean independent of confounders for the coefficients to have a causal interpretation. This appendix uses data simulations to illustrate that the RQR model provides unbiased coefficients despite the confounders affecting the variance of the treatment residuals. We define a confounder $x_i \sim N(10, 2)$ and let the treatment be equal to $t_i = x_i + \epsilon_i$ where $\epsilon_i \sim N(0, 1 * x_i)$. Then, we define the pre-treatment outcome as $y_i^0 = x_i * 1 + \epsilon_i$, where $\epsilon_i \sim N(0, 2)$. Finally, we define the observed outcome as $y_i = ((r_i - 0.50) * .1)t_i + y_i^0$, where $r_i \sim U[0, 1]$ is the rank in the pre-treatment distribution. Appendix Figure E1 clearly shows that although the treatment residuals are mean independent of the confounder, the variance increases as a function of the confounder (i.e., heteroscedasticity). However, this is not a concern since the RQR model only requires that the treatment is mean independent of the confounders: the RQR coefficients are unbiased, as Appendix Figure E2 shows.

Appendix Figure E1: Scatter plot between treatment residuals and the confounder x_i in a single draw (N=10,000).

Appendix Figure E2: Estimated bias with RQR and GQR when confounder influences the variance of the treatment (2,000 draws of N=2,000).

Supplementary Online Appendix F: Heterogeneous effects of confounders on the outcome

The Supplementary Appendices C, D, and E have dealt with various forms of heterogeneity in how confounders influence the treatment. This appendix simulates data where the confounder has a curvilinear effect on the outcome and data where the confounder has a quantile treatment effect on the outcome. Concerning the curvilinear scenario, we start by defining a confounder $x_i \sim N(0,2)$ and let the treatment be equal to $t_i = x_i + \epsilon_i$ where $\epsilon_i \sim N(0,1)$. Then, we define the pre-treatment outcome as $y_i^0 = x_i * 1 - x_i^2 * .5 + \epsilon_i$, where $\epsilon_i \sim N(0,2)$. Finally, we define the observed outcome as $y_i = ((r_i - 0.50) * .1)t_i + y_i^0$, where $r_i \sim U[0,1]$ is the rank in the pre-treatment distribution. In the quantile treatment effects scenario, we define the confounder, the treatment, and the outcome in the same way as in the curvilinear scenario. However, the pre-treatment outcome is defined as $y_i^0 = u_i * x_i + \epsilon_i$, where $\epsilon_i \sim N(0,2)$ and $u_i \sim U[0,1]$ is the rank in ϵ_i . Appendix Figure F1 shows that the RQR model provides unbiased estimates in both simulation scenarios. The GQR model is slightly more biased in this specific scenario, although the absolute bias is small.

Appendix Figure F1: Estimated bias with RQR and GQR when the confounder has a curvilinear effect on the outcome (panel a) and a quantile treatment effect on the outcome (panel b) (10,000 draws of $N=2,000$).

Supplementary Online Appendix G: Unobserved time-invariant confounding

The data simulation in this appendix deviates from the simulations in the main text and the other appendices in that it simulates panel data with time-invariant unobserved confounding. Thus, we need to apply a fixed effects estimator to estimate the QTEs. We create panel data for individuals i measured at time point j . Then, we define a time-invariant unobserved confounder that varies at the individual level only: $u_i \sim N(0,1)$ and an observed confounder that is time-variant $x_{ij} \sim N(0,1)$. The treatment is defined as $t_{ij} = x_{ij} + time_{ij} * .5 + u_i + \epsilon_{ij}$ where $\epsilon_{ij} \sim N(0,2)$. Then, we define the pre-treatment outcome as $y_{ij}^0 = x_{ij} * 1 + u_i * 1 + time_{ij} * 1 + \epsilon_{ij}$, where $\epsilon_{ij} \sim N(0,1)$. Finally, we define the observed outcome as $y_{ij} = ((r_{ij} - 0.50) * .4)t_{ij} + y_{ij}^0$, where $r_{ij} \sim U[0,1]$ is the rank in the pre-treatment distribution.

The results in Appendix Figure G1 show that the RQR and GQR models are considerably biased when the time-invariant unobserved confounding is not accounted for. However, using a fixed effects estimator in the first-step regression in the RQR model provides unbiased QTE estimates.

Appendix Figure G1: Estimated quantile treatment effects in panel data with time-invariant unobserved confounding with fixed effects (RQR-FE) and without fixed effects (RQR and GQR) (10,000 reps of 1,000 individuals measured at 4 time points).

Supplementary Online Appendix H: Rank invariance assumption

Rank invariance means that an individual's rank in the potential outcome distribution is the same across treatment states. Consider a binary treatment variable t and let y_i^0 and y_i^1 be the potential outcomes for $t = 0, 1$. Further, let $r_i^0 \sim U[0, 1]$ and $r_i^1 \sim U[0, 1]$ be the potential ranks for $t = 0, 1$. Rank invariance is then defined as $r_i^0 = r_i^1$. This appendix uses data simulations to illustrate that failure to meet the rank invariance assumption biases QTE models' ability to identify individual-level QTEs, including the RQR model.

Building on the data simulations in Frandsen and Lefgren (2018), we define a bivariate confounder x_i (25%=1) and let the probability of being treated depend on x_i : $P(t_i = 1|x_i = 0) = 0.067$ and $P(t_i = 1|x_i = 1) = 0.20$. Further, we define the outcome under treatment as $y_i^1 = (1 - \zeta)y_i^0 + \zeta n_i^0 + 1$, where y_i^0 and $n_i^0 \sim N(0, .75)$ and ζ determines the degree of departure from rank invariance. Setting $\zeta = 0$ corresponds to rank invariance, while setting $\zeta = 1$ implies that the potential ranks are orthogonal. Finally, the outcome is defined as $y_i = t_i y_i^1 + (1 - t_i) y_i^0 + x_i$.

Using this setup, we estimate quantile regression coefficients at the 10th, 25th, 75th, and 90th quantile, where we vary ζ from 0 to 0.50. In each of the simulation scenarios, we calculate the correlation between the ranks in the potential outcome distributions (i.e., $cor(r_i^0, r_i^1)$) and plot the bias in the estimates by this correlation. Appendix Figure H1 shows the bias on the y-axis and the correlation between potential ranks on the x-axis. When the rank invariance assumption holds (i.e., $cor(r_i^0, r_i^1) = 1$), then the RQR model is unbiased. However, when the rank invariance assumption is not met, the RQR model provides biased estimates of the individual-level QTE. Note that deviations from the rank invariance assumption affect the GQR and PS-QTE approach similarly, with the bias seemingly somewhat stronger in the matching approach in this example.

Appendix Figure H1: Estimated bias in simulation scenarios where rank invariance does not hold (10,000 draws of $N=2,000$).

Supplementary Online Appendix I: Quantile crossing

This appendix illustrates that the RQR coefficients cannot be interpreted as QTEs unless the quantile function is strictly monotonically increasing. We simulate two data sets. First, we generate a data set without quantile crossing. We define a confounder $x_i \sim N(-1, 2)$ and let the treatment be equal to $t_i = x_i + \epsilon_i$ where $\epsilon_i \sim N(0, 2.5)$. Then, we define the pre-treatment outcome as $y_i^0 = x_i * 1 + \epsilon_i$, where $\epsilon_i \sim N(0, 2)$. Finally, we define the observed outcome as $y_i = ((r_i - 0.50) * .4)t_i + y_i^0$, where $r_i \sim U[0, 1]$ is the rank in the pre-treatment distribution. The data set where quantiles crossing is similar, except that the variation in treatment effect across quantiles is larger: $((r_i - 0.50) * 1.5)t_i$.

Appendix Figure I1: An example of a quantile crossing problem in simulated data from one draw (N=5,000). Panel a is without quantile crossing (within the observed variable space), while panel b is with quantile crossing.

Note: The solid lines show predicted quantile values for different levels of the treatment residuals, based on a linear quantile regression model. The scatter plots are observed unit values on treatment residuals and the outcome.

Quantile crossing implies that the estimated quantile curves cross one another within the bounds of the observations (He, 1997), as illustrated in panel b in Appendix Figure I1.¹¹ For units with treatment residual values below -7 in this example, the estimated quantile value at the 10th percentile is estimated to be higher than the corresponding quantile value at higher quantiles. Such quantile crossing violates the requirement that distribution functions and their associated inverse functions should be monotone increasing (Koenker, 2005, p. 56). The implication is biased estimates, as can be seen from panel b in Figure I2: when the quantiles cross, the RQR and GQR coefficients are considerably biased. However, both models identify the QTE without bias when the quantile function is strictly monotonically increasing, as in panel a.

Appendix Figure I2: Estimated bias in scenarios with and without quantile crossing (2,000 reps of N=2,000).

¹¹ Unless the quantile lines are parallel, the quantile functions must cross somewhere. However, it should be sufficiently far away from the mean of the predictor.

Supplementary Online Appendix J: Discrete outcome variable

Quantile regressions estimators may fail to identify treatment effects with discrete outcome variables, such as count data (Machado & Silva, 2005). As shown by this appendix, this also applies to the RQR model, since it estimates a bivariate CQR model with the treatment residuals as the only independent variable. Consider the following setup, where we have a confounder $x_i \sim N(0,2)$ and let the treatment be equal to $t_i = x_i + \epsilon_i$ where $\epsilon_i \sim N(0,2)$. Then, we define the pre-treatment outcome as $y_i^0 = x_i * 1 + \epsilon_i$, where $\epsilon_i \sim N(4,7)$. Finally, we define the observed outcome as $y_i = ((r_i - 0.50) * .4)t_i + y_i^0$, where $r_i \sim U[0,1]$ is the rank in the pre-treatment distribution.

Appendix Figure J1: Quantile plot of the rounded outcome variable.

The crux is that we round the outcome to the nearest integer value. In our single draw of $N=100,000$ observations, this rounding results in only 62 unique values of the outcome distribution, meaning that not all quantiles are separable in the data, as shown by the quantile plot in Appendix Figure J1. The RQR model fails to estimate QTEs with this data, as can be seen by the RQR coefficients

for the discrete outcome in panel a in Appendix Figure J2 (red line with circle). The GQR model is slightly less biased, as can be seen in panel b. However, given this setup, where the outcome is the discretized result of a continuous underlying process, we could easily solve this problem by adding uniformly distributed noise to the outcome variable (Machado & Silva, 2005): the blue line with squared symbols is close to the true QTE.

Appendix Figure J2: Estimated RQR and GQR coefficients from one single draw with a discrete outcome variable and a jittered outcome variable (N=50,000).

Supplementary Online Appendix K: Replication files

Stata 16.0 is used throughout this paper, and this supplementary online appendix includes datasets and Stata do-files to replicate all the results. For the review process, data and code to replicate the results are made available on SocArXiv using an anonymized view-only link. Reviewers and editors can access Appendix K by using the following link (unzip the AppendixK.zip folder and read the instructions in the README.txt file):

https://osf.io/ewntp/?view_only=a9260eed331e4987938dec81c328f8c8

If our paper is accepted for publication, we will make data and code available as online supplemental material.