

TTPI

Tax and Transfer Policy Institute

Quantile treatment effects revisited: Uncovering the distributional consequences of a welfare experiment

TTPI - Working Paper 18/2023 November 2023

Mathias Sinning

Tax and Transfer Policy Institute

Crawford School of Public Policy, Australian National University

Abstract

Heterogeneous effects of welfare reforms on earnings, transfers, and income have been established theoretically and empirically. Evaluation studies often focus on quantile treatment effects (QTE), which rely on the marginal distributions of potential treatment and control outcomes. Parameters that depend on the joint distribution of potential outcomes, such as quantiles of the distribution of treatment effects (QDTE), receive less attention. We propose a strategy to identify these parameters. We leverage the property that, under random assignment, rank correlation coefficients between actual treatment and predicted control state outcomes are identical, irrespective of whether predictions are based on treatment or control units. To identify QDTE, we assume that all permutations of observation units satisfying this property are equally likely. Rearranging quantiles yields a generalized version of quantile treatment effects (GQTE). We employ a reweighting approach for identification under strong ignorability. We test the predictor strength and demonstrate that highly predictive covariates yield unbiased, consistent, and asymptotically normal estimators. Our analysis of Connecticut's Jobs First program reveals initial income increases for a larger fraction of participants than previously recognized. Long-term gains were at least twice as large as those derived from conventional QTE and concentrated at the lower end of the distribution.

JEL Codes: C13, C21, D04, I38**Keywords:** Treatment effect heterogeneity, distribution of treatment effects, rank correlation, welfare experiments

** I thank Josh Angrist, Colin Cameron, Ashley Craig, Christian Gillitzer, Tue Gørgens, James Heckman, Kosuke Imai, Gary King, Myoung-Jae Lee, Seojeong Lee, Thomas Lemieux, Sangsoo Park, Jeffrey Smith and Xiang Zhou for helpful conversations and comments on earlier drafts. I also thank participants of the Applied Statistics Seminar at the Harvard University Institute of Quantitative Social Science, the Vancouver School of Economics Seminar at the University of British Columbia, and the 2019 Asia and China Meetings of the Econometric Society. This work was supported by the Australian Research Council under Grants DP150104247, LP160100810, and LP220100291. I report there are no competing interests to declare. All correspondence to Mathias Sinning, E-mail: mathias.sinning@anu.edu.au*

Tax and Transfer Policy Institute

Crawford School of Public Policy

College of **Asia and the Pacific**

+61 2 6125 9318

tax.policy@anu.edu.au

The Australian National University

Canberra ACT 0200 Australia

www.anu.edu.au

The Tax and Transfer Policy Institute (TTPI) is an independent policy institute that was established in 2013 with seed funding from the federal government. It is supported by the Crawford School of Public Policy of the Australian National University.

TTPI contributes to public policy by improving understanding, building the evidence base, and promoting the study, discussion and debate of the economic and social impacts of the tax and transfer system.

The Crawford School of Public Policy is the Australian National University's public policy school, serving and influencing Australia, Asia and the Pacific through advanced policy research, graduate and executive education, and policy impact.

1 Introduction

Understanding the distributional consequences of welfare reforms is important for the design of tax and transfer policies. Transfer programs, which aim to provide social security and enhance equity, can affect labor supply and earnings (Kline and Tartari, 2016). Some design features of these programs, such as a generous earnings disregard, can lead to increased earnings for certain segments of the distribution while simultaneously reducing earnings for other segments (Bitler et al., 2006). A key challenge when studying the effectiveness of transfer programs is that even ideal experiments can only be used to identify features of the distribution of treatment effects that depend on separate marginal distributions of treatment and control outcomes, such as the average treatment effect (ATE) or quantile treatment effects (QTE).¹

Bitler et al. (2006) analyze data from Connecticut’s Jobs First welfare experiment and find that QTE estimates display heterogeneous treatment responses consistent with predictions from labor supply theory. Using the same data source, Bitler et al. (2017) formally test and reject the hypothesis that subgroup-specific ATE can explain important features of the treatment effect heterogeneity revealed by QTE. However, the authors emphasize that QTE are different from quantiles of the distribution of treatment effects (QDTE) if the rank invariance assumption is violated. Bitler et al. (2005) provide evidence of rank reversals in a related study based on Canadian data, suggesting that their QTE estimates may not be interpreted as QDTE.²

This paper proposes an approach to identify QDTE. The approach is used to explore the distributional consequences of Connecticut’s Jobs First welfare experiment on women’s earnings, transfers, and income. We begin by considering rank correlation coefficients (Kendall, 1938) between actual treatment outcomes and predicted control

¹Experiments are informative about ATE without any further identifying assumptions. In contrast, a rank invariance assumption is required to identify QTE as defined by Doksum (1974) and Lehmann (1974). Rank invariance implies that observation units maintain their relative position in the potential outcome distribution regardless of whether they are being assigned to the treatment group or the control group. Although rank invariance is a strong assumption, QTE estimation has become a standard approach in the treatment effects literature. Important methodological contributions include Chernozhukov and Hansen (2005), Firpo (2007) and Frölich and Melly (2013).

²QDTE are only identical to QTE if the rank invariance assumption holds and if QTE are monotonically increasing along the distribution. In general, the interpretation of QTE is quite different from that of QDTE (Lee, 2006). For example, consider a welfare reform that aims to increase individual income. In this context, a positive median treatment effect implies that the income of the median person has increased. In contrast, a positive median of the distribution of treatment effects indicates that the incomes of at least 50 percent of the study population have increased.

state outcomes in a scenario where treatment and control group sizes are equal. Under random assignment, we obtain identical rank correlation coefficients, regardless of whether we use treatment or control units to calculate predicted outcomes. We refer to this property as ‘Rank Correlation Property’ (RCP). We show that the rank correlation coefficient between potential treatment and control outcomes is identified under the assumption that all permutations of observation units satisfying the RCP are equally likely. We use the same assumption to identify QDTE as quantiles of average unit-level treatment effects derived from permutations of observation units satisfying the RCP. Rearranging QDTE yields a generalized version of quantile treatment treatment effects (GQTE), which do not require a rank invariance assumption but align with conventional QTE under rank invariance.³

Our identification strategy rests on modest assumptions because we focus on situations in which relevant control outcome predictors are observed. Many experimental studies collect pre-treatment outcomes along with other highly predictive covariates. We propose a simulation-based test to assess the strength of model predictors, measured by the rank correlation coefficient between actual and predicted control outcomes. We use simulated data to estimate specific target values of the rank correlation coefficient between potential treatment and control outcomes. We test whether the estimated rank correlation coefficients deviate significantly from these target values. We also employ a Kolmogorov-Smirnov test to ascertain whether deviations of estimated rank correlation coefficients from target values have an impact on our GQTE estimates. While we cannot test whether all permutations of observation units satisfying the RCP are equally likely, we demonstrate that the distribution of permutation-specific GQTE is approximately normal for each quantile, strongly supporting the use of averages associated with the assumption of equally likely permutations.

We extend our approach to identify our parameters of interest under strong ignorability. We employ a reweighting approach akin to Firpo (2007) to obtain reweighted versions of actual and predicted outcomes. We use these reweighted outcomes to specify a ‘conditional RCP’. The identification of QDTE under strong ignorability is based

³The distribution of treatment effects also serves as a basis for deriving various other parameters (Heckman et al., 1997; Heckman and Vytlačil, 2007). In our empirical application, we will determine the proportion of women who benefit from Jobs First and calculate the total gains and losses resulting from the program.

on the assumption that all permutations of observation units satisfying the conditional RCP are equally likely. Finally, we extend our approach to situations in which treatment and control group sizes are unequal. Monte Carlo simulations confirm that our estimators are unbiased, consistent, and asymptotically normal if relevant predictors are available. We also find that estimated bootstrap standard errors are a good representation of the true sampling variation.

Following Bitler et al. (2006), we estimate separate effects for the periods preceding and following the implementation of time limits, which led to the exclusion of Jobs First women from welfare benefits. During the pre-time limit period (quarters 1-7 after the initiation of Jobs First), our estimated rank correlation coefficients between potential treatment and control outcomes range from approximately 0.7 to 0.8. During the post-time limit period (quarters 8-16), this range shifts to approximately 0.6 to 0.7 due to mobility over time. We test the relevance of our predictors and find that they are sufficiently strong to produce unbiased results.

Our findings confirm substantial heterogeneity in response to the Jobs First welfare experiment. Our GQTE estimates exhibit qualitative similarity to conventional QTE estimates during the pre-time limit period. In contrast, when we account for rank changes during the post-time limit period, we uncover both positive effects at the lower end and negative effects at the higher end of the earnings distribution. Negative effects at the higher end of the earnings distribution suggest that some women in the Jobs First group were either unable or unwilling to increase their earnings, even though incentives to reduce earnings to receive higher welfare payments disappeared with the implementation of time limits. Instead, our findings indicate that the initially observed negative effects of Jobs First on the earnings distribution were more persistent than previously recognized. Regarding income, we find that women with lower incomes are most responsive to reductions in transfers following the implementation of time limits. We also observe negative income effects at the higher end of the distribution, which align with the corresponding negative earnings effects.

We use QTE and GQTE estimates to assess the proportions of positive and negative treatment effects, revealing substantial disparities between the two methods. During the pre-time limit period, QTE estimates yield a lower proportion of positive income effects (85.9%) than GQTE estimates (99.0%). When focussing on effects that reach statistical

significance at a 0.1% level, the proportion of positive income effects based on QTE estimates drops to 52.5%, whereas the proportion based on GQTE estimates remains higher at 71.7%. During the post-time limit period, QTE and GQTE estimates exhibit similar proportions of statistically significant income effects, although the locations of these effects differ substantially between the two methods.

We also use our estimates to quantify the total gains and losses emanating from the Jobs First program. We find that the total gains and losses during the post-time limit period are preliminary driven by the impact of Jobs First on the earnings distribution. Moreover, our findings indicate that the increase in earnings successfully offset the majority of losses in transfers stemming from the implementation of time limits. We observe total gains in income of approximately \$190,000 based on QTE estimates and around \$383,000 based on GQTE estimates. When we apply a significance level of 0.1%, these gains decrease to approximately \$43,000 based on QTE estimates, whereas the corresponding figure based on GQTE estimates is approximately \$260,000, more than six times higher.

The remainder of this paper is structured as follows: Section 2 provides a brief overview of related literature. Section 3 introduces some basic ideas and defines our parameters of interest. Identification is covered in Section 4, followed by estimation in Section 5. The results of our empirical application are presented in Section 6. Section 7 concludes.

2 Related literature

The study of heterogeneous treatment effects has traditionally relied on estimating ATE within subgroups that share the same observed characteristics (Crump et al., 2008; Lee and Shaikh, 2013; List et al., 2019). In recent years, more adaptable approaches for estimating heterogeneous treatment effects have emerged, often leveraging machine learning techniques. Examples include Su et al. (2009), Hill (2011), Imai and Ratkovic (2013), Athey and Imbens (2016), Shalit et al. (2017), Chernozhukov et al. (2018), Wager and Athey (2018), Powers et al. (2018), Künzel et al. (2019), Hahn et al. (2020), and Nie and Wager (2021). While subgroup analysis can yield valuable insights, it may overlook critical aspects of treatment effect heterogeneity when essential information

is unobserved. Most importantly, subgroup analysis cannot be used to determine the distribution of treatment effects and the parameters derived from it. Our approach addresses this issue.

Our work builds on the growing literature on the distribution of treatment effects. A major strand of this literature focuses on the identification of features of the distribution of treatment effects without making any identifying assumptions. Because point identification of features beyond the average cannot be achieved without further assumptions, this literature uses partial identification and proposes methods to bound features of the distribution of treatment effects (Makarov, 1982; Rüschenendorf, 1982; Frank et al., 1987; Williamson and Downs, 1990; Heckman et al., 1997; Fan and Wu, 2010; Fan and Park, 2010; Firpo and Ridder, 2019; Russel, 2021). Unfortunately, the bounds arising from this literature are often too wide to be informative.

Another strand imposes restrictions to tighten the bounds. Manski (1997) establishes bounds under the assumption of monotone treatment response. Heckman et al. (1997) explore bounds within the framework of a Roy model. Kim (2018) proposes a generalization of these restrictions by considering support restrictions on potential outcomes. Using the same data as in our analysis, Kline and Tartari (2016) derive bounds from restrictions on revealed preferences. Frandsen and Lefgren (2021) assume that potential outcomes are mutually stochastically increasing. Lee (2022) imposes stochastic dominance assumptions to bound the conditional distribution of treatment effects. While restrictions can substantially tighten the bounds on the distribution of treatment effects, there is no guarantee that these bounds will necessarily provide useful insights in empirical applications. Moreover, restrictions are often implausible in real-world settings.⁴

Our approach shares a conceptual link with Frandsen and Lefgren (2021), who consider the correlation between predicted treatment and control outcomes to draw inferences about the correlation between potential outcomes. We show that the rank correlation coefficient between potential outcomes is identified if highly predictive covariates are available and if all permutations of observation units satisfying the RCP are

⁴For example, Kline and Tartari (2016) refer to instances where decision makers have been found violating revealed preference restrictions. Frandsen and Lefgren (2021) advise against applying their approach if, for instance, unobserved characteristics are expected to be advantageous in the control state but detrimental in the treatment state.

equally likely. There is also a connection between our approach and the work of Fan and Park (2009), who show that knowledge of the specific value of Kendall’s (1938) rank correlation coefficient between potential treatment and control outcomes does not improve the bounds on the distribution of treatment effects if the rank correlation coefficient is positive. We do not consider bounds on the distribution of treatment effects because we observe that the distribution of permutation-specific GQTE is approximately normal for each quantile. This finding strongly supports our focus on averages associated with equally likely permutations, rather than on bounds linked to permutations of observation units with near-zero probability of occurrence.

Our approach is more directly related to the work of Heckman et al. (1997), who derive features of the ‘impact distribution’ assuming knowledge of the rank correlation coefficient between potential outcomes.⁵ Heckman et al. (1997) pair percentiles of the marginal distributions of treatment and control outcomes to estimate the mean of features of the impact distribution. They use random permutations of percentiles that produce a given rank correlation coefficient and consider each permutation as equally likely. We employ a similar approach by using random permutations of observation units satisfying the RCP. While Heckman et al. (1997) obtain average features of the impact distribution from permutations of percentiles, we obtain features of average unit-level treatment effects from permutations of observation units.⁶

3 Basic ideas and parameters of interest

Before defining our parameters of interest, we introduce some notation, drawing upon the potential outcomes framework introduced by Rubin (1974). Following Imbens and Rubin (2015), we consider a super-population of size N_{sp} , where N_{sp} is large but countable. Throughout the paper, we use the subscript ‘sp’ to refer to the super-population. For each unit i in the super-population, $i \in \{1, \dots, N_{\text{sp}}\}$, we define the potential out-

⁵Several other studies impose restrictions on the dependence between potential outcomes and unobservables to identify the joint outcome distribution. Abbring and Heckman (2007) provide an overview of this literature.

⁶We illustrate in Section 6.2 that the impact distribution proposed by Heckman et al. (1997) captures the full amount of heterogeneity in the data by averaging over permutation-specific treatment effects that were arranged in ascending order, regardless of the location of control outcomes. In contrast, our identification strategy yields a distribution of treatment effects that is made up of average unit-level treatment effects associated with unit-level control outcomes.

comes $Y_i(1)$ and $Y_i(0)$, which are viewed as fixed. We define a treatment indicator, denoted as D_i , which takes on the value 1 if unit i is assigned to the treatment group, and 0 otherwise. For each unit i , only the potential outcome corresponding to their actual treatment received can be observed: $Y_i = Y_i(1)D_i + Y_i(0)(1 - D_i)$. We also consider a k -dimensional set of covariates X_i .

We present two commonly used assumptions regarding the assignment of units to treatment and control groups. One assumption asserts that units are assigned randomly:

ASSUMPTION 1 (Random Assignment): $Y_i(1), Y_i(0) \perp\!\!\!\perp D_i$.

Random assignment ensures that any differences between treatment and control outcomes are solely attributable to the treatment. Another commonly used assumption is the ‘strong ignorability’ assumption introduced by Rosenbaum and Rubin (1983). Under strong ignorability, the assignment of observation units to treatment and control groups is viewed as random (or as good as random), conditional on a set of covariates. The strong ignorability assumption can be stated as follows:

ASSUMPTION 2 (Strong Ignorability):

- (a) (Unconfoundedness): $Y_i(1), Y_i(0) \perp\!\!\!\perp D_i | X_i$.
- (b) (Common Support): $0 < Pr(D_i = 1 | X_i = x) < 1$.

Assumption 2(a) is commonly referred to as ‘unconfoundedness’ (Rubin, 1990) because it implies that comparing units that share the same characteristics is sufficient to eliminate all confounding.^{7,8} There are many situations in which this assumption is violated because selection is based on unobservables.⁹ Assumption 2(b) requires the presence of both treated and untreated units for each value of $X = x$.

⁷Assumption 2(a) is also known as ‘ignorability’ (Rosenbaum and Rubin, 1983), ‘selection on observables’ (Heckman and Robb, 1985) and ‘conditional independence assumption’ (Lechner, 1999).

⁸Many studies assume unconfoundedness to identify average treatment effects although their identification only requires a conditional mean independence assumption (i.e., $\mathbb{E}[Y_i(d) | X_i, D_i] = \mathbb{E}[Y_i(d) | X_i]$, with $d = \{0, 1\}$), which is considerably weaker than unconfoundedness (Heckman et al., 1998). In our case, unconfoundedness is relevant because we are not just interested in the mean but also in the quantiles of the distribution of treatment effects.

⁹Heckman and Vytlačil (2007) and Abbring and Heckman (2007) provide an overview of possible solutions associated with selection on unobservables.

Throughout the paper, we use $F_A(a)$ to denote the distribution function of a random variable A and $F_{A|B=b}(a|B=b)$ to denote the distribution function of A conditional on $B=b$. Moreover, for $u \in (0, 1)$, we define $q_u(F_A(a)) = \inf\{q : \Pr(A \leq q) \geq u\}$ as the quantile function of the distribution function $F_A(a)$. We are particularly interested in the distribution functions of the potential outcomes, $F_{Y(1)}(y)$ and $F_{Y(0)}(y)$. Following Rubin (1974), we consider the difference between potential outcomes, $\Delta_i = Y_i(1) - Y_i(0)$, which allows us to specify the distribution function of treatment effects, $F_\Delta(\delta)$. Our parameters of interest are formally defined as follows.

DEFINITION 1 (Population Quantiles of the Distribution of Treatment Effects): *For $u \in (0, 1)$, the Population Quantiles of the Distribution of Treatment Effects are given by*

$$q_{\Delta, u} = q_u(F_\Delta(\delta)) = \inf\left\{\delta : \frac{1}{N_{\text{sp}}} \sum_{i=1}^{N_{\text{sp}}} \mathbf{1}\{\Delta_i \leq \delta\} \geq u\right\}. \quad (1)$$

While this paper focuses on the Population Quantiles of the Distribution of Treatment Effects (PQDTE), similar results can be obtained for the Population Quantiles of the Distribution of Treatment Effects on the Treated (PQDTE), $q_u(F_{\Delta|D=1}(\delta|D=1))$. We consider two different estimands: Quantiles of the Impact Distribution (QID) and Quantiles of the Distribution of Treatment Effects (QDTE). Our use of QID is inspired by Heckman et al. (1997), who present estimates of quantiles of the ‘impact distribution’. Our QID estimator relies on quantiles of the distribution of unit-level treatment effects. While QID capture the full spectrum of heterogeneity in the data, they ignore that unit-level treatment effects associated with units having similar or identical control outcomes may be vastly different. We demonstrate that rearranging unit-level treatment effects by control outcome yields a generalized version of quantile treatment effects (GQTE), which do not require imposing a rank invariance assumption but align with conventional quantile treatment effects (QTE) under rank invariance. If the GQTE are monotonically increasing, they coincide with QDTE.

4 Identification

4.1 Measuring rank dependence

We use Kendall's (1938) rank correlation coefficient to specify the degree of rank dependence between potential outcomes in the treatment and control state. We consider the potential outcomes in the super-population. For any two pairs $(Y_i(1), Y_i(0))$ and $(Y_j(1), Y_j(0))$, $i, j \in \{1, \dots, N_{\text{sp}}\}$, we define the product $\pi_{ij} = (Y_i(1) - Y_j(1))(Y_i(0) - Y_j(0))$ as concordant if $\pi_{ij} > 0$ and as discordant if $\pi_{ij} < 0$. The total number of comparable pairs is $N_{\text{sp}}(N_{\text{sp}} - 1)/2$. Given the number of concordant and discordant pairs, p_c and p_d , Kendall's (1938) rank correlation coefficient between potential treatment and control outcomes in the super-population may be written as

$$\tilde{\tau}_{\text{sp}} = (p_c - p_d)/(N_{\text{sp}}(N_{\text{sp}} - 1)/2). \quad (2)$$

Throughout this paper, we consider a modified version of the rank correlation coefficient, which adjusts for ties,

$$\tau_{\text{sp}} = (p_c - p_d)/\sqrt{((N_{\text{sp}}(N_{\text{sp}} - 1)/2) - p_a)((N_{\text{sp}}(N_{\text{sp}} - 1)/2) - p_b)}, \quad (3)$$

where $p_a = \sum_{i=1}^{s_1} a_i(a_i - 1)/2$ and $p_b = \sum_{j=1}^{s_0} b_j(b_j - 1)/2$, and where a_i (b_j) is the number of tied $Y(1)$ ($Y(0)$) values in the i th (j th) set of ties, with s_1 (s_0) denoting the number of sets (Kendall and Gibbons, 1990).

In the following, we will use the notation ' $\tau(\mathbf{V}_1, \mathbf{V}_0)$ ' to denote the rank correlation coefficient between two vectors \mathbf{V}_1 and \mathbf{V}_0 of equal length. Given the vectors $\mathbf{Y}_{\text{sp}}(1) = (Y_1(1), \dots, Y_{N_{\text{sp}}}(1))'$ and $\mathbf{Y}_{\text{sp}}(0) = (Y_1(0), \dots, Y_{N_{\text{sp}}}(0))'$, we refer to the case in which $\tau_{\text{sp}} = \tau(\mathbf{Y}_{\text{sp}}(1), \mathbf{Y}_{\text{sp}}(0)) = 1$ as perfect positive rank dependence. In this case, we can compare the highest ranked outcome in $\mathbf{Y}_{\text{sp}}(1)$ to the highest ranked outcome in $\mathbf{Y}_{\text{sp}}(0)$, the second-highest ranked outcome in $\mathbf{Y}_{\text{sp}}(1)$ to the second-highest ranked outcome in $\mathbf{Y}_{\text{sp}}(0)$, and so on. Similarly, we refer to the case in which $\tau_{\text{sp}} = \tau(\mathbf{Y}_{\text{sp}}(1), \mathbf{Y}_{\text{sp}}(0)) = -1$ as perfect negative rank dependence. This case involves comparing the lowest ranked outcome in $\mathbf{Y}_{\text{sp}}(1)$ to the highest ranked outcome in $\mathbf{Y}_{\text{sp}}(0)$, the second-lowest ranked outcome in $\mathbf{Y}_{\text{sp}}(1)$ to the second-highest ranked

outcome in $\mathbf{Y}_{\text{sp}}(0)$, and so on. Heckman et al. (1997) use perfect positive and negative rank dependence to obtain bounds on the distribution of treatment effects. They conclude that these bounds are too wide to provide meaningful insights.

4.2 Rank invariance

Rank invariance refers to the case of perfect positive rank dependence between potential treatment and control outcomes. The rank invariance assumption can be stated as follows.

ASSUMPTION 3 (Rank Invariance): $\tau_{\text{sp}} = 1$.

Under Assumption 3, the identification of our parameters of interest relies on knowledge of the separate marginal distributions of potential treatment and control outcomes. For $v \in (0, 1)$, the QTE under rank invariance as defined by Doksum (1974) and Lehmann (1974) may be written as

$$QTE_v = q_v(F_{Y(1)}(y)) - q_v(F_{Y(0)}(y)).$$

For $u \in (0, 1)$, identification of $q_{\Delta, u}$ under Assumptions 1 and 3 follows from identification of the marginal distributions $F_{Y(1)}(y)$ and $F_{Y(0)}(y)$ under random assignment. Fan and Park (2009) establish that the $q_{\Delta, u}$ are identified if the QTE_v are monotonically increasing in v . If the QTE_v are not monotonically increasing in v , then they have to be rearranged to coincide with the $q_{\Delta, u}$. Formally, if the QTE_v are monotonically increasing in v , then

$$\begin{aligned} q_{\Delta, u} &= q_u(F_{\Delta}(\delta)) & (4) \\ &= q_v(F_{Y(1)}(y)) - q_v(F_{Y(0)}(y)) \\ &= q_v(F_{Y(1)|D=1}(y|D=1)) - q_v(F_{Y(0)|D=0}(y|D=0)) \\ &= q_v(F_{Y|D=1}(y|D=1)) - q_v(F_{Y|D=0}(y|D=0)). \end{aligned}$$

The first equality in (4) follows from Definition 1. The second equality follows from Assumption 3 and from the proof of Lemma 4 in Fan and Park (2009). The third

equality follows from Assumption 1. The fourth equality follows from $Y_i = Y_i(1)D_i + Y_i(0)(1 - D_i)$.

Our parameters of interest are also identified when we replace Assumptions 1 and 3 by Assumptions 2 and 3. Under Assumption 2, the marginal distributions $F_{Y(d)}(y)$, with $d = \{0, 1\}$, are identified through

$$F_{Y(d)}(y) = \int F_{Y(d)|X=x, D=d}(y|X=x, D=d) dF_X(x). \quad (5)$$

Identification of our parameters of interest follows from the identification of $QTE_v = q_v(F_{Y(1)}(y)) - q_v(F_{Y(0)}(y))$. If the QTE_v are monotonically increasing in v , then $q_{\Delta, u} = q_v(F_{Y(1)}(y)) - q_v(F_{Y(0)}(y))$. Similar to (4), if the QTE_v are not monotonically increasing in v , they have to be rearranged to coincide with our parameters of interest.

4.3 Imperfect rank dependence

In this section, we consider identification under Assumption 1 in situations where $0 \leq \tau_{sp} \leq 1$. We ignore situations where $\tau_{sp} < 0$ because they are expected to be uncommon in real-world applications. Extending our approach to these situations is straightforward. Moreover, we focus on the special case in which treatment and control group sizes are equal. Specifically, in the super-population of size N_{sp} , N_{sp}^1 units are randomly assigned to the treatment group and $N_{sp}^0 = N_{sp} - N_{sp}^1$ units are randomly assigned to the control group, with $N_{sp}^1 = N_{sp}^0 = N_{sp}/2$. We discuss more general cases in Section 4.4 below.

4.3.1 Rank correlation coefficient

We define $Y_{id} = Y_i$ if $D_i = d$, $d \in \{0, 1\}$, and consider the $((N_{sp}/2) \times 1)$ -vectors of observable treatment and control outcomes $\mathbf{Y}_{sp, d} = (Y_{1d}, \dots, Y_{(N_{sp}/2)d})'$. Because the ranks of $\mathbf{Y}_{sp, 1}$ relative to the ranks of $\mathbf{Y}_{sp, 0}$ are unknown, we consider all possible permutations of observable treatment outcomes and define the $((N_{sp}/2) \times (N_{sp}/2))$ -permutation matrix $\mathbf{\Pi}_{sp, p}$, $p \in \mathcal{P}_{sp} = \{1, \dots, (N_{sp}/2)!\}$.¹⁰

¹⁰Instead of keeping the ranks of $\mathbf{Y}_{sp, 0}$ fixed and considering permutations of $\mathbf{Y}_{sp, 1}$, an alternative approach involves keeping the ranks of $\mathbf{Y}_{sp, 1}$ fixed and considering permutations of $\mathbf{Y}_{sp, 0}$. We have chosen the former approach because it enables us to determine the average unit-level treatment effect on control group members if they had received the treatment. Moreover, this approach allows us to sort the average unit-level treatment effects by control outcome to obtain the GQTE. See Section 6.2 for details.

Under Assumption 1, the rank correlation coefficient between potential treatment and control outcomes may be written as

$$\tau_{\text{sp}} = \sum_{p \in \mathcal{P}_{\text{sp}}} \Pr[\tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \mathbf{Y}_{\text{sp},0}) = \tau_{\text{sp}}] \tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \mathbf{Y}_{\text{sp},0}). \quad (6)$$

Equation (6) illustrates that knowledge of the probabilities $\Pr[\tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \mathbf{Y}_{\text{sp},0}) = \tau_{\text{sp}}], p \in \mathcal{P}_{\text{sp}}$, would solve the identification problem. To make progress towards this goal, we use the set of covariates to obtain predicted outcomes. We define the $(N_{\text{sp}} \times k)$ -covariate matrix \mathbf{X}_{sp} and the corresponding $((N_{\text{sp}}/2) \times k)$ -covariate matrix $\mathbf{X}_{\text{sp},0}$ for members of the control group. We obtain the rank correlation coefficients between actual and predicted outcomes in three steps:

1. We calculate the predicted outcomes $\widehat{\mathbf{Y}}_{\text{sp}} = \mathbf{X}_{\text{sp}}(\mathbf{X}'_{\text{sp},0}\mathbf{X}_{\text{sp},0})^{-1}\mathbf{X}'_{\text{sp},0}\mathbf{Y}_{\text{sp},0}$.¹¹ We use the elements of the $(N_{\text{sp}} \times 1)$ -vector $\widehat{\mathbf{Y}}_{\text{sp}} = (\widehat{Y}_1, \dots, \widehat{Y}_{N_{\text{sp}}})'$ to define $\widehat{Y}_{id} = \widehat{Y}_i$ if $D_i = d$, $d \in \{0, 1\}$, and consider the $((N_{\text{sp}}/2) \times 1)$ -vectors of predicted treatment and control outcomes $\widehat{\mathbf{Y}}_{\text{sp},d} = (\widehat{Y}_{1d}, \dots, \widehat{Y}_{(N_{\text{sp}}/2)d})'$.
2. We obtain the rank correlation coefficient between actual and predicted outcomes for members of the control group, $\tau(\mathbf{Y}_{\text{sp},0}, \widehat{\mathbf{Y}}_{\text{sp},0})$.
3. We obtain the rank correlation coefficient between actual and predicted outcomes for members of the treatment group, $\tau(\mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},1})$.

In the following, we focus on situations where $|\tau(\mathbf{Y}_{\text{sp},0}, \widehat{\mathbf{Y}}_{\text{sp},0})| \geq |\tau(\mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},1})|$. Focusing on these situations greatly facilitates the presentation of our approach. We ignore situations where $|\tau(\mathbf{Y}_{\text{sp},0}, \widehat{\mathbf{Y}}_{\text{sp},0})| < |\tau(\mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},1})|$ because they are rare in empirical applications. We leave the study of these situations for future work.¹²

¹¹Alternative methods, including machine learning algorithms, can be used to predict the values or ranks of the observed outcomes. We obtain predicted outcomes in a linear regression framework and consider continuous outcome variables because we find that this approach is sufficient to fit the data used in our empirical application. To mitigate potential problems associated with overfitting, we use the same set of covariates as in Bitler et al. (2006).

¹²The experimental study of Gillitzer and Sinning (2020) provides an example in which the out-of-sample prediction of ranks is stronger than the in-sample prediction. The authors investigate the effectiveness of a reminder from the tax office on tax payments. Instead of observing pre-intervention outcomes, they observe the initial amount of tax owed, which is a strong predictor of tax payments of the treatment group (who receive a reminder) but a relatively weak predictor of tax payments of the control group (who do not receive a reminder).

We make the following assumption about the relevance of predictors:

ASSUMPTION 4 (Relevance of predictors): $\tau(\mathbf{Y}_{\text{sp},0}, \widehat{\mathbf{Y}}_{\text{sp},0})$ is positive and sufficiently large.

Assumption 4 requires the availability of relevant control outcome predictors. The set of covariates used in our analysis contains pre-intervention outcomes, which constitute strong predictors of post-intervention control outcomes. We are able to formally test the strength of predictors, as discussed below.

It is useful to consider the rank correlation coefficient $\tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},0})$ associated with the rank correlation coefficient $\tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \mathbf{Y}_{\text{sp},0})$. Under Assumptions 1 and 4, permutations of $\mathbf{Y}_{\text{sp},1}$ must satisfy the condition $\tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},0}) = \tau(\mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},1})$ because the rank correlation coefficient between actual and predicted outcomes must be the same, regardless of whether its calculation is based on predicted treatment or control outcomes. We refer to this property as ‘Rank Correlation Property’ (RCP). Permutations of $\mathbf{Y}_{\text{sp},1}$ that do not satisfy the RCP occur with a probability of zero.¹³

Therefore,

$$\tau_{\text{sp}} = \sum_{p \in \mathcal{S}_{\text{sp}}} \Pr[\tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \mathbf{Y}_{\text{sp},0}) = \tau_{\text{sp}}] \tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \mathbf{Y}_{\text{sp},0}), \quad (7)$$

where $\mathcal{S}_{\text{sp}} = \{p \in \mathcal{P}_{\text{sp}} \mid \tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},0}) = \tau(\mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},1})\}$. Equation (7) follows from equation (6), Assumptions 1 and 4, and from

$$\begin{aligned} & \sum_{p \in \mathcal{P}_{\text{sp}}} \Pr[\tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \mathbf{Y}_{\text{sp},0}) = \tau_{\text{sp}}] \\ &= \sum_{p \in \mathcal{S}_{\text{sp}}} \Pr[\tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \mathbf{Y}_{\text{sp},0}) = \tau_{\text{sp}} \mid \tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},0}) = \tau(\mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},1})] \\ &\times \Pr[\tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},0}) = \tau(\mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},1})] \\ &+ \sum_{p \in \mathcal{S}'_{\text{sp}}} \Pr[\tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \mathbf{Y}_{\text{sp},0}) = \tau_{\text{sp}} \mid \tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},0}) \neq \tau(\mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},1})] \\ &\times \Pr[\tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},0}) \neq \tau(\mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},1})], \end{aligned}$$

¹³It is possible to impose a more stringent version of the RCP by considering permutations that satisfy both $\tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},0}) = \tau(\mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},1})$ and $\tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},0}, \widehat{\mathbf{Y}}_{\text{sp},1}) = \tau(\mathbf{Y}_{\text{sp},0}, \widehat{\mathbf{Y}}_{\text{sp},0})$. While implementing this stricter version increases complexity, sensitivity checks indicate that it yields very similar results. We leave the use of the stricter RCP for future research.

where $\mathcal{S}'_{\text{sp}} = \{p \in \mathcal{P}_{\text{sp}} \mid \tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},0}) \neq \tau(\mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},1})\}$. Permutations that do not satisfy the RCP occur with a probability of zero due to random assignment under Assumption 1, $\Pr[\tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},0}) \neq \tau(\mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},1})] = 0$ for all $p \in \mathcal{P}_{\text{sp}}$. Using Bayes' law, $\Pr[\tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \mathbf{Y}_{\text{sp},0}) = \tau_{\text{sp}}] = \Pr[\tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \mathbf{Y}_{\text{sp},0}) = \tau_{\text{sp}} \mid \tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},0}) = \tau(\mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},1})] \Pr[\tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},0}) = \tau(\mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},1})]$ for all $p \in \mathcal{S}_{\text{sp}}$.

We make the following assumption about permutations with a positive probability of occurrence:

ASSUMPTION 5 (Equally likely permutations): *All permutations of treatment group outcomes with a positive probability of occurrence are equally likely to occur.*

Assumption 5 implies that

$$\Pr[\tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \mathbf{Y}_{\text{sp},0}) = \tau_{\text{sp}}] = \frac{\mathbf{1}\{\Pr[\tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \mathbf{Y}_{\text{sp},0}) = \tau_{\text{sp}}] > 0\}}{\sum_{p \in \mathcal{P}_{\text{sp}}} \mathbf{1}\{\Pr[\tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \mathbf{Y}_{\text{sp},0}) = \tau_{\text{sp}}] > 0\}}$$

for all $p \in \mathcal{P}_{\text{sp}}$. In the absence of further information, all permutations of $\mathbf{Y}_{\text{sp},1}$ that satisfy the RCP have a positive probability of occurrence under Assumptions 1 and 4. Therefore, under Assumptions 1, 4, and 5, the rank correlation coefficient between potential treatment and control outcomes is identified through

$$\tau_{\text{sp}} = \frac{1}{n_{\text{sp},p}} \sum_{p \in \mathcal{S}_{\text{sp}}} \tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \mathbf{Y}_{\text{sp},0}), \quad (8)$$

with $n_{\text{sp},p} = \sum_{p \in \mathcal{P}_{\text{sp}}} \mathbf{1}\{\Pr[\tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \mathbf{Y}_{\text{sp},0}) = \tau_{\text{sp}}] > 0\} = \sum_{p \in \mathcal{S}_{\text{sp}}} \mathbf{1}\{\tau(\mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},0}) = \tau(\mathbf{Y}_{\text{sp},1}, \widehat{\mathbf{Y}}_{\text{sp},1})\}$.

4.3.2 Quantiles of the distribution of treatment effects

In this section, we apply the same reasoning as in the previous section to identify our parameters of interest under Assumptions 1, 4 and 5. We define $\Delta_{\text{sp},p} = \mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1} - \mathbf{Y}_{\text{sp},0}$. Under Assumption 1, the distribution of treatment effects may be written as

$$F_{\Delta}(\delta) = F_{\Delta'}(\delta'), \quad (9)$$

where

$$\Delta' = \sum_{p \in \mathcal{P}_{\text{sp}}} \Pr[F_{\Delta_{\text{sp},p}}(\delta_{\text{sp},p}) = F_{\Delta}(\delta)] \Delta_{\text{sp},p}.$$

Knowledge of the probabilities $\Pr[F_{\Delta_{\text{sp},p}}(\delta_{\text{sp},p}) = F_{\Delta}(\delta)]$, $p \in \mathcal{P}_{\text{sp}}$, would solve the identification problem. Permutations of $\mathbf{Y}_{\text{sp},1}$ that do not satisfy the RCP occur with a probability of zero under Assumption 1. Therefore,

$$\Delta' = \sum_{p \in \mathcal{S}_{\text{sp}}} \Pr[F_{\Delta_{\text{sp},p}}(\delta_{\text{sp},p}) = F_{\Delta}(\delta)] \Delta_{\text{sp},p}.$$

In the absence of further information, all permutations of $\mathbf{Y}_{\text{sp},1}$ that satisfy the RCP have a positive probability of occurrence under Assumptions 1 and 4. Therefore, under Assumptions 1, 4, and 5, the distribution of treatment effects is identified through

$$F_{\Delta}(\delta) = F_{\Delta'}(\delta'), \quad (10)$$

where $\Delta' = \frac{1}{n_{\text{sp},p}} \sum_{p \in \mathcal{S}_{\text{sp}}} \Delta_{\text{sp},p}$. Identification of our parameters of interest follows from

$$q_{\Delta,u} = q_u(F_{\Delta}(\delta)) = q_u(F_{\Delta'}(\delta')).$$

It is important to note that $\Delta' = \frac{1}{n_{\text{sp},p}} \sum_{p \in \mathcal{S}_{\text{sp}}} \mathbf{\Pi}_{\text{sp},p} \mathbf{Y}_{\text{sp},1} - \mathbf{Y}_{\text{sp},0}$ involves averaging over permutations of $\mathbf{Y}_{\text{sp},1}$ while keeping the ranks of $\mathbf{Y}_{\text{sp},0}$ fixed. Computing quantiles of the distribution of permutation-specific treatment effects before averaging over permutations would produce a different result. Under Assumptions 1, 4, and 5, we may write

$$\begin{aligned} q_{\Delta,u}^{QID} &= \sum_{p \in \mathcal{P}_{\text{sp}}} \Pr[q_u(F_{\Delta_{\text{sp},p}}(\delta_{\text{sp},p})) = q_u(F_{\Delta}(\delta))] q_u(F_{\Delta_{\text{sp},p}}(\delta_{\text{sp},p})) \\ &= \frac{1}{n_{\text{sp},p}} \sum_{p \in \mathcal{S}_{\text{sp}}} q_u(F_{\Delta_{\text{sp},p}}(\delta_{\text{sp},p})). \end{aligned} \quad (11)$$

The parameters specified by (11) bear resemblance to those of Heckman et al. (1997) who estimate quantiles and other parameters of the ‘impact distribution’ based on assumed values of τ_{sp} . Averaging quantiles of permutation-specific distributions of treatment effects ignores that the location of control outcomes changes each time the quan-

tiles are computed. Consequently, this approach does not yield QDTE. We provide a detailed discussion of this issue in Section 6.2.

4.4 Controlling for covariates

In this section, we extend our approach in two important ways. Firstly, we consider identification of our parameters of interest under strong ignorability. Secondly, we proof identification of our parameters of interest for both equal and unequal group sizes.

4.4.1 Equal group sizes

We first consider the case in which $N_{\text{sp}}^1 = N_{\text{sp}}^0$ to identify our parameters of interest under Assumptions 2, 4, and 5. Using results presented in Firpo (2007), we may write the marginal distributions $F_{Y^{(d)}}(y)$, $d = \{0, 1\}$, in terms of inverse propensity score weighted averages,

$$F_{Y^{(d)}}(y) = \frac{1}{N_{\text{sp}}} \sum_{i=1}^{N_{\text{sp}}} w_{id} \mathbf{1}\{Y_i \leq y\}, \quad (12)$$

where $w_{i1} = \frac{D_i}{Pr(D_i=1|X_i)}$ and $w_{i0} = \frac{1-D_i}{1-Pr(D_i=1|X_i)}$. Following Parzen (1979), we consider the order statistics $Z_{(1)d} \leq \dots \leq Z_{(N_{\text{sp}}^d)d}$, where

$$Z_{(i)d} = F_{Y^{(d)}}^{-1}[\theta_d] = \inf\{y : \frac{1}{N_{\text{sp}}} \sum_{i=1}^{N_{\text{sp}}} w_{id} \mathbf{1}\{Y_i \leq y\} \geq \theta_d\}, \quad (13)$$

with $(i-1)/N_{\text{sp}}^d < \theta_d \leq i/N_{\text{sp}}^d$, $i \in \{1, \dots, N_{\text{sp}}^d\}$. We define the $(N_{\text{sp}}^d \times 1)$ -vectors $\mathbf{Z}_{\text{sp},d} = (Z_{1d}, \dots, Z_{N_{\text{sp}}^d d})'$. Reweighting can impact the within-group ranks of both observed and predicted outcomes. The identification of our parameters of interest under Assumptions 2, 4, and 5 relies on using reweighted outcomes and adopting the approach outlined in Section 4.3.2. Theorem 1 formally states this result.

THEOREM 1 (Identification of $q_{\Delta,u}$ with Equal Group Sizes): *Let $N_{\text{sp}}^1 = N_{\text{sp}}^0$. Then, under Assumptions 2, 4, and 5, the Population Quantiles of the Distribution of Treatment Effects, $q_{\Delta,u}$, are identified from data on (Y, D, X) .*

All proofs are provided in Appendix C. Theorem 1 states that our parameters of interest are identified from data on (Y, D, X) under Assumptions 2, 4, and 5 if $N_{\text{sp}}^1 = N_{\text{sp}}^0$.

4.4.2 Unequal group sizes

Since Assumption 2 is applicable to N_{sp}^d order statistics, as demonstrated in the previous section, it also extends to multiples of N_{sp}^d order statistics. We leverage this property to describe the quantile functions of treatment and control outcomes using $N_{\text{sp}}^1 N_{\text{sp}}^0$ values.¹⁴ Formally, we consider the order statistics $Z_{(1)d} \leq \dots \leq Z_{(N_{\text{sp}}^1 N_{\text{sp}}^0)d}$, where $Z_{(i')d} = \inf\{y : \frac{1}{N_{\text{sp}}} \sum_{i=1}^{N_{\text{sp}}} w_{id} \mathbf{1}\{Y_i \leq y\} \geq \theta_d\}$ with $(i' - 1)/(N_{\text{sp}}^1 N_{\text{sp}}^0) < \theta_d \leq i'/(N_{\text{sp}}^1 N_{\text{sp}}^0)$, $i' \in \{1, \dots, N_{\text{sp}}^1 N_{\text{sp}}^0\}$. Moreover, we define the $(N_{\text{sp}}^1 N_{\text{sp}}^0 \times n_d)$ -transformation matrix $\mathbf{M}_{\text{sp},d}$ that transforms the $(N_{\text{sp}}^d \times 1)$ -vectors $\mathbf{Z}_{\text{sp},d}$ into the $(N_{\text{sp}}^1 N_{\text{sp}}^0 \times 1)$ -vectors

$$\mathbf{M}_{\text{sp},d} \mathbf{Z}_{\text{sp},d} = \left(\underbrace{Z_{1d}, \dots, Z_{1d}}_{N_{\text{sp}}^{1-d}}, \dots, \underbrace{Z_{N_{\text{sp}}^d d}, \dots, Z_{N_{\text{sp}}^d d}}_{N_{\text{sp}}^{1-d}} \right)'.$$

Considering $N_{\text{sp}}^1 N_{\text{sp}}^0$ values of the order statistics allows us to obtain a generalized version of Theorem 1, which involves using the transformed outcome variables $\mathbf{M}_{\text{sp},d} \mathbf{Z}_{\text{sp},d}$ instead of $\mathbf{Z}_{\text{sp},d}$. While using $N_{\text{sp}}^1 N_{\text{sp}}^0$ order statistics addresses the problem of unequal group sizes, the challenges resulting from uncertainty about the ranking of treatment outcomes relative to control outcomes remain unchanged. The identification of our parameters of interest continues to be based on permutations of $\mathbf{Z}_{\text{sp},1}$. Theorem 2 states our main result.

THEOREM 2 (Identification of $q_{\Delta,u}$): *Under Assumptions 2, 4, and 5, the Population Quantiles of the Distribution of Treatment Effects, $q_{\Delta,u}$, are identified from data on (Y, D, X) .*

Theorem 2 states that our parameters of interest are identified from data on (Y, D, X) under Assumptions 2, 4 and 5. Because Theorem 2 holds for any $N_{\text{sp}}^1, N_{\text{sp}}^0 \in \mathbb{N}_+$, it includes Theorem 1 as a special case. In the case of rank invariance, our parameters of

¹⁴Applying this approach to large samples may pose computational challenges. However, the use of frequency weights can substantially alleviate this issue by reducing the necessary number of rows in a transformed dataset.

interest simplify to $q_{\Delta,u} = q_v(F_{Z_1}(z_1)) - q_v(F_{Z_0}(z_0))$, where $Z_1 = \mathbf{M}_{\text{sp},1} \mathbf{\Pi}_{\text{sp},p} \mathbf{Z}_{\text{sp},1}$ and $Z_0 = \mathbf{M}_{\text{sp},0} \mathbf{Z}_{\text{sp},0}$.¹⁵ As noted earlier, the differences between quantiles have to be monotonically increasing in v to obtain this result.

5 Estimation

Using the framework introduced by Imbens and Rubin (2015), we consider a finite Simple Random Sample (SRS) of fixed size N drawn from the super-population and define R_i as a sampling indicator, which takes on the value 1 if unit i is sampled, and 0 otherwise. We define the treatment indicator W_i , which takes on the value 1 if unit i is sampled and assigned to the treatment group, and 0 if unit i is assigned to the control group. We simplify the presentation by assigning all unsampled units to the control group, implying $W_i = 0$ if $R_i = 0$.

We observe N_1 units randomly assigned to the treatment group and $N_0 = N - N_1$ units randomly assigned to the control group. We define the $(N \times 1)$ -vectors of potential outcomes in the finite sample, $\mathbf{Y}(d) = (Y_1(d), \dots, Y_N(d))'$, $d = \{0, 1\}$. We also define the $(N_d \times 1)$ -vectors of observed treatment and control outcomes, $\mathbf{Y}_d = (Y_{1d}, \dots, Y_{N_d d})'$. We consider the $(N \times k)$ -covariate matrix \mathbf{X} and the corresponding $(N_0 \times k)$ -covariate matrix \mathbf{X}_0 of the control group sample. Using the control group sample, we estimate a linear regression model to obtain the parameter vector $\widehat{\gamma}_0 = (\mathbf{X}'_0 \mathbf{X}_0)^{-1} \mathbf{X}'_0 \mathbf{Y}_0$. The vector of predicted values is given by $\widehat{\mathbf{Y}} = \mathbf{X} \widehat{\gamma}_0$. Using the elements of $\widehat{\mathbf{Y}}$, we define $\widehat{Y}_{id} = \widehat{Y}_i$ if $W_i = d$ and consider the $(N_d \times 1)$ -vectors of predicted values $\widehat{\mathbf{Y}}_d = (\widehat{Y}_{id}, \dots, \widehat{Y}_{N_d d})$. Moreover, we define the $(N_1 \times N_1)$ - permutation matrix $\mathbf{\Pi}_{p, p} \in \mathcal{P} = \{1, \dots, N_1!\}$.

Given the sample Y_1, \dots, Y_N , the empirical distribution functions of $F_{Y(d)}(y)$ may be written as $\widehat{F}_{Y(d)}(y) = \frac{1}{N} \sum_{i=1}^N w_{id} \mathbf{1}\{Y_i \leq y\}$, where $w_{i1} = \frac{W_i}{\Pr(W_i=1|X_i)}$ and $w_{i0} = \frac{1-W_i}{1-\Pr(W_i=1|X_i)}$. For

$$i' \in \{\{j_{11}, \dots, j_{1N_{(1-d)}}\}, \dots, \{j_{N_d 1}, \dots, j_{N_d N_{(1-d)}}\}\} = \{1, \dots, N_1 N_0\},$$

we define the order statistics $Z_{(1)d} \leq \dots \leq Z_{(N_1 N_0)d}$ as values of the empirical quantile function of continuous and monotonically increasing empirical distribution functions

¹⁵Figure A3d demonstrates that our approach produces very similar results to the approach of Firpo (2007) under rank invariance.

$\widehat{F}_{Y(d)}(y)$,

$$Z_{(i')d} = \widehat{F}_{Y(d)}^{-1}[\theta_d] = \inf\left\{y : \frac{1}{N} \sum_{i=1}^N w_{id} \mathbf{1}\{Y_i \leq y\} \geq \theta_d\right\},$$

with $(i' - 1)/(N_0N_1) < \theta_d \leq i'/(N_0N_1)$, $i' \in \{1, \dots, N_0N_1\}$. We define the $(N_1N_0 \times 1)$ -vectors

$$\begin{aligned} \mathbf{V}_d &= (Z_{(j_{11})d}, \dots, Z_{(j_{1N(1-d)})d}, \dots, Z_{(j_{N_d1})d}, \dots, Z_{(j_{N_dN(1-d)})d})' \\ &= \underbrace{(Z_{(1)d}, \dots, Z_{(1)d})}_{N(1-d)}, \dots, \underbrace{(Z_{(N_d)d}, \dots, Z_{(N_d)d})}_{N(1-d)}', \end{aligned}$$

where

$$Z_{(i)d} = \widehat{F}_{Y(d)}^{-1}[\theta_d] = \inf\left\{y : \frac{1}{N} \sum_{i=1}^N w_{id} \mathbf{1}\{Y_i \leq y\} \geq \theta_d\right\},$$

with $(i - 1)/N_d < \theta_d \leq i/N_d$, $i \in \{1, \dots, N_d\}$. We define the $(N_d \times 1)$ -vectors $\mathbf{Z}_d = (Z_{1d}, \dots, Z_{N_d d})'$. Similarly, we define

$$\widehat{Z}_{(i)d} = \inf\left\{\widehat{y} : \frac{1}{N} \sum_{i=1}^N w_{id} \mathbf{1}\{\widehat{Y}_i \leq \widehat{y}\} \geq \theta_d\right\}$$

and the $(N_d \times 1)$ -vectors $\widehat{\mathbf{Z}}_d = (\widehat{Z}_{1d}, \dots, \widehat{Z}_{N_d d})'$. We consider the permutations $\mathbf{\Omega}_d \mathbf{V}_d = \mathbf{M}_d \mathbf{Z}_d$, where $\mathbf{\Omega}_d$ are $(N_1N_0 \times N_1N_0)$ -permutation matrices that transform the $(N_d \times 1)$ -vectors \mathbf{Z}_d into the vectors $(N_1N_0 \times 1)$ -vectors

$$\mathbf{M}_d \mathbf{Z}_d = \underbrace{(Z_{1d}, \dots, Z_{1d})}_{N(1-d)}, \dots, \underbrace{(Z_{N_d d}, \dots, Z_{N_d d})}_{N(1-d)}'.$$

5.1 Rank correlation coefficient

Our estimator of the rank correlation coefficient between potential treatment and control outcomes τ_{sp} is given by

$$\widehat{\tau} = \frac{1}{n_p} \sum_{p \in \mathcal{S}} \tau(\mathbf{M}_1 \mathbf{\Pi}_p \mathbf{Z}_1, \mathbf{M}_0 \mathbf{Z}_0),$$

where $\mathcal{S} = \{p \in \mathcal{P} | \tau(\mathbf{M}_1 \mathbf{\Pi}_p \mathbf{Z}_1, \mathbf{M}_0 \widehat{\mathbf{Z}}_0) = \tau(\mathbf{Z}_1, \widehat{\mathbf{Z}}_1)\}$ and $n_p = \sum_{p \in \mathcal{S}} \mathbf{1}\{\tau(\mathbf{M}_1 \mathbf{\Pi}_p \mathbf{Z}_1, \mathbf{M}_0 \widehat{\mathbf{Z}}_0) = \tau(\mathbf{Z}_1, \widehat{\mathbf{Z}}_1)\}$.

Under Assumption 1, our estimator simplifies to $\widehat{\tau} = \frac{1}{n_p} \sum_{p \in \mathcal{S}} \tau(\mathbf{M}_1 \mathbf{\Pi}_p \mathbf{Y}_1, \mathbf{M}_0 \mathbf{Y}_0)$. When treatment and control group sample sizes are equal, $N_1 = N_0 = N/2$, it further simplifies to $\widehat{\tau} = \frac{1}{n_p} \sum_{p \in \mathcal{S}} \tau(\mathbf{\Pi}_p \mathbf{Y}_1, \mathbf{Y}_0)$.

Due to the large number of possible permutations, we use the following steps to estimate τ_{sp} in our empirical application:

1. We search for a random permutation of \mathbf{Z}_1 that satisfies $\tau(\mathbf{M}_1 \mathbf{\Pi}_p \mathbf{Z}_1, \mathbf{M}_0 \widehat{\mathbf{Z}}_0) = \tau(\mathbf{Z}_1, \widehat{\mathbf{Z}}_1)$ and calculate $\tau(\mathbf{\Pi}_p \mathbf{Z}_1, \mathbf{Z}_0)$.
2. We repeat the previous step P times.
3. We calculate $\widehat{\tau}_P = \frac{1}{P} \sum_{p=1}^P \tau(\mathbf{M}_1 \mathbf{\Pi}_p \mathbf{Z}_1, \mathbf{M}_0 \mathbf{Z}_0)$.

Table 1 presents estimates of τ_{sp} for alternative values of $\tau(\mathbf{Y}_0, \widehat{\mathbf{Y}}_0)$ and $\tau(\mathbf{Y}_1, \widehat{\mathbf{Y}}_1)$ based on simulation data. Each estimate involves 100 random permutations of $2 \times 16,772$ observation units ($N_1 = N_0 = N/2 = 16,772$), consistent with the restricted analysis sample using in our application below. For simplicity, we consider a model with a single continuous covariate and exclude the possibility of ties. We obtain bootstrap standard errors for the estimated parameters by drawing a random sample with replacement, stratified by treatment group to ensure that $N_1 = N_0$, for each permutation.¹⁶

The results in Table 1 illustrate that it is possible to obtain precise estimates of τ_{sp} , even if the strength of the model predictors, measured by $\tau(\mathbf{Y}_0, \widehat{\mathbf{Y}}_0)$, is only moderate. For $\tau(\mathbf{Y}_0, \widehat{\mathbf{Y}}_0) \geq 0.3$, our estimates are identical to the target values if we round to one decimal place. For $\tau(\mathbf{Y}_0, \widehat{\mathbf{Y}}_0) = 0.2$, we observe reasonably precise estimates of τ_{sp} if $\tau_{\text{sp}} < 0.9$. However, we obtain an estimate of 0.83 if $\tau_{\text{sp}} = 0.9$, indicating that it is more challenging to estimate τ_{sp} when the ranks of the potential treatment and control outcomes are highly correlated.¹⁷

<Table 1 about here.>

¹⁶The estimation of bootstrap standard errors involves the use of ties due to random sampling with replacement.

¹⁷The reason for this result is that the underlying values of $\tau(\mathbf{Y}_1, \widehat{\mathbf{Y}}_1)$, which are presented in Table A1 of Appendix A, are very close to each other when τ_{sp} is close to 1 and $\tau(\mathbf{Y}_0, \widehat{\mathbf{Y}}_0)$ is small.

The use of simulated data to estimate target values of τ_{sp} for given values of $\tau(\mathbf{Y}_0, \widehat{\mathbf{Y}}_0)$ allows us to develop a statistical test for assessing the strength of model predictors, measured by $\tau(\mathbf{Y}_0, \widehat{\mathbf{Y}}_0)$. We test whether our estimated rank correlation coefficients deviate significantly from these target values. We also employ a Kolmogorov-Smirnov test to ascertain whether deviations in estimated rank correlation coefficients from target values impact our QDTE estimates. Our test results indicate that our predictors are sufficiently strong to yield unbiased QDTE estimates. A more detailed discussion of our test results is provided in Appendix A.

Results of Monte Carlo simulations (presented in Tables B1-B4 of Appendix B) are in line with the findings presented in Table 1. In particular, we find that our estimator of τ_{sp} is fairly precise even for relatively small sample sizes and unbiased if $\tau(\mathbf{Y}_0, \widehat{\mathbf{Y}}_0)$ is sufficiently large. Biases observed for low values of $\tau(\mathbf{Y}_0, \widehat{\mathbf{Y}}_0)$ and large values of τ_{sp} decline as the sample size increases. Our simulation results also confirm the validity of estimated bootstrap standard errors.

5.2 Quantiles of the Distribution of Treatment Effects

For $u \in (0, 1)$, the PQDTE estimator is given by

$$\widehat{q}_{\Delta, u} = q_u(F_{\Delta^{\text{obs}}}(\delta^{\text{obs}})),$$

where $\Delta^{\text{obs}} = \frac{1}{n_p} \sum_{p \in \mathcal{S}} \mathbf{M}_1 \mathbf{\Pi}_p \mathbf{Z}_1 - \mathbf{M}_0 \mathbf{Z}_0$. We impose Assumption 1 and consider the case in which $N_1 = N_0 = N/2$ to show that $\widehat{q}_{\Delta, u}$ is an unbiased estimator of $q_{\Delta, u}$. In this case, our estimator may be written as

$$\widehat{q}_{\Delta, u} = \inf\{\delta^{\text{obs}} : \frac{1}{N_1} \sum_{i=1}^{N_{\text{sp}}} R_i (1 - W_i) \mathbf{1}\{\Delta_i^{\text{obs}} \leq \delta^{\text{obs}}\} \geq u\},$$

where Δ_i^{obs} is the i th element of $\Delta^{\text{obs}} = \frac{1}{n_p} \sum_{p \in \mathcal{S}} \mathbf{\Pi}_p \mathbf{Y}_1 - \mathbf{Y}_0$. Following Imbens and Rubin (2015), we can take the expectation of our estimator over the randomization distribution

by conditioning on the $(N_{\text{sp}} \times 1)$ -vector $\mathbf{R} = (R_1, \dots, R_{N_{\text{sp}}})'$, and on $\mathbf{Y}_{\text{sp}}(1)$ and $\mathbf{Y}_{\text{sp}}(0)$:

$$\begin{aligned}
\mathbb{E}_W[\widehat{q}_{\Delta,u} | \mathbf{R}, \mathbf{Y}_{\text{sp}}(1), \mathbf{Y}_{\text{sp}}(0)] &= \inf\{\delta^{\text{obs}} : \frac{1}{N_0} \sum_{i=1}^{N_{\text{sp}}} R_i \mathbb{E}_W[1 - W_i] \mathbf{1}\{\Delta_i^{\text{obs}} \leq \delta^{\text{obs}}\} \geq u\} \\
&= \inf\{\delta^{\text{obs}} : \frac{1}{N_0} \sum_{i=1}^{N_{\text{sp}}} R_i \frac{N_0}{N} \mathbf{1}\{\Delta_i^{\text{obs}} \leq \delta^{\text{obs}}\} \geq u\} \quad (14) \\
&= \inf\{\delta : \frac{1}{N} \sum_{i=1}^{N_{\text{sp}}} R_i \mathbf{1}\{\Delta_i \leq \delta\} \geq u\} \\
&= q_{\Delta,u}^{\text{fs}},
\end{aligned}$$

where $q_{\Delta,u}^{\text{fs}}$ are the finite-population quantiles of the distribution of treatment effects.

Taking expectations over the distribution generated by the random sampling yields

$$\begin{aligned}
\mathbb{E}_{\text{sp}}[q_{\Delta,u}^{\text{fs}} | \mathbf{Y}_{\text{sp}}(1), \mathbf{Y}_{\text{sp}}(0)] &= \inf\{\delta : \frac{1}{N} \sum_{i=1}^{N_{\text{sp}}} \mathbb{E}_{\text{sp}}[R_i] \mathbf{1}\{\Delta_i \leq \delta\} \geq u\} \\
&= \inf\{\delta : \frac{1}{N} \sum_{i=1}^{N_{\text{sp}}} \frac{N}{N_{\text{sp}}} \mathbf{1}\{\Delta_i \leq \delta\} \geq u\} \quad (15) \\
&= \inf\{\delta : \frac{1}{N_{\text{sp}}} \sum_{i=1}^{N_{\text{sp}}} \mathbf{1}\{\Delta_i \leq \delta\} \geq u\} \\
&= q_{\Delta,u}.
\end{aligned}$$

We may use a similar approach to study the QID estimator

$$\widehat{q}_{\Delta,u}^{\text{QID}} = \frac{1}{n_p} \sum_{p \in \mathcal{S}} \widehat{q}_{\Delta_p,u}^{\text{QID}} = \frac{1}{n_p} \sum_{p \in \mathcal{S}} q_u(F_{\Delta_p^{\text{obs}}}(\delta_p^{\text{obs}})), \quad (16)$$

where $\Delta_p^{\text{obs}} = \mathbf{\Pi}_p \mathbf{Y}_1 - \mathbf{Y}_0$. It is possible to take expectations over the randomization distribution by conditioning on \mathbf{R} , $\mathbf{Y}_{\text{sp}}(1)$, and $\mathbf{Y}_{\text{sp}}(0)$ to show that

$$\mathbb{E}_W[\widehat{q}_{\Delta_p,u}^{\text{QID}} | \mathbf{R}, \mathbf{Y}_{\text{sp}}(1), \mathbf{Y}_{\text{sp}}(0)] = q_{\Delta_p,u}^{\text{fs}},$$

where $q_{\Delta_p,u}^{\text{fs}}$ are the permutation-specific finite-population quantiles of the distribution of treatment effects. Taking expectations over the distribution generated by the random sampling yields

$$\mathbb{E}_{\text{sp}}[q_{\Delta_p,u}^{\text{fs}} | \mathbf{Y}_{\text{sp}}(1), \mathbf{Y}_{\text{sp}}(0)] = q_{\Delta_p,u},$$

where $q_{\Delta_p, u}$ are the permutation-specific quantiles of the distribution of treatment effects in the super-population. While the permutation-specific elements of the QID estimator are unbiased, averaging over these elements does not yield $q_{\Delta, u}$. We illustrate this issue in our empirical application. Due to the large number of possible permutations, we use Steps 1 and 2 presented in the previous section to obtain QDTE and QID estimates.

6 Application

6.1 Background and data

Connecticut introduced the Jobs First program in response to the nationwide mandate for welfare reform initiated by the Personal Responsibility and Work Opportunity Reconciliation Act in 1996. To assess the program’s effectiveness, the state conducted a randomized experiment in collaboration with the Manpower Development Research Corporation (MDRC). The experiment compared the Jobs First program with the previous Aid to Families with Dependent Children (AFDC) welfare program for low-income single parents with children. Jobs First incorporates key elements commonly found in US welfare programs, including earnings disregards, work requirements, and financial sanctions for non-compliance with these requirements. The program introduced a combination of both positive and negative work incentives. Detailed descriptions of the key elements of Jobs First and AFDC can be found in Bitler et al. (2006) and Kline and Tartari (2016).

Our analysis relies on public use files from MDRC’s experimental evaluation of Connecticut’s Jobs First waiver from AFDC rules. The files contain rounded data on quarterly earnings, as well as monthly welfare and food stamps income, covering a period of two years leading up to program assignment and extending for a minimum of four years thereafter. Following Bitler et al. (2006), our analysis focuses on three outcome variables: earnings, transfers (cash welfare plus food stamps), and income (the sum of earnings and transfers). MDRC’s dataset contains a sample of 4,803 cases, with 2,396 assigned to Jobs First and 2,407 to AFDC. To obtain an analysis sample with equal treatment and control group sizes, we randomly exclude 11 observations from the AFDC sample to create a restricted AFDC sample of 2,396 observations. This sample

restriction enables us to illustrate our approach for equal group sizes. Subsequently, we expand our approach to accommodate scenarios in which treatment and control group sizes differ.

Table 2 replicates the summary statistics of pre-treatment characteristics presented in Bitler et al. (2006) and provides additional results for the restricted AFDC sample. We observe that the exclusion of 11 randomly chosen observations from the AFDC sample has no significant impact on the sample means, with all p -values associated with the comparison of group means of unrestricted and restricted AFDC samples exceeding 0.9.

<Table 2 about here.>

Because our approach relies on predicting values or ranks of observed outcomes, we employ the same set of covariates as used in Bitler et al. (2006) to mitigate potential problems associated with overfitting. A complete list of covariates can be found in Appendix B of Bitler et al. (2006). We use the same set of covariates to estimate our parameters of interest under strong ignorability.

6.2 QID, QDTE and GQTE

Heckman et al. (1997) pair percentiles of marginal distributions of treatment and control outcomes to calculate the mean of parameters of the impact distribution. They use random permutations of percentiles to achieve a given value of τ_{sp} .¹⁸ We modify their approach in two ways that have minimal impact on our results. Firstly, we use restricted analysis samples with equal treatment and control group sizes. This change enables us to consider random permutations of unit-level outcomes instead of relying on percentiles of marginal distributions of treatment and control outcomes. This modification produces slightly more accurate results. The impact of imposing sample restrictions is negligible in our application, as demonstrated below (see Figure A3 in Appendix A).

Secondly, instead of assuming a specific value of τ_{sp} , we focus on random permutations that satisfy the RCP. Figures 1a and 1b depict the distributions of rank correlation coefficients resulting from 100 random permutations of \mathbf{Y}_1 that satisfy the RCP. We observe that the rank correlation distributions are relatively narrow, suggesting that QID

¹⁸See the discussion related to Table 5A in Heckman et al. (1997) for details.

estimates based on the RCP will closely resemble QID estimates based on a fixed value of τ_{sp} that coincides with the support of the corresponding rank correlation distribution.

Given these modifications, our QID estimator may be written as

$$QID = \frac{1}{P} \sum_{p=1}^P q_u(F_{\Delta_p^{\text{obs}}}(\delta_p^{\text{obs}})), \quad (17)$$

where $\Delta_p^{\text{obs}} = \mathbf{\Pi}_p \mathbf{Y}_1 - \mathbf{Y}_0$. Figures 1c-1h present QID estimates for earnings, transfers, and income in pre- and post-time limit periods based on our restricted analysis samples. The figures confirm that the discrepancy between the percentile approach, which relies on permutations consistent with a fixed value $\hat{\tau}$, and the use of unit-level treatment effects in conjunction with the RCP is negligible.

<Figure 1 about here.>

Even though our data is entirely different from that used by Heckman et al. (1997), we observe similar overall patterns in our findings. Specifically, we notice substantially larger extremes (both positive and negative) in the tails of the impact distributions for lower values of $\hat{\tau}$. For example, in the case of earnings during the pre-time limit period (Figure 1c) based on $\hat{\tau} = 0.71$, the values of the impact distribution range from approximately $-\$1,300$ to $\$1,400$. In the case of earnings during the post-time limit period (Figure 1d) based on $\hat{\tau} = 0.62$, the corresponding values of the impact distribution range from approximately $-\$3,100$ to $\$3,100$. We observe similar patterns for transfers (Figures 1e-1f) and income (Figures 1g-1h).¹⁹

Heckman et al. (1997) conclude that impact distributions require a high degree of positive rank dependence because low values of τ_{sp} produce negative values of the impact distribution that are too large to be plausible.²⁰ In the following, we offer an alternative explanation. Using our restricted analysis sample for the pre-time limit period, we obtain permutation-specific QID estimates for selected random permutations of income. We present three permutation-specific impact distributions in Figures 2a-2c.

¹⁹We find that the rank correlation coefficients during the post-time limit period are lower than those observed during the pre-time limit period due to greater mobility with time after the intervention.

²⁰Using data on earnings of women aged 22 and above from an experimental evaluation of training programmes financed under Title IIA of the US Job Training Partnership Act (JTPA), they argue that their results are only reasonable for $\tau_{sp} \geq 0.8$.

We also select a control group member whose income matches the median control group income and pinpoint their location on each permutation-specific impact distribution.

We observe that the selected control group member is located near the lower end of the first permutation-specific impact distribution (Figure 2a), near the median of the second distribution (Figure 2b), and near the upper end of the third distribution (Figure 2c). This finding illustrates that our parameters of interest cannot be obtained by averaging over quantiles of permutation-specific impact distributions because this approach fails to capture the underlying unit-level treatment effect we would expect to observe for any given member of the control group.

<Figure 2 about here.>

Figure 2d presents both average unit-level treatment effects and QDTE estimates. Our findings indicate that the Jobs First program led to income increases across the entire income distribution during the pre-time limit period. This finding is broadly consistent with the results of Bitler et al. (2006), who report non-negative QTE estimates across the income distribution during the pre-time limit period. However, while Bitler et al. (2006) find zero QTE estimates for the bottom 10 quantiles due to zero income in both groups, our QDTE estimates are strictly positive. This difference arises because we compare each control group member to a set of treatment group members associated with random permutations that satisfy the RCP. In the absence of rank invariance, we compare control group members with zero income to treatment group members whose income is either zero or positive. Consequently, averaging over the distribution of permutation-specific treatment effects produces positive QDTE estimates.²¹

Rearranging the average unit-level treatment effects shown in Figure 2d yields GQTE estimates, which are presented in Figure 2e. GQTE estimates are obtained by averaging over average unit-level treatment effects within quantile groups of the control outcome. This approach provides us with a generalized summary measure of quantile treatment effects that does not rely on a rank invariance assumption. The GQTE provides an estimate of the treatment effect that would have been observed if control group members at a given point of the control outcome distribution, such as the

²¹An important challenge in estimating QDTE in real-world applications is the handling of ties. We address this challenge by maintaining a random but fixed ranking of ties in control group outcomes, while allowing for varying ranks of treatment group outcomes through alternative permutations.

median, had received the treatment. Calculating the average within quantile groups is essential to obtain a measure that captures the heterogeneity around the ATE.²²

Another perspective on GQTE estimates is to view them as averages of permutation-specific GQTE estimates. Permutation-specific GQTE estimates are obtained by averaging unit-level treatment effects within quantile groups of the control outcome. They differ from permutation-specific QID estimates in terms of how the unit-level treatment effects are aggregated. While permutation-specific QIDs capture the full extent of treatment effect heterogeneity in the data, permutation-specific GQTEs aggregate both positive and negative unit-level treatment effects within quantile groups. In contrast to permutation-specific QIDs, permutation-specific GQTEs have the property that control group members at the median of their outcome distribution are always located at the median, regardless of the chosen permutation.

We use permutation-specific GQTEs for income during the pre-time limit period to illustrate the implications of imposing Assumption 5. Figure A2 in Appendix A presents the density functions of permutation-specific GQTEs at selected quantiles, along with the GQTE and associated density functions across quantiles. We observe that the density functions approximate normal distributions as the number of permutations increases (Figures A2a-A2e). We also find that 50 percent of the probability mass of the density functions concentrates near the GQTE (Figure A2f), providing strong support for calculating averages in accordance with Assumption 5. Additionally, our findings illustrate the potential for developing tests to compare alternative identifying assumptions, such as comparing GQTE based on Assumption 5 to GQTE based on alternative weighting schemes.

Figure 2f demonstrates that our approach yields nearly identical results to the conventional QTE approach under rank invariance. While the conventional QTE approach involves calculating quantiles of separate marginal distributions of treatment and control outcomes, our approach is based on averaging unit-level treatment effects within quantile groups of control outcomes. In cases where the rank invariance assumption holds, we have no basis for favoring one approach over another. However, in contrast to our approach, the conventional QTE approach can only be applied if the rank invariance

²²Using the median instead of the average within quantile groups, for example, could potentially yield non-positive treatment effects across the entire distribution, even though the ATE is positive.

assumption holds.

6.3 Quantiles of the distribution of treatment effects

Figure 3 presents QDTE estimates for earnings, transfers, and income, both before and after time limits take effect. In line with the results of Bitler et al. (2006) presented in Figure A1, we find substantial heterogeneity in responses to the welfare reform. During the pre-time limit period, we observe a range of effects on earnings, spanning from approximately $-\$470$ to $\$370$ (Figure 3a). This range broadly aligns with the range of $-\$300$ to $\$500$ depicted in Figure A1a. However, after time limits take effect, we discover a much wider range of effects on earnings, varying from approximately $-\$1,020$ to $\$900$ (Figure 3b), which differs considerably from the corresponding range of $-\$100$ to $\$800$ presented in Figure A1b.

<Figure 3 about here.>

The effects of Jobs First on transfers during the pre-time limit period span a range of approximately $-\$30$ to $\$540$ (Figure 3c), compared to the range of $\$0$ to $\$700$ observed in Figure A1c. The effects on transfers during the post-time limit period range from approximately $-\$470$ to $\$200$ (Figure 3d), which broadly aligns with the range of $-\$550$ to $\$100$ presented in Figure A1d. The effects on income are largely driven by the patterns observed for earnings. The effects on income during the pre-time limit period (Figure 3e), ranging from approximately $\$40$ to $\$570$, are similar to the range of $\$0$ to $\$800$ observed in Figure A1e. In contrast, the effects on income, ranging from approximately $-\$940$ to $\$880$ (Figure 3f), deviate substantially from the range of $-\$300$ to $\$300$ presented in Figure A1f.

In contrast to the results of Bitler et al. (2006), which rely on rank invariance, our findings explicitly account for rank changes. We obtain rank correlation coefficients between potential treatment and control outcomes ranging from approximately 0.7 to 0.8 during the pre-time limit period. While Bitler et al. (2006) maintain the rank invariance assumption across time periods, our analysis reveals a considerable decline in the estimated rank correlation between potential treatment and control outcomes over time, reaching approximately 0.6 to 0.7 during the post-time limit period. The dramatic

differences in the effects on earnings and income during the post-time limit period can be attributed to mobility over time following the intervention.

6.4 Generalized quantile treatment effects

Figure 4 depicts GQTE estimates for earnings, transfers, and income during pre- and post-time limit periods. Bitler et al. (2006) find that Jobs First had no impact at the bottom of the earnings distribution, increased earnings in the middle, and reduced earnings at the top before time limits took effect (Figure A1a).²³ Our GQTE estimates largely support this overall conclusion (Figure 4a). However, when accounting for the possibility of rank changes, we observe small positive effects at the bottom of the earnings distribution, with a gradual increase from \$55 at quantile 1 to \$250 at quantile 45, and a further rise to \$355 at quantile 54. These effects can be attributed to underlying positive average unit-level treatment effects. Specifically, we compare women in the AFDC group with zero earnings to women in the Jobs First group who have either zero or positive earnings.²⁴ In contrast, Bitler et al. (2006) report effects that are exactly zero for the first 48 quantiles of the earnings distribution because earnings are equal to zero for 48 percent of person quarters in the Jobs First sample and 55 percent of person quarters in the AFDC sample.

After time limits take effect, Bitler et al. (2006) observe relatively large positive effects in the middle of the earnings distribution but find no impact at the bottom and top segments (Figure A1b). This finding suggests that the implementation of time limits, which render Jobs First women ineligible for welfare benefits (with possible extensions for some), reinforces positive earnings effects while eliminating negative ones. In contrast, our GQTE estimates reveal increases in both positive effects at the lower

²³Zero effects observed at the bottom of the earnings distribution can be attributed to the comparison of women without earnings in both groups. According to static labor supply theory, assignment to Jobs First should lead to increased earnings, transfers, and income in the middle of the distribution if substitution effects dominate income effects. Negative effects of Jobs First on earnings at the top of the distribution may be attributed to the behavioral-induced eligibility effect (Ashenfelter, 1983). This effect arises from individual decisions to reduce labor supply (and earnings) in order to increase welfare payments.

²⁴This approach is consistent with static labor supply theory. Bitler et al. (2006) note that if a woman with zero earnings in the AFDC group was assigned to the Jobs First group, she could either continue to have zero earnings or enter the labor market and increase her income while still receiving the same amount of transfers. Therefore, the expected unit-level treatment effect for this woman is positive.

end of the earnings distribution, ranging from \$153 at quantile 1 to \$840 at quantile 46, and negative effects at the higher end, ranging from $-\$18$ at quantile 60 to $-\$1,076$ at quantile 98 (Figure 4b). Although incentives to reduce earnings to receive higher welfare payments disappeared with the implementation of time limits, our findings suggest that women at the higher end of the earnings distribution in the Jobs First group were either unable or unwilling to increase their earnings. Instead, we find that negative effects become more pronounced over time when we account for rank changes, indicating that the initially observed negative effects of Jobs First on the earnings distribution were more persistent than previously recognized.

We observe positive effects of Jobs First on transfers for a substantial fraction of the distribution (Figure 4c), reflecting that the Jobs First program's earnings disregard up to the poverty line was more generous than that of AFDC. Bitler et al. (2006) report zero effects for the bottom 20 quantiles of the transfer distribution (Figure A1c) because women with no earnings under AFDC and those with earnings below the poverty line under Jobs First receive maximum benefit payments. However, when accounting for rank changes, we observe small positive effects at the lower end of the distribution, ranging from \$119 at quantile 1 to \$338 at quantile 18. Moreover, we find substantial positive effects (exceeding \$350) for all quantiles ranging from 19 to 50, while observing small or statistically insignificant effects towards the higher end of the distribution. Overall, these findings are similar to those of Bitler et al. (2006) and align with the predictions of static labor supply theory.

In the post-time limit period, Bitler et al. (2006) find that the effects of Jobs First on transfers are precisely zero for the first 47 quantiles, and negative for quantiles 48 to 96 (Figure A1d). They highlight that reductions in the rate of cash assistance are the main contributing factor to the negative effects on transfers. Interestingly, they also report a positive cash assistance gap conditional on eligibility, which can be attributed to the generous earnings disregard provided by Jobs First. When considering rank changes, we uncover small positive effects of Jobs First on transfers. These effects gradually increase from \$13 at quantile 1 to \$200 at quantile 48, and then decrease to \$4 at quantile 54 (Figure 4d). From quantile 55 to 98, we observe negative effects, with a minimum of $-\$452$ at quantile 72. Overall, our findings regarding transfers during the post-time limit period are similar to those of Bitler et al. (2006), indicating that rank

changes had a limited impact on the observed effects.

<Figure 4 about here.>

Bitler et al. (2006) report that Jobs First had no impact on the first 10 quantiles of the income distribution. From quantile 11 to quantile 37, the effects were modest or not statistically significant. However, they observed substantial increases in incomes in the middle and at the top of the distribution, reaching up to \$800 at quantile 82 (Figure A1e). Our findings are largely consistent with these results, indicating that the contribution of rank changes to the observed effects was moderate (Figure 4e). Specifically, our analysis reveals positive (although not always statistically significant) effects at the bottom of the distribution. The effects range from \$66 at quantile 1 to \$379 at quantile 10, gradually declining to \$117 at quantile 24. For the remainder of the distribution, we observe positive effects ranging from \$147 at quantile 25 to \$513 at quantile 89.

Bitler et al. (2006) find that the effects of Jobs First on income during the post-time limit period exhibit a distinct pattern. Specifically, they report that the effects are zero for the first 18 quantiles, negative for the subsequent 22 quantiles, and then turn positive from quantile 46 to quantile 89, before eventually returning to zero or becoming negative (Figure A1f). Bitler et al. (2006) predict that while some women in the Jobs First group may experience an increase in earnings that exceeds the loss in transfers due to the implementation of time limits, this is not the case for others. They also predict that income increases will occur at higher income levels than income decreases if offered wages remain constant. Our estimates validate the first prediction but do not support the second (Figure 4f). When considering rank changes, our analysis reveals substantial positive effects that gradually increase from \$150 at quantile 1 to \$863 at quantile 18, and then gradually decrease to \$60 at quantile 31. Between quantile 32 and quantile 67, we observe small negative effects of approximately $-\$100$ or less (in absolute terms), followed by a decline from $-\$113$ at quantile 68 to $-\$950$ at quantile 98. These findings suggest that women with lower incomes are most responsive to reductions in transfers following the implementation of time limits. Moreover, the negative income effects observed at the higher end of the distribution align with the corresponding negative effects of Jobs First on earnings.

In sum, the discrepancies between our findings and those presented in Bitler et al. (2006) can be attributed entirely to the explicit inclusion of rank changes in our analysis. While static labor supply theory would not have predicted certain effects stemming from rank changes, our results align well with theoretical predictions, particularly during the pre-time limit period. Deviations from static labor supply theory during the post-time limit period can be attributed to persistent effects observed during the pre-time limit period and increased mobility over time.

6.5 Additional parameters of interest

The treatment effect distributions presented in Figures 3 and 4 permit inferences about the proportion of women who experience improvements, setbacks, or remain unaffected by the Jobs First program. Table 3 presents the proportion of positive and negative treatment effects on earnings, transfers, and income during the pre- and post-time limit periods, as derived from QTE estimates and GQTE estimates. The table includes the overall proportion of treatment effects, regardless of their significance levels, as well as the proportions of treatment effects that are statistically significant at the 5%, 1% and 0.1% level, respectively.²⁵

When examining earnings during the pre-time limit period, our findings reveal substantial disparities in the proportions of positive and negative effects derived from QTE and GQTE estimates. Based on QTE estimates, 34.3% of the effects of Jobs First on earnings are positive, while 12.1% are negative. In contrast, the proportion of positive effects based on GQTE estimates is 68.7%, whereas the proportion of negative effects is 31.3%. A considerable gap in the proportion of positive effects remains when we focus on effects that are statistically significant at the 0.1% level (18.2% based on QTE versus 58.6% based on GQTE). However, the gap in the proportion of negative effects disappears when we consider a 0.1% significance level. A similar pattern can be observed for the post-time limit period. These findings highlight substantial disparities between QTE and GQTE estimates in determining the proportions of ‘winners’ and ‘losers’ of the Jobs First program.

²⁵We present alternative significance levels to address the issue of inflated type I error rates associated with conducting the same hypothesis test at each percentile of the distribution. We defer the application of alternative tests, such as the bootstrap test of distributional treatment effects proposed in Abadie (2002), to future work.

<Table 3 about here.>

According to QTE estimates, the overall proportion of positive effects on transfers during the pre-time limit period is 77.8%, while GQTE estimates indicate a higher proportion of 97.0%. At a 0.1% significance level, the proportions of positive effects are 55.6% based on QTE estimates and 78.8% based on GQTE estimates. Both QTE and GQTE reveal no significantly negative effects of Jobs First on transfers before the implementation of time limits. During the post-time limit period, we observe notable differences in the proportions of positive effects. While QTE estimates indicate no significantly positive effects, GQTE estimates show that 49.5% of effects are positive at a 5% significance level, and 10.1% are positive at a 0.1% significance level. In contrast, the proportions of negative effects do not differ much between QTE and GQTE estimates. Overall, we find considerable disparities in the proportion of positive effects obtained from QTE and GQTE estimates. Nonetheless, in addition to assessing the proportions of treatment effects, it is equally important to examine their magnitude. We will discuss this issue below.

Our analysis uncovers notable discrepancies in the proportion of positive treatment effects on income between QTE and GQTE estimates during the pre-time limit period. Specifically, the overall proportion of positive effects based on QTE estimates is considerably lower at 85.9% compared to the corresponding proportion of 99.0% based on GQTE estimates. When focusing on effects that reach statistical significance at a 0.1% level, the proportion of positive effects based on QTE estimates decreases to 52.5%, while the proportion based on GQTE estimates remains higher at 71.7%. We observe no significantly negative effects on income, regardless of whether we consider QTE or GQTE estimates.

Interestingly, QTE and GQTE estimates yield similar proportions of statistically significant effects on income during the post-time limit period. Specifically, we find that the proportion of positive effects at a 5% significance level is 23.2% based on QTE estimates and 24.2% based on GQTE estimates. When considering a 0.1% significance level, these proportions decrease to 6.1% based on QTE estimates and 15.2% based on GQTE estimates. Moreover, we observe that the proportion of negative effects at a 5% significance level is 7.1% based on QTE estimates and 6.1% based on GQTE estimates. We do not find any negative effects on income when applying a 0.1% signifi-

cance level. Despite the observed similarities in the proportions of positive and negative treatment effects on income, the locations of these effects differ substantially between QTE estimates (Figure A1f) and GQTE estimates (Figure 4f).

Table 4 provides an overview of the total gains and losses, and presents the standard deviation of treatment effects for earnings, transfers, and income during the pre- and post-time limit periods. In the following, we focus on the most pronounced disparities observed between QTE estimates and GQTE estimates, which occur for earnings and income, particularly in the post-time limit period. Specifically, when considering the overall gains in earnings during quarters 8-16, QTE estimates yield total gains of approximately \$382,000, while GQTE estimates reveal higher gains of approximately \$649,000. When applying a 0.1% significance level, these gains decrease slightly to approximately \$301,000 based on QTE estimates and to approximately \$616,000 based on GQTE estimates. These findings indicate that the total gains resulting from the Jobs First program were more than twice as large as previously recognized. Figure 4b reveals that a considerable fraction of these gains occurred at the lower end of the earnings distribution, contrasting with the results presented in Figure A1b, which suggests a concentration of gains in the middle of the distribution.

<Table 4 about here.>

In addition to higher total gains, we also observe higher total losses in earnings when accounting for rank changes. However, total losses decrease substantially when we apply a stringent significance level of 0.1%. In this case, we observe total losses of approximately \$31,000 based on GQTE estimates, and no losses based on QTE estimates. Figure 4b illustrates that the observed losses occur at the higher end of the earnings distribution. Accounting for rank changes more than doubles the standard deviation of treatment effects on earnings in quarters 8-16, from 272.6 based on QTE estimates to 578.2 based on GQTE estimates.

Turning to income during the post-time limit period, we observe total gains in income of approximately \$190,000 based on QTE estimates when statistical significance is ignored. Based on GQTE estimates, total gains in income are more than twice as large, reaching around \$383,000. When considering a 0.1% significance level, total gains in income decrease to approximately \$43,000 based on QTE estimates. The

corresponding figure based on GQTE estimates is more than six times higher, totaling approximately \$260,000. We also observe total losses in income during the post-time limit period, which amount to approximately \$51,000 based on QTE estimates when a 5% significance level is applied. The corresponding figure based on GQTE estimates is approximately \$107,000. However, total losses disappear when a 0.1% significance level is applied.

Taken together, our findings suggest that the total gains and losses observed during the post-time limit period are primarily driven by the impact of Jobs First on the earnings distribution. Gains in earnings were successful in mitigating the majority of losses in transfers stemming from the implementation of time limits, which led to the exclusion of Jobs First women from welfare benefits. While the QTE estimates (Figure A1f) indicate that losses occur below the median of the income distribution and gains occur above it, our GQTE estimates (Figure 4f) reveal that gains occur at the lower end of the distribution, while losses, to the extent that they are statistically significant, occur at the higher end.

7 Conclusions

This paper introduces an approach to identify quantiles of the distribution of treatment effects (QDTE) and examines the distributional impacts of Connecticut’s Jobs First welfare experiment on women’s earnings, transfers, and income. Our focus on situations in which relevant outcome predictors are available enables us to identify our parameters of interest under modest assumptions. We leverage a property we term the ‘Rank Correlation Property’ (RCP), which establishes that, under random assignment, rank correlation coefficients between actual treatment outcomes and predicted control state outcomes are identical, regardless of whether predictions are based on treatment or control units. We show that the rank correlation coefficient between potential outcomes is identified under the assumption of equally likely permutations of observation units satisfying the RCP. We use the same assumption to identify QDTE as quantiles of average unit-level treatment effects. Rearranging quantiles yields generalized quantile treatment effects (GQTE), which do not rely on a rank invariance assumption but align with conventional QTE under rank invariance. We employ reweighting methods to

identify our parameters of interest under strong ignorability. Monte Carlo simulations validate the unbiasedness, consistency, and asymptotic normality of our estimators when relevant outcome predictors are available.

We find substantial heterogeneity in the effects of Jobs First on earnings, transfers, and income. We estimate the effects separately for the periods before and after the implementation of time limits. During the pre-time limit period, our GQTE estimates align qualitatively with conventional QTE estimates. In contrast, accounting for rank changes during the post-time limit period reveals both positive effects at the lower end and negative effects at the higher end of the earnings distribution, indicating that the initially observed negative effects on earnings were more persistent than previously recognized. Our findings also suggest that women with lower incomes are most responsive to reductions in transfers following the implementation of time limits. Negative income effects at the higher end of the distribution are consistent with the corresponding negative earnings effects.

We use QTE and GQTE estimates to determine the proportions of positive and negative treatment effects. During the pre-time limit period, GQTE estimates show a higher proportion of positive income effects compared to QTE estimates. During the post-time limit period, QTE and GQTE estimates exhibit similar proportions of income effects, although the locations of these effects differ substantially between the two methods. Total gains and losses resulting from Jobs First are mainly driven by changes in the earnings distribution, with GQTE-based estimates showing considerably higher gains compared to QTE-based estimates. These findings highlight the importance of accounting for rank changes when estimating distributional effects.

Tables and Figures

TABLE 1: ESTIMATION OF τ_{sp} FOR ALTERNATIVE VALUES OF $\tau(\mathbf{Y}_0, \widehat{\mathbf{Y}}_0)$ AND $\tau(\mathbf{Y}_1, \widehat{\mathbf{Y}}_1)$

Target	τ_{sp}								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$\tau(\mathbf{Y}_0, \widehat{\mathbf{Y}}_0)$									
0.9	0.90 (0.00)	0.80 (0.00)	0.70 (0.00)	0.60 (0.00)	0.50 (0.01)	0.40 (0.01)	0.30 (0.01)	0.20 (0.01)	0.10 (0.01)
0.8	0.91 (0.00)	0.80 (0.00)	0.70 (0.01)	0.60 (0.01)	0.50 (0.00)	0.40 (0.01)	0.30 (0.01)	0.20 (0.01)	0.10 (0.01)
0.7	0.92 (0.01)	0.80 (0.00)	0.70 (0.01)	0.60 (0.01)	0.50 (0.01)	0.40 (0.01)	0.30 (0.01)	0.20 (0.01)	0.10 (0.01)
0.6	0.89 (0.01)	0.80 (0.01)	0.70 (0.01)	0.60 (0.01)	0.50 (0.01)	0.40 (0.01)	0.30 (0.01)	0.20 (0.01)	0.10 (0.01)
0.5	0.88 (0.01)	0.79 (0.01)	0.70 (0.01)	0.60 (0.01)	0.50 (0.01)	0.40 (0.01)	0.30 (0.01)	0.20 (0.01)	0.10 (0.01)
0.4	0.87 (0.01)	0.78 (0.01)	0.70 (0.01)	0.61 (0.01)	0.50 (0.01)	0.40 (0.01)	0.30 (0.01)	0.21 (0.01)	0.11 (0.01)
0.3	0.85 (0.01)	0.77 (0.01)	0.69 (0.01)	0.62 (0.02)	0.51 (0.02)	0.41 (0.02)	0.31 (0.02)	0.21 (0.02)	0.11 (0.02)
0.2	0.83 (0.01)	0.75 (0.02)	0.69 (0.01)	0.63 (0.03)	0.52 (0.03)	0.42 (0.02)	0.32 (0.03)	0.22 (0.02)	0.12 (0.02)
0.1	0.76 (0.06)	0.71 (0.02)	0.69 (0.01)	0.66 (0.03)	0.57 (0.05)	0.47 (0.04)	0.37 (0.05)	0.26 (0.05)	0.17 (0.05)

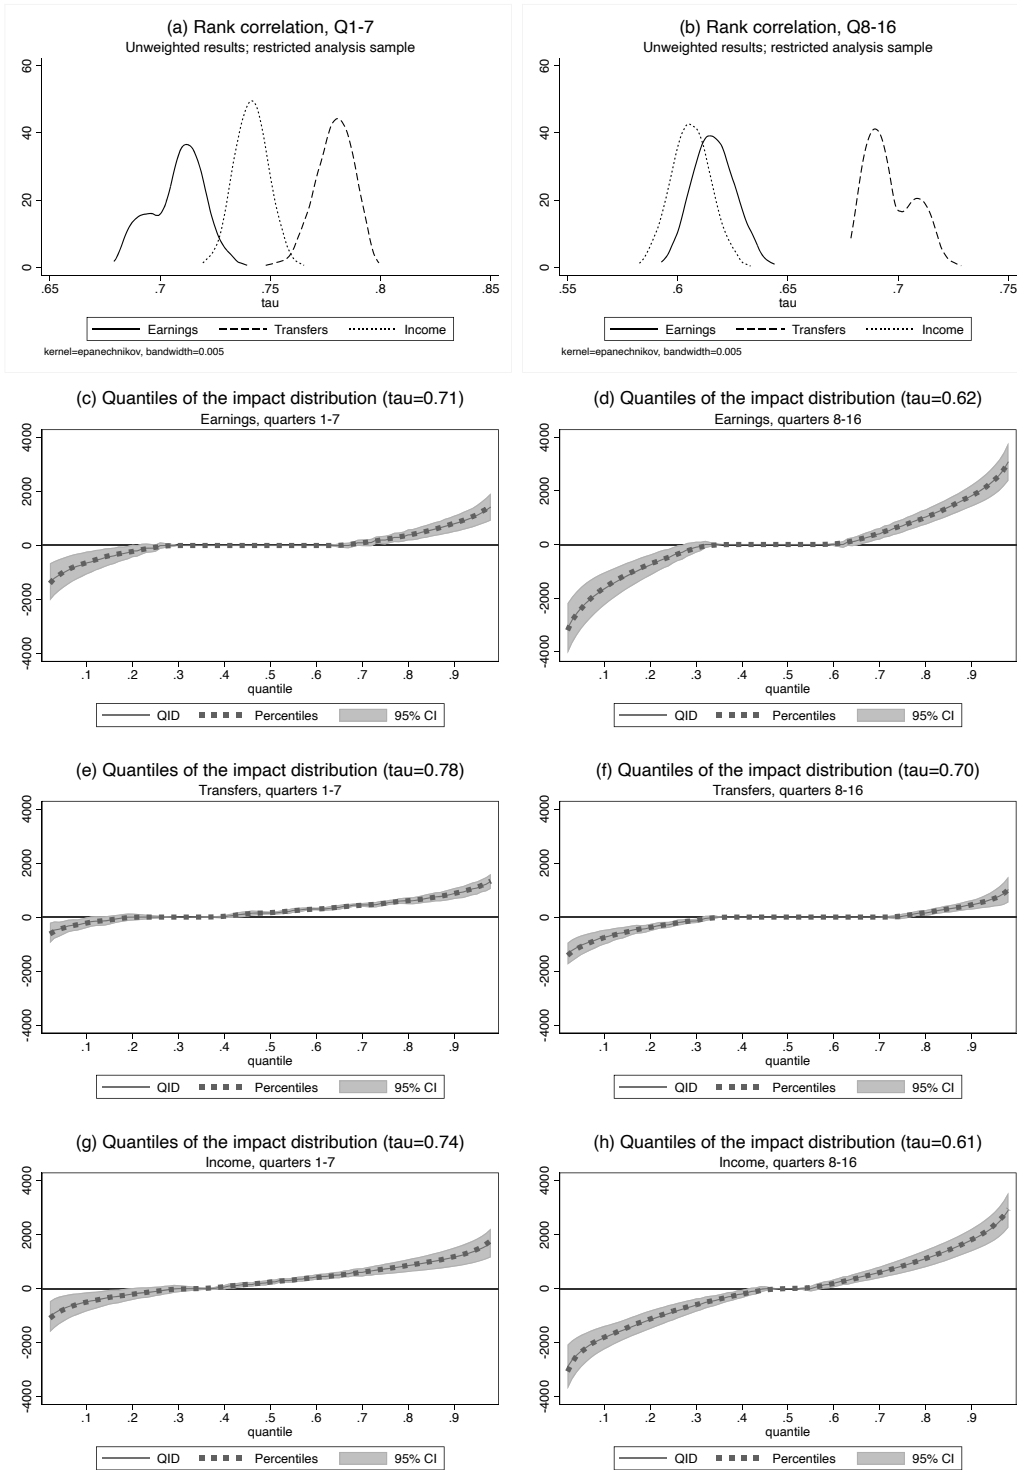
Note: This table presents estimated values of τ_{sp} using simulation data for alternative values of $\tau(\mathbf{Y}_0, \widehat{\mathbf{Y}}_0)$ and $\tau(\mathbf{Y}_1, \widehat{\mathbf{Y}}_1)$. Each estimate is based on 100 random permutations of $2 \times 16,772$ observation units. Bootstrap standard errors (100 replications), which are presented in parentheses, were estimated by drawing a random sample with replacement stratified by group membership for each permutation.

TABLE 2: PRE-TREATMENT CHARACTERISTICS

	Means			<i>p</i> -values		
	Jobs First	AFDC		<i>p_U</i>	<i>p_R</i>	<i>p_{UR}</i>
		<i>U</i>	<i>R</i>			
White	0.362	0.348	0.348	0.307	0.305	0.996
Black	0.368	0.371	0.370	0.836	0.857	0.978
Hispanic	0.207	0.216	0.217	0.423	0.416	0.989
Never married	0.654	0.661	0.660	0.624	0.663	0.957
Div/wid/sep/living apart	0.332	0.327	0.328	0.715	0.753	0.960
HS dropout	0.350	0.334	0.335	0.242	0.279	0.933
HS diploma/GED	0.583	0.604	0.603	0.155	0.168	0.967
More than HS diploma	0.066	0.062	0.061	0.557	0.504	0.935
More than two children	0.235	0.214	0.213	0.089	0.079	0.953
Mother younger than 25	0.289	0.297	0.297	0.573	0.547	0.968
Mother age 25-34	0.410	0.418	0.417	0.609	0.639	0.967
Mother older than 34	0.301	0.286	0.286	0.264	0.267	0.996
Recipient (stock) sample	0.624	0.593	0.592	0.029	0.024	0.942
Average quarterly values						
Earnings	678.91	785.90	785.97	0.009	0.010	0.999
AFDC payments	890.82	835.11	832.47	0.015	0.011	0.907
Food stamps	352.12	339.35	338.45	0.156	0.129	0.918
Fraction of quarters with						
Any earnings	0.322	0.351	0.351	0.006	0.006	0.977
Any AFDC payments	0.573	0.544	0.543	0.026	0.021	0.934
Any food stamps	0.607	0.598	0.597	0.485	0.443	0.943
<i>N</i>	2,396	2,407	2,396			

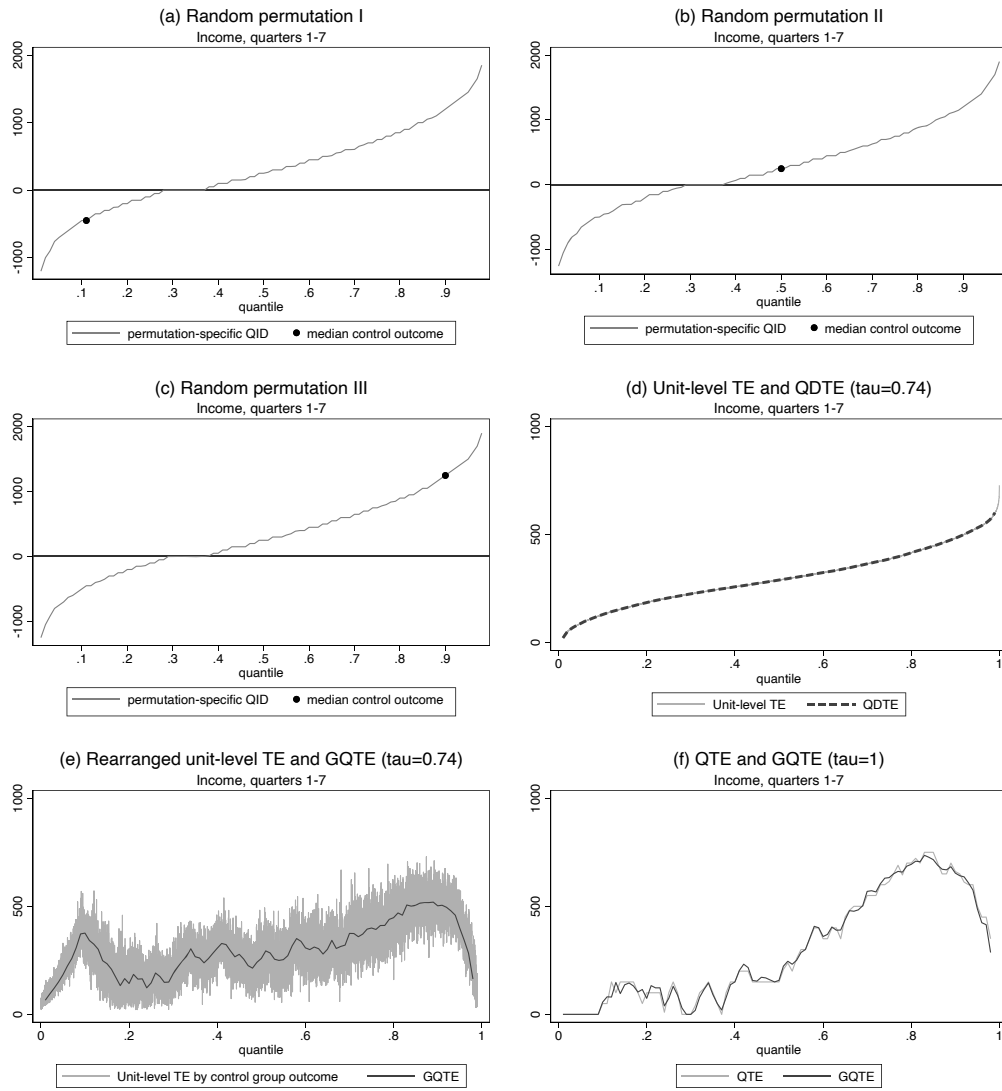
Note: AFDC-*U*: unrestricted AFDC sample; AFDC-*R*: restricted AFDC sample. *p*-values refer to the comparison of means between groups. *p_U*: Jobs First sample vs. AFDC-*U* sample. *p_R*: Jobs First sample vs. AFDC-*R* sample. *p_{UR}*: AFDC-*U* sample vs. AFDC-*R* sample. Several variables have missing values for a small number of observations (marital status: 220 observations, educational attainment: 288 observations, number of children: 161 observations).

FIGURE 1: RANK CORRELATION DISTRIBUTIONS AND QUANTILES OF THE IMPACT DISTRIBUTION (QID)



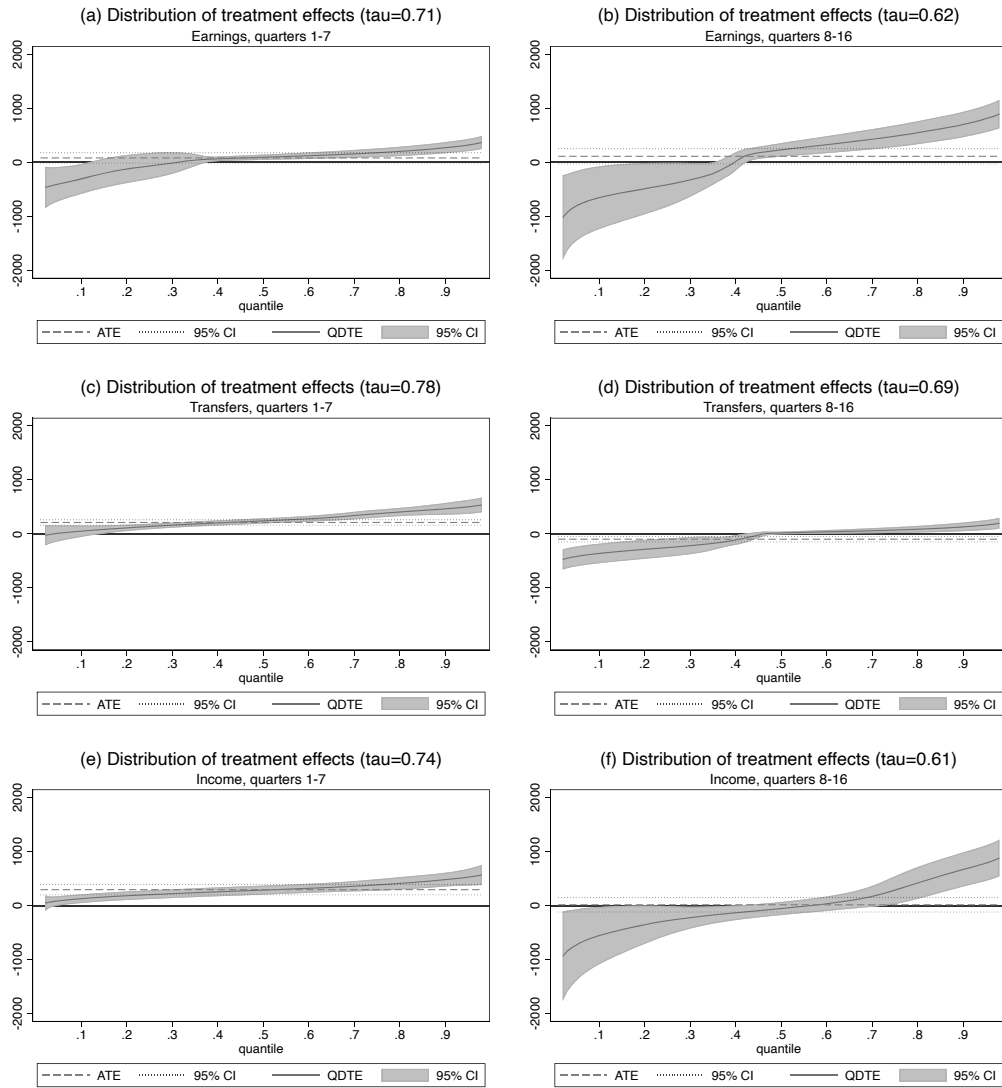
Note: Unweighted results; restricted analysis sample. Figures 1a and 1b: rank correlation distributions based on 100 permutations of treatment outcomes that satisfy the RCP. Kernel density estimates were plotted over a range that exceeds the minimum and maximum values observed by ± 0.0075 . Figures 1c-1h: Confidence intervals are based on clustered bootstrap standard errors (100 replications) to account for repeated observations. Top and bottom percentiles not included due to high sampling variability.

FIGURE 2: RELATION BETWEEN QID, QDTE AND GQTE,
INCOME, QUARTERS 1-7



Note: Unweighted results; restricted analysis sample. Figures 2a-2c: QID for selected random permutations of treatment outcomes. Figures 2d-2e: Average unit-level treatment effects based on 100 random permutations of treatment outcomes. Figure 2f: Results based on rank invariance assumption. Figures 2a-2c and 2e-2f: Top percentile not included due to high sampling variability.

FIGURE 3: QUANTILES OF THE DISTRIBUTION OF TREATMENT EFFECTS (QDTE)



Note: Weighted results; unrestricted analysis sample. Confidence intervals based on clustered bootstrap standard errors (100 replications) to account for repeated observations. Top and bottom percentiles not included due to high sampling variability.