

Quantile Treatment Effects in Difference in Differences Models with Panel Data*

Brantly Callaway[†]

Tong Li[‡]

First Version: April 2015

This Version: September 2018

Abstract

This paper considers identification and estimation of the Quantile Treatment Effect on the Treated (QTT) under a straightforward distributional extension of the most commonly invoked Mean Difference in Differences assumption used for identifying the Average Treatment Effect on the Treated (ATT). Identification of the QTT is more complicated than the ATT though because it depends on the unknown dependence (or copula) between the change in untreated potential outcomes and the initial level of untreated potential outcomes for the treated group. To address this issue, we introduce a new Copula Stability Assumption that says that the missing dependence is constant over time. Under this assumption and when panel data is available, the missing dependence can be recovered, and the QTT is identified. Second, we allow for identification to hold only after conditioning on covariates and provide very simple estimators based on propensity score re-weighting for this case. We use our method to estimate the effect of increasing the minimum wage on quantiles of local labor markets' unemployment rates and find significant heterogeneity.

JEL Codes: C14, C20, C23

Keywords: Quantile Treatment Effect on the Treated, Difference in Differences, Copula, Panel Data, Propensity Score Re-weighting

*We are grateful to the editor Christopher Taber, a co-editor, and three referees for their constructive comments that have greatly improved the paper. We would also like to thank Don Andrews, Stephane Bonhomme, Sergio Firpo, Antonio Galvao, Federico Gutierrez, John Ham, James Heckman, Arthur Lewbel, Magne Mogstad, Derek Neal, John Pepper, Peter Phillips, Pedro Sant'Anna, Azeem Shaikh, Youngki Shin, Steve Stern, Pierre Perron, Ed Vytlačil, Kaspar Wuthrich, and participants in seminars at the Bank of Canada, Beijing Institute of Technology, Boston College, Boston University, Korea University, Lancaster University, National Sun Yat-Sen University, the National University of Singapore, Seoul National University, the University of Chicago, the University of Iowa, the University of Sydney, the University of Texas, the University of Virginia, UC San Diego, Vanderbilt University, Yale University, at the conference in honor of Takeshi Amemiya in Xiamen, China, June 2015, at the 11th World Congress of the Econometric Society, at the 2017 China Meeting of the Econometric Society held in Wuhan, China, June 2017, at the 2017 International Conference on Challenges and Perspectives of Econometrics on Data Analysis held in Hsinchu, Taiwan, June 2017, at the Conference on Econometric Theory and Applications held in Taipei, June 2018, and at the 4th Dongbei Econometrics Workshop held in Dalian, June 2018 for their comments and suggestions. Li acknowledges gratefully the hospitality and support of the Becker Friedman Institute at the University of Chicago. Code for the method proposed in the paper is available as the `panel.qtet` method in the R `qte` package.

[†]Department of Economics, Temple University, 1301 Cecil B. Moore Avenue, Ritter Annex 867, Philadelphia, PA 19122. Phone: (215) 204-8881, Email: brantly.callaway@temple.edu

[‡]Department of Economics, Vanderbilt University, VU Station B #351819, Nashville, TN 37235-1819. Phone: (615) 322-3582, Fax: (615) 343-8495, Email: tong.li@vanderbilt.edu

1 Introduction

Although most research using program evaluation techniques focuses on estimating the average effect of participating in a program or treatment, in some cases a researcher may be interested in understanding the distributional impacts of treatment participation. For example, for two labor market policies with the same mean impact, policymakers are likely to prefer a policy that tends to increase income in the lower tail of the income distribution to one that tends to increase income in the middle or upper tail of the income distribution. In contrast to the standard linear model, the treatment effects literature explicitly recognizes that the effect of treatment can be heterogeneous across different individuals (Heckman and Robb (1985) and Heckman, Smith, and Clements (1997)). Recently, many methods have been developed that identify distributional treatment effect parameters under common identifying assumptions such as selection on observables (Firpo (2007)), access to an instrumental variable (Abadie, Angrist, and Imbens (2002), Chernozhukov and Hansen (2005), Carneiro and Lee (2009), and Frolich and Melly (2013)), or access to repeated observations over time (Athey and Imbens (2006), Bonhomme and Sauder (2011), Chernozhukov, Fernandez-Val, Hahn, and Newey (2013), and Jun, Lee, and Shin (2016)). This paper focuses on identifying and estimating a particular distributional treatment effect parameter called the Quantile Treatment Effect on the Treated (QTT) using a Difference in Differences assumption for identification.

Empirical researchers commonly employ Difference in Differences assumptions to credibly identify the Average Treatment Effect on the Treated (ATT) (early examples include Card (1990) and Card and Krueger (1994)). Despite the prevalence of DID methods in applied work, there has been very little empirical work studying the distributional effects of a treatment with identification that exploits having access to repeated observations over time (Recent exceptions include Meyer, Viscusi, and Durbin (1995), Finkelstein and McKnight (2008), Pomeranz (2015), and Havnes and Mogstad (2015)).

The first contribution of the current paper is to provide identification and estimation results for the QTT under a straightforward extension of the most common Mean Difference in Differences Assumption (Heckman and Robb (1985), Heckman, Ichimura, Smith, and Todd (1998), and Abadie (2005)). In particular, we strengthen the assumption of mean independence between (i) the change in untreated potential outcomes over time and (ii) whether or not an individual is treated to full independence. We call this assumption the Distributional Difference in Differences Assumption.

For empirical researchers, methods developed under the Distributional Difference in Differences Assumption are valuable precisely because the identifying assumptions are straight-

forward extensions of the Mean Difference in Differences assumptions that are frequently employed in applied work. This means that almost all of the intuition for applying a Difference in Differences method for the ATT will carry over to identifying the QTT using our method.

Although applying a Mean Difference in Differences assumption leads straightforwardly to identification of the ATT, using the Distributional Difference in Differences Assumption to identify the QTT faces some additional challenges. The reason for the difference is that Mean Difference in Differences exploits the linearity of the expectation operator. In fact, with only two periods of data (which can be either repeated cross sections or panel) and under the same Distributional Difference in Differences assumption considered in the current paper, the QTT is known to be partially identified (Fan and Yu (2012)) without further assumptions. In practice, these bounds tend to be quite wide. Lack of point identification occurs because the dependence (or copula) between (i) the change in untreated potential outcomes for the treated group and (ii) the initial level of untreated potential outcomes for the treated group is unknown. For identifying the ATT, knowledge of this dependence is not required and point identification results can be obtained.

To move from partial identification back to point identification, we introduce a new assumption which we call the Copula Stability Assumption. This assumption says that the copula, which captures the unknown dependence mentioned above, does not change over time. To give an example, consider the case where the outcome of interest is earnings. The Copula Stability Assumption says that if we observe in the past that the largest earnings increases tended to go to those with the highest earnings, then, in the present (and in the absence of treatment), the largest earnings increases would have gone to those with the highest earnings. Importantly, this does not place any restrictions on the marginal distributions of outcomes over time allowing, for example, the outcomes to be non-stationary. There are two additional requirements for invoking this assumption relative to the Mean Difference in Differences Assumption: (i) access to panel data (repeated cross sections is not enough) and (ii) access to at least three periods of data (rather than at least two periods of data) where two of the periods must be pre-treatment periods and the third period is post-treatment. We show that the additional requirements that the Copula Stability Assumption places on the type of model that is consistent with the Distributional Difference in Differences Assumption are small.

The second contribution of the paper is to extend the results to the case where the identifying assumptions hold conditional on covariates. There are several reasons why this is an important contribution. First, we show that, for many models where an unconditional Mean Difference in Differences assumption holds, the Distributional Difference in Differences

Assumption is likely to require conditioning on covariates. Second, conditional on covariates versions of our assumptions can allow the path of untreated potential outcomes to depend on observed characteristics.

Having simple identification results when identification holds conditional on some covariates stands in contrast to existing methods for estimating QTTs. The methods are either (i) unavailable or at least computationally challenging when the researcher desires to make the identifying assumptions conditional on covariates or (ii) require strong parametric assumptions on the relationship between the covariates and outcomes. Because the ATT can be obtained by integrating the QTT and is available under weaker assumptions, a researcher’s primary interest in studying the QTT is likely to be in the shape of the QTT rather than the location of the QTT. In this regard, the parametric assumptions required by other methods to accommodate covariates may be restrictive because nonlinearities or misspecification of the parametric model could easily be confused with the shape of the QTT. This difference between our method and other methods appears to be fundamental. To our knowledge, there is no work on nonparametrically allowing for conditioning on covariates in alternative methods; and, at the least, doing so would be computationally challenging. Moreover, a similar propensity score re-weighting technique to the one used in the current paper does not appear to be available for existing methods.

Based on our identification results, estimation of the QTT is straightforward and computationally fast. Without covariates, estimating the QTT relies only on estimating unconditional moments, empirical distribution functions, and empirical quantiles. When the identifying assumptions require conditioning on covariates, we estimate the propensity score in a first step, but second step estimation is simple and fast. We show that our estimator of the QTT converges to a Gaussian process at the parametric rate \sqrt{n} even when the propensity score is estimated nonparametrically. This result allows us to conduct uniform inference over a range of quantiles and to test, for example, whether the distribution of treated potential outcomes stochastically dominates the distribution of untreated potential outcomes.

We conclude the paper by analyzing the effect of increasing the minimum wage on quantiles of the unemployment rates of local labor markets. Despite the average effect of increasing the minimum wage on the unemployment rate being close to 0, using our method, we find that the average effect masks substantial heterogeneity. The difference between the 10th percentile of unemployment among counties that had higher minimum wages and the 10th percentile of counterfactual unemployment had they not had higher minimum wages is negative. However, the effect is quite different elsewhere in the distribution. At the median and upper quantiles, the effect is positive. As long as counties do not change their ranks (or

at least do not change their ranks too much) in the distribution of unemployment rates due to the increase in the minimum wage, these results indicate that counties with tight labor markets experienced decreases in the unemployment rate following the minimum wage increase while counties with higher unemployment rates experienced more unemployment due to the increase in the minimum wage. We find similar results using alternative methods such as Quantile Difference in Differences and Change in Changes (Athey and Imbens (2006)).

Because we focus on nonparametric identifying assumptions, the current paper is related to the literature on nonseparable panel data models (Altonji and Matzkin (2005), Evdokimov (2010), Bester and Hansen (2012), Graham and Powell (2012), Hoderlein and White (2012), and Chernozhukov, Fernandez-Val, Hahn, and Newey (2013)). The most similar of these is Chernozhukov, Fernandez-Val, Hahn, and Newey (2013) which considers a nonseparable model and, similarly to our paper, obtains point identification for observations that are observed in both treated and untreated states. Relative to Chernozhukov, Fernandez-Val, Hahn, and Newey (2013), we exploit having access to a control group much more and our setup is compatible with more complicated distributional shifts in outcomes over time such as the top of the income distribution increasing more than the bottom of the income distribution.

Perhaps the most similar work to ours is Athey and Imbens (2006). Their Change in Changes model identifies the QTT for models that are monotone in a scalar unobservable. They assume that the distribution of unobservables does not change over time (though the distribution of unobservables can be different for the treated group and untreated group) but allow for the return to unobservables to change over time. One advantage of their approach relative to ours is that it only requires two periods of data. However, our main assumptions are more closely related to DID assumptions that are frequently invoked in empirical work.

2 Background

The setup and notation used in this paper is common in the statistics and econometrics literature. We consider a panel data case where the researcher has access to at least three periods of data for all agents in the sample; we denote the three periods by t , $t - 1$, and $t - 2$. We focus on the case of a binary treatment. We also focus, as is common in the Difference in Differences literature, on the case where no one receives treatment before the final period which simplifies the exposition; a similar result for a subpopulation of the treated group could be obtained with little modification in the more general case. Let $D = 1$ for individuals that are treated at time t (we suppress an individual subscript i throughout much of the paper to minimize notation) – these individuals form the treated group – and

let $D = 0$ for individuals that are never treated. The researcher observes outcomes Y_t , Y_{t-1} , and Y_{t-2} for each individual in each time period. The researcher also possibly observes some covariates X which we assume are constant over time as in Heckman, Ichimura, Smith, and Todd (1998) and Abadie (2005).

Following the treatment effects literature, we assume that individuals have potential outcomes in the treated or untreated state: Y_{1t} and Y_{0t} , respectively. The fundamental problem is that exactly one (never both) of these outcomes is observed for a particular individual. Using the above notation, the observed outcome Y_t can be expressed as follows:

$$Y_t = DY_{1t} + (1 - D)Y_{0t}$$

Because no one is treated in previous periods, untreated potential outcomes are observed for both the treated group and untreated group.¹ That is,

$$Y_{t-1} = Y_{0t-1} \text{ and } Y_{t-2} = Y_{0t-2}$$

For any particular individual, the unobserved potential outcome is called the counterfactual. The individual's treatment effect, $Y_{1t} - Y_{0t}$ is therefore never available because only one of the potential outcomes is observed for a particular individual. Instead, the literature has focused on identifying and estimating various functionals of treatment effects and the assumptions needed to identify them.

In cases where (i) the effect of a treatment is thought to be heterogeneous across individuals and (ii) understanding this heterogeneity is of interest to the researcher, estimating distributional treatment effects such as quantile treatment effects is likely to be important. Comparing the distribution of observed outcomes to a counterfactual distribution of untreated potential outcomes is a very important ingredient for evaluating the effect of a program or policy (Sen (1997) and Carneiro, Hansen, and Heckman (2001)) and provides more information than the average effect of the program alone. For example, a policy maker may be in favor of implementing a job training program that increases the lower tail of the distribution of earnings while decreasing the upper tail of the distribution of earnings even if the average effect of the program is zero.

For a random variable W , the τ -quantile, w_τ , of W is defined as

$$w_\tau = G_W^{-1}(\tau) := \inf\{w : F_W(w) \geq \tau\} \tag{1}$$

¹To clearly distinguish between treated and untreated potential outcomes, we use a potential outcomes notation where Y_{1t} , Y_{0t-1} , Y_{0t-2} are observed outcomes for the treated group (but Y_{0t} is not an observed outcome for the treated group) and Y_{0t} , Y_{0t-1} , and Y_{0t-2} are observed outcomes for the untreated group.

An example is the 0.5-quantile – the median.² Researchers interested in program evaluation may be interested in other quantiles as well. For example, researchers studying a job training program may be interested in the effect of the program on low income individuals. In this case, they may study the 0.05 or 0.1-quantile. Similarly, researchers studying the effect of a policy on high earners may look at the 0.95-quantile.

Let $F_{Y_{1t}|D=1}$ and $F_{Y_{0t}|D=1}$ denote the distributions of Y_{1t} and Y_{0t} conditional on being in the treated group, respectively. Then, the Quantile Treatment Effect on the Treated (QTT)³ is defined as

$$\text{QTT}(\tau) = F_{Y_{1t}|D=1}^{-1}(\tau) - F_{Y_{0t}|D=1}^{-1}(\tau) \quad (2)$$

The QTT is the parameter studied in this paper. Identification under Difference in Differences assumptions exploits the fact that, for the treated group, both a treated potential outcome and an untreated potential outcome are observed (though at different points in time). A treated potential outcome is never observed for the untreated group. For this reason, Difference in Differences techniques are typically used to identify treatment effect parameters that are conditional on being part of the treated group; for example, Difference in Differences methods for the average effect of participating in a treatment also identify the Average Treatment Effect on the Treated, not the average treatment effect for the population at large.

3 Identification

Let $\Delta Y_{0t} = Y_{0t} - Y_{0t-1}$ denote the time difference in untreated potential outcomes. The most common nonparametric assumption used to identify the ATT in Difference in Differences models is the following:

Assumption 3.1 (Mean Difference in Differences).

$$E[\Delta Y_{0t}|D = 1] = E[\Delta Y_{0t}|D = 0]$$

This is the “parallel trends” assumptions that is common in applied research. It states that, on average, the unobserved change in untreated potential outcomes for the treated group is equal to the observed change in untreated outcomes for the untreated group. To study the QTT, Assumption 3.1 needs to be strengthened because the QTT depends on the

²In this paper, we study Quantile Treatment Effects. A related topic is quantile regression. See Koenker (2005).

³Quantile Treatment Effects were first studied by Doksum (1974) and Lehmann (1974)

entire distribution of untreated outcomes for the treated group rather than only the mean of this distribution.

The next assumption strengthens Assumption 3.1 and this is the assumption maintained throughout the paper.

Distributional Difference in Differences Assumption.

$$\Delta Y_{0t} \perp\!\!\!\perp D$$

The Distributional Difference in Differences Assumption says that the distribution of the change in untreated potential outcomes does not depend on whether or not the individual belongs to the treated or the untreated group. Intuitively, it generalizes the idea of “parallel trends” holding on average to the entire distribution. In applied work, the validity of using a Difference in Differences approach to estimate the ATT hinges on whether the unobserved trend for the treated group can be replaced with the observed trend for the untreated group. This is exactly the same sort of thought experiment that needs to be satisfied for the Distributional Difference in Differences Assumption to hold. Being able to invoke a standard assumption to identify the QTT stands in contrast to the existing literature on identifying the QTT in similar models which generally require less familiar assumptions on the relationship between observed and unobserved outcomes.

Using statistical results on the distribution of the sum of two random variables with known marginal distributions but unknown copula, Fan and Yu (2012) show that this assumption is not strong enough to point identify the counterfactual distribution $F_{Y_{0t}|D=1}$, but it does partially identify it. In practice, these bound tend to be very wide – too wide to be useful in most applications.

3.1 Main Results: Identifying QTT in Difference in Differences Models

The main theoretical contribution of this paper is to impose a Distributional Difference in Differences Assumption plus additional data requirements and an additional assumption that may be plausible in many applications to identify the QTT. The additional data requirement is that the researcher has access to at least three periods of panel data with two periods preceding the period where individuals may first be treated. This data requirement is stronger than is typical in most Difference in Differences setups which usually only require two periods of repeated cross-sections (or panel) data. The additional assumption is that the

dependence – that is, the copula⁴ – between (i) the change in untreated potential outcomes for the treated group and (ii) the initial level of untreated potential outcomes for the treated group is stable over time. This assumption says that if, in the past, the largest increases in outcomes tend to go to those initially at the top of the distribution, then in the present, the largest increases in outcomes will tend to go to those who start out at the top of the distribution. It does not restrict what the distribution of the change in outcomes over time is nor does it restrict the distribution of outcomes in the previous period; instead, it restricts the dependence between these two random variables. We discuss this assumption in more detail and show how it can be used to point identify the QTT below.

Intuitively, the reason why a restriction on the dependence between $(\Delta Y_{0t}|D=1)$ and $(Y_{0t-1}|D=1)$ is useful is the following. If the joint distribution $(\Delta Y_{0t}, Y_{0t-1}|D=1)$ were known, then $F_{Y_{0t}|D=1}$ (the distribution of interest) could be derived from it. The marginal distributions $F_{\Delta Y_{0t}|D=1}$ (through the Distributional Difference in Differences assumption) and $F_{Y_{0t-1}|D=1}$ (from the data) are both identified. However, because observations are observed separately for untreated and treated individuals, even though each of these marginal distributions are identified, the joint distribution is not identified. Since, from Sklar’s Theorem (Sklar (1959)), joint distributions can be expressed as the copula function (capturing the dependence) of the two marginal distributions, the only piece of information that is missing is the copula.⁵ We use the idea that the dependence is the same between period t and period $(t-1)$. With this additional information, $F_{\Delta Y_{0t}, Y_{0t-1}|D=1}$ is identified and therefore the counterfactual distribution of untreated potential outcomes for the treated group, $F_{Y_{0t}|D=1}$, is identified.

The time invariance of the dependence between $(\Delta Y_{0t}|D=1)$ and $(Y_{0t-1}|D=1)$ can be expressed in the following way. Let $F_{\Delta Y_{0t}, Y_{0t-1}|D=1}$ be the joint distribution of ΔY_{0t} and Y_{0t-1} for the treated group. By Sklar’s Theorem

$$F_{\Delta Y_{0t}, Y_{0t-1}|D=1}(\delta, y') = C_{\Delta Y_{0t}, Y_{0t-1}|D=1}(F_{\Delta Y_{0t}|D=1}(\delta), F_{Y_{0t-1}|D=1}(y')) \quad (3)$$

where $C_{\Delta Y_{0t}, Y_{0t-1}|D=1}(\cdot, \cdot)$ is a copula function.⁶ Next, we state the second main assumption which replaces the unknown copula with the copula for the same outcomes but in the previous period which is identified because no one is treated in the periods before t .

⁴The copula of two random variables is the joint distribution of the ranks of the two random variables. It contains all the information about the dependence between the two random variables, but it does not contain any information about the marginal distributions of the random variables.

⁵For a continuous distribution, the copula representation is unique. Joe (1997), Nelsen (2007), and Joe (2015) are useful references for more details on copulas.

⁶The bounds in Fan and Yu (2012) arise by replacing the unknown copula function $C_{\Delta Y_{0t}, Y_{0t-1}|D=1}(\cdot, \cdot)$ with those that make the upper bound the largest and lower bound the smallest.

Copula Stability Assumption.

$$C_{\Delta Y_{0t}, Y_{0t-1}|D=1}(\cdot, \cdot) = C_{\Delta Y_{0t-1}, Y_{0t-2}|D=1}(\cdot, \cdot)$$

The Copula Stability Assumption says that the dependence between $(\Delta Y_{0t}|D = 1)$ and $(Y_{0t-1}|D = 1)$ is the same as the dependence between $(\Delta Y_{0t-1}|D = 1)$ and $(Y_{0t-2}|D = 1)$. It is important to note that this assumption does not require any *particular* dependence structure, such as independence or perfect positive dependence; rather, it requires that whatever the dependence structure is in the past, one can recover it and reuse it in the current period. It also does not require choosing any parametric copula. However, it may be helpful to consider a simple, more parametric example. If the copula of $(\Delta Y_{0t-1}|D = 1)$ and $(Y_{0t-2}|D = 1)$ is Gaussian with parameter ρ , the Copula Stability Assumption says that the copula continues to be Gaussian with parameter ρ in period t but the marginal distributions are allowed to change in unrestricted ways. Likewise, if the copula is Archimedean, the Copula Stability Assumption requires the generator function to be constant over time but the marginal distributions can change in unrestricted ways.

One of the key insights of this paper is that, in some particular situations such as the panel data case considered in the paper, we are able to observe the historical dependence between the marginal distributions. There are many applications in economics where the missing piece of information for identification is the dependence between two random variables. In those cases, previous research has resorted to (i) assuming some dependence structure such as independence or perfect positive dependence or (ii) varying the copula function over some or all possible dependence structures to recover bounds on the joint distribution of interest. To our knowledge, we are the first to use historical observed outcomes to obtain a historical dependence structure and then assume that the dependence structure is stable over time.

Before presenting the identification result, we need some additional assumptions. First, let $\Delta Y_t = Y_t - Y_{t-1}$. Let $\Delta \mathcal{Y}_{t|D=0}$ denote the support of the change in outcomes for the untreated group in period t . Let $\Delta \mathcal{Y}_{t-1|D=1}$, $\mathcal{Y}_{t-1|D=1}$, and $\mathcal{Y}_{t-2|D=1}$ denote the support of the change in outcomes for the treated group in period $(t - 1)$, the support of outcomes for the treated group in period $(t - 1)$, and the support of outcomes for the treated group in period $(t - 2)$, respectively.

Assumption 3.2. *Each of the random variables ΔY_t for the untreated group and ΔY_{t-1} , Y_{t-1} , and Y_{t-2} for the treated group are continuously distributed on their support with densities that are uniformly bounded from above and bounded away from 0.*

Assumption 3.3. *The observed data $\{Y_{it}, Y_{it-1}, Y_{it-2}, X_i, D_i\}_{i=1}^n$ are independent and identically distributed draws from the joint distribution $F_{Y_t, Y_{t-1}, Y_{t-2}, X, D}$; and $Y_{it} = D_i Y_{1it} + (1 -$*

$D_i)Y_{0it}, Y_{it-1} = Y_{0it-1}, \text{ and } Y_{it-2} = Y_{0it-2}.$

Assumption 3.2 says that outcomes are continuously distributed. Copulas are unique on the range of their marginal distributions; thus, continuously distributed outcomes guarantee that the copula is unique. However, for the Copula Stability Assumption, one could weaken this assumption to $Range(F_{\Delta Y_{0t}|D=1}) \subseteq Range(F_{\Delta Y_{t-1}|D=1})$ and $Range(F_{Y_{t-1}|D=1}) \subseteq Range(F_{Y_{t-2}|D=1})$ and still obtain point identification. On the other hand, although neither our DDID Assumption nor the standard mean DID Assumption explicitly require continuously distributed outcomes, it should be noted that standard limited dependent variable models with unobserved heterogeneity would not generally satisfy either of these DID assumptions. Assumption 3.3 says that we are in the case with panel data, and that no one is treated in the first two periods. Assumption 3.3 could potentially be relaxed in several ways. More periods of data could be available – our method requires at least three periods of data, but more periods could be incorporated (for example, it seems possible to extend the approach of Callaway and Sant’Anna (2018) for the ATT to our case for the QTT). Also, our setup could allow for some individuals to be treated in earlier periods than the last one and our results would continue to go through for the group of individuals that are first treated in the last period; considering the case where no one is treated before the last period is standard in DID setups. Assumption 3.3 also says that other covariates X are time invariant. This assumption can be relaxed by focusing on the subset of individuals whose covariates do not change over time.⁷

Theorem 1. *Under the Distributional Difference in Differences Assumption, the Copula Stability Assumption, and Assumptions 3.2 and 3.3*

$$\begin{aligned} & F_{Y_{0t}|D=1}(y) \\ &= E \left[\mathbb{1}\{F_{\Delta Y_i|D=0}^{-1}(F_{\Delta Y_{t-1}|D=1}(\Delta Y_{t-1})) \leq y - F_{Y_{t-1}|D=1}^{-1}(F_{Y_{t-2}|D=1}(Y_{t-2}))\} | D = 1 \right] \quad (4) \end{aligned}$$

and

$$QTT(\tau) = F_{Y_i|D=1}^{-1}(\tau) - F_{Y_{0t}|D=1}^{-1}(\tau)$$

which is identified

Theorem 1 is the main identification result of the paper. It says that the counterfactual

⁷Appendix A also discusses the possibility of including time varying covariates though they must enter our model in a more restrictive way than time invariant covariates. Essentially, the problem with time varying covariates is that that one cannot separate individuals changing ranks in the distribution of outcomes over time due to changes in covariates or due to unobservables.

distribution of untreated potential outcomes for the treated group is identified.⁸ To provide some intuition, we provide a short outline of the proof (the full proof is contained in Appendix B). First, notice that $P(Y_{0t} \leq y|D = 1) = E[\mathbb{1}\{\Delta Y_{0t} + Y_{0t-1} \leq y\}|D = 1]$. This expression is an integral over the joint distribution of $(\Delta Y_{0t}, Y_{0t-1}|D = 1)$ and will be identified when the joint distribution is identified. Under the Distributional DID assumption, this joint distribution is not identified (though the marginals are), but the Copula Stability Assumption replaces the unknown copula in Equation (3) with the observed copula for the treated group in the previous period which leads to the identification result. Replacing the unknown copula with a copula from the past is what increases the required number of periods from two to three.⁹ The particular form of the result in Theorem 1 arises from using the dependence structure in period $t - 1$ (notice that the expectation is over ΔY_{t-1} and Y_{t-2}). The terms of the form $F^{-1}(F(\cdot))$ “adjust” forward outcomes from the previous period and account for the marginal distributions changing over time. Finally, the Distributional Difference in Differences Assumption allows us to replace $F_{\Delta Y_{0t}|D=1}^{-1}(\cdot)$ with $F_{\Delta Y_{0t}|D=0}^{-1}(\cdot)$ which is just the quantiles of the distribution of the change in (observed) untreated outcomes for the untreated group.

The following example shows what additional conditions need to be satisfied for our model to be valid in a standard DID setup.

Example 1. *Consider the following baseline model for Mean DID.*

$$Y_{0it} = \theta_t + \eta_i + v_{it}$$

where θ_t is a time fixed effect that is common for the treated and untreated groups, η_i is individual heterogeneity that may be distributed differently across the treated group and untreated group, and v_{it} are time varying unobservables.¹⁰ For Mean DID to identify the ATT, it must be the case that $E[\Delta v_t|D = 1] = E[\Delta v_t|D = 0]$. Sufficient conditions for the assumptions in our model to hold are (i) $\Delta v_t \perp\!\!\!\perp D$ and (ii) $C_{\Delta v_t, \eta + v_{t-1}|D=1} = C_{\Delta v_{t-1}, \eta + v_{t-2}|D=1}$.

Condition (i) just strengthens Mean DID to Distributional DID. Condition (ii) implies

⁸Although we focus on the QTT, Theorem 1 says that the counterfactual distribution of untreated potential outcomes for the treated group is identified. This also implies that any functional of this distribution is identified. Examples include the variance of untreated potential outcomes, the Gini coefficient, and Lorenz curves; these might be of interest for social welfare calculations. See, for example, Barrett and Donald (2009) and Firpo and Pinto (2016). We thank the editor for pointing this out.

⁹Adding and subtracting Y_{0t-1} is also the first step for showing that the Mean Difference in Differences Assumption identifies $E[Y_{0t}|D = 1]$; the problem is much easier in the mean case though due to the linearity of expectations and no indicator function which implies that only the marginal distributions need to be identified.

¹⁰One other thing to note in this model is that it does not restrict how treated outcomes are generated at all which is standard in the Mean DID case and holds in our case as well.

that the Copula Stability Assumption will hold. An interesting sufficient condition for Condition (ii) is $(v_t, v_{t-1} | \eta, D = 1)$ and $(v_{t-1}, v_{t-2} | \eta, D = 1)$ follow the same distribution (this implies Condition (ii) because it implies that the joint distributions $F_{v_t, v_{t-1}, \eta}$ and $F_{v_{t-1}, v_{t-2}, \eta}$ are equal). Condition (ii) will also hold automatically if the time varying unobservables are iid. Condition (ii) allows for the distribution of the time varying unobservables to change over time, it allows for serial correlation in the time varying unobservables, and it allows for the time varying unobservables to be correlated with the individual heterogeneity. Each of these are realistic possibilities in applied work.

We prove the validity of the claims in Example 1 in Appendix B. Some comments on Example 1 are in order. For identifying the ATT, the setup in Example 1 is straightforward. However, obtaining quantile treatment effects is much more challenging because the model is nonlinear in this case.

There are several main existing approaches to identifying quantile effects in this sort of model. First, one could try to estimate the individual fixed effects, which is the approach generally taken in the fixed effects quantile regression literature.¹¹ Relative to our approach, this would require a large number of time periods and the resulting estimates would have a different interpretation.¹² Another idea would be to impose additional independence conditions among the unobservables (e.g. independence between η and the time varying unobservables and that the time varying unobservables are independent over time) and use results that come primarily from the measurement error literature (e.g. Li and Vuong (1998), Evdokimov (2010), Arellano and Bonhomme (2016), and Freyberger (2018)). Our approach does not require any of these additional conditions. Finally, under the additional condition that $(v_t, v_{t-1} | \eta, D = 1)$ follows the same distribution as $(v_{t-1}, v_{t-2} | \eta, D = 1)$, the approach of Chernozhukov, Fernandez-Val, Hahn, and Newey (2013) as well as the Change in Changes model (Athey and Imbens (2006)) would hold.¹³ But this extra condition is substantially

¹¹The work on panel quantile regression includes Koenker (2004), Abrevaya and Dahl (2008), Lamarche (2010), Canay (2011), Rosen (2012), Galvao, Lamarche, and Lima (2013), Chen (2015), Li and Oka (2015), and Arellano and Bonhomme (2016), among others. Another difference is that, because we do not impose a parametric model, our method allows for the effect of treatment to vary across individuals with different covariates in an unspecified way. On the other hand, our method only applies to the case where the researcher is interested only in the effect of a binary treatment; quantile regression methods can deliver estimates for multiple, possibly continuous variables.

¹²We focus on an unconditional QTT whereas the quantile treatment effects identified in panel QR models are conditional – both on covariates and on unobserved heterogeneity. This means that the results from our method should be interpreted in the same way as the difference between treated and untreated quantiles if individuals were randomly assigned to treatment. See Frolich and Melly (2013) for a good discussion of the difference between conditional and unconditional quantile treatment effects.

¹³The Change in Changes model also requires an additional support condition in this type of model that is not required using our approach. In particular, our approach is likely to perform better when the distribution of η is quite different across the treated and untreated groups and especially in the tails of the distribution.

stronger; it implies that the distribution of the outcomes can only shift location over time. Condition (ii) is substantially weaker than this and can allow the distribution of untreated potential outcomes to shift in arbitrary ways over time.

4 Allowing for covariates

In our view, one of the main reasons that there has been little use of distributional methods with panel data is that existing work has focused primarily on the case without conditioning on other covariates.¹⁴ First, note that all of our identification results in the previous section continue to go through after conditioning on covariates X . However, we do not pursue this approach here (Appendix A discusses this possibility in more detail); the main cost of this approach is that nonparametric estimation would be very challenging in realistic applications where the number of covariates is moderately large due to needing to estimate multiple conditional distributions. Instead, this section extends the previous results to the case where a Conditional DDID assumption holds, but the Copula Stability Assumption continues to hold unconditionally. This setup generalizes the case where Unconditional DDID also holds, but imposes more restrictions on how covariates enter the model. The primary restriction is that the effect of covariates should not change over time. The primary benefit is that estimation is very simple, involving only an additional first-step estimate of the propensity score. This setup strikes a good balance between generality of assumptions and ease of implementation that is likely to be beneficial in applied work. We make the following assumption.

Conditional Distributional Difference in Differences Assumption.

$$\Delta Y_{0t} \perp\!\!\!\perp D|X$$

This assumption says that, after conditioning on covariates X , the distribution of the change in untreated potential outcomes for the treated group is equal to the distribution of the change in untreated potential outcomes for the untreated group. The next example shows that having the Conditional DDID assumption may be important even in cases where an unconditional mean DID assumption holds and would identify the ATT

We demonstrate this difference in the Supplementary Appendix using Monte Carlo simulations.

¹⁴Recent work such as Melly and Santangelo (2015) and Callaway, Li, and Oka (2018) has begun relaxing this restriction.

Example 2. Consider the following model

$$Y_{it} = q(U_{it}, D_i, X_i) + \eta_i$$

with $(U_t, U_{t-1}|X, \eta, D)$ and $(U_{t-1}, U_{t-2}|X, \eta, D) \sim F_{U_1, U_2}$ where F_{U_1, U_2} is a bivariate distribution with uniform marginals, η is time invariant unobserved heterogeneity that may be correlated with observables and distributed differently for the treated and untreated groups, and $q(\tau, d, x)$ is strictly increasing in τ for all $(d, x) \in \{0, 1\} \times \mathcal{X}$.

In this model,

- The Unconditional Mean Difference in Differences Assumption holds
- The Unconditional Distributional Difference in Differences Assumption does not hold
- The Conditional Distributional Difference in Differences Assumption holds
- The Unconditional Copula Stability Assumption holds

In Appendix B, we show that the claims in Example 2 hold. The model in Example 2 is a Skorohod representation for panel quantile regression while allowing for serial correlation among U . This model allows the effect of covariates to be different at different parts of the conditional distribution. For example, if Y is earnings, it is well known that the effect of education is different at different parts of the conditional distribution. One sufficient condition for the Unconditional DDID assumption is that X has only a location effect on outcomes. Another sufficient condition is that the distribution of X is the same for the treated and untreated groups. Neither of these conditions seem likely to hold in the types of applications where a researcher is interested in understanding the distributional effect of a program or policy.

Example 2 is a leading case for using distributional methods to understand heterogeneity in the effect of a treatment, and the main conclusion to be reached from this example is that even when an unconditional mean DID assumption holds, one is likely to need to condition on covariates to justify the DDID assumption. On the other hand, in this model, the Unconditional Copula Stability Assumption continues to hold.

By invoking the Conditional Distributional Difference in Differences Assumption rather than the Unconditional Distributional Difference in Differences Assumption, it is important to note that, for the purpose of identification, the only part of Theorem 1 that needs to be adjusted is the identification of $F_{\Delta Y_{0t}|D=1}$. Under the Unconditional Distributional Difference in Differences Assumption, this distribution could be replaced directly by $F_{\Delta Y_{1t}|D=0}$; however, now we utilize a propensity score re-weighting technique to replace this distribution with

another object (discussed more below). Importantly, all other objects in Theorem 1 can be handled in exactly the same way as they were previously. Particularly, the Copula Stability Assumption continues to hold without needing any adjustment such as conditioning on X .

With covariates, we also require an additional standard assumption for identification.

Assumption 4.1. $p := P(D = 1) > 0$ and $p(x) := P(D = 1|X = x) < 1$.

The first part of this assumption says that there is some positive probability that individuals are treated. The second part says that for an individual with any possible value of covariates x , there is some positive probability that he will not be treated. This is a standard overlap assumption used in the treatment effects literature.

Theorem 2. *Under the Conditional Distributional Difference in Differences Assumption, the Copula Stability Assumption, and Assumptions 3.2, 3.3 and 4.1*

$$\begin{aligned} F_{Y_{0t}|D=1}(y) &= \mathbb{E} \left[\mathbb{1} \left\{ F_{\Delta Y_{0t}|D=1}^{p, -1} (F_{\Delta Y_{t-1}|D=1}(\Delta Y_{t-1})) \leq y - F_{Y_{t-1}|D=1}^{-1} (F_{Y_{t-2}|D=1}(Y_{t-2})) \right\} | D = 1 \right] \end{aligned}$$

where

$$F_{\Delta Y_{0t}|D=1}^p(\delta) = \mathbb{E} \left[\frac{1-D}{p} \frac{p(X)}{1-p(X)} \mathbb{1} \{ \Delta Y_t \leq \delta \} \right] \quad (5)$$

and

$$\text{QTT}(\tau) = F_{Y_{1t}|D=1}^{-1}(\tau) - F_{Y_{0t}|D=1}^{-1}(\tau)$$

which is identified

This result is very similar to the main identification result in Theorem 1. The only difference is that $F_{\Delta Y_{0t}|D=1}$ is no longer identified by the distribution of untreated potential outcomes for the untreated group; instead, it is replaced by the re-weighted distribution in Equation 5. Equation 5 can be understood in the following way. It is a weighted average of the distribution of the change in outcomes experienced by the untreated group. The $\frac{p(X)}{1-p(X)}$ term weights up untreated observations that have covariates that make them more likely to be treated. Equation 5 is almost exactly identical to the re-weighting estimators given in Hirano, Imbens, and Ridder (2003), Abadie (2005), and Firpo (2007); the only difference is the term $\mathbb{1} \{ \Delta Y_t \leq \delta \}$ in our case is given by Y_t , ΔY_t , and $\mathbb{1} \{ Y_t \leq y \}$ in each of the other cases, respectively.

5 Estimation

In this section, we discuss the estimation procedure as well as outlining an inference procedure to conduct uniformly valid inference over a range of quantiles using the empirical bootstrap. We provide formal theoretical results for the limiting process of our estimator in Appendix C as well as a formal justification for the use of the empirical bootstrap in Appendix C.

We estimate the QTT by

$$\widehat{\text{QTT}}(\tau) = \hat{F}_{Y_{1t}|D=1}^{-1}(\tau) - \hat{F}_{Y_{0t}|D=1}^{-1}(\tau)$$

The first term is estimated directly from the data by inverting the estimated empirical distribution of observed outcomes for the treated group.

$$\hat{F}_{Y_{1t}|D=1}^{-1}(\tau) = \inf\{y : \hat{F}_{Y_{1t}|D=1}(y) \geq \tau\}$$

We estimate counterfactual quantiles by

$$\hat{F}_{Y_{0t}|D=1}^{-1}(\tau) = \inf\{y : \hat{F}_{Y_{0t}|D=1}(y) \geq \tau\}$$

where

$$\hat{F}_{Y_{0t}|D=1}(y) = \frac{1}{n_D} \sum_{i \in \mathcal{D}} \mathbb{1}\{\hat{F}_{\Delta Y_t|D=0}^{-1}(\hat{F}_{\Delta Y_{t-1}|D=1}(\Delta Y_{it-1})) \leq y - \hat{F}_{Y_{t-1}|D=1}^{-1}(\hat{F}_{Y_{t-2}|D=1}(Y_{it-2}))\}$$

which follows from the identification result in Theorem 1 and where distribution functions are estimated by empirical distribution functions and quantile functions are estimated by inverting empirical distribution functions; here n_D is the number of observations in the treated group and \mathcal{D} is the set of treated individuals.

The final issue is estimating $F_{\Delta Y_{0t}|D=1}^{p,-1}(\delta)$ when identification depends on covariates. Using the identification result in Theorem 2, we can easily construct an estimator of the distribution function

$$\hat{F}_{\Delta Y_{0t}|D=1}^p(\delta) = \frac{1}{n} \sum_{i=1}^n \frac{(1 - D_i)}{p} \frac{\hat{p}(X_i)}{(1 - \hat{p}(X_i))} \mathbb{1}\{\Delta Y_{t,i} \leq \delta\} \Bigg/ \frac{1}{n} \sum_{i=1}^n \frac{(1 - D_i)}{p} \frac{\hat{p}(X_i)}{(1 - \hat{p}(X_i))}$$

where the last term in the denominator normalizes the weights to sum to one in finite samples; it ensures that $\hat{F}_{\Delta Y_{0t}|D=1}^p$ is a distribution function and this term is asymptotically

negligible. One can invert this distribution to obtain its quantiles.

To conduct inference, we propose using the empirical bootstrap to construct uniform confidence bands that cover $QTT(\tau)$ with fixed probability for all values $\tau \in \mathcal{T} = [\epsilon, 1 - \epsilon] \subset (0, 1)$ for some small, positive ϵ . We derive formal results on the limiting process and the validity of the bootstrap for our estimator of the QTT in Appendix C in the unconditional case and in the Supplementary Appendix in the case where the Distributional DID Assumption holds after conditioning on covariates. In the latter case, we consider both parametric and nonparametric estimators of the propensity score and provide high level conditions for the estimator of the propensity score that can be satisfied by other estimators (e.g. semiparametric estimators) under some regularity conditions.

Let $\widehat{QTT}^*(\tau)$ denote an estimate of the QTT using the same steps as above but with a bootstrapped sample (i.e. a sample with n observations drawn from the original sample with equal probabilities and with replacement). Theorem 4, in the Appendix, shows that the empirical bootstrap can be used to approximate the limiting process of our estimator.¹⁵ To obtain uniform confidence bands, let B be the number of bootstrap iterations and for $b = 1, \dots, B$ calculate

$$I^b = \sup_{\tau \in \mathcal{T}} \hat{\Sigma}(\tau)^{-1/2} \left| \sqrt{n}(\widehat{QTT}^b(\tau) - \widehat{QTT}(\tau)) \right|$$

where $\widehat{QTT}^b(\tau)$ is estimated using a bootstrapped sample and where $\hat{\Sigma}(\tau)^{1/2} = (q_{0.75}(\tau) - q_{0.25}(\tau))/(z_{0.75} - z_{0.25})$, which is the bootstrapped interquartile range divided by the interquartile range of a standard normal random variable; this is a uniformly consistent of $\Sigma(\tau)^{1/2}$ with $\Sigma(\tau)$ being the asymptotic variance function of the QTT. Then, a $(1 - \alpha)$ confidence band is given by

$$\hat{C}_{QTT}(\tau) = \widehat{QTT}(\tau) \pm c_{1-\alpha}^B \hat{\Sigma}(\tau)^{1/2} / \sqrt{n}$$

for $\tau \in \mathcal{T}$ and where $c_{1-\alpha}^B$ is the $(1 - \alpha)$ quantile of $\{I^b\}_{b=1}^B$.

6 Application

In this section, we use our method to study the effect of increasing the minimum wage on county-level unemployment rates. There is a wide body of research that studies the effect of the minimum wage on employment exploiting policy level changes across states (for example,

¹⁵In the case where a researcher is interested in an “extreme” quantile such as the 0.01 quantile or perhaps the 0.05 quantile, then alternative inference procedures may need to be used (see, for example, Chernozhukov, Fernandez-Val, and Kaji (2018)).

Neumark and Wascher (1992) and Dube, Lester, and Reich (2010), among many others). Like most of the literature, we use variation in state-level changes in the minimum wage. Also, we suppose that there may be time invariant differences in the unemployment rate across counties that cannot be accounted for by observable differences in county characteristics. This implies that a DID approach should be used and is in line with much of the literature on minimum wage increases.

The aim of this section is different from most research on the effect of increasing the minimum wage. The literature almost exclusively looks at the average effect, or the coefficient in a linear regression model, of increasing the minimum wage on employment for teenagers, restaurant workers, or some other subgroup. Instead, by looking at the QTT, we examine how the effect of increasing the minimum wage varies by the strength of a county's local labor market. In other words, we ask the question: What is the distribution of unemployment rates across counties following a minimum wage increase relative to what it would have been if the minimum wage had not been increased? This goal is also different from trying to understand the effect of minimum wage increases at different parts of the individual income distribution as in Dube (2017).

Unlike most of the literature on minimum wages, instead of using a long panel of counties, states, and many changes in minimum wage policy across states; we focus on a particular period where the federal minimum wage was flat while there was variation in state minimum wages. The U.S. federal minimum wage increased from \$4.25 to \$5.15 between 1996 and 1997. It did not increase again until the Fair Minimum Wage Act was proposed on January 5, 2007 and enacted on May 25, 2007. The Act increased the federal minimum wage to \$5.85 on July 24, 2007 and increased the minimum wage in two more increments, settling at \$7.25 in July of 2009.

In 2006, there were 33 states for whom the federal minimum wage was the binding minimum wage in the state. Of these, we drop two states – New Hampshire and Pennsylvania – because they are located in the Northern census region; census region is an important control in the minimum wage literature (Dube, Lester, and Reich (2010)) and almost all states in the Northern census region had minimum wages higher than the federal minimum wage by 2006. Of the remaining states, 11 increased their minimum wage by the first quarter of 2007 – these states form our treated group.¹⁶ The other 20 did not increase their minimum wage until the federal minimum wage increased in July of 2007.¹⁷

¹⁶The amount that these states increased their minimum wage does vary across states, but we lump them all into the same category and ignore heterogeneity with respect to the amount that they increased the minimum wage. Among these states, the new average minimum wage was \$6.41 (roughly a 25% increase in the minimum wage). The largest increase was in Michigan which increased its state minimum wage to \$6.95, and the smallest increase was in West Virginia which increased its minimum wage to \$5.85.

¹⁷The states that increased their minimum wage were: Arizona, Arkansas, Colorado, Maryland, Michigan,

Table 1: Summary Statistics

	Treated Counties	Untreated Counties	Diff	P-val on Diff
Unem. Rate 2007	6.10	5.07	1.028	0.00
Unem. Rate 2006	6.25	5.34	0.904	0.00
Unem. Rate 2005	7.09	6.10	0.984	0.00
South	0.37	0.64	-0.274	0.00
North-Central	0.42	0.28	0.135	0.00
West	0.21	0.07	0.140	0.00
Log Med. Inc.	10.35	10.32	0.033	0.00
Log Pop.	10.34	9.91	0.437	0.00

Notes: Summary statistics for counties by whether or not their minimum wage increased in Q1 of 2007 (treated) or not (untreated). Unemployment rates are calculated using February unemployment and labor force estimates from the Local Area Unemployment Database. Median income is the county's median income from 1997 and comes from the 2000 County Data Book. Population is the county's population in 2000 and comes from the 2000 County Data Book.

Sources: Local Area Unemployment Statistics Database from the BLS and 2000 County Data Book

County level unemployment rates are the outcome variable. We obtain these from the Local Area Unemployment Statistics Database from the Bureau of Labor Statistics. Unemployment rates are available monthly and we use unemployment rates in February as the outcome variable. We choose February instead of January because it does not overlap with the holidays and choose it over later months because it is further away from the federal minimum wage change in July. We also merge in county characteristics from the 2000 County Data Book. In our application, these include 2000 county population and 1997 county median income. We collected data for each year from 2000-2007. Our method requires three periods of data, but the earlier periods allow us to pre-test our model in earlier periods.

Table 1 provides summary statistics. From 2005-2007, the level of unemployment rates is higher for treated counties than for untreated counties. The gap narrows from 2005 to

Missouri, Montana, Nevada, North Carolina, Ohio, and West Virginia. The states that did not increase their minimum wage were: Alabama, Georgia, Idaho, Indiana, Iowa, Kansas, Kentucky, Louisiana, Mississippi, Nebraska, New Mexico, North Dakota, Oklahoma, South Carolina, South Dakota, Tennessee, Texas, Utah, Virginia, and Wyoming.

2006, the period before any counties have increased minimum wages, and then expands again from 2006 to 2007; this may provide some suggestive evidence that the minimum wage is increasing unemployment rates on average. Counties that are treated are also different from untreated counties in terms of their observable characteristics. Treated counties are more likely to be in the West and North-Central regions while untreated counties are more likely to be in the South. Median incomes are very similar (though statistically different) across treated and untreated counties. And treated counties tend to be more populated; log population of 10.34 for treated counties is almost 31,000 while log population of 9.91 for untreated counties is just over 20,000.


The main results from using our method are presented in Figure 1. The upper panel provides estimates without conditioning on covariates. The lower panel provides estimates that condition on county characteristics; the specification for the propensity score interacts region with quadratic terms in log median income and log population as well as their interaction. The results are very similar whether or not covariates are included.

On average, we find that increasing the minimum wage has a small positive effect on the unemployment rate. Both with and without covariates, we estimate that increasing the minimum wage increases the unemployment rate by 0.12 percentage points. Without covariates, the effect is statistically significant. With covariates, the effect is not statistically significant. However, there is much heterogeneity. At the low end of the unemployment rate distribution, the effect of increasing the minimum wage on the unemployment rate appears to be negative. For example, at the 10th percentile, the unemployment rate is estimated to be 0.44 (p-value: 0.000) percentage points lower following the minimum wage increase than it would have been without the minimum wage increase (with covariates the estimate is 0.45 (p-value: 0.008)). However, in the middle and upper parts of the unemployment rate distribution, increasing the minimum wage appears to increase unemployment. The difference between the medians of unemployment rates in the presence or absence of the minimum wage increase is 0.31 (p-value: 0.000) percentage points (with covariates the estimate is 0.32 (p-value: 0.029)). The estimated difference between the 90th percentiles is 0.36 (p-value: 0.029) percentage points (with covariates the estimate is 0.27 (p-value: 0.216)).

For comparison, Figure 2 plots bounds on the QTT when no assumption is made about the copula between the change in untreated potential outcomes and the initial level of untreated potential outcomes for the treated group as in Fan and Yu (2012). These bounds are very wide – they cover 0 at all values of τ – and they do not include additional sampling uncertainty. For example, the difference between the median unemployment rate for treated counties and their counterfactual unemployment rate is bounded between -1.01 and 1.41.

Neither our Distributional Difference in Differences Assumption nor the Copula Stability


Figure 1: QTT estimates of the effect of increasing the minimum wage on county-level unemployment rates



Notes: The top panel provides estimates of the QTT using the no-covariates version of the method proposed in the current paper. The lower panel provides QTT estimates when the DDID assumption holds only after conditioning on covariates. 95% pointwise confidence intervals are computed using the bootstrap with 1000 iterations.

Sources: Local Area Unemployment Statistics Database from the BLS and 2000 County Data Book

Figure 2: Bounds for QTT with unknown copula




Notes: The figure shows bounds on QTTs when the copula between the change in untreated potential outcomes and the initial level of untreated potential outcomes for the treated group is treated as being completely unknown. The results are obtained using the authors' implementation of the method in Fan and Yu (2012). The figure displays point estimates of the bounds and does not include standard errors or any uncertainty due to sampling.

Sources: Local Area Unemployment Statistics Database from the BLS

Assumptions are directly testable, but, like existing Difference in Differences methods, our assumptions can be pre-tested when additional pre-treatment periods are available. The simplest way to implement a pre-test is to estimate the model in the period (or periods) before treatment and test that the QTT is 0 for all values of τ . Also, because our Copula Stability Assumption is new, we provide an additional test for only the Copula Stability Assumption. The idea of this test is to compute Kendall's Tau (a standard dependence measure that depends only on the copula (see Nelsen (2007))) in each pre-treatment year and test whether or not it changes over time. We perform both of these tests on the minimum

Figure 3: Kendall's Tau estimates for treated counties by year



Notes: The figure contains estimates of Kendall's Tau for states that increased their minimum wages in the first quarter of 2007. 95% pointwise confidence intervals are computed using the empirical bootstrap with 1000 iterations.

Sources: Local Area Unemployment Statistics Database from the BLS

wage data next.

Figure 3 plots Kendall's Tau for the change in unemployment rates and the initial level of unemployment rates for treated counties from 2001 to 2006. Kendall's Tau varies very little over this period and is always somewhat less than 0 indicating slight negative dependence between the change and initial level of unemployment. A Wald test fails to reject the equality of Kendall's Tau in all periods (p-value: 0.524). This provides suggestive empirical evidence in favor of the Copula Stability Assumption in this application. Second, we compute QTTs in each pre-treatment period from 2002 to 2006. In these periods, the QTTs should be equal to 0 everywhere. These are available in Supplementary Appendix Figure 2 and our method tends to perform very well in the earlier periods. Finally, as an additional robustness check,

we compute QTTs using the Change in Changes method with and without covariates and with the Quantile Difference in Differences method (these are available in Supplementary Appendix Figure 1). These other methods show very similar patterns as our main results.

Taken together, these results suggest that there is a great deal of heterogeneity of the effect of increasing the minimum wage across local labor markets. If we impose the additional assumption that counties maintain their rank in the distribution of unemployment when the minimum wage increases, the results indicate that counties with tight labor markets experience decreases in unemployment while counties with high unemployment see fairly large increases in unemployment. Even in the absence of such an assumption, our results indicate that increasing the minimum wage can have negative consequences for some local labor markets although the average effect may be fairly small.

7 Conclusion

This paper has considered identification and estimation of the QTT under a distributional extension of the most common Mean Difference in Differences Assumption used to identify the ATT. Even under this Distributional Difference in Differences Assumption, the QTT is still only partially identified because it depends on the unknown dependence between the change in untreated potential outcomes and the initial level of untreated potential outcomes for the treated group. We introduced the Copula Stability Assumption which says that the missing dependence is constant over time. Under this assumption and when panel data is available, the QTT is point identified. We show that the Copula Stability Assumption is likely to hold in exactly the type of models that are typically estimated using Difference in Differences techniques under mild additional conditions. This idea of a time invariant copula may also be valuable in other areas of microeconomic research especially when a researcher has access to panel data.

In many applications it is important to invoke identifying assumptions that hold only after conditioning on some covariates. We developed simple estimators of the QTT using propensity score re-weighting. In an application on the effect of minimum wage increases on local unemployment rates, we found that increasing the minimum wage tended to widen the distribution of local unemployment rates. Using pre-treatment periods, we also found suggestive empirical evidence in favor of the Copula Stability Assumption.

A Identification and Estimation under a Conditional Copula Stability Assumption

Some of our main results dealt with the case where the Distributional Difference in Differences Assumption held conditional on covariates, but the Copula Stability Assumption held unconditionally. We showed that this combination of assumptions is likely to hold in the most common type of model where empirical researchers use Difference in Differences to identify the ATT. We also provided some empirical evidence in favor of the Unconditional Copula Stability Assumption.

However, in some applications, a researcher may wish to make the Copula Stability Assumption hold after conditioning on covariates. This assumption says that the copula between the change in untreated potential outcomes and the initial level of untreated potential outcomes for the treated group does not change over time after conditioning on some covariates X .

Conditional Copula Stability Assumption.

$$C_{\Delta Y_{0t}, Y_{0t-1}|X, D=1}(\cdot, \cdot|x) = C_{\Delta Y_{0t-1}, Y_{0t-2}|X, D=1}(\cdot, \cdot|x)$$

Importantly, the QTT continues to be identified under the Conditional Copula Stability Assumption.

Proposition 1. *Assume that, for all $x \in \mathcal{X}$, ΔY_t for the untreated group, ΔY_{t-1} , Y_{t-1} , and Y_{t-2} for the treated group are continuously distributed conditional on x . Under the Conditional Distributional Difference in Differences Assumption, the Conditional Copula Stability Assumption, and Assumptions 3.2 and 3.3*

$$\begin{aligned} & P(Y_{0t} \leq y|X = x, D = 1) \\ &= E \left[\mathbb{1} \left\{ F_{\Delta Y_{0t}|X, D=0}^{-1} \left(F_{\Delta Y_{0t-1}|X, D=1}(\Delta Y_{0t-1}|x) \right) \right. \right. \\ & \quad \left. \left. \leq y - F_{Y_{0t-1}|X, D=1}^{-1} \left(F_{Y_{0t-2}|X, D=1}(Y_{0t-2}|x) \right) \right\} | X = x, D = 1 \right] \end{aligned}$$

and

$$\text{QTT}(\tau; x) = F_{Y_{1t}|X, D=1}^{-1}(\tau|x) - F_{Y_{0t}|X, D=1}^{-1}(\tau|x)$$

which is identified, and

$$P(Y_{0t} \leq y|D = 1) = \int_{\mathcal{X}} P(Y_{0t} \leq y|X = x, D = 1) dF(x|D = 1)$$

and

$$\text{QTT}(\tau) = F_{Y_{1t}|D=1}^{-1}(\tau) - F_{Y_{0t}|D=1}^{-1}(\tau)$$

which is identified.

There are several advantages to this approach. First, under the Conditional Copula Stability Assumption, the path of untreated potential outcomes can depend on the covariates. This could be important in applications where the return to some covariate – for example, the return to education – changes over time. Conditional Difference in Differences assumptions for the ATT (Heckman, Ichimura, Smith, and Todd (1998) and Abadie (2005)) allow for this pattern. Second, under the Conditional Copula Stability Assumption, it is possible to allow for time varying covariates; however, the effect of time varying covariates must be a location-shift. Finally, under the Conditional Copula Stability Assumption, one can obtain estimates of conditional quantile treatment effects.

On the other hand, there are some costs associated with the Conditional Copula Stability Assumption. Primarily, estimation becomes potentially much more challenging. Nonparametric estimation would require estimating five conditional distribution functions and conditional quantile functions which is likely to be quite challenging in practice. One could replace nonparametric estimation by assuming a parametric model for each conditional quantile function though parametric assumptions are unattractive in our model because it is not clear how misspecification in any of the first step conditional distribution/quantile functions would affect our estimates of the QTT.

In subsequent work (Callaway, Li, and Oka (2018)), we considered a conditional copula assumption in a related model. Those results are likely to go through with minor adaptations to the current model. Melly and Santangelo (2015) use parametric quantile regressions to estimate a conditional version of the Change in Changes model (Athey and Imbens (2006)); Wuthrich (2015) uses a similar approach to estimate quantile treatment effects with endogeneity. One could also adapt those types of results to our setup in a straightforward way.

B Proofs

B.1 Identification

B.1.1 Identification without covariates

In this section, we prove Theorem 1. Namely, we show that the counterfactual distribution of untreated outcome $F_{Y_{0t}|D=1}(y)$ is identified. First, we state two well known results without proof used below that come directly from Sklar’s Theorem.

Lemma B.1. *For two continuously distributed random variables X and Y , their joint density in terms of the copula pdf is given by*

$$f_{X,Y}(x, y) = c_{X,Y}(F_X(x), F_Y(y))f_X(x)f_Y(y)$$

Lemma B.2. *For two continuously distributed random variables X and Y , their copula pdf in terms of their joint density is given by*

$$c_{X,Y}(u, v) = f_{X,Y}(F_X^{-1}(u), F_Y^{-1}(v)) \frac{1}{f_X(F_X^{-1}(u))} \frac{1}{f_Y(F_Y^{-1}(v))}$$

Proof of Theorem 1. To minimize notation, let $f_t(\cdot, \cdot) = f_{\Delta Y_{0t}, Y_{0t-1}|D=1}(\cdot, \cdot)$ be the joint pdf of the change in untreated potential outcome and the initial untreated potential outcome for the treated group, and let $f_{t-1}(\cdot, \cdot) = f_{\Delta Y_{0t-1}, Y_{0t-2}|D=1}(\cdot, \cdot)$ be the joint pdf in the previous period. Similarly, let $c_t(\cdot, \cdot) = c_{\Delta Y_{0t}, Y_{0t-1}|D=1}(\cdot, \cdot)$ and $c_{t-1}(\cdot, \cdot) = c_{\Delta Y_{0t-1}, Y_{0t-2}|D=1}(\cdot, \cdot)$ be the copula pdfs for the change in untreated potential outcomes and initial level of untreated outcomes for the treated group at period t and $t - 1$, respectively. And, finally, let $\Delta \mathcal{Y} = \Delta \mathcal{Y}_{0t|D=1}$ (the support of the change in untreated potential outcomes for the treated group) and $\mathcal{Y} = \mathcal{Y}_{t-1|D=1}$ (the support of outcomes for the treated group in period $t - 1$). Then,

$$P(Y_{0t} \leq y | D = 1)$$

$$= P(\Delta Y_{0t} + Y_{0t-1} \leq y | D = 1)$$

$$= E[\mathbb{1}\{\Delta Y_{0t} \leq y - Y_{0t-1}\} | D = 1]$$

$$= \int_{\mathcal{Y}} \int_{\Delta \mathcal{Y}} \mathbb{1}\{\delta \leq y - y'\} f_t(\delta, y') \, d\delta \, dy'$$

$$= \int_{\mathcal{Y}} \int_{\Delta \mathcal{Y}} \mathbb{1}\{\delta \leq y - y'\} c_t(F_{\Delta Y_{0t}|D=1}(\delta), F_{Y_{0t-1}|D=1}(y')) f_{\Delta Y_{0t}|D=1}(\delta) f_{Y_{0t-1}|D=1}(y') \, d\delta \, dy' \quad (6)$$

$$= \int_{\mathcal{Y}} \int_{\Delta \mathcal{Y}} \mathbb{1}\{\delta \leq y - y'\} c_{t-1}(F_{\Delta Y_{0t}|D=1}(\delta), F_{Y_{0t-1}|D=1}(y')) f_{\Delta Y_{0t}|D=1}(\delta) f_{Y_{0t-1}|D=1}(y') \, d\delta \, dy' \quad (7)$$

$$\begin{aligned} &= \int_{\mathcal{Y}} \int_{\Delta \mathcal{Y}} \mathbb{1}\{\delta \leq y - y'\} f_{t-1}\left(F_{\Delta Y_{0t-1}|D=1}^{-1}(F_{\Delta Y_{0t}|D=1}(\delta)), F_{Y_{0t-2}|D=1}^{-1}(F_{Y_{0t-1}|D=1}(y'))\right) \\ &\quad \times \frac{f_{\Delta Y_{0t}|D=1}(\delta)}{f_{\Delta Y_{0t-1}|D=1}(F_{\Delta Y_{0t-1}|D=1}^{-1}(F_{\Delta Y_{0t}|D=1}(\delta)))} \times \frac{f_{Y_{0t-1}|D=1}(y')}{f_{Y_{0t-2}|D=1}(F_{Y_{0t-2}|D=1}^{-1}(F_{Y_{0t-1}|D=1}(y')))} \, d\delta \, dy' \quad (8) \end{aligned}$$

Equation 6 rewrites the joint distribution in terms of the copula pdf using Lemma B.1; Equation 7 uses the Copula Stability Assumption; Equation 8 rewrites the copula pdf as the joint distribution (now in period $t - 1$) using Lemma B.2.

Now, make a change of variables: $u = F_{\Delta Y_{0t-1}|D=1}^{-1}(F_{\Delta Y_{0t}|D=1}(\delta))$ and $v = F_{Y_{0t-2}|D=1}^{-1}(F_{Y_{0t-1}|D=1}(y'))$. This implies the following:

1. $\delta = F_{\Delta Y_{0t}|D=1}^{-1}(F_{\Delta Y_{0t-1}|D=1}(u))$
2. $y' = F_{Y_{0t-1}|D=1}^{-1}(F_{Y_{0t-2}|D=1}(v))$
3. $d\delta = \frac{f_{\Delta Y_{0t-1}|D=1}(u)}{f_{\Delta Y_{0t}|D=1}(F_{\Delta Y_{0t}|D=1}^{-1}(F_{\Delta Y_{0t-1}|D=1}(u)))} du$

$$4. \quad dy' = \frac{f_{Y_{0t-2}|D=1}(v)}{f_{Y_{0t-1}|D=1}(F_{Y_{0t-1}|D=1}^{-1}(F_{Y_{0t-2}|D=1}(v)))} dv$$

Plugging in (1)-(4) in Equation 8 and noticing that the substitutions for $d\delta$ and dy' cancel out the fractional terms in the third and fourth lines of Equation 8 implies

$$\text{Equation 8} = \int_{\mathcal{Y}_{t-2}|D=1} \int_{\Delta\mathcal{Y}_{t-1}|D=1} \mathbb{1}\{F_{\Delta Y_{0t}|D=1}^{-1}(F_{\Delta Y_{0t-1}|D=1}(u)) \leq y - F_{Y_{0t-1}|D=1}^{-1}(F_{Y_{0t-2}|D=1}(v))\} \\ \times f_{t-1}(u, v) dudv \quad (9)$$

$$= \text{E} \left[\mathbb{1}\{F_{\Delta Y_{0t}|D=1}^{-1}(F_{\Delta Y_{0t-1}|D=1}(\Delta Y_{0t-1})) \leq y - F_{Y_{0t-1}|D=1}^{-1}(F_{Y_{0t-2}|D=1}(Y_{0t-2}))\} | D = 1 \right] \quad (10)$$

$$= \text{E} \left[\mathbb{1}\{F_{\Delta Y_{0t}|D=0}^{-1}(F_{\Delta Y_{0t-1}|D=1}(\Delta Y_{0t-1})) \leq y - F_{Y_{0t-1}|D=1}^{-1}(F_{Y_{0t-2}|D=1}(Y_{0t-2}))\} | D = 1 \right] \quad (11)$$

where Equation 9 follows from the discussion above, Equation 10 follows by the definition of expectation, and Equation 11 follows from the Distributional Difference in Differences Assumption. Equation (11) implies the result because each of the distributions of potential outcomes are directly identified by their observed counterparts. \square

B.1.2 Identification with covariates

In this section, we prove Theorem 2.

Proof. All of the results from the proof of Theorem 1 will still go through with the exception of the last step which uses the Unconditional Distributional Difference in Differences Assumption. Therefore, all that needs to be shown is that $F_{\Delta Y_{0t}|D=1} = F_{\Delta Y_{0t}|D=1}^p$ under the conditions in Theorem 2. Notice,

$$\begin{aligned} \text{P}(\Delta Y_{0t} \leq \delta | D = 1) &= \frac{\text{P}(\Delta Y_{0t} \leq \delta, D = 1)}{p} \\ &= \text{E} \left[\frac{\text{P}(\Delta Y_{0t} \leq \delta, D = 1 | X)}{p} \right] \\ &= \text{E} \left[\frac{p(X)}{p} \text{P}(\Delta Y_{0t} \leq \delta | X, D = 1) \right] \\ &= \text{E} \left[\frac{p(X)}{p} \text{P}(\Delta Y_{0t} \leq \delta | X, D = 0) \right] \quad (12) \end{aligned}$$

$$= \text{E} \left[\frac{p(X)}{p} \text{E}[(1 - D) \mathbb{1}\{\Delta Y_t \leq \delta\} | X, D = 0] \right] \quad (13)$$

$$\begin{aligned}
&= \mathbb{E} \left[\frac{p(X)}{p(1-p(X))} \mathbb{E}[(1-D)\mathbb{1}\{\Delta Y_t \leq \delta\} | X] \right] \\
&= \mathbb{E} \left[\frac{1-D}{1-p(X)} \frac{p(X)}{p} \mathbb{1}\{\Delta Y_t \leq \delta\} \right] \tag{14}
\end{aligned}$$

where Equation 12 holds by the Conditional Distributional Difference in Differences Assumption. Equation 13 holds by replacing $P(\cdot)$ with $\mathbb{E}(\mathbb{1}\{\cdot\})$ and then multiplying by $(1-D)$ which is permitted because the expectation conditions on $D=0$. Additionally, conditioning on $D=0$ allows us to replace the potential outcome ΔY_{0t} with the actual outcome ΔY_t because ΔY_t is the observed change in potential untreated outcomes for the untreated group. Finally, Equation 14 simply applies the Law of Iterated Expectations to conclude the proof. \square

B.2 Proof of the results in Example 1

For the first part, notice that $\Delta Y_{0it} = \theta_t - \theta_{t-1} + \Delta v_{it}$. This has the same distribution for the treated group and untreated group under Condition (i).

For the second part, first note that,

$$\begin{aligned}
F_{\Delta Y_{0t}|D=1}(\delta) &= P(\Delta Y_{0t} \leq \delta | D=1) \\
&= P(\Delta v_t \leq \delta - (\theta_t - \theta_{t-1}) | D=1) \\
&= P\left(\Delta v_{t-1} \leq F_{\Delta v_{t-1}|D=1}^{-1}(F_{\Delta v_t}(\delta - (\theta_t - \theta_{t-1}))) | D=1\right) \\
&= P\left(\Delta Y_{0t-1} \leq F_{\Delta v_{t-1}|D=1}^{-1}(F_{\Delta v_t}(\delta - (\theta_t - \theta_{t-1}))) + (\theta_{t-1} - \theta_{t-2}) | D=1\right) \\
&= F_{\Delta Y_{0t-1}|D=1}\left(F_{\Delta v_{t-1}|D=1}^{-1}(F_{\Delta v_t}(\delta - (\theta_t - \theta_{t-1}))) + (\theta_{t-1} - \theta_{t-2})\right)
\end{aligned}$$

where the third equality holds by Condition (ii). Similarly,

$$\begin{aligned}
F_{Y_{0t-1}|D=1}(y) &= P(Y_{0t-1} \leq y | D=1) \\
&= P(\eta + v_{t-1} \leq y - \theta_{t-1} | D=1) \\
&= P(\eta + v_{t-2} \leq F_{\eta+v_{t-2}|D=1}^{-1}(F_{\eta+v_{t-1}}(y - \theta_{t-1})) | D=1) \\
&= P(Y_{0t-2} \leq F_{\eta+v_{t-2}|D=1}^{-1}(F_{\eta+v_{t-1}}(y - \theta_{t-1})) + \theta_{t-2} | D=1) \\
&= F_{Y_{0t-2}|D=1}\left(F_{\eta+v_{t-2}|D=1}^{-1}(F_{\eta+v_{t-1}}(y - \theta_{t-1})) + \theta_{t-2}\right)
\end{aligned}$$

where the third equality holds by Condition (ii). Finally, consider

$$\begin{aligned}
C_{\Delta Y_{0t}, Y_{0t-1}|D=1}(u, v) &= P(F_{\Delta Y_{0t}|D=1}(\Delta Y_{0t}) \leq u, F_{Y_{0t-1}|D=1}(Y_{0t-1}) \leq v | D=1) \\
&= P\left(F_{\Delta Y_{0t-1}|D=1}\left(F_{\Delta v_{t-1}|D=1}^{-1}(F_{\Delta v_t}(\Delta Y_{0t} - (\theta_t - \theta_{t-1}))) + (\theta_{t-1} - \theta_{t-2})\right) \leq u, \right.
\end{aligned}$$

$$\begin{aligned}
& F_{Y_{0t-2}|D=1}(F_{\eta+v_{t-2}|D=1}^{-1}(F_{\eta+v_{t-1}}(Y_{0t-1} - \theta_{t-1})) + \theta_{t-2}) \leq v | D = 1) \\
= & P\left(F_{\Delta Y_{0t-1}|D=1}\left(F_{\Delta v_{t-1}|D=1}^{-1}(F_{\Delta v_t}(\Delta v_t)) + (\theta_{t-1} - \theta_{t-2})\right) \leq u, \right. \\
& \left. F_{Y_{0t-2}|D=1}(F_{\eta+v_{t-2}|D=1}^{-1}(F_{\eta+v_{t-1}}(\eta + v_{t-1})) + \theta_{t-2}) \leq v | D = 1\right) \\
= & P\left(F_{\Delta Y_{0t-1}|D=1}(\Delta v_{t-1} + (\theta_{t-1} - \theta_{t-2})) \leq u, F_{Y_{0t-2}|D=1}(\eta + v_{t-2} + \theta_{t-2}) \leq v | D = 1\right) \\
= & P(F_{\Delta Y_{0t-1}|D=1}(\Delta Y_{0t-1}) \leq u, F_{Y_{0t-2}|D=1}(Y_{0t-2}) \leq v | D = 1) \\
= & C_{\Delta Y_{0t-1}, Y_{0t-2}|D=1}(u, v)
\end{aligned}$$

which proves the result. The second equality holds by the first two results of this section. And the third equality substitutes for ΔY_{0t} and Y_{0t-1} . The fourth equality holds by Condition (ii).

B.3 Proof of the results in Example 2

The nonseparable model $Y_{it} = q(U_{it}, X_i, D_i) + \eta_i$ can be equivalently written in terms of potential outcomes:

$$\begin{aligned}
Y_{1it} &= q_1(U_{it}, X_i) + \eta_i \\
Y_{0it} &= q_0(U_{it}, X_i) + \eta_i
\end{aligned}$$

Unconditional Mean Difference in Differences Holds

$$\begin{aligned}
E[Y_{0t}|D = d] &= \int q_0(u, x) + \eta \, dF_{U_t, X, \eta|D=d}(u, x, \eta) \\
&= \int q_0(u, x) + \eta \, dF_{U_t} \, dF_{X, \eta|D=d}(u, x, \eta) \\
&= \int q_0(u, x) + \eta \, dF_{U_{t-1}} \, dF_{X, \eta|D=d}(u, x, \eta) \\
&= \int q_0(u, x) + \eta \, dF_{U_{t-1}, X, \eta|D=d}(u, x, \eta) \\
&= E[Y_{0t-1}|D = d]
\end{aligned}$$

which implies that for the treated group and untreated group the average change in untreated potential outcomes is 0.

Conditional Distributional Difference in Differences Holds

$$\begin{aligned}
P(\Delta Y_{0t} \leq \delta | X = x, D = 1) &= \int \mathbb{1}\{q_0(u, x) - q_0(\tilde{u}, x) \leq \delta\} \, dF_{U_t, U_{t-1}|X, D=1}(u, \tilde{u}|x) \\
&= \int \mathbb{1}\{q_0(u, x) - q_0(\tilde{u}, x) \leq \delta\} \, dF_{U_t, U_{t-1}|X, D=0}(u, \tilde{u}|x)
\end{aligned}$$

$$= \mathbb{P}(\Delta Y_{0t} \leq \delta | X = x, D = 0)$$

where the second equality holds because $(U_t, U_{t-1}) \perp\!\!\!\perp (X, D)$.

Unconditional Distributional Difference in Differences Does Not Hold

$$\begin{aligned} \mathbb{P}(\Delta Y_{0t} \leq \delta | D = 1) &= \mathbb{E}[\mathbb{P}(\Delta Y_{0t} \leq \delta | X, D = 1) | D = 1] \\ &= \mathbb{E}[\mathbb{P}(\Delta Y_{0t} \leq \delta | X, D = 0) | D = 1] \end{aligned}$$

where the second equality holds by the result for the Conditional Distributional Difference in Differences Assumption holding. The last quantity is, in general, not equal to $\mathbb{P}(\Delta Y_{0t} \leq \delta | D = 0)$ because the distribution of X can be different across the two groups.

Unconditional Copula Stability Holds

$$\begin{aligned} \mathbb{P}(\Delta Y_{0t} \leq \delta, Y_{0t-1} \leq y | D = 1) &= \mathbb{P}(q_0(U_t, X) - q_0(U_{t-1}, X) \leq \delta, q_0(U_{t-1}, X) \leq y | D = 1) \\ &= \mathbb{P}(q_0(U_{t-1}, X) - q_0(U_{t-2}, X) \leq \delta, q_0(U_{t-2}, X) \leq y | D = 1) \\ &= \mathbb{P}(\Delta Y_{0t-1} \leq \delta, Y_{0t-2} \leq y | D = 1) \end{aligned}$$

which implies that the Copula Stability Assumption holds.

C Asymptotic Normality and Inference

This appendix considers the asymptotic properties of our estimator of the QTT. We show that our estimator of the QTT converges uniformly to a Gaussian process. Our results essentially follow because empirical distribution functions converge uniformly to Gaussian processes and because we show the Hadamard differentiability of the map from distribution functions to the QTT. We also provide formal justification for using the empirical bootstrap to conduct inference as discussed in the main text. We provide similar results for the case where the Distributional DID Assumptions holds after conditioning on covariates in the Supplementary Appendix.

Before proving the main results, we state an additional assumption.

Assumption C.1. For $(s, d) \in \{t, t-1, t-2\} \times \{0, 1\}$, $\mathcal{Y}_{t|D=d} \subseteq \mathcal{Y} \subset \mathbb{R}$ and \mathcal{Y} is compact.

We denote empirical processes by

$$\hat{G}_{\Delta Y_s | D=d}(\delta) = \sqrt{n}(\hat{F}_{\Delta Y_s | D=d}(\delta) - F_{\Delta Y_s | D=d}(\delta)) \quad \text{and} \quad \hat{G}_{Y_s | D=d}(y) = \sqrt{n}(\hat{F}_{Y_s | D=d}(y) - F_{Y_s | D=d}(y))$$

for $s \in \{t, t-1, t-2\}$ and $d \in \{0, 1\}$. Next, let $\tilde{Y}_{it} = F_{\Delta Y_t | D=0}^{-1}(F_{\Delta Y_{t-1} | D=1}(\Delta Y_{it-1})) + F_{Y_{t-1} | D=1}^{-1}(F_{Y_{t-2} | D=1}(Y_{it-2}))$ with support $\mathcal{Y}_{0t|D=1}$; these are pseudo-observations if each distribution and quantile function were known. Let $\tilde{F}_{Y_{0t}|D=1}(y) = \frac{1}{n_D} \sum_{i \in \mathcal{D}} \mathbb{1}\{\tilde{Y}_{it} \leq y\}$. Then, define

$$\tilde{G}_{Y_{0t}|D=1}(y) = \sqrt{n}(\tilde{F}_{Y_{0t}|D=1}(y) - F_{Y_{0t}|D=1}(y))$$

As a first step, we establish a functional central limit theorem for the empirical processes of each of the terms used in our identification result. Let $l^\infty(S)$ denote the space of all uniformly bounded functions on the set S that are equipped with the supremum norm $\|\cdot\|_\infty$ and let $C(S)$ denote the space of all continuous functions on the set S .

Proposition 2. *Under the Distributional Difference in Differences Assumption, Copula Stability Assumption, and Assumptions 3.2, 3.3 and C.1*

$$(\hat{G}_{\Delta Y_t|D=0}, \hat{G}_{\Delta Y_{t-1}|D=1}, \tilde{G}_{Y_{0t}|D=1}, \hat{G}_{Y_t|D=1}, \hat{G}_{Y_{t-1}|D=1}, \hat{G}_{Y_{t-2}|D=1}) \rightsquigarrow (\mathbb{W}_1, \mathbb{W}_2, \mathbb{V}_0, \mathbb{V}_1, \mathbb{W}_3, \mathbb{W}_4)$$

in the space $\mathcal{S} = l^\infty(\Delta \mathcal{Y}_{t|D=0}) \times l^\infty(\Delta \mathcal{Y}_{t-1|D=1}) \times l^\infty(\mathcal{Y}_{0t|D=1}) \times l^\infty(\mathcal{Y}_{t|D=1}) \times l^\infty(\mathcal{Y}_{t-1|D=1}) \times l^\infty(\mathcal{Y}_{t-2|D=1})$ where $(\mathbb{W}_1, \mathbb{W}_2, \mathbb{V}_0, \mathbb{V}_1, \mathbb{W}_3, \mathbb{W}_4)$ is a tight Gaussian process with mean 0 and block diagonal covariance function $V(y, y') = \text{diag}(V_1(y, y'), V_2(y, y'))$ where

$$V_1(y, y') = (F_{\Delta Y_t|D=0}(y_1 \wedge y'_1) - F_{\Delta Y_t|D=0}(y_1)F_{\Delta Y_t|D=0}(y'_1)) / (1 - p)$$

and

$$V_2(y, y') = \text{E}[\psi(y)\psi(y)']$$

for $y = (y_1, y_2, y_3, y_4, y_5, y_6) \in \mathcal{S}$ and $y' = (y'_1, y'_2, y'_3, y'_4, y'_5, y'_6) \in \mathcal{S}$ and where

$$\psi(y) = 1/\sqrt{p} \begin{pmatrix} \mathbb{1}\{\Delta Y_{t-1} \leq y_2\} - F_{\Delta Y_{t-1}|D=1}(y_2) \\ \mathbb{1}\{\tilde{Y}_t \leq y_3\} - F_{Y_{0t}|D=1}(y_3) \\ \mathbb{1}\{Y_t \leq y_4\} - F_{Y_t|D=1}(y_4) \\ \mathbb{1}\{Y_{t-1} \leq y_5\} - F_{Y_{t-1}|D=1}(y_5) \\ \mathbb{1}\{Y_{t-2} \leq y_6\} - F_{Y_{t-2}|D=1}(y_6) \end{pmatrix}$$

Proof of Proposition 2. The result follows immediately from the functional central limit theorem for empirical distribution functions (see, for example, van der Vaart and Wellner (1996)). \square

Next, we work to establish the joint limiting distribution of observed treated outcomes and counterfactual untreated potential outcomes for the treated group. The key step in showing this result is to establish the Hadamard Differentiability of the counterfactual distribution of untreated potential outcomes for the treated group. Before stating the main result, we provide several helpful lemmas.

Let $F_0 = (F_{10}, F_{20}, F_{30}, F_{40})$ where F_{j0} , for $j = 1, \dots, 4$, are distribution functions; we assume that each F_{j0} has compact support $\mathcal{U}_j \subset \mathbb{R}$. We also suppose that each F_{j0} has a density function f_{j0} that is uniformly bounded away from 0 and ∞ on its supports. Let (U_2, U_4) be two random variables on $\mathcal{U}_2 \times \mathcal{U}_4$ with joint distribution F_{U_2, U_4} . We assume that $U_2 \sim F_{20}$ and that $U_4 \sim F_{40}$ and that the conditional distribution $F_{U_2|U_4}$ has a continuous density function $f_{U_2|U_4}$ that is uniformly bounded from 0 and ∞ .

For $F = (F_1, F_2, F_3, F_4) \in C(U_1) \times C(U_2) \times C(U_3) \times C(U_4)$, let

$$\phi_n(F) = \frac{1}{n_D} \sum_{i \in \mathcal{D}} \mathbb{1}\{F_1^{-1}(F_2(U_{i2})) \leq y - F_3^{-1}(F_4(U_{i4}))\}$$

and

$$\phi_0(F) = \mathbb{E} \left[\mathbb{1} \{ F_1^{-1}(F_2(U_2)) \leq y - F_3^{-1}(F_4(U_4)) \} \middle| D = 1 \right]$$

both taking values in $l^\infty(\mathcal{Y})$ and let $\hat{F} = (\hat{F}_1, \hat{F}_2, \hat{F}_3, \hat{F}_4)$ denote a vector of consistent estimators of F_0 .

As a first step, we establish the Hadamard Differentiability of $\phi_0(F)$. We do this in several steps. First, we use the following result due to Callaway, Li, and Oka (2018)¹⁸

Lemma C.1. *Let $\mathbb{D} = C(\mathcal{U}_3) \times C(\mathcal{U}_4)$ and define the map $\Psi : \mathbb{D}_\Psi \subset \mathbb{D} \mapsto l^\infty(\mathcal{U}_4)$ as*

$$\Psi(F) := F_3^{-1} \circ F_4$$

where $\mathbb{D}_\Psi := \mathbb{E} \times \mathbb{E}$ where \mathbb{E} is the set of all distribution functions with strictly positive, bounded densities. Then, the map Ψ is Hadamard Differentiable at (F_{30}, F_{40}) tangentially to \mathbb{D} with derivative at (F_{30}, F_{40}) in $\psi := (\psi_1, \psi_2) \in \mathbb{D}$

$$\Psi'_{(F_{30}, F_{40})}(\psi) = \frac{\psi_2 - \psi_1 \circ F_{30}^{-1} \circ F_{40}}{f_{30} \circ F_{30}^{-1} \circ F_{40}}$$

Lemma C.2. *Let $\mathbb{A} = C(\mathcal{U}_1) \times C(\mathcal{U}_4)$. Define the map $\Lambda : \mathbb{A}_\Lambda \subset \mathbb{A} \mapsto l^\infty(\mathcal{Y}\mathcal{U}_4)$, given by*

$$\Lambda(\Gamma)(y) = \Gamma_1(y - \Gamma_2)$$

for $\Gamma = (\Gamma_1, \Gamma_2) \in \mathbb{A}$. Then, the map Λ is Hadamard differentiable at $(F_{10}, F_{30}^{-1} \circ F_{40})$ tangentially to \mathbb{A} with derivative in $\alpha := (\alpha_1, \alpha_2) \in \mathbb{A}$ given by

$$\Lambda'_{(F_{10}, F_{30}^{-1} \circ F_{40})}(\alpha)(y) = \alpha_1(y - F_{30}^{-1} \circ F_{40}) - f_{10}(y - F_{30}^{-1} \circ F_{40} \circ \alpha_2)$$

Proof. Let $\Lambda_1 : \mathbb{A}_\Lambda \mapsto \mathbb{A}_{\Lambda_2} \subset C(\mathcal{U}_1) \times C(\mathcal{Y}\mathcal{U}_4)$ be given by $\Lambda_1(\Xi) = (\Xi_1, \cdot - \Xi_2)$. Lemma 3.9.25 of van der Vaart and Wellner (1996) implies that the map Λ_1 is Hadamard differentiable at Ξ tangentially to $C(\mathcal{U}_1) \times C(\mathcal{Y}\mathcal{U}_4)$ with derivative in $\xi = (\xi_1, \xi_2) \in C(\mathcal{U}_1) \times C(\mathcal{Y}\mathcal{U}_4)$ given by

$$\Lambda'_{1, \Xi}(\xi) = (\xi_1, -\xi_2)$$

Let $\Lambda_2 : \mathbb{A}_{\Lambda_2} \mapsto l^\infty(\mathcal{Y}\mathcal{U}_4)$ be given by $\Lambda_2(\Upsilon) = \Upsilon_1 \circ \Upsilon_2$. Lemma 3.9.27 of van der Vaart and Wellner (1996) implies that Λ_2 is Hadamard differentiable at Υ tangentially to \mathbb{A} with derivative at Υ in $v = (v_1, v_2) \in \mathbb{A}$ given by

$$\Lambda'_{2, \Upsilon}(v) = v_1 \circ \Upsilon_2 + \Upsilon'_{1, \Upsilon_2} \circ v_2$$

By the chain rule for Hadamard differentiable maps

$$\begin{aligned} \Lambda'_{(F_{10}, y - F_{30}^{-1} \circ F_{40})}(\alpha) &= \Lambda'_{2, (F_{10}, y - F_{30}^{-1} \circ F_{40})} \circ \Lambda'_{1, (F_{10}, y - F_{30}^{-1} \circ F_{40})}(\alpha) \\ &= \Lambda'_{2, (F_{10}, y - F_{30}^{-1} \circ F_{40})} \circ (\alpha_1, -\alpha_2) \end{aligned}$$

¹⁸We use the notation “ \circ ” to indicate the composition of functions, e.g. $f \circ g \circ h(x) = f(g(h(x)))$.

$$= \alpha_1(y - F_{30}^{-1} \circ F_{40}) - f_{10}(y - F_{30}^{-1} \circ F_{40} \circ \alpha_2)$$

which is the result. \square

Lemma C.3. Let $\mathbb{B} = C(\mathcal{U}_2) \times C(\mathcal{Y}\mathcal{U}_4)$. Define the map $\Phi : \mathbb{B}_\Phi \subset \mathbb{B} \mapsto l^\infty(\mathcal{Y}\mathcal{U}_4)$ given by

$$\Phi(\Omega) = \Omega_1^{-1} \circ \Omega_2$$

Then, the map Φ is Hadamard differentiable at $(F_{20}, F_{10}(\cdot - F_{30}^{-1} \circ F_{40}))$ tangentially to \mathbb{B} with derivative at $(F_{20}, F_{10}(\cdot - F_{30}^{-1} \circ F_{40}))$ in $\omega := (\omega_1, \omega_2) \in \mathbb{B}$ given by

$$\Phi'_{(F_{20}, F_{10}(\cdot - F_{30}^{-1} \circ F_{40}))}(\omega) = \frac{\omega_2 - \omega_1 \circ F_{20}^{-1} \circ F_{10} \circ (\cdot - F_{30}^{-1} \circ F_{40})}{f_{20} \circ F_{20}^{-1} \circ F_{10} \circ (\cdot - F_{30}^{-1} \circ F_{40})}$$

Proof. The proof follows by the same argument as in Lemma C.1. \square

Lemma C.4. Let $\mathbb{D} = C(\mathcal{U}_1) \times C(\mathcal{U}_2) \times C(\mathcal{U}_3) \times C(\mathcal{U}_4)$ and consider the map $\phi : \mathbb{D}_\phi \subset \mathbb{D} \mapsto l^\infty(\mathcal{Y})$ given by

$$\phi(F)(y) = P(F_1^{-1}(F_2(U_2)) + F_3^{-1}(F_4(U_4)) \leq y)$$

for $F = (F_1, F_2, F_3, F_4) \in \mathbb{D}_\phi$ where $\mathbb{D}_\phi = \mathbb{E}^4$ where \mathbb{E} is the set of all distribution functions with strictly positive and bounded densities. Then, the map ϕ is Hadamard Differentiable at F_0 tangentially to \mathbb{D} with derivative in $\gamma = (\gamma_1, \gamma_2, \gamma_3, \gamma_4) \in \mathbb{D}$ given by

$$\begin{aligned} \phi'_{F_0}(\gamma) = & \int \left\{ \gamma_1(y - F_{30}^{-1} \circ F_{40}) - f_{10} \left(y - F_{30}^{-1} \circ F_{40} \circ \frac{\gamma_4 - \gamma_3 \circ F_{30}^{-1} \circ F_{40}}{f_{30} \circ F_{30}^{-1} \circ F_{40}} \right) \right. \\ & \left. - \gamma_2 \circ F_{20}^{-1} \circ F_{10}(y - F_{30}^{-1} \circ F_{40}) \right\} \times \frac{f_{U_2|U_4}(F_{20}^{-1} \circ F_{10}(y - F_{30}^{-1} \circ F_{40}(u_4))|u_4)}{f_{20} \circ F_{20}^{-1} \circ F_{10}(y - F_{30}^{-1} \circ F_{40})} dF_{U_4}(u_4) \end{aligned}$$

Proof. First, notice that

$$\begin{aligned} \phi(F)(y) &= P(U_2 \leq F_2^{-1} \circ F_1(y - F_3^{-1} \circ F_4(U_4))) \\ &= P(U_2 \leq \Phi(F_2, \Lambda(F_1, \Psi(F_3, F_4))(U_4)(y))) \end{aligned}$$

Define the map $\pi : \mathbb{D}_\pi \mapsto l^\infty(\mathcal{Y})$ where \mathbb{D}_π is the set of all functions $F_2^{-1}(F_1(\cdot - F_3^{-1}(F_4)))$ for $(F_1, F_2^{-1}, F_3^{-1}, F_4) \in \mathbb{E} \times \mathbb{E}^- \times \mathbb{E}^- \times \mathbb{E}$ as

$$\pi(\chi)(y) = \int F_{U_2|U_4}(\chi(u_4)(y)|u_4) dF_{U_4}(u_4) \quad (15)$$

Then, for $F \in \mathbb{D}$ and $y \in \mathcal{Y}$, $\phi = \pi \circ \Phi \circ \Lambda \circ \Psi$

Using the same arguments as in Callaway, Li, and Oka (2018, Lemma A2), π is Hadamard differentiable tangentially to \mathbb{D} with derivative at $F_{20}^{-1} \circ F_{10}(y - F_{30}^{-1} \circ F_{40})$ in $\zeta \in \mathbb{D}$ given by

$$\pi'_{F_{20}^{-1} \circ F_{10}(y - F_{30}^{-1} \circ F_{40})}(\zeta)(y) = \int \zeta(u_4) f_{U_2|U_4}(F_{20}^{-1} \circ F_{10}(y - F_{30}^{-1} \circ F_{40}(u_4))|u_4) dF_{U_4}(u_4) \quad (16)$$

By the chain rule for Hadamard differentiable functions (cf. van der Vaart and Wellner (1996, Lemma 3.9.3)),

$$\phi'_{F_0}(\gamma) = \pi'_{F_{20}^{-1} \circ F_{10}(\cdot - F_{30}^{-1} \circ F_{40})} \circ \Phi'_{(F_{20}, F_{10}(\cdot - F_{30}^{-1} \circ F_{40}))}(\gamma_2, \Lambda'_{(F_{10}, F_{30}^{-1} \circ F_{40})}(\gamma_1, \Psi'_{(F_{30}, F_{40})}(\gamma_3, \gamma_4))$$

Plugging in the results from Lemmas C.1 to C.3 and Equation (16) implies the result. \square

The following proposition is a main theoretical result of the paper and establishes the joint limiting distribution of the distribution of (observed) treated potential outcomes for the treated group and counterfactual untreated potential outcomes for the treated group.

Proposition 3. *Let $\hat{G}_0(y) = \sqrt{n}(\hat{F}_{Y_{0t}|D=1}(y) - F_{Y_{0t}|D=1}(y))$ and let $\hat{G}_1(y) = \sqrt{n}(\hat{F}_{Y_{1t}|D=1}(y) - F_{Y_{1t}|D=1}(y))$. Under the Distributional Difference in Differences Assumption, Copula Stability Assumption, and Assumptions 3.2, 3.3 and C.1*

$$(\hat{G}_0, \hat{G}_1) \rightsquigarrow (\mathbb{G}_0, \mathbb{G}_1)$$

where \mathbb{G}_0 and \mathbb{G}_1 are tight Gaussian processes with mean 0 with almost surely uniformly continuous paths on the space $\mathcal{Y}_{0t|D=1} \times \mathcal{Y}_{1t|D=1}$ given by

$$\mathbb{G}_1 = \mathbb{V}_1$$

and

$$\begin{aligned} \mathbb{G}_0 = \mathbb{V}_0 + \int \left\{ \mathbb{W}_1 \circ K_2(y, v) - f_{\Delta Y_t|D=0} \left(y - F_{Y_{t-1}|D=1}^{-1} \circ F_{Y_{t-2}|D=1} \circ \frac{\mathbb{W}_4 - \mathbb{W}_3 \circ K_1(v)}{f_{Y_{t-1}|D=1} \circ K_1(v)} \right) \right. \\ \left. - \mathbb{W}_2 \circ K_3(y, v) \right\} \times \frac{f_{\Delta Y_{t-1}|Y_{t-2}, D=1}(K_3(y, v)|v)}{f_{\Delta Y_{t-1}|D=1}(K_3(y, v))} dF_{Y_{t-2}|D=1}(v) \end{aligned}$$

where

$$K_1(v) := F_{Y_{t-1}|D=1}^{-1} \circ F_{Y_{t-2}|D=1}(v)$$

$$K_2(y, v) := y - K_1(v)$$

$$K_3(y, v) := F_{\Delta Y_{t-1}|D=1}^{-1} \circ F_{\Delta Y_t|D=0}(K_2(y, v))$$

Here, \mathbb{V}_0 is the variance that would obtain for estimating the counterfactual distribution of untreated potential outcomes for the treated group if each distribution and quantile function were known. The second term comes from having to estimate each of these distribution and quantile functions in a first step.

Proof of Proposition 3. First, notice that, uniformly in y

$$\begin{aligned}
\sqrt{n}(\hat{F}_{Y_{0t}|D=1}(y) - F_{Y_{0t}|D=1}(y)) &= \sqrt{n}(\phi_n(\hat{F}) - \phi_0(F_0)) \\
&= \sqrt{n}\left((\phi_n(\hat{F}) - \phi_0(\hat{F})) - (\phi_n(F_0) - \phi_0(F_0))\right) \\
&\quad + \sqrt{n}(\phi_n(F_0) - \phi_0(F_0)) - \sqrt{n}(\phi_0(\hat{F}) - \phi_0(F_0)) \\
&= \sqrt{n}(\phi_n(F_0) - \phi_0(F_0)) - \phi'_{F_0}\sqrt{n}(\hat{F} - F_0) + o_p(1)
\end{aligned}$$

where the last equality holds by a stochastic equicontinuity argument similar to those in van der Vaart and Wellner (2007) or Ferreira, Firpo, and Galvao (2017) and by Lemma C.4. Then, the result holds by Proposition 2 and an application of the functional central limit theorem. \square

With Proposition 3 in hand, our main result for the QTT follows straightforwardly by the Hadamard Differentiability of quantiles. We next state as a theorem the limiting process for our estimator of the QTT.

Theorem 3. *Suppose $F_{Y_{0t}|D=1}$ admits a positive continuous density $f_{Y_{0t}|D=1}$ on an interval $[a, b]$ containing an ε -enlargement of the set $\{F_{Y_{0t}|D=1}^{-1}(\tau) : \tau \in \mathcal{T}\}$. Under the Distributional Difference in Differences Assumption, the Copula Stability Assumption, and Assumptions 3.2, 3.3 and C.1*

$$\sqrt{n}(\widehat{QTT}(\tau) - QTT(\tau)) \rightsquigarrow \bar{\mathbb{G}}_1(\tau) - \bar{\mathbb{G}}_0(\tau)$$

where $(\bar{\mathbb{G}}_0(\tau), \bar{\mathbb{G}}_1(\tau))$ is a stochastic process in the metric space $(l^\infty(\mathcal{T}))^2$ with

$$\bar{\mathbb{G}}_0(\tau) = \frac{\mathbb{G}_0(F_{Y_{0t}|D=1}^{-1}(\tau))}{f_{Y_{0t}|D=1}(F_{Y_{0t}|D=1}^{-1}(\tau))} \quad \text{and} \quad \bar{\mathbb{G}}_1(\tau) = \frac{\mathbb{G}_1(F_{Y_t|D=1}^{-1}(\tau))}{f_{Y_t|D=1}(F_{Y_t|D=1}^{-1}(\tau))}$$

Proof of Theorem 3. Under the conditions stated in Theorem 3, the result follows from the Hadamard differentiability of the quantile map (van der Vaart and Wellner (1996, Lemma 3.9.23(ii))) and by Proposition 3. \square

Finally, for this section, we state a result on the validity of the empirical bootstrap for our procedure.

Theorem 4. *Under the Distributional Difference in Differences Assumption, Copula Stability Assumption, and Assumptions 3.2, 3.3 and C.1,*

$$\sqrt{n}\left(\widehat{QTT}^*(\tau) - \widehat{QTT}(\tau)\right) \rightsquigarrow_* \bar{\mathbb{G}}_0(\tau) - \bar{\mathbb{G}}_1(\tau)$$

where $(\bar{\mathbb{G}}_0, \bar{\mathbb{G}}_1)$ are as in Theorem 3 and \rightsquigarrow_* indicates weak convergence in probability under the bootstrap law (Gine and Zinn (1990)).

Proof. The result holds because our estimate of the QTT is Donsker and by Theorem 3.6.1 in van der Vaart and Wellner (1996). \square

References

- [1] Abadie, Alberto. “Semiparametric difference-in-differences estimators”. *The Review of Economic Studies* 72.1 (2005), pp. 1–19.
- [2] Abadie, Alberto, Joshua Angrist, and Guido Imbens. “Instrumental variables estimates of the effect of subsidized training on the quantiles of trainee earnings”. *Econometrica* 70.1 (2002), pp. 91–117.
- [3] Abrevaya, Jason and Christian Dahl. “The effects of birth inputs on birthweight: Evidence from quantile estimation on panel data”. *Journal of Business & Economic Statistics* 26.4 (2008), pp. 379–397.
- [4] Altonji, Joseph and Rosa Matzkin. “Cross section and panel data estimators for non-separable models with endogenous regressors”. *Econometrica* (2005), pp. 1053–1102.
- [5] Arellano, Manuel and Stéphane Bonhomme. “Nonlinear panel data estimation via quantile regressions”. *The Econometrics Journal* 19.3 (2016).
- [6] Athey, Susan and Guido Imbens. “Identification and inference in nonlinear difference-in-differences models”. *Econometrica* 74.2 (2006), pp. 431–497.
- [7] Barrett, Garry F and Stephen G Donald. “Statistical inference with generalized Gini indices of inequality, poverty, and welfare”. *Journal of Business & Economic Statistics* 27.1 (2009), pp. 1–17.
- [8] Bester, C Alan and Christian Hansen. “Identification of marginal effects in a nonparametric correlated random effects model”. *Journal of Business & Economic Statistics* (2012).
- [9] Bonhomme, Stéphane and Ulrich Sauder. “Recovering distributions in difference-in-differences models: A comparison of selective and comprehensive schooling”. *Review of Economics and Statistics* 93.2 (2011), pp. 479–494.
- [10] Callaway, Brantly, Tong Li, and Tatsushi Oka. “Quantile treatment effects in difference in differences models under dependence restrictions and with only two time periods”. *Journal of Econometrics* 206.2 (2018), pp. 395–413.
- [11] Callaway, Brantly and Pedro HC Sant’Anna. “Difference-in-differences with multiple time periods and an application on the minimum wage and employment”. Working Paper. 2018.
- [12] Canay, Ivan. “A simple approach to quantile regression for panel data”. *The Econometrics Journal* 14.3 (2011), pp. 368–386.
- [13] Card, David. “The impact of the Mariel boatlift on the Miami labor market”. *Industrial & Labor Relations Review* 43.2 (1990), pp. 245–257.
- [14] Card, David and Alan Krueger. “Minimum wages and employment: A case study of the fast-food industry in New Jersey and Pennsylvania”. *The American Economic Review* 84.4 (1994), p. 772.
- [15] Carneiro, Pedro, Karsten Hansen, and James Heckman. “Removing the veil of ignorance in assessing the distributional impacts of social policies”. *Swedish Economic Policy Review* 8 (2001), pp. 273–301.

- [16] Carneiro, Pedro and Sokbae Lee. “Estimating distributions of potential outcomes using local instrumental variables with an application to changes in college enrollment and wage inequality”. *Journal of Econometrics* 149.2 (2009), pp. 191–208.
- [17] Chen, Heng. “Within-group estimators for fixed effects quantile models with large N and large T”. Working Paper, Bank of Canada. 2015.
- [18] Chernozhukov, Victor, Ivan Fernandez-Val, Jinyong Hahn, and Whitney Newey. “Average and quantile effects in nonseparable panel models”. *Econometrica* 81.2 (2013), pp. 535–580.
- [19] Chernozhukov, Victor, Ivan Fernandez-Val, and Tetsuya Kaji. “Extremal quantile regression: An overview”. *Handbook of Quantile Regression*. Ed. by Koenker, Roger, Victor Chernozhukov, Xuming He, and Limin Peng. Chapman and Hall, 2018.
- [20] Chernozhukov, Victor and Christian Hansen. “An IV model of quantile treatment effects”. *Econometrica* 73.1 (2005), pp. 245–261.
- [21] Doksum, Kjell. “Empirical probability plots and statistical inference for nonlinear models in the two-sample case”. *The Annals of Statistics* (1974), pp. 267–277.
- [22] Dube, Arindrajit. “Minimum wages and the distribution of family incomes”. Working Paper. 2017.
- [23] Dube, Arindrajit, T William Lester, and Michael Reich. “Minimum wage effects across state borders: Estimates using contiguous counties”. *The Review of Economics and Statistics* 92.4 (2010), pp. 945–964.
- [24] Evdokimov, Kirill. “Identification and estimation of a nonparametric panel data model with unobserved heterogeneity”. Working Paper, Princeton University. 2010.
- [25] Fan, Yanqin and Zhengfei Yu. “Partial Identification of distributional and quantile treatment effects in difference-in-differences models”. *Economics Letters* 115.3 (2012), pp. 511–515.
- [26] Ferreira, Francisco HG, Sergio Firpo, and Antonio F Galvao. “Estimation and inference for actual and counterfactual growth incidence curves”. Working Paper. 2017.
- [27] Finkelstein, Amy and Robin McKnight. “What did Medicare do? The initial impact of Medicare on mortality and out of pocket medical spending”. *Journal of Public Economics* 92.7 (2008), pp. 1644–1668.
- [28] Firpo, Sergio. “Efficient semiparametric estimation of quantile treatment effects”. *Econometrica* 75.1 (2007), pp. 259–276.
- [29] Firpo, Sergio and Cristine Pinto. “Identification and estimation of distributional impacts of interventions using changes in inequality measures”. *Journal of Applied Econometrics* 31.3 (2016), pp. 457–486.
- [30] Freyberger, Joachim. “Nonparametric panel data models with interactive fixed effects”. *Review of Economic Studies* 85.3 (2018), pp. 1824–1851.
- [31] Frolich, Markus and Blaise Melly. “Unconditional quantile treatment effects under endogeneity”. *Journal of Business & Economic Statistics* 31.3 (2013), pp. 346–357.

- [32] Galvao, Antonio F, Carlos Lamarche, and Luiz Renato Lima. “Estimation of censored quantile regression for panel data with fixed effects”. *Journal of the American Statistical Association* 108.503 (2013), pp. 1075–1089.
- [33] Gine, Evarist and Joel Zinn. “Bootstrapping general empirical measures”. *The Annals of Probability* (1990), pp. 851–869.
- [34] Graham, Bryan and James Powell. “Identification and estimation of average partial effects in irregular correlated random coefficient panel data models”. *Econometrica* (2012), pp. 2105–2152.
- [35] Havnes, Tarjei and Magne Mogstad. “Is universal child care leveling the playing field?” *Journal of Public Economics* 127 (2015), pp. 100–114.
- [36] Heckman, James, Hidehiko Ichimura, Jeffrey Smith, and Petra Todd. “Characterizing selection bias using experimental data”. *Econometrica* 66.5 (1998), pp. 1017–1098.
- [37] Heckman, James and Richard Robb. “Alternative methods for evaluating the impact of interventions”. *Longitudinal Analysis of Labor Market Data*. Ed. by Heckman, James and Burton Singer. Cambridge: Cambridge University Press, 1985, pp. 156–246.
- [38] Heckman, James, Jeffrey Smith, and Nancy Clements. “Making the most out of programme evaluations and social experiments: Accounting for heterogeneity in programme impacts”. *The Review of Economic Studies* 64.4 (1997), pp. 487–535.
- [39] Hirano, Keisuke, Guido Imbens, and Geert Ridder. “Efficient estimation of average treatment effects using the estimated propensity score”. *Econometrica* 71.4 (2003), pp. 1161–1189.
- [40] Hoderlein, Stefan and Halbert White. “Nonparametric identification in nonseparable panel data models with generalized fixed effects”. *Journal of Econometrics* 168.2 (2012), pp. 300–314.
- [41] Joe, Harry. *Multivariate Models and Multivariate Dependence Concepts*. CRC Press, 1997.
- [42] Joe, Harry. *Dependence Modeling with Copulas*. CRC Press, Boca Raton, FL, 2015.
- [43] Jun, Sung Jae, Yoonseok Lee, and Youngki Shin. “Treatment effects with unobserved heterogeneity: A set identification approach”. *Journal of Business & Economic Statistics* 34.2 (2016), pp. 302–311.
- [44] Koenker, Roger. “Quantile regression for longitudinal data”. *Journal of Multivariate Analysis* 91.1 (2004), pp. 74–89.
- [45] Koenker, Roger. *Quantile Regression*. Cambridge University Press, 2005.
- [46] Lamarche, Carlos. “Robust penalized quantile regression estimation for panel data”. *Journal of Econometrics* 157.2 (2010), pp. 396–408.
- [47] Lehmann, Erich. *Nonparametrics: Statistical Methods Based on Ranks*. Holden-Day, San Francisco, 1974.
- [48] Li, Tong and Tatsushi Oka. “Set identification of the censored quantile regression model for short panels with fixed effects”. *Journal of Econometrics* 188.2 (2015), pp. 363–377.

- [49] Li, Tong and Quang Vuong. “Nonparametric estimation of the measurement error model using multiple indicators”. *Journal of Multivariate Analysis* 65.2 (1998), pp. 139–165.
- [50] Melly, Blaise and Giulia Santangelo. “The changes-in-changes model with covariates”. Working Paper. 2015.
- [51] Meyer, Bruce, Kip Viscusi, and David Durbin. “Workers’ compensation and injury duration: Evidence from a natural experiment”. *The American Economic Review* (1995), pp. 322–340.
- [52] Nelsen, Roger. *An Introduction to Copulas*. 2nd ed. Springer, 2007.
- [53] Neumark, David and William Wascher. “Employment effects of minimum and subminimum wages: panel data on state minimum wage laws”. *Industrial & Labor Relations Review* 46.1 (1992), pp. 55–81.
- [54] Pomeranz, Dina. “No taxation without information: Deterrence and self-enforcement in the value added tax”. *The American Economic Review* 105.8 (2015), pp. 2539–2569.
- [55] Rosen, Adam. “Set identification via quantile restrictions in short panels”. *Journal of Econometrics* 166.1 (2012), pp. 127–137.
- [56] Sen, Amartya. *On Economic Inequality*. Clarendon Press, 1997.
- [57] Sklar, Abe. *Fonctions de répartition à n dimensions et leurs marges*. Publications de L Institut de Statistique de L Universite de Paris, 1959.
- [58] van der Vaart, Aad W and Jon A Wellner. *Weak Convergence*. Springer, 1996.
- [59] van der Vaart, Aad W and Jon A Wellner. “Empirical processes indexed by estimated functions”. *Lecture Notes-Monograph Series* (2007), pp. 234–252.
- [60] Wuthrich, Kaspar. “Semiparametric estimation of quantile treatment effects with endogeneity”. Working Paper. 2015.