

Quantile Treatment Effects in R: The qte Package

Brantly Callaway

2019-06-09

Introduction

When evaluating the effect of a program, policymakers may be interested in more than just the average impact of the program. For example, there may be cases where the mean net impact (benefit minus cost) is negative, but the program is still implemented because of the distributional impacts of the program. One example from economics is the effect of a job training program. Consider two job training programs that both have the same mean net impact which is negative. If the first job training program tends to increase wages at the top of the wage distribution, it is likely to be rejected by policymakers. On the other hand, if the second job training program tends to increase wages at the bottom of the wage distribution, it may be approved by policymakers. Traditional methods, such as linear regression where the effect of participating in a program is constrained to be the same across all individuals are not suitable for distinguishing between these two types of programs. Neither are methods that allow for treatment effect heterogeneity but identify only the average effect of participating in a program. The distributional methods available in the R package QTE are useful in cases such as the above example where the researcher is interested in both allowing for treatment effect heterogeneity and understanding some aspect of the treatment effect heterogeneity.

The idea that participating in a program can have a different effect on even observationally identical individuals has a long history in the econometrics literature [11]. There has been much recent interest in studying treatment effect heterogeneity [12, 2]. Although the QTE might not be the first choice parameter of interest for studying treatment effect heterogeneity, it is easier to identify as it only depends on the marginal distributions rather than the joint distribution. In some cases the QTE is identical to the Quantiles of the Treatment Effect [7] when individuals maintain their rank in the treated and untreated distributions and the QTE is increasing in τ . On the other hand, one can write down simple examples where the effect of treatment is heterogeneous across agents and the QTE is constant or where the QTE is not constant but there is no treatment effect heterogeneity.

This article deals with estimating QTEs in several situations that are likely to be common in applied research. Models where treatment is randomly assigned or *as good as random* after conditioning on some covariates.¹ Methods for the case where the researcher has access to repeated cross sectional data or panel data are also available in the methods `panel.qtet` [5], `MDiD`, `QDiD`, and `CiC` [4]. It should also be mentioned that methods for estimating the QTE when identification is based on an instrumental variable [10] are not currently available in the `ci.qte` package though the author's of that method do have some code available in Stata with the command `ivqte`.

Because the QTE and the Average Treatment Effect (ATE)² (or the QTET and ATT) are both functions of the marginal distributions $F_{\tau}(y_1)$ and $F_{\tau_0}(y_0)$, estimating the QTE implies the ATE can be recovered from it. Usually, the ATE is identified under weaker conditions than the QTE; e.g. the ATE usually requires some type of mean independence assumption while the QTE requires some type of full independence assumption. As such, researchers that study the QTE are usually interested in the shape of the QTE rather than the location as the location is already available under the weaker assumptions that identify the QTE. In the job training example, the ATE is useful for thinking about whether or not participating in a job training program tends to increase wages; the QTE is useful for thinking about whether the job training program affects wages of workers at the top of the wage distribution more or less than workers at the bottom of the wage distribution.

¹This type of assumption is typically called either Selection on Observables or a Conditional Independence assumption are handled by the `ci.qte` or `ci.qtet` methods using the method developed in [9].

²The R package `Matching` is available to estimate the ATE and the ATT using matching techniques. In the `ci.qte` package, their are estimators of the ATE or ATT available as a by-product of estimating quantile treatment effects. These are typically based on re-weighting.




Figure 1: test caption

Background on treatment effects

There is a large literature on treatment effects. [13] provides a thorough introduction. The QTE package follows the most common setup. First, we are considering the case of a binary instrument. All individuals in the population either participate in the treatment or not. The set of participating individuals is called the treatment group; the set of non-participating individuals is called the untreated or control group.

There are many examples of programs that fit into this category. The example that we use throughout this article is a job training program. We are interested in whether or not participating in a job training program has an effect on subsequent wages. There are many other examples in different fields. In medicine, one might be interested in studying the effect of a drug on some health outcome. In a randomized experiment, this evaluation is relatively easy. A more challenging example would be the effect of cigarette smoking on longevity. Because an individual cannot be randomly assigned to smoke, the researcher must invoke some assumption to identify an effect of smoking. In political science, an example would be the effect of being the incumbent on an election outcome.

Notation

We use a common notation throughout the paper. Let Y_1 and Y_0 denote the potential treated outcome and potential untreated outcome for individual i (the i subscript is suppressed throughout for notational simplicity), respectively. Let D be a binary variable that takes the value 1 if an individual is treated and takes the value 0 if the individual is untreated. For each individual, only one of the potential outcomes is observed. We denote the observed outcome Y and note that it can be written as $Y = DY_1 + (1 - D)Y_0$. In some cases, we may also observe some covariates X . These covariates may affect both potential outcomes and the decision of whether or not to be treated.

The Difference in Differences models involve either repeated cross section or panel data and therefore, in those models, an additional subscript denoting time t is required. The notation is otherwise the same except, for example, Y_{1t} would denote the treated potential outcome at time t and Y_{0t-1} would denote the untreated potential outcome at time $t - 1$.

Quantiles

Figure 1: A plot of a standard normal distribution truncated from the left at zero. The lower dotted line represents the median; the upper dotted line represents the 95th percentile. The figure on the right inverts the same distribution to show the quantiles. Each dotted line once again corresponds to the median and the 95th percentile.

For a random variable Z with distribution function $F_Z(\cdot)$, the τ -th quantile, z_τ is defined as

$$z_\tau = \inf\{z : F_Z(z) \geq \tau\}$$

When Z is continuously distributed, the definition can be simplified to

$$z_\tau = F_Z^{-1}(\tau)$$

In either case, the quantile z_τ is identified whenever the distribution function $F_Z(\cdot)$ is identified.

Many quantiles are quite familiar. The 0.5-quantile is the median. The 0.25- and 0.75-quantiles are the 1st and 3rd quartile. The 0.01-, 0.02-, 0.03-... are the 1st, 2nd, 3rd, and so on percentiles. Figure 1(a) plots the distribution function of random draws from a $N(1, 1)$ distribution truncated at 0. From the distribution function, the τ -th quantile can be found by (i) finding τ on the y-axis, (ii) finding where a horizontal line from that point intersects the distribution function, and (iii) drawing a vertical line down to the y-axis. In Figure 1(a), this procedure is followed to find the median and the 95-th percentile of the distribution. In Figure 1(b), the entire distribution from Figure 1(a) is inverted to obtain all the quantiles. The dotted lines in Figure 1(b) correspond to the median and the 95-th percentile.

Quantile Treatment Effects

QTEs are defined as the difference between the quantiles (for a particular value of τ) of the treated potential outcome distribution and the untreated potential outcome distribution.




Figure 2: The figure is two cumulative distributions. The red distribution is the same truncated standard normal distribution as in Figure 1. The blue distribution is an adjusted truncated normal distribution. The lower dotted line represents the Quantile Treatment Effect for $\tau = 0.5$; the upper dotted line represents the Quantile Treatment Effect for $\tau = 0.95$. The figure on the right inverts both distribution functions and displays all the Quantile Treatment Effects for $\tau \in [0, 1]$

For $\tau \in [0, 1]$

$$QTE(\tau) = F_{Y_1}^{-1}(\tau) - F_{Y_0}^{-1}(\tau)$$

Likewise, QTETs are defined as the difference between the quantiles of (i) the distribution of treated potential outcomes for the treated group and (ii) the distribution of untreated potential outcomes for the treated group.

$$QTET(\tau) = F_{Y_1|D=1}^{-1}(\tau) - F_{Y_0|D=1}^{-1}(\tau)$$

For identification, the key issue is whether or not each of the distributions are identified. For the QTE, some assumption needs to be invoked to identify both $F_{Y_1}(y_1)$ and $F_{Y_0}(y_0)$. For the QTET, $F_{Y_1|D=1}(y_1)$ is identified directly from the data because we observe at the distribution of treated outcomes for the treated group. However, to identify $F_{Y_0|D=1}(y_0)$ will require some identifying assumption because we do not observe untreated potential outcomes for the treated group.

Figure 2 shows the QTE for two simulated distributions. The purple distribution in Figure 2(a) is the same as the distribution in Figure 1; the green distribution is an adjusted truncated normal distribution. The horizontal distance between the two distributions gives the QTE. For example, QTE(0.5) and QTE(0.95) are represented by the two dotted lines in the figure. The QTE for all values of τ is plotted in Figure 2(b).

Identification and Estimation

This section covers identification and estimation under the most commonly invoked assumptions for identifying the QTE or QTET. But first, we give a brief description of the dataset used in the examples.

Data

We use data from a job training program in the 1970s. This program provided extensive job training to workers over a 9-18 month period. The program randomly assigned some applicants to the program and randomly denied some other applicants. The dataset also includes non-experimental observations from the Panel Study of Income Dynamics (PSID). None of the individuals from the PSID participated in the job training program, but they can be combined with the treated group from the job training program to form an “observational” dataset. Then, various techniques that can be applied in non-experimental settings can be applied to this new dataset, and the results can be compared to the results obtained from using the full experimental dataset as a check on the performance of various non-experimental methods [14]. We use a subset that consists of males with three (2 pre-training and 1 post-training) years of earnings data (the outcome of interest). This subset has been considered in [6, 15, 9].

Using the QTE package

After installing the `qte` package, the package can be loaded by

```
library(qte)
```

```
## Registered S3 methods overwritten by 'ggplot2':  
## method      from  
## [.quosures  rlang  
## c.quosures  rlang  
## print.quosures rlang
```

and the job training datasets loaded by the following command

```
data(lalonde)
```

which loads the four datasets `lalonde.exp`, `lalonde.psid`, `lalonde.exp.panel`, and `lalonde.psid.panel`. The `exp` datasets contain the experimental datasets, and the `psid` datasets contain the observational datasets. The two datasets with the suffix `panel` contain exactly the same data as their respective counterparts, but the data is formatted in a panel (three observations per individual in the dataset rather than one) to facilitate estimation using the Difference in Differences methods mentioned below.

The main methods in the `ci.qte` package are `ci.qte`, `ci.qtet`, `panel.qtet`, `QDiD`, `CIc`, and `MDiD`. These are useful for estimating the QTE and QTET and are each based on specific identifying assumptions (discussed in more detail below). Each method requires the user to supply a formula for the outcome on the treatment status, a vector of names of conditioning X variables, and a data frame. The user should also specify for which values of τ the QTE or QTET should be calculated. There are other available parameters discussed in more detail in the manual. The most common uses are detailed in the rest of this section. Each method returns a QTE object. The QTE object contains the estimated QTE or QTET, ATE or ATT, and some other information such as standard errors and confidence intervals. Once a QTE object has been estimated, the user can call `summary` and `ggqte` methods to obtain results (a `plot` method is available, but we recommend the `ggqte` method for plotting which builds on the `ggplot2` package).

Randomization

When an individual is randomly assigned to treatment, then QTE is identified. Random treatment assignment is common in medical studies. It is sometimes available in the social sciences in cases such as lab

experiments, field experiments, and has sometimes been built into some programs such as job training. Randomization is sometimes called the 'gold standard' of empirical research, but though there are some questions that randomization cannot answer and randomized experiments in practice may differ from randomized experiments in theory [heckman-smith-1995; manski-1996]. For example, individuals may not comply with their random assignment. [Without perfect compliance, the random assignment to the treated or untreated group can be used as an instrument for actual treatment status. This setup can lead to identification of the ATT [imbens-angrist-1994] or a conditional QTET [abadie-angrist-imbens-2002]. This approach is not currently available in the QTE package.]

Under randomization, the following independence condition holds

$$(Y_1, Y_0) \perp D$$

Under this assumption, the marginal distributions $F_{Y_1}(y_1)$ and $F_{Y_0}(y_0)$ are both identified. The reason they are identified is that $F_{Y_1}(y_1) = F_{Y|D=1}(y_1)$ and $F_{Y_0}(y_0) = F_{Y|D=0}(y_0)$ under the above condition. And each of the two distributions $F_{Y|D=1}(y_1)$ and $F_{Y|D=0}(y_1)$ can be obtained by looking at the distribution of outcomes for the treated group and the untreated group, respectively. In other words, under randomization of treatment, the partially unobserved distribution of treated potential outcomes is the same as the fully observed distribution of realized outcomes for the treated group; and likewise, the partially unobserved distribution of untreated potential outcomes is the same as the fully observed distribution of realized outcomes for the untreated group.

Estimation

Under randomization, the QTE can be easily estimated using the `ci.qte` command with no covariates. In this section, we use the experimental portion of the job training set. The `ci.qte` method also computes an estimate of the Average Treatment Effect as a byproduct.

In the case of the job training program, the parameter being identified is the QTET. Individuals who desired to be treated were randomized into or out of the program. To see this, let D^* be a binary variable indicating whether an individual desires to be treated ($D^* = 1$) or does not desire to be treated ($D^* = 0$). Because individuals that do not desire to be treated never attempt to join the job training program, they cannot be randomized. Amongst those individuals who desire to be treated, some are assigned to the job training program ($R = 1$) and some are randomized out of the program ($R = 0$). Under this setup, $(Y_1, Y_0) \perp R|D^* = 1$. The distribution $F_{Y_1|D=1}(y_1) = F_{Y_1|D^*=1,R=1}(y_1)$ is identified using the outcomes from the group with $D^* = 1$ and $R = 1$. The distribution $F_{Y_0|D=1}(y_0) = F_{Y_0|D^*=1,R=1}(y_0) = F_{Y_0|D^*=1,R=0}(y_0)$ where the last equality holds by the assumption. This last distribution can be constructed using observations of untreated outcomes for the group that was randomized out of treatment.

The results from the `ci.qte` or `ci.qtet` command will be identical in this case, but it is important to note what the correct interpretation of the results is. The QTET can be estimated with the following command

```
jt.rand <- ci.qtet(re78 ~ treat, data=lalonde.exp, probs=seq(0.05, 0.95, 0.05),
  se=T, iters=10)
```

The `ci.qtet` command returns a QTE object

```
class(jt.rand)
```

```
## [1] "QTE"
```

and the summary command provides the results

```
summary(jt.rand)
```

```
##
## Quantile Treatment Effect:
##
## tau QTE Std. Error
```


```

## 0.05  0.00  0.00
## 0.1   0.00  0.00
## 0.15  0.00  0.00
## 0.2   0.00  0.00
## 0.25  338.65 298.15
## 0.3   846.40 372.83
## 0.35  1451.51 499.85
## 0.4   1177.72 955.75
## 0.45  1396.08 934.09
## 0.5   1123.55 626.54
## 0.55  1181.54 564.96
## 0.6   1466.51 692.59
## 0.65  2115.04 719.54
## 0.7   1795.12 913.00
## 0.75  2347.49 976.70
## 0.8   2278.12 982.33
## 0.85  2178.28 1225.28
## 0.9   3239.60 1751.00
## 0.95  3979.62 2760.70
##
## Average Treatment Effect: 1794.34
## Std. Error: 429.80

```

The QTET can also be plotted with the `ggqte` command.

```
ggqte(jt.rand)
```



Conditional Independence Assumption

A second assumption that can be used to identify the QTE is the following

$$(Y_1, Y_0) \perp D | X$$

where X is a vector of covariates. This assumption says that after conditioning on these X covariates, selection into treatment is as good as random. The Conditional Independence Assumption is sometimes also called Selection on Observables. In the example of the job training program, X may contain variables such as age, education, and other background characteristics. In practice, under the Conditional Independence Assumption, individuals with "similar" characteristics can be compared across treatment status to build potential outcome distributions. If individuals with the same covariates select into treatment based on the expected gain from treatment (e.g. if individuals that experience larger gains from job training are more likely to participate in job training than observationally similar individuals that would experience smaller gains), then this assumption is necessarily violated. Also, note that this assumption is a full independence assumption. It does identify the Average Treatment Effect, but the ATE can be identified under a weaker mean independence assumption.

When the researcher has access to panel data, lags of the outcome or lags of the other covariates could be included in X . This may make the identifying assumption more plausible in applications. In the job training example X could contain lags of earnings or unemployment histories which may be important to condition on to correctly evaluate job training programs [Ashenfelter [3]; heckman-smith-clements-1997].

The Conditional Independence Assumption implies that the distributions $F_{Y_1|X}(y_1)$ and $F_{Y_0|X}(y_0)$ are identified. This follows because $F_{Y_j|X}(y_j) = F_{Y_j|X, D=j}(y_j)$ under the Conditional Independence Assumption, and the latter quantity can be obtained by calculating the conditional distribution of observed outcomes for members of each group. Then, $F_{Y_0}(y_1)$ and $F_{Y_0}(y_0)$ can be obtained by integrating out X . For example,

$$F_{Y_1}(y_1) = E_X[F_{Y_1|X}(y_1)]$$

This distribution can be easily inverted to obtain the quantiles. Although this identification result is straightforward, implementing it may be challenging in practice. The most natural estimator is $\hat{F}_{Y_1}(y_1) = \frac{1}{n} \sum_{i=1}^n \hat{F}_{Y_1|X=x_i}(y_1)$. However, in general (for example when at least one of the elements of X is continuous), an estimator for $\hat{F}_{Y_1|X=x}(y_1)$ will be computationally difficult as it requires some smoothing procedure such as kernel or series methods. These methods have the additional difficulty of requiring the researcher to choose a bandwidth in the case of kernels or the number of terms in the case of series methods.

Instead, the QTE package implements the propensity score re-weighting approach developed in [9]. The propensity score is $p(x) = P(D=1|X=x)$, and the propensity score re-weighting approach rewrites $E_X[F_{Y_1|X}(y_1)]$ as a single unconditional moment that depends on y_1 and the propensity score. The propensity score is estimated in a first step and then a sum of weighted check functions is minimized in a second step to obtain an estimate of the QTET. This technique allows some flexibility in trading off stronger assumptions with computational complexity. For example, a parametric specification for the propensity score will be computationally easy, but will rely on distributional and functional form assumptions that may not be correct. On the other hand, nonparametric methods may be used that rely on less restrictive assumptions (though they do require some smoothness conditions on the propensity score), but they are likely to be more computationally complex.

Estimation

Estimation of QTEs under the Conditional Independence Assumption can be performed using the command `ci.qte`, and estimation of the QTET can be performed with the command `ci.qtet`. With covariates, this command estimates unconditional QTEs using the propensity score re-weighting procedure of [9]. Without modification, the command estimates in a first stage the propensity score using a logit model where all the covariates enter in levels.³

³One future extension is to allow the user the option of passing in a previously estimated propensity score


The following command estimates the QTE of the job training program using the observational data and conditioning on several available variables.

```
jt.cia <- ci.qte(re78 ~ treat,
  xformula=age + education + black + hispanic + married + nodedegree,
  data=lalonde.psid,
  probs=seq(0.05,0.95,0.05), se=T, iters=10)
summary(jt.cia)
```

```
##
## Quantile Treatment Effect:
##
## tau   QTE Std. Error
## 0.05   0.00         0.00
## 0.1    0.00       482.67
## 0.15  -2955.45    1516.04
## 0.2   -7447.74    1558.03
## 0.25  -8754.13    1598.37
## 0.3   -9632.81    1754.95
## 0.35 -10499.12    1630.75
## 0.4   -11313.14   2463.08
## 0.45 -12583.22   3129.43
## 0.5   -13397.72   3585.21
## 0.55 -13619.18   3666.03
## 0.6   -12500.73   3202.99
## 0.65 -12703.25   2973.33
## 0.7   -14180.98   2558.12
## 0.75 -16545.34   1748.21
## 0.8   -18614.16   1052.99
## 0.85 -21397.01    650.86
## 0.9   -25253.39   2143.69
## 0.95 -29856.92   6036.10
##
## Average Treatment Effect:  -13115.59
## Std. Error:                1201.25
```

and the QTE can be plotted

```
ggqte(jt.cia)
```

and next the QTET of the job training is computed under the same setup


```
jt.ciat <- ci.qtet(re78 ~ treat,
  xformula=age + education + black + hispanic + married + nodedgree,
  data=lalonde.psid,
  probs=seq(0.05,0.95,0.05), se=T, iters=10)
summary(jt.ciat)
```

```
##
## Quantile Treatment Effect:
##
## tau   QTE Std. Error
## 0.05    0.00      0.00
## 0.1     0.00     370.72
## 0.15  -218.70     962.08
## 0.2   -1034.41    1152.60
## 0.25  -2025.71    2143.52
## 0.3   -3586.78    2174.47
## 0.35  -4459.40    1822.95
## 0.4   -5067.52    1382.80
## 0.45  -4988.25    1319.41
## 0.5   -4659.23    1069.71
## 0.55  -4612.01    1092.07
## 0.6   -5144.90     978.44
## 0.65  -4715.92     877.22
## 0.7   -5150.22    1268.21
## 0.75  -5884.24    1151.09
## 0.8   -6985.32    1209.84
## 0.85  -8235.10    1474.37
```

```
## 0.9 -10575.45 2317.42
## 0.95 -10782.01 3015.20
##
## Average Treatment Effect: 9708.90
## Std. Error: 1051.30
```

Finally, the QTET is plotted

```
ggqte(jt.ciat)
```



Difference in Differences Methods

In some applications the Conditional Independence Assumption may not hold. An alternative identifying assumption that the researcher may be interested in making is the following Distributional Difference in Differences Assumption

$$\Delta Y_{0t} \perp D_t | X$$

This is the identifying assumption made in [8, 5]. This is a direct extension of the “parallel trends” assumption made in the Difference in Differences literature [1].

Note that these types of methods only identify the QTET (not the QTE). [5] show that this assumption only identifies the QTET when the researcher has access to three periods of panel data.

Estimation

Estimation using the method in [5] is similar to the preceding methods. However, the user must now additionally specify some extra parameters: (i) t , $t_{\min 1}$, and $t_{\min 2}$ are the three time periods, and (ii) the names of the columns in data containing the time period and an individual id are required to be passed in as $tname$ and $idname$. data

itself must now be a panel dataset containing individual-time period observations. For this section the `lalonde.psid` dataset (which contains one row for each individual and with two lags of earnings available) has been converted into the `lalonde.psid.panel` (which contains three rows of observations for each individual).


```
head(lalonde.psid.panel)[,c("id", "year", "re", "treat",
"age", "education", "black", "hispanic", "married", "nodegree", "u75")]

## id year re treat age education black hispanic married nodegree u75
## 1 1 1978 9930.05 1 37 11 1 0 1 1 1
## 2 2 1978 3595.89 1 22 9 0 1 0 1 1
## 3 3 1978 24909.50 1 30 12 1 0 0 0 1
## 4 4 1978 7506.15 1 27 11 1 0 0 1 1
## 5 5 1978 289.79 1 33 8 1 0 0 1 1
## 6 6 1978 4056.49 1 22 9 1 0 0 1 1
```

The [5] method can be invoked with the `panel.qtet` command:

```
jt.pqtet <- panel.qtet(re ~ treat, t=1978, tmin1=1975, tmin2=1974,
tname="year", idname="id",
xformula=age + education + black + hispanic + married + nodegree,
data=lalonde.psid.panel,
probs=seq(0.05,0.95,0.05), se=T, iters=10)
summary(jt.pqtet)
```

```
##
## Quantile Treatment Effect:
##
## tau QTE Std. Error
## 0.05 0.00 0.00
## 0.1 0.00 276.70
## 0.15 -218.70 749.00
## 0.2 -2321.31 750.64
## 0.25 -3209.09 739.00
## 0.3 -3723.19 955.08
## 0.35 -3606.03 1002.72
## 0.4 -3107.37 1050.09
## 0.45 -2199.44 1144.72
## 0.5 -1935.37 834.20
## 0.55 -1506.79 1115.08
## 0.6 -595.06 984.51
## 0.65 334.80 1000.03
## 0.7 184.35 1127.29
## 0.75 473.71 1213.02
## 0.8 948.60 945.46
## 0.85 1497.56 1173.38
## 0.9 1167.20 1464.92
## 0.95 1965.56 2722.20
##
## Average Treatment Effect: -316.57
## Std. Error: 867.55
ggqte(jt.pqtet)
```



In the most common Difference in Differences case with two periods of panel or repeated cross sectional data, [8] shows that the Distributional Difference in Differences Assumption leads to partial identification of the QTET. In this case, a user can estimate these bounds using the method bounds.⁴

```
b1 <- bounds(re ~ treat, t=1978, tmin1=1975, data=lalonde.psid.panel,
  idname="id", tname="year")
summary(b1)
```


```
##
## Bounds on the Quantile Treatment Effect on the Treated:
##
## tau Lower Bound Upper Bound
## tau Lower Bound Upper Bound
## 0.05 -51.72 0
## 0.1 -1220.84 0
## 0.15 -1881.9 0
## 0.2 -2601.32 0
## 0.25 -2916.38 485.23
## 0.3 -3080.16 943.05
## 0.35 -3327.89 1505.98
## 0.4 -3240.59 2133.59
## 0.45 -2982.51 2616.84
## 0.5 -3108.01 2566.2
## 0.55 -3342.66 2672.82
## 0.6 -3491.4 3065.7
## 0.65 -3739.74 3349.74
## 0.7 -4647.82 2992.03
## 0.75 -4826.78 3219.32
```

⁴Currently, no standard errors or confidence intervals are computed using the bounds method.

```
##      0.8      -5801.7    2702.33
##      0.85     -6588.61   2499.41
##      0.9      -8953.84   2020.84
##      0.95    -14283.61    397.04
##
## Average Treatment Effect on the Treated: 2326.51
```

and the bounds can be plotted

```
plot(b1)
```



Finally, the QTE package implements several other methods that base identification on a Difference in Differences-type assumption. Each of these methods achieve point identification when only two periods of repeated cross sectional data is available. First, the identifying assumption for Quantile Difference in Differences (implemented in the QDiD method) is

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

From this expression, we notice that $F_{Y_{0t}|D_t=1}^{-1}(\tau)$ is the term that needs to be identified for identification of the QTE. The other terms are identified because (i) no one is treated (by assumption) in the first period and (ii) untreated potential outcomes are observed for the untreated group. Invoking a conditional on covariates version of this assumption is not straightforward. We do provide a method for invoking this assumption with covariates. First, a linear model with all the terms in X is run and the residuals plus the predicted value of the group effect are recovered. Then, the unconditional version of the model is run on these terms.

Next, the identifying assumption for the Change in Changes model of [4] can be stated nonparametrically as

$$F_{Y_{0t}|D_t=1}^{-1}(\tau) - F_{Y_{0t-1}|D_t=1}^{-1}(\tau) = F_{Y_{0t}|D_t=0}(F_{Y_{0t-1}|D_t=0}(y)) - F_{Y_{0t-1}|D_t=0}(F_{Y_{0t-1}|D_t=0}(y))$$

Whereas QDiD treats time and group symmetrically, CiC treats time and group asymmetrically. The other advantage of computing QTE with this method is that it is invariant to the scale of the dependent variable. The Change in Changes model is implemented by the method CiC.

Finally, the identifying assumption for the Mean Difference in Differences model is

$$F_{Y_{0t}|D_t=1}^{-1}(\tau) - F_{Y_{0,t-1}|D_t=1}^{-1}(\tau) = E[Y_{0t} - Y_{0,t-1} | D_t = 0]$$


This assumption is restrictive in that it restricts the difference in the quantiles of the distribution of untreated potential outcomes for the treated group in periods t and $t - 1$ be the same for all values of τ .

Estimation

The QTET for participating in the job training program is estimated using each of the three methods ddid2, QDiD, and CiC next.

```
jt.ddid2 <- ddid2(re ~ treat, t=1978, tmin1=1975,
  tname="year", idname="id",
  xformula=-age + I(age^2) + education + black + hispanic + married + nodegree,
  data=lalonde.psid.panel,
  probs=seq(0.05,0.95,0.05), se=T, iters=10)
summary(jt.ddid2)
```

```
##
## Quantile Treatment Effect:
##
## tau   QTE Std. Error
## 0.05  9832.19    1264.64
## 0.1   4544.63     812.46
## 0.15  2068.77     503.28
## 0.2   849.12      364.80
## 0.25  485.23      360.72
## 0.3   943.05      371.16
## 0.35  817.00      542.71
## 0.4   681.37      689.55
## 0.45  1010.12     747.71
## 0.5   992.30      712.86
## 0.55  938.13      827.91
## 0.6   950.31      993.47
## 0.65  977.04      972.30
## 0.7   366.22     1042.61
## 0.75  563.49     1078.98
## 0.8   -352.59    1020.80
## 0.85  -825.48    1241.85
## 0.9   -1978.06   1995.85
## 0.95  -3333.16   2420.89
##
## Average Treatment Effect: 1665.63
## Std. Error: 826.59
ggqte(jt.ddid2)
```




```
jt.qdid <- QDiD(re ~ treat, t=1978, tmin1=1975,
  tname="year", idname="id",
  xformla=-age + I(age^2) + education + black + hispanic + married + nodegree,
  data=lalonde.psid.panel,
  probs=seq(0.05,0.95,0.05), se=T, iters=10, panel=T)
summary(jt.qdid)
```

```
##
## Quantile Treatment Effect:
##
## tau   QTE Std. Error
## 0.05   306.34   1139.64
## 0.1    -606.23   929.07
## 0.15   -856.17   864.60
## 0.2    -1430.94   859.51
## 0.25   -1200.28   926.79
## 0.3    -902.21   993.97
## 0.35   -500.36   1161.23
## 0.4     94.13   1238.04
## 0.45   930.04   1188.89
## 0.5   1538.92   1050.97
## 0.55  2197.63   1062.75
## 0.6   3051.61   849.21
## 0.65  3536.42   862.18
## 0.7   2994.79   860.56
## 0.75  3930.74   1266.09
## 0.8   4385.78   1283.07
## 0.85  4373.68   1747.75
## 0.9   3921.55   1646.41
```

```
## 0.95 4380.63 2824.86
##
## Average Treatment Effect: 2033.09
## Std. Error: 1022.81
```

```
ggqte(jt.qdid)
```




```
jt.cic <- CiC(re ~ treat, t=1978, tmin1=1975,
  tname="year", idname="id",
  xformula=-age + I(age^2) + education + black + hispanic + married + nodegree,
  data=lalonde.psid.panel,
  probs=seq(0.05,0.95,0.05), se=T, iters=10, panel=T)
summary(jt.cic)
```

```
##
## Quantile Treatment Effect:
##
## tau QTE Std. Error
## 0.05 0.00 0.00
## 0.1 0.00 0.00
## 0.15 0.00 0.00
## 0.2 0.00 0.00
## 0.25 485.23 292.92
## 0.3 929.88 255.64
## 0.35 1460.36 353.92
## 0.4 2321.11 681.79
## 0.45 3462.56 724.39
## 0.5 4232.31 606.60
## 0.55 5010.34 635.86
## 0.6 6210.67 833.74
```



```
## 0.65 7458.11 798.16
## 0.7 7508.97 792.62
## 0.75 8210.84 1071.29
## 0.8 8332.69 1108.25
## 0.85 8515.98 1273.78
## 0.9 8822.00 1973.07
## 0.95 8273.74 2221.75
##
## Average Treatment Effect: 4629.24
## Std. Error: 669.31
```

```
ggqte(jt.cic)
```



References

- [1] Alberto Abadie. "Semiparametric difference-in-differences estimators". In: *The Review of Economic Studies* 72.1 (2005), pp. 1–19.
- [2] Jaap Abbring and James Heckman. "Econometric evaluation of social programs, part III: Distributional treatment effects, dynamic treatment effects, dynamic discrete choice, and general equilibrium policy evaluation". In: *Handbook of Econometrics* 6 (2007), pp. 5145–5303.
- [3] Orley Ashenfelter. "Estimating the effect of training programs on earnings". In: *The Review of Economics and Statistics* (1978), pp. 47–57.
- [4] Susan Athey and Guido Imbens. "Identification and inference in nonlinear difference-in-differences models". In: *Econometrica* 74.2 (2006), pp. 431–497.
- [5] Brantly Callaway and Tong Li. "Quantile treatment effects in difference in differences models with panel data". Working Paper. 2017.

- [6] Rajeev Dehejia and Sadek Wahba. "Causal effects in nonexperimental studies: Reevaluating the evaluation of training programs". In: *Journal of the American Statistical Association* 94.448 (1999), pp. 1053–1062.
- [7] Yanqin Fan and Sang Soo Park. "Partial identification of the distribution of treatment effects and its confidence sets". In: *Advances in Econometrics* 25 (2009), pp. 3–70.
- [8] Yanqin Fan and Zhengfei Yu. "Partial Identification of distributional and quantile treatment effects in difference-in-differences models". In: *Economics Letters* 115.3 (2012), pp. 511–515.
- [9] Sergio Firpo. "Efficient semiparametric estimation of quantile treatment effects". In: *Econometrica* 75.1 (2007), pp. 259–276.
- [10] Markus Frolich and Blaise Melly. "Unconditional quantile treatment effects under endogeneity". In: *Journal of Business & Economic Statistics* 31.3 (2013), pp. 346–357.
- [11] James Heckman and Richard Robb. "Alternative methods for evaluating the impact of interventions". In: *Longitudinal Analysis of Labor Market Data*. Ed. by James Heckman and Burton Singer. Cambridge: Cambridge University Press, 1985, pp. 156–246.
- [12] James Heckman, Jeffrey Smith, and Nancy Clements. "Making the most out of programme evaluations and social experiments: Accounting for heterogeneity in programme impacts". In: *The Review of Economic Studies* 64.4 (1997), pp. 487–535.
- [13] Guido Imbens and Jeffrey Wooldridge. "Recent developments in the econometrics of program evaluation". In: *Journal of Economic Literature* 47.1 (2009), pp. 5–86.
- [14] Robert LaLonde. "Evaluating the econometric evaluations of training programs with experimental data". In: *The American Economic Review* (1986), pp. 604–620.
- [15] Jeffrey Smith and Petra Todd. "Does matching overcome LaLonde's critique of nonexperimental estimators?" In: *Journal of Econometrics* 125.1 (2005), pp. 305–353.