

Heterogeneous Effects of Job Displacement on Earnings

Afrouz Azadikhah Jahromi* Brantly Callaway†

June 13, 2019

Abstract

This paper considers how the effect of job displacement varies across different individuals. In particular, our interest centers on features of the distribution of the *individual-level* effect of job displacement. Identifying features of this distribution is particularly challenging – e.g., even if we could randomly assign workers to be displaced or not, many of the parameters that we consider would not be point identified. We exploit our access to panel data, and our approach relies on comparing outcomes of displaced workers to outcomes the same workers would have experienced if they had not been displaced and if they maintained the same rank in the distribution of earnings as they had before they were displaced. Using data from the Displaced Workers Survey, we find that displaced workers earn about \$157 per week less than they would have earned if they had not been displaced. We also find that there is substantial heterogeneity. We estimate that 42% of workers have higher earnings than they would have had if they had not been displaced and that a large fraction of workers have substantially lower earnings than the average effect of displacement. Finally, we also document major differences in the distribution of the effect of job displacement across education levels, sex, age, and counterfactual earnings levels. Throughout the paper, we rely heavily on quantile regression. First, we use quantile regression as a flexible (yet feasible) first step estimator of conditional distributions and quantile functions that our main results build on. We also use quantile regression to study how covariates affect the distribution of the individual-level effect of job displacement.

JEL Codes: J63, C21

Keywords: Job Displacement, Joint Distribution of Potential Outcomes, Distribution of the Treatment Effect, Quantile Regression, Heterogeneous Effects, Rank Invariance

*Ph.D. Candidate, Department of Economics, Temple University, Email: afrouz@temple.edu.

†Assistant Professor, Department of Economics, Temple University, Email: brantly.callaway@temple.edu

‡We thank Michael Bognanno, Bernd Fitzenberger, Weige Huang, Blaise Melly, Moritz Ritter, Jungmo Yoon, and participants at the Economic Applications of Quantile Regression 2.0 Conference for helpful comments.

1 Introduction

A typical finding in the job displacement literature is that the “effect” of job displacement is large and persistent (see, for example, Jacobson, LaLonde, and Sullivan (1993) and much subsequent work). However, considering a single effect of job displacement ignores potential heterogeneity across individuals. Treatment effect heterogeneity is a general feature of many program evaluation problems in economics, but treatment effect heterogeneity is likely to be particularly important in the context of job displacement. In particular, some individuals may be able to find similar jobs quickly after they are displaced, but other individuals may experience larger losses due to decreased match quality with their employer (e.g., Jovanovic (1979)), loss of firm-, occupation-, or industry-specific human capital (Topel (1991) and Neal (1995)), loss of a job at a high paying firm (Abowd, Kramarz, and Margolis (1999)), or failure to find full time work (Farber (2017)).¹ In the current paper, we are interested in trying to distinguish between the cases where all individuals experience roughly the same effect of job displacement or where the “effect” of job displacement mixes together individuals who have relatively small effects of job displacement with individuals who experience large negative effects of job displacement. Distinguishing between these cases is important for our understanding of job displacement as well as being useful for targeting policy responses to job displacement.

Learning about treatment effect heterogeneity requires learning about the distribution of the treatment effect itself.² And the distribution of the treatment effect itself depends, in turn, on the joint distribution of treated and untreated potential outcomes. Learning about the joint distribution of treated and untreated potential outcomes is quite challenging – e.g., even if we could randomly assign workers’ displacement status, this would only identify

¹See Carrington and Fallick (2017) for a good discussion of reasons for why earnings may decline following job displacement.

²A related (though distinct) idea is to compare the distribution of observed outcomes for displaced workers to their counterfactual distribution of outcomes if they had not been displaced. Researchers often compare the quantiles of these two distributions – this class of parameters is called quantile treatment effects. However, these types of parameters are likely to have severe limitations for understanding heterogeneous effects of job displacement which we discuss in more detail in the next section.

the marginal distributions not the joint – and has been the subject of a large literature in econometrics. This literature includes Heckman, Smith, and Clements (1997), Fan and Park (2010), Fan, Guerre, and Zhu (2017), Frandsen and Lefgren (2017), and Callaway (2018), among others. This literature has largely focused on partial identification results under relatively weak assumptions or invoking stronger assumptions to obtain point identification. Here, we consider an assumption that results in point identification. The assumption is that, in the absence of job displacement, an individual would have maintained their rank in the distribution of earnings across time periods. This sort of assumption is a kind of rank-invariance assumption though it is a rank invariance over time assumption rather than rank invariance across treatment states which is more common. One can also think of this less as an assumption and more as trying to understand the question: What is an individual’s earnings following job displacement relative to what they would have earned if they had not been displaced *and* had maintained their rank in the distribution of earnings? This interpretation, as we discuss below, does not require the aforementioned rank invariance assumption. Instead, it subtly changes the parameter of interest.

The first step to understanding heterogeneous effects of job displacement is to estimate the (conditional) distributions of treated and untreated potential outcomes. This is a non-trivial task (though one we have largely abstracted from so far because these sorts of arguments are more standard). For identification, we use the approach of Callaway, Li, and Oka (2018). The two main identifying assumptions are that the distribution of the change in untreated potential outcomes over time is independent of treatment status after conditioning on covariates (this is a Difference in Differences-type assumption) and that the conditional copula for the change in untreated potential outcomes and initial level of untreated potential outcomes is the same for the treated group and untreated group.³ What these assumptions mean in the context of job displacement is that (i) the path of earnings that displaced work-

³We note that other approaches could be used for identification here; for example, one could assume that treatment is as good as randomly assigned after conditioning on covariates and lagged outcomes or one could assume a conditional version of the Change in Changes model (Athey and Imbens (2006)) as in Melly and Santangelo (2015).

ers would have experienced had they not been displaced is the same as the path of earnings that non-displaced workers with the same characteristics did experience, and (ii) if, for example, the largest increases in earnings for non-displaced workers tended to go to individuals with the highest earnings in the previous period, the same would have occurred for displaced workers if they had not been displaced from their job. One challenge is that we cannot implement the estimation procedure proposed in Callaway, Li, and Oka (2018) because their approach is for the case with discrete covariates only and involves sample splitting. This is not feasible in the current application where there are some continuous covariates (e.g., age) and too many discrete covariates relative to the number of observations to split the sample on every possible combination of the discrete covariates. Instead, we use quantile regression (Koenker and Bassett Jr (1978) and Koenker (2005)) to estimate each required conditional distribution / quantile function. Quantile regression provides a flexible way to model the conditional distributions and quantiles that we need to estimate in the first step.⁴

After the conditional distributions of treated and untreated potential outcomes have been obtained in the first step, we use the rank invariance assumption to construct the joint distribution of treated and untreated potential outcomes. In particular, under the assumption of rank invariance, an individual's counterfactual outcome can be obtained by taking their (conditional) rank in the first time period and evaluating the conditional quantile of counterfactual outcomes (which is identified) at that particular rank. This step results in every observed treated potential outcome being matched with an untreated potential outcome. Once this step has been completed, we can obtain any feature of the distribution of the treatment effect. This also implies that the distribution of the treatment effect conditional on covariates is identified as well. In a final step, we again use quantile regression to summarize how the effect of job displacement varies with covariates at different parts of the conditional distribution.

We use the Displaced Workers Survey (DWS) which is a supplement to the Current

⁴Other papers that have used quantile regression as a first-step estimator include Melly and Santangelo (2015) and Wüthrich (2019).

Population Survey (CPS) to estimate the entire distribution of the effect of job displacement on weekly earnings for displaced workers and evaluate how the effects of job displacement vary across individuals with different observable characteristics (e.g., age, sex, education, marital status) and as a function of what a displaced worker's earnings would have been in the absence of job displacement. Our main results use the 2016 DWS. We find that, on average, displaced workers lose about \$157 per week relative to what they would have earned if they had not been displaced (this corresponds to 18% of pre-displacement earnings). We also find that there is substantial heterogeneity. We estimate that 42% of displaced workers have *higher* earnings following displacement than they would have had if they had not been displaced. On the other hand, we also estimate that 21% of workers earn at least \$500 less per week than they would have earned if they had not been displaced. In a final step, we use quantile regression to study the effect of covariates on the distribution of the individual-level effect of job displacement. We find substantial heterogeneity across observed characteristics. The largest effects of job displacement appear to be concentrated among older, male, college graduates as well as those who would have had high earnings if they had not been displaced. As an additional check, we compare our results to those obtained under the condition that a worker who is displaced from their job would have earnings that are exactly equal to their pre-displacement earnings.⁵ We also extend our results to different time periods using the 2010 DWS (corresponding to workers displaced during the Great Recession) and the 1998 DWS (corresponding to workers displaced in a substantially earlier time period when the U.S. economy was strong). We find larger effects of job displacement during the Great Recession, but, interestingly, we continue to find a substantial fraction, 28%, that have higher earnings following displacement than they would have had if they had not been displaced.

⁵An alternative way to think about this is asking the question: What is an individual's earnings following job displacement relative to what they earned before they were displaced? This sort of question does not require any identifying assumption (both of these are observed outcomes), but it may not be the research question of interest.

2 Identification

In this section, we discuss the parameters that we are interested in identifying and the assumptions required to identify them. But first, we introduce some notation. We consider the case where a researcher has access to two periods of panel data which corresponds to the available data that we use from the Displaced Workers Survey. We denote the two periods by t and $t - 1$. We define outcomes in each period by Y_t and Y_{t-1} . Individuals have treated and untreated potential outcomes, $Y_s(1)$ and $Y_s(0)$, for $s \in \{t, t - 1\}$; these correspond to the outcomes that a particular individual would experience if they were displaced or not displaced at a particular point in time. We consider the case where no one is treated in the first period and some individuals become treated in the second period.⁶ Let the variable D denote whether or not an individual is treated. Thus, we observe $Y_t = DY_t(1) + (1 - D)Y_t(0)$ and $Y_{t-1} = Y_{t-1}(0)$. We also observe some covariates X .

Our interest centers on features of the distribution of $(Y_t(1) - Y_t(0)|D = 1)$. This is the distribution of the treatment effect conditional on being part of the treated group. In other words, it is the distribution of the difference between outcomes that displaced workers experienced and what they would have experienced if they had not been displaced. However, identifying features of this distribution is challenging – primarily because, for each individual, either $Y_t(1)$ or $Y_t(0)$ is observed but $Y_t(1) - Y_t(0)$ is not observed for any individual. Parameters that depend on this distribution include many treatment effect heterogeneity-type parameters: (i) the fraction of displaced workers that have higher earnings following displacement than they would have had if they had not been displaced, and (ii) how much lower earnings are for individuals who are most affected by job displacement, among others. We are also interested in features of the distribution of $(Y_t(1) - Y_t(0)|Y_t(0), D = 1)$. These

⁶This is essentially the setup in the case of our application on job displacement due to the Displaced Workers Survey being administered only every other year. In addition, part of our identification argument utilizes a Difference in Differences-type identification strategy. This sort of identification strategy identifies effects for individuals who move from being untreated to treated over time. It does not identify effects for individuals who are treated in the first period; moreover, this group is not used for identifying counterfactual outcome distributions (see discussion below) and therefore can be dropped from the analysis.

sorts of parameters include the distribution of the treatment effect for displaced workers as a function of what a displaced worker’s earnings would have been if they had not been displaced. Finally, we are interested in the distribution of the treatment effect for displaced workers as a function of covariates; i.e., features of the distribution of $(Y_t(1) - Y_t(0) | X, D = 1)$. In particular, these parameters allow us to think about what individual-level characteristics are associated with the largest effects of job displacement. Our approach allows one to look at different parts of the conditional distribution of the effect of job displacement which is substantially more informative than looking at average effects across different values of covariates which is commonly done in applied work (see, for example, the discussion in Bitler, Gelbach, and Hoynes (2017)).

What each of these parameters have in common is that they depend on the joint distribution $(Y_t(1), Y_t(0) | X, D = 1)$. By Sklar’s Theorem (Sklar (1959)), joint distributions can be written as the copula (which captures the dependence of two random variables) of the marginal distributions. Thus, we take a two step approach to identifying this joint distribution. In the first step, we identify the marginal distributions $F_{Y_t(1) | X, D=1}$ and $F_{Y_t(0) | X, D=1}$. In the second step, we impose additional assumptions to additionally identify their joint distribution.

2.1 Step 1: Identifying Marginal Distributions

The first step in our analysis is to show that $F_{Y_t(1) | X, D=1}$ and $F_{Y_t(0) | X, D=1}$ are identified. $F_{Y_t(1) | X, D=1}$ is immediately identified from the sampling process. It is just the conditional distribution of displaced potential earnings for the group of individuals that are displaced from their job – these are the earnings that are observed for displaced workers.

On the other hand, identifying $F_{Y_t(0) | X, D=1}$ is substantially more challenging. This is the distribution of earnings that displaced workers would have experienced if they had not been displaced. We make the following two assumptions:

Assumption 1 (Distributional Difference in Differences).

$$\Delta Y_t(0) \perp\!\!\!\perp D|X$$

Assumption 2 (Copula Invariance Assumption). *Let $C_{\Delta Y_t(0), Y_{t-1}(0)|X, D=d}$ denote the copula between the change in untreated potential outcomes and the initial level of untreated potential outcomes conditional on covariates and treatment status. Then,*

$$C_{\Delta Y_t(0), Y_{t-1}(0)|X, D=1} = C_{\Delta Y_t(0), Y_{t-1}(0)|X, D=0} \quad a.s.$$

Assumptions 1 and 2 are the assumptions made in Callaway, Li, and Oka (2018) to identify the counterfactual distribution $F_{Y_t(0)|X, D=1}$. The first is a Difference in Differences assumption. It says that the distribution of the path of earnings that individuals who are displaced from their job would have experienced if they had not been displaced from their job is the same as the distribution of the path of earnings that individuals that were not displaced from their job actually experienced (after conditioning on covariates). The second assumption is a Copula Invariance assumption. It says that if, for example, we observe the largest earnings increases over time going to those at the top of the earnings distribution for the group of non-displaced workers, then we would have observed the same thing for displaced workers if they had not been displaced (again, this is conditional on covariates). Callaway, Li, and Oka (2018) show that, under these conditions

$$F_{Y_t(0)|X, D=1}(y|x) = P(\Delta Y_t + F_{Y_{t-1}|X, D=1}^{-1}(F_{Y_{t-1}|X, D=0}(Y_{t-1}|x)|x) \leq y|x, D = 0) \quad (1)$$

where every term in Equation (1) is identified. This implies that both marginal distributions of interest are identified.

2.2 Step 2: Identifying the Joint Distribution

Identifying the joint distribution of treated and untreated potential outcomes is a notoriously difficult challenge in the microeconometrics literature. Standard identifying assumptions (or even random assignment of the treatment) do not identify this joint distribution. This section describes our approach to identifying this distribution. Define

$$\tilde{Y}_t(0) = Q_{Y_t(0)|X, D=1}(F_{Y_{t-1}|X, D=1}(Y_{t-1}|X)|X) \quad (2)$$

Notice that $F_{Y_{t-1}|X, D=1}(Y_{t-1}|X)$ is a displaced worker's rank in the conditional earnings distribution in the period before anyone is displaced. Thus, conditional on our identification arguments in the previous section for the counterfactual distribution, the transformation on the right hand side of Equation (2) is available and $\tilde{Y}_t(0)$ is the outcome that a displaced worker would experience if they had not been displaced and if they had the same rank in the earnings distribution as they did in the previous period.

Immediately, Equation (2) implies that the joint distribution $(Y_t(1), \tilde{Y}_t(0)|X, D = 1)$ is identified. One idea is to just call features of this joint distribution as the parameters of interest. In other words, without any additional assumptions, we can identify parameters such as the fraction of displaced individuals that have higher earnings than they would have had if they had not been displaced and if they had maintained the same rank in the earnings distribution as they did in the pre-treatment period.

Alternatively, we can make the additional assumption:

Assumption 3 (Rank Invariance Over Time (RIOT)).

$$F_{Y_t(0)|X, D=1}(Y_t(0)|X) = F_{Y_{t-1}|X, D=1}(Y_{t-1}|X)$$

Assumption 3 says that, in the absence of job displacement, displaced workers would have maintained their rank in the distribution of earnings over time. Assumption 3 implies

that $Y_t(0) = \tilde{Y}_t(0)$ and, thus, that the joint distribution $(Y_t(1), Y_t(0)|X, D = 1)$ is identified, and we can obtain any features of this joint distribution. For comparison, identifying this distribution allows one to ask: What fraction of displaced workers actually benefit from job displacement without the additional caveat mentioned above. At any rate, employing Assumption 3 leads to exactly the same results as focusing on features of the joint distribution $(Y_t(1), \tilde{Y}_t(0)|X, D = 1)$ – only the interpretation is different. We use the notation $Y_t(0)$ rather than $\tilde{Y}_t(0)$ throughout the rest of the paper for simplicity but we interpret our results carefully in the application as the caveat of being conditional on having the same rank over time likely applies.

Remark 1. An alternative approach is to define $\tilde{Y}_t(0) = Y_{t-1}$ for individuals in the treated group.⁷ In this case, one can immediately identify the joint distribution $(Y_t(1), \tilde{Y}_t(0)|X, D = 1)$. Relative to the previous case, there are some advantages and disadvantages to this approach. The main advantage is that it does not require Assumptions 1 and 2 as the distribution of $(Y_t(1), \tilde{Y}_t(0)|X, D = 1)$ is identified from the sampling process rather than requiring identification assumptions. Without further assumptions, one can immediately identify things like: the fraction of displaced workers who have higher earnings following job displacement than they had before they were displaced, etc. The main disadvantage is that, in this case, $Y_t(0) = \tilde{Y}_t(0)$ seems less plausible. But, at any rate, this approach serves as a useful comparison.

Remark 2. Assumption 3 is distinct from the more commonly invoked “cross-sectional” rank invariance assumption – namely, that a displaced worker’s rank in the distribution of $Y_t(1)$ (which is observed) is the same as their rank would be in the distribution of $Y_t(0)$ if they had not been displaced. This is a very strong (and implausible) assumption in the current context as it rules out things like individuals not finding work following displacement.

⁷Farber (2017, Section 7) considers the fraction of displaced workers who have higher earnings following displacement than they did in their previous job which is similar to what we propose here.

2.3 Comparison to Quantile Treatment Effects

A common alternative strategy for trying to analyze heterogeneous effects of experiencing some treatment is to compare the distribution of treated potential outcomes to the distribution of untreated potential outcomes; in our case, this would mean comparing the distribution of outcomes that displaced workers actually experienced to the distribution of outcomes that they would have experienced in the absence of job displacement.⁸ Unlike our main parameters of interest, comparing these two distributions does not require identifying the joint distribution of treated and untreated potential outcomes – in fact, these are the two distributions that are identified using only Step 1 of our identification argument above. Often researchers invert these two distributions and take their difference. The result of this transformation is a class of parameters called quantile treatment effects which are important parameters in the program evaluation literature.

In the context of job displacement, these types of parameters may have severe limitations. These limitations stem from the fact that individuals may have substantial differences in their “rank” if they were displaced relative to if they were not displaced.⁹ This means that, although it is tempting to interpret differences between, for example, the 10th percentile of displaced potential earnings and non-displaced potential earnings as the effect of job displacement for individuals at the lower part of the earnings distribution, this sort of interpretation is unlikely to be correct. The reason is that some individuals in the lower

⁸Identifying these distributions typically either requires experimental data or some identifying argument. However, there are many approaches that are available in the econometrics literature for identifying these distributions; e.g., Firpo (2007) under selection on observables, Abadie (2003), Frolich and Melly (2013), and Wüthrich (2019) when a researcher has access to an instrument, Athey and Imbens (2006), Bonhomme and Sauder (2011), Melly and Santangelo (2015), Callaway, Li, and Oka (2018), and Callaway and Li (2019) when a researcher has access to more than one period of data. The latter set of papers is most relevant in the current application.

⁹This sort of issue is well-known in the econometrics literature (e.g., Heckman, Smith, and Clements (1997)). In other cases, comparing these two distributions may be very useful for policy evaluation or to perform social welfare calculations (Sen (1997) and Carneiro, Hansen, and Heckman (2001)). However, this does not apply for job displacement primarily because job displacement is not a “policy” (if it were, our results would indicate that it is a really bad policy). Instead, the reason why researchers would be interested in distributions here is to examine heterogeneous effects as we do in the current paper, and simply comparing the marginal distributions is of limited usefulness in this case.

part of the earnings distribution might have had relatively high earnings if they had not been displaced. When some individuals remain unemployed or have substantially reduced hours following job displacement, this can systematically lead to larger differences between the observed distribution and counterfactual distribution occurring in the lower part of the distribution irrespective of whether the effect of job displacement is bigger or smaller for low or high earnings workers. Instead, to think about these sorts of effects, we propose to estimate the distribution of the individual-level effect of job displacement for displaced workers conditional on what earnings would have been in the absence of displacement. Because of our above results on identifying the joint distribution of treated and untreated potential outcomes, these sorts of parameters are identified and estimable in our framework.

3 Estimation

Given the identification results above, determining the best estimation strategy is still challenging. Like many applications in economics, we have a relatively large number of covariates that we would like to condition on to make the identifying assumptions more plausible and only a moderate number of observations. The first step is to estimate the counterfactual distribution of untreated potential outcomes conditional on covariates for the treated group as in Equation (1). This requires estimating three conditional distribution/quantile functions. These are challenging. Distributional assumptions about the distribution of outcomes conditional on covariates are likely to be invalid. On the other hand, implementing fully non-parametric estimators is infeasible in our case. Callaway, Li, and Oka (2018) consider the case with only discrete covariates and propose nonparametric estimators based on splitting the sample according to each possible combination of discrete covariates. This is infeasible in our case because we have some continuous covariates and the number of discrete covariates is too large relative to the number of observations to split the sample. Instead, we propose to estimate these conditional distributions/quantiles using quantile regression. In particular,

we suppose that

Assumption 4 (Quantile Regression). *For all $(s, d) \in \{t, t - 1\} \times \{0, 1\}$,*

$$Y_s(d) = Q_{Y_s|X, D=d}(U_{sd}|X) = P_{sd}(X)' \beta_{sd}(U_{sd})$$

where $U_{sd}|X, D \sim U[0, 1]$ and $P_{sd}(X)$ are some transformations of the vector of covariates X .

Assumption 4 suggests that we can estimate each of the terms in Equation (1) using QR. In particular, consider the following algorithm to estimate the counterfactual distribution in Equation (1) and obtain $Y_t(0)$ for each individual in our sample.

Algorithm 1. Let τ denote a fine grid of equally spaced elements between 0 and 1,

1. Estimate $\hat{Q}_{Y_{t-1}|X, D=1}(\tau)$ by quantile regression of Y_{t-1} on X for individuals in the treated group
2. Estimate $\hat{Q}_{Y_{t-1}|X, D=0}(\tau)$ by quantile regression and invert to obtain $\hat{F}_{Y_{t-1}|X, D=0}$.
3. Set $\tilde{Y}_{t-1} := \hat{Q}_{Y_{t-1}|X, D=1}(\hat{F}_{Y_{t-1}|X, D=0}(Y_{t-1}|X)|X)$
4. Estimate $\hat{Q}_{Y_t(0)|X, D=1}(\tau)$ by quantile regression of $\Delta Y_t + \tilde{Y}_{t-1}$ on X for individuals in the untreated group
5. Set $Y_t(0) := \hat{Q}_{Y_t(0)|X, D=1}(\hat{F}_{Y_{t-1}|X, D=1}(Y_{t-1}|X)|X)$

Algorithm 1 provides $Y_{it}(0)$ for all individuals in the treated group. Thus, we can immediately obtain features of the unconditional distribution of $(Y_t(1), Y_t(0)|D = 1)$ using the draws of $(Y_{it}, Y_{it}(0))$ for displaced workers. For example, we can estimate the fraction of individuals that have higher earnings following job displacement by

$$\frac{1}{n_1} \sum_{i \in \mathcal{D}} \mathbb{1}\{Y_{it} - Y_{it}(0) > 0\}$$

where n_1 denotes the number of displaced workers and \mathcal{D} denotes the set of displaced workers. We are also interested in how features of the distribution of $Y_t(1) - Y_t(0)$ vary with covariates. It is straightforward to run OLS regressions of $Y_t(1) - Y_t(0)$ on covariates. These results are interesting, and we report them in the application. However, our identification results allow us to go substantially further than this. In particular, we can also estimate how covariates affect the entire distribution of the effect of job displacement. To estimate this, we again can use quantile regression. In particular, we suppose that

$$Q_{(Y_t(1)-Y_t(0))|X,D=1}(\tau|x) = P_{\Delta_X}(x)' \beta_{\Delta_X}(\tau)$$

where P_{Δ_X} denotes transformations of the covariates. The estimates of these conditional quantiles and quantile regression parameters are of interest. For example, if we set τ to be small (e.g., 0.05), $\beta_{\Delta_X}(\tau)$ provides the effect of covariates on the individual-level effect of job displacement for individuals who are most negatively affected by job displacement.

Finally, we are interested in features of the distribution of $Y_t(1) - Y_t(0)$ across different values of $Y_t(0)$; that is, how the effect of job displacement changes across different values of what a displaced worker's earnings would have been if they had not been displaced. Once again, one can estimate OLS regressions of $Y_t(1) - Y_t(0)$ on $Y_t(0)$; however, we also use quantile regression to study how the entire distribution of the effect of job displacement is related to non-displaced potential earnings. Here we suppose that

$$Q_{(Y_t(1)-Y_t(0))|Y_t(0),D=1}(\tau|y) = P_{\Delta_Y}(y)' \beta_{\Delta_Y}(\tau)$$

where $P_{\Delta_Y}(y)$ are transformations of non-displaced potential earnings, and β_{Δ_Y} gives the effect of non-displaced potential earnings on particular quantiles of the individual-level effect of job displacement.

4 Results

Our work is related to a large literature on job displacement. We only briefly survey the most relevant research here. Almost all work on job displacement considers a single “effect” of job displacement. This effect usually corresponds to an estimated parameter in an OLS regression. Much of the literature (e.g., Jacobson, LaLonde, and Sullivan (1993)) has been interested in the long-term effect of job displacement. This is particularly true for papers that have used widely available panel data sets such as the NLSY and PSID (e.g., Ruhm (1991), Stevens (1997), and Kletzer and Fairlie (2003)) as well as administrative data (e.g., von Wachter, Song, and Manchester (2009) and Couch and Placzek (2010)). On the other hand, research that has used the Displaced Workers Survey has focused on short run effects of job displacement because the DWS does not contain follow up interviews with displaced workers. Thus, here, we consider only short-run heterogeneous effects of job displacement though it would be interesting to extend our approach to the long-term case though this would require an alternative data source.

Two strands of the literature are most related to the current paper. First, in a series of work (Farber (1997), Farber (2005), and Farber (2017)), Henry Farber has used the DWS to study the effect of job displacement. Broadly, the main concerns of these papers are the employment and earnings effects of job displacement and how these have varied over time. As mentioned above, Farber (2017) also considers the fraction of displaced workers whose earnings are higher following job displacement than they were in their pre-displacement job. A primary difference between our main results and those in Farber (2017) is that we try to adjust for how earnings for displaced workers would have evolved over time in the absence of job displacement. Our results that also use raw changes in earnings over time are also somewhat different; in particular, Farber (2017, Section 7) focuses mostly on displaced workers that have moved back into full time work while we focus on all displaced workers that are still in the labor force. In addition, we think about the entire distribution of the effect of job displacement and how the distribution of the effect of job displacement varies

across covariates and varies for individuals who would have had high or low earnings if they had not been displaced from their job. On the other hand, Farber (2017) carefully considers the role of measurement error and shows that this may somewhat reduce the fraction of displaced workers who actually have higher earnings following job displacement than they had in their pre-displacement job. Second, the application in Callaway (2018) also considers heterogeneous effects of job displacement. The assumptions in that paper are weaker than in the current paper (particularly in the case when we invoke Assumption 3) though his approach leads to bounds on heterogeneous treatment effect parameters rather than point identification which implies that we are able to learn more about the heterogeneous effect of job displacement. In addition, using our approach, we are also able to relate heterogeneous effects to covariates as well as to what earnings would have been in the absence of job displacement.

4.1 Data

We use data from the Displaced Workers Survey (DWS). The DWS has been administrated every two years since 1984 as a supplement to the Current Population Survey (CPS) (Flood, King, Ruggles, and Warren (2017)) and contains information on earnings and employment status for displaced workers. Respondents are asked whether they lost their jobs at any point in the last three years due to: (i) plant closing or moving, (ii) insufficient work, or (iii) position being abolished. These are sometimes referred to as the “big three” reasons for job displacement (Farber (2017)). Our main results below focus on 2015-2016 which is the most recent period when DWS data is available. We also provide similar results for 2009-2010 (during the Great Recession) and 1997-1998 (an earlier period when the U.S. economy was quite strong).

Although the DWS provides useful information on job displacement status and the current weekly earnings of displaced workers and their weekly earnings on their lost job, it does not provide any information about the earnings of non-displaced workers (though non-

displaced workers can be identified). After identifying non-displaced workers using the DWS, we obtain their earnings by matching the DWS to CPS outgoing rotation group data using individual-level identification numbers which are available from IPUMS.¹⁰ This matching process helps us to find information on the current weekly earnings as well as weekly earnings in the previous year for non-displaced workers. Effectively, this turns our untreated group into a two period panel dataset. We also drop all non-displaced workers whose earnings are missing in either period. For workers that are displaced, we drop those that are not currently in the labor force and assign 0 earnings to those in the labor force but who are not currently employed. The timing for available earnings is also somewhat different for displaced and non-displaced workers. For displaced workers, the DWS asks for their earnings in their pre-displacement job. But displacement could have occurred any time in the last three years. On the other hand, for non-displaced workers, we only have their current weekly earnings and their earnings one year ago.

The summary statistics for key variables of displaced and non-displaced workers in 2016 are presented in Table 1. Weekly earnings are about \$74 dollars higher for non-displaced workers than displaced workers in the period before they are displaced. Following job displacement, this gap in average weekly earnings increases to \$313. This occurs due to a combination of average weekly earnings increasing by \$44 for non-displaced workers and decreasing by \$195 for displaced workers (on top of the existing gap in the previous period). For covariates, we include an individual's sex, race, marital status, age, and education. Displaced workers are more likely to be male, not married, younger, and not have a college degree; both groups are roughly equally likely to be white or have less than a high school education.

¹⁰Households in the CPS are usually interviewed for four consecutive months, are out of the sample for the next eight months, and then interviewed for another four months. After the eighth time participating, they leave the sample. Households that are interviewed in the fourth or eighth month (which are called outgoing rotation groups) are asked additional labor market questions. Thus, using our approach, we observe earnings for the untreated group in two consecutive years, but they can be different months for different individuals.

Table 1: Summary Statistics

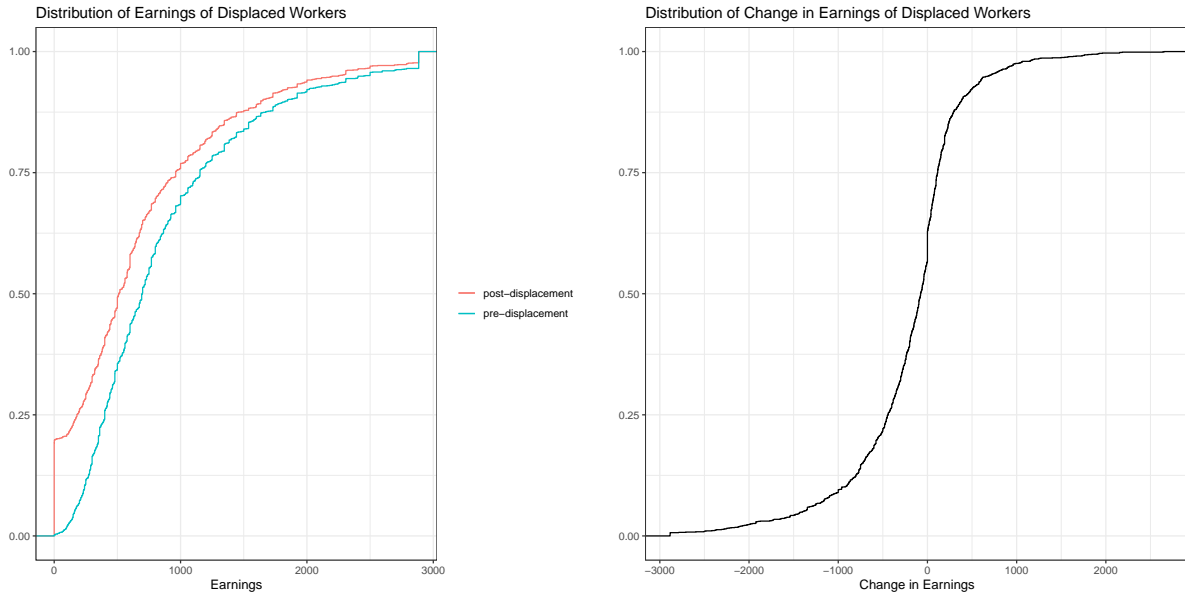
	Displaced	Non-Displaced	Difference	P-val on Difference
Earnings				
2016 Earnings	687.82	1000.43	-312.604	0.00
2015 Earnings	882.56	956.43	-73.875	0.00
Change Earnings	-194.74	43.99	-238.729	0.00
Covariates				
Male	0.57	0.51	0.057	0.00
White	0.81	0.82	-0.008	0.48
Married	0.51	0.63	-0.113	0.00
College	0.33	0.40	-0.066	0.00
Less HS	0.07	0.06	0.009	0.21
Age	42.11	45.4	-3.294	0.00
N	1633	3915		

Sources: CPS and Displaced Workers Survey

4.2 Results

As a first step, Figure 1 plots the distributions of weekly earnings for displaced workers pre- and post-displacement. Clearly, the distribution of post-displacement earnings is shifted to the left. Notably, there is also a mass point in the distribution of post-displacement earnings – about 20% of displaced workers have 0 earnings. These distributions also exhibit a pattern that is consistent with very heterogeneous effects – namely, the difference between the distributions is largest in the lower part of the earnings distributions. The right panel of Figure 1 plots the distribution of the change in weekly earnings over time for displaced workers. Recall that this distribution does not require any identifying assumption; it is just the difference in earnings over time. Interestingly, our estimates in the figure indicate that about 37% of displaced workers have higher weekly earnings than they had before they were displaced. This number may overstate how many workers actually have higher earnings than they *would have had* if they had not been displaced especially in the case where earnings would have increased in the absence of job displacement. On the other hand, around 22%

Figure 1: Distributions of Earnings of Displaced Workers




Notes: The left panel contains plots of the distributions of earnings before and after displacement for displaced workers. The right panel plots the distribution of the change in weekly earnings for workers who were displaced from their jobs.

of workers earn at least \$500 less per week than before displacement and about 10% earn at least \$1000 less per week than before displacement – these are substantial differences. Thus, it is immediately clear that the difference in the average decrease in earnings for displaced workers relative to non-displaced workers of \$239 per week (see Table 1) masks substantial heterogeneity.

The next step in our analysis is to construct the counterfactual distribution of weekly earnings that displaced workers would have experienced if they had not been displaced from their jobs. We do this using the approach of Callaway, Li, and Oka (2018) as described in Equation (1) and Algorithm 1. Figure 2 plots this counterfactual distribution.¹¹ Notice that the counterfactual distribution sits to the right of the actual distribution of earnings. Also, the biggest difference between the two distributions occurs in the lower part of the distribution. Once the counterfactual distribution is identified, this also implies that the

¹¹The distribution depends on the value of the covariates; the figure reports the average of the distribution over all values of covariates for displaced workers in the dataset. This is a standard way of converting conditional effects into unconditional effects.

Figure 2: Counterfactual Distribution of Earnings for Displaced Workers




Notes: Plots of the actual and counterfactual distribution of earnings for displaced workers.

average effect of job displacement is identified. We estimate that job displacement reduces earnings by \$157 per week on average, which is about 18% of pre-displacement earnings.

Remark 3. The results in Figure 2 are also closely related to quantile treatment effects. The Quantile Treatment Effect on the Treated (QTT) is the difference between the quantiles of observed outcomes for displaced workers and the quantiles of the counterfactual distribution of outcomes that they would have experienced if they had not been displaced. The QTT comes from inverting the distributions in Figure 2. Because the difference between these two distributions is largest at the lower part of the distribution, it is tempting to interpret this as saying that the effect of job displacement is largest for workers in the lower part of the earnings distribution. This is an incorrect interpretation though because it ignores the possibility that workers change their ranks in the distribution of earnings due to job displacement (e.g., workers who would have had high earnings if they had not been displaced remaining unemployed following displacement). In fact, we find evidence (discussed below) that points in exactly the opposite direction – that the effect of job displacement appears to

Figure 3: Distribution of the Treatment Effect under Assumption 3



Notes: The black line plots the distribution of the change in outcomes over time for displaced workers. The blue line plots the distribution of the treatment effect for displaced workers.

be larger for high earnings workers than low earnings workers.

Under Assumption 3, we can map each value of $Y_t(1)$ to $Y_t(0)$ and create a pair of $(Y_t(1), Y_t(0))$ for each displaced worker. Figure 3 plots the distribution of $Y_t(1) - Y_t(0)$ for displaced workers under Assumption 3. Interestingly, the distribution of the treatment effect is remarkably close to the distribution of the change in outcomes over time.¹² This implies that our results for the distribution of the actual effect of job displacement are essentially identical to our earlier estimates that instead use the change in earnings over

¹²It is worth explaining more why the distributions are so similar. The black line is the distribution of $Y_t - Y_{t-1}$ for displaced workers. The blue line plots the distribution of $Y_t - Y_t(0)$ for displaced workers. Thus, any differences between the two distributions comes from differences between Y_{t-1} and $Y_t(0)$. From Equation (2), we generate $Y_t(0) = Q_{Y_t(0)|X, D=1}(F_{Y_{t-1}|X, D=1}(Y_{t-1}|X)|X)$; in practice, this amounts to finding an individual's rank in the conditional distribution of outcomes in time period $t - 1$ and applying the conditional quantile function for untreated potential outcomes to that rank. Also, notice that $Y_{t-1} = Q_{Y_{t-1}|X, D=1}(F_{Y_{t-1}|X, D=1}(Y_{t-1}|X)|X)$; in other words, an alternative way to obtain an individual's outcome in period $t - 1$ is to find their rank in the conditional distribution in time period $t - 1$ and apply the conditional quantile for outcomes in period $t - 1$ to that rank. Thus, any difference between the distribution of $Y_t - Y_{t-1}$ and $Y_t - Y_t(0)$ for displaced workers comes down to differences in the conditional quantiles (or distributions) of $Y_t(0)$ and Y_{t-1} . In our case, these two distributions are very similar – suggestive evidence of this can be found in Figures 1 and 2 – and this explains the reason for the high degree of similarity in Figure 3.

time for displaced workers.¹³ Here, we estimate that (i) 42% of displaced workers would have had higher earnings if they had not been displaced,¹⁴ (ii) 21% have at least \$500 less in weekly earnings than they would have had if they had not been displaced, and (iii) 8% have at least \$1000 less weekly earnings than they would have had if they had not been displaced.

Taken together, our results so far indicate that there is substantial heterogeneity in the effect of job displacement. Even though there is a fairly large average effect of job displacement, a substantial fraction of workers appear to be no worse off due to job displacement than they would have been if they had not been displaced. On the other hand, the negative effects of job displacement seem to be largely concentrated among 10-20% of displaced workers who experience large negative effects.

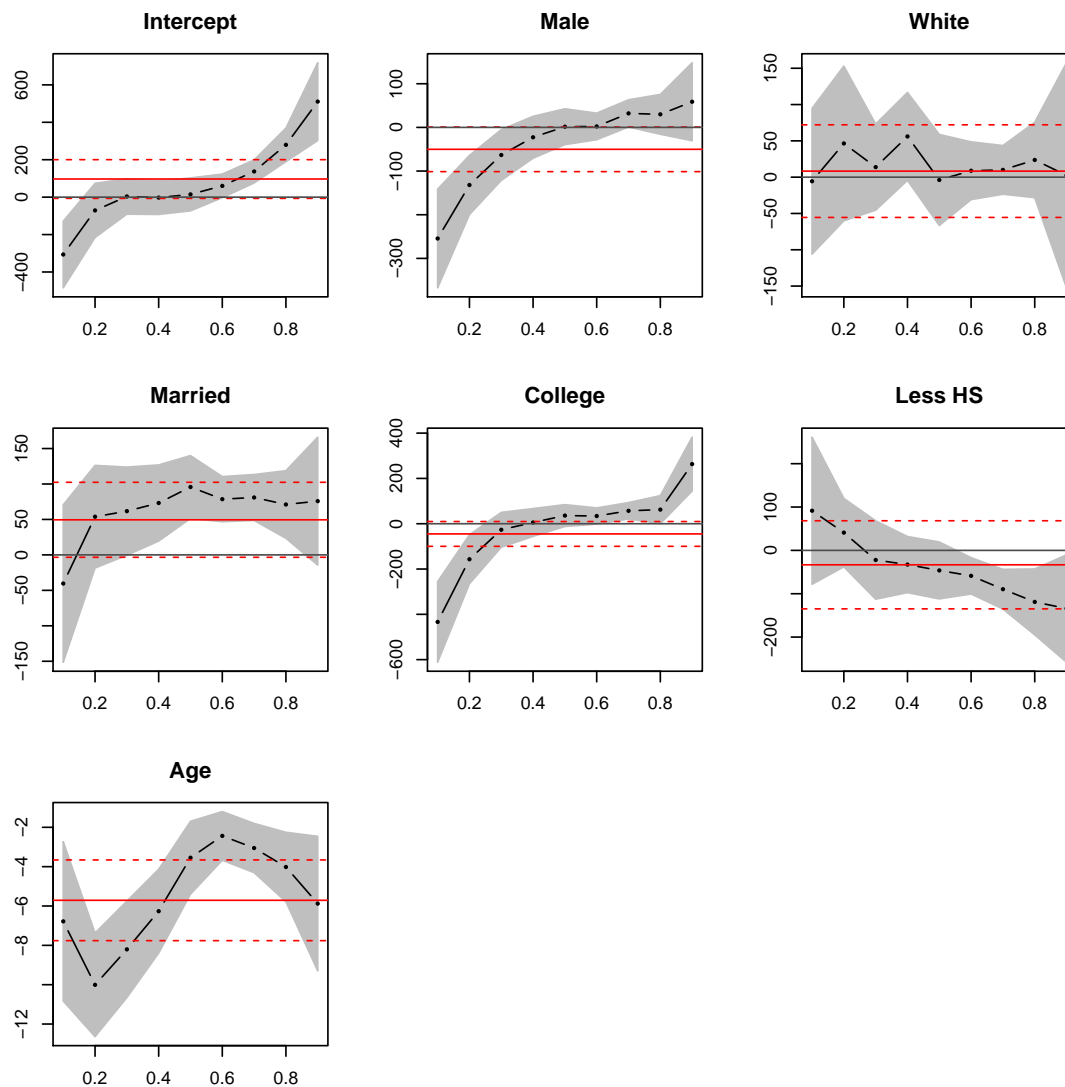
Next, we use quantile regression to see how covariates affect the distribution of the treatment effect for displaced workers. Figure 4 provides estimates of the effect of each covariate across different quantiles of the conditional distribution of the treatment effect. There are some quite interesting patterns here. On average, the effect of job displacement is somewhat more negative for men than for women, but the effect is very heterogeneous. For individuals that experience large negative effects, the effects for men are much more negative than the effects for women.¹⁵ On the other hand, for individuals that are towards

¹³As discussed earlier, the appropriate interpretation of these results depends on whether or not one believes Assumption 3. Under Assumption 3, the results that we report should be considered as differences between actual earnings of displaced workers and what they would have earned had they not been displaced. Otherwise, these are differences between actual outcomes for displaced workers and the outcome that they would have experienced if (i) they had not been displaced from their job and (ii) they maintained the same rank in the conditional earnings distribution over time.

¹⁴It is perhaps surprising that such a large fraction of displaced workers appear to have higher earnings than they would have had if they had not been displaced. Comparing pre- and post-displacement earnings of displaced workers, Farber (2017) similarly finds large percentages of workers earning more in their post-displacement job than in their pre-displacement job. Some possible explanations for this finding that he offers (and that likely apply to our case as well) are: (i) some displaced workers may have higher earnings but lower utility in their new job, (ii) search costs may keep workers from looking for new jobs that could potentially offer higher earnings, and (iii) risk aversion in the sense that workers have more information about their current job and may be reluctant to leave it for the possibility of a higher paying job when they have less information about other aspects of the new job. Another possible explanation is measurement error in earnings though Farber (2017) considers a sensitivity analysis related to measurement error and finds only fairly moderate changes to these results even under conservative assumptions on measurement error in earnings.

¹⁵Recall that the outcome in these quantile regressions is $Y_t - Y_t(0)$ for displaced individuals. This means that “large negative effects” occur at the lower quantiles in Figure 4.

Figure 4: Quantile Regression Estimates of the Treatment Effect on Covariates



Notes: QR estimates of the effect of covariates on the quantiles of the treatment effect. The solid horizontal line provides OLS estimates of the effect of each covariate, and the horizontal dashed line contains a 90% confidence interval. The dotted line provides QR estimates of the effect of each covariate at particular quantiles from 0.1, 0.2, ..., 0.9. The shaded area contains pointwise 90% confidence intervals from the quantile regressions.


the top of the distribution of the treatment effect (i.e., experience small or positive effects of job displacement), men tend to have larger benefits than women. Differences in the effect of job displacement by race tend to be small and not exhibit strong patterns at different parts of the conditional distribution. The effect of job displacement for college graduates is also very heterogeneous. Some college graduates experience tremendously large negative

effects from job displacement, while college graduates (conditional on doing well following job displacement) tend to do better than high school graduates. For individuals with less than a high school education, these results are almost exactly reversed. Finally, our results indicate that being older increases the negative effects of job displacement among those most affected by job displacement. Once again, though, the effects are quite heterogeneous – the effect of age for those less affected by job displacement appears to be either positive or close to 0.

These results suggest that the largest negative effects of job displacement tend to be for older, male, college graduates. This should probably not be a surprising result as the effect of job displacement depends on both earnings following job displacement, Y_t , as well as what earnings *would have been* in the absence of job displacement, $Y_t(0)$. Older, male, college graduates may have large earnings losses due to job displacement particularly in the case where if, had they not been displaced, they would have had high earnings.

Finally, in this section, we investigate how the effect of job displacement varies with $Y_t(0)$ – i.e., how the effect varies across individuals as a function of what their earnings would have been if they had not been displaced. We plot quantile regression estimates from regressing $(Y_t - Y_t(0))$ on $Y_t(0)$ in Figure 5. On average, the effect of job displacement is larger for individuals who would have had higher earnings in the absence of job displacement. Notably, this is exactly the opposite interpretation as one would be tempted to make if one only compared the marginal distributions (see discussion in Remark 3). In addition, there is substantial heterogeneity. For individuals that experience the largest negative effects of job displacement, these negative effects appear to be substantially larger for those who would have had high earnings in the absence of job displacement. On the other hand, for those who are less affected by job displacement, the difference between the effect for those who would have had high earnings relative to those who would have had low earnings in the absence of job displacement is small.

Figure 5: Quantile Regression Estimates of the Treatment Effect on $Y_t(0)$



Notes: QR estimates of the effect of what earnings would have been in the absence of job displacement, $Y_t(0)$, on the quantiles of the treatment effect. The solid horizontal line provides OLS estimates of the effect of non-displaced potential earnings, and the horizontal dashed line contains a 90% confidence interval. The dotted line provides QR estimates of the effect of non-displaced potential earnings at particular quantiles from 0.1, 0.2, \dots , 0.9. The shaded area contains pointwise 90% confidence intervals from the quantile regressions.

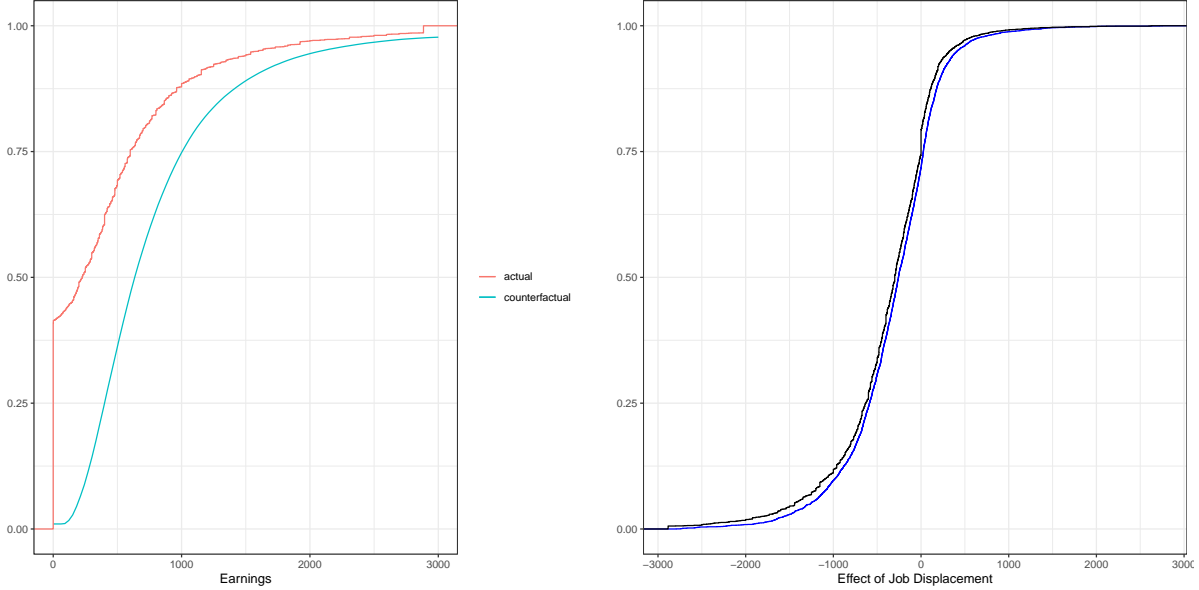
4.3 Results in Other Time Periods

To conclude this section, we provide analogous results for 2009-2010 (corresponding to the Great Recession) and for 1997-1998 (an earlier time period when the U.S. economy was quite strong). We briefly discuss the main results from these other periods below and provide additional details in the appendix (in particular, we provide summary statistics and quantile regression estimates of the effect of covariates and the effect of non-displaced potential earnings on the distribution of the treatment effect). All of the results in this section use the same identification arguments and estimation approaches that were applied to the 2015-2016 data above.

Results for 2009-2010

These results are for workers who were displaced at the height of the Great Recession.

Figure 6: Main Results for 2009-2010



Notes: The left panel contains plots of the distributions of earnings for displaced workers as well as their counterfactual distribution of earnings if they had not been displaced. The right panel contains the distribution of the treatment effect for displaced workers.

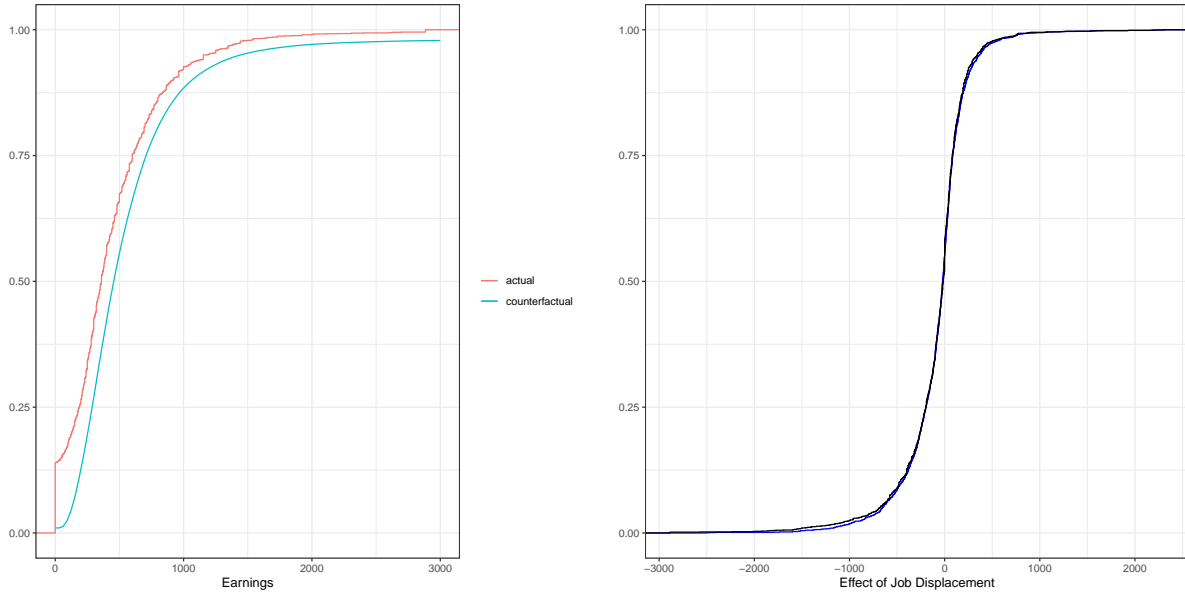
Figure 6 plots the observed and counterfactual distributions of outcomes for displaced workers. The effects of job displacement are much larger, on average, than in 2015-2016. We estimate that workers who were displaced from their job earn \$308 less on average than they would have earned if they had not been displaced; this corresponds to 39% lower earnings on average relative to pre-displacement earnings. Much of these effects appear to be driven by not being able to find work. 41% of displaced workers continued to be unemployed in 2010 (for 2015-2016, only 20% of displaced workers continued to be unemployed in 2016). Interestingly, we estimate that 28% of displaced workers have higher earnings than they would have had if they had not been displaced. Even when we just compare current earnings to pre-displacement earnings, we find that 21% of displaced workers had higher earnings following displacement than they had in their pre-displacement job. This indicates that even though the effects of job displacement appear to be substantially larger during the Great Recession, there is still substantial heterogeneity and a non-trivial fraction of displaced workers seem to have higher earnings than they otherwise would have had.

The effect of covariates on the distribution of the treatment effect in 2009-2010 is broadly similar to the effect in 2015-2016 (see Figure 8 in the appendix). In particular, among those most negatively affected by job displacement, older, male, college graduates tend to be the most affected (here, there is also some evidence that among those most affected by job displacement, married displaced workers are also more negatively affected by job displacement). Finally, similarly to 2015-2016, the effect of job displacement tends to be negatively related to earnings that individuals would have experienced if they had not been displaced (see Figure 9 in the appendix). These effects tend to be biggest for individuals who experience the largest negative effects of job displacement.

Results for 1997-1998

Finally, we examine the effect of job displacement in 1997-1998. This is an earlier period than the ones we have considered so far, and the U.S. economy was quite strong during this period. Figure 7 contains our main estimates from this period. We estimate that displaced workers earned on average \$68 less per week than they would have earned if they had not been displaced. This estimate corresponds to 13% lower earnings on average relative to pre-displacement earnings. This estimate is somewhat smaller than the corresponding estimate for 2015-2016 and much smaller than the corresponding estimate for 2009-2010. Part of the reason for the smaller effects here is that only 14% of displaced workers are unemployed in 1998. In addition, we estimate that 45% of displaced workers have higher earnings than they would have had if they had not been displaced. Similarly, 42% of displaced workers had higher earnings in their current job than they did in their pre-displacement job. Finally, our estimates for the effect of covariates and untreated potential outcomes are broadly the same as in the other periods (see Figures 10 and 11 in the appendix). For displaced workers that are most affected by job displacement, these effects once again tend to be largest for older, male, college graduates and for those who would have had high earnings if they had not been displaced. Given our results in earlier sections, these results seem to indicate that

Figure 7: Main Results for 1997-1998



Notes: The left panel contains plots of the distributions of earnings for displaced workers as well as their counterfactual distribution of earnings if they had not been displaced. The right panel contains the distribution of the treatment effect for displaced workers.

the finding that the largest effects of job displacement are concentrated among these groups is robust across time periods and across various states of the macroeconomy.

5 Conclusion

In this paper, we have studied how the effect of job displacement varies across different individuals. To do this, we compared the earnings of displaced workers to what their earnings would have been if they had maintained their rank in the (conditional) distribution of weekly earnings (though we allow for the distribution of weekly earnings to change over time). A key requirement of our approach was to be able to estimate several conditional distributions in a first step in a way that was both not too restrictive while also being feasible with a moderate sized data set and a fairly large number of covariates. To do this, we relied heavily on first-step quantile regression estimators.

We found that displaced workers lose about \$157 per week due to job displacement, on

average. However, that average effect masks a tremendous amount of heterogeneity. A large fraction of workers appear to be no worse off (or even better off) in terms of their earnings following job displacement than they would have been if they had not been displaced. On the other hand, another large fraction loses substantially more than the average earnings loss. Once we had obtained the distribution of the effect of job displacement, we used quantile regression to study how the distribution of the effect of job displacement depends on covariates. We also showed that this distribution varies substantially across sex, education levels, and age, as well as across different amounts of earnings that displaced workers would have had if they had not been displaced.

References

- [1] Abadie, Alberto. “Semiparametric instrumental variable estimation of treatment response models”. *Journal of Econometrics* 113.2 (2003), pp. 231–263.
- [2] Abowd, John M, Francis Kramarz, and David N Margolis. “High wage workers and high wage firms”. *Econometrica* 67.2 (1999), pp. 251–333.
- [3] Athey, Susan and Guido Imbens. “Identification and inference in nonlinear difference-in-differences models”. *Econometrica* 74.2 (2006), pp. 431–497.
- [4] Bitler, Marianne P, Jonah B Gelbach, and Hilary W Hoynes. “Can variation in subgroups’ average treatment effects explain treatment effect heterogeneity? Evidence from a social experiment”. *Review of Economics and Statistics* 99.4 (2017), pp. 683–697.
- [5] Bonhomme, Stephane 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.
- [6] Callaway, Brantly. “Tight bounds on distributional treatment effect parameters using panel data with an application on job displacement”. Working Paper. 2018.
- [7] Callaway, Brantly and Tong Li. “Quantile treatment effects in difference in differences models with panel data”. Working Paper. 2019.
- [8] 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.
- [9] 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.
- [10] Carrington, William J and Bruce Fallick. “Why do earnings fall with job displacement?” *Industrial Relations: A Journal of Economy and Society* 56.4 (2017), pp. 688–722.
- [11] Couch, Kenneth A and Dana W Placzek. “Earnings losses of displaced workers revisited”. *The American Economic Review* (2010), pp. 572–589.
- [12] Fan, Yanqin, Emmanuel Guerre, and Dongming Zhu. “Partial identification of functionals of the joint distribution of potential outcomes”. *Journal of Econometrics* 197.1 (2017), pp. 42–59.

- [13] Fan, Yanqin and Sang Soo Park. “Sharp bounds on the distribution of treatment effects and their statistical inference”. *Econometric Theory* 26.03 (2010), pp. 931–951.
- [14] Farber, Henry S. “The changing face of job loss in the United States, 1981-1995”. *Brookings Papers on Economic Activity. Microeconomics* (1997), pp. 55–142.
- [15] Farber, Henry S. “What do we know about job loss in the United States? Evidence from the Displaced Workers Survey, 1984-2004”. *Economic Perspectives* 29.2 (2005), pp. 13–29.
- [16] Farber, Henry S. “Employment, hours, and earnings consequences of job loss: US evidence from the displaced workers survey”. *Journal of Labor Economics* 35.S1 (2017), S235–S272.
- [17] Firpo, Sergio. “Efficient semiparametric estimation of quantile treatment effects”. *Econometrica* 75.1 (2007), pp. 259–276.
- [18] Flood, Sarah, Miriam King, Steven Ruggles, and J Robert Warren. *Integrated Public Use Microdata Series, Current Population Survey: Version 5.0*. Minneapolis: University of Minnesota. 2017.
- [19] Frandsen, Brigham R and Lars J Lefgren. “Partial identification of the distribution of treatment effects”. Working Paper. 2017.
- [20] Frolich, Markus and Blaise Melly. “Unconditional quantile treatment effects under endogeneity”. *Journal of Business & Economic Statistics* 31.3 (2013), pp. 346–357.
- [21] 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.
- [22] Jacobson, Louis S, Robert J LaLonde, and Daniel G Sullivan. “Earnings losses of displaced workers”. *The American Economic Review* (1993), pp. 685–709.
- [23] Jovanovic, Boyan. “Job matching and the theory of turnover”. *Journal of Political Economy* 87.5, Part 1 (1979), pp. 972–990.
- [24] Kletzer, Lori G and Robert W Fairlie. “The long-term costs of job displacement for young adult workers”. *ILR Review* 56.4 (2003), pp. 682–698.
- [25] Koenker, Roger. *Quantile Regression*. Cambridge University Press, 2005.
- [26] Koenker, Roger and Gilbert Bassett Jr. “Regression quantiles”. *Econometrica* (1978), pp. 33–50.
- [27] Melly, Blaise and Giulia Santangelo. “The changes-in-changes model with covariates”. Working Paper. 2015.

- [28] Neal, Derek. “Industry-specific human capital: Evidence from displaced workers”. *Journal of Labor Economics* 13.4 (1995), pp. 653–677.
- [29] Ruhm, Christopher J. “Are workers permanently scarred by job displacements?” *The American Economic Review* 81.1 (1991), pp. 319–324.
- [30] Sen, Amartya. *On Economic Inequality*. Clarendon Press, 1997.
- [31] Sklar, Abe. *Fonctions de répartition à n dimensions et leurs marges*. Publications de L Institut de Statistique de L Universite de Paris, 1959.
- [32] Stevens, Ann Huff. “Persistent effects of job displacement: The importance of multiple job losses”. *Journal of Labor Economics* (1997), pp. 165–188.
- [33] Topel, Robert. “Specific capital, mobility, and wages: Wages rise with job seniority”. *Journal of Political Economy* 99.1 (1991), pp. 145–176.
- [34] von Wachter, Till, Jae Song, and Joyce Manchester. “Long-term earnings losses due to mass layoffs during the 1982 recession: An analysis using US administrative data from 1974 to 2004”. Working Paper. 2009.
- [35] Wüthrich, Kaspar. “A closed-form estimator for quantile treatment effects with endogeneity”. *Journal of Econometrics* 210.2 (2019), pp. 219–235.

A Additional Tables and Figures

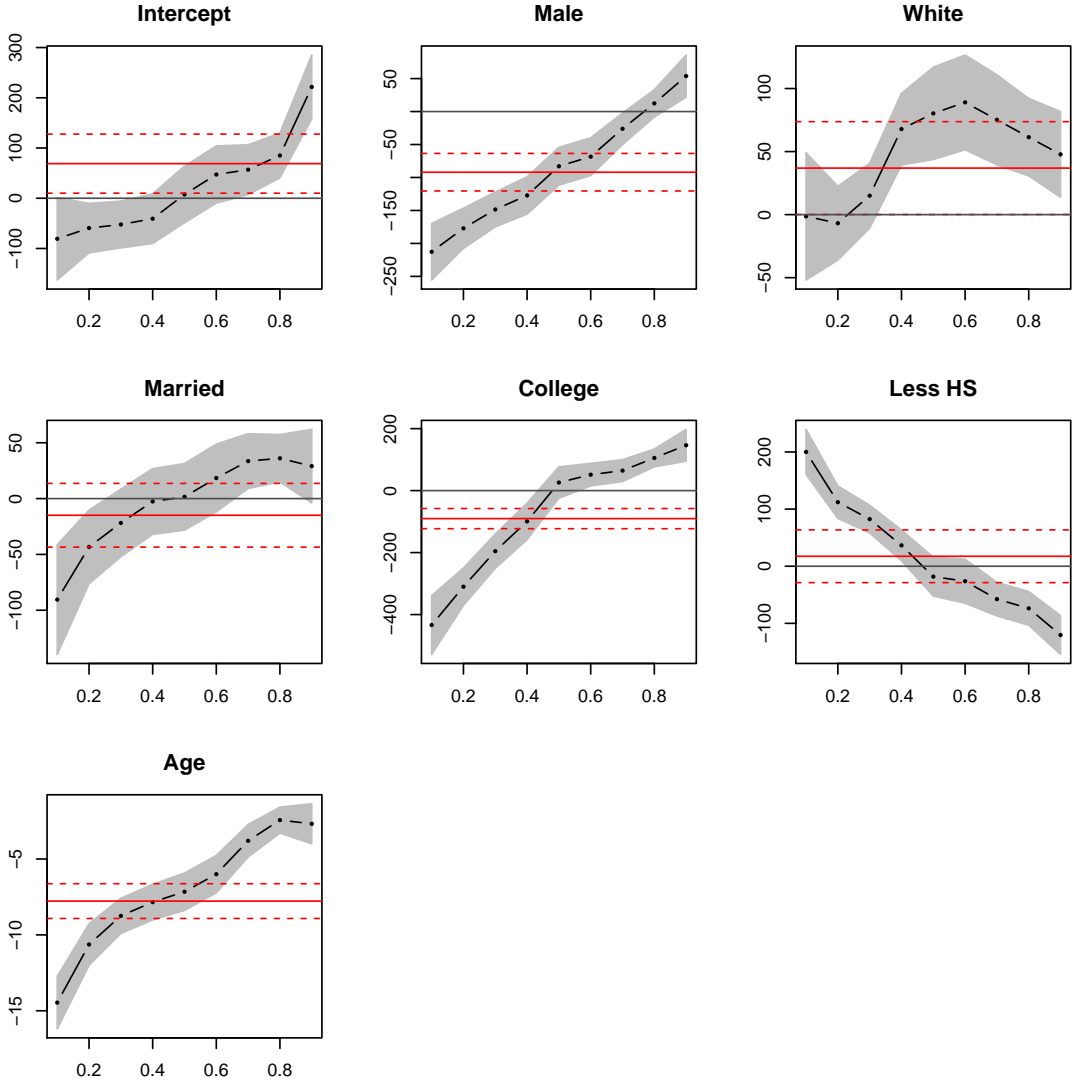
A.1 2009-2010 Results

Table 2: 2009-2010 Summary Statistics

	Displaced	Non-Displaced	Difference	P-val on Difference
Earnings				
2010 Earnings	418.37	887.54	-469.169	0.00
2009 Earnings	797.24	876.66	-79.424	0.00
Change Earnings	-378.87	10.87	-389.745	0.00
Covariates				
Male	0.60	0.48	0.122	0.00
White	0.83	0.85	-0.017	0.04
Married	0.53	0.62	-0.096	0.00
College	0.25	0.38	-0.123	0.00
Less HS	0.10	0.06	0.044	0.00
Age	41.50	44.87	-3.371	0.00
N	4088	4000		


Sources: CPS and Displaced Workers Survey

Figure 8: 2009-2010 Quantile Regression Estimates of the Treatment Effect on Covariates



Notes: QR estimates of the effect of covariates on the quantiles of the treatment effect in 2009-2010. The solid horizontal line provides OLS estimates of the effect of each covariate, and the horizontal dashed line contains a 90% confidence interval. The dotted line provides QR estimates of the effect of each covariate at particular quantiles from 0.1, 0.2, ..., 0.9. The shaded area contains pointwise 90% confidence intervals from the quantile regressions.

Figure 9: 2009-2010 Quantile Regression Estimates of the Treatment Effect on $Y_t(0)$



Notes: QR estimates of the effect of what earnings would have been in the absence of job displacement, $Y_t(0)$, on the quantiles of the treatment effect in 2009-2010. The solid horizontal line provides OLS estimates of the effect of non-displaced potential earnings, and the horizontal dashed line contains a 90% confidence interval. The dotted line provides QR estimates of the effect of non-displaced potential earnings at particular quantiles from 0.1, 0.2, ..., 0.9. The shaded area contains pointwise 90% confidence intervals from the quantile regressions.

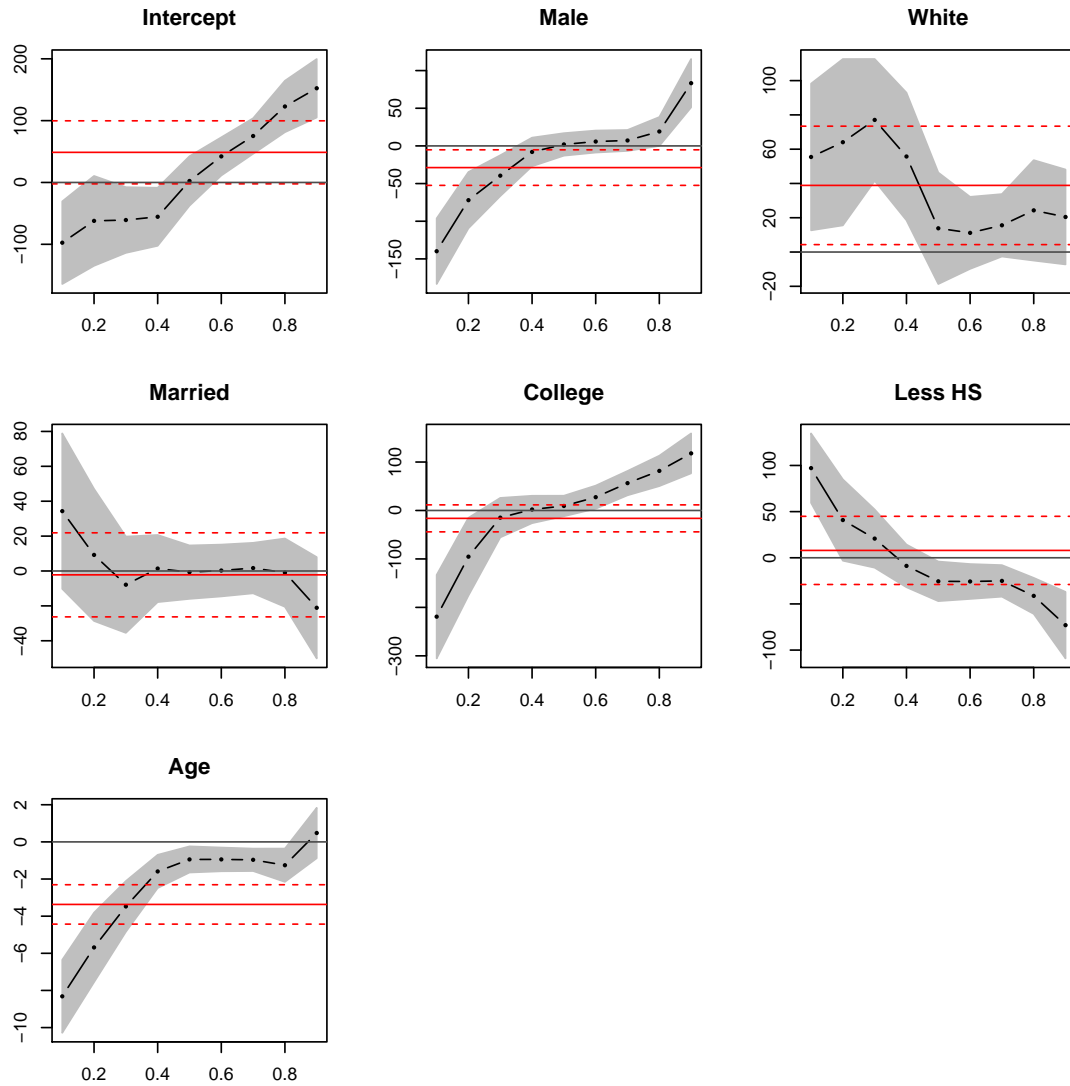
A.2 1997-1998 Results

Table 3: 1997-1998 Summary Statistics

	Displaced	Non-Displaced	Difference	P-val on Difference
Earnings				
1998 Earnings	443.38	621.07	-177.692	0.00
1997 Earnings	526.52	582.40	-55.886	0.00
Change Earnings	-83.14	38.67	-121.806	0.00
Covariates				
Male	0.56	0.50	0.057	0.00
White	0.87	0.88	-0.008	0.34
Married	0.55	0.66	-0.108	0.00
College	0.24	0.31	-0.069	0.00
Less HS	0.12	0.08	0.038	0.00
Age	38.69	42.03	-3.34	0.00
N	2387	3842		


Sources: CPS and Displaced Workers Survey

Figure 10: 1997-1998 Quantile Regression Estimates of the Treatment Effect on Covariates



Notes: QR estimates of the effect of covariates on the quantiles of the treatment effect in 1997-1998. The solid horizontal line provides OLS estimates of the effect of each covariate, and the horizontal dashed line contains a 90% confidence interval. The dotted line provides QR estimates of the effect of each covariate at particular quantiles from 0.1, 0.2, ..., 0.9. The shaded area contains pointwise 90% confidence intervals from the quantile regressions.

Figure 11: 1997-1998 Quantile Regression Estimates of the Treatment Effect on $Y_t(0)$



Notes: QR estimates of the effect of what earnings would have been in the absence of job displacement, $Y_t(0)$, on the quantiles of the treatment effect in 1997-1998. The solid horizontal line provides OLS estimates of the effect of non-displaced potential earnings, and the horizontal dashed line contains a 90% confidence interval. The dotted line provides QR estimates of the effect of non-displaced potential earnings at particular quantiles from 0.1, 0.2, ..., 0.9. The shaded area contains pointwise 90% confidence intervals from the quantile regressions.