

Strittmatter, Anthony

Conference Paper

What is the Value Added by using Causal Machine Learning Methods in a Welfare Experiment Evaluation?

Beiträge zur Jahrestagung des Vereins für Socialpolitik 2019: 30 Jahre Mauerfall - Demokratie und Marktwirtschaft - Session: Labor Economics VIII, No. E09-V1

Provided in Cooperation with:

Verein für Socialpolitik / German Economic Association

Suggested Citation: Strittmatter, Anthony (2019) : What is the Value Added by using Causal Machine Learning Methods in a Welfare Experiment Evaluation?, Beiträge zur Jahrestagung des Vereins für Socialpolitik 2019: 30 Jahre Mauerfall - Demokratie und Marktwirtschaft - Session: Labor Economics VIII, No. E09-V1, ZBW - Leibniz-Informationszentrum Wirtschaft, Kiel, Hamburg

This Version is available at:

<https://hdl.handle.net/10419/203499>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

What is the Value Added by using Causal Machine Learning Methods in a Welfare Experiment Evaluation?

Anthony Strittmatter
University of St. Gallen
First Draft: September 17, 2018

December 16, 2018

Abstract

I investigate causal machine learning (CML) methods to estimate effect heterogeneity by means of conditional average treatment effects (CATEs). In particular, I study whether the estimated effect heterogeneity can provide evidence for the theoretical labour supply predictions of Connecticut's Jobs First welfare experiment. For this application, [Bitler, Gelbach, and Hoynes \(2017\)](#) show that standard CATE estimators fail to provide evidence for theoretical labour supply predictions. Therefore, this is an interesting benchmark to showcase the value added by using CML methods. I report evidence that the CML estimates of CATEs provide support for the theoretical labour supply predictions. Furthermore, I document some reasons why standard CATE estimators fail to provide evidence for the theoretical predictions. However, I show the limitations of CML methods that prevent them from identifying all the effect heterogeneity of Jobs First.

Keywords: Labour supply, individualized treatment effects, conditional average treatment effects, random forest.

JEL classification: H75, I38, J22, J31, C21.

The Manpower Demonstration Research Corporation (MDRC) provided me with the experimental data used in this article. I remain solely responsible for how the data have been used and interpreted. The study is part of the Swiss National Science Foundation (SNSF) project "Causal Analysis with Big Data", grant number SNSF 407540_166999, and is included in the Swiss National Research Programme "Big Data" (NRP 75). A previous version of the paper was presented at UC Berkeley and University of St. Gallen. I thank participants, in particular Ulrich Glogowsky, Bryan Graham, Pat Kline, Michael Knaus, Michael Lechner, Davud Rostam-Afschar, and Michael Zimmert for helpful comments and suggestions. The usual disclaimer applies.

Addresses for correspondence: Anthony Strittmatter, Swiss Institute for Empirical Economic Research (SEW), University of St. Gallen, Varnbuelstr. 14, CH-9000 St. Gallen, Switzerland, Anthony.Strittmatter@unisg.ch, www.anthonystrittmatter.com.

1 Introduction

Many empirical studies pursue the estimation of heterogeneous policy effects with respect to exogenous covariates. Often, we call these effects conditional average treatment effects (CATEs). Recently, it has been proposed to use causal machine learning (CML) methods to estimate CATEs (see, e.g., [Athey, 2018](#), [Belloni, Chernozhukov, and Hansen, 2014](#), for reviews of the CML literature). A possible advantage of using CML methods is that they are able to incorporate many covariates that are potentially responsible for effect heterogeneity. Using relatively flexible CML methods to deal with these covariates makes it less likely that effect heterogeneity will be overlooked (e.g., with respect to interactions between covariates) compared to manual modelling. However, CML methods are to a large extent black-box approaches. It is unclear how much value CML methods can add to economic applications compared to more standard estimation methods.

In this study, I revisit the effects of Connecticut’s Jobs First welfare experiment on the labour supply. Well-established labour supply theory gives us clear predictions about the heterogeneity margins of this experiment (see, e.g., [Kline and Tartari, 2016](#), for a comprehensive summary). However, [Bitler, Gelbach, and Hoynes \(2017\)](#) document the limitations of a standard CATE estimator in terms of its ability to provide evidence for these theoretical predictions. This is puzzling because [Bitler, Gelbach, and Hoynes \(2006\)](#) show that quantile treatment effects (QTEs) can uncover evidence for the theoretical labour supply model. It appears that the Jobs First data contains relevant information that can support labour supply theory, but it is not straightforward how to uncover the appropriate heterogeneity by means of standard CATE estimators.¹

This study contributes to the aforementioned literature in at least four ways. First, I investigate whether CML methods can reveal more effect heterogeneity than standard CATE estimators. Second, I verify whether the revealed effect heterogeneity provides evidence for the theoretical labour supply predictions of the Jobs First programme. Third, I reveal some modelling restrictions that prevent the standard CATE estimators from revealing more effect heterogeneity. Fourth, I test whether the estimates of the CATEs and QTEs are nested.

[Bitler, Gelbach, and Hoynes \(2017\)](#) consider local constant models, which are one of the workhorse methods used to estimate CATEs in empirical economics. Local constant models stratify the sample in different groups defined by the covariates and report subgroups’ average treatment effects. Local constant models uncover effect heterogeneity across groups but report constant effects within groups. There are

¹Moreover, the surveys of [Grogger and Karoly \(2005\)](#) and [Michalopoulos and Schwartz \(2001\)](#) document that several other welfare programme evaluations, that rely on standard CATE estimators, find only little effect heterogeneity.

three potential reasons why local constant models fail to support labour supply theory.² First, the choice of the subgroups could be suboptimal. Second, constant effects within groups do not accurately approximate the (possibly continuous) treatment effects. Third, the covariates used in the local constant models may be insufficient for explaining effect heterogeneity. Suitable CML methods can overcome all three possible disadvantages of local constant models. It is essential to understand why local constant models fail to support theoretical predictions in the Jobs First case because they are widely used and the unreliability of local constant models may carry over to other applications.

My main analysis is based on the double machine learning approach proposed in [Chernozhukov, Chetverikov, Demirer, Duflo, Hansen, Newey, and Robins \(2018\)](#). This is a generic approach that can incorporate many different machine learning estimators, which means it can be used to compare machine learning estimators with different modelling restrictions. I consider off-the-shelf machine learning estimators that can be used with standard personal computers. Accordingly, these machine learning estimators can be accessed by a widespread audience. Furthermore, no special IT infrastructure is required, in contrast to machine learning methods that require cloud computing infrastructure (e.g., [TensorFlow](#)). This could be useful when dealing with confidential data, such as the MDRC data, because it is often easier to comply with data security regulations and laws using a personal computer than using applications on the cloud.

In particular, I consider the "tree" and the "random forest" machine learning estimators (see, e.g., [Hastie, Tibshirani, and Friedman, 2009](#)). Tree estimators split the data into mutually exclusive groups defined by the covariates and report effect heterogeneity by means of subgroups' average treatment effects. Similar to local constant models, tree estimators uncover effect heterogeneity across groups but report constant effects within groups. However, tree estimators employ data driven algorithms to select the subgroups, whereas subgroups have to be manually selected for local constant models. Both methods could, in principle, incorporate many covariates. However, it is more convenient to use tree estimators when the covariate space becomes large. Random forest estimators are ensemble methods. They estimate many trees based on different subsets of the data and covariates, and then, they report the average of the different tree estimates. This implies, random forest estimators relax additionally the restriction of constant effects within subgroups. Accordingly, the tree and random forest estimators are suitable for relaxing, in a stepwise fashion, the modelling restrictions of local constant models.

²I assume that the theoretical predictions are at least good approximations of the labour supply effects and measurement errors or other data problems do not prevent us from finding empirical support for the labour supply predictions.

My results suggest that tree and random forest estimators can provide evidence for the theoretical labour supply model when they incorporate at least a decent selection of covariates. This suggests including many relevant covariates is an important ingredient to establish the match between theoretical prediction and empirical results in the Jobs First application. Solely using data driven ways to select subgroups, without incrementing the considered relevant covariates, does not seem to improve the match. In particular, the way how [Bitler, Gelbach, and Hoynes \(2017\)](#) manually stratify the sample in groups cannot be outperformed by the CML algorithms that are based on the same covariates. The random forests show more stable results than the tree estimators.

Furthermore, I provide evidence that the CML estimates of the CATEs and QTEs provide disparate sets of information for evaluating the Jobs First case. This suggests the estimated CATEs do not uncover all the inherent effect heterogeneity of the Jobs First experiment. Both CATEs and QTEs have advantages and disadvantages. CATEs could be useful to design tailor made welfare schemes that optimise the labour supply response of specific target groups or to create assignment rules. QTEs enable to study responses on the entire labour supply distribution (in a fully flexible way), but it is difficult to assign these responses to specific groups.

In the next section, I provide some background information about the Jobs First welfare experiment. In [Section 3](#), I introduce the MDRC data. In [Section 4](#), I describe the empirical framework of this study. In [Section 5](#), I document the empirical results. In [Section 6](#), I discuss the relations between CATEs and QTEs. The final section concludes. The [Online Appendices A-F](#) provide supplementary descriptive statistics and results.

2 The Jobs First welfare programme

In 1996, Connecticut replaced the Aid for Families with Dependent Children (AFDC) with the Jobs First programme. The Jobs First programme created work incentives for people on assistance that differed from those offered by the AFDC programme.

[Figure 1](#) shows the earnings and welfare transfers in a stylised way. The maximum welfare payment W is similar under both welfare schemes, but additional earnings lead to different welfare payment deductions. AFDC recipients had a fixed earnings disregard B of \$120 per month during the first 12 months of employment while on assistance and \$90 per month afterwards. Furthermore, 51% of any additional earnings was disregarded during the first four months of employment and 27% of any additional earnings afterwards. In contrast, the Jobs First programme disregarded all earnings below the federal poverty line (FPL). Earnings above the poverty line would terminate all the welfare benefit payments of the Jobs First pro-

gramme (which is like a cliff in the benefit payment scheme). The Jobs First and AFDC programmes differ in other aspects besides the financial work incentives. The additional changes are summarised in Online Appendix [A](#).

Figure [1](#) around here

[Bitler, Gelbach, and Hoynes \(2006\)](#) use a static labour supply model to develop four hypotheses. First, the Jobs First programme deducts less earnings from welfare payments than the AFDC programme. Thus, Jobs First should have a positive effect on the extensive margin of the labour supply. Second, the Jobs First recipients with relatively low earnings (between B and E in Figure [1](#)) can keep more of their additional income than AFDC recipients. This should create positive work incentives when the substitution effects dominate the income effects. Third, the FPL is considerably higher than the earnings amount E, at which participants lose their eligibility for AFDC welfare benefit payments. Accordingly, this provides a lump-sum transfer to Jobs First welfare recipients with earnings between E and the FPL. This reduces the optimal earnings in the presence of negative income effects. Furthermore, the cliff construction of the Jobs First payment scheme creates incentives to reduce earnings to just below the FPL, which might increase income or leisure time. Fourth, for recipients sufficiently above the FPL, the AFDC and Jobs First programmes provide the same labour supply incentives.

To summarize, static labour supply theory predicts that Jobs First assignment has (a) positive effects on the extensive margin of labour supply, (b) positive effects far below the FPL, (c) negative effects slightly below and above the FPL, and (d) no effects far above the FPL. Accordingly, the two welfare programmes create a mix of positive and negative work incentives which makes it hard to empirically address effect heterogeneity.

3 Experimental data

The Connecticut Department of Social Services required the MDRC to conduct a randomized control trial to evaluate the Jobs First programme. Experimental participants were single-parent welfare applicants and recipients who lived in Manchester or New Haven. Between January 1996 and February 1997, 4,803 experimental participants were randomly assigned to either the AFDC (control group) or Jobs First (treatment group) programmes.³

³I drop one experimental participant that had extraordinary high earnings. However, this did not change the point estimates much but makes the estimation of the confidence intervals more stable.

3.1 Variable definitions

The MDRC’s public use files for the Jobs First programme contain baseline data on demographic and family composition variables merged with longitudinal administrative information on welfare and food stamps payments and earnings provided by the state unemployment insurance system. The outcome variable is earnings per quarter in US dollars.⁴ The treatment is assignment to the Jobs First programme. I follow 4,802 experimental participants for seven quarters after RA to the Jobs First programme or the AFDC programme. The total sample contains 33,614 observations.

Being above or below the FPL is one of the major factors that drives effect heterogeneity according to the theoretical considerations. Similar to previous Jobs First studies, I cannot observe the administrative assistance unit size, which determines the FPL of the experimental participant. Following the suggestions of [Kline and Tartari \(2016\)](#), I calculate the assistance unit size based on the number of children at RA.⁵ This may lead to the underestimation of the FPL because mothers can have more children during the seven quarters after RA. To account for this potential disadvantage, I inflate the number of children by one for all mothers as a robustness check (in the following, I call this calculation of the FPL ”extra child”).

I distinguish between three sets of covariates which I use for the heterogeneity analysis. I label them ”baseline”, ”decent”, and ”kitchen sink” covariates. [Table 1](#) summarizes the different covariate categories. The baseline covariates contain the elapsed quarters since RA and the earnings seven quarters prior to RA. The selection of the baseline covariates follows [Bitler, Gelbach, and Hoynes \(2017\)](#). They use these two covariates in some of their main results.

Table 1 around here

The decent covariates include 13 variables. In addition to the baseline variables, they include age, education, information about children, and more information about the earnings and welfare history. This is still a relatively small set of covariates; however, it may be difficult to consider even these relatively few variables with a basic estimator when including many non-linearities and interactions between the covariates.

The kitchen sink covariates include 68 variables. I include all exogenous covariates for which data are available from the MDRC and that possibly matter for effect

⁴I do not observe the earnings that are reported by the experimental participants to the welfare agency, which could matter when misreporting is a major practice (see discussion in, e.g., [Kline and Tartari, 2016](#)).

⁵I cannot calculate the assistance unit size for 160 experimental participants because the number of children is not reported.

heterogeneity. The kitchen sink covariates include different measures of the variables that are already included in the decent covariates. The additional variables are ethnicity, marital status, information about the residence, information on previous participation in education or labour market programmes, and more information on the earnings and welfare history.

3.2 Descriptive statistics

Table 2 reports the descriptive statistics of the main variables. The mean earnings do not differ greatly between the Jobs First and AFDC participants. Thus, the average effects of the Jobs First experiment do not differ significantly from zero (see, e.g., Bitler, Gelbach, and Hoynes, 2006). However, fewer participants in the Jobs First programme are unemployed compared to the participants in the AFDC programme. More participants in the Jobs First programme have employment with earnings below the FPL than the participants in the AFDC programme. The group of participants with earnings above the FPL is relatively small regardless of whether I inflate the number of children used to calculate the FPL or not. The share of participants with earnings above the FPL does not differ much between the Jobs First and AFDC participants.

Table 2 documents the standardised difference in the baseline and decent covariates between the Jobs First and AFDC participants.⁶ Table B.1 in the Online Appendix B shows the descriptive statistics of the kitchen sink covariates. If the RA to the Jobs First and AFDC programmes was appropriate, then we expect that all pre-RA covariates are balanced. Table 2 shows that there are no large differences between the pre-RA covariates. However, there are some small differences. Jobs First participants have slightly more children, less previous earnings, and have received more welfare than AFDC recipients.

Table 2 around here

⁶The standardized difference in variable X between samples A and B is defined as

$$SD = \frac{|\bar{X}_A - \bar{X}_B|}{\sqrt{\frac{1}{2}(\text{Var}(X_A) + \text{Var}(X_B))}} \cdot 100,$$

where \bar{X}_A denotes the mean of sample A and \bar{X}_B denotes the mean of sample B . Rosenbaum and Rubin (1983) consider an absolute standardized difference higher than 20 to be "large."

4 Empirical approach

4.1 Estimation target

The binary dummy D_i equals one when an experimental participant is assigned to the Jobs First programme and zero when she is assigned to the AFDC welfare scheme. Following Rubin's (1974) potential outcome framework, $Y_{it}(1)$ denotes the potential earnings outcome under Jobs First for individual i in quarter t (for $i = 1, \dots, N$ and $t = 1, \dots, 7$). Correspondingly, $Y_{it}(0)$ denotes the potential earnings outcome under AFDC for individual i in quarter t . Each individual can either receive welfare under the Jobs First or AFDC programmes but not simultaneously under both welfare schemes. This implies that only one potential outcome is observable. Under the stable unit treatment value assumption (SUTVA), the observed outcome equals

$$Y_{it} = Y_{it}(1)D_i + Y_{it}(0)(1 - D_i). \quad (1)$$

Individual i 's causal effect of being assigned to the Jobs First programme instead of the AFDC welfare scheme on earnings is

$$\delta_{it} = Y_{it}(1) - Y_{it}(0).$$

We cannot identify individual causal effects without assumptions that are implausible in many applications (e.g., the assumption of effect homogeneity). Nevertheless, group averages of δ_{it} may be identifiable under plausible assumptions. For example, the identification of the average treatment effect (ATE), $\rho = E[\delta_{it}]$, and the average treatment effect on the treated (ATET), $\theta = E[\delta_{it}|D_i = 1]$, are standard in policy evaluations (see, e.g., [Imbens and Wooldridge, 2009](#)). CATEs can potentially uncover effect heterogeneity based on exogenous pretreatment variables X_{it} . The CATEs are

$$\bar{\delta}(x) = E[\delta_{it}|X_{it} = x].$$

Under the random treatment assignment and SUTVA, the CATEs

$$\bar{\delta}(x) = E[Y_{it}|D_i = 1, X_{it} = x] - E[Y_{it}|D_i = 0, X_{it} = x],$$

are identified from observable data on (Y_{it}, D_i, X_{it}) .

The CATEs are often labelled individualised or personalised treatment effects. This is to some extent misleading because these labels might suggest that the CATEs can get very close to the individual causal effects. To achieve this, the individual causal effects have to be (almost) deterministic and we have to observe all relevant determining variables. In many applications, those requirements are incredible. Nev-

ertheless, even when the CATEs are not equal to the individual causal effects, they have the potential to give a more complete picture about the effect heterogeneity than the ATEs or ATETs.

4.2 CML approach

Several CML approaches to model effect heterogeneity have been proposed in the literature (see, e.g., [Chen, Tian, Cai, and Yu, 2017](#), [Imai and Ratkovic, 2013](#), [Knaus, Lechner, and Strittmatter, 2018a](#), [Nie and Wager, 2017](#), among others). In this study, I use the efficient score

$$Y_{it}^{\text{eff}} = \mu_1(X_{it}) - \mu_0(X_{it}) + \frac{D_i(Y_{it} - \mu_1(X_{it}))}{p(X_{it})} - \frac{(1 - D_i)(Y_{it} - \mu_0(X_{it}))}{1 - p(X_{it})},$$

proposed in [Robins and Rotnitzky \(1995\)](#). The efficient score includes the three nuisance parameters $\mu_1(X_{it}) = E[Y_{it}|D_i = 1, X_{it}]$, $\mu_0(X_{it}) = E[Y_{it}(0)|X_{it}] = E[Y_{it}|D_i = 0, X_{it}]$, and $p(X_{it}) = E[D_i|X_{it}]$. Each of these nuisance parameters can be estimated with methods that are suited to make predictions, such as machine learning estimators. The expected value of the efficient score is the ATE, $\rho = E[Y_{it}^{\text{eff}}]$. Because of orthogonality properties of the efficient score, estimates of the ATE remain consistent even when either $\mu_0(x)$ and $\mu_1(x)$ or $p(x)$ is misspecified. This double robustness to misspecification could be crucial in the case of overfitting.

[Chernozhukov et al. \(2018\)](#) call this the double machine learning (DML) approach, because it combines first step auxiliary predictions of nuisance parameters to estimate the ATE in the second step. They discuss how we can obtain \sqrt{N} -consistent and asymptotically normal estimates of the ATE. One important finding is, that even if the estimates of the nuisance parameter have a slow convergence rate, the ATE estimates can still converge with \sqrt{N} . However, when it comes to the function of the efficient score, such as the CATEs $\bar{\delta}(X_{it}) = E[Y_{it}^{\text{eff}}|X_{it}]$, we know much less about the asymptotic properties of the DML approach. [Lee, Okui, and Whang \(2017\)](#) show that the consistency results maintain for functionals, but they have to impose restrictive assumptions about the estimation of the nuisance parameters.

Another popular CML approach is the causal forest estimator (see [Athey, Tibshirani, and Wager, 2018](#), [Wager and Athey, 2018](#)). Under certain conditions, we know that causal forest estimators remain consistent and asymptotically normal even for the estimation of CATEs. The DML and causal forest approaches have similar good finite sample properties for estimating CATEs (see, e.g., [Knaus, Lechner, and Strittmatter, 2018b](#), and references therein). I use the DML approach in the main specifications. Additionally, I show that the results do not change considerably when considering the causal forest estimator.

Chernozhukov et al. (2018) suggest using a cross-fitting procedure to estimate the nuisance parameters. To implement this procedure, I partition the data in two random samples. I use the first sample to estimate the nuisance parameters. I extrapolate the fitted values of the nuisance parameters for the second sample. Then, I use the second sample to estimate the CATEs.⁷ For the forest estimators, I additionally switch the first and second samples and repeat the cross-fitting procedure. Then, I report the average CATEs obtained from the first and second sample.

4.3 Machine learning estimators

Many different machine learning estimators can be combined with the DML approach. These machine learning estimators include regularized regression (e.g., lasso, ridge, and elastic net), boosting, (deep) neural nets, and Bayesian learners (see, e.g., [Hastie, Tibshirani, and Friedman, 2009](#), for an overview). I focus on classification and regression tree (CART) and random forest estimators (e.g., [Breiman, 2001](#)) because they mimic the modelling restrictions of the local constant model. I use the R-packages `rpart` and `grf` to implement these estimators.

4.3.1 Tree estimator

CARTs partition the sample into mutually exclusive leaves. Let $\pi = \{l_1, \dots, l_{\#(\pi)}\}$ be a specific tree or sample partition, let $l_j \equiv l_j(x, \pi)$ be the respective leaf (for $j = 1, \dots, \#(\pi)$), and let $\#(\pi)$ be the number of leaves in tree π . The leaf $l_j(x, \pi)$ of tree π is a function of the covariate space of X_{it} such that $x \in l_j$. For an explicit example, consider that X_{it} contains only a binary indicator for gender. Then, we can choose between two possible trees; either we make no sample split, $\pi' = \{l_1\} = \{\text{men}, \text{women}\}$, or we partition men and women into two separate leaves, $\pi'' = \{l_1, l_2\} = \{\{\text{men}\}, \{\text{women}\}\}$.

Given tree π , we can estimate the CATEs by

$$\hat{\delta}(x, \pi) = \frac{1}{\sum_{i=1}^N \sum_{t=1}^7 1\{X_{it} \in l_j(x, \pi)\}} \sum_{i=1}^N \sum_{t=1}^7 1\{X_{it} \in l_j(x, \pi)\} \cdot \hat{Y}_{it}^{\text{eff}},$$

where $\hat{Y}_{it}^{\text{eff}}$ is the estimate of Y_{it}^{eff} , which I extrapolate from the retained cross-fitting sample. CARTs are built with a greedy algorithm, i.e., by recursively adding leaves to the tree that minimise the mean-squared-error (MSE) (e.g., [Breiman, Friedman, Stone, and Olshen, 1984](#)). Accordingly, CARTs aim to partition the sample into leaves with homogeneous outcomes. I select the optimal tree π^* and $\bar{\delta}_{\text{tree}}(x) =$

⁷In the clustered bootstrap procedure that I use to compute confidence intervals, I make sure that each individual can only enter one cross-fitting sample, but never both.

$\delta(x, \pi^*)$ based on the out-of-sample MSE, which I calculate with a 10-fold cross-validation procedure. To stabilise the trees, I impose the restriction that each leaf should contain at least 50 observations.

Following the suggestions of [Athey and Imbens \(2016\)](#), I use the so-called honest inference procedure, which means I split the sample in two parts of equal size. Then, I use the first partition to build the tree (training sample) and the second partition to estimate the CATEs (estimation sample).⁸

4.3.2 Generalised random forest estimator

Generalised random forests are assembled out of G decorrelated honest trees $\delta(x, \pi_g)$ (for $g = 1, \dots, G$). The decorrelated honest trees are estimated using different subsamples of the data and subsets of the covariates. The random forest estimator of the CATEs is the average of these honest trees,

$$\hat{\delta}_{RF}(x) = \frac{1}{G} \sum_{g=1}^G \hat{\delta}_{\text{tree}}(x, \pi_g).$$

The honest trees of a random forest are often build deep (i.e., with small terminal leaves). This means we no longer try to optimize the leaf size of the trees with the cross-validation procedure. Instead, we build many deep honest trees, which have a small bias but a large variance. The averaging across different honest trees reduces the variance (this is often called "bagging"). [Athey, Tibshirani, and Wager \(2018\)](#) explore the consistency and asymptotic normality of generalised random forests.

I build random forests with 1,000 decorrelated trees, each having a minimum leaf size of 10 observations.⁹ In each subsample, I select randomly 50% of the individuals and a random selection of two-thirds of the covariates.

5 Results

5.1 Baseline results

To create a benchmark, I first document some of the findings from the local constant model. Following one of the main specifications in [Bitler, Gelbach, and Hoynes \(2017\)](#), I stratify the data by previous earnings and elapsed quarters since RA. I classify the previous earnings seven quarters before RA into zero earnings ($earn_0$),

⁸In the clustered bootstrap procedure that I use to compute confidence intervals, I make sure that each individual can either enter the training or estimation sample, but never both.

⁹The number of trees G is an important tuning parameter for random forests. [Table C.1](#) in the [Online Appendix C](#) shows that the out-of-sample MSE improves when we increase the number of trees. However, with 100 trees, the prediction power of random forests is already almost saturated, and the MSE improvements are marginal.

earnings below the median among those with positive earnings ($earn_1$), and earnings above the median ($earn_2$). Furthermore, I create dummies for each elapsed quarter since RA (q_1, \dots, q_7). Then, I interact these dummy variables (including their interactions) with the treatment dummy (D), such that the model

$$Y = \sum_{k=0}^3 \sum_{t=1}^7 \gamma_{kt} \cdot earn_k \cdot q_t + \sum_{k=0}^3 \sum_{t=1}^7 \delta_{kt} \cdot D \cdot earn_k \cdot q_t,$$

is fully stratified. The CATEs are δ_{kt} .

It would be complex to report CATEs for every group, especially when I use many covariates. Therefore, I reduce the dimension of the CATEs. Figure 2 documents the CATEs of the local constant model by earnings percentile τ under AFDC. This is $\omega(\tau) = E[\bar{\delta}(X_{it}) | c_{floor}(\tau) \leq Y_{it}(0) < c_{ceil}(\tau)]$ with the percentile interval delimited by $c_{floor}(\tau)$ and $c_{ceil}(\tau)$. A closely related approach is endogeneous stratification which estimates $\omega^*(\tau) = E[Y_{it}(1) - Y_{it}(0) | c_{floor}(\tau) \leq E[Y_{it}(0) | X_{it}] < c_{ceil}(\tau)]$ (e.g., [Alberto, Chingos, and West, 2018](#)). While this approach uses variation in X_{it} to extrapolate $Y_{it}(0)$ to the treated observations, I use the variation in X_{it} to estimate the CATEs directly. Using the entire data to estimate $\bar{\delta}(X_{it})$, but only the control observations to estimate $\omega(\tau)$ prevents me from the additional extrapolation step.

Figure 2 starts at the 55th percentile because of the mass point at zero earnings (Figure B.1 in the Online Appendix B reports the cumulative potential earnings distributions). The two vertical lines report the location of the minimum and maximum FPL of the individuals in the sample. The theoretical predictions suggest that the CATEs are positive far below the FPL and negative above the FPL. The point estimates of the local constant model support these theoretical predictions, but we are unable to statistically distinguish the effects from zero.

Figure 2 around here

5.2 Results of the CML approach

5.2.1 Estimation of the nuisance parameters

For the estimation of the nuisance parameters, I always use the random forest estimator with the kitchen sink covariates. In this way, the results of the different CATE estimators do not depend on the specification of the nuisance parameters.

Figure B.2 in the Online Appendix B documents the histogram of the estimated Jobs First assignment probability. Under RA, we would expect no variation in the assignment probability. The assignment probability varies between 41% and 61%, with an average assignment probability of 50%. Even though the range of the es-

estimated treatment probability is narrow, it is far from homogeneous. However, the DML approach accounts for the differences in the treatment probabilities. Furthermore, we do not have to worry about common support problems, because the propensity score is far away from zero and one (see discussion in, e.g., [Lechner and Strittmatter, 2017](#)).

Figure [B.3](#) in the Online Appendix [B](#) reports the density of the estimated earnings under the AFDC and Jobs First programmes. The estimated earnings are always higher than zero, i.e., we fail to estimate the mass points at zero earnings (see Figure [B.1](#) in the Online Appendix [B](#)).

5.2.2 Tree estimates

Figure [3](#) documents the findings of the tree estimator with the baseline covariates. At all percentiles, the CATEs do not differ significantly from zero. This results suggests that the data-driven way to stratify the sample does not necessarily improve the match between the theoretical predictions and empirical results.

Figure [3](#) around here

For the CATEs in Figures [4](#) and [5](#), I stepwise increase the number of considered covariates. When I include at least the decent selection of covariates, the estimated effects have the sign that is predicted by the labour supply theory. I find positive effects far below the FPL and negative effects slightly below and above the FPL. However, the tree estimates are very unstable, and the point estimates are not centred in the confidence interval.

Figures [4](#) and [5](#) around here

In Figures [C.1](#), [C.2](#), and [C.3](#) in the Online Appendix [C](#), I report the relative MSE of the cross-validation samples. For the tree with the baseline controls, the relative MSE is almost flat. Regardless of how the tree estimator stratifies the sample, it never significantly outperforms the benchmark of homogeneous effects. The selected tree minimises the MSE with 26 final leaves or 25 splits (as opposed to 21 groups in the local constant model). Figure [C.4](#) in the Online Appendix [C](#) shows the structure of the tree with baseline controls. Of the 25 splits, 24 splits are based on the previous earnings and only one is based on the elapsed quarters. The relative MSE of the trees with the decent and kitchen sink covariates is never saturated (see Figures [C.2](#) and [C.3](#) in the Online Appendix [C](#)). Eventually, the tree estimators do not add additional leaves because I impose the restriction that each leaf must have 50 observations. This is an indication that the modelling restriction of the within-group constant effects is not appropriate in the Jobs First case. Figures

C.5 and C.6 in the Online Appendix C show the complex structure of the trees with the decent and kitchen sink covariates. The earnings and welfare history, as well as information about the children, are important split variables.

5.2.3 Random forest estimates

Figure 6 documents the results of the random forest estimator with the baseline controls. In comparison to the tree estimators, the confidence intervals are narrower and the point estimates more in the centre of the confidence intervals. The random forest finds positive effects far below the FPL, which aligns with the labour supply theory. However, the results lack clear evidence for negative CATEs above the FPL.

Figure 6 around here

Figure 7 reports the random forest estimated with the decent covariates. These findings clearly support the theoretical labour supply model. We find negative effects far below the FPL and positive effects above the FPL. Furthermore, the CATEs turn negative between the minimum and maximum FPL. Figure 8 indicates that the point estimates are very similar when we consider the kitchen sink covariates. However, it appears that the results are slightly less volatile when we use the kitchen sink covariates instead of the decent covariates. Overall, the random forest estimates are much more stable than the tree estimates.

Figures 7 and 8 around here

While the results in Figures 4, 5, 7, and 8 provide some evidence for the labour supply model, they do not provide support for all aspects of the model. In particular, these figures do not show positive effects on the extensive margin because of the mass point at zero (see Figure B.1 in the Online Appendix B). Furthermore, these figures do not provide evidence for negative effects slightly below the FPL and zero effects far below the FPL because the FPL is different for each individual.

5.2.4 Sorted CATEs

To go one step further, I report the sorted CATEs by different earnings categories. For this purpose, I follow the sorted predicted effects (SPEs) approach of Chernozhukov, Fernández-Val, and Luo (2018). The SPEs ($\kappa(\tau) = Q_{\hat{\delta}(X)}(\tau)$) report the quantiles τ of the CATEs. The SPE is the minimum CATE value such that at least the share τ of the CATE distribution lies below this value. In the following, I focus on the random forest estimator with the kitchen sink controls because this specification has shown the best performance until now.

Figure 9 shows the SPEs for the entire sample. The CATEs vary between -\$600 and \$400 from the 1% to the 99% quantile. Approximately 25% of the experimental participants have negative CATEs. Kline and Tartari (2016) estimate that at least 20% of the experimental participants opt into welfare under the Jobs First programme by decreasing their labour supply below the FPL. The CATE estimates are consistent with their findings.

Figure 9 around here

Figure 10 reports the SPEs of the CATEs for participants with zero earnings under AFDC. Accordingly, this figure enables us to investigate the effects on the extensive margin of the labour supply. The CATEs are dominantly positive for the experimental participants who would have zero earnings under AFDC. Consistent with the theory, this result suggests positive effects on the extensive margin of the labour supply. However, there is a very small percentage of individuals with negative CATEs, which lead to predicted negative earnings. Furthermore, it appears that the group with no effects on the extensive margin is too small. The potential earnings distributions in Figure B.1 in the Online Appendix B suggest that the share of people without an extensive margin response is much larger. It could be that we overlook important effect heterogeneity, because we do not have any covariates for labour demand. Possibly, the share of Jobs First participants who want to provide labour is larger than the share that actually finds a job.

Figure 10 around here

Figure 11 documents the SPEs for participants with positive earnings below the FPL under AFDC. For this group, the labour supply model predicts a mix of positive and negative CATEs. The empirical findings support the theoretical predictions.

Figure 11 around here

Figure 12 reports the SPEs of the CATEs for participants with earnings above the FPL under AFDC. Consistent with labour supply theory, the empirical results suggest negative and non-significant effects for this group. There is also a small fraction of experimental participants with positive CATEs. It could be that we find positive effects because we mismeasure the FPL. Therefore, Figure D.1 in the Online Appendix D reports additional results, which use the inflated the number of children per mother to calculate the FPL. The share of people with positive effects declines, but they do not disappear completely.

Figure 12 around here

To summarize, the CML approaches provide evidence for the labour supply model, but we have to incorporate at least a decent selection of covariates. Furthermore, the random forests appear more stable than the tree estimates. The results suggest positive effects on the extensive margin of the labour supply. The CATEs are positive far below the FPL, negative below and above the FPL, and zero above the FPL, which is consistent with the intensive margin effects predicted by the labour supply model. However, the CML approach reports a few CATEs that are definitely out of bounds. Furthermore, the results probably overestimates the effects on the extensive margin of the labour supply.

5.2.5 Additional results

Figure D.2 in the Online Appendix D reports the CATEs obtained from the causal forest estimator (e.g., [Athey, Tibshirani, and Wager, 2018](#)) with the kitchen sink controls. The results do not differ strongly from the estimates of the DML approach with the random forest.

The early CML literature suggests that the modified outcome method (without covariate adjustments)

$$Y_{it}^* = \frac{D_i - Pr(D_i = 1)}{Pr(D_i = 1)Pr(D_i = 0)} Y_{it}$$

is sufficient to estimate the CATEs, $\delta(X_{it}) = E[Y_i^*|X_{it}]$, in randomised experiments. Figure D.3 in the Online Appendix D documents that the modified outcome method (without covariate adjustments) fails to provide evidence for the theoretical predictions, even for the random forest estimator and kitchen sink controls. This finding is consistent with previous studies that document the poor properties of the modified outcome method (e.g., [Athey and Imbens, 2016](#)).

6 Relation between CATEs and QTEs

[Bitler, Gelbach, and Hoynes \(2006\)](#) use a QTE approach to evaluate the Jobs First programme. Their results also support the theoretical labour supply predictions (Figure D.4 in the Online Appendix D replicates their findings). This raises the following question: do CATEs and QTEs contain the same information?

6.1 Empirically testable conditions

In contrast to the CATEs, QTEs do not identify heterogeneity by subgroups of the population but rather on the potential earnings distributions.¹⁰ The potential outcome distributions are defined by

$$\begin{aligned} F_{Y(1)}(y) &= Pr(Y_{it}(1) \leq y) = Pr(Y_{it}(0) + \delta_{it} \leq y), \text{ and} \\ F_{Y(0)}(y) &= Pr(Y_{it}(0) \leq y) = Pr(Y_{it}(1) - \delta_{it} \leq y), \end{aligned} \quad (2)$$

with $Y_{it}(1) = Y_{it}(0) + \delta_{it}$. The potential quantile $Q_{Y(d)}(\tau)$ is the minimum value of $Y_{it}(d)$ such that at least the share τ of the earnings distribution lies below this value. QTEs are defined as

$$\delta_{QTE}(\tau) = Q_{Y(1)}(\tau) - Q_{Y(0)}(\tau),$$

the difference between the potential quantiles.

Bitler, Gelbach, and Hoynes (2017) propose the simulated potential outcome distributions,

$$\begin{aligned} F_{Y(1)}^S(y) &= Pr(Y_{it}(0) + \bar{\delta}(X_{it}) \leq y), \text{ and} \\ F_{Y(0)}^S(y) &= Pr(Y_{it}(1) - \bar{\delta}(X_{it}) \leq y). \end{aligned} \quad (3)$$

When

$$F_{Y(1)}^S(y) = F_{Y(1)}(y), \text{ and} \quad (4)$$

$$F_{Y(0)}^S(y) = F_{Y(0)}(y), \quad (5)$$

then CATEs and QTEs carry the same information and just differ in how they report the effects.¹¹ I exploit conditions (4) and (5) to determine whether CATEs and QTEs are nested. Moreover, (4) and (5) are necessary (but not sufficient) conditions for $\delta_{it} = \bar{\delta}(X_{it})$, as can be observed when comparing (2) and (3).

Bitler, Gelbach, and Hoynes (2017) exploit conditions (4) and (5) to test the null of within-group constant treatment effects, that vary across a limited number of observable subgroups. They suggest a testing procedure based on within-group permutations. This procedure cannot be applied to the random forest estimator, because it does not have a clear group structure. Therefore, I test (4) and (5) across all subgroups using a bootstrap based testing procedure. This procedure enables us to test the null hypothesis that CATEs and QTEs are nested. I provide the details

¹⁰Under very strong assumptions, the effects on the potential outcome distribution coincide with the individual causal effects. These assumption imply that individuals do not systematically change their rank in the potential outcome distributions as a result of the treatment status (see discussion in, e.g., Chernozhukov and Hansen, 2005, Firpo, 2007).

¹¹When CATEs and QTEs are nested, then it is always possible to calculate the QTEs from the CATEs using the simulated potential outcome distributions. However, it is not necessarily possible to calculate CATEs from QTEs.

of the distributional Kolmogorov-Smirnov (KS) test in the Online Appendix E.

6.2 Test results

Figure 13 reports the simulated earnings distribution using the random forest estimator with the kitchen sink covariates. The simulated earnings distributions are sometimes above and sometimes below the potential earnings distributions. However, there are two striking differences. First, the simulated distributions cannot detect the mass point at zero. Second, the simulated earnings for the lower ranks are negative. These issues are disadvantages of the CATEs compared to QTEs, that by construction cannot have these problems. Figures F.1-F.7 in the Online Appendix F report the simulated distributions of all other CML estimators.

Figure 13 around here

Table 3 shows the results of a Kolmogorov-Smirnov test for the equality of the potential and simulated earnings distributions. KS_1 tests condition (4), KS_0 tests condition (5), and KS_{\max} tests both conditions jointly. For all tests and all estimators, we have to reject the null that the two distributions are equal, which suggests that the QTEs and CATEs contain different information. Furthermore, this is evidence that the CATEs do not equal the individual causal effects.

Table 3 around here

Whether it is more appropriate to use CATEs or QTEs depends on the concrete research questions. For example, CATEs are more appropriate for developing assignment rules for programmes. QTEs are more appropriate for investigating earnings inequality when we are not concerned about the exact location of specific individuals in the earnings distribution.

7 Conclusions

I study the value added by using CML methods in a case study style using the example of Connecticut's Jobs First randomised welfare experiment. In this application, standard CATE estimators fail to find supporting evidence for the theoretical labour supply predictions. I provide evidence that CML methods can overcome this disadvantage of standard CATE estimators. Accordingly, CML methods add value to the welfare evaluation, in the sense that they provide evidence for the theoretical labour supply predictions. However, this works only when the CML methods incorporate many important heterogeneity variables. The random forest estimator show more

stable results than the tree estimator.

However, the CML methods are not able to uncover the entire effect heterogeneity of the Jobs First programme. Furthermore, CML methods have two shortcomings compared to QTEs. First, CML methods cannot detect mass points in the earnings distributions. Second, CML methods predict negative earnings for a small group. These shortcomings open directions for the further development of CML methods.

References

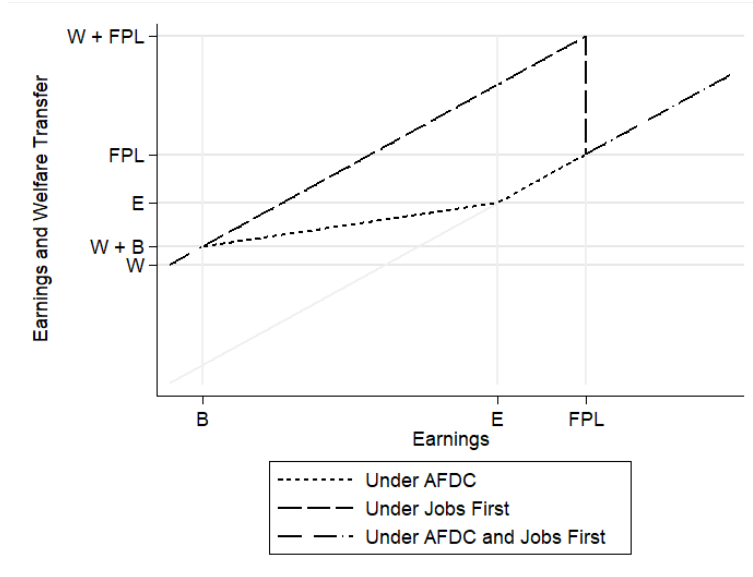
- ALBERTO, A., M. M. CHINGOS, AND M. R. WEST (2018): “Endogenous Stratification in Randomized Experiments,” *Review of Economics and Statistics*, 100(4), 567–580.
- ATHEY, S. (2018): “The Impact of Machine Learning on Economics,” in *The Economics of Artificial Intelligence: An Agenda*, ed. by A. Agrawal, J. Gans, and A. Goldfarb. NBER, Proceedings of Economics AI Conference.
- ATHEY, S., AND G. IMBENS (2016): “Recursive Partitioning for Heterogeneous Causal Effects,” *Proceedings of the National Academy of Science of the United States of America (PNAS)*, 113(27), 7353–7360.
- ATHEY, S., J. TIBSHIRANI, AND S. WAGER (2018): “Generalized Random Forests,” *Annals of Statistics*, forthcoming.
- BELLONI, A., V. CHERNOZHUKOV, AND C. HANSEN (2014): “High-Dimensional Methods and Inference on Treatment and Structural Effects in Economics,” *Journal of Economic Perspectives*, 28(2), 29–50.
- BITLER, M., J. GELBACH, AND H. HOYNES (2006): “What Mean Impacts Miss: Distributional Effects of Welfare Reform Experiments,” *American Economic Review*, 96(4), 988–1012.
- (2017): “Can Variation in Subgroups’ Average Treatment Effects Explain Treatment Effect Heterogeneity? Evidence from a Social Experiment,” *Review of Economics and Statistics*, 99(4), 683–697.
- BLOOM, D., S. SCRIVENER, C. MICHALOPOULOS, P. MORRIS, R. HENDRA, D. ADAMS-CIARDULLO, J. WALTER, AND W. VARGAS (2002): *Jobs First: Final Report on Connecticut’s Welfare Reform Initiative*. Manpower Demonstration Research Corporation, New York.
- BREIMAN, L. (2001): “Random Forests,” *Machine Learning*, 45(1), 5–32.

- BREIMAN, L., J. FRIEDMAN, C. STONE, AND R. OLSHEN (1984): *Classification and Regression Trees*. CRC Press.
- CHEN, S., L. TIAN, T. CAI, AND M. YU (2017): “A General Statistical Framework for Subgroup Identification and Comparative Treatment Scoring,” *Biometrics*, 73(4), 1199–1209.
- CHERNOZHUKOV, V., D. CHETVERIKOV, M. DEMIRER, E. DUFLO, C. HANSEN, W. NEWEY, AND J. ROBINS (2018): “Double/Debiased Machine Learning for Treatment and Structural Parameters,” *The Econometrics Journal*, 21(1), C1–C68.
- CHERNOZHUKOV, V., I. FERNÁNDEZ-VAL, AND Y. LUO (2018): “The Sorted Effects Method: Discovering Heterogeneous Effects Beyond their Averages,” *Econometrica*, forthcoming.
- CHERNOZHUKOV, V., AND C. HANSEN (2005): “An IV Model of Quantile Treatment Effects,” *Econometrica*, 73(1), 245–261.
- FIRPO, S. (2007): “Efficient Semiparametric Estimation of Quantile Treatment Effects,” *Econometrica*, 75(1), 259–276.
- GROGGER, J. T., AND L. A. KAROLY (2005): *Welfare Reform: Effects of a Decade of Change*. Harvard University Press.
- HASTIE, T., R. TIBSHIRANI, AND J. FRIEDMAN (2009): *Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, 2nd edn.
- IMAI, K., AND M. RATKOVIC (2013): “Estimating Treatment Effect Heterogeneity in Randomized Program Evaluation,” *Annals of Applied Statistics*, 7(1), 443–470.
- IMBENS, G., AND J. WOOLDRIDGE (2009): “Recent Developments in the Econometrics of Program Evaluation,” *Journal of Economic Literature*, 47(1), 5–86.
- KLINE, P., AND M. TARTARI (2016): “Bounding the Labor Supply Responses to a Randomized Welfare Experiment: A Revealed Preference Approach,” *American Economic Review*, 106(4), 972–1014.
- KNAUS, M., M. LECHNER, AND A. STRITTMATTER (2018a): “Heterogeneous Employment Effects of Job Search Programs: A Machine Learning Approach,” *arXiv:1709.10279*.
- (2018b): “Machine Learning Estimation of Heterogeneous Causal Effects: Empirical Monte Carlo Evidence,” *arXiv:1810.13237*.

- LECHNER, M., AND A. STRITTMATTER (2017): “Practical Procedures to Deal with Common Support Problems in Matching Estimation,” *Econometric Reviews*, forthcoming.
- LEE, S., R. OKUI, AND Y.-J. WHANG (2017): “Doubly Robust Uniform Confidence Band for the Conditional Average Treatment Effect Function,” *Journal of Applied Econometrics*, 32(7), 1207–1225.
- MICHALOPOULOS, C., AND C. SCHWARTZ (2001): *What Works Best for Whom: Impacts of 20 Welfare-to-Work Programs by Subgroup*. Manpower Demonstration Research Corporation.
- NIE, X., AND S. WAGER (2017): “Quasi-Oracle Estimation of Heterogeneous Treatment Effects,” *Working Paper*.
- ROBINS, J. M., AND A. ROTNITZKY (1995): “Semiparametric Efficiency in Multivariate Regression Models with Missing Data,” *Journal of the American Statistical Association*, 90(429), 122–129.
- ROSENBAUM, P. R., AND D. B. RUBIN (1983): “The Central Role of Propensity Score in Observational Studies for Causal Effects,” *Biometrika*, 70(1), 41–55.
- RUBIN, D. (1974): “Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies,” *Journal of Educational Psychology*, 66(5), 688–701.
- WAGER, S., AND S. ATHEY (2018): “Estimation and Inference of Heterogeneous Treatment Effects Using Random Forests,” *Journal of the American Statistical Association*, 113(523), 1228–1242.

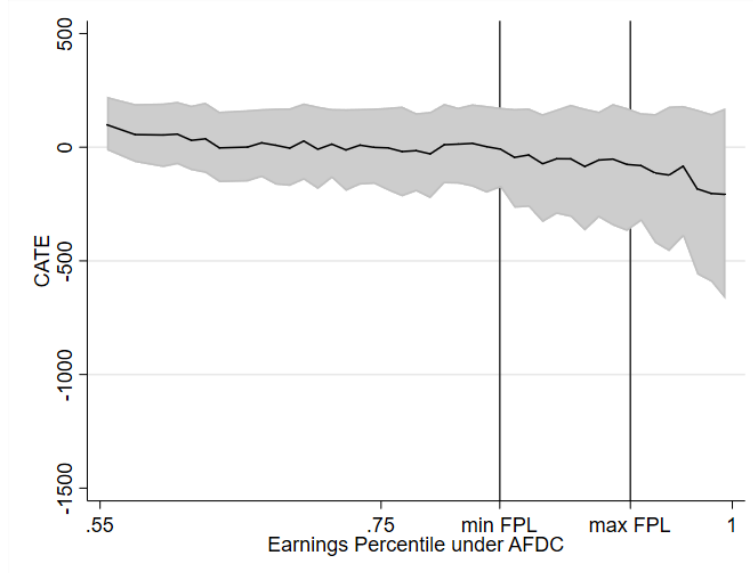
Figures

Figure 1: Earnings and welfare transfer under AFDC and Jobs First.



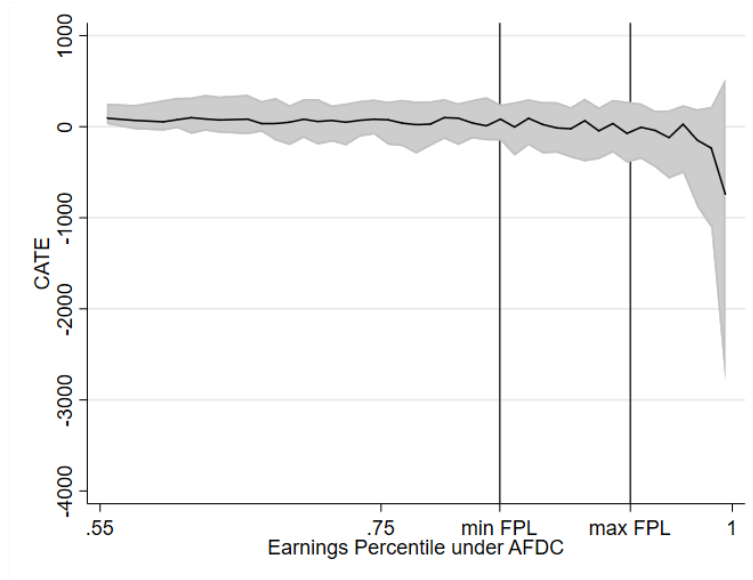
Note: Unemployed persons receive the welfare amount W under AFDC and Jobs First. AFDC welfare recipients have a fixed earnings disregard B . Any earnings above B reduce the welfare amount proportionally. The welfare is completely terminated at earnings E . Under Jobs First, all earnings below the FPL are disregarded. Any earnings above the FPL terminates welfare payments completely.

Figure 2: Local constant model CATEs by earnings percentile under AFDC.



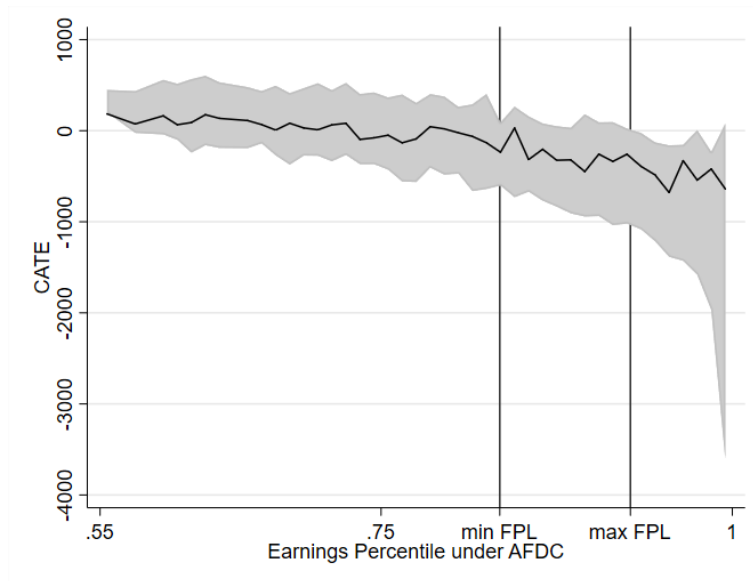
Note: I estimate the distribution of the CATEs using an individual-level clustered bootstrap approach. The gray area reports the 99%-confidence interval based on the percentile method. I control for the baseline variables following the sample stratification proposed in [Bitler, Gelbach, and Hoynes \(2017\)](#).

Figure 3: Tree CATEs (baseline controls) by earnings percentile under AFDC.



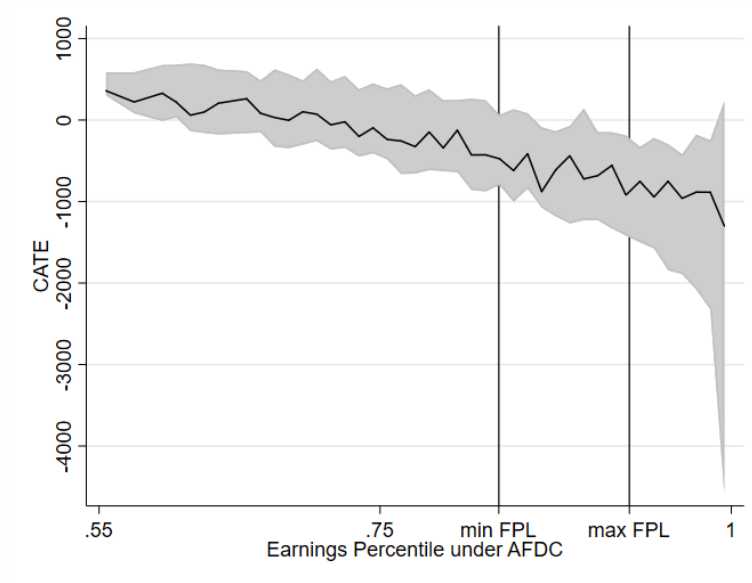
Note: I estimate the distribution of the CATEs using an individual-level clustered bootstrap approach with 2,000 replications. The gray area reports the 99%-confidence interval based on the percentile method.

Figure 4: Tree CATEs (decent controls) by earnings percentile under AFDC.



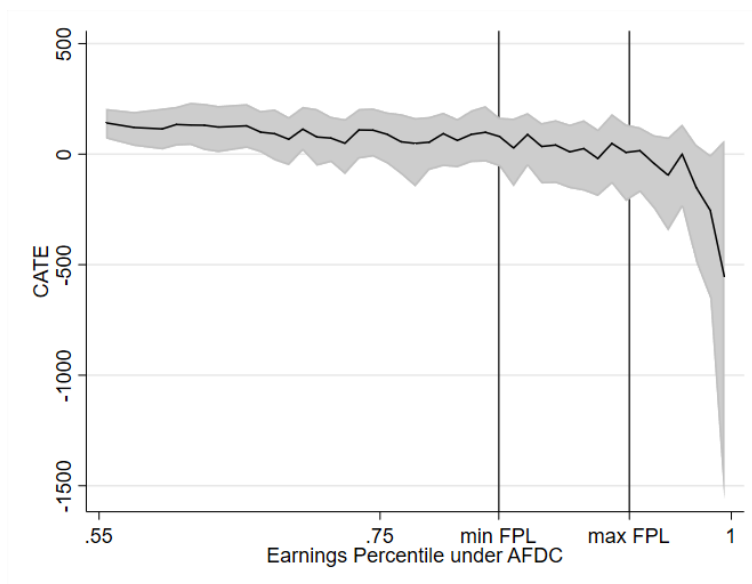
Note: I estimate the distribution of the CATEs using an individual-level clustered bootstrap approach with 2,000 replications. The gray area reports the 99%-confidence interval based on the percentile method.

Figure 5: Tree CATEs (kitchen sink controls) by earnings percentile under AFDC.



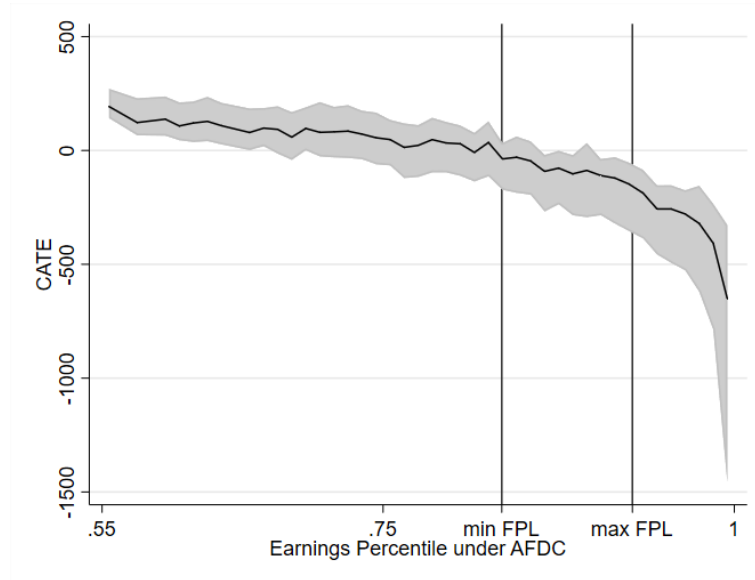
Note: I estimate the distribution of the CATEs using an individual-level clustered bootstrap approach with 2,000 replications. The gray area reports the 99%-confidence interval based on the percentile method.

Figure 6: Random forest CATEs (baseline controls) by earnings percentile under AFDC.



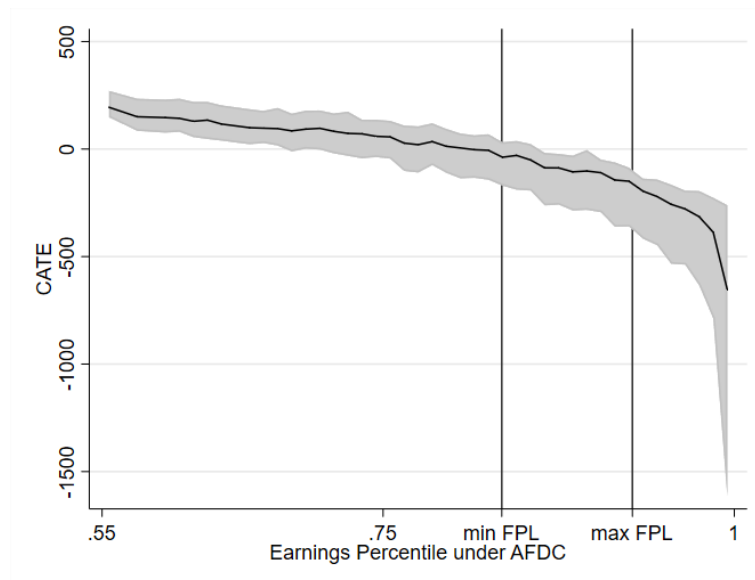
Note: I estimate the distribution of the CATEs using an individual-level clustered bootstrap approach with 2,000 replications. The gray area reports the 99%-confidence interval based on the percentile method.

Figure 7: Random forest CATEs (decent controls) by earnings percentile under AFDC.



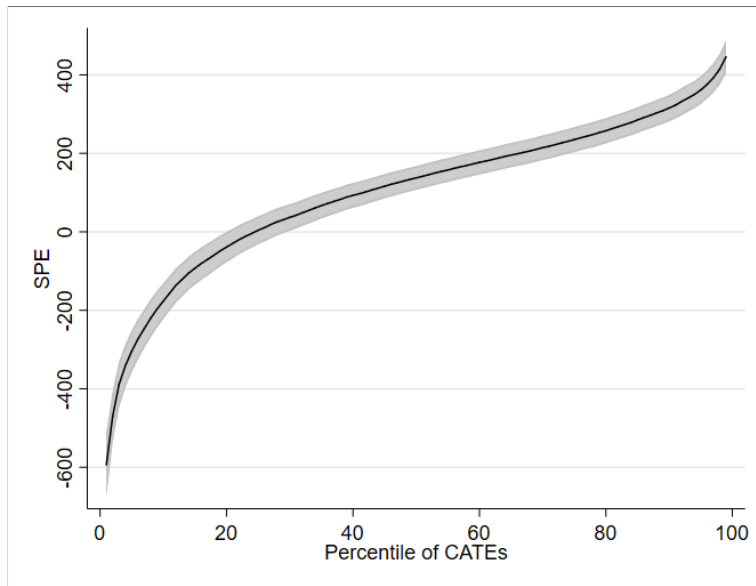
Note: I estimate the distribution of the CATEs using an individual-level clustered bootstrap approach with 2,000 replications. The gray area reports the 99%-confidence interval based on the percentile method.

Figure 8: Random forest CATEs (kitchen sink controls) by earnings percentile under AFDC.



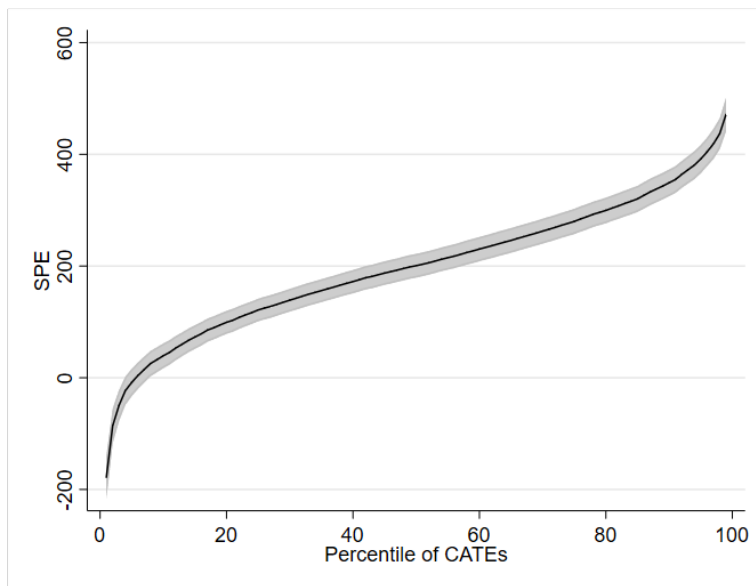
Note: I estimate the distribution of the CATEs using an individual-level clustered bootstrap approach with 2,000 replications. The gray area reports the 99%-confidence interval based on the percentile method.

Figure 9: Sorted predicted effects of the CATEs.



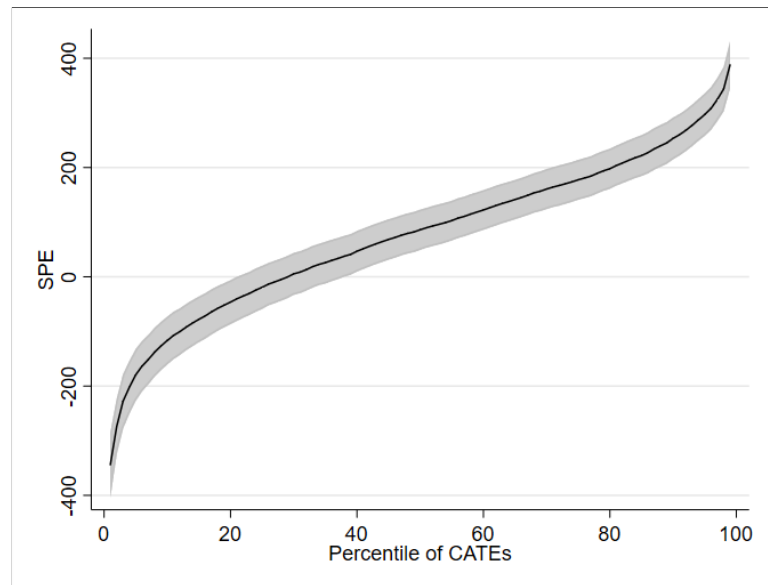
Note: I estimate the 99%-confidence interval (gray area) using the bootstrap Algorithm 2.1 of [Chernozhukov, Fernández-Val, and Luo \(2018\)](#) with 2,000 replications.

Figure 10: Sorted predicted effects of the CATEs for participants with zero earnings under AFDC.



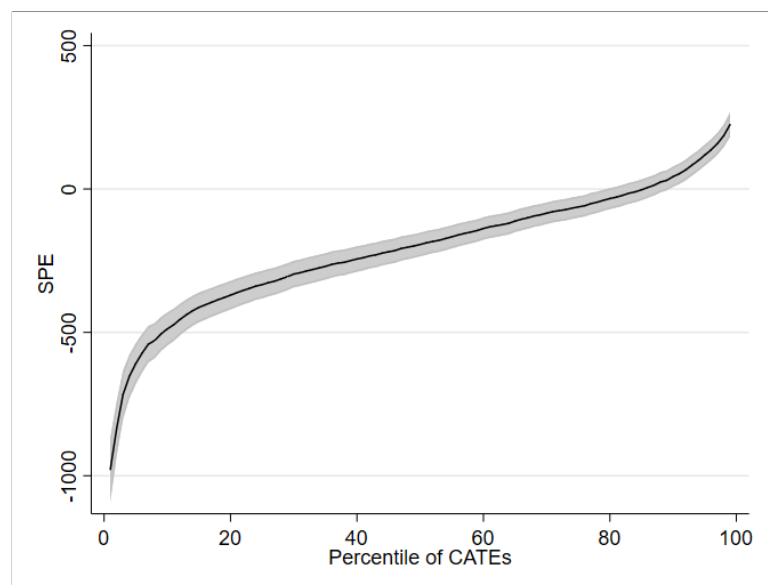
Note: I estimate the 99%-confidence interval (gray area) using the bootstrap Algorithm 2.1 of [Chernozhukov, Fernández-Val, and Luo \(2018\)](#) with 2,000 replications.

Figure 11: Sorted predicted effects of the CATEs for participants with positive earnings below the FPL under AFDC.



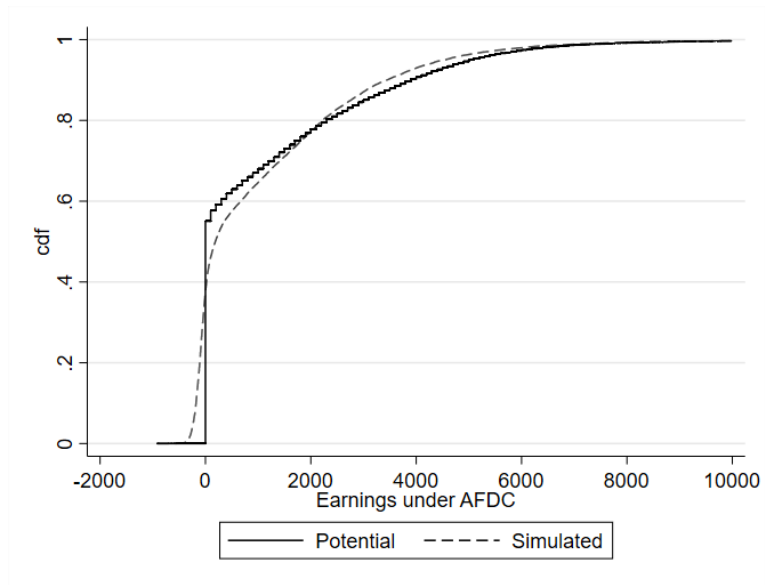
Note: I estimate the 99%-confidence interval (gray area) using the bootstrap Algorithm 2.1 of [Chernozhukov, Fernández-Val, and Luo \(2018\)](#) with 2,000 replications.

Figure 12: Sorted predicted effects of the CATEs for participants with earnings above the FPL under AFDC.

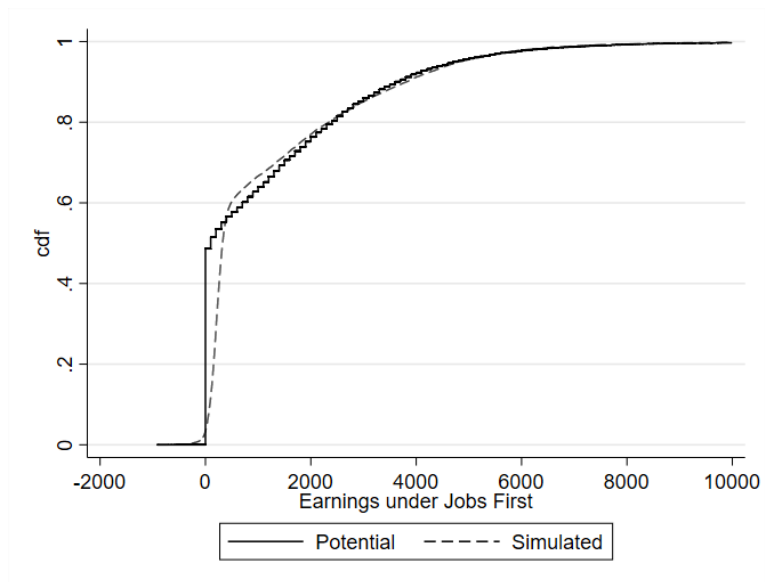


Note: I estimate the 99%-confidence interval (gray area) using the bootstrap Algorithm 2.1 of [Chernozhukov, Fernández-Val, and Luo \(2018\)](#) with 2,000 replications.

Figure 13: Potential and simulated earnings distributions obtained from the random forest estimator and kitchen sink controls.



(a) Under AFDC



(b) Under Jobs First

Tables

Table 1: List of covariates.

Baseline:	elapsed quarters since RA, earnings seven quarters prior to RA
Decent:	age, education, number of children, age of the youngest child, amount of AFDC assistance received seven quarters prior to RA, amount of food stamps received seven quarters prior to RA, dummy variable indicating a positive amount of earnings in at least one of the seven quarter prior to RA, dummy variable indicating a positive amount of AFDC assistance received in at least one of the seven quarters prior to RA, dummy variable indicating a positive amount of food stamps in at least one of the seven quarters prior to RA
Kitchen sink:	ethnicity, marital status, city of residence, information on living in a publicly subsidised home, information on relocations, participation in different types of education and labour market programmes in the 12 months prior to RA (e.g., English as a Secondary Language (ESL), Adult Basic Education (ABE), General Education Development (GED), job readiness skills, work experience, vocational education, post secondary education, and high school), earnings for each of the seven quarters prior to the RA, amount of AFDC assistance received for each of the seven quarters prior to the RA, amount of food stamps received for each of the seven quarters prior to RA, the number of quarters on AFDC, dummy variable indicating whether the family received AFDC during childhood, dummy variable indicating whether work was never recorded, dummy variable indicating whether work was recorded at RA

Note: The decent covariates also include the baseline covariates. The kitchen sink covariates also include the decent covariates. I include dummies for missing values whenever necessary (see Table B.2 in the Online Appendix B for details).

Table 2: Descriptive statistics of the main variables.

	Jobs First		AFDC		SD
	Mean	St. Dev.	Mean	St. Dev.	
	(1)	(2)	(3)	(4)	(5)
Earnings per quarter (in \$)	1173	1789	1125	1868	2.6
Share of participants with					
no earnings	0.49	0.50	0.55	0.50	13.3
earn. below FPL	0.39	0.49	0.31	0.46	17.0
earn. above FPL	0.13	0.33	0.14	0.35	4.1
earn. above FPL (extra child)	0.09	0.28	0.10	0.30	4.4
Baseline covariates					
Quarters since RA	4.0	2.0	4.0	2.0	0.0
Earnings in pre-Q7 (in \$)	682	1552	774	1781	5.5
Decent covariates					
Age categories					
< 20 years	0.09	0.28	0.09	0.28	1.2
20-24 years	0.20	0.40	0.21	0.41	2.7
25-34 years	0.41	0.49	0.42	0.49	1.5
35-45 years	0.25	0.43	0.23	0.42	4.1
> 44 years	0.05	0.22	0.06	0.23	1.2
Education categories					
No degree	0.33	0.47	0.31	0.46	3.8
High school	0.55	0.50	0.57	0.50	3.2
More than high school	0.06	0.24	0.06	0.23	1.9
Age youngest child (in years)	4.6	4.7	4.5	4.8	2.2
Number of children	1.6	1.0	1.5	1.0	6.0
AFDC pre-Q7 (in \$)	920	925	865	896	6.0
Food stamps pre-Q7 (in \$)	306	319	293	301	4.4
Any earnings pre-Q1/7	0.33	0.37	0.36	0.38	7.9
Any AFDC pre-Q1/7	0.57	0.45	0.54	0.45	6.5
Any food stamps pre-Q1/7	0.61	0.44	0.60	0.43	2.1
Participants	2,396		2,406		
Observations	16,772		16,842		

Note: The last column reports the standardised difference (SD). Earnings in pre-Q7 refers to earnings in the seven quarters before RA. Any earnings pre-Q1/7 is a dummy variable indicating that earnings were positive in at least one of the seven quarters prior to RA. FPL (extra child) means that I inflate the number of children per mother by one when calculating the assistance unit size.

Table 3: Kolmogorov-Smirnov test for equality between potential and simulated earnings distributions.

	KS_{\max}		KS_1		KS_0	
	test stat.	p-value	test stat.	p-value	test stat.	p-value
	(1)	(2)	(3)	(4)	(5)	(6)
Baseline	0.412	0.000	0.364	0.000	0.412	0.000
Tree						
Baseline	0.467	0.010	0.438	0.020	0.412	0.000
Decent	0.301	0.000	0.270	0.000	0.309	0.000
Kitchen sink	0.311	0.000	0.243	0.000	0.311	0.000
Random forest						
Baseline	0.464	0.000	0.432	0.000	0.463	0.000
Decent	0.405	0.000	0.404	0.000	0.329	0.000
Kitchen sink	0.452	0.000	0.452	0.000	0.375	0.000
Causal forest						
Kitchen sink	0.443	0.000	0.443	0.000	0.418	0.000

Note: See the details of the test in the Online Appendix E. The p-values are bootstrapped with 2,000 replications. KS_1 tests condition (4), KS_0 tests condition (5), and KS_{\max} tests both conditions jointly.

Online Appendix to "What is the Value Added by using Causal Machine Learning Methods in a Welfare Experiment Evaluation?"

Anthony Strittmatter
University of St.Gallen

Sections:

- A. Additional changes of Jobs First
- B. Additional descriptive statistics
- C. Specification of the tuning parameters
- D. Similarity between CATEs and QTEs
- E. Simulated earnings distributions

A Additional changes of Jobs First

The Jobs First and AFDC programmes differ in other aspects besides the financial work incentives. AFDC participants receive welfare benefits for an unlimited amount of time, whereas the lifetime benefit duration of Jobs First recipients is limited to 21 months. However, Bloom, Scrivener, Michalopoulos, Morris, Hendra, Adams-Ciardullo, Walter, and Vargas (2002) document that extensions of this time limit were possible and common. To account for possible time limit differences, I focus on the short-term effects of the Jobs First programme on earnings during the first 21 months after random assignment (RA).

The recipients of both the Jobs First and AFDC programmes had work requirements. Unemployed AFDC recipients could meet the work requirements by participating in an employment-related programme. Bloom et al. (2002) argue that AFDC work requirements were only casually enforced. Unemployed Jobs First recipients could meet the work requirements by participating in employment programmes meant to reduce the time taken for job placement. Moreover, Jobs First implemented stricter sanction rules in the case of non-compliance with the work requirement.

Furthermore, Jobs First recipients had different access to child care subsidies and Medicaid than AFDC recipients. However, Bloom et al. (2002) argue that these differences had little impact on actual child and health care availability because contemporaneous state-level programmes compensated for the differences. Finally, the reform affected how alimony payments from fathers were deducted from welfare payments but only when the alimony payments were between \$50 and \$100. Kline and Tartari (2016) argue that these differences are negligible because they only affect a small fraction of mothers.

B Additional descriptive statistics

Table B.1: Descriptive statistics of the kitchen sink covariates.

	Jobs First		AFDC		SD
	Mean	St. Dev.	Mean	St. Dev.	
	(1)	(2)	(3)	(4)	
Kitchen sink covariates					
Ethnicity					
Black	0.37	0.48	0.37	0.48	0.6
White	0.36	0.48	0.35	0.48	2.9
Hispanic	0.21	0.41	0.22	0.41	2.3
Other ethnicity	0.06	0.24	0.06	0.25	0.6
Marital status					
Never married	0.62	0.48	0.63	0.48	1.4
Married	0.33	0.47	0.32	0.47	1.4
Separated	0.06	0.23	0.05	0.22	2.5
Divorced	0.11	0.32	0.12	0.33	3.2
Widowed	0.01	0.09	0.01	0.11	3.2
City of residence					
New Haven	0.75	0.43	0.76	0.43	0.9
Manchester	0.25	0.43	0.24	0.43	0.9
Housing information					
Public housing	0.11	0.31	0.10	0.30	1.4
Subsidised housing	0.23	0.42	0.22	0.42	1.3
Temporary housing	0.01	0.11	0.01	0.10	1.2
Other housing	0.65	0.48	0.66	0.47	2.3
Num. moves last 2 years	0.91	1.23	0.95	1.22	3.0
Any programme or educ. during last year	0.20	0.40	0.20	0.40	0.8
ESL	0.01	0.08	0.01	0.09	2.9
ABE	0.01	0.08	0.01	0.09	1.4
GED	0.03	0.16	0.03	0.16	0.5
Vocational education	0.03	0.16	0.03	0.17	3.0
Post secondary educ.	0.06	0.23	0.05	0.21	4.8
Job search programme	0.01	0.10	0.01	0.10	0.8
Work experience programme	0.06	0.24	0.07	0.25	3.5
High school	0.02	0.14	0.02	0.14	1.2
Job readiness skills	0.01	0.08	0.01	0.08	0.5

< Table continues on next page. >

Table B.1 continued.

	Jobs First		AFDC		SD
	Mean	St. Dev.	Mean	St. Dev.	
	(1)	(2)	(3)	(4)	(5)
Employment history					
Earnings in pre-Q6	727	1932	814	1759	4.7
Earnings in pre-Q5	683	1506	840	1769	9.6
Earnings in pre-Q4	695	1522	845	1776	9.1
Earnings in pre-Q3	682	1440	791	1620	7.1
Earnings in pre-Q2	678	1513	726	1495	3.1
Earnings in pre-Q1	632	1368	689	1399	4.1
Never worked	0.12	0.32	0.11	0.31	2.3
Employed at RA	0.21	0.41	0.23	0.42	5.3
Welfare history					
AFDC pre-Q6	916	924	862	904	5.9
AFDC pre-Q5	899	905	837	895	6.9
AFDC pre-Q4	887	879	820	872	7.7
AFDC pre-Q3	862	845	801	824	7.3
AFDC pre-Q2	869	812	822	802	5.8
AFDC pre-Q1	884	786	837	795	6.0
Food stamps pre-Q6	309	326	301	311	2.3
Food stamps pre-Q5	323	337	312	321	3.3
Food stamps pre-Q4	349	351	333	340	4.7
Food stamps pre-Q3	370	361	357	349	3.7
Food stamps pre-Q2	394	366	383	360	3.3
Food stamps pre-Q1	413	360	396	353	4.8
Quarters on AFDC at RA	4.1	2.3	4.1	2.2	3.6
AFDC during childhood	0.23	0.42	0.24	0.43	2.3

Note: The last column reports the standardised difference (SD). Earnings, AFDC and food stamp payments are measured in US dollars. Earnings in pre-Q6 means earnings six quarters before RA. ESL is the abbreviation for English Secondary Language. ABE is the abbreviation for Adult Basic Education. GED is the abbreviation for General Education Development.

Table B.2: Balance of missing dummies.

	Jobs First		AFDC		SD
	Mean	St. Dev.	Mean	St. Dev.	
	(1)	(2)	(3)	(4)	
Missing Dummies					
Education	0.06	0.23	0.06	0.24	2.5
Children	0.12	0.32	0.14	0.35	7.7
Marital status	0.05	0.21	0.05	0.21	0.3
Num. moves last 2 years	0.05	0.21	0.04	0.21	0.7
programme or educ. during last year	0.06	0.23	0.05	0.22	2.1
Never worked	0.03	0.18	0.03	0.18	0.9
Employed at RA	0.04	0.20	0.05	0.21	1.5
Quarters on AFDC at RA	0.05	0.23	0.05	0.22	1.0
AFDC during childhood	0.07	0.25	0.06	0.25	0.3

Note: The last column reports the standardised difference (SD).

Figure B.1: Cumulative potential earnings distributions under AFDC and Jobs First.

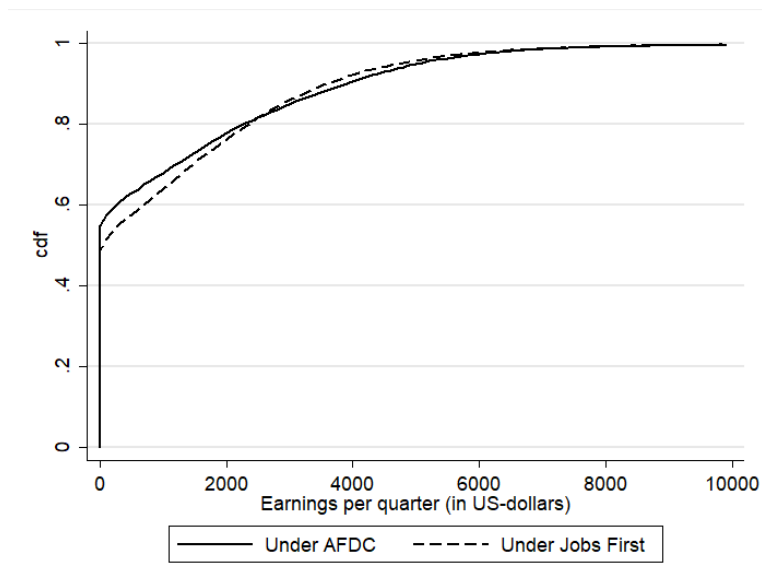
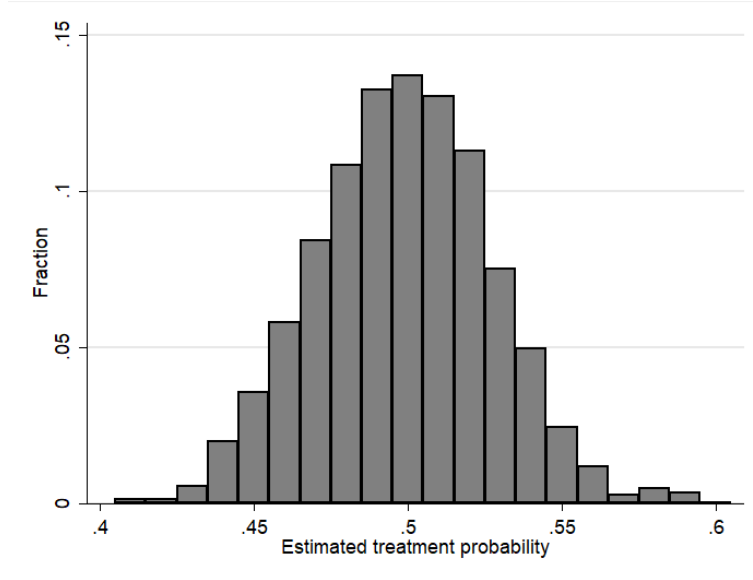
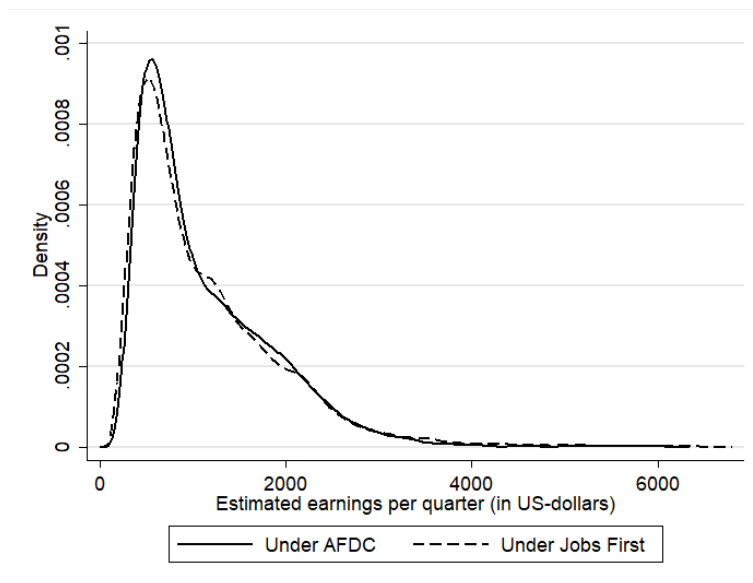


Figure B.2: Histogram of the estimated treatment probability.



Note: I use the random forest estimator with the kitchen sink covariates.

Figure B.3: Density of estimated earnings under AFDC and Jobs First.



C Specification of the tuning parameters

Table C.1: MSE by forest size.

	Tree size of forest							
	Single tree	10	50	100	500	1,000	5,000	10,000
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
P-score		0.506	0.502	0.500	0.500	0.500	0.500	0.500
$Y(0)$		1551	1540	1537	1532	1532	1531	1532
$Y(1)$		1563	1537	1536	1531	1531	1530	1531
Y^{eff} , baseline	3059	3190	3013	3026	3002	3001	2999	3000
Y^{eff} , decent	3252	3228	3015	3029	3002	2999	2997	2998
Y^{eff} , kitchen sink	3353	3221	3013	3024	2995	2995	2991	2992

Figure C.1: Relative MSE for the tree with baseline covariates.

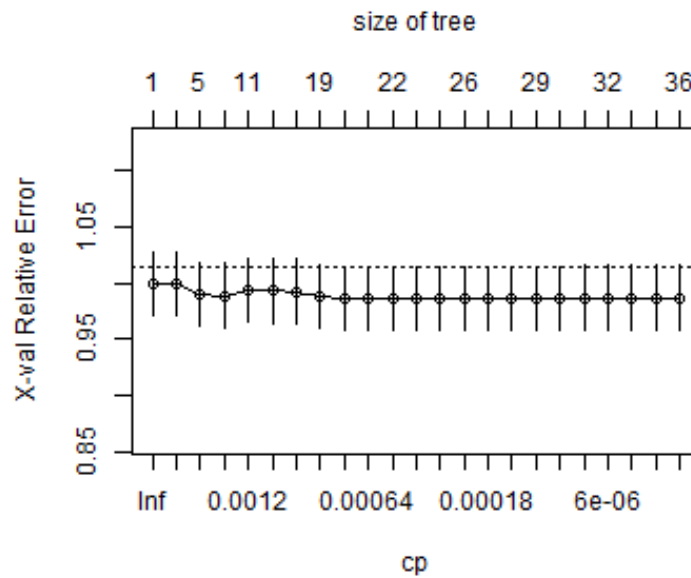


Figure C.2: Relative MSE for the tree with decent covariates.

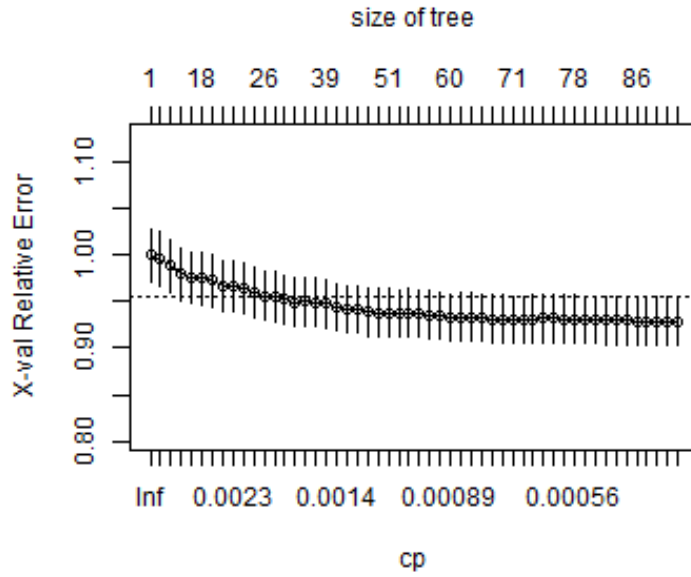


Figure C.3: Relative MSE for the tree with kitchen sink covariates.

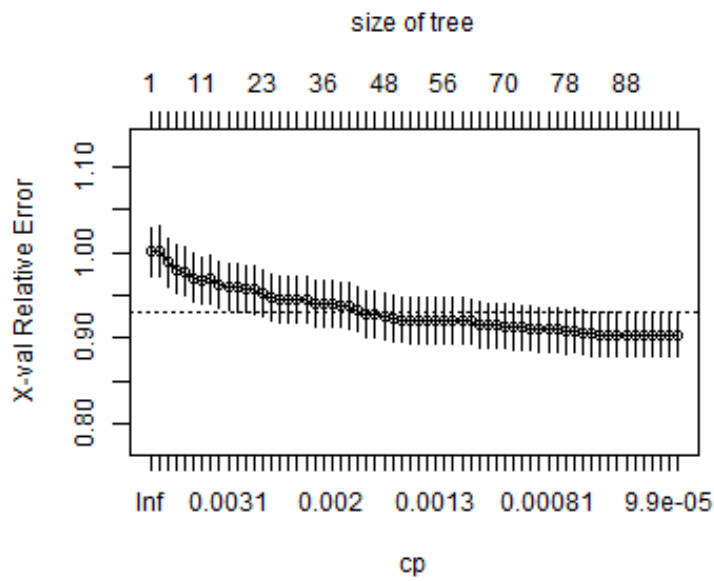
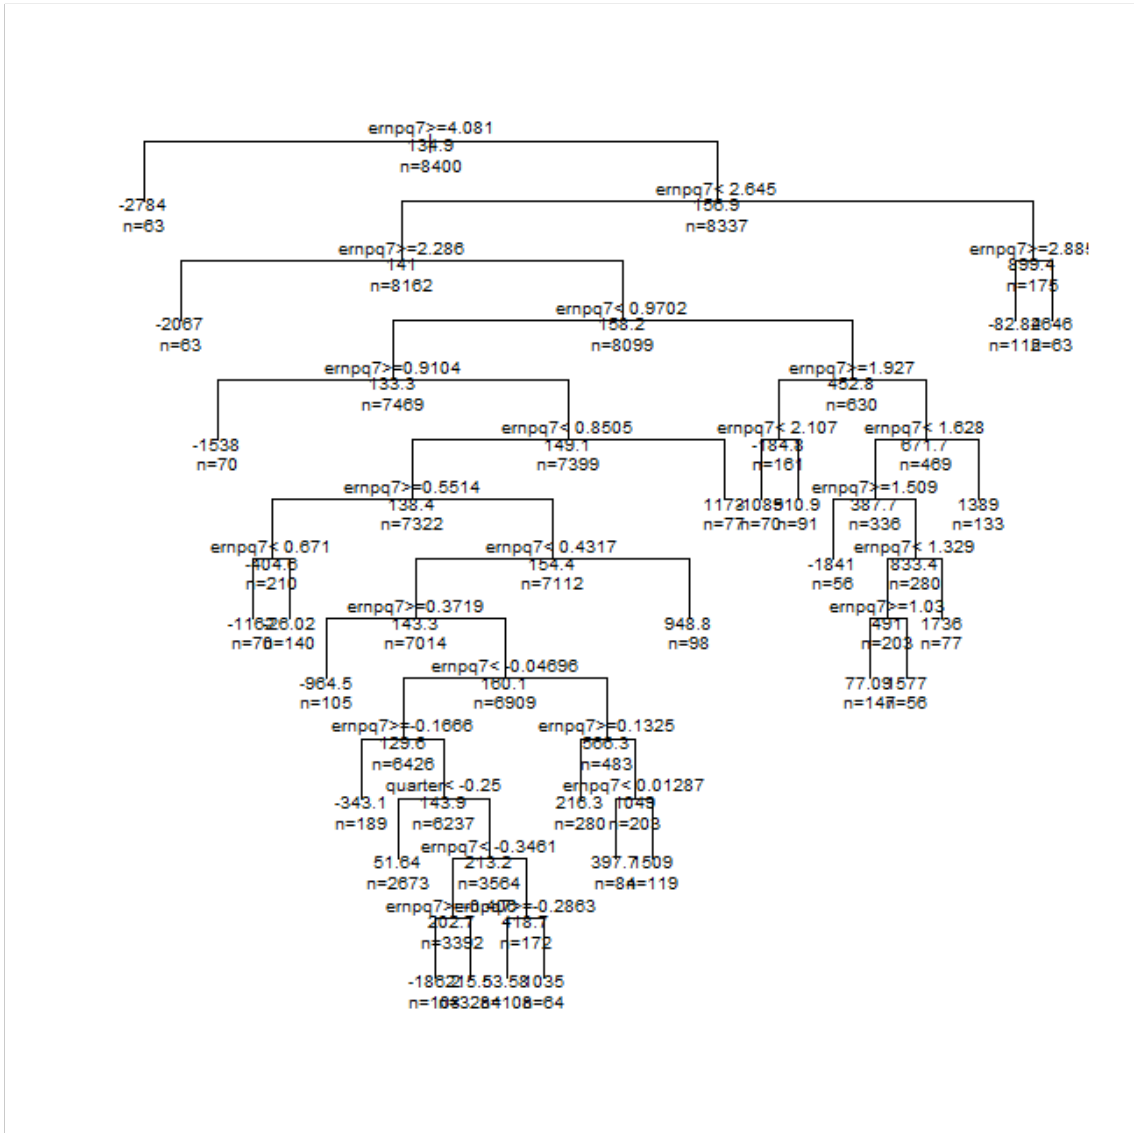
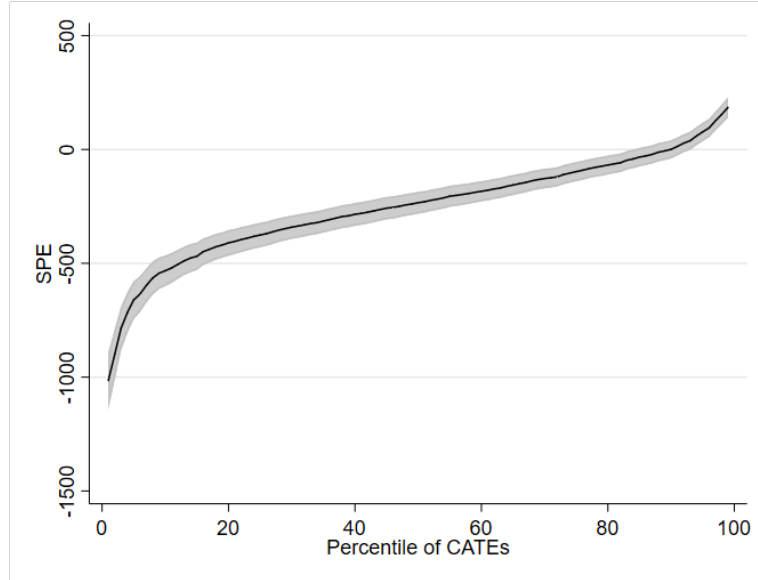


Figure C.4: Structure of the tree with baseline covariates.



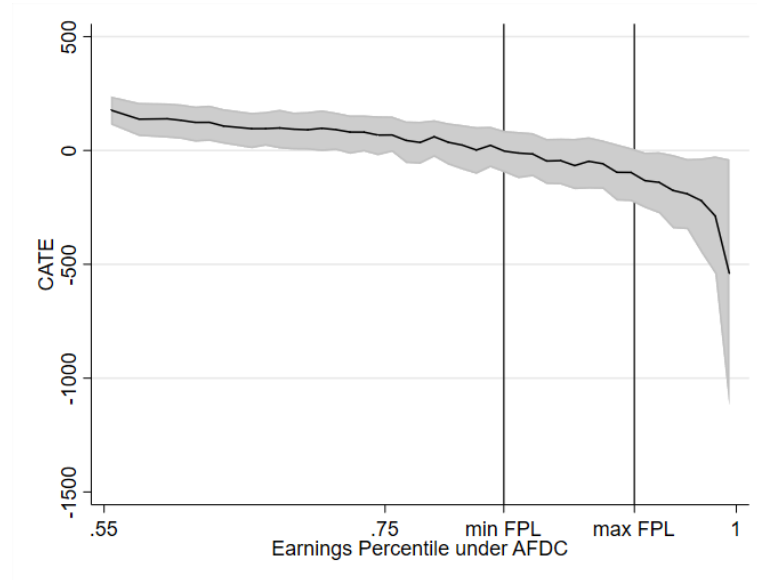
D Additional results

Figure D.1: Sorted predicted effects of the CATEs for participants with earnings above the FPL (extra child) under AFDC.



Note: FPL (extra child) means that I inflate the number of children per mother by one when calculating the assistance unit size. I estimate the 99%-confidence interval (gray area) using the bootstrap Algorithm 2.1 of [Chernozhukov, Fernández-Val, and Luo \(2018\)](#) with 2,000 replications.

Figure D.2: Causal forest CATEs (kitchen sink controls) by earnings percentile under AFDC.



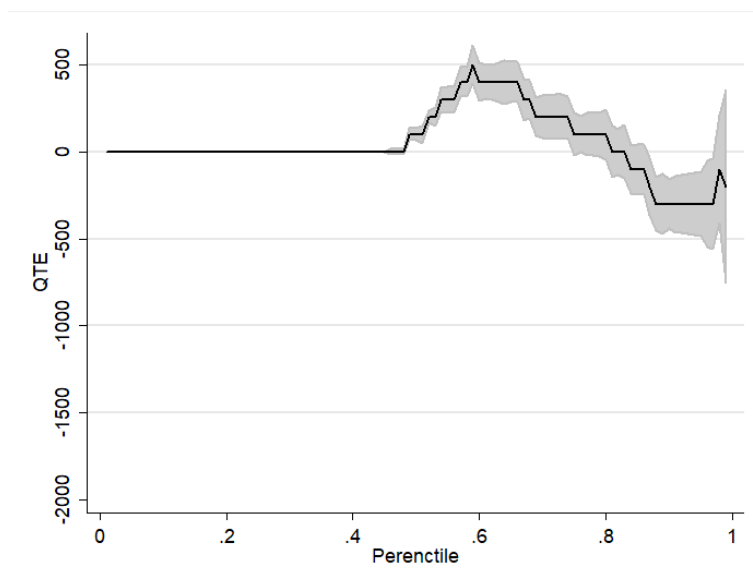
Note: I estimate the distribution of the CATEs using an individual-level clustered bootstrap approach. The gray area reports the 99%-confidence interval based on the percentile method.

Figure D.3: Modified outcome method random forest CATEs (kitchen sink controls) by earnings percentile under AFDC.



Note: I estimate the distribution of the CATEs using an individual-level clustered bootstrap approach. The gray area reports the 99%-confidence interval based on the percentile method.

Figure D.4: Quantile treatment effects (QTEs).



Note: Replication of [Bitler, Gelbach, and Hoynes \(2006\)](#). The gray area shows the 99%-confidence interval.

E Similarity between CATEs and QTEs

I propose using a bootstrap procedure to compute the p-values for a Kolmogorov-Smirnov (KS) test for the equality of the potential and simulated outcome distributions. The potential and simulated outcome distributions

$$\begin{aligned} F_{Y(1)}(y) &= Pr(Y_{it} \leq y | D_i = 1), \\ F_{Y(0)}(y) &= Pr(Y_{it} \leq y | D_i = 0), \\ F_{Y(1)}^S(y) &= Pr(Y_{it} + \delta(X_{it}) \leq y | D_i = 0), \\ F_{Y(0)}^S(y) &= Pr(Y_{it} - \delta(X_{it}) \leq y | D_i = 1), \end{aligned}$$

are identified from observable data. The empirical distribution functions are denoted by $\hat{F}_{Y(1)}(y)$, $\hat{F}_{Y(0)}(y)$, $\hat{F}_{Y(1)}^S(y)$, and $\hat{F}_{Y(0)}^S(y)$.

Based on the testable conditions (4) and (5), the possible Kolmogoroff-Smirnov (KS) test statistics are

$$\begin{aligned} \widehat{KS}_1 &= \sup_y |\hat{F}_{Y(1)}(y) - \hat{F}_{Y(1)}^S(y)|, \\ \widehat{KS}_0 &= \sup_y |\hat{F}_{Y(0)}(y) - \hat{F}_{Y(0)}^S(y)|, \text{ and} \\ \widehat{KS}_{\max} &= \max(\widehat{KS}_1, \widehat{KS}_0), \end{aligned}$$

where T_{\max} is a joint test of both conditions.

To obtain the critical values for these test statistics, I compute the bootstrap distributions of the KS test statistics (see, e.g., Giné and Zinn, 1990). I construct B bootstrap resamples with a procedure that accounts for clustering on the individual level. Let $\hat{F}_{Y(1)}^{(b)}(y)$, $\hat{F}_{Y(0)}^{(b)}(y)$, $\hat{F}_{Y(1)}^{S(b)}(y)$, and $\hat{F}_{Y(0)}^{S(b)}(y)$ denote the respective bootstrapped distributions of the bootstrap sample b (for $b = 1, \dots, B$). The re-centered KS test statistics are

$$\begin{aligned} \widehat{KS}_1^{(b)} &= \sup_y |\hat{F}_{Y(1)}^{(b)}(y) - \hat{F}_{Y(1)}^{S(b)}(y) - (\hat{F}_{Y(1)}(y) - \hat{F}_{Y(1)}^S(y))|, \\ \widehat{KS}_0^{(b)} &= \sup_y |\hat{F}_{Y(0)}^{(b)}(y) - \hat{F}_{Y(0)}^{S(b)}(y) - (\hat{F}_{Y(0)}(y) - \hat{F}_{Y(0)}^S(y))|, \text{ and} \\ \widehat{KS}_{\max}^{(b)} &= \max(\widehat{KS}_1^{(b)}, \widehat{KS}_0^{(b)}). \end{aligned}$$

The estimated p-values $\hat{\alpha}(\widehat{KS})$ for the null hypotheses, that the potential and sim-

ulated outcome distributions are equal, are

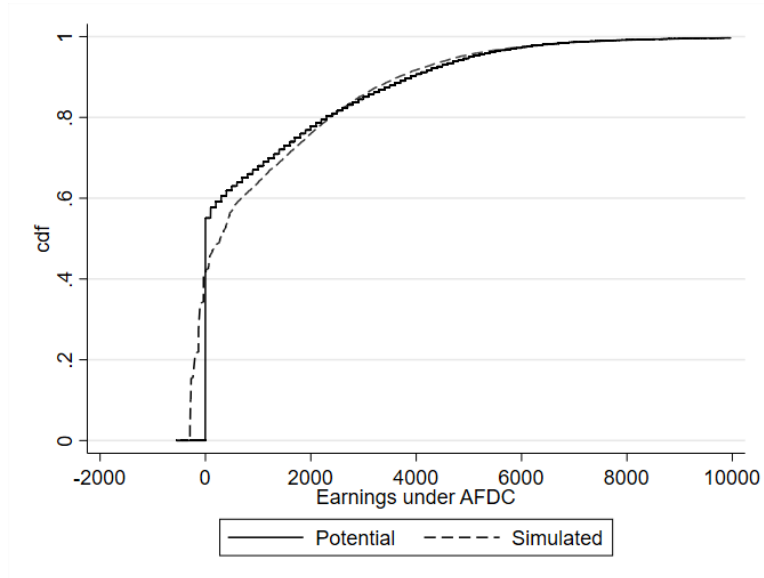
$$\hat{\alpha}(\widehat{KS}_1) = \frac{1}{B} \sum_{b=1}^B 1 \left\{ \widehat{KS}_1^{(b)} > \widehat{KS}_1 \right\},$$

$$\hat{\alpha}(\widehat{KS}_0) = \frac{1}{B} \sum_{b=1}^B 1 \left\{ \widehat{KS}_1^{(b)} > \widehat{KS}_1 \right\}, \text{ and}$$

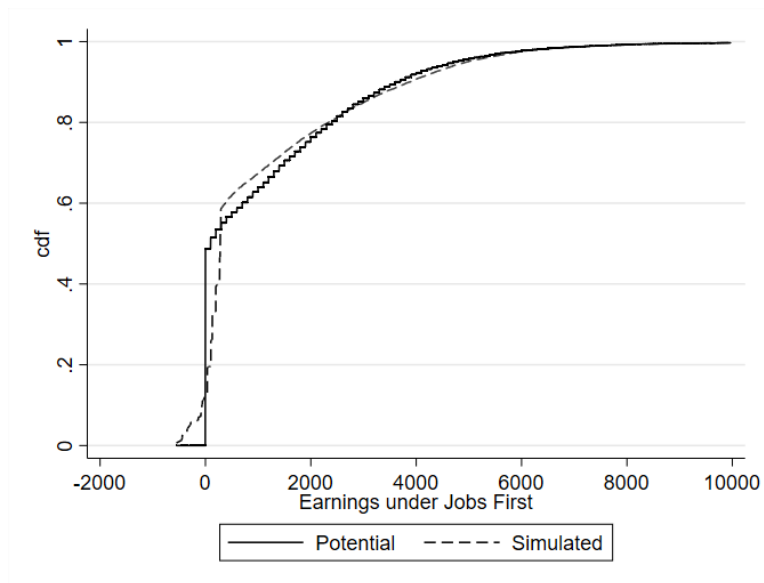
$$\hat{\alpha}(\widehat{KS}_{\max}) = \frac{1}{B} \sum_{b=1}^B 1 \left\{ \widehat{KS}_{\max}^{(b)} > \widehat{KS}_{\max} \right\}.$$

F Simulated earnings distributions

Figure F.1: Potential and simulated earnings distributions obtained from the local constant model (baseline specification).

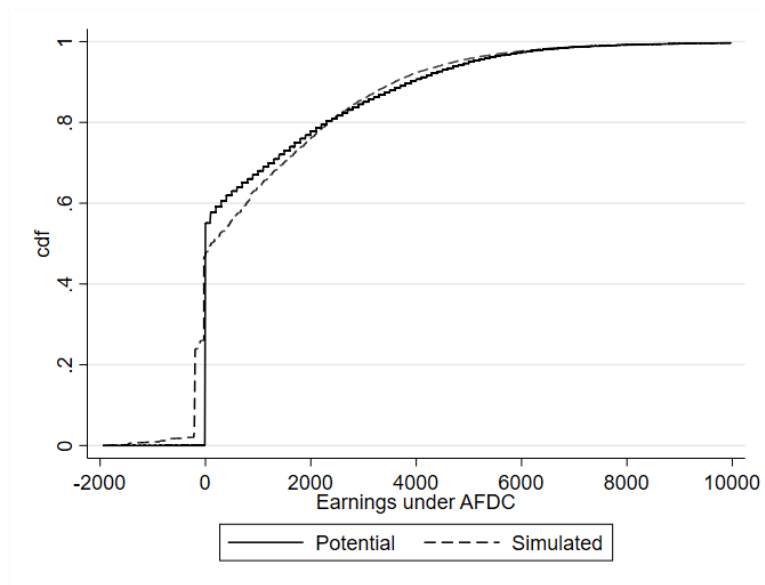


(a) Under AFDC

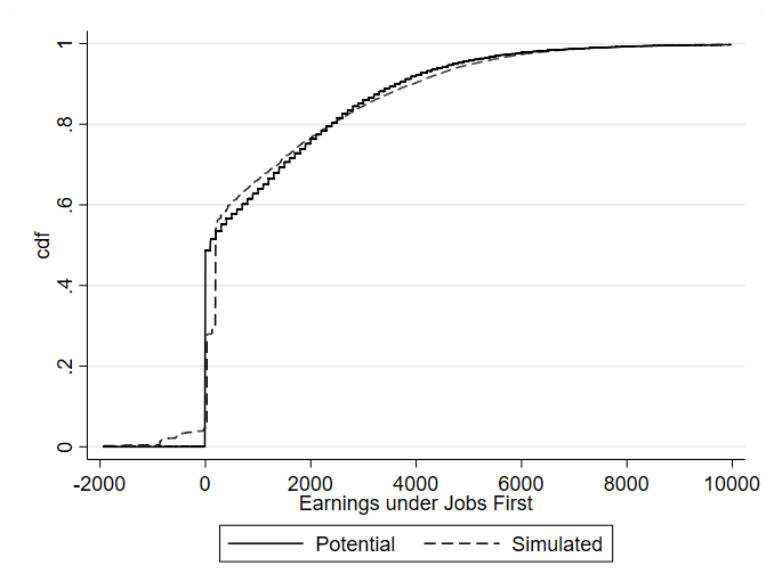


(b) Under Jobs First

Figure F.2: Potential and simulated earnings distributions obtained from the tree estimator and the baseline controls.

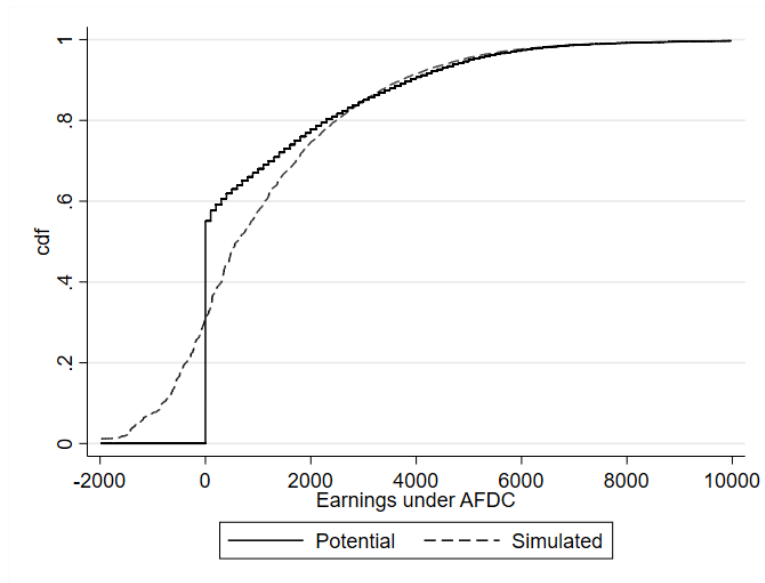


(a) Under AFDC

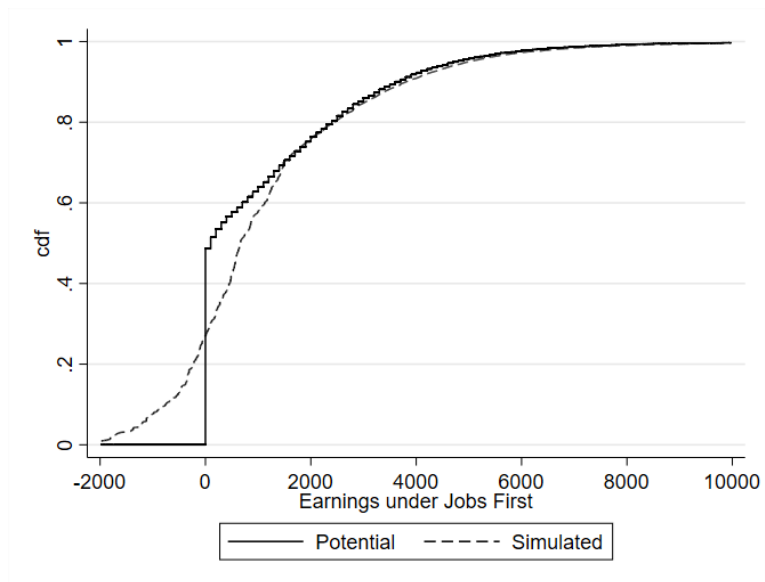


(b) Under Jobs First

Figure F.3: Potential and simulated earnings distributions obtained from the tree estimator and the decent controls.

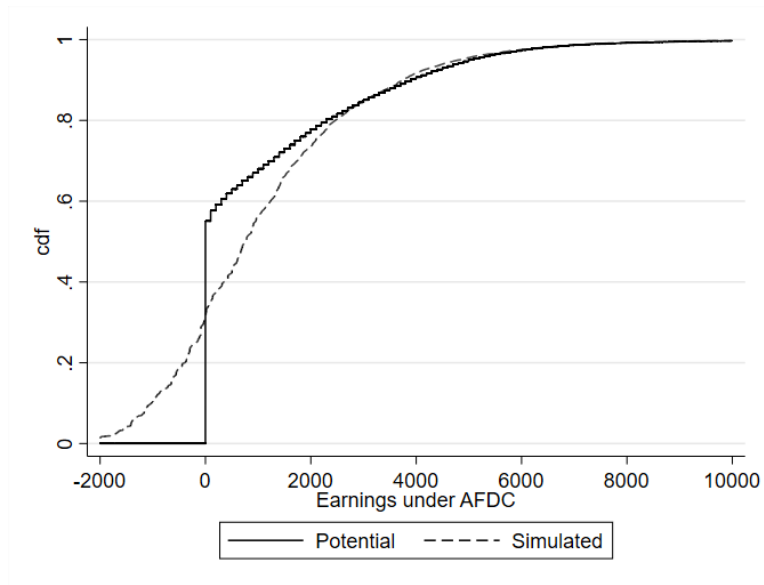


(a) Under AFDC

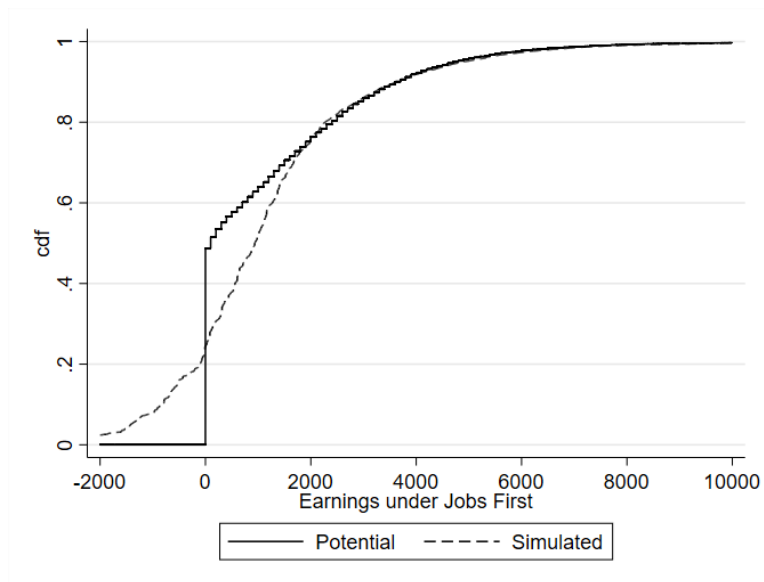


(b) Under Jobs First

Figure F.4: Potential and simulated earnings distributions obtained from the tree estimator and the kitchen sink controls.

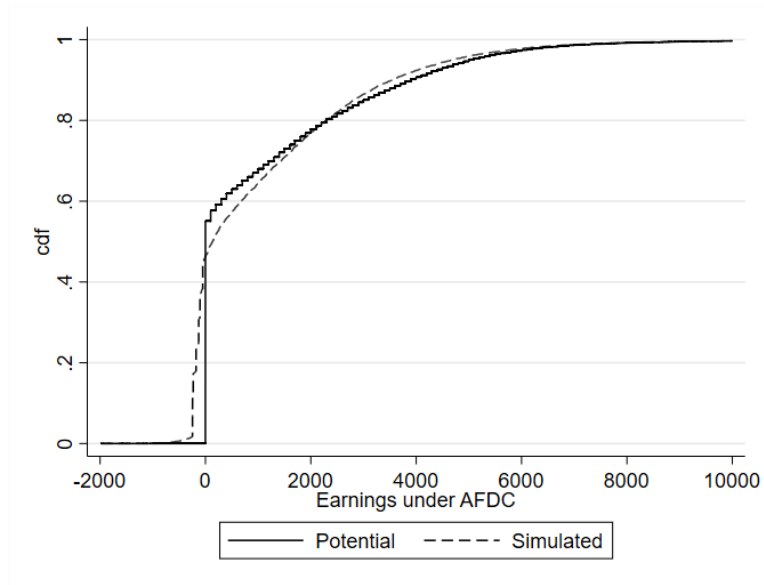


(a) Under AFDC

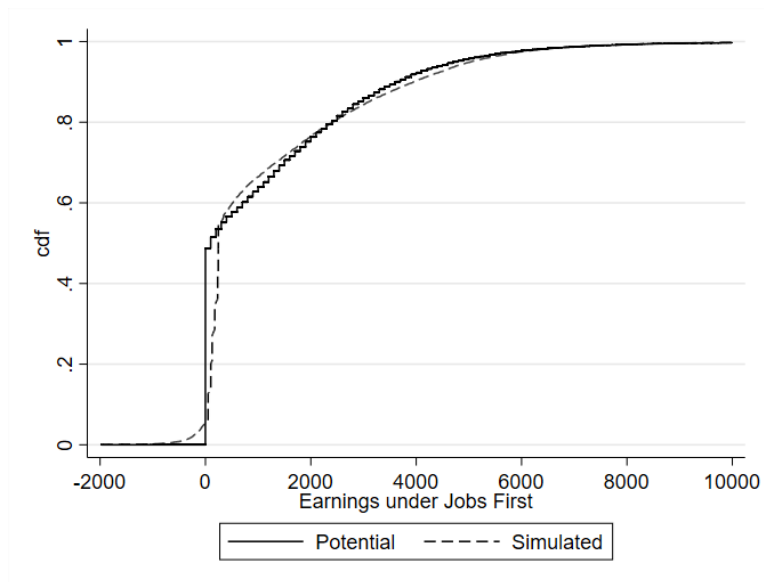


(b) Under Jobs First

Figure F.5: Potential and simulated earnings distributions obtained from the random forest estimator and the baseline controls.

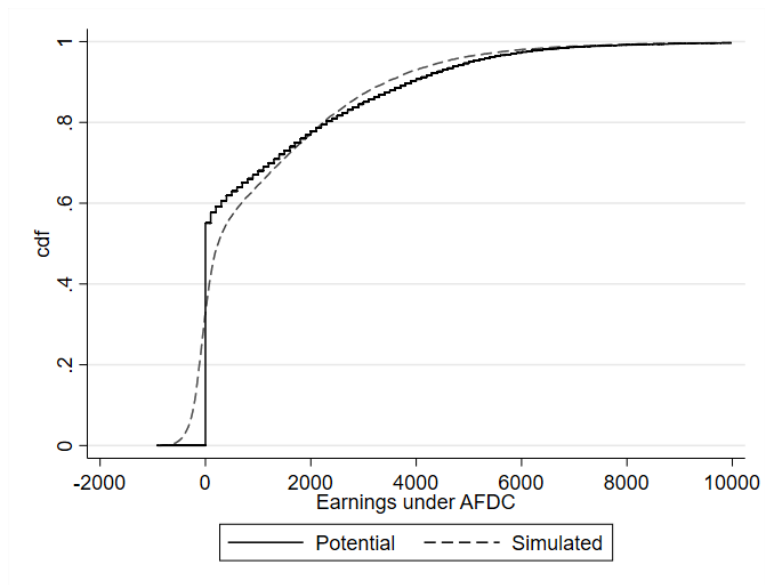


(a) Under AFDC

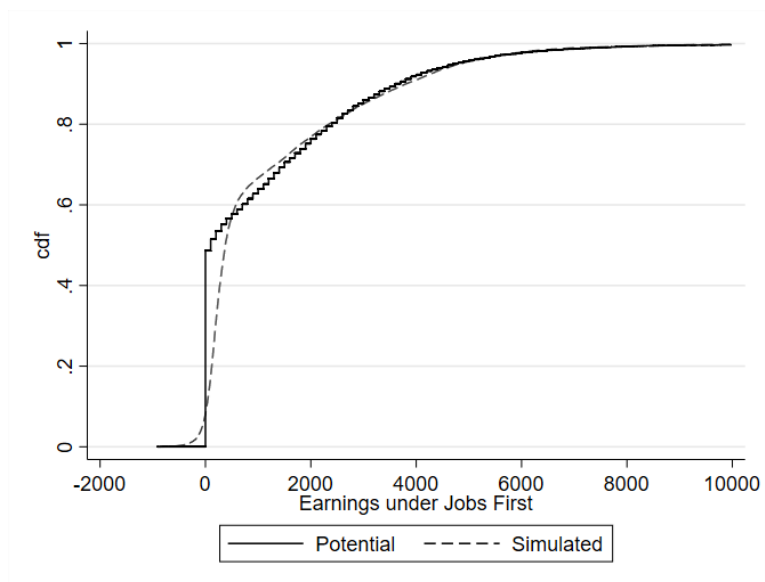


(b) Under Jobs First

Figure F.6: Potential and simulated earnings distributions obtained from the random forest estimator and the kitchen sink controls.

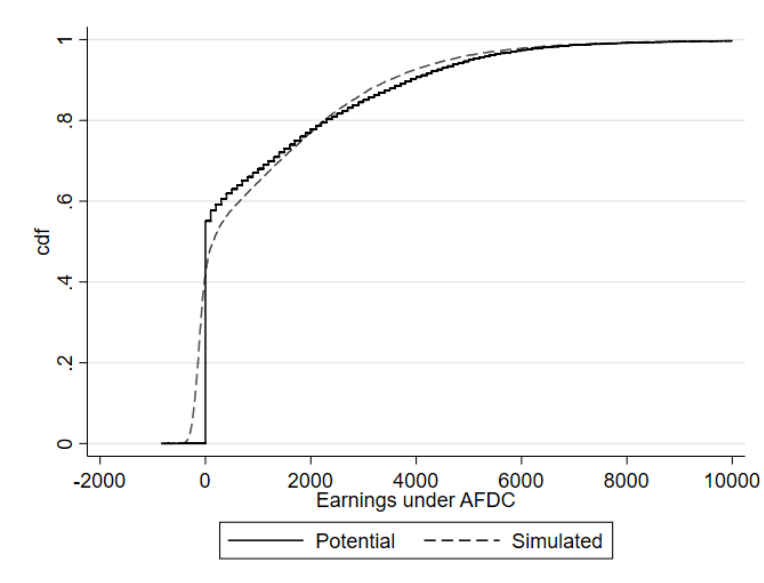


(a) Under AFDC

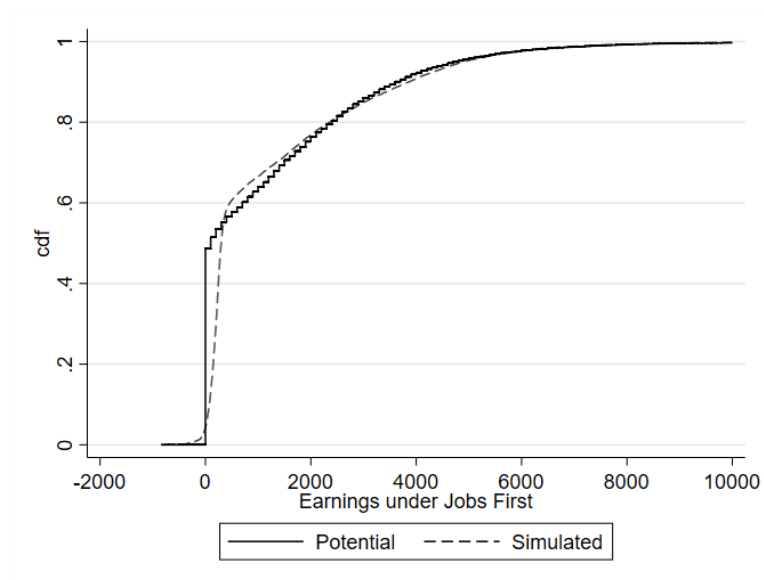


(b) Under Jobs First

Figure F.7: Potential and simulated earnings distributions obtained from the causal forest estimator and the kitchen sink controls.



(a) Under AFDC



(b) Under Jobs First