Front-door Difference-in-Differences Estimators*

Adam Glynn†
aglynn@fas.harvard.edu

Konstantin Kashin‡
kkashin@fas.harvard.edu

April 9, 2014

Abstract

In this paper, we develop front-door difference-in-differences estimators that utilize information from post-treatment variables in addition to information from pre-treatment covariates. Even when the front-door criterion does not hold, these estimators allow the identification of causal effects by utilizing assumptions that are analogous to standard difference-in-differences assumptions. We also demonstrate that causal effects can sometimes be bounded by front-door and front-door difference-in-differences estimators under relaxed assumptions. We illustrate these points with an application to the National JTPA (Job Training Partnership Act) Study and with an application to Florida's early in-person voting program. For the JTPA study, we show that an experimental benchmark can be bracketed with front-door and front-door difference-in-differences estimates. Surprisingly, neither of these estimates uses control units. For the Florida program, we find some evidence that early in-person voting had small positive effects on turnout in Florida in 2008 and 2012. This provides a counterpoint to recent claims that early voting had a negative effect on turnout in 2008.

*We thank Justin Grimmer, Manabu Kuroki, Kevin Quinn, and seminar participants at Emory, Harvard, Notre Dame, NYU, Ohio State, UC Davis, and UMass Amherst for comments and suggestions. An earlier version of this paper was presented at the 2014 Asian Political Methodology Meeting in Tokyo.

†Department of Government and Institute for Quantitative Social Science, Harvard University, 1737 Cambridge Street, Cambridge, MA 02138 (http://scholar.harvard.edu/aglynn).

‡Department of Government and Institute for Quantitative Social Science, Harvard University, 1737 Cambridge Street, Cambridge, MA 02138 (http://konstantinkashin.com).

1 Introduction

One of the main tenets of observational studies is that post-treatment variables should not be included in an analysis because naively conditioning on these variables can block some of the effect of interest, leading to post-treatment bias (King, Keohane and Verba, 1994). While this is usually sound advice, it seems to contradict recommendations from the process tracing literature that information about mechanisms can be used to assess the plausibility of an effect (Collier and Brady, 2004; George and Bennett, 2005; Brady, Collier and Seawright, 2006).

The front-door criterion (Pearl, 1995) and its extensions (Kuroki and Miyakawa, 1999; Tian and Pearl, 2002*a*,*b*; Shpitser and Pearl, 2006) resolve this apparent contradiction, providing a means for nonparametric identification of treatment effects using post-treatment variables. Importantly, the front-door approach can identify causal effects even when there are unmeasured common causes of the treatment and the outcome (i.e., the total effect is confounded). The basic idea is the following. Suppose the total effect of a treatment on an outcome can be partitioned into a set of pathways and constituent effects. Further suppose that although the total effect is confounded, the constituent effects are identified.[1] When the constituent effects are identified, the front-door adjustment provides a formula for combining these constituent effects into the total effect. Furthermore, unlike traditional path analysis, the front-door adjustment allows for heterogeneous effects and does not require parametric assumptions.

While front-door adjustment seems a powerful tool for data analysts, it has been used infrequently (VanderWeele, 2009) due to concerns that the assumptions required for point identification will rarely hold (Cox and Wermuth, 1995; Imbens and Rubin, 1995). A number of papers have proposed weaker and more plausible sets of assumptions (Joffe, 2001; Kaufman, Kaufman and MacLehose, 2009; Glynn and Quinn, 2011) that tend to correspond to conceptions of process trac-

[1]This can happen when there is an unmeasured common cause of the treatment and outcome, but there are no unmeasured common causes of the treatment and the mediators or of the mediators and the outcome (precise conditions are discussed below).

ing. However, these approaches rely on binary or bounded outcomes, and even in large samples these methods only provide bounds on causal effects (i.e., partial instead of point identification). Additionally, these bounds on effects typically include zero. Recently, Glynn and Kashin (2013) developed bias formulas that allow the front-door assumptions to be weakened via sensitivity analysis. This allows for any type of outcome variable and increases the possibility that the front-door approach will be informative (e.g., establishing that zero is not a plausible value for the effect).

In this paper, we demonstrate that the bias described in Glynn and Kashin (2013) can sometimes be removed by a difference-in-differences approach when there is one-sided noncompliance. Glynn and Kashin (2013) showed that with one-sided noncompliance, the front-door estimator implies substituting treated noncompliers for controls. For example, if you want to study the effect of signing up for a program, the front-door estimator compares the outcomes of those that sign up (the treated) to the subset of those that sign up but do not show up (the treated noncompliers). Contrast this with standard approaches (e.g., matching and regression) that would compare those that sign up (the treated) with those that do not sign up (the controls).

The front-door difference-in-differences approach extends the front-door approach in the following manner. First, we identify the treated units of interest, which we will refer to as the group of interest. Second, if we can identify a group of treated units distinct from our group of interest for which we believe the treatment should have no effect, then a non-zero front-door estimate for this group can be attributed to bias. We will refer to this group as the differencing group. For example, in the context of the early voting application to follow, we consider the effects of an early in-person (EIP) voting program on turnout for elections in 2008 and 2012. One differencing group we consider is potential voters that used an absentee ballot in the previous election. EIP was unlikely to have an effect on these voters, as they had already demonstrated their ability to vote by another means of early voting. Therefore, we consider non-zero front-door estimates of the turnout effect for this group to be evidence of bias.

If we further assume that the bias for the differencing group is equal to the bias for our group of interest, then by subtracting the front-door estimator for this group from the front-door estimator for the group of interest, we can remove the bias from our front-door estimate for the group of interest. Note that if all effects and bias are positive, then when the bias from the differencing group is larger than the bias for the group of interest and/or the treatment has an effect for the differencing group, then this differencing approach can provide a lower bound on the effect of the program. We demonstrate this within the context of a job training study. However, we also demonstrate that the bias for each group is related to the proportion of compliers in the group, and therefore, an equal bias assumption is untenable without an additional adjustment. This will be described in detail below.

The paper is organized as follows. Section 2 presents the bias formulas for the front-door approach to estimating average treatment effects on the treated (ATT), both for the general case and the simplification for nonrandomized program evaluations with one-sided noncompliance. Section 3 presents the difference-in-differences approach for front-door estimators for the simplified case and discusses the required assumptions. Section 4 presents an application of the front-door difference-in-differences estimator to the National JTPA (Job Training Partnership Act) Study. Section 5 presents an application of the front-door difference-in-differences estimator to election law: assessing the effects of early in-person voting on turnout in Florida. Section 6 concludes.

2 Bias for the Front-Door Approach for ATT

In this section, we present large-sample bias formulas for the front-door approach for estimating the average treatment effect on the treated (ATT). This is often the parameter of interest when assessing the effects of a program or a law. The beginning of this section parallels the early discussion in Glynn and Kashin (2013), but is important for the development of the difference-in-differences estimator presented in this paper. For an outcome variable Y and a binary treatment/action A, we define the potential outcome under active treatment as $Y(a_1)$ and the potential outcome under control as

$Y(a_0)$.[2] Our parameter of interest is the ATT, defined as $\tau_{att} = E[Y(a_1)|a_1] - E[Y(a_0)|a_1] = \mu_{1|a_1} - \mu_{0|a_1}$. We assume consistency, $E[Y(a_1)|a_1] = E[Y|a_1]$, so that the mean potential outcome under active treatment for the treated units is equal to the observed outcome for the treated units such that $\tau_{att} = E[Y|a_1] - E[Y(a_0)|a_1]$. The ATT is therefore the difference between the mean outcome for the treated units and mean counterfactual outcome for these units, had they not received the treatment.

We also assume that $\mu_{0|a_1}$ is potentially identifiable by conditioning on a set of observed covariates X and unobserved covariates U. To clarify, we assume that if the unobserved covariates were actually observed, the ATT could be estimated by standard approaches (e.g., matching). For simplicity in presentation we assume that X and U are discrete, such that

$$\mu_{0|a_1} = \sum_x \sum_u E[Y|a_0, x, u] \cdot P(u|a_1, x) \cdot P(x|a_1),$$

but continuous variables can be handled analogously. However, even with only discrete variables we have assumed that the conditional expectations in this equation are well-defined, such that for all levels of X and U amongst the treated units, all units had a positive probability of receiving either treatment or control (i.e., positivity holds).

The front-door adjustment for a set of measured post-treatment variables M can be written as the following:

$$\mu^{fd}_{0|a_1} = \sum_x \sum_m P(m|a_0, x) \cdot E[Y|a_1, m, x] \cdot P(x|a_1).$$

We can thus define the large-sample front-door estimator of ATT as:

$$\tau^{fd}_{att} = \mu_{1|a_1} - \mu^{fd}_{0|a_1}.$$

[2]Note that we must assume that these potential outcomes are well defined for each individual, and therefore we are making the stable unit treatment value assumption (SUTVA).

The large-sample bias in the front-door estimate of ATT, which is entirely attributable to the bias in the front-door estimate of $\mu_{0|a_1}$, is the following (see Appendix A.1 for a proof):

$$B^{fd}_{att} = \sum_x P(x|a_1) \sum_m \sum_u P(m|a_0, x, u) \cdot E[Y|a_0, m, x, u] \cdot P(u|a_1, x)$$
$$- \sum_x P(x|a_1) \sum_m \sum_u P(m|a_0, x) \cdot E[Y|a_1, m, x, u] \cdot P(u|a_1, m, x).$$

As discussed in Glynn and Kashin (2013), it is possible for the front-door approach to provide a large-sample unbiased estimator for the ATT even in the presence of an unmeasured confounder that would bias traditional covariate adjustment techniques such as matching and regression. Specifically, the front-door bias will be zero when three conditions hold: (1) $E[Y|a_1, m, x, u] = E[Y|a_0, m, x, u]$, (2) $P(m|a_0, x) = P(m|a_0, x, u)$, and (3) $P(u|a_1, m, x) = P(u|a_1, x)$. The first will hold when Y is mean independent of A conditional on U, M, and X, while the latter two will hold if U is independent of M conditional on X and a_0 or a_1.

For the difference-in-differences estimators we consider in this paper, we use the special case of nonrandomized program evaluations with one-sided noncompliance. Following the literature in econometrics on program evaluation, we define the program impact as the ATT where the active treatment (a_1) is assignment into a program (Heckman, LaLonde and Smith, 1999), and when M indicates whether the active treatment (a_1) was actually received. We use the short-hand notation m_1 to denote that active treatment was received and m_0 if it was not.

Assumption 1 (One-sided noncompliance)
$P(m_0|a_0, x) = P(m_0|a_0, x, u) = 1$ *for all* x, u.

Assumption 1 implies that only those assigned to treatment can receive treatment, and the front-door large-sample estimator reduces to the following under this assumption:

$$
\begin{aligned}
\tau_{att}^{fd} &= \mu_{1|a_1} - \mu_{0|a_1}^{fd} \\
&= E[Y|a_1] - \sum_x \sum_m P(m|a_0, x) \cdot E[Y|a_1, m, x] \cdot P(x|a_1) \\
&= E[Y|a_1] - \sum_x \underbrace{E[Y|a_1, m_0, x]}_{\text{treated non-compliers}} \cdot P(x|a_1) &(1) \\
&= \sum_x P(x|a_1) \cdot P(m_1|x, a_1) \cdot \left\{ \underbrace{E[Y|a_1, m_1, x] - E[Y|a_1, m_0, x]}_{\text{“effect” of receiving treatment}} \right\} &(2)
\end{aligned}
$$

The formulas in (1) and (2) are interesting because they do not rely upon outcomes of control units in the construction of proxies for the potential outcomes under control for treated units (see Appendix A.2 for the derivation of (2)). This is a noteworthy point that has implications for research design that we will revisit subsequently. The formula in (1) can be compared to the standard large-sample covariate adjustment for ATT:

$$
\begin{aligned}
\tau_{att}^{std} &= \mu_{1|a_1} - \mu_{0|a_1}^{std} \\
&= E[Y|a_1] - \sum_x \underbrace{E[Y|a_0, x]}_{\text{controls}} \cdot P(x|a_1). &(3)
\end{aligned}
$$

Roughly speaking, standard covariate adjustment matches units that were assigned treatment to similar units that were assigned control. On the other hand, front-door estimates match units that were assigned treatment to similar units that were assigned treatment but did not receive treatment. This sort of comparison is not typical, so it is helpful to consider the informal logic of the procedure before presenting the formal statements of bias. The fundamental question is whether the treated noncompliers provide reasonable proxies for the missing counterfactuals: the outcomes that would have occurred if the treated units had not been assigned treatment. Therefore, in order for the front-door approach to be unbiased in large samples, we are effectively assuming that 1) assignment to treatment has no effect if treatment is not received and 2) those that are assigned but don't receive treatment are comparable in some sense to those that receive treatment. This will be made more precise below.

The front-door formula in (2), with the observable proportions $P(x|a_1)$ and $P(m_1|a_1, x)$ multiplying the estimated effect of receiving the treatment, is helpful when considering the simplified front-door ATT bias, which can be written in terms of the same observable proportions (see Appendices A.3 and A.4 for proofs):

$$B_{att}^{fd} = \sum_x P(x|a_1)P(m_1|a_1,x)\sum_u \Bigg[E[Y|a_0,m_0,x,u]\cdot[P(u|a_1,m_1,x)-P(u|a_1,m_0,x)] \\ + \Bigg\{E[Y(a_0)|a_1,m_1,x,u]\frac{P(m_1|a_1,x,u)}{P(m_1|a_1,x)} - E[Y(a_0)|a_1,m_0,x,u]\cdot\frac{\frac{E[Y|a_1,m_0,x,u]}{E[Y(a_0)|a_1,m_0,x,u]}-P(m_0|a_1,x,u)}{P(m_1|a_1,x)}\Bigg\}\cdot P(u|a_1,m_0,x)\Bigg]$$

The unobservable portion of this bias formula (i.e., everything after the $\sum_u$), can be difficult to interpret, but there are a number of assumptions that allow us to simplify the formula. For example, we might assume that treatment does not have an effect on the outcome for noncompliers.

Assumption 2 (Exclusion restriction)

No direct effect for noncompliers: $E[Y|a_1, m_0, x, u] = E[Y(a_0)|a_1, m_0, x, u]$.

When combined with the consistency assumption, Assumption 2 can also be written as $E[Y(a_1)|a_1, m_0, x, u] = E[Y(a_0)|a_1, m_0, x, u]$. If this exclusion restriction holds, then the bias simplifies to the following:

$$B_{att}^{fd} = \sum_x P(x|a_1)P(m_1|a_1,x)\sum_u \Bigg[E[Y|a_0,m_0,x,u]\cdot[P(u|a_1,m_1,x)-P(u|a_1,m_0,x)] \\ + \Bigg\{E[Y(a_0)|a_1,m_1,x,u]\frac{P(m_1|a_1,x,u)}{P(m_1|a_1,x)} - E[Y(a_0)|a_1,m_0,x,u]\cdot\frac{P(m_1|a_1,x,u)}{P(m_1|a_1,x)}\Bigg\}\cdot P(u|a_1,m_0,x)\Bigg]$$

If instead we assume that compliance rates are constant across levels of u within levels of x,

Assumption 3 (Constant compliance rates across values of u within levels of x)

$P(m_1|a_1, x, u) = P(m_1|a_1, x)$ *for all* x *and* u*,* then due to the binary measure of treatment received, we know that $P(u|a_1, m_1, x) = P(u|a_1, m_0, x)$ (see Appendix A.5), and the bias simplifies to the following:

$$B^{fd}_{att} = \sum_x P(x|a_1)P(m_1|a_1, x) \sum_u \Bigg[\Bigg\{ E[Y(a_0)|a_1, m_1, x, u] \\ - E[Y(a_0)|a_1, m_0, x, u] \cdot \frac{\frac{E[Y|a_1, m_0, x, u]}{E[Y(a_0)|a_1, m_0, x, u]} - P(m_0|a_1, x, u)}{P(m_1|a_1, x)} \Bigg\} \cdot P(u|a_1, m_0, x) \Bigg]$$

Assumption 3 can be strengthened and the bias simplified further in some cases of clustered treatment assignment. Because the front-door estimator uses only treated units under Assumption 1, it is possible that all units within levels of x were assigned in clusters such that U is actually measured at the cluster level. We present an example of this in the application, where treatment (the availability of early in-person voting) is assigned at the state level, and therefore all units within a state (e.g., Florida) have the same value of u. Formally, the assumption can be stated as the following:

Assumption 4 (u is constant among treated units within levels of x)
For any two units with a_1 and covariate values (x, u) and (x', u'), $x = x' \Rightarrow u = u'$.

When Assumption 4 holds, the u notation is redundant, and can be removed from the bias formula which simplifies as the following:

$$B^{fd}_{att} = \sum_x P(x|a_1)P(m_1|a_1, x) \Bigg\{ E[Y(a_0)|a_1, m_1, x] - E[Y(a_0)|a_1, m_0, x] \cdot \frac{\frac{E[Y|a_1, m_0, x]}{E[Y(a_0)|a_1, m_0, x]} - P(m_0|a_1, x)}{P(m_1|a_1, x)} \Bigg\} \tag{4}$$

Finally, it can be instructive to consider the formula when both Assumption 2 and Assumption 4 hold. In this scenario, the remaining bias is due to an unmeasured common cause of the mediator and the outcome.

$$B^{fd}_{att} = \sum_x P(x|a_1)P(m_1|a_1, x)\{E[Y(a_0)|a_1, m_1, x] - E[Y(a_0)|a_1, m_0, x]\}$$

In some applications, the bias B_{att}^{fd} may be small enough for the front-door estimator to provide a viable approach. For others, we may want to remove the bias. In the next section, we discuss a difference-in-differences approach to removing the bias.

3 Front-door Difference-in-Differences Estimators

If we define the front-door estimator within levels of a covariate x as $\tau_{att,x}^{fd}$, then the front-door estimator can be written as a weighted average of strata-specific front-door estimators where the weights are relative strata sizes for treated units:

$$\tau_{att}^{fd} = \sum_x P(x|a_1)\tau_{att,x}^{fd}.$$

If we further define the group of interest as the stratum g_1 and the differencing group as the stratum g_2, then the front-door estimators within levels of x for these groups can be written as:

$$\begin{aligned}
\tau_{att,x,g_1}^{fd} &= P(m_1|x,a_1,g_1)\{E[Y|a_1,m_1,x,g_1] - E[Y|a_1,m_0,x,g_1]\}, \\
\tau_{att,x,g_2}^{fd} &= P(m_1|x,a_1,g_2)\{E[Y|a_1,m_1,x,g_2] - E[Y|a_1,m_0,x,g_2]\}.
\end{aligned}$$

Using these components, the front-door difference-in-differences estimator can be written as

$$\begin{aligned}
\tau_{att}^{fd-did} &= \sum_x P(x|a_1,g_1)\left[\tau_{att,x,g_1}^{fd} - \frac{P(m_1|a_1,x,g_1)}{P(m_1|a_1,x,g_2)}\tau_{att,x,g_2}^{fd}\right] \qquad (5) \\
&= \sum_x P(x|a_1,g_1)P(m_1|x,a_1,g_1)\left[\{E[Y|a_1,m_1,x,g_1] - E[Y|a_1,m_0,x,g_1]\}\right. \\
&\quad \left. - \{E[Y|a_1,m_1,x,g_2] - E[Y|a_1,m_0,x,g_2]\}\right]. \qquad (6)
\end{aligned}$$

Hence, (5) shows that within levels of x, the front-door difference-in-differences estimator is the difference between the front-door estimator from the group of interest and a scaled front-door estimator from the differencing group, where the scaling factor is the ratio of the compliance rates in the two groups. Then, the overall front-door difference-in-differences estimator is a weighted average of the estimators within levels of x, where the weights are determined by the group of interest proportions of x for treated units. Intuitively, the scaling factor is necessary because it places the front-door estimate for the differencing group on the same compliance scale as the front-door estimate for the group of interest. The necessity of this adjustment can be most easily seen in (6), where we see that the main goal is to remove the bias from the $\{E[Y|a_1, m_1, x, g_1] - E[Y|a_1, m_0, x, g_1]\}$ component of group 1 with the $\{E[Y|a_1, m_1, x, g_1] - E[Y|a_1, m_0, x, g_1]\}$ component of group 2.

In order for the front-door difference-in-differences estimator to give us an unbiased estimate of the ATT for the group of interest in large samples, we need the following two assumptions to hold. If we further define bias within levels of x for a generic group g as

$$B^{fd}_{att,x,g} = P(m_1|a_1, x, g) \sum_u \Bigg[E[Y|a_0, m_0, x, u, g] \cdot [P(u|a_1, m_1, x, g) - P(u|a_1, m_0, x, g)] + \{E[Y(a_0)|a_1, m_1, x, u, g] \frac{P(m_1|a_1, x, u, g)}{P(m_1|a_1, x, g)} - E[Y(a_0)|a_1, m_0, x, u, g] \cdot \frac{\frac{E[Y|a_1, m_0, x, u, g]}{E[Y(a_0)|a_1, m_0, x, u, g]} - P(m_0|a_1, x, u, g)}{P(m_1|a_1, x, g)} \} P(u|a_1, m_0, x, g) \Bigg],$$

then the assumption we need for the differencing group is the following:

Assumption 5 (No effect for the differencing group)
$\tau^{fd}_{att,x,g_2} = B^{fd}_{att,x,g_2}$ *for all x.*

Note that Assumption 5 can often be weakened. If we believe there are effects for the differencing group, but these have the same sign for the group of interest, then subtracting the scaled estimated effect from the differencing group will remove too much from the estimated effect in the group of interest. For example, if we believe that effects for the group of interest and the differencing group would be positive, then the front-door difference-in-differences estimator would tend to be understated. If we additionally believe that the bias in the front-door estimator is positive prior to the differencing, then the front-door and front-door difference-in-differences estimator will bracket the truth in large samples.

We also need to assume that the bias in the group of interest (g_1) can be removed using the bias from the differencing group (g_2):

Assumption 6 (Bias for g_1 equal to scaled bias for g_2 within levels of x)
$B^{fd}_{att,x,g_1} = \frac{P(m_1|a_1,x,g_1)}{P(m_1|a_1,x,g_2)} B^{fd}_{att,x,g_2}$ *for all x.*

If Assumptions 1, 5, and 6 hold, then τ^{fd-did}_{att} has no large-sample bias for τ_{att} (see Appendix B.1 for a proof). However, the interpretation of Assumption 6 will often be simplified when Assumptions 2, 3, or 4 hold. This will be discussed in the context of the applications, but one special case is useful to consider for illustrative purposes. When Assumptions 1 through 5 hold, then Assumption 6 is equivalent to the following:

$$\{E[Y(a_0)|a_1, m_1, x, g_1] - E[Y(a_0)|a_1, m_0, x, g_1]\} = \{E[Y(a_0)|a_1, m_1, x, g_2] - E[Y(a_0)|a_1, m_0, x, g_2]\}$$

Note that this equality is analogous to the parallel trends assumption for standard difference-in-differences estimators.

4 Illustrative Application: National JTPA Study

We now illustrate how front-door and front-door difference-in-differences estimates for the average treatment effect on the treated (ATT) can bracket the experimental truth in the context of the National JTPA Study, a job training evaluation with both experimental data and a nonexperimental comparison group. This section builds upon Glynn and Kashin (2013), which demonstrates the superior performance of front-door adjustment compared to standard covariate adjustments like regression and matching when estimating the ATT for nonrandomized program evaluations with one-sided noncompliance. Specifically for the National JTPA Study, matching adjustments using the nonexperimental comparison group can come close to the experimental estimates only when one has "detailed retrospective questions on labor force participation, job spells, earnings" (Heckman et al., 1998). However, in the absence of detailed labor force histories, Glynn and Kashin (2013) show that it is possible to create a comparison group that more closely resembles an experimental control group using the front-door approach. Nonetheless, while the front-door approach was shown to be preferable to standard covariate adjustments for the National JTPA Study, front-door estimates for adult males appeared to exhibit positive bias. In this section, we attempt to address this bias using the front-door difference-in-differences approach developed in this paper.

In order to implement the difference-in-differences approach, we focus on currently or once married adult men as the group of interest (henceforth referred to as simply married men).[3] We measure program impact as the ATT on 18-month earnings in the post-randomization or post-eligibility period. For married males, the experimental benchmark is $703.27 .[4] This focus on married men enables us to use single adult men as the differencing group in a front-door difference-in-difference approach. It is likely that Assumption 5 is violated for this differencing group, because the job training program should have effects for single men. However, we anticipate the effects of the training program will be smaller for single men than for married men, due to evidence that marriage improves men's productivity (e.g., see Korenman and Neumark (1991)). Therefore the front-door difference-in-differences approach should provide a lower bound on the job training effect for married men.

The Department of Labor implemented the National JTPA Study between November 1987 and September 1989 in order to gauge the efficacy of the Job Training Parternship Act (JTPA) of 1982. The Study randomized JTPA applicants into treatment and control groups at 16 study sites (referred to as service delivery areas, or SDAs) across the United States. Participants randomized into the treatment group were allowed to receive JTPA services, whereas those in the control group were prevented from receiving program services for an 18-month period following random assignment

[3]Age for adult men ranges from 22 to 54 at random assignment / eligibility screening. Once married men comprises individuals who report that they are widowed, divorced, or separated.

[4]See discussion of how we created our sample and the earnings data in Appendix C.

(Bloom et al., 1993; Orr et al., 1994). Crucially for our analysis, 61.4% of married men allowed to receive JTPA services actually utilized at least one of those services. Moreover, the Study also collected a nonexperimental comparison group of individuals who met JTPA eligibility criteria but chose not to apply to the program in the first place.[5] Since this sample of eligible nonparticipants (ENPs) was limited to 4 service delivery areas, we restrict our entire analysis to only these 4 sites.

4.1 Results

The front-door and front-door difference-in-differences estimates for the effect of the JTPA program on married males - our group of interest - are presented in Figure 1 across a range of covariate sets. Additionally, we present the standard covariate adjusted estimates for comparison. We use OLS separately within experimental treated and observational control groups (the ENPs) for the standard estimates. For front-door estimates, we use OLS separately within the "experimental treated and received treatment" and "experimental treated and didn't receive treatment" groups. Therefore, these estimates assume linearity and additivity within these comparison groups when conditioning on covariates, albeit we note that we obtain similar results when using more flexible methods that relax these parametric assumptions. The experimental benchmark (dashed line), is the only estimate that uses the experimental control units.

For the empty conditioning set, the front-door estimate is slightly above the experimental benchmark. Even without seeing the experimental benchmark, this estimate is likely affected by positive bias because those that fail to show up to the job training program are likely to be less diligent individuals than those that show up. Given the anticipated positive bias in the front-door estimates, we use the front-door difference-in-differences estimator to either recover an unbiased point estimate or obtain a lower bound, depending on our assumptions as to the effect of the program in the differencing group. If we believe that the JTPA program had no effect for single males (i.e., Assumption 5

[5]See Appendix C for additional information regarding the ENP sample. See Smith (1994) for details of ENP screening process.

holds), and we also believe that Assumptions 1 and 6 hold, then the difference-in-differences estimator will return an unbiased estimate of the effect for the group of interest in large samples. If, on the other hand, we believe there might be a non-negative effect for single males, then we would obtain a lower bound for the effect for the group of interest. In this application, it is more likely that there was positive effect of the JTPA program for single males, albeit one smaller than for married males. Hence, the front-door difference-in-differences estimator will likely give us a lower bound for the effect of the JTPA program for married males.

Figure 1: Comparison of standard covariate adjusted estimates, front-door, and front-door difference-in-differences estimates for the JTPA effect for married adult males. The dashed line is the experimental benchmark. 95% bootstrapped confidence intervals are based on 10,000 replicates.

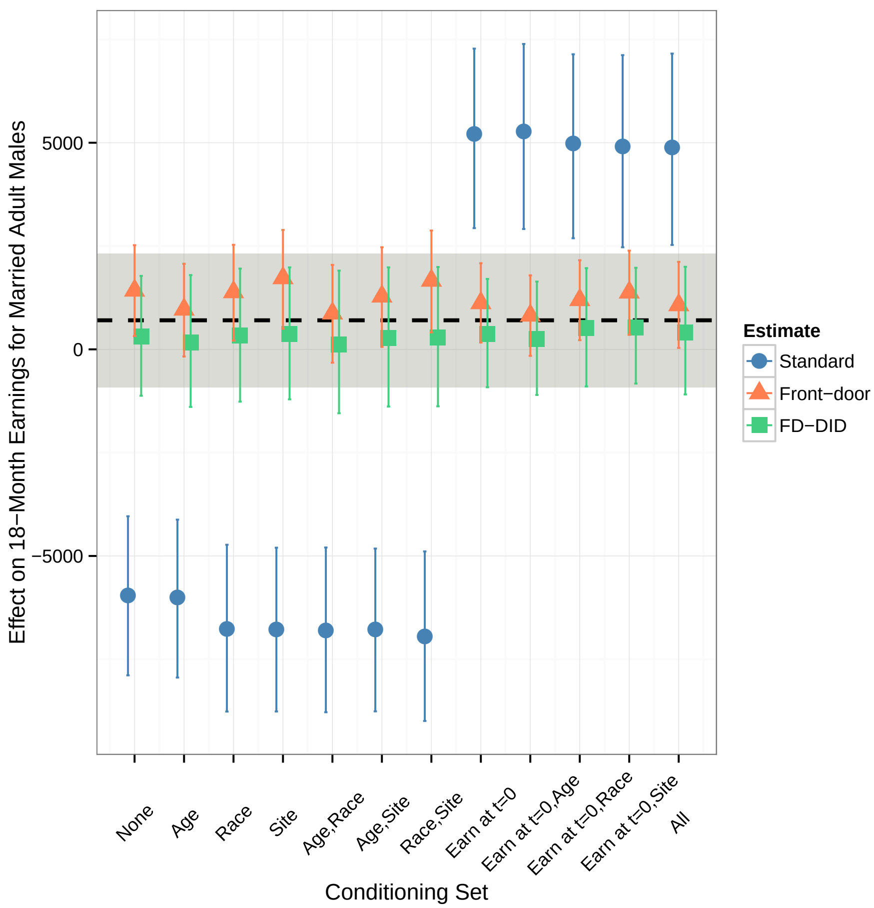

The front-door estimate that we obtain for single males is $946.09 when examining the empty conditioning set. In order to construct the front-door difference-in-differences estimator, we have to scale this estimate by the ratio of compliance for married males to compliance for single males, which is equal to $0.614/0.524 \approx 1.172$. Subtracting the scaled front-door estimate for single males from the front-door estimate for married males as shown in (5), we obtain an estimate of \$315.41. This is slightly below the experimental benchmark and thus indeed functions as lower bound. In sharp contrast to the front-door and front-door difference-in-differences estimates that rather tightly bound the truth, the bias in the standard estimate is -\$6661.90. It is noteworthy that the front-door estimate acts as an upper bound and the front-door difference-in-differences estimate acts as a lower bound across all conditioning sets presented in Figure 1.

5 Illustrative Application: Early Voting

In this section, we present front-door and front-door difference-in-differences estimates for the average treatment effect on the treated (ATT) of an early in-person voting program in Florida. We want to evaluate the impact that the presence of early voting had upon turnout for some groups in the 2008 and 2012 presidential elections in Florida. In traditional regression or matching approaches (either cross sectional or difference-in-differences), data from Florida would be compared to data from states that did not implement early in-person voting. These approaches are potentially problematic because there may be unmeasured differences between the states, and these differences may change across elections. One observable manifestation of this is that the candidates on the ballot will be different for different states in the same election year and for different election years in the same state. The front-door and front-door difference-in-differences approaches allows us to solve this problem by confining analysis to comparisons made amongst modes of voting within a single presidential election in Florida.

Additionally, by restricting our analysis to Florida, we are able to use individual-level data from the Florida Voter Registration Statewide database, maintained since January 2006 by the Florida

Department of State's Division of Elections. This allows us to avoid the use of self-reported turnout, provides a very large sample size, and makes it possible to implement all of the estimators discussed in earlier sections because we observe the mode of voting for each individual. The data contains two types of records by county: registration records of voters contained within *voter extract files* and voter history records contained in *voter history files*. The former contains demographic information - including, crucially for this paper, race - while the latter details the voting mode used by voters in a given election. The two records can be merged using a unique voter ID available in both file types. However, voter extract files are snapshots of voter registration records, meaning that a given voter extract file will not contain all individuals appearing in corresponding voter history file because individuals move in and out of the voter registration database. We therefore use voter registration files from four time periods to match our elections of interest: 2006, 2008, and 2010 book closing records, and the 2012 post-election registration record. Our total population, based on the total unique voter IDs that appear in any of the voter registration files, is 16.4 million individuals. Appendix D provides additional information regarding the pre-processing of the Florida data.

Information on mode of voting in the voter history files allows us to define compliance with the program for the front-door estimator (i.e., those that utilize EIP voting in the election for which we are calculating the effect are defined as compliers). Additionally, we use information on previous mode of voting to partition the population into a group of interest and differencing groups. In order to maximize data reliability, we define our group of interest as individuals that used EIP in a previous election (e.g., 2008 EIP voters are the group of interest when analyzing the turnout effect for the 2012 election). In other words, we are assessing what would have happened to these 2008 EIP voters in 2012 if the EIP program had not been available in 2012. To calculate the EIP effect on turnout for the 2012 election, we separately consider 2008 and 2010 EIP voters as our groups of interest. For the 2008 EIP effect on turnout, we rely upon 2006 EIP voters as our group of interest. An attempt to define the group of interest more broadly (e.g., including non-voters) or in terms of earlier elections (e.g., the 2004 election) would involve the use of less reliable data, and would therefore introduce methodological complications that are not pertinent to the illustration presented here.[6] Therefore, the estimates presented in this application are confined only to those individuals that utilized EIP in a previous election and hence we cannot comment on the overall turnout effect.

We consider two differencing groups for each analysis: those who voted absentee and those that voted on election day in a previous election. When considering the 2012 EIP effect for 2008 EIP voters, for example, we use 2008 absentee and election day voters as our differencing groups. It is likely that the 2012 EIP program had little or no effect for 2008 absentee voters and perhaps only a minimal effect for 2008 election day voters, as these groups had already demonstrated an ability to vote by other means. Therefore, any apparent effects estimated for these groups will be primarily due to bias, and this bias can then be removed from the estimates for the group of interest. As discussed in earlier sections, the estimates from the differencing groups must be scaled according to the level of compliance for the group of interest. Finally, the existence of two differencing groups allows us to conduct a placebo test by using election day voters as the group of interest and the absentee voters as the differencing group in each case. This analysis is explored below.

Despite the limited scope of the estimates presented here, these results have some bearing on the recent debates regarding the effects of early voting on turnout. There have been a number of papers that find null results for the effects of early voting on turnout (Gronke, Galanes-Rosenbaum and Miller, 2007; Gronke et al., 2008; Fitzgerald, 2005; Primo, Jacobmeier and Milyo, 2007; Wolfinger, Highton and Mullin, 2005), and Burden et al. (2014) finds a surprising negative effect of early voting on turnout in 2008.[7] However, identification of turnout effects from observational data using traditional statistical approaches such as regression or matching are hampered by unobserved

[6]Following Gronke and Stewart (2013), we restrict our analysis to data starting in 2006 due to its greater reliability than data from 2004. We also might like to extend the group of interest to those that did not vote in a previous election, but we avoid assessing either 2008 or 2012 EIP effects for these voters because it is difficult to calculate the eligible electorate and consequently the population of non-voters. In their analysis of the prevalence of early voting, Gronke and Stewart (2013) use all voters registered for at least one general election between 2006 and 2012, inclusive, as the total eligible voter pool. However, using registration records as a proxy for the eligible electorate may be problematic (McDonald and Popkin, 2001). By focusing on the 2008 voting behavior of individuals who voted early in 2006, we avoid the need to define the eligible electorate and the population of non-voters.

[7]Burden et al. (2014) examine a broader definition of early voting that includes no excuse absentee voting.

confounders that affect both election laws and turnout (Hanmer, 2009). If these unobserved confounders vary across elections, then traditional difference-in-differences estimators will also be biased.

The front-door estimators presented here provide an alternative approach with useful properties. First, front-door adjustment can identify the effect of EIP on turnout in spite of the endogeneity of election laws that can lead to bias when using standard approaches. Second, unlike traditional regression, matching, or difference-in-differences based estimates, the front-door estimators considered here only require data from Florida within a given year. This means that we can effectively include a Florida/year fixed effect in the analysis, and we do not have to worry about cross-state or cross-time differences skewing turnout numbers across elections. We also include county fixed effects in the analysis in order to control for within-Florida differences.

However, in addition to the limited scope of our analysis, it is important to note that the exclusion restriction is likely violated for this application. Since early in-person voting decreases waiting times on election day, it is possible that it actually increases turnout among those that only consider voting on election day. This would mean that front-door estimates would understate the effect if all other assumptions held because the front-door estimator would be ignoring a positive component of the effect. Alternatively, Burden et al. (2014) suggest that campaign mobilization for election day may be inhibited, such that early voting hurts election day turnout. This would mean that front-door estimates would overstate the effect because the front-door estimator would be ignoring a negative component of the effect. This can also be seen by examining the bias formula (4) (because the EIP treatment is assigned at the state level, Assumptions 1 and 4 will hold).

Taken together, the overall effect of these exclusion restrictions is unclear and would depend on the strength of the two violations. The predictions also become less clear once we consider the front-door difference-in-differences approach, where additional bias in front-door estimates might cancel with bias in the estimates for the differencing group. For the remainder of this analysis, we will assume that all such violations of the exclusion restriction cancel out in the front-door difference-in-differences estimator. This is implicit in Assumption 6.

5.1 Results

The front-door and front-door difference-in-differences estimates for the 2012 EIP program are presented in Figure 2. The estimates all utilize county fixed effects and are calculated separately across the racial categories.[8] The orange estimates are the front-door estimates of the 2012 EIP effect for voters that used EIP in 2008. We anticipate that these estimates exhibit large positive bias because 2008 EIP voters would be more likely to vote in 2012, even in the absence of EIP, than the 2008 non-EIP group (this group includes individuals that did not vote in 2008). In terms of the bias formula in (4), this is equivalent to saying that $E[Y(a_0)|a_1, m_1, x] > E[Y(a_0)|a_1, m_0, x]$.

In order to address this bias, we also present front-door difference-in-differences estimates for the same group of interest (2008 EIP voters) in green, with 2008 absentee voters (triangles) and 2008 election day voters (squares) as the differencing groups. The former, for example, is constructed as the difference between front-door estimates for 2008 early voters and the front-door estimates for 2008 absentee voters, with the front-door estimates for the differencing group scaled by the ratio of early voter compliance to absentee voter compliance as shown in (5). The purple estimates in Figure 2 represent the placebo test, with 2008 election day voters standing in as the group of interest and the absentee voters as the differencing group.

For the 2008 EIP voters, the 2012 EIP front-door difference-in-differences estimates (green) are positive and significant at the 99% level (based on 10,000 block bootstraps at the county level). There is some evidence of differences between the racial categories, but these differences change depending on which differencing group is used. The purple estimates are for the most part indistinguishable from zero, indicating that the placebo tests have mostly been passed. The slightly negative purple

[8]We calculate the FD-DID estimates within each county and then average using the population of the group of interest as the county weight. Due to very small sample sizes in a few counties, we are occasionally unable to calculate front-door estimates. In these cases, we omit the counties from the weighted average when calculating the front-door estimates with fixed effects. We note that due to their small size, these counties are unlikely to exert any significant impact upon the estimates regardless.

estimate for whites again indicates either bias, or perhaps a negative effect of the 2012 EIP program for white 2008 election day voters. Note that even if we believe this estimate, the weighted average of the green and the purple effects for whites (i.e., the 2012 EIP effect for the 2008 EIP and election day voters together) produces a slightly positive estimate, albeit this estimate is indistinguishable from zero. In sum, the evidence points to a slightly positive turnout effect of the 2012 EIP program on the 2008 EIP users.

Figure 2: Front-door estimates for the turnout effect in 2012 for voters who voted early in 2008 (by race). All estimates include county fixed effects. 99% block bootstrapped confidence intervals are based on 10,000 replicates.

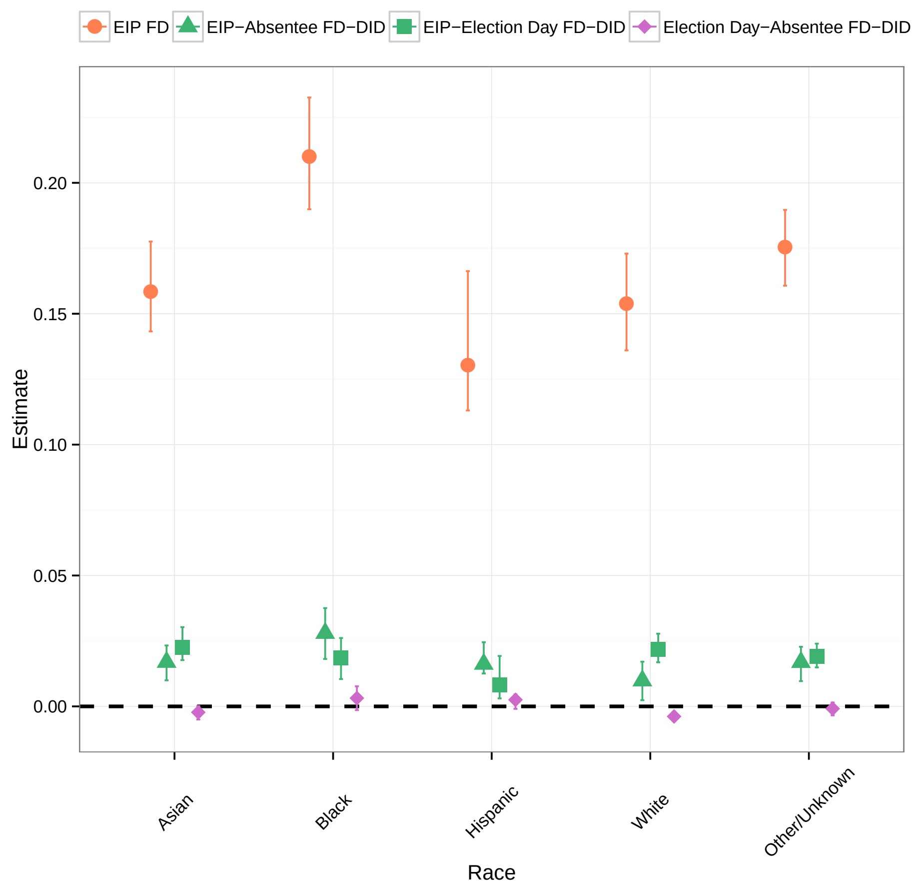

In order to make a more direct comparison to the results of Burden et al. (2014), which finds a negative EIP effect for the 2008 election, we present front-door and front-door difference-in-differences estimates for the 2008 EIP program in Figure 3. We present the estimates for 2006 EIP voters as the group of interest (orange for front-door and green for front-door difference-in-differences), using 2006 absentee voters (triangles) and 2006 election day voters (squares) as the differencing groups. Unfortunately, we cannot present the estimates with the group of interest and the differencing groups defined in terms of 2004 behavior because the data from 2004 are not reliable (as mentioned above). The placebo test is constructed as the difference in front-door estimates between 2006 election day voters and 2006 absentee voters (purple). As before, the EIP program estimates are positive and significant at the 99% level. All placebo tests, with the exception of the white estimate, are indistinguishable from zero, giving us confidence in the estimated EIP effects. Even if the slightly negative placebo estimate for whites indicates a true negative effect of the 2008 EIP program, and not bias, the weighted average of the green and the purple effects (i.e., the 2008 EIP effect for the 2006 EIP and election day voters together), again produces a slightly positive estimate. Therefore, as opposed to Burden et al. (2014), we do not find any evidence that the presence of an EIP program in Florida decreased turnout in 2008. On the contrary, we find evidence that early voting increased turnout for the subset of individuals who voted early in 2006. Moreover, comparing the point estimates across races, we find some evidence that the program had a disproportionate benefit for African-Americans. It is also notable that the size of the estimated EIP effect for 2008 is more than double the estimated EIP effect for 2012 when looking at EIP voters as the group of interest across all races. There are two potential reasons for this. First, since our estimates for the 2008 EIP program are obtained using groups defined by 2006 midterm election behavior, we note that midterm election early voters are likely different than presidential election year voters. Second, the nature of the early voting program changed between the 2008 and 2012 elections, as described in Gronke and Stewart (2013). One of the most significant alterations to the program was a near-halving of the early voting period from 14 days to 8 days. This change might possibly reduce the effect of the EIP program in 2012 when compared to 2008.

In order to isolate the consequences of the change in the early voting program from changes

Figure 3: Front-door estimates for the turnout effect in 2008 for voters who voted early in 2006 (by race). All estimates include county fixed effects. 99% block bootstrapped confidence intervals are based on 10,000 replicates.

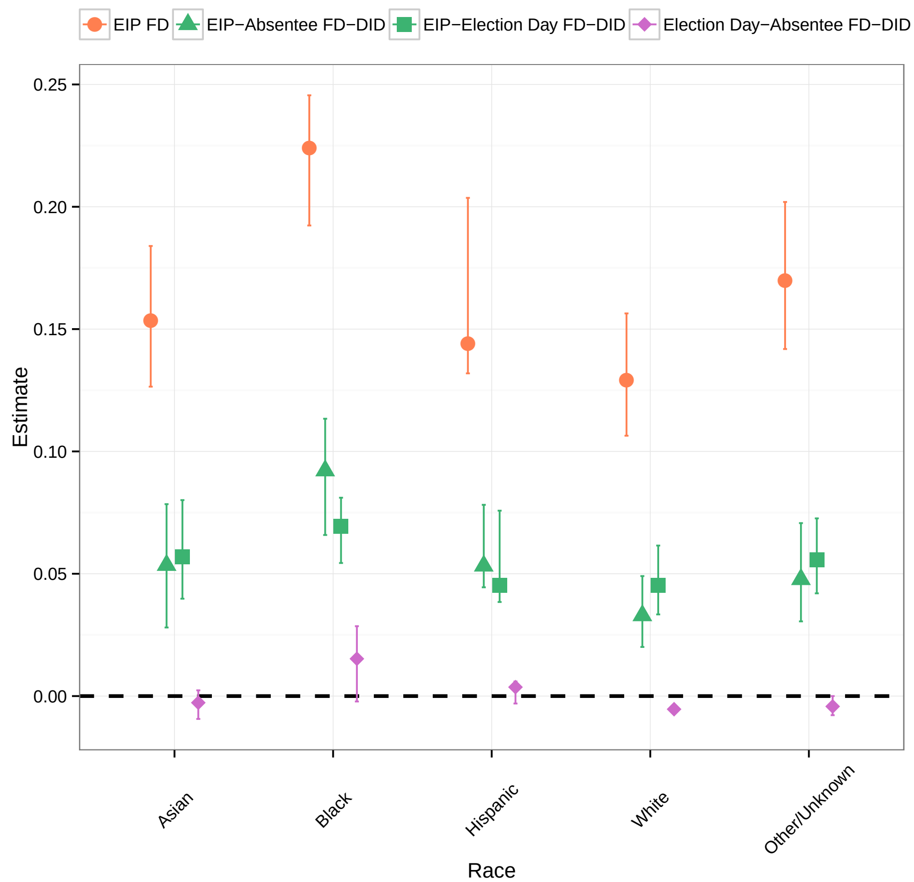

in the construction of the group of interest and differencing groups, we re-estimate the effects of the 2012 EIP program using 2010 EIP voters as the group of interest (orange and green), and using 2010 absentee (triangles) and election day voters (squares) as the differencing groups. Placebo tests are reported using 2010 election day voters as the group of interest and 2010 absentee voters as the differencing group (purple). These results are presented in Figure 4, and they are quite similar to the results in Figure 3. This provides some evidence that if we were able to obtain reliable data from the 2004 election, our estimates for the 2008 EIP program would likely have produced something similar to Figure 2 when using 2004 EIP voters as the group of interest, and 2004 absentee (triangles) and election day voters (squares) as the differencing groups,. However, the estimates in Figure 3 are slightly larger than estimates in Figure 4. This is consistent with the reduction in the early voting window for the 2012 election.

Figure 4: Front-door estimates for the turnout effect in 2012 for voters who voted early in 2010 (by race). All estimates include county fixed effects. 99% block bootstrapped confidence intervals are based on 10,000 replicates.

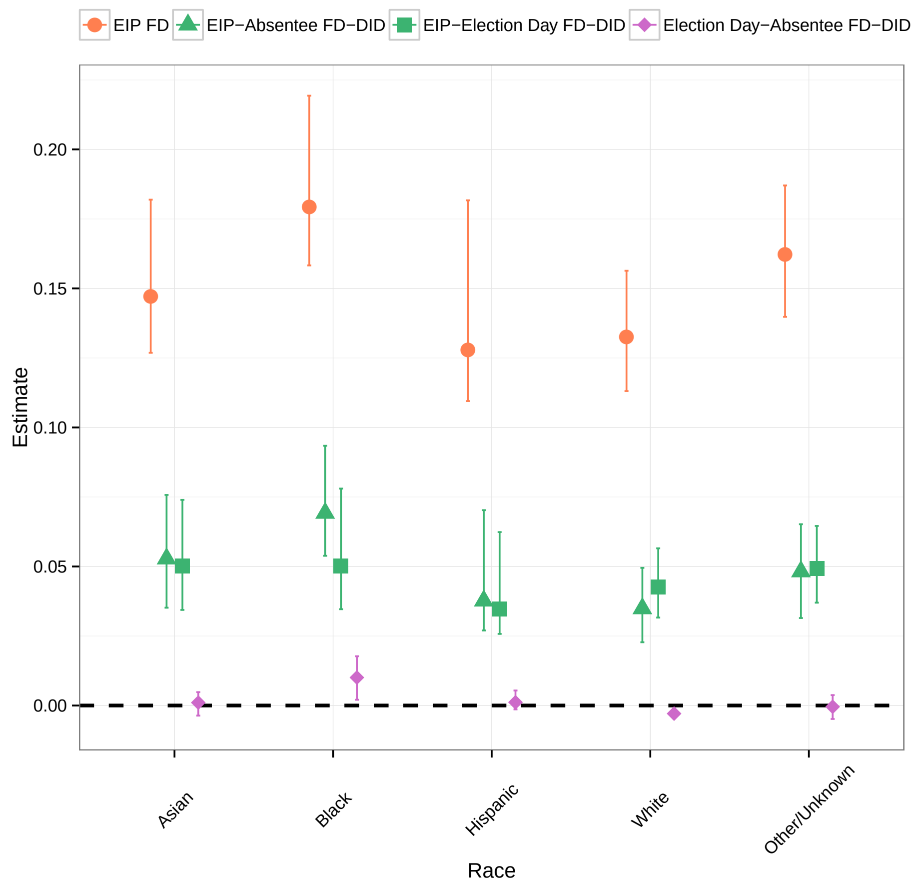

6 Conclusion

In this paper, we have developed front-door difference-in-differences estimators for nonrandomized program evaluations with one-sided noncompliance and an exclusion restriction. These estimators allow for asymptotically unbiased estimation via front-door techniques, even when front-door estimators have significant bias. Furthermore, this allows for program evaluation when all of the relevant units have been assigned to treatment.

We illustrated this technique with an application to the National JTPA (Job Training Partnership Act) Study and with an application to the effects of Florida's early in-person voting program on turnout. For the job training application, we showed that front-door and front-door difference-in-differences could be used to bracket the experimental benchmark. For the application to the effects of an early in-person (EIP) voting program on turnout in Florida in 2008 and 2012, we found that for two separate differencing groups, the program had small but significant positive effects. While the scope of the analysis is limited, this result provides some evidence to counter previous results in the literature that early voting programs had either no effect or negative effects.

Finally, the results in this paper have implications for research design and analysis. First, the examples demonstrate the importance of collecting post-treatment variables that represent compliance with, or uptake of, the treatment. Such information allows front-door and front-door difference-in-differences analyses to be carried out as a robustness check on standard approaches. Second, the bracketing of the experimental benchmark in the JTPA application show that control units are not always necessary for credible causal inference. This is a remarkable finding that should make a number of previously infeasible studies possible (e.g., when it is unethical or impossible to withhold treatment from individuals).

References

Bloom, Howard S, Larry L Orr, George Cave, Stephen Bell and Fred Doolittle. 1993. "The National JTPA Study: Title IIA Impacts on Earnings and Employment at 18 Months." Bethesda, MD: . 14

Brady, Henry E., David Collier and Jason Seawright. 2006. "Toward a Pluralistic Vision of Methodology." *Political Analysis* 14:353–368. 2

Burden, Barry C., David T. Canon, Kenneth R. Mayer and Donald P. Moynihan. 2014. "Election Laws, Mobilization, and Turnout: The Unanticipated Consequences of Election Reform." *American Journal of Political Science* 58(1):95–109. 18, 19, 21, 22

Collier, David and Henry E. Brady. 2004. *Rethinking Social Inquiry: Diverse Tools, Shared Standards*. Lanham, MD: Rowman & Littlefield. 2

Cox, D. R. and Nanny Wermuth. 1995. "Discussion of 'Causal diagrams for empirical research.'" *Biometrika* 82:688–689. 2

Fitzgerald, Mary. 2005. "Greater Convenience but not Greater Turnout: The Impact of Alternative Voting Methods on Electoral Participation in the United States." *American Politics Research* 33:842–867. 18

George, A.L. and A. Bennett. 2005. *Case studies and theory development in the social sciences*. Mit Press. 2

Glynn, A. and K. Kashin. 2013. "Front-door Versus Back-door Adjustment with Unmeasured Confounding: Bias Formulas for Front-door and Hybrid Adjustments." Working Paper, Winner of the Gosnell Prize for Excellence in Political Methodology. 3, 4, 6, 12, 13, 34

Glynn, A. and K. Quinn. 2011. "Why Process Matters for Causal Inference." *Political Analysis* 19(3):273–286. 2

Gronke, P., E. Galanes-Rosenbaum and P. Miller. 2007. "Early Voting and Turnout." *PS: Political Science and Politics* XL. 18

Gronke, Paul and Charles Stewart. 2013. "Early Voting in Florida." Working Paper. 18, 22, 35, 36

Gronke, Paul, Eva Galanes-Rosenbaum, Peter A. Miller and Daniel Toffey. 2008. "Convenience Voting." *Annual Review of Political Science* 11:437–455. 18

Hanmer, Michael J. 2009. *Discount Voting: Voter Registration Reforms and Their Effects*. Cambridge University Press. 19

Heckman, James, Hidehiko Ichimura, Jeffrey Smith and Petra Todd. 1998. "Characterizing selection bias using experimental data." *Econometrica* 66:1017–1098. 13, 34, 35

Heckman, James J and Jeffrey A Smith. 1999. "The pre-programme earnings dip and the determinants of participation in a social programme: implications for simple programme evaluation strategies." *Economic Journal* . 34

Heckman, James J, Robert J LaLonde and Jeffrey A Smith. 1999. The Economics and Econometrics of Active Labor Market Programs. In *Handbook of Labor Economics, Volume III*, ed. O. Ashenfelter and D. Card. Elsevier Science North-Holland. 6

Imbens, Guido and Donald Rubin. 1995. "Discussion of 'Causal diagrams for empirical research.'" *Biometrika* 82:694–695. 2

Joffe, M.M. 2001. "Using information on realized effects to determine prospective causal effects." *Journal of the Royal Statistical Society. Series B, Statistical Methodology* pp. 759–774. 2

Kaufman, Sol, Jay S. Kaufman and Richard F. MacLehose. 2009. "Analytic bounds on causal risk differences in directed acyclic graphs with three observed binary variables." *Journal of Statistical Planning and Inference* 139:3473–87. 2

King, Gary, Robert O. Keohane and Sidney Verba. 1994. *Designing Social Inquiry: Scientific Inference in Qualitative Research*. 1 ed. Princeton University Press. 2

Korenman, Sanders and David Neumark. 1991. "Does Marriage Really Make Men More Productive?" *The Journal of Human Resources* 26(2):282–307. 13

Kuroki, Manabu and Masami Miyakawa. 1999. "Identifiability Criteria for Causal Effects of Joint Interventions." *J. Japan Statist. Soc.* 29(2):105–117. 2

McDonald, Michael P. and Samuel L. Popkin. 2001. "The Myth of the Vanishing Voter." *American Political Science Review* 95:963–974. 18

Orr, Larry L, Howard S Bloom, Stephen H Bell, Winston Lin, George Cave and Fred Doolittle. 1994. "The National JTPA Study: Impacts, Benefits, And Costs of Title IIA." Bethesda, MD: . 14, 35

Pearl, Judea. 1995. "Causal diagrams for empirical research." *Biometrika* 82:669–710. 2

Primo, David M., Matthew L. Jacobmeier and Jeffrey Milyo. 2007. "Estimating the Impact of State Policies and Institutions with Mixed-Level Data." *State Politics & Policy Quarterly* 7:446–459. 18

Shpitser, Ilya and Judea Pearl. 2006. "Identification of Conditional Interventional Distributions." Proceedings of the Twenty Second Conference on Uncertainty in Artificial Intelligence (UAI). 2

Smith, Jeffrey A. 1994. "Sampling Frame for the Eligible Non-Participant Sample." *Mimeo* . 14

Stewart, Charles. 2012. "Declaration of Dr. Charles Stewart III." State of Florida vs. United States of America. 35

Tian, J. and J. Pearl. 2002*a*. A general identification condition for causal effects. In *Proceedings of the National Conference on Artificial Intelligence*. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999 pp. 567–573. 2

Tian, J. and J. Pearl. 2002*b*. On the identification of causal effects. In *Proceedings of the American Association of Artificial Intelligence*. 2

VanderWeele, Tyler J. 2009. "On the relative nature of overadjustment and unnecessary adjustment." *Epidemiology* 20(4):496–499. 2

Wolfinger, Raymond E., Benjamin Highton and Megan Mullin. 2005. "How Postregistration Laws Affect the Turnout of Citizens Registered to Vote." *State Politics & Policy Quarterly* 5:1–23. 18

A ATT Proofs

A.1 Large-Sample Bias

The bias in the front-door estimate of $E[Y(a_0)|a_1]$ is the following:

$$
\begin{aligned}
B^{fd}_{a_1} &= \mu^{fd}_{0|a_1} - \mu_{0|a_1} \\
&= \sum_x \sum_m P(m|a_0, x) \cdot E[Y|a_1, m, x] \cdot P(x|a_1) - \sum_x \sum_u E[Y|a_0, x, u] \cdot P(u|x, a_1) \cdot P(x|a_1) \\
&= \sum_x \sum_m P(m|a_0, x) \sum_u E[Y|a_1, m, x, u] \cdot P(u|a_1, m, x) \cdot P(x|a_1) \\
&- \sum_x \sum_u \sum_m E[Y|a_0, m, x, u] \cdot P(m|a_0, x, u) \cdot P(u|x, a_1) \cdot P(x|a_1) \\
&= \sum_x P(x|a_1) \sum_m \sum_u P(m|a_0, x) \cdot E[Y|a_1, m, x, u] \cdot P(u|a_1, m, x) \\
&- \sum_x P(x|a_1) \sum_m \sum_u P(m|a_0, x, u) \cdot E[Y|a_0, m, x, u] \cdot P(u|a_1, x)
\end{aligned}
$$

Note that the bias will be zero when Y is mean independent of A conditional on U, M, and X (i.e., $E[Y|a_1, m, x, u] = E[Y|a_0, m, x, u]$) and U is independent of M conditional on X and a_0 or a_1 (i.e., $P(m|a_0, x) = P(m|a_0, x, u)$ and $P(m|a_1, x) = P(m|a_1, x, u)$). Hence, again it is possible for the front-door approach to provide an unbiased estimator when there is an unmeasured confounder.

A.2 Front-door Adjustment with One-Sided Noncompliance

In the special case of one-sided noncompliance, the front-door estimator can be written as the following:

$$
\begin{aligned}
\tau_{att}^{fd} &= E[Y|a_1] - \sum_x E[Y|a_1, m_0, x] \cdot P(x|a_1) \\
&= \sum_x E[Y|a_1, x] \cdot P(x|a_1) - \sum_x E[Y|a_1, m_0, x] \cdot P(x|a_1) \qquad \text{(by LTP)} \\
&= \sum_x P(x|a_1) \left\{E[Y|a_1, x] - E[Y|a_1, m_0, x]\right\} \qquad \text{(by rearranging and factoring)} \\
&= \sum_x P(x|a_1) \left\{E[Y|a_1, m_1, x] \cdot P(m_1|x, a_1) + E[Y|a_1, m_0, x] \cdot P(m_0|x, a_1) - E[Y|a_1, m_0, x]\right\} \qquad \text{(by LTP)} \\
&= \sum_x P(x|a_1) \left\{E[Y|a_1, m_1, x] \cdot P(m_1|x, a_1) + E[Y|a_1, m_0, x] \cdot [P(m_0|x, a_1) - 1]\right\} \qquad \text{(by factoring)} \\
&= \sum_x P(x|a_1) \left\{E[Y|a_1, m_1, x] \cdot P(m_1|x, a_1) - E[Y|a_1, m_0, x] \cdot [1 - P(m_0|x, a_1)]\right\} \qquad \text{(by factoring a -1)} \\
&= \sum_x P(x|a_1) \left\{E[Y|a_1, m_1, x] \cdot P(m_1|x, a_1) - E[Y|a_1, m_0, x] \cdot P(m_1|x, a_1)\right\} \qquad \text{(by prob axiom)} \\
&= \sum_x P(x|a_1) P(m_1|x, a_1) \left\{E[Y|a_1, m_1, x] - E[Y|a_1, m_0, x]\right\} \qquad \text{(by factoring)}
\end{aligned}
$$

A.3 Large-Sample Bias Under One-Sided Noncompliance

The front-door and standard covariate adjustment ATT bias can be written as the following, utilizing the fact that $P(m_0|a_0) = 1$ and $P(m_0|a_1) = 0$:

$$
\begin{aligned}
B_{att}^{fd} &= \mu_1 - \mu_{0|a_1}^{fd} - (\mu_1 - \mu_{0|a_1}) \\
&= \mu_{0|a_1} - \mu_{0|a_1}^{fd} \\
&= -B_{a_1}^{fd} \\
&= \sum_x P(x|a_1) \sum_u E[Y|u, a_0, m_0, x] P(u|a_1, x) \\
&\quad - \sum_x P(x|a_1) \sum_u E[Y|u, a_1, m_0, x] P(u|a_1, x, m_0)
\end{aligned}
$$

Adding and subtracting $\sum_x P(x) \sum_u E[Y|a_0, m_0, u] \cdot P(u|a_1, m_0)$:

$$\begin{aligned}
&= \sum_x P(x|a_1) \sum_u E[Y|u, a_0, m_0, x] \cdot [P(u|a_1, x) - P(u|a_1, x, m_0)] \\
&- \sum_x P(x|a_1) \sum_u \{E[Y|u, a_1, m_0, x] - E[Y|u, a_0, m_0, x]\} \cdot P(u|a_1, m_0, x)
\end{aligned}$$

$$\begin{aligned}
B^{std}_{att} &= \mu_1 - \mu^{std}_{0|a_1} - (\mu_1 - \mu_{0|a_1}) \\
&= \mu_{0|a_1} - \mu^{std}_{0|a_1} \\
&= -B^{std}_{a_1} \\
&= \sum_x P(x|a_1) \sum_u E[Y|u, a_0, m_0, x] \cdot [P(u|a_1, x) - P(u|a_0, x)]
\end{aligned}$$

A.4 Front-door Bias Simplification

The front-door bias under one-sided noncompliance can be written as:

$$B^{fd}_{att} = \sum_x P(x|a_1) \sum_u E[Y|a_0, m_0, x, u][\underbrace{P(u|a_1, x) - P(u|a_1, m_0, x)}_{\varepsilon}] \tag{7}$$

$$+ \sum_x P(x|a_1) \sum_u \{\underbrace{E[Y|a_0, m_0, x, u] - E[Y|a_1, m_0, x, u]}_{\eta}\} P(u|a_1, m_0, x). \tag{8}$$

ε can be rewritten as:

$$\begin{aligned}
\varepsilon &= P(u|a_1, x) - P(u|a_1, m_0, x) \\
&= P(u|a_1, m_1, x)P(m_1|a_1, x) + P(u|a_1, m_0, x)P(m_0|a_1, x) - P(u|a_1, m_0, x) \\
&= P(u|a_1, m_1, x)P(m_1|a_1, x) + P(u|a_1, m_0, x)[P(m_0|a_1, x) - 1] \\
&= P(u|a_1, m_1, x)P(m_1|a_1, x) - P(u|a_1, m_0, x)P(m_1|a_1, x) \\
&= P(m_1|a_1, x)[P(u|a_1, m_1, x) - P(u|a_1, m_0, x)].
\end{aligned}$$

We can also expand η as:

$$
\begin{aligned}
\eta &= E[Y|a_0, m_0, x, u] - E[Y|a_1, m_0, x, u] \\
&= E[Y|a_0, x, u] - E[Y|a_1, m_0, x, u] \text{ due to one-sided noncompliance} \\
&= E[Y(a_0)|a_0, x, u] - E[Y|a_1, m_0, x, u] \text{ due to consistency} \\
&= E[Y(a_0)|a_1, x, u] - E[Y|a_1, m_0, x, u] \text{ due to ignorability conditional on } u \text{ and } x \\
&= E[Y(a_0)|a_1, m_1, x, u]P(m_1|a_1, x, u) + E[Y(a_0)|a_1, m_0, x, u]P(m_0|a_1, x, u) - E[Y|a_1, m_0, x, u] \\
&= E[Y(a_0)|a_1, m_1, x, u]P(m_1|a_1, x, u) - E[Y(a_0)|a_1, m_0, x, u] \cdot \{\frac{E[Y|a_1, m_0, x, u]}{E[Y(a_0)|a_1, m_0, x, u]} - P(m_0|a_1, x, u)\} \\
&= P(m_1|a_1, x) \left\{ E[Y(a_0)|a_1, m_1, x, u]\frac{P(m_1|a_1, x, u)}{P(m_1|a_1, x)} - E[Y(a_0)|a_1, m_0, x, u] \cdot \frac{\frac{E[Y|a_1, m_0, x, u]}{E[Y(a_0)|a_1, m_0, x, u]} - P(m_0|a_1, x, u)}{P(m_1|a_1, x)} \right\}.
\end{aligned}
$$

We note that the bias can be written as scaled by the compliance proportion within levels of x ($P(m_1|a_1, x)$).

We can thus rewrite front-door bias under one-sided noncompliance as:

$$
\begin{aligned}
B_{att}^{fd} = &\sum_x P(x|a_1)P(m_1|a_1, x) \sum_u \Bigg[E[Y|a_0, m_0, x, u] \cdot [P(u|a_1, m_1, x) - P(u|a_1, m_0, x)] \\
&+ \left\{ E[Y(a_0)|a_1, m_1, x, u]\frac{P(m_1|a_1, x, u)}{P(m_1|a_1, x)} - E[Y(a_0)|a_1, m_0, x, u] \cdot \frac{\frac{E[Y|a_1, m_0, x, u]}{E[Y(a_0)|a_1, m_0, x, u]} - P(m_0|a_1, x, u)}{P(m_1|a_1, x)} \right\} P(u|a_1, m_0, x) \Bigg].
\end{aligned}
$$

A.5 Front-door Bias Under Assumption 3

Assumption 3 and binary M implies that $\varepsilon = 0$:

$$
\begin{aligned}
P(m_1|a_1, x, u) &= \frac{P(u|a_1, m_1, x) \cdot P(m_1|a_1, x)}{P(u|a_1, x)} \qquad \text{(by Bayes' Rule)} \\
1 &= \frac{P(u|a_1, m_1, x)}{P(u|a_1, x)} \qquad \text{(by Assumption 3)} \\
P(u|a_1, x) &= P(u|a_1, m_1, x)
\end{aligned}
$$

Since M is binary, by similar logic as above we know that $P(u|a_1, x) = P(u|a_1, m_0, x)$.

Therefore:

$$\begin{aligned}
\varepsilon &= P(m_1|a_1,x)[P(u|a_1,m_1,x) - P(u|a_1,m_0,x)] \\
&= P(m_1|a_1,x)[P(u|a_1,x) - P(u|a_1,x)] \\
&= 0
\end{aligned}$$

Under Assumption 3, we can simplify front-door bias to:

$$\begin{aligned}
B^{fd}_{att} = \sum_x P(x|a_1)P(m_1|a_1,x)\sum_u \Bigg[& E[Y(a_0)|a_1,m_1,x,u]\frac{P(m_1|a_1,x,u)}{P(m_1|a_1,x)} \\
& - E[Y(a_0)|a_1,m_0,x,u]\cdot \frac{\frac{E[Y|a_1,m_0,x,u]}{E[Y(a_0)|a_1,m_0,x,u]} - P(m_0|a_1,x,u)}{P(m_1|a_1,x)}\Bigg]\cdot P(u|a_1,x).
\end{aligned}$$

B Front-Door Difference-in-Differences Proofs

B.1 No Large-Sample Bias in the Front-door Difference-in-Differences Estimator

First define $\tau_{att,x} = E[Y(a_1)|a_1,x] - E[Y(a_0)|a_1,x]$. It is well known that $\tau_{att} = \sum_x \tau_{att,x}P(x|a_1)$. Therefore in order to show that τ^{fd-did}_{att} has no bias, we need only show a lack of bias for $\tau_{att,x}$ within levels of x. If Assumptions 5 and 6 hold, then the front-door difference-in-differences estimator has no large-sample bias:

$$\begin{aligned}
\tau^{fd-did}_{att,x} &= \tau^{fd}_{att,x,g_1} - \frac{P(m_1|a_1,x,g_1)}{P(m_1|a_1,x,g_2)}\tau^{fd}_{att,x,g_2} \\
&= \tau_{att,x} + B^{fd}_{att,x,g1} - \frac{P(m_1|a_1,x,g_1)}{P(m_1|a_1,x,g_2)}\tau^{fd}_{att,x,g_2} \\
&= \tau_{att,x} + B^{fd}_{att,x,g1} - \frac{P(m_1|a_1,x,g_1)}{P(m_1|a_1,x,g_2)}B^{fd}_{att,x,g_2} \text{ by Assumption 5} \\
&= \tau_{att,x} + B^{fd}_{att,x,g1} - B^{fd}_{att,x,g1} \text{ by Assumption 6} \\
&= \tau_{att,x}
\end{aligned}$$

C National JTPA Study Data

Our analysis makes use of the following samples in the National JTPA Study: experimental active treatment group, experimental control group, and the nonexperimental / eligible nonparticipant (ENP) group. Note that the active treatment group for our purposes means receving any JTPA service, even though the services actually received from the JTPA varied across individuals.[9] We restrict our attention to the 4 *service delivery areas* at which the ENP sample was collected: Fort Wayne, IN; Corpus Christi, TX; Jackson, MS, and Providence, RI. In this analysis, we only examine adult males and follow the the sample restrictions in Appendix B1 of Heckman et al. (1998) and Glynn and Kashin (2013) to construct our final sample. The final sample sizes (by marital status) are presented in Table 1.

Table 1: Sample sizes for adult males by marital status. The treated units are broken up into compliers (C) and noncompliers (NC). Control denotes experimental control and ENP denotes the eligible nonparticipants.

	Treated		Control	ENP
	C	*NC*		
Non-single	484	304	274	292
Single	350	318	266	92

The raw data and edited analysis files are available as part of the National JTPA Study Public Use Data from the Upjohn Institute. The covariates for the experimental sample are available through the background information form (BIF) and the covariates for ENPs are available through the long baseline survey (LBS). The experimental samples completed the BIF, which contains demographic information, social program participation, and training and education histories, at the time of random assignment. The ENPs completed the LBS anywhere from 0 to 24 months following eligiblity screening. Unlike the BIF which mostly covers the previous year in terms of labor market experiences, the LBS covers the past 5 years prior to the survey date and thus provides a much richer portrait of labor market participation. Moreover, experimental control units at the 4 ENP sites also received the long baseline survey, completed 1-2 months after random assignment. Heckman et al. (1998), Heckman and Smith (1999), and related works rely on the detailed labor force participation data and earnings histories in LBS to identify selection bias by comparing the experimental control units to the nonexperimental control units. Unfortunately, treated units were never administered the LBS and we have no detailed labor force participation data for multiple years prior to random assignment. Moreover, no one survey instrument was administered to all three of the samples we are using in this analysis, yielding issues of noncomparability. The limited set of covariates we use in the conditioning sets in our analysis have all been established to be comparable by verifying their values across the BIF and LBS for the experimental control group, which completed both surveys at the 4 ENP sites.

The dataset we use was obtained in communication with Jeffrey Smith and Petra Todd. It is the dataset used in the estimates presented in Section 11 of Heckman et al. (1998) and contains all three samples we use in our analysis. It also contains compliance information for the experimental treated group sample. The covariates we utilize in our analysis have been cross-checked against the raw data

[9]The National JTPA Study classified services received into the following 6 categories: classroom training in occupational skills, on-the-job training, job search assistance, basic education, work experience, and miscellaneous.

from the Upjohn Institute. There are also additional covariates in the Heckman et al. (1998) data that have been imputed using a linear regression as described in Appendix B3 of their paper.

The outcome variable we use in the analysis is total 18-month earnings in the period following random assignment (for experimental units) or eligiblity screening (for ENPs). The monthly total earnings variable available from the public use data files is the `totearn` variable. The data covers months 1-30 after random assignment (denoted as $t+1$ to $t+30$, where t is the time of random assignment). The data also includes data for t, the month of random assignment. Note that this variable is not raw earnings data, but was constructed by Abt Associates from the First and Second Follow-up Surveys, as well as based on data from state unemployment agencies, for the initial JTPA report.[10] Please consult Appendix A of Orr et al. (1994) for description of the First Follow-up Survey, Second Follow-up Survey, and earnings data from state unemployment insurance agencies and Appendix B of the same report for construction and imputation of the 30-month earnings variables. The Narrative Description of the National JTPA Study Public Use Files also contains description of the imputation process (see `http://www.upjohninst.org/erdc/njtpa.html`).

In our analysis, we rely upon the monthly total earnings variable in the dataset we obtained from Jeffrey Smith and Petra Todd. We have verified the earnings data used in the calculation of the program impact from this dataset against the earnings variables in the public use data and they match exactly except for a few individuals where Heckman et al. (1998) have imputed missing monthly data. This applies to around 1% of observations and thus is unlikely to substantively change any results. A unit-by-unit comparison of earnings across the raw data and the data we are using can be obtained from us upon request. Note also that some individuals had missing earnings data for some months. In the construction of the 18-month total earnings variable, we mean impute the missing months using the average of the individual's available monthly earnings. Details on the extent of missingness are available from authors upon request.

D Florida Voting Data

To construct our population of eligible voters, we examine individuals that have appeared in one of four voter registration snapshots: book closing records from 10/10/2006, 10/20/2008, and 10/18/2010, as well as a 2012 election recap record from 1/4/2013. This yields a total population of 16,371,725 individuals that we are able to subset by race (Asian / Pacific Islander, Black (not Hispanic), Hispanic, White (Not Hispanic), and Other). Note that the Other category contains individuals who self-identify as American Indian / Alaskan Native, Multiracial, or Other, as well as individuals for whom race is unknown. In cases where race changes across the voter registration records, we use the latest available self-reported race. Such changes affect only 1.1% of observations. The breakdown of the the population by race is presented in the rightmost column of Table 2.

We use voter history files from 08/03/2013 to subset the population by voting mode in each election. The voter history files required pre-processing before we could use them for estimation. As mentioned in Gronke and Stewart (2013) and Stewart (2012), there is an issue of duplication of voter identification numbers within the same election. In some cases, this duplication is rather

[10] One of the major imputations was a decision to divide raw earnings by a `shares` variable which adjust earnings reported for incomplete months (due to the timing of the interviews) to full monthly earnings.

innocuous because the voting mode is identical across records. In these cases, we simply remove duplicate records and include the voter in our analysis. In other cases, voters are recorded as both having voted in a given election and not having voted (code "N"). In these cases, we assume that the voter did indeed cast a ballot and use that code. Finally, there are a few instances in which a voter is recorded to have voted in multiple ways. For example, a voter history file may indicate that a voter voted both absentee and early at a given election. While Gronke and Stewart (2013) indicates that voters may legitimately appear multiple times in the voter history file, this makes the task of stratifying by voting mode difficult. As a result, we choose to exclude individuals who are recorded to have voted using more than one mode. When analyzing the 2008 election subsetting by 2006 voting modes, we exclude 385 individuals. The corresponding numbers for analysis of the 2012 election subsetting using 2008 and 2010 voting groups are 1951 and 2998, respectively. These figures are dwarfed by the sample sizes and thereby highly unlikely to exert any serious effect upon our estimates.

We also made several choices regarding the definition of voting modes. Specifically, we classified anyone who voted absentee (code "A") and whose absentee ballot was not counted (code "B") as having voted absentee. We classified anyone who voted early (code "E") and anyone cast a provisional ballot early (code "F") as having voted early. Finally, we classify anyone who voted at the polls (code "Y") and cast a provisional ballot at the polls (code "Z") as having voted on election day. We do not use the code "P", which indicates that an individual cast a provisional ballot that was not counted since we cannot ascertain whether it was cast on election day, early, or as an absentee voter.

Another difficulty with the data is defining the eligible electorate and thus individuals who did not vote. While the voter history files have a code "N" for did not vote, most individuals who do not vote are not present in the voter history files at all. For example, for the 2008 election there were no "N" codes at all in the voter history files. Therefore, we count an individual as not having voted in a given election if they appeared in the voter registration files at one point but are either not present in the voter history file for that election or are coded as "N".

Table 2: Voting modes as percent of population in 2006, 2008, and 2010 elections. Note that percentages of individuals who did not vote, whose provisional ballots were not counted, or who are dropped due to conflicting voting modes are not shown.

	2006			2008			2010			
Race	Early	Absentee	Election Day	Early	Absentee	Election Day	Early	Absentee	Election Day	Total
Asian	2.87	2.61	12.21	15.15	10.45	20.43	4.75	5.69	13.76	233664
Black	2.81	1.85	16.86	27.67	7.14	16.97	6.41	4.24	18.57	2159473
Hispanic	2.46	2.83	12.46	15.05	8.60	22.72	4.11	6.01	13.37	2049683
White	5.89	5.60	23.16	14.48	13.57	25.32	7.39	9.05	20.63	11179293
Other	2.60	2.43	11.97	13.67	8.12	19.31	3.87	4.19	12.26	749612

Table 3: Compliance as percent of voting groups in 2006-2008, 2008-2012, and 2010-2012 transitions. Note that percentages of individuals who did not vote, whose provisional ballots were not counted, or who are dropped due to conflicting voting modes are not shown.

	2006-2008			2008-2012			2010-2012		
Race	Early	Absentee	Election Day	Early	Absentee	Election Day	Early	Absentee	Election Day
Asian	55.43	12.14	27.59	37.48	8.50	12.68	57.16	9.12	24.65
Black	71.53	18.66	53.70	49.12	12.35	19.48	67.91	15.05	43.65
Hispanic	56.70	8.98	29.50	33.18	6.81	11.02	53.03	7.31	21.63
White	53.87	9.19	22.98	40.54	7.69	12.33	56.69	7.11	20.41
Other/Unknown	55.42	10.43	28.10	37.15	7.33	11.06	57.43	8.28	24.50