

Front-door Difference-in-Differences Estimators^{*}

Adam Glynn[†] Konstantin Kashin[‡]

aglynn@emory.edu kvkashin@gmail.com

This article has now been published at the *American Journal of Political Science*, Article DOI: 10.1111/ajps.12311

Abstract

We develop front-door difference-in-differences estimators as an extension of front-door estimators. Under one-sided noncompliance, an exclusion restriction, and assumptions analogous to parallel trends assumptions, this extension allows identification when the front-door criterion does not hold. Even if the assumptions are relaxed, we show that the front-door and front-door difference-in-differences estimators may be combined to form bounds. Finally, we show that under one-sided noncompliance, these techniques do not require the use of control units. We illustrate these points with an application to a job training study and with an application to Florida's early in-person voting program. For the job training study, we show that these techniques can recover an experimental benchmark. For the Florida program, we find some evidence that early in-person voting had small positive effects on turnout in 2008. This provides a counterpoint to recent claims that early voting had a negative effect on turnout in 2008.

Replication Materials: The data, code, and any additional materials required to replicate all analyses in this article are available on the American Journal of Political Science Dataverse within the Harvard Dataverse Network, at: Glynn, Adam; Kashin, Konstantin, 2016, "Replication Data for: Front-door Difference-in-Differences Estimators", doi:10.7910/DVN/XZFHCP, Harvard Dataverse, DRAFT VERSION [UNF:6:jBjpD+5b0iC10Cx0GkYKw==]

^{*}We thank Barry Burden, 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. Earlier versions of this paper were presented at the 2014 MPSA Conference and the 2014 Asian Political Methodology Meeting in Tokyo.

[†]Department of Political Science, Emory University, 327 Tarbutton Hall, 1555 Dickey Drive, Atlanta, GA 30322 (<http://scholar.harvard.edu/aglynn>).

[‡]Data Scientist, The Institute for Quantitative Social Science, Harvard University, 1737 Cambridge St, Cambridge MA 02138.

Introduction

In this paper, we develop front-door difference-in-differences estimators as an extension of front-door estimators (Pearl, 1995). This extension removes bias when there is a violation of the key front-door assumption. In this sense, front-door difference-in-differences estimators are analogous to standard difference-in-differences estimators (Abadie, 2005). Importantly, under a certain set of assumptions detailed below, front-door difference-in-differences estimators allow credible causal inference without the use of traditional control units. As we demonstrate with assessments of a job training program and an early in-person voting program, the ability to make inferences in the absence of comparable control units provides an important research strategy.

Review of Front-door

The front-door criterion (Pearl, 1995) and its extensions (Kuroki and Miyakawa, 1999; Tian and Pearl, 2002a,b; Shpitser and Pearl, 2006) provide 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. Figure 1 presents the directed acyclic graph associated with the front-door criterion. The formal definition of this graph can be found in Pearl (1995, 2009), but for our purposes, it will suffice to note the following: A represents the treatment/action variable, M represents a set of mediating variables (often a singleton), Y represents the outcome, X represents covariates, U and V represent sets of unobserved variables, and arrows represent the possible existence of effects from one set of variables to another.¹ Solid arrows are allowed for the front-door criterion to hold; dashed arrows are not allowed for the front-door criterion to hold. Note the existence of solid arrows from U to both

¹To simplify presentation, we have not included arrows between X , U , and V . While the graph implies that these sets of variables are independent, this is not required for the techniques below.

A and Y . Hence, unmeasured common causes of the treatment and outcome are allowed. As seen in Figure 1, the front-door approach works by identifying the effects of A on M and the effects of M on Y , and then putting them back together into the effect of interest: A on Y .

[Figure 1 about here.]

However, the front-door adjustment 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). These assumptions are represented by the dashed arrows in Figure 1. Hence, while common causes of A and Y are allowed for the front-door criterion to hold, common causes of M and Y (not mediated by A) are not allowed. Additionally, the front-door criterion will not hold when A has a direct effect on Y .

Difference-in-Differences with Front-Door

A standard difference-in-differences (DD) estimator uses observations for which there should be no effect (often pre-treatment observations for treatment and control units) to estimate and remove the bias from a selection on observables approach (often post-treatment observations for treatment and control observations). The front-door difference-in-differences (front-door DD) approach developed in this paper works in a similar manner to DD with two major differences. First, the goal of front-door DD is to remove bias due to common causes of M and Y . Therefore the differencing is done with respect to the effect of the mediator, and then the estimated “mediator effect” is scaled to estimate the effect of the treatment. Second, front-door DD with pre-treatment observations is only possible when mediator information is available from the pre-treatment period. This information may not be available in repeated cross-section designs.

With these differences in mind, the front-door DD proceeds analogously to the DD approach. In the over-time version, the estimated effect of the mediator M in the pre-treatment period is subtracted from the estimated effect of the mediator M in the post-treatment period (after adjusting for covariates). This corrected mediator effect is then scaled according to the estimated effect of A

on M , so the effect of A on Y can be estimated. In non-over-time examples, we may find units of observation for which we know there can be no effect. These are analogous to the pre-treatment observations from an over-time approach. Generally, we refer to such observations, for which there can be no effect, as the differencing group and observations for which there can be an effect (e.g., post-treatment observations) as the group of interest. For an over-time example, we consider a job training program with the pre-program observations on individuals as the differencing group. In a non-over-time example, we estimate the effects of an early in-person (EIP) voting program on turnout for elections by leveraging voters that used an absentee ballot in the previous election as a differencing group.²

As with DD estimation, front-door DD estimation will only provide credible estimates when the assumptions hold. Outside of perhaps over-time front-door DD analysis, it will often be hard to make this case. However, we demonstrate in the applications that credible bounds can sometimes be estimated when the assumptions don't hold. In particular, we demonstrate that front-door and front-door DD estimates can be used in a bracketing approach under certain circumstances (although great care must be taken that such circumstances hold).

Bias for the Front-Door Approach for ATT

In this section, we simplify and restate some of the results from [Glynn and Kashin \(2013\)](#) on large-sample bias formulas for the front-door approach to estimating the average treatment effect on the treated (ATT). Throughout this paper, all references to bias will mean large-sample bias in the context of nonparametric estimation. This allows us to avoid questions of modeling.

ATT is often the parameter of interest when assessing the effects of a program or law. For an

²EIP was unlikely to have a large effect *on turnout* for these voters, as they had already demonstrated their ability to vote by another means. The application discusses this in detail.

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 $Y(a_0)$.³ 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 (perhaps with additional data) 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 potentially 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, identification would assume 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 (positivity holds).

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

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

³Note 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).

Conditioning on a_1 is a slight adjustment from the original front-door formula (Pearl, 1995) that targets the average for the treated units instead of all units. We can thus define the large-sample front-door estimator of ATT as:

$$\tau_{att}^{fd} = \mu_{1|a_1} - \mu_{0|a_1}^{fd}.$$

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 (Heckman, LaLonde and Smith, 1999) on program evaluation, we define the program impact as the ATT where the active treatment (a_1) is assignment into a program, and when $M = m_1$ indicates if active treatment was actually received and $M = m_0$ if it was not.

Assumption 1 (One-sided noncompliance)

$$P(m_0|a_0, x) = P(m_0|a_0, x, u) = 1 \text{ for all } x, u.$$

Assumption 1 implies that only those assigned to treatment can receive treatment.⁴ The front-door

⁴One example where one-sided noncompliance will hold trivially is when, due to logistical or ethical concerns, a treatment cannot be withheld from any individual. Additionally, we might wonder whether the effect of treatment assignment would still be of interest in this circumstance. The effect of treatment assignment (often known as the intent-to-treat effect) is often of interest when assignment is manipulable as a policy variable and compliance is not (Heckman, Ichimura, Smith and Todd, 1998).

large-sample estimator can be re-written in the following manner.

$$\begin{aligned}
\tau_{att}^{fd} &= \mu_{1|a_1} - \mu_{0|a_1}^{fd} \\
&= E[Y|a_1] - \sum_x \sum_m P(m|a_0, x) \cdot E[Y|a_1, m, x] \cdot P(x|a_1) \\
&= E[Y|a_1] - \sum_x \underbrace{E[Y|a_1, m_0, x]}_{\text{treated non-compliers}} \cdot P(x|a_1) \tag{1}
\end{aligned}$$

$$= \sum_x P(x|a_1) \cdot P(m_1|x, a_1) \cdot \left\{ \underbrace{E[Y|a_1, m_1, x] - E[Y|a_1, m_0, x]}_{\text{"effect" of receiving treatment}} \right\} \tag{2}$$

The formulas in (1) and (2) do not rely upon outcomes of control units in the construction of proxies for the potential outcomes under control for treated units (see Section 1.1 of the Supporting Information, henceforth SI, 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:

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

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, one must effectively assume that 1) assignment to

treatment has no effect if treatment is not received and 2) after conditioning, those that are assigned but don't receive treatment are comparable in some sense to those that receive treatment.⁵ This will be made more precise below.

The front-door formula in (2), with the 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 as the following (see Sections 1.2 and 1.3 of the SI for proofs):

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

The unobservable portion of this bias formula (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:

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

⁵This second assumption will only be tenable in limited circumstances because we often believe those that receive treatment are fundamentally incomparable to those that don't receive treatment. This motivates our extension to the front-door DD estimator.

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)

$$P(m_1|a_1, x, u) = P(m_1|a_1, x) \text{ for all } x \text{ and } u,$$

then due to the binary measure of treatment received, $P(u|a_1, m_1, x) = P(u|a_1, m_0, x)$ (see Section 1.4 of the SI), and the bias simplifies:

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

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 measured at the cluster level. We present an example of this in the early voting 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 to:

$$B_{att}^{fd} = \sum_x P(x|a_1)P(m_1|a_1, x) \left\{ E[Y(a_0)|a_1, m_1, x] - E[Y(a_0)|a_1, m_0, x] \cdot \frac{\frac{E[Y|a_1, m_0, x]}{E[Y(a_0)|a_1, m_0, x]} - P(m_0|a_1, x)}{P(m_1|a_1, x)} \right\} \quad (4)$$

Finally, when both Assumption 2 and Assumption 4 hold, the remaining bias is due to an unmea-

shared common cause of compliance and the outcome:

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

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.

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 , and we maintain Assumption 1 (one-sided noncompliance), 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 DD estimator can be written as:

$$\tau_{att,g_1}^{fd-did} = \sum_x P(x|a_1, g_1) \left[\tau_{att,x,g_1}^{fd} - \frac{P(m_1|a_1, x, g_1)}{P(m_1|a_1, x, g_2)} \tau_{att,x,g_2}^{fd} \right] \quad (5)$$

$$\begin{aligned} &= \sum_x P(x|a_1, g_1) P(m_1|x, a_1, g_1) \left[\{E[Y|a_1, m_1, x, g_1] - E[Y|a_1, m_0, x, g_1]\} \right. \\ &\quad \left. - \{E[Y|a_1, m_1, x, g_2] - E[Y|a_1, m_0, x, g_2]\} \right]. \end{aligned} \quad (6)$$

Hence, (5) shows that within levels of x , the front-door DD estimator for the group of interest 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 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 $\{E[Y|a_1, m_1, x, g_2] - E[Y|a_1, m_0, x, g_2]\}$ component of group 2 (remove bias from the “mediator effect”).

In order for the front-door DD estimator to remove the large sample bias from the front-door estimator of the ATT for the group of interest, we need the following assumption to hold (where we denote bias within levels of x for the interest group g_1 as B_{att,x,g_1}^{fd}):

Assumption 5 (Bias for g_1 equal to scaled front-door formula for g_2 within levels of x)

$$B_{att,x,g_1}^{fd} = \frac{P(m_1|a_1,x,g_1)}{P(m_1|a_1,x,g_2)} \tau_{att,x,g_2}^{fd} \text{ for all } x.$$

There are two things to note about Assumption 5. First, when using an over-time approach, the compliance rates of the two groups will be equal ($P(m_1|a_1, x, g_1) = P(m_1|a_1, x, g_2)$), because time does not alter an individual’s definition as a complier. Hence, Assumption 5 simplifies to $B_{att,x,g_1}^{fd} =$

τ_{att,x,g_2}^{fd} for all x in the over-time case. Second, Assumption 5 can often be weakened if only a bound is needed. For example, if the estimated effect for the differencing group is positive, and we believe the front-door bias for the group of interest is also positive, but smaller than the scaled estimated effect for the differencing group, then subtracting the scaled estimated effect from the differencing group will remove too much from the estimated effect in the group of interest. Hence the front-door difference-in-differences approach will produce a lower bound.

Therefore, if we believe that the front-door estimator and front-door DD estimator have bias of different signs, then these can be used in a bracketing approach. For example, if we believe the bias in the front-door estimator is positive prior to the differencing, and we believe the bias of the front-door DD estimator is negative, then the front-door and front-door DD estimator can be used together to bracket the truth in large samples. This will be discussed in the context of the illustrative applications in the following sections.

If Assumptions 1 and 5 hold, then τ_{att}^{fd-did} has no large-sample bias for τ_{att} (see SI 2.1 for proof). However, the interpretation of Assumption 5 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. When Assumptions 1 through 4 hold, then Assumption 5 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.

Illustrative Application: National JTPA Study

We now illustrate how front-door and front-door DD estimates for the average treatment effect on the treated (ATT) can be used to estimate and 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 (see SI Section 3 for details). We measure program impact as the ATT on 18-month earnings in the post-randomization or post-eligibility period, where active treatment is assignment into the program (perhaps self-selected assignment). We focus on the effect of sign-up on earnings for three reasons: 1) we can compare front-door estimates to the experimental benchmark, 2) this effect is the same parameter of interest as in much of the econometrics literature utilizing JTPA data (Heckman, Ichimura and Todd, 1997; Heckman and Smith, 1999), and 3) this is often the policy-relevant causal effect when considering whether or not to extend the opportunity for job training. Furthermore, Heckman et al. (1998) showed that 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.” In the following, we discuss the use of front-door DD estimators to provide similar information in the absence of detailed labor force histories, and in fact, in the absence of any individuals that did not sign up for the program.

As discussed below, the simple front-door estimator is anticipated to exhibit positive bias when estimating ATT for the JTPA program for adult males. In the following subsections, we consider two front-door DD approaches to correcting this bias. First, we consider using an over-time approach to remove positive bias from the front-door estimator. Second, we consider the more conservative approach of using single adult males as a differencing group, which allows us to provide a lower bound on the effect of the program for married adult males (due to the job training effect being smaller for single males (Korenman and Neumark, 1991)). Because the front-door estimator provides an upper bound, these two estimators can be used in a bracketing approach.

Results: Over-Time Differencing

The most simple front-door estimator for the effects of the JTPA program takes the mean 18-month earnings of those that both signed up for the program and showed up for their training and subtracts

the mean 18-month earnings of those that signed up for the program but failed showed up for their training,⁶ and then scales this estimate by the rate at which those that signed up actually showed up. Because we have not yet used covariates, this estimator can be written as a simplified version of (2):

$$\tau_{att}^{fd} = P(m_1|a_1) \cdot \left\{ \underbrace{E[Y|a_1, m_1] - E[Y|a_1, m_0]}_{\text{"effect" of receiving treatment}} \right\},$$

where a_1 indicates signing up for the program, m_1 indicates showing up for the program, m_0 indicates failing to show up for the program, and Y denotes 18-month earnings. Because those that show up are likely to be more diligent/disciplined than those that fail to show up, we expect this estimator to be positively biased.

In an attempt to remove the anticipated positive bias, we can use the baseline earnings of these individuals as a differencing group. The most simple version of this estimator does the following: a) takes the mean 18-month earnings of those that both signed up for the program and showed up for their training and subtracts the mean 18-month earnings of those that signed up for the program but failed showed up for their training, b) takes the mean baseline (0-month) earnings of those that both signed up for the program and showed up for their training and subtracts the mean baseline earnings of those that signed up for the program but failed to show up for their training, c) takes the difference between these two estimates, and d) scales this difference by the proportion that showed up among those that signed up. As above, because we have not used covariates at this point, this estimator can be written as a simplified version of (6):

$$\tau_{att,g_1}^{fd-did} = P(m_1|a_1, g_1) \left[\{E[Y|a_1, m_1, x, g_1] - E[Y|a_1, m_0, x, g_1]\} \right. \\ \left. - \{E[Y|a_1, m_1, x, g_2] - E[Y|a_1, m_0, x, g_2]\} \right],$$

⁶Those who failed to show up for training (dropouts) refers to those that signed up and were accepted to the program, but failed to formally enroll in the program (42.7% of adult males).

where a_1 indicates signing up for the program, m_1 indicates showing up for the program, m_0 indicates failing to show up for the program, g_1 indicates a post-treatment measurement (at 18 months), g_2 indicates a baseline measurement (at 0 months), and Y now can denote either 18-month or 0-month earnings, depending on whether g_1 or g_2 is in the conditioning set.

The front-door and front-door DD estimates for the effect of the JTPA program on adult males are presented in Figure 2. The experimental benchmark (solid black line) is the only estimate that uses the experimental control units. While the front-door estimator appears to exhibit some of the anticipated positive bias, the estimate lies within the 95% confidence interval from the experiment. The front-door DD estimator gets a bit closer to the experimental benchmark and its 95% interval more clearly covers the benchmark.

[Figure 2 about here.]

Although the improvement from the front-door DD estimate is minimal here, this may be due to the relatively good quality of the front-door estimate. If we didn't see the experimental results (as would be true for non-illustrative applications), the similarity between the front-door and front-door DD estimates would give us some confidence as to the robustness of the findings (and this confidence would not be misplaced for this example). However, if even after seeing these results we prefer a more conservative estimate of the effect of sign-up, we can define a different differencing group using the observed covariates.

Results: Single Males as a Differencing Group

If we didn't have the experimental benchmark, we might not be confident that the bias in the pre-treatment period is equal to the bias in the post-treatment period, and hence we may want to use an additional differencing group as a robustness strategy. In this subsection, we discuss the use of never married men (henceforth referred to simply as single men) as the differencing group and currently or once married adult men as the group of interest (henceforth referred to simply as married men).

The use of a differencing group that is a subset of the individuals (single men as subset of adult

men) adds an additional complication: we must consider whether the effect of interest is the average effect of the program for all individuals or just the average over the individuals in the group of interest. Fortunately, conversion between the two effects is straightforward due to the assumption that the effect of the program is zero for the differencing group. Specifically, the average effect over all individuals is the average effect for the group of interest times the proportion of individuals in the group of interest. In order to simplify the presentation and because this conversion is straightforward, we continue this section focusing on the effect for the group of interest. All of the following results are substantively replicated when we convert to the analysis for all individuals.

With single men as the differencing group, we include baseline earnings as a covariate, which rules out the use of the simplified versions of (2) and (6) from the previous subsection. However, the use of covariates in the analysis also allows us to compare the performance of the front-door and front-door DD estimators to standard covariate adjustments like regression and matching.

The front-door and front-door DD estimates for the effect of the JTPA program on married males are presented in Figure 3 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 eligible non-participants or 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 obtain similar results when using more flexible methods that relax these parametric assumptions. The experimental benchmark (dashed line) is the only estimate that uses experimental control units.

First, note that the front-door estimates exhibit uniformly less estimation error than estimates from standard covariate adjustments across all conditioning sets in Figure 3. The error in the standard estimates for the null conditioning set and conditioning sets that are combinations of age, race, and site are negative. The error becomes positive when we include baseline earnings in the condi-

tioning set. In sharp contrast, the stability of front-door estimates is remarkable. Thus front-door estimates are preferable to standard covariate adjustment when more detailed information on labor force participation and historic earnings is not available.

[Figure 3 about here.]

In spite of the superior performance of front-door estimates compared to standard covariate adjustment, the front-door estimates are slightly above the experimental benchmark across all covariate sets. As mentioned above, without seeing the experimental benchmark, we might believe these estimates are 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 DD 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, and that Assumptions 1 and 5 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 a positive effect of the JTPA program for single males, albeit one smaller than for married males (Korenman and Neumark, 1991). Hence, the front-door DD estimator will likely give us a lower bound for the effect of the JTPA program for married males. In fact, in many applications we may be unable to find a differencing group with no effect, yet still be able to use front-door and front-door DD approaches to bound the causal effect of interest given our beliefs about the sign and relative scale of effects in the group of interest and the differencing group. Figure 3 shows exactly this result. The front-door DD estimator forms a lower bound for across all conditioning sets and the front-door provides an upper bound.

Illustrative Application: Early Voting

In this section, we present front-door DD estimates for the average treatment effect of an early in-person voting program in Florida. We want to evaluate the impact that the presence of early voting had on 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 DD 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. Section 4 of the SI 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 (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. In other words, we are assessing what would have happened to these previous EIP voters in 2008 if the EIP program had not been available in 2008. 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.⁷ 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 2008 EIP effect for 2006 EIP voters, for example, we use 2006 absentee and election day voters as our differencing groups. The existence of an EIP program in 2008 might have induced some 2006 absentee ballot users to change their mode of voting in 2008 (e.g., from absentee to EIP), but it is unlikely to have caused them to vote. This is because 2006 absentee voters who voted EIP in 2008, would likely have just voted absentee in 2008 if the EIP program did not exist in 2012. For example, experimental evidence suggests that while mobilizing people to vote early increases turnout, it does not significantly alter the proportion of people that vote by mail and slightly reduces the proportion voting on election

⁷Following [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.

day (Mann and Mayhew, 2012). It thus seems reasonable to assume that EIP offers alternative, not additional, opportunities for voting to past absentee and election day voters. In this case, any apparent effects on turnout estimated for these groups will be primarily due to bias, which can then be removed from the estimates for the group of interest. If in fact these apparent effects represent real effects for these groups, then our results will produce a lower bound. 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. 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 using cross-state comparisons that find null results for the effects of early voting on turnout (Gronke, Galanes-Rosenbaum, Miller and Toffey, 2008) and (Gronke, Galanes-Rosenbaum and Miller, 2007; Fitzgerald, 2005; Primo, Jacobmeier and Milyo, 2007; Wolfinger, Highton and Mullin, 2005), and Burden, Canon, Mayer and Moynihan (2014) finds a surprising negative effect of early voting on turnout in 2008.⁸ However, identification of turnout effects from observational data using traditional statistical approaches such as regression or matching rely on the absence of unobserved 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. See Keele and Minozzi (2013) for a discussion within the context of election laws and turnout. Additionally, a reduction in Florida's early voting program between 2008 and 2012 provided evidence that early voting may encourage voter turnout (Herron and Smith, 2014).

The front-door estimators presented here provide an alternative approach to estimating turnout

⁸Burden et al. (2014) examine a broader definition of early voting that includes no excuse absentee voting.

effects 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, the exclusion restriction is 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 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 5.

Results

In order to construct the front-door estimate of the 2008 EIP effect for our group of interest, we calculate the turnout rate in 2008 for all individuals who voted early in 2006. We also calculate the non-complier turnout rate in 2008 by excluding all individuals who voted early in 2008 from the previous calculation. The front-door estimate of the 2008 EIP effect for 2006 early voters is thus the difference between the former and latter turnout rates. Quite intuitively, the counterfactual turnout rate without EIP for the group of interest is the observed turnout rate of non-compliers in that group. We do not devote much attention to the front-door estimates seeing as they are implausibly large.⁹ The positive bias stems from the fact that 2006 EIP voters would be more likely to vote in 2008, even in the absence of EIP, than the 2006 non-EIP group (this group includes individuals that did not vote in 2006). 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 present front-door DD estimates for the 2008 EIP program in Figure 4. The estimates all utilize county fixed effects and are calculated separately across the racial categories.¹⁰ The front-door difference-in-differences estimates for the group of interest (2006 EIP voters) are 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 2006 early voters and the front-door estimates for 2006 absentee voters, with the front-door estimates for the differencing group scaled by the ratio of early voter compliance to

⁹Front-door estimates are available in Table 4 of the SI.

¹⁰We calculate front-door DD 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. Due to their small size, these counties are unlikely to exert any significant impact upon the estimates regardless.

absentee voter compliance as shown in (5). The purple estimates in Figure 4 represent the placebo test, with 2006 election day voters standing in as the group of interest and the absentee voters as the differencing group. In general, we note that if there exists more than one plausible differencing group, then one should conduct the analysis using each differencing group separately, as well as a placebo test to verify the plausibility of Assumption 5.

[Figure 4 about here.]

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, we generally 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.

Our methodology uses voting behavior in 2006 only to define groups and does not compare turnout of voters across elections. Thus any differences between presidential election and midterm election voters (see e.g. [Gronke and Toffey \(2008\)](#)) does not pose a prima facie problem for the analysis. However, robustness checks using the 2012 election are presented in the Section 5 of the SI.

Conclusion

In this paper, we have developed front-door DD estimators for nonrandomized program evaluations with one-sided noncompliance and an exclusion restriction. These estimators allow for asymptotically unbiased estimation, even when front-door estimators are biased. Additionally, even when

the front-door DD assumptions do not hold exactly, these estimators sometimes allow for informative bounds.

We illustrated front-door DD 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 DD could be used to recover 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 at least 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.

More broadly, this approach is most likely to be helpful in one of three scenarios. First, with longitudinal/panel data, the over-time front-door DD approach can be used (as in the JTPA study). Second, for some cross-sectional applications a differencing group will be apparent. The absentee voters provide one example, while others may derive from eligibility cut-offs. Third, even if a perfect differencing group is not available, a bound may be possible if we are willing to make assumptions about heterogeneity of effects between the group of interest and differencing group. We showed an example of this with the married versus single analysis of the JTPA study, and there are likely to be a number of applications where this is possible (such as when heterogeneity has been studied in prior randomized studies). If we also have beliefs about the direction of bias for the front-door approach then we can use the front-door and front-door DD in a bracketing approach.

Finally, the results in this paper have implications for research design and analysis. The bracketing of the experimental benchmark in the JTPA application shows 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

- Abadie, Alberto. 2005. "Semiparametric Difference-in-Differences Estimators." *Review of Economic Studies* 72:1–19.
- 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.
- Cox, David R. and Nanny Wermuth. 1995. "Discussion of 'Causal diagrams for empirical research.'" *Biometrika* 82:688–689.
- 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.
- Glynn, Adam and Konstantin Kashin. 2013. "Front-door Versus Back-door Adjustment with Unmeasured Confounding: Bias Formulas for Front-door and Hybrid Adjustments." 71st Annual Conference of the Midwest Political Science Association.
- Gronke, Paul and Charles Stewart. 2013. "Early Voting in Florida." Paper presented at the Annual Meeting of the Midwest Political Science Association, Chicago, IL.
- Gronke, Paul and Daniel Krantz Toffey. 2008. "The Psychological and Institutional Determinants of Early Voting." *Journal of Social Issues* 64(3):503–524.
- Gronke, Paul, Eva Galanes-Rosenbaum, Peter A. Miller and Daniel Toffey. 2008. "Convenience Voting." *Annual Review of Political Science* 11:437–455.
- Gronke, Paul, Eva Galanes-Rosenbaum and Peter Miller. 2007. "Early Voting and Turnout." *PS: Political Science and Politics* XL.
- Hanmer, Michael J. 2009. *Discount Voting: Voter Registration Reforms and Their Effects*. Cambridge University Press.
- Heckman, James, Hidehiko Ichimura, Jeffrey Smith and Petra Todd. 1998. "Characterizing selection bias using experimental data." *Econometrica* 66:1017–1098.
- Heckman, James, Hidehiko Ichimura and Petra Todd. 1997. "Matching as an econometric evaluation estimator evidence from evaluating a job training program." *Review of Economic Studies* 64:605–654.
- 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* .

- 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.
- Herron, Michael C. and Daniel A. Smith. 2014. "Race, Party, and the Consequences of Restricting Early Voting in Florida in the 2012 General Election." *Political Research Quarterly* .
- Imbens, Guido and Donald Rubin. 1995. "Discussion of 'Causal diagrams for empirical research.'" *Biometrika* 82:694–695.
- Keele, Luke and William Minozzi. 2013. "How Much Is Minnesota Like Wisconsin? Assumptions and Counterfactuals in Causal Inference with Observational Data." *Political Analysis* 21(2):193–216.
- Korenman, Sanders and David Neumark. 1991. "Does Marriage Really Make Men More Productive?" *The Journal of Human Resources* 26(2):282–307.
- Kuroki, Manabu and Masami Miyakawa. 1999. "Identifiability Criteria for Causal Effects of Joint Interventions." *J. Japan Statist. Soc.* 29(2):105–117.
- Mann, Christopher B. and Genevieve Mayhew. 2012. "Multiple Voting Methods, Multiple Mobilization Opportunities? Voting Behavior, Institutional Reform, and Mobilization Strategy." Paper presented at the Annual Meeting of the Southern Political Science Association, New Orleans, LA.
- McDonald, Michael P. and Samuel L. Popkin. 2001. "The Myth of the Vanishing Voter." *American Political Science Review* 95:963–974.
- Pearl, Judea. 1995. "Causal diagrams for empirical research." *Biometrika* 82:669–710.
- Pearl, Judea. 2009. *Causality: Models, Reasoning, and Inference*. 2 ed. Cambridge University Press.
- 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.
- Shpitser, Ilya and Judea Pearl. 2006. "Identification of Conditional Interventional Distributions." Proceedings of the Twenty Second Conference on Uncertainty in Artificial Intelligence (UAI).
- Tian, Jin and Judea Pearl. 2002a. 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.
- Tian, Jin and Judea Pearl. 2002b. On the identification of causal effects. In *Proceedings of the American Association of Artificial Intelligence*.
- VanderWeele, Tyler J. 2009. "On the relative nature of overadjustment and unnecessary adjustment." *Epidemiology* 20(4):496–499.
- 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.

Figure 1: Front-door Directed Acyclic Graph (DAG). A represents the treatment/action variable, M represents a set of mediating variables, Y represents the outcome, X represents covariates, and U and V represent sets of unobserved variables. To simplify presentation, we have assumed that X , U , and V are independent (this is implied by the lack of arrows between them), but this is not required. Solid arrows are allowed for the front-door criterion to hold within this group. Dashed arrows are not allowed for the front-door criterion to hold in this group.




Figure 2: Comparison of front-door and over-time front-door difference-in-differences estimates for the JTPA effect for adult males. The solid line is the experimental benchmark and the dashed lines represent the confidence interval. All intervals are 95% bootstrapped confidence intervals based on 10,000 replicates.




Figure 3: Comparison of standard covariate adjusted estimates, front-door, and front-door difference-in-differences estimates for the JTPA effect for married adult males. Solid lines represent the experimental benchmark. 95% bootstrapped confidence intervals are based on 10,000 replicates.




Figure 4: Front-door difference-in-differences 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.

