

ORIGINAL ARTICLE

Estimating controlled direct effects through marginal structural models

Michelle Torres*

Department of Political Science, Rice University, 6100 Main Street, MS-24, Houston, TX 77005, USA

*Corresponding author. Email: smtorres@rice.edu

(Received 8 December 2017; revised 30 June 2018; accepted 7 September 2018; first published online 13 February 2020)

Abstract

When working with panel data, many researchers wish to estimate the direct effects of time-varying factors on future outcomes. However, when a baseline treatment affects both the confounders of further stages of the treatment and the outcome, the estimation of controlled direct effects (CDEs) using traditional regression methods faces a bias trade-off between confounding bias and post-treatment control. Drawing on research from the field of epidemiology, in this article I present a marginal structural modeling (MSM) approach that allows scholars to generate unbiased estimates of CDEs. Further, I detail the characteristics and implementation of MSMs, compare the performance of this approach under different conditions, and discuss and assess practical challenges when conducting them. After presenting the method, I apply MSMs to estimate the effect of wealth in childhood on political participation, highlighting the improvement in terms of bias relative to traditional regression models. The analysis shows that MSMs improve our understanding of causal mechanisms especially when dealing with multi-categorical time-varying treatments and non-continuous outcomes.

Keywords: Models for panel data; causal inference

In recent years, considerable progress has been made in providing methodological tools that allow political scientists to better estimate the causal effects of treatments on outcomes. However, in many cases, we are interested not in identifying the effect of a variable at one period, but rather in assessing effects in a dynamic setting. We might, for instance, observe units in multiple time periods and wish to estimate the independent effect of treatments at each stage on some future outcome. In estimating these effects, researchers can better understand not only *how* and *why* political phenomena are linked, but also the potential consequences of changing a treatment of interest that varies through time. Yet, standard tools in the literature are often ill suited for making valid causal claims in dynamic settings.

To provide some clarity to this discussion, consider the following example, which is depicted visually in [Figure 1](#). Past research on political participation identifies wealth as a key factor that influences citizens' political participation (Verba *et al.*, 1978; Almond and Verba, 1989). Typically, scholars emphasize the effect of wealth in adulthood (measured through self-reported income of adult respondents) as a provider of resources that ease participation (e.g., a car that helps a citizen to reach a polling station). However, a separate question is how wealth in early stages of life (measured through the income of a citizen's parents) can affect political participation *independent of the effect of wealth in adulthood*. Finding such an effect would suggest that, for instance, children being raised in wealthier homes receive a lifelong boost in terms of socialization, cognitive skills, or psychological orientations toward politics that affect participation regardless of their own economic success later in life (Beck and Jennings, 1982; Brady *et al.*, 1995; Currie, 2008). [Figure 1](#) shows a simplified version of this example and highlights this unmediated effect through the bolded red arrows (path *b* and path *a-d*).

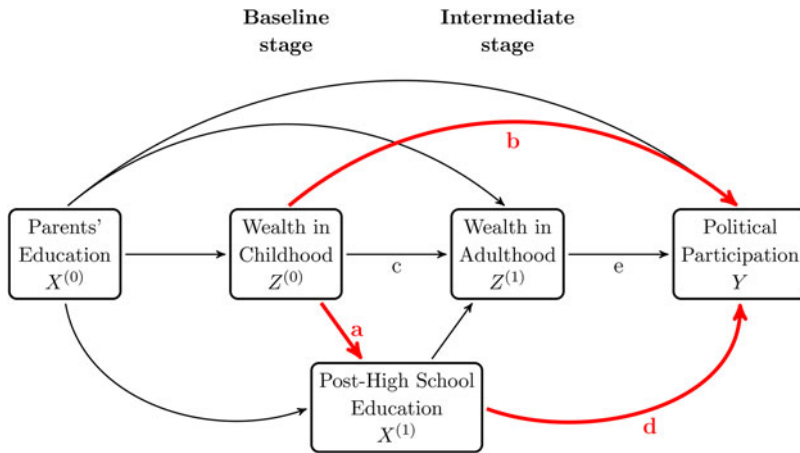


Figure 1. DAG showing the relationship between a time-varying treatment (wealth) and outcome (political participation). *Note:* The bold paths represent the controlled direct effect of the baseline treatment (wealth in childhood) on the outcome (political participation).

Although dynamic treatments in panel datasets abound in the political science literature, applied scholars have thus far been given little guidance as to how to proceed in such settings. Indeed, in situations such as the one depicted in Figure 1 traditional regression techniques offer no way to consistently estimate causal effects. Specifically, researchers wishing to estimate the effect of treatments at various time points in dynamic processes using traditional regression confront a bias trade-off between *confounding bias* and *post-treatment bias*. Confounding bias, sometimes termed omitted variable bias, results from failure to control for important common causes of treatment and outcome when estimating causal effects—a confounder. Post-treatment bias arises from controlling for an intermediate variable that has been affected by the treatment—a post-treatment variable.

The key to understanding this trade-off is that in order to estimate the unmediated effect of a treatment at the baseline stage (e.g., wealth in childhood/parents' income), one must simultaneously correctly estimate the effect of the treatment at intermediate stages (e.g., wealth/income in adulthood). Yet, to estimate the effect at the intermediate stage (path *e* in Figure 1), researchers must either fail to account for important confounders (e.g., attending college) or include such confounders as “control” variables. The former approach will induce confounding bias into our estimates by failing to control for a variable that is causally prior to both income in adulthood and political participation. The latter will introduce post-treatment bias by controlling for a variable that is itself affected by parents' income. In such cases, both controlling for intermediate confounders and failing to control for them will result in biased estimates.

While this problem has certainly not gone unrecognized in the broader statistics literature, these issues have received relatively little attention in political science. Some seemingly plausible approaches, such as mediation analysis (Imai *et al.*, 2010, 2011), are unsuited to handling the dynamic relationship between treatment and confounding variables. Other approaches are more difficult to implement and offer little flexibility. Structural nested mean models, as presented by Acharya *et al.* (2016), for instance, are only suitable when the outcomes are continuous.

In this article, I draw on research from the field of epidemiology (Hernán *et al.*, 2000; Robins *et al.*, 2000), to outline a marginal structural models (MSMs) framework for estimating controlled direct effects (CDEs) of multi-valued treatments at different time periods that is both easy to implement and suitable for use with several data types. This class of models was introduced to political science by Blackwell (2013) and later discussed by Imai and Ratkovic (2014). Still, their studies mainly focus on the estimation of *cumulative effects* of dynamic treatments, and offer little discussion on the applicability of these models to the estimation of CDEs. Further, I extend previous work

by testing and addressing practical challenges of this method, such as the tools for weight estimation, implications and use of weights, and consequences of the violation of the main assumptions.

MSMs overcome the bias trade-off dilemma described above by using an inverse probability of treatment weighted (IPTW) estimator. This allows researchers to account for confounders while avoiding directly controlling for post-treatment variables (Robins *et al.*, 2000; Blackwell, 2013; Blackwell and Glynn, 2014). By estimating correct weights, researchers are able to create pseudo-samples that are balanced with respect to confounders and therefore allow for consistent estimation of causal quantities of interest. Importantly, unlike previous methods in the political science literature, I present and detail the implementation of MSMs to estimate CDEs with multi-categorical treatments as well as non-continuous outcomes.

In the next section, I define CDEs and discuss the challenges that researchers face when estimating them in a dynamic setting. I then provide an overview of MSMs—and the assumptions that undergird them—and explain how they allow for unbiased estimates of treatment effects. As part of this presentation, I not only detail important elements of the implementation of these models but also provide guidance for the weighting process, and for the cases in which assumptions are not fully fulfilled. The section also includes a discussion of the advantages of MSMs over other alternatives, especially traditional regression models. Finally, I present an application that compares the inferences reached by MSMs and traditional models regarding the causes of political participation. This application focuses on the estimation of the CDE of wealth in childhood, as measured by parents' income, on political activism using a panel survey that spans over 30 years.

1. Controlled direct effects and bias trade-off

To formally articulate the difficulty of estimating causal effects in a dynamic setting, I return to the example depicted in Figure 1. We are interested in calculating the effect of wealth in childhood on political participation that is not mediated by wealth in adulthood. The effect of economic resources on political participation has been widely studied. However, recent studies have recognized and focused on the cumulative and long-term effects that economic conditions in childhood may have on participation in later stages of life. For example, Ojeda (2018) finds that it is possible to identify two participations gaps with different sizes and implications: one that childhood economic history generates, and another caused by income in adulthood.

For the illustration and application presented below, I measure wealth of an individual using her own income, and her parents' income.¹ In Figure 1, the unmediated effect of the latter is represented by the highlighted paths (*a-d* and *b*). Substantively, this will allow us to explore the impact of early economic conditions on adult political participation independent of the level of affluence later in life.² In other words, if we could fix (or control) respondents' income in adulthood to a specific level, what would be the effect of changes in parents' income on adult political activity? This quantity is known as the CDE which I define formally below (Pearl, 2001; VanderWeele, 2009; Pearl, 2011). This estimand is useful to (1) understand the mechanism through which treatments affect the outcome, and (2) explore the different effects that treatment regimes have on an outcome. The estimation of CDEs is relevant to address several social science questions: the analysis of the effects of historical institutions on current economic and political conditions (e.g., “zoning” on political participation), the study of issues related to public policy (e.g., the impact of welfare programs on economic development), or the exploration of early conditions of citizens on their current political attitudes and behavior.

¹ *Wealth* is the treatment of interest in two different stages: childhood and adulthood. While “wealth” can imply multiple factors, there is a strong correlation between wealth indices and income (Córdova, 2009).

² It is important to highlight that while this case considers a “sequence” of conceptually similar treatments, researchers can also use this method for sequences of semantically different variables as long as they hold a clear causal relationship (i.e., one precedes and affects the other).

1.1. Defining controlled direct effects

Our goal is to estimate the causal effect of a treatment Z (income) at different “stages” in time. Although the model can be easily extended to allow for multiple stages, I focus on only two stages of treatment: parents’ income ($t=0$) and income in adulthood ($t=1$). For this discussion I assume that the measurement of wealth of an individual i , income,³ in both stages can be either low ($Z_i^{(t)} = 0$), middle ($Z_i^{(t)} = 1$), or high ($Z_i^{(t)} = 2$).⁴ Finally, I assume that the education of each subject’s parents and their level of post-High School education are the sole confounders, which means that these variables are affecting both treatments and the outcome.

The outcome of interest is an individual’s level of political participation, denoted Y . Let $Y_{Z^{(0)}=a}$ be the subject’s level of political participation if parents’ income ($Z^{(0)}$) is set to a value a . Thus, $Y_{Z^{(0)}=0}$ represents the outcome when the respondent’s parents have a low income, whereas $Y_{Z^{(0)}=1}$ and $Y_{Z^{(0)}=2}$ represent the response of the same respondent if her parents’ income was medium and high, respectively. Since only one of the possible values will be observed for each individual, then two of the values of Y are *potential* outcomes while the other is the *observed* outcome. Similarly, the intermediate treatment stage $Z^{(1)}$, income in adulthood, can also take on three values. Therefore, let $Y_{Z^{(0)}=a, Z^{(1)}=b}$ denote the level of political participation of a subject if her parents’ income and income in adulthood were set to values of a and b , respectively.

With this notation, we define the CDE by “fixing” the second-stage of treatment to a specific value (Pearl, 2001; VanderWeele, 2009). It is important to highlight that this “fixing” assumes that the researcher has the capacity to artificially manipulate the intermediate stages. In practice, the estimation of the CDE is especially useful for policy design and experimental settings where researchers have the chance of manipulating the treatment stages. For example, Akee *et al.* (2018) manipulates the assignment of unconditional money transfers at different stages of life to study the effect that income has on civic participation. Although this option is not easily available for social scientists, and especially for those dealing with observational data, this quantity is still useful to have a better understanding of the potential outcomes that different treatment combinations generate. CDEs aid with the operationalization and analysis of the core concept of causal inference: the definition and modeling of *counterfactuals*. Therefore, the value of such estimand should not be underrated, even in the cases where the manipulation of any of the treatment stages is not possible.

We formally define the CDE as:

$$CDE = Y_{Z^{(0)}=a, Z^{(1)}=b} - Y_{Z^{(0)}=a', Z^{(1)}=b}. \tag{1}$$

Conceptually, the CDE estimand represents the effect of a treatment at a specific time period while controlling the level of treatment at different stages. In this example, we are interested in the CDE for the baseline treatment ($t=0$). Of course, we cannot directly calculate the CDE since the counterfactual values are not observed. However, with standard regularity assumptions, we can provide an unbiased estimate of the CDE by calculating the *average controlled direct effect* (ACDE)

$$\begin{aligned} ACDE &= \mathbb{E}(Y_{Z^{(0)}=a, Z^{(1)}=b} - Y_{Z^{(0)}=a', Z^{(1)}=b}) \\ &= \mathbb{E}(Y_{Z^{(0)}=a, Z^{(1)}=b}) - \mathbb{E}(Y_{Z^{(0)}=a', Z^{(1)}=b}), \end{aligned}$$

³The categorization of a continuous variable such as income is a common practice in multiple fields. One important reason is measurement. In order to decrease non-response and increase perceptions of privacy, respondents generally choose their income from multiple categorical options defined by the researcher. Also, conceptually, researchers are generally interested in the differences between levels of income rather than in its unitary nature (Moore and Welniak, 2000; Córdova, 2009).

⁴Under the assumption that $Z_i^{(t)}$, $X_i^{(t)}$ and Y_i are sampled iid from a population, I treat them as random and therefore avoid the subindex i .

where $\mathbb{E}(\cdot)$ refers to the expectation over the individuals in the sample. This is simply the difference between the average outcomes for units that received different treatments (a and a') at stage $t = 0$ while holding the second stage constant at b .

1.2. The bias trade-off

Although in theory the ACDE seems relatively straightforward to calculate, in practice it is not. In fact, there is actually no way to correctly estimate the ACDE using standard regression techniques. The dilemma is the following: since we want to estimate the effect of the treatment at each stage of the treatment sequence separately, we must estimate a coefficient representing the effect of parents' income and another one for the effect of income in adulthood. In order to generate unbiased estimates we must control for all confounders—the set of variables that affect both the treatment and the outcome—in order to avoid *confounding bias*. However, some of the confounding variables for the intermediate-level treatments are themselves affected by the baseline treatment. Therefore, controlling for these covariates will introduce post-treatment bias into our estimates (Rubin, 1977; Rosenbaum, 1984; Elwert and Winship, 2014; Montgomery *et al.*, 2018). As a consequence, we have a situation where *both* controlling for and *not* controlling for confounders will result in biased estimates of the ACDE.

To make this trade-off clearer, I return to the example depicted in Figure 1. In this instance, the problematic variable is *post-High School education*. Why is it necessary to include this variable in the model in the first place? The answer is that the assignment of the treatment in observational studies is not random. In this example, both having a high levels of wealth and political participation are dependent on other factors such as levels of educational attainment. The implication of non-random assignment to treatment is that the observed differences in the outcomes between treated and untreated groups cannot only be attributed to the presence of the treatment but potentially also to inherent differences between the two groups. Therefore, once we identify all confounders, a necessary step is to account for this imbalance. In a standard regression, this would be done by including education as a control variable.

However, including post-High School education as a control variable results in a different problem: post-treatment bias. In our example, whether or not respondents seek post-secondary education is itself caused (in part) by the baseline treatment (wealth in childhood). In the language of causal inference, education is therefore a “collider” (Elwert and Winship, 2014), and controlling for it in a regression will bias estimates of causal effects.

In summary, when confounders are affected by a baseline treatment we face an inevitable bias trade-off: excluding problematic confounders leads to omitted variable bias, but including them leads to post-treatment control bias. Although not always recognized, this trade-off and its consequences are frequently encountered in political science research. If we are dealing with panel or longitudinal data, then it is natural to identify treatments varying through time and complex interactions between those treatment stages and confounders that are not static. In the next section, I explain how adopting a marginal structural modeling framework allows us to address confounding bias *without* introducing post-treatment bias.

2. Estimating CDE using marginal structural models

MSMs are a class of models used to estimate the causal effect of time-varying treatments such as medicine prescription or medical procedure histories (Robins, 1999; Robins *et al.*, 2000; Hernán *et al.*, 2001). Classic applications have focused on estimating the cumulative effects of these time-varying treatments on future outcomes, and previous applications of MSM in political science focused on estimating these cumulative effects (Blackwell, 2013). My presentation below builds on more recent work by researchers who have extended the MSM framework to also estimate CDEs and, under certain conditions, natural direct and indirect effects (VanderWeele, 2009;

Nandi *et al.*, 2012). I cover and detail cases where the treatment is multi-valued and the outcome is non-continuous, to address questions relevant to political science using panel survey data.

In general, MSMs are useful when dealing with cases where (1) the treatment takes few values, (2) there exists a covariate that acts both as determinant of the outcome of interest and as a predictor of an intermediate stage of treatment, and (3) past exposure to baseline treatment predicts subsequent levels of this covariate. As I reviewed above, the decision to control or not control for these covariates inevitably leads to either confounding or post-treatment bias. However, through an IPTW estimator, MSMs provide unbiased estimates once we meet certain assumptions. The core idea of these models is that through the weights estimated via IPTW, we create a “pseudo-population” consisting of copies of each subject in the sample. This pseudo-population has two important features: first, the probability of receiving the second stage of the treatment is unconditional on the confounders affected by the baseline treatment eliminating the necessity of controlling for them in the final model. Second, the potential outcomes are the same as in the true population allowing the estimation of unbiased causal effects (Robins *et al.*, 2000). The number of replicas in the pseudo-sample is calculated based on the probability of observing a particular sequence of treatment conditional on relevant confounders (Robins, 1999).

Before providing the details of the method, it is important to note that several previous scholars have applied models closely related to MSMs in political science. Perhaps the earliest example is Glynn and Quinn (2010), who introduced and extended the IPTW approach for estimating causal effects in a cross-sectional setting. After Blackwell (2013) formally introduced MSMs to political science, Imai and Ratkovic (2015) generalized the covariate balancing propensity score (CBPS) to dynamic settings to achieve a more balanced pseudo-population. More recently, Samii *et al.* (2017) applied the IPTW framework for estimating causal effects using a machine learning approach for assigning treatment weights. However, the method most closely related to the objective described here, the estimation of CDE, are structural nested mean models (SNMMs), which were recently introduced to political science by Acharya *et al.* (2016). I provide a brief discussion comparing and contrasting the MSM and SNMM approaches for estimating CDEs in the Appendix.

2.1. Assumptions

Going back to our example, MSMs allow us to model levels of political activism of individuals receiving each of the potential *Parents’ income–Income in adulthood* sequences: low–low, low–middle, low–high, middle–low, middle–middle, middle–high, high–low, high–middle, and high–high. However, modeling these unconditional (or marginal) distributions requires the fulfillment of two assumptions.

The first is the *sequential ignorability condition*, which guarantees the necessary statistical exogeneity for the identification of causal effects (Robins, 1999).⁵ In essence, this assumption is an extension of a general condition for the estimation of causal effects in single-stage settings: controlling for confounders $X^{(t)}$ assures independence ($\perp\!\!\!\perp$) of the potential outcomes $Y_{Z^{(0)}, Z^{(1)}, \dots, Z^{(T)}}$ from the treatment $Z^{(t)}$. For the multi-stage setting, we need to meet this same condition for each treatment stage. In our example, this would mean controlling for education of the parents (denoted here as $X^{(0)}$) to avoid confounding of parents’ income—the first treatment stage $Z^{(0)}$. Formally,

$$Y_{Z^{(0)}, Z^{(1)}} \perp\!\!\!\perp Z^{(0)} | X^{(0)}. \quad (3)$$

For the second stage, it is necessary not only to control for education of the parents, $X^{(0)}$, and post-High School education, $X^{(1)}$, to avoid confounding bias, but also to include parents’ income, $Z^{(0)}$, as another confounder of wealth in adulthood, $Z^{(1)}$, and participation, Y . In other words, the

⁵Note that if there are multiple confounders, the values of $X^{(0)}$ and $X^{(1)}$ are going to be in matrix form rather than vectors.

outcome needs to be independent of any stage in the treatment sequence, conditional on past confounders and treatments,⁶

$$Y_{Z^{(0)}, Z^{(1)}} \prod Z^{(1)} | Z^{(0)}, X^{(0)}, X^{(1)}. \tag{4}$$

The second assumption is the *positivity assumption* which states that a treatment value should not be limited to a single level l of the control variables. Intuitively, this means that all subjects in the sample must have a non-zero probability of getting exposure to the different levels of treatment. In our example, the assumption implies that an individual that did not attend college and whose parents had a low income should still have a non-zero chance of receiving a middle or high income as an adult.⁷ Formally,

$$\text{If } \Pr(Z^{(0)} = z^{(0)}, (X^{(0)}, X^{(1)}) = (x^{(0)}, x^{(1)})) > 0, \text{ then} \tag{5}$$

$$\Pr(Z^{(1)} = z^{(1)} | (X^{(0)}, X^{(1)}) = (x^{(0)}, x^{(1)}), Z^{(0)} = z^{(0)}) > 0. \tag{6}$$

Once we meet these assumptions, we can use MSMs to estimate the ACDE.^{8,9}

2.2. Benefits of the pseudo-sample

MSMs aim to model the potential outcomes for the different sequences of treatment. This strategy allows for the estimation of CDEs. For example, consider the following model:

$$\mathbb{E}[Y_{Z^{(0)}, Z^{(1)}}] = \alpha_0 + \alpha_1 Z^{(0)} + \alpha_2 Z^{(1)}. \tag{7}$$

The ACDE in model 7 is the expected value of the differences in Y when $Z^{(0)}$ is 1 and when $Z^{(0)}$ is 0, while fixing $Z^{(1)}$ to b . Then,

$$\begin{aligned} \mathbb{E}[Y_{Z^{(0)}=1, Z^{(1)}=b} - Y_{Z^{(0)}=0, Z^{(1)}=b}] &= \alpha_0 + \alpha_1 \cdot 1 + \alpha_2 \cdot b \\ &\quad - (\alpha_0 + \alpha_1 \cdot 0 + \alpha_2 \cdot b) \\ &= \alpha_1(1 - 0) = \alpha_1. \end{aligned} \tag{8}$$

In other words, when the second treatment stage is set to b , the baseline stage has a causal effect of α_1 on the outcome. This estimation only holds if the differences we observe in Y are only related to the treatment and not to other confounders. From previous sections, we know that in our example, as in all observational studies, this is not true. Wealth in each of the two stages is not randomized: the levels of this “treatment” are not independent from past economic

⁶We can define this assumption more generally as $Y \rightarrow \prod Z^{(t)} | Z^{(t-1)}, X^{(t)}$. Where \rightarrow indicates the treatment or covariate regime up to the time indicated in parentheses. $z^{(t)}$

⁷More generally, if $\Pr(Z^{(t-1)} = z^{(t-1)}, X^{(t)} = x^{(t)}) > 0$, then $\Pr(Z^{(t)} = z^{(t)} | X^{(t)} = x^{(t)}, Z^{(t-1)} = z^{(t-1)}) > 0$.

⁸On the one hand, the fulfillment of the first condition can be difficult given that there is no technique that allows us to diagnose the degree to which it is met. However, this is a classic (and necessary) assumption in any causal analysis. *Naïve* regression estimators are not exempt from meeting the ignorability assumption either. Furthermore, previous work by Blackwell (2013) and VanderWeele (2010) includes the development of sensitivity analyses that allow us to assess the strength of the inferences made from MSMs. On the other hand, the fulfillment of the positivity assumption can be difficult in cases where there is a continuous treatment and confounders, and then alternatives like SNMMs are preferred (VanderWeele, 2009).

⁹The simulations in the Appendix show how the bias and variance of the ACDE change depending on mild to strong violations of these assumptions. See discussion below.

conditions or education. The implication is that each income sequence has different probabilities of being observed given the values of the confounding factors (e.g., a subject with a college degree is more likely to have a higher income than one that only completed High School). MSMs use these probabilities to build weights that balance the sample across treatment groups. The weights are the product of two components, one per treatment stage, defined as follows:

$$\mathcal{W}(t) = W_{Z^{(0)}} \times W_{Z^{(1)}} = \frac{f(Z^{(0)}|X^{(0)})}{f(Z^{(0)})} \times \frac{f(Z^{(1)}|Z^{(0)}, X^{(0)}, X^{(1)})}{f(Z^{(1)}|Z^{(0)})}. \quad (9)$$

The numerator in each of the components of the $\mathcal{W}(t)$ term is the probability that an individual received his own observed treatment at time t , $Z^{(t)}$, given his own past treatment (up to $t-1$) and covariate history up to point t (Robins, 1999). For example, in the income example, the numerator of $W_{Z^{(1)}}$ is simply the probability that an individual has her own observed income in adulthood conditional on her observed parents' income, and educational attainment after High School.¹⁰ At the same time, the denominator is the probability that a subject received her observed treatment at time t but only conditional on her treatment sequence until $t-1$. In the example, the denominator of $W_{Z^{(1)}}$ is the probability of observing the actual income in adulthood but only conditional on parents' income.¹¹

Once we obtain these weights, we estimate the parameters in Equation 8 using a weighted least squares regression in which each subject is given as a weight the inverse of her corresponding $\mathcal{W}(t)$. For illustrative purposes we implement a weighted regression, however the researcher has full flexibility to model the outcome as long as it is applied to the weighted sample. Thus, the model can range from a simple weighted mean to a complex non-parametric weighted model.¹² The weighted model handles confounding while avoiding explicit conditioning and post-treatment bias. How does weighting achieve this? Recall that treatment sequences have different probabilities of being observed given the values of the confounders. By weighting, we are "leveling the field" and breaking the link between the second treatment stage and its confounders: the problematic variables affected by the baseline treatment. To have a more intuitive understanding of this, we can view the pseudo-population as a sample composed of each individual in the original population plus $(\mathcal{W}(t)^{-1} - 1)$ copies of themselves.

Consider the following hypothetical example based on the income case.¹³ The first panel of Figure 2 shows the distribution of subjects in the original sample across the different levels of parents' income and adulthood income as well as education. Each human figure represents 1000 individuals. Each cell represents a potential combination of income in childhood and adulthood: low (\$), middle (\$\$), or high (\$\$\$). Furthermore, the level of post-High School education is indicated by the hat and color of the figures: black symbols wearing a hat attended college while gray figures did not. In this example, we assume that sequential ignorability holds and, as the picture shows, the positivity assumption is met (there is at least one human figure in all possible combinations of parents' income, income in adulthood and education).

Just by visual inspection, it is clear from the figure that the probability of, for example, receiving a high income in adulthood is strongly determined by both levels of parents' income and

¹⁰Note that if it is the beginning of the sequence, $t=0$, then the numerator would only be conditional on the confounders of $Z^{(0)}$ and Y . That is, $f(Z^{(0)}|X^{(0)})$.

¹¹The denominator of this quantity can be replaced with another function of treatment history. This would not affect the consistency or unbiasedness of the estimator. The numerator is introduced as a "stabilizer" of weights in order to avoid extreme values. The efficiency of the estimator can be influenced by the decision for the numerator. However, the selected function should not include the intermediate or confounding variables in the model.

¹²The comparison of weighting methods used to model the outcome escapes the scope of this paper. However, as in any other study, researchers should select the appropriate modeling technique based on a deep understanding of the data and full awareness of the assumptions and trade-offs that the different methods convey.

¹³This example is based on one designed by Robins (1997).

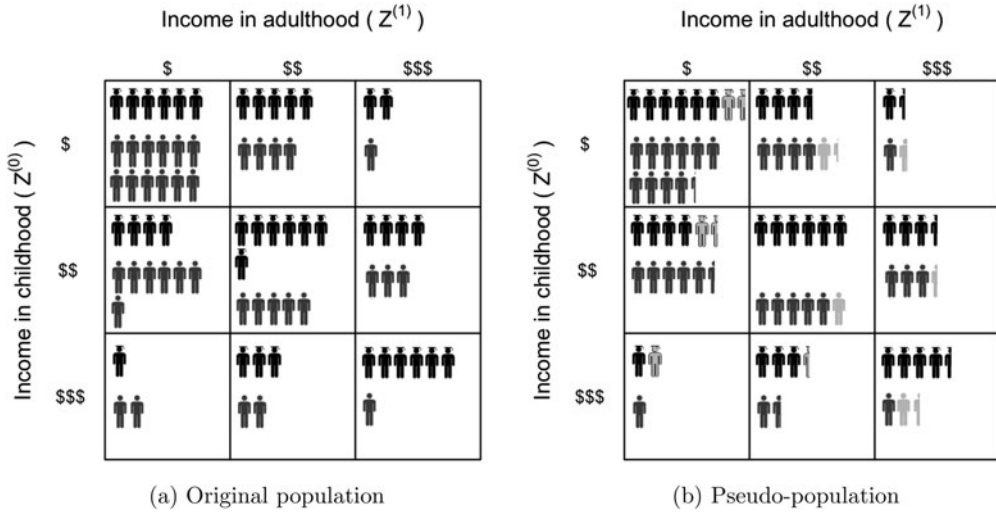


Figure 2. Distribution of individuals based on treatment sequence (parents’ income and income in adulthood) and confounder affected by treatment (post-high school education). *Note:* Each figure represents 1,000 individuals. Black figures with hat indicate that those subjects attended college, while gray figures only completed high school. The panels show the distribution of respondents across treatment conditions.

education. Table 1 presents the information by stratum and actual probabilities of receiving a particular income in adulthood $Z^{(1)}$ given parents’ income and education.¹⁴ For example, the probability of having a high income in adulthood if a subject has a high income in childhood but does not attend college is $1000/5000 = 0.2$ (bold cell over sum of light-shaded cells in column 7 of Table 1). However, the probability of having a high income in adulthood when parents’ income is high but the subject also attends college is much higher, $6000/10,000 = 0.6$. In other words, we have unbalance across levels of educational attainment. From column 3 of Table 2 (labeled as “Original”), where we can see a summary of these probabilities for all strata, we can conclude that income in adulthood is not independent of levels of education, but that it acts as a confounder of this variable and political participation.¹⁵

However, we can eliminate this unbalance by creating a pseudo-population based on copies (or reductions) of the subjects in the original sample using the inverse of the weights $\mathcal{W}(t)$. The column labeled as $\mathcal{W}(t)^{-1}$ in Table 2 presents this quantity and all the information necessary to construct it. Based on this information we can build the pseudo-sample shown in the second panel of Figure 2. We can now repeat the same exercise of calculating the probabilities of having a high income in adulthood for the individuals in the new sample. Column 4 of Table 2 (labeled as “Pseudo”) presents these new estimated probabilities. It is important to highlight that while the calculation of these probabilities involved a simple stratification approach, cases with multiple confounders will require more intensive modeling techniques.¹⁶

Once we weight the sample, the probability of having a high income in adulthood is equal for both levels of education within each parents’ income strata—the second treatment stage is

¹⁴For sake of space, the table with the full set of strata is presented in the Appendix.

¹⁵If we ignore this confounder, the potential differences that we could observe in levels of political participation between groups defined by the different levels of income could not be attributed to the effect of this variable but to the differences in education levels.

¹⁶Given that the unbiasedness of MSMs rely on an accurate estimation of the weights, different models will lead to different estimations of ACDEs. A brief comparison of different modeling tools for the estimation of weights is presented in Section 2.3.

Table 1. Calculation of weights for each stratum in sample

$Z^{(0)}$	$Z^{(1)}$	$X^{(1)}$	$f(Z^{(1)} Z^{(0)})$	$f(Z^{(1)} Z^{(0)}, X^{(1)})$	$\mathcal{W}(t)^{-1}$	Original-pop N	Pseudo-pop N
0	0	0	0.600	0.706	0.850	12,000	10,200
0	0	1	0.600	0.462	1.299	6,000	7,794
0	1	0	0.300	0.235	1.277	4,000	5,108
0	1	1	0.300	0.385	0.779	5,000	3,895
0	2	0	0.100	0.059	1.695	1,000	1,695
0	2	1	0.100	0.154	0.649	2,000	1,298
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
2	0	0	0.200	0.400	0.500	2,000	1,000
2	0	1	0.200	0.100	2.000	1,000	2,000
2	1	0	0.333	0.400	0.832	2,000	1,664
2	1	1	0.333	0.300	1.110	3,000	3,330
2	2	0	0.467	0.200	2.335	1,000	2,335
2	2	1	0.467	0.600	0.778	6,000	4,668

Note: $Z^{(0)}$ is parents' income where 0 is low, 1 is middle, and 2 is high. $Z^{(1)}$ is income in adulthood where 0 is low, 1 is middle, and 2 is high. $X^{(1)}$ is post-High School education where 0 is no college and 1 is college. The full table with the probabilities and weights for the full set of treatment and covariate combinations can be found in the Appendix.

Table 2. Probabilities of having a high income in adulthood

$Z^{(0)}$	$X^{(1)}$	Original $\Pr(Z^{(1)} = 2 X^{(1)}, Z^{(0)})$	Pseudo $\Pr(Z_p^{(1)} = 2 X_p^{(1)}, Z_p^{(0)})$
Low income	No college	0.059	0.100
Low income	College	0.154	0.100
Middle income	No college	0.200	0.233
Middle income	College	0.267	0.233
High income	No college	0.200	0.467
High income	College	0.600	0.467

Note: $Z^{(0)}$ is parents' income, $Z^{(1)}$ is income in adulthood, and $X^{(1)}$ is post-high education.

balanced within parents' income and education groups. For example, a subject who did not attend college and whose parents had a high income has a probability of having a high income in adulthood of $2335/4999 = 0.467$. Similarly, a subject that reports that her parents had a high income but that attended college has a probability of $4668/9998 = 0.467$ of having a high income in adulthood. Thus, in the pseudo-population, the confounder $X^{(1)}$ does not predict the treatment at $t = 1$ given the baseline treatment. Post-High School education is no longer a confounder and we can assess the CDE of early income $Z^{(0)}$ on political participation.

The last step of this process consists of fitting a weighted regression of the outcome variable on both the baseline and intermediate treatments using the vector of weights $\mathcal{W}(t)^{-1}$. Other covariates can be included in this regression but these have to be strictly pre-treatment.¹⁷

2.3. Weighting: methods and implications

2.3.1. Estimation of weights

As I reviewed in the previous section, creating a balanced pseudo-sample involves an accurate estimation of the probabilities of observing the multiple treatment sequences conditional on

¹⁷This decision should be strongly motivated by substantive and theoretical knowledge of the question under analysis, as well as by a deep understanding of the data. This is, if there are pre-treatment confounders, they should be included as part of the weight estimation as in any other model aiming to support causal claims. If fulfilled, the sequential ignorability assumption and subsequent weighting guarantee a pseudo-random assignment of the treatment stages. Therefore, covariate adjustment is not necessary. However, it tends to improve precision and reduce standard errors *if* the covariates are predictive of the outcome (Miratrix *et al.*, 2013).



Figure 3. Mean bias of predicted probabilities by model. *Note:* Each corner shows one of the nine possible treatment sequences. The axes show the difference between true ACDE and the estimated ACDE by a given model: CDE 1 = $P(Y_{Z^{(0)}=1, Z^{(1)}=0}) - P(Y_{Z^{(0)}=0, Z^{(1)}=0})$, CDE 2 = $P(Y_{Z^{(0)}=2, Z^{(1)}=0}) - P(Y_{Z^{(0)}=0, Z^{(1)}=0})$, CDE 3 = $P(Y_{Z^{(0)}=2, Z^{(1)}=0}) - P(Y_{Z^{(0)}=1, Z^{(1)}=0})$, CDE 4 = $P(Y_{Z^{(0)}=1, Z^{(1)}=1}) - P(Y_{Z^{(0)}=0, Z^{(1)}=1})$, CDE 5 = $P(Y_{Z^{(0)}=2, Z^{(1)}=1}) - P(Y_{Z^{(0)}=0, Z^{(1)}=1})$, CDE 6 = $P(Y_{Z^{(0)}=2, Z^{(1)}=1}) - P(Y_{Z^{(0)}=1, Z^{(1)}=1})$, CDE 7 = $P(Y_{Z^{(0)}=1, Z^{(1)}=2}) - P(Y_{Z^{(0)}=0, Z^{(1)}=2})$, CDE 8 = $P(Y_{Z^{(0)}=2, Z^{(1)}=2}) - P(Y_{Z^{(0)}=0, Z^{(1)}=2})$, CDE 9 = $P(Y_{Z^{(0)}=2, Z^{(1)}=2}) - P(Y_{Z^{(0)}=1, Z^{(1)}=2})$.

covariate history. This implies an appropriate model specification of treatment assignment, and a suitable method to estimate probabilities.

In general, the most common alternative to estimate the assignment of a particular multi-category treatment sequence is a generalized linear model for categorical data. The simplicity of the model is its most attractive feature, but the trade-off between parsimony and strong predictive power that could potentially reduce bias has not been fully explored. Therefore, in this section I present a comparison of three different approaches to estimate weights, and an analysis that each of them yields when used in a MSM framework.

The main objective of this exercise is to compare the magnitude of the mean bias of the estimates of the ACDEs that come from four different models: a *naïve*, or saturated model that includes post-treatment covariates, and three MSMs using weights that were obtained using three different methods—an ordered logistic regression (ologit), a generalized additive model (GAM), and a random forest (RF).

I simulate 500 datasets and for each of the four models I record the differences between the estimated ACDEs and the *true* ones to obtain measures of bias.¹⁸ Figure 3 presents the average bias for each model of the nine potential ACDEs. In this figure, each corner of the polygon

¹⁸Details about the specific simulation setting can be found in the Appendix.

represents the mean difference between the true and estimated ACDE of the baseline treatment on the outcome, when fixing the intermediate treatment to a certain level. For example, CDE 1 represents the difference in probabilities of attending a rally between subjects that had a middle income in childhood ($Z^{(0)} = 1$) and others that had low income in childhood ($Z^{(0)} = 0$), when the income in adulthood is fixed to low ($Z^{(1)} = 0$). Further, the colored lines represent each of the four different models: the naïve model, and the three MSMs.

The analysis confirms that all of the MSMs perform significantly better than the saturated model (in blue), and provides useful information about the weighting methods. First, there are no substantive differences between them. The RF shows a slightly better performance than the GAM or the ordered logit, but it does not seem to be a substantive difference. This is due to the simplicity of the example, where both the treatment assignment and outcome model are not complex. However, it is worth noting that these differences might be higher in cases with larger sets of confounders, more complex interactions and relationships between the variables, or where distributional assumptions are harder to meet. For example Montgomery and Olivella (2018) show that regression trees yield better estimates of probability of treatment sequences in cases with multiple confounders, which in turn improve the pseudo-sample balance and the overall performance of the MSM. Second, although the mean bias for all treatment sequences is very close to zero, there are instances where this is not the case. The simulation setting is purposely designed to allow for samples where the positivity assumption is not fulfilled. In this case, although the expected bias is not zero, it (1) is small even under settings where the positivity assumption is violated, and (2) performs better than traditional regressions regardless of the method used for the weights estimation.

2.3.2. Practical considerations about weights

The act of weighting motivates multiple questions with potentially strong implications for the estimation of ACDEs. How should we proceed if the weights have extreme values? How do we account for uncertainty when estimating the weights? What implications does weighting have in terms of variance? What is the correct “modeling approach” when dealing with weighted samples? While all of these are important questions which merit thoughtful answers, their discussion escape the scope of this piece. However, this section aims to serve as a brief guide for researchers interested in the implementation of MSMs and a starting point for further exploration of these topics.

First, it is common to encounter cases where the pseudo-sample is constructed using very extreme weights. This occurs when the treatment and covariate combinations have very few observations. Since the weighting process aims to “level and balance” the different treatment and covariate sequences, then those with few individuals will be compensated with higher weights for its members to “represent” those that we cannot observe. Extreme weights may result in unstable estimators with high variance (Kang and Schafer, 2007). To account for this issue, researchers should consider trimming or truncating the weights, as well as assessing the sensitivity of the estimates to this alternative (mainly by exploring the changes in the distribution of weights). Although these alternatives do not completely eliminate the bias, they help to reduce it (Platt *et al.*, 2012) and also improve the variance of the estimator. Another strategy is to restrict the analysis to cases with moderate weights. While this will not lead to an unbiased estimate of the ATE or ACDE in the *full* sample, it provides information about these effects among the population exposed to the treatment combinations which in practice may be more realistic to observe (Platt *et al.*, 2012).

Second, it is important to consider that a consequence of the use of weights, regardless of the method used to calculate them, is that it induces within-subject correlation (by “duplicating” individuals), and therefore the standard error estimates reported by standard programs may be invalid. To account for this issue users should use bootstrap methods when assessing the reliability of the estimates (Hernán *et al.*, 2000). It is crucial that in order for the weights to remain useful, they must be estimated in each bootstrapped sample. This will not only help to improve the

estimation of standard errors, but will partially ameliorate concerns related to the inclusion of uncertainty in the estimation of weights. Current applications of MSMs do not include information on the uncertainty of the predicted probabilities used for the derivation of the weights. Further studies should address this issue in order to reach better inferences of the object under analysis.

2.4. Advantages and disadvantages of MSMs

MSMs have multiple strengths and advantages but also some weaknesses. First, even though MSMs can theoretically handle any type of outcome and treatment variables, their use is mainly restricted to categorical or binary treatments. The reason is that a large number of values complicates the fulfillment of the positivity assumption. In cases where the treatment is continuous SNMMs should be favored (Acharya *et al.*, 2016).

Further, MSMs estimates are sensitive to misspecification of the treatment assignment model. This is due to the reliability of IPTW on the calculation of probabilities of treatment sequences. However, there are alternatives that help to alleviate and diagnose this issue. As reviewed in the previous section, there are multiple methods that might aid to achieve more accurate weights in the presence of multiple covariates (Watkins *et al.*, 2013). Further, Imai and Ratkovic (2015) generalize the CBPS methodology (Imai and Ratkovic, 2014) to time-varying treatments and confounder settings such that the covariate balance is improved in each stage. In addition, authors like Robins (1997), VanderWeele (2010), Blackwell (2013) have developed and implemented tools to conduct sensitivity analysis in order to assess the robustness of the estimates of ACDEs in multiple scenarios where the sequential ignorability assumption is violated.

Finally, although IPTW estimators remain unbiased even in cases with small samples, the standard errors tend to be larger than in naïve models. Weights induce higher variance and higher standard errors of the estimates under study (see simulations in Appendix). Depending on each particular case and data, researchers should consider the bias-efficiency trade-off when using MSMs (Westreich *et al.*, 2012).

3. MSMs and controlled direct effects in practice: the effect of parents' income on political participation

In this section I present an extension of the example outlined above regarding the CDE of income in youth on political participation that is not mediated by income in adulthood. This application illustrates the differences between MSMs and traditional regression models in terms of inferences, and second, it extends the analysis of CDE to non-binary treatments and non-continuous outcomes.

To illustrate the consequences that confounding and post-treatment bias have on results and inferences, I compare the estimates from MSMs to two naïve models: the *overcontrol* or saturated model and the *undercontrol* model. A common approach is to include all relevant confounders in a regression regardless of whether these are affected by the treatment—the *overcontrol* model. A less common but still plausible practice is to avoid problematic confounders and limit the analysis to the baseline and intermediate treatments—the *undercontrol* model.¹⁹

Wealth is assumed to affect several factors in early stages of life, such as motivation, abilities, skills, and favorable social environments. However, the causal effect of these early economic conditions on political participation that is not mediated by economic status in adulthood is understudied.

I provide evidence that the CDE of parents' income on participation is positive. That is, if we set income in adulthood to a certain level (e.g., by providing subsidies or stipends), there would

¹⁹This model also includes strictly pre-intermediate treatment and outcome confounders.

still be an effect of early economic conditions on political participation. However, the magnitude and reliability of this effect varies depending on the specific type of political activity that a subject pursues. While some activities require actual monetary resources, others are more likely to require skills developed in early stages of life (Verba and Nie, 1972; Verba *et al.*, 1978). For example, Lipset (1959) finds that middle-life practices contribute to the development of democratic political orientations and these, in turn, are associated with engagement in activities such as rallies or protests. However, other activities, such as donating to a campaign, are more likely to be influenced solely by the availability of resources associated with income at the moment of the event (e.g., money, time, transportation means, context).

The data to test these effects comes from the Youth-Parent Socialization Panel Study (Jennings *et al.*, 2005). This is a panel study in which a sample of students and their parents were interviewed for four waves in 1965, 1973, 1982, and 1997. I use the models below to study two different outcomes measured in 1982: attending a political rally, and giving money to a candidate or campaign. For the treatment sequence, I measure parents' income, the first stage, as the family income reported by each student's parents in 1965. Income in adulthood, the second stage, is the family income reported by the student in 1982. The treatment variable is a four-category variable (based on income quartiles) that ranges from low to high income. The confounders of participation and income included in the model were selected based on findings in the previous literature.²⁰

I estimate the ACDE of income in youth on political participation using a stabilized inverse-probability-weighted MSM as described above. In brief, I fit a weighted logistic regression model of the form:

$$\begin{aligned} \Pr(\text{Political event}_{i,1982} = 1) = & \text{logit}(\alpha_0 + \alpha_1 \text{Income}_{i,1965} \\ & + \alpha_2 \text{Income}_{i,1982} \\ & + \alpha_3 \text{Race} + \alpha_4 \text{Gender}), \end{aligned}$$

for the events *attending a rally* and *giving money to a campaign*, where 1 indicates that the respondent engaged in that activity in the period between 1973 and 1982.

I account for potential confounding of time-varying items by fitting the earlier models with stabilized inverse probability weights of the form:

$$\begin{aligned} \mathcal{W}_t^{-1} &= w_{(1965)}^{-1} \times w_{(1982)}^{-1} \\ &= \frac{f(Z^{(1965)}|X^{(1965)})}{f(Z^{(1965)})} \times \frac{f(Z^{(1982)}|Z^{(1965)}, X^{(1965)}, X^{(1982)})}{f(Z^{(1982)}|Z^{(1965)})}, \end{aligned}$$

where $f(\cdot)$ is the inverse of the ordered categorical logistic regression to estimate probabilities.

The predicted probabilities for the numerator and denominator were assigned based on the income category that each panelist reported.²¹

Table 3 presents the results for each of the two main outcomes of interest: attending a rally and donating money. For comparison purposes, I implement three modeling strategies: a weighted MSM model, an *overcontrol* model that explicitly controls for all covariates regardless of whether these are post-treatment, and a third *undercontrol* model that excludes relevant confounders.

²⁰Full description in Appendix.

²¹For the weight estimation model and bootstrapping I used my own code in R. Existing packages do not handle treatments with multiple categories. The functions are available upon request. The weights were re-estimated in each of the 500 bootstrapped samples.

Table 3. Controlled direct effects of early income on participation

	Attend a rally			Donate money		
	MSM (1)	Over control (2)	Under control (3)	MSM (4)	Over control (5)	Under control (6)
Parents' income (1965)						
Quartile 2	0.076 ⁺ (0.045)	-0.019 (0.036)	0.015 (0.036)	0.021 (0.053)	-0.036 (0.036)	0.008 (0.036)
Quartile 3	0.106 ⁺ (0.058)	0.023 (0.044)	0.050 (0.043)	0.063 (0.065)	0.008 (0.042)	0.058 (0.041)
Quartile 4	0.197 (0.092)	0.048 (0.055)	0.082 (0.049)	0.063 (0.099)	0.008 (0.054)	0.081 (0.050)
Income adulthood (1982)						
Quartile 2	0.005 (0.057)	-0.004 (0.039)	0.033 (0.040)	0.0004 (0.078)	0.008 (0.042)	0.044 (0.042)
Quartile 3	-0.058 (0.060)	-0.037 (0.045)	0.018 (0.045)	-0.001 (0.066)	0.006 (0.047)	0.058 (0.048)
Quartile 4	0.081 (0.066)	0.086 ⁺ (0.044)	0.185 (0.044)	0.196 (0.078)	0.189 (0.048)	0.288 (0.048)
N	963	963	963	963	963	963
Post-treatment controls		✓			✓	

Note: Coefficient estimates for covariates/controls omitted. **Bolded** coefficients reliable at more than 95 percent, ⁺ at more than 90 percent. Controls include education of both mother and father, political interest and race of the head of the household, student's education, political efficacy, political interest and knowledge, gender, and race. Regressions include gender and race of the student as strictly pre-treatment covariates. Cut-points and constant terms omitted.

There are significant differences in the magnitude and reliability of the coefficients. First, the results from the MSM in column 1 indicate that parents' income has a reliable and positive impact on the propensity of an individual to attend a rally once she is an adult.²² Being in the fourth quartile of parents' income increases participation in a rally independent of the effect of income in adulthood. However, the results from models 2 and 3 do not support this finding. The overcontrol model does not yield any reliable coefficient (at 95 percent level), while the undercontrol model indicates that belonging to the fourth quartile of income in adulthood increases the probability of attending a rally.

Intuitively and substantively, the results from the MSM are in line with theoretical expectations: activities like a rally are less dependent on economic resources acquired in late stages of life and more likely to be affected by other traits such as group consciousness (Miller *et al.*, 1981) or cross-cutting networks (Mutz, 2002) that are influenced by socio-economic conditions in childhood. The results support the idea that even if all adults manage to close the income gap, there would still be a pervasive effect of inequality on the propensity to participate in rallies. However, the results of the traditional regression models fail to recover this effect.

The results for the "Donate money" outcome are also consistent with theoretical expectations. The effect of parents' income that is not mediated by income in adulthood cannot be distinguished from zero in any of the models considered. However, the effect of income in adulthood is positive and reliable for the fourth income quartile. This suggests that the contemporaneous effect of income in adulthood is more relevant in determining monetary contributions to a candidate or a campaign than any other resources acquired in early stages. The factual resources that income provides, as well as other determinants such as the network that the professional environment and higher income in adulthood might affect, impact the likelihood of engaging in this activity.

²²The baseline category is the lowest income quartile (Quartile 1), so all interpretations of coefficients are made with respect to this category.

Altogether, the results confirm the hypothesis that for certain activities, there is a positive effect of parents' income on political participation. Although the association of income in adulthood and political participation has been widely supported by many authors, it is important to isolate its effect from that of early economic conditions. Although for certain activities, such as donating money, the effect of income in adulthood proves to be stronger (probably due to the necessity of resources provided by income in later stages to complete the task), there are other activities such as attending a rally that are more influenced by other traits and characteristics that are highly likely to be developed in (and shaped by) early economic conditions. These effects are accurately captured and estimated through MSMs in contrast to regression techniques that may lead us to substantively different conclusions.

4. Conclusion

The estimation of direct effects is increasingly receiving attention from scholars in multiple fields. Disentangling causal paths is a strategy that has the potential to improve our understanding of a wide variety of political phenomena. Moreover, the analysis of complex structures, such as those in which there are time-varying treatments and confounders, motivates several research questions in multiple fields that can be answered with the estimation of CDEs. More generally, we can explore the effect of a baseline treatment on an outcome, when we assume that the intermediate treatment stage is set to a particular level. For example, we may be interested in assessing the effect of zoning criteria on political and community engagement that is not mediated by the area's subsequent capital gain, evaluating the effect of welfare support on approval ratings before and after a policy reform, exploring contemporaneous implications of historical variables, or examining similar dynamic relationships.

The analysis of these cases is challenging in methodological terms. The estimation of the ACDE is complicated when there are time-varying treatments and time-varying confounders affected by the treatment. In dynamic social settings, we have reasons to believe that this is the rule rather than the exception. In this article, I explain the two sources of bias that we are likely to encounter when there are confounders affected by the treatment: first, when these confounders are omitted from a regular regression, the causal effect of an intermediate stage cannot be identified due to confounding bias. Nevertheless, controlling for those confounders may induce bias in the estimation of early treatment stages due to post-treatment control. Under these settings, CDEs cannot be estimated using conventional regression approaches because they do not solve the trade-off between confounding and post-treatment control biases.

In order to solve this bias trade-off, I have used MSMs and IPTW estimators as an alternative for the estimation of the ACDE. Through the calculation of weights that "balance" the marginal distributions of potential outcomes, MSMs account for confounding variables while avoiding post-treatment control bias. I described MSMs' characteristics and presented a detailed description of their implementation especially when dealing with multi-valued treatments. I also illustrated some of the differences in terms of bias between distinct methods for the weights estimation. After the application of this class of models to the analysis of the effects of inequality on political outcomes, I examined the different estimates that we can get from the MSM approach and other common *naïve* models that either under or overcontrol for problematic confounders. More specifically, the results show that the estimates of the effect of parents' income on participation from regular regression techniques differ from those yielded by MSMs. This in turn has an impact on the inferences that we make.

Despite the wide applicability and accessibility of MSMs, there are issues related to these models that motivate several further questions. Possibilities for future research include the implementation of inverse probability of treatment and censored weighting estimators in samples as a way of accounting for panel attrition. This would improve the efficiency and accuracy of the estimates by taking into account a problem that is likely to affect the variables under analysis: attrition and

non-response. In summary, MSMs are a feasible alternative when dealing with panel/time-series structures and time-varying treatments. They offer a straightforward method for the estimation of CDEs, under a small number of assumptions and can be implemented using off-the-shelf software. The application of this method to political questions will certainly lead to a better understanding of the causal associations that exist in the complex systems in which we live.

Supplementary material. To view supplementary material for this article, please visit <https://doi.org/10.1017/psrm.2020.3>.

Acknowledgments. I would like to thank Matt Blackwell, Dave Carlson, Jeff Gill, Erin Hartman, Jonathan Homola, Kosuke Imai, Jacob Montgomery, Marc Ratkovic, Jon Rogowski, Leslie Schwindt-Bayer, Betsy Sinclair, Margit Tavits, Tess Wise, and members of the Testing Political Theories seminar and the Data Science Lab at Washington University in St. Louis for helpful comments.

References

- Acharya A, Blackwell M and Sen M (2016) Explaining causal findings without bias: detecting and assessing direct effects. *American Political Science Review* **110**, 512–29.
- Akee R, Copeland W, Costello EJ, Holbein JB and Simeonova E (2018) Family income and the intergenerational transmission of voting behavior: evidence from an income intervention. NBER Working Paper 24770.
- Almond G and Verba S (1989) *The Civic Culture: Political Attitudes and Democracy in Five Nations*. Newbury Park, CA: Sage.
- Beck PA and Jennings MK (1982) Pathways to participation. *American Political Science Review* **76**, 94–108.
- Blackwell M (2013) A framework for dynamic causal inference in political science. *American Journal of Political Science* **57**, 504–520.
- Blackwell M and Glynn A (2014) How to Make Causal Inferences with Time-Series Cross-Sectional Data. Technical Report Working Paper. Harvard University.
- Brady HE, Verba S and Schlozman KL (1995) Beyond SES: a resource model of political participation. *American Political Science Review* **89**, 271–294.
- Córdova A (2009) Methodological note: measuring relative wealth using household asset indicators. *Americas Barometer Insights* **6**, 1–9.
- Currie J (2008) *Healthy, Wealthy, and Wise: Socioeconomic Status, Poor Health in Childhood, and Human Capital Development*. Technical Report. National Bureau of Economic Research.
- Elwert F and Winship C (2014) Endogenous selection bias: the problem of conditioning on a collider variable. *Annual Review of Sociology* **40**, 31–53.
- Glynn AN and Quinn KM (2010) An introduction to the augmented inverse propensity weighted estimator. *Political Analysis* **18**, 36–56.
- Hernán MÁ, Brumback B and Robins JM (2000) Marginal structural models to estimate the causal effect of zidovudine on the survival of HIV-positive men. *Epidemiology* **11**, 561–570.
- Hernán MÁ, Brumback B and Robins JM (2001) Marginal structural models to estimate the joint causal effect of nonrandomized treatments. *Journal of the American Statistical Association* **96**, 440–448.
- Imai K and Ratkovic M (2014) Covariate balancing propensity score. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **76**, 243–263.
- Imai K and Ratkovic M (2015) Robust estimation of inverse probability weights for marginal structural models. *Journal of the American Statistical Association* **110**, 1013–1023.
- Imai K, Keele L and Tingley D (2010) A general approach to causal mediation analysis. *Psychological Methods* **15**, 309–334.
- Imai K, Keele L, Tingley D and Yamamoto T (2011) Unpacking the black box of causality: learning about causal mechanisms from experimental and observational studies. *American Political Science Review* **105**, 765–789.
- Jennings MK, Markus GB, Niemi RG and Stoker L (2005) Youth-parent socialization panel study, 1965–1997: four waves combined.
- Kang JDY and Schafer JL (2007) Demystifying double robustness: a comparison of alternative strategies for estimating a population mean from incomplete data. *Statistical Science* **22**, 523–539.
- Lipset SM (1959) Some social requisites of democracy: economic development and political legitimacy. *American Political Science Review* **53**, 69–105.
- Miller AH, Gurin P, Gurin G and Malanchuk O (1981) Group consciousness and political participation. *American Journal of Political Science* **25**, 494–511.
- Miratrix LW, Sekhon JS and Yu B (2013) Adjusting treatment effect estimates by post-stratification in randomized experiments. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **75**, 369–396.
- Montgomery JM and Olivella S (2018) Tree-based models for political science data. *American Journal of Political Science* **62**(3), 729–744.

- Montgomery JM, Nyhan B and Torres M** (2018) How conditioning on posttreatment variables can ruin your experiment and what to do about it. *American Journal of Political Science* **62**, 760–775.
- Moore JC and Welniak EJ** (2000) Income measurement error in surveys: a review. *Journal of Official Statistics* **16**, 331.
- Mutz DC** (2002) The consequences of cross-cutting networks for political participation. *American Journal of Political Science* **46**, 838–855.
- Nandi A, Glymour MM, Kawachi I and VanderWeele TJ** (2012) Using marginal structural models to estimate the direct effect of adverse childhood social conditions on onset of heart disease, diabetes, and stroke. *Epidemiology* **23**, 223–232.
- Ojeda C** (2018) The two income-participation gaps. *American Journal of Political Science* **62**, 813–829.
- Pearl J** (2001) Direct and indirect effects. *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence (UAI '01)*. San Francisco, CA: Morgan Kaufmann Publishers Inc., pp. 411–420.
- Pearl J** (2011) *The Mediation Formula: A Guide to the Assessment of Causal Pathways in Nonlinear Models*. Technical Report DTIC Document.
- Platt RW, Delaney JAC and Suissa S** (2012) The positivity assumption and marginal structural models: the example of warfarin use and risk of bleeding. *European Journal of Epidemiology* **27**, 77–83.
- Robins JM** (1997) Causal inference from complex longitudinal data. In Berkane M. (eds) *Latent Variable Modeling and Applications to Causality*. New York, NY: Springer, pp. 69–117.
- Robins JM** (1999) Association, causation, and marginal structural models. *Synthese* **121**, 151–179.
- Robins JM, Hernán MÁ and Brumback B** (2000) Marginal structural models and causal inference in epidemiology. *Epidemiology* **11**, 550–560.
- Rosenbaum PR** (1984) The consequences of adjustment for a concomitant variable that has been affected by the treatment. *Journal of the Royal Statistical Society. Series A (General)* **147**, 656–666.
- Rubin DB** (1977) Assignment to treatment group on the basis of a covariate. *Journal of Educational and Behavioral Statistics* **2**, 1–26.
- Samii C, Paler L and Daly S** (2017) Retrospective causal inference with machine learning ensembles: an application to Anti-Recidivism Policies in Colombia. *Political Analysis* **24**, 434–456.
- VanderWeele TJ** (2009) Marginal structural models for the estimation of direct and indirect effects. *Epidemiology* **20**, 18–26.
- VanderWeele TJ** (2010) Bias formulas for sensitivity analysis for direct and indirect effects. *Epidemiology* **21**, 540–551.
- Verba S and Nie NH** (1972) *Participation in America: Political Democracy and Social Equality*. Chicago, IL: University of Chicago Press.
- Verba S, Nie NH and Kim J** (1978) *Participation and Political Equality: A Seven-Nation Comparison*. Chicago, IL: University of Chicago Press.
- Watkins S, Jonsson-Funk M, Brookhart MA, Rosenberg SA, O'shea TM and Daniels J** (2013) An empirical comparison of tree-based methods for propensity score estimation. *Health Services Research* **48**, 1798–1817.
- Westreich D, Cole SR, Schisterman EF and Platt RW** (2012) A simulation study of finite-sample properties of marginal structural Cox proportional hazards models. *Statistics in Medicine* **31**, 2098–2109.