

Identifying causal mechanisms (primarily) based on inverse probability weighting

Martin Huber

University of St. Gallen, Dept. of Economics

Abstract: This paper demonstrates the identification of causal mechanisms of a binary treatment under selection on observables, (primarily) based on inverse probability weighting. I.e., we consider the average indirect effect of the treatment, which operates through an intermediate variable (or mediator) that is situated on the causal path between the treatment and the outcome, as well as the (unmediated) direct effect. Even under random treatment assignment, subsequent selection into the mediator is generally non-random such that causal mechanisms are only identified when controlling for confounders of the mediator and the outcome. To tackle this issue, units are weighted by the inverse of their conditional treatment propensity given the mediator and observed confounders. We show that the form and applicability of weighting depend on whether some confounders are themselves influenced by the treatment or not. A simulation study gives the intuition for these results and an empirical application to the direct and indirect health effects (through employment) of the U.S. Job Corps program is also provided.

Keywords: causal mechanisms, mediation analysis, direct and indirect effects, inverse probability weighting.

JEL classification: C14, C21, I38.

An earlier version of this paper was circulated under the title “Identifying causal mechanisms in experiments (primarily) based on inverse probability weighting”. I have benefited from comments by Guido Imbens, Michael Lechner, Giovanni Mellace, Andreas Steinmayr, Teppei Yamamoto, participants at “Frontiers in the Analysis of Causal Mechanisms” (Harvard, Mar 2012), “Symposium on Causality” (Jena, Jul 2012), the EALE Annual Meetings (Bonn, Sep 2012), the IAB conference “Field Experiments in Policy Evaluation” (Nuremberg, Oct 2012), and the IFAU/IZA Conference on Labor Market Policy Evaluation (Bonn, Oct 2012), and three anonymous referees. Financial support from the Swiss National Science Foundation grant PBSGP1_138770 is gratefully acknowledged. Address for correspondence: Martin Huber, University of St. Gallen, Varnbühlstrasse 14, CH-9000 St. Gallen, martin.huber@unisg.ch.

1 Introduction

A vast literature in economics and other social sciences is concerned with the evaluation of average treatment effects (ATE), both in randomized experiments and in observational studies. In many evaluations not only the (total) ATE appears interesting, but also the causal mechanisms through which it operates. In this case, one would like to disentangle the *direct* effect of the treatment on the outcome as well as the *indirect* ones that run through one or more intermediate variables, so-called mediators. E.g., when assessing the employment or earnings effects of an active labor market policy, policy makers might want to know to which extent the total impact comes from increased search effort, human capital, or other mediators that are themselves affected by the policy. However, even in experiments, causal mechanisms are not easily identified. As discussed in Robins and Greenland (1992), random treatment assignment does not imply randomness of the mediator, which may be regarded as intermediate outcome. Therefore, the total effect cannot be disentangled by simply conditioning on a mediator, because this generally introduces selection bias coming from variables influencing both the mediator and the outcome, see Rosenbaum (1984).¹

The main contribution of this paper is to show that an easily implemented version of inverse probability weighting² (IPW) identifies causal mechanisms under both discrete or continuous mediators, given that a sequential selection on observables (or conditional independence) assumption holds. The latter requires (i) that the treatment is either random or exogenous given the covariates and (ii) that the mediator is exogenous given the covariates and the treatment, see for instance Imai, Keele, and Yamamoto (2010). Then, direct and indirect effects can be identified by weighting observations by their inverse conditional propensities to be in a particular treatment state (i) given the observed covariates and (ii) given the mediator and the observed covariates. Furthermore, we also discuss that the identification results for the indirect effects change if some of the covariates are themselves a function of the treatment. If the latter is the case, the identifi-

¹For this reason, already the early work on mediation analysis of Judd and Kenny (1981) highlights the importance of controlling for such confounders. It therefore seems surprising that this issue has been ignored in so many applications in social sciences that claim to identify direct and indirect effects.

²The idea of IPW goes back to Horvitz and Thompson (1952), who first proposed an estimator of the population mean in the presence of non-randomly missing data.

cation of the “total” indirect effect, which also accounts for correlations between those covariates affected by the treatment and the mediator, requires additional restrictions, see Robins (2003), Avin, Shpitser, and Pearl (2005), and Imai and Yamamoto (2011). In contrast, the “partial” indirect effect, which only considers the immediate link between the treatment and the mediator (and no “detour” via any covariates), is identified under weaker assumptions. We provide a simulation study that gives the intuition for these identification issues and apply our methods to experimental data on Job Corps, a U.S. educational program for disadvantaged youths.

The evaluation of direct and indirect effects, often referred to as mediation analysis, is widespread in social sciences such as epidemiology, political sciences, and psychology, see MacKinnon (2008). While many studies follow Baron and Kenny (1986) and rely on linear mediation models, more general identification under conditional exogeneity of the treatment and the mediator has been considered by Pearl (2001), Robins (2003), Petersen, Sinisi, and van der Laan (2006), VanderWeele (2009), Imai, Keele, and Yamamoto (2010), Albert and Nelson (2011), and Imai and Yamamoto (2011), among others. One of the rare studies in economics is Flores and Flores-Lagunes (2009), who evaluate the direct earnings effect of Job Corps when controlling for the mediator “work experience”. The issue is that participating in a training likely decreases work experience shortly after program start compared to nonparticipation due to decreased job search effort during training participation (“locking-in effect”). Assuming mediator exogeneity conditional on pre-treatment covariates, Flores and Flores-Lagunes (2009) estimate a positive direct effect on earnings based on a regression approach.³

If conditional exogeneity does not hold (and plausible instruments are not available), point identification is lost, but partial identification based on deriving upper and lower bounds on the direct and indirect effects might still be useful. E.g., Kaufman, Kaufman, MacLenose, Greenland, and Poole (2005), Cai, Kuroki, Pearl, and Tian (2008), and Sjölander (2009) focus on partial identification in randomized medical trials with binary treatments and impose specific restric-

³As a further example, Simonsen and Skipper (2006) use a semiparametric identification strategy based on matching to assess the direct wage effect of motherhood in Denmark by controlling for several mediators through which motherhood may have an influence on wages. They find negative direct effects which vary little across different sectors.

tions such as monotonicity of the mediator in the treatment in order to tighten the bounds. In economics, Flores and Flores-Lagunes (2010) consider additional assumptions (e.g., a particular order of the mean potential outcomes of various subpopulations) and assess the effectiveness of various components of the Job Corps program.

This paper focusses on point identification and makes four contributions to the literature on causal mechanisms in economics: Firstly, it derives identification results based on IPW by the treatment propensity score that are straightforward to implement by semi- or nonparametric estimation. If the mediator is exogenous conditional on pre-treatment covariates, our approach allows relaxing one functional form assumption imposed in Flores and Flores-Lagunes (2009) (their Assumption 3). It is also easier to implement than the nonparametric estimators of Imai, Keele, and Yamamoto (2010), which require estimating the conditional mean of the outcome and the conditional density of the mediator. Secondly and in contrast to Flores and Flores-Lagunes (2009), we also discuss identification when mediator exogeneity only holds conditional on post-treatment covariates which are themselves a function of the treatment, such that pre-treatment variables do not fully capture mediator selection. This appears realistic in most applications including Job Corps, where the treatment likely affects variables that potentially confound the mediator and the outcome, e.g., intermediate health shortly prior to the mediator. While direct effects are still identified by IPW in this set up after a modification of the initial assumptions, the identification of indirect effects requires additional restrictions. We present a functional form restriction allowing us to do so, which, however, is less general than the entirely nonparametric identification under IPW. Thirdly, we show that IPW still identifies a partial indirect effect when keeping the confounders fixed, i.e., the part of the indirect effect not working through post-treatment confounders. Fourthly, as an empirical contribution, the present work appears to be the first which assesses the direct and indirect health effects of the Job Corps program.

The remainder of this paper is organized as follows. Section 2 defines the parameters of interest (the average direct and indirect effects) and shows identification (mostly) based on IPW. Section 3 briefly discusses estimation. Section 4 presents a simulation study which provides the

intuition for various identification issues. In Section 5, we apply our methods to the experimental study of the Job Corps program. Section 6 concludes.

2 Parameters of interest and identification

2.1 Definition of parameters

Suppose we are interested in the average treatment effect (ATE) of a binary treatment indicator D on some outcome variable Y . Furthermore, assume that we would like to disentangle the ATE into a direct component and an indirect effect operating through the mediator M which has bounded support and may be discrete or continuous. To define the parameters of interest, we use the potential outcome framework advocated by Rubin (1974) (among many others) and considered in the direct and indirect effects framework for instance by Rubin (2004), Ten Have, Joffe, Lynch, Brown, Maisto, and Beck (2007), and Albert (2008). Let $Y(d)$, $M(d)$ denote the potential outcome and the potential mediator state under treatment $d \in \{0, 1\}$. For each unit only one of the two potential outcomes and mediator states, respectively, is observed, because the realized outcome and mediator values are $Y = D \cdot Y(1) + (1 - D) \cdot Y(0)$ and $M = D \cdot M(1) + (1 - D) \cdot M(0)$.

The ATE is defined by $\Delta = E[Y(1) - Y(0)]$. To disentangle this total effect into a direct and indirect (through M) causal channel, first note that the potential outcome can be rewritten as a function of both the treatment and the intermediate variable M : $Y(d) = Y(d, M(d))$. It follows that the (average) direct effect is identified by

$$\theta(d) = E[Y(1, M(d)) - Y(0, M(d))], \quad d \in \{0, 1\}, \quad (1)$$

i.e., by exogenously varying the treatment but keeping the mediator fixed at its potential value for $D = d$. Equivalently, the (average) indirect effects is defined as

$$\delta(d) = E[Y(d, M(1)) - Y(d, M(0))], \quad d \in \{0, 1\}, \quad (2)$$

i.e., by exogenously shifting the mediator to its potential values under treatment and non-treatment but keeping the treatment fixed at $D = d$.⁴ Note that the ATE is the sum of the direct and indirect effects defined upon opposite treatment states:

$$\begin{aligned}\Delta &= E[Y(1, M(1)) - Y(0, M(0))] \\ &= E[Y(1, M(1)) - Y(0, M(1))] + E[Y(0, M(1)) - Y(0, M(0))] = \theta(1) + \delta(0) \\ &= E[Y(1, M(0)) - Y(0, M(0))] + E[Y(1, M(1)) - Y(1, M(0))] = \theta(0) + \delta(1),\end{aligned}\quad (3)$$


which follows from adding and subtracting $E[Y(0, M(1))]$ after the first and $E[Y(1, M(0))]$ after the third equality. The notation $\theta(1), \theta(0)$ and $\delta(1), \delta(0)$ highlights the possibility of effect heterogeneity w.r.t. the treatment state, i.e., the presence of interaction effects between the treatment and the mediator. However, it is obvious that these effects cannot be identified without further assumptions, as either $Y(1, M(1))$ or $Y(0, M(0))$ is observed for any unit, whereas $Y(1, M(0))$ and $Y(0, M(1))$ are never observed. Therefore, identification of direct and indirect effects hinges on the existence of exogenous variation in the treatment and the mediator.

2.2 Identification given observed confounders not affected by the treatment

We now introduce our identifying assumptions, maintaining an i.i.d. framework throughout the paper. We start with the framework of conditional mediator exogeneity given the treatment and observed covariates (denoted by X) which are themselves *not* a function of D , with the leading case being pre-treatment covariates (evaluated prior to treatment assignment). Figure 1 provides a graphical illustration using a directed acyclic graph, where each arrow represents a causal path. Further below we will consider another set of restrictions assuming conditional mediator exogeneity given the treatment and covariates that are (at least partially) themselves a function of D , and thus, post-treatment variables, which makes identification more difficult.

⁴Pearl (2001) named these parameters the natural direct and indirect effects, whereas Robins (2003) referred to them as the pure direct and indirect effects and Flores and Flores-Lagunes (2009) as net and mechanism average treatment effects, respectively.

Figure 1: Causal paths under conditional exogeneity of the mediator



Our first assumption requires the treatment to be conditionally independent (given X) of any potential post-treatment variable, i.e., the potential mediator states and the potential outcomes. This is referred to as conditional independence, selection on observables, or exogeneity in the treatment evaluation literature, see for instance Imbens (2004).

Assumption 1 (conditional independence of the treatment):

$\{Y(d', m), M(d)\} \perp D | X$ for all $d', d \in \{0, 1\}$ and m in the support of M .

Assumption 1 implies that there are no unobserved confounders jointly affecting the treatment on the one hand and the mediator and/or the outcome on the other hand conditional on the covariates X . In observational studies the plausibility of this assumption, which has been criticized among others by Heckman and Navarro-Lozano (2004), critically hinges on the richness of the data. In experiments, it is satisfied if the treatment is either randomized within strata defined on X or randomized unconditionally, i.e., independent of X . (In the latter case, even the stronger condition $\{Y(d', m), M(d), X\} \perp D$ holds.)

The second assumption imposes conditional independence of the mediator given the treatment and the covariates along with a common support restriction on the conditional treatment probability:

Assumption 2 (conditional independence of the mediator):

- (a) $Y(d', m) \perp M | D = d, X = x$ for all $d', d \in \{0, 1\}$ and m, x in the support of M, X ,
- (b) $\Pr(D = d | M = m, X = x) > 0$ for all $d \in \{0, 1\}$ and m, x in the support of M, X .

Assumption 2(a) states that conditional on D and X , the effect of the mediator on the outcome

is unconfounded. This implies that there exist no unobserved confounders jointly causing the mediator and the outcome given the treatment and the covariates. Assumption 2(a) is for instance violated if unobserved pre-treatment variables affect both M and Y directly, i.e., not only through D and X . Neither does it hold if unobserved post-treatment variables influence M and Y and are not fully determined by X and/or D . Therefore, 2(a) is a very strong restriction and only appears realistic if detailed information on potential confounders of the mediator is available in the data (even in experiments with random treatment assignment where Assumption 1 is naturally satisfied). Furthermore, the described issue of post-treatment confounding must be plausibly ruled out. Note that alternatively to 2(a), identification might also be based on an instrument for the mediator as for instance discussed in Imai, Keele, Tingley, and Yamamoto (2010).

Assumption 2(b) is a common support restriction requiring that the conditional probability to be treated given M, X , henceforth referred to as propensity score, is larger than zero in either treatment state. It follows that $\Pr(D = d|X = x) > 0$ must hold, too. Note that by Bayes' theorem, Assumption 2(b) also implies that $\Pr(M = m|D = d, X = x) > 0$ (or in the case of M being continuous, that the conditional density of M given D, X is larger than zero: $f_{M|D,X}(m, d, x) > 0$). I.e., conditional on X , the mediator state must not be a deterministic function of the treatment, otherwise identification is infeasible due to the lack of comparable units in terms of the mediator across treatment states. Assumptions 1 and 2 correspond to the sequential ignorability assumption of Imai, Keele, and Yamamoto (2010) (their Assumption 1) and Tchetgen Tchetgen and Shpitser (2011b), among others (see also the closely related Theorem 2 of Pearl (2001)). Also Flores and Flores-Lagunes (2009) use a similar set of restrictions in their Assumptions 2, 5, and 6, but in addition impose a functional form assumption (their Assumption 3) which is not necessary for nonparametric identification.

Assumptions 1 and 2 allow us to identify $E[Y(d, M(d))]$ and $E[Y(d, M(1-d))]$. Starting with

the former, note that

$$\begin{aligned}
E[Y(d, M(d))] &= E[E[Y(d, M(d))|X = x]] = E[E[Y|D = d, X = x]] \\
&= E \left[E \left[\frac{Y \cdot I\{D = d\}}{\Pr(D = d|X)} \middle| X = x \right] \right] = E \left[\frac{Y \cdot I\{D = d\}}{\Pr(D = d|X)} \right]. \tag{4}
\end{aligned}$$

where the first equality follows from the law of iterated expectations, the second from Assumption 1, the third from basic probability theory, and the last from the law of iterated expectations (see also Hirano, Imbens, and Ridder (2003)). Concerning the latter,

$$\begin{aligned}
&E[Y(d, M(1 - d))] \\
&= \int \int E[Y(d, m)|M(1 - d) = m, X = x] dF_{M(1-d)|X=x}(m) dF_X(x) \\
&= \int \int E[Y(d, m)|D = d, M = m, X = x] dF_{M|D=1-d, X=x}(m) dF_X(x) \\
&= \int \int E[Y|D = d, M = m, X = x] \cdot \frac{\Pr(D = 1 - d|M, X)}{\Pr(D = 1 - d|X)} dF_{M|X=x}(m) dF_X(x) \\
&= E \left[E \left[E \left[\frac{Y \cdot I\{D = d\}}{\Pr(D = d|M, X)} \middle| M = m, X = x \right] \cdot \frac{\Pr(D = 1 - d|M, X)}{\Pr(D = 1 - d|X)} \middle| X = x \right] \right] \\
&= E \left[\frac{Y \cdot I\{D = d\}}{\Pr(D = d|M, X)} \cdot \frac{\Pr(D = 1 - d|M, X)}{\Pr(D = 1 - d|X)} \right]. \tag{5}
\end{aligned}$$

The first equality follows from the law of iterated expectations and from replacing the outer expectations by integrals, the second from Assumptions 1 and 2, the third from Bayes' theorem, the fourth from basic probability theory and from replacing the integrals by expectations, and the last from the law of iterated expectations. Therefore, $\theta(d)$ and $\delta(d)$ are identified by either subtracting (5) from (4) or vice versa, depending on whether d is one or zero. It follows by simple algebra that the direct and indirect effects are obtained from Propositions 1 and 2.⁵

⁵Propositions 1 and 2 can also be derived by starting from the mediation formulae (see Pearl (2001)) provided in (8), which identify the direct and indirect effects under Assumptions 1 and 2. E.g., considering the direct effect, note that

$$\begin{aligned}
\theta(d) &= \int \int \{E[Y|D = 1, M = m, X = x] - E[Y|D = 0, M = m, X = x]\} dF_{M|D=d, X=x}(m) dF_X(x) \\
&= \int \int \{E[Y|D = 1, M = m, X = x] - E[Y|D = 0, M = m, X = x]\} \cdot \frac{\Pr(D = d|M, X)}{\Pr(D = d|X)} dF_{M|X=x}(m) dF_X(x) \\
&= E \left[E \left[E \left[\frac{Y \cdot D}{\Pr(D = 1|M, X)} - \frac{Y \cdot (1 - D)}{1 - \Pr(D = 1|M, X)} \right] \cdot \frac{\Pr(D = d|M, X)}{\Pr(D = d|X)} \middle| X = x \right] \right] = E \left[\left(\frac{Y \cdot D}{\Pr(D = 1|M, X)} - \frac{Y \cdot (1 - D)}{1 - \Pr(D = 1|M, X)} \right) \cdot \frac{\Pr(D = d|M, X)}{\Pr(D = d|X)} \right].
\end{aligned}$$

Proposition 1:

Under Assumptions 1 and 2, the average direct effect is identified by

$$\theta(d) = E \left[\left(\frac{Y \cdot D}{\Pr(D = 1|M, X)} - \frac{Y \cdot (1 - D)}{1 - \Pr(D = 1|M, X)} \right) \cdot \frac{\Pr(D = d|M, X)}{\Pr(D = d|X)} \right]. \quad (6)$$

Proposition 1 implies that by propensity score-based weighting, the distributions of both M and X are balanced between treatment and control groups such that the direct effect is identified. In particular, the distribution of the mediator in both groups corresponds to that of $M(d)$ in the total population. Concerning the indirect effect, note that by (3), it corresponds to the difference between the average and the direct effect defined on the opposite treatment state: $\delta(d) = \Delta - \theta(1 - d)$. Proposition 2 provides the representation of the indirect effect based on IPW, which is numerically identical to this difference:

Proposition 2:

Under Assumptions 1 and 2, the average indirect effect is identified by

$$\delta(d) = E \left[\frac{Y \cdot I\{D = d\}}{\Pr(D = d|M, X)} \cdot \left(\frac{\Pr(D = 1|M, X)}{\Pr(D = 1|X)} - \frac{1 - \Pr(D = 1|M, X)}{1 - \Pr(D = 1|X)} \right) \right]. \quad (7)$$

From a practitioner's perspective, a nice feature of these identification results is that they are straightforward to implement. They only involve the (possibly parametric or nonparametric) estimation of two binary choice models for the propensity scores which are then plugged into the (normalized) sample analogs of Propositions 1 and 2, as outlined in Section 3. No parametric restrictions are imposed on the models of the outcome and the mediator such that arbitrary nonlinearities are allowed for. In contrast, the standard approach in the literature consists in estimating the ingredients of the following alternative representations of the parameters of interest, see for instance equations (8) and (26) in Pearl (2001) and Theorem 1 in Imai, Keele, and

Yamamoto (2010):

$$\theta(d) = \int \int \{E[Y|D = 1, M = m, X = x] - E[Y|D = 0, M = m, X = x]\} dF_{M|D=d, X=x}(m) dF_X(x), \quad (8)$$

$$\delta(d) = \int \int E[Y|D = d, M = m, X = x] \{dF_{M|D=1, X=x}(m) - dF_{M|D=0, X=x}(m)\} dF_X(x). \quad (9)$$

This requires estimators for the conditional mean of Y given D, M, X and the conditional density of M given D, X . In the literature, parametric methods have been most commonly used, see for instance Pearl (2011) and VanderWeele (2009).⁶ They, however, appear unattractive due to their severe functional form restrictions and the potentially difficult interpretability of direct and indirect effects under non-linear modeling (e.g., when both the outcome and the mediator are binary). Nonparametric estimation as recently proposed in Imai, Keele, and Yamamoto (2010) avoids these shortcomings, but might be cumbersome in empirical applications if X is high dimensional and/or M is continuous. In contrast, estimation based on Propositions 1 and 2 is less prone to such issues as it just relies on two propensity score models. Our IPW-based results are also more general than the regression approach of Flores and Flores-Lagunes (2009). The latter does not require the estimation of the conditional density of the mediator, but imposes a functional form restriction (their Assumption 3) on the expected potential outcomes across potential mediator states (for the treatment fixed) which we need not invoke here.⁷

However, it is important to note that IPW also has its drawbacks: If the common support assumption 2(b) is close to being violated, estimation may be unstable and the variance may explode, see Frölich (2004) and Khan and Tamer (2010), among others. Furthermore, IPW is less robust to propensity score misspecification than other classes of estimators, as documented for instance in Kang and Schafer (2007) and Waernbaum (2012). Therefore, matching on the

⁶Furthermore, Tchetgen Tchetgen and Shpitser (2011a) and Zheng and van der Laan (2012) (among others) provide multiply robust parametric estimators based on conditional mean estimation of the outcome and conditional density estimation of both the treatment and the mediator.

⁷Only if the mediator is exogenous conditional on post-treatment confounders which are themselves influenced by the treatment, we have to rely on a similar assumption, as outlined below.

propensity score (see for instance Rosenbaum and Rubin (1983)), which remains consistent under particular forms of propensity score misspecification, might represent a viable alternative to IPW.

2.3 Identification given observed confounders affected by the treatment

It appears unlikely in many applications that conditioning on pre-treatment covariates is sufficient to control for mediator endogeneity, given that the mediator is itself a post-treatment variable. Equivalent to the treatment evaluation literature, where potential confounders of the treatment are measured at or shortly before the treatment, potential confounders of the mediator should be controlled for just before the selection into the mediator takes place. Then, however, it appears likely that at least some of these covariates are a function of the treatment, too, implying that they are themselves mediators that affect the mediator of interest. Therefore, Robins (2003) suspects that the set up relying on Assumptions 1 and 2 is of limited practical relevance. This most likely also applies to our application presented in Section 5, where we are interested in the effect of the Job Corps program on health. The mediator is employment and clearly, some potential confounders affecting both employability and health (such as the labor market state shortly prior to employment) are most likely a function of the treatment.

As in Robins (2003) and Imai and Yamamoto (2011), we therefore also consider a framework in which D is permitted to have an effect on post-treatment confounders of the mediator, which we denote by W . In this case, mediation analysis becomes more complicated and requires us to introduce additional notation by rewriting the potential mediator and potential outcome also as functions of W : $M(d) = M(d, W(d))$ and $Y(d, M(d)) = Y(d, M(d, W(d)), W(d))$, where $W(d)$ is the vector of potential values of W for $D = d$. Then, the total indirect effect is defined as

$$\delta^t(d) = E[Y(d, M(1, W(1)), W(d)) - Y(d, M(0, W(0)), W(d))]. \quad (10)$$

We refer to $\delta^t(d)$ as the total indirect effect because it comprises all effects via M which either come from D directly or “take a devious route” through W . I.e., this parameter accounts for the

fact that M is affected by D both directly and indirectly through a change in W . In contrast, the partial indirect effect only identifies the effect through M directly coming from D , but not going through W :

$$\delta^p(d) = E[Y(d, M(1, W(d)), W(d)) - Y(d, M(0, W(d)), W(d))]. \quad (11)$$

I.e., $\delta^p(d)$ is the ceteris paribus indirect effect via the mediator when holding W constant at the level implied by d such that any channel through the post-treatment covariates is shut down. This is what regressing Y on $(1, D, M, W)$ and multiplying the coefficient on M with the first stage effect of D on M obtains as “indirect effect” in the limit, given that no further confounders are present. Obviously, this effect neglects any correlations between M and W . We therefore argue that the total indirect effect is the more interesting parameter,⁸ but nevertheless discuss the identification of both parameters. However, it will be shown further below that $\delta^p(d)$ is more easily identified than $\delta^t(d)$. As a further remark, note that $\delta^t(d)$ corresponds to $\delta(d)$ (which also refers to the total of indirect effects) in Section 2.2, where we, however, did not distinguish between total and partial effects due to the absence of post-treatment confounders.

The direct effect is defined as


$$\theta(d) = E[Y(1, M(d, W(d)), W(d)) - Y(0, M(d, W(d)), W(d))], \quad (12)$$

i.e., it corresponds to the change in the mean potential outcome due to an exogenous change in the treatment, while keeping the mediator and the post-treatment covariates fixed. Note that this definition differs from Imai and Yamamoto (2011), who consider the difference between the ATE and the total indirect effect to be the “direct” effect: $E[Y(d, M(d, W(d)), W(d)) - Y(1 - d, M(d, W(d)), W(1 - d))]$. However, this includes changes in the mean potential outcome which are due to a change in W which is not mediated by M . Here, we define the direct effect in a narrower sense that also excludes (inherently

⁸Also Imai and Yamamoto (2011) focus on $\delta^t(d)$ as indirect effect and do not consider $\delta^p(d)$ at all.

indirect) channels via W . For this reason, $\theta(d)$ and $\delta^t(d)$ or $\delta^p(d)$, respectively, do not add up to the ATE, as either $E[Y(d, M(d, W(1)), W(1)) - Y(d, M(d, W(0)), W(0))]$ or $E[Y(d, M(d, W(d)), W(1)) - Y(d, M(d, W(d)), W(0))]$ are not accounted for, respectively.

Figure 2: Causal paths with pre-treatment covariates (X) and post-treatment covariates (W)



The directed acyclic graph in Figure 2 displays a set up where the treatment affects the observed confounders W of the mediator. Because we need to condition on W , identification requires ruling out unobserved confounders that jointly cause W on the one hand and M and/or Y on the other hand. Furthermore, just as in Section 2.2, we also have to control for pre-treatment covariates that jointly affect the treatment and the mediator/outcome. A further identification issue arises if all or some pre-treatment covariates also have an impact on the post-treatment covariates W . Then, conditioning on the latter changes the distribution of pre-treatment covariates even if they were initially balanced across treatment states as in randomized experiments. For this reason, we in addition to the confounders of the treatment also need to control for all those pre-treatment covariates that are jointly related to W on the one hand and directly to the outcome and/or the mediator on the other hand. Otherwise, conditioning on W would introduce a correlation between D and those pre-treatment covariates predicting both W and M or Y and thus, treatment endogeneity. We denote by X the vector of all pre-treatment covariates that may cause one or several of the forms of confounding just discussed and replace Assumptions 1 and 2 by Assumptions 3 and 4 to incorporate post-treatment confounders of the mediator.

Assumption 3 (conditional independence of the treatment):

- (a) $\{Y(d'', m, w'), M(d', w), W(d)\} \perp D | X = x$ for all $d'', d', d \in \{0, 1\}$ and m, w', w, x in the

support of M, W, X ,

(b) $\{Y(d'', m, w''), M(d', w')\} \perp D | W = w, X = x$ for all $d', d \in \{0, 1\}$ and m, w'', w', w, x in the support of M, W, X .

Similar to Assumption 1, Assumption 3(a) says that D is conditionally independent of any potential post-treatment variables, namely the potential outcomes, the potential mediator states, and the potential values of the post-treatment confounders of the mediator. Assumption 3(b) is new and states that the conditional independence of the treatment and the potential outcomes or mediator states must also hold when jointly conditioning on W and X , implying that all pre-treatment covariates affecting both W and M or Y are included in X .

Assumption 4 (conditional independence of the mediator):

(a) $Y(d'', m, w') \perp M | D = d, W = w, X = x$ for all $d'', d' \in \{0, 1\}$ and m, w', w, x in the support of M, W, X ,

(b) $\Pr(D = d | M = m, W = w, X = x) > 0$ for all $d \in \{0, 1\}$ and m, w, x in the support of M, W, X .

Assumption 4(a) is related to, but somewhat weaker than Assumption 2(a), as conditional independence is now only required to hold given pre- and post-treatment covariates. This for instance rules out that unobserved pre-treatment variables affect both M and Y directly, i.e., not only through D , X , and/or W . Furthermore, it requires that any unobserved post-treatment variables which are not a deterministic function of D and/or X must not jointly cause M and Y conditional on W . Assumption 4(b) says that the common support restriction on the treatment propensity must hold when conditioning on M and both W and X .

When comparing our assumptions to others made in the literature, it turns out that Assumptions 3 and 4(a) are similar to FRCISTG (fully randomized causally interpretable structural tree graph) in Robins (2003) (see also Robins (1986)) and Assumption 2 in Imai and Yamamoto (2011), with one important difference: Robins (2003) and Imai and Yamamoto (2011) merely assume that $Y(d, m, W(d))$ and $M(d)$ defined on the same treatment are conditionally independent. In contrast, Assumption 4(a) imposes conditional independence of $Y(d'', m, W(d'))$

and $M(d)$ for possibly distinct d'', d', d (i.e., also for potential outcomes and potential mediator states defined on opposite treatments).⁹ This stronger conditional independence allows for nonparametric identification of the direct and partial indirect effects without further restrictions (other than Assumptions 3 and 4), but not of the total indirect effect, as discussed below. In contrast, FRCISTG does not identify any of the parameters without further functional form restrictions (as imposed in Robins (2003) and Imai and Yamamoto (2011)).

Propositions 3 and 4 concern the identification of the direct effect, which hinges on the identifiability of $E[Y(1 - d, M(d, W(d)), W(d))]$, and of the partial indirect effect, respectively, which requires $E[Y(d, M(1 - d, W(d)), W(d))]$. The respective proofs, which are similar to (5), are provided in Appendix A.2.

Proposition 3:

Under Assumptions 3 and 4, the average direct effect is identified by

$$\theta(d) = E \left[\left(\frac{Y \cdot D}{\Pr(D = 1|M, W, X)} - \frac{Y \cdot (1 - D)}{1 - \Pr(D = 1|M, W, X)} \right) \cdot \frac{\Pr(D = d|M, W, X)}{\Pr(D = d|X)} \right]. \quad (13)$$

Proposition 4:

Under Assumptions 3 and 4, the average partial indirect effect is identified by

$$\delta^p(d) = E \left[\frac{Y \cdot I\{D = d\}}{\Pr(D = d|M, W, X)} \cdot \frac{\Pr(D = d|W, X)}{\Pr(D = d|X)} \cdot \left(\frac{\Pr(D = 1|M, W, X)}{\Pr(D = 1|W, X)} - \frac{1 - \Pr(D = 1|M, W, X)}{1 - \Pr(D = 1|W, X)} \right) \right]. \quad (14)$$

The identification of the total indirect effect requires identifying $E[Y(d, M(1 - d, W(1 - d)), W(d))]$. Unfortunately, this is not feasible without further assumptions, see the proof in Avin, Shpitser, and Pearl (2005) and the results of Albert and Nelson (2011), who consider a sequential conditional independence assumption that is similar to ours. The reason is that conditional on X (when the treatment is as good as random), the identification of $E[Y(d, M(1 - d, W(1 - d)), W(d))]$ requires exogenously adjusting the distribution of M given $D = d$ to that of M given $D = 1 - d$,

⁹Robins and Richardson (2010) present a DGP in their Appendix B where indeed FRCISTG holds while our stronger restriction is not satisfied. However, Robins (2003) argues that it seems hard to construct realistic scenarios where one set of assumptions holds while the other one does not.

while at the same time keeping the distribution of W fixed (given $D = d$). Obviously, this is impossible if W and M are not independent conditional on X .

However, Robins (2003) shows under FRCISTG that the total indirect effect is identified under an additional restriction, namely the absence of interaction effects between D and M . Formally, his assumption implies that the unit-level treatment effect (for any unit i) for the mediator fixed is constant across different values of the mediator:

$$Y_i(1, m, X_i(1)) - Y_i(0, m, X_i(0)) = Y_i(1, m', X_i(1)) - Y_i(0, m', X_i(0)) = B_i,$$

where B_i is an unit-level constant. Unfortunately, this assumption appears unattractive in empirical applications (see for instance the discussion in Imai, Tingley, and Yamamoto (2012), Section 3.1) and restricts the usefulness of nonparametric identification advocated in recent work. However, Imai and Yamamoto (2011) demonstrate that under FRCISTG the assumption of no interaction effect can be relaxed to assuming a homogeneous interaction effect:

$$Y_i(1, m, X_i(1)) - Y_i(0, m, X_i(0)) = B_i + Cm,$$

where C is constant for any m . I.e., the interaction between the treatment and the mediator varies homogeneously for all observations.

Here, we propose a functional form restriction on the relation of the mean potential outcome and the mediator across treatment states that is comparable to Assumption 3 in Flores and Flores-Lagunes (2009) (who, however, use it in the set up where D does not affect W).

Assumption 5 (functional form of mean potential outcome-mediator relation):

- (a) $E[Y(d, m, w)|D = d, M(d, W(d)) = m, W(d) = w, X = x] = \mu(d, m, w, x)$ for all $d \in \{0, 1\}$ and m, w, x in the support of M, W, X , with the function $\mu(D, M, W, X)$ being linear in M ,
- (b) $E[Y(d, m, w)|D = d, M(d, W(d)) = m, W(d) = w, X = x] = E[Y(d, m, w)|D = d, M(1 - d, W(1 - d)) = m, W(d) = w, X = x]$ for all $d \in \{0, 1\}$ and m, w, x in the support of M, W, X .

By Assumption 5(a) the conditional mean potential outcome under $D = d$ can be characterized by the regression model μ , which is restricted to be a linear function of the mediator. Assumption 5(b) states that one can use this very same regression model to predict $E[Y(d, m, x) | M(1 - d, W(1 - d)) = m, W(d) = w, X = x]$. I.e., the functional form of $\mu(d, M(d, W(d)), W(d), X)$ is assumed to be identical to $\mu(d, M(1 - d, W(1 - d)), W(d), X)$ whenever $M(d, W(d)) = M(1 - d, W(1 - d))$. This restriction implies that the mean interaction effect of the treatment and the mediator is the same for any $M(1, W(1)) = M(0, W(0)) = m$ conditional on $W(d)$ and X , otherwise the regression functions necessarily differ. Together with Assumptions 3 and 4, Assumption 5 permits the identification of $E[Y(d, M(1 - d, W(1 - d)), W(d))]$, even though $M(1 - d, W(1 - d))$ is not known for units with $D = d$. The trick is to use $E[M(1 - d, W(1 - d))]$, which can be identified in the population with $D = 1 - d$, in our regression model: $\mu(d, E[M(1 - d, W(1 - d))], W(d), x)$. By the linearity assumption 5(b) and the law of iterated expectations it follows that $E[\mu(d, E[M(1 - d, W(1 - d))], w, x) | D = d] = E[\mu(d, M(1 - d, W(1 - d)), w, x) | D = d]$, which allows identifying $E[Y(d, M(1 - d, W(1 - d)), W(d))]$. However, it has to be stressed that Assumption 5 is far from being innocuous. Firstly, it requires a correctly specified regression model that allows making predictions across mediator states. Secondly, the linearity assumption rules out more general relations between the outcome and the mediator.

Given that these considerably stronger functional form assumptions are satisfied, $E[Y(d, M(1 - d, W(1 - d)), W(d))]$, which is required for identifying the total indirect effect, is

obtained from the following result:

$$\begin{aligned}
& E[Y(d, M(1-d, W(1-d)), W(d))] \\
&= \int E[E[Y(d, M(1-d, W(1-d)), w)|W(d)=w, X=x|X=x]dF_X(x) \\
&= \int E[E[Y(d, M(1-d, W(1-d)), w)|D=d, W=w, X=x|D=d, X=x]dF_X(x) \\
&= \int E[E[Y(d, E[M(1-d, W(1-d))], w)|D=d, W=w, X=x|D=d, X=x]dF_X(x) \\
&= \int E[E[Y(d, E[M(1-d, W(1-d))], w)|D=d, M=E[M(1-d, W(1-d))], W=w, X=x|D=d, X=x]dF_X(x) \\
&= \int E[\mu(d, E[M(1-d, W(1-d))], W, x)|D=d, X=x]dF_X(x) \\
&= \int E\left[\frac{\mu(d, E\left[E\left[\frac{M \cdot I\{D=1-d\}}{\Pr(D=1-d|X)}|X=x\right], W, x) \cdot I\{D=d\}}{\Pr(D=d|X)}\middle|X=x\right]dF_X(x) \\
&= E\left[\frac{\mu(d, E\left[\frac{M \cdot I\{D=1-d\}}{\Pr(D=1-d|X)}\right], W, X) \cdot I\{D=d\}}{\Pr(D=d|X)}\right]. \tag{15}
\end{aligned}$$

The first equality follows from the law of iterated expectations and from replacing the outer expectation by an integral, the second from Assumption 3, the third from Assumption 5, the fourth from Assumption 4, the fifth from Assumption 5, the sixth from Assumption 3 and from basic probability theory, and the last from replacing the integral by an expectation and from the law of iterated expectations.

Proposition 5:

Under Assumptions 3, 4, and 5, the average total indirect effect is identified by

$$\delta^t(1) = E\left[\frac{\left\{Y - \mu(1, E\left[\frac{M \cdot (1-D)}{1 - \Pr(D=1|X)}\right], W, X)\right\} \cdot D}{\Pr(D=1|X)}\right], \delta^t(0) = E\left[\frac{\left\{\mu(0, E\left[\frac{M \cdot D}{\Pr(D=1|X)}\right], W, X) - Y\right\} \cdot (1-D)}{1 - \Pr(D=1|X)}\right]. \tag{16}$$

It is interesting to compare this result to Albert and Nelson (2011), who use a similarly strong sequential ignorability assumption and assume parametric models for D , W , M , and Y which are estimated by maximum likelihood methods. However, they do not identify $E[Y(d, M(1-d, W(1-d)), W(d))]$ and $\delta^t(d)$ without further assumptions, because the linearity restriction (5a) is not imposed therein. Therefore, recovering the total indirect effect comes at the cost of ruling out models that are nonlinear in the mediator, which are permitted in Albert and Nelson (2011).

3 Estimation

For estimation in Sections 4 and 5, we use normalized versions of the sample analogs of the IPW-based identification results in Propositions 1 to 4, such that the weights of the observations in either treatment state add up to unity, as advocated in Imbens (2004) and Busso, DiNardo, and McCrary (2009b).¹⁰ E.g., the normalized estimators of the direct effects under treatment and non-treatment are given by

$$\begin{aligned}\hat{\theta}(1) &= \frac{\sum Y_i \cdot D_i / \hat{p}(X_i)}{\sum D_i / \hat{p}(X_i)} - \frac{\sum Y_i \cdot (1 - D_i) \cdot \hat{p}(M_i, X_i) / [(1 - \hat{p}(M_i, X_i)) \cdot \hat{p}(X_i)]}{\sum (1 - D_i) \cdot \hat{p}(M_i, X_i) / [(1 - \hat{p}(M_i, X_i)) \cdot \hat{p}(X_i)]}, \\ \hat{\theta}(0) &= \frac{\sum Y_i \cdot D_i \cdot (1 - \hat{p}(M_i, X_i)) / [\hat{p}(M_i, X_i) \cdot (1 - \hat{p}(X_i))]}{\sum D_i \cdot (1 - \hat{p}(M_i, X_i)) / [\hat{p}(M_i, X_i) \cdot (1 - \hat{p}(X_i))]} - \frac{\sum Y_i \cdot (1 - D_i) / (1 - \hat{p}(X_i))}{\sum (1 - D_i) / (1 - \hat{p}(X_i))}.\end{aligned}$$

i is the index of the observations in the i.i.d. sample and $\hat{p}(M_i, X_i)$, $\hat{p}(X_i)$ denote the respective estimates of the propensity scores $\Pr(D = 1|M_i, X_i)$, $\Pr(D = 1|X_i)$, which we estimate by probit specifications. Our semiparametric IPW estimators (into which the propensity scores enter parametrically) can be expressed as sequential GMM estimators where propensity score estimation represents the first step and effect estimation the second step, see Newey (1984). It follows from his results that our methods are $\sqrt{n}$ -consistent under standard regularity conditions. Furthermore, IPW estimators are sufficiently smooth for the bootstrap being consistent for inference. We therefore estimate the standard errors in the application in Section 5 by 1999 bootstrap draws.

Concerning the estimation of the total indirect effects based on Proposition 5, μ is specified as linear function of the covariates and the mediator within each treatment state, such that two linear models (for $D = 1, 0$) have to be estimated. I.e., $\mu(d, M, X, W) = \beta_0^d + M\beta_M^d + X\beta_X^d + W\beta_W^d$, where $\beta_0^d, \beta_M^d, \beta_X^d, \beta_W^d$ represent the constant and the slope coefficients on M, X, W for $D = d$. Note that estimating two separate models for each treatment state allows for interactions of D and each of M, X, W . For inference, we again use the bootstrap in the application.

¹⁰We do not use any propensity score trimming as considered in Busso, DiNardo, and McCrary (2009a), Crump, Hotz, Imbens, and Mitnik (2009), and Huber, Lechner, and Wunsch (2012), because propensity scores close to the boundaries 0 and 1 do neither occur in our simulation study in Section 4, nor in our application in Section 5.

4 Simulations

This section presents a simulation study that provides some intuition for the identification results and the issues related to imposing the wrong set of assumptions. For ease of exposition, we consider a data generating process (DGP) which is based on linear equations:

$$Y = 0.5D + M + \beta DM + V + 0.25U + \epsilon_1, \quad (17)$$

$$M = 0.5D + 0.5V + 0.25U + \epsilon_2, \quad (18)$$

$$V = \gamma D + 0.25U + \epsilon_3, \quad (19)$$

$$D = I\{0.25U + \epsilon_4 > 0\}, \text{ with } U, \epsilon_1, \epsilon_2, \epsilon_3, \epsilon_4 \sim N(0, 1), \text{ independently of each other.} \quad (20)$$

(17) is the outcome equation, in which the observed Y is a function of the observed variables D, M, V, U and an unobserved term ϵ_1 . β gauges the interaction effect between D and M such that $\beta = 0$ satisfies the assumption of no interaction discussed in Robins (2003). By (18), the mediator is a function of D, V, U and the unobservable ϵ_2 .¹¹ The parameter γ in (19) determines whether V is caused by D and which set of assumptions is valid. If $\gamma = 0$, Assumptions 1 and 2 hold if X therein is defined as $\{V, U\}$. Otherwise, identification has to be based on Assumptions 3 to 5, with W therein corresponding to V and X corresponding to U . By (20), the treatment is conditionally independent given U with an unconditional treatment probability of 0.5. In the simulations, we set γ to 0 and 0.2 and β to 0 and 0.5. Table 1 provides the true direct and indirect effects for the various scenarios, which are explained in more detail in Appendix A.3.

We run 5000 Monte Carlo simulations with 2000 observations and estimate the models by the normalized sample analogs of either Propositions 1 and 2 (for $\gamma = 0$) or Propositions 3 to 5 (for $\gamma = 0.2$), respectively. In addition, we also consider OLS regression, which estimates the direct effect as the coefficient on D in a regression of Y on $(1, D, M, V, U)$. The indirect effect corresponds to the coefficient on M in the latter regression multiplied with the coefficient on D

¹¹Note that the potential mediator states are $M(1) = 0.5 + 0.5V + 0.25U + \epsilon_2$ and $M(0) = 0.5V + 0.25U + \epsilon_2$, the potential outcomes are $Y(1, M(d)) = 0.5 + (1 + \beta)M(d) + V + 0.25U + \epsilon_1$ and $Y(0, M(d)) = M(d) + V + 0.25U + \epsilon_1$ for $d \in \{1, 0\}$.

Table 1: True direct and indirect effects for various scenarios

Effect	$\gamma=0$				$\gamma=0.2$			
	$\beta=0$		$\beta=0.5$		$\beta=0$		$\beta=0.5$	
	$d=1$	$d=0$	$d=1$	$d=0$	$d=1$	$d=0$	$d=1$	$d=0$
$\theta(d)$	0.5	0.5	0.75	0.5	0.5	0.5	0.8	0.5
$\delta(d)$	0.5	0.5	0.75	0.5	-	-	-	-
$\delta^t(d)$	-	-	-	-	0.6	0.6	0.9	0.6
$\delta^p(d)$	-	-	-	-	0.5	0.5	0.75	0.5

in a regression of M on $(1, D, V, U)$. This approach omits interactions between D and M so that direct and indirect effects are implicitly assumed to be homogeneous w.r.t. the treatment.¹² For this reason we also estimate a linear model that includes the treatment mediator interaction as modeled in (17) for $\beta \neq 0$. See for instance Imai, Keele, and Yamamoto (2010) for a definition of the direct and indirect effects in terms of regression coefficients in this somewhat more general framework. Finally, we also include a naive OLS estimator where the D - M interaction and V, U are omitted so that the confounders of the mediator/treatment are not controlled for.

Table 2: Bias, variance, and MSE of various estimators under Assumptions 1 and 2 ($\gamma = 0$)

Est.	$\beta=0$						$\beta=0.5$					
	$d=1$			$d=0$			$d=1$			$d=0$		
	bias	var	MSE	bias	var	MSE	bias	var	MSE	bias	var	MSE
$\hat{\theta}_{IPW}(d)$	-0.002	0.004	0.004	-0.002	0.004	0.004	-0.002	0.004	0.004	-0.003	0.005	0.005
$\hat{\theta}_{OLS.ia}(d)$	-0.001	0.003	0.003	-0.001	0.003	0.003	-0.001	0.004	0.004	-0.000	0.004	0.004
$\hat{\theta}_{OLS}(d)$	-0.000	0.002	0.002	-0.000	0.002	0.002	-0.125	0.002	0.018	0.125	0.002	0.018
$\hat{\theta}_{naive}(d)$	-0.124	0.004	0.020	-0.124	0.004	0.020	-0.249	0.004	0.066	0.001	0.004	0.004
$\hat{\delta}_{IPW}(d)$	0.000	0.008	0.008	0.000	0.008	0.008	0.000	0.013	0.013	0.000	0.008	0.008
$\hat{\delta}_{OLS.ia}(d)$	-0.000	0.003	0.003	-0.000	0.003	0.003	-0.001	0.006	0.006	-0.000	0.003	0.003
$\hat{\delta}_{OLS}(d)$	-0.001	0.003	0.003	-0.001	0.003	0.003	-0.126	0.004	0.020	0.124	0.004	0.020
$\hat{\delta}_{naive}(d)$	0.461	0.006	0.219	0.461	0.006	0.219	0.373	0.009	0.147	0.623	0.009	0.396

Table 2 presents the bias, variance, and mean squared error (MSE) of the various estimators for $\gamma = 0$. The IPW estimators based on Propositions 1 and 2 ($\hat{\theta}_{IPW}(d)$, $\hat{\delta}_{IPW}(d)$) are almost unbiased and their MSEs are moderate in any scenario. The same applies to OLS regression including the D - M interaction term ($\hat{\theta}_{OLS.ia}(d)$, $\hat{\delta}_{OLS.ia}(d)$), which is somewhat more efficient because it is fully

¹²This issue is often encountered in empirical work, see Section 4.1 of Imai, Keele, Tingley, and Yamamoto (2010) for a revision and discussion of the shortcomings of the standard linear framework.

parametric. For $\beta=0$, also the OLS estimators without interaction ($\hat{\theta}_{\text{OLS}}(d)$, $\hat{\delta}_{\text{OLS}}(d)$) perform decently. However, they are biased if $\beta=0.5$ due to omitting the interaction term. The naive estimators ($\hat{\theta}_{\text{naive}}(d)$, $\hat{\delta}_{\text{naive}}(d)$) are biased in general, as they disregard confounding.

Table 3: Bias, variance, and MSE of various estimators under Assumptions 3 to 5 ($\gamma = 0.2$)

Est.	$\beta=0$						$\beta=0.5$					
	$d=1$			$d=0$			$d=1$			$d=0$		
	bias	var	MSE	bias	var	MSE	bias	var	MSE	bias	var	MSE
$\hat{\theta}_{\text{IPW}}(d)$	-0.004	0.004	0.004	-0.003	0.004	0.004	-0.004	0.004	0.004	-0.004	0.006	0.006
$\hat{\theta}_{\text{OLS.ia}}(d)$	-0.001	0.003	0.003	-0.001	0.003	0.003	-0.026	0.004	0.004	0.025	0.004	0.004
$\hat{\theta}_{\text{OLS}}(d)$	-0.001	0.002	0.002	-0.001	0.002	0.002	-0.150	0.002	0.025	0.150	0.002	0.025
$\hat{\theta}_{\text{naive}}(d)$	0.026	0.004	0.005	0.026	0.004	0.005	-0.123	0.004	0.020	0.177	0.004	0.036
$\hat{\delta}^t(d)$	-0.001	0.003	0.003	-0.001	0.003	0.003	-0.001	0.006	0.006	-0.001	0.003	0.003
$\hat{\delta}_{\text{IPW}}(d)$ for $\delta^t(d)$	-0.099	0.009	0.018	0.202	0.009	0.049	0.202	0.014	0.055	0.202	0.009	0.049
$\hat{\delta}_{\text{OLS.ia}}(d)$ for $\delta^t(d)$	-0.401	0.003	0.163	-0.100	0.003	0.013	-0.151	0.006	0.029	-0.100	0.003	0.013
$\hat{\delta}_{\text{OLS}}(d)$ for $\delta^t(d)$	-0.100	0.002	0.012	-0.100	0.002	0.012	-0.275	0.003	0.079	0.025	0.003	0.004
$\hat{\delta}_{\text{naive}}(d)$ for $\delta^t(d)$	0.511	0.006	0.267	0.511	0.006	0.267	0.397	0.009	0.166	0.697	0.009	0.494
$\hat{\delta}_{\text{IPW}}^p(d)$	0.002	0.003	0.003	0.002	0.003	0.003	0.002	0.007	0.007	0.002	0.003	0.003
$\hat{\delta}_{\text{IPW}}(d)$ for $\delta^p(d)$	0.051	0.009	0.011	0.302	0.009	0.100	0.352	0.014	0.138	0.302	0.009	0.100
$\hat{\delta}_{\text{OLS.ia}}(d)$ for $\delta^p(d)$	-0.000	0.003	0.003	-0.000	0.003	0.003	-0.001	0.006	0.006	-0.000	0.003	0.003
$\hat{\delta}_{\text{OLS}}(d)$ for $\delta^p(d)$	-0.000	0.002	0.002	-0.000	0.002	0.002	-0.125	0.003	0.019	0.125	0.003	0.019
$\hat{\delta}_{\text{naive}}(d)$ for $\delta^p(d)$	0.611	0.006	0.379	0.611	0.006	0.379	0.547	0.009	0.308	0.797	0.009	0.644

Table 3 provides the results for $\gamma = 0.2$. Concerning the direct effects, we observe a similar pattern in terms of the relative performance across estimators as before. Taking a look at the total indirect effect reveals that the estimator based on Proposition 4, $\hat{\delta}^t(d)$, is almost unbiased. In contrast, when using IPW based on Proposition 2 ($\hat{\delta}_{\text{IPW}}(d)$) for estimating $\delta^t(d)$, the bias is non-negligible. Also OLS both with and without treatment-mediator interaction is generally biased under any β , because V is controlled for in the regression as if it was not affected by D . However, for $\beta=0$, $\hat{\delta}_{\text{OLS}}(d)$ still consistently estimates the partial indirect effect $\delta^p(d)$, i.e., the ceteris paribus impact going through M for V, U fixed. $\hat{\delta}_{\text{OLS.ia}}(d)$ does so even if $\beta=0.5$, because it accounts for the interaction effect. Also $\hat{\delta}_{\text{IPW}}^p(d)$, the IPW estimator of the indirect effect based on Proposition 3, performs well when estimating the partial indirect effect, while $\hat{\delta}_{\text{IPW}}(d)$ is again biased. The bias of the naive estimator is large for both the partial and total indirect effects.

In conclusion, the methods proposed in this paper perform well in the simulations if chosen

Table 4: Bias, variance, and MSE of various estimators under $\beta=0.5$ and $\gamma = 0.2$

Est.	interaction DU added to eq. (17)						interactions $DU + DV$ added to eq. (17)					
	$d=1$			$d=0$			$d=1$			$d=0$		
	bias	var	MSE	bias	var	MSE	bias	var	MSE	bias	var	MSE
$\hat{\theta}_{IPW}(d)$	-0.002	0.005	0.005	-0.003	0.006	0.006	-0.002	0.006	0.006	-0.002	0.008	0.008
$\hat{\theta}_{OLS.ia}(d)$	0.064	0.005	0.009	-0.065	0.005	0.009	0.171	0.007	0.036	-0.172	0.007	0.036
$\hat{\theta}_{OLS}(d)$	-0.125	0.003	0.019	0.125	0.003	0.019	-0.125	0.004	0.020	0.125	0.004	0.020
$\hat{\theta}_{naive}(d)$	-0.140	0.006	0.026	0.110	0.006	0.018	-0.050	0.006	0.008	0.200	0.006	0.046
$\hat{\delta}_{IPW}(d)$	0.001	0.014	0.014	0.000	0.008	0.008	0.001	0.014	0.014	0.001	0.003	0.003
$\hat{\delta}_{OLS.ia}(d)$	0.064	0.007	0.011	-0.065	0.003	0.007	0.170	0.009	0.039	-0.172	0.002	0.031
$\hat{\delta}_{OLS}(d)$	-0.125	0.003	0.019	0.125	0.003	0.019	-0.125	0.003	0.019	0.125	0.003	0.019
$\hat{\delta}_{naive}(d)$	0.456	0.010	0.218	0.706	0.010	0.509	0.319	0.008	0.110	0.569	0.008	0.332

according to the appropriate identifying assumptions. However, so does $\hat{\theta}_{OLS.ia}(d)$, $\hat{\delta}_{OLS.ia}(d)$ for $\gamma = 0$. Therefore, the question arises whether there exist scenarios in which IPW outperforms a parametric estimator that is flexible enough to allow for treatment-mediator interactions. This may generally be the case if the parametric model is misspecified. To demonstrate this, we reconsider the DGP with $\gamma = 0$ and $\beta = 0.5$, but add the interaction DU to the right hand side of equation (17) such that $\hat{\theta}_{OLS.ia}(d)$, $\hat{\delta}_{OLS.ia}(d)$ are based on a wrong outcome model. In a second set up, we additionally include the interaction DV . The results in Table 4 show that IPW is still consistent, while the bias of $\hat{\theta}_{OLS.ia}(d)$, $\hat{\delta}_{OLS.ia}(d)$ increases as the misspecification of the outcome becomes more severe. A more comprehensive investigation of the finite sample performance of IPW relative to alternative estimators is left to future research.

5 Application

We apply the estimators resulting from Propositions 1 to 5 to a welfare policy experiment with a binary treatment assignment (D) which was conducted in the mid-1990s to assess the publicly funded U.S. Job Corps program.¹³ The program targets young individuals (aged 16-24 years)

¹³Note that the compliance with the treatment assignment was not perfect. According to Schochet, Burghardt, and McConnell (2008) only 73 % of eligible individuals actually enrolled at Job Corps centers. Here, we abstract from this issue and consider the assignment as treatment variable. Strictly speaking, we therefore consider (direct and indirect) “intention to treat” effects rather than treatment effects.

who legally reside in the U.S. and come from a low-income household. It provides participants with approximately 1200 hours of vocational training and education, housing, and board over an average duration of 8 months. Participants also receive health education as well as health and dental care. Schochet, Burghardt, and Glazerman (2001) and Schochet, Burghardt, and McConnell (2008) discuss in detail the experimental design¹⁴ and the ATEs on a broad range of outcomes. Their findings suggest that Job Corps increases educational attainment, reduces criminal activity, and increases employment and earnings (at least for some years after the program).

Flores and Flores-Lagunes (2009) appear to be the first to assess the causal mechanisms of the program and find a positive direct effect on earnings after controlling for the mediator work experience which they assume to be conditionally exogenous given pre-treatment variables. Flores and Flores-Lagunes (2010) bound the indirect effects of Job Corps on employment and earnings which are mediated by the achievement of a GED, high school degree, or vocational degree as well as the direct effects based on a partial identification approach that allows for mediator endogeneity. In contrast to these studies which are concerned with labor market outcomes, we focus on the program's effects on general health. To be precise, we consider a binary health indicator (Y) evaluated 2.5 years after randomization, which is equal to one if self-assessed general health is stated to be very good and zero otherwise. In this context, employment appears to be an interesting mediator, as it is affected by Job Corps and may itself have an impact on health.

In line with this idea, Huber, Lechner, and Wunsch (2011) find that entering employment increases self-assessed mental health when investigating a sample of German welfare recipients. Furthermore, several studies in medicine and social sciences conclude that there is a negative association between unemployment and health, see for instance the surveys by Jin, Shah, and Svoboda (1997), Björklund and Eriksson (1998), and Mathers and Schofield (1998). We therefore disentangle the total health effect into a direct and an indirect component that is due to a change in the likelihood to work. If there existed a positive total effect which, however, only

¹⁴In particular, Schochet, Burghardt, and Glazerman (2001) report that the randomization of the program was successful: Of 94 observed pre-treatment covariates, only 5 were statistically significantly different across treatment groups at the 5 % level, which is what one would expect by chance.

operated through employment, this could imply that health care and health education were less decisive for general health than the human capital related interventions of Job Corps which affect employability. In this context, the analysis of causal mechanisms may help to assess the usefulness of different components of a program in place.

We define employment in the first half of the second year after randomization (i.e., half way between the treatment assignment and the measurement of the outcome) as our mediator of interest (M). I.e., $M = 1$ in case of any kind of employment and $M = 0$ otherwise. We argue that the covariates to be controlled for should include potential confounders that are measured shortly before the mediator, as they may change over time, in particular as a function of the treatment. In contrast to Flores and Flores-Lagunes (2009), we therefore do not exclusively rely on pre-treatment covariates, but also use variables measured in the year after treatment assignment, just before the assessment of the mediator. Nevertheless, we also condition on a rich set of pre-treatment variables, not only to control for mediator endogeneity, but also to avoid confounding of the treatment induced by conditioning on post-treatment variables only.

The empirical literature, see for instance Mulatu and Schooler (2002) and Llena-Nozal, Lindelboom, and Portrait (2004) among many others, suggests that socio-economic factors such as education, age, and income are strongly correlated with health while they also determine an individual's employment perspectives. As discussed in Huber, Lechner, and Wunsch (2011), similar arguments are likely to hold for the labor market history. E.g., previous jobs might have a positive or negative effect on health depending on an individual's level of stress, willingness/reluctance to work, or physical strain. Furthermore, as acknowledged in Böckerman and Ilmakunnas (2009), it appears important to condition on initial (in our case: pre-mediator) health, which allows controlling for time-constant unobservable confounders. In the data, we do not only observe initial health, but also health behavior prior to the mediator period such as alcohol and drug abuse.

We analyze the direct and indirect effects of the program separately by gender, in order to account for potential effect heterogeneity. We restrict the initial data set (14,327 youths with completed baseline survey prior to the treatment assignment) to the 4,352 females and 5,673

Table 5: Descriptives

Variable	<i>Females</i>				<i>Males</i>			
	$M = 1$	$M = 0$	diff	p-val	$M = 1$	$M = 0$	diff	p-val
<i>socio-economic factors</i>								
age at assignment	18.750	18.377	0.373	0.000	18.506	17.880	0.627	0.000
years of education at ass.	10.523	10.047	0.476	0.000	10.242	10.116	0.126	0.397
in school in yr. before ass.	0.629	0.635	-0.006	0.662	0.644	0.696	-0.052	0.000
number of own kids in household after 1st yr	0.496	0.671	-0.175	0.000	0.129	0.092	0.037	0.001
ethnicity: black*	0.509	0.587	-0.078	0.000	0.415	0.529	-0.114	0.000
ethnicity: white*	0.247	0.152	0.094	0.000	0.354	0.223	0.131	0.000
household size at ass.	4.503	4.825	-0.323	0.000	4.351	4.492	-0.141	0.011
low household income in yr before ass.**	0.313	0.382	-0.069	0.000	0.227	0.288	-0.062	0.000
high household income in yr before ass.**	0.329	0.393	-0.063	0.000	0.366	0.405	-0.040	0.003
number of times in welfare before ass.	2.079	2.326	-0.247	0.000	1.915	2.134	-0.218	0.000
food stamps in yr before ass.	0.506	0.593	-0.086	0.000	0.337	0.423	-0.086	0.000
public assistance in yr before ass.	0.251	0.283	-0.032	0.020	0.240	0.262	-0.022	0.067
in public housing 1 yr after ass.	0.164	0.231	-0.067	0.000	0.118	0.171	-0.053	0.000
transfer payments in 1st yr after ass.	0.507	0.621	-0.114	0.000	0.276	0.376	-0.099	0.000
ever arrested before ass.	0.160	0.172	-0.013	0.276	0.317	0.329	-0.012	0.357
number of arrests in 1st yr after ass.	0.071	0.056	0.015	0.123	0.242	0.340	-0.098	0.000
<i>pre-mediator labor market state</i>								
ever worked before ass.	0.132	0.183	-0.051	0.000	0.123	0.155	-0.032	0.001
worked in yr before ass.	0.707	0.480	0.227	0.000	0.728	0.526	0.202	0.000
worked in 1st yr after ass.	0.812	0.389	0.423	0.000	0.828	0.424	0.404	0.000
worked some time in months 9-12 after ass.	0.720	0.200	0.520	0.000	0.743	0.225	0.518	0.000
worked all the time in months 9-12	0.384	0.009	0.375	0.000	0.394	0.014	0.381	0.000
in training in yr. before ass.	0.017	0.017	-0.000	0.903	0.019	0.019	0.000	0.911
vocational training in months 9-12	0.208	0.257	-0.048	0.000	0.186	0.231	-0.045	0.000
academic training in months 9-12	0.323	0.411	-0.089	0.000	0.317	0.430	-0.112	0.000
<i>pre-mediator health (behavior)</i>								
health at ass. (1=very good, 4=bad)	1.721	1.762	-0.041	0.074	1.619	1.648	-0.029	0.143
health after 1 year (1=very good, 4=bad)	1.847	1.868	-0.020	0.387	1.721	1.747	-0.026	0.213
phys./emot. problems at ass.	0.055	0.056	-0.001	0.852	0.045	0.049	-0.004	0.468
phys./emot. problems 1 yr after ass.	0.157	0.140	0.017	0.115	0.126	0.114	0.012	0.200
alcohol abuse before ass.	0.555	0.460	0.094	0.000	0.654	0.553	0.101	0.000
alcohol abuse 1 yr after ass.	0.239	0.159	0.080	0.000	0.366	0.245	0.121	0.000
illegal drugs before ass.	0.004	0.004	0.000	0.998	0.005	0.006	-0.001	0.664
illegal drugs 1 yr after ass.	0.010	0.008	0.002	0.413	0.020	0.013	0.007	0.053
very good health after 2.5 yrs (Y)	0.333	0.333	0.000	0.999	0.428	0.413	0.015	0.287

Note: *: Baseline category is neither black, nor white. **: Baseline category is intermediate household income.

males for which the post-treatment variables M and Y are observed in the follow-up survey after 2.5 years. Table 5 presents descriptive evidence that the selection into the mediator is indeed selective for females and males in our evaluation sample. Individuals entering employment 1 to 1.5 years after randomization are on average slightly older, (in the case of females) more educated, (in the case of males) less often arrested, more likely to be white, less likely to receive on public housing, transfer payments, and food stamps, and living in smaller households at assignment. The association with household income is non-monotonic, whereas the number of kids is (as expected) negatively associated with female employment and positively with male employment. Concerning

the labor market history, we see a strong positive correlation between previous employment and the mediator and a negative association of the latter with being in a training activity in the year before the mediator assessment, pointing to locking-in effects. In contrast, pre-mediator health is not strongly correlated with the mediator employment. Both the differences in general health (evaluated on a scale) and the incidence of physical or emotional problems (dummy variable) are insignificant. Maybe surprisingly, alcohol abuse is higher among the working than among the non-working, while differences in illegal drug use are mostly insignificant.

In the estimation of the direct effects as well as the total and partial indirect effects based on Assumptions 3 to 5, we control for both the pre-assignment and post-assignment (but pre-mediator) values of these potential confounders in separate propensity score specifications (probit) for females and males.¹⁵ We test the models of $\Pr(D = 1|M, W, X)$ by the nonparametric specification test for propensity score models proposed by Shaikh, Simonsen, Vytlacil, and Yildiz (2009), which does not reject the specifications at any conventional significance level.¹⁶ Furthermore, we also estimate the indirect effects based on Assumptions 1 and 2, where we only condition on the pre-assignment covariates when estimating the propensity score $\Pr(D = 1|M, X)$.

Table 6: Effects on the incidence of very good general health after 2.5 years

	$\hat{\Delta}$	$\hat{\theta}(1)$	$\hat{\theta}(0)$	$\hat{\delta}^t(1)$	$\hat{\delta}^t(0)$	$\hat{\delta}^p(1)$	$\hat{\delta}^p(0)$	$\hat{\delta}(1)$	$\hat{\delta}(0)$
<i>Females</i>									
Effect	0.028	0.029	0.028	-0.000	0.002	-0.000	-0.000	0.003	-0.000
S.E.	0.014	0.016	0.016	0.001	0.002	0.000	0.001	0.004	0.001
p-value	0.044	0.068	0.071	0.953	0.261	0.827	0.830	0.352	0.770
<i>Males</i>									
Effect	0.022	0.002	0.002	0.000	0.003	0.000	0.001	0.011	0.001
S.E.	0.013	0.016	0.015	0.001	0.002	0.000	0.001	0.004	0.001
p-value	0.100	0.876	0.909	0.904	0.180	0.717	0.434	0.002	0.376

Note: Standard errors (S.E.) are estimated based on 1999 bootstrap draws.

Table 6 presents the estimated effects on females and males. The second column gives the

¹⁵The distributions of the propensity scores across treatment states and gender are provided in Appendix A.1.

¹⁶Shaikh, Simonsen, Vytlacil, and Yildiz (2009) show that $f_{\Pr(D=1|M,W,X)|D=1}(\rho|D=1) = \frac{\Pr(D=0)}{\Pr(D=1)} \frac{\rho}{1-\rho} f_{\Pr(D=1|M,W,X)|D=0}(\rho|D=0)$ for all $\rho \in (0,1)$, with $f_{\Pr(D=1|M,W,X)|D=d}(\cdot|D=d)$ being the pdf of $\Pr(D=1|M,W,X)$ conditional on $D=d$, is a testable implication of a correctly specified propensity score. Using their test based on kernel density estimation, where the bandwidth is chosen according to the Silverman (1986) rule of thumb, we obtain p-values of 0.657 and 0.431 for the propensity scores in the female and male samples, respectively. The non-rejection of the models is insensitive to using twice or half the bandwidth.

ATE, i.e., the mean difference between treated and non-treated outcomes, along with the standard error (s.e.) and the p-value. Taking a look at the females, the estimate suggests that Job Corps increases the incidence of a very good general health state by 2.8 % points.¹⁷ The direct effects under treatment (column 3) and non-treatment (column 4) are almost identical to the ATE and significant at the 10 % level. Therefore, the program appears to have a sizeable effect that is not mediated by employment. In contrast, the total and partial indirect effects based on Assumptions 1 to 3 (columns 5 to 8) as well as the indirect effects based on Assumptions 1 and 2 (columns 9 and 10) are close to zero and insignificant. Therefore, employment does not seem to mediate the effectiveness of the program in any important way. For males, the ATE amounts to 2.2 % points and is significant at the 10 % level.¹⁸ In contrast to the females, however, we do not find any sizeable direct effects, which points to effect heterogeneity w.r.t. gender. An interesting picture arises when looking at the indirect effects. While the total and partial indirect effects based on Assumptions 3 to 5 are all close to zero, estimation based on Assumptions 1 and 2 leads to partly conflicting results: $\hat{\delta}(1)$ is significantly positive,¹⁹ which is at odds with $\hat{\delta}^t(1)$, as both estimators target the same parameter (the total of indirect effects). This again demonstrates the importance of carefully considering the choice of identifying assumptions.

To check the sensitivity of our results to potential attrition bias due to restricting our sample to individuals with observed post-treatment variables, we consider the response behavior in the follow-up period to be a function of the observed variables D, W, X . This corresponds to the missing at random assumption of Rubin (1976).²⁰ The latter allows correcting for attrition bias by weighting observations in the estimation by $R/\Pr(R = 1|D, W, X)$ with R being the binary response indicator, see for instance Wooldridge (2002, 2007). We estimate the response propensity $\Pr(R = 1|D, W, X)$ using a probit model and find that controlling for attrition substantially decreases the precision of the estimates, but does not overthrow our results. We therefore con-

¹⁷The mean outcome is 0.343 among the treated and 0.315 among the non-treated such that the ATE amounts to roughly 8 to 9 % of the mean outcomes.

¹⁸The mean outcomes are 0.432 under treatment and 0.410 under non-treatment.

¹⁹In contrast, the direct effects $\hat{\theta}(1)$, $\hat{\theta}(0)$ remain virtually unchanged when conditioning on pre-treatment covariates only, both for males and females.

²⁰For a discussion of alternative forms of missingness in experiments, see Huber (2012).

clude that for our sample of disadvantaged youths in the U.S., the health effects mediated by employment appear to be negligible. In contrast, our estimates point to a considerable direct effect of the program on the subjective health state of females.


6 Conclusion

This paper has demonstrated how to identify average direct and indirect effects of a binary treatment under selection on observables, (mainly) based on inverse probability weighting (IPW) using the treatment propensity score. Identification relies on the assumption of conditional exogeneity of the treatment as well as the mediator (the intermediate outcome of interest through which the indirect effect operates) given observed variables. We have considered two sets of assumptions: Mediator exogeneity (i) given the treatment and covariates which are not influenced by the treatment (with the leading case being pre-treatment variables) and (ii) given the treatment and covariates which are themselves a function of the treatment. It has been shown that direct effects can be straightforwardly identified in either case, whereas the identification of indirect effects becomes more cumbersome in the latter case, which, however, appears more realistic in empirical applications. The identification issues for either set of assumptions have been demonstrated in a simulation study. Finally, we have provided an application to the experimental evaluation study of the Job Corps program. As the results are partly sensitive to the choice of assumptions, the importance of carefully considering the plausibility of the imposed identifying restrictions in the analysis of causal mechanisms cannot be overemphasized.

A Appendix

A.1 Propensity score distributions

Figure 3: Distribution of the propensity scores across treatment states and gender



Note: Kernel density estimation of the estimated $\Pr(D = 1|M, W, X)$ is based on the Gaussian kernel and the Silverman (1986) rule of thumb for bandwidth selection.

A.2 Proofs of Propositions 3 and 4

The identification of the direct effect $\theta(d)$ under Assumptions 3 and 4 requires identifying $E[Y(1 - d, M(d, W(d)), W(d))]$:

$$\begin{aligned}
& E[Y(1 - d, M(d, W(d)), W(d))] \\
&= \int \int E[Y(1 - d, m, w) | M(d, W(d)) = m, W(d) = w, X = x] dF_{M(d, W(d)), W(d) | X=x}(m, w) dF_X(x) \\
&= \int \int E[Y(1 - d, m, w) | D = 1 - d, M = m, W = w, X = x] dF_{M, W | D=d, X=x}(m, w) dF_X(x) \\
&= \int E[Y | D = 1 - d, M = m, W = w, X = x] \cdot \frac{\Pr(D = d | M, W, X)}{\Pr(D = d | X)} dF_{M, W, X}(m, w, x) \\
&= E \left[E \left[\frac{Y \cdot I\{D = 1 - d\}}{\Pr(D = 1 - d | M, W, X)} \middle| M = m, W = w, X = x \right] \cdot \frac{\Pr(D = d | M, W, X)}{\Pr(D = d | X)} \right] \\
&= E \left[\frac{Y \cdot I\{D = 1 - d\}}{\Pr(D = 1 - d | M, W, X)} \cdot \frac{\Pr(D = d | M, W, X)}{\Pr(D = d | X)} \right]. \tag{A.1}
\end{aligned}$$

The first equality follows from the law of iterated expectations and from replacing the outer expectation by an integral, the second one from Assumptions 3 and 4, the third from Bayes' theorem, the fourth from basic probability theory and from replacing integrals by expectations, and the last from the law of iterated expectations. From comparing (A.1) to (4) in the paper it becomes obvious that the identification result for the direct effect under Assumptions 3 and 4 is similar to that under Assumptions 1 and 2, which gives rise to Proposition 3.

The identification of the partial indirect effect $\delta^p(d)$ under Assumptions 3 and 4 requires identifying $E[Y(d, M(1 - d, W(d)), W(d))]$:

$$\begin{aligned}
& E[Y(d, M(1 - d, W(d)), W(d))] \\
&= \int \int \int E[Y(d, m, w) | M(1 - d, W(d)) = m, W(d) = w, X = x] dF_{M(1-d, w) | W(d)=w, X=x}(m) dF_{W(d) | X=x}(w) dF_X(x) \\
&= \int \int \int E[Y(d, m, w) | D = d, M = m, W = w, X = x] dF_{M | D=d, W=w, X=x}(m) dF_{W | D=d, X=x}(w) dF_X(x) \\
&= \int \int E[Y | D = d, M = m, W = w, X = x] \cdot \frac{\Pr(D = 1 - d | M, W, X)}{\Pr(D = 1 - d | W, X)} dF_{M | X=x, W=w}(m) \frac{\Pr(D = d | W, X)}{\Pr(D = d | X)} dF_{W, X}(w, x) \\
&= E \left[E \left[E \left[\frac{Y \cdot I\{D = d\}}{\Pr(D = d | M, W, X)} \middle| M = m, W = w, X = x \right] \cdot \frac{\Pr(D = 1 - d | M, W, X)}{\Pr(D = 1 - d | W, X)} \middle| X = x, W = w \right] \cdot \frac{\Pr(D = d | W, X)}{\Pr(D = d | X)} \right] \\
&= E \left[\frac{Y \cdot I\{D = d\}}{\Pr(D = d | M, W, X)} \cdot \frac{\Pr(D = 1 - d | M, W, X)}{\Pr(D = 1 - d | W, X)} \cdot \frac{\Pr(D = d | W, X)}{\Pr(D = d | X)} \right]. \tag{A.2}
\end{aligned}$$

The first equality follows from the law of iterated expectations and from replacing the outer expectations by integrals, the second one from Assumptions 3 and 4, the third from Bayes' theorem, the fourth from

basic probability theory and from replacing integrals by expectations, and the last from the law of iterated expectations. By simple algebra it can be shown that the partial indirect effect is therefore identified as outlined in Proposition 4.

A.3 Details about the direct and indirect effects in the simulations

We first consider $\gamma = 0$ (such that Assumptions 1 and 2 hold) and define X to be $\{V, U\}$. For $\beta = 0$, the average direct effect is homogeneous and simply corresponds to the coefficient on D , i.e., $\theta(1) = \theta(0) = 0.5$, because there is no interaction between D and M . Likewise, the average indirect effect is $\delta(1) = \delta(0) = E[M(1)] - E[M(0)] = 0.5$, because by equation (18) in the paper, $E[M(1)] = 0.5$ (as $E[X(1)] = E[X(0)] = 0$ and $\gamma = 0$) and $E[M(0)] = 0$. The ATE is the sum of both effects and thus equal to one. In the simulations, we also consider $\beta = 0.5$. This implies that the direct effect for $D = 1$ is $\theta(1) = E[0.5 + M(1) + \beta M(1) + V + 0.25U + \epsilon_1] - E[M(1) + V + 0.25U + \epsilon_1] = E[0.5 + \beta M(1)] = 0.5 + \beta E[M(1)]$. Note that 0.5 is the direct effect not accounting for the interaction with M , whereas $\beta E[M(1)] = 0.25$ gives the interaction effect. Therefore, $\theta(1) = 0.75$. Likewise, $\theta(0) = 0.5 + \beta E[M(0)] = 0.5$. Concerning the indirect effects, $\delta(1) = E[0.5D + M(1) + \beta DM(1) + V + 0.25U + \epsilon_1 | D = 1] - E[0.5D + M(0) + \beta DM(0) + V + 0.25U + \epsilon_1 | D = 1] = E[(1 + \beta) \cdot (M(1) - M(0))] = 0.75$, whereas $\delta(0) = E[(M(1) - M(0))] = 0.5$. Again, $\theta(d) + \delta(1 - d)$ yields the ATE, which is now equal to 1.25.

Secondly, we set $\gamma = 0.2$ (such that Assumptions 1 and 2 are violated) and define W to be V and X to be U . This implies that $E[W(1)] = 0.2$, while $E[W(0)] = 0$. Therefore, the mean potential mediator states under treatment and non-treatment are, respectively, $E[M(1)] = E[M(1, W(1))] = 0.5 + 0.5 \cdot 0.2 = 0.6$ and $E[M(0)] = E[M(0, W(0))] = 0$. For $\beta = 0$, the average direct effect is again constant: $\theta(d) = E[0.5 + M(d, W(d)) + W(d) + 0.25X + \epsilon_1] - E[M(d, W(d)) + W(d) + 0.25X + \epsilon_1] = 0.5 + E[M(d, W(d))] + E[W(d)] - E[M(d, W(d))] - E[W(d)] = 0.5$. The average total indirect effect is $\delta^t(d) = E[M(1, W(1))] + E[W(d)] - E[M(0, W(0))] - E[W(d)] = 0.6$. Note that the ATE is equal to 1.3, i.e., 0.2 higher than the sum of the total indirect effect and the direct effect, which corresponds to the change in the outcome due to the change in W which does not go through M . Concerning the average partial indirect effect, $E[M(1, W(0))] = 0.5 + 0.5 \cdot E[W(0)] = 0.5$ and $E[M(0, W(1))] = 0.5 \cdot E[W(1)] = 0.1$. Therefore, $\delta^p(1) = E[M(1, W(1))] + E[W(1)] - E[M(0, W(1))] - E[W(1)] = 0.6 - 0.1 = 0.5$, whereas $\delta^p(0) = E[M(1, W(0))] + E[W(0)] - E[M(0, W(0))] - E[W(0)] = 0.5 - 0 = 0.5$. Finally, we consider the case that $\gamma = 0.2$ and

$\beta = 0.5$. Now, $\theta(1) = 0.5 + \beta E[M(1, W(1))] = 0.8$ and $\theta(0) = 0.5$, $\delta^t(1) = E[(1 + \beta) \cdot (M(1, W(1)) - M(0, W(0)))] = 1.5 \cdot (0.6 - 0) = 0.9$, and $\delta^t(0) = E[(M(1, W(1)) - M(0, W(0)))] = 0.6$ and $\delta^p(1) = 1.5 \cdot (E[M(1, W(1))] - E[M(0, W(1))]) = 0.75$, $\delta^p(0) = E[M(1, W(0))] - E[M(0, W(0))] = 0.5$. Again, the sum of the direct effect and the total indirect effect is 0.2 lower than the ATE (1.4).

References

- ALBERT, J. M. (2008): “Mediation analysis via potential outcomes models,” *Statistics in Medicine*, 27, 1282–1304.
- ALBERT, J. M., AND S. NELSON (2011): “Generalized causal mediation analysis,” *Biometrics*, 67, 1028–1038.
- AVIN, C., I. SHPITSER, AND J. PEARL (2005): “Identifiability of path-specific effects,” in *IJCAI-05, Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence*, pp. 357–363, Edinburgh, UK.
- BARON, R. M., AND D. A. KENNY (1986): “The Moderator-Mediator Variable Distinction in Social Psychological Research: Conceptual, Strategic, and Statistical Considerations,” *Journal of Personality and Social Psychology*, 51, 1173–1182.
- BJÖRKLUND, A., AND T. ERIKSSON (1998): “Unemployment and mental health: evidence from research in the Nordic countries,” *Scandinavian Journal of Social Welfare*, 7, 219–235.
- BÖCKERMAN, P., AND P. ILMAKUNNAS (2009): “Unemployment and self-assessed health: evidence from panel data,” *Health Economics*, 18, 161–179.
- BUSO, M., J. DIÑARDO, AND J. MCCRARY (2009a): “Finite Sample Properties of Semiparametric Estimators of Average Treatment Effects,” *unpublished manuscript*.
- (2009b): “New Evidence on the Finite Sample Properties of Propensity Score Matching and Reweighting Estimators,” *IZA Discussion Paper No. 3998*.
- CAI, Z., M. KUROKI, J. PEARL, AND J. TIAN (2008): “Bounds on Direct Effects in the Presence of Confounded Intermediate Variables,” *Biometrics*, 64, 695–701.
- CRUMP, R. K., V. J. HOTZ, G. W. IMBENS, AND O. A. MITNIK (2009): “Dealing with limited overlap in estimation of average treatment effects,” *Biometrika*, 96, 187–199.
- FLORES, C. A., AND A. FLORES-LAGUNES (2009): “Identification and Estimation of Causal Mechanisms and Net Effects of a Treatment under Unconfoundedness,” *IZA DP No. 4237*.
- (2010): “Nonparametric Partial Identification of Causal Net and Mechanism Average Treatment Effects,” *mimeo, University of Florida*.
- FRÖLICH, M. (2004): “Finite Sample Properties of Propensity-Score Matching and Weighting Estimators,” *The Review of Economics and Statistics*, 86, 77–90.

- HECKMAN, J., AND S. NAVARRO-LOZANO (2004): “Using Matching, Instrumental Variables, and Control Functions to Estimate Economic Choice Models,” *The Review of Economics and Statistics*, 86, 30–57.
- HIRANO, K., G. W. IMBENS, AND G. RIDDER (2003): “Efficient Estimation of Average Treatment Effects Using the Estimated Propensity Score,” *Econometrica*, 71, 1161–1189.
- HORVITZ, D., AND D. THOMPSON (1952): “A Generalization of Sampling without Replacement from a Finite Population,” *Journal of American Statistical Association*, 47, 663–685.
- HUBER, M. (2012): “Identification of average treatment effects in social experiments under alternative forms of attrition,” *Journal of Educational and Behavioral Statistics*, 37, 443–474.
- HUBER, M., M. LECHNER, AND C. WUNSCH (2011): “Does leaving welfare improve health? Evidence for Germany,” *Health Economics*, 20, 484–504.
- (2012): “The performance of estimators based on the propensity score,” *forthcoming in the Journal of Econometrics*.
- IMAI, K., L. KEELE, D. TINGLEY, AND T. YAMAMOTO (2010): “Unpacking the Black Box: Learning about Causal Mechanisms from Experimental and Observational Studies,” *mimeo*.
- IMAI, K., L. KEELE, AND T. YAMAMOTO (2010): “Identification, Inference and Sensitivity Analysis for Causal Mediation Effects,” *Statistical Science*, 25, 51–71.
- IMAI, K., D. TINGLEY, AND T. YAMAMOTO (2012): “Experimental Designs for Identifying Causal Mechanisms,” *forthcoming in the Journal of the Royal Statistical Society, Series A*.
- IMAI, K., AND T. YAMAMOTO (2011): “Identification and Sensitivity Analysis for Multiple Causal Mechanisms: Revisiting Evidence from Framing Experiments,” *unpublished manuscript*.
- IMBENS, G. W. (2004): “Nonparametric Estimation of Average Treatment Effects under Exogeneity: A Review,” *The Review of Economics and Statistics*, 86, 4–29.
- JIN, R. L., C. P. SHAH, AND T. J. SVOBODA (1997): “The impact of unemployment on health: a review of the evidence,” *Journal of Public Health Policy*, 18, 275–301.
- JUDD, C. M., AND D. A. KENNY (1981): “Process Analysis: Estimating Mediation in Treatment Evaluations,” *Evaluation Review*, 5, 602–619.
- KANG, J. D. Y., AND J. L. SCHAFER (2007): “Demystifying Double Robustness: A Comparison of Alternative Strategies for Estimating a Population Mean from Incomplete Data,” *Statistical Science*, 22, 523–539.
- KAUFMAN, S., J. KAUFMAN, R. MACLENOSE, S. GREENLAND, AND C. POOLE (2005): “Improved Estimation of Controlled Direct Effects in the Presence of Unmeasured Confounding of Intermediate Variables,” *Statistics in Medicine*, 24, 1683–1702.

- KHAN, S., AND E. TAMER (2010): “Irregular Identification, Support Conditions, and Inverse Weight Estimation,” *Econometrica*, 78, 2021–2042.
- LLENA-NOZAL, A., M. LINDEBOOM, AND F. PORTRAIT (2004): “The effect of work on mental health: does occupation matter?,” *Health Economics*, 13, 1045–1062.
- MACKINNON, D. P. (2008): *Introduction to Statistical Mediation Analysis*. Taylor and Francis, New York.
- MATHERS, C. D., AND D. J. SCHOFIELD (1998): “The health consequences of unemployment: the evidence,” *The Medical Journal of Australia*, 168, 178–182.
- MULATU, S., AND C. SCHOOLER (2002): “Causal Connections between Socio-Economic Status and Health: Reciprocal Effects and Mediating Mechanisms,” *Journal of Health and Social Behavior*, 43, 22–41.
- NEWBY, W. K. (1984): “A method of moments interpretation of sequential estimators,” *Economics Letters*, 14, 201–206.
- PEARL, J. (2001): “Direct and indirect effects,” in *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*, pp. 411–420, San Francisco. Morgan Kaufman.
- (2011): “The Causal Mediation Formula - A practitioner guide to the assessment of causal pathways,” *Technical report R-379*.
- PETERSEN, M. L., S. E. SINISI, AND M. J. VAN DER LAAN (2006): “Estimation of Direct Causal Effects,” *Epidemiology*, 17, 276–284.
- ROBINS, J. M. (1986): “A new approach to causal inference in mortality studies with sustained exposure periods - application to control of the healthy worker survivor effect,” *Mathematical Modelling*, 7, 1393–1512.
- (2003): “Semantics of causal DAG models and the identification of direct and indirect effects,” in *In Highly Structured Stochastic Systems*, ed. by P. Green, N. Hjort, and S. Richardson, pp. 70–81, Oxford. Oxford University Press.
- ROBINS, J. M., AND S. GREENLAND (1992): “Identifiability and Exchangeability for Direct and Indirect Effects,” *Epidemiology*, 3, 143–155.
- ROBINS, J. M., AND T. RICHARDSON (2010): “Alternative graphical causal models and the identification of direct effects,” in *Causality and Psychopathology: Finding the Determinants of Disorders and Their Cures*, ed. by P. Shrout, K. Keyes, and K. Omstein. Oxford University Press.
- ROSENBAUM, P. (1984): “The consequences of adjustment for a concomitant variable that has been affected by the treatment,” *Journal of Royal Statistical Society, Series A*, 147, 656–666.
- ROSENBAUM, P., AND D. RUBIN (1983): “The Central Role of the Propensity Score in Observational Studies for Causal Effects,” *Biometrika*, 70, 41–55.

- RUBIN, D. B. (1974): “Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies,” *Journal of Educational Psychology*, 66, 688–701.
- (1976): “Inference and Missing Data,” *Biometrika*, 63, 581–592.
- (2004): “Direct and Indirect Causal Effects via Potential Outcomes,” *Scandinavian Journal of Statistics*, 31, 161–170.
- SCHOCHET, P. Z., J. BURGHARDT, AND S. GLAZERMAN (2001): “National Job Corps Study: The Impacts of Job Corps on Participants Employment and Related Outcomes,” *Report (Washington, DC: Mathematica Policy Research, Inc.)*.
- SCHOCHET, P. Z., J. BURGHARDT, AND S. MCCONNELL (2008): “Does Job Corps Work? Impact Findings from the National Job Corps Study,” *The American Economic Review*, 98, 1864–1886.
- SHAIKH, A. M., M. SIMONSEN, E. J. VYTLACIL, AND N. YILDIZ (2009): “A specification test for the propensity score using its distribution conditional on participation,” *Journal of Econometrics*, 151, 33–46.
- SILVERMAN, B. (1986): *Density estimation for statistics and data analysis*. Chapman and Hall, London.
- SIMONSEN, M., AND L. SKIPPER (2006): “The Costs of Motherhood: An Analysis Using Matching Estimators,” *Journal of Applied Econometrics*, 21, 919–934.
- SJÖLANDER, A. (2009): “Bounds on natural direct effects in the presence of confounded intermediate variables,” *Statistics in Medicine*, 28, 558–571.
- TCHETGEN TCHETGEN, E. J., AND I. SHPITSER (2011a): “Semiparametric Estimation of Models for Natural Direct and Indirect Effects,” *Harvard University Biostatistics Working Paper 129*.
- (2011b): “Semiparametric theory for causal mediation analysis: Efficiency bounds, multiple robustness, and sensitivity analysis,” *Technical report, Harvard University School of Public Health*.
- TEN HAVE, T. R., M. M. JOFFE, K. G. LYNCH, G. K. BROWN, S. A. MAISTO, AND A. T. BECK (2007): “Causal mediation analyses with rank preserving models,” *Biometrics*, 63, 926–934.
- VANDERWEELE, T. J. (2009): “Marginal Structural Models for the Estimation of Direct and Indirect Effects,” *Epidemiology*, 20, 18–26.
- WAERNBAUM, I. (2012): “Model misspecification and robustness in causal inference: comparing matching with doubly robust estimation,” *Statistics in Medicine*, 31, 1572–1581.
- WOOLDRIDGE, J. (2002): “Inverse Probability Weighted M-Estimators for Sample Selection, Attrition and Stratification,” *Portuguese Economic Journal*, 1, 141–162.
- (2007): “Inverse probability weighted estimation for general missing data problems,” *Journal of Econometrics*, 141, 1281–1301.

ZHENG, W., AND M. J. VAN DER LAAN (2012): “Targeted Maximum Likelihood Estimation of Natural Direct Effects,” *The International Journal of Biostatistics*, 8, 1–40, Article 3.