

Assessing the Causal Effect of Policies: An Approach Based on Stochastic Interventions

Iván Díaz*

Mark J. van der Laan[†]

*Division of Biostatistics, University of California, Berkeley, USA, idiuz@jhu.edu

[†]Division of Biostatistics, University of California, Berkeley, USA, laan@berkeley.edu

This working paper is hosted by The Berkeley Electronic Press (bepress) and may not be commercially reproduced without the permission of the copyright holder.

<http://biostats.bepress.com/ucbbiostat/paper298>

Copyright ©2012 by the authors.

Assessing the Causal Effect of Policies: An Approach Based on Stochastic Interventions

Iván Díaz and Mark J. van der Laan

Abstract

Stochastic interventions are a powerful tool to define parameters that measure the causal effect of a realistic intervention that intends to alter the population distribution of an exposure. In this paper we follow the approach described in Díaz and van der Laan (2011) to define and estimate the effect of an intervention that is expected to cause a truncation in the population distribution of the exposure. The observed data parameter that identifies the causal parameter of interest is established, as well as its efficient influence function under the non parametric model. Inverse probability of treatment weighted (IPTW), augmented IPTW and targeted minimum loss based estimators (TMLE) are proposed, their consistency and efficiency properties are determined. An extension to longitudinal data structures is presented and its use is demonstrated with a real data example.

1 Introduction

Current approaches to causal inference Rubin (1974, 1978), Pearl (2000, 2009) define causal parameters as functions of the distribution of random variables generated by a system in which the stochastic nature of a set of variables is intervened on, leading to changes in the stochastic nature of the variables that depend causally on them. Such interventions may be defined in various ways: static, dynamic or stochastic. A static intervention is one in which the treatment is set to a given fixed value deterministically, while a dynamic intervention allows such value to depend on variables that precede it causally. Static interventions have also been called deterministic (Korb, Hope, Nicholson, and Axnick, 2004) or atomic (Pearl, 2000).

In spite of their wide use, deterministic interventions (whether static or dynamic) do not provide an appropriate framework to answer causal questions about phenomena that are not subject to direct intervention. Feasible interventions often interact with other factors (e.g., a medication has impact in several organs), fail to put the exposure of interest into a deterministic state (e.g., it is unrealistic to set an individuals' exercising regime according to a deterministic function), or are the result of implementing policies that target stochastic changes in the behavior of a population (e.g., the use of mass media messages advertising condom use as a means of prevention of HIV infection is a deterministic treatment at the community level that renders a stochastic one at the individual level, because each individual will react stochastically to the intervention depending upon exogenous observed or non observed factors (McAlister, 1991)).

In general (Korb et al., 2004), an intervention can be simply defined as an external manipulation of a causal system, whether that manipulation is deterministic or stochastic. A static intervention corresponds to an alteration of the causal system in which the density of the exposure is changed to a degenerate one. One can also intervene in the exposure by changing its density in any arbitrary way, which leads to a natural generalization of the counterfactual framework of (Rubin, 1978). This general approach is perhaps of more interest from a policy making standpoint: if the counterfactual distribution of the exposure reflects the expected changes induced by a hypothetical intervention policy, the intervened model contains all the information about the causal effect of the intervention in the distribution of the outcome.

Stochastic interventions also provide a new, natural way of non-parametrically defining causal parameters for any type of exposure (e.g., continuous ones), regardless of its support and dominating measure. Thus far this was only possible through the use of misspecified parametric models or the use of marginal structural models (Neugebauer and van der Laan, 2007). Some advantages of stochastic interventions with respect to marginal structural models include weakening the positivity assumption and robustness with respect to misspecification of the model for the treatment

mechanism.

Because stochastic interventions generalize static and dynamic interventions, and since several intervention policies are not representable in terms of either static or dynamic interventions, the development of methods for identification and estimation of parameters defined in terms of stochastic interventions is of main interest to the causal inference research community.

Among the few works dealing with the mathematical formalization of stochastic interventions figure (Didelez, Dawid, and Geneletti, 2006) and (Dawid and Didelez, 2010), who provide a systematic and comprehensive discussion of identification of parameters of stochastic, dynamic and static interventions, studying them from a decision-theoretic viewpoint, exploiting representations of causal systems in terms of regime indicators and influence diagrams, and presenting a parallel between their theory and existing theory for dynamic, non-stochastic regimes. (Tian, 2008) shows that the identification of sequential intervention, whether stochastic or not, can be reduced to identification of a specific set of sequential static interventions, for which there are complete identification algorithms available in the literature. It is therefore no surprise that identification of our parameter in section 2.2 requires no further assumptions than those required for identification of a static intervention.

Stochastic interventions arise in applications either inspired by a deterministic intervention, or because they are of interest in themselves. The most popular example of the former situation is given by the definition of natural direct effects, in which the effect of A on Y is confounded by W and mediated by a variable Z . If A and Z are binary, one can define the counterfactual Y_{1,Z_0} (Robins and Greenland, 1992, Pearl, 2001, Zheng and van der Laan, 2011, Hafeman and VanderWee, 2011) as the outcome under a model in which A has been set to $a = 1$ with probability one, and the distribution of Z has been changed to that of Z_0 , the latter being the counterfactual of Z obtained when A is set to $a = 0$ with probability one. This setting provides an example in which the intervention of interest is performed in two nodes, using a static intervention for A , and a stochastic intervention for Z . (Didelez et al., 2006) and (Robins and Richardson, 2010) discuss in detail the case in which several direct and indirect effects are defined and studied in the context of stochastic interventions. (Taubman, Robins, Mittleman, and Hernn, 2009) considered an intervention in the BMI defined by a truncation of the original exposure distribution, which, contrary to the truncation that we will use in this paper, relocates the mass originally located above the threshold across all the values below the threshold. As explained by (Stitelman, Hubbard, and Jewell, 2010), such intervention is usually the result of dichotomizing a continuous variable and considering a static intervention in the dichotomous version of the treatment variable. This dichotomization represents current common practice, in section 3 we will discuss the differences

with the approach presented in this paper. (Cain, Robins, Lanoy, Logan, Costagliola, and Hernán, 2010) briefly discuss a stochastic intervention in the context of comparing dynamic treatment regimes for HIV infected patients. The regimes they discuss are of the type “initiate treatment within m months after the recorded CD4 cell count first falls below x ”, and they are interested in an atomic intervention in the CD4 cell count X , and a discrete uniform $\{0, m\}$ post-intervention distribution for the number of months before treatment M . Such intervention is discussed in more detail by (Young, Cain, Robins, OReilly, and Hernn, 2011).

Among the applications in which stochastic interventions arise as an interest in themselves, (Diaz and van der Laan, 2011) considered the effect of an intervention in a population of people over 55 years of age that aimed to change the distribution of the amount of energy spent in leisure time physical activity on all cause mortality. In the present manuscript we will analyze the effect of an intervention that intends to reduce air pollution levels below a certain threshold, but allows a stochastic distribution of air pollutants below such threshold. The claims about identifiability and properties of the estimators presented in this paper are valid only for this stochastic intervention, although they can be generalized to a broader class of interventions.

Consistent and efficient estimation of statistical parameters in semi parametric models has been studied by Bickel, Klaassen, Ritov, and Wellner (1997), van der Laan and Robins (2003), Rose and van der Laan (2011), Tsiatis (2006), among others. In particular Rose and van der Laan (2011) provide a very valuable link between efficient estimation theory in semiparametric models and causal inference, empowering researchers with tools to define a causal parameter of interest, truthfully propose a model for the distribution of the data, and compute an efficient, targeted estimate of the parameter of interest under that model. By a truthful definition of the statistical model we mean that the start point is a completely non parametric model, that can only be reduced in size if real knowledge about the distribution of the data is obtained. Parametric and other assumptions often made for the sake of computational convenience are not allowed: they do not represent knowledge about the phenomena under study and therefore result in biased estimates.

In this article we use efficiency theory in semiparametric models, and in particular the targeted minimum loss based estimation road map as described by Rose and van der Laan (2011) to assess the effect of a (hypothetical) law that enforces pollution levels below a certain cutoff point.

The paper is organized as follows. In Section 2 we define the observed and counterfactual data, as well as the causal and statistical parameter and its efficient influence function. In Section 3 we discuss how this problem would be tackled with existing methods, and argue that the conclusions of such methods are

misleading. In Section 4 we present three estimators of the statistical parameter of interest: an inverse probability of treatment weighted estimator (IPTW), an augmented IPTW that solves the efficient influence curve equation, and a targeted minimum loss based estimator (TMLE). Section 5 provides an extension to longitudinal data settings and illustrates its use to measure the effect of NO_2 concentrations in the air on asthma symptoms in children between 6 and 11 years of age. Finally, Section 6 provides some concluding remarks and directions of future research.

2 Observed data, counterfactuals and parameter of interest

2.1 Causal and Statistical Models

Consider an experiment in which an exposure variable A , a continuous or binary outcome Y and a set of covariates W are measured for n randomly sampled subjects, and the outcome is measured subject to an indicator of missingness denoted by C . Let $O = (W, A, C, CY)$ represent a random variable with distribution P_0 , and O_1, \dots, O_n represent n i.i.d. observations of O . Assume that the following non parametric structural equation model (Pearl, 2000, NPSEM) holds:

$$W = f_W(U_W); \quad A = f_A(W, U_A); \quad C = f_C(A, W, U_C); \quad CY = C f_Y(A, W, U_Y), \quad (1)$$

where U_W, U_A, U_C and U_Y are exogenous random variables assumed to satisfy the randomization assumption $(U_C, U_A) \perp\!\!\!\perp U_Y | W$. The true distribution of O can be factorized as

$$P_0(O) = P_0(W)P_0(A|W)P_0(C|A, W)\{P_0(Y|A, W, C)\}^C\{I(CY = 0)\}^{1-C}, \quad (2)$$

and we denote $g_0(A|W) \equiv P_0(A|W)$, $\phi_0(A, W) \equiv P_0(C = 1|A, W)$, and $\bar{Q}_0(A, W, C) \equiv E_0(Y|A, W, C)$.

In the next subsections we will use this data structure to define a causal and statistical parameter of interest, find its efficient influence curve (Bickel et al., 1997, van der Laan and Robins, 2003), and establish the asymptotic properties of estimators that solve the efficient curve equation.

2.2 Causal and statistical parameters

Assume that the interest of the researcher relies in estimating the effect of a policy that will cause a truncation on the exposure, relocating the probability mass originally located above certain threshold δ_2 in an interval (δ_1, δ_2) , where $\delta_1 = \delta_2 - \varepsilon$

for some small ε . Formally put within the causal framework of Pearl (2000), such policy can be described by considering the modified system

$$W = f_W(U_W); \quad A_{P_\delta} = T(g_I)\{f_A(W, U_A), W\}; \quad C_{P_\delta} = 1; \quad Y_{P_\delta,1} = f_Y(A_{P_\delta}, W, U_Y), \quad (3)$$

where g_I denotes a user-given (but possibly unknown, e.g. one could set $g_I = g_0$) conditional distribution of A given W ,

$$T(g_I)(A, W) = \begin{cases} G_I^{-1}\{G_0(A)\} & \text{if } A < \delta_1 \\ G_I^{-1}\left\{\frac{G_I(A|W) - G_I(\delta_1|W)}{K(g_I)(W)} + G_I(\delta_1|W)\right\} & \text{if } A \geq \delta_1, \end{cases} \quad (4)$$

and G_I denotes the distribution function corresponding to g_I . The distribution of A_{P_δ} is given by

$$P_\delta(g_I)(A_{P_\delta} = a|W) = \begin{cases} g_I(a|W) & \text{if } a < \delta_1 \\ g_I(a|W)K(g_I)(W) & \text{if } \delta_1 \leq a \leq \delta_2 \\ 0 & \text{otherwise,} \end{cases} \quad (5)$$

where

$$K(g)(W) = \frac{1 - G\{\delta_1|W\}}{G\{\delta_1, \delta_2|W\}},$$

and in an abuse of notation $G\{\delta_1, \delta_2|W\} \equiv \int_{\delta_1 \leq a \leq \delta_2} g(a|W)d\mu(a)$. This intervention has two consequences on the distribution of the exposure: (1) it changes the distribution of values of A below δ_1 from g_0 to g_I ; and (2) it relocates the values of A above δ_1 between δ_1 and δ_2 according to distribution (5). As special case we will consider the case $g_I = g_0$, which is of particular interest when we want to assess the effect of policies that enforce the value of certain exposure below a pre-specified level. In such cases the distribution of the set of individuals that already comply with the enforced cut-off is expected to remain unchanged, making consequence (1) void.

Under the randomization assumption, the expectation of the outcome $Y_{P_\delta,1}$ is identified as a function of the observed data generating mechanism P_0 as

$$\Psi(P_0) = E(Y_{P_\delta,1}) = E_{g_I, Q_W} \left\{ \bar{Q}_0(A, W, 1) \times M(g_I)(A, W) \right\}, \quad (6)$$

where $M(g)(A, W) = I_{\delta_1}(A) + I_{\delta_1, \delta_2}(A) \times K(g)(W)$, $I_{\delta_1}(A) = I(A < \delta_1)$ and $I_{\delta_1, \delta_2}(A) = I(\delta_1 \leq A \leq \delta_2)$. This identification result follows from the following argument. The usual consistency assumption $(A = a, C = 1) \Rightarrow Y_{a,1} = Y$ implies $(A_{P_\delta} = a, C = 1) \Rightarrow Y_{P_\delta,1} = Y_{a,1}$, therefore $P(Y_{P_\delta,1} = y | A_{P_\delta} = a, C = 1, W = w) = P(Y_{a,1} = y | A_{P_\delta} = a, C = 1, W = w)$. It is easy to verify that $(U_A, U_C) \perp\!\!\!\perp U_Y | W$ implies $(A_{P_\delta}, C) \perp\!\!\!\perp Y_{a,c} | W$ for

all a, c . Thus $P(Y_{a,1} = y | A_{P_\delta} = a, C = 1, W = w) = P(Y_{a,1} = y | W = w)$, which from standard arguments for identification of static interventions (see for example Pearl (2000)) can be shown to be identified by $P(Y = y | A = a, C = 1, W = w)$. This result can also be derived using the “G-recursion” formula, presented by Dawid and Didelez (2010), which generalizes the G-computation formula for dynamic regimes (Robins, 1986). It can also be shown that the assumptions stated here are equivalent to the assumption of “simple stability” as defined by Dawid and Didelez (2010), which generalizes the (sequential) randomization assumption to the case of stochastic interventions.

The parameter in (6) is a weighted mean of $\bar{Q}_0(A, W, 1)$ (with respect to the joint distribution of A and W), in which values of $\bar{Q}_0(A, W, 1)$ for which $A < \delta_1$ receive weight one, values for which $\delta_1 \leq A \leq \delta_2$ receive weight $K(g_I)(W)$, and values for which $A > \delta_2$ receive weight 0. This makes intuitive sense; if the portion of the population whose exposure is originally above δ_2 is relocated in exposure levels in $[\delta_1, \delta_2]$, the expected outcome of individuals in $[\delta_1, \delta_2]$ should be reweighed by $K(g_I)(W)$, and the portion above δ_2 should be reweighed by zero, given that no portion of the population will fall in that region after the intervention.

As a consequence of the formal equivalence between the counterfactual and the non-parametric structural equation model frameworks (Pearl, 2000, section 7.4.4.), all the results presented in this paper can be derived under either paradigm. Furthermore, parameter (6) is a purely statistical parameter defined as the expectation of the outcome under a different distribution of A given W , and can therefore be of interest in itself, without any underlying causal assumption or interpretation. In the following subsections we deal with estimation of (6) under a non-parametric model.

2.3 Efficient Influence Curve

The efficient influence curve is a key element in semi-parametric efficient estimation, since it defines the linear approximation of any efficient and regular asymptotically linear estimator, and therefore provides an asymptotic bound for the variance of all regular asymptotically linear estimators (Bickel et al., 1997). We limit the discussion to efficient estimation of parameter (6) when $g_I = g_0$; the case of a user given function g_I is easier and can be studied using similar arguments.

Lemma 1. *The efficient influence curve of parameter (6) when $g_I = g_0$ is given by*

$$D(P_0)(O) = \frac{C}{\phi_0(A, W)} M(g_0)(A, W) \{Y - \bar{Q}_0(A, W, 1)\} \quad (7)$$

$$+ 0 \quad (8)$$

$$+ M(g_0)(A, W) \left\{ Q_0(A, W, 1) - \frac{E_{P_0}\{\bar{Q}_0(A, W, 1)I_{\delta_1, \delta_2}(A)|W\}}{G_0\{\delta_1, \delta_2|W\}} \right\} \\ + \frac{E_{P_0}\{\bar{Q}_0(A, W, 1)I_{\delta_1, \delta_2}(A)|W\}}{G_0\{\delta_1, \delta_2|W\}} - E_{P_0}\{\bar{Q}_0(A, W, 1)M(g_0)(A, W)|W\} \quad (9)$$

$$+ E_{P_0}\{\bar{Q}_0(A, W, 1)M(g_0)(A, W)|W\} - \Psi(P_0). \quad (10)$$

where the terms (7)-(10) are denoted by $D_1(P_0)$, $D_2(P_0)$, $D_3(P_0)$, and $D_4(P_0)$; respectively, and correspond to the orthogonal decomposition of the efficient influence curve implied by the factorization of the likelihood in (2).

This decomposition of the score is going to be useful later on during the construction of a targeted maximum likelihood estimator of ψ_0 , to define the correct parametric fluctuations. The following lemma provides the conditions under which an estimator that solves the efficient influence curve equation is consistent.

Lemma 2. *Let $D(O|\bar{Q}, g, \phi, \psi_0)$ be the estimating equation implied by the efficient influence function of Lemma 1:*

$$D(O|\bar{Q}, g, \phi, \psi_0) = \frac{C}{\phi(A, W)} M(g)(A, W) \{Y - \bar{Q}(A, W, 1)\} + M(g)(A, W) \times \\ \left\{ \bar{Q}(A, W, 1) - \frac{E_P\{\bar{Q}(A, W, 1)I_{\delta_1, \delta_2}(A)|W\}}{G\{\delta_1, \delta_2|W\}} \right\} + \frac{E_P\{\bar{Q}(A, W, 1)I_{\delta_1, \delta_2}(A)|W\}}{G\{\delta_1, \delta_2|W\}} - \psi_0. \quad (11)$$

We have that $E_{P_0}D(O|\bar{Q}, g, \phi, \psi_0) = 0$ if and only if $K(g) = K(g_0)$ and either $\bar{Q} = \bar{Q}_0$ or $\phi = \phi_0$.

As a consequence of Lemma 2, a substitution estimator of $\Psi(P_0)$ that solves the efficient influence curve equation will be consistent if and only if $K(g_0)$ and either \bar{Q}_0 or ϕ_0 are estimated consistently, and it will be efficient if and only if all of the estimators for $K(g_0)$, \bar{Q}_0 and ϕ_0 are consistent. The robustness of this estimating equation is then tied to robustness of the estimator for $K(g_0)$. This consistency condition on the initial estimator g_n is weaker than the conditions needed for other methods for continuous exposures (e.g., marginal structural models (Neugebauer and van der Laan, 2007)); we only need an estimator g_n that is consistent in the sense that $K(g_n) \rightarrow K(g_0)$, which is much weaker than the condition of $g_n \rightarrow g_0$ required for marginal structural models. This is because $K(g_0)$ only depends on

the conditional probabilities $G\{\delta_1|W\}$ and $G\{\delta_1, \delta_2|W\}$, which can be consistently estimated by a misspecified estimator of the density.

An additional advantage with respect to marginal structural models and other methods for continuous variables (Petersen, Porter, Gruber, Wang, and van der Laan, 2010) is given by the positivity assumption needed to identify and estimate the parameter of interest. The positivity assumption required to estimate marginal structural models is

$$\sup_{a \in \mathcal{A}} \frac{h(a)}{g_0(a|W)} < \infty, - a.e.,$$

for a user-specified weight function h . The function $h(a) = 1$ is commonly used, since it implies giving equal weights to all the possible treatment values. The positivity assumption needed to identify and estimate our parameter of interest is given by

$$G_0\{\delta_1, \delta_2|W\} > 0, - a.e.,$$

which is a condition that depends on the choice of the interval (δ_1, δ_2) and its probability under G_0 , and is thus more likely to be true than positivity of the density g_0 for all the values $a \in \mathcal{A}$.

3 Common Practice

An alternative formulation of the causal problem of assessing the effect of a truncation in the exposure, which is the current standard in applications of causal inference methods (e.g., Brotman, Klebanoff, Nansel, Andrews, Schwebke, Zhang, Yu, Zenilman, and Scharfstein (2008), Bryan, Yu, and van der Laan (2004), Joffe, Have, Feldman, and Kimmel (2004), Tager, Haight, Sternfeld, Yu, and Laan (2004)), is given by the use of a dichotomous version $A^* = I(A < \delta_2)$ of the continuous treatment variable. The effect of a truncation of A is evaluated in terms of the static intervention $A^* = 1$, and the parameter is defined as $E\{E(Y|A^* = 1, W)\}$, which corresponds (as proven by Stitelman et al. (2010)) with a stochastic intervention on A in which g_0 is changed to

$$P_{\delta}(g_0)(A_{P_{\delta}} = a|W) = \begin{cases} g_0(a|W)/G\{\delta_2|W\} & \text{if } a < \delta_2 \\ 0 & \text{otherwise,} \end{cases} \quad (12)$$

which is equal to (5) only if $G\{\delta_1|W\} = 0$. This means that $E\{E(Y|A^* = 1, W)\}$ measures the effect of a policy that will cause a truncation in the exposure, but will relocate the mass of the non-compliers (i.e., $G\{\delta_2|W\}$) across all the values below δ_2 . As a consequence, the two parameters assess policies with different hypothetical effects on the density of the exposure; it is the researcher's responsibility to judge

which option is a more likely post-intervention distribution for the policy that is being evaluated.

For instance, in section 5.2 we estimate the effect of a policy that enforces pollutant levels below a predefined threshold. Under such a policy, individuals polluting above the threshold will only have an incentive to reduce their pollution levels to a value that is in accordance with the policy, having no further incentive to go below the enforced cut-off point once they have reached it. Therefore, the most likely post intervention distribution for this policy is one that locates the probability mass associated to the non-compliers around the cut-off point, i.e., intervention (5). The use of intervention (12) in this example could lead to misleading conclusions. As an example, consider the following data generating mechanism

$$W_1 \sim U\{0, 1\}; \quad W_2 \sim Ber\{0.7\}; \quad W_3 \sim N\{W_1, .25 \exp(2W_1)\}$$

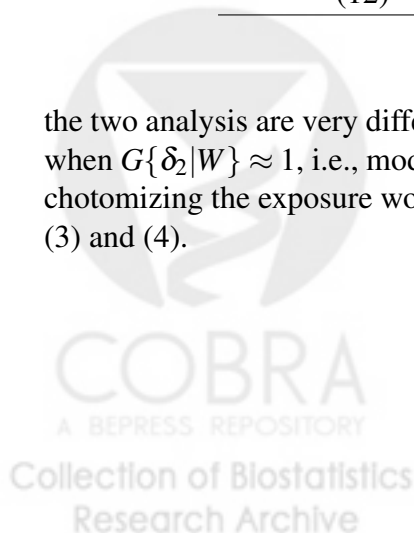
$$A \sim Beta\{S_1(W), S_2(W)\}$$

$$\bar{Q}(A, W) = \text{expit}\{1 + W_1 + 1.5A + 2AW_1 + .5AW_2 - 2W_1W_2 + .2W_1W_3\},$$

where we consider four different values for S_1 and S_2 : (1) $S_1(W) = S_2(W) = S(W)$, (2) $S_1(W) = S(W)$ and $S_2(W) = \text{expit}\{S(W)\}$, (3) $S_1(W) = \text{expit}\{S(W)\}$ and $S_2(W) = S(W)$; and (4) $S_1(W) = \text{expit}\{S(W)\}$ and $S_2(W) = \text{expit}\{S(W)\}$; for $S(W) = 2.5 + .6W_1 + .3W_2W_3 - .2W_1W_3 - .1(1 - W_2)W_3$. This four scenarios provide four different shapes of the beta distribution: (1) symmetric bell-shaped, (2) skewed to the left, (3) skewed to the right; and (4) symmetric U-shaped. For these four scenarios, table 3 shows the parameter $E(Y_{P_\delta} - Y)$ under interventions (5) and (12) for $(\delta_1, \delta_2) = (0.8, 0.9)$, providing a situation in which the conclusions obtained from

Intervention	(1)	(2)	(3)	(4)
(5)	-0.0001	-0.0001	-0.0019	-0.0009
(12)	-0.0009	-0.0002	-0.0112	-0.0111

the two analysis are very different. In this example the two effects are fairly similar when $G\{\delta_2|W\} \approx 1$, i.e., models (1) and (2). The use of the standard practice of dichotomizing the exposure would lead to misleading results, particularly for models (3) and (4).



4 Estimators

4.1 Initial Estimators

In this section we present three estimators for the parameter defined in (6). The TMLE and the A-IPTW estimators solve the efficient influence curve equation and inherit the properties derived from Lemma 2. The IPTW is inefficient, and will be consistent only if the estimator of ϕ_0 is consistent. The TMLE is expected to perform better than the A-IPTW if the positivity assumption $\sup_{a \in \mathcal{A}} \phi_0(A, W) > 0$, $-a.e.$ is violated. The finite sample properties of these estimators have been studied elsewhere (Porter, Gruber, van der Laan, and Sekhon, 2011, Diaz and van der Laan, 2011, Rose and van der Laan, 2011).

The estimators presented in this section require initial estimates of \bar{Q}_0 , g_0 and ϕ_0 , which can be obtained through machine learning techniques, parametric or semi-parametric models. The consistency of these initial estimators will determine the consistency and efficiency of the estimators of ψ_0 , as discussed previously. Parametric models are commonly used for the sole sake of their convenient analytical properties, but they encode assumptions on the distribution of the data that are not legitimate knowledge about the phenomenon under study and usually cause a large amount of bias in the estimated parameter. As an alternative, we recommend the use of machine learning techniques such as the super learner (van der Laan, Polley, and Hubbard, 2007). Super learner is a methodology that uses cross-validated risks to find an optimal estimator among a library defined by the convex hull of a user-supplied list of candidate estimators. One of its most important theoretical properties is that its solution converges to the oracle estimator (i.e., the candidate in the library that minimizes the loss function with respect to the true probability distribution). Proofs and simulations regarding these and other asymptotic properties of the super learner can be found in van der Laan, Dudoit, and Keles (2004) and van der Laan and Dudoit (2003). We will assume that g_0 is estimated consistently in the sense that $K(g_n) \rightarrow K(g_0)$.

Influence curve based variance estimators are provided for these three estimators. Consistency of the variance estimators also depends on the consistency of the initial estimators of \bar{Q}_0 , and ϕ_0 . These dependency can be avoided at the cost of computational time and effort by using bootstrapped estimates of the variance.

4.2 IPTW

Given an estimator g_n^0 of the exposure density g_0 , and an estimator ϕ_n^0 of the missing mechanism, the IPTW estimator of ψ_0 is defined as

$$\psi_{n,1} = \frac{1}{n} \sum_{i=1}^n \frac{C_i}{\phi_n^0(A_i, W_i)} M(g_n^0)(A_i, W_i) Y_i.$$

The IPTW is an asymptotically linear estimator with influence curve

$$D_{IPTW}(O|g_0, \phi_0, \psi_0) = \frac{C}{\phi_0(A, W)} M(g_0)(A, W) Y - \psi_0,$$

therefore the variable $\sqrt{n}(\psi_{n,1} - \psi_0)$ converges in distribution to $N(0, P_0 D_{IPTW}^2(g_0))$, whose variance can be estimated as the empirical variance of $D_{IPTW}^2(O|g_n^0, \phi_n^0, \psi_{n,1})$. This is a conservative estimator of the variance of the IPTW, as proven in van der Laan and Robins (2003).

4.3 Augmented IPTW

The augmented IPTW is the value $\psi_{n,2}$ that solves the equation $\sum_{i=1}^n D(O_i|\bar{Q}_n^0, g_n^0, \phi_n^0, \psi_0) = 0$, for initial estimates \bar{Q}_n^0 , g_n^0 and ϕ_n^0 of \bar{Q}_0 , g_0 and ϕ_0 .

$$\psi_{n,2} = \frac{1}{n} \sum_{i=1}^n \left[\frac{C_i}{\phi_n^0(A_i, W_i)} M(g_n^0)(A_i, W_i) \{Y_i - \bar{Q}_n^0(A_i, W_i, 1)\} + M(g_n^0)(A_i, W_i) \times \left\{ \bar{Q}_n^0(A_i, W_i, 1) - \frac{E_{g_n^0}\{\bar{Q}_n^0(A, W, 1)I_{\delta_1, \delta_2}(A)|W_i\}}{G_n^0\{\delta_1, \delta_2|W_i\}} \right\} + \frac{E_{g_n^0}\{\bar{Q}(A, W, 1)I_{\delta_1, \delta_2}(A)|W_i\}}{G_n^0\{\delta_1, \delta_2|W_i\}} \right]. \quad (13)$$

If the initial estimators are consistent, the A-IPTW is an asymptotically linear estimator with influence curve $D(O|\bar{Q}_0, g_0, \phi_0, \psi_0)$. As in the case of the IPTW, the variable $\sqrt{n}(\psi_{n,2} - \psi_0)$ converges in law to a random variable with distribution $N\{0, P_0 D^2(\cdot|\bar{Q}_0, g_0, \phi_0, \psi_0)\}$, whose variance can be estimated as the empirical variance of $D^2(O|\bar{Q}_n^0, g_n^0, \phi_n^0, \psi_{n,2})$. (Rose and van der Laan, 2011, Appendix 18) show that inference based on this variance estimator is valid only if ϕ_n^0 is consistent, providing exact inference when \bar{Q}_n^0 is consistent, and conservative inference when \bar{Q}_n^0 is inconsistent.

4.4 Targeted Minimum Loss Based Estimator

Targeted maximum likelihood estimation (van der Laan and Rubin, 2006) is a loss-based semiparametric estimation method that yields a substitution estimator of a target parameter of the probability distribution of the data that solves the efficient influence curve estimating equation, and thereby yields a double robust locally efficient estimator of the parameter of interest, under regularity conditions.

In order to define a targeted maximum likelihood estimator for ψ_0 , we need first to define three elements: (1) A loss function $L(Q)$ for the relevant part of the likelihood required to evaluate $\Psi(P)$, which in this case is $Q = (\bar{Q}, g, Q_W)$. This function must satisfy $Q_0 = \arg \min_Q E_{P_0} L(Q)(O)$, where Q_0 denotes the true value of Q ; (2) An initial estimator Q_n^0 of Q_0 ; (3) A parametric fluctuation $Q(\varepsilon)$ through Q_n^0 such that the linear span of $\frac{d}{d\varepsilon} L\{Q(\varepsilon)\}|_{\varepsilon=0}$ contains the efficient influence curve $D(P)$ defined in Lemma (1). These elements are defined below:

Loss Function

As loss function for Q , we will consider $L(Q) = L_Y(\bar{Q}) + L_A(g) + L_W(Q_W)$, where for continuous Y we set $L_Y(\bar{Q}) = \{Y - \bar{Q}(A, W, C)\}^2$, for binary Y we set $L_Y(\bar{Q}) = Y \log\{\bar{Q}(A, W, C)\} + (1 - Y) \log\{1 - \bar{Q}(A, W, C)\}$, $L_A(g) = -\log g(A|W)$, and $L_W(Q_W) = -\log Q_W(W)$. It can be easily verified that this function satisfies $Q_0 = \arg \min_Q E_{P_0} L(Q)(O)$.

Parametric Fluctuation

Given an initial estimator Q_n^k of Q_0 , with components $(\bar{Q}_n^k, g_n^k, Q_{W,n}^k)$, and an initial estimator ϕ_n^0 of ϕ_0 , we define the $(k + 1)$ th fluctuation of Q_n^k as follows:

$$\begin{aligned} m\{\bar{Q}_n^{k+1}(\varepsilon_1)(A, W)\} &= m\{\bar{Q}_n^k(A, W)\} + \varepsilon_1 H_1^k(A, W) \\ g_n^{k+1}(\varepsilon_1)(A|W) &\propto \exp\{\varepsilon_1 H_3^k(A, W)\} g_n^k(A|W) \\ Q_{W,n}^{k+1}(\varepsilon_2)(W) &\propto \exp\{\varepsilon_2 H_4^k(W)\} Q_{W,n}^k(W), \end{aligned}$$

where

$$H_1^k(A, W) = \frac{C}{\phi_n^0(A, W)} M(g_n^k)(A, W), \quad H_3^k(A, W) = D_3(P^k)(O), \quad \text{and} \quad H_4(W) = D_4(P^k)(O),$$

with D_3 and D_4 defined as in Lemma 1, and m is the identity or logit function depending on whether the outcome is continuous or binary. Note that this fluctuation satisfies the condition $D(P) \in \langle \frac{d}{d\varepsilon} L\{Q(\varepsilon)\}|_{\varepsilon=0} \rangle$, which is a key element of targeted minimum loss based estimation.

Targeted Maximum Likelihood Estimator

The TMLE is defined by the following iterative process:

1. Initialize $k = 0$.
2. Estimate ε as $\varepsilon_n^k = \arg \min_{\varepsilon} P_n L\{Q_n^k(\varepsilon)\}$.
3. Compute $Q_n^{k+1} = Q_n^k(\varepsilon_n^k)$.
4. Update $k = k + 1$ and iterate steps 2 through 4 until convergence (i.e., until $\varepsilon_n^k = 0$)

First of all, note that the value of ε_2 that minimizes the part of the loss function corresponding to the marginal distribution of W in the first step (i.e., $-P_n \log Q_{W,n}^1(\varepsilon_2)$) is $\varepsilon_2^1 = 0$. Therefore, the iterative estimation of ε only involves the estimation of ε_1 . The k th step estimation of ε_1 is obtained by numerically minimizing $P_n(L_Y(\bar{Q}_n^k(\varepsilon_1)) + L_A(g_n^k(\varepsilon_1)))$.

The TMLE of ψ_0 is defined as $\psi_{n,3} \equiv \lim_{k \rightarrow \infty} \Psi(P_n^k)$, assuming this limit exists. In practice, the iteration process is carried out until convergence in the values of ε_k is achieved, and an estimator Q_n^* is obtained. The variance of $\psi_{n,3}$ can be estimated by the empirical variance of $D^2(O|\bar{Q}_n^*, g_n^*, \phi_n^0, \psi_{n,3})$, which is a consistent estimator only if both ϕ_n^0 and \bar{Q}_n^* are consistent, is conservative if ϕ_n^0 is consistent but \bar{Q}_n^* is not, and is inconsistent in any other case.

5 Extension to Longitudinal Data and Application

5.1 Longitudinal Interventions

Assume now that the observed data structure is the same presented in Section 2, but now we have repeated measures in the sense that for each subject the observed variables were recorded at time points $t = 1, \dots, T$. That is, the observed data in this case can be described as a vector $O = (W_t, A_t, C_t, C_t Y_t : t = 1, \dots, T) = (O_t : t = 1, \dots, T)$. We can now define a time specific counterfactual outcome given by $Y_{t,P_{t,\delta}}$, where the stochastic intervention of interest is performed by changing each time-specific exposure mechanism $g_{t,0}$ to $P_{t,\delta}$, with $P_{t,\delta}$ analogous to P_δ in (5). The parameter of interest can be defined now as a causal effect based on a Marginal Structural Model (Neugebauer and van der Laan, 2007) with only intercept:

$$\beta_0 = \arg \min_{\beta} \sum_{t=1}^T \{E_0(Y_{t,P_{t,\delta}}) - m_{\beta}(t)\}^2 w(t),$$

where we set $m_{\beta}(t) = \beta$, and $w(t)$ is a weight function initially set to $1/T$. For this case (usually called intercept only model), our parameter of interest reduces to

$$\beta_0 = \sum_{t=1}^T w(t) E_0(Y_{t,P_{t,\delta}}), \quad (14)$$

which is the weighted average of the time specific causal effects. We acknowledge that this parameter does not provide a measure of the trend in the counterfactual process, and that any trend due to accumulated effects of the application of the policy will be hidden in this effect. However, this parameter does provide a measure of the overall effect of a policy when applied repetitively at every time point. Another parameter that may be of interest is given by the expectation of the counterfactual outcome at the last time point in the study, which will provide a measure of the final effect of implementing a given policy during T units of time. Parameters given by more complex marginal structural models can also be defined and estimated this way.

The efficient influence curve of parameter (14) is given by the weighted average of the time point specific influence curves:

$$D_{\beta}(O|\bar{Q}, g, \phi, \beta_0) = \sum_{t=1}^T w(t)D(O_t|\bar{Q}_t, g_t, \phi_t, \psi_{t,0}), \quad (15)$$

where D is defined in (11) and \bar{Q}_t, g_t, ϕ_t and $\psi_{t,0}$ denote the conditional expectation of the outcome, exposure mechanism, missingness mechanism and expectation of the counterfactual outcome for each time specific data structure. Estimators that solve the efficient influence curve equation

$$\sum_{i=1}^n \sum_{t=1}^T w(t)D(O_{it}|\bar{Q}_t, g_t, \phi_t, \psi_{t,0}), \quad (16)$$

inherit the consistency and efficiency properties of estimators mentioned in Lemma 2, where the consistency conditions are now replaced by consistency in the estimation of all the time specific mechanisms $\bar{Q}_{t,0}, g_{t,0}$ and $\phi_{t,0}$. To estimate each of these initial parameters we can choose to fit different estimators for each time point, or we can also choose to do smoothing over t , by including it as a covariate in each of the conditional expectations and probabilities involved.

Estimation of the parameter in (14) can now be performed by applying the estimators presented in section 4 to a pooled dataset in which time has been added as a covariate and each row corresponds to a specific subject time point combination. The IPTW estimator, for example, would now be given by

$$\psi_{n,1} = \frac{1}{nT} \sum_{i=1}^n \sum_{t=1}^T \frac{C_{it}}{\phi_{n,t}^0(A_{it}, W_{it})} M(g_{n,t}^0)(A_{it}, W_{it}) Y_{it},$$

and the augmented IPTW by

$$\psi_{n,2} = \frac{1}{nT} \sum_{i=1}^n \sum_{t=1}^T \left[\frac{C_{it}}{\phi_n^0(A_{it}, W_{it})} M(g_{n,t}^0)(A_{it}, W_{it}) \{Y_{it} - \bar{Q}_{n,t}^0(A_{it}, W_{it}, 1)\} + M(g_{n,t}^0)(A_{it}, W_{it}) \times \right. \\ \left. \left\{ \bar{Q}_{n,t}^0(A_{it}, W_{it}, 1) - \frac{E_{g_{n,t}^0} \{ \bar{Q}_{n,t}^0(A, W, 1) I_{\delta_1, \delta_2}(A) | W_{it} \}}{G_{n,t}^0 \{ I_{\delta_1, \delta_2}(A) | W_{it} \}} \right\} + \right. \\ \left. \frac{E_{g_{n,t}^0} \{ \bar{Q}(A, W, 1) I_{\delta_1, \delta_2}(A) | W_{it} \}}{G_{n,t}^0 \{ I_{\delta_1, \delta_2}(A) | W_{it} \}} \right], \quad (17)$$

which can be seen to solve equation (16). The TML estimator is defined analogous to the definition given in the previous section, with \bar{Q}_n^k , H_1^k , ϕ_n^k , H_3^k , Q_W^k , and H_4^k replaced by their t -dependent counterparts. However, the same parameters ε_1 and ε_2 are used to fluctuate all these t -dependent estimates. Estimation of ε in step 2 of the iterative process that defines the TMLE is performed now with respect to the empirical distribution P_{nT} given by the pooled dataset, and the estimating equation in Lemma (2) is replaced by its counterpart summing also over t and with t -dependent estimated values of \bar{Q}_n^k , H_1^k , ϕ_n^k , H_3^k , Q_W^k , and H_4^k . The estimators of the variance of these estimators presented in the previous section can also be adapted to these longitudinal estimators. Remarks about consistency of the variance estimators of section 4 carry on to these variance estimators.

5.2 Application

In this section we present the results of applying the method for longitudinal data described in the previous section to assess the effect of a program that constrains air pollution levels on wheezing in children with asthma. These data were originally analyzed by Mann, Balmes, Bruckner, Mortimer, Margolis, Pratt, Hammond, Lurmann, and Tager (2010) as part of the Fresno Asthmatic Childrens Environment Study (FACES). In the original paper whose objective was to evaluate whether exposure to ambient pollution is associated with increased respiratory symptoms, wheeze was found to be associated with short-terms exposures to NO_2 with an odds ratio of 1.10 (C.I. (1.02, 1.20)) for a 8.7 parts per billion increase. The data consisted of a sample of 315 children between 6 and 11 years of age who have active asthma. Reports of morning wheeze were collected for 14 days, up to three times a year, from December 2000 through March 2005, which lead to approximately 12 data panels for each child. For a comprehensive description of the study, the interested reader is referred to the original paper.

We are interested in investigating the effect of NO_2 concentrations measured 24 hours before each visit on the current presence of wheezing. The confounders

we considered (i.e., W variables) are: gender, age, race, height, low birth weight, born prematurely, atopy, presence of eczema, rhinitis, mother smoked during pregnancy, whether child was ever breastfed, presence of asthma in father and mother, no smoking policy in the house, anyone smokes in the house, relative humidity, temperature, season of the year, whether the house is rented or owned and income.

We estimated the effect of a policy that enforces NO_2 levels below 28.15 ppb. We assume that such intervention will produce a change in the population distribution of the exposure corresponding to a relocation of the probability mass originally above 28.15 ppb between 26.05 and 28.15 ppb in the intervened population. The values 28.15 and 26.05 ppb correspond with the 85th and 80th percentile of the distribution of NO_2 , respectively.

If the objective is to perform a comparison of the prevalence of wheezing in the hypothetical intervened population with the prevalence in the current population, we can define a population intervention parameter ψ_0^1 as $\psi_0^1 = \psi_0 - \mu_0$, where $\mu_0 = E_{P_0}(Y)$. This parameter compares the expectation of the outcome under the policy of interest with its current expectation, and therefore provides a measure of the gain obtained by implementing the policy.

Since we observed a coarsened version of Y , we cannot use the empirical mean as an estimator of μ_0 . Because estimation of this expectation is equivalent to estimation of the expectation of the outcome under the intervention $C = 1$, we suggest the use of the TMLE for static interventions as described in (Rose and van der Laan, 2011, chapter 4). Such estimator also utilizes initial estimators of $\phi_0(A, W)$ and $\bar{Q}_0(A, W, C)$, and is double robust under misspecification of either model. For further details about the properties and implementation the TMLE for μ_0 , the reader is referred to the original sources.

For a given estimator ψ_n of ψ_0 , and an asymptotically linear estimator μ_n of $E_{P_0}(Y)$ with influence curve $D_\mu(P)$, an asymptotically linear estimator of ψ_0^1 is given by $\psi_n^1 = \psi_n - \mu_n$. Its influence curve can be computed as $D_{\psi^1}(P)(O) = D_\psi(P)(O) - D_\mu(P)(O)$, and its variance can be estimated through the sample variance of $D_{\psi^1}(P)(O)$. Here $D_\psi(P)(O)$ represents the influence curve of each of the estimators defined in Section 4. The estimates of ψ_0^1 and their standard errors are presented in Table 1. Confidence intervals and p-values for hypothesis testing can be computed based on the normal approximations for asymptotically linear estimators described in Section 4. In light of the theoretical properties of these estimators, we rely on the TMLE and A-IPTW to measure the effect of the intervention of interest. The estimated value of ψ_n^1 means that under a policy that enforces places with NO_2 levels above 28.15 to decrease their levels to some value in the interval (26.05, 28.15), the prevalence of wheezing in children with asthma between 6 and 11 years of age would be reduced by 0.50%. However, our estimated effect

Table 1: Estimates of ψ_0^1 and ψ_0 (in %).

	TMLE	A-IPTW	IPTW
ψ_0^1	0.50 (0.40)	0.15 (0.89)	0.99 (1.04)
ψ_0	13.53 (0.56)	13.17 (0.99)	14.63 (1.14)

is not significant at a 95% confidence level, which does not mean that the effect is inexistent or epidemiologically irrelevant.

6 Conclusion

In this paper we propose a specific type of causal parameter defined by a stochastic intervention in terms of a truncation of the original distribution of the exposure. We present an application example in which the effect of a potential policy enforcing pollution levels under certain threshold is measured. Our approach allows the estimation of the effect of potential policies that result in stochastic interventions (for example because they fail to put every subject in a predefined exposure level). We argue that our parameter makes more sense from a policy- and decision-making point of view as compared to current practice.

The stochastic interventions framework allowed us to naturally define an effect for a continuous exposure, which is a topic that has received little attention in the causal inference literature. Assumptions like the positivity assumption and the consistency of an initial estimator for the exposure mechanism are weakened as compared to those required for estimating other causal parameters for continuous or categorical exposures. Two consistent and efficient estimators for the parameter of interest were proposed, and their use was illustrated with an example.

Acknowledgements

We would like to thank Dr. Ira Tager from the Division of Epidemiology at the UC Berkeley School of Public Health for kindly making available the dataset that was used in our data analysis, as well as Katia Eliseeva for her support in providing information about the dataset.

References

- P.J. Bickel, C.A.J. Klaassen, Y. Ritov, and J. Wellner. *Efficient and Adaptive Estimation for Semiparametric Models*. Springer-Verlag, 1997.
- Rebecca M. Brotman, Mark A. Klebanoff, Tonja R. Nansel, William W. Andrews, Jane R. Schwebke, Jun Zhang, Kai F. Yu, Jonathan M. Zenilman, and Daniel O. Scharfstein. A longitudinal study of vaginal douching and bacterial vaginosis marginal structural modeling analysis. *American Journal of Epidemiology*, 168(2):188–196, 2008. doi: 10.1093/aje/kwn103. URL <http://aje.oxfordjournals.org/content/168/2/188.abstract>.
- Jenny Bryan, Zhuo Yu, and Mark J. van der Laan. Analysis of longitudinal marginal structural models. *Biostatistics*, 5(3):361–380, 2004. doi: 10.1093/biostatistics/kxg041. URL <http://biostatistics.oxfordjournals.org/content/5/3/361.abstract>.
- Lauren E Cain, James M Robins, Emilie Lanoy, Roger Logan, Dominique Costagliola, and Miguel A. Hernán. When to start treatment? a systematic approach to the comparison of dynamic regimes using observational data. *The International Journal of Biostatistics*, 6, 2010. URL <http://www.bepress.com/ijb/vol6/iss2/18>.
- A. Philip Dawid and Vanessa Didelez. Identifying the consequences of dynamic treatment strategies: A decision-theoretic overview. *CoRR*, abs/1010.3425, 2010.
- Ivan Diaz and Mark van der Laan. Population intervention causal effects based on stochastic interventions. *Biometrics*, page In press., 2011. ISSN 1541-0420. doi: 10.1111/j.1541-0420.2011.01685.x. URL <http://dx.doi.org/10.1111/j.1541-0420.2011.01685.x>.
- Vanessa Didelez, A. Philip Dawid, and Sara Geneletti. Direct and indirect effects of sequential treatments. In *UAI*. AUAI Press, 2006. ISBN 0-9749039-2-2.
- Danella M. Hafeman and Tyler J. VanderWeele. Alternative assumptions for the identification of direct and indirect effects. *Epidemiology*, 22(6), 2011.
- Marshall M. Joffe, Thomas R. Ten Have, Harold I. Feldman, and Stephen E. Kimmel. Model selection, confounder control, and marginal structural models: Review and new applications. *The American Statistician*, 58(4):pp. 272–279, 2004. ISSN 00031305. URL <http://www.jstor.org/stable/27643582>.
- Kevin. Korb, Lucas. Hope, Ann. Nicholson, and Karl. Axnick. Varieties of causal intervention. In Chengqi Zhang, Hans W. Guesgen, and Wai-Kiang Yeap, editors, *PRICAI 2004: Trends in Artificial Intelligence*, volume 3157 of *Lecture Notes in Computer Science*, pages 322–331. Springer Berlin / Heidelberg, 2004.
- Jennifer K. Mann, John R. Balmes, Tim A. Bruckner, Kathleen M. Mortimer, Helene G. Margolis, Boriana Pratt, S. Katharine Hammond, Frederick W. Lurmann, and Ira B. Tager. Short-term effects of air pollution on wheeze in asthmatic

- children in fresno, california. *Environ Health Perspect*, 118(10), 06 2010. doi: 10.1289/ehp.0901292. URL <http://dx.doi.org/10.1289/ehp.0901292>.
- Alfred L. McAlister. Population behavior change: A theory-based approach. *Journal of Public Health Policy*, 12(3):pp. 345–361, 1991. ISSN 01975897. URL <http://www.jstor.org/stable/3342846>.
- Romain Neugebauer and Mark van der Laan. Nonparametric causal effects based on marginal structural models. *Journal of Statistical Planning and Inference*, 137(2):419 – 434, 2007. ISSN 0378-3758. doi: DOI: 10.1016/j.jspi.2005.12.008. URL <http://www.sciencedirect.com/science/article/pii/S0378375806000334>.
- J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge, 2000.
- Judea Pearl. Direct and indirect effects. In *Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence*, UAI '01, pages 411–420, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc. ISBN 1-55860-800-1. URL <http://dl.acm.org/citation.cfm?id=647235.720084>.
- Judea Pearl. Causal inference in statistics: An overview. *Statistics Surveys*, page 350, 2009.
- Maya L Petersen, Kristin E. Porter, Susan. Gruber, Yue. Wang, and Mark J. van der Laan. Diagnosing and responding to violations in the positivity assumption. *Stat Methods Med Res*, 2010. ISSN 1477-0334. URL <http://www.biomedsearch.com/nih/Diagnosing-responding-to-violations-in/21030422>.
- Kristin E. Porter, Susan Gruber, Mark J. van der Laan, and Jasjeet S. Sekhon. The relative performance of targeted maximum likelihood estimators. *The International Journal of Biostatistics*, 7(1):pp. 1–34, 2011.
- James Robins and Thomas Richardson. Alternative graphical causal models and the identification of direct effects. Working Paper 100, Harvard School of Public Health, 2010.
- J.M. Robins. 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, 1986.
- J.M. Robins and S. Greenland. Identifiability and exchangeability for direct and indirect effects. *Epidemiology*, 3(0):143–155, 1992.
- S. Rose and M.J. van der Laan. *Targeted Learning: Causal Inference for Observational and Experimental Data*. Springer, New York, 2011.
- D. B. Rubin. Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies. *Journal of Educational Psychology*, 1974. URL <http://www.eric.ed.gov/ERICWebPortal/detail?accno=EJ118470>.
- D.B. Rubin. Bayesian inference for causal effects: the role of randomization. *Annals of Statistics*, 6:34–58, 1978.

- Ori M Stitelman, Alan E Hubbard, and Nicholas P. Jewell. The impact of coarsening the explanatory variable of interest in making causal inferences: Implicit assumptions behind dichotomizing variables. 2010. URL <http://www.bepress.com/ucbbiostat/paper264>.
- Ira B. Tager, Thaddeus Haight, Barbara Sternfeld, Zhuo Yu, and Mark van Der Laan. Effects of physical activity and body composition on functional limitation in the elderly: Application of the marginal structural model. *Epidemiology*, 15(4):pp. 479–493, 2004. ISSN 10443983. URL <http://www.jstor.org/stable/20485932>.
- Sarah L Taubman, James M Robins, Murray A Mittleman, and Miguel A Hernn. Intervening on risk factors for coronary heart disease: an application of the parametric g-formula. *International Journal of Epidemiology*, 38(6):1599–1611, 2009. doi: 10.1093/ije/dyp192. URL <http://ije.oxfordjournals.org/content/38/6/1599.abstract>.
- Jin Tian. Identifying dynamic sequential plans. In *Proceedings of the Twenty-Fourth Conference Annual Conference on Uncertainty in Artificial Intelligence (UAI-08)*, pages 554–561, Corvallis, Oregon, 2008. AUAI Press.
- A.A. Tsiatis. Information based monitoring of clinical trials. *Statistics in Medicine*, 2006.
- M.J. van der Laan and S. Dudoit. Unified cross-validation methodology for selection among estimators and a general cross-validated adaptive epsilon-net estimator: Finite sample oracle inequalities and examples. Technical report, Division of Biostatistics, University of California, Berkeley, November 2003.
- M.J. van der Laan and J.M. Robins. *Unified methods for censored longitudinal data and causality*. Springer, New York, 2003.
- M.J. van der Laan and D. Rubin. Targeted maximum likelihood learning. *The International Journal of Biostatistics*, 2(1), 2006.
- M.J. van der Laan, S. Dudoit, and S. Keles. Asymptotic optimality of likelihood-based cross-validation. *Statistical Applications in Genetics and Molecular Biology*, 3, 2004.
- M.J. van der Laan, E. Polley, and A. Hubbard. Super learner. *Statistical Applications in Genetics and Molecular Biology*, 6(25), 2007. ISSN 1.
- Jessica Young, Lauren Cain, James Robins, Eilis O'Reilly, and Miguel Hernn. Comparative effectiveness of dynamic treatment regimes: An application of the parametric g-formula. *Statistics in Biosciences*, 3:119–143, 2011. ISSN 1867-1764. URL <http://dx.doi.org/10.1007/s12561-011-9040-7>. 10.1007/s12561-011-9040-7.
- Wenjing Zheng and Mark J. van der Laan. Targeted maximum likelihood estimation of natural direct effect. Working Paper 288 <http://www.bepress.com/ucbbiostat/paper288>, U.C. Berkeley Division of

Biostatistics Working Paper Series, 2011.

