University of California, Berkeley
U.C. Berkeley Division of Biostatistics Working Paper Series

Year 2011 *Paper* 289

Population Intervention Causal Effects Based on Stochastic Interventions

Ivan Diaz Munoz* Mark J. van der Laan†

*Division of Biostatistics, School of Public Health, University of California, Berkeley, idiazm@berkeley.edu

†Division of Biostatistics, School of Public Health, University of California, Berkeley, laan@berkeley.edu

http://biostats.bepress.com/ucbbiostat/paper289

Population Intervention Causal Effects Based on Stochastic Interventions

Ivan Diaz Munoz and Mark J. van der Laan

Abstract

Estimating the causal effect of an intervention on a population typically involves defining parameters in a nonparametric structural equation model (Pearl, 2000, NPSEM) in which the treatment or exposure is deter- ministically assigned in a static or dynamic way. We define a new causal parameter that takes into account the fact that intervention policies can result in stochastically assigned exposures. The statistical parameter that identifies the causal parameter of interest is established. Inverse probability of treatment weighting (IPTW), augmented IPTW (A-IPTW), and targeted maximum likelihood estimators (TMLE) are developed. A simulation study is performed to demonstrate the properties of these estimators, which include the double robustness of the A-IPTW and the TMLE. An application example using physical activity data is presented.

1 Introduction

Most causal inference problems are addressed by defining parameters of the distribution of the counterfactual outcome that one would obtain in a controlled experiment in which an exposure variable A is set to some pre-specified value a deterministically. A widely used example of this framework is the causal effect for a binary treatment, in which the expectation of the outcome in a hypothetical world in which everybody receives treatment is compared with its counterpart in a world in which nobody does. Other common way of addressing causal problems consists in considering parameters that reflect the difference between the distribution of a counterfactual outcome in such hypothetical intervened world and the distribution of the actual outcome; these parameters are often referred to as population intervention parameters (Hubbard and van der Laan, 2005).

In order to estimate such exposure-specific counterfactual parameters from observational data, one has to assume that all subjects in the population have a positive probability of receiving the exposure level a under consideration. This assumption is often referred to as experimental treatment assignment (ETA) and can be a highly unrealistic assumption in most cases. Additionally, when the exposure of interest is not a variable that can be directly manipulated (e.g., social or behavioral phenomena), any intervention targeting a change in the exposure distribution will result in a population whose exposure is stochastic rather than deterministic. Furthermore, in most practical cases, deterministic interventions are not feasible and their causal effect on the outcome as described in the previous paragraph loses its appeal as a measure of the gain obtained by implementing a given policy that intends to indirectly modify the exposure mechanism.

An example that illustrates these ideas is presented in Section 6. These data were collected by Tager, Hollenberg, and Satariano (1998) and analyzed by Bembom and van der Laan (2007) with the main goal of assessing the effect of vigorous physical activity on mortality in the elderly. Firstly, as argued by Bembom and van der Laan (2007), ETA assumptions as needed to identify the causal effect of a static treatment are quite unrealistic since health problems are expected to prevent an important proportion of the population from high levels of physical activity. Secondly, it is clear that it is not possible to intervene in the population in a way such that each subject is enforced to a pre-specified physical activity level, even if that level is a deterministic function of measured covariates such as health status or socioeconomic level. Therefore, any intervention on the population that targets changes in physical activity level will induce a random post-intervention exposure. These and other reasons why deterministic interventions are not always the best approach to estimate causal effects are discussed in Korb, Hope, Nicholson, and Axnick (2004) and Eberhardt and Scheines (2006). Korb et al. (2004) define an intervention on a variable A in a causal model as an action that intends to change the distribution of A. This general definition includes as special cases static and dynamic deterministic interventions (through degenerate distributions), but it also allows the definition of the causal effect in terms of a non degenerate distribution, as exploited in this article.

In our example, the question of whether higher levels of Leisure-Time Physical Activity (LTPA) cause a reduction in mortality rates in the elderly can be better addressed by considering the effect of a policy that aims to cause an increase in the mean of LTPA, possibly depending on covariates such as health status or socioeconomic level. As we will see in Section 2, this problem corresponds to considering the effect of an intervention that shifts the location of the treatment mechanism.

Despite these considerations, current developments and applications have almost exclusively focused on deterministic interventions. Among the few works using stochastic interventions figure Cain, Robins, Lanoy, Logan, Costagliola, and Hernán (2010), who used a stochastic intervention in the context

of comparing dynamic treatment regimes with a grace period; and Taubman, Robins, Mittleman, and Hernán (2009), who consider an intervention in the BMI defined by a truncation of the original exposure distribution.

In this paper, we focus the discussion on the definition and estimation of the effect of interventions that are intended to cause a shift in the conditional mean of the exposure given the covariates, such as the LTPA example. Other type of stochastic interventions of interest arises in applications in which the interest relies in estimating the effect of a policy that enforces the level of exposure below a certain threshold. Such policies can modify the distribution of the exposure in various ways. For example, if the interest relies on estimating the effect of a policy that constrains air pollution emissions below a cutoff point, it is reasonable to think that the probability mass associated with values above that cutoff in the original exposure mechanism will be relocated around the cutoff after the intervention. This is because under such a policy, high-polluting companies will not have any incentive to go below the enforced cutoff point.

Alternative threshold-like interventions can lead to a distribution of the exposure that acts like a truncation (i.e., relocating the mass across all values of the exposure distribution below the threshold). In fact, as proven by Stitelman, Hubbard, and Jewell (2010), the intervention obtained by considering a dichotomous version of a continuous treatment and defining a usual static intervention (e.g., the BMI intervention in Taubman et al. (2009)), corresponds to a stochastic intervention on the original continuous treatment that truncates the exposure below the value defining the dichotomization.

Our major goal is to introduce stochastic intervention causal parameters as a way of measuring the effect that certain policies have on the outcome of interest. As we will see, estimation of the these parameters requires weaker assumptions than estimation of other causal parameters (e.g., marginal structural models), relaxing assumptions about positivity and consistency of the initial estimators, and thus providing a more flexible way of estimating causal effects. We will start in Section 2 by defining the parameter of interest, in Section 3 we present its efficient influence curve, and discuss the double robustness of estimators that solve the efficient influence curve equation. This section also provides the tools for defining the targeted maximum likelihood estimators in Section 4.3. In Section 5, we present a simulation study, and in Section 6 we present an application example.

2 Data and Parameter of Interest

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. Let $O = (W,A,Y)$ represent a random variable with distribution P_0, and $O_1,\ldots,O_n$ represent n i.i.d. observations of O. Assume that the following NPSEM holds:

$$\begin{aligned} W &= f_W(U_W) \\ A &= f_A(W,U_A) \\ Y &= f_Y(A,W,U_Y), \end{aligned} \tag{1}$$

where U_W, U_A and U_Y are exogenous random variables such that $U_A \perp\!\!\!\perp U_Y$ holds, and either $U_W \perp\!\!\!\perp U_Y$ or $U_W \perp\!\!\!\perp U_A$ holds (randomization assumption). The true distribution P_0 of O can be factorized as

$$P_0(O) = P_0(Y|A,W)P_0(A|W)P_0(W),$$

where we denote $g_0(A|W) \equiv P_0(A|W)$, $\bar{Q}_0(A,W) \equiv E_0(Y|A,W)$, and $Q_{W,0}(W) \equiv P_0(W)$.

Counterfactual outcomes under stochastic interventions are denoted by Y_{P_δ}, and are defined as the outcome of a causal model in which the equation in the SCM (1) corresponding to A is removed, and A is set equal to a with probability $P_\delta(g_0)(A=a|W)$. The latter is called the intervention distribution, which we allow to depend on the true exposure mechanism g_0. Although any stochastic intervention of interest can be defined in this way, in this paper we focus on the discussion of the intervention distribution:

$$P_\delta(g_0)(A=a|W) = g_0(A-\delta(W)|W), \tag{2}$$

for a known function $\delta(W)$. Note that this is a shifted version of the current treatment mechanism, where the shifting value is allowed to vary across strata defined by the covariates. As discussed in Section 6, one can be interested in the effect of a policy that encourages people to exercise more, leading to a population where the distribution of physical activity is shifted according to certain health and socioeconomic variables. As implicitly stated in (2), we will assume that the functional form of the exposure mechanism induced by the intervention differs from the original exposure mechanism only through its conditional expectation given the covariates.

2.1 Identification

Let A_{P_δ} denote the exposure variable under the intervened system (i.e., A_{P_δ} is distributed according to $P_\delta(g)$). We have that

$$P(Y_{P_\delta}=y) = \sum_{a\in\mathscr{A}}\sum_{w\in\mathscr{W}} P(Y_{P_\delta}=y|A_{P_\delta}=a,W=w)P_\delta(g)(A=a|W=w)P(W=w),$$

where $\mathscr{A}$ and $\mathscr{W}$ are the support of A and W respectively. From the NPSEM (1) we have that $P(Y_{P_\delta}=y|A_{P_\delta}=a,W=w) = P(Y_a=y|A_{P_\delta}=a,W=w)$, where Y_a is the counterfactual outcome when the exposure is set to level a with probability one. Note also that the usual randomization assumption $A\perp\!\!\!\perp Y_a|W$ implies $A_{P_\delta}\perp\!\!\!\perp Y_a|W$, and therefore $P(Y_a=y|A_{P_\delta}=a,W=w) = P(Y_a=y|W=w)$. Under the consistency assumption ($A=a$ implies $Y_a=Y$) the latter quantity is identified by $P(Y=y|A=a,W=w)$. Our counterfactual distribution can be written as

$$P(Y_{P_\delta}=y) = \sum_{a\in\mathscr{A}}\sum_{w\in\mathscr{W}} P(Y=y|A=a,W=w)P_\delta(g)(A=a|W=w)P(W=w).$$

We define the parameter of interest as mapping $\Psi:\mathscr{M}\to R$ that takes any element in a statistical model $\mathscr{M}$ and maps it into a number in the reals. The true value of the parameter is given by the mapping evaluated at the true distribution $P_0\in\mathscr{M}$. The parameter of the counterfactual distribution that we are interested in estimating is $E(Y_{P_\delta})$. This parameter can be written as a function of the distribution of the observed data as

$$E(Y_{P_\delta}) = \Psi(P) = \sum_{A\in\mathscr{A}}\sum_{W\in\mathscr{W}} \bar{Q}(A,W)P_\delta(g)(A|W)Q_W(W).$$

Note that this parameter only depends on $Q=(\bar{Q},g,Q_W)$, and therefore can also be written as some mapping $\Psi^1:\mathscr{Q}\to\mathbb{R}$. In an abuse of notation, we will denote these two mappings indistinctly by Ψ. We are interested in estimating $\psi_0=\Psi(Q_0)$.

3 Efficient Influence Curve

In this section we derive the efficient influence curve for the parameter of interest presented in the previous section. 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 (see Appendix A), and therefore provides an asymptotic bound for the variance of all regular asymptotically linear estimators (Bickel, Klaassen, Ritov, and Wellner, 1997).

For the particular case of the stochastic intervention defined in (2) the parameter of interest is

$$\Psi(P) = \sum_{A\in\mathscr{A}}\sum_{W\in\mathscr{W}} \bar{Q}(A,W)g(A-\delta(W)|W)Q_W(W) = E_P\{\bar{Q}(A+\delta(W),W)\}, \tag{3}$$

evaluated at $P = P_0$. The last equality can be checked by changing the index in the summation to $A-\delta(W)$. Equation (3) corresponds exactly with computing the marginal mean of Y from the joint distribution of (W,A,Y) with A replaced by $A+\delta(W)$. Note also that if $\delta(W)=0$, equation (3) is equal to the expectation of Y.

Result 1. *The efficient influence curve of (3) is*

$$D(P)(O) = \frac{g(A-\delta(W)|W)}{g(A|W)}\{Y-\bar{Q}(A,W)\} + \bar{Q}(A+\delta(W),W) - \Psi(P). \tag{4}$$

Since this influence curve as well as the parameter of interest depend only on Q, we will use the notations $D(P)(O)$ and $D(Q)(O)$ interchangeably.

Proof. First of all, notice that the nonparametric estimator of (3) is given by

$$\begin{aligned}\hat{\Psi}(P_n) &= \sum_{y\in\mathscr{Y}}\sum_{a\in\mathscr{A}}\sum_{w\in\mathscr{W}} yP_n(y|a,w)P_n(a-\delta(w)|w)P_n(w)\\ &= \sum_{y\in\mathscr{Y}}\sum_{a\in\mathscr{A}}\sum_{w\in\mathscr{W}} y\frac{P_nf_{y,a,w}}{P_nf_{a,w}}P_nf_{a-\delta(w),w},\end{aligned} \tag{5}$$

where $P_n = \frac{1}{n}\sum_{i=1}^n \delta_{o_i}$ is the empirical measure, $f_{y,a,w} = I(Y=y,A=a,W=w)$, $f_{a,w} = I(A=a,W=w)$, $f_{a-\delta(w),w} = I(A=a-\delta(w),W=w)$, and $I(\cdot)$ denotes the indicator function. Here Pf denotes $\int fdP$.

Recall that the efficient influence curve in a non-parametric model corresponds with the influence curve of the non-parametric estimator. This is true because the influence curve of any regular estimator is also a gradient, and a non-parametric model has only one gradient. Rose and van der Laan (2011) show that if $\hat{\Psi}(P_n)$ is a substitution estimator such that $\psi_0 = \hat{\Psi}(P_0)$, and $\hat{\Psi}(P_n)$ can be written as $\hat{\Psi}^*(P_nf : f\in\mathscr{F})$ for some class of functions $\mathscr{F}$ and some mapping Ψ^*, the influence curve of $\hat{\Psi}(P_n)$ is equal to

$$IC(P_0)(O) = \sum_{f\in\mathscr{F}} \frac{d\hat{\Psi}(P_0)}{dP_0f}(f(O)-P_0f).$$

Applying this result to (9) with $\mathscr{F} = \{f_{y,a,w}, f_{a,w}, f_{a-\delta(w),w}\}$ gives the desired result. □

Note that this efficient influence curve can be decomposed in three parts corresponding to the orthogonal decomposition of the tangent space implied by the factorization of the likelihood:

$$\begin{aligned} D_1(P)(O) &= \frac{g(A-\delta(W)|W)}{g(A|W)}\{Y-\bar{Q}(A,W)\} \\ D_2(P)(O) &= \bar{Q}(A+\delta(W),W)-E_P\{\bar{Q}(A+\delta(W),W)|W\} \\ D_3(P)(O) &= E_P\{\bar{Q}(A+\delta(W),W)|W\}-\Psi(P). \end{aligned} \tag{6}$$

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

Result 2. *Let $D(O|\psi_0,\bar{Q},g)$ be the estimating function implied by the efficient influence curve $D(P)(O)$:*

$$D(O|\psi_0,\bar{Q},g)=\frac{g(A-\delta(W)|W)}{g(A|W)}\{Y-\bar{Q}(A,W)\}+\bar{Q}(A+\delta(W),W)-\psi_0,$$

let $w(g)(a,w)=g(a-\delta(w)|w)/g(a|w)$, and let $g(a|w)>0$ for all $a\in\mathscr{A}$ and $w\in\mathscr{W}$. We have that $E_{P_0}D(O|\psi_0,\bar{Q},g)=0$ if either g is such that $w(g)=w(g_0)$, or $\bar{Q}=\bar{Q}_0$

Proof. Conditioning first on (A,W) and then on W we get

$$\begin{aligned} E_{P_0}D(O|\psi_0,\bar{Q},g) = {} & E_{P_0}\left\{\sum_{a\in\mathscr{A}}\frac{g_0(a|W)}{g(a|W)}g(a-\delta(W)|W)\{\bar{Q}_0(a,W)-\bar{Q}(a,W)\}\right\} \\ & +E_{P_0}\left\{\sum_{a\in\mathscr{A}}g_0(a-\delta(W)|W)\bar{Q}(a,W)\right\}-E_{P_0}\left\{\sum_{a\in\mathscr{A}}g_0(a-\delta(W)|W)\bar{Q}_0(a,W)\right\}. \end{aligned}$$

which completes the proof. □

As a consequence of result 2, under regularity conditions stated in Theorem 1 of van der Laan and Rubin (2006), a substitution estimator of $\Psi(P_0)$ that solves the efficient influence curve equation $P_nD(\cdot|\psi_0,\bar{Q},g)$ will be consistent if either one of $w(g_0)$ and Q_0 is estimated consistently, and it will be efficient if and only if both $w(g_0)$ and Q_0 are estimated consistently. We only rely on consistent estimation of the weight function $w(g_0)$, which can be easier to obtain than consistent estimation of the density g_0, which is required for double robustness of parameters in marginal structural models (Neugebauer and van der Laan, 2007). This double robustness is a very interesting result, since $\Psi(P)$ depends on both $\bar{Q}$ and g. Intuition on this double robustness can be obtained by looking at the definition of the parameter in (3): if $\bar{Q}_0$ is known, a consistent estimator can always be obtained by computing the empirical mean of $\bar{Q}_0(A+\delta(W),W)$; on the other hand, if the weight function $w(g_0)$ is known, a consistent estimate of ψ_0 would be given by a weighted average of the outcome, with weights $w(g_0)(A,W)$.

3.1 Positivity Assumption

Alternatives to definition and estimation of causal effects in the context of continuous or categorical multilevel treatments are given by marginal structural models (MSM) and risk differences like the parameters

presented in Petersen, Porter, Gruber, Wang, and van der Laan (2010). One of the assumptions required to estimate those parameters (the positivity assumption) is given by

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

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.

From the formula of the efficient influence curve, the positivity assumption needed to identify and estimate our parameter of interest is given by

$$\sup_{a \in \mathscr{A}} \frac{g_0(a - \delta(W)|W)}{g_0(a|W)} < \infty, - a.e. \tag{7}$$

This means that if $\inf_{a \in \mathscr{A}} g_0(a|W) > \varepsilon$ for some small ε, we can always choose a function δ that while being of interest to the research problem, is never large enough to produce unstable weights. As a result, the positivity assumption as needed to estimate our parameter of interest is more easily achievable than the positivity assumption as required to estimate other causal parameter for continuous exposures.

4 Estimators

In this section we present three possible estimators for the parameter of interest. A brief review of concepts in semiparametric efficient estimation can be found in the Appendix A. The TMLE and the A-IPTW estimators solve the efficient influence curve equation, and therefore, from Result 2, are consistent estimators if either one of $Q_0(A,W)$ and $g_0(A|W)$ is estimated consistently. Also from Result 2, the TMLE and the A-IPTW are efficient if and only if both of these quantities are estimated consistently. The IPTW is inefficient, and will be consistent only if the estimator of $g_0(A|W)$ is consistent. The TMLE is expected to perform better than the A-IPTW in situations in which the positivity assumption as described in (7) is violated, which will be the case, for example, if δ takes on very large values. The TMLE is also a better alternative than the A-IPTW when the efficient estimating equation has multiple solutions, or the solution of the influence curve goes out of the natural bounds for the parameter of interest.

The estimators presented in this section require initial estimates of $Q_0(A,W)$ and $g_0(A|W)$, 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 nice analytical properties, but they encode assumptions about 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 linear combination of candidate estimators in a user-supplied library. 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).

4.1 IPTW

Given an estimator g_n^0 of the exposure density, the IPTW estimator of ψ_0 is defined as

$$\psi_{n,1} = \frac{1}{n}\sum_{i=1}^{n} \frac{g_n^0(A_i - \delta(W_i)|W_i)}{g_n^0(A_i|W_i)} Y_i.$$

The IPTW is an asymptotically linear estimator with influence curve

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

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

$$\frac{1}{n}\sum_{i=1}^{n} D_{IPTW}^2(O_i|\psi_{n,1}, g_n^0).$$

This variance estimator is conservative, as proved in van der Laan and Robins (2003) and corroborated in the simulation section.

4.2 Augmented IPTW

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

$$\psi_{n,2} = \frac{1}{n}\sum_{i=1}^{n} \frac{g_n^0(A_i - \delta|W_i)}{g_n^0(A_i|W_i)}\{Y_i - \bar{Q}_n^0(A_i, W_i)\} + \bar{Q}_n^0(A_i + \delta(W_i), W_i).$$

The A-IPTW is an asymptotically linear estimator with influence curve $D(O|\psi_0, \bar{Q}_0, g_0)$. 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 $N\{0, P_0 D^2(\cdot|\psi_0, \bar{Q}_0, g_0)\}$, whose variance can be estimated as

$$\frac{1}{n}\sum_{i=1}^{n} D^2(O_i|\psi_{n,2}, \bar{Q}_n^0, g_n^0).$$

4.3 Targeted Maximum Likelihood 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 (4).

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 $L_Y(\bar{Q}) = \{Y - \bar{Q}(A,W)\}^2$, for binary Y we set $L_Y(\bar{Q}) = Y \log\{\bar{Q}(A,W)\} + (1-Y)\log\{1-\bar{Q}(A,W)\}$, $L_A(g) = -\log g(A|W)$, 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 $(\bar{Q}_n^k, g_n^k, Q_{W,n}^k)$, we define the $(k+1)$th fluctuation of Q_n^k as follows:

$$\begin{aligned}
\bar{Q}_n^{k+1}(\varepsilon_1)(A,W) &= \bar{Q}_n^k(A,W) + \varepsilon_1 H_1^k(A,W) \\
g_n^{k+1}(\varepsilon_1)(A|W) &= \frac{\exp\{\varepsilon_1 H_2^k(A,W)\} g_n^k(A|W)}{\int_{\mathscr{A}} \exp\{\varepsilon_1 H_2^k(A,W)\} g_n^k(A|W)} \\
Q_{W,n}^{k+1}(\varepsilon_2)(W) &= \frac{\exp\{\varepsilon_2 H_3^k(W)\} Q_{W,n}^k(W)}{\int_{\mathscr{W}} \exp\{\varepsilon_2 H_3^k(W)\} Q_{W,n}^k(W)},
\end{aligned}$$

where $H_1^k(A,W) = g_n^k(A - \delta(W)|W)/g_n^k(A|W)$, $H_2^k(A,W) = D_2(P^k)(O)$ and $H_3(W) = D_3(P^k)(O)$, with D_2 and D_3 defined as in (6). We define these fluctuations using a two-dimensional ε with two different parameters ε_1 and ε_2, though it is theoretically correct to define these fluctuations using any dimension for ε, as far as the condition $D(P) \in < \frac{d}{d\varepsilon} L\{Q(\varepsilon)\}|_{\varepsilon=0} >$ is satisfied, where $< \cdot >$ denotes linear span. The convenience of the particular choice made here will be clear once the targeted maximum likelihood estimator (TMLE) is defined.

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 kth step estimation of ε_1 is obtained by minimizing $P_n(L_Y(\bar{Q}_n^k(\varepsilon_1)) + L_A(g_n^k(\varepsilon_1)))$, which implies solving the estimating equation

$$S^k(\varepsilon_1) = \sum_{i=1}^n \left[Y_i - \{\bar{Q}_n^k(A_i,W_i) + \varepsilon_1 H_1^k(O_i)\} \right] H_1^k(O_i) + D_2(P_n^k)(O_i) - \frac{\sum_{A \in \mathscr{A}} D_2(P_n^k)(O_i)\ \exp\{\varepsilon_1 D_2(P_n^k)(O_i)\}\ g_n^k(A_i|W_i)}{\sum_{A \in \mathscr{A}} \exp\{\varepsilon_1 D_2(P_n^k)(O_i)\}\ g_n^k(A_i|W_i)} \quad (8)$$

where

$$D_2(P_n^k)(O) = Q_n^k(A+\delta(W),W) - \sum_{A\in\mathscr{A}} Q_n^k(A+\delta(W_i),W_i)g_n^k(A|W_i).$$

The TMLE of ψ_0 is defined as $\psi_{n,3} \equiv \lim_{k\to\infty}\Psi(P_n^k)$, 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. Under the conditions of Theorem 2.3 of van der Laan and Robins (2003), a conservative estimator of the variance of $\psi_{n,3}$ is given by

$$\frac{1}{n}\sum_{i=1}^{n} D^2(O_i|\psi_{n,3},\bar{Q}_n^*,g_n^*).$$

5 Simulation Study

In order to assess the finite sample properties of the proposed estimators, a simulation study was performed. Consider the following data generating distribution:

$$\begin{aligned} W_1 &\sim U\{0,1\} \\ W_2 &\sim Ber\{0.7\} \\ A|W_1,W_2 &\sim Poisson\{\exp(3+.3\log(W_1)-.2\exp(W_1)W_2\} \\ Y|A,W_1,W_2 &\sim N\{1+.5A-.2AW_2+2A\tan(W_1^2)-2W_1W_2+AW_1W_2,\ 1\}. \end{aligned}$$

Assuming that we are interested in estimating the effect of a constant shift of $\delta(W_1,W_2)=2$, the true parameter value for this data generating distribution is $\psi_0 = 22.95$, and the efficiency bound equals $\{Var_{P_0}D(P_0)(O)\}^{1/2} = 17.81$.

For sample sizes $n = 50, 100, 200$ and 500, we simulated 2000 samples from the previous data generating distribution, and estimated ψ_0 using the three estimators proposed in the previous section. As initial estimators of $\bar{Q}_0(A,W)$ and $g_0(A|W)$ we considered four cases: 1) correctly specified model for both $\bar{Q}_0(A,W)$ and $g_0(A|W)$, 2) incorrectly specified model for $\bar{Q}_0(A,W)$ but correctly specified for $g_0(A|W)$, 3) correctly specified model for $\bar{Q}_0(A,W)$ but incorrectly specified for $g_0(A|W)$, and 4) incorrectly specified model for both $\bar{Q}_0(A,W)$ and $g_0(A|W)$; where misspecification of the models was performed by considering the correct distribution and link function but only main terms in the linear predictor.

TML estimation of ψ_0 was performed using the R `tmle.shift()` function presented in Appendix B. The average and variance of the estimates across the 2000 samples was computed as an approximation to the expectation and variance of the estimator (Table 1), respectively.

The results in Table 1 confirm the double robustness of the TMLE and A-IPTW, which had been proven analytically in Result 2. The TMLE and A-IPTW are unbiased even for small sample sizes, whereas the IPTW needs larger sample sizes to achieve unbiasedness.

Regarding the variance of the estimators, Table 2 shows that the IPTW estimator is inefficient, and its influence-curve-based variance estimator is very conservative.

The variances of the TMLE and A-IPTW are approximately equal to the efficiency bound if the models for $\bar{Q}_0$ and g_0 are correctly specified, although the same equality is observed if only $\bar{Q}_n^0$ is misspecified. This is because, as stated in Result 2, we only need consistent estimation of the weights $w(g_0)(A,W)$, which can be obtained through a possibly misspecified estimator of g_0. On the other hand,

Table 1: Expectation of the estimators for different sample sizes and model specifications. True value is 22.95.

n	Model	TMLE	IPTW	A-IPTW
50	1	22.99	22.66	22.99
	2	22.99	22.49	22.99
	3	22.88	22.66	22.91
	4	22.01	22.49	22.04
100	1	22.95	22.81	22.95
	2	22.96	22.61	22.95
	3	22.89	22.81	22.92
	4	21.97	22.61	22.00
200	1	22.99	22.89	22.99
	2	22.99	22.68	22.99
	3	22.94	22.89	22.96
	4	21.99	22.68	22.02
500	1	22.97	22.93	22.97
	2	22.97	22.71	22.97
	3	22.93	22.93	22.96
	4	21.97	22.71	22.00

the variance of these estimators is considerably affected by misspecification of the model for $\bar{Q}_0$ (models 3 and 4), even if g_n^0 is correctly specified. Influence-curve-based estimators of the variance seem to do a good job for these two estimators.

Since all estimators considered are asymptotically linear, 95% normal-based confidence intervals can be computed. Their coverage probabilities are presented in Table (3). The conservativeness of the IPTW can also be appreciated here. The consistent TMLE and A-IPTW based confidence intervals have perfect asymptotic coverage probability. In this simulation we do not observe significant differences between the TMLE and the A-IPTW.

6 Application

With the objective of illustrating the procedure described in the previous sections, we revisit the problem analyzed by Bembom and van der Laan (2007) of assessing the extent to which physical activity in the elderly is associated with reductions in cardiovascular morbidity and mortality, and improvement in, or prevention of metabolic abnormalities. Tager et al. (1998) followed a group of people over 55 years of age living around Sonoma, CA, over a time period of about ten years as part of a longitudinal study of physical activity and fitness (Study of Physical Performance and Age Related Changes in Sonomans - SPPARCS). The goal in analyzing the data that were collected as part of this study is to examine the effect of baseline vigorous LTPA (Leisure Time Physical Activity) on subsequent five-year all-cause mortality.

In this paper, we use the same measure of LTPA used by Bembom and van der Laan (2007), which is a continuous score based on the number of hours that the participants were engaged in vigorous

Table 2: Standard error of the estimator (times $\sqrt{n}$). Expectation of the influence curve based estimator of the variance (times $\sqrt{n}$) in parentheses. Efficiency bound is 17.81

n	Model	TMLE	IPTW	A-IPTW
50	1	17.94 (17.66)	20.33 (26.80)	17.94 (17.66)
	2	17.94 (17.67)	19.16 (25.03)	17.94 (17.66)
	3	18.92 (17.81)	20.33 (26.80)	18.94 (18.08)
	4	18.21 (18.07)	19.16 (25.03)	18.25 (17.77)
100	1	17.93 (17.74)	20.36 (27.63)	17.93 (17.74)
	2	17.93 (17.75)	19.04 (25.72)	17.93 (17.75)
	3	18.96 (18.14)	20.36 (27.63)	18.98 (18.45)
	4	18.34 (18.37)	19.04 (25.72)	18.35 (18.06)
200	1	17.77 (17.77)	20.17 (28.00)	17.77 (17.77)
	2	17.77 (17.78)	18.93 (25.97)	17.77 (17.77)
	3	18.62 (18.35)	20.17 (28.00)	18.64 (18.68)
	4	17.98 (18.57)	18.93 (25.97)	18.00 (18.24)
500	1	17.38 (17.79)	20.40 (28.37)	17.39 (17.79)
	2	17.38 (17.80)	18.94 (26.24)	17.39 (17.80)
	3	18.50 (18.49)	20.40 (28.37)	18.52 (18.84)
	4	17.74 (18.71)	18.94 (26.24)	17.76 (18.36)

physical activities such as jogging, swimming, bicycling on hills, or racquetball in the last seven days, and the standard intensity values in metabolic equivalents (MET: Metabolic Equivalent of Task) of such activities, where one MET is approximately equal to the oxygen consumption required for sitting quietly.

The primary confounding factors that we adjust for are described in Table 4. Age and gender are natural confounders, and the rest of the variables intend to account for the subject's underlying level of general health. Of the 2092 subjects enrolled in the SPPARCS study, 40 were missing information in at least on of this variables; our analysis is based on the remaining 2052 subjects.

In the sequel of this section, the vector containing the confounders will be denoted by W, the continuous MET score by A, and the indicator of five-year all-cause mortality by Y. In this paper, we are interested in estimating the effect of a policy that will produce an increase of 12 METs (corresponding, for instance, to bicycling during three hours at less than 10mph per week) in the average of the conditional distribution physical activity, given the covariates. Note that our intervention could also be defined by using different values of MET in each strata defined by the covariates W.

Initial estimators of the conditional density $g_0(A|W)$ and the conditional expectation $\bar{Q}_0(A,W)$ are presented below.

6.1 Initial estimator of g_0

For the estimation of the density $g_0(A|W)$, we consider the estimator presented in Díaz and van der Laan (2011). We now provide a summary of the rationale behind this estimator. Consider $k+1$ values $\alpha_0, \alpha_1, \ldots, \alpha_k$ spanning the range of the data and defining k bins. Now, consider the following class of

Table 3: .

n	Model	TMLE	IPTW	A-IPTW
50	1	0.93	0.97	0.93
	2	0.93	0.96	0.93
	3	0.92	0.97	0.92
	4	0.90	0.96	0.89
100	1	0.94	0.98	0.94
	2	0.94	0.98	0.94
	3	0.93	0.98	0.94
	4	0.89	0.98	0.89
200	1	0.95	0.98	0.95
	2	0.95	0.97	0.95
	3	0.94	0.98	0.95
	4	0.87	0.97	0.87
500	1	0.95	0.99	0.95
	2	0.95	0.98	0.95
	3	0.94	0.99	0.95
	4	0.78	0.98	0.78

histogram-like candidate estimators of the conditional density $g_0(A|W)$

$$g_{n,\alpha}(A=a|W) = \frac{Pr_n\{A \in [\alpha_{m-1}, \alpha_m)|W\}}{\alpha_m - \alpha_{m-1}}, \text{ for } \alpha_{m-1} \leq a < \alpha_{m-1},$$

where the choice of the location of α values and the number of bins index the candidates in the class, and the probabilities in the numerator are estimated through super learner. The final estimator of the density consists of a convex combination of these estimators found through minimization of cross-validated empirical risks. One of the most important properties of this method 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). For further reference and properties of this estimator in the context of estimation of causal effects, the reader is referred to the original paper.

As an example, Figure 1 shows an estimated density $g_n(A|W)$ for a particular profile W. As pointed out in Díaz and van der Laan (2011), we note that this methodology allows the detection of high density areas in the exposure mechanism, like the one detected at zero in Figure 1. This spike appears because this is a "zero-inflated" exposure, in which a large proportion of the population do not practice any amount of physical activity.

6.2 Initial estimator of $\bar{Q}_0$

For the initial estimator of $\bar{Q}_0$ we used the super learner (van der Laan et al., 2007). Super learner is a machine learning technique that uses cross-validation to choose a convex combination of estimators in a

Table 4: Confounders.

Variable	Description
Gender	Female Male
Age	Age in years
Health	Self-rated health status: Excellent Fair Poor
NRB	Score of self-reported physical functioning rescaled between 0 and 1
Card	Previous occurrence of any of the following cardiac events: Angina, myocardial infarction, congestive heart failure, coronary by-pass surgery, and coronary angioplasty
Chron	Presence of any of the following chronic health conditions: stroke, cancer, liver disease, kidney disease, Parkinson's disease, and diabetes mellitus
Smoking	Never smoked Current smoker Ex-smoker
Decline	Activity decline compared to 5 or 10 years earlier

library of candidate estimators. This estimator was also proven to perform asymptotically as well as the oracle selector.

Table 5 shows the candidates used, their cross-validated risks and the weight that the super learner assigns to each of them. It is worth to note that in order to get a consistent estimator of $\bar{Q}_0$ (sufficient condition for the TMLE of ψ_0 to be asymptotically unbiased), the library of candidate estimators should be as large as possible. Since this is an illustrating example, we allow ourselves to use this small library.

Table 5: Super learner output for estimation of $\bar{Q}_0$.

	Cross-validated Risk	Coef.
GLM main effects	0.1079	0.0000
GLM main eff. and two way interactions	0.1143	0.0835
GAM degree 2	0.1073	0.0000
GAM degree 3	0.1071	0.9165
Bayes GLM main effects	0.1078	0.0000

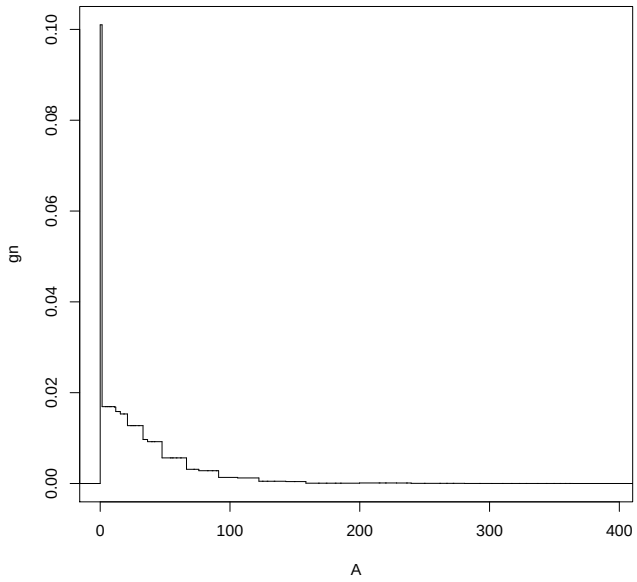

Figure 1: Estimated conditional density of A given the profile age = 73, gender = male, health = fair, nrb = 1, card = yes, smoke = never smoked, decline = yes, and chron = yes.

6.3 Estimators of ψ_0

Table 6 shows the three estimates of ψ_0 with their standard errors, as described in Section 4.

As an example, the TML estimated value of $\psi_{n,3} = 0.16$ indicates that if a policy that increases the average leisure time physical activity by the equivalent of 12 METs is implemented, the estimated risk of death in the intervened will be 16%.

If the objective is to perform a comparison with the current risk of death, we can define a population intervention parameter ψ_0^1 as

$$\psi_0^1 = \psi_0 - E_{P_0}(Y).$$

This is a parameter that compares the expected risk of death in the intervened population with the current risk of death, and therefore describes the gain obtained by carrying out the intervention of interest. For a given estimator ψ_n of ψ_0, an asymptotically linear estimator of ψ_0^1 is given by $\psi_n^1 = \psi_n - \bar{Y}$. Its influence curve can be computed as

$$D^1(P)(O) = D(P)(O) - \{Y - E_P(Y)\},$$

and its variance is estimated through the sample variance of $D^1(P)(O)$. Here $D(P)(O)$ is 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 6. Confidence intervals and p-values for hypothesis testing can be computed based on the normal approximations for asymptotically linear estimator described in Section 4 and Appendix A. In light of the results from the simulation section and the theoretical properties of the 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 if a policy increasing the average time of physical activity by the equivalent of 12 METs

Table 6: Estimates of ψ_0.

	TMLE	A-IPTW	IPTW
ψ_0	0.1600(0.0104)	0.1599(0.0105)	0.1454(0.0135)
ψ_0^1	−0.0179(0.0071)	−0.0179(0.0071)	−0.0324(0.0117)

(corresponding, for instance, to bicycling during three hours per week at less than 10mph) is put in place, the risk of all-cause mortality in the elderly would be reduced by 1.79%. These results are consistent with the findings of Bembom and van der Laan (2007).

7 Discussion

In this paper we define a new parameter for the causal effect of a population intervention that, opposed to most of the parameters presented in the literature, accounts for the fact that in most cases, even after the implementation of the intervention, the exposure continues to be a random variable. We argue that this parameter makes more intuitive sense when the objective is to assess the causal effect of policies intending to modify an exposure variable that cannot be directly intervened upon. For example, as argued in Bembom and van der Laan (2007), it makes little sense to talk about the effect of a static intervention in which every subject in a population of elderly people is required to increase their levels of physical activity to the maximum. It is well known that such intervention will never be possible due to health status and physical functioning constraints, and therefore the causal effect of of such intervention will over estimated the effect of any realistic intervention.

An alternative to overcome this issue, which deals with defining realistic individualized treatment and intention to treat rules is presented in van der Laan and Petersen (2007), and is used in Bembom and van der Laan (2007) to analyze the physical activity data used in Section 6. The choice between that alternative and the one presented in this paper depends on the type of policy for which the effect needs to be estimated. For example, assume that the exposure under study is air pollutants. In that case, we can define individualized pollution regimes for each type of factory, and design a policy intervention that enforces them by law. Under that situation, the effect of individualized deterministic treatment regimes might be more appealing as a way of measuring the effect that such intervention will have in a given outcome. However in examples like the one presented in Section 6, since no deterministic intervention is possible in practice, the causal effect of any population intervention might be better reflected by a causal parameter that takes into account the randomness of the intervention. Three estimators of the parameter were proposed, two of which are double robust to misspecification of the models for the treatment mechanism g_0 and the conditional expectation $\bar{Q}_0$, even when the parameter depends on these two quantities. This double robustness property is proven analytically, and corroborated in a simulation study.

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 Tad

Haight and Oliver Bembom for their support in providing information about the dataset.

Appendices

A Review of Efficiency Estimation in Semiparametric Models

The objective of this section is to provide an intuitive explanation of certain elements of efficient estimation in semiparametric models. We do not pretend to give a comprehensive or rigorous review, instead we intend to provide the non trained reader with the basic intuition for understanding why the methods described in the paper work. Careful and detailed definitions of the concepts described here, and rigorous proofs of most of the claims can be found in Bickel et al. (1997) and van der Vaart (1998).

A.1 Asymptotically Linear Estimators

Let $X \sim P_0 \in \mathscr{M}$, where $\mathscr{M}$ is a statistical (semi or non parametric) model, and let $\Psi : \mathscr{M} \mapsto \mathbb{R}$ be a parameter defined as a mapping that takes elements in the model and maps them into the reals (e.g., the mean $\Psi(P) = \int x dP(x)$). An estimator ψ_n of $\psi_0 = \Psi(P_0)$ is called asymptotically linear if there exist a function $IC : \mathscr{X} \times \mathscr{M} \mapsto \mathbb{R}$ such that $IC(\cdot, P_0) \in L_2(P_0)$, $\int IC(x, P_0) dP_0(x) = 0$ (Bickel et al., 1997), and

$$\psi_n - \psi_0 = \frac{1}{n} \sum_{i=1}^{n} IC(X_i, P_0) + o_P(n^{-1/2}).$$

The function IC is called the influence function of the estimator, and plays an important role in estimation and inference, since it defines the asymptotic variance of the estimator. From the central limit theorem, we conclude that if ψ_n is asymptotically linear with influence curve IC, then $\sqrt{n}(\psi_n - \psi_0) \xrightarrow{d} N\{0, P_0 IC^2(\cdot, P_0)\}$.

A.2 Efficiency

Consider a family of parametric submodels $\mathscr{M}_\varepsilon = \{P_\varepsilon : \varepsilon\} \subset \mathscr{M}$ that covers $\mathscr{M}$ and satisfies $P_{\varepsilon=0} = P_0$. A typical choice of family of parametric submodels is $\{\{p_\varepsilon(x) = [1 + \varepsilon s(x)] p_0(x) : \varepsilon\} : P_0 s = 0\}$, where each parametric submodel is indexed by a function s, which is also its score. The tangent space is defined as the closed linear span of the scores of all parametric submodels. A parameter Ψ is called pathwise differentiable if there exists a function ν such that for each submodel

$$\left. \frac{d\Psi(P_\varepsilon)}{d\varepsilon} \right|_{\varepsilon=0} = P_0 \nu s.$$

The function ν is called a gradient of the pathwise derivative. The only gradient D that is an element of the tangent space is called the efficient influence function, and corresponds with the influence function

of any regular asymptotically linear (RAL) efficient estimator, i.e., any RAL estimator whose asymptotic variance equals the efficiency bound (van der Vaart, 1998), which is the semiparametric generalization to the Cramer-Rao lower bound. The efficiency bound is then equal to $E_{P_0}D^2(X)$.

The efficient influence function has been used by several authors (Bickel et al., 1997, van der Laan and Robins, 2003, Scharfstein, Tsiatis, and Robins, 1997, van der Laan and Rubin, 2006) to construct RAL efficient estimators. The basic idea to optimally estimate the parameter of interest is to find estimators that solve the efficient influence curve equation. The properties of estimators that solve a system of equations have been extensively studied in the literature and are provided by the theory of M-estimators. Important references in M-estimation include Bickel et al. (1997), van der Laan and Robins (2003), Tsiatis (2006) and Rose and van der Laan (2011).

A.3 Proofs

Proof. Result 1. First of all, notice that the nonparametric estimator of ψ_0 is given by

$$\begin{aligned}\hat{\Psi}(P_n) &= \sum_{y\in\mathscr{Y}}\sum_{a\in\mathscr{A}}\sum_{w\in\mathscr{W}} yP_n(y|a,w)P_n(a-\delta(w)|w)P_n(w)\\ &= \sum_{y\in\mathscr{Y}}\sum_{a\in\mathscr{A}}\sum_{w\in\mathscr{W}} y\frac{P_nf_{y,a,w}}{P_nf_{a,w}}P_nf_{a-\delta(w),w}, \end{aligned} \tag{9}$$

where $P_n=\frac{1}{n}\sum_{i=1}^n\delta_{o_i}$ is the empirical measure, $f_{y,a,w}=I(Y=y,A=a,W=w)$, $f_{a,w}=I(A=a,W=w)$, $f_{a-\delta(w),w}=I(A=a-\delta(w),W=w)$, and $I(\cdot)$ denotes the indicator function. Here Pf denotes $\int fdP$.

Recall that the efficient influence curve in a non-parametric model corresponds with the influence curve of the non-parametric estimator. This is true because the influence curve of any regular estimator is also a gradient, and a non-parametric model has only one gradient. Rose and van der Laan (2011) show that if $\hat{\Psi}(P_n)$ is a substitution estimator such that $\psi_0=\hat{\Psi}(P_0)$, and $\hat{\Psi}(P_n)$ can be written as $\hat{\Psi}^*(P_nf:f\in\mathscr{F})$ for some class of functions $\mathscr{F}$ and some mapping Ψ^*, the influence curve of $\hat{\Psi}(P_n)$ is equal to

$$IC(P_0)(O)=\sum_{f\in\mathscr{F}}\frac{d\hat{\Psi}^*(P_0)}{dP_0f}\{f(O)-P_0f\}.$$

Applying this result to (9) with $\mathscr{F}=\{f_{y,a,w},f_{a,w},f_{a-\delta(w),w}\}$ gives the desired result. □

Proof. Result 2. Conditioning first on (A,W) and then on W we get

$$\begin{aligned}E_{P_0}D(O|\psi_0,\bar{Q},g) = E_{P_0}&\left[\sum_{a\in\mathscr{A}}\frac{g_0(a|W)}{g(a|W)}g(a-\delta(W)|W)\{\bar{Q}_0(a,W)-\bar{Q}(a,W)\}\right]\\ &+E_{P_0}\left[\sum_{a\in\mathscr{A}}g_0(a-\delta(W)|W)\bar{Q}(a,W)\right]-E_{P_0}\left[\sum_{a\in\mathscr{A}}g_0(a-\delta(W)|W)\bar{Q}_0(a,W)\right],\end{aligned}$$

which completes the proof. □

B R function `tmle.shift()`

B.1 Arguments

Argument	Description
`Y`	Outcome vector.
`A`	Treatment vector.
`W`	Covariates matrix.
`Qn`	An initial estimator of $\bar{Q}_0$ in the form of a function that takes a vector **A** and a matrix **W** and returns the vector of conditional expectations of Y given **A** and **W**.
`gn`	An initial estimator g_0 that takes as input a vector **A** and a matrix **W** and returns the density of A conditional on W at points **A**.
`delta`	A function of W defining the parameter of interest.
`tol`	Tolerance value for the convergence of ε.
`max.iter`	Maximum of iterations allowed.
`Aval`	A vector with equally spaced values indicating a partition of the support of A over which to compute Riemann sums to approximate the integrals involved in the estimation process.

Table 7: Arguments of the R function `tmle.shift`

B.2 Code

```
tmle.shift <- function(Y, A, W, Qn, gn, delta, tol = 1e-5, iter.max = 5, Aval){
  # interval partition length
  h.int <- Aval[3]-Aval[2]
  # this function takes as input initial estimator of Q and g and returns
  # their updated value
  f.iter <- function(Qn, gn, gn0d = NULL, prev.sum = 0, first = FALSE){
      # numerical integrals and equation (7)
      Qnd <- t(sapply(1:nrow(W), function(i)Qn(Aval + delta, W[i,])))
      gnd <- t(sapply(1:nrow(W), function(i)gn(Aval, W[i,])))
      gnd <- gnd/rowSums(gnd)
      if(first) gn0d <- gnd
      EQnd <- rowSums(Qnd*gnd)*h.int
      D2   <- Qnd - EQnd
      QnAW <- Qn(A, W)
      H1   <- gn(A - delta, W)/gn(A, W)
      # equation (8)
      est.equation <- function(eps){
        sum((Y - (QnAW + eps*H1)) * H1 + (Qn(A + delta, W) - EQnd) -
        rowSums(D2*exp(eps*D2 + prev.sum)*gn0d)/rowSums(exp(eps*D2 + prev.sum)*gn0d))
      }
      eps  <- uniroot(est.equation, c(-1, 1))$root
      # updated values
```

```
      gn.new   <- function(a, w)exp(eps*Qn(a + delta, w)) * gn(a, w)
      Qn.new   <- function(a, w)Qn(a, w) + eps * gn(a - delta, w)/gn(a, w)
      prev.sum <- prev.sum + eps*D2
      return(list(Qn = Qn.new, gn = gn.new, prev.sum =
                      prev.sum, eps = eps, gn0d = gn0d))
  }
  ini.out <- f.iter(Qn, gn, first = TRUE)
  gn0d    <- ini.out$gn0d
  iter = 0
  # iterative procedure
  while(abs(ini.out$eps) > tol & iter <= iter.max){
    iter = iter + 1
    new.out <- f.iter(ini.out$Qn, ini.out$gn, gn0d, ini.out$prev.sum)
    ini.out <- new.out
  }
Qnd <- t(sapply(1:nrow(W), function(i)ini.out$Qn(Aval + delta, W[i,])))
gnd <- t(sapply(1:nrow(W), function(i)ini.out$gn(Aval, W[i,])))
gnd <- gnd/rowSums(gnd)
# plug in tmle
psi.hat <- mean(rowSums(Qnd*gnd)*h.int)
# influence curve of tmle
IC      <- (Y - ini.out$Qn(A, W))*ini.out$gn(A - delta, W)/ini.out$gn(A, W) +
           ini.out$Qn(A + delta, W) - psi.hat
var.hat <- var(IC)/length(Y)
return(c(psi.hat = psi.hat, var.hat = var.hat, IC = IC))
}
```

B.3 Example

Here is an example of how to use the previous function based on the data generating mechanism presented in the simulation

```
n  <- 100
W <- data.frame(W1 = runif(n), W2 = rbinom(n, 1, 0.7))
A  <- rpois(n, lambda = exp(3 + .3*log(W$W1) - .2*exp(W$W1)*W$W2))
Y  <- rbinom(n, 1, plogis(-1 + .05*A - .02*A*W$W2 + .2*A*tan(W$W1^2) -
      .02*W$W1*W$W2 + 0.1*A*W$W1*W$W2))
fitA.0 <- glm(A ~ I(log(W1)) + I(exp(W1)):W2, family = poisson, data = data.frame(A, W))
fitY.0 <- glm(Y ~ A + A:W2 + A:I(tan(W1^2)) + W1:W2 + A:W1:W2, family =
          binomial, data = data.frame(A, W))
gn.0  <- function(A = A, W = W)dpois(A, lambda = predict(fitA.0, newdata = W,
         type = "response"))
Qn.0 <- function(A = A, W = W)predict(fitY.0, newdata = data.frame(A, W,
       row.names = NULL), type = "response")
tmle00 <- tmle.shift(Y, A, W, Qn.0, gn.0, delta=2, tol = 1e-4, iter.max = 5,
          Aval = seq(1, 60, 1))
```

References

O. Bembom and M.J. van der Laan. A practical illustration of the importance of realistic individualized treatment rules in causal inference. *Electronic Journal of Statistics*, 2007.

P.J. Bickel, C.A.J. Klaassen, Y. Ritov, and J. Wellner. *Efficient and Adaptive Estimation for Semiparametric Models*. Springer-Verlag, 1997.

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.

Iván Díaz and Mark J. van der Laan. Super learner based conditional density estimation with application to marginal structural models. 2011. URL http://www.bepress.com/ucbbiostat/paper282.

F. Eberhardt and R. Scheines. Interventions and causal inference. *Department of Philosophy. Paper 415*, 2006. URL http://repository.cmu.edu/philosophy/415.

A.E. Hubbard and M.J. van der Laan. Population intervention models in causal inference. Technical report 191, Division of Biostatistics, University of California, Berkeley, 2005.

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.

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.

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.html.

S. Rose and M.J. van der Laan. *Targeted Learning: Causal Inference for Observational and Experimental Data*. Springer, New York, 2011.

D.O. Scharfstein, A.A. Tsiatis, and J.M. Robins. Semiparametric efficiency and its implications on the design and analysis of group-sequential studies. *Journal of the American Statistical Association*, 92 (440):1342–1350, 1997.

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.

I.B. Tager, M. Hollenberg, and W. Satariano. Self-reported leisure-time physical activity and measures of cardiorespiratory fitness in an elderly population. 1998.

Sarah L Taubman, James M Robins, Murray A Mittleman, and Miguel A Hernán. 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.

A.A. Tsiatis. *Semiparametric theory and missing data*. Springer series in statistics. Springer, 2006. ISBN 9780387324487. URL http://books.google.com/books?id=xqZFi2EMB40C.

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 M.L. Petersen. Causal effect models for realistic individualized treatment and intention to treat rules. *International Journal of Biostatistics*, 3(1), 2007.

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.

A. W. van der Vaart. *Asymptotic Statistics*. Cambridge University Press, 1998.