

Targeted Estimation and Inference for the Sample Average Treatment Effect

Laura B. Balzer*

Maya L. Petersen[†]

Mark J. van der Laan[‡]

*Division of Biostatistics, University of California, Berkeley - the SEARCH Consortium, lb-balzer@hsph.harvard.edu

[†]Division of Biostatistics, University of California, Berkeley - the SEARCH Consortium, may-aliv@berkeley.edu

[‡]Division of Biostatistics, University of California, Berkeley - the SEARCH Consortium, 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/paper334>

Copyright ©2015 by the authors.

Targeted Estimation and Inference for the Sample Average Treatment Effect

Laura B. Balzer, Maya L. Petersen, and Mark J. van der Laan

Abstract

While the population average treatment effect has been the subject of extensive methods and applied research, less consideration has been given to the sample average treatment effect: the mean difference in the counterfactual outcomes for the study units. The sample parameter is easily interpretable and is arguably the most relevant when the study units are not representative of a greater population or when the exposure's impact is heterogeneous. Formally, the sample effect is not identifiable from the observed data distribution. Nonetheless, targeted maximum likelihood estimation (TMLE) can provide an asymptotically unbiased and efficient estimate of both the population and sample parameters. In this paper, we study the asymptotic and finite sample properties of the TMLE for the sample effect and provide a conservative variance estimator. In most settings, the sample parameter can be estimated more efficiently than the population parameter. Finite sample simulations illustrate the potential gains in precision and power from selecting the sample effect as the target of inference. As a motivating example, we discuss the Sustainable East Africa Research in Community Health (SEARCH) study, an ongoing cluster randomized trial for HIV prevention and treatment.

1 Introduction

In many studies, the goal is to estimate the impact of an exposure on the outcome of interest. Often the target causal parameter is the population average treatment effect (PATE): the expected difference in the counterfactual outcomes if all members of some population were exposed and if all members of that population were unexposed. If there are no unmeasured confounders and there is sufficient variability in the exposure assignment (i.e. if the randomization and positivity assumptions hold), then we can identify the causal parameter as a function of the observed data distribution (Rosenbaum and Rubin, 1983; Robins, 1986). The resulting statistical parameter can be estimated with a variety of algorithms.

Alternate causal parameters, receiving less attention, include the sample average treatment effect (SATE) and the conditional average treatment effect (CATE). The sample effect is the average difference in the counterfactual outcomes for the actual study units (Neyman, 1923). In other words, the SATE is the intervention effect for the sample at hand. The SATE remains interpretable if the study units are not reflective of a larger population and is responsive to variations in the intervention effect by measured and unmeasured covariates. The conditional effect is the average difference in the expected counterfactual outcomes, treating the baseline covariates of the study units as fixed (Abadie and Imbens, 2002). In other words, the CATE is the intervention effect, given the measured covariates of the sample units. The exact interpretation and the variability of the CATE depend on the conditioning set. Another key difference between the three parameters is in the variance of common estimators. As shown by Imbens (2004) and elaborated here, an efficient estimator of the sample parameter is often more precise than the same estimator of the conditional parameter, which is, in turn, often more precise than the same estimator of the population parameter.

To the best of our knowledge, this is the first paper to propose using targeted maximum likelihood estimation (TMLE) for the SATE. TMLE is a general algorithm for constructing double robust, semiparametric efficient, substitution estimators (van der Laan and Rubin, 2006; van der Laan and Rose, 2011). Even though the SATE is not identified, we prove that the TMLE, presented here, is an asymptotically linear estimator of the SATE and provide a conservative approximation of its influence curve. Our results generalize the variance derivations of Imbens (2004) to allow misspecification of the outcome regression (i.e. the conditional mean outcome, given the exposure and covariates) and estimation of the propensity score (i.e. the conditional probability of the receiving the exposure, given the covariates). Simulations are used to evaluate the finite sample performance of our point estimator and proposed variance estimator. The simulations also serve to highlight the differences between the three causal parameters and the potential gains in power from selecting the sample effect as the target of inference. We begin by reviewing the structural causal model of Pearl (1995) and motivate our discussion with the Sustainable East Africa Research in Community Health (SEARCH) trial for HIV prevention and treatment (NCT01864603) (University of California, San Francisco, 2013).

2 Causal Model & Causal Parameters

SEARCH is an ongoing cluster randomized trial to evaluate the effect of a community-based strategy for HIV prevention and treatment in rural East Africa. In intervention communities, all individuals testing HIV+ are immediately eligible for antiretroviral therapy (ART) with streamlined delivery, including enhanced services for initiation, linkage, and retention in care. In control communities, all individuals testing HIV+ are offered ART according to in-country guidelines, largely based on CD4+ T cell counts. The study hypothesis is that ART initiation at any CD4 count and with

streamlined delivery will reduce the five-year cumulative HIV incidence. The primary outcome as well as other health, educational and economic outcomes will be measured among approximately 320,000 individuals, enrolled in the study. For the purposes of discussion, we focus on the community-level data. Thereby, our results are equally applicable to clustered and non-clustered data structures.

Consider the following data generating process for a randomized trial with two arms. First, the study units are selected. While some trials obtain a simple random sample from a well-defined population, in many other studies the selection of units is more systematic. In the SEARCH trial, for example, 32 communities were selected from Western Uganda (Mbarara region), Eastern Uganda (Tororo region) and the Southern Nyanza Province in Kenya. These communities satisfied the study’s inclusion criteria, including community size, health care infrastructure and accessibility by a maintained transportation route. Next, the baseline covariates W are measured. Throughout we use “baseline” to refer to covariates measured prior to implementation of the intervention. For the SEARCH trial, these include region, occupational mix, migration index, male circumcision coverage and measures of HIV prevalence.

Next, the intervention is randomized to the study units. Balanced allocation of the intervention can be guaranteed by randomly assigning the intervention to $n/2$ units and the control to remaining units or by randomizing within matched pairs. In the SEARCH trial, for example, communities were first matched on baseline covariates and then the intervention randomized within the resulting 16 matched pairs (Balzer et al., 2015). For ease of exposition, we present the causal model for the simple scenario, where the intervention is completely randomized, but our results are general. Let A be a binary variable, reflecting the assigned level of the intervention. For the SEARCH trial, A equals one if the community was assigned to the treatment (all individuals testing positive for HIV are immediately offered ART with streamlined care) and equals zero if the community was assigned to the control (all individuals testing positive for HIV are offered ART according to in-country guidelines). At the end of followup, the outcome Y is measured. For the SEARCH trial, Y is the five-year cumulative incidence of HIV and will be measured through longitudinal follow-up. The observed data for a given study unit are then

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

We observe n independent, identically distributed (i.i.d.) copies of O with distribution P_0 . We note that for estimation and inference of the sample and conditional effects, we can weaken the i.i.d. assumption. In particular, we do not need any assumptions on the joint distribution of covariates $P_0(W_1, \dots, W_n)$. For further details, see Balzer et al. (2015).

This data generating process can be described by the following structural causal model (SCM) (Pearl, 1995, 2000). Each component of the observed data is assumed to be a deterministic function of its parents (variables that may influence its value) and unobservable background factors:

$$\begin{aligned} W &= f_W(U_W) \\ A &= \mathbb{I}(U_A < 0.5) \\ Y &= f_Y(W, A, U_Y) \end{aligned}$$

Let $X = (W, A, Y)$ denote the set of endogenous factors and $U = (U_W, U_A, U_Y)$ denote the set of the background factors with joint distribution P_U . By design, the random error determining the intervention assignment U_A is independent from the unmeasured factors contributing the baseline covariates U_W and the outcome U_Y :

$$U_A \perp\!\!\!\perp (U_W, U_Y)$$

Specifically, U_A is independently drawn from a Uniform(0,1). The causal model \mathcal{M}^F provides the set of allowed distributions for (U, X) and implies the statistical model \mathcal{M} for the set of possible distributions of the observed data O . The true joint distribution of the background and endogenous factors $P_{U,X,0}$ is an element of \mathcal{M}^F , and the true distribution of the observed data P_0 is an element of \mathcal{M} . In a randomized trial, the statistical model is semiparametric.

Through interventions on the SCM, we can generate the counterfactual outcome $Y(a)$, which is the outcome if possibly contrary-to-fact the unit was assigned $A = a$:

$$\begin{aligned} W &= f_W(U_W) \\ A &= a \\ Y(a) &= f_Y(W, a, U_Y) \end{aligned}$$

The distribution of the counterfactuals can then be used to define the causal parameter of interest. Often, the target of inference is the population average treatment effect (PATE):

$$\Psi^P(P_{U,X}) = E_{U,X}[Y(1) - Y(0)]$$

where the subscript (U, X) denotes the expectation over the distribution of $P_{U,X}$ (which implies the distribution of the counterfactual outcomes). This causal parameter is the expected difference in the counterfactual outcomes for underlying target population from which the units were sampled. For the SEARCH trial, $\Psi^P(P_{U,X})$ is the difference in the expected counterfactual cumulative incidence of HIV if possibly contrary-to-fact all communities in some hypothetical target population implemented the test-and-treat strategy, and expected counterfactual cumulative incidence of HIV if possibly contrary-to-fact all communities in that hypothetical target population continued with the standard of care. From the SCM, we see that the expectation is over the measured factors W and unmeasured factors U_Y , which determine the counterfactual outcomes for the population. In other words, the true value of $\Psi^P(P_{U,X})$ does not depend on the sampled values of W or U_Y .

An alternative causal estimand is the sample average treatment effect (SATE), which was first proposed in Neyman (1923):

$$\Psi^S(P_{U,X}) = \frac{1}{n} \sum_{i=1}^n Y_i(1) - Y_i(0)$$

This is simply the intervention effect for the study units. For the SEARCH trial, $\Psi^S(P_{U,X})$ is the average difference in the counterfactual cumulative incidence of HIV under the test-and-treat strategy and under the standard of care for the $n = 32$ study communities. The parameter is data-adaptive; its value changes with each new selection or sample of units. The SATE remains interpretable if the study units were systematically selected and is responsive to variation in intervention effect by measurable and unmeasurable factors (i.e. $\{W, U_Y\}$).

An intermediate between the population and sample parameters is the conditional average treatment effect (CATE), which was first proposed in Abadie and Imbens (2002):

$$\begin{aligned} \Psi^C(P_{U,X}) &= \frac{1}{n} \sum_{i=1}^n E_{U,X}[Y_i(1) - Y_i(0) | W_1, \dots, W_n] \\ &= \frac{1}{n} \sum_{i=1}^n E_{U,X}[Y_i(1) - Y_i(0) | W_i] \end{aligned}$$

where the second equality holds under our assumption that study units are causally independent (i.e. the baseline covariates and intervention assignment of one unit do not affect the outcome of

another unit). This parameter is the difference in the expected counterfactual outcomes, treating the measured covariates of the study units as fixed. For the SEARCH trial, $\Psi^C(P_{U,X})$ is interpreted as the average difference in the expected counterfactual cumulative incidence of HIV under the test-and-treat strategy and under the standard of care, given the measured covariates of the $n = 32$ study communities. From the SCM, we see that the expectation is over the unmeasured factors U_Y that determine the counterfactual outcomes.

The exact interpretation, the true value and the variability of the CATE depend on the conditioning set. As an extreme example, suppose that in the SEARCH trial the only measured covariate were region. Then we would interpret $\Psi^C(P_{U,X})$ as the treatment effect, given the regional distribution of communities. If the regional distribution were set by design (e.g. 10 communities in Eastern Uganda, 10 communities in Western Uganda and 12 communities in Kenya), then we would obtain the same value $\Psi^C(P_{U,X})$ over repeated studies. In other words, we would be averaging out all the other factors contributing to HIV incidence. Now suppose the set of measured covariates included both region (set by design) and baseline HIV prevalence (varying from community to community). Then we would interpret the CATE as the treatment effect, given the regional distribution and baseline prevalence of the study communities. Over repeated studies, the value of the CATE would change due the differences in the sampled values of baseline prevalence.

3 Identifiability

To identify the above causal effects, we must write them as some function of the observed data distribution. Under the randomization and positivity assumptions, we can identify the mean counterfactual outcome within strata of covariates (Rosenbaum and Rubin, 1983; Robins, 1986):

$$E_{U,X,0}[Y(a)|W] = E_{U,X,0}[Y(a)|A = a, W] = E_0[Y|A = a, W]$$

where the subscript 0 denotes the expectation over the true distribution. (Recall $P_{U,X,0}$ is the true joint distribution of the background and endogenous factors and P_0 is the true distribution of the observed data.) Briefly, the randomization assumption states that the counterfactual outcome is independent of the exposure, given the measured covariates: $A \perp\!\!\!\perp Y(a)|W$. This is equivalent to the no unmeasured confounders assumption (Rosenbaum and Rubin, 1983). The positivity assumption states that there is sufficient variability in the exposure assignment within strata of covariates. Both assumptions hold by design in a randomized trial. As a well known result, the PATE $\Psi^P(P_{U,X,0})$ is easily identified as

$$\begin{aligned} \Psi_0^P(P_0) &= E_0 \left[E_0(Y|A = 1, W) - E_0(Y|A = 0, W) \right] \\ &= E_0[\bar{Q}_0(1, W) - \bar{Q}_0(0, W)] \end{aligned}$$

where $\bar{Q}_0(A, W) = E_0(Y|A, W)$ denotes the conditional mean outcome, given the exposure and covariates. This statistical estimand is also called the G-computation identifiability result (Robins, 1986). For the SEARCH trial, $\Psi_0^P(P_0)$ is the difference in expected cumulative incidence of HIV, given the treatment and measured covariates, and the expected cumulative incidence of HIV, given the control and measured covariates, averaged (standardized) with respect to the covariate distribution in the population. As with the causal parameter, there is one true value $\Psi_0^P(P_0) = \psi_0^P$ for the population. In a randomized trial, conditioning on the covariates W is not needed for identifiability, but can provide efficiency gains during estimation (e.g. Fisher (1932); Cochran (1957); Cox and McCullagh (1982); Tsiatis et al. (2008); Moore and van der Laan (2009)).

Analogously, we can identify the CATE $\Psi^C(P_{U,X,0})$ as

$$\begin{aligned}\Psi_0^C(P_0) &= \frac{1}{n} \sum_{i=1}^n [E_0(Y_i|A_i = 1, W_i) - E_0(Y_i|A_i = 0, W_i)] \\ &= \frac{1}{n} \sum_{i=1}^n [\bar{Q}_0(1, W_i) - \bar{Q}_0(0, W_i)]\end{aligned}\tag{1}$$

This statistical estimand is the difference in the conditional expectation of the outcome, given the intervention and measured covariates, evaluated at the treatment vs. control level of the intervention, but now averaged over the sampled values of the measured covariates. For the SEARCH study, $\Psi_0^C(P_0)$ is the sample average of the difference in expected cumulative incidence of HIV, given the treatment and measured covariates, and the expected cumulative incidence of HIV, given the control and measured covariates. As with the CATE, the interpretation, the true value and the variability of $\Psi_0^C(P_0)$ depend on both the conditioning set and the sample.

Unlike the other two causal parameters, the SATE is *non-identifiable*. We cannot write the causal parameter as a function of the observed data distribution. This point has not received much attention. Instead, researchers have largely focused on the lack of identifiability of the variance of standard estimators (Neyman, 1923; Rubin, 1990; Imbens, 2004; Imai, 2008). To elaborate, let us use the structural causal model \mathcal{M}^F to rewrite the SATE in terms of the CATE:

$$\begin{aligned}\Psi^S(P_{U,X}) &= \frac{1}{n} \sum_{i=1}^n Y_i(1) - Y_i(0) \\ &= \frac{1}{n} \sum_{i=1}^n f_Y(W_i, 1, U_{Y_i}) - f_Y(W_i, 0, U_{Y_i}) \\ &= \frac{1}{n} \sum_{i=1}^n E_{U,X} [Y_i(1) - Y_i(0) | W_i, U_{Y_i}]\end{aligned}$$

The second equality is from the definition of counterfactuals as interventions on the causal model. The final equality is the CATE, given the measured baseline covariates as well as the unmeasured factors. If we had access to all pre-intervention covariates impacting the outcome (i.e. $\{W, U_Y\}$), then we could apply the results for estimation and inference for the conditional parameter, as detailed in Balzer et al. (2015). In reality, we only measure a subset of these covariates (i.e. W) and only this subset is available for estimation and inference. Nonetheless, we show below that the TMLE is asymptotically linear for the SATE and the corresponding variance estimator is asymptotically conservative.

4 Estimation & Inference

There are many well-established algorithms for estimation of the population parameter $\Psi_0^P(P_0)$. For example, matching and inverse weighting estimators rely on knowledge or estimation of the propensity score, which is the conditional probability of being exposed, given the covariates $P_0(A = 1|W)$ (e.g. Horvitz and Thompson (1952); Rosenbaum and Rubin (1983); Robins (1993); Hernán et al. (2006)). Simple substitution estimators rely on estimation of the outcome regression, which is the conditional mean outcome given the exposure and covariates $\bar{Q}_0(A, W) = E_0(Y|A, W)$ (e.g. Robins (1986); Ahern et al. (2009); Snowden et al. (2011)). A third class of estimators requires estimation of both the propensity score and the outcome regression. This class includes augmented

inverse probability of treatment weighting (AIPTW) (e.g. Robins et al. (1994); Rotnitzky et al. (1998); Robins et al. (2000); van der Laan and Robins (2003)) and TMLE (e.g. van der Laan and Rubin (2006); van der Laan and Rose (2011)). These estimators are double robust in that they will be consistent if either the propensity score or the outcome regression is consistently estimated. If both functions are consistently estimated at a fast enough rate and there is sufficient variability in the propensity score, these estimators are also asymptotically efficient in that they attain the lowest possible variance among a large class of regular, asymptotically linear estimators. An important distinction between AIPTW and TMLE is that the former is based on solving an estimating equation, while the latter is a substitution estimator, providing stability in the context of sparsity (Gruber and van der Laan, 2010; Balzer et al., 2014). We focus our discussion on TMLE in a randomized trial and provide generalizations to an observational setting.

4.1 For the Population Parameter

For the population estimand, a TMLE can be implemented with the following steps. First, we obtain an initial estimate of the outcome regression $\bar{Q}_0(A, W)$. This function can be estimated with maximum likelihood or with an *a priori* specified data-adaptive procedure, such as Super Learner (van der Laan et al., 2007). In a randomized trial, the propensity score $g_0(1|W) = P_0(A = 1|W)$ is known and does not need to be estimated. In the two-armed trial, for example, we have $g_0(1|W) = g_0(1) = 0.5$. Estimation of the propensity score, however, can improve efficiency by capturing chance imbalances in the covariate distribution between treatment groups (e.g. van der Laan and Robins (2003); Moore and van der Laan (2009)). We could, for example, obtain an estimator $g_n(1|W)$ by running logistic regression of the intervention A on the measured covariates W .

Next, we target the initial estimator of the outcome regression $\bar{Q}_n(A, W)$. This targeting step uses information in the propensity score to obtain the optimal bias-variance tradeoff for the parameter of interest and to solve the efficient score equation. It is accomplished by running logistic regression¹ of the outcome Y on the covariate $H_n(A, W) = \left(\frac{\mathbb{I}(A=1)}{g_n(1|W)} - \frac{\mathbb{I}(A=0)}{g_n(0|W)} \right)$ with the $\text{logit}(x) = \log\{x/(1-x)\}$ of the initial estimator $\bar{Q}_n(A, W)$ as offset. The estimated coefficient ϵ_n is then plugged into the fluctuation model to yield targeted updates of the outcome regression under the treatment and under the control:

$$\begin{aligned}\bar{Q}_n^*(1, W) &= \text{expit} \left[\text{logit} [\bar{Q}_n(1, W)] + \epsilon_n H_n(1, W) \right] \\ \bar{Q}_n^*(0, W) &= \text{expit} \left[\text{logit} [\bar{Q}_n(0, W)] + \epsilon_n H_n(0, W) \right]\end{aligned}$$

where expit is the inverse of the logit function and where the $*$ denotes the targeted estimator. In a randomized trial, if the propensity score is treated as known (i.e. not estimated) and the regression model used for initial estimation of $\bar{Q}_0(A, W)$ contains an intercept and a main term for the exposure, then this targeting step will not yield an update and can be skipped (Moore and van der Laan, 2009; Rosenblum and van der Laan, 2010). Lastly, the targeted estimates are substituted into the parameter mapping:

$$\Psi_n(P_n) = \frac{1}{n} \sum_{i=1}^n \left[\bar{Q}_n^*(1, W_i) - \bar{Q}_n^*(0, W_i) \right]$$

¹As detailed in Gruber and van der Laan (2010), the same procedure can be applied to a bounded continuous outcome and adds robustness in the context of sparsity.

where P_n denotes the empirical distribution, placing mass $1/n$ on each observation O_i . The sample mean is the nonparametric maximum likelihood estimator of the marginal distribution of baseline covariates.

Under regularity conditions, the TMLE is a consistent and asymptotically linear estimator of the population parameter (van der Laan and Rubin, 2006):

$$\Psi_n(P_n) - \Psi_0^P(P_0) = \frac{1}{n} \sum_{i=1}^n D^P(\bar{Q}, g_0)(O_i) + o_P(1/\sqrt{n})$$

with influence curve

$$\begin{aligned} D^P(\bar{Q}, g_0)(O) &= D_Y(\bar{Q}, g_0)(O) + D_W(\bar{Q}, g_0)(O) \\ D_Y(\bar{Q}, g_0)(O) &= \left(\frac{\mathbb{I}(A=1)}{g_0(1|W)} - \frac{\mathbb{I}(A=0)}{g_0(0|W)} \right) (Y - \bar{Q}(A, W)) \\ D_W(\bar{Q}, g_0)(O) &= \bar{Q}(1, W) - \bar{Q}(0, W) - \psi_0^P \end{aligned}$$

where $\bar{Q}(A, W)$ denotes the limit of the TMLE $\bar{Q}_n^*(A, W)$ and we are assuming the propensity score is known or consistently estimated, as will always be true when A is randomized. The first term of the influence curve D_Y is the weighted residuals (i.e. weighted deviations between the observed outcome and the limit of the predicted outcome). The second term D_W is deviation between the limit of the estimated strata-specific association and the marginal association.

The standardized estimator is asymptotically normal with variance given by the variance of its influence curve, divided by sample size n . Under consistent estimation of the outcome regression (i.e. when $\bar{Q}(A, W) = \bar{Q}_0(A, W)$), the TMLE will be asymptotically efficient and achieve the lowest possible variance among a large class of estimators (Bickel et al., 1993). In other words, its influence curve equals the efficient influence curve. Thereby, improved estimation of conditional mean outcome leads to more precise estimators of the population effect. In a randomized trial, adjusting for measured baseline covariates with TMLE can lead to substantial efficiency gains without risk of bias due to regression model misspecification (Moore and van der Laan, 2009; Rosenblum and van der Laan, 2010). In finite samples, the variance of the TMLE is well-approximated by the sample variance of the estimated influence curve, divided by sample size. The algorithm is available in the `tmle` (Gruber and van der Laan, 2012) and `ltmle` (Schwab et al., 2014) packages in R (R Core Team, 2014).

4.2 For the Conditional Parameter

The TMLE for the population parameter $\Psi_0^P(P_0)$ also serves as an estimator of the conditional parameter $\Psi_0^C(P_0)$. The steps are analogous with one important exception. In the final step of substituting in the targeted estimates, the empirical mean is now considered part of the parameter mapping (Eq. 1) and not an estimator of the covariate distribution, which is considered fixed. As a result, there is no contribution to the variance from the covariate distribution. Thereby, estimators of the conditional parameter are often more efficient than those of the population parameter (Imbens, 2004; Abadie and Imbens, 2008; Imbens, 2011).

The TMLE is also a consistent and asymptotically linear estimator of the conditional estimand (Balzer et al., 2015):

$$\Psi_n(P_n) - \Psi_0^C(P_0) = \frac{1}{n} \sum_{i=1}^n D^C(\bar{Q}, g_0)(O_i) + o_P(1/\sqrt{n})$$

with influence curve given by

$$D^{\mathcal{C}}(\bar{Q}, g_0)(O) = D_Y(\bar{Q}, g_0)(O) - E_0[D_Y(\bar{Q}, g_0)(O)|\mathbf{W}]$$

where $\mathbf{W} = (W_1, \dots, W_n)$ denotes the vector of baseline covariates for the study units. The influence curve of the TMLE for $\Psi_0^{\mathcal{C}}(P_0)$ depends on the true conditional mean outcome $\bar{Q}_0(A, W)$. In particular, the conditional expectation of the D_Y component, given the vector of baseline covariates, equals the deviation between the true mean and the limit of the estimated mean:

$$E_0[D_Y(\bar{Q}, g_0)(O)|\mathbf{W}] = [\bar{Q}_0(1, W) - \bar{Q}(1, W)] - [\bar{Q}_0(0, W) - \bar{Q}(0, W)]$$

Under consistent estimation of the outcome regression (i.e. when $\bar{Q}(A, W) = \bar{Q}_0(A, W)$), this term is zero and the TMLE for $\Psi_0^{\mathcal{C}}(P_0)$ efficient. In this setting, the TMLE for the conditional estimand will often have a smaller asymptotic variance than the same TMLE for the population estimand:

$$\begin{aligned}\sigma^{2,\mathcal{P}} &= \text{Var}[D_Y(\bar{Q}_0, g_0)(O)] + \text{Var}[D_W(\bar{Q}_0, g_0)(O)] \\ \sigma^{2,\mathcal{C}} &= \text{Var}[D_Y(\bar{Q}_0, g_0)(O)]\end{aligned}$$

They will only have the same efficiency bound when there is no variability in the treatment effect across strata of covariates (i.e. when $\text{Var}[\bar{Q}_0(1, W) - \bar{Q}_0(0, W) - \psi_0^{\mathcal{P}}] = 0$). In many settings, there will be effect modification, and focusing on estimation of the conditional parameter will yield more precision and power.

In practice, there are likely to be deviations between the true outcome regression and the limit of our estimator. Nonetheless, we can conservatively approximate the influence curve of the TMLE for the conditional estimand $\Psi_0^{\mathcal{C}}(P_0)$ as

$$D_n^{\mathcal{C}}(O_i) = D_{Y,n}(O_i) = \left(\frac{\mathbb{I}(A_i = 1)}{g_0(1|W_i)} - \frac{\mathbb{I}(A_i = 0)}{g_0(0|W_i)} \right) (Y_i - \bar{Q}_n^*(A_i, W_i)) \quad (2)$$

(Balzer et al., 2015). Thereby, we obtain an asymptotically conservative variance estimator with the sample variance of the weighted residuals, divided by sample size n . As estimation of the outcome regression improves, the deviations between the true and estimated means are reduced and we get closer to approaching the efficiency bound $\sigma^{2,\mathcal{C}}$. Thereby, in a randomized trial, adjusting for baseline covariates, predictive of the outcome, can substantially improve power by reducing variability in the estimator and resulting in a less conservative variance estimator (Balzer et al., 2015).

4.3 For the Sample Parameter

For a randomized trial, Neyman (1923) proposed estimating the SATE $\Psi^{\mathcal{S}}(P_{U,X})$ with the unadjusted estimator, which is the difference in the average outcomes among the treated units and the average outcomes among the control units:

$$\Psi_{n,unadj}(P_n) = \frac{\sum_{i=1}^n \mathbb{I}(A_i = 1)Y_i}{\sum_{i=1}^n \mathbb{I}(A_i = 1)} - \frac{\sum_{i=1}^n \mathbb{I}(A_i = 0)Y_i}{\sum_{i=1}^n \mathbb{I}(A_i = 0)}$$

In this setting, the difference-in-means estimator will be unbiased for the SATE, conditional on the vector of counterfactual outcomes $\mathbf{Y}(\mathbf{a}) = \{Y_i(a) : i = 1, \dots, n, a = 0, 1\}$. However, its variance remains unidentifiable as it relies on the correlation of the counterfactual outcomes $\{Y_i(1), Y_i(0)\}$ (Neyman, 1923). Imbens (2004) later generalized this work for an efficient estimator (i.e. a regular, asymptotically linear estimator, whose influence curve equals the efficient influence curve) in

an observational setting. In particular, he showed that an efficient estimator for the population parameter was unbiased for the sample parameter, conditional on the vector of baseline covariates \mathbf{W} and the set of counterfactual outcomes $\mathbf{Y}(\mathbf{a})$. Further he expressed the variance of an efficient estimator of the SATE in terms of the variance of the the same estimator of the PATE minus the variance of the unit-specific treatment effect across the population. We now extend these results to the TMLE when the estimator of $\bar{Q}_0(A, W)$ converges to a possibly misspecified limit and suggest alternate methods for variance estimation.

The TMLE for the population and conditional parameters ($\Psi_0^C(P_0)$ and $\Psi_0^C(P_0)$) is also a consistent and asymptotically linear estimator of the SATE:

$$\Psi_n(P_n) - \Psi^S(P_{U,X}) = \frac{1}{n} \sum_{i=1}^n D^S(\bar{Q}, g_0)(U_i, X_i) + o_P(1/\sqrt{n})$$

with influence curve

$$\begin{aligned} D^S(\bar{Q}, g_0)(U, X) &= D^C(\bar{Q}, g_0)(O) - D^F(U, X) \\ D^F(U, X) &= Y(1) - Y(0) - [\bar{Q}_0(1, W) - \bar{Q}_0(0, W)] \end{aligned}$$

The proof is given in Appendix A. Recall $U = (U_W, U_A, U_Y)$ denotes the set of background factors and $X = (W, A, Y)$ denotes the endogenous factors in our SCM. In words, the influence curve of the TMLE for the sample parameter D^S is given by the influence curve for TMLE of the conditional parameter D^C minus a non-identifiable piece, which captures the deviations between the unit-specific treatment effect and expected effect within covariate strata:

$$\begin{aligned} D^F(U_i, X_i) &= Y_i(1) - Y_i(0) - [E_0(Y_i|A_i = 1, W_i) - E_0(Y_i|A_i = 0, W_i)] \\ &= Y_i(1) - Y_i(0) - [E_{U,X,0}(Y_i(1)|W_i) - E_{U,X,0}(Y_i(0)|W_i)] \\ &= Y_i(1) - Y_i(0) - E_{U,X,0}[Y_i(1) - Y_i(0)|W_i] \end{aligned}$$

In the last line, the expectation is over the unmeasured factors U_Y that determine the counterfactual outcomes.

The standardized estimator of the SATE is asymptotically normal with mean zero and variance

$$\begin{aligned} Var[D^S(U, X)] &= Var[D^C(O)] + Var[D^F(U, X)] - 2Cov[D^C(O), D^F(U, X)] \\ &= Var[D^C(O)] - Var[D^F(U, X)] \end{aligned}$$

The proof is given in Appendix B. Since the variance of the D^F component must be greater than or equal to zero, the asymptotic variance of the TMLE as an estimator of the sample parameter will be less than or equal to the asymptotic variance of the same estimator of the conditional parameter. They will only have the same precision when there is no variability in the treatment effect within strata of covariates W . In many settings, however, there will be heterogeneity, and TMLE for the SATE will be more precise and powerful.

Along the same lines, we can conservatively approximate the influence curve of the TMLE for SATE by ignoring the non-identifiable piece D^F . Specifically, we obtain a conservative variance estimator with the sample variance of the estimated D_n^C component (Eq. 2), divided by sample size n . This variance estimator is easy to implement as the relevant pieces are known or already estimated. As a result, this may provide an attractive alternative to the matching estimator of the variance, proposed by Abadie and Imbens (2002) and discussed in Imbens (2004). We note that the bootstrap is inappropriate as the parameter changes with each sample.

4.4 Generalization for Observational Studies

Thus far, we have focused on a randomized trial. In an observational setting, TMLE can be implemented in an analogous manner. In the first step, we estimate both the outcome regression $\bar{Q}_0(A, W)$ and the propensity score $g_0(1|W)$. Again, we could use parametric regression or data-adaptive algorithms. In the targeting step, we run logistic regression of the outcome Y on the estimated covariate $H_n(A, W)$ with the *logit* of the initial estimator $\bar{Q}_n(A, W)$ as offset. We then plug-in the estimated coefficient ϵ_n to obtain the targeted estimates $\bar{Q}_n^*(1, W)$ and $\bar{Q}_n^*(0, W)$. The targeted estimates are then substituted into the parameter mapping.

In an observational setting, TMLE also exhibits desirable asymptotic properties. TMLE is double robust: if either the outcome regression $\bar{Q}_0(A, W)$ or the propensity score $g_0(1|W)$ are consistently estimated, we will have a consistent estimate of the parameter of interest. If both functions are consistently estimated at a fast enough rate and the positivity assumption holds, then the TMLE will be asymptotically efficient. As before, the TMLE for the sample parameter will be at least as precise as the TMLE for the conditional parameter, which will be at least as precise as the TMLE for the population parameter. Furthermore, if the outcome regression is not consistently estimated but the propensity score is consistently estimated with maximum likelihood, then $D^{\mathcal{P}}(\bar{Q}, g_0)$ provides an asymptotically conservative approximation of the influence curve for the TMLE of the population estimand $\Psi_0^{\mathcal{P}}(P_0)$ (van der Laan and Robins, 2003). Likewise, under these conditions, $D^{\mathcal{C}}(\bar{Q}, g_0)$ provides an asymptotically conservative approximation of the influence curve for the TMLE of the conditional estimand $\Psi_0^{\mathcal{C}}(P_0)$ and thereby SATE. Further details are given in Appendix C.

5 Simulation Study

We present the following simulation study to (1) further illustrate the differences between the causal parameters, (2) demonstrate implementation of the TMLE, and (3) understand the impact of the parameter specification on the estimator's true variance, on variance estimation and on attained power. We focus on a randomized trial to illustrate the potential gains in efficiency with adjustment during the analysis. All simulations were carried out in R v3.1.0 (R Core Team, 2014).

5.1 Data generating process and estimators

Consider the following data generating process for unit $i = \{1, \dots, n\}$. First, we generated the background error $U_{Y,i}$ by drawing from a standard normal distribution. Then we generated three baseline covariates $W = (W1, W2, W3)$ by drawing independently from a standard normal distribution. The exposure A_i was randomized such that the treatment allocation was balanced overall. Recall A_i is a binary indicator, equaling 1 if the unit is randomized to the treatment and 0 if the unit is randomized to the control. The outcome Y_i was generated as

$$Y_i = \text{expit}[A_i + 0.5(W1_i + W2_i + W3_i) + U_{Y,i} + 1.5A_i(W1_i - W2_i) - AU_{Y,i}]/5$$

We also generated the counterfactual outcomes $Y_i(a)$ by intervening to set $A_i = a$. For sample sizes $n = \{50, 70, 100\}$, this data generating process was repeated 2,500 times. For each sample, the SATE was calculated as the average difference in the counterfactual outcomes, and the CATE was calculated as the average difference in the expected counterfactual outcomes, given the baseline covariates for the study units. (The conditional expectation $E_{U,X}[Y_i(a)|W_i]$ was approximated by fixing (a, W_i) and averaging the counterfactual outcomes over 75,000 units.) The PATE was

calculated by averaging the difference in the counterfactual outcomes over a population of 500,000 units.

We compared the performance of the unadjusted estimator to the TMLE with two methods for initial estimation of the outcome regression $\bar{Q}_0(A, W)$. Specifically, we estimated $\bar{Q}_0(A, W)$ with logistic regression, including as main terms the exposure A , the covariate $W1$ and an interaction $A*W1$. We also estimated $\bar{Q}_0(A, W)$ with Super Learner, an optimal machine-learning approach (van der Laan et al., 2007). In particular, we used 10-fold cross-validation to create the best convex combination of algorithm-specific estimates from the following library: logistic regression with main terms for the exposure A and a single covariate, logistic regression with main terms for the exposure A , a single covariate and their interaction, as well as stepwise logistic regression with and without interactions. The unadjusted estimator can be considered as a special case of the TMLE, where $\bar{Q}_n(A, W) = \bar{Q}_n(A)$. Inference was based on the estimated influence curve. We constructed Wald-type 95% confidence intervals and tested the null hypothesis of no effect.

5.2 Simulation Results

Table 1 gives a summary of the parameter values across 2,500 samples. Recall the true values of the SATE and CATE depend on the units included in the study. The sample effect ranged from -1.04% to 6.69%; the conditional effect ranged from -0.12% to 5.46%, while the population effect was constant at 2.73%. By averaging out the unmeasured factors contributing to the counterfactual outcomes U_Y , the CATE was less variable than the SATE. Likewise, by averaging out the measured and unmeasured factors contributing to the counterfactual outcomes (W, U_Y) across the population, the PATE is less variable than the CATE. As expected, the variability in the SATE and CATE decrease with increasing sample size.

	SATE				CATE				PATE			
	min	mean	max	var	min	mean	max	var	min	mean	max	var
$n = 50$	-1.04	2.73	5.81	1.03E-2	-0.12	2.74	5.46	7.29E-3	2.73	2.73	2.73	0
$n = 70$	-0.34	2.71	6.69	7.06E-3	0.29	2.71	5.34	4.94E-3	2.73	2.73	2.73	0
$n = 100$	0.21	2.72	5.49	4.77E-3	0.46	2.72	5.11	3.53E-3	2.73	2.73	2.73	0

Table 1: Summary of the causal parameters over 2,500 simulations of size $n = \{50, 70, 100\}$. All values are in percent.

Table 2 illustrates the performance of the estimators over the 2,500 simulated data sets. Specifically, we give the bias as the average deviation between the point estimate and (sample-specific) true value, the standard deviation σ as the square root of the variance of estimator relative to its target, and the average standard error estimate $\hat{\sigma}$, based on the influence curve. We also show the “true” power, which is the proportion of times the false null hypothesis would be rejected if the estimator’s variance σ^2 were known, and the attained power, which is the proportion of times the false null hypothesis was rejected when the variance was estimated. The 95% confidence interval coverage is also included.

Since the exposure was randomized, all estimators are unbiased. There was no risk of bias due to misspecification the regression model for $\bar{Q}_0(A, W)$ (e.g. Moore and van der Laan (2009); Rosenblum and van der Laan (2010)). As expected, the variance of the estimators decreased and thereby the “true” power increased with increasing sample size and with adjustment (e.g. Fisher (1932); Cochran (1957); Cox and McCullagh (1982)). For example, the true power of the unadjusted estimator for the SATE was 57% with $n = 50$, 69% with $n = 70$ and 80% with $n = 100$. After adjusting for a single covariate with TMLE, the true power for the SATE increased to 78% with $n = 50$, 89% with $n = 70$, and 96% with $n = 100$. There were minimal differences in the variance

and true power of the TMLE using logistic regression for initial estimation of the outcome regression and TMLE using Super Learner for initial estimation.

For all sample sizes and algorithms, the impact of the target parameter specification on precision and true power was notable. As predicted by theory, the variance was lowest and thereby true power highest for the SATE. Consider, for example, the TMLE using logistic regression and a sample of 50 units. The standard deviation σ of this estimator of the population effect was 29% higher than the standard deviation of this estimator of the conditional effect and 53% higher than the standard deviation of this estimator of the sample effect. Furthermore, if we knew the variance of this estimator, then we would have 78% power to detect the sample effect, 71% power to detect the conditional effect and only 55% power to detect the population effect.

In practice, however, we must estimate the variance. When our target of inference is the SATE or the CATE, the sample variance of weighted residuals (Eq. 2), divided by sample size n , provides an asymptotically conservative variance estimator. When our target of inference is the PATE, we must also account for estimation of the covariate distribution. In the finite sample simulations, the impact of having a conservative variance estimator on inference for SATE was considerable. In all settings, the standard deviation was over-estimated: $\hat{\sigma} > \sigma$. As a result, the attained power was less than the true power and the confidence interval coverage was conservative (i.e. greater than the nominal rate of 95%). Likewise, when the CATE was the target of inference, the standard deviation was conservatively approximated and thereby the attained power was less than the true power. For both the sample and conditional effects, the TMLE using Super Learner was able to obtain a more precise fit of $\bar{Q}_0(A, W)$ and thereby a less conservative variance estimator. As a result, this TMLE was able to achieve the most power. We note that the attained power is the same for the SATE and CATE, because we used the same point and variance estimator for both parameters.

Despite the conservative variance estimator, the TMLE for the SATE or CATE achieved higher power than the TMLE for the PATE at all sample sizes. With 50 units, for example, the attained power for the TMLE with Super Learner was 66% for the sample/conditional effect and only 57% for the population effect. Notably, the attained power was the same for the unadjusted estimator of the 3 parameters. The attained power of the unadjusted estimator did not vary, because the estimated D_W component of influence curve and thereby its variance were zero:

$$\bar{Q}_n(1) - \bar{Q}_n(0) - \Psi_{n,unadj}(P_n) = 0$$

where $\bar{Q}_n(A)$ denotes the treatment-specific mean. Thus, using the unadjusted estimator sacrificed any potential gains in attained power by specifying the SATE or the CATE as the target of inference.



		Variability			Power		CI
		Bias	σ	$\hat{\sigma}$	True	Att.	Cover.
SATE							
$n = 50$	Unadj	1.1E-4	1.3E-2	1.6E-2	0.57	0.41	0.98
	TMLE	7.7E-4	8.8E-3	1.2E-2	0.78	0.63	0.99
	TMLE+SL	5.3E-4	8.8E-3	1.1E-2	0.78	0.66	0.98
$n = 70$	Unadj	2.7E-4	1.1E-2	1.4E-2	0.69	0.52	0.99
	TMLE	6.4E-4	7.2E-3	1.0E-2	0.89	0.75	0.99
	TMLE+SL	4.5E-4	7.2E-3	9.7E-3	0.88	0.78	0.99
$n = 100$	Unadj	3.5E-5	9.0E-3	1.1E-2	0.80	0.66	0.98
	TMLE	2.7E-4	5.9E-3	8.7E-3	0.96	0.87	0.99
	TMLE+SL	1.4E-4	6.0E-3	8.2E-3	0.95	0.89	0.99
CATE							
$n = 50$	Unadj	3.3E-5	1.4E-2	1.6E-2	0.51	0.41	0.97
	TMLE	6.9E-4	1.1E-2	1.2E-2	0.71	0.63	0.97
	TMLE+SL	4.5E-4	1.0E-2	1.1E-2	0.71	0.66	0.95
$n = 70$	Unadj	2.0E-4	1.2E-2	1.4E-2	0.62	0.52	0.97
	TMLE	5.8E-4	8.7E-3	1.0E-2	0.83	0.75	0.97
	TMLE+SL	3.9E-4	8.7E-3	9.7E-3	0.82	0.78	0.96
$n = 100$	Unadj	8.2E-6	9.7E-3	1.1E-2	0.76	0.66	0.97
	TMLE	2.4E-4	7.0E-3	8.7E-3	0.93	0.87	0.98
	TMLE+SL	1.1E-4	7.1E-3	8.2E-3	0.93	0.89	0.97
PATE							
$n = 50$	Unadj	9.5E-5	1.6E-2	1.6E-2	0.40	0.41	0.94
	TMLE	7.5E-4	1.4E-2	1.3E-2	0.55	0.58	0.94
	TMLE+SL	5.1E-4	1.4E-2	1.3E-2	0.54	0.57	0.93
$n = 70$	Unadj	-8.1E-6	1.4E-2	1.4E-2	0.52	0.52	0.94
	TMLE	3.7E-4	1.1E-2	1.1E-2	0.69	0.70	0.94
	TMLE+SL	1.8E-4	1.1E-2	1.1E-2	0.68	0.69	0.94
$n = 100$	Unadj	-1.3E-4	1.1E-2	1.1E-2	0.67	0.66	0.95
	TMLE	9.8E-5	9.0E-3	9.2E-3	0.86	0.85	0.95
	TMLE+SL	-2.8E-5	9.1E-3	9.3E-3	0.84	0.83	0.95

Table 2: Summary of estimator performance over 2,500 simulations. The rows denote the sample sizes $n = \{50, 70, 100\}$ and the estimator: unadjusted, TMLE with logistic regression and TMLE with Super Learner (“TMLE+SL”). Bias is the average deviation between the point estimate and (sample-specific) true value; σ is the square root of the variance of the estimator, and $\hat{\sigma}$ is the average standard error estimate, based on the influence curve. The true power (“Power True”) is the proportion of times the false null hypothesis would be rejected if the estimator’s variance σ^2 were known, while the attained power (“Power Att.”) is the proportion of times the false null hypothesis was rejected when estimating the variance. The confidence interval coverage (“CI Cover.”) is the proportion the 95% confidence intervals that contained the true parameter value.

6 Discussion

To our knowledge, this is the first paper to propose using TMLE for estimation and inference of the SATE. Despite lack of identifiability, we proved that the TMLE was an asymptotically linear estimator of the SATE. If there is heterogeneity in the intervention effect within strata of covariates, the sample parameter will be estimated with more precision than the conditional parameter. Furthermore, if there is heterogeneity in the intervention effect across strata of covariates, the sample and conditional parameters will be estimated with more precision than the population parameter. In practice, we can estimate the variance of the TMLE for the SATE with the sample variance of the weighted residuals, divided by sample size. This is an intuitive variance estimator and straightforward to implement. We also discussed generalizations to an observational setting. The TMLE will provide at least as much precision and power to detect the impact of a non-randomized exposure on the study units than in the overall population. Our conclusions should also extend to a trial with adaptive pair-matching (Balzer et al., 2015). Formal study is warranted and an area of future work, but we hypothesize that a trial targeting the sample effect and implementing adaptive pair-matching will be more efficient than a trial targeting the sample effect and not implementing pair-matching.

Finite sample simulations highlighted the differences between the causal parameters and the impact of the target parameter specification on the estimator's variance and attained power. We also compared the unadjusted estimator (i.e. difference-in-means estimator) to the TMLE with various methods for initial estimation of $\bar{Q}_0(A, W)$. As predicted by theory, all estimators were unbiased and adjustment lead to greater power. An estimator of the SATE was less variable than the same estimator of the CATE, which was less variable the same estimator of the PATE. While the differences in the true power (the proportion of times the false null hypothesis would be rejected if we knew the estimator's variance) were substantial, the difference in the attained power were attenuated due to the conservative variance estimator. Greater differences in the attained power were seen with a more aggressive fit of conditional mean outcome. As estimation of $\bar{Q}_0(A, W)$ improves, the TMLE becomes a more precise estimator (i.e. smaller true variance) and the variance estimator becomes less conservative. In small trials (e.g. $n \leq 30$) such as early phase clinical trials or cluster randomized trials, obtaining a precise estimate of $\bar{Q}_0(A, W)$ is likely to be challenging. In practice, many baseline covariates are predictive of the outcome, but adjusting for too many covariates can result in over-fitting. Future work will investigate the use of cross-validation to data-adaptively select the optimal adjustment set in trials with limited sample sizes.

Overall, we believe the SATE is an under-utilized causal parameter. It is simply the intervention effect for the study units. The SATE avoids assumptions about representative sampling (e.g. a simple random sample) from some target population. Furthermore, the SATE is responsive to heterogeneity in the treatment effect and avoids assumptions that the observed impact is generalizable to other contexts (Stuart et al., 2011; Hartman et al., 2015). These generalizations can be made with the formal methods for transportability (Pearl and Bareinboim, 2013) and do not have to be assumed during the parameter specification. To obtain a point estimate, the implementation of the TMLE is identical to that of the conditional and population estimands. To obtain conservative inference, we only need to take the sample variance of the weighted residuals, divided by sample size. Thereby, estimation and inference for the SATE does not require any extra work and is likely to give us more power to detect the impact of the exposure on the outcome.

Appendix A: The TMLE is an asymptotically linear estimator of the SATE

Consider the statistical parameter corresponding to the population average treatment effect (PATE):

$$\begin{aligned}\Psi_0^{\mathcal{P}}(P_0) &= E_0[E_0(Y|A=1, W) - E_0(Y|A=0, W)] \\ &= E_0[\bar{Q}_0(1, W) - \bar{Q}_0(0, W)]\end{aligned}$$

where $\bar{Q}_0(A, W) = E_0(Y|A, W)$ denotes the conditional expectation of the outcome, given the exposure and covariates. The TMLE for $\Psi_0^{\mathcal{P}}(P_0)$ is defined by the following substitution estimator:

$$\Psi_n(P_n) = \frac{1}{n} \sum_{i=1}^n [\bar{Q}_n^*(1, W_i) - \bar{Q}_n^*(0, W_i)]$$

where P_n denotes the empirical distribution, putting mass $1/n$ on each $O_i = (W_i, A_i, Y_i)$ and $\bar{Q}_n^*(A, W)$ denotes the targeted estimator.

Suppose the exposure mechanism, denoted $g_0(A|W) = P_0(A|W)$, is known as in a randomized trial. Under the following regularity conditions, the TMLE of $\Psi_0^{\mathcal{P}}(P_0)$ is asymptotically linear (van der Laan and Rubin, 2006):

$$\Psi_n(P_n) - \Psi_0^{\mathcal{P}}(P_0) = \frac{1}{n} \sum_{i=1}^n D^{\mathcal{P}}(\bar{Q}, g_0)(O_i) + o_P(1/\sqrt{n})$$

with influence curve

$$D^{\mathcal{P}}(\bar{Q}, g_0)(O) = \left(\frac{\mathbb{I}(A=1)}{g_0(1|W)} - \frac{\mathbb{I}(A=0)}{g_0(0|W)} \right) (Y - \bar{Q}(A, W)) + \bar{Q}(1, W) - \bar{Q}(0, W) - \Psi_0^{\mathcal{P}}(P_0)$$

where $\bar{Q}(A, W)$ denotes the limit of the TMLE $\bar{Q}_n^*(A, W)$. Specifically, we assume the positivity assumption holds: for some $\delta > 0$, $\delta < g_0(1|W) < 1 - \delta$. We also assume that $P_0[D_n^{\mathcal{P}}(\bar{Q}_n^*, g_0) - D^{\mathcal{P}}(\bar{Q}, g_0)]^2 \rightarrow 0$ in probability and that $D_n^{\mathcal{P}}(\bar{Q}_n^*, g_0)$ is in the P_0 -Donsker class with probability tending to 1. Here we used notation $P_0 f = \int f(o) dP_0(o)$ for some function f .

Theorem 1. *Suppose we have n i.i.d. observations of random variable $O = (W, A, Y) \sim P_0$, where W denotes the baseline covariates, A denotes the exposure, and Y denotes the outcome. Consider the sample average treatment effect (SATE) $\Psi^{\mathcal{S}}(P_{U,X}) = \frac{1}{n} \sum_{i=1}^n Y_i(1) - Y_i(0)$, where $P_{U,X}$ denote the joint distribution of the background factors $U = (U_W, U_A, U_Y)$ and exogenous factors $X = (W, A, Y)$. Under the above regularity conditions, the TMLE $\Psi_n(P_n) = \frac{1}{n} \sum_{i=1}^n \bar{Q}_n^*(1, W_i) - \bar{Q}_n^*(0, W_i)$ is an asymptotically linear estimator of the SATE:*

$$\Psi_n(P_n) - \Psi^{\mathcal{S}}(P_{U,X}) = \frac{1}{n} \sum_{i=1}^n D^{\mathcal{S}}(U_i, X_i) + o_P(1/\sqrt{n})$$

with influence curve

$$\begin{aligned}D^{\mathcal{S}}(U, X) &= D^{\mathcal{C}}(\bar{Q}, g_0)(O) - D^{\mathcal{F}}(U, X) \\ D^{\mathcal{C}}(\bar{Q}, g_0)(O) &= \left(\frac{\mathbb{I}(A=1)}{g_0(1|W)} - \frac{\mathbb{I}(A=0)}{g_0(0|W)} \right) (Y - \bar{Q}(A, W)) - [\bar{Q}_0(W) - \bar{Q}(W)] \\ D^{\mathcal{F}}(U, X) &= Y(1) - Y(0) - [\bar{Q}_0(1, W) - \bar{Q}_0(0, W)]\end{aligned}$$

where $\bar{Q}(W) = \bar{Q}(1, W) - \bar{Q}(0, W)$ denotes the difference in the treatment-specific means.

We note that D^C is the influence curve of the TMLE for the conditional estimand $\Psi_0^C(P_0) = \frac{1}{n} \sum_{i=1}^n \bar{Q}_0(1, W_i) - \bar{Q}_0(0, W_i)$, which corresponds to the conditional average treatment effect (CATE) under the necessary causal assumptions (Balzer et al., 2015). The remaining non-identifiable piece D^F is difference between the unit-specific effect and the effect within strata of covariates.

Proof. Let $\bar{Q}_0(W) = \bar{Q}_0(1, W) - \bar{Q}_0(0, W)$ denote the true difference in treatment-specific means. We can write the deviation between the TMLE $\Psi_n(P_n)$ for the population estimand $\Psi_0^P(P_0)$ and the SATE as

$$\begin{aligned}
 \Psi_n(P_n) - \Psi^S(P_{U,X}) &= \Psi_n(P_n) - \Psi_0^P(P_0) - [\Psi^S(P_{U,X}) - \Psi_0^P(P_0)] \\
 &= \frac{1}{n} \sum_{i=1}^n D^P(O_i) - [\Psi^S(P_{U,X}) - \Psi_0^P(P_0)] + o_P(1/\sqrt{n}) \\
 &= \frac{1}{n} \sum_{i=1}^n D^P(O_i) - \left[\frac{1}{n} \sum_{i=1}^n Y_i(1) - Y_i(0) - \bar{Q}_0(W_i) + \bar{Q}_0(W_i) - \Psi_0^P(P_0) \right] + o_P(1/\sqrt{n}) \\
 &= \frac{1}{n} \sum_{i=1}^n \left(\frac{\mathbb{I}(A_i = 1)}{g_0(1|W_i)} - \frac{\mathbb{I}(A_i = 0)}{g_0(0|W_i)} \right) (Y_i - \bar{Q}(A_i, W_i)) + \bar{Q}(W_i) - \Psi_0^P(P_0) \\
 &\quad - \left[\frac{1}{n} \sum_{i=1}^n Y_i(1) - Y_i(0) - \bar{Q}_0(W_i) + \bar{Q}_0(W_i) - \Psi_0^P(P_0) \right] + o_P(1/\sqrt{n}) \\
 &= \frac{1}{n} \sum_{i=1}^n \left(\frac{\mathbb{I}(A_i = 1)}{g_0(1|W_i)} - \frac{\mathbb{I}(A_i = 0)}{g_0(0|W_i)} \right) (Y_i - \bar{Q}(A_i, W_i)) - [\bar{Q}_0(W_i) - \bar{Q}(W_i)] \\
 &\quad - \left[\frac{1}{n} \sum_{i=1}^n Y_i(1) - Y_i(0) - \bar{Q}_0(W_i) \right] + o_P(1/\sqrt{n}) \\
 &= \frac{1}{n} \sum_{i=1}^n D^C(O_i) - D^F(U_i, X_i) + o_P(1/\sqrt{n})
 \end{aligned}$$

where the influence curve of the TMLE for the conditional estimand $\Psi_0^C(P_0)$ is

$$D^C(\bar{Q}, g_0)(O) = \left(\frac{\mathbb{I}(A = 1)}{g_0(1|W)} - \frac{\mathbb{I}(A = 0)}{g_0(0|W)} \right) (Y - \bar{Q}(A, W)) - [\bar{Q}_0(W) - \bar{Q}(W)]$$

and where

$$D^F(U, X) = Y(1) - Y(0) - \bar{Q}_0(1, W) - \bar{Q}_0(0, W)$$

Thus, we have shown the TMLE is an asymptotically linear estimator of the SATE:

$$\sqrt{n} \left[\Psi_n(P_n) - \Psi^S(P_{U,X}) \right] = \frac{1}{n} \sum_{i=1}^n D^S(U_i, X_i) + o_P(1/\sqrt{n})$$

with influence curve

$$D^S(U, X) = D^C(O) - D^F(U, X)$$

□

Appendix B: Variance and variance estimation for the TMLE of the SATE

Theorem 2. *The standardized TMLE for the SATE is asymptotically normal:*

$$\begin{aligned} \sqrt{n} \left[\Psi_n(P_n) - \Psi^S(P_{U,X}) \right] &\xrightarrow{D} N(0, \sigma^{2,S}) \\ \text{with } \sigma^{2,S} &= \text{Var}[D^C] + \text{Var}[D^F] - 2\text{Cov}[D^C, D^F] \\ &= \text{Var}[D^C] - \text{Var}[D^F] \end{aligned}$$

Proof. The covariance term is

$$\begin{aligned} \text{Cov}[D^C, D^F] &= E_{U,X,0}[D^C \times D^F] \\ &= E_{U,X,0} \left[\left\{ \left(\frac{\mathbb{I}(A=1)}{g_0(1|W)} - \frac{\mathbb{I}(A=0)}{g_0(0|W)} \right) (Y - \bar{Q}(A, W)) - (\bar{Q}_0(W) - \bar{Q}(W)) \right\} \right. \\ &\quad \left. \times \{Y(1) - Y(0) - \bar{Q}_0(W)\} \right] \\ &= E_{U,X,0} \left[\left(\frac{\mathbb{I}(A=1)}{g_0(1|W)} - \frac{\mathbb{I}(A=0)}{g_0(0|W)} \right) (Y - \bar{Q}(A, W)) \times \{Y(1) - Y(0) - \bar{Q}_0(W)\} \right] \\ &\quad - E_{U,X,0} \left[\{ \bar{Q}_0(W) - \bar{Q}(W) \} \times \{Y(1) - Y(0) - \bar{Q}_0(W)\} \right] \end{aligned}$$

Under the randomization assumption, the D^F component has conditional mean zero, given the baseline covariates W :

$$\begin{aligned} E_{U,X,0}[Y(1) - Y(0) - \bar{Q}_0(W)|W] &= E_{U,X,0}[Y(1)|W] - E_{U,X,0}[Y(0)|W] - \bar{Q}_0(W) \\ &= E_0(Y|A=1, W) - E_0(Y|A=0, W) - \bar{Q}_0(W) = 0 \end{aligned}$$

Therefore, the second term is zero.

For the first term, we have

$$\begin{aligned} &E_{U,X,0} \left[\left(\frac{\mathbb{I}(A=1)}{g_0(1|W)} - \frac{\mathbb{I}(A=0)}{g_0(0|W)} \right) (Y - \bar{Q}(A, W)) \times [Y(1) - Y(0) - \bar{Q}_0(W)] \right] \\ &= E_{U,X,0} \left[\left(\frac{\mathbb{I}(A=1)}{g_0(1|W)} - \frac{\mathbb{I}(A=0)}{g_0(0|W)} \right) (Y - \bar{Q}(A, W) + \bar{Q}_0(A, W) - \bar{Q}_0(A, W)) \times [Y(1) - Y(0) - \bar{Q}_0(W)] \right] \\ &= E_{U,X,0} \left[\left(\frac{\mathbb{I}(A=1)}{g_0(1|W)} - \frac{\mathbb{I}(A=0)}{g_0(0|W)} \right) (Y - \bar{Q}_0(A, W)) \times [Y(1) - Y(0) - \bar{Q}_0(W)] \right] \\ &+ E_{U,X,0} \left[\left(\frac{\mathbb{I}(A=1)}{g_0(1|W)} - \frac{\mathbb{I}(A=0)}{g_0(0|W)} \right) (\bar{Q}_0(A, W) - \bar{Q}(A, W)) \times [Y(1) - Y(0) - \bar{Q}_0(W)] \right] \end{aligned}$$

It follow that this equals:

$$\begin{aligned}
& E_{U,X,0} \left[\frac{I(A=1)}{g_0(1|W)} (Y(1) - \bar{Q}_0(1, W)) \times [Y(1) - Y(0) - \bar{Q}_0(W)] \right] \\
& - E_{U,X,0} \left[\frac{I(A=0)}{g_0(0|W)} (Y(0) - \bar{Q}_0(0, W)) \times [Y(1) - Y(0) - \bar{Q}_0(W)] \right] \\
& + E_{U,X,0} \left[\frac{I(A=1)}{g_0(1|W)} (\bar{Q}_0(1, W) - \bar{Q}(1, W)) \times [Y(1) - Y(0) - \bar{Q}_0(W)] \right] \\
& - E_{U,X,0} \left[\frac{I(A=0)}{g_0(0|W)} (\bar{Q}_0(0, W) - \bar{Q}(0, W)) \times [Y(1) - Y(0) - \bar{Q}_0(W)] \right]
\end{aligned}$$

Under the randomization assumption, we have

$$E_0 \left[\frac{I(A=a)}{g_0(a|W)} \middle| Y(1), Y(0), W \right] = 1$$

Therefore, the sum of first two terms reduce to the variance of the D^F component:

$$E_0 \left[[Y(1) - Y(0) - \bar{Q}_0(W)] \times [Y(1) - Y(0) - \bar{Q}_0(W)] \right] = E_0 \left[[Y(1) - Y(0) - \bar{Q}_0(W)]^2 \right]$$

The sum of last two terms equals zero, using that the conditional mean of D^F component, given W equals zero. Therefore, we have that the covariance term equals the variance of the non-identifiable component D^F :

$$2Cov[D^C, D^F] = 2Var[D^F]$$

Thus, the asymptotic variance of the standardized estimator for the SATE is

$$\sigma^{2,S} = Var[D^C] - Var[D^F]$$

□

The asymptotic variance $\sigma^{2,S}$ is always less than or equal to $Var[D^C]$. We can estimate the upper bound as

$$\sigma_n^{2,S} = Var_n[D_n^C]$$

where Var_n is the sample variance and D_n^C is the (conservative) estimate of the influence curve for the TMLE for the conditional parameter:

$$D_n^C(O_i) = \left(\frac{\mathbb{I}(A_i=1)}{g_0(1|W_i)} - \frac{\mathbb{I}(A_i=0)}{g_0(0|W_i)} \right) (Y_i - \bar{Q}_n^*(A_i, W_i))$$

Appendix C: Generalization to allow for estimation of the exposure mechanism

Suppose our target of inference is the population estimand $\Psi_0^P(P_0)$ and the exposure mechanism is consistently estimated with maximum likelihood: $g_n(A|W)$. Then the TMLE is asymptotically linear with influence curve given by the influence curve at the possibly misspecified limit $\bar{Q}(A, W)$ minus its projection on the tangent space T_g of the model for $g_0(A|W)$ (van der Laan and Robins, 2003):

$$D^{\mathcal{P}, g_n}(\bar{Q}, g_0) = D^{\mathcal{P}}(\bar{Q}, g_0) - \prod [D^{\mathcal{P}}(\bar{Q}, g_0) | T_g]$$

This projection is a function of (A, W) with conditional mean zero, given W . Analogously, when we target the conditional estimand $\Psi_0^C(P_0)$, the influence curve of the TMLE is

$$D^{C,g_n}(\bar{Q}, g_0) = D^C(\bar{Q}, g_0) - \prod [D^C(\bar{Q}, g_0)|T_g]$$

and when we target the SATE $\Psi^S(P_{U,X})$, the influence curve of the TMLE is

$$D^{S,g_n}(\bar{Q}, g_0) = D^{C,g_n}(\bar{Q}, g_0) - D^F$$

The proof is analogous to the above and thus omitted.

The standardized estimator of the SATE then is asymptotically normal with mean 0 and variance given by the variance of influence curve:

$$\sigma^{2,S,g_n} = Var[D^{C,g_n}] + Var[D^F] - 2Cov[D^{C,g_n}, D^F]$$

The covariance of the projection $\prod [D^C(\bar{Q}, g_0)|T_g]$ and D^F is zero. (If we take the expectation given (A, W) , then the projection term is constant and the D^F term is zero.) Thus, the asymptotic variance of standardized estimator (when the exposure mechanism is estimated according to a correctly specified model) is

$$\sigma^{2,S,g_n} = Var[D^{C,g_n}] - Var[D^F]$$

We will to have a conservative variance estimator by ignoring the projection term and the non-identifiable piece D^F .

7 Acknowledgements

The authors would like to acknowledge and thank the entire SEARCH collaboration for their helpful comments and discussion. This work was supported, in part, by the National Institute Of Allergy And Infectious Diseases of the National Institutes of Health under award numbers R01-AI074345 and UM1AI069502. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. Maya Petersen is a recipient of a Doris Duke Clinical Scientist Development Award.

References

- A. Abadie and G. Imbens. Simple and bias-corrected matching estimators for average treatment effects. Technical Report 283, NBER technical working paper, 2002.
- A. Abadie and G.W. Imbens. Estimation of the conditional variance in paired experiments. *Annales d'Economie et de Statistique*, 91/92:175–187, 2008.
- J. Ahern, A. Hubbard, and S. Galea. Estimating the effects of potential public health interventions on population disease burden: A step-by-step illustration of causal inference methods. *American Journal of Epidemiology*, 169(9):1140–1147, 2009.
- L. Balzer, J. Ahern, S. Galea, and M.J. van der Laan. Estimating Effects on Rare Outcomes: Knowledge is Power. *Under Review*, 2014. Available at <http://biostats.bepress.com/ucbbiostat/paper310/>.

- L.B. Balzer, M.L. Petersen, M.J. van der Laan, and the SEARCH Consortium. Adaptive pair-matching in the SEARCH trial and estimation of the intervention effect. *Statistics in Medicine*, 34(6):999–1011, 2015. doi: 10.1002/sim.6380.
- P.J. Bickel, C.A.J. Klaassen, Y. Ritov, and J.A. Wellner. *Efficient and Adaptive Estimation for Semiparametric Models*. Johns Hopkins University Press, Baltimore, 1993.
- W.G. Cochran. Analysis of covariance: its nature and uses. *Biometrics*, 13:261–281, 1957.
- D.R. Cox and P. McCullagh. Some aspects of analysis of covariance. *Biometrics*, 38(3):541–561, 1982.
- R.A. Fisher. *Statistical methods for research workers*. Oliver and Boyd Ltd., Edinburgh, 4th edition, 1932.
- S. Gruber and M.J. van der Laan. A targeted maximum likelihood estimator of a causal effect on a bounded continuous outcome. *The International Journal of Biostatistics*, 6(1):Article 26, 2010. doi: 10.2202/1557-4679.1260.
- S. Gruber and M.J. van der Laan. tmle: An R package for targeted maximum likelihood estimation. *Journal of Statistical Software*, 51(13):1–35, 2012.
- E. Hartman, R. Grieve, R. Ramsahai, and J.S. Sekhon. From sample average treatment effect to population average treatment effect on the treated: combining experimental with observational studies to estimate population treatment effects. *Journal of the Royal Statistical Society: Series A*, pages 2–22, Jan 2015.
- M.A. Hernán, E. Lanoy, D. Costagliola, and J.M. Robins. Comparison of dynamic treatment regimes via inverse probability weighting. *Basic Clin Pharmacol Toxicol*, 98(3):237–242, 2006.
- D.G. Horvitz and D.J. Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association*, 47:663–685, 1952.
- K. Imai. Variance identification and efficiency analysis in randomized experiments under the matched-pair design. *Statistics in Medicine*, 27(24):4857–4873, 2008. doi: 10.1002/sim.3337.
- G.W. Imbens. Nonparametric estimation of average treatment effects under exogeneity: a review. *Review of Economics and Statistics*, 86(1):4–29, 2004. doi: 10.1162/003465304323023651.
- G.W. Imbens. Experimental design for unit and cluster randomized trials. Technical report, NBER technical working paper, 2011.
- K.L. Moore and M.J. van der Laan. Covariate adjustment in randomized trials with binary outcomes: Targeted maximum likelihood estimation. *Statistics in Medicine*, 28(1):39–64, 2009. doi: 10.1002/sim.3445.
- J. Neyman. Sur les applications de la theorie des probabilités aux expériences agricoles: Essai des principes (In Polish). English translation by D.M. Dabrowska and T.P. Speed (1990). *Statistical Science*, 5:465–480, 1923.
- J. Pearl. Causal diagrams for empirical research. *Biometrika*, 82:669–710, 1995.
- J. Pearl. *Causality: Models, Reasoning and Inference*. Cambridge University Press, New York, 2000. second ed., 2009.

- J. Pearl and E. Bareinboim. Transportability across studies: A formal approach. Technical Report R-372, Computer Science Department, University of California, Los Angeles, 2013. Available at ftp://www.mmss.cs.ucla.edu/pub/stat_ser/r372.pdf.
- R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2014. URL <http://www.R-project.org>.
- 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. *Math Mod*, 7:1393–1512, 1986.
- J.M. Robins. Analytic methods for estimating hiv treatment and cofactor effects. In D.G. Ostrow and R. Kessler, editors, *Methodological issues of AIDS Mental Health Research*, pages 213–290. Plenum Publishing, New York, 1993.
- J.M. Robins, A. Rotnitzky, and L.P. Zhao. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association*, 89:846–866, 1994.
- J.M. Robins, M.A. Hernán, and B. Brumback. Marginal structural models and causal inference in epidemiology. *Epidemiology*, 11(5):550–560, 2000.
- P.R. Rosenbaum and D.B. Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70:41–55, 1983.
- M. Rosenblum and M.J. van der Laan. Simple, efficient estimators of treatment effects in randomized trials using generalized linear models to leverage baseline variables. *The International Journal of Biostatistics*, 6(1):Article 13, 2010. doi: 10.2202/1557-4679.1138.
- A. Rotnitzky, J.M. Robins, and D.O. Scharfstein. Semiparametric regression for repeated outcomes with nonignorable nonresponse. *Journal of the American Statistical Association*, 93:1321–1339, 1998.
- Donald B. Rubin. Comment: Neyman (1923) and causal inference in experiments and observational studies. *Stat Sci*, 5(4):472–480, 1990.
- J. Schwab, S. Lendle, M. Petersen, and M. van der Laan. *ltmle: Longitudinal Targeted Maximum Likelihood Estimation*, 2014. R package version 0.9.3-1. Available at <http://CRAN.R-project.org/package=ltmle>.
- J.M. Snowden, S. Rose, and K.M. Mortimer. Implementation of g-computation on a simulated data set: demonstration of a causal inference technique. *Am J Epidemiol*, 173(7):731–738, 2011.
- E.A. Stuart, S.R. Cole, C.P. Bradshaw, and P.J. Leaf. The use of propensity scores to assess the generalizability of results from randomized trials. *J.R. Statist. Soc. A*, 174(Part 2):369–386, 2011.
- A.A. Tsiatis, M. Davidian, M. Zhang, and X. Lu. Covariate adjustment for two-sample treatment comparisons in randomized clinical trials: A principled yet flexible approach. *Stat Med*, 27(23):4658–4677, 2008. doi: 10.1002/sim.3113.
- University of California, San Francisco. Sustainable East Africa Research in Community Health (SEARCH). ClinicalTrials.gov, 2013. Available at <http://clinicaltrials.gov/show/NCT01864603>.

- M. van der Laan and S. Rose. *Targeted Learning: Causal Inference for Observational and Experimental Data*. Springer, Berlin Heidelberg New York, 2011.
- M.J. van der Laan and J.M. Robins. *Unified Methods for Censored Longitudinal Data and Causality*. Springer-Verlag, New York Berlin Heidelberg, 2003.
- M.J. van der Laan and D.B. Rubin. Targeted maximum likelihood learning. *The International Journal of Biostatistics*, 2(1):Article 11, 2006. doi: 10.2202/1557-4679.1043.
- M.J. van der Laan, E.C. Polley, and A.E. Hubbard. Super learner. *Statistical Applications in Genetics and Molecular Biology*, 6(1):25, 2007.

