UC Berkeley
UC Berkeley Electronic Theses and Dissertations

Title

Design and Analysis of Cluster Randomized Trials with Application to HIV Prevention and Treatment

Permalink

https://escholarship.org/uc/item/7nv62814

Author

Balzer, Laura

Publication Date

2015

Peer reviewed|Thesis/dissertation

Design and Analysis of Cluster Randomized Trials
with Application to HIV Prevention and Treatment

by

Laura Balzer

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Biostatistics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Mark Van Der Laan, Co-chair
Professor Maya Petersen, Co-chair
Professor Jennifer E Ahern

Spring 2015

Design and Analysis of Cluster Randomized Trials with Application to HIV Prevention and Treatment

Abstract

Design and Analysis of Cluster Randomized Trials
with Application to HIV Prevention and Treatment

by

Laura Balzer

Doctor of Philosophy in Biostatistics

University of California, Berkeley

Professor Mark Van Der Laan, Co-chair

Professor Maya Petersen, Co-chair

This dissertation is focused on the development of the optimal design and analysis for cluster randomized trials. Specifically, we tackle three common questions: whether or not to pair-match clusters, which causal parameter best captures the intervention effect, and how to select the adjustment set for the analysis. We begin by introducing a formal framework for causal inference in Chapter 1. Throughout, the Sustainable East Africa Research in Community Health (SEARCH) trial serves as the motivating example (NCT01864603). SEARCH is an ongoing community randomized trial to evaluate the impact of immediate and streamlined antiretroviral therapy on HIV incidence in rural East Africa.

In Chapter 2, we consider pair-matching, an intuitive design strategy to protect study validity and to potentially increase power in randomized trials. In a common design, candidate units are identified, and their baseline characteristics are used to create the best $n/2$ matched pairs. Within the resulting pairs, the intervention is randomized, and the outcomes are measured at the end of follow-up. We consider this design to be adaptive, because the construction of the matched pairs depends on the baseline covariates of all candidate units. As a consequence, the observed data cannot be considered as $n/2$ independent, identically distributed (i.i.d.) pairs of units, as common practice assumes. Instead, the observed data consist of n dependent units. Chapter 2 explores the consequences of adaptive pair-matching in randomized trials for estimation of the conditional average treatment effect (CATE): the intervention effect, given the measured covariates of the n study units. We contrast the unadjusted estimator with TMLE and show substantial efficiency gains from matching and further gains with adjustment.

In Chapter 3, we compare three causal parameters: the population, conditional and sample average treatment effects. Using a structural causal model, we explicitly define each parameter, discuss interpretation, and formally examine identifiability. To the best of our knowledge, Chapter 3 is the first to propose using TMLE for estimation and inference of the sample effect. In most settings, the sample parameter will be estimated more efficiently

than the conditional parameter, which will, in turn, 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.

Finally in Chapter 4, we discuss adjustment for measured covariates during the analysis to reduce variance and increase power in randomized trials. To avoid misleading inference, the analysis plan must be pre-specified. However, it is often unclear *a priori* which baseline covariates (if any) should be included in the analysis. In the SEARCH trial, for example, there are 16 matched pairs of communities and many potential adjustment variables, including region, HIV prevalence, male circumcision coverage and measures of community-level viral load. In Chapter 4, we propose a rigorous procedure to data-adaptively select the adjustment set, which maximizes the efficiency of the analysis. Specifically, we use cross-validation to select from a pre-specified library the candidate TMLE that minimizes the estimated variance. For further gains in precision, we also propose a collaborative procedure for estimating the known exposure mechanism. Our small sample simulations demonstrate the promise of the methodology to maximize study power, while maintaining nominal confidence interval coverage. Our procedure is tailored to the scientific question (sample vs. population treatment effect) and study design (pair-matched or not) and alleviates many of the common concerns.

To Pops and Nick - Thank you.

Contents

Acknowledgments

Above all, my favorite thing about U.C. Berkeley is the people. From day one, I was struck by the passion and compassion of my professors and colleagues. My Ph.D. would not have been possible without the unfaltering guidance and mentorship of the following people (to name a few):

- Mark van der Laan, whose patience and love of learning know no bounds

- Maya Petersen, who first introduced me to causal inference and has since become my role model

- Jennifer Ahern, who tirelessly helps me bridge methods and application

- The SEARCH team, who are always striving for the best

- The Berkeley Biostat and Epi family for supporting me through personal and professional tragedies and triumphs

- My family for the love and laughs

Chapter 1

Introduction: A Roadmap for Causal Inference

This chapter introduces the methods, as well as their limitations, to learn causal relationships from data. Consider, for example, the following questions:

1. What proportion of patients taking drug X suffered adverse side effects?
2. Which patients taking drug X are more likely to suffer adverse side effects?
3. Would the risk of adverse effects be lower if all patients took drug X instead of drug Y?

The first question is purely descriptive; the second can be characterized as a prediction problem, while the last is causal. Causal inference is distinct from statistical inference in that it seeks to make conclusions about the world under changed conditions [1]. In the third example, our goal is to make inferences about how the distribution of patient outcomes would differ if all patients had taken drug X vs. if the same patients, over the same time frame and under the same conditions, had taken drug Y. Purely statistical analyses are sometimes endowed with causal interpretations. Furthermore, many of our non-causal questions have causal elements. For example, Geng *et al.* [2] sought to assess whether sex was an independent predictor of mortality among patients initiating drug therapy (i.e. describe a non-causal association) but in the absence of loss to follow up (i.e. a change to the existing conditions).

In this chapter, we review a formal framework for causal inference to (1) state the scientific question, (2) express our causal knowledge and limits of that knowledge, (3) specify the causal parameter, (4) specify the observed data and their link to the causal model, (5) assess identifiability of our causal parameter as some function of the observed data distribution, (6) estimate the corresponding statistical parameter, and (7) interpret our results [3–5]. As illustrated in Figure 1.1, there are many sources of association between two variables, including direct effects, indirect effects, measured confounding, unmeasured confounding and selection bias [6]. Methods to delineate causation from correlation are perhaps more pressing now than ever [7, 8].
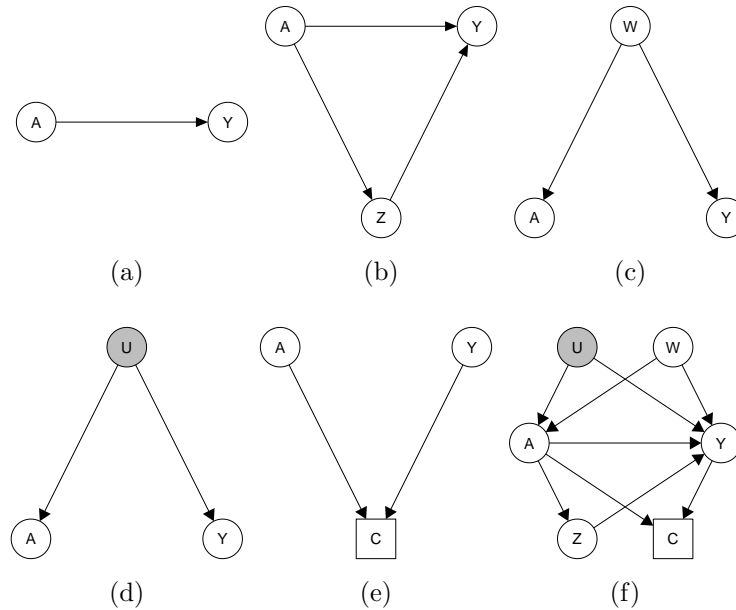
Figure 1.1: Some of the sources of dependence between an exposure A and an outcome Y: (a) the exposure A directly affects the outcome Y; (b) the exposure A directly affects the outcome Y as well as indirectly affects it through the mediator Z; (c) the exposure A has no effect on the outcome Y, but an association is induced by a measured common cause W; (d) the exposure A has no effect on the outcome Y, but an association is induced by an unmeasured common cause U; (e) the exposure A has no effect on the outcome Y, but an association is induced by only examining data among those not censored C; (f) all these sources of dependence are present. Please note this not *an exhaustive list.*

1.1 The Scientific Question

The first step in the causal "roadmap" is to specify the scientific objective. As a running example, we will consider the timing of antiretroviral therapy (ART) initiation and its impact on outcomes among HIV+ individuals. Early ART initiation has been been shown to improve patient outcomes as well as reduce transmission between discordant couples [9–11]. Suppose we want to learn the effect of immediate ART initiation (i.e. irrespective of CD4+ T cell count) on mortality. Large consortiums, such as the International epidemiologic Databases to Evaluate AIDS (IeDEA) and Sustainable East Africa Research in Community Health (SEARCH), are providing unprecedented quantities of data to answer this and other questions [12, 13].

To sharply frame our scientific aim, we need to further specify the system, including the target population (e.g. patients and context), the exposure (e.g. criteria and timing), and the outcome. As a second try, consider our goal as learning the impact of initiating ART within

one month of diagnosis on five-year all-cause mortality among adults, recently diagnosed with HIV in Sub-Saharan Africa. This might seem like an insurmountable task, and it may seem safer to frame our question in terms of an association. Indeed, there seems to be a tendency to shy away from causal language when stating the scientific objective. However, we are not fundamentally interested in the correlation between early ART initiation and mortality among HIV+ adults. Instead, we want to isolate the effect of interest from the spurious sources of dependence (e.g. confounding, selection bias, informative censoring) as shown in Figure 1.1. The framework, discussed in this chapter, provides a pathway from our scientific aim to estimation of a statistical parameter that best approximates our causal effect, while keeping any assumptions transparent.

1.2 The Causal Model

The second step of the roadmap is to specify our causal model. Causal inference is distinct from statistics in that it requires something more than a sample from the observed data distribution. In particular, causal inference requires specification of background knowledge, and causal models provide a rigorous language for expressing this knowledge and its limits. In this chapter, we introduce *structural causal models* [14] to formally represent which variables potentially affect one another, the roles of unmeasured factors, and the functional form of those relationships. Structural causal models unify causal graphs [15, 16], structural equations [17, 18] and counterfactuals. We also briefly introduce the Neyman-Rubin potential outcomes framework [19–21] and discuss its relation to the structural causal model.

Consider again our running example. Let W denote the set of baseline covariates, including socio-demographics, clinical measurements and social constructs. The exposure A is an indicator, equalling 1 if the patient initiated ART within one month of diagnosis and equalling 0 otherwise (i.e. initiation took longer than one month). Finally, the outcome Y is an indicator that the patient did not survive five years of follow up. These factors have scientific meaning to the question and comprise the set of *endogenous variables*: $X = \{W, A, Y\}$. They can be measurable (e.g. age and sex) or unmeasurable and are affected by other variables in the model.

Each endogenous variable is associated with a set of background factors $U = (U_W, U_A, U_Y)$ with some joint distribution P_U. These represent all the unmeasured factors, affecting other variables in the model but not included in X. For example, U_A could include unknown clinic-level factors, influencing whether or not a patient initiates early ART. Likewise, U_Y may include a patient's genetic risk profile. Furthermore, there might be shared unmeasured causes between the endogenous variables. For example, socio-economic status may impact both whether a patient initiates early ART as well as his/her five-year mortality.

Each endogenous variable is also associated with a structural equation. These functions help encode our causal knowledge. Suppose, for example, we believe that the set of baseline covariates possibly impact whether a patient initiates early ART, and that both the covariates and exposure may affect subsequent morality. Then we write each endogenous variable as a

deterministic function of its "parents", variables that *may* impact its value:

$$W = f_W(U_W)$$
$$A = f_A(W, U_A)$$
$$Y = f_Y(W, A, U_Y) \tag{1.1}$$

These functions $F = \{f_W, f_A, f_Y\}$ are left unspecified (non-parametric). For example, the third equation f_Y encodes that the covariates W and the exposure A may have influenced the value taken by the outcome Y. We have not, however, restricted their relationships: A and any member of W may interact on an additive (or any other) scale to affect Y and the impacts of A and W on Y may be nonlinear.

The structural causal model, denoted $\mathcal{M}^{\mathcal{F}}$, is defined by all possible distributions of P_U and all possible sets of functions F, which are compatible with our assumptions (if any). For the above example, there is some true joint distribution $P_{U,0}$ of health care access, personal preferences for ART use, socio-economic factors, etc. Randomly sampling a patient from the population corresponds to drawing a particular realization u from $P_{U,0}$. Likewise, there are some true structural equations F_0 that would deterministically generate the endogenous variables $X = x$ if given input $U = u$. For a given distribution P_U and set of functions F, the structural causal model $\mathcal{M}^{\mathcal{F}}$ describes the following data generating process for (U, X):

1. Drawing the background factors U from some joint probability distribution P_U,
2. Generating the baseline covariates W as some deterministic function f_W of U_W,
3. Generating the exposure A as some deterministic function f_A of covariates W and U_A,
4. Generating the outcome Y as some deterministic function f_Y of covariates W, the exposure A and U_Y.

Thus, the model $\mathcal{M}^{\mathcal{F}}$ is the collection of all possible probability distributions $P_{U,X}$ for the exogenous and endogenous variables (U, X). The true joint distribution is an element of the causal model: $P_{U,X,0} \in \mathcal{M}^{\mathcal{F}}$. The structural causal model is also sometimes also called a non-parametric structural equation model (NPSEM) [14, 16].

In other settings, we may have more in-depth knowledge about the data generating process. This knowledge is generally encoded in two ways. First, excluding a variable from the parent set of X_j encodes that this variable does not directly impact the value X_j takes. These assumptions are known as *exclusion restrictions*. Second, restricting the set of allowed distributions for P_U encodes that some variables do not have any unmeasured common causes. These assumptions are known as *independence assumptions*. Suppose, for example, that patients were randomized R to early ART initiation, but adherence A was imperfect. Then the treatment assignment R would only be determined by chance (e.g. a coin flip) and not influenced by baseline covariates W. The unmeasured factors determining treatment assignment would be independent from all other unmeasured factors:

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

This is an independence assumption that restricts the allowed distribution of background factors P_U. Furthermore, suppose that randomization R only affects the mortality Y through its effect on adherence A. The resulting structural equations are then

$$
\begin{aligned}
W &= f_W(U_W) \\
R &= f_R(U_R) \\
A &= f_A(W, R, U_A) \\
Y &= f_Y(W, A, U_Y)
\end{aligned}
\tag{1.2}
$$

We have made two exclusion restrictions: (1) the baseline covariates W do not influence randomization R, and (2) randomization R has no direct effect on the outcome Y. The structural causal model is then defined by all probability distributions for U that are compatible with our independence assumptions and all sets of functions $F = (f_W, f_R, f_A, f_Y)$ that are compatible with our exclusion restrictions.

A causal graph can be drawn from the structural causal model [14]. Each endogenous variable (node) is connected to its parents and background error term with a directed arrow. The potential dependence between the background factors are encoded by the inclusion of a node representing any unmeasured common cause. Exclusion restrictions are encoded by absence of a directed arrow. Likewise, independence assumptions are encoded with the absence of a node representing an unmeasured common cause. The corresponding causal graphs for the two examples are given in Figure 1.2.

1.3 The Target Causal Quantity

The structural causal model $\mathcal{M}^{\mathcal{F}}$ describes not only the system as it currently exists but also as it would exist under changed conditions. The structural equations are autonomous; an intervention on one equation does not affect the remaining ones. Therefore, we can modify a function and see how changes are transmitted through the system. For example, modifying the treatment decision does not change the effect of the treatment on the outcome. Thereby, we can make a targeted modification to represent our intervention of interest. In our running example (Eq. 1.1 and Figure 1.2a), a self-selected group of patients initiated early ART. To answer our scientific question, we need to modify how this exposure variable was generated. Specifically, we can intervene to start all patients on ART within one month of testing HIV+ (i.e. deterministically set $A = 1$), and we can intervene to delay all patients from starting ART until one month after testing HIV+ (i.e. deterministically set $A = 0$):

$$
\begin{aligned}
W &= f_W(U_W) & W &= f_W(U_W) \\
A &= 1 & A &= 0 \\
Y_1 &= f_Y(W, 1, U_Y) & Y_0 &= f_Y(W, 0, U_Y)
\end{aligned}
$$

(a)   (b)
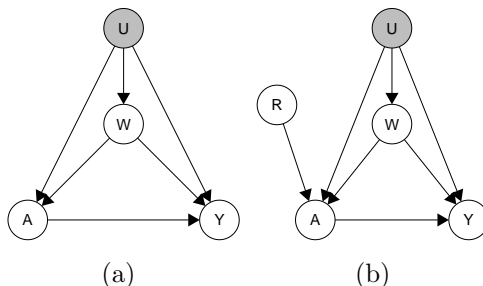
Figure 1.2: Directed acyclic graphs representing the structural causal model for our study (Eq. 1.1) and for the hypothetical randomized trial (Eq. 1.2): (a) This graph only encodes the time-ordering between baseline covariates W, the exposure A and the outcome Y. A single node U represents the unmeasured common causes of the endogenous variables. (b) This graph encodes the randomization R of some treatment with incomplete adherence A. There are two exclusion restrictions: the baseline covariates W do not impact the randomization R, and the randomization R has no direct effect on the outcome Y. There is also an independence assumption: the unmeasured factors contributing to randomization are independent of the unmeasured factors, contributing to the other variables.

Alternative exposure mechanisms include dynamic interventions [22–25], which are responsive to patient characteristics, and stochastic interventions[1] [26], which are non-deterministic.

The *counterfactual outcome* Y_a is then the outcome a patient would have had, if possibly contrary to fact, he or she had received exposure level $A = a$. More formally, $Y_a = Y_a(u)$ is defined as the solution to the equation f_Y under an intervention to set $A = a$ (with input $U = u$). Thereby $Y_a(U)$ is a post-intervention random variable, whose probability distribution is induced by the set of structural equations F and the joint distribution of the background factors P_U. In other words, the structural causal model $\mathcal{M}^{\mathcal{F}}$ is also a model on the distribution of counterfactuals. In the Neyman-Rubin causal framework, these quantities are known as *potential outcomes* and are assumed to exist for all units under the treatment levels of interest [19–21, 27]. For this example, the "full data" would consist of baseline covariates and the outcomes under all possible exposures: $X^{\mathcal{F}} = \big(W, (Y_a : a \in \{0, 1\})\big)$. The structural causal model $\mathcal{M}^{\mathcal{F}}$ also serves as a model for the set of possible full data distributions, each corresponding to a different intervention on the endogenous variables.

The distribution of these counterfactuals (potential outcomes) can then be used to define

[1]For simplicity, we have been considering the time scale to be in months. Depending our scientific question and the data resolution, we might be interested in shorter or longer intervals. If our time interval were days, then an intervention to start by day 30 (i.e. within one month) is actually stochastic intervention. Alternatively, we could consider an intervention to initiate therapy on each day or not. For further discussion of longitudinal treatment regimes, see Appendix A.

the target causal parameter. Consider, for example, the average treatment effect:

$$\Psi^{\mathcal{F}}(P_{U,X}) = \mathbb{E}_{U,X}(Y_1) - \mathbb{E}_{U,X}(Y_0)$$

where the subscript (U, X) denotes the expectation over the distribution $P_{U,X}$ (which implies the distribution of the counterfactual random variables (Y_1, Y_0)). In words, $\Psi^{\mathcal{F}}(P_{U,X})$ is the difference in the expected counterfactual outcome if everyone in the population were exposed and the expected counterfactual outcome if everyone in the population were not exposed. Formally, $\Psi^{\mathcal{F}}$ is a mapping from a distribution $P_{U,X}$ in the causal model $\mathcal{M}^{\mathcal{F}}$ to the real number line. For our example, $\Psi^{\mathcal{F}}(P_{U,X})$ is the difference in the counterfactual risk of mortality if all patients immediately initiated ART and if all patients delayed ART initiation. For a binary outcome, this causal quantity corresponds the causal risk difference. We could also specify this contrast on the relative scale, within a certain strata of the population (e.g. those with baseline CD4 counts above 350 cells/mm^3), for the actual study units (i.e. the sample average treatment effect [19]), or for some other population (i.e. transportability [28–30]).

Marginal structural models provide an alternative way to define our target parameter [31]. They are a summary measure of how the counterfactual outcome changes as a function of the exposure and possibly pre-treatment covariates. Consider, for example, the impact of reducing the time (in months) between HIV diagnosis and treatment initiation. The intervention variable A would then be continuous. (An alternative approach would be to treat the exposure as a time-dependent binary variable as discussed in Appendix A.) To generate the relevant counterfactual outcomes[2], we would repeatedly intervene on the structural causal model to set $A = a$ for all levels of a in the exposure set of interest $\mathcal{A} = \{1, 2, 3, \ldots\}$. If we knew the true shape of the relationship between the expected counterfactual outcome $\mathbb{E}_{U,X}(Y_a)$ and the treatment level a, we could summarize it with a parametric model [31], such as the following

$$logit\big[\mathbb{E}_{U,X}(Y_a)\big] = m(a|\beta)$$
$$\text{with } m(a|\beta) = \beta_0 + \beta_1 a$$

where $logit(x) = log(x/(1 - x))$. This model assumes that the counterfactual mortality risk is a function linear on the logistic scale of time to treatment initiation a. This marginal structural model restricts the set of possible counterfactual distributions and thereby places an assumption on our causal model $\mathcal{M}^{\mathcal{F}}$.

In many cases, we do not have sufficient information to confidently specify a parametric model for this dose-response curve. Instead, we can use a *working marginal structural model* as a summary of the causal relationship of interest [32]. The target causal parameter is then the projection of the true causal curve onto a working model. Consider for example

$$\beta(P_{U,X}|m) = argmin_{\beta} \; \mathbb{E}_{U,X}\left[\sum_{a \in \mathcal{A}} -log\big[m(a|\beta)^{Y_a}\big(1 - m(a|\beta)\big)^{(1-Y_a)}\big]\right]$$

---

[2] Under the Neyman-Rubin framework, we would assume the existence of the potential outcomes Y_a for all exposures $a \in \mathcal{A}$.

where our projection is the negative log-likelihood loss function. Intuitively, we can think of this projection as summarizing the full data (i.e. all counterfactuals) with a parametric regression curve. As usual, the quality of the summary depends on the underlying causal curve and the question of interest.

1.4 The Observed Data and their Link to the Causal Model

Thus far, we have not specified the data that will be or have been collected in our study. Instead, we have discussed endogenous variables X (observable and possibly unobservable), background factors U (unobservable), and set of counterfactuals $(Y_a : a \in \mathcal{A})$. In this step, we specify the observed data, their link to the causal model and the resulting statistical model.

Suppose we have a simple random sample of n patients from our target population. On each patient, we measure some baseline covariates W, including sex, age and CD4 count, the exposure A (whether or not the patient initiates ART within one month of diagnosis), and the outcome Y as the patient's five-year mortality. Then the observed data for a given patient are $O = (W, A, Y)$, which have some true, but unknown distribution P_0. We assume that the observed data are generated by sampling n times from a distribution compatible with (contained in) the structural causal model. Recall the structural causal model provides a description of the data generating system under existing conditions as well as under specific interventions. The distribution of the background factors P_U and the structural equations F identify the distribution of the endogenous variables X as well as the distribution of the observed data O. The observed data O are a subset of (U, X). Suppose, for example, we observe all the endogenous nodes (i.e. $O = X$). Then we have

$$P(O = o) = \sum_u P_{U,X}(X = x | U = u) P_U(U = u) = \sum_u \mathbb{I}(X(u) = x) P_U(U = u)$$

where the summation generalizes to an integral for continuous valued variables. This framework naturally accommodates more complicated links, such as case-control sampling and matched sampling [33, 34].

Thereby, the structural causal model $\mathcal{M}^{\mathcal{F}}$, which is the set of possible distributions for (U, X), implies our statistical model $\mathcal{M}$, which is the set of possible distributions for the observed data O. The true distribution of the observed data P_0 is implied by the true distribution $P_{U,X,0}$ of (U, X) and is an element of the statistical model: $P_0 \in \mathcal{M}$. The causal model may, but often does not, place any restrictions on the statistical model. For example, the causal model, describing the data generating process for our observational study (Figure 1.2a), implies a *non-parametric* statistical model. There are no restrictions on the possible observed data distributions. In contrast, the causal model, corresponding to the randomized trial (Figure 1.2b), will only generate distributions where the randomization R is

independent of the baseline covariates W. This is a testable assumption and implies a *semi-parametric* statistical model. We refer the reader to Pearl [14, 15] for further discussion of a graphical criteria to evaluate independence between two variables as implied by a structural causal model or its corresponding directed acyclic graph.

Suppose that instead of specifying a structural causal model, we chose to follow the Neyman-Rubin framework. Specifically, we assumed the existence of the potential outcomes $Y_a : a \in \mathcal{A}$ in Step 3. To relate these potential outcomes to the observed data, we need the stable unit treatment value assumption (SUTVA) [35]. First, the potential outcomes for one unit must not be impacted by the treatment assignment of another unit (i.e. no interference)[3]. Secondly, there must not be multiple versions of the treatment $A = a$. With this assumption, we can map the potential outcomes to the observed outcomes:

$$Y_i = A_i Y_{i,1} + (1 - A_i) Y_{i,0}$$

For unit i, we only get to see the outcome Y_i, corresponding to his or her observed exposure A_i. As a result, causal inference can be treated as a missing data problem.

1.5 Assessment of Identifiability

In Step 3, we specified our scientific question as a causal parameter $\Psi^{\mathcal{F}}(P_{U,X})$, a function of the distribution of counterfactuals (potential outcomes). In Step 4, we specified the observed data O and the statistical model $\mathcal{M}$. In this step, we establish whether our causal parameter can be written as some function of the observed data distribution. More formally, for each $P_{U,X}$ compatible with the structural causal model $\mathcal{M}^{\mathcal{F}}$, we want to establish the equivalence between the causal parameter $\Psi^{\mathcal{F}}(P_{U,X})$ and the statistical parameter $\Psi(P)$. If so, we state that the causal parameter is *identified*. If not, we explicitly state the additional assumptions needed to make inferences about the causal parameter using the observed data distribution. We keep these convenience-based assumptions separate from our knowledge-based assumptions, reflected in the structural causal model $\mathcal{M}^{\mathcal{F}}$.

Consider a simplified example, where we want to learn the five-year mortality risk if, possibly contrary to fact, all HIV+ adults initiated ART within one month of diagnosis: $P_{U,X}(Y_1 = 1)$. Suppose we have not collected any baseline covariates; thereby, the observed data are simply $O = (A, Y)$. Then the causal parameter will only equal the observed mortality risk among exposed if the *only* source of association is due to the effect of interest:

$$P(Y = y | A = 1) = P_{U,X}(Y_1 = y | A = 1)$$
$$\overset{?}{=} P_{U,X}(Y_1 = y)$$

---

[3]The structural causal model, given in Eq. 1.1, implicitly assumes independence between study units. Recent work relaxing this assumption and considering a network of interacting units is given in van der Laan [36].

The first equality is by the definition of counterfactuals and then second holds if the counterfactual outcome Y_a is independent of the exposure A. In the absence of measured baseline covariates, the outcome is only a function of the exposure and its background factors: $Y = f_Y(A, U_Y)$. Once we intervene to set $A = a$, then the counterfactual outcome is only a function of its error: $Y_a(U) = f_Y(a, U_Y)$. If the unmeasured factors contributing to the outcome U_Y are independent of those contributing the exposure U_A, then the randomization assumption holds $Y_a \perp\!\!\!\perp A$, and the counterfactual risk $P_{U,X}(Y_1 = 1)$ is identified as the observed risk among those exposed $P(Y = 1 | A = 1)$. The randomization assumption is equivalent to stating that there are no unmeasured confounders of the exposure-outcome relation. Intuitively, this assumption holds by design a randomized trial.

In most observational settings, the assumption of no common (measured or unmeasured) causes of the exposure and the outcome will not hold. We can weaken the randomization assumption by conditioning on a set of measured baseline covariates: $Y_a \perp\!\!\!\perp A | W$. The adjustment set W needs to block all spurious sources of association without creating any new sources of dependence or blocking any of the effect of A on Y. As illustrated in Figure 1.3, the *back-door criterion* can aid the evaluation of the randomization assumption [14]. A set of variables W satisfies the back-door criterion for the relationship of (A, Y) if (1) no node in W is a descendant of A and (2) W blocks all back-door paths from A to Y, where "back-door" refers to a path with an arrow into A. The rationale for condition 1 is to avoid blocking the path of interest or introducing spurious associations (i.e. conditioning on a collider). The rationale for condition 2 is to block any remaining spurious sources of association. For the basic structure (Figure 1.3), the randomization assumption will hold if the following independence assumptions are true

$$U_A \perp\!\!\!\perp U_Y \text{ and } U_A \perp\!\!\!\perp U_W \text{ or } U_Y \perp\!\!\!\perp U_W$$

There must not be any unmeasured common causes of the exposure and the outcome, and of the exposure and covariates or of the outcome and covariates. As illustrated in Figure 1.4, this graphical criteria can aid in the selection of an appropriate adjustment set.

When the randomization assumption holds, we can identify the distribution of counterfactuals within strata of covariates. Specifically, we have that for each $P_{U,X} \in \mathcal{M}^{\mathcal{F}}$

$$P_{U,X}(Y_a = y | W = w) = P_{U,X}(Y_a = y | A = a, W = w)$$
$$= P(Y = y | A = a, W = w)$$

where the distribution P of the observed data is implied by $P_{U,X}$. This gives us the G-computation identifiability result [27] for the true distributions $P_{U,X,0}$ and P_0:

$$E_{U,X,0}(Y_a) = \sum_w E_0(Y | A = a, W = w) P_0(W = w)$$

where the summation generalizes to an integral for continuous covariates. Likewise, we can identify the difference in the expected counterfactual outcomes (i.e. the average treatment
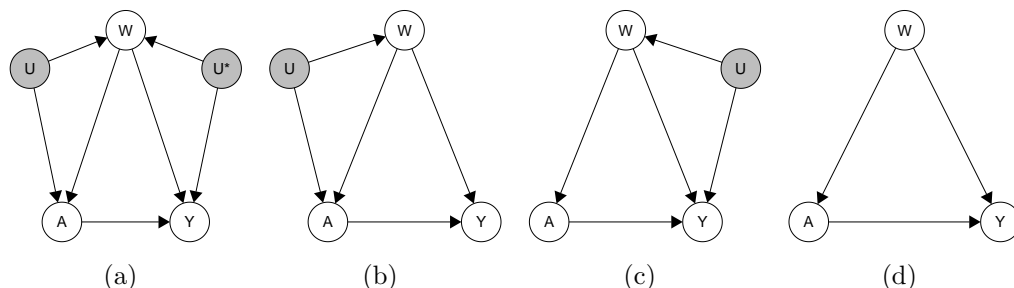
Figure 1.3: *Considering the back-door criterion for the basic structure: For all the graphs, the exposure A and the outcome Y do not share an unmeasured common cause. (a) The covariates W are* not *sufficient to block all back-door paths. Conditioning on the covariates W blocks the path $Y \to W \to A$. However, conditioning on W (a collider of U and U^*) opens a new path: $Y \to U^* \to U \to A$. (b) The covariates W and the outcome Y also do not share an unmeasured common cause. The covariates W are sufficient to block all back-door paths. (c) The exposure A and the covariates W also do not share an unmeasured common cause. The covariates W are sufficient to block all back-door paths. (d) All the unmeasured background factors are independent. The covariates W are sufficient to block all back-door paths.*

effect) in terms of the difference in the conditional mean outcomes, averaged (standardized) with respect to the covariate distribution:

$$\underbrace{E_{U,X,0}(Y_1 - Y_0)}_{\Psi^{\mathcal{F}}(P_{U,X,0})} = \underbrace{\sum_w \big[E_0(Y|A=1,W=w) - E_0(Y|A=0,W=w)\big]P_0(W=w)}_{\Psi(P_0)}$$

Identifiability also relies on having sufficient support in the data. The G-computation formula requires that the conditional mean $E_0(Y|A=a,W=w)$ is well-defined for all possible values of w and levels of a of interest. In a non-parametric statistical model, each exposure of interest must occur with some positive probability for each possible covariate strata:

$$\min_{a \in \mathcal{A}} P_0(A=a|W=w) > 0, \text{ for all } w \text{ for which } P_0(W=w) > 0$$

This condition is known as the *positivity assumption* and as the experimental treatment assignment assumption.

Suppose, for example, that the randomization assumption held conditionally on a single binary baseline covariate. Then our statistical estimand could be rewritten as

$$\Psi(P_0) = \big[E_0(Y|A=1,W=1) - E_0(Y|A=0,W=1)\big]P_0(W=1)$$
$$+ \big[E_0(Y|A=1,W=0) - E_0(Y|A=0,W=0)\big]P_0(W=0)$$
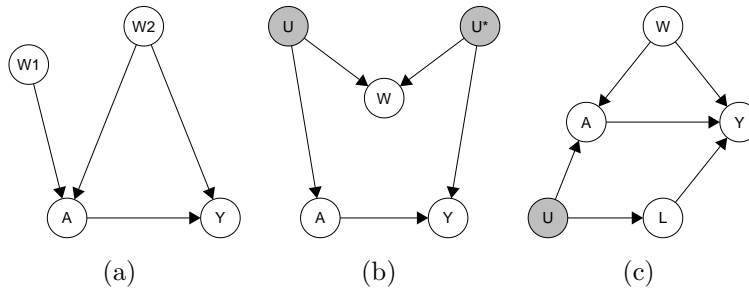
Figure 1.4: Considering the back-door criterion: (a) The set of covariates $W2$ are sufficient to block the back-door path from $Y \to W2 \to A$. Further adjustment for $W1$ is unnecessary and potentially harmful. (b) The randomization assumption holds conditionally on $\emptyset$. Adjusting for W (i.e. conditioning on a collider of U and U^*) open a back-door path and induces a spurious association between A and Y. (c) The randomization assumption holds conditionally on (W, L). The covariates L are needed to block the back-door path from $Y \to L \to U \to A$, even though L occurs temporally after the exposure A.

As an extreme, suppose that in the population, there are zero exposed patients with this covariate: $P_0(A = 1|W = 1) = 0$. Then there would be no information about outcomes under the exposure for this subpopulation. To identify the treatment effect, we could consider a different target parameter (e.g. the effect among those with $W = 0$) or consider additional modeling assumptions (e.g. the effect is the same among those with $W = 1$ and $W = 0$). Both options are a bit dissatisfying and other approaches may be taken [37]. The risk of violating the positivity assumption is exacerbated with higher dimensional data (i.e. as the number of covariates or their levels grow).

In many cases, our initial assumptions, encoded in the structural causal model $\mathcal{M}^{\mathcal{F}}$, are not sufficient to identify the causal effect $\Psi^{\mathcal{F}}(P_{U,X})$. Indeed, for our running example (Figure 1.2a), the set of baseline covariates are not sufficient to block the back-door paths from the outcome to the exposure. The question then becomes how to proceed? Possible options include giving up, gathering more data, or continuing to estimation while clearly acknowledging the lack of identifiability during the interpretation step. To facilitate the third option, we can use $\mathcal{M}^{\mathcal{F}*}$ to denote the structural causal model augmented with additional convenience-based assumptions needed for identifiability. This gives us a way to proceed, while separating our real knowledge $\mathcal{M}^{\mathcal{F}}$ from our wished identifiability assumptions $\mathcal{M}^{\mathcal{F}*}$.

Overall, identifiability assumptions and the resulting estimands are specific to the causal parameter $\Psi^{\mathcal{F}}(P_{U,X})$. We are focusing on a point treatment effect (i.e. distribution of counterfactuals under interventions on a single node or variable). Different identifiability results are needed for interventions on more than one node (e.g. longitudinal treatment effects and direct effects) and interventions responding to patient characteristics (e.g. dynamic regimes). Furthermore, a given causal parameter may have more than one identifiability result (e.g.

instrumental variables and the front door criterion). See, for example, Pearl [14].

1.6   Estimation and Inference

In the previous step, we defined the parameter of interest as a mapping from the statistical model to the parameter space: $\Psi : \mathcal{M} \rightarrow \mathbb{R}$. In other words, the statistical parameter is a function, whose input is any distribution P compatible with the statistical model and whose output is a real number. The parameter mapping applied to the true observed data distribution P_0 is called the *estimand* and denoted $\Psi(P_0)$. Recall we have n independent, identically distributed (i.i.d.) copies of the random variable $O = (W, A, Y)$. The empirical distribution P_n corresponds to putting a weight $1/n$ on each copy of O_i. An *estimator* is a function, whose input is the observed data (a realization of P_n) and output a value in the parameter space.

In this chapter, we consider *substitution estimators* based on the G-Computation identifiability result [27]:

$$\Psi(P_0) = E_0\big[E_0(Y|A = 1, W) - E_0(Y|A = 0, W)\big] \tag{1.3}$$

A simple substitution estimator for $\Psi(P_0)$ can be implemented as follows.

1. Estimate of the conditional expectation of the outcome, given the exposure and co-variates, denoted $\hat{E}(Y|A, W)$.
2. Use this estimate to generate the predicted outcomes for each unit, setting $A = 1$ and $A = 0$.
3. Take the sample average of the difference in these predicted outcomes:

$$\hat{\Psi}(P_n) = \frac{1}{n} \sum_{i=1}^{n} \hat{E}(Y_i|A_i = 1, W_i) - \hat{E}(Y_i|A_i = 0, W_i)$$

The last step corresponds with estimating the marginal covariate distribution $P_0(W)$ with the sample proportion: $\frac{1}{n} \sum_i \mathbb{I}(W_i = w)$.

There are many options available for estimating the conditional expectation $E_0(Y|A, W)$. Often, parametric models are used to relate the conditional mean outcome to the possible predictor variables and the exposure. Suppose, for example, we knew that the conditional expectation of a continuous outcome could be described by the following parametric regression model:

$$E_0(Y|A, W) = \beta_0 + \beta_1 A + \beta_2 W1 + \beta_3 W2 + \beta_4 A^* W1 + \beta_5 A^* W2$$

where $W = \{W1, W2\}$ denotes the set of covariates, needed for identifiability. Then this knowledge should have been encoded in our structural causal model $\mathcal{M}^{\mathcal{F}}$ with implied restrictions on our statistical model $\mathcal{M}$. (In other words, we avoid introducing new assumptions during the analysis.) The coefficients in this regression model could be estimated with maximum likelihood or with ordinary least squares regression. The estimate $\hat{\beta}_1$ does not, however,

provide an estimate of the G-computation identifiability result. The exact interpretation of $\hat{\beta}_1$ depends on which variables and which interactions are included in the parametric model. To obtain an estimate of $\Psi(P_0)$, we need to average the predicted outcomes with respect to the distribution of covariates:

$$\hat{\Psi}(P_n) = \frac{1}{n} \sum_{i=1}^{n} \hat{E}(Y|A=1, W_i) - \hat{E}(Y|A=0, W_i)$$

$$= \frac{1}{n} \sum_{i=1}^{n} \left(\hat{\beta}_1 + \hat{\beta}_4 W1_i + \hat{\beta}_5 W2_i \right)$$

As a second example, suppose we knew that the conditional risk of a binary outcome could be described by the following parametric model:

$$logit\left[E_0(Y|A, W)\right] = \beta_0 + \beta_1 A + \beta_2 W1 + \ldots + \beta_{10} W9$$

where $W = \{W1, \ldots, W9\}$ denotes the set of covariates, needed for identifiability. Then the estimate $\hat{\beta}_1$ would provide an estimate of the logarithm of the conditional odds ratio. An estimate of the G-computation identifiability result is given by averaging the expected outcomes under the exposure $A = 1$ and the control $A = 0$:

$$\hat{\Psi}(P_n) = \frac{1}{n} \sum_{i=1}^{n} \left(\frac{1}{1 + exp^{-(\hat{\beta}_0 + \hat{\beta}_1 + \hat{\beta}_2 W1_i + \ldots + \hat{\beta}_{10} W9_i)}} - \frac{1}{1 + exp^{-(\hat{\beta}_0 + \hat{\beta}_2 W1_i + \ldots + \hat{\beta}_{10} W9_i)}} \right)$$

In most cases, our background knowledge is inadequate to describe the conditional expectation $E_0(Y|A, W)$ with such parametric models. Indeed, with high dimensional data, the sheer number of potential covariates will likely make it impossible to correctly specify the functional form. If the assumed parametric model is incorrect, the point estimates will often be biased and inference misleading. In other words, the structural causal model $\mathcal{M}^{\mathcal{F}}$, representing our knowledge of the underlying data generating process, often implies a non-parametric statistical model $\mathcal{M}$. Our estimation approach should respect the statistical model.

To avoid unsubstantiated assumptions about functional form, it is sometimes possible to estimate $E_0(Y|A, W)$ with the empirical mean in each exposure-covariate strata. Unfortunately, even when all covariates are discrete valued, non-parametric maximum likelihood estimators quickly become ill-defined due to the curse of dimensionality; the number of possible exposure-covariate combinations far exceed the number of observations.

Various regression model selection routines can help alleviate these problems. For example, stepwise regression will add and subtract variables in hopes of minimizing the Akaike information criterion (AIC) or the Bayesian information criterion (BIC). Other data-adaptive methods, based on cross-validation, involve splitting the data into training and validation sets. Each possible algorithm (e.g. various parametric models or semiparametric methods) is then fit on the training set and its performance assessed on the validation set. The measure of performance can be defined by a loss function, such as the L2-squared error or the

negative log-likelihood. Super learner, for example, uses cross-validation to select the candidate algorithm with the best performance or to build the optimal (convex) combination of estimates from candidate algorithms [38, 39]. A point estimate could then be obtained by averaging the difference in predicted outcomes for each unit under the exposure and under the control.

While these data-adaptive methods avoid betting on one *a priori* specified parametric regression model and are amenable to semiparametric algorithms, there is no reliable way to obtain statistical inference for parameters, such as the G-Computation estimand $\Psi(P_0)$. Treating the final algorithm as if it were pre-specified ignores the selection process. Furthermore, the selected algorithm was tailored to maximize/minimize some criterion with regards to the conditional expectation $E_0(Y|A, W)$ and will, in general, not provide the best bias-variance tradeoff for estimating the statistical parameter $\Psi(P_0)$. Indeed, estimating the conditional mean outcome Y in every strata of (A, W) is a much more ambitious task than estimating one number (the difference in conditional means, averaged with respect to the covariate distribution). Thus without an additional step, the resulting estimator will be overly biased relative to its standard error, preventing accurate inference.

Targeted maximum likelihood estimation (TMLE) provides a way forward [3, 40]. TMLE is a general algorithm for the construction of double robust, semiparametric, efficient, substitution estimators. TMLE allows for data-adaptive estimation while obtaining valid statistical inference. For the G-computation estimand, the TMLE algorithm uses information in the estimated exposure mechanism $\hat{P}(A|W)$ to update the initial estimator of the conditional mean $E_0(Y|A, W)$. The targeted estimates are then substituted into the parameter mapping. The updating step achieves a targeted bias reduction for the parameter of interest $\Psi(P_0)$ and serves to solve the efficient score equation. As a result, TMLE is a double robust estimator; it will be consistent for $\Psi(P_0)$ if either the conditional expectation $E_0(Y|A, W)$ or the exposure mechanism $P_0(A|W)$ is estimated consistently. When both functions are consistently estimated at a fast enough rate, the TMLE will be efficient in that it achieves the lowest asymptotic variance among a large class of estimators. These asymptotic properties typically translate into lower bias and variance in finite samples. The advantages of TMLE have been repeatedly demonstrated in both simulation studies and applied analyses (e.g. [37, 41–43]). The procedure is available with standard software such as the `tmle` and `ltmle` packages in `R` [44–46].

Thus far, we have discussed obtaining a point estimate from a simple or targeted substitution estimator. To create confidence intervals and test hypotheses, we also need to quantify uncertainty. A simple substitution estimator based a correctly specified parametric model is asymptotically linear, and its variance can be approximated by the variance of its influence curve, divided by sample size n. It is worth emphasizing that our estimand $\Psi(P_0)$ often does not correspond to a single coefficient, and therefore we usually cannot read off the reported standard error from common software. Under reasonable conditions, the TMLE is also asymptotically linear, and inference can be based on an estimate of its influence curve. Further discussion of influence curve-based inference is given in subsequent chapters.

Overall, we focused on substitution estimators (simple and targeted) of the G-computation

identifiability result [27]. The simple substitution estimator only requires an estimate of the marginal distribution of baseline covariates $P_0(W)$ and the conditional expectation of the outcome, given the exposure and covariates $E_0(Y|A, W)$. TMLE also requires an estimate of the exposure mechanism $P_0(A|W)$. There are many other algorithms available for estimation of $\Psi(P_0)$. A popular class of estimators rely only on estimation of the exposure mechanism [22, 47–51]. Inverse probability of treatment weighting (IPTW) estimators, for example, control for measured confounders by up-weighting exposure-covariate groups that are under-represented and down-weighting exposure-covariate groups that are over-represented (relative to what would be seen were the exposure randomized). Its double robust counterpart, augmented-IPTW, shares many of the same properties as TMLE [52–55]. A key distinction is that IPTW and augmented-IPTW are solutions to estimating equations and thereby respond differently in the face of challenges due to strong confounding and rare outcomes [37, 56]. Throughout, we maintain that estimators should respect the knowledge encoded in the statistical model and not introduce new assumptions. An estimator should be selected for analysis based on its performance (e.g. bias, variance, robustness) as opposed to convenience or habit.

1.7 Interpretation of the Results

The last step of the roadmap is interpreting the results. In our running example, the identifiability assumptions did not hold. Nonetheless, the statistical estimand (Eq. 1.3) always has a statistical interpretation as the difference in the expected outcome, given the exposure and covariates in the adjustment set, and the expected outcome, given the control and covariates in the adjustment set, standardized with respect to the covariate distribution in the population. For our example, $\Psi(P_0)$ can be interpreted as the marginal risk difference: the difference in the mortality risk among patients with early vs. delayed ART initiation but the same values of the measured covariates (e.g. baseline CD4 count, age and sex), averaged with respect to the distribution of these covariates. This estimand can be considered as the best approximation to the causal quantity of interest, given the limitations in the observed data. If the identifiability assumptions hold, our estimate would be endowed with a causal interpretation: a summary of how the distribution of the data would change under a specific intervention. For our example, the causal interpretation would be the difference in the five-year counterfactual mortality risk if all patients initiated early ART vs. if all patients delayed ART initiation. Further interpretation in terms of the impact of a "real world" intervention or in terms of a randomized trial requires additional assumptions.

1.8 Discussion

In this chapter, we introduced a formal framework for causal inference [3, 4]. Our running example was to estimate the effect of early ART initiation (within one month of diagnosis) on

the five-year mortality risk among HIV+ adults in Sub-Saharan Africa. Our structural causal model $\mathcal{M}^{\mathcal{F}}$ only reflected the causal-ordering of our variables; we did not make any exclusion restrictions, independence assumptions or functional form assumptions. Counterfactuals were generated by deterministically intervening on the data generating system, described by the structural causal model, to set $A = 1$ (i.e. early initiation) and also to set $A = 0$ (i.e. delayed initiation). We focused on the average treatment effect for this static exposure. The observed data $O = (W, A, Y)$ were assumed to be generated by sampling n independent times from a probability distribution compatible with the structural causal model $\mathcal{M}^{\mathcal{F}}$, which implied a non-parametric statistical model $\mathcal{M}$. Although our identifiability assumptions did not hold, we still defined a statistical estimand $\Psi(P_0)$ as a best approximation of our wished for causal quantity. We briefly discussed a simple (parametric) substitution estimator and a targeted substitution estimator (TMLE), which allows for data-adaptive estimation while obtaining valid inference. Since our needed identifiability assumptions were not met, we interpreted our estimate as the marginal difference in the mortality risk, given early ART initiation and the measured covariates, and the mortality risk, given delayed ART initiation and the measured covariates, standardized with respect to the covariate distribution.

This framework is easily extended to more complicated data structures. Consider, for example, the following scientific questions, corresponding to interventions on multiple exposure nodes and to alternate counterfactual treatment assignment mechanisms:

- *Longitudinal treatment effects* [31, 55, 57–69]: How does cumulative time until ART initiation affect mortality among recently diagnosed HIV+ adults? What is the effect of routine HIV RNA viral load monitoring, as compared to routine CD4+ T cell count monitoring, on mortality among patients initiating early ART? What would be impact of early ART initiation on the five-year mortality if there were no losses to follow up?

- *Dynamic regimes* (individualized treatment rules) [22–25, 60, 70–72]: How would mortality of have differed if HIV+ adults initiated ART based on HIV RNA viral loads as opposed to CD4+ T cell counts?

- *Direct and indirect effects* [73–76]: What is the direct effect of early ART initiation on five-year mortality that is not mediated through changes in HIV RNA viral load?

- *Stochastic interventions* (non-deterministic interventions) [26]: What would be the five-year mortality if the distribution of time until ART initiation shifted towards shorter wait times? What is the impact of early ART initiation on five-year mortality if HIV RNA viral load, the intermediate, remained at the value it would have been in the absence of the exposure (i.e. the natural direct effect [77–79])?

Overall, access to unprecedented amounts of data does not undo the age-old adage: "correlation is not causation". Indeed, there are numerous sources of association (dependence) between two variables: direct effects, indirect effects, measured confounding, unmeasured confounding and selection bias. The methods, introduced here, allow researchers to move

from saying drug X is associated with an adverse side effect to saying (under the necessary and transparently stated assumptions) an adverse side effect is caused by drug X. Even if the needed identifiability assumptions are not expected to hold, this framework helps us to estimate a statistical parameter, coming as close to the wished causal parameter. In other words, this framework ensures that the scientific question is driving the analysis and not the other way around.

Chapter 2

Adaptive Pair-matching in Randomized Trials with Unbiased and Efficient Effect Estimation

Pair-matching helps balance treatment groups with respect to important determinants of the outcome at baseline [80, 81]. In observational studies, matching can help control for confounding. In randomized trials, there is no confounding; the probability of receiving the intervention or the control is a known constant. Nonetheless, covariate imbalance is common in small trials, and data sparsity may limit our ability to adjust for these characteristics during the analysis. Thereby, matching is sometimes implemented in randomized trials to protect study credibility. For example, the "face validity" [82] of a randomized trial for violence prevention could be compromised if neighborhoods with highest baseline violence were all randomized, by chance, to the control level of the intervention. Matching is also implemented to improve study power. By decreasing variation in the outcome within pairs, matching may, but is not guaranteed to, increase study efficiency. The conflicting recommendations on pair-matching have inspired a heated debate in the literature for over sixty years [82–96].

Much of the work in the design and the analysis of pair-matched trials has assumed that the observed data consist of $n/2$ independent and identically distributed (i.i.d.) units (e.g. [82, 96–100]). Such a data structure could arise by randomly sampling $n/2$ matched pairs from some target population of pre-existing matched units. Often, however, there may be substantial logistical or financial barriers to practical implementation of this design. Alternatively, this data structure could arise by (i) sampling a unit from an infinite target population, (ii) measuring its baseline covariates, (iii) repeatedly sampling units until the baseline covariates of the second were sufficiently close to the first, (iv) randomizing the intervention within the matched pair, (v) measuring the outcomes, and (vi) repeating this process $n/2$ times. This pair-matching scheme may also be impractical and is likely to be resource intensive. Theoretically, this design also yields less information for estimating the (population) average treatment effect than a design randomly pairing two sampled units

[101].

A different pair-matching scheme was implemented in the Sustainable East Africa Research in Community Health (SEARCH) trial (NCT01864603) [13, 102]. SEARCH is a multinational, multidisciplinary consortium to evaluate the health, economic and educational impacts of a community-based strategy for immediate and streamlined antiretroviral therapy (ART) for all HIV-positive persons. In the trial, 54 candidate communities were identified from rural Uganda and Kenya. These clusters satisfied the study's inclusion criteria, which included community size, health care infrastructure and sufficient distance from other potential study units. Thirty-two communities were then pair-matched within region and on baseline predictors of HIV transmission and health care delivery. The intervention has been randomized within the resulting 16 matched pairs and the 5-year cumulative incidence of HIV will be measured at the conclusion of the trial. We consider this design to be *adaptive*, because partitioning of the study communities into matched pairs was a function of the baseline covariates of all candidates. Thereby, the observed data do not consist of $n = 32$ i.i.d. random variables or of $n/2 = 16$ i.i.d. paired random variables. Instead, the observed data consist of n dependent units. For examples of other types of adaptive designs, see [103–106].

To the best of our understanding, adaptive pair-matching has been implemented in several other cluster randomized trials. Examples include the Mwanza trial to prevent HIV [107], the PRISM trial to prevent postpartum depression [108], and the SPACE study to promote physical activity [109]. The process of selecting $n/2$ pairs based on the covariates of n candidates is also known in other literature as "nonbipartite matching" [93, 110] and has motivated the development of "optimal multivariate matching" algorithms to pair units based on several covariates simultaneously [111–114]. Previously, van der Laan *et al.* [115] explored the consequences of adaptive pair-matching for estimation of the population average treatment effect. This chapter explores the consequences of adaptive pair-matching for estimation of the average treatment effect, conditional on the measured baseline covariates of the n study units. For brevity, we will refer to this causal parameter as the conditional average treatment effect (CATE). This parameter was initially proposed in Abadie and Imbens [116], can be interpreted as the intervention effect, given the measured covariates of the sample at hand, and often leads to more precise estimators [95, 117, 118].

Adjustment for baseline covariates during the analysis can help control for chance imbalances in important determinants of the outcome and can also increase study efficiency (e.g. [119–125]). Nonetheless, the recommendations on whether and how to adjust in pair-matched trials have been conflicting (e.g. [82, 94–96, 100, 126, 127]). The intervention effect can be estimated with the average of the differences in the outcomes within matched pairs. Alternatively, one could take a multi-step approach of first fitting a regression model with terms for the pairs and covariates (but not the intervention) and then contrasting the observed versus predicted outcomes within matched pairs [96, 107, 128]. In all cases, the estimation approach should be tailored to the parameter of interest (i.e population vs. conditional average treatment effect). To the best of our knowledge, this is the first work to propose targeted minimum loss-based estimation (TMLE) for the CATE in a randomized

trial. Without risking bias due to regression model misspecification [122, 124, 125], TMLE allows for further adjustment for baseline characteristics (beyond that attained by matching alone) and thereby can provide an efficient estimate of the intervention effect.

The remainder of the chapter is outlined as follows. We first describe the adaptive design and the resulting data structure. Second, we motivate the use of the CATE as the causal parameter of interest. Third, we discuss two estimators of the corresponding statistical parameter: the unadjusted difference in outcomes within matched pairs and targeted minimum loss-based estimation (TMLE). The latter estimator allows for further adjustment of important baseline covariates, beyond that attained with matching, and is thereby more powerful under reasonable scenarios. We also provide asymptotically conservative variance estimators and finite sample simulations. We conclude with some practical recommendations. While the SEARCH trial serves as the motivating example, our conclusions are applicable to other randomized trials and also general to other study outcomes beyond incidence. Moreover, we focus on data at the level of the experimental unit (i.e. the unit of randomization). Thereby, our results are applicable to both individually randomized trials as well as cluster randomized trials. Detailed proofs are given in the Appendix B. This chapter was reproduced with permission from Balzer *et al.* [129].

2.1 The Estimation Problem

The SEARCH consortium will estimate the impact of immediate antiretroviral therapy (ART), initiated at all CD4+ T cell counts and delivered by a streamlined care system, on the 5-year cumulative incidence of HIV [13]. The trial began enrolling communities in 2013, and data collection is ongoing. In communities randomized to the intervention, all individuals testing positive for HIV will be immediately eligible for ART with streamlined delivery, which includes enhanced services for initiation, linkage and retention in care. In communities randomized to the control, all individuals testing positive for HIV will be offered ART according to in-country guidelines, which are primarily based on CD4+ T cell counts. HIV incidence, as well as other health, economic and educational outcomes, will be measured among approximately 320,000 individuals, followed longitudinally for the 5 years of the trial. The SEARCH study aims to understand the impact this community-based "test-and-treat" program on both HIV-positive individuals and their greater communities [9, 10, 130–135].

For the purposes of understanding the adaptive design, we focus on the cluster-level data. Let N denote the number of candidate communities considered for inclusion in the study, n denote the number of communities selected for the SEARCH trial, and $n/2$ denote the number of matched pairs. Let W represent the pre-intervention community-level covariates, which include region, proximity to trucking routes, occupational mix and baseline population HIV RNA levels [136]. A subset of these baseline covariates were used to select the $n/2$ best matched pairs of communities from the N possible candidates. Within the resulting pairs, the intervention was randomized. The treatment variable A is a binary indicator, equalling one if the community was assigned to the intervention (all individuals testing positive for

HIV are immediately offered ART with streamlined care delivery) and equalling zero if the community was assigned to the control (all individuals testing positive for HIV are offered ART according to in-country guidelines). Finally, the outcome Y is the 5-year cumulative incidence of HIV, which will be measured through longitudinal follow-up. Thereby, the data structure for a SEARCH community is $O = (W, A, Y)$.

The adaptive design has important implications for estimation and inference [115]. Mainly, the partitioning of the sample into $n/2$ pairs is a function of the baseline covariates of all N candidates. Adaptive pair-matching results in n dependent copies of O. Nonetheless, given the covariates of all candidate communities $W^N = (W_1, \ldots, W_N)$, the observed data can be represented as $n/2$ conditionally independent random variables:

$$\bar{O}_j = \big(O_{j1}, O_{j2}\big) = \big((W_{j1}, A_{j1}, Y_{j1}), (W_{j2}, A_{j2}, Y_{j2})\big)$$

where the index $j = 1, \ldots, n/2$ denotes the partitioning of the candidates $\{1, \ldots, N\}$ into matched pairs according to similarity on their baseline covariates W^N. Throughout the subscripts $j1$ and $j2$ denote the first and second communities within matched pair j. We place no assumptions on the joint distribution of covariates $P_0(W^N)$, where subscript 0 denotes the true but unknown distribution. The treatment assignment mechanism is known; with probability 0.5, the first unit is randomized to the intervention and the second to the control:

$$P_0(A_{j1} = 1, A_{j2} = 0 \mid W^N) = P_0(A_{j1} = 0, A_{j2} = 1 \mid W^N) = 0.5$$

Study communities are assumed to be causally independent (i.e. no contamination or spillover effects). In other words, we assume that the baseline covariates and intervention assignment of one community do not affect the outcome of another study community. Recent work, relaxing these assumptions and considering a network of interacting units, is elaborated in van der Laan [36]. Under these assumptions, the conditional distribution of the observed data, given the baseline covariates of the candidate units, factorizes as

$$P_0(O_1, \ldots, O_n \mid W_1, \ldots, W_N) = \prod_{j=1}^{n/2} \left\{ P_0(A_{j1}, A_{j2} \mid W^N) P_0(Y_{j1} \mid A_{j1}, W_{j1}) P_0(Y_{j2} \mid A_{j2}, W_{j2}) \right\}$$

$$= 0.5 \prod_{j=1}^{n/2} \left\{ P_0(Y_{j1} \mid A_{j1}, W_{j1}) P_0(Y_{j2} \mid A_{j2}, W_{j2}) \right\}$$

$$= P_0(O_1, \ldots, O_n \mid W_1, \ldots, W_n) = P_0^n(O^n \mid W^n)$$

Throughout, P_0^n denotes the true conditional distribution of the observed data, given the baseline covariates of the n study units $W^n = (W_1, \ldots, W_n)$. There are no other restrictions on the set of possible observed data distributions, and the resulting statistical model $\mathcal{M}$ is semiparametric.

2.2 The Conditional Average Treatment Effect (CATE)

The goal of the SEARCH trial is to estimate the effect of a strategy for immediate and streamlined ART for all HIV diagnosed persons on the 5-year cumulative HIV incidence in rural East African communities. A common target of inference is the population average treatment effect $E[Y(1)] - E[Y(0)]$ or its relative counterpart $E[Y(1)]/E[Y(0)]$, where $Y(a)$ denotes the counterfactual cumulative incidence under treatment level $A = a$. This causal parameter is the difference in the expected outcomes if all communities (in some hypothetical target population) were to receive the intervention and if all communities (in some hypothetical target population) were to receive the control.

An alternative estimand involves conditioning on the measured baseline covariates of the study communities [95, 116–118]:

$$\psi^F = \frac{1}{n} \sum_{i=1}^{n} E[Y_i(1) - Y_i(0) \mid W^n]$$

where $Y_i(a)$ denotes the counterfactual cumulative incidence under treatment level $A = a$ for unit i. This parameter is the difference in the expected counterfactual outcomes, treating the baseline covariates of the study communities as fixed. As a result, the parameter is data-adaptive; its value changes with the sample of study units. Nonetheless, ψ^F can be interpreted as the intervention effect, given the covariates the sample units. Greater generalizability is up to the reader and not implicitly assumed in the parameter specification. Furthermore, by obviating estimation of the covariate distribution, estimators of the conditional parameter will also often be more precise than those of the population parameter [95, 116–118].

2.3 Estimation

Since the intervention is randomized within matched pairs, the causal parameter is readily identifiable from the conditional distribution of the observed data. The statistical estimand is

$$\Psi(P_0^n) = \frac{1}{n} \sum_{i=1}^{n} \left[E_0(Y_i \mid A_i = 1, W_i) - E_0(Y_i \mid A_i = 0, W_i) \right]$$
$$= \frac{1}{n} \sum_{i=1}^{n} \left[\bar{Q}_0(1, W_i) - \bar{Q}_0(0, W_i) \right]$$

where $\bar{Q}_0(A, W)$ denotes the conditional mean outcome, given the intervention A and covariates W. In other words, the target parameter is the average difference in the strata-specific expected HIV incidence under the intervention and control for the n study communities.

This estimand is still random through the vector of covariates $W^n = (W_1, \ldots, W_n)$. The
true value ψ_0 depends on the sample of n units.

An intuitive estimator of ψ_0 is the average difference in outcomes within matched pairs:

$$\hat{\psi}_{unadj} = \frac{1}{n/2} \sum_{j=1}^{n/2} \left(Y_{j1} - Y_{j2} \right)$$

where the observations within matched pair j have been ordered such that the first corre-
sponds to the intervention, $A_{j1} = 1$, and the second the control, $A_{j2} = 0$. This estima-
tor is equivalent to taking the difference in the average outcomes among intervention units
$\bar{Q}_n(1) = E_n(Y|A = 1)$ and the average outcomes among control units $\bar{Q}_n(0) = E_n(Y|A = 0)$.
Since the intervention is randomized, the unadjusted estimator is unbiased for the parame-
ter of interest, given the vector of covariates W^n. (See Appendix B.1 for the accompanying
proof.) When the measured covariates are predictive of the outcome, this simple difference-in-
means estimator tends to be *inefficient* as it fails to adjust for measured covariates. Despite
recent advances in matching algorithms [93, 113, 114], there is likely to be some residual
imbalance on pre-intervention determinants of the outcome within matched pairs. Further-
more, even if we succeeded in matching well on all available characteristics, there might be
additional baseline covariates that are predictive of the outcome, but were unavailable during
the matching process. In the SEARCH trial, for example, baseline population HIV RNA
levels are thought to be a major driver of incidence but were unavailable during matching.

An alternative approach is to use TMLE, which can provide an unbiased and efficient
estimate of the intervention effect. A TMLE for $\Psi(P_0^n)$ is given by the following substitution
estimator:

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

where $\bar{Q}_n^*(A, W)$ denotes a targeted estimate of the conditional mean function $E_0(Y|A, W)$.
In general, this targeting step is used to achieve the optimal bias-variance trade-off for the
parameter of interest and to solve the efficient score equation [40]. We refer the reader to
van der Laan and Rose [3] for a detailed discussion and worked examples of TMLE. In an
adaptive pair-matched trial, a TMLE for $\Psi(P_0^n)$ can be implemented as follows.

1. Estimate the conditional mean function $\bar{Q}_0(A, W)$ by regressing the outcome Y on the
   treatment A and covariates W, while ignoring the dependence in the data.

   - For a binary outcome or a bounded continuous outcome, the negative log likeli-
     hood is a valid loss function and provides stability in the context of sparsity [137].
     Specifically, the boundedness property of the logistic function guarantees the pre-
     dicted outcomes are within the appropriate range (e.g. [0,1] for a proportion).
   - For a continuous outcome, initial estimation of the conditional mean $\bar{Q}_0(A, W)$
     can also be based on linear regression, which can yield more power than non-linear
     (logistic) regression in randomized trials. In particular, Rubin and van der Laan

[122] detail the use of least squares regression to optimize the fit of $\bar{Q}_0(A, W)$ to achieve the lowest possible variance.

- Initial estimation can also be based on an *a priori* specified data-adaptive method, such as Super Learner [39]. In all cases, there is no risk of bias due to model misspecification [122, 124, 125].

2. If the initial regression model included an intercept and a main term for the exposure, the estimator of the conditional mean outcome $\bar{Q}_n(A, W)$ is already targeted. Skip to step 3. Otherwise, update the initial estimator as follows.

- If logistic regression was used for initial estimation, then the following fluctuation sub-model is appropriate:

$$logit\big[\bar{Q}_n(A, W)(\epsilon)\big] = logit\big[\bar{Q}_n(A, W)\big] + \epsilon H(A),$$
$$\text{where } H(A) = \left(\frac{\mathbb{I}(A = 1)}{P_0(A = 1)} - \frac{\mathbb{I}(A = 0)}{P_0(A = 0)}\right)$$

and ϵ is the univariate parameter. If linear regression was used, then the following fluctuation sub-model is appropriate:

$$\bar{Q}_n(A, W)(\epsilon) = \bar{Q}_n(A, W) + \epsilon H(A)$$

with ϵ and $H(A)$ are defined as above. In practice, run logistic (linear) regression of the outcome Y on the covariate $H(A)$, using the initial estimate as offset. Then plug the estimated coefficient ϵ_n into the fluctuation model to yield the targeted estimates $\bar{Q}_n^*(A, W) = \bar{Q}_n(A, W)(\epsilon_n)$.

3. Take the sample average of the differences in the expected outcomes:

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

where $\bar{Q}_n^*(1, W_i)$ denotes the expected outcome for unit i under the intervention and $\bar{Q}_n^*(0, W_i)$ denotes the expected outcome for unit i under the control. It is worth emphasizing that empirical mean, here, is part of target parameter mapping; we are not estimating the covariate distribution as would be required for the population average treatment effect.

In practice, many cluster randomized trials have a limited number of (conditionally) independent units. For example, there are only 16 conditionally independent pairs in the SEARCH trial. As a result, the number of parameters in the regression model for $\bar{Q}_0(A, W)$ can quickly approach the number of observations. Therefore, the curse of dimensionality can prevent adjustment for all the measured covariates W or the inclusion of multiple interaction terms. Nonetheless, it is often possible to adjust for a single or few covariates and obtain

efficiency gains without risk [124, 125]. Furthermore, when the regression model for $\bar{Q}_0(A, W)$
includes an intercept and the exposure A as a main term, the initial estimator is already
targeted. Thus, we can obtain an unbiased and more efficient estimator in two steps: estimate
$\bar{Q}_0(A, W)$ with main terms linear or logistic regression, and take the sample average of the
differences in the expected outcomes under the treatment and control.

2.4 Statistical Inference

As established in Appendix B.2, both the unadjusted estimator and the TMLE are asymp-
totically linear and normally distributed. Briefly, an estimator is asymptotically linear if
the difference between the estimator and the estimand behaves (in first order) as an em-
pirical mean of a function, known as the influence curve, of the unit data [3]. Then the
limit distribution of the standardized estimator is normal with mean 0 and variance given
by the variance of its influence curve. With an estimate of the influence curve and thereby
an estimate of the variance, the standard normal distribution can be used for confidence
interval construction and hypothesis testing in large studies. For trials with limited num-
bers of (conditionally) independent units, the Student's t-distribution with $n/2$-1 degrees of
freedom is an appropriate alternative to the standard normal distribution. Randomization
inference, in contrast, may not be appropriate, as it is testing a different null hypothesis of
a constant treatment effect (e.g. $Y_i(0) = Y_i(1) \; \forall i$) [138, 139]. The causal and statistical
estimands, considered here, are in terms of a sample average effect over the study units.

The influence curve for the TMLE of $\Psi(P_0^n)$ in a trial with adaptive pair-matching is the
following function of the paired data (proof in Appendix B.2):

$$IC(\bar{O}_j) = \bar{D}^*(\bar{O}_j) - E_0\big[\bar{D}^*(\bar{O}_j)\big|W^n\big]$$

$$\bar{D}^*(\bar{O}_j) = \frac{1}{2}\bigg\{D^*(O_{j1}) + D^*(O_{j2})\bigg\}$$

$$D^*(O_i) = \left(\frac{\mathbb{I}(A_i = 1)}{P_0(A_i)} - \frac{\mathbb{I}(A_i = 0)}{P_0(A_i)}\right)\big(Y_i - \bar{Q}(A_i, W_i)\big)$$

where $\bar{Q}(A, W)$ denotes the limit of the targeted estimator of the conditional mean function
$\bar{Q}_0(A, W)$ and where the marginal probability of being assigned the treatment or the control
is known: $P_0(A) = 0.5$. Through the conditional expectation of $\bar{D}^*(\bar{O}_j)$, given the vector of
covariates W^n, the influence curve relies on the true but unknown conditional mean outcome
$\bar{Q}_0(A, W)$:

$$E_0\big[\bar{D}^*(\bar{O}_j)\big|W^n\big] = \frac{1}{2}\bigg\{\big(\bar{Q}_0(1, W_{j1}) - \bar{Q}(1, W_{j1})\big) - \big(\bar{Q}_0(0, W_{j1}) - \bar{Q}(0, W_{j1})\big)$$

$$+ \big(\bar{Q}_0(1, W_{j2}) - \bar{Q}(1, W_{j2})\big) - \big(\bar{Q}_0(0, W_{j2}) - \bar{Q}(0, W_{j2})\big)\bigg\}$$

This term captures deviations between the true and estimated mean outcomes for observations within a matched pair. The influence curve for the unadjusted estimator $\hat{\psi}_{unadj}$ is analogous, but with $\bar{Q}(A, W)$ replaced with the limit of the treatment-specific mean $\bar{Q}_n(A) = E_n(Y|A)$. For either estimator, there is no contribution from the covariate distribution, which is considered fixed.

The asymptotic variance of the unadjusted estimator or the TMLE is then given by the variance of its influence curve, divided by $n/2$. Improved estimation of the conditional mean outcome $\bar{Q}_0(A, W)$ leads to more precise estimators of intervention effect $\Psi(P_0^n)$. Specifically, if this conditional mean is consistently estimated (i.e. if $\bar{Q}(A, W) = \bar{Q}_0(A, W)$), then the term, involving deviations between the true and estimated means, is zero, and the estimator of $\Psi(P_0^n)$ is asymptotically efficient. In other words, the estimator's influence curve equals the efficient influence curve, and the estimator has lowest possible variance among a large class of estimators [3]. Otherwise, the estimator is still be unbiased, but does not achieve the efficiency bound. When the baseline covariates W impact the outcome, the targeted estimator of the conditional mean outcome $\bar{Q}_n^*(A, W)$ is expected to be closer to the true mean $\bar{Q}_0(A, W)$ than the unadjusted estimator $\bar{Q}_n(A)$. As a result, the asymptotic variance of the TMLE $\hat{\psi}_{adj}$ is often smaller than that of the unadjusted estimator $\hat{\psi}_{unadj}$. Thus, for both individual and cluster randomized trials, TMLE is often a more efficient estimator of the CATE than the unadjusted estimator.

Consistent estimation of the influence curve and thereby the asymptotic variance rely on consistent estimation of this conditional mean $\bar{Q}_0(A, W)$, which might be particularly challenging when n is small, as common in cluster randomized trials. Nonetheless, we can conservatively approximate the influence curve of the unadjusted estimator $\hat{\psi}_{unadj}$ or the TMLE $\hat{\psi}_{adj}$ by the difference in residuals within matched pairs (proof in Appendix B.2):

$$\hat{IC}_{unadj}(\bar{O}_j) = \left(Y_{j1} - \bar{Q}_n(1)\right) - \left(Y_{j2} - \bar{Q}_n(0)\right)$$
$$\hat{IC}_{adj}(\bar{O}_j) = \left(Y_{j1} - \bar{Q}_n^*(1, W_{j1})\right) - \left(Y_{j2} - \bar{Q}_n^*(0, W_{j2})\right)$$

respectively. Again, $\bar{Q}_n(A)$ denotes an unadjusted estimate of the treatment-specific mean, $\bar{Q}_n^*(A, W)$ denotes a targeted estimate of the conditional mean outcome, and observations in matched pair j have been ordered such that the first corresponds to intervention ($A_{j1} = 1$) and the second to the control ($A_{j2} = 0$). An asymptotically conservative variance estimator is then given by the sample variance of the estimated influence curve, divided by $n/2$. For $\hat{\psi}_{unadj}$, this is equivalent to the sample variance of the within pair differences, divided by $n/2$, and is commonly recommended for pair-matched randomized trials [96] even though it is known to be conservative if the conditional parameter is the target of inference [95, 116–118]. To obtain a less conservative variance estimator for $\hat{\psi}_{unadj}$, Abadie and Imbens [118] proposed a matching estimator, involving the variance of pairs-of-pairs with similar covariates. Our approach to reduce the true variance of the estimator and obtain a less conservative variance estimate is through adjustment with TMLE. In most practical settings, the sum of squared adjusted residuals is smaller than the sum of squared unadjusted residuals. Thereby, the estimated variance of the TMLE is often smaller than the estimated variance of

the unadjusted algorithm. In summary, this implies that covariate adjustment with TMLE results in a more precise estimator (i.e. smaller true variance) and a less conservative variance estimator.

We also briefly note that a randomized trial with adaptive pair-matching will often be more efficient for estimation of the CATE than a randomized trial without matching. The designs will only have the same efficiency bound if the conditional mean outcome is consistently estimated (i.e. $\bar{Q}(A, W) = \bar{Q}_0(A, W)$). In practice, we expect there to be some deviations between the true and estimated means. If these deviations are positively correlated within matched pairs, the asymptotic variance of the TMLE will be smaller in the adaptive trial than in the completely randomized trial. In finite samples, we also expect there to be an efficiency gain from adaptive pair-matching. Mainly, if we succeed in matching pairs on predictive covariates, then the sample covariance of the residuals within matched pairs will be positive and the adaptive design will yield more power. We refer the reader to Appendix B.3 for further details and associated proofs.

2.5 Simulation Study

We present the following set of simulations to demonstrate (1) implementation of the above estimators, (2) the potential gain in efficiency with adaptive pair-matching, and (3) the further gain with adjustment during the analysis due to having a more precise estimator and a less conservative variance estimator. These simulations are not intended to represent the full complexities of a cluster randomized trial. (To be clear, these simulations were not the ones used when developing the design and analysis of the SEARCH trial.) Nonetheless, they explore some of the challenges faced, such as rare outcomes, the inability to match on all baseline covariates, and limited numbers of conditionally independent units. All simulations were done in `R` v3.0.1 [46].

Data Generating Process & Estimators

For $n = 32$ units, three baseline covariates $W = (W1, W2, W3)$ were independently drawn from a normal distribution with mean 0 and standard deviation 1. A fourth covariate Z was generated as a function of these baseline covariates and random noise U_Z:

$$Z = expit\big[-0.25 + 0.5^*W1 + W2 + 2^*W3 + 0.5^*U_Z\big]/4$$

where the *expit* function is the inverse of the *logit* function and U_Z was drawn independently from a normal with mean 0 and standard deviation 1. To imitate adaptive pair-matching, the nonbipartite matching algorithm (`nbpMatching` v1.3.6 [111]) was applied to the set of n covariates $W^n = (W_1, \ldots, W_n)$ with $W_i = (W1_i, W2_i, W3_i)$. Within the resulting 16 matched pairs, the exposure A was randomized. As before, A is binary indicator, equaling 1 if the unit was randomized to the intervention and 0 otherwise. Finally, the outcome Y was

generated as

$$Y = expit\big[\beta_0 + 0.5^*W1 + 0.5^*W2 + 0.5^*W3 + 7^*Z - A + 0.25^*A^*Z\big]/15 + U_Y$$

where random noise U_Y was drawn independently from a uniform distribution with minimum 0 and maximum 0.025. Dividing by 15 was done to scale the outcome Y, representing a proportion, to be within plausible ranges for the cumulative incidence of HIV. The term β_0 was set to either -2 or 0.5 to examine the performance of the estimators when the outcome was rare ("Simulation A") or more common ("Simulation B"). To simulate the null scenario, the treatment was randomly assigned within pairs but the outcomes generated as if all communities received the control ($A = 0$). For comparison, we also simulated equivalent data for a non-matched randomized trial with balanced allocation of the treatment.

Over 5000 data sets, we examined the performance of the unadjusted estimator and TMLE. For the latter, we compared linear to logistic main terms regression with various adjustment sets. Linear regression can result in more efficient estimation, by minimizing the empirical variance of the influence curve [122]. With rare outcomes, however, logistic regression can provide stability, by guaranteeing the predicted outcomes respect the model bounds (i.e. are in [0,1]). Therefore, we expected the TMLE with logistic regression to result in better performance when the outcome was rare (Simulation A) and the TMLE with linear regression to result in better performance when the outcome was more common (Simulation B). In terms of adjustment sets, we compared regression models with main terms for the exposure A and the covariate Z as well as regression models with main terms for the exposure A, the matching covariates W and the remaining covariate Z. Recall Z was an important determinant of the outcome but not used in matching. We expected that the fully adjusted estimator (TMLE with main terms for (A, W, Z)) would suffer from over-fitting. Since main terms regression models were used, the fluctuation step of the TMLE algorithm did not provide an update. In all cases, there was no risk of bias due to regression model misspecification [124, 125]. Inference was based on the sample variance of the estimated influence curve and the Student's t-distribution with 15 degrees of freedom. The corresponding TMLE implementation and proof of statistical inference for the non-matched randomized trial are given in Appendix B.3.

Results

Recall the true value of the statistical estimand depends on the $n = 32$ communities in the sample. Table 2.1 shows the minimum, mean and maximum value of the intervention effect ψ_0 over the 5,000 simulated data sets. For comparison, the table also gives the corresponding summaries of the exposure-specific effects: $\psi_0(a) = \frac{1}{n}\sum_{i=1}^{n} E_0(Y|A_i = a, W_i, Z_i)$. This estimand is the sample average of the conditional mean outcome, setting the exposure $A = a$ and given the covariates (W, Z). For Simulation A, representing a rare outcome, the average values of the effect under the exposure $\psi_0(1)$ and the control $\psi_0(0)$ were 0.024 and 0.032, respectively. The corresponding mean value of the intervention effect ψ_0 was

	$\psi_0(1)$			$\psi_0(0)$			$\psi_0 = \psi_0(1) - \psi_0(0)$		
	min	mean	max	min	mean	max	min	mean	max
Simulation A	0.018	0.024	0.031	0.023	0.032	0.043	-0.012	-0.009	-0.005
Simulation B	0.038	0.050	0.061	0.050	0.061	0.069	-0.013	-0.011	-0.007

Table 2.1: *Summary of the true value of the exposure-specific effects $\psi_0(a) = 1/n \sum_i E_0(Y_i \mid A_i = a, W_i, Z_i)$ and the target parameter ψ_0 over 5,000 simulations of $n = 32$ communities. The rows indicate the setting with Simulation A corresponding to a rare outcome and Simulation B corresponding to a more common outcome. Recall the true value is dependent on the sample.*

-0.009, translating to 26.41% reduction in the incidence of the outcome (on average). For Simulation B, representing a more common outcome, the average values of the conditional effect under the exposure $\psi_0(1)$ and the control $\psi_0(0)$ were 0.05 and 0.061, respectively. The corresponding average value of the target parameter ψ_0 was -0.011, translating to a 17.90% reduction in the incidence of the outcome (on average).

For Simulation A, Table 2.2 illustrates the performance of the estimators over 5,000 simulated data sets. All estimators were unbiased. As expected, there was an efficiency gain with matching. The standard deviation (square root of the variance of the point estimates) of the unadjusted estimator was 1.58 times higher without matching than with matching. Likewise, the attained power (proportion of simulated trials where the null hypothesis was correctly rejected) jumped from 34% to 64% with matching. As expected, adaptive pair-matching on the three covariates W reduced variability in the outcomes within matched pairs. The coefficient of variation, measuring of the variability in outcomes between units in the absence of the intervention, was $k = 0.53$, while the matched-pair coefficient of variation, measuring of the variability in outcomes within matched pairs in the absence of the intervention, was $k_m = 0.29$ [96].

There was also an efficiency gain from adjustment. For the non-matched design, the standard deviation of the unadjusted estimator was 1.58 times higher than the standard deviation of the TMLE, using linear regression to adjust for Z. The corresponding power increased from 34% to 72%. For the adaptive design, the standard deviation of the unadjusted estimator was 1.13 times higher than the standard deviation of the TMLE, using linear regression to adjust for Z. The corresponding attained power increased from 64% to 74%. For both designs, there was a further precision gain by using logistic regression to adjust for Z. Under sparsity, logistic regression can be more stable than linear regression and is guaranteed to yield parameter estimates within the appropriate range (i.e. [0,1] for proportions) [137]. While there was some power gain from adjusting for all four covariates (W, Z), there was also a risk in over-fitting the regression model and under-estimating the variance. Recall the variance estimators in Section 2.3 are asymptotically conservative, and the simulations represent finite samples. Indeed, with a main terms regression model for the conditional mean outcome, there were 5 parameters with only 16 conditionally indepen-

	Bias	Std. Dev.	Std. Error	t-stat	CI Cov.	Power
Simulation A			No Matching			
Unadj.	-0.00011	0.0054	0.0053	-1.6	96	34
TMLE linear for Z	-0.00008	0.0034	0.0032	-2.7	94	72
TMLE logit for Z	-0.00008	0.0033	0.0030	-2.9	94	78
TMLE linear for (W, Z)	-0.00011	0.0033	0.0027	-3.2	91	82
TMLE logit for (W, Z)	-0.00013	0.0031	0.0024	-3.6	90	88
			Adaptive Pair-Matching			
Unadj.	-0.00004	0.0034	0.0035	-2.5	96	64
TMLE linear for Z	-0.00004	0.0030	0.0030	-2.9	96	74
TMLE logit for Z	-0.00006	0.0030	0.0028	-3.2	94	80
TMLE linear for (W, Z)	-0.00004	0.0030	0.0029	-3.1	95	79
TMLE logit for (W, Z)	-0.00008	0.0029	0.0026	-3.5	93	84
Simulation B			No Matching			
Unadj.	-0.00007	0.0063	0.0062	-1.8	95	38
TMLE linear for Z	-0.00009	0.0035	0.0033	-3.4	94	88
TMLE logit for Z	-0.00013	0.0037	0.0036	-3.1	95	84
TMLE linear for (W, Z)	-0.00015	0.0032	0.0026	-4.2	91	96
TMLE logit for (W, Z)	-0.00037	0.0036	0.0031	-3.7	91	91
			Adaptive Pair-Matching			
Unadj.	-0.00007	0.0036	0.0036	-3.1	96	80
TMLE linear for Z	-0.00008	0.0030	0.0030	-3.8	96	92
TMLE logit for Z	-0.00011	0.0031	0.0033	-3.4	97	89
TMLE linear for (W, Z)	-0.00010	0.0029	0.0028	-4.1	95	95
TMLE logit for (W, Z)	-0.00023	0.0031	0.0032	-3.6	96	90

Table 2.2: For Simulation A (rare outcome) and Simulation B (more common outcome), summary of the estimator performance over 5,000 simulations of $n = 32$ communities. The rows indicate the estimator and the columns the performance metric: bias as the average deviation between the point estimate and sample-specific true value; standard deviation as the square root of the variance of the point estimates; standard error as the average standard error estimate based on the influence curve; t-statistic as the average value of the test statistic (point estimate divided by standard error estimate); confidence interval coverage as the proportion of intervals containing the true parameter value (in %), and power as the proportion of studies correctly rejecting the null hypothesis (in %).

dent units. As a result, the confidence interval coverage (proportion of studies containing the true parameter value) was less than the nominal rate of 95% for the fully adjusted estimator. Likewise, the type I error rate (proportion of studies falsely rejecting the null hypothesis) was greater than $\alpha = 0.05$ for the fully adjusted estimator, as shown in Table 1 of Appendix B.4.

Conversely, for both the unadjusted estimator and the TMLE only adjusting for Z, there was good confidence interval coverage and control of type I error rates. Indeed, there was some evidence of over-coverage of confidence intervals and conservative Type I error rates for the unadjusted estimator in both designs, as predicted by theory.

The results for Simulation B, representing a more common outcome, are also given in Table 2.2 and largely echoed the above findings. Because the exposure was randomized, all estimators were unbiased. As before, there was a substantial efficiency gain with matching. Adaptive matching on the three covariates $W = (W1, W2, W3)$ reduced variability in the outcomes within pairs. The coefficient of variation was $k = 0.27$, while the matched-pair coefficient of variation was $k_m = 0.14$. Again, there was also a substantial precision gain from adjustment. With a more common outcome, however, there was a greater gain in power from adjusting for Z with linear regression than logistic regression for both designs. Here, minimizing the sum of squared residuals helped to minimize the empirical variance of the influence curve and thereby maximize the empirical efficiency [122]. With the fully adjusted estimator, again there was some risk of over-fitting and inference was optimistic. In contrast, for both the unadjusted estimator and the TMLE adjusting only for Z, there was good confidence interval coverage as well as Type I error control (Table 1 of Appendix B.4). In summary, our finite sample simulations support our theoretical results: adaptive pair-matching yields more power than complete randomization, and further efficiency gains can be attained through adjustment during the analysis.

2.6 Discussion

To our knowledge, this is the first work to study and articulate the consequences of adaptive pair-matching for estimation of the average treatment effect, given the baseline covariates of the n study units. This work was motivated by SEARCH trial, which aims to estimate the effect of immediate ART, delivered in a streamlined fashion, on the five-year cumulative incidence of HIV. The decision to pair-match communities in the trial was motivated by a desire to protect study credibility and by the potential to increase study power. Through careful definition of the data generating experiment, we recognized that the design would not yield $n/2$ i.i.d. paired units, as current practice assumes. Instead, by constructing the matched pairs as a function of the baseline covariates of all candidate communities, the adaptive design results in n dependent units and $n/2$ conditionally independent units, given the baseline covariates of the study communities.

To the best of our understanding, adaptive pair-matching is a common design and has been implemented in other cluster randomized trials (e.g. [107–109]). In practice, adaptive pair-matching (a.k.a. "nonbipartite matching") can be carried out with standard software. For example, the `nbpMatching` package [111] in `R` and the corresponding web application will generate the set of optimal matched pairs as function of a user-supplied matrix of covariates [113, 114]. These tools allow the user to weight covariates differently (e.g. on importance or relevance to the outcome) and to specify the maximum number of matches - choices, which

should be driven by subject matter knowledge as well as resource constraints.

We focused on estimation of the CATE. By obviating estimation of the covariate distribution, estimators of the conditional parameter will often be less variable than estimators of the population parameter [95, 116–118]. We contrasted the unadjusted estimator with TMLE adjusting for baseline covariates. We provided a step-by-step implementation of the latter estimator and detailed proofs of inference. Both estimators can be implemented ignoring the dependence in the data and with standard software, such as the `tmle` [44] and `ltmle` [45] packages in `R`. Asymptotically conservative inference can obtained with the sample variance of the pairwise differences in residuals, divided by $n/2$. When the baseline covariates are predictive of the outcome, the unadjusted estimator will be less efficient than the TMLE. Furthermore, the estimated variance of the TMLE will often be less conservative than that of the unadjusted estimator.

Finite sample simulations were used to evaluate estimator performance and verify our theoretical results. Since the intervention was randomized, all estimators were unbiased [122, 124, 125]. There was an efficiency gain with matching and a further gain with adjustment. When the outcome was quite rare, adjusting for a single baseline covariate with logistic regression yielded more power than adjustment with linear regression. When the outcome was more common, the converse was observed. While the variance estimators are asymptotically conservative, there was some risk of over-adjusting in small trials. Indeed, with only 16 (conditionally) independent units, adjusting for all 4 baseline covariates resulted in under-coverage of the confidence intervals and higher than nominal Type I error rates.

Previously, Imai *et al.* [94] suggested, "randomization by cluster without prior construction of matched pairs, when pairing is feasible, is an exercise in self-destruction." Our work also suggests that asymptotically and in finite samples, a randomized trial with adaptive pair-matching will often be more efficient for estimation of the CATE than its completely randomized counterpart. The trials will only have the same efficiency bound when the conditional mean outcome, given the exposure and covariates, is consistently estimated. In practice, we expect there to be some deviations between the true and estimated means. When these deviations are positively correlated within matched pairs, the design with adaptive pair-matching will be more efficient (Appendix B.3). In finite samples, pair-matching will also often result in a positive covariance of the residuals (deviations between the observed and predicted outcomes) within matched pairs and thereby smaller finite sample variance.

Overall, adaptive pair-matching is an intuitive strategy to group candidate units on similarity in their baseline covariates. Pair-matching will protect study credibility. Combining subject matter knowledge with modern matching algorithms (e.g. `nbpMatching` [111]) is likely to result in studies, where pair-matching substantially improves study power. We recommend specifying the intervention effect in terms of the conditional parameter, which considers the covariate distribution as fixed and obviates its estimation, resulting in less variable estimators. We also recommend adjusting for baseline variables as the data allow. Simulations, such as those presented here, can help inform the practitioner as to the optimal adjustment set. Future work will involve the use of cross-validation to data-adaptively select for the adjustment set. We also plan to formally study the asymptotic and finite sam-

ple properties of analysis approaches based on covariate-adjusted residuals for estimation
and inference of the CATE [96, 107, 128]. We will also investigate the impact of adaptive
stratification on estimation and inference for both the population and conditional average
treatment effect. While our work was motivated by a cluster randomized trial with the out-
come of cumulative incidence, the results are generally applicable to other trials with binary
or continuous outcomes.

Chapter 3

Targeted Estimation and Inference for the Sample Average Treatment Effect

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 [27, 48]. 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 n study units [19]. In other words, the SATE is the intervention effect for the sample at hand. The conditional effect is the average difference in the expected counterfactual outcomes, treating the measured baseline covariates of the study units as fixed [116]. In other words, the CATE is the intervention effect, averaging out the unmeasured factors contributing to the counterfactual outcomes but conditional on the measured factors contributing to the counterfactual outcomes. As detailed below, the exact interpretation and the variability of the CATE depend on the conditioning set (i.e. on the set of measured covariates). Briefly, the CATE will be constant across repeated studies if the measured covariates (e.g. region) are constant and will change across repeated studies if the measured covariates (e.g. HIV prevalence) are not constant. In contrast, the SATE changes with each new selection or sample of units. The sample effect will only equal the conditional effect if all factors impacting the outcome are measured. Another key difference between the three parameters is in the variance of common estimators. As shown by Imbens [117] 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 work to propose using targeted maximum likelihood estimation (TMLE) for the SATE. TMLE is a general algorithm for constructing double robust, semiparametric efficient, substitution estimators [3, 40]. 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 [117] 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 [16] and motivate our discussion with the Sustainable East Africa Research in Community Health (SEARCH) trial for HIV prevention and treatment (NCT01864603) [13]. This chapter was reproduced with permission from Balzer *et al.* [140].

3.1 Causal Model and 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 [129]. 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 $P_0(W_1, \ldots, W_n)$. For further details, see Balzer *et al.* [129].

This data generating process can be described by the following structural causal model (SCM) [14, 16]. 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:

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

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 $A = a$:

$$W = f_W(U_W)$$
$$A = a$$
$$Y(a) = f_Y(W, a, U_Y)$$

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):

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

where the subscript (U, X) denotes the expectation over the distribution $P_{U,X}$ (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, $\Psi^{\mathcal{P}}(P_{U,X})$ 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 $\Psi^{\mathcal{P}}(P_{U,X})$ 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 [19]:

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

This is simply the intervention effect for the study units. For the SEARCH trial, $\Psi^{\mathcal{S}}(P_{U,X})$ 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 the 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 [116]:

$$\Psi^{\mathcal{C}}(P_{U,X}) = \frac{1}{n}\sum_{i=1}^{n} E_{U,X}\big[Y_i(1) - Y_i(0)\big|W_1, \ldots, W_n\big]$$
$$= \frac{1}{n}\sum_{i=1}^{n} E_{U,X}\big[Y_i(1) - Y_i(0)\big|W_i\big]$$

where the second equality holds under our assumption that study units are causally indepen-
dent (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, $\Psi^{\mathcal{C}}(P_{U,X})$ is interpreted as the average difference in the expected counterfactual cumu-
lative 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 $\Psi^{\mathcal{C}}(P_{U,X})$ 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 of the CATE 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.

3.2 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 [27, 48]:

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

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 distribu-
tion of the observed data.) Briefly, the randomization assumption states that the counterfac-
tual 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 [48]. The positivity as-
sumption 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 $\Psi^{\mathcal{P}}(P_{U,X,0})$ is easily identified as

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

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 [27]. For the SEARCH trial, $\Psi_0^{\mathcal{P}}(P_0)$ 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 $\Psi_0^{\mathcal{P}}(P_0) = \psi_0^{\mathcal{P}}$ 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. [119–121, 123, 124]).

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

$$\Psi_0^{\mathcal{C}}(P_0) = \frac{1}{n} \sum_{i=1}^{n} \left[E_0(Y_i|A_i = 1, W_i) - E_0(Y_i|A_i = 0, W_i) \right]$$

$$= \frac{1}{n} \sum_{i=1}^{n} \left[\bar{Q}_0(1, W_i) - \bar{Q}_0(0, W_i) \right] \tag{3.1}$$

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, $\Psi_0^{\mathcal{C}}(P_0)$ 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 $\Psi_0^{\mathcal{C}}(P_0)$ 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 [19, 117, 126, 141]. To elaborate, let us use the structural causal model $\mathcal{M}^F$ to rewrite the SATE in terms of the CATE:

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

$$= \frac{1}{n} \sum_{i=1}^{n} f_Y(W_i, 1, U_{Y_i}) - f_Y(W_i, 0, U_{Y_i})$$

$$= \frac{1}{n} \sum_{i=1}^{n} E_{U,X} \left[Y_i(1) - Y_i(0) \big| W_i, U_{Y_i} \right]$$

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.* [129]. 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.

3.3 Estimation and Inference

There are many well-established algorithms for estimation of the population parameter $\Psi_0^{\mathcal{P}}(P_0)$. 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 measured covariates $P_0(A = 1|W)$ (e.g. [22, 47–49]). 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. [27, 142, 143]). 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. [31, 52, 53, 55]) and TMLE (e.g. [3, 40]). 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 [56, 137]. We focus our discussion on TMLE in a randomized trial and provide generalizations to an observational setting in Appendix C.3.

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 [39]. 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. [55, 124]). 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 run-

ning logistic regression[1] 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 $logit(x) = log\{x/(1-x)\}$ 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:

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

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

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 [124, 125]. Lastly, the targeted estimates are substituted into the parameter mapping:

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

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 [40]:

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

with influence curve

$$D^{\mathcal{P}}(\bar{Q}, g_0)(O) = D_Y(\bar{Q}, g_0)(O) + D_W(\bar{Q}, g_0)(O)$$

$$D_Y(\bar{Q}, g_0)(O) = \left(\frac{\mathbb{I}(A=1)}{g_0(1|W)} - \frac{\mathbb{I}(A=0)}{g_0(0|W)}\right)\left(Y - \bar{Q}(A, W)\right)$$

$$D_W(\bar{Q}, g_0)(O) = \bar{Q}(1, W) - \bar{Q}(0, W) - \psi_0^{\mathcal{P}}$$

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

---

[1]As detailed in Gruber and van der Laan [137], the same procedure can be applied to a bounded continuous outcome and adds robustness in the context of sparsity.

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 [144]. 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 [124, 125]. 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` [44] and `ltmle` [45] packages in `R` [46].

For the Conditional Parameter

The TMLE for the population parameter $\Psi_0^{\mathcal{P}}(P_0)$ also serves as an estimator of the conditional parameter $\Psi_0^{\mathcal{C}}(P_0)$. 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. 3.1) 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 [95, 117, 118, 129].

The TMLE is also a consistent and asymptotically linear estimator of the conditional parameter [129]:

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

with influence curve given by

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

where $\mathbf{W} = (W_1, \ldots, W_n)$ 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:

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

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:

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

They will only have the same efficiency bound when there is no variability in the treatment effect across strata of covariates (i.e. when $Var[\bar{Q}_0(1, W) - \bar{Q}_0(0, W) - \psi_0^{\mathcal{P}}] = 0$). 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

$$D_n^{\mathcal{C}}(O_i) = D_{Y,n}(O_i) = \left(\frac{\mathbb{I}(A_i = 1)}{g_0(1|W_i)} - \frac{\mathbb{I}(A_i = 0)}{g_0(0|W_i)}\right)\left(Y_i - \bar{Q}_n^*(A_i, W_i)\right) \tag{3.2}$$

[129]. 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 [129].

For the Sample Parameter

For a randomized trial, Neyman [19] 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 $\{Y_i(1), Y_i(0)\}$ [19]. Imbens [117] 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 effects 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 an alternate method for variance estimation.

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

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

with influence curve

$$D^{\mathcal{S}}(\bar{Q}, g_0)(U, X) = D^{\mathcal{C}}(\bar{Q}, g_0)(O) - D^F(U, X)$$
$$D^F(U, X) = Y(1) - Y(0) - \left[\bar{Q}_0(1, W) - \bar{Q}_0(0, W)\right]$$

The proof is given in Appendix C.1. 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 $D^{\mathcal{S}}$ is given by the influence curve for TMLE of the conditional parameter $D^{\mathcal{C}}$ minus a non-identifiable piece, which captures the deviations between the unit-specific treatment effect and expected effect within covariate strata:

$$D^F(U_i, X_i) = Y_i(1) - Y_i(0) - \left[E_0(Y_i|A_i = 1, W_i) - E_0(Y_i|A_i = 0, W_i)\right]$$
$$= Y_i(1) - Y_i(0) - \left[E_{U,X,0}(Y_i(1)|W_i) - E_{U,X,0}(Y_i(0)|W_i)\right]$$
$$= Y_i(1) - Y_i(0) - E_{U,X,0}\left[Y_i(1) - Y_i(0)\big|W_i\right]$$

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

$$Var[D^{\mathcal{S}}(U, X)] = Var\left[D^{\mathcal{C}}(O)\right] + Var\left[D^F(U, X)\right] - 2Cov\left[D^{\mathcal{C}}(O), D^F(U, X)\right]$$
$$= Var\left[D^{\mathcal{C}}(O)\right] - Var\left[D^F(U, X)\right]$$

The proof is given in Appendix C.2. 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. 3.2), 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 [116] and discussed in Imbens [117]. We note that the bootstrap is inappropriate as the parameter changes with each sample.

3.4 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 [46].

Data generating process and estimators

Consider the following data generating process for unit $i = \{1, \ldots, n\}$. 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

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

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 [39]. 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.

Simulation Results

Table 3.1 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			
	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
$n = 70$	-0.34	2.71	6.69	7.06E-3	0.29	2.71	5.34	4.94E-3
$n = 100$	0.21	2.72	5.49	4.77E-3	0.46	2.72	5.11	3.53E-3

Table 3.1: Summary of the causal parameters over 2,500 simulations of size $n = \{50, 70, 100\}$. All values are in percent. The true value of the PATE was 2.73%.

Table 3.2 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. [124, 125]). As expected, the variance of the estimators decreased and thereby the "true" power increased with increasing sample size and with adjustment (e.g. [119–121]). 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. 3.2), 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.

| | | Bias | Variability | | Power | | CI |
			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 3.2: Summary of estimator performance over 2,500 simulations. The rows denote the sample sizes 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% CI that contained the true parameter value.

3.5 Discussion

To our knowledge, this is the first work 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. In 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 [129]. 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 [145, 146]. These generalizations can be made with the formal methods for transportability [30] 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.

Chapter 4

Adaptive Pre-specification in Randomized Trials With and Without Pair-Matching

The objective of a randomized trial is to evaluate the effect of an intervention on the outcome of interest. In this setting, the difference in the average outcomes among the treated units and the average outcomes among the control units provides a simple and unbiased estimate of the intervention effect. Adjusting for measured covariates during the analysis can substantially reduce the estimator's variance and thereby increase study power (e.g. [119–121, 123, 124]). Nonetheless, the recommendations on adjustment in randomized trials have been conflicting [96, 147–150]. The advice seems to depend on the study design, the unit of randomization, the application, and the sample size. As a result, many researchers are left wondering how to adjust for baseline covariates, if at all.

Consider a trial, where the treatment is randomly allocated to $n/2$ units and the remaining units are assigned to the control. There is a rich literature on locally efficient estimation in this setting (e.g. [122–124, 151, 152].) For example, parametric regression can be used to obtain an unbiased and more precise estimate of the intervention effect. Briefly, the outcome is regressed on the exposure and covariates according to a working model. Following Rosenblum and van der Laan [125], we use "working" to emphasize that the regression function need not be and often is not correctly specified. This working model can include interaction terms and can be linear or non-linear. The estimated coefficients are then used to obtain the predicted outcomes for all units under the treatment and the control. The difference or ratio in the average of the predicted outcomes provides an estimate of the intervention effect.

For continuous outcomes and linear working models without interaction terms, this procedure is known as analysis of covariance (ANCOVA), and the coefficient for the exposure is equal to the estimate of the intervention effect. For binary outcomes, Moore and van der Laan [124] detailed the potential gains in precision from adjustment via logistic regression for estimating the treatment effect on the absolute or relative scale (i.e. risk difference, risk ratio or odds ratio). Furthermore, the authors showed that parametric maximum likelihood

estimation (MLE) was equivalent to targeted maximum likelihood estimation (TMLE) in this setting [3, 40]. As a result, the asymptotic properties of the TMLE, including double robustness and asymptotic linearity, hold even if the working model for outcome regression is misspecified. Furthermore, this approach is locally efficient in that the TMLE will achieve the lowest possible variance among a large class of estimators if the working model is correctly specified. Rosenblum and van der Laan [125] expanded these results for a large class of general linear models. Indeed, the parametric MLE and TMLE can be considered special cases of the double robust estimators of Scharfstein *et al.* [153] and semiparametric approaches of Tsiatis *et al.* [123] and Zhang *et al.* [151]. For a recent and detailed review of these estimation approaches, we refer the reader to Colantuoni and Rosenblum [154].

Now consider a pair-matched trial, where the intervention is randomly allocated within the $n/2$ matched pairs. The proposed estimation strategies have been more limited in this setting. Indeed, the perceived "analytical limitations" of pair-matched trials have led some researchers to shy away from this design [82, 95, 150]. As with a completely randomized trial, the unadjusted difference in treatment-specific means provides an unbiased but inefficient estimate of the intervention effect. To include covariates in the analysis and to potentially increase power, Hayes and Moulton [96] suggested regressing the outcome on the covariates (but not on the exposure) and then contrasting the observed versus predicted outcomes within matched pairs. Alternatively, TMLE can provide an unbiased and locally efficient approach in pair-matched trials [115, 129]. Specifically, the algorithm can be implemented as if the trial were completely randomized: (1) fit a working model for the mean outcome, given the exposure and covariates, (2) obtain predicted outcomes for all units under the treatment and control, and (3) contrast the average of the predicted outcomes on the relevant scale. Inference, however, must respect the pair-matching scheme [115, 129].

A common challenge to the both designs is the selection of the covariates for inclusion in the analysis. Many variables are measured prior to implementation of the intervention, and it is difficult to *a priori* specify an appropriate working model. For a completely randomized trial, covariate adjustment will lead to gains in precision if (i) the covariates are predictive of the outcome and (ii) the covariates are imbalanced between treatment groups (e.g. [155]). Balance is guaranteed as sample size goes to infinity, but rarely seen in practice. Analogously, in a pair-matched trial, covariate adjustment will improve precision if there is an imbalance on predictive covariates after matching.

Limited sample sizes pose an additional challenge to covariate selection. A recent review of randomized clinical trials reported that the median number of participants was 58 with an interquartile range of 27-161 [156]. Likewise, a recent review of cluster randomized trials reported that the median number of units was 31 with an interquartile range of 13-60 [157]. In small trials, adjusting for too many covariates can lead to overfitting and inflated Type I error rates (e.g. [129, 152, 155]). Finally, *ad hoc* selection of the adjustment set leads to concerns that researchers will go on a "fishing expedition" to find the covariates resulting in most power and again risking inflation of the Type I error rate (e.g. [123, 147, 158]).

In sum, covariate adjustment in randomized trials can provide meaningful improvements in precision and thereby statistical power. To preserve inference, the working model, includ-

ing the adjustment variables, must be specified *a priori*. In practice, sample size often limits the size of the adjustment set, and best set is unclear before the trial's conclusion. This results in an important challenge: the need to learn from the data to realize precision gains, but doing so in pre-specified and rigorous way to maintain valid statistical inference.

In this paper, we apply the principle of *empirical efficiency maximization* to data-adaptively select from a pre-specified library the candidate TMLE, which minimizes variance and thereby maximizes the precision of the analysis [122, 159]. We contribute to the existing methodology by modifying this strategy for pair-matched trials. To our knowledge, such a data-adaptive procedure has not been proposed or implemented for this study design. We further contribute to the literature by collaboratively estimating the exposure mechanism for additional gains in precision [160, 161]. We also generalize the results for estimation and inference of both the population and sample average treatment effects [140]. Our finite sample simulations demonstrate the practical performance with limited numbers of independent units, as is common in early phase clinical trials and in cluster randomized trials. 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 (NCT01864603) [13].

4.1 Motivating Example and Causal Parameters

SEARCH is a community randomized trial to estimate the effect of immediate and stream-lined antiretroviral therapy (ART) on HIV incidence as well as other health, economic and educational outcomes. The trial is being conducted in 32 rural communities in Uganda and Kenya. Extensive baseline characteristics were collected through ethnographic mapping and community-wide censuses. Examples include region, occupational mix, measures of mobility, HIV prevalence and community-level HIV RNA viral load. A subset of these characteristics was used to create the 16 best matched pairs of communities [129]. The intervention was randomized within matched pairs. In treatment communities, HIV testing is expanded, and all individuals testing HIV+ are immediately eligible for ART with enhanced services for initiation, linkage and retention in care. In control communities, all individuals testing HIV+ are eligible for ART, according to in-country guidelines. The primary outcome is the five-year cumulative incidence of HIV and will be measured through longitudinal follow-up. The observed data for a given SEARCH community can be denoted

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

where W represents the vector of baseline covariates, A represents the intervention assignment, and Y denotes the outcome. Specifically, A is a binary indicator, equalling one if the community was randomized to the treatment and zero if the community was randomized to the control.

In this paper, we consider estimation and inference for the population average treatment effect (PATE) and the sample average treatment effect (SATE). Let $Y(a)$ denote the out-

come if possibly contrary-to-fact the unit were assigned intervention-level $A = a$. The causal parameters are a function of the distribution P_X of the full data, comprised of the baseline covariates and the counterfactual outcomes of interest: $X = (W, Y(1), Y(0))$ [19, 20]. Specifically, the PATE is the expected difference in the counterfactual outcomes if all members of the population were assigned the intervention and if all members of that population were assigned the control:

$$\Psi^{\mathcal{P}}(P_X) = E_X\big[Y(1) - Y(0)\big] \tag{4.1}$$

where the expectation is over the full data distribution P_X. There is one true value of $\Psi^{\mathcal{P}}(P_X)$ for the target population. For the SEARCH trial, the PATE is the expected difference in the counterfactual cumulative incidence of HIV if all communities in the hypothetical target population implemented the test-and-treat strategy and the counterfactual cumulative incidence of HIV if all communities in that target population maintained the standard of care.

The sample parameter is the average difference in the counterfactual outcomes for the study units [19]:

$$\Psi^{\mathcal{S}}(P_X) = \frac{1}{n} \sum_{i=1}^{n} \big[Y_i(1) - Y_i(0)\big] \tag{4.2}$$

were $Y_i(a)$ denotes the outcome if possibly contrary-to-fact unit i were assigned intervention-level $A = a$. The SATE is data-adaptive; its true value depends on the n units in the sample. The SATE is easily interpretable, responsive to heterogeneity in the intervention effect, and arguably the most relevant when the study units are not representative of a greater population. For the SEARCH trial, the SATE 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 32 study communities.

4.2 Targeted Estimation in a Randomized Trial Without Matching

In this section, we ignore the pair-matching scheme in the SEARCH trial and assume the observed data consist of n independent, identically distribution (i.i.d.) copies of $O = (W, A, Y)$ with some true, but unknown distribution P_0, which factorizes as

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

We do not make any assumptions about the common covariate distribution $P_0(W)$ or about the common conditional distribution of the outcome, given the intervention and covariates $P_0(Y|A, W)$. By design, the intervention A is randomized with probability 0.5. Therefore,

the exposure mechanism is known: $P_0(A = 1|W) = g_0(1|W) = 0.5$. The statistical model $\mathcal{M}$, describing the set of possible observed data distributions, is semiparametric.

Since the intervention is randomized, we can easily identify the population effect $\Psi^{\mathcal{P}}(P_X)$ (Eq. 4.1) from the observed data distribution. Our statistical estimand is the difference in the expected outcome, given the treatment and covariates, and the expected outcome, given the control and covariates, averaged (standardized) with respect to the covariate distribution in the population [27]:

$$\Psi(P_0) = E_0\big[E_0(Y|A = 1, W) - E_0(Y|A = 0, W)\big]$$
$$= E_0\big[\bar{Q}_0(1, W) - \bar{Q}_0(0, W)\big]$$

where $\bar{Q}_0(A, W) = E_0(Y|A, W)$ denotes the conditional mean outcome, given the exposure and covariates. As discussed in the introduction, there are many algorithms available for unbiased and locally efficient estimation of this statistical parameter in a randomized trial (e.g. [122–124, 151, 152]). Throughout, our focus is on TMLE, a general methodology for the construction of double robust, semiparametric efficient substitution estimators [3, 40].

A TMLE for population effect also serves as a consistent and asymptotically linear estimator of the sample effect $\Psi^{\mathcal{S}}(P_X)$ (Eq. 4.2) [140]. The estimator can be implemented in three steps.

Step 1. Initial estimation: Estimate the expected outcome, given the exposure and covariates $\bar{Q}_0(A, W) = E_0(Y|A, W)$. We could use a pre-specified parametric working model (as discussed above) or a more data-adaptive approach (as discussed below).

Step 2. Targeting: Update the initial estimator $\bar{Q}_n(A, W)$.

    i. Calculate the "clever" covariate based on the known or estimated exposure mechanism $g_n(A|W) = P_n(A|W)$:

$$H_n(A, W) = \left(\frac{\mathbb{I}(A = 1)}{g_n(1|W)} - \frac{\mathbb{I}(A = 0)}{g_n(0|W)}\right)$$

    ii. If the outcome is continuous and unbounded, run linear regression of the outcome Y on the covariate $H_n(A, W)$ with the initial estimator as offset. Plug in the estimated coefficient ϵ_n to yield the targeted update: $\bar{Q}_n^*(A, W) = \bar{Q}_n(A, W) + \epsilon_n H_n(A, W)$.

    iii. If the outcome is binary or bounded in $[0, 1]$[1], run logistic regression of the outcome Y on the covariate $H_n(A, W)$ with the $\text{logit}(x) = log\{x/(1 - x)\}$ of the initial estimator $\bar{Q}_n(A, W)$ as offset. Plug in the estimated coefficient ϵ_n to yield the targeted update: $\bar{Q}_n^*(A, W) = \text{expit}\big\{\text{logit}[\bar{Q}_n(A, W)] + \epsilon_n H_n(A, W)\big\}$, where expit is the inverse-logit.

---

[1]In greater generality, the logistic fluctuation can also be used for a continuous outcome that is bounded in $[a, b]$ by first applying the following transformation to the outcome: $Y^* = (Y - a)/(b - a)$. For further details, see Gruber and van der Laan [137].

Step 3. Parameter estimation: Obtain the predicted outcomes for all observations under the treatment $\bar{Q}_n^*(1, W)$ and control $\bar{Q}_n^*(0, W)$. Average the difference in predicted outcomes:

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

If the initial estimator for $\bar{Q}_0(A, W)$ is based on a working regression model with an intercept and a main term for the exposure and if the exposure mechanism is treated as known, then the updating step can be skipped [125]. Further precision, however, can be attained by using a data-adaptive algorithm for initial estimation of $\bar{Q}_0(A, W)$ and by estimating the exposure mechanism $g_0(A|W)$ [55].

Under regularity conditions, the standardized estimator converges to a normal distribution with mean 0 and variance given by the variance of its influence curve, divided by sample size n. The influence curve for the TMLE of the population parameter (PATE) can be estimated from the observed data distribution by

$$D_n^{\mathcal{P}}(g_n, \bar{Q}_n^*)(O) = \left(\frac{\mathbb{I}(A = 1)}{g_n(1|W)} - \frac{\mathbb{I}(A = 0)}{g_n(0|W)} \right) \left(Y - \bar{Q}_n^*(A, W) \right)$$
$$+ \bar{Q}_n^*(1, W) - \bar{Q}_n^*(0, W) - \Psi_n(\bar{Q}_n^*) \tag{4.3}$$

[3]. The influence curve for the TMLE of the sample parameter (SATE) can be *conservatively* estimated from the observed data distribution by

$$D_n^{\mathcal{S}}(g_n, \bar{Q}_n^*)(O) = \left(\frac{\mathbb{I}(A = 1)}{g_n(1|W)} - \frac{\mathbb{I}(A = 0)}{g_n(0|W)} \right) \left(Y - \bar{Q}_n^*(A, W) \right) \tag{4.4}$$

[140]. For the SATE, there is no variance contribution from the covariate distribution, which is considered fixed. Asymptotically, the SATE will often be estimated with more precision than the PATE [19, 117]. However, as shown by Balzer *et al.* [140] and in the simulations that follow, the gains in precision from specifying the SATE as the target of inference can be attenuated in small trials, because this influence curve-based variance estimator is conservative.

Adaptive Pre-specified Approach for Step 1. Initial Estimation

Consider again the SEARCH trial for HIV prevention and treatment. Recall that the outcome Y is the five-year cumulative incidence of HIV and bounded between 0 and 1. Suppose that we want to use logistic regression for initial estimation of the expected outcome, given the exposure and measured covariates $\bar{Q}_0(A, W)$. It is unclear *a priori* which covariates should be included in the working model and in what form. For example, baseline HIV prevalence is a known predictor of the outcome and may be imbalanced between the treatment and control groups. Likewise, there might be substantial heterogeneity in the treatment

effect by region and allowing for an interaction between region and the intervention may reduce the variance of the TMLE. Including all the covariates and the relevant interactions in the working model is likely to result in overfitting and misleading inference. To facilitate selection between candidate initial estimators and thereby candidate TMLEs, we propose the following cross-validation selector.

First, we propose a library of candidate working models for initial estimation of the conditional mean outcome $\bar{Q}_0(A, W)$. This library should be pre-specified in the protocol or the analysis plan. A possible library could consist of the following logistic regression working models:

$$\text{logit}[\bar{Q}^{(a)}(A, W, \beta)] = \beta_0 + \beta_1 A$$
$$\text{logit}[\bar{Q}^{(b)}(A, W, \beta)] = \beta_0 + \beta_1 A + \beta_2 W1$$
$$\text{logit}[\bar{Q}^{(c)}(A, W, \beta)] = \beta_0 + \beta_1 A + \beta_2 W2 + \beta_3 A \times W2$$

where, for example, $W1$ denotes baseline prevalence and $W2$ denotes region. Of course, there are many more candidate algorithms, and we are considering this simple set for pedagogic purposes. We also note that the first working model corresponds to the unadjusted estimator.

Second, we need to pre-specify a loss function to measure the performance of candidate estimators. Following the principle of empirical efficiency maximization [159], we propose using the variance of the estimated influence curve of the TMLE for the parameter of interest. Specifically, if the target of inference is the population effect, our loss function is

$$\mathcal{L}^{\mathcal{P}}(\bar{Q}) = Var\big[D_n^{\mathcal{P}}(g_0, \bar{Q})\big],$$

and if the target of inference is the sample effect, our loss function is

$$\mathcal{L}^{\mathcal{S}}(\bar{Q}) = Var\big[D_n^{\mathcal{S}}(g_0, \bar{Q})\big].$$

For SATE, this corresponds to the L2 squared error loss function $\mathcal{L}^{\mathcal{S}}(\bar{Q}) = \big(Y - \bar{Q}(A, W)\big)^2$.

Next, we need to pre-specify our cross-validation scheme, used to generate an estimate of the expected loss (i.e. the "risk") for each of the candidate estimators. For generality, we present V-fold cross-validation, where the data are randomly split into V partitions, called "folds", of size $\approx n/V$. To respect the limited sample sizes common in early phase clinical trials and in cluster randomized trials, leave-one-out cross-validation may be appropriate. Leave-one-out cross-validation corresponds with $V = n$-fold cross-validation, where each fold corresponds to one observation. The cross-validation procedure for initial estimation of the conditional mean $\bar{Q}_0(A, W)$ can be implemented as follows.

i. For each fold $v = \{1, \ldots, V\}$ in turn,

    a. Set the observation(s) in fold v to be the validation set and the remaining observations to be the training set.

b. Fit each algorithm using only data in the training set. For the above library, we would run logistic regression of the outcome Y on the exposure A and covariates W, according to the working model. Denote the initial regression fits as $\bar{Q}_n^{(a)}(A, W)$, $\bar{Q}_n^{(b)}(A, W)$ and $\bar{Q}_n^{(c)}(A, W)$, respectively.

c. For each algorithm, use the estimated fit to predict the outcome(s) for the observation(s) in the validation set under the treatment and the control. For the first algorithm, for example, we would have $\bar{Q}_n^{(a)}(1, W_k)$ and $\bar{Q}_n^{(a)}(0, W_k)$ for observation O_k in the validation set.

ii. For each algorithm, estimate the risk with the sample variance of the cross-validated estimate of the influence curve. If the target of inference is the PATE, our risk estimates would be

$$\text{CV-risk}^{\mathcal{P},(a)} = Var_n\big[D_n^{\mathcal{P}}\big(g_0, \bar{Q}_n^{(a)}\big)\big]$$
$$\text{CV-risk}^{\mathcal{P},(b)} = Var_n\big[D_n^{\mathcal{P}}\big(g_0, \bar{Q}_n^{(b)}\big)\big]$$
$$\text{CV-risk}^{\mathcal{P},(c)} = Var_n\big[D_n^{\mathcal{P}}\big(g_0, \bar{Q}_n^{(c)}\big)\big]$$

where Var_n denotes the sample variance and we are treating the exposure mechanism as known: $g_0(A|W) = 0.5$. If instead the target of inference is SATE, our risk estimates would be

$$\text{CV-risk}^{\mathcal{S},(a)} = Var_n\big[D_n^{\mathcal{S}}(g_0, \bar{Q}_n^{(a)})\big]$$
$$\text{CV-risk}^{\mathcal{S},(b)} = Var_n\big[D_n^{\mathcal{S}}\big(g_0, \bar{Q}_n^{(b)}\big)\big]$$
$$\text{CV-risk}^{\mathcal{S},(c)} = Var_n\big[D_n^{\mathcal{S}}\big(g_0, \bar{Q}_n^{(c)}\big)\big]$$

iii. Select the algorithm with the smallest cross-validated risk.

The selected working model is then used for initial estimation of $\bar{Q}_0(A, W)$ in Step 1 of the TMLE algorithm. Since our library was limited to parametric working models and the exposure mechanism was treated as known, the updating step can be skipped. In other words, the chosen estimator was already targeted $\bar{Q}_n(A, W) = \bar{Q}_n^*(A, W)$ and can be used for Step 3 parameter estimation.

4.3 Targeted Estimation in a Randomized Trial With Matching

Recall the pair-matching scheme briefly described in Section 4.1 for the SEARCH trial. First, the potential study units were selected. Then baseline covariates, such as region, occupational mix and measures of migration, were collected. A matching algorithm was applied to the baseline covariates of candidate units to create the best 16 matched pairs. The

intervention was randomized within the resulting pairs, and the outcome will be measured with longitudinal follow-up. This pair-matching scheme is considered to be *adaptive*, because the resulting matched pairs are a function of the baseline covariates of all the candidate units [115, 129]. This design has also been called "nonbipartite matching" and "optimal multivariate matching" [93, 112, 113].

The adaptive design creates a dependence in the data. Since the construction of the matched pairs is a function of the baseline covariates of all n study units, the observed data do not consist of $n/2$ i.i.d. paired observations, as current practice sometimes assumes (e.g. [82, 96, 97, 100]). Instead, we have n dependent copies of $O = (W, A, Y)$. Nonetheless, there is a lot of conditional independence in the data. Mainly, once we consider the baseline covariates of the study units as fixed, then we recover $n/2$ conditionally independent units:

$$\bar{O}_j = \big(O_{j1}, O_{j2}\big) = \big((W_{j1}, A_{j1}, Y_{j1}), (W_{j2}, A_{j2}, Y_{j2})\big)$$

where the index $j = 1, \ldots, n/2$ denotes the partitioning of the candidate units $\{1, \ldots n\}$ into matched pairs according to similarity in their baseline covariates $(W_1, \ldots, W_n)$. Throughout subscripts $j1$ and $j2$ index the observations within matched pair j. The conditional distribution of the observed data, given the baseline covariates of the study units, factorizes as

$$P_0(O_1, \ldots, O_n | W_1, \ldots W_n) = \prod_{j=1}^{n/2} P_0(A_{j1}, A_{j2} | W_1, \ldots, W_n) P_0(Y_{j1} | A_{j1}, W_{j1}) P_0(Y_{j2} | A_{j2}, W_{j2})$$

$$= 0.5 \prod_{j=1}^{n/2} P_0(Y_{j1} | A_{j1}, W_{j1}) P_0(Y_{j2} | A_{j2}, W_{j2})$$

where the second line follows from randomization of the intervention within matched pairs. For estimation and inference of the PATE, we need to assume that each community's baseline covariates W_i are independently drawn from some common distribution $P_0(W)$. For estimation and inference of the SATE, this assumption on the covariate distribution can be weakened. (See the Appendix D.1 for further details.)

Despite the dependence in the data, a TMLE for the population effect (PATE) can be implemented as if the sample were n i.i.d. units [115]. In Step 1, we obtain an initial estimator of $\bar{Q}_0(A, W)$ with an *a priori*-specified parametric working model or with a more data-adaptive method. In Step 2, we target the initial estimator $\bar{Q}_n(A, W)$ by using information in the known or estimated exposure mechanism. Finally in Step 3, we obtain the predicted outcomes for all observations under the treatment $\bar{Q}_n^*(1, W)$ and the control $\bar{Q}_n^*(0, W)$, and then take the sample average of the difference in these predicted outcomes. Furthermore, the variance of the TMLE can be estimated by treating the sample as n i.i.d. units. In other words, inference can be based on the sample variance of the estimated influence curve in the non-matched trial $D_n^{\mathcal{P}}$ (Eq. 4.3), divided by n [115]. This variance estimator ignores any gains in precision from pair-matching and will be conservative under reasonable

assumptions. A less conservative variance estimator can be obtained by accounting for the potential correlations of the residuals within matched pairs:

$$\rho_n(\bar{Q}_n^*)(\bar{O}_j) = \frac{1}{n/2} \sum_{j=1}^{n/2} \left(Y_{j1} - \bar{Q}_n^*(A_{j1}, W_{j1})\right)\left(Y_{j2} - \bar{Q}_n^*(A_{j2}, W_{j2})\right) \tag{4.5}$$

[115]. An estimate of the asymptotic variance of the TMLE is then given by the sample variance of $D_n^{\mathcal{P}}$ minus $2\rho_n$, all divided by n.

In a pair-matched trial, a TMLE for the population effect is also a consistent and asymptotically linear estimator of the sample effect. The proof is given in Appendix D.1. Furthermore, the influence curve for the TMLE of the SATE can be *conservatively* estimated by

$$\bar{D}_n^{\mathcal{S}}(g_n, \bar{Q}_n^*)(\bar{O}_j) = \frac{1}{2}\left[D_n^{\mathcal{S}}(g_n, \bar{Q}_n^*)(O_{j1}) + D_n^{\mathcal{S}}(g_n, \bar{Q}_n^*)(O_{j2}) \right] \tag{4.6}$$

where $D_n^{\mathcal{S}}(g_n, \bar{Q}_n^*)(O)$ is the estimated influence curve for observation O in the non-matched trial (Eq. 4.4). The proof is given in Appendix D.2. Inference can be based on the sample variance of the estimated (paired) influence curve $\bar{D}_n^{\mathcal{S}}$, divided by $n/2$. If we order observations within matched pairs such that first corresponds to the intervention ($A_{j1} = 1$) and the second to the control ($A_{j2} = 0$) and treat the exposure mechanism as known $g_0(A|W) = 0.5$, we have

$$\bar{D}_n^{\mathcal{S}}(g_0, \bar{Q}_n^*)(\bar{O}_j) = \left(Y_{j1} - \bar{Q}_n^*(1, W_{j1})\right) - \left(Y_{j2} - \bar{Q}_n^*(0, W_{j2})\right)$$

In this setting, the sample variance of the pairwise differences in residuals, divided by $n/2$, provides a conservative variance estimator.

Adaptive Pre-specified Approach for Step 1. Initial Estimation

By balancing intervention groups with respect to baseline determinants of the outcome, pair-matching increases the efficiency of the study (e.g. [94, 115, 129]). Nonetheless, residual imbalance on the baseline predictors often remains, and adjusting for these covariates during the analysis can further increase efficiency. In the SEARCH trial, for example, the matched pairs were created before baseline HIV prevalence was measured. As a result, there is likely to be variation across pairs in baseline prevalence, which is a known driver of HIV incidence. Adjusting for baseline prevalence during the analysis is likely to reduce the variance of the TMLE and result in a less conservative variance estimator. Unfortunately, it is unclear *a priori* whether adjusting for prevalence will yield more power than than adjusting for other baseline covariates, such as male circumcision prevalence or measures of community-level HIV RNA viral load. With only 16 (conditionally) independent units, we are limited as to the size of the adjustment set. Adjusting for too many covariates can result in over-fitting. As before, we want to data-adaptively select the candidate TMLE (i.e. working regression model), which maximizes the empirical efficiency.

The data-adaptive procedure for Step 1 initial estimation of conditional mean outcome $\bar{Q}_0(A, W)$, outlined in Sec. 4.2 for a non-matched trial, can be modified for a pair-matched trial. As before, we need to pre-specify our library of candidate estimators, our measure of performance and the cross-validation scheme. We can use the same library of candidate working models for initial estimation of the conditional mean outcome $\bar{Q}_0(A, W)$. For the loss function, however, we want to use the estimated variance of the TMLE under pair-matching. To elaborate, consider the loss function for the SATE in a non-matched trial. Minimizing the sum of squared residuals (i.e. minimizing the variance of $D_n^{\mathcal{S}}$ (Eq. 4.4)) targets the conditional mean outcome $\bar{Q}_0(A, W)$. As a result, the algorithm could select a working model adjusting for a covariate that is highly predictive of the outcome but on which we matched perfectly. In the SEARCH trial, for example, communities were paired within region, because HIV incidence is expected to be highly heterogeneous across regions. Therefore, minimizing the empirical variance of $D_n^{\mathcal{S}}$ might lead to the selection of the candidate TMLE with main terms for the intervention and region. This selection would not improve the precision of the analysis over the unadjusted algorithm. Instead, we want to select the TMLE minimizing the conservative estimator of the variance in the pair-matched design. Thereby, our loss functions for the PATE and SATE are

$$\mathcal{L}^{\mathcal{P}}(\bar{Q}) = Var\left[D_n^{\mathcal{P}}(g_0, \bar{Q}) - 2\rho_n(\bar{Q})\right]$$
$$\mathcal{L}^{\mathcal{S}}(\bar{Q}) = Var\left[\bar{D}_n^{\mathcal{S}}(g_0, \bar{Q})\right]$$

respectively. Finally, the pair should be treated as the unit of (conditional) independence in the cross-validation scheme. In other words, when the data are split into V-folds, the pairing should be preserved. In small trials, leave-one-pair-out cross-validation may be appropriate. With these modifications, we can implement the cross-validation scheme, outlined in Sec. 4.2, to data-adaptively select the candidate working model, which minimizes the estimated variance of the TMLE in a pair-matched trial.

4.4 Collaborative Estimation of the Exposure Mechanism

Even though the intervention A is randomized with balanced allocation, estimating the known exposure mechanism $g_0(A|W)$ can increase the precision of the analysis [55]. As before, we want to respect the study design (i.e. pair-matched or not) as well as adjust for a covariate only if its inclusion improves the empirical efficiency. For example, we may not want to include a covariate that is imbalanced between the intervention groups (i.e. predictive of A) but not predictive of the outcome. Likewise, if a given covariate (e.g. $W1$) was included in the working model for $\bar{Q}_0(A, W)$, further adjusting for this covariate when estimating the exposure mechanism may not increase precision. To this end, we propose to incorporate the Collaborative-TMLE (C-TMLE) approach into our algorithm [160, 161].

Adaptive Pre-specified Approach for Step 2. Targeting

First, we propose a library of candidate estimators of the exposure mechanism $g_0(A|W)$.
As before, this library should be pre-specified in the protocol or analysis plan. A possible
library could consist of the following logistic regression working models:

$$\text{logit}[g^{(a)}(W, \beta)] = \beta_0$$
$$\text{logit}[g^{(b)}(W, \beta)] = \beta_0 + \beta_1 W1$$
$$\text{logit}[g^{(c)}(W, \beta)] = \beta_0 + \beta_1 W2$$

where, for example, $W1$ is baseline prevalence and $W2$ is region. Each algorithm would yield
a different update to a given initial estimator of the conditional mean outcome $\bar{Q}_n(A, W)$,
selected by the data-adaptive procedure for Step 1 (Sec. 4.2 and 4.3). In other words, each
candidate estimator of $g_0(A|W)$ results in a different targeted estimator $\bar{Q}_n^*(A, W)$. We also
note that the first working model corresponds to the unadjusted estimator.

To choose between candidate algorithms, we need to pre-specify a loss function. As before,
we propose using the estimated variance of the TMLE, appropriate for the scientific question
(i.e. population or sample effect) and study design (i.e. pair-matched or not). Finally, we
need to pre-specify our cross-validation scheme, used to obtain an honest measure of risk and
to reduce the potential for over-fitting. As before, we present V-fold cross-validation, where
the data are partitioned into V folds of size $\approx n/V$. If matching was used, the partitioning
should preserve the pairs. The cross-validation selector for collaborative estimation of the
exposure mechanism can be implemented as follows.

i. For each fold $v = \{1, \ldots, V\}$ in turn,

   a. Set the observation(s) in fold v to be the validation set and the remaining observations
   to be the training set.

   b. Using only data in the training set, fit each algorithm for estimating the exposure
   mechanism. For the above library, we would run logistic regression of the exposure A
   on the covariates W, according to the working model. Denote the estimated exposure
   mechanisms as $g_n^{(a)}(A|W)$, $g_n^{(b)}(A|W)$ and $g_n^{(c)}(A|W)$, respectively.

   c. For each algorithm, use the estimated fit of the exposure mechanism to target the
   initial estimator $\bar{Q}_n(A, W)$, also fit with the training set.

   d. For each algorithm, obtain targeted predictions of the outcome(s) for the observa-
   tion(s) in the validation set under the treatment and the control. For the first algo-
   rithm, for example, we would have $\bar{Q}_n^{(a),*}(1, W_k)$ and $\bar{Q}_n^{(a),*}(0, W_k)$ for observation O_k
   in the validation set.

ii. For each algorithm, estimate the risk with the cross-validated variance estimator, ap-
propriate for the target parameter and study design.

iii. Select the algorithm with the smallest cross-validated risk.

The chosen estimator is then used for targeting in Step 2 of the TMLE algorithm.

4.5 Obtaining Inference

In summary, we have proposed the following data-adaptive TMLE to maximize the precision and power of a randomized trial.

> Step 1. Initial estimation of the conditional mean outcome with the working model $\bar{Q}_n(A, W)$, which was data-adaptively selected to maximize the empirical efficiency of the analysis (Sec. 4.2 for a non-matched trial and Sec. 4.3 for a matched trial).

> Step 2. Targeting the initial estimator using the estimated exposure mechanism $g_n(A|W)$, which was data-adaptively selected to further maximize the empirical efficiency of the analysis (Sec. 4.4).

> Step 3. Obtaining a point estimate by averaging the difference in the targeted predicted outcomes:
> $$\Psi_n(\bar{Q}_n^*) = \frac{1}{n} \sum_{i=1}^n \left[\bar{Q}_n^*(1, W_i) - \bar{Q}_n^*(0, W_i) \right]$$

We now need a variance estimator that accounts for this selection process. For this, we propose using a cross-validated variance estimator. As before, the data are split into validation and training sets, respecting the unit of (conditional) independence. The selected TMLE is fit using the data in the training set and used to estimate the influence curve[2] for the observation(s) in the validation set. The step-by-step instructions are given in the Appendix D.3. The sample variance of the cross-validated estimate of the influence curve can then be used for hypothesis testing and the construction of Wald-type confidence intervals. For trials with a limited number of independent units, the Student's t-distribution is an appropriate alternative to the standard normal distribution.

4.6 Simulation Study

We present the following simulation studies to demonstrate (1) implementation of the proposed methodology, (2) the potential gains in precision and power from data-adaptive estimation of the conditional mean outcome, (3) the additional gains in precision and power from collaborative estimation of the exposure mechanism, and (4) maintenance of nominal confidence interval coverage. All simulations were conducted in R v3.1.2 [46].

---

[2] For the TMLE of the population effect in a matched trial, we also need a cross-validated estimate of the correction term ρ_n (Eq. 4.5.) This term is a function of the residuals, which can be estimated for each pair in the validation set based on targeted estimator $\bar{Q}_n^*(A, W)$, fit with the training set.

Study 1

For each unit $i = \{1, \ldots, n\}$, we generated the nine baseline covariates by drawing from a
multivariate normal with mean 0 and variance 1. The correlation between the first three
covariates $\{W1, W2, W3\}$ and between the second three covariates $\{W4, W5, W6\}$ was 0.5,
while the correlation between the remaining covariates $\{W7, W8, W9\}$ was 0. The expo-
sure A was randomized such that the treatment allocation was balanced overall. For the
non-matched trial, we randomly assigned the intervention to $n/2$ units and the control to
the remaining $n/2$ units. For the pair-matched trial, we used the non-bipartite matching
algorithm `nbpMatch` to pair units on covariates $\{W1, \ldots, W6\}$ [111]. The exposure A was
randomized within the resulting matched pairs. Recall A is a binary indicator, equalling 1
if the unit was assigned the treatment and 0 if the unit was assigned the control. For each
unit, the outcome Y was then generated as

$$Y = 0.4A + 0.25(W1 + W2 + W4 + W5 + U_Y) + 0.25A(W1 + U_Y)$$

where U_Y was drawn from a standard normal. We also generated the counterfactual outcomes
$Y(1)$ and $Y(0)$ by intervening to set $A = a$. To reflect the limited sample sizes common
in early phase clinical trials and in cluster randomized trials, we selected a sample size of
$n = 40$. This resulted in $n/2 = 20$ conditionally independent units in the pair-matched trial.

For each study design (non-matched or matched), this data generating process was re-
peated 2,500 times. Recall that the sample effect $\Psi^{\mathcal{S}}(P_X)$ (Eq. 4.2) is data-adaptive param-
eter; its value changes with each new selection of units. Thereby, for each repetition, the
SATE was calculated as the sample average of the difference in the counterfactual outcomes.
The SATE ranged from 0.23 to 0.60 with a mean of 0.40. In contrast, the population effect
$\Psi^{\mathcal{P}}(P_X)$ (Eq. 4.1) is constant and was calculated by averaging the difference in the coun-
terfactual outcomes over a population of 900,000 units. The true value of the PATE was
0.40.

We compared the performance of the unadjusted estimator to TMLE with various ap-
proaches to covariate adjustment. Specifically, we implemented the TMLE algorithm, where
the initial estimation of the conditional mean outcome $\bar{Q}_0(A, W)$ was based on a linear work-
ing model with main terms for the intervention A and irrelevant covariate $W9$ and where the
exposure mechanism was treated as known: $g_0(A|W) = 0.5$. This approach was equivalent to
standard maximum likelihood estimation (MLE) and represented the unfortunate scenario
where the researcher pre-specified adjustment for a covariate that was not predictive of the
outcome.

We also implemented a TMLE with the data-adaptive approach for Step 1 initial es-
timation of the conditional mean outcome (Sec. 4.2 and 4.3). Our library consisted of 10
working linear regression models, each with an intercept, a main term for the exposure A
and a main term for one baseline covariate: $\{\emptyset, W1, \ldots, W9\}$, where $\emptyset$ corresponds to the
unadjusted estimator. Our loss function was the estimated variance of the TMLE, appro-
priate for the target parameter and study design. We chose the candidate working model
with the lowest estimated risk, based on leave-one-out cross-validation for the non-matched

trial and leave-one-pair-out cross-validation for the matched trial. We also implemented Collaborative-TMLE (C-TMLE), which couples the data-adaptive approach for Step 1 initial estimation of the conditional mean outcome (Sec. 4.2 and 4.3) with the data-adaptive approach for Step 2 targeting (Sec. 4.4). For the latter, our library of candidates to estimate the exposure mechanism consisted of 10 working logistic regression models, each with an intercept and a main term for one baseline covariate: $\{\emptyset, W1, \ldots, W9\}$. The same loss function and cross-validation scheme were used for C-TMLE.

For the unadjusted estimator and the MLE, inference was based on the estimated influence curve. For the data-adaptive TMLEs, inference was based on the cross-validated estimate of the influence curve (Sec. 4.5). We assumed the standardized estimator followed the Student's t-distribution with $n - 2 = 38$ degrees of freedom for the non-matched trial and with $n/2 - 1 = 19$ degrees of freedom for the matched trial.

Results

Table 4.1 illustrates the performance of the estimators over the 2,500 simulated data sets. Specifically, we show the mean squared error (MSE), the relative MSE (rMSE), the average standard error estimate $\hat{\sigma}$, the attained power and the 95% confidence interval coverage. As expected, matching improved efficiency. The MSE of the unadjusted estimator, for example, was approximately 2 times larger in the non-matched trial than in the pair-matched trial. Furthermore, for the pair-matched trial, targeting the sample effect, as opposed to the population effect, resulted in substantial gains in attained power: 38% with the unadjusted estimator for the PATE and 53% with the same estimator for the SATE. For the non-matched trial, targeting the sample parameter increased efficiency, but did not directly translate into increased power due to the conservative variance estimator for the SATE.

In all scenarios, the TMLE with data-adaptive selection of the initial estimator of $\bar{Q}_0(A, W)$ improved precision over the unadjusted estimator and the MLE. Collaborative estimation of the exposure mechanism $g_0(A|W)$ led to further gains in precision. Consider, for example, estimation of the PATE in a trial without matching. The MSE of the unadjusted estimator was 1.44 times larger than the TMLE and 1.51 times larger than the C-TMLE. The attained power was 36%, 51% and 52%, respectively. Furthermore, the precision of the MLE, adjusting for the irrelevant covariate $W9$, was worse than the estimators in all scenarios. This demonstrates the potential peril of relying on one pre-specified adjustment variable. As a second example, consider the attained power to detect that the SATE was different from zero in the pair-matched trial. We would have 53% power with the unadjusted estimator and 54% power with the MLE, adjusting for the irrelevant covariate $W9$. By incorporating the cross-validation selector for initial estimation of $\bar{Q}_0(A, W)$, the TMLE achieved 68% power. By further incorporating collaborative estimation of the exposure mechanism $g_0(A|W)$, the C-TMLE achieved 70% power. Overall, the greatest efficiency was achieved with C-TMLE for the SATE in the pair-matched trial. Indeed, the MSE of the unadjusted estimator for the population parameter in the trial without matching was nearly 3 times larger than the MSE

	PATE					**SATE**				
	MSE[a]	rMSE[b]	$\hat{\sigma}^c$	Pow[d]	Cov[e]	MSE[a]	rMSE[b]	$\hat{\sigma}^c$	Pow[d]	Cov[e]
					Non-Matched Trial					
Unadj.	6.3E-2	1.00	0.25	0.36	0.95	6.0E-2	1.05	0.25	0.36	0.95
MLE	6.6E-2	0.96	0.25	0.37	0.94	6.3E-2	1.00	0.25	0.37	0.95
TMLE	4.4E-2	1.44	0.20	0.51	0.94	4.2E-2	1.52	0.20	0.49	0.95
C-TMLE	4.2E-2	1.51	0.20	0.52	0.95	3.9E-2	1.63	0.20	0.50	0.96
					Matched Trial					
Unadj.	3.3E-2	1.93	0.22	0.38	0.99	3.0E-2	2.11	0.18	0.53	0.97
MLE	3.4E-2	1.84	0.22	0.38	0.98	3.2E-2	1.99	0.18	0.54	0.96
TMLE	2.5E-2	2.53	0.18	0.56	0.98	2.4E-2	2.63	0.16	0.68	0.95
C-TMLE	2.4E-2	2.66	0.18	0.58	0.98	2.3E-2	2.78	0.15	0.70	0.95

[a]Mean squared error: the bias (average deviation between the point estimate and sample-specific true value) - squared plus the variance
[b]Relative MSE: the MSE of the unadjusted estimator for the PATE in a non-matched trial relative to (divided by) the MSE of another estimator
[c]Average standard error estimate, based on the estimated influence curve
[d]Attained power: proportion of times the false null hypothesis was rejected
[e]Confidence interval (CI) coverage: proportion of times the true value was contained in the 95% CI

Table 4.1: Summary of estimator performance for Simulation 1. The rows denote the study design and the estimator: unadjusted, MLE adjusting for $W9$, TMLE with data-adaptive selection of the initial estimator, and Collaborative-TMLE (C-TMLE) with data-adaptive selection of the initial estimator paired with data-adaptive estimation of the exposure mechanism.

of the C-TMLE for the sample effect in the pair-matched trial. Throughout the confidence interval coverage was maintained near or above the nominal rate of 95%.

Further insight into the efficiency gains with the proposed TMLE and C-TMLE is provided by Table 4.2, which shows the proportion of times a working model was selected for initial estimation of the conditional mean outcome and for collaborative estimation of the exposure mechanism. When targeting the PATE, the selection for $\bar{Q}_0(A, W)$ was similar with and without pair-matching. This was not surprising, because our measure of performance (i.e. the loss function) was the estimated variance of the TMLE, and the variance estimator in a pair-matched trial is given by the estimated variance in the non-matched trial minus a correction term ρ_n, which was close to 0. When targeting the SATE, however, the selection procedure was more optimized to the study design. For example, the working model with main terms for the intervention and $W1$ was selected in 57% of the studies without matching and in only 38% of the studies with matching. Instead, working models adjusting for other predictive covariates were selected more frequently. Furthermore, the collaborative procedure for estimation of the exposure mechanism was able to identify settings where the no adjustment would yield the greatest gains in efficiency. Specifically, the unadjusted estimator $g_n(A|W) = 0.5$ was selected in nearly 80% of the studies without matching and in less than

	Adjustment variable	$\emptyset$	$W1$	$W2$	$W3$	$W4$	$W5$	$W6$	$W7$	$W8$	$W9$
		\multicolumn{11}{c}{Working model for initial estimation of $\bar{Q}_0(A,W)$}									
PATE	non-matched	0	57	19	1	11	11	1	0	0	0
	matched	0	54	20	1	12	12	1	0	0	0
SATE	non-matched	0	57	19	1	11	11	1	0	0	0
	matched	0	38	21	5	15	15	3	0	1	1
		\multicolumn{11}{c}{Working model for estimation of $g_0(A|W)$}									
	Adjustment variable	$\emptyset$	$W1$	$W2$	$W3$	$W4$	$W5$	$W6$	$W7$	$W8$	$W9$
PATE	non-matched	79	2	3	2	3	3	2	2	2	2
	matched	29	9	10	9	9	9	8	5	6	6
SATE	non-matched	77	2	3	2	3	4	2	2	3	2
	matched	16	10	11	10	11	10	10	7	7	7

Table 4.2: For Simulation 1, the proportion of times a covariate was selected in the working linear regression model for initial estimation of $\bar{Q}_0(A,W)$ and in the working logistic regression model for collaborative estimation of the exposure mechanism $g_0(A|W)$.

30% of the studies with matching.

Study 2

For the second simulation study, we increased the complexity of the data-generating process and reduced the sample size to $n = 30$. As before, we generated nine baseline covariates from a multivariate normal with mean 0, variance 1 and the same correlation structure. We also generated a binary variable R, equalling 1 with probability 0.5 and equalling -1 with probability 0.5. The final covariate Z was generated as a function of these baseline covariates and random noise U_Z:

$$Z = R \times \text{expit}(W1 + W4 + W7 + 0.5U_Z)$$

where the expit is the inverse of the logit function and U_Z was drawn independently from a standard normal. As before, the intervention A was randomized with balanced allocation. For a non-matched trial, the treatment was randomly assigned to $n/2$ units and the control to the remaining $n/2$ units. For the pair-matched trial, we used the non-bipartite matching algorithm `nbpMatch` to explore two matching sets [111]. In the first, units were matched on R, a baseline covariate strongly impacting Z. In the second, units were matched on $\{R, W2, W5, W8\}$. The intervention A was randomized within the matched pairs. For each unit, the outcome Y was then generated as

$$Y = \text{expit}[0.75A + 0.5(W2 + W5 + W8) + 1.5Z + 0.25U_Y + 0.75A(W2 - W5) + 0.5AZ]/7.5$$

where U_Y was drawn from a standard normal. The outcome is now a continuous variable bounded in $[0, 1]$ (e.g. a proportion). We also generated the counterfactual outcomes $Y(1)$

	R	correlation 0.5			correlation 0.5			correlation 0			Z
		$W1$	$W2$	$W3$	$W4$	$W5$	$W6$	$W7$	$W8$	$W9$	
Parents of covariate Z	✓	✓			✓			✓			
Parents of the outcome Y	✓		✓			✓			✓		✓
Matching set 1	✓										
Matching set 2	✓		✓			✓			✓		

Table 4.3: For Simulation 2, the relationships between baseline covariates and the outcome as well as the adaptive pair-matching schemes.

and $Y(0)$ by intervening to set $A = a$. For each study design, this data generating process was repeated 2,500 times. The SATE and PATE were calculated as before. The SATE ranged from 0.25% to 3.1% with a mean of 1.6%. The true value of the PATE was 1.6%. Table 4.3 depicts the relationship between the baseline covariates and the outcome as well as the adaptive pair-matching schemes.

We compared the same algorithms: the unadjusted estimator, the MLE adjusting for the irrelevant covariate $W9$, the TMLE with data-adaptive initial estimation of the conditional mean outcome, and the C-TMLE pairing data-adaptive initial estimation of the conditional mean outcome with data-adaptive targeting. Our library for initial estimation of the conditional mean outcome $\bar{Q}_0(A, W)$ consisted of 12 working logistic regression models, each with an intercept and a main term for the exposure A and a main term for one candidate adjustment variable $\{\emptyset, R, W1, \ldots, W9, Z\}$. Our library for collaborative estimation of the exposure mechanism $g_0(A|W)$ included 12 working logistic regression models, each with an intercept and a main term for one candidate adjustment variable: $\{\emptyset, R, W1, \ldots, W9, Z\}$. We used the same measure of performance and cross-validation scheme. As before, inference was based on the estimated influence curve for the unadjusted estimator and the MLE and on the cross-validated estimate of the influence curve for the data-adaptive TMLEs (Sec. 4.5). We assumed the standardized estimator followed the Student's t-distribution with $n-2 = 28$ degrees of freedom for the non-matched trial and with $n/2 - 1 = 14$ degrees of freedom for the matched trial.

Results

The results for the second simulation study are given in Table 4.4 and largely echoed the above findings. Pair-matching, even on a single covariate (i.e. matching set 1), improved the precision of the analysis. Targeting the sample effect instead of the population effect further improved efficiency. Allowing for data-adaptive selection of the working model for initial estimation of $\bar{Q}_0(A, W)$ yielded even greater precision, and the most efficient analysis was with C-TMLE. Indeed, the MSE of the unadjusted estimator for the PATE in the non-matched trial was nearly 5 times higher than the MSE of the C-TMLE when matching on

predictive covariates (i.e. matching set 2). This resulted in over 30% more power to detect the intervention effect.

	PATE					SATE				
	MSE[a]	rMSE[b]	$\hat{\sigma}^c$	Pow[d]	Cov[e]	MSE[a]	rMSE[b]	$\hat{\sigma}^c$	Pow[d]	Cov[e]
	Non-Matched									
Unadj.	1.7E-4	1.00	0.013	0.21	0.94	1.5E-4	1.16	0.013	0.21	0.96
MLE	1.8E-4	0.97	0.013	0.23	0.94	1.5E-4	1.12	0.013	0.23	0.95
TMLE	1.1E-4	1.54	0.010	0.33	0.93	9.0E-5	1.92	0.010	0.31	0.96
C-TMLE	1.1E-4	1.57	0.010	0.34	0.93	8.6E-5	1.99	0.010	0.32	0.97
	Matching Set 1									
Unadj.	1.2E-4	1.43	0.012	0.21	0.96	9.9E-5	1.74	0.011	0.28	0.97
MLE	1.3E-4	1.34	0.011	0.23	0.96	1.1E-4	1.63	0.011	0.29	0.96
TMLE	9.9E-5	1.74	0.009	0.33	0.95	7.5E-5	2.30	0.009	0.38	0.96
C-TMLE	9.6E-5	1.79	0.009	0.36	0.94	7.5E-5	2.31	0.008	0.43	0.95
	Matching Set 2									
Unadj.	6.4E-5	2.70	0.011	0.18	0.99	4.5E-5	3.85	0.009	0.36	0.99
MLE	7.1E-5	2.43	0.011	0.21	0.99	5.2E-5	3.28	0.009	0.37	0.99
TMLE	5.1E-5	3.39	0.009	0.31	0.99	3.5E-5	4.86	0.008	0.46	0.98
C-TMLE	5.1E-5	3.35	0.009	0.35	0.98	3.6E-5	4.78	0.007	0.52	0.98

[a]Mean squared error: the bias (average deviation between the point estimate and sample-specific true value) - squared plus the variance
[b]Relative MSE: the MSE of the unadjusted estimator for the PATE in a non-matched trial relative to (divided by) the MSE of another estimator
[c]Average standard error estimate, based on the estimated influence curve
[d]Attained power: proportion of times the false null hypothesis was rejected
[e]Confidence interval (CI) coverage: proportion of times the true value was contained in the 95% CI

Table 4.4: Summary of estimator performance for Simulation 2. The rows denote the study design and the estimator: unadjusted, MLE adjusting for $W9$, TMLE with data-adaptive selection of the initial estimator, and Collaborative-TMLE (C-TMLE) with data-adaptive selection of the initial estimator paired with data-adaptive estimation of the exposure mechanism.

For these simulations, there was a notable impact of parameter specification on estimator performance. We first focus on the estimation of the PATE and then on estimation of the SATE. When the population effect was the target of inference, the gains in attained power from pair-matching were attenuated despite the gains in MSE. This was likely due to the slight underestimation of the standard error in the non-matched trial and overestimation in the pair-matched trial. Indeed, the confidence interval coverage in the non-matched trial was less than nominal (93-94%), while the coverage when matching well (i.e. set 2) approached 100%. For this set of simulations, the correction factor ρ_n (Eq. 4.5) used in variance estimation for the pair-matched design was approximately 0. As a result, the variance estimator in the pair-matched trial was quite conservative, and the cross-validation

selection scheme was more optimized for the non-matched trial. The latter point is evidenced by Table 4.5, which shows the proportion of times a candidate working model was selected. The logistic regression model adjusting for R was selected for initial estimation of $\bar{Q}_0(A, W)$ in 10% of the studies without matching and in 8% of the studies when matching well on R (i.e. set 1). Furthermore, when matching on several covariates (i.e. set 2), the selection of working models for $\bar{Q}_0(A, W)$ was very similar to the selection in the non-matched trial.

| | | \multicolumn{12}{c}{Selected in the working model for initial estimation of $\bar{Q}_0(A, W)$} |
Adjustment variable		$\emptyset$	R	$W1$	$W2$	$W3$	$W4$	$W5$	$W6$	$W7$	$W8$	$W9$	Z
PATE	non-matched	0	10	0	16	1	0	0	0	0	3	0	70
	matched set1	0	8	1	30	1	0	0	0	0	4	0	56
	matched set2	0	8	0	16	1	0	0	0	0	2	0	72
SATE	non-matched	0	11	1	18	1	0	0	0	0	3	0	67
	matched set1	0	2	3	55	3	1	1	1	1	11	1	21
	matched set2	0	6	2	30	3	2	1	2	1	7	2	44
		\multicolumn{12}{c}{Selected in the working model for estimation of $g_0(A	W)$}
Adjustment variable		$\emptyset$	R	$W1$	$W2$	$W3$	$W4$	$W5$	$W6$	$W7$	$W8$	$W9$	Z
PATE	non-matched	79	2	2	2	2	1	1	1	2	3	2	2
	matched set1	37	9	5	6	5	4	3	3	5	5	5	13
	matched set2	25	10	4	9	6	5	9	5	4	8	4	11
SATE	non-matched	78	2	2	3	2	1	1	1	2	3	2	2
	matched set1	25	9	5	7	6	6	5	5	6	7	6	11
	matched set2	13	10	6	12	6	5	9	6	6	10	5	11

Table 4.5: For Simulation 2, the proportion of times a covariate was selected in the working logistic regression model for initial estimation of $\bar{Q}_0(A, W)$ and in the working logistic regression model for collaborative estimation of the exposure mechanism $g_0(A|W)$.

In contrast, when estimating the SATE, smaller MSE directly translated to greater attained power, while maintaining nominal, if not conservative, confidence interval coverage. For example, the attained power of the TMLE was 31% in the non-matched trial, 38% when matching on a single covariate and 46% when matching on several covariates. Likewise, the attained power of the C-TMLE was 32% in the non-matched trial, 43% in the trial pair-matching on a single covariate and 52% in trial matching on several covariates. From Table 4.5, we see that the working model adjusting for R was selected for initial estimation of $\bar{Q}_0(A, W)$ in 11% of the studies without matching and only in 2% of the studies when matching well on R (i.e. set 1). In the latter, more weight was given to other predictive baseline covariates, such as $W2$ and $W8$.

4.7 Discussion

This paper builds on the rich history of covariate adjustment in randomized trials (e.g. [119–121, 123, 151, 152, 154, 155, 162]). In particular, Rubin and van der Laan [122] proposed the principle of *empirical efficiency maximization* as a strategy to select the estimator of $\bar{Q}_0(A, W)$ that minimized the empirical variance of the estimated efficient influence curve. Their procedure, however, relied on solving a weighted nonlinear least squares problem. Our approach only requires researchers to take the sample variance of the estimated influence curve. More recently, van der Laan and Gruber [160] proposed collaborative estimation of the exposure mechanism to achieve the greatest bias reduction in the targeting step of TMLE in a observational study. In randomized trials, there is no risk of bias from regression model misspecification (e.g. [125]). Thereby, the collaborative approach, implemented here, serves only to increase precision by estimating the known exposure mechanism. To our knowledge, this is the first research into C-TMLE in a randomized trial setting. Most recently, van der Laan [159] suggested selection of the candidate (C-)TMLE based on minimizing the estimated variance of its influence curve. Our paper generalizes this scheme for estimation and inference of both the population and sample average treatment effects in randomized trials with and without pair-matching.

Our simulations illustrate the performance of the proposed procedure in realistically-sized (i.e. small) trials. In particular, with only 15 (conditionally) independent units, our procedure was able to identify the optimal working model for initial estimation of $\bar{Q}_0(A, W)$ from a library of 12 candidates as well as for collaborative estimation of $g_0(A|W)$ from a library of 12 candidates, while maintaining close to nominal confidence interval coverage. The simulations also indicated the most efficient approach was estimating the sample effect with C-TMLE in pair-matched trial. Indeed, this approach was nearly 5 more efficient than targeting the population effect with the unadjusted estimator in the non-matched trial. Thereby, our procedure dispels the common concerns of "analytical limitations" to pair-matched trials (e.g. [82, 95, 150]).

There are several areas of future work. First, our library of candidate estimators was limited to simple parametric working models. This choice was made for pedagogic purpose and to avoid over-fitting in small trials. In larger trials, we can expand the library to include working models with multiple adjustment variables and interactions as well as selection procedures (e.g. stepwise regression) and semiparametric algorithms. Future work will involve simulations to evaluate the methodology in larger trials. The application to matched triplets, as opposed to matched pairs, should be straightforward. However, the impact of adaptive stratification on estimation and inference merits additional consideration. Finally, we focused on two causal parameters: the population and sample average treatment effects. TMLE is a general methodology for the construction of double robust, semiparametric, efficient substitution estimators for a wide range of parameters. Our proposed strategy for covariate selection should extend to other causal parameters, such as the conditional average treatment effect, the average treatment effect among the treated, and the natural direct effect.

Overall, we proposed a general strategy to increase power in randomized trials. Specifically, we used cross-validation to select the candidate TMLE that optimized the efficiency of the analysis. Since the step-by-step algorithm (including the library definition) was pre-specified, there was no risk of bias or misleading inference from *ad hoc* analytic decisions. In other words, we have proposed a black box procedure to data-adaptively select the most powerful analysis. Furthermore, including the unadjusted estimator as a candidate obviates the need for guidelines on whether or not to adjust (e.g. [154, 155]). Finally, our procedure is tailored to the scientific question (population vs. sample effect) and study design (with or without pair-matching). Decisions about whether to adjust and how to adjust are made with a rigorous and principled approach, removing some of the "human art" from statistics.

Bibliography

[1]   J. Pearl. "An Introduction to Causal Inference". *The International Journal of Biostatistics* 6.2 (2010), Article 7.

[2]   E.H. Geng et al. "A Causal Framework for Understanding the Effect of Losses to Follow-up on Epidemiologic Analyses in Clinic-based Cohorts: The Case of HIV-infected Patients on Antiretroviral Therapy in Africa". *American Journal of Epidemiology* 175.10 (2012), pp. 1080–1087.

[3]   M. van der Laan and S. Rose. *Targeted Learning: Causal Inference for Observational and Experimental Data.* New York Dordrecht Heidelberg London: Springer, 2011.

[4]   M.L. Petersen and M.J. van der Laan. "Causal Models and Learning from Data: Integrating Causal Modeling and Statistical Estimation". *Epidemiology* 25.3 (2014), pp. 418–426.

[5]   J. Ahern and A. Hubbard. "A roadmap for estimating and interpreting population intervention parameters". In: *Methods in Social Epidemiology.* Ed. by J.M. Oakes and J.S. Kaufman. San Francisco: Jossey-Bass, Forthcoming.

[6]   M.A. Hernán, S. Hernández-Díaz, and J.M. Robins. "A structural approach to selection bias". *Epidemiology* 15.5 (2004), pp. 615–625.

[7]   S. Rose. "Big data and the future". *Significance* 9.4 (2012), pp. 47–48.

[8]   G. Marcus and E. Davis. "Eight (No, Nine!) Problems With Big Data". *The New York Times* (2014). URL: http://www.nytimes.com/2014/04/07/opinion/eight-no-nine-problems-with-big-data.html.

[9]   S.D. Lawson and A.D. Harries. "Reducing tuberculosis-associated early mortality in antiretroviral treatment programmes in sub-Saharan Africa". *AIDS* 25.12 (2011), pp. 1554–1555. DOI: 10.1097/QAD.0b013e328348fb61.

[10] M.S. Cohen et al. "Prevention of HIV-1 infection with early antiretroviral therapy". *New England Journal of Medicine* 365.6 (2011), pp. 493–505. DOI: 10.1056/NEJMoa1105243.

[11] World Health Organization. *Consolidated Guidelines on the Use of Antiretroviral Drugs for Treating and Preventing HIV Infection.* Geneva, 2013.

[12] M. Egger et al. "Cohort Profile: The international epidemiological databases to evaluate AIDS (IeDEA) in Sub-Saharan Africa". *International Journal of Epidemiology* 41.5 (2011), pp. 1256–1264.

[13] University of California, San Francisco. *Sustainable East Africa Research in Community Health (SEARCH)*. ClinicalTrials.gov. 2013. URL: `http://clinicaltrials.gov/show/NCT01864603`.

[14] J. Pearl. *Causality: Models, Reasoning and Inference.* Second ed., 2009. New York: Cambridge University Press, 2000.

[15] J. Pearl. *Probabilistic Reasoning in Intelligent Systems.* San Mateo, A: Morgan Kaufmann, 1988.

[16] J. Pearl. "Causal diagrams for empirical research". *Biometrika* 82 (1995), pp. 669–710.

[17] A. Goldberger. "Structural equation models in the social sciences". *Econometrica: Journal of the Econometric Society* 40 (1972), pp. 979–1001.

[18] O. Duncan. *Introduction to Structural Equation Models.* New York: Academic Press, 1975.

[19] J. Neyman. "Sur les applications de la theorie des probabilites aux experiences agricoles: Essai des principes (In Polish). English translation by D.M. Dabrowska and T.P. Speed (1990)". *Statistical Science* 5 (1923), pp. 465–480.

[20] D.B. Rubin. "Estimating causal effects of treatments in randomized and nonrandomized studies." *Journal of Educational Psychology* 66.5 (1974), pp. 688–701. DOI: `10.1037/h0037350`.

[21] P.W. Holland. "Statistics and Causal Inference". *Journal of the American Statistical Association* 81.396 (1986), pp. 945–960.

[22] M.A. Hernán et al. "Comparison of dynamic treatment regimes via inverse probability weighting". *Basic & Clinical Pharmacology & Toxicology* 98.3 (2006), pp. 237–242.

[23] M.J. van der Laan and M.L. Petersen. "Causal Effect Models for Realistic Individualized Treatment and Intention to Treat Rules". *The International Journal of Biostatistics* 3.1 (2007), Article 3.

[24] L.E. Cain et al. "When to start treatment? A systematic approach to the comparison of dynamic regimes using observational data". *The International Journal of Biostatistics* 6.2 (2010), Article 18.

[25] M.L. Petersen and M.J. van der Laan. "Case Study: Longitudinal HIV Cohort Data". In: *Targeted Learning: Causal Inference for Observational and Experimental Data.* Ed. by M.J. van der Laan and S. Rose. New York Dordrecht Heidelberg London: Springer, 2011.

[26] I. Díaz Muñoz and M. van der Laan. "Population Intervention Causal Effects Based on Stochastic Interventions". *Biometrics* 68.2 (2012), pp. 541–549.

[27] J.M. Robins. "A new approach to causal inference in mortality studies with sustained exposure periods–application to control of the healthy worker survivor effect". *Mathematical Modelling* 7 (1986), pp. 1393–1512.

[28] M.A. Hernán and T.J. VanderWeele. "Compound treatments and transportability of causal inference". *Epidemiology* 22 (2011), pp. 368–377.

[29] M.L. Petersen. "Compound treatments, transportability, and the structural causal model: the power and simplicity of causal graphs". *Epidemiology* 22 (2011), pp. 378–381.

[30] J. Pearl and E. Bareinboim. *Transportability across studies: A formal approach.* Tech. rep. R-372. Computer Science Department, University of California, Los Angeles, 2013. URL: ftp://www.mmss.cs.ucla.edu/pub/stat_ser/r372.pdf.

[31] J.M. Robins, M.A. Hernán, and B. Brumback. "Marginal structural models and causal inference in epidemiology". *Epidemiology* 11.5 (2000), pp. 550–560.

[32] R. Neugebauer and M. J. van der Laan. "Nonparametric causal effects based on marginal structural models". *Journal of Statistical Planning and Inference* 137.2 (2007), pp. 419–434.

[33] M.J. van der Laan. "Estimation based on case-control designs with known prevalence probability". *The International Journal of Biostatistics* 4.1 (2008), Article 17.

[34] S. Rose and M.J. van der Laan. "Why match? Investigating matched case-control study designs with causal effect estimation". *The International Journal of Biostatistics* 5.1 (2009), Article 1.

[35] D. Rubin. "Randomization Analysis of Experimental Data: The Fisher Randomization Test Comment". *Journal of the American Statistical Association* 75.371 (1980), pp. 591–593.

[36] M.J. van der Laan. "Causal Inference for a Population of Causally Connected Units". *Journal of Causal Inference* 0.0 (2014), pp. 1–62. DOI: 10.1515/jci-2013-0002.

[37] M.L. Petersen et al. "Diagnosing and responding to violations in the positivity assumption". *Statistical Methods in Medical Research* 21.1 (2012), pp. 31–54. DOI: 10.1177/0962280210386207.

[38] E.C. Polley, S. Rose, and M.J. van der Laan. "Super Learner". In: *Targeted Learning: Causal Inference for Observational and Experimental Data.* Ed. by M.J. van der Laan and S. Rose. New York Dordrecht Heidelberg London: Springer, 2011.

[39] M.J. van der Laan, E.C. Polley, and A.E. Hubbard. "Super learner". *Statistical Applications in Genetics and Molecular Biology* 6.1 (2007), p. 25.

[40] M.J. van der Laan and D.B. Rubin. "Targeted Maximum Likelihood Learning". *The International Journal of Biostatistics* 2.1 (2006), Article 11. DOI: 10.2202/1557-4679.1043.

[41] S. Rose and M.J. van der Laan. "Why TMLE?" In: *Targeted Learning: Causal Inference for Observational and Experimental Data*. Ed. by M.J. van der Laan and S. Rose. New York Dordrecht Heidelberg London: Springer, 2011.

[42] K.E. Porter et al. "The relative performance of targeted maximum likelihood estimators". *The International Journal of Biostatistics* 7.1 (2011), Article 31.

[43] O.M. Stitelman, V. De Gruttola, and M.J. van der Laan. "A General Implementation of TMLE for Longitudinal Data Applied to Causal Inference in Survival Analysis". *The International Journal of Biostatistics* 8.1 (2012).

[44] S. Gruber and M.J. van der Laan. "tmle: An R Package for Targeted Maximum Likelihood Estimation". *Journal of Statistical Software* 51.13 (2012), pp. 1–35.

[45] J. Schwab et al. *ltmle: Longitudinal Targeted Maximum Likelihood Estimation*. R package version 0.9.3-1. 2014. URL: http://CRAN.R-project.org/package=ltmle.

[46] 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.

[47] D.G. Horvitz and D.J. Thompson. "A generalization of sampling without replacement from a finite universe". *Journal of the American Statistical Association* 47 (1952), pp. 663–685.

[48] P.R. Rosenbaum and D.B. Rubin. "The central role of the propensity score in observational studies for causal effects". *Biometrika* 70 (1983), pp. 41–55.

[49] J.M. Robins. "Analytic methods for estimating HIV treatment and cofactor effects". In: *Methodological issues of AIDS Mental Health Research*. Ed. by D.G. Ostrow and R. Kessler. New York: Plenum Publishing, 1993, pp. 213–290.

[50] J.M. Robins and M.A. Hernán. "Estimation of the causal effects of time-varying exposures". In: *Longitudinal Data Analysis*. Ed. by G. Fitzmaurice et al. Boca Raton, FL: Chapman & Hall/CRC, 2009. Chap. 23.

[51] M.L. Petersen et al. "Assessing the effectiveness of antiretroviral adherence interventions. Using marginal structural models to replicate the findings of randomized controlled trials". *JAIDS* 43.Suppl 1 (2006), S96–S103.

[52] 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 (1994), pp. 846–866.

[53] A. Rotnitzky, J.M. Robins, and D.O. Scharfstein. "Semiparametric regression for repeated outcomes with nonignorable nonresponse". *Journal of the American Statistical Association* 93 (1998), pp. 1321–1339.

[54] J.M. Robins. "Robust estimation in sequentially ignorable missing data and causal inference models". In: *1999 Proceedings of the American Statistical Association*. Alexandria, VA: American Statistical Association, 2000, pp. 6–10.

[55] M.J. van der Laan and J.M. Robins. *Unified Methods for Censored Longitudinal Data and Causality*. New York Berlin Heidelberg: Springer-Verlag, 2003.

[56] L.B. Balzer and M.J. van der Laan. *Estimating Effects on Rare Outcomes: Knowledge is Power*. Tech. rep. 310. Division of Biostatistics, University of California, Berkeley, 2014. URL: http://biostats.bepress.com/ucbbiostat/paper310/.

[57] M.A. Hernán, B. Brumback, and J.M. Robins. "Marginal structural models to estimate the causal effect of zidovudine on the survival of HIV-positive men". *Epidemiology* 11.5 (2000), pp. 561–570.

[58] L.M. Bodnar et al. "Marginal Structural Models for Analyzing Causal Effects of Time-dependent Treatments: An Application in Perinatal Epidemiology". *American Journal of Epidemiology* 159.10 (2004), pp. 926–934.

[59] S.R. Cole and M.A. Hernán. "Constructing Inverse Probability Weights for Marginal Structural Models". *American Journal of Epidemiology* 168.6 (2008), pp. 656–664.

[60] J.M. Robins, L. Orellana, and A. Rotnitzky. "Estimation and extrapolation of optimal treatment and testing strategies". *Statistics in Medicine* 27.23 (2008), pp. 4678–4721.

[61] S.L. Taubman et al. "Intervening on risk factors for coronary heart disease: an application of the parametric G-formula". *International Journal of Epidemiology* 38.6 (2009), pp. 1599–1611.

[62] E.C. Polley and M.J. van der Laan. "Super Learning for Right-Censored Data". In: *Targeted Learning: Causal Inference for Observational and Experimental Data*. Ed. by M.J. van der Laan and S. Rose. New York Dordrecht Heidelberg London: Springer, 2011.

[63] K.L. Moore and M.J. van der Laan. "RCTs with Time-to-Event Outcomes". In: *Targeted Learning: Causal Inference for Observational and Experimental Data*. Ed. by M.J. van der Laan and S. Rose. New York Dordrecht Heidelberg London: Springer, 2011.

[64] O.M. Stitelman and M.J. van der Laan. "Collaborative targeted maximum likelihood for time-to-event data". *The International Journal of Biostatistics* 6.1 (2010), Article 21.

[65] M.J. van der Laan and S. Gruber. "Targeted minimum loss based estimation of causal effects of multiple time point interventions". *The International Journal of Biostatistics* 8.1 (2012).

[66] M.E. Schnitzer, E.E. Moodie, and R.W. Platt. "Targeted maximum likelihood estimation for marginal time-dependent treatment effects under density misspecification". *Biostatistics* (2012).

[67] M.E. Schnitzer et al. "Effect of breastfeeding on gastrointestinal infection in infants: a targeted maximum likelihood approach for clustered longitudinal data". *Annals of Applied Statistics* 8.2 (2014), pp. 703–725.

[68] M.L. Petersen et al. "Targeted Maximum Likelihood Estimation for Dynamic and Static Longitudinal Marginal Structural Working Models". *Journal of Causal Inference* 2.2 (2014). DOI: `10.1515/jci-2013-0007`.

[69] A.L. Decker et al. "Semiparametric Estimation of the Impacts of Longitudinal Interventions on Adolescent Obesity using Targeted Maximum-Likelihood: Accessible Estimation with the ltmle Package". *Journal of Causal Inference* 2.1 (2014), pp. 95–108.

[70] 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* 1 (2007), pp. 574–596.

[71] M.M. Kitahata et al. "Effect of early versus deferred antiretroviral therapy for HIV on survival". *New England Journal of Medicine* 360.18 (2009), pp. 1815–1826.

[72] M.A. Hernán and J.M. Robins. "Comment on: Early versus deferred antiretroviral therapy for HIV on survival". *New England Journal of Medicine* 361.8 (2009), pp. 823–824.

[73] M.L. Petersen, S.E. Sinisi, and M.J. van der Laan. "Estimation of direct causal effects". *Epidemiology* 17.3 (2006), pp. 276–284.

[74] M.J. van der Laan and M.L. Petersen. "Direct effect models". *The International Journal of Biostatistics* 4.1 (2008), Article 23.

[75] T.J. VanderWeele. "Marginal structural models for the estimation of direct and indirect effects". *Epidemiology* 20.1 (2009), pp. 18–26.

[76] A.E. Hubbard, N.P. Jewell, and M.J. van der Laan. "Direct Effects and Effect Among the Treated". In: *Targeted Learning: Causal Inference for Observational and Experimental Data*. Ed. by M.J. van der Laan and S. Rose. New York Dordrecht Heidelberg London: Springer, 2011.

[77] W. Zheng and M.J. van der Laan. "Targeted maximum likelihood estimation for natural direct effects". *The International Journal of Biostatistics* 8.1 (2012), pp. 1–40.

[78] W. Zheng and M.J. van der Laan. *Causal Mediation in a Survival Setting with Time-Dependent Mediators*. Tech. rep. 295. Division of Biostatistics, University of California, Berkeley, 2012. URL: `http://biostats.bepress.com/ucbbiostat/paper295/`.

[79] S. Lendle, M.S. Subbaraman, and M.J. van der Laan. "Identification and efficient estimation of the natural direct effect among the untreated". *Biometrics* 69 (2013), pp. 310–317.

[80] M.C. Costanza. "Matching". *Preventive Medicine* 24.5 (1995), pp. 425–433. DOI: `10.1006/pmed.1995.1069`.

[81] K.J. Rothman, S. Greenland, and T.L. Lash. *Modern Epidemiology*. Phildelphia: Lippincott Williams & Wilkins, 2008.

[82] N. Klar and A. Donner. "The merits of matching in community intervention trials: a cautionary tale". *Statistics in Medicine* 16.15 (1997), pp. 1753–1764. DOI: `10.1002/(SICI)1097-0258(19970815)16:15<1753::AID-SIM597>3.0.CO;2-E`.

[83] W.G. Cochran. "Matching in Analytical Studies". *American Journal of Public Health and the Nations Health* 43.6 (1953), pp. 684–691. DOI: `10.2105/AJPH.43.6_Pt_1.684`.

[84] W.Z. Billewicz. "Matched samples in medical investigations". *British Journal of Preventive & Social Medicine* 18 (1964), pp. 167–173. DOI: `10.1136/jech.18.4.167`.

[85] L.H. Youkeles. "184 Note: Loss of Power through Ineffective Pairing of Observations in Small Two-Treatment All-or-None Experiments". *Biometrics* 19.1 (1963), pp. 175–180. DOI: `10.2307/2527582`.

[86] K.K. Mathen. "Matching in comparative studies in public health". *Indian Journal of Public Health* 8.4 (1963), pp. 161–169.

[87] W.Z. Billewicz. "The Efficiency of Matched Samples: An Empirical Investigation". *Biometrics* 21.3 (1965), pp. 623–644. DOI: `10.2307/2528546`.

[88] O.S. Miettinen. "The matched pairs design in the case of all-or-none responses". *Biometrics* 24.2 (1968), pp. 339–352. DOI: `10.2307/2528039`.

[89] S.M. McKinlay. "Pair-Matching–A Reappraisal of a Popular Technique". *Biometrics* 33.4 (1977), pp. 725–735. DOI: `10.2307/2529471`.

[90] S. Wacholder and C.R. Weinberg. "Paired versus Two-Sample Design for a Clinical Trial of Treatments with Dichotomous Outcome: Power considerations". *Biometrics* 38.3 (1982), pp. 801–812. DOI: `10.2307/2530059`.

[91] D.C. Martin et al. "The effect of matching on the power of randomized community intervention studies". *Statistics in Medicine* 12.3-4 (1993), pp. 329–338. DOI: `10.1002/sim.4780120315`.

[92] P. Diehr et al. "Breaking the matches in a paired *t*-test for community interventions when the number of pairs is small". *Statistics in Medicine* 14.13 (1995), pp. 1491–1504. DOI: `10.1002/sim.4780141309`.

[93] R. Greevy et al. "Optimal multivariate matching before randomization". *Biostatistics* 5.2 (2004), pp. 263–275. DOI: `10.1093/biostatistics/5.2.263`.

[94] K. Imai, G. King, and C. Nall. "The Essential Role of Pair Matching in Cluster-Randomized Experiments, with Application to the Mexican Universal Health Insurance Evaluation". *Statistical Science* 24.1 (2009), pp. 29–53. DOI: `10.1214/08-STS274`.

[95] G.W. Imbens. *Experimental Design for Unit and Cluster Randomized Trials*. Tech. rep. NBER Technical Working Paper, 2011.

[96] R.J. Hayes and L.H. Moulton. *Cluster Randomised Trials*. Boca Raton: Chapman & Hall/CRC, 2009.

[97] L.S. Freedman et al. "The Efficiency of the Matched-Pairs Design of the Community Intervention Trial for Smoking Cessation (COMMIT)". *Controlled Clinical Trials* 18.2 (1997), pp. 131–139. DOI: 10.1016/S0197-2456(96)00115-8.

[98] D.M. Murray. *Design and Analysis of Group-Randomized Trials*. New York: Oxford University Press, 1998.

[99] A. Donner and N. Klar. *Design and Analysis of Cluster Randomization Trials in Health Research*. London: Arnold, 2000.

[100] M.J. Campbell, A. Donner, and N. Klar. "Developments in cluster randomized trials and *Statistics in Medicine*". *Statistics in Medicine* 26.1 (2007), pp. 2–19. DOI: 10.1002/sim.2731.

[101] L.B. Balzer, M.L. Petersen, and M.J. van der Laan. *Why Match in Individually and Cluster Randomized Trials?* Tech. rep. 294. Division of Biostatistics, University of California, Berkeley, 2012. URL: http://www.bepress.com/ucbbiostat/paper294.

[102] R. Granich et al. "Antiretroviral therapy in prevention of HIV and TB: update on current research efforts". *Current HIV Research* 9.6 (2011), pp. 446–469. DOI: 10.2174/157016211798038597.

[103] L.J. Wei et al. "Statistical inference with data-dependent treatment allocation rules". *Journal of the American Statistical Association* 85.156-162 (1990). DOI: 10.1080/01621459.1990.10475319.

[104] R.N. Tamura et al. "A Case Study of an Adaptive Clinical Trial in the Treatment of Out-Patients with Depressive Disorder". *Journal of the American Statistical Association* 89.427 (1994), pp. 768–776. DOI: 10.2307/2290902.

[105] L.M. Collins, S.A. Murphy, and K.L. Bierman. "A conceptual framework for adaptive preventive interventions". *Prevention Science* 5.3 (2004), pp. 185–196. DOI: 10.1023/B:PREV.0000037641.26017.00.

[106] A. Chambaz and M.J. van der Laan. "Targeting the optimal design in randomized clinical trials with binary outcomes and no covariate: Theoretical Study". *The International Journal of Biostatistics* 7.1 (2011), Article 10. DOI: 10.2202/1557-4679.1247.

[107] H. Grosskurth et al. "Impact of improved treatment of sexually transmitted diseases on HIV infection in rural Tanzania: randomised controlled trial". *The Lancet* 346.8974 (1995), pp. 530–536. DOI: 10.1016/S0140-6736(95)91380-7.

[108] L. Watson et al. "Mounting a community-randomized trial: sample size, matching, selection, and randomization issues in PRISM". *Controlled Clinical Trials* 25.3 (2004), pp. 235–250. DOI: 10.1016/j.cct.2003.12.002.

[109] M. Toftager et al. "SPACE for physical activity-a multicomponent intervention study: study design and baseline findings from a cluster randomized controlled trial". *BMC Public Health* 11 (2011), p. 777. DOI: 10.1186/1471-2458-11-777.

[110] G. King et al. "A "Politically Robust" Experimental Design for Public Policy Evaluation, with Application to the Mexican Universal Health Insurance Program". *Journal of Policy Analysis and Management* 26.3 (2007), pp. 479–506. DOI: `10.1002/pam.20279`.

[111] B. Lu, R. Greevy, and C. Beck. *nbpMatching: functions for non-bipartite optimal matching.* R package version 1.3.6. 2012. URL: `http://CRAN.R-project.org/package=nbpMatching`.

[112] K. Zhang and D.S. Small. "Comment: The Essential Role of Pair Matching in Cluster-Randomized Experiments, with Application to the Mexican Universal Health Insurance Evaluation". *Statistical Science* 25.1 (2009), pp. 59–64. DOI: `10.1214/09-STS274B`.

[113] B. Lu et al. "Optimal Nonbipartite Matching and its Statistical Applications". *American Statistician* 65.1 (2011), pp. 21–30. DOI: `10.1198/tast.2011.08294`.

[114] Robert A. Greevy et al. "Reweighted Mahalanobis distance matching for cluster-randomized trials with missing data". *Pharmacoepidemiology and Drug Safety* 21.148–154 (2012). DOI: `10.1002/pds.3260`.

[115] M.J. van der Laan, L.B. Balzer, and M.L. Petersen. "Adaptive Matching in Randomized Trials and Observational Studies". *Journal of Statistical Research* 46.2 (2012), pp. 113–156.

[116] A. Abadie and G. Imbens. *Simple and Bias-Corrected Matching Estimators for Average Treatment Effects.* Tech. rep. 283. NBER technical working paper, 2002.

[117] G.W. Imbens. "Nonparametric estimation of average treatment effects under exogeneity: a review". *Review of Economics and Statistics* 86.1 (2004), pp. 4–29. DOI: `10.1162/003465304323023651`.

[118] A. Abadie and G.W. Imbens. "Estimation of the Conditional Variance in Paired Experiments". *Annales d'Economie et de Statistique* 91/92 (2008), pp. 175–187.

[119] R.A. Fisher. *Statistical methods for research workers.* 4th. Edinburgh: Oliver and Boyd Ltd., 1932.

[120] W.G. Cochran. "Analysis of Covariance: its Nature and Uses." *Biometrics* 13 (1957), pp. 261–281.

[121] D.R. Cox and P. McCullagh. "Some aspects of analysis of covariance". *Biometrics* 38.3 (1982), pp. 541–561.

[122] Daniel B. Rubin and M.J. van der Laan. "Empirical efficiency maximization: improved locally efficient covariate adjustment in randomized experiments and survival analysis". *The International Journal of Biostatistics* 4.1 (2008), Article 5. DOI: `10.2202/1557-4679.1084`.

[123] A.A. Tsiatis et al. "Covariate adjustment for two-sample treatment comparisons in randomized clinical trials: A principled yet flexible approach". *Statistics in Medicine* 27.23 (2008), pp. 4658–4677. DOI: 10.1002/sim.3113.

[124] 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 (2009), pp. 39–64. DOI: 10.1002/sim.3445.

[125] 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 (2010), Article 13. DOI: 10.2202/1557-4679.1138.

[126] K. Imai. "Variance identification and efficiency analysis in randomized experiments under the matched-pair design". *Statistics in Medicine* 27.24 (2008), pp. 4857–4873. DOI: 10.1002/sim.3337.

[127] J. Hill and M. Scott. "Comment: The Essential Role of Pair Matching". *Statistical Science* 24.1 (2009), pp. 54–58. DOI: 10.1214/09-STS274A.

[128] The COMMIT Research Group. "Community Intervention Trial for Smoking Cessation (COMMIT): I. Cohort Results from a Four-Year Community Intervention". *American Journal of Public Health* 85 (1995), pp. 183–192. DOI: 10.2105/ajph.85.2.183.

[129] L.B. Balzer et al. "Adaptive pair-matching in randomized trials with unbiased and efficient effect estimation". *Statistics in Medicine* 34.6 (2015), pp. 999–1011.

[130] H. Thirumurthy, M. Goldstein, and J. Graff Zivin. "The economic impact of AIDS treatment: labor supply in Western Kenya". *Journal of Human Resources* 43 (2008), pp. 511–552. DOI: 10.1353/jhr.2008.0009.

[131] G. J. Zivin, H. Thirumurthy, and M. Goldstein. "AIDS Treatment and Intrahousehold Resource Allocation: Children's Nutrition and Schooling in Kenya". *Journal of Public Economics* 93.7-8 (2009), pp. 1008–1015. DOI: 10.1016/j.jpubeco.2009.03.003.

[132] K. Middelkoop et al. "Antiretroviral therapy and TB notification rates in a high HIV prevalence South African community". *Journal of Acquired Immune Deficiency Syndromes* 56.3 (2011), pp. 263–269. DOI: 10.1097/QAI.0b013e31820413b3.

[133] C. Flateau, G. Le Loup, and G. Pialoux. "Consequences of HIV infection on malaria and therapeutic implications: a systematic review". *The Lancet Infectious Diseases* 11.7 (2011), pp. 541–556. DOI: 10.1016/S1473-3099(11)70031-7.

[134] M.G. Fowler et al. "Perinatal HIV and its prevention: progress toward an HIV-free generation". *Clinics in Perinatology* 37.4 (2012), pp. 699–719, vii. DOI: 10.1016/j.clp.2010.09.002.

[135] E.J. Schouten et al. "Prevention of mother-to-child transmission of HIV and the health related Millennium Development Goals: time for a public health approach". *The Lancet* 378.9787 (2011), pp. 282–284. DOI: `10.1016/S0140-6736(10)62303-3`.

[136] V. Jain et al. "Assessment of Population-Based HIV RNA Levels in a Rural East African Setting Using a Fingerprick-Based Blood Collection Method". *Clinical Infectious Diseases* 56.4 (2013), pp. 598–605. DOI: `10.1093/cid/cis881`.

[137] 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 (2010), Article 26. DOI: `10.2202/1557-4679.1260`.

[138] D.S. Small, T.R. Ten Have, and P.R. Rosenbaum. "Randomization Inference in a Group–Randomized Trial of Treatments for Depression: Covariate Adjustment, Noncompliance, and Quantile Effects". *Journal of the American Statistical Association* 103.481 (2008), pp. 271–279. DOI: `10.1198/016214507000000897`.

[139] K. Zhang, M. Traskin, and D.S. Small. "A Powerful and Robust Test Statistic for Randomization Inference in Group-Randomized Trials with Matched Pairs of Groups". *Biometrics* 68 (2012), pp. 75–84. DOI: `10.1111/j.1541-0420.2011.01622.x`.

[140] L.B. Balzer, M.L. Petersen, and M.J. van der Laan. *Targeted estimation and inference of the sample average treatment effect*. Tech. rep. 334. Division of Biostatistics, University of California at Berkeley, 2015. URL: `http://biostats.bepress.com/ucbbiostat/paper334/`.

[141] Donald B. Rubin. "Comment: Neyman (1923) and Causal Inference in Experiments and Observational Studies". *Statistical Science* 5.4 (1990), pp. 472–480.

[142] 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 (2009), pp. 1140–1147.

[143] J.M. Snowden, S. Rose, and K.M. Mortimer. "Implementation of G-Computation on a Simulated Data set: demonstration of a Causal Inference Technique". *American Journal of Epidemiology* 173.7 (2011), pp. 731–738.

[144] P.J. Bickel et al. *Efficient and Adaptive Estimation for Semiparametric Models*. Baltimore: Johns Hopkins University Press, 1993.

[145] E.A. Stuart et al. "The use of propensity scores to assess the generalizability of results from randomized trials". *Journal of the Royal Statistical Society: Series A* 174.Part 2 (2011), pp. 369–386.

[146] E. Hartman et al. "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* (2015), pp. 2–22.

[147] S.J. Pocock et al. "Subgroup analysis, covariate adjustment and baseline comparisons in clinical trial reporting: current practice and problems". *Statistics in Medicine* 21.19 (2002), pp. 2917–2930. DOI: 10.1002/sim.1296.

[148] P.C. Austin et al. "A substantial and confusing variation exists in handling of baseline covariates in randomized controlled trials: a review of trials published in leading medical journals". *Journal of Clinical Epidemiology* 63.142-153 (2010).

[149] B.C. Kahn et al. "The risks and rewards of covariate adjustment in randomized trials: an assessment of 12 outcomes from 8 studies". *Trials* 15.139 (2014), pp. 1–7.

[150] M.J. Campbell. "Cluster Randomized Trials". In: *Handbook of Epidemiology, 2nd edition*. Ed. by W. Ahrens and I. Pigeot. Springer, 2014.

[151] M. Zhang, A.A. Tsiatis, and M. Davidian. "Improving Efficiency of Inferences in Randomized Clinical Trials Using Auxiliary Covariates". *Biometrics* 64.3 (2008), pp. 707–715. DOI: 10.1111/j.1541-0420.2007.00976.x.

[152] C. Shen, X. Li, and L. Li. "Inverse probability weighting for covariate adjustment in randomized studies". *Statistics in Medicine* 33 (2014), pp. 555–568.

[153] D.O. Scharfstein, A. Rotnitzky, and J.M. Robins. "Adjusting for Nonignorable Drop-Out Using Semiparametric Nonresponse Models (with Rejoiner)". *Journal of the American Statistical Association* 94.448 (1999), 1096–1120 (1135–1146). DOI: 10.2307/2669930.

[154] E. Colantuoni and M. Rosenblum. *Leveraging Prognostic Baseline Variables to Gain Precision in Randomized Trials*. Tech. rep. 263. Johns Hopkins University, Dept. of Biostatistics Working Papers, 2015. URL: http://biostats.bepress.com/jhubiostat/paper263.

[155] K.L. Moore et al. "Robust extraction of covariate information to improve estimation efficiency in randomized trials". *Statistics in Medicine* 30.19 (2011), pp. 2389–2408.

[156] R.M. Califf et al. "Characteristics of Clinical Trials Registered in ClinicalTrials.gov, 2007-2010". *JAMA* 307.17 (2012), pp. 1838–1847.

[157] S. Selvaraj and V. Prasad. "Characteristics of Cluster Randomized Trials: Are They Living Up to the Randomized Trial?" *JAMA Internal Medicine* 173.23 (2013), p. 313.

[158] B.A. Olken. *Pre-Analysis Plans in Economics*. Tech. rep. Massachusetts Institute of Technology Department of Economics, 2015. URL: http://economics.mit.edu/files/10399.

[159] M.J. van der Laan. "Appendix A.19: Efficiency maximization and TMLE". In: *Targeted Learning: Causal Inference for Observational and Experimental Data*. Ed. by M.J. van der Laan and S. Rose. New York Dordrecht Heidelberg London: Springer, 2011.

[160] M.J. van der Laan and S. Gruber. "Collaborative double robust targeted maximum likelihood estimation". *The International Journal of Biostatistics* 6.1 (2010).

[161] S. Gruber and M.J. van der Laan. "C-TMLE of an Additive Point Treatment Effect". In: *Targeted Learning: Causal Inference for Observational and Experimental Data*. Ed. by M.J. van der Laan and S. Rose. New York Dordrecht Heidelberg London: Springer, 2011.

[162] S. Yuan, H.H. Zhang, and M. Davidian. "Variable selection for covariate-adjusted semiparametric inference in randomized clinical trials". *Statistics in Medicine* 31 (2012), pp. 3789–3804.

[163] A.W. van der Vaart and J.A. Wellner. *Weak convergence and empirical processes*. Berlin Heidelberg New York: Springer, 1996.

Appendix A

Technical Appendix for Chapter 1

In the Introduction (Chapter 1), we focused on causal parameters corresponding to a static intervention on a single node. In this Appendix, we step through the causal roadmap for an example of a longitudinal effect, corresponding to a multiple time point intervention.

Step 1 - Specify the scientific question: What is the effect of delayed ART initiation on patient outcomes? As before, we want to be specific about the target population: recently diagnosed HIV+ adults in Sub-Saharan Africa. We also need to be clear about the definition and timing of the exposures. For simplicity, let us assume the patients have monthly clinic visits and therefore could initiate ART or not each month. (This framework could easily be extended to shorter or longer time intervals.) Suppose the outcome is viral suppression after 12 months of follow up.

Step 2 - Specify the causal model: Let baseline ($t = 0$) be the time the patient is diagnosed with HIV. Let $L(0)$ represent the vector of baseline covariates, including socio-demographics, clinical measurements and social constructs. Likewise, let $L(t)$ represent the vector of time-updated covariates (e.g. clinical measurements). Let $A(t)$ be an indicator that the patient initiated ART at time t. For example, $A(0) = 1$ represents starting ART on the same day as diagnosis (i.e. month 0), while $A(1) = 1$ represents initiation at the first month clinic visit. Finally, let Y be an indicator that the patient had undetectable HIV RNA viral load at the end of follow up. For simplicity, let us consider only three time points and assume complete follow up. Our structural causal model $\mathcal{M}^{\mathcal{F}}$, only reflecting the causal-ordering, is given by

- Endogenous nodes: $X = (L(0), A(0), L(1), A(1), Y)$

- Exogenous nodes: $U = (U_{L(0)}, U_{A(1)}, U_{L(1)}, U_{A(1)}, U_Y)$ with some true joint distribution $P_{U,0}$. We place no assumptions on set of possible distributions for U. (During the identifiability step, we will need to make some independence assumptions. However, we want to keep our true knowledge, as specified by structural causal model $\mathcal{M}^{\mathcal{F}}$, separate from the additional assumptions needed for identifiability.)
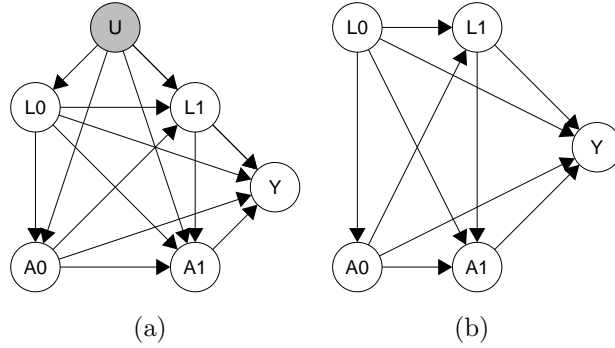
(a)   (b)

Figure A.1: Directed acyclic graph corresponding to the longitudinal effect (a) when we make no independence assumptions on background factors and (b) when we assume the background factors are all independent. $L(0)$ denotes baseline covariates; $A(0)$ denotes whether the patient initiated ART at $t = 0$; $L(1)$ denotes time-updated covariates; $A(1)$ denotes whether the patient initiated ART at $t = 1$, and Y denotes undetectable viral load.

- Structural equations:

$$L(0) = f_{L(0)}(U_{L(0)})$$
$$A(0) = f_{A(0)}(L(0), U_{A(0)})$$
$$L(1) = f_{L(1)}(L(0), A(0), U_{L(1)})$$
$$A(1) = f_{A(1)}(L(0), A(0), L(1), U_{A(1)})$$
$$Y = f_Y(L(0), A(0), L(1), A(1), U_Y)$$

- We have not made any exclusion restrictions or independence assumptions. The corresponding directed acyclic graph is given in Figure A.1a.

Step 3 - Specify the target causal quantity. Let $Y_{a(0)a(1)}$ denote the counterfactual outcome (viral suppression) if a patient, possibly contrary to fact, had treatment history $(a(0), a(1))$. Counterfactuals are generated by intervening on the structural causal model:

$$L(0) = f_{L(0)}(U_{L(0)})$$
$$A(0) = a(0)$$
$$L(1) = f_{L(1)}(L(0), a(0), U_{L(1)})$$
$$A(1) = a(1)$$
$$Y_{a(0)a(1)} = f_Y(L(0), a(0), L(1), a(1), U_Y)$$

For the two binary exposures (initiate or not at time t), the set of possible exposure combinations are $\mathcal{A} = \{10, 01, 00\}$. For example, Y_{01} corresponds to preventing ART initiation

at month 0 and starting ART at the one month clinic visit. Suppose our goal is to contrast expected counterfactual outcome if, possibly contrary to fact, all patients immediately initiated ART with the expected counterfactual outcome if, possibly contrary to fact, all patient delayed ART initiation until one month after diagnosis:

$$\Psi^{\mathcal{F}}(P_{U,X,0}) = E_{U,X,0}(Y_{10} - Y_{01})$$

Step 4 - Specify the observed data and their link to the causal model. The observed data consist of n i.i.d. copies of

$$O = \big(L(0), A(0), L(1), A(1), Y\big) \sim P_0$$

We assume the observed data were generated by sampling n independent times from a data generating process compatible with $\mathcal{M}^{\mathcal{F}}$. The resulting statistical model $\mathcal{M}$, describing the possible observed data distributions, is non-parametric.

Step 5 - Assessment of identifiability. For the purposes of discussion, suppose that the unmeasured factors $U = (U_{L(0)}, U_{A(0)}, U_{L(1)}, U_{A(1)}, U_Y)$ are all independent (Figure A.1b). Even if this assumption held, there is not one set of covariates that simultaneously satisfy the back-door criterion for all intervention nodes. The baseline covariates $L(0)$ alone fail, because there is an unblocked back-door path from Y through $L(1)$ to $A(1)$. In other words, the effect of initiation at one month $A(1)$ on the outcome Y is confounded by time-updated covariates $L(1)$. The baseline and time-updated covariates $(L(0), L(1))$ jointly fail, because we are losing (blocking) the effect of early ART initiation $A(0)$ on the outcome Y that goes through the covariates $L(1)$. This challenge is generally known as *time-dependent confounding* [27, 31, 50]: time-varying covariates confound the effect of future exposures on the outcome, but are affected by past exposures.

To identify the effects of longitudinal interventions, we consider the problem sequentially. For each $A(k)$ in sequence, we ask if its effect on Y can be identified by conditioning on some subset of the observed past. This leads to the *sequential randomization assumption* [27]:

$$Y_{a(0)a(1)} \perp\!\!\!\perp A(0) \big| L(0) \text{ and } Y_{a(0)a(1)} \perp\!\!\!\perp A(1) \big| \big(L(0), A(0), L(1)\big)$$

In words, we assume the counterfactual outcome $Y_{a(0)a(1)}$ is independent from the intervention $A(k)$ at time k, given the observed past. With the sequential randomization assumption as well a longitudinal version of the positivity assumption, the expectation of the counterfactual outcome, indexed by multiple interventions, can be identified by the longitudinal G-Computation formula [27]:

$$E_{U,X,0}\big(Y_{a(0)a(1)}\big) = \sum_{l_0, l_1} \bigg[E_0\big(Y | A(1) = a(1), L(1) = l(1), A(0) = a(0), L(0) = l(0)\big)$$

$$\times P_0\big(L(1) = l(1) | A(0) = a(0), L(0) = l(0)\big) \times P_0\big(L(0) = l(0)\big) \bigg] = \Psi(P_0)$$

Now we are averaging with respect to the appropriate distribution of covariates and thereby capturing the effect of both exposures $(a(1), a(0))$ on the outcome Y through the covariates $(L(0), L(1))$.

Step 6 - Estimation and Inference: As with single time-point interventions, there are a variety of methods to estimate statistical parameters, corresponding under the necessary assumptions to longitudinal causal effects. Examples include longitudinal IPTW, "parametric G-computation" (maximum likelihood estimation of the longitudinal G-computation formula) and TMLE [31, 55, 57–69].

Step 7 - Interpretation of the Results: As with the single time-point setting, the strength of our interpretations depend on rigorous evaluation of the needed assumptions. Even when the identifiability assumptions do not hold, we always have a statistical interpretation of $\Psi(P_0)$.

Appendix B

Technical Appendix for Chapter 2

B.1 The unadjusted estimator is unbiased

Recall the unadjusted estimator is the average difference in the outcomes within matched pairs:

$$\hat{\psi}_{unadj} = \frac{1}{n/2} \sum_{j=1}^{n/2} \left[A_{j1}Y_{j1} - (1 - A_{j1})Y_{j1} + A_{j2}Y_{j2} - (1 - A_{j2})Y_{j2} \right]$$

If observations within matched pairs have been ordered such that the first corresponds to treatment and the second to the control, the estimator can be expressed $\frac{1}{n/2} \sum_{j=1}^{n/2} (Y_{j1} - Y_{j2})$. Given the vector of covariates $W^n = (W_1, \ldots, W_n)$, the unadjusted estimator is unbiased for the statistical estimand:

$$
\begin{aligned}
E_0\left[\hat{\psi}_{unadj} \mid W^n\right] = & \frac{1}{n/2} \sum_{j=1}^{n/2} \left[E_0[A_{j1}Y_{j1} \mid W^n] - E_0[(1 - A_{j1})Y_{j1} \mid W^n] \right. \\
& \left. + E_0[A_{j2}Y_{j2} \mid W^n] - E_0[(1 - A_{j2})Y_{j2} \mid W^n] \right] \\
= & \frac{1}{n/2} \sum_{j=1}^{n/2} \left[\bar{Q}_0(1, W_{j1})E_0(A_{j1} \mid W^n) - \bar{Q}_0(0, W_{j1})E_0((1 - A_{j1}) \mid W^n) \right. \\
& \left. + \bar{Q}_0(1, W_{j2})E_0(A_{j2} \mid W^n) - \bar{Q}_0(0, W_{j2})E_0((1 - A_{j2}) \mid W^n) \right] \\
= & \frac{1}{n/2} \sum_{j=1}^{n/2} \frac{1}{2} \left[\bar{Q}_0(1, W_{j1}) - \bar{Q}_0(0, W_{j2}) + \bar{Q}_0(1, W_{j2}) - \bar{Q}_0(0, W_{j2}) \right] \\
= & \frac{1}{n} \sum_{i=1}^{n} \bar{Q}_0(1, W_i) - \bar{Q}_0(0, W_i) = \Psi(P_0^n).
\end{aligned}
$$

Thus, $\hat{\psi}_{unadj}$ is an unbiased estimator of $\Psi(P_0^n)$, conditional on W^n.

B.2 Statistical inference for the TMLE

In this subsection, we establish that the proposed TMLE is an asymptotically linear estimator of the conditional average treatment effect (CATE) in an adaptive pair-matched trial, where $n/2$ matched pairs are created as a function of baseline covariates of n candidate units. We then consider the adaptive design, where $n/2$ matched pairs are created as function of the baseline covariates of $N > n$ candidate units and the remaining $(N - n)$ units discarded. The latter adaptive design is a generalization of the first and the derived theorems are applicable. The theoretical results also apply to the unadjusted estimator $\hat{\psi}_{unadj}$, which can be considered a special case.

Let P_0^n denote the conditional distribution of $O^n = (O_1, \dots, O_n)$, given the vector of covariates $W^n = (W_1, \dots, W_n)$. The statistical estimand is a function of this conditional distribution:

$$\Psi(P_0^n) = \frac{1}{n} \sum_{i=1}^n \bar{Q}_0(1, W_i) - \bar{Q}_0(0, W_i),$$

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

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

where $\bar{Q}_n^*(A, W)$ denotes targeted estimates of the conditional mean function $\bar{Q}_0(A, W)$. Let ψ_0 denote the true parameter value and ψ_n^* denote the estimate.

Let us define the following function of $O = (W, A, Y)$:

$$D^*(\bar{Q}, g_0)(O) \equiv \left(\frac{\mathbb{I}(A = 1)}{g_0(A)} - \frac{\mathbb{I}(A = 0)}{g_0(A)} \right) (Y - \bar{Q}(A, W)),$$

where the marginal probability of receiving the intervention or the control is $g_0(A) = P_0(A) = 0.5$ in a randomized trial with two arms. By construction, TMLE solves $D^*(\bar{Q}, g_0)(O)$ at the targeted update $\bar{Q}_n^*$:

$$P_n D^*(\bar{Q}_n^*, g_0) = \frac{1}{n} \sum_{i=1}^n D^*(\bar{Q}_n^*, g_0)(O_i) = 0,$$

where P_n denotes the empirical distribution, placing mass $(1/n)$ on each $O_i, \; i = 1, \dots, n$. It is of interest to note that this equality can be rewritten as

$$\frac{1}{n/2} \sum_{j=1}^{n/2} \left\{ \bar{Q}_n^*(1, W_{j1}) - \bar{Q}_n^*(0, W_{j2}) \right\} = \frac{1}{n/2} \sum_{j=1}^{n/2} \left\{ Y_{j1} - Y_{j2} \right\},$$

where observations in pair j have again been ordered such that the first corresponds to the intervention $A_{j1} = 1$ and the second to the control $A_{j2} = 0$. Thus, the TMLE has the

interesting property that if it is used to predict the counterfactual effect $Y(1) - Y(0)$ for each pair j, then the average of these j-specific effects equals the unadjusted estimator.

Let $P_0^n f = E[f(O^n) \mid W^n]$ denote the conditional expectation of a function f of the data O^n, given the covariate vector W^n. For all $\bar{Q}(A, W)$, we have

$$P_0^n D^*(\bar{Q}, g_0) = \big(\bar{Q}_0(1, W_i) - \bar{Q}_0(0, W_i)\big) - \big(\bar{Q}(1, W_i) - \bar{Q}(0, W_i)\big).$$

Therefore, the statistical estimand $\Psi(P_0^n)$ minus the TMLE $\Psi(\bar{Q}_n^*)$ can be written as the empirical mean of the above conditional expectation:

$$\Psi(P_0^n) - \Psi(\bar{Q}_n^*) = \frac{1}{n} \sum_{i=1}^{n} P_0^n D^*(\bar{Q}_n^*, g_0).$$

Combining the latter equality with $P_n D^*(\bar{Q}_n^*, g_0) = 0$ yields

$$(\psi_n^* - \psi_0) = P_n \left\{ D^*(\bar{Q}_n^*, g_0) - P_0^n D^*(\bar{Q}_n^*, g_0) \right\}$$

$$= \frac{1}{n} \sum_{i=1}^{n} \left\{ D^*(\bar{Q}_n^*, g_0)(O_i) - P_0^n D^*(\bar{Q}_n^*, g_0) \right\}.$$

We can re-write this equality in terms of the empirical distribution $P_{n/2}$, which puts mass $1/(n/2)$ on each paired data point $\bar{O}_j = (O_{j1}, O_{j2})$:

$$(\psi_n^* - \psi_0) = P_{n/2} \left\{ \bar{D}^*(\bar{Q}_n^*, g_0) - P_0^n \bar{D}^*(\bar{Q}_n^*, g_0) \right\}$$

$$= \frac{1}{n/2} \sum_{j=1}^{n/2} \left\{ \bar{D}^*(\bar{Q}_n^*, g_0)(\bar{O}_j) - P_0^n \bar{D}^*(\bar{Q}_n^*, g_0) \right\}$$

$$\text{where } \bar{D}^*(\bar{Q}_n^*, g_0)(\bar{O}_j) = \frac{1}{2} \left\{ D^*(\bar{Q}_n^*, g_0)(O_{j1}) + D^*(\bar{Q}_n^*, g_0)(O_{j2}) \right\}$$

Now let $\mathcal{F}$ be a set of multivariate real valued functions so that $\bar{Q}_n^*(A, W)$ is an element of $\mathcal{F}$ with probability 1. Define the process $(Z_n(\bar{Q}) : \bar{Q} \in \mathcal{F})$ by

$$Z_n(\bar{Q}) = \frac{1}{\sqrt{n/2}} \sum_{j=1}^{n/2} \left\{ \bar{D}^*(\bar{Q}, g_0)(\bar{O}_j) - P_0^n \bar{D}^*(\bar{Q}, g_0) \right\}$$

Conditional on the covariate vector $W^n = (W_1, \ldots, W_n)$, $Z_n(\bar{Q})$ is a sum of $n/2$ independent mean zero random variables $\bar{D}^*(\bar{Q}, g_0)(\bar{O}_j) - P_0^n \bar{D}^*(\bar{Q}, g_0)$, $j = 1, \ldots, n/2$. Below we establish asymptotic equicontinuity of $(Z_n(\bar{Q}) : \bar{Q} \in \mathcal{F})$ so that $Z_n(\bar{Q}_n^*) - Z_n(\bar{Q}) \to 0$ in probability. Then, we can conclude that

$$\sqrt{n/2}(\psi_n^* - \psi_0) = \frac{1}{\sqrt{n/2}} \sum_{j=1}^{n/2} \left\{ \bar{D}^*(\bar{Q}, g_0)(\bar{O}_j) - P_0^n \bar{D}^*(\bar{Q}, g_0) \right\} + o_P(1).$$

Since the main term on the right-hand side, conditional on W^n, is a sum of independent mean zero random variables, we can apply the central limit theorem for sums of independent random variables.

Let us define the following function of the paired data $\bar{O}_j = (O_{j1}, O_{j2})$:

$$IC_j(\bar{Q}, \bar{Q}_0, g_0) \equiv \bar{D}^*(\bar{Q}, g_0)(\bar{O}_j) - P_0^n \bar{D}^*(\bar{Q}, g_0),$$

where the notation recognizes that $P_0^n \bar{D}^*(\bar{Q}, g_0)$ also depends on the true conditional mean $\bar{Q}_0(A, W) = E_0(Y|A, W)$. We assume that

$$\Sigma_0 = \lim_{n \to \infty} \frac{1}{n/2} \sum_{j=1}^{n/2} P_0^n IC_j(\bar{Q}, \bar{Q}_0, g_0)^2$$

exists as a limit. Then, we have shown $\sqrt{n/2}(\psi_n^* - \psi_0) \Rightarrow_d N(0, \Sigma_0)$.

To establish the asymptotic equicontinuity result, we use a few fundamental building blocks. Let $\mathcal{F}_d = \{f_1 - f_2 : f_1, f_2 \in \mathcal{F}\}$. Let $\sigma_n^2(f) = P_0^n Z_n(f)^2$ be the conditional variance. Note that $Z_n(f)/\sigma_n(f)$ is a sum of $n/2$ independent mean zero bounded random variables and the variance of this sum equals 1. Bernstein's inequality states that $P(|\sum_j Y_j| > x) \le 2\exp\left(-\frac{1}{2}\frac{x^2}{v+Mx/3}\right)$, where $v \ge \mathrm{VAR}\sum_j Y_j$. Thus, by Bernstein's inequality, conditional on W^n, we have

$$P\left(\frac{|Z_n(f)|}{\sigma_n(f)} > x\right) \le 2\exp\left(-\frac{1}{2}\frac{x^2}{1+Mx/3}\right) \le K\exp(-Cx^2),$$

for a universal K and C. This implies $\|Z_n(f)/\sigma_n(f)\|_{\psi_2} \le (1 + K/C)^{0.5}$, where for a given convex function ψ with $\psi(0) = 0$, $\|X\|_\psi \equiv \inf\{C > 0 : E\psi(|X|/C) \le 1\}$ is the so called Orlics norm, and $\psi_2(x) = \exp(x^2) - 1$. Thus $\|Z_n(f)\|_{\psi_2} \le C_1\sigma_n(f)$ for $f \in \mathcal{F}^d$. This result allows us to apply Theorem 2.2.4 in van der Vaart and Wellner [163]: for each $\delta > 0$ and $\eta > 0$, we now have

$$\|\sup_{\sigma_n(f) \le \delta} |Z_n(f)|\|_{\psi_2} \le K\left\{\int_0^\eta \psi_2^{-1}(N(\epsilon, \sigma_n, \mathcal{F}_d))d\epsilon + \delta\psi_2^{-1}(N^2(\eta, \sigma_n, \mathcal{F}_d)), \right\}, \qquad \text{(B.1)}$$

where $N(\epsilon, \sigma_n, \mathcal{F}_d)$ is the number of balls of size ϵ w.r.t. norm $\|f\| = \sigma_n(f)$ to cover $\mathcal{F}_d$.

Convergence of a sequence of random variables to zero with respect to ψ_2-orlics norm implies convergence in expectation to zero and thereby convergence of that sequence of random variables to zero in probability. Let δ_n be a sequence converging to zero, and let η_n also converge to zero but slowly enough so that the term $\delta_n\psi_2^{-1}(N^2(\eta_n, \sigma_n, \mathcal{F}^d))$ converges to zero as $n \to \infty$. By assumption, $\int_0^{\delta_n} \psi_2^{-1}(N(\epsilon, \sigma_n, \mathcal{F}^d)d\epsilon$ converges to zero. Thus,

$$\lim_{\delta_n \to 0}\left\{\int_0^{\delta_n} \psi_2^{-1}(N(\epsilon, \sigma_n, \mathcal{F}^d)d\epsilon + \delta_n\psi_2^{-1}(N^2(\eta_n, \sigma_n, \mathcal{F}^d))\right\} = 0.$$

This proves that

$$E\left(\sup_{\{f:\sigma_n(f)\leq\delta_n\}} \mid Z_n(f)\mid\right) \to 0.$$

Thus, if $\sigma_n(\bar{Q}_n^* - \bar{Q}) \to 0$ in probability, then $Z_n(\bar{Q}_n^* - \bar{Q}) \to 0$ in probability. This proves the following theorem.

Theorem 1. *Consider the TMLE $\Psi(\bar{Q}_n^*)$ of the statistical estimand $\Psi(P_0^n) = 1/n \sum_{i=1}^{n}\{\bar{Q}_0(1, W_i) - \bar{Q}_0(0, W_i)\}$. Let $P_0^n f$ represent the conditional expectation of a function f of O^n, given the vector of covariates W^n. This conditional expectation, $P_0^n f$, is thus still random through W^n. Let $\mathcal{F}$ be a set of multivariate real valued functions so that $\bar{Q}_n^*$ is an element of $\mathcal{F}$ with probability 1. Define*

$$Z_n(\bar{Q}) = \frac{1}{\sqrt{n/2}} \sum_{j=1}^{n/2} IC_j(\bar{Q}, \bar{Q}_0, g_0)$$

where

$$IC_j(\bar{Q}, \bar{Q}_0, g_0) \equiv \bar{D}^*(\bar{Q}, g_0)(\bar{O}_j) - P_0^n \bar{D}^*(\bar{Q}, g_0)$$

$$\bar{D}^*(\bar{Q}, g_0)(\bar{O}_j) = \frac{1}{2}\left\{D^*(\bar{Q}, g_0)(O_{j1}) + D^*(\bar{Q}, g_0)(O_{j2})\right\}$$

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

where $g_0(A) = P_0(A)$ is known. We make the following assumptions.

Uniform bound: *Assume $\sup_{\bar{Q}\in\mathcal{F}} \sup_O \mid D^*(\bar{Q}, g_0) \mid < M < \infty$, where the second supremum is over a set that contains the support of each O_i.*

Convergence of variances: *Assume that for a specified $\{\sigma_0^2(\bar{Q}) : \bar{Q} \in \mathcal{F}\}$, for any $\bar{Q} \in \mathcal{F}$, $\frac{1}{n/2}\sum_{j=1}^{n/2} P_0^n IC_j(\bar{Q}, \bar{Q}_0, g_0)^2 \to \sigma_0^2(\bar{Q})$ a.s (i.e, for almost every $(W^n, n \geq 1)$).*

Convergence of $\bar{Q}_n^*$ to some limit: *For any $\bar{Q}_1, \bar{Q}_2 \in \mathcal{F}$, we define*

$$\sigma_n^2(\bar{Q}_1 - \bar{Q}_2) = \frac{1}{n/2}\sum_{j=1}^{n/2} P_0^n\{IC_j(\bar{Q}_1, \bar{Q}_0, g_0) - IC_j(\bar{Q}_2, \bar{Q}_0, g_0)\}^2,$$

where we note that the right-hand side indeed only depends on $\bar{Q}_1, \bar{Q}_2$ through its difference $\bar{Q}_1 - \bar{Q}_2$.

Assume that for a particular $\bar{Q}^* \in \mathcal{F}$, $\sigma_n^2(\bar{Q}_n^* - \bar{Q}^*) \to 0$ in probability as $n \to \infty$.

Entropy condition: *Let* $\mathcal{F}^d = \{f_1 - f_2 : f_1, f_2 \in \mathcal{F}\}$. *Let* $N(\epsilon, \sigma_n, \mathcal{F}^d)$ *be the covering number of the class* $\mathcal{F}^d$ *w.r.t norm/dissimilarity* $\| f \| = \sigma_n(f)$. *Assume that the class* $\mathcal{F}$ *satisfies*

$$\lim_{\delta_n \to 0} \int_0^{\delta_n} \sqrt{\log N(\epsilon, \sigma_n, \mathcal{F}^d)} d\epsilon = 0$$

Asymptotic equicontinuity of process: *Then,*

$$Z_n(\bar{Q}_n^*) - Z_n(\bar{Q}^*) \text{ converges to zero in probability, as } n \to \infty.$$

First order linear approximation: *As a consequence,*

$$\sqrt{n/2}(\psi_n^* - \psi_0) = Z_n(\bar{Q}^*) + o_P(1).$$

Asymptotic normality: *In addition,* $Z_n(\bar{Q}^*)$ *converges to* $N(0, \sigma_0^2(\bar{Q}^*))$, *so that*

$$\sqrt{n/2}(\psi_n^* - \psi_0) \text{ converges in distribution to } N(0, \sigma_0^2(\bar{Q}^*)).$$

The asymptotic variance $\sigma_0^2(\bar{Q}^*)$ *equals the limit of*

$$\sigma_{0,n}^2 = \frac{1}{n/2} \sum_{j=1}^{n/2} P_0^n \left\{ IC_j(\bar{Q}_n^*, \bar{Q}_0, g_0) \right\}^2$$

If Y_i is d-dimensional outcome, then the application of the above theorem to each component of ψ_n^* yields the desired asymptotic linearity for the d-dimensional ψ_n^* and thereby the asymptotic normality as well.

Conservative variance estimation

The above result suggests the following estimator of the asymptotic variance of the standardized TMLE:

$$\hat{\Sigma} = \frac{1}{n/2} \sum_{j=1}^{n/2} \left\{ IC_j(\bar{Q}_n^*, \bar{Q}_{n,np}, g_0)(\bar{O}_j) \right\}^2$$

where $\bar{Q}_{n,np}$ is a consistent estimator of $\bar{Q}_0$. Unfortunately, such a variance estimator relies upon consistent estimation of the conditional mean function $\bar{Q}_0$, which is particular concerning when n is small. However, we will now show that one can obtain a conservative variance estimate, which does not rely on a consistent estimator of the conditional mean function $\bar{Q}_0$.

The asymptotic variance of the standardized estimator $\sqrt{n/2}(\psi_n^* - \psi_0)$ can be expressed as

$$\Sigma_0 = \lim_{n \to \infty} \frac{1}{n/2} \sum_{j=1}^{n/2} P_0^n \left[IC_j(\bar{Q}^*, \bar{Q}_0, g_0) \right]^2$$

$$= \lim_{n \to \infty} \frac{1}{n/2} \sum_{j=1}^{n/2} P_0^n \left[\bar{D}^*(\bar{Q}^*, g_0) \right]^2 - \left[P_0^n \bar{D}^*(\bar{Q}^*, g_0) \right]^2.$$

The latter term is zero when $\bar{Q}^*(A, W) = \bar{Q}_0(A, W)$:

$$P_0^n \bar{D}^*(\bar{Q}^*, g_0) = \frac{1}{2} \left\{ \bar{Q}_0(1, W_{j1}) - \bar{Q}_0(0, W_{j1}) - \left(\bar{Q}^*(1, W_{j1}) - \bar{Q}^*(0, W_{j1}) \right) \right.$$

$$\left. + \bar{Q}_0(1, W_{j2}) - \bar{Q}_0(0, W_{j2}) - \left(\bar{Q}^*(1, W_{j2}) - \bar{Q}^*(0, W_{j2}) \right) \right\}$$

Thus, the true variance Σ_0 is always less than or equal to an upper bound Σ_0^u, where

$$\Sigma_0^u = \lim_{n \to \infty} \frac{1}{n/2} \sum_{j=1}^{n/2} P_0^n \left\{ \bar{D}^*(\bar{Q}^*, g_0) \right\}^2$$

Again, if the conditional mean is consistently estimated $\bar{Q}^*(A, W) = \bar{Q}_0(A, W)$, $\Sigma_0^u = \Sigma_0$.
We can consistently estimate the upper bound Σ_0^u with

$$\hat{\Sigma}^u = \frac{1}{n/2} \sum_{j=1}^{n/2} \left\{ \bar{D}^*(\bar{Q}_n^*, g_0)(\bar{O}_j) \right\}^2$$

Recall

$$\bar{D}^*(\bar{Q}, g_0)(\bar{O}_j) = \frac{1}{2} \left\{ D^*(\bar{Q}, g_0)(O_{j1}) + D^*(\bar{Q}, g_0)(O_{j2}) \right\}$$

$$= \frac{1}{2} \left[\left(\frac{\mathbb{I}(A_{j1} = 1)}{g_0(A_{j1})} - \frac{\mathbb{I}(A_{j1} = 0)}{g_0(A_{j1})} \right) \left(Y_{j1} - \bar{Q}(A_{j1}, W_{j1}) \right) \right.$$

$$\left. + \left(\frac{\mathbb{I}(A_{j2} = 1)}{g_0(A_{j2})} - \frac{\mathbb{I}(A_{j2} = 0)}{g_0(A_{j2})} \right) \left(Y_{j2} - \bar{Q}(A_{j2}, W_{j2}) \right) \right]$$

Ordering the observations within pairs, such that index $j1$ corresponds to the unit randomized to the intervention ($A_{j1} = 1$) and $j2$ corresponds to the unit randomized to the control ($A_{j2} = 0$), it follows that

$$\bar{D}^*(\bar{Q}, g_0)(\bar{O}_j) = Y_{j_1} - \bar{Q}(1, W_{j1}) - \left(Y_{j_2} - \bar{Q}(0, W_{j2}) \right),$$

allowing us to represent the conservative variance estimator $\hat{\Sigma}^u$ as the difference in residuals within matched pairs:

$$\hat{\Sigma}^u = \frac{1}{n/2} \sum_{j=1}^{n/2} \left\{ Y_{j_1} - \bar{Q}_n^*(1, W_{j1}) - \left(Y_{j_2} - \bar{Q}_n^*(0, W_{j2}) \right) \right\}^2.$$

Generalization to $N > n$ candidate units

Now consider the common adaptive design, where first N candidate units are selected, the best $n/2$ matched pairs selected as a function of the covariate vector $W^N = (W_1, \ldots, W_n, \ldots, W_N)$, and the remaining $N - n$ units discarded. In the SEARCH trial, for example, 16 matched pairs were formed as a function of the baseline covariates of 54 candidate communities. As a result of this adaptive design, the treatment assignment mechanism depends on the N candidate communities. Nonetheless, in a randomized trial, the conditional likelihood of the observed data factorizes as

$$P_0(O_1, \ldots, O_n | W_1, \ldots, W_N) = \prod_{j=1}^{n} g_0(A_{j1}, A_{j2} | W_1, \ldots, W_N) P_0(Y_{j1} | A_{j1}, W_{j1}) P_0(Y_{j2} | A_{j2}, W_{j2})$$

$$= 0.5 \prod_{j=1}^{n} P_0(Y_{j1} | A_{j1}, W_{j1}) P_0(Y_{j2} | A_{j2}, W_{j2})$$

$$= P_0(O_1, \ldots, O_n \mid W_1, \ldots, W_n) = P_0^n(O^n | W^n)$$

Therefore, given the baseline covariates of the n study units $W^n = (W_1, \ldots W_n)$, we still have $n/2$ conditionally independent observations. Furthermore, recall that the statistical estimand corresponds to the average treatment effect, conditional on the baseline covariates of the n study units:

$$\Psi(P_0^n) = \frac{1}{n} \sum_{i=1}^{n} \bar{Q}_0(1, W_i) - \bar{Q}_0(0, W_i)$$

Since we condition on $W^n = (W_1, \ldots, W_n)$ in the target parameter and corresponding TMLE, the actual distribution that generated these n covariates is not important. Recall we make no assumptions about the joint distribution of $P_0(W^N)$. We only need to assume that the conditional variance still converges. As a result, we can apply the same TMLE and asymptotics. As detailed in van der Laan *et al.* [115], this is a much different result than when the target parameter is the marginal (population) average treatment effect. In the latter case, the so-called adaptive missingness has important implications for estimation and inference to a target population of units.

B.3 Comparison with a non-matched trial (i.e. a completely randomized trial)

In this section, we consider estimation and inference for the conditional average treatment effect (CATE) in a trial, where the intervention is completely randomized. We consider implementation of the TMLE and the corresponding asymptotics. We conclude with an efficiency comparison between a trial within adaptive pair-matching and a trial without pair-matching.

TMLE for the CATE in a non-matched trial

Let $\bar{Q}_n(A, W)$ be an initial estimator of $\bar{Q}_0(A, W)$, which can be obtained by regressing the outcome Y_i on exposure A_i and covariates W_i, $i = 1, \ldots, n$. For a binary or bounded continuous outcome, the negative log-likelihood is a valid loss function:

$$-L(\bar{Q})(O) = Y \log \bar{Q}(A, W) + (1 - Y) \log(1 - \bar{Q}(A, W))$$

Now consider the logistic fluctuation submodel:

$$\text{logit}\left[\bar{Q}_n(A, W)(\epsilon)\right] = \text{logit}\left[\bar{Q}_n(A, W)\right] + \epsilon H(A)$$
$$\text{where } H(A) = \left(\frac{\mathbb{I}(A = 1)}{g_0(A)} - \frac{\mathbb{I}(A = 0)}{g_0(A)}\right)$$

In a randomized trial with two arms, the probability of receiving the intervention or control is $g_0(A = a) = P_0(A = a) = 0.5$. Let ϵ_n be the minimizer of the empirical mean of the loss function:

$$\epsilon_n = \arg\min_{\epsilon} P_n L\left(\bar{Q}_n(A, W)(\epsilon)\right) = \frac{1}{n} \sum_{i=1}^{n} L\left(\bar{Q}_n(A, W)(\epsilon)\right)(O_i)$$

The TMLE of the conditional mean outcome $\bar{Q}_0(A, W)$ is defined by plugging in the estimated coefficient ϵ_n into the fluctuation model $\bar{Q}_n^*(A, W) = \bar{Q}_n(A, W)(\epsilon_n)$. The TMLE of $\Psi(P_0^n)$ is defined as the corresponding plug-in estimator:

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

As before, initial estimation of the conditional mean function $\bar{Q}_0(A, W)$ can also be based on least squares regression and targeting achieved with the following fluctuation submodel:

$$\bar{Q}_n(A, W)(\epsilon) = \bar{Q}_n(A, W) + \epsilon H(A)$$

Recall the definition of $D^*(\bar{Q}, g_0)(O)$ as the following function of the observed data $O = (W, A, Y)$:

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

where the probability of receiving the intervention or the control is $g_0(A) = P_0(A) = 0.5$ in a randomized trial. By construction, TMLE solves $D^*(\bar{Q}, g_0)(O)$ at the targeted update $\bar{Q}_n^*$:

$$P_n D^*(\bar{Q}_n^*, g_0) = \frac{1}{n} \sum_{i=1}^{n} D^*(\bar{Q}_n^*, g_0)(O_i) = 0$$

where P_n denotes the empirical distribution, placing mass $(1/n)$ on each O_i, $i = 1, \ldots, n$. For all $\bar{Q}(A, W)$, we also have

$$P_0^n D^*(\bar{Q}, g_0) = \big(\bar{Q}_0(1, W_i) - \bar{Q}_0(0, W_i)\big) - \big(\bar{Q}(1, W_i) - \bar{Q}(0, W_i)\big),$$

where $P_0^n f = E\big[f(O^n)|W^n\big]$ denotes the conditional expectation of the function f of the data O^n, given the covariate vector W^n. Therefore, the statistical estimand $\Psi(P_0^n)$ minus the TMLE $\Psi(\bar{Q}_n^*)$ can be written as the empirical mean of the above conditional expectation:

$$\Psi(P_0^n) - \Psi(\bar{Q}_n^*) = \frac{1}{n} \sum_{i=1}^n P_0^n D^*(\bar{Q}_n^*, g_0).$$

Combining the latter equality with $P_n D^*(\bar{Q}_n^*, g_0) = 0$ yields

$$\sqrt{n}(\psi_n^* - \psi_0) = \frac{1}{\sqrt{n}} \sum_{i=1}^n \left\{ D^*(\bar{Q}_n^*, g_0)(O_i) - P_0^n D^*(\bar{Q}_n^*, g_0) \right\}.$$

Recall $\mathcal{F}$ is the set of multivariate real-valued functions such that $\bar{Q}_n^*(A, W)$ is an element of $\mathcal{F}$ with probability 1. Define the process $(Z_n(\bar{Q}) : \bar{Q} \in \mathcal{F})$ by

$$Z_n(\bar{Q}) = \frac{1}{\sqrt{n}} \sum_{i=1}^n \left\{ D^*(\bar{Q}, g_0)(O_i) - P_0^n D^*(\bar{Q}, g_0) \right\},$$

Conditional on the covariate vector $W^n = (W_1, \ldots, W_n)$, $Z_n(\bar{Q})$ is a sum of n independent mean zero random variables $D^*(\bar{Q}, g_0)(O_i) - P_0^n D^*(\bar{Q}, g_0)$, $i = 1, \ldots, n$. Below we establish asymptotic equicontinuity of $(Z_n(\bar{Q}) : \bar{Q} \in \mathcal{F})$ so that $Z_n(\bar{Q}_n^*) - Z_n(\bar{Q}) \to 0$ in probability. Then, we can conclude that

$$\sqrt{n}(\psi_n^* - \psi_0) = \frac{1}{\sqrt{n}} \sum_{i=1}^n \left\{ D^*(\bar{Q}, g_0)(O_i) - P_0^n D^*(\bar{Q}, g_0) \right\} + o_P(1).$$

Since the main term on the right-hand side, conditional on W^n, is a sum of independent mean zero random variables, we can apply the central limit theorem for sums of independent random variables.

For a completely randomized trial, let us define the following function of the unit data $O_i = (W_i, A_i, Y_i)$:

$$IC_i(\bar{Q}, \bar{Q}_0, g_0) \equiv D^*(\bar{Q}, g_0)(O_i) - P_0^n D^*(\bar{Q}, g_0),$$

where the notation recognizes that $P_0^n D^*(\bar{Q}, g_0)$ also depends on the true conditional mean $\bar{Q}_0(A, W) = E_0(Y|A, W)$. We assume that

$$\Sigma_0 = \lim_{n \to \infty} \frac{1}{n} \sum_{i=1}^n P_0^n IC_i(\bar{Q}, \bar{Q}_0, g_0)^2$$

exists as a limit. Then, we have shown $\sqrt{n}(\psi_n^* - \psi_0) \Rightarrow_d N(0, \Sigma_0)$.

To establish the asymptotic equicontinuity result, we use a few fundamental building blocks. Let $\mathcal{F}_d = \{f_1 - f_2 : f_1, f_2 \in \mathcal{F}\}$. Let $\sigma_n^2(f) = P_0^n Z_n(f)^2$ be the conditional variance. Note that $Z_n(f)/\sigma_n(f)$ is a sum of n independent mean zero bounded random variables and the variance of this sum equals 1. Bernstein's inequality states that $P(|\sum_j Y_j| > x) \leq 2\exp\left(-\frac{1}{2}\frac{x^2}{v+Mx/3}\right)$, where $v \geq \text{VAR}\sum_j Y_j$. Thus, by Bernstein's inequality, conditional on W^n, we have

$$P\left(\frac{|Z_n(f)|}{\sigma_n(f)} > x\right) \leq 2\exp\left(-\frac{1}{2}\frac{x^2}{1+Mx/3}\right) \leq K\exp(-Cx^2),$$

for a universal K and C. This implies $\|Z_n(f)/\sigma_n(f)\|_{\psi_2} \leq (1+K/C)^{0.5}$, where for a given convex function ψ with $\psi(0) = 0$, $\|X\|_{\psi} \equiv \inf\{C > 0 : E\psi(|X|/C) \leq 1\}$ is the so called Orlics norm, and $\psi_2(x) = \exp(x^2) - 1$. Thus $\|Z_n(f)\|_{\psi_2} \leq C_1\sigma_n(f)$ for $f \in \mathcal{F}^d$. This result allows us to apply Theorem 2.2.4 in van der Vaart and Wellner [163]: for each $\delta > 0$ and $\eta > 0$, we now have

$$\|\sup_{\sigma_n(f)\leq\delta}|Z_n(f)|\|_{\psi_2} \leq K\left\{\int_0^\eta \psi_2^{-1}(N(\epsilon, \sigma_n, \mathcal{F}_d)d\epsilon + \delta\psi_2^{-1}(N^2(\eta, \sigma_n, \mathcal{F}_d)),\right\}, \qquad \text{(B.2)}$$

where $N(\epsilon, \sigma_n, \mathcal{F}_d)$ is the number of balls of size ϵ w.r.t. norm $\|f\| = \sigma_n(f)$ to cover $\mathcal{F}_d$.

Convergence of a sequence of random variables to zero with respect to ψ_2-orlics norm implies convergence in expectation to zero and thereby convergence of that sequence of random variables to zero in probability. Let δ_n be a sequence converging to zero, and let η_n also converge to zero but slowly enough so that the term $\delta_n\psi_2^{-1}(N^2(\eta_n, \sigma_n, \mathcal{F}^d))$ converges to zero as $n \to \infty$. By assumption, $\int_0^{\delta_n} \psi_2^{-1}(N(\epsilon, \sigma_n, \mathcal{F}^d)d\epsilon$ converges to zero. Thus,

$$\lim_{\delta_n\to 0}\left\{\int_0^{\delta_n} \psi_2^{-1}(N(\epsilon, \sigma_n, \mathcal{F}^d)d\epsilon + \delta_n\psi_2^{-1}(N^2(\eta_n, \sigma_n, \mathcal{F}^d))\right\} = 0.$$

This proves that

$$E\left(\sup_{\{f:\sigma_n(f)\leq\delta_n\}}|Z_n(f)|\right) \to 0.$$

Thus, if $\sigma_n(\bar{Q}_n^* - \bar{Q}) \to 0$ in probability, then $Z_n(\bar{Q}_n^* - \bar{Q}) \to 0$ in probability. This proves the following theorem.

Theorem 2. *Consider the TMLE $\Psi(\bar{Q}_n^*)$ for the statistical estimand $\Psi(P_0^n) = 1/n\sum_{i=1}^n\{\bar{Q}_0(1, W_i) - \bar{Q}_0(0, W_i)\}$ defined above for a trial without pair-matching (i.e. a trial with complete randomization). Let $P_0^n f$ represents a conditional expectation of a function f of O^n, given W^n. This conditional expectation is thus still random through W^n. Let $\mathcal{F}$ be a set of multivariate real valued functions so that $\bar{Q}_n^*$ is an element of $\mathcal{F}$ with probability 1. Define*

$$Z_n(\bar{Q}) = \frac{1}{\sqrt{n}}\sum_{i=1}^n IC_i(\bar{Q}, \bar{Q}_0, g_0),$$

where

$$IC_i(\bar{Q}, \bar{Q}_0, g_0) \equiv D^*(\bar{Q}, g_0)(O_i) - P_0^n D^*(\bar{Q}, g_0)$$

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

We make the following assumptions.

Uniform bound: *Assume* $\sup_{\bar{Q} \in \mathcal{F}} \sup_O \mid D^*(\bar{Q}, g_0) \mid < M < \infty$, *where the second supremum is over a set that contains the support of each O_i.*

Convergence of variances: *Assume that for a specified $\{\sigma_0^2(\bar{Q}) : \bar{Q} \in \mathcal{F}\}$, for any $\bar{Q} \in \mathcal{F}$, $\frac{1}{n} \sum_{i=1}^n P_0^n IC_i(\bar{Q}, \bar{Q}_0, g_0)^2 \to \sigma_0^2(\bar{Q})$ a.s (i.e, for almost every $(W^n, n \geq 1)$).*

Convergence of $\bar{Q}_n^*$ to some limit: *For any $\bar{Q}_1, \bar{Q}_2 \in \mathcal{F}$, we define*

$$\sigma_n^2(\bar{Q}_1 - \bar{Q}_2) = \frac{1}{n} \sum_{i=1}^n P_0^n \big\{ IC_i(\bar{Q}_1, \bar{Q}_0, g_0) - IC_i(\bar{Q}_2, \bar{Q}_0, g_0) \big\}^2,$$

where we note that the right-hand side indeed only depends on $\bar{Q}_1, \bar{Q}_2$ through its difference $\bar{Q}_1 - \bar{Q}_2$.

Assume that for a particular $\bar{Q}^* \in \mathcal{F}$, $\sigma_n^2(\bar{Q}_n^* - \bar{Q}^*) \to 0$ in probability as $n \to \infty$.

Entropy condition: *Let $\mathcal{F}^d = \{f_1 - f_2 : f_1, f_2 \in \mathcal{F}\}$. Let $N(\epsilon, \sigma_n, \mathcal{F}^d)$ be the covering number of the class $\mathcal{F}^d$ w.r.t norm/dissimilarity $\| f \| = \sigma_n(f)$. Assume that the class $\mathcal{F}$ satisfies*

$$\lim_{\delta_n \to 0} \int_0^{\delta_n} \sqrt{\log N(\epsilon, \sigma_n, \mathcal{F}^d)} d\epsilon = 0$$

Asymptotic equicontinuity of process: *Then,*

$$Z_n(\bar{Q}_n^*) - Z_n(\bar{Q}^*) \text{ converges to zero in probability, as } n \to \infty.$$

First order linear approximation: *As a consequence,*

$$\sqrt{n}(\psi_n^* - \psi_0) = Z_n(\bar{Q}^*) + o_P(1).$$

Asymptotic normality: *In addition, $Z_n(\bar{Q}^*)$ converges to $N(0, \sigma_0^2(\bar{Q}^*))$, so that*

$$\sqrt{n}(\psi_n^* - \psi_0) \text{ converges in distribution to } N(0, \sigma_0^2(\bar{Q}^*)).$$

The asymptotic variance $\sigma_0^2(\bar{Q}^*)$ equals the limit of

$$\sigma_{0,n}^2 = \frac{1}{n} \sum_{i=1}^n P_0^n \left\{ IC_i(\bar{Q}^*, \bar{Q}_0, g_0) \right\}^2. \tag{B.3}$$

If Y_i is a d-dimensional outcome, then application of the above theorem to each component of ψ_n^* yields the desired asymptotic linearity for the d-dimensional ψ_n^* and thereby the asymptotic normality as well.

Conservative variance estimation

As before, we can obtain a conservative variance estimator, which does not rely on a consistent estimator of the conditional mean function $\bar{Q}_0(A, W)$. The asymptotic variance of the standardized estimator in the design with complete randomization can be expressed as

$$\Sigma_0 = \lim_{n \to \infty} \frac{1}{n} \sum_{i=1}^{n} P_0^n \{D^*(\bar{Q}^*, g_0)\}^2 - \{P_0^n D^*(\bar{Q}^*, g_0)\}^2$$

The latter term is zero when $\bar{Q}^*(A, W) = \bar{Q}_0(A, W)$:

$$P_0^n D^*(\bar{Q}^*, g_0) = \big(\bar{Q}_0(1, W_i) - \bar{Q}_0(0, W_i)\big) - \big(\bar{Q}^*(1, W_i) - \bar{Q}^*(0, W_i)\big)$$

Thus, the true variance Σ_0 is always less than or equal to an upper bound Σ_0^u, where

$$\Sigma_0^u = \lim_{n \to \infty} \frac{1}{n} \sum_{i=1}^{n} P_0^n \{D^*(\bar{Q}^*, g_0)\}^2$$

We can consistently estimate the upper bound Σ_0^u with

$$\hat{\Sigma}^u = \frac{1}{n} \sum_{i=1}^{n} \big\{D^*(\bar{Q}_n^*, g_0)(O_i)\big\}^2$$

$$= \frac{4}{n} \sum_{i=1}^{n} \big(Y_i - \bar{Q}_n^*(A_i, W_i)\big)^2.$$

where we have used that the treatment assignment mechanism $g_0(A) = P_0(A) = 0.5$ in a randomized trial.

Comparison of asymptotic variances of the TMLEs in a trial without pair-matching (i.e. complete/independent randomization) and a trial with pair-matching

The above two theorems give us the following approximations for the TMLEs $\psi_{n,I}^*$ under independent randomization and $\psi_{n,M}^*$ under adaptive pair-matching:

$$\sqrt{n}(\psi_{n,I}^* - \psi_0) = \frac{1}{\sqrt{n}} \sum_{i=1}^{n} \left\{ D^*(\bar{Q}, g_0)(O_i) - P_0^n D^*(\bar{Q}, g_0) \right\} + o_P(1)$$

$$\sqrt{n/2}(\psi_{n,M}^* - \psi_0) = \frac{1}{\sqrt{n/2}} \sum_{j=1}^{n/2} \left\{ \bar{D}^*(\bar{Q}, g_0)(\bar{O}_j) - P_0^n \bar{D}^*(\bar{Q}, g_0) \right\} + o_P(1)$$

where

$$\bar{D}^*(\bar{Q}, g_0)(\bar{O}_j) = \frac{1}{2} \left\{ D^*(\bar{Q}, g_0)(O_{j1}) + D^*(\bar{Q}, g_0)(O_{j2}) \right\}$$

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

The corresponding asymptotic variances are

$$\Sigma_{0,I} = \lim_{n \to \infty} \frac{1}{n} \sum_{i=1}^{n} P_0^n \left[D^*(\bar{Q}, g_0)^2 \right] - \left[P_0^n D^*(\bar{Q}, g_0) \right]^2$$

$$\Sigma_{0,M} = \lim_{n \to \infty} \frac{1}{n/2} \sum_{j=1}^{n/2} P_0^n \left[\bar{D}^*(\bar{Q}^*, g_0)^2 \right] - \left[P_0^n \bar{D}^*(\bar{Q}^*, g_0) \right]^2,$$

respectively. Expanding out the squared terms and simplifying, the asymptotic variance of the standardized estimator in the independent design is

$$\Sigma_{0,I} = \lim_{n \to \infty} \frac{1}{n} \sum_{i=1}^{n} \left\{ 2E_0 \left[(Y_i - \bar{Q}_0(1, W_i))^2 \middle| A_i = 1, W^n \right] \right.$$

$$+ 2E_0 \left[(Y_i - \bar{Q}_0(0, W_i))^2 \middle| A_i = 0, W^n \right]$$

$$\left. + \left[\bar{Q}_0(1, W_i) - \bar{Q}(1, W_i) + \bar{Q}_0(0, W_i) - \bar{Q}(0, W_i) \right]^2 \right\}$$

Likewise, the asymptotic variance of the standardized estimator in the adaptive design is

$$\Sigma_{0,M} = \lim_{n \to \infty} \frac{1}{2n} \sum_{i=1}^{n} \left\{ 2E_0 \left[(Y_i - \bar{Q}_0(1, W_i))^2 \middle| A_i = 1, W^n \right] \right.$$

$$+ 2E_0 \left[(Y_i - \bar{Q}_0(0, W_i))^2 \middle| A_i = 0, W^n \right]$$

$$\left. + \left[\bar{Q}_0(1, W_i) - \bar{Q}(1, W_i) + \bar{Q}_0(0, W_i) - \bar{Q}(0, W_i) \right]^2 \right\} - \rho_0$$

$$= 0.5\Sigma_{0,I} - \rho_0$$

where ρ_0 is the following pairwise product

$$\rho_0 = \lim_{n \to \infty} \frac{1}{n} \sum_{j=1}^{n/2} \left\{ \left[\bar{Q}_0(1, W_{j1}) - \bar{Q}(1, W_{j1}) + \bar{Q}_0(0, W_{j1}) - \bar{Q}(0, W_{j1}) \right] \times \right.$$

$$\left. \left[\bar{Q}_0(1, W_{j2}) - \bar{Q}(1, W_{j2}) + \bar{Q}_0(0, W_{j2}) - \bar{Q}(0, W_{j2}) \right] \right\}$$

The proof is omitted here, but readily available upon request from the authors.

Thus, the asymptotic variance of the TMLE in the independent design is $\Sigma_{0,I}/n$ whereas the asymptotic variance of the TMLE in the adaptive design is $\Sigma_{0,M}/(n/2) = \Sigma_{0,I}/n - 2\rho_0/n$. When we match well on measured and unmeasured factors, the product of the deviations between the true conditional means and the estimated means within matched pairs is expected to be positive:

$$\rho_0 \geq 0$$

Under this condition, the adaptive design will be more efficient than the completely randomized trial. As an example, consider the unadjusted estimator and suppose we match perfectly on W, which is predictive of the outcome. Then the relevant term is

$$\left[\bar{Q}_0(1, W_j) - \bar{Q}_n(1) + \bar{Q}_0(0, W_j) - \bar{Q}_n(0)\right]^2 > 0$$

If we consistently estimate $\bar{Q}_0(A, W)$, then the cross-term ρ_0 is zero and the efficiency bound of the two designs is the same:

$$\Sigma_{0,M}/(n/2) = \Sigma_{0,I}/n$$

In finite samples, we also expect there to be an efficiency gain from pair-matching. Comparing the proposed variance estimators, we have

$$\hat{\Sigma}_M^u = \frac{1}{n/2} \sum_{j=1}^{n/2} \left[\left(Y_{j1} - \bar{Q}_n^*(1, W_{j1})\right) - \left(Y_{j2} - \bar{Q}_n^*(0, W_{j2})\right)\right]^2$$

$$= \frac{1}{n/2} \sum_{j=1}^{n/2} \left[\left(Y_{j1} - \bar{Q}_n^*(1, W_{j1})\right)^2 + \left(Y_{j2} - \bar{Q}_n^*(0, W_{j2})\right)^2\right.$$

$$\left. - 2\left(Y_{j1} - \bar{Q}_n^*(1, W_{j1})\right)\left(Y_{j2} - \bar{Q}_n^*(0, W_{j2})\right)\right]$$

$$\hat{\Sigma}_I^u = \frac{4}{n} \sum_{i=1}^{n} (Y_i - \bar{Q}_n^*(A_i, W_i))^2$$

$$= \frac{4}{n} \sum_{j=1}^{n/2} \left(Y_{j1} - \bar{Q}_n^*(1, W_{j1})\right)^2 + \left(Y_{j2} - \bar{Q}_n^*(1, W_{j2})\right)^2$$

Then, the difference is

$$\frac{\hat{\Sigma}_I^u}{n} - \frac{\hat{\Sigma}_M^u}{n/2} = \frac{2}{(n/2)^2} \sum_{j=1}^{n/2} \left(Y_{j1} - \bar{Q}_n^*(1, W_{j1})\right)\left(Y_{j2} - \bar{Q}_n^*(0, W_{j2})\right)$$

If we succeed in matching pairs on predictive covariates W, then the sample covariance of residuals within matched pairs will be positive. Under this condition (expected to hold in practice), adaptive pair-matching will yield more precise estimates in finite samples.

B.4 Simulation results under the null

The following table gives the simulation results when there is no effect. The null scenario was simulated by randomly assigning the intervention, but generating the outcomes under the control ($A = 0$). Recall Simulation A represents a rare outcome and Simulation B represents a more common outcome.

	Bias	Std. Dev.	Std. Error	t-stat	CI Cov.	α
Simulation A			No Matching			
Unadj.	0.00015	0.0061	0.0060	0.0	95	5
TMLE linear for Z	0.00001	0.0033	0.0032	0.0	94	6
TMLE logit for Z	0.00003	0.0032	0.0030	0.0	94	6
TMLE linear for (W, Z)	0.00003	0.0030	0.0026	0.0	91	9
TMLE logit for (W, Z)	0.00005	0.0029	0.0024	0.0	90	10
			Adaptive Pair-Matching			
Unadj.	0.00002	0.0034	0.0034	0.0	96	4
TMLE linear for Z	0.00005	0.0028	0.0028	0.0	95	5
TMLE logit for Z	0.00005	0.0027	0.0027	0.0	95	5
TMLE linear for (W, Z)	0.00005	0.0027	0.0026	0.0	94	6
TMLE logit for (W, Z)	0.00005	0.0027	0.0025	0.0	94	6
Simulation B			No Matching			
Unadj.	0.00017	0.0058	0.0057	0.0	95	5
TMLE linear for Z	0.00006	0.0035	0.0033	0.0	94	6
TMLE logit for Z	0.00007	0.0036	0.0035	0.0	94	6
TMLE linear for (W, Z)	0.00007	0.0031	0.0027	0.0	91	9
TMLE logit for (W, Z)	0.00007	0.0035	0.0030	0.0	91	9
			Adaptive Pair-Matching			
Unadj.	0.00002	0.0034	0.0033	0.0	95	5
TMLE linear for Z	0.00006	0.0029	0.0028	0.0	95	5
TMLE logit for Z	0.00005	0.0030	0.0029	0.0	95	5
TMLE linear for (W, Z)	0.00004	0.0028	0.0026	0.0	94	6
TMLE logit for (W, Z)	0.00004	0.0030	0.0028	0.0	94	6

Table B.1: For Simulation A (rare outcome) and Simulation B (more common outcome) with no treatment effect, summary of the estimator performance over 5,000 simulations of $n = 32$ communities. The rows indicate the estimator and the columns the performance metric: bias as the average deviation between the point estimate and sample-specific true value; standard deviation as the square root of the variance of the point estimates; standard error as the average standard error estimate based on the influence curve; t-statistic as the average value of the test statistic (point estimate divided by standard error estimate); confidence interval coverage as the proportion of intervals containing the true parameter value (in %), and the type I error (α) as the proportion of studies falsely rejecting the null hypothesis (in %)

Appendix C

Technical Appendix for Chapter 3

C.1 The TMLE is an asymptotically linear estimator of the SATE

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

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

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:

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

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 [40]:

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

with influence curve

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

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

$P_0\big[D_n^{\mathcal{P}}(\bar{Q}_n^*, g_0) - D^{\mathcal{P}}(\bar{Q}, g_0)\big]^2 \to 0$ 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 $P_0 f = \int f(o) dP_0(0)$ for some function f.

Theorem 3. *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

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

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

We note that $D^{\mathcal{C}}$ is the influence curve of the TMLE for the conditional estimand $\Psi_0^{\mathcal{C}}(P_0) = \frac{1}{n}\sum_{i=1}^n \bar{Q}_0(1, W_i) - \bar{Q}_0(0, W_i)$, which corresponds to the conditional average treatment effect (CATE) under the necessary causal assumptions [129]. 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

$\Psi_0^{\mathcal{P}}(P_0)$ and the SATE as

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

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

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

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:

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

with influence curve

$$
D^{\mathcal{S}}(U,X) = D^{\mathcal{C}}(O) - D^F(U,X)
$$

$\square$

C.2 Variance and variance estimation for the TMLE of the SATE

Theorem 4. *The standardized TMLE for the SATE is asymptotically normal:*

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

Proof. The covariance term is

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

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

$$E_{U,X,0}\left[Y(1) - Y(0) - \bar{Q}_0(W)\big|W\right] = E_{U,X,0}[Y(1)|W] - E_{U,X,0}[Y(0)|W] - \bar{Q}_0(W)$$
$$= E_0(Y|A=1,W) - E_0(Y|A=0,W) - \bar{Q}_0(W) = 0$$

Therefore, the second term is zero.

For the first term, we have

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

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

$$\left.\times\left[Y(1)-Y(0)-\bar{Q}_0(W)\right]\right]$$

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

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

It follow that this equals:

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

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

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

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

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:

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

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

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

$\square$

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

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

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

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

C.3 Generalization to allow for estimation of the exposure mechanism

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 $\Psi_0^{\mathcal{P}}(P_0)$ [55]. 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.

More formally, suppose our target of inference is the population estimand $\Psi_0^{\mathcal{P}}(P_0)$ 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 $g_0(A|W)$ [55]:

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

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

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

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

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

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:

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

The covariance of the projection $\prod \left[D^{\mathcal{C}}(\bar{Q}, g_0) \big| T_g \right]$ 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

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

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

Appendix D

Technical Appendix for Chapter 4

D.1 The TMLE is an asymptotically linear estimator of the SATE in an adaptive pair-matched trial

In this Appendix we first review the asymptotic linearity results of Balzer *et al.* [129] for estimation and inference of the the statistical parameter corresponding to the conditional average treatment effect (CATE) [116]:

$$\Psi_0^{\mathcal{C}}(P_0) = \frac{1}{n} \sum_{i=1}^{n} \left[\bar{Q}_0(1, W_i) - \bar{Q}_0(0, W_i) \right].$$

We then provide a theorem showing that the TMLE for the SATE is asymptotically normal in a trial with adaptive pair-matching, which results in $n/2$ conditionally independent copies of $\bar{O}_j = (O_{j1}, O_{j2}) = \big((W_{j1}, A_{j1}, Y_{j1}), (W_{j2}, A_{j2}, Y_{j2})\big)$.

As discussed in Balzer *et al.* [129], the TMLE for conditional estimand $\Psi_0^{\mathcal{C}}(P_0)$ is defined by the following substitution estimator:

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

where $\bar{Q}_n^*(A, W)$ denotes the targeted estimator. Under the following assumptions, the TMLE for $\Psi_0^{\mathcal{C}}(P_0)$ is asymptotically linear:

$$\Psi_n(P_n) - \Psi_0^{\mathcal{C}}(P_0) = \frac{1}{n/2} \sum_{j=1}^{n/2} \bar{D}^{\mathcal{C}}(\bar{Q}, \bar{Q}_0, g_0)(\bar{O}_j) + o_P(n^{-1/2})$$

with influence curve

$$\bar{D}^{\mathcal{C}}(\bar{Q}, \bar{Q}_0, g_0)(\bar{O}_j) = \frac{1}{2} \left[D^{\mathcal{C}}(\bar{Q}, \bar{Q}_0, g_0)(O_{j1}) + D^{\mathcal{C}}(\bar{Q}, \bar{Q}_0, g_0)(O_{j2}) \right]$$

$$D^{\mathcal{C}}(\bar{Q}, \bar{Q}_0, g_0)(O_i) = \left(\frac{\mathbb{I}(A_i = 1)}{g_0(A_i)} - \frac{\mathbb{I}(A_i = 0)}{g_0(A_i)} \right) \left(Y_i - \bar{Q}(A_i, W_i) \right)$$

$$- \left[\left(\bar{Q}_0(1, W_i) - \bar{Q}(1, W_i) \right) - \left(\bar{Q}_0(0, W_i) - \bar{Q}(0, W_i) \right) \right]$$

where $\bar{Q}(A, W)$ denotes the limit of the targeted estimator of the conditional mean function $\bar{Q}_0(A, W)$ and where the marginal probability of being assigned the treatment or the control is known: $g_0(A) = P_0(A) = 0.5$ [129]. Specifically, we assume

- Uniform bound: Assume $\sup_{\bar{Q} \in \mathcal{F}} \sup_O | \left(\frac{\mathbb{I}(A_i=1)}{g_0(A_i)} - \frac{\mathbb{I}(A_i=0)}{g_0(A_i)} \right) \left(Y_i - \bar{Q}(A_i, W_i) \right) | < M < \infty$
  where $\mathcal{F}$ is the set of multivariate real valued functions so that $\bar{Q}_n^*$ is an element of $\mathcal{F}$ with probability 1 and where the second supremum is over a set that contains the support of each O_i.

- Convergence of variances: Assume that for a specified $\{\sigma^{2,\mathcal{C}}(\bar{Q}) : \bar{Q} \in \mathcal{F}\}$, for any $\bar{Q} \in \mathcal{F}$, $\frac{1}{n/2} \sum_{j=1}^{n/2} P_0^n \bar{D}^{\mathcal{C}}(\bar{Q}, \bar{Q}_0, g_0)^2 \to \sigma^{2,\mathcal{C}}(\bar{Q})$ a.s (i.e., for almost every $(W^n, n \geq 1)$). Throughout $P_0^n f = E_0[f|W^n]$ denotes the conditional expectation of a function f of $O^n = (O_1, \ldots, O_n)$, given the vector of baseline covariates $W^n = (W_1, \ldots, W_n)$. We will relax this assumption below.

- Convergence of $\bar{Q}_n^*$ to some limit: For any $\bar{Q}_1, \bar{Q}_2 \in \mathcal{F}$, we define $\sigma_n^2(\bar{Q}_1 - \bar{Q}_2) = \frac{1}{n/2} \sum_{j=1}^{n/2} P_0^n \{\bar{D}^{\mathcal{C}}(\bar{Q}_1, \bar{Q}_0, g_0) - \bar{D}^{\mathcal{C}}(\bar{Q}_2, \bar{Q}_0, g_0)\}^2$. Assume that for a particular $\bar{Q} \in \mathcal{F}$, $\sigma_n^2(\bar{Q}_n^* - \bar{Q}) \to 0$ in probability as $n \to \infty$.

- Entropy condition: Let $\mathcal{F}^d = \{f_1 - f_2 : f_1, f_2 \in \mathcal{F}\}$. Let $N(\epsilon, \sigma_n, \mathcal{F}^d)$ be the covering number of the class $\mathcal{F}^d$ w.r.t norm/dissimilarity $\| f \| = \sigma_n(f)$. Assume that the class $\mathcal{F}$ satisfies $\lim_{\delta_n \to 0} \int_0^{\delta_n} \sqrt{\log N(\epsilon, \sigma_n, \mathcal{F}^d)} d\epsilon = 0$.

Theorem 5. *Let W denote the measured baseline covariates, A the intervention assignment and Y the outcome. A randomized trial with adaptive pair-matching results in $n/2$ conditionally independent copies of paired random variable*

$$\bar{O}_j = \left(O_{j1}, O_{j2} \right) = \left((W_{j1}, A_{j1}, Y_{j1}), (W_{j2}, A_{j2}, Y_{j2}) \right)$$

where index $j = \{1, \ldots, n/2\}$ denotes the partitioning of the study units $\{1, \ldots n\}$ into matched pairs according to similarity on their baseline covariates $W^n = (W_1, \ldots, W_n)$. Our target of inference is the sample average treatment effect (SATE) [19]:

$$\Psi^{\mathcal{S}}(P_X) = \frac{1}{n} \sum_{i=1}^{n} Y_i(1) - Y_i(0)$$

where P_X denotes the distribution of the full data $X = \big(W, Y(1), Y(0)\big)$. Under the above 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:

$$\Psi_n(P_n) - \Psi^{\mathcal{S}}(P_X) = \frac{1}{n/2} \sum_{j=1}^{n/2} \bar{D}^{\mathcal{S}}(\bar{Q}, \bar{Q}_0, g_0)(\bar{X}_j, \bar{O}_j) + o_P(n^{-1/2})$$

with influence curve

$$\bar{D}^{\mathcal{S}}(\bar{Q}, \bar{Q}_0, g_0)(\bar{X}_j, \bar{O}_j) = \bar{D}^{\mathcal{C}}(\bar{Q}, \bar{Q}_0, g_0)(\bar{O}_j) - \bar{D}^{\mathcal{F}}(\bar{Q}_0)(\bar{X}_j, \bar{O}_j)$$

where $\bar{Q}(A, W)$ denotes the limit of the targeted estimator of the conditional mean function $\bar{Q}_0(A, W)$ and where the marginal probability of being assigned the treatment or the control is known $g_0(A) = P_0(A)$.

The first component $\bar{D}^{\mathcal{C}}(\bar{O}_j)$ is the influence curve for the TMLE targeting the conditional estimand $\Psi_0^{\mathcal{C}}(P_0) = \frac{1}{n} \sum_{i=1}^{n} \big[\bar{Q}_0(1, W_i) - \bar{Q}_0(0, W_i)\big]$ in a trial with adaptive pair-matching:

$$\bar{D}^{\mathcal{C}}(\bar{Q}, \bar{Q}_0, g_0)(\bar{O}_j) = \frac{1}{2} \bigg[D^{\mathcal{C}}(\bar{Q}, \bar{Q}_0, g_0)(O_{j1}) + D^{\mathcal{C}}(\bar{Q}, \bar{Q}_0, g_0)(O_{j2}) \bigg]$$

$$\text{with } D^{\mathcal{C}}(\bar{Q}, \bar{Q}_0, g_0)(O_i) = \left(\frac{\mathbb{I}(A_i = 1)}{g_0(A_i)} - \frac{\mathbb{I}(A_i = 0)}{g_0(A_i)} \right) \big(Y_i - \bar{Q}(A_i, W_i) \big)$$

$$- \bigg[\big(\bar{Q}_0(1, W_i) - \bar{Q}(1, W_i)\big) - \big(\bar{Q}_0(0, W_i) - \bar{Q}(0, W_i)\big) \bigg]$$

The second component $\bar{D}^{\mathcal{F}}(\bar{X}_j, \bar{O}_j)$ is the following function of the paired full data $\bar{X}_j = \big(X_{j1}, X_{j2}\big)$:

$$\bar{D}^{\mathcal{F}}(\bar{Q}_0)(\bar{X}_j, \bar{O}_j) = \frac{1}{2} \bigg[D^{\mathcal{F}}(\bar{Q}_0)(X_{j1}, O_{j1}) + D^{\mathcal{F}}(\bar{Q}_0)(X_{j2}, O_{j2}) \bigg]$$

$$\text{with } D^{\mathcal{F}}(\bar{Q}_0)(X_i, O_i) = Y_i(1) - Y_i(0) - \big[\bar{Q}_0(1, W_i) - \bar{Q}_0(0, W_i) \big]$$

The standardized TMLE for the SATE is asymptotically normal with mean 0 and variance $\sigma^{2, \mathcal{S}}$ given by the limit of

$$\sigma_n^{2, \mathcal{S}} = \frac{1}{n/2} \sum_{j=1}^{n/2} P_0^n \bigg\{ \bar{D}^{\mathcal{S}}(\bar{Q}, \bar{Q}_0, g_0)(\bar{X}_j, \bar{O}_j) \bigg\}^2$$

where $P_0^n f = E_0[f | W^n]$ denotes the conditional expectation of a function f of $O^n = (O_1, \ldots, O_n)$, given the vector of baseline covariates $W^n = (W_1, \ldots, W_n)$.

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 conditional estimand $\Psi_0^{\mathcal{C}}(P_0)$ and the SATE as

$$
\begin{aligned}
\Psi_n&(P_n) - \Psi^{\mathcal{S}}(P_X) \\
&= \Psi_n(P_n) - \Psi_0^{\mathcal{C}}(P_0) - \left[\Psi^{\mathcal{S}}(P_X) - \Psi_0^{\mathcal{C}}(P_0)\right] \\
&= \frac{1}{n/2} \sum_{j=1}^{n/2} \bar{D}^{\mathcal{C}}(\bar{O}_j) - \left[\Psi^{\mathcal{S}}(P_X) - \Psi_0^{\mathcal{C}}(P_0)\right] + o_P(n^{-1/2}) \\
&= \frac{1}{n/2} \sum_{j=1}^{n/2} \bar{D}^{\mathcal{C}}(\bar{O}_j) - \left[\frac{1}{n}\sum_{i=1}^{n} Y_i(1) - Y_i(0) - \bar{Q}_0(W_i)\right] + o_P(n^{-1/2}) \\
&= \frac{1}{n/2} \sum_{j=1}^{n/2} \left[\bar{D}^{\mathcal{C}}(\bar{O}_j) - \frac{1}{2}\Big(Y_{j1}(1) - Y_{j1}(0) - \bar{Q}_0(W_{j1}) + Y_{j2}(1) - Y_{j2}(0) - \bar{Q}_0(W_{j2})\Big)\right] \\
&\quad + o_P(n^{-1/2}) \\
&= \frac{1}{n/2} \sum_{j=1}^{n/2} \left[\bar{D}^{\mathcal{C}}(\bar{O}_j) - \bar{D}^{\mathcal{F}}(\bar{X}_j, \bar{O}_j)\right] + o_P(n^{-1/2})
\end{aligned}
$$

where $\bar{D}^{\mathcal{C}}(\bar{O}_j)$ is the influence curve of the TMLE for the conditional estimand $\Psi_0^{\mathcal{C}}(P_0)$ under adaptive pair-matching and where $\bar{D}^{\mathcal{F}}(\bar{X}_j, \bar{O}_j)$ is the following function of the paired full data $\bar{X}_j = (X_{j1}, X_{j2})$:

$$
\bar{D}^{\mathcal{F}}(\bar{X}_j, \bar{O}_j) = \frac{1}{2}\left[D^{\mathcal{F}}(X_{j1}, O_{j1}) + D^{\mathcal{F}}(X_{j2}, O_{j2})\right]
$$
$$
\text{with } D^{\mathcal{F}}(X_i, O_i) = Y_i(1) - Y_i(0) - \left[\bar{Q}_0(1, W_i) - \bar{Q}_0(0, W_i)\right]
$$

Thus, we have shown the TMLE is an asymptotically linear estimator of the SATE in a trial with adaptive pair-matching:

$$
\Psi_n(P_n) - \Psi^{\mathcal{S}}(P_X) = \frac{1}{n/2} \sum_{j=1}^{n/2} \bar{D}^{\mathcal{S}}(\bar{X}_j, \bar{O}_j) + o_P(n^{-1/2})
$$

with influence curve
$$
\bar{D}^{\mathcal{S}}(\bar{X}_j, \bar{O}_j) = \bar{D}^{\mathcal{C}}(\bar{O}_j) - \bar{D}^{\mathcal{F}}(\bar{X}_j, \bar{O}_j).
$$

$\square$

Strictly speaking, the influence curve must only be a function of the observed data. Nonetheless, the theorem is sufficient to prove asymptotic normality and consistency of the TMLE.

D.2 Variance and variance estimation for the TMLE of the SATE in an adaptive pair-matched trial

Theorem 6. *The asymptotic variance of the standardized estimator is given by the limit of*

$$\sigma_n^{2,\mathcal{S}} = \frac{1}{n/2} \sum_{j=1}^{n/2} P_0^n \left\{ \bar{D}^{\mathcal{S}}(\bar{Q}, \bar{Q}_0, g_0)(\bar{X}_j, \bar{O}_j) \right\}^2$$

$$= \frac{1}{n/2} \sum_{j=1}^{n/2} \left[P_0^n \left\{ \bar{D}^{\mathcal{C}}(\bar{Q}, \bar{Q}_0, g_0)(\bar{O}_j) \right\}^2 - \frac{1}{4} P_0^n \left\{ D^{\mathcal{F}}(\bar{Q}_0)(X_{j1}, O_{j1}) \right\}^2 \right.$$

$$\left. - \frac{1}{4} P_0^n \left\{ D^{\mathcal{F}}(\bar{Q}_0)(X_{j2}, O_{j2}) \right\}^2 \right]$$

Proof. The conditional variance can be expressed as

$$\sigma_n^{2,\mathcal{S}} = \frac{1}{n/2} \sum_{j=1}^{n/2} P_0^n \left\{ \bar{D}^{\mathcal{S}}(\bar{X}_j, \bar{O}_j) \right\}^2$$

$$= \frac{1}{n/2} \sum_{j=1}^{n/2} P_0^n \left\{ \bar{D}^{\mathcal{C}}(\bar{O}_j) - \bar{D}^{\mathcal{F}} \bar{X}_j, \bar{O}_j) \right\}^2$$

$$= \frac{1}{n/2} \sum_{j=1}^{n/2} \left[P_0^n \left\{ \bar{D}^{\mathcal{C}}(\bar{O}_j) \right\}^2 + P_0^n \left\{ \bar{D}^{\mathcal{F}}(\bar{X}_j, \bar{O}_j) \right\}^2 - 2 P_0^n \left\{ \bar{D}^{\mathcal{C}}(\bar{O}_j) \times \bar{D}^{\mathcal{F}}(\bar{X}_j, \bar{O}_j) \right\} \right]$$

The conditional variance of the $\bar{D}^{\mathcal{F}}(\bar{X}_j, \bar{O}_j)$ component is

$$P_0^n \left\{ \bar{D}^{\mathcal{F}}(\bar{X}_j, \bar{O}_j) \right\}^2 = P_0^n \left\{ \frac{1}{2} \left(D^{\mathcal{F}}(X_{j1}, O_{j1}) + D^{\mathcal{F}}(X_{j2}, O_{j2}) \right) \right\}^2$$

$$= \frac{1}{4} P_0^n \left\{ D^{\mathcal{F}}(X_{j1}, O_{j1}) \right\}^2 + \frac{1}{4} P_0^n \left\{ D^{\mathcal{F}}(X_{j2}, O_{j2}) \right\}^2$$

$$+ \frac{1}{2} P_0^n \left\{ D^{\mathcal{F}}(X_{j1}, O_{j1}) \times D^{\mathcal{F}}(X_{j2}, O_{j2}) \right\}$$

Under the randomization assumption, each $D^{\mathcal{F}}(X_i, O_i)$ component has conditional mean zero, given its baseline covariates W_i:

$$E[D^{\mathcal{F}}(X_i, O_i) | W_i] = E \left[Y_i(1) - Y_i(0) - \bar{Q}_0(W_i) | W_i \right]$$

$$= E[Y_i(1) | W_i] - E[Y_i(0) | W_i] - \left[\bar{Q}_0(1, W_i) - \bar{Q}_0(0, W_i) \right]$$

$$= E_0(Y_i | A_i = 1, W_i) - E_0(Y_i | A_i = 0, W_i) - \left[\bar{Q}_0(1, W_i) - \bar{Q}_0(0, W_i) \right] = 0$$

Therefore, the conditional covariance of the $D^{\mathcal{F}}$ components within a matched pair is zero:

$$P_0^n\left\{D^{\mathcal{F}}(X_{j1}, O_{j1}) \times D^{\mathcal{F}}(X_{j2}, O_{j2})\right\}$$

$$= E\left[D^{\mathcal{F}}(X_{j1}, O_{j1}) \times D^{\mathcal{F}}(X_{j2}, O_{j2})\big|W^n\right]$$

$$= E[E[Y_{j1}(1) - Y_{j1}(0) - \bar{Q}_0(W_{j1})|W_{j1}] \times D^{\mathcal{F}}(X_{j2}, O_{j2})] = 0$$

The conditional variance of the $\bar{D}^{\mathcal{F}}(\bar{X}_j, \bar{O}_j)$ component simplifies to

$$P_0^n\left\{\bar{D}^{\mathcal{F}}(\bar{X}_j, \bar{O}_j)\right\}^2 = \frac{1}{4}P_0^n\left\{D^{\mathcal{F}}(X_{j1}, O_{j1})\right\}^2 + \frac{1}{4}P_0^n\left\{D^{\mathcal{F}}(X_{j2}, O_{j2})\right\}^2$$

The conditional covariance of the $\bar{D}^{\mathcal{C}}(\bar{O}_j)$ and $\bar{D}^{\mathcal{F}}(\bar{X}_j, \bar{O}_j)$ components is

$$P_0^n\left\{\bar{D}^{\mathcal{C}}(\bar{O}_j) \times \bar{D}^{\mathcal{F}}(\bar{X}_j, \bar{O}_j)\right\}$$

$$= \frac{1}{4}P_0^n\left\{\left[D^{\mathcal{C}}(O_{j1}) + D^{\mathcal{C}}(O_{j2})\right] \times \left[D^{\mathcal{F}}(X_{j1}, O_{j1}) + D^{\mathcal{F}}(X_{j2}, O_{j2})\right]\right\}$$

$$= \frac{1}{4}\bigg[P_0^n\{D^{\mathcal{C}}(O_{j1}) \times D^{\mathcal{F}}(X_{j1}, O_{j1})\} + P_0^n\{D^{\mathcal{C}}(O_{j1}) \times D^{\mathcal{F}}(X_{j2}, O_{j2})\}$$

$$+ P_0^n\{D^{\mathcal{C}}(O_{j2}) \times D^{\mathcal{F}}(X_{j1}, O_{j1})\} + P_0^n\{D^{\mathcal{C}}(O_{j2}) \times D^{\mathcal{F}}(X_{j2}, O_{j2})\}\bigg]$$

As shown in Appendix of [140], the covariance of the $D^{\mathcal{C}}(O_i)$ and $D^{\mathcal{F}}(O_i)$ components is equal to the variance of $D^{\mathcal{F}}(O_i)$. Therefore, we have

$$P_0^n\left\{D^{\mathcal{C}}(O_{j1}) \times D^{\mathcal{F}}(X_{j1}, O_{j1})\right\} = P_0^n\left\{D^{\mathcal{F}}(X_{j1}, O_{j2})\right\}^2$$

$$P_0^n\left\{D^{\mathcal{C}}(O_{j2}) \times D^{\mathcal{F}}(X_{j2}, O_{j2})\right\} = P_0^n\left\{D^{\mathcal{F}}(X_{j2}, O_{j2})\right\}^2$$

Under the randomization assumption, the other terms are zero:

$$P_0^n\left\{D^{\mathcal{C}}(O_{j1}) \times D^{\mathcal{F}}(X_{j2}, O_{j2})\right\} = E\left[D^{\mathcal{C}}(O_{j1}) \times E\left[D^{\mathcal{F}}(X_{j2}, O_{j2})|W_{j2}\right]\right] = 0$$

$$P_0^n\left\{D^{\mathcal{C}}(O_{j2}) \times D^{\mathcal{F}}(X_{j1}, O_{j1})\right\} = E\left[D^{\mathcal{C}}(O_{j2}) \times E\left[D^{\mathcal{F}}(X_{j1}, O_{j1})|W_{j1}\right]\right] = 0$$

We have that the conditional covariance of the $\bar{D}^{\mathcal{C}}(\bar{O}_j)$ and $\bar{D}^{\mathcal{F}}(\bar{X}_j, \bar{O}_j)$ components equals

$$P_0^n\left\{\bar{D}^{\mathcal{C}}(\bar{O}_j) \times \bar{D}^{\mathcal{F}}(\bar{X}_j, \bar{O}_j)\right\} = \frac{1}{4}P_0^n\left\{D^{\mathcal{F}}(X_{j1}, O_{j2})\right\}^2 + \frac{1}{4}P_0^n\left\{D^{\mathcal{F}}(X_{j2}, O_{j2})\right\}^2$$

Combining the terms, we have

$$\sigma_n^{2,\mathcal{S}} = \frac{1}{n/2} \sum_{j=1}^{n/2} \left[P_0^n \left\{ \bar{D}^\mathcal{C}(\bar{O}_j) \right\}^2 - \frac{1}{4} P_0^n \left\{ D^\mathcal{F}(X_{j1}, O_{j1}) \right\}^2 - \frac{1}{4} P_0^n \left\{ D^\mathcal{F}(X_{j2}, O_{j2}) \right\}^2 \right]$$

$\square$

The asymptotic variance of the TMLE for the SATE $\sigma^{2,\mathcal{S}}$ is always less than or equal to $\sigma^{2,\mathcal{C}}$, which is the asymptotic variance of the TMLE for the conditional parameter. As shown in [129], we can estimate the upper bound as

$$\hat{\sigma}^{2,\mathcal{S}} = \hat{\sigma}^{2,\mathcal{C}} = \frac{1}{n/2} \sum_{j=1}^{n/2} \left\{ \hat{\bar{D}}^\mathcal{C}(\bar{Q}_n^*, g_0)(\bar{O}_j) \right\}^2$$

$$\hat{\bar{D}}^\mathcal{C}(\bar{Q}_n^*, g_0)(\bar{O}_j) = \frac{1}{2} \left[\hat{D}^\mathcal{C}(\bar{Q}_n^*, g_0)(O_{j1}) + \hat{D}^\mathcal{C}(\bar{Q}_n^*, g_0)(O_{j2}) \right]$$

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

Ordering the observations within matched pairs, such that the first corresponds to the unit randomized to the intervention ($A_{j1} = 1$) and the second to the control ($A_{j2} = 0$), it follows that

$$\hat{\bar{D}}^\mathcal{C}(\bar{Q}_n^*, g_0)(\bar{O}_j) = \hat{\bar{D}}^\mathcal{C}(\bar{Q}_n^*, g_0)(\bar{O}_j) = \left(Y_{j1} - \bar{Q}_n^*(1, W_{j1}) \right) - \left(Y_{j2} - \bar{Q}_n^*(0, W_{j2}) \right)$$

allowing us to represent the variance estimator as the sample variance of the difference in residuals within matched pairs:

$$\hat{\sigma}^{2,\mathcal{S}} = \hat{\sigma}^{2,\mathcal{C}} = \frac{1}{n/2} \sum_{j=1}^{n/2} \left\{ \left(Y_{j1} - \bar{Q}_n^*(1, W_{j1}) \right) - \left(Y_{j2} - \bar{Q}_n^*(0, W_{j2}) \right) \right\}^2$$

This variance estimator will be consistent if there is no heterogeneity in the treatment effect within strata of covariates (i.e. if the variance of the $D^\mathcal{F}$ component is zero) *and* if the conditional mean function $\bar{Q}_0(A, W)$ is consistently estimated. Otherwise, the variance estimator will be conservative.

We can relax the assumption that the conditional variance converges to some limit. Specifically, we have that the standardized and scaled estimator converges to a normal distribution with mean 0 and variance given by the ratio of the true conditional variance $\sigma_n^{2,\mathcal{S}}$ divided by our conservative estimator $\hat{\sigma}^{2,\mathcal{S}}$:

$$\frac{\Psi(P_n) - \Psi^\mathcal{S}(P_X)}{\hat{\sigma}^{2,\mathcal{S}}/\sqrt{n/2}} \to N\left(0, \frac{\sigma_n^{2,\mathcal{S}}}{\hat{\sigma}^{2,\mathcal{S}}} \right)$$

Since the ratio of variances is always ≤ 1. The standard normal distribution $N(0,1)$ provides a conservative approximation to the asymptotic distribution.

D.3 Step-by-step instructions to obtain a cross-validated variance estimator

Let $\bar{Q}_n(A, W)$ denote the initial estimator for conditional mean outcome, which was selected through the data-adaptive procedure (Sec. 4.2 for a non-matched trial and Sec. 4.3 for a pair-matched trial). Let $g_n(A|W)$ denote the estimator of the exposure mechanism, which was collaboratively selected through the data-adaptive procedure (Sec. 4.4). A cross-validated estimate of the variance of the data-adaptive TMLE can be implemented as follows. As before, we present V-fold cross-validation, where the data are partitioned into V folds of size $\approx n/V$. If matching was used, the partitioning should preserve the pairs.

i. For each fold $v = \{1, \ldots, V\}$ in turn,

   a. Set the observation(s) in fold v to be the validation set and the remaining observations to be the training set.

   b. Using observations in the training set, fit the selected TMLE.
      - Fit the selected working model for the conditional mean outcome $\bar{Q}_n(A, W)$.
      - Fit the selected working model for the exposure mechanism $g_n(A|W)$.
      - Target the initial estimator. Denote the estimated fluctuation coefficient ϵ_n.

   c. For each observation O_k in the validation set, estimate the influence curve (and correction factor ρ_n if relevant).
      - Use the initial fit $\bar{Q}_n(A, W)$, based on the training data, to obtain initial predictions of the outcome under the treatment $\bar{Q}_n(1, W_k)$ and under the control $\bar{Q}_n(0, W_k)$.
      - Use the the fit of the exposure mechanism $g_n(A, W)$, based on the training set, to calculate the clever covariate $H_n(A_k, W_k)$.
      - Update the initial estimates with the estimated fluctuation parameter ϵ_n. Denote the targeted predictions of the outcome under the treatment $\bar{Q}_n^*(1, W_k)$ and the control $\bar{Q}_n^*(0, W_k)$.
      - Plug-in the relevant components to estimate influence curve, appropriate for the target parameter and study design.

ii. Estimate the variance of the data-adaptive TMLE with the sample variance of the estimated influence curve, normalized by the appropriate sample size.