

Efficient Principally Stratified Treatment Effect Estimation in Crossover Studies with Absorbent Binary Endpoints

Alex Luedtke^{1,2} and Jiacheng Wu³

Abstract: We consider the estimation of the effect of a binary treatment on a binary endpoint conditional on a post-randomization quantity in a counterfactual world where all individuals received treatment. It is generally difficult to identify this effect without strong, untestable assumptions. It has been shown that identifiability assumptions become weaker under a crossover design where individuals not receiving treatment are later provided treatment. Under the assumption that the post-treatment biomarker observed in these crossover individuals is the same as would have been observed had they received treatment at the start of the study, the treatment effect can be identified with only mild additional assumptions. This remains true if the endpoint is absorbent, that is, if the post-crossover treatment biomarker is not meaningful if the endpoint has already occurred. Examples of absorbent endpoints include death and HIV infection. We provide identifiability conditions for the principally stratified treatment effect of interest when the data arise from a crossover design and describe situations where these conditions would be falsifiable given a large sample from the observed data distribution. We then introduce a nonparametric estimator for this effect. When the biomarker is discrete, this estimator is efficient among all regular and asymptotically linear estimators.

Keywords: absorbent binary endpoints, closeout design, crossover design, principal stratification

AMS 2000 subject classifications: 62K99, 62G20

1. Introduction

Suppose one wishes to assess the effect of a binary treatment on a binary endpoint. For simplicity, we refer to the levels of the treatment as “treated” and “untreated”, but one could alternatively study a contrast between two active treatments. The objective is to develop a biomarker that explains much of the treatment effect variation, i.e. a variable for which the stratified treatment effect is highly variable. Sometimes, the most predictive biomarkers are defined using the counterfactual values that a post-treatment biomarker would take had an individual been treated or untreated. Conditioning on these biomarkers was termed principal stratification by [Frangakis and Rubin \(2002\)](#). Principally stratified analyses are especially interesting in situations where one wishes to bridge a treatment’s efficacy from one population to another, and so treats a small subset of the new population and observes their biomarkers in order to predict what the treatment effect would be in this population ([Gilbert et al., 2011](#)). Though additional transportability assumptions ([Bareinboim and Pearl, 2012](#)) are needed to be able to bridge a treatment effect, these assumptions become more plausible if a strong effect modifier is available. In this paper, we focus only on counterfactual post-randomization biomarkers that occur under treatment. This ensures that one could observe the counterfactual biomarker under an appropriate intervention – this is in contrast to stratifying on the biomarker’s counterfactual value both after treatment and after a lack of treatment. Furthermore, in vaccine studies, which represent an important application area, the post-treatment biomarkers of most interest are immune responses, and in many settings untreated individuals who have not experienced the endpoint will have no immune response, so that studying treatment effects conditional on unvaccinated immune responses is uninteresting. We assume that the endpoint is absorbent

¹ Department of Statistics, University of Washington, 4060 E Stevens Way NE, Seattle, WA 98195, USA.
E-mail: aluedtke@uw.edu

² Vaccine and Infectious Disease Division, Fred Hutchinson Cancer Research Center, 1100 Fairview Ave N, Seattle, WA 98109, USA.

³ Department of Biostatistics, University of Washington, 1705 NE Pacific St, Seattle, WA 98195, USA.
E-mail: wuj92@uw.edu

(Nason and Follmann, 2010), so that the biomarker has no substantive meaning once the endpoint has occurred. For example, if the treatment is vaccination, the biomarker is vaccine-induced immune response, and the endpoint is HIV infection, then a measured immune response is only of interest if it precedes HIV infection. For most biomarkers of interest, death is also an absorbent endpoint.

Under most study designs, identifying principally stratified treatment effects requires strong, untestable assumptions. In the context of vaccine trials, Follmann (2006) demonstrated the utility of crossover designs (e.g., Woods et al., 1989) for dramatically weakening these needed assumptions, where here the designs were referred to as closeout placebo vaccination designs and featured a crossover of uninfected placebos to the vaccine. There have since been numerous works that study closeout designs (Gilbert and Hudgens, 2008; Wolfson and Gilbert, 2010; Gabriel and Gilbert, 2013; Huang et al., 2013). Others have demonstrated that crossover designs are also of interest for estimating principally stratified treatment effects for treatments other than vaccines (Wolfson and Henn, 2014; Gabriel and Follmann, 2016). All of these earlier works either relied on baseline covariates being discrete or on correctly specified (semi)parametric models. In this work, we provide an estimator of principally stratified treatment effects that is efficient within the nonparametric model that at most makes assumptions on the probability of receiving treatment given covariates. The efficiency of this estimator relies on the biomarker being discrete, since otherwise the semiparametric efficiency bound (Pfanzagl, 1990; Bickel et al., 1993) will not be defined. We then generalize this nonparametric estimator to the case that the biomarker is continuous.

From a practical standpoint, closeout placebo vaccinations are easy to perform in a vaccine clinical trial setting because uninfected placebo recipients are still under follow-up at the end of the trial and are thus able to be vaccinated. Nonetheless, a downside to performing a vaccinating uninfected placebo recipients at the end of the study is that they are no longer available for additional follow-up as placebo recipients, making it difficult to assess the long-term efficacy of the vaccine. One option is to only vaccinate a (random) subset of uninfected placebo recipients at the end of a trial so that some placebo recipients are still available for longer-term follow-up.

We note that studying principally stratified treatment effects answers fundamentally different questions than does studying direct effects (Robins and Greenland, 1992; Pearl, 2001; VanderWeele, 2015), which would aim to study the effect of treatment on the outcome if the post-treatment biomarker were set to some fixed value or to the value it would have taken had the individual been untreated. This is in contrast to principal stratification analyses, which are useful for bridging vaccine efficacy from one population to another, provided one collects the post-treatment biomarker on a small subset of individuals in the new population and certain transportability assumptions hold. Identifying counterfactual direct effect estimands with the observed data distribution requires strong assumptions, including that there are no unmeasured confounders of the effect of the biomarker on the outcome (Cole and Hernán, 2002). In crossover studies, this no-unmeasured-confounders assumption does not need to hold to obtain valid inference for the principally stratified parameters that we study in this work. Nonetheless, we emphasize that estimation should be focused on the quantity that best reflects the scientific question. Furthermore, we note that, though principally stratified analyses focus on a different quantity than direct effect analyses, it has been shown that the presence of a nonzero principally stratified effect, conditional equality of the counterfactual biomarker under treatment and the counterfactual biomarker under a lack of treatment, implies the presence of a nonzero direct effect (VanderWeele, 2008); the reverse implication does not hold. In this work, we study effects stratified only on the counterfactual biomarker under treatment, so the result of VanderWeele (2008) does not generally apply. Nonetheless, in settings in which the counterfactual biomarker under a lack of treatment is degenerate at some value s_0 , e.g. in HIV vaccine studies where the biomarker is an immune response and there is no immune response if the vaccine is not administered, the presence of a nonzero treatment effect conditional on the counterfactual biomarker value under treatment equal to s_0 indicates the presence of a direct effect, i.e. an effect in a counterfactual world where treatment was administered and the biomarker had been set to s_0 .

Organization of Manuscript. Section 2 introduces the notation, parameters of the counterfactual distribution, identifiability of these parameters with parameters of the observed data distribution, and the

statistical estimation problem. Section 3 describes our proposed estimator when the biomarker is discrete. Section 4 presents an extension of this estimation scheme to two-phase sampling studies. Section 5 presents an estimator that can be used when the biomarker is continuous. For simplicity, this estimator is presented in a single-phase sampling setup, but the extension to two-phase sampling is straightforward. Section 6 presents a simulation study. Section 7 closes with a discussion.

Appendix B provides proofs of the results from the main text. Appendix C presents theorems and proofs for the validity of the estimator for continuous biomarkers presented in Section 5. Appendix D gives a brief review of the derivation of efficient influence functions, which are used to define our estimators. Appendix E presents a TMLE that can be used when the data was derived from a two-phase sampling design. Appendix F gives additional simulation results.

2. Notation and Identifiability

Consider the counterfactual data structure $(W, A, S_1, Y_0, Y_1, S_0^c)$ drawn from a distribution P^F , where W is a baseline covariate, A is a treatment indicator, S_1 is a post-treatment biomarker, Y_a is a counterfactual outcome of interest under treatment a that takes on values zero and one, and S_0^c is the counterfactual post-crossover biomarker had the individual not been treated at baseline and subsequently been crossed over to treatment. In a closeout placebo vaccination study, S_1 represents a post-vaccination immune response, and S_0^c denotes the closeout biomarker measurements for placebos with $Y = 0$, and $S_0^c = 0$ for all other individuals. To simplify presentation, we assume that S_1 and S_0^c are measured immediately following treatment. One could alternatively assume that no events occur before the time at which S_1 is measured (Follmann, 2006), or could assume some form of equal clinical risk, which roughly states that the risk of the event occurring before this time point is equal between treated and untreated individuals (Wolfson and Gilbert, 2010).

The objective is to estimate a contrast between the principally stratified mean outcome on treatment and in the absence of treatment, i.e. between $\mathbb{E}^F[Y_1|S_1 = s_1^*]$ and $\mathbb{E}^F[Y_0|S_1 = s_1^*]$, where here and throughout we let \mathbb{E}^F denote an expectation under P^F and we define s_1^* as a particularly interesting value of the biomarker S_1 . For example, we may be interested in estimating the principally stratified relative risk

$$\text{RR}^F(s_1^*) \equiv \frac{\mathbb{E}^F[Y_1|S_1 = s_1^*]}{\mathbb{E}^F[Y_0|S_1 = s_1^*]}$$

or vaccine efficacy

$$\text{VE}^F(s_1^*) \equiv 1 - \frac{\mathbb{E}^F[Y_1|S_1 = s_1^*]}{\mathbb{E}^F[Y_0|S_1 = s_1^*]} \quad (1)$$

One may also be interested in RR^F or VE^F as a curve, across all values of s_1^* . The developments in this work can also be used to estimate an additive contrast, thereby giving the conditional additive treatment effect.

In practice, we observe $(W, A, S, Y, S^c) \sim P$. We wish to identify the principally stratified mean outcomes with a parameter of this observed data distribution P . We make the following identifiability assumptions, where we note that here and throughout we use \mathbb{E} to denote an expectation under P :

- (A1) A draw of $O = (W, A, S, Y, S^c)$ from P has the same distribution as a draw of $(W, A, \mathbb{1}_{A=1} S_1, Y_A, \mathbb{1}_{A=0} S_0^c)$ from P^F .

The above is implied by the consistency assumption, which states that $S_1 = S$ if $A = 1$, $Y = Y_A$, and $S^c = S_0^c$ if $A = 0$. We also make the following assumptions:

- (A2) Ignorable treatment assignment: $A \perp (S_1, Y_0, Y_1) | W$, and
 (A3) Crossover assumption: for P^F almost all w , $S_1 | Y_0 = 0, W = w$ under P^F has the same distribution as $S_0^c | Y_0 = 0, W = w$ under P^F .

This assumption is ill-defined when $P^F(Y_0 = 0|W = w) = 0$, but we note that it could be replaced by equality of the conditional subdistribution functions $S_1, Y_0 = 0|W = w$ and $S_0^c, Y_0 = 0|W = w$ under P^F . If only the effect at a single s_1^* is of interest, then the final assumption above could be replaced by the weaker assumption that, for P^F almost all w , $P^F(S_1 = s_1^*|Y_0 = 0, W = w) = P^F(S_0^c = s_1^*|Y_0 = 0, W = w)$. Note that, in either case, the third assumption above is implied by $S_0^c = S_1$ almost surely, but does not require this. As an example of the greater generality of the above assumption, it allows for a true underlying biomarker to be measured with error, where the conditional probability statement is then interpreted for the noised measurement of the underlying biomarker. Note also that this assumption becomes weaker as the baseline covariate W become more predictive of S_1 and S_0^c . Therefore, using the richest possible baseline covariate by including all available individual-level information is expected to yield the most plausible crossover assumption.

Throughout we also require the strong positivity assumption that, for some $\delta > 0$, the treatment mechanism $P(A = 1|W)$ falls between δ and $1 - \delta$ with probability one over $W \sim P$. We will also assume that there exists some fixed $\delta > 0$ so that any estimate of the treatment mechanism discussed in this paper satisfies the strong positivity assumption with probability approaching one. In Appendix B.1, we prove the following identifiability result.

Theorem 1. *If A1, A2, and A3 hold, then:*

1. $P^F(S_1 = s_1^*) = \mathbb{E}[P(S = s_1^*|A = 1, W)]$,
2. $P^F(Y_1 = 1, S_1 = s_1^*) = \mathbb{E}[P(Y = 1, S = s_1^*|A = 1, W)]$, and
3. $P^F(Y_0 = 0, S_1 = s_1^*) = \mathbb{E}[P(Y = 0, S^c = s_1^*|A = 0, W)]$.

These identifiability results enable the identification of $P^F(Y_1 = 1|S_1 = s_1^*) = \frac{P^F(Y_1=1, S_1=s_1^*)}{P^F(S_1=s_1^*)}$ and of $P^F(Y_0 = 1|S_1 = s_1^*) = 1 - \frac{P^F(Y_0=0, S_1=s_1^*)}{P^F(S_1=s_1^*)}$, and therefore also enable the identification of the principally stratified relative risk or any other contrast of these two quantities.

It will be convenient to define parameters mapping from our statistical model \mathcal{M} , i.e. the nonparametric model that at most places restrictions on the probability of receiving treatment $A = 1$ given covariates W , to the real line. In particular, for an arbitrary distribution $P' \in \mathcal{M}$, we define parameters corresponding to each of the above identifiability results, where to ease notation we omit the dependence of these parameters on the choice of s_1^* :

1. $\Psi_1(P') \equiv \mathbb{E}_{P'}[P'(S = s_1^*|A = 1, W)]$,
2. $\Psi_2(P') \equiv \mathbb{E}_{P'}[\mathbb{E}_{P'}[Y|A = 1, S = s_1^*, W]P'(S = s_1^*|A = 1, W)]$, and
3. $\Psi_3(P') \equiv \mathbb{E}_{P'}[P'(S^c = s_1^*|A = 0, Y = 0, W)P'(Y = 0|A = 0, W)]$.

Note that, under our identifiability conditions, the principally stratified relative risk is equal to

$$RR^P(s_1^*) \equiv \frac{\Psi_2(P')}{\Psi_1(P') - \Psi_3(P')} \tag{2}$$

While the numerator above is bounded between 0 and 1, the denominator may be negative. Therefore, it is possible that the above, which is supposedly identified with a counterfactual relative risk, is negative. This is of course impossible, and therefore indicates a failure of our identifiability assumptions. In Theorem A.1 of Appendix B.1.2 we show that, under A1 and A2, the assumption A3 is not falsifiable from the observed data distribution if $\Psi_4(P) = 0$, where

$$\Psi_4(P') \equiv \int \int \left(P'(S^c = s_1, Y = 0|A = 0, w) - P'(S = s_1|A = 1, w) \right)^+ dP'(S = s_1|A = 1, w)dP'(w).$$

Above x^+ denotes the positive part of x . If, on the other hand, A1 and A2 hold and $\Psi_4^F(P^F) > 0$, then it is easy to show that A3 cannot hold. Our statistical inference for the principally stratified relative risk will hinge on A3. The fact that $\Psi_4(P) > 0$ implies that A3 is false suggests that one could provide a test of A3, and perform this test as a sanity check before proceeding with inference for the principally stratified relative risk. We do not further consider such a test here.

3. Estimation for discrete biomarker

Until otherwise stated, we work at fixed s_1^* . Suppose we observe n i.i.d. samples $O_1, \dots, O_n \sim P$, where $O_i = (W_i, A_i, S_i, Y_i, S_i^c)$. The notational overload on subscripts (indexing individual i and counterfactuals) will not be problematic because the remainder of this work focuses only on observed data quantities.

We will present an estimator of $\underline{\Psi}(P) \equiv (\Psi_k(P) : k = 1, 2, 3) \equiv (\psi_k : k = 1, 2, 3)$ and show that it has a normal limiting distribution. This will enable both a test of A3 and the construction of confidence intervals for the quantity $\text{RR}^P(s_1^*)$ that is, under conditions discussed in the previous section, identified with the principally stratified relative risk. Before the presentation of this estimator, whose construction is somewhat involved, Section 3.1 motivates the theoretical development needed to define the estimator. Section 3.2 presents the estimator.

3.1. Motivation

We will derive the efficient influence function $\underline{D}^{P'} : \mathcal{O} \rightarrow \mathbb{R}^3$ of the parameter $\underline{\Psi}$, where we note that, as the notation suggests, the influence function depends on the distribution P' at which $\underline{\Psi}$ is evaluated. While a more thorough review of the definition of efficient influence functions is given in Appendix D, the key result is that they yield a first-order expansion of the form:

$$\underline{\Psi}(P') - \underline{\Psi}(P) = -\mathbb{E}_{P'}[\underline{D}^{P'}(O)] + \text{Rem}(P, P'), \quad (3)$$

where above $\text{Rem}(P, P') \in \mathbb{R}^3$ is small (in Euclidean norm) relative to the leading term on the right whenever P' is close to P . In our setting, P' will denote an estimate of P , where one really only needs to estimate the conditional expectations and probabilities under P needed to evaluate $\underline{\Psi}$ and \underline{D} , rather than the entire distribution. Once this result has been established, we will use it to develop a targeted minimum loss-based estimator (TMLE) of the three-dimensional estimand $\underline{\psi}$. We can then estimate any smooth function of this quantity, e.g. a relative risk.

Before presenting the TMLE, which is somewhat complex, we provide intuition as to its strong theoretical properties. Suppose we have estimated the components of P needed to evaluate $\underline{\Psi}$ and \underline{D} . Denote the estimate by \hat{P} . Further suppose that this initial estimate is close to P in the sense that $\text{Rem}(P, \hat{P}) \approx 0$, where we later give an explicit expression for this remainder and also give an exact rate requirement to make the approximation symbol precise. Then, by (3), one has that $\underline{\Psi}(\hat{P}) - \underline{\Psi}(P) \approx -\mathbb{E}_P[\underline{D}^{\hat{P}}(O)]$. Because \underline{D}^P is mean zero when applied to draws of $O \sim P$, if \hat{P} is close to the true data generating distribution then it is expected that the right-hand side is close to zero. Nonetheless, when W is continuous it is difficult to quantify how close to zero this expectation is. The TMLE overcomes this challenge by obtaining an estimate \hat{P}^* that is a slight fluctuation of \hat{P} , where this fluctuated estimate now satisfies

$$\frac{1}{n} \sum_{i=1}^n \underline{D}^{\hat{P}^*}(O_i) = \underline{0}. \quad (4)$$

In fact, it would suffice for the left-hand side to converge to zero in probability faster than the inverse of the square root of the sample size. Arguments from M-estimation can be used to show that the fluctuation step does not destroy the convergence properties of the initial estimate \hat{P} , so that $\text{Rem}(P, \hat{P}^*) \approx 0$ whenever this statement holds for the initial, unfluctuated estimator. Combining the above with (3), we have that

$$\underline{\Psi}(\hat{P}^*) - \underline{\Psi}(P) \approx \frac{1}{n} \sum_{i=1}^n \left\{ \underline{D}^{\hat{P}^*}(O_i) - \mathbb{E}_P[\underline{D}^{\hat{P}^*}(O)] \right\}.$$

Multiplying both sides by the square root of sample size, and noting that, under some minor regularity conditions, $\underline{D}^{\hat{P}^*}$ can be replaced by its mean-square limit (often, though not always, this will be \underline{D}^P), we obtain the convergence in distribution result $n^{1/2}[\underline{\Psi}(\hat{P}^*) - \underline{\Psi}(P)] \xrightarrow{d} N(\underline{0}, \Sigma)$, where the covariance

matrix Σ can be estimated using the empirical covariance of $\underline{D}^{\hat{P}^*}(O)$. Confidence intervals for (2) can be derived via the delta method.

Having presented the high-level arguments for our TMLE, we devote the remainder of this section to formally establishing the validity of this estimator. First, we derive the efficient influence functions of the three parameters of interest. Next, we present a TMLE that fluctuates an initial estimate of P so that (4) is satisfied. We then present a theorem giving regularity conditions for the convergence in distribution to a multivariate normal, and we also observe that these regularity conditions can be weakened via cross-validation. Finally, we combine a log transformation with the delta method to demonstrate how to construct confidence intervals for contrasts between the treatment-specific principally stratified risks.

3.2. Presentation of estimator

Per the above Motivation section, the efficient influence function of $\underline{\Psi}$ should play a major role in estimation. Therefore, the following result, which presents an expression for the efficient influence function, will be useful. Before presenting the result, we note that throughout we will use the following conventions for any distribution P' , function f , and realizations (a, w) of (A, W) : $\mathbb{E}_{P'}[f(O)|a, w] = \mathbb{E}_{P'}[f(O)|A = a, W = w]$ and $P'(a|w) = P'(A = a|W = w)$.

Theorem 2. *Within the model \mathcal{M} that at most places restrictions on the probability of receiving treatment given baseline covariates, the parameter $\underline{\Psi}$ has efficient influence function $\underline{D}^{P'} \equiv (D_k^{P'} : k = 1, 2, 3)$ at P' , where, for $k = 1, 2, 3$,*

$$D_k^{P'}(o) \equiv \frac{\mathbb{1}\{a = a_k\}}{P'(a|w)} \{f_k(o) - \mathbb{E}_{P'}[f_k(O)|a_k, w]\} + \mathbb{E}_{P'}[f_k(o)|a_k, w] - \Psi_k(P').$$

Above, $f_1(o) \equiv \mathbb{1}_{S=S_1^*}$ and $a_1 = 1$; $f_2(o) \equiv \mathbb{1}_{Y=1, S=S_1^*}$ and $a_2 = 1$; and $f_3(o) \equiv \mathbb{1}_{Y=0, S^c=S_1^*}$ and $a_3 = 0$.

The proof of this result is in Appendix B.2. It easy to verify that the remainder $\text{Rem}(P, P') \equiv (\text{Rem}_k(P, P') : k = 1, 2, 3)$ in (3) takes the following form:

$$\text{Rem}_k(P, P') = \mathbb{E}_P \left[\left(1 - \frac{P(a_k|W)}{P'(a_k|W)} \right) (\mathbb{E}_{P'}[f_k(O)|a_k, W] - \mathbb{E}_P[f_k(O)|a_k, W]) \right]. \tag{5}$$

Using this efficient influence function, it remains to develop an estimator satisfying (4). This estimator is given in Algorithm 1, which independently estimates each parameter $\Psi_k(P)$, $k = 1, 2, 3$, by invoking a TMLE for estimating $\mathbb{E} \mathbb{E}[f(O)|A = a, W]$ for an arbitrary function f and treatment a . Many variants of this k -specific TMLE have been presented elsewhere (e.g., van der Laan and Rose, 2011). We note here that we are in fact implementing three separate TMLEs, and so it is not immediately obvious that there exists a single distribution \hat{P}^* such that our estimator is equal to $\Psi(\hat{P}^*)$. We can show that a unique distribution exists if $\hat{P}_2^*(f_2(O) = 1|a_2, w) \leq \hat{P}_1^*(f_1(O) = 1|a_1, w)$ for all w .

We now present the limiting distribution of the proposed estimator. This result relies on the following conditions:

- \hat{D} belongs to a fixed Donsker class (van der Vaart and Wellner, 1996) with probability approaching one;
- $\int [\hat{D}(o) - D(o)]^2 dP(o)$ converges to zero in probability as $n \rightarrow \infty$;
- all remainders are negligible: $\max_k \text{Rem}_k(P, \hat{P}_k^*) = o_P(n^{-1/2})$.

Theorem 3. *Under the above listed conditions, there exists a covariance matrix Σ such that*

$$n^{1/2} \left[\hat{\underline{\Psi}} - \underline{\Psi} \right] = n^{-1/2} \sum_{i=1}^n \underline{D}^P(O_i) + o_P(1) \xrightarrow{d} \text{Normal}(\underline{0}, \Sigma),$$

The proof of this result is omitted because very similar proofs have already appeared in the literature many times (e.g., van der Laan and Rose, 2011; Kennedy, 2016, for overviews). The limiting covariance matrix Σ can be consistently estimated using the empirical covariance of $\underline{D}^{\hat{P}^*}(O)$.

Algorithm 1 TMLE for Estimating $\Psi_1(P)$, $\Psi_2(P)$, and $\Psi_3(P)$

Takes as input n observations $\text{Obs} \equiv \{O_i : i = 1, \dots, n\}$.

function TMLE(Obs)

Initial Estimates: Define an initial estimator \hat{P} of P :

- The marginal distribution of W under \hat{P} should be the empirical.
- Only the estimates of the components of $P(O|W)$ needed to evaluate Ψ_k and D_k , $k = 1, 2, 3$, are needed: $\hat{P}(A = 1|W = \cdot)$, $\hat{P}(Y = 1|S = s_1^*, A = 1, W = \cdot)$, $\hat{P}(S = s_1^*|A = 1, W = \cdot)$, and $\hat{P}(Y = 0, S^c = s_1^*|A = 0, W = \cdot)$.

▷ Conditional probability estimates should fall in $(0, 1)$.

▷ Recall the definitions of f_k and a_k from Theorem 2.

for $k = 1, 2, 3$ **do**

Fluctuation for Targeting Step: Using observations $i = 1, \dots, n$, fit the intercept $\hat{\epsilon}_k$ using an intercept-only logistic regression with outcome $f_k(O_i)$, offset logit $\hat{P}\{f_k(O) = 1|a, w\}$, and weights $\frac{\mathbb{1}_{A_i=a_k}}{\hat{P}(A_i|W_i)}$.

Targeting Step: We now define a fluctuation of \hat{P}_k^* of \hat{P} , targeted towards estimation of ψ_k . Define

$$\hat{P}_k^* \{f_k(O) = 1|a_k, w\} \equiv \text{expit} \{ \text{logit} \hat{P} \{f_k(O) = 1|a_k, w\} + \hat{\epsilon}_k \}.$$

The marginal distribution of $W \sim \hat{P}$ does not need to be fluctuated.

Plug-In Estimator: The estimator of ψ_k is $\hat{\psi}_k \equiv \Psi_k(\hat{P}_k^*)$, and the estimate of the corresponding component of the influence function is $\hat{D}_k \equiv D_k^{\hat{P}_k^*}$.

return Estimate $\hat{\Psi} \equiv (\hat{\psi}_k : k = 1, 2, 3)$ and estimated influence function $\hat{D} \equiv (\hat{D}_k : k = 1, 2, 3)$.

To develop sufficient conditions for $\max_k \text{Rem}_k(P, \hat{P}_k^*)$ to be negligible, one can use that each $\text{Rem}_k(P, \hat{P}_k^*)$ is upper bounded by a constant times the product of the root-mean-square distance between $\hat{P}_k^*(a_k|W = \cdot)$ and $P(a_k|W = \cdot)$ and the root-mean-square distance between $\hat{P}_k^*(f_k(O) = 1|a_k, W = \cdot)$ and $P(f_k(O) = 1|a_k, W = \cdot)$. In a parametric model, each of these root-mean-square distances is $O_P(n^{-1/2})$, and outside of a parametric model sufficient smoothness will enable several estimators to ensure that the product of these two rates is $o_P(n^{-1/2})$. Examples include generalized additive models (Hastie and Tibshirani, 1990; Horowitz, 2009) and the highly adaptive lasso estimator (van der Laan, 2017). If the data are generated from a randomized clinical trial, and the known treatment assignment probabilities are used, then each $\text{Rem}_k(P, \hat{P}_k^*)$ is exactly zero.

The Donsker condition in the theorem above restricts the flexibility of the initial distribution estimate \hat{P} . This restriction may be unpleasant in situations where the dependence of $P\{f_k(O) = 1|a, w\}$ on w is believed to be complicated. In these cases, one can modify the above estimator via a V -fold sample splitting approach, resulting in a cross-validated TMLE (Zheng and van der Laan, 2011). For ease of exposition, we focus on the case that $V = 10$. Here, one partitions the data into ten mutually exclusive and exhaustive folds of approximately equal size. One then fits ten initial estimates of \hat{P} , where each initial estimate leaves out a different one of these ten folds. For each observation i , let $v(i)$ denote the fold that contains i . Let \hat{P}^{v} denote the initial estimate based on the nine out of ten folds that do not include fold v . The sample splitting procedure then replaces each evaluation of (a conditional distribution of) \hat{P} at O_i within the for loop over k by $\hat{P}^{v(i)}$, where we recall that k indexes the three dimensions of the output of the parameter $\underline{\Psi}$. A single, non-fold-specific fluctuation $\hat{\epsilon}_k$ is still returned from the logistic regression in this for loop. The fluctuated initial estimates are still fold-specific, with

$$\hat{P}_k^{v,*} \{f_k(O) = 1|a, w\} \equiv \text{expit} \{ \text{logit} \hat{P}^v \{f_k(O) = 1|a, w\} + \hat{\epsilon}_k \}.$$

The estimator of ψ_k is $\hat{\psi}_k^{CV} \equiv \frac{1}{n} \sum_{i=1}^n \Psi_k(\hat{P}_k^{v(i),*})$. The estimator $\hat{\Psi}^{CV}$ should have the same multivariate normal limit as $\hat{\Psi}$ would have if the Donsker condition is satisfied, but this cross-validated TMLE will still have this multivariate normal limit even if this Donsker condition fails.

As discussed in the introduction, the delta method can be used to derive confidence intervals for any smooth mapping $\Gamma : \mathbb{R}^3 \rightarrow \mathbb{R}$ of $\underline{\psi}$, where the efficient influence function is then $\langle D^P, \dot{\Gamma}(\underline{\psi}) \rangle$, with $\dot{\Gamma}$ the gradient of Γ . For a discrete biomarker, as is studied in this section, this will require that the biomarker takes on the value s_1^* with positive probability, i.e. that $\mathbb{E}[P(S = s_1^*|A = 1, W)] > 0$. For example, for the log of the relative risk defined in (2), $\Gamma(\underline{x}) \equiv \log(x_2) - \log(x_1 - x_3)$ and $\dot{\Gamma}(\underline{x}) = (-[x_1 - x_3]^{-1}, x_2^{-1}, [x_1 - x_3]^{-1})$. For the risk difference on treatment $A = 1$ versus on treatment $A = 0$,

$\Gamma(x) = (x_2 - x_3)/x_1$ and $\dot{\Gamma}(x) = x_1^{-2}(x_3 - x_2, x_1, -x_1)$. In either case, the influence function of our estimator $\hat{\psi}$ can be estimated using $\langle \hat{D}, \dot{\Gamma}(\hat{\psi}) \rangle$. Letting $\hat{\sigma}^2$ denote the empirical variance of $\langle \hat{D}(O), \dot{\Gamma}(\hat{\psi}) \rangle$, a 95% confidence interval for $\Gamma(\hat{\psi})$ is given by $\Gamma(\hat{\psi}) \pm 1.96n^{-1/2}\hat{\sigma}$.

4. Estimation Under Two-Phase Sampling

In some settings, such as in many vaccine studies, the post-treatment biomarker S and the post-crossover biomarker S^c will only be measured on a subset of the individuals in the study. For example, the biomarker S may be measured on all treated ($A = 1$) cases ($Y = 1$) and all members of a prespecified cohort randomly drawn from the set of all individuals, whereas the crossover biomarker S^c may be measured on a random subset of untreated ($A = 0$) controls ($Y = 0$). In this section, we study estimation in the setting where the biomarkers are discrete.

Section 4.1 gives an overview of the notation and assumptions needed in this setting. Section 4.2 presents an inverse probability weighted TMLE based on the framework of [Rose and van der Laan \(2011\)](#). This estimator is appealing because it only requires minor modifications to the estimator from Section 3.2, namely in the inclusion of sampling weights. When the baseline covariates are not discrete, the inverse probability weighted TMLE presented in this section may not be efficient, because it does not fully leverage the predictive power of the baseline covariates for the missing biomarker, which is especially important for individuals for whom these biomarkers are not measured. Though [Rose and van der Laan \(2011\)](#) also presented a TMLE that can attain efficiency in settings where the baseline covariates are continuous, we instead present a one-step estimator for these settings, that has the advantage of being easier to describe and implement, but the disadvantage of not being a plug-in estimator.

4.1. Overview

For treated individuals, let Δ denote the indicator of having the biomarker S measured. For untreated individuals, let Δ denote the indicator of having the crossover biomarker S^c measured. We use the convention that Δ is always one for untreated cases, which does not cause a problem because S^c is defined to be zero for all untreated cases.

We make the missing at random assumptions that the indicator Δ is independent of the biomarkers (S, S^c) given phase-one information (W, A, Y) , where we recall that S is degenerate for untreated individuals and S^c is degenerate for treated individuals and untreated controls. For each individual i , the indicator of having the biomarker measured is given by Δ_i . We suppose that $\Delta_1, \dots, \Delta_n$ is an i.i.d. sequence given (O_1, \dots, O_n) , where we remind the reader that $O_i \equiv (W_i, A_i, S_i, Y_i, S_i^c)$. We let $\tilde{O} \equiv (W, A, \Delta S, Y, \Delta S^c)$ denote the censored data structure, which is equal to O except that the biomarkers are unobserved for all individuals with $\Delta = 0$. Let \tilde{P} denote the corresponding distribution of (Δ, \tilde{O}) . Our estimation scheme will be a function of the observed n i.i.d. draws of $\tilde{O} \sim \tilde{P}$.

Let $\Pi(\tilde{O})$ denote the probability that $\Delta = 1$ given \tilde{O} . For brevity, we let $\pi \equiv \Pi(\tilde{O})$ for an arbitrary $\tilde{O} \sim \tilde{P}$ and, for each individual i , we let $\pi_i \equiv \Pi(O_i)$. We note that $\pi_i = 1$ for all untreated cases. Furthermore, in a case-cohort sampling scheme, $\pi_i = 1$ for all treated cases.

4.2. Inverse Probability Weighted TMLE

In this section, we will assume that π_i is known for each individual. If π_i only relies on W through a discrete coarsening V of W , then the inference that we propose will be valid, albeit conservative, if π_i is replaced by a nonparametric maximum likelihood estimator $\hat{\pi}_i$ of the probability that $\Delta = 1$ conditional on (V, A, Y) . In fact, using this nonparametric maximum likelihood estimator typically leads to asymptotic efficiency gains even in our setting where π_i is known, and never leads to a loss of asymptotic efficiency in this setting (Theorem 2.3 of [van der Laan and Robins, 2003](#)).

Defining the inverse probability weighted (IPW) TMLE requires only minor modifications to Algorithm 1. The modifications to our estimation scheme hinge on the fact that, for any f , $\mathbb{E}_P[f(O)] =$

$\mathbb{E}_{\tilde{P}}[\Delta\pi^{-1}f(\tilde{O})]$. This fact is useful both when f is a loss function, in which case it suggests to include weights $\Delta\pi^{-1}$ in the estimation procedure, and when f is an influence function, in which case it suggests to weight \underline{D}^P in the leading term on the right-hand side of (3). The estimation scheme, which represents a two-phase sampling TMLE as presented in Rose and van der Laan (2011), is presented in Algorithm 3 in Section 4. When obtaining the initial estimate \hat{P} of P , one can use loss-based approaches, e.g. minimizing the empirical mean-squared error or Kullback-Leibler divergence to estimate the needed components of P . The key in the two-phase sampling case is to include the inverse of the (estimated) individual-level probabilities for belonging to the second phase as weights, i.e. including as weights $\Delta_i\pi_i^{-1}$, in whatever loss-based estimation procedure that is used to estimate the needed components of P (see Section 3.2.2 of Rose and van der Laan, 2011).

Confidence intervals for the IPW TMLE can again be constructed using that, under regularity conditions closely related to those of Theorem 3, $n^{1/2}[\hat{\Psi} - \Psi] \xrightarrow{d} \text{Normal}(\mathbf{0}, \tilde{\Sigma})$. The covariance matrix $\tilde{\Sigma}$ can be estimated by the empirical covariance of the inverse weighted influence function $(\Delta_i\tilde{\pi}_i^{-1}\hat{D}_k(\tilde{O}_i) : k = 1, 2, 3)$, where $\tilde{\pi}_i$ is defined in Algorithm 3 as a rescaling of π_i that ensures that $\Delta\tilde{\pi}^{-1}$ has empirical mean one within both the treated and untreated strata of the observations. The only difference relative to the estimation of Σ is that the influence function is now weighted by $\Delta_i\tilde{\pi}_i^{-1}$. Confidence intervals for the relative risk can again be generated using the delta method.

4.3. One-Step Estimator that Fully Leverages Continuous Baseline Covariates

We now present a one-step estimator that fully leverages the predictive power of baseline covariates, treatment status, and endpoint status for the missing values of the biomarker. Under some regularity conditions, this estimator is efficient in the model in which the probability of treatment given covariates is unknown, and is only slightly inefficient in settings where the probability of treatment given covariates is known (Marsh, 2016). In the setting where the treatment mechanism is unknown, the efficient influence function is given by

$$\tilde{D}^{P'}(\tilde{o}) \equiv \delta\pi^{-1}\underline{D}^{P'}(\tilde{o}) + (1 - \delta\pi^{-1})\mathbb{E}_{P'}\left[\underline{D}^{P'}(\tilde{O}) \mid \Delta = 1, a, w, y\right],$$

where $\tilde{o} \equiv (w, a, \delta s, y, \delta s^c)$ and we abuse notation and let $\pi = \Pi(\tilde{o})$. The elements of the above vector rewrite as

$$\begin{aligned} \tilde{D}_k^{P'}(\tilde{o}) &\equiv \frac{\mathbf{1}\{a = a_k\}}{P'(a|w)} \left\{ \delta\pi^{-1}f_k(o) + (1 - \delta\pi^{-1})\mathbb{E}_{P'}[f_k(O) \mid \Delta = 1, a, w, y] - \mathbb{E}_{P'}[f_k(O) \mid a_k, w] \right\} \\ &\quad + \mathbb{E}_{P'}[f_k(o) \mid a_k, w] - \Psi_k(P'), \end{aligned} \quad (6)$$

To define our estimator, we first suppose that we have an estimator \hat{P} of the distribution P of the data that would have been observed had the probability of membership in the second phase of sampling been equal to one for all individuals. This estimator can be obtained via loss-based learning using inverse probability weighting (see Section 3.2 of Rose and van der Laan, 2011). Next, we define our estimator as $\hat{\Psi} \equiv \Psi(\hat{P}) + \frac{1}{n}\sum_{i=1}^n \tilde{D}^{P'}(\tilde{O}_i)$. Let $\hat{\Sigma}_n$ denote the empirical covariance matrix of $\tilde{D}^{P'}(\tilde{O}_i)$, $i = 1, \dots, n$, where here we acknowledge the notational overload with the covariance matrix $\hat{\Sigma}_n$ from Section 3.2 (single-phase sampling). Under some conditions, one can show that $n^{1/2}\hat{\Sigma}_n^{-1/2}[\hat{\Psi} - \Psi(P)] \xrightarrow{d} N(\mathbf{0}, \text{Id})$, where Id denotes the 3×3 identity matrix. One can then use the delta method to develop confidence intervals for any contrast based on the principally stratified means $\underline{\Psi}$, such as the log relative risk.

To gain some intuition on this estimator, note that, if the baseline covariates are strongly predictive of the biomarker among treated individuals and of the crossover biomarker among uninfected untreated individuals, then the estimate of $\mathbb{E}_P[f_k(O) \mid \Delta = 1, a, w, y]$ will be approximately equal to $f_k(o)$. Therefore, $\delta\pi^{-1}f_k(o) + (1 - \delta\pi^{-1})\mathbb{E}_{P'}[f_k(O) \mid \Delta = 1, a, w, y]$ is approximately equal to $f_k(o)$, so that $\tilde{D}_k^{P'}(\tilde{o})$ is approximately equal to the efficient influence function in the setting where phase two information is observed on everyone.

Finally, we note that, though we have assumed that Π is known, the proposed one-step estimator will also yield valid inference if Π is replaced by an estimate and a doubly robust term is negligible.

5. Continuous Biomarker

5.1. Algorithm and Theoretical Guarantees

Thus far we have assumed that the biomarker is discrete. Suppose now that the biomarkers S_1, S_0^c are continuous, with support in \mathbb{R} . We will show the methods that we have proposed immediately extend to this case once one replaces the indicators that $S = s_1^*$ and $S^c = s_1^*$ by kernels. The extension to the case that the biomarkers have support in \mathbb{R}^d , $d > 1$, is straightforward by using kernels for \mathbb{R}^d -valued data, though, given that kernel smoothers are highly susceptible to the curse of dimensionality, we expect that this method will only yield informative estimates when d is small. For simplicity, we consider the full sampling case where all treated individuals and all untreated individuals without the event have their biomarker measured. As in the previous section, our objective will be to estimate the effect of treatment conditional on a biomarker value s_1^* .

This treatment effect will be a function of three parameters, which are defined with respect to the (conditional) Lebesgue densities of the biomarkers S and S^c , which are assumed to exist. To avoid introducing new notation for these densities, we will abuse notation so that writing “ $S = s$ ” or “ $S^c = s$ ” in a conditional probability statement refers to the conditional density of S or S^c at s . Once we have made this abuse of notation, the definition of $\underline{\Psi}(P)$ from Section 2 does not need to be changed to study the setting where the biomarkers are continuous. So, for example, $\Psi_1(P') \equiv \mathbb{E}_{P'}[P'(S = s_1^* | A = 1, W)]$ represents the Lebesgue density of S at s_1^* , conditional on $A = 1$ and W , averaged across values of $W \sim P'$.

Let K denote a kernel in \mathbb{R} , where this function is centered at zero and integrates to one. Examples of kernels include $K(x) = \mathbb{1}\{-1 \leq 2x \leq 1\}$ (uniform kernel), $K(x) \propto \exp(-x^2/2)$ (Gaussian kernel), and $K(x) \propto \frac{1}{2}(3 - x^2)\exp(-x^2/2)$ (fourth-order Gaussian kernel), where the constants in the ‘ \propto ’ statements are chosen so that the kernels integrate to one. The first two kernels are second-order in the sense that $\int x^t K(x) dx = 0$ for $t = 1$, whereas the final kernel is fourth-order in the sense that $\int x^t K(x) dx = 0$ for $t = 1, 2, 3$. Generally, a kernel is of order r if $\int x^t K(x) dx = 0$ for all $t = 1, \dots, r - 1$. We define K_h as the function $x \mapsto h^{-1}K[xh^{-1}]$, where $h > 0$ is the bandwidth.

Before presenting our algorithm, we introduce the kernel-dependent pseudo-outcomes that we will use to replace f_k , $k = 1, 2, 3$. In particular, we define $f_{1,h}(O_i) \equiv K_h(S - s_1^*)$, $f_{2,h}(O_i) \equiv \mathbb{1}_{Y=1} K_h(S - s_1^*)$, and $f_{3,h}(O_i) \equiv \mathbb{1}_{Y=0} K_h(S^c - s_1^*)$. We continue to use the definitions $a_1 = 1$, $a_2 = 1$, and $a_3 = 0$. For an arbitrary P' , we also define $D_{k,h}^{P'}$ analogously to $D_k^{P'}$ from Theorem 2, but with each instance of f_k replaced by $f_{k,h}$. We also define a smoothed version of our parameter. To do this, it will be useful to make the dependence of $\underline{\Psi}(P)$ on s_1^* explicit in the notation: in particular, we define $\underline{\Psi}(P; s) \equiv \{\Psi_k(P; s) : k = 1, 2, 3\}$ for all s in the support of $S|A = 1$ under P . We now define and give an equivalent expression for the kernel-smoothed version of $\underline{\Psi}(P; s_1^*)$:

$$\underline{\Psi}_h(P) \equiv \int \underline{\Psi}(P; s) K_h(s - s_1^*) ds = \mathbb{E}[\mathbb{E}[f_{k,h}(O) | a_k, W]]. \tag{7}$$

The equivalent expression relies on changing the order of integration between the integral over s and the integral over $W \sim P$. We write the three components of $\underline{\Psi}_h(P)$ as $\Psi_{k,h}(P)$ for $k = 1, 2, 3$. It will be convenient to define $q_{k,h}$ as the function $w \mapsto \mathbb{E}[f_{k,h}(O) | a_k, w]$.

Our proposed estimation procedure is presented in Algorithm 2. In our procedure, we estimate $q_{k,h}$ for each $k = 1, 2, 3$, and so by the latter identity in (7) it follows that we can use these expectations to evaluate $\underline{\Psi}_h$ at our estimated distribution.

Our estimator can be analyzed using standard arguments for kernel estimators. We will break our analysis into a study of the bias and a study of the variance of our estimator. We will show that our estimator has small bias for the smoothed principally stratified mean vector $\underline{\Psi}_h(P)$ at a user-defined value of h . Therefore, the bias of our estimator for $\underline{\Psi}(P)$ will be driven by the convergence of a kernel-smoothed parameter to the true target parameter. As is standard for kernel estimators, the rate of convergence of the bias will be driven by the differentiability of $s \mapsto \underline{\Psi}(P; s)$ at s and by the order of the kernel that we select.

Algorithm 2 CV-TMLE for Estimating $\Psi_1(P)$, $\Psi_2(P)$, and $\Psi_3(P)$ when the biomarker is continuous

Takes as input n observations $\text{Obs} \equiv \{O_i : i = 1, \dots, n\}$ and a kernel K_h .

For simplicity, focus on 10-fold cross-validation.

function CV-TMLE(Obs)

for $v = 1, \dots, 10$ **do**

Initial Estimates: Using data from training set v , define an initial estimator \hat{P}^v of P :

- The marginal distribution of W under \hat{P}^v should be the empirical of the observations in training set v .
- Using only observations in the training set v , compute an estimate $w \mapsto \hat{P}^v(A = 1|w)$ of $w \mapsto P^v(A = 1|w)$, and an estimate $w \mapsto \hat{q}_{k,h}^v(w)$ of $w \mapsto q_{k,h}(w)$.

▷ Each $q_{k,h}$ should have range in $(0, \infty)$.

for $k = 1, 2, 3$ **do**

Fluctuation for Targeting Step: Define $\hat{\epsilon}_k$ as a minimizer (in real-valued ϵ) of

$$\left(\frac{1}{n} \sum_{i=1}^n \frac{\mathbb{1}_{A_i=a_k}}{\hat{P}^v(A_i|W_i)} \left[f_{k,h}(O_i) - \exp \left\{ \log \hat{q}_{k,h}^v(W_i) + \epsilon \right\} \right] \right)^2.$$

Targeting Step: For each fold v , define

$$\hat{q}_{k,h}^{v,*}(w) \equiv \exp \left\{ \log \hat{q}_{k,h}^v(w) + \hat{\epsilon}_k \right\}.$$

The marginal distribution of $W \sim \hat{P}$ does not need to be fluctuated.

Estimator: The estimator of $\psi_{k,h}$ is $\frac{1}{n} \sum_{i=1}^n \Psi_{k,h}(\hat{P}_{k,h}^{v(i),*})$, and, for each individual i , the estimate of the corresponding component of the smoothed influence function is $\hat{D}_{k,h,i} \equiv D_{k,h}^{\hat{P}_{k,h}^{v(i),*}}$.

return Estimate $\hat{\Psi}_h \equiv (\hat{\Psi}_{k,h} : k = 1, 2, 3)$ and estimated smoothed influence function $\hat{D}_{h,i} \equiv (\hat{D}_{k,h,i} : k = 1, 2, 3)$.

Theorem 4 (Bias for Unsmoothed Parameter). *Let K be an r^{th} -order kernel and suppose that, for each $k \in \{1, 2, 3\}$, $s \mapsto \Psi_{k,h}(P; s)$ is $(t+1)^{\text{th}}$ -order differentiable on the real line with uniformly bounded $(t+1)^{\text{th}}$ derivative. If $\int |u^{\min\{r,t+1\}} K(u)| du < \infty$, then $\|\underline{\Psi}_h(P) - \underline{\Psi}(P)\| = O(h^{\min\{r,t+1\}})$.*

Next, we study the variance of our estimator, which is (asymptotically) unbiased for the parameter evaluated at the kernel-smoothed data generating distribution. Before presenting this result, we will define the covariance matrix Σ_n to be used in the upcoming theorem. For \bar{q}_k equal to the limit of the estimates of q_{k,\hat{h}_n} as $n \rightarrow \infty$ (see the upcoming Condition A4) and $\epsilon \in \mathbb{R}$, we let

$$g_{k,\epsilon,\bar{h}_n}(o) \equiv \frac{\mathbb{1}_{a=a_k}}{P(a_k|w)} \left[f_{k,\bar{h}_n}(o) - \exp \{ \log \bar{q}_k(w) + \epsilon \} \right] + \exp \{ \log \bar{q}_k(w) + \epsilon \} \quad (8)$$

and define Σ_n as the 3×3 matrix representing the covariance matrix under $O \sim P$ of $(\bar{h}_n^{1/2} g_{k,\bar{\epsilon}_k,\bar{h}_n}(O) : k = 1, 2, 3)$, where $\bar{\epsilon}_k$, $k = 1, 2, 3$, is defined in the upcoming Condition A9. The following theorem relies on conditions that are stated in the upcoming Section 5.2.

Theorem 5 (Estimation of Smoothed Parameter). *If Conditions A4 through A10 hold, then*

$$(n\hat{h}_n)^{1/2} \Sigma_n^{-1/2} \left[\hat{\Psi}_{\hat{h}_n} - \underline{\Psi}_{\hat{h}_n}(P) \right] \xrightarrow{d} N(\underline{0}, \text{Id}). \quad (9)$$

The proof of this result is given in Appendix C. Under mild conditions, the Σ_n matrices will be consistently estimable via the cross-validated empirical covariance matrix $\hat{\Sigma}_n \equiv \frac{1}{n} \sum_{i=1}^n \hat{D}_{\hat{h}_n,i} [\hat{D}_{\hat{h}_n,i}]^T$.

Note that $\underline{\Psi}_{\hat{h}_n}(P)$ is a data adaptive parameter (van der Laan et al., 2013), in the sense that it is random through the bandwidth \hat{h}_n . One of the conditions for the theorem, listed in the next subsection, is that there exists a deterministic sequence \bar{h}_n such that $\hat{h}_n/\bar{h}_n \rightarrow 1 + o_P(\bar{h}_n^{1/2})$. In fact, many selection procedures satisfy the (typically stronger) condition that $\hat{h}_n/\bar{h}_n - 1 = O_P(n^{-1/2})$ for a deterministic sequence \bar{h}_n (Chiu, 1991; Hall et al., 1991; Fan and Marron, 1992). Under these stronger conditions, the difference between the data adaptive smoothed parameter at bandwidth \hat{h}_n and the sample size dependent, but deterministic, smoothed parameter at bandwidth \bar{h}_n is of the order $n^{-1/2}$, so that one can replace the data adaptive parameter by $\underline{\Psi}_{\bar{h}_n}(P)$ in the above theorem without changing the result.

An estimator of the smoothed relative risk $RR_h^P(s_1^*) \equiv \Psi_{2,h}(P)/[\Psi_{3,h}(P) - \Psi_{1,h}(P)]$ is given by $\widehat{RR}_{\hat{h}_n}(s_1^*) \equiv \hat{\Psi}_{2,\hat{h}_n}/[\hat{\Psi}_{3,\hat{h}_n} - \hat{\Psi}_{1,\hat{h}_n}]$. Similar delta method arguments used to develop confidence intervals for discrete biomarkers can be used in the continuous setting. Then, Wald-type confidence intervals for the log smoothed relative risk can be defined by using that, under the conditions of Theorem 5, $(n\hat{h}_n)^{1/2} \hat{\sigma}_{\hat{h}_n}^{-1} [\log \widehat{RR}_{\hat{h}_n}(s_1^*) - \log RR_h^P(s_1^*)]$ converges to a standard normal distribution, where, for $\dot{\Gamma}(\underline{x}) = (-x_2^{-1}, -[x_1 - x_3]^{-1}, [x_1 - x_3]^{-1})$, we define

$$\hat{\sigma}_{\hat{h}_n}^2 \equiv \frac{1}{n} \sum_{i=1}^n \left[\langle \hat{D}_{\hat{h}_n,i}, \dot{\Gamma}(\hat{\Psi}_{\hat{h}_n}) \rangle - \frac{1}{n} \sum_{j=1}^n \langle \hat{D}_{\hat{h}_n,j}, \dot{\Gamma}(\hat{\Psi}_{\hat{h}_n}) \rangle \right]^2.$$

Combining the above two theorems demonstrates that there are rates for \bar{h}_n such that asymptotically normal inference is possible for the unsmoothed parameter $RR^P(s_1^*)$. Indeed, balancing the standard error of $\hat{\Psi}_{\hat{h}_n}$, i.e. $O((n\hat{h}_n)^{-1/2})$, and its bias for estimating the unsmoothed parameter, i.e. $O(h^{\min\{r,t+1\}})$ for r, t defined in Theorem 4, shows that the bias converges more quickly than the standard error if $h = o(n^{-1/(2\min\{r,t+1\}+1)})$. Nonetheless, selecting these rates for \bar{h}_n requires undersmoothing, i.e. choosing a rate for \bar{h}_n that is slower than the rate that is optimal according to a criterion such as the mean-squared error of a density estimator at $S = s_1$, conditional on $A = 1$. Following the estimation of the causal effect of continuous treatments in Kennedy et al. (2016), who himself followed the suggestion of Wasserman (2006), we focus our inference solely on the smoothed relative risk parameter, rather than on the original, unsmoothed parameter. The rate of decay on the bias from the above theorem is still interesting: it shows that this smoothed parameter is getting close to the unsmoothed parameter as the sample size grows. Nonetheless, we do not attempt to quantify the proximity of the smoothed parameter to the unsmoothed parameter in the confidence interval that we construct. Once one has adopted this perspective, one has flexibility in selecting the bandwidth h . For example, one could use a cross-validated mean integrated squared error criterion, where this criterion is selected for estimating the marginal density of S among all treated individuals.

An alternative approach would be to use the recent work of (Bibaut and van der Laan, 2017), that enables principled selection of an undersmoothed bandwidth for non-pathwise-differentiable parameters such as $\underline{\Psi}$. Consideration of the performance of this method for the estimation of $\underline{\Psi}$ is beyond the scope of this work.

5.2. Conditions for Asymptotically Normal Estimation of the Smoothed Parameter

Theorem 5 used the following conditions to prove the asymptotic normality of the estimator for continuous biomarkers:

- (A4) The treatment mechanism estimates satisfy $\max_v \|\hat{P}^v(a|W = \cdot) - P(a|W = \cdot)\| = o_P(1)$. Furthermore, for $k = 1, 2, 3$, there exists a function \bar{q}_k such that $\max_v \|\hat{q}_{k,\hat{h}_n}^v - \bar{q}_k\| = o_P(1)$.
- (A5) There exists a deterministic positive sequence $\{\bar{h}_n\}_{n=1}^\infty$ such that $\hat{h}_n/\bar{h}_n \rightarrow 1 + o_P(\bar{h}_n^{-1/2})$ and $n\bar{h}_n$ diverges to infinity.
- (A6) $\max_v \|\hat{q}_{k,h}^v - q_{k,h}\| \|\hat{P}^v(a_k|W = \cdot) - P(a_k|W = \cdot)\| = o_P(n^{-1/2}\hat{h}_n^{-1/2})$.
- (A7) The kernel K can be evaluated using a composition of a finite number of arithmetic operations (+, −, ÷, ×), indicator functions ($x \mapsto \mathbb{1}\{x > c\}$ for a constant c), and exponential functions ($x \mapsto e^x$). Furthermore, the kernel K is bounded and uniformly continuous.
- (A8) There exists a constant $c > 0$ such that $P\{c^{-1} < \min_{k,v} \hat{q}_{k,h}^v(W) \leq \max_{k,v} \hat{q}_{k,h}^v(W) < c\} = 1 - o_P(1)$ and a constant $\delta > 0$ such that $P\{\max_{k,v,a} 1/\hat{P}_k^v(a|W) < \delta^{-1}\} = o_P(1)$.
- (A9) For each $k \in \{1, 2, 3\}$, there exists a fixed $\bar{e}_k \in \mathbb{R}$ such that $\hat{e}_{k,n} \rightarrow \bar{e}_k$ in probability as $n \rightarrow \infty$, where here (and here only) we make the dependence of \hat{e}_k on n explicit in the notation by denoting this quantity by $\hat{e}_{k,n}$.

- (A10) There exists some fixed $\delta > 0$ such that, with probability approaching one, all of the eigenvalues of Σ_n are bounded below by δ .

We now discuss these conditions.

Condition A4. This condition requires mean-square consistency of the needed nuisance parameter estimates, where we require this consistency for the estimators within each training fold. If the same estimators are used within each training fold, then this is equivalent to requiring mean-square consistency of the estimation of the needed nuisance parameters on the full sample. This condition does not require any rate for this convergence, and so is a minor assumption relative to A6, though we note that A4 is not necessarily implied by A6 (e.g., if the treatment mechanism is known and one uses this information in estimation). Note that here we have assumed that the treatment mechanism estimates converge to the true treatment mechanism in mean-square, whereas we have only required that \hat{q}_{k,\hat{h}_n}^y has some limit. The condition on the treatment mechanism seems minor given that the treatment mechanism is known in most settings where crossover designs would be employed.

Condition A5. This condition requires that the bandwidth selection procedure that is employed will eventually return a (nearly) deterministic result. The procedures of (Chiu, 1991; Hall et al., 1991) yield data-driven bandwidth selections \hat{h}_n that satisfy $\hat{h}_n/h_n^\circ = 1 + O_P(n^{-1/2})$, where here h_n° is the optimal bandwidth in terms of mean integrated squared error. In fact, Fan and Marron (1992) demonstrates that the $O_P(n^{-1/2})$ term quantifying the multiplicative deviation of these bandwidth selection procedures from the optimal bandwidth is optimal even in terms of the leading constant. Therefore, if one selects \hat{h}_n to optimize an estimate of the density of S among treated individuals, we expect that the selected bandwidth will eventually behave as a deterministic sequence and therefore that A5 will be satisfied.

Condition A6. This condition requires that product of the estimates of the treatment mechanism and the smoothed densities $q_{k,h}$ converge to zero faster than does the standard error of our estimator. If the treatment mechanism is known, as is the case in a randomized trial, then one can ensure that this product is exactly zero at every sample size. Otherwise, the plausibility of this condition will typically rely on the dimensionality of W . As a reference point, we note that, if a parametric model could be correctly specified for an unknown treatment mechanism, then this condition would also automatically be satisfied if $\hat{h}_n \rightarrow 0$ in probability, since in this case the product from this condition would be $O_P(n^{-1/2})$.

Condition A7. This condition does not rely on the data or the data generating distribution and therefore can be verified in applications. We use this condition to ensure certain empirical process and envelope conditions that we use in our proof are satisfied. This condition is satisfied by most kernels used in practice. For example, it is satisfied by $K(x) = \mathbb{1}\{-1/2 \leq x \leq 1/2\}$ (uniform kernel), $K(x) = \exp(-x^2/2)$ (Gaussian kernel), and $K(x) = \frac{1}{2}(3 - x^2)\exp(-x^2/2)$ (fourth-order Gaussian kernel).

Condition A8. This assumption ensures that our estimates both of the density $q_{k,h}$ and of the treatment mechanism do not fall too close to zero or one, and can be enforced in practice via truncation. For the truncation not to damage the chance of the consistency of these estimates in A6, Condition A8 requires that the true treatment mechanism does not fall too close to zero or one. This will be plausible for the treatment mechanism in a randomized trial in which all individuals have a non-negligible probability of receiving or not receiving treatment. The assumption on the density $q_{k,h}$ requires that, within each stratum of covariates, any given realization of the biomarker is not very rare or very likely. Furthermore, the assumption on $q_{k,h}$ requires that the probability of having $Y = 1$ within baseline covariate strata of individuals not receiving treatment at baseline is never too common (this is to be expected for rare events), and also that the event probability is bounded away from zero among treated individuals within each stratum of baseline covariates.

Condition A9. The existence of a limit \bar{e}_k for \hat{e}_k can be ensured under mild conditions using arguments for Z-estimators (van der Vaart and Wellner, 1996). We omit these arguments here for brevity.

Condition A10. This condition aims to ensure that the limit is non-degenerate. Its plausibility can be checked in practice by checking the size of the eigenvalues of the estimator $\hat{\Sigma}_n$ of Σ_n , given by $\hat{\Sigma}_n \equiv \frac{1}{n} \sum_{i=1}^n \hat{D}_{\hat{h}_n}(O_i) [\hat{D}_{\hat{h}_n}(O_i)]^T$, and ensuring that the smallest eigenvalue is not close to zero.

6. Simulation

We conduct simulations similar to the settings in [Gilbert and Hudgens \(2008\)](#), though with continuous rather than discrete baseline covariates and biomarker. Here we focus on a univariate post-treatment biomarker because such biomarkers are often of interest in vaccine studies (see [Corey et al., 2015](#), for examples). We consider sample size 5000 with equal number of people in treatment $A = 1$ and control $A = 0$. We generate (W, S_1) from bivariate normal distribution with mean $(0.41, 0.41)^T$ and covariance matrix

$$\begin{pmatrix} 0.55^2 & 0.55^2 \cdot 0.5 \\ 0.55^2 \cdot 0.5 & 0.55^2 \end{pmatrix}.$$

For $a = 0, 1$, we generate the potential outcome from the logistic model

$$P^F(Y_a = 1 | W, S_1) = \text{expit}(\beta_0 + \beta_1 a + \beta_2 W + \beta_3 S_1 + \beta_4 a S_1)$$

where $\text{expit}(x) \equiv \exp(x)/(1 + \exp(x))$ and $\beta_0 = 1.2, \beta_1 = -0.6, \beta_2 = -0.5, \beta_3 = -0.1, \beta_4 = -0.9$, such that the average disease rate is 9.5% among individuals receiving $A = 1$ and is 18.8% among individuals receiving $A = 0$. We consider a full sampling setting where all treated individuals ($A = 1$) and all untreated controls ($A = 0, Y = 0$) have the biomarker measured. We consider Gaussian kernel $K_h(x) = \frac{1}{\sqrt{2\pi}h} \exp(-\frac{x^2}{2h^2})$ to smooth $\mathbb{1}\{S_1 = s_1^*\}$. We use R package SuperLearner ([Polley and van der Laan, 2012](#)) to estimate the kernel-smoothed $\Psi_h(P)$ in equation (7). We consider GLM, GLM with interaction, stepwise regression, neural network and sample mean in our SuperLearner library, whose corresponding wrapper functions in the SuperLearner package are respectively given by `SL.glm`, `SL.glm.interaction`, `SL.step`, `SL.nnet`, `SL.mean`. All of these functions were run at their default settings in version 2.0-22 of the SuperLearner package.

We calculate two versions of the true parameters: one without smoothing and one with smoothing. The true parameter without smoothing can be derived from the following identity

$$\begin{aligned} \Psi_1(P; s_1^*) &= P^F(S_1 = s_1^*) \\ \Psi_2(P; s_1^*) &= \int P^F(Y_1 = 1 | S_1 = s_1^*, W = w) P^F(S_1 = s_1^*, W = w) dw \\ \Psi_3(P; s_1^*) &= \int P^F(Y_0 = 0 | S_1 = s_1^*, W = w) P^F(S_1 = s_1^*, W = w) dw \end{aligned}$$

The target parameter is log relative risk $\Psi_h = \log(\frac{\Psi_{2,h}}{\Psi_{1,h} - \Psi_{3,h}})$. The smoothed version of the true parameter is calculated by (7). We calculate the 95% Wald-type confidence interval and its coverage over the unsmoothed true parameter and smoothed true parameter across 1000 simulations with different kernel bandwidth h . Figure 1 shows the confidence interval coverage for a fixed $s_1^* = 0.6$ with different choices of h . The coverage probability for the unsmoothed vaccine efficacy parameter drops below the nominal level when h is small or large. The coverage probability for the smoothed parameter drops below the nominal level when h is small but does not drop below nominal level when h is large. The results are similar for other values of s_1^* .

The primary results that we present in the main text will be based on a bandwidth of 0.2. We selected this bandwidth by first noting the three places that the bandwidth is used in our procedure: first, to estimate the density of $S|A = 1, W$ that appears in the definition of Ψ_1 ; next, to estimate the density of $S|A = 1, Y = 1, W$ that appears in the definition of Ψ_2 ; and, finally, to estimate the density of $S^c|A = 0, Y = 0, W$ that appears in the definition of Ψ_3 . To simplify selection of the bandwidth, we considered the estimation of the corresponding densities when the covariate W was not conditioned on, namely the estimation of the densities of $(S|A = 1)$, $(S|A = 1, Y = 1)$, and $(S^c|A = 0, Y = 0)$. Since Y is a rare outcome, the sparsest stratum in these conditioning statements will occur in the second listed density, namely the density of $S|A = 1, Y = 1$. Consequently, we expect that the bandwidth selection should be driven by the estimation of Ψ_2 . To select the bandwidth for our simulations, we generated many data sets of size 5000, estimated the density of $S|A = 1, Y = 1$, and evaluated the default bandwidth selected

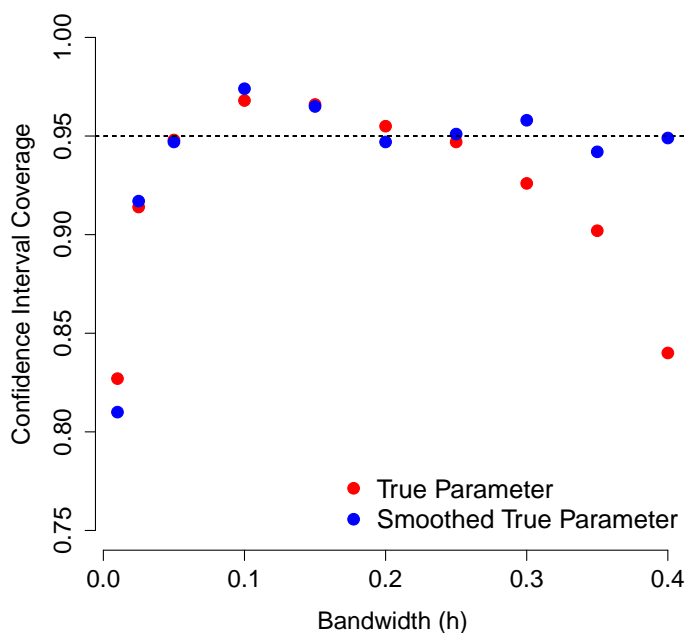


FIGURE 1. The coverage probability for the unsmoothed vaccine efficacy parameter drops below the nominal level when h is small or large. The coverage probability for the smoothed vaccine efficacy parameter drops below the nominal level when h is small but does not drop below nominal level when h is large.

by the density function in R ; this default bandwidth is selected based on Eq. 3.31 of Silverman (1986). We found that values of approximately 0.15 were typically selected. Because the densities used in the definitions of Ψ_1 , Ψ_2 , and Ψ_3 are defined conditionally on W , we inflated this selected bandwidth to 0.2 to define the primary bandwidth considered in the main text. In Appendix F we also present results at bandwidths of 0.1, 0.3, and 0.4.

Table 1 shows the bias, standard error, and coverage probability of the unsmoothed and smoothed truth for different s_1^* with bandwidth chosen to be 0.2. The proposed TMLE estimator has low bias and is asymptotically unbiased. As expected, the coverage probability for the smoothed vaccine efficacy parameter is better than the coverage probability for the unsmoothed vaccine efficacy parameter. Our additional simulation results for bandwidths 0.1, 0.3, and 0.4 are given in Appendix F. These results continue to show appropriate coverage of the smoothed parameter, though, as expected, when the bandwidth is large we see poorer coverage of the unsmoothed parameter.

s_1^*	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Bias, Truth	-0.02	-0.01	-0.00	0.01	0.02	0.03	0.04	0.05	0.07	0.09
Bias, Smoothed	-0.01	-0.02	-0.02	-0.01	-0.01	-0.01	-0.01	-0.00	0.01	0.03
Standard Error	0.18	0.17	0.17	0.17	0.17	0.18	0.19	0.21	0.23	0.27
Coverage, Truth	0.95	0.95	0.95	0.95	0.94	0.95	0.95	0.96	0.96	0.97
Coverage, Smoothed	0.96	0.95	0.95	0.94	0.93	0.94	0.95	0.95	0.96	0.96

TABLE 1. Bias, standard error, and coverage probability of the log relative risk point estimator and confidence intervals for different s_1^* values with bandwidth chosen to be 0.2.

7. Discussion

We have presented a targeted minimum loss-based estimator (TMLE) for estimating contrasts between the principally stratified means of absorbent endpoints in a crossover design, where we recall that an absorbent endpoint is an endpoints whose occurrence renders any subsequent measures of the biomarker of

interest scientifically uninteresting. For most biomarkers, death is an absorbent endpoint. In HIV vaccine trials, HIV infection is an absorbent endpoint when the biomarker of interest is the immune response. We have established sufficient conditions for the nonparametric identifiability of treatment-specific principally stratified means. Our identifiability conditions do not require that the baseline covariates be discrete. We also established a necessary and sufficient condition for the falsifiability of our assumption that relates the conditional distribution of the biomarker in the population that was treated at baseline to the conditional distribution of the biomarker in the population that was treated at the crossover stage. An implication of this result is that the crossover assumption can be tested: if the test fails to reject, then there is no evidence in the data that the crossover assumption is falsifiable. If the test rejects, then there is evidence in the data that the crossover assumption is false.

Our proposed method does not handle right-censoring on the absorbent endpoint. When dropout is rare, this will at most induce a small amount of bias in an estimate of the principally stratified mean outcomes. When dropout is more common but rare enough that a reasonable proportion of individuals have follow-up to the end of the study, the methods in this work can be readily extended by incorporating inverse probability of censoring weights to the proposed estimators. When few individuals have follow-up until the end of the study, further methodological development is needed to develop low-variance estimators that account for dropout.

When the biomarker is continuous, our proposed method relies on a choice of bandwidth. Though our theoretical results provide guarantees for a wide range of sequences of selected bandwidths, we have given only limited discussion to how the bandwidth should be selected in practice. We did present one heuristic strategy in our Section 6, which involves using existing bandwidth selection approaches for estimating the conditional density of S given that $A = 1$ and $Y = 1$ – this density is related to a density that appears in the definition of the parameter Ψ_2 . We leave the theoretical analysis of this and other bandwidth selection strategies to future work.

Our proposed estimators are especially advantageous in settings where (i) the baseline covariates are predictive of the post-treatment biomarker and (ii) the relationship between these baseline covariates and the biomarker is complex. The predictiveness of the baseline covariates improves the plausibility of the identifiability assumption linking the biomarker measured in treated individuals to the biomarker measured in untreated individuals following their crossover to treatment. In these settings, obtaining statistical inference for the contrast of interest using existing methods relies on either ignoring the baseline covariates, discretizing the baseline covariates, or assuming a parametric relationship between the biomarker and these baseline covariates. Our proposed TMLE is especially compelling in these settings because it enables the incorporation of flexible estimation techniques for the conditional distribution of the biomarker given the baseline covariates, while still admitting statistical inference.

Acknowledgements

This work was partially supported by the National Institute of Allergy and Infectious Disease at the National Institutes of Health under award number UM1 AI068635. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. Alex Luedtke gratefully acknowledges the support of the New Development Fund of the Fred Hutchinson Cancer Research Center. The authors thank Peter Gilbert for helpful comments that greatly improved the quality of the manuscript.

References

- Anthony, M. and Bartlett, P. L. (1999). Neural network learning: theoretical foundations.
- Bareinboim, E. and Pearl, J. (2012). Transportability of Causal Effects: Completeness Results. In *Twenty-Sixth AAAI Conference on Artificial Intelligence*.
- Bibaut, A. F. and van der Laan, M. J. (2017). Data-adaptive smoothing for optimal-rate estimation of possibly non-regular parameters. *arXiv preprint arXiv:1706.07408*.
- Bickel, P. J., Klaassen, C. A. J., Ritov, Y., and Wellner, J. A. (1993). *Efficient and adaptive estimation for semiparametric models*. Johns Hopkins University Press, Baltimore.

- Chiu, S.-T. (1991). Bandwidth selection for kernel density estimation. *The Annals of Statistics*, pages 1883–1905.
- Cole, S. R. and Hernán, M. A. (2002). Fallibility in estimating direct effects. *International journal of epidemiology*, 31(1):163–165.
- Corey, L., Gilbert, P. B., Tomaras, G. D., Haynes, B. F., Pantaleo, G., and Fauci, A. S. (2015). Immune correlates of vaccine protection against hiv-1 acquisition. *Science translational medicine*, 7(310):310rv7–310rv7.
- Fan, J. and Marron, J. S. (1992). Best possible constant for bandwidth selection. *The Annals of Statistics*, pages 2057–2070.
- Follmann, D. (2006). Augmented designs to assess immune response in vaccine trials. *Biometrics*, 62(4):1161–1169.
- Frangakis, C. E. and Rubin, D. B. (2002). Principal stratification in causal inference. *Biometrics*, 58(1):21–29.
- Gabriel, E. E. and Follmann, D. (2016). Augmented trial designs for evaluation of principal surrogates. *Biostatistics*, 17(3):453–467.
- Gabriel, E. E. and Gilbert, P. B. (2013). Evaluating principal surrogate endpoints with time-to-event data accounting for time-varying treatment efficacy. *Biostatistics*, 15(2):251–265.
- Gilbert, P. B. and Hudgens, M. G. (2008). Evaluating candidate principal surrogate endpoints. *Biometrics*, 64(4):1146–1154.
- Gilbert, P. B., Hudgens, M. G., and Wolfson, J. (2011). Commentary on "Principal Stratification—a Goal or a Tool?" by Judea Pearl. *Int J Biostat*, 7(1):1–15.
- Hall, P., Sheather, S. J., Jones, M. C., and Marron, J. S. (1991). On optimal data-based bandwidth selection in kernel density estimation. *Biometrika*, 78(2):263–269.
- Hastie, T. J. and Tibshirani, R. J. (1990). *Generalized additive models*. Chapman & Hall, London.
- Horowitz, J. L. (2009). *Semiparametric and nonparametric methods in econometrics*, volume 12. Springer.
- Huang, Y., Gilbert, P. B., and Wolfson, J. (2013). Design and estimation for evaluating principal surrogate markers in vaccine trials. *Biometrics*, 69(2):301–309.
- Kennedy, E. H. (2016). Semiparametric theory and empirical processes in causal inference. In *Statistical Causal Inferences and Their Applications in Public Health Research*, pages 141–167. Springer.
- Kennedy, E. H., Ma, Z., McHugh, M. D., and Small, D. S. (2016). Non-parametric methods for doubly robust estimation of continuous treatment effects. *J. Royal Stat. Soc Series B*.
- Marsh, T. (2016). Efficient inference for an additive gene-treatment interaction from a nested two-phase study.
- Nason, M. and Follmann, D. (2010). Design and analysis of crossover trials for absorbing binary endpoints. *Biometrics*, 66(3):958–965.
- Pearl, J. (2001). Direct and indirect effects. In *Proceedings of the Seventeenth Conference on Uncertainty in artificial intelligence*, pages 411–420, San Francisco. Morgan Kaufmann.
- Pfanzagl, J. (1990). *Estimation in semiparametric models*. Springer.
- Polley, E. and van der Laan, M. (2012). Super Learner Prediction.
- Robins, J. M. and Greenland, S. (1992). Identifiability and exchangeability for direct and indirect effects. *Epidemiol*, 3:143–155.
- Rose, S. and van der Laan, M. J. (2011). A Targeted Maximum Likelihood Estimator for Two-Stage Designs. *Int J Biostat*, 7(17).
- Silverman, B. W. (1986). *Density Estimation for Statistics and Data analysis*. Chapman & Hall.
- van der Laan, M. (2017). A Generally Efficient Targeted Minimum Loss Based Estimator based on the Highly Adaptive Lasso. *Int J Biostat*.
- van der Laan, M. J., Hubbard, A. E., and Kherad, S. (2013). Statistical inference for data adaptive target Parameters. Technical Report 314, Division of Biostatistics, University of California, Berkeley.
- van der Laan, M. J. and Robins, J. M. (2003). *Unified methods for censored longitudinal data and causality*. Springer, New York Berlin Heidelberg.
- van der Laan, M. J. and Rose, S. (2011). *Targeted Learning: Causal Inference for Observational and Experimental Data*. Springer, New York, New York.
- van der Vaart, A. W. and Wellner, J. A. (1996). *Weak convergence and empirical processes*. Springer, Berlin Heidelberg New York.
- VanderWeele, T. (2015). *Explanation in causal inference: methods for mediation and interaction*. Oxford University Press.
- VanderWeele, T. J. (2008). Simple relations between principal stratification and direct and indirect effects. *Statistics & Probability Letters*, 78(17):2957–2962.
- Wasserman, L. (2006). *All of Nonparametric Statistics*. Springer-Verlag New York, Inc., Secaucus, NJ, USA.
- Wolfson, J. and Gilbert, P. (2010). Statistical identifiability and the surrogate endpoint problem, with application to vaccine trials. *Biometrics*, 66(4):1153–1161.
- Wolfson, J. and Henn, L. (2014). Hard, harder, hardest: principal stratification, statistical identifiability, and the inherent difficulty of finding surrogate endpoints. *Emerging themes in epidemiology*, 11(1):14.
- Woods, J. R., Williams, J. G., and Tavel, M. (1989). The two-period crossover design in medical research. *Ann Intern Med*, 110(7):560–566.
- Zheng, W. and van der Laan, M. J. (2011). Cross-validated targeted minimum-loss-based estimation. In *Targeted Learning*, pages 459–474. Springer.

Appendix

Appendix B: Proofs of Results from Main Text

B.1. Identifiability

B.1.1. Identifiability under A1, A2, and A3

Proof of Theorem 1. Proof of 1. Using that $P^F(S_1 = s_1^*) = \mathbb{E}^F[P^F(S_1 = s_1^*|W)]$, ignorability shows that $P^F(S_1 = s_1^*|A = 1, W)$, and consistency then shows that the right-hand side is equal to $\mathbb{E}[P(S = s_1^*|A = 1, W)]$.

Proof of 2. Using that $P^F(Y_1 = 1, S_1 = s_1^*) = \mathbb{E}^F[P^F(Y_1 = 1, S_1 = s_1^*|W)]$, ignorability shows that $P^F(Y_1 = 1, S_1 = s_1^*|W) = P^F(Y_1 = 1, S_1 = s_1^*|A = 1, W)$, and consistency shows that the right-hand side is equal to $\mathbb{E}[P(Y = 1, S = s_1^*|A = 1, W)]$.

Proof of 3. Note that

$$\begin{aligned}
 P^F(Y_0 = 0, S_1 = s_1^*) &= \mathbb{E}^F \left[\mathbb{E}^F \left[\mathbb{1}_{\{Y_0=0, S_1=s_1^*\}} | W \right] \right] && \text{(Tower Property)} \\
 &= \mathbb{E}^F [P^F(S_1 = s_1^*|Y_0 = 0, W) P^F(Y_0 = 0|W)] && \text{(Simplification)} \\
 &= \mathbb{E}^F [P^F(S_0^c = s_1^*|Y_0 = 0, W) P^F(Y_0 = 0|W)] && \text{(A3)} \\
 &= \mathbb{E}^F [P^F(Y_0 = 0, S_0^c = s_1^*|W)] && \text{(Simplification)} \\
 &= \mathbb{E}^F [P^F(Y_0 = 0, S_0^c = s_1^*|A = 0, W)] && \text{(A2)} \\
 &= \mathbb{E}[P(Y = 0, S^c = s_1^*|A = 0, W)]. && \text{(A1)}
 \end{aligned}$$

□

B.1.2. Non-Falsifiability of A3

We now show that A3 is not falsifiable from the observed data if $\Psi_4(P) = 0$. We do this by constructing a distribution that is consistent with the observed data (in the sense of A1), satisfies the ignorability assumption A2, and satisfies the crossover assumption A3. This shows that, if $\Psi_4(P) = 0$, then we cannot falsify A3 from the observed data, even as the sample size approaches infinity.

Theorem A.1. *If $\Psi_4(P) = 0$, then condition A3 is not falsifiable from the observed data.*

Proof. Suppose that $\Psi_4(P) = 0$. We establish that A3 is not falsifiable by showing that there exists a distribution P_1^F such that, if $P^F = P_1^F$, then A1, A2, and A3 hold. Consider the distribution P_1^F , defined by its conditional densities (with respect to appropriate dominating measures):

$$\begin{aligned}
 p_1^F(y_1|s_0^c, y_0, s_1, a, w) &= p(y|A = 1, s_1, w), \\
 p_1^F(s_0^c|s_1, y_0, a, w) &= p(S^c = s_0^c|Y = 0, A = 0, w) \\
 p_1^F(s_1, y_0|a, w) &= \begin{cases} p(S^c = s_1, Y = 0|A = 0, w), & \text{if } y_0 = 0 \\ p(S = s_1|A = 1, w) - p(S^c = s_1, Y = 0|A = 0, w), & \text{if } y_0 = 1, \end{cases} \\
 p_1^F(a, w) &= p(a, w).
 \end{aligned}$$

Given that P is a probability measure, clearly, the first, second, and fourth conditional densities are everywhere nonnegative and sum to one. The third conditional density is nonnegative because $\Psi_4(P) = 0$ and sums to one since

$$\sum_{s_1} \sum_{y_0} p_1^F(s_1, y_0|a, w) = \sum_{s_1} p(S = s_1|A = 1, w) = 1.$$

Condition A2 is easily established for the above since a does not appear on the right-hand side of any of the density definitions except for the definition of $p_1^F(a, w)$. To establish A1, the only challenge is in showing that $p_1^F(s_1|w) = p(S = s_1|A = 1, w)$ and $p_1^F(y_0|w) = p(Y = y_0|A = 0, w)$, but these results both follow by respectively summing $p_1^F(s_1, y_0|a, w)$ over y_0 and s_1 . Furthermore, note that A3 holds if $P^F = P_1^F$, since, for all s_1 , the fact that S_1 is independent of A, S_0^c given W shows that $p_1^F(S_1 = s_1|Y_0 = 0, w)$ equals $p(S^c = s_1|Y = 0, A = 0, W)$, and A1 and A2 show that $p(S^c = s_1|Y = 0, A = 0, W) = p_1^F(S_0^c = s_1|Y_0 = 0, w)$. \square

B.2. Estimation when Biomarker is Discrete

Proof of Theorem 2. It is well known (van der Laan and Robins, 2003) that, for an arbitrary bounded function $f : \mathcal{O} \rightarrow \mathbb{R}$, $a \in \{0, 1\}$, and parameter $\Phi^{f,a}$ defined by $\Phi^{f,a}(P') \equiv \mathbb{E}_{P'}[\mathbb{E}_{P'}[f(O)|A = a, W]]$, the gradient at P' is

$$o \mapsto \frac{\mathbb{1}\{A = a\}}{P'(a|w)} \{f(o) - \mathbb{E}_{P'}[f(o)|A, W]\} + \mathbb{E}_{P'}[f(o)|A = a, W] - \Phi^{f,a}(P').$$

The proof concludes by noting that $\Psi_k(P') = \Phi^{f_k, a_k}(P')$ for all $k = 1, 2, 3$. \square

Appendix C: Estimation when Principal Biomarker is Continuous

We first present several lemmas used to prove Theorem 5 from the main text, which establishes an asymptotically normal distribution for the estimate of the smoothed parameter. The proof of these lemmas is deferred until after that of Theorem 5. We then prove Theorem 4 from the main text, which shows that, under some conditions, the smoothed parameter converges to the true, unsmoothed parameter as the bandwidth shrinks to zero.

We now present the lemmas used in the above proof of Theorem 5. The first lemma uses fixed functions $w \mapsto \check{q}_k(w)$ and $w \mapsto \check{\pi}_k(w)$. When we invoke this theorem, we will do so at \hat{q}_{k, \hat{h}_n}^v and $\hat{P}(a_k|W = \cdot)$, where we will do this conditionally on training sample v so that we can treat these functions as known.

Lemma A.2. Fix k and functions $w \mapsto \check{q}_k(w)$ and $w \mapsto \check{\pi}_k(w)$. The class $\mathcal{G}_k(\check{q}_k, \check{\pi}_k)$ defined in (A.3) is VC-subgraph (van der Vaart and Wellner, 1996). Furthermore, the VC dimension of this class does not depend on $\check{q}_k, \check{\pi}_k$.

The next lemma is useful for establishing that our estimator ensures that the empirical mean of the smoothed influence function is zero.

Lemma A.3. Fix $k \in \{1, 2, 3\}$. Define

$$Z_{k, \hat{h}_n}(\varepsilon) \equiv \frac{1}{n} \sum_{i=1}^n \frac{\mathbb{1}_{A_i = a_k}}{\hat{P}(A_i|W_i)} \left[f_{k, \hat{h}_n}(O_i) - \exp \left\{ \log \hat{q}_{k, \hat{h}_n}^{v(i)}(W_i) + \varepsilon \right\} \right].$$

If A8, then $Z_{k, \hat{h}_n}(\hat{\varepsilon}_k) = 0$ with probability approaching one.

The final lemma presents a technical condition controlling a class used in the proof of Theorem 5.

Lemma A.4. Fix $\delta > 0$. Under the conditions of Theorem 5, the class $\tilde{\mathcal{G}}_{k, \delta, \hat{h}_n}^v$ defined in the proof of Theorem 5 has envelope $c\delta$ for a constant c that does not depend on the sample realization O_1, \dots, O_n .

Proof of Theorem 5. To simplify notation, we assume that n is divisible by 10 so that validation fold v is of size $n/10$ – note that the results of this proof go through without this assumption. We let P_n^v denote the empirical distribution of the observations in validation set v . For a function g and a distribution Q , we let $Qg = \mathbb{E}_Q[g(O)]$.

Fix $k \in \{1, 2, 3\}$. By the same expansion used to establish (5),

$$\Psi_{k,h}(P') - \Psi_{k,h}(P) = -PD_{k,h}^{P'} + \text{Rem}_{k,h}(P, P'), \tag{A.1}$$

where, with $q'_{k,h}(w) \equiv \mathbb{E}_{P'}[f_{k,h}(O)|a_k, w]$,

$$\text{Rem}_{k,h}(P, P') = \mathbb{E}_P \left[\frac{(P'(a_k|W) - P(a_k|W)) (q'_{k,h}(W) - q_{k,h}(W))}{P'(a_k|W)} \right].$$

By Cauchy-Schwarz and A8, this remainder upper bounds as

$$|\text{Rem}_{k,h}(P, P')| \leq \delta^{-1} \|P'(a_k|W = \cdot) - P(a_k|W = \cdot)\| \|q'_{k,h} - q_{k,h}\|.$$

For a given training fold v , the above is $o_P(n^{-1/2}\hat{h}_n^{-1/2})$ by A6 when evaluated at $h = \hat{h}_n$ and $P' = \hat{P}_{k,\hat{h}_n}^{v,*}$. Multiplying by n^{-1} and summing the identity (A.1) over $v = 1, \dots, 10$ and $i \in v$ at $h = \hat{h}_n$ and $P' = \hat{P}_{k,\hat{h}_n}^{v,*}$, shows that

$$\hat{\Psi}_{k,h} - \Psi_{k,h}(P) = -\frac{1}{n} \sum_{v=1}^{10} \sum_{i:v(i)=1} \mathbb{E} \left[D_{k,\hat{h}_n}^{\hat{P}_{k,\hat{h}_n}^{v,*}}(O) \right] + o_P(n^{-1/2}\hat{h}_n^{-1/2}).$$

Using Lemma A.3, with probability approaching 1 the above shows that

$$\hat{\Psi}_{k,h} - \Psi_{k,h}(P) = \frac{1}{10} \sum_{v=1}^{10} (P_n^v - P) D_{k,\hat{h}_n}^{\hat{P}_{k,\hat{h}_n}^{v,*}} + o_P(n^{-1/2}\hat{h}_n^{-1/2}). \tag{A.2}$$

Consider the following \mathbb{R}^6 -valued data structure, which can be derived from the observed data structure: $\underline{O} \equiv (\check{q}_k(W), \check{\pi}_k(W), A, S, S^c, Y)$. For $i = 1, \dots, n$, we will let \underline{O}_i denote their corresponding \mathbb{R}^6 -valued data structure. Noting that $f_{k,h}(o)$ can be evaluated from $\underline{o} = (\check{q}_k(w), \check{\pi}_k(w), a, s, s^c, y)$, we define $\tilde{f}_{k,h}$ so that $\tilde{f}_{k,h}(\underline{o}) = f_{k,h}(o)$. For fixed functions $w \mapsto \check{q}_k(w)$ and $w \mapsto \check{\pi}_k$, define the class

$$\mathcal{G}_k(\check{q}_k, \check{\pi}_k) \equiv \{ \underline{o} \mapsto g_{k,\varepsilon,h}^v(\underline{o}) : \varepsilon \in \mathbb{R}, h \in \mathbb{R} \}, \tag{A.3}$$

where

$$g_{k,\varepsilon,h}^v(\underline{o}) \equiv \frac{\mathbb{1}_{a=a_k}}{\check{\pi}_k(w)} [\tilde{f}_{k,h}(\underline{o}) - \exp\{\log \check{q}_k(w) + \varepsilon\}] + \exp[\log \check{q}_k(w) + \varepsilon]. \tag{A.4}$$

Lemma A.2 demonstrates that $\mathcal{G}_k(\check{q}_k, \check{\pi}_k)$ is VC-subgraph (van der Vaart and Wellner, 1996), and that the VC dimension does not depend on $\check{q}_k, \check{\pi}_k$. For each v and for the class $\mathcal{G}_k(\hat{q}_{k,h}^v, w \mapsto \hat{P}^v(a_k|w))$, let $g_{k,\varepsilon,h}^v$ denote the element of this class indexed by ε and h . Note that (A.2) rewrites as

$$\hat{\Psi}_{k,h} - \Psi_{k,h}(P) = \frac{1}{10} \sum_{v=1}^{10} (P_n^v - P) g_{k,\hat{\varepsilon}_k,\hat{h}_n}^v + o_P(n^{-1/2}\hat{h}_n^{-1/2}). \tag{A.5}$$

Consider the following subclass of $\mathcal{G}_k(\hat{q}_{k,h}^v, w \mapsto \hat{P}^v(a_k|w))$:

$$\mathcal{G}_{k,\delta}^v \equiv \left\{ \underline{o} \mapsto g_{k,\varepsilon,h}^v(\underline{o}) : |\varepsilon - \bar{\varepsilon}_k| \leq \delta, |1 - \hat{h}_n/\bar{h}_n| \leq \delta \bar{h}_n^{1/2} \right\}.$$

The VC dimension of this subclass is no greater than that of $\mathcal{G}_k(\hat{q}_{k,h}^v, w \mapsto \hat{P}^v(a_k|w))$ and so is bounded by a constant that does not depend on the sample. Combining this with the fact that the VC dimension of this subclass is bounded by a universal constant c_1 shows that the VC dimension of the class $\tilde{\mathcal{G}}_{k,\delta,\bar{h}_n}^v \equiv \left\{ \bar{h}_n^{1/2} [g_{k,\varepsilon,\bar{h}_n}^v - g_{k,\bar{\varepsilon}_k,\bar{h}_n}^v] : g_{k,\varepsilon,\bar{h}_n}^v \in \mathcal{G}_{k,\delta}^v \right\}$ is also bounded by a constant that does not rely on the sample.

Introducing $\tilde{\mathcal{G}}_{k,\delta,\bar{h}_n}^v$ is useful because, for any fixed $\delta > 0$, (A.5) combined with A5 and A9 yield that, with probability approaching one,

$$\begin{aligned} & \left| \hat{\Psi}_{k,h} - \Psi_{k,h}(P) - \frac{1}{10} \sum_{v=1}^{10} (P_n^v - P) g_{k,\bar{\epsilon}_k,\bar{h}_n}^v - o_P(n^{-1/2}\hat{h}_n^{-1/2}) \right| \\ & \leq \left| \frac{1}{10} \sum_{v=1}^{10} (P_n^v - P) (g_{k,\bar{\epsilon}_k,\hat{h}_n}^v - g_{k,\bar{\epsilon}_k,\bar{h}_n}^v) \right| \leq \frac{1}{10\bar{h}_n^{1/2}} \sum_{v=1}^{10} \sup_{\tilde{g}_\delta^v \in \tilde{\mathcal{G}}_{k,\delta,\bar{h}_n}^v} |(P_n^v - P) \tilde{g}_\delta^v|. \end{aligned} \tag{A.6}$$

Lemma A.4 shows that the class $\tilde{\mathcal{G}}_{k,\delta,\bar{h}_n}^v$ has an envelope upper bounded by $c\delta$ for a constant c that does not depend on the sample, i.e. $\sup_{g \in \tilde{\mathcal{G}}_{k,\delta,\bar{h}_n}^v} |g(\underline{o})| \leq c\delta$ for all \underline{o} . Now, for each v we know that $\tilde{\mathcal{G}}_{k,\delta,\bar{h}_n}^v$ is a VC-class and there exists an upper bound on the VC-dimension that does not depend on O_1, \dots, O_n , and so Theorem 2.14.1 in van der Vaart and Wellner (1996) can be used to show that, conditionally on the training set v ,

$$n^{1/2} \mathbb{E}_{P_0^v} \left[\sup_{\tilde{g}_\delta^v \in \tilde{\mathcal{G}}_{k,\delta,\bar{h}_n}^v} |(P_n^v - P) \tilde{g}_\delta^v|^2 \right]^{1/2} \lesssim \delta,$$

where P_0^v is the distribution of the data in validation fold v and here and throughout we use “ \lesssim ” to denote “less than or equal to up to a positive multiplicative constant that does not depend on the observations O_1, \dots, O_n ”. The above holds for all δ . Therefore, for any deterministic sequence $\delta_n \downarrow 0$, the above shows that the right-hand side of (A.6) is $o_P(n^{-1/2}\bar{h}_n^{-1/2})$ for this choice of $\delta = \delta_n$. Moreover, by A5, $o_P(n^{-1/2}\bar{h}_n^{-1/2}) = o_P(n^{-1/2}\hat{h}_n^{-1/2})$.

We will choose the sequence $\delta_n \downarrow 0$ in such a way that the display in (A.6) at $\delta = \delta_n$ holds with probability approaching one. We show that such a sequence exists as follows. Let $\mathcal{E}_{\delta,n}$ denote the event that (A.6) does not hold. Because, for any fixed δ , the complement of $\mathcal{E}_{\delta,n}$ holds with probability approaching one, we know that there exists a sequence $\gamma_{\delta,n} \downarrow 0$ so that $P\{\mathcal{E}_{\delta,n}\} \leq \gamma_{\delta,n}$ for all n . As this is true for all δ , there exists a sequence $\delta_n \downarrow 0$ so that $P\{\mathcal{E}_{\delta_n,n}\} \rightarrow 0$ as $n \rightarrow \infty$. One way to construct such a sequence $\{\delta_n\}$ is as follows: first, let $n_1 = 1$; then recursively, for $j = 1, 2, \dots$, choose $\delta_n = 1/2^j$ for all $n = n_j + 1, \dots, n_{j+1}$, where n_{j+1} is the smallest natural number that is greater than n_j such that $\max\{\gamma_{1/2^j, n_{j+1}}, \gamma_{1/2^{j+1}, n_{j+1}}\} \leq 1/2^j$ — notably, the fact that $\max\{\gamma_{1/2^j, n}, \gamma_{1/2^{j+1}, n}\} \rightarrow 0$ as $n \rightarrow \infty$ implies that n_{j+1} is finite. The event $\mathcal{E}_{\delta_n,n}$ occurs with probability approaching zero along this deterministic sequence, and so the display in (A.6) holds with probability approaching one when $\delta = \delta_n$. Plugging this choice of $\{\delta_n\}$ sequence into the discussion from the previous paragraph, which controlled the right-hand side of (A.6) for a given sequence $\delta_n \downarrow 0$, we see that

$$\hat{\Psi}_{k,h} - \Psi_{k,h}(P) = \frac{1}{10} \sum_{v=1}^{10} (P_n^v - P) g_{k,\bar{\epsilon}_k,\bar{h}_n}^v + o_P(n^{-1/2}\hat{h}_n^{-1/2}).$$

For each v , the function $g_{k,\bar{\epsilon}_k,\bar{h}_n}^v$ does not depend on validation sample v . Therefore, Chebyshev’s inequality applied over each $(P_n^v - P) g_{k,\bar{\epsilon}_k,\bar{h}_n}^v$ can be used to show that, if there exists a function $g_{k,\bar{\epsilon}_k,\bar{h}_n}$ such that $\bar{h}_n P_0^v[(g_{k,\bar{\epsilon}_k,\bar{h}_n}^v - g_{k,\bar{\epsilon}_k,\bar{h}_n})^2] = o_P(1)$, then $(P_n^v - P) g_{k,\bar{\epsilon}_k,\bar{h}_n}^v$ is equal to $(P_n^v - P) g_{k,\bar{\epsilon}_k,\bar{h}_n} + o_P(n^{-1/2}\bar{h}_n^{-1/2})$, where under A5 the o_P term can also be expressed as $o_P(n^{-1/2}\hat{h}_n^{-1/2})$. Under A4, the convergence to a fixed function $g_{k,\bar{\epsilon}_k,\bar{h}_n}$ does indeed hold, with $g_{k,\bar{\epsilon}_k,\bar{h}_n}$ defined in (8). Hence,

$$\begin{aligned} \hat{\Psi}_{k,h} - \Psi_{k,h}(P) &= \frac{1}{10} \sum_{v=1}^{10} (P_n^v - P) g_{k,\bar{\epsilon}_k,\bar{h}_n} + o_P(n^{-1/2}\hat{h}_n^{-1/2}) \\ &= (P_n - P) g_{k,\bar{\epsilon}_k,\bar{h}_n} + o_P(n^{-1/2}\hat{h}_n^{-1/2}). \end{aligned}$$

The above asymptotically linear expansion holds for each k .

Let $g_{\bar{\varepsilon}, \bar{h}_n} : o \mapsto (g_{k, \bar{\varepsilon}_k, \bar{h}_n}(o) : k = 1, 2, 3)$. Using A4 and A8, each component in the \mathbb{R}^3 -valued vector $g_{\bar{\varepsilon}, \bar{h}_n}(O)$ is almost surely bounded by a universal constant times \bar{h}_n^{-1} for $O \sim P$. Using that \bar{h}_n and \hat{h}_n are essentially equivalent under A5, our result will follow if we establish asymptotic normality of the sequence $n^{1/2}(P_n - P)\bar{h}_n^{1/2}\Sigma_n^{-1/2}g_{\bar{\varepsilon}, \bar{h}_n}$. We do this via a Cramér-Wold device and a central limit theorem for triangular arrays. Let b be an \mathbb{R}^3 -valued column vector. Observe that A10 ensures that, with probability approaching one as $n \rightarrow \infty$, $\bar{h}_n^{1/2}b^T\Sigma_n^{-1/2}g_{\bar{\varepsilon}, \bar{h}_n}(O)$ is almost surely bounded by a constant c_b times $\bar{h}_n^{-1/2}$. Furthermore, $\text{Var}_P[\bar{h}_n^{1/2}\Sigma_n^{-1/2}b^Tg_{\bar{\varepsilon}, \bar{h}_n}(O)]$ is equal to $\|b\|^2$. Using that $n\bar{h}_n \rightarrow \infty$ (A5), one can show that the Lindeberg condition for triangular arrays is satisfied and the sequence $\langle b, n^{1/2}(P_n - P)\bar{h}_n^{1/2}\Sigma_n^{-1/2}g_{\bar{\varepsilon}, \bar{h}_n} \rangle$ converges in distribution to a normal random variable with variance $\|b\|^2$. As b was arbitrary, $n^{1/2}(P_n - P)\bar{h}_n^{1/2}\Sigma_n^{-1/2}g_{\bar{\varepsilon}, \bar{h}_n}$ converges to a $N(\underline{0}, \text{Id})$ random variable. \square

Proof of Lemma A.2. The subgraph of a class $\mathcal{G}' \equiv \{o \mapsto g(o) : g\}$ is defined as the set $\{(o, z) : z < g(o)\}$. For a pair (o, z) , membership to this set can be computed by first computing $g(o)$, and subsequently returning an indicator that $z < g(o)$. Now, for \mathcal{G}_k , these functions g are indexed only by parameters $(\varepsilon, h) \in \mathbb{R}^2$. Furthermore, for properly defined t that only depends on the choice of kernel K , A7 ensures that all of these these functions can be computed using no more than a total of t arithmetic operations $(+, -, \div, \times)$, indicator functions $(x \mapsto \mathbb{1}\{x > c\}$ for a constant c), and exponential functions $(x \mapsto e^x)$. By Theorem 8.14 in Anthony and Bartlett (1999), the VC-dimension of this class is no more than $m(t+2)(t+2+19\log_2[9(t+2)])$, hence is finite and does not depend on $\check{q}_k, \check{\kappa}_k$. \square

Proof of Lemma A.3. Fix $k \in \{1, 2, 3\}$. By the strong positivity assumption, with probability approaching one there exists at least one i such that $A_i = a_k$. Without loss of generality, suppose that this holds. Combining this with the bound on \hat{q}_{k, \hat{h}_n}^v in A8, Z_{k, \hat{h}_n} is a monotonically strictly decreasing. Furthermore, $Z_{k, \hat{h}_n}(\varepsilon)$ respectively diverges to $+\infty$ and $-\infty$ as ε tends to $-\infty$ and $+\infty$. Furthermore, Z_{k, \hat{h}_n} is continuous, and therefore there exists a (unique) solution in ε to $Z_{k, \hat{h}_n}(\varepsilon) = 0$, and this solution must coincide with $\hat{\varepsilon}_k$, defined as the minimizer of $Z_{k, \hat{h}_n}(\varepsilon)^2$ in ε . \square

Proof of Lemma A.4. Note that

$$\begin{aligned} & \bar{h}_n^{1/2} |g_{k, \varepsilon, h}^v(o) - g_{k, \bar{\varepsilon}_k, \bar{h}_n}^v(o)| \\ & \leq \bar{h}_n^{1/2} \frac{\mathbb{1}_{a=a_k}}{\hat{P}_k(a_k|w)} |f_{k, h}(o) - f_{k, \bar{h}_n}(o)| + \bar{h}_n^{1/2} \left| \frac{\mathbb{1}_{a=a_k}}{\hat{P}_k(a_k|w)} - 1 \right| \left| e^{\log \hat{q}_{k, h}^v(w) + \varepsilon} - e^{\log \hat{q}_{k, h}^v(w) + \bar{\varepsilon}_k} \right| \\ & \leq \bar{h}_n^{1/2} \hat{P}_k(a_k|w)^{-1} \left[|f_{k, h}(o) - f_{k, \bar{h}_n}(o)| + \left| e^{\log \hat{q}_{k, h}^v(w) + \varepsilon} - e^{\log \hat{q}_{k, h}^v(w) + \bar{\varepsilon}_k} \right| \right] \end{aligned}$$

By (A7), K is bounded and so $|f_{k, h}(o) - f_{k, \bar{h}_n}(o)| \lesssim |h^{-1} - \bar{h}_n^{-1}|$. Using a Taylor expansion and the bound on $\hat{q}_{k, h}^v$ from (A8), we also have that, with probability approaching one, $\left| e^{\log \hat{q}_{k, h}^v(w) + \varepsilon} - e^{\log \hat{q}_{k, h}^v(w) + \bar{\varepsilon}_k} \right| \lesssim |\varepsilon - \bar{\varepsilon}_k|$. Furthermore, by (A8), $\hat{P}_k(a_k|w)^{-1} \lesssim 1$. Hence,

$$\begin{aligned} |g_{k, \varepsilon, h}^v(o) - g_{k, \bar{\varepsilon}_k, \bar{h}_n}^v(o)| & \lesssim \bar{h}_n^{1/2} |h^{-1} - \bar{h}_n^{-1}| - \bar{h}_n^{1/2} |\varepsilon - \bar{\varepsilon}_k| \\ & = \bar{h}_n^{-1/2} \left| 1 - \frac{h}{\bar{h}_n} \right| \left(1 + \left[\frac{\bar{h}_n}{h} - 1 \right] \right) - \bar{h}_n^{1/2} |\varepsilon - \bar{\varepsilon}_k| \end{aligned}$$

If $g_{k, \varepsilon, h}^v$ falls in $\mathcal{G}_{k, \delta}^v$, then the right-hand side upper bounds by $c\delta$ for a constant c that does not depend on O_1, \dots, O_n . Therefore, the envelope of the class $\tilde{\mathcal{G}}_{k, \delta, \bar{h}_n}^v$ is of the order δ . \square

Proof of Theorem 4. Fix $k \in \{1, 2, 3\}$ and suppose that $r \geq t + 1$. A Taylor series expansion shows that

$$\begin{aligned} \Psi_{k,h}(P) - \Psi_k(P) &= \int [\Psi_k(P; s) - \Psi_k(P)] K_h(s - s_1^*) ds \\ &= \sum_{j=1}^t \left[\frac{1}{j!} \frac{\partial^j}{\partial s^j} \Psi_k(P; s) \Big|_{s=s_1^*} \int (s - s_1^*)^j K_h(s - s_1^*) ds \right] \\ &\quad + \int \frac{(s - s_1^*)^{t+1}}{(t+1)!} \frac{\partial^{t+1}}{\partial s^{t+1}} \Psi_k(P; s) \Big|_{s=\tilde{s}_s} K_h(s - s_1^*) ds, \end{aligned}$$

where each \tilde{s}_s is an intermediate value falling between s and s_1^* . So, the right-hand side is equal to the final term. By the uniform bound on the $(t + 1)^{\text{th}}$ derivative of $s \mapsto \Psi_k(P; s)$, the magnitude of the final term is upper bounded by a t -dependent constant times

$$\int |(s - s_1^*)^{t+1} K_h(s - s_1^*)| ds = h^{t+1} \int \left| \left(\frac{s - s_1^*}{h} \right)^{t+1} K \left(\frac{s - s_1^*}{h} \right) \right| \frac{ds}{h} = h^{t+1} \int |u^{t+1} K(u)| du.$$

The right-hand side is $O(h^{t+1})$ by assumption. The result follows because $k \in \{1, 2, 3\}$ was arbitrary. \square

Appendix D: Efficient Influence Functions

We now review the definition of efficient influence functions, which is used in Theorem 2. Define the following fluctuation submodel through P :

$$dP_\varepsilon(o) \equiv [1 + \varepsilon h(o)] dP(o), \text{ where } \mathbb{E}_P[h(O)] = 0 \text{ and } \sup_o |h(o)| < \infty.$$

The function h is a score, and the closure of the linear span of all scores yields the tangent space. It is the resulting tangent space that is important, as pathwise differentiability is equivalent for any set of functions h that yield the same tangent space. Hence, the restriction that $\sup_{o \in \mathcal{O}} |h(o)| < \infty$, while convenient, will have no impact on the resulting differentiability properties.

The parameter Ψ is called pathwise differentiable at P if there exists a $D^P \in L_0^2(P)$, i.e. a D^P such that $D^P(O)$ has mean zero and finite variance under $O \sim P$, such that

$$\Psi(P_\varepsilon) - \Psi(P) = \varepsilon \int D^P(o) h(o) dP(o) + o(\varepsilon).$$

We call D^P a gradient of Ψ at P . The efficient influence function is the gradient D^P for which $D^P(O)$, $O \sim P$, has minimal variance. In a nonparametric model, $D^P(O)$ is almost surely unique.

Appendix E: Two-Phase Sampling Algorithms

Algorithm 3 TMLE for Estimating $\Psi_1(P)$, $\Psi_2(P)$, and $\Psi_3(P)$ under Two-Phase Sampling

Takes as input n observations $\text{Obs} \equiv \{O_i : i = 1, \dots, n\}$, and also the weights $\Delta_i \pi_i^{-1}$.

function TMLE(Obs)

Stabilization: For $a = 0, 1$, let $c(a) \equiv \frac{\sum_{i:A_i=a} \Delta_i / \pi_i}{\#\{i:A_i=a\}}$. For each i , let $\bar{\pi}_i \equiv c(A_i) \pi_i$.

▷ For $a = 0, 1$, the choice of $c(a)$ ensures that $\frac{1}{\#\{i:A_i=a\}} \sum_{i:A_i=a} \Delta_i \bar{\pi}_i^{-1} = 1$.

Initial Estimates: Define an initial estimator \hat{P} of P :

— The marginal distribution of W under \hat{P} should be the weighted empirical distribution, for which each observations puts a mass of $\Delta_i \bar{\pi}_i^{-1}$ at each observation \tilde{O}_i .

— Only the estimates of the components of $P(O|W)$ needed to evaluate Ψ_k and D_k , $k = 1, 2, 3$, are needed: $\hat{P}(A = 1|W = \cdot)$, $\hat{P}(Y = 1|S = s_1^*, A = 1, W = \cdot)$, $\hat{P}(S = s_1^*|A = 1, W = \cdot)$, and $\hat{P}(Y = 0, S^c = s_1^*|A = 0, W = \cdot)$.

▷ If estimated using loss-based learning, include observation weights $\Delta_i \bar{\pi}_i^{-1}$.

▷ Recall the definitions of f_k and a_k from Theorem 2.

for $k = 1, 2, 3$ **do**

Fluctuation for Targeting Step: Using observations $i = 1, \dots, n$, fit the intercept $\hat{\epsilon}_k$ using an intercept-only logistic regression with outcome $f_k(O_i)$, offset $\text{logit } \hat{P}\{f_k(O) = 1|a, w\}$, and weights $\frac{\mathbb{1}_{A_i=a_k}}{\hat{P}(A_i|W_i)} \Delta_i \bar{\pi}_i^{-1}$.

Targeting Step: We now define a fluctuation of \hat{P}_k^* of \hat{P} , targeted towards estimation of ψ_k . Define

$$\hat{P}_k^* \{f_k(O) = 1|a_k, w\} \equiv \text{expit} \{ \text{logit } \hat{P} \{f_k(O) = 1|a_k, w\} + \hat{\epsilon}_k \}.$$

The marginal distribution of $W \sim \hat{P}$ does not need to be fluctuated.

Plug-In Estimator: The estimator of ψ_k is $\hat{\psi}_k \equiv \Psi_k(\hat{P}_k^*)$, and the estimate of the corresponding component of the (full data) influence function is $\hat{D}_k \equiv D_k^{\hat{P}_k^*}$.

▷ Here, by “full data” influence function, we mean that \hat{D}_k is the influence function that we would have had if the phase two covariates had been measured on all individuals. When subsequently developing confidence intervals, $\hat{D}_k(O_i)$ will be weighted by $\Delta_i \bar{\pi}_i^{-1}$.

return Estimate $\hat{\Psi} \equiv (\hat{\psi}_k : k = 1, 2, 3)$ and estimated full data influence function $\hat{D} \equiv (\hat{D}_k : k = 1, 2, 3)$.

Appendix F: Additional Simulation Results

s_1^*	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Bias, Truth	0.03	0.04	0.04	0.03	0.03	0.03	0.03	0.05	0.06	0.08
Bias, Smoothed	0.04	0.05	0.02	-0.01	-0.01	0.02	0.05	0.08	0.07	0.06
Standard Error	0.28	0.27	0.26	0.26	0.27	0.29	0.31	0.34	0.39	0.49
Coverage, Truth	0.97	0.97	0.98	0.98	0.97	0.97	0.97	0.97	0.97	0.96
Coverage, Smoothed	0.97	0.97	0.97	0.96	0.96	0.97	0.97	0.98	0.97	0.96

TABLE 2. Bias, standard error, and coverage probability of the log relative risk point estimator and confidence intervals for different s_1^* values with bandwidth chosen to be $h = 0.1$.

s_1^*	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Bias, Truth	-0.03	-0.02	0.00	0.02	0.04	0.06	0.08	0.10	0.13	0.15
Bias, Smoothed	0.03	0.02	0.02	0.02	0.01	0.01	0.00	-0.00	-0.00	0.00
Standard Error	0.14	0.13	0.13	0.13	0.13	0.14	0.14	0.15	0.17	0.19
Coverage, Truth	0.93	0.94	0.94	0.95	0.95	0.94	0.93	0.93	0.92	0.92
Coverage, Smoothed	0.96	0.95	0.94	0.95	0.95	0.95	0.96	0.96	0.95	0.95

TABLE 3. Bias, standard error, and coverage probability of the log relative risk point estimator and confidence intervals for different s_1^* values with bandwidth chosen to be $h = 0.3$.

s_1^*	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Bias, Truth	-0.07	-0.04	-0.01	0.02	0.05	0.08	0.12	0.15	0.18	0.22
Bias, Smoothed	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	-0.00	-0.00
Standard Error	0.11	0.11	0.11	0.11	0.11	0.11	0.12	0.13	0.14	0.15
Coverage, Truth	0.90	0.93	0.95	0.95	0.93	0.91	0.84	0.79	0.73	0.71
Coverage, Smoothed	0.95	0.96	0.95	0.95	0.95	0.95	0.95	0.95	0.95	0.96

TABLE 4. Bias, standard error, and coverage probability of the log relative risk point estimator and confidence intervals for different s_1^* values with bandwidth chosen to be $h = 0.4$.