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

Year 2016 *Paper* 356

Doubly-robust Nonparametric Inference on the Average Treatment Effect

David Benkeser[*] Marco Carone[†]
Mark J. van der Laan[‡] Peter Gilbert[**]

[*]Division of Biostatistics, University of California, Berkeley, benkeser@berkeley.edu

[†]Department of Biostatistics, University of Washington and the Vaccine and Infectious Disease Division, Fred Hutchinson Cancer Research Institute, mcarone@uw.edu

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

[**]Department of Biostatistics, University of Washington and the Vaccine and Infectious Disease Division, Fred Hutchinson Cancer Research Institute, pgilbert@fhcrc.org

Doubly-robust Nonparametric Inference on the Average Treatment Effect

David Benkeser, Marco Carone, Mark J. van der Laan, and Peter Gilbert

Abstract

Doubly-robust estimators are widely used to draw inference about the average effect of a treatment. Such estimators are consistent for the effect of interest if either one of two nuisance parameters is consistently estimated. However, if flexible, data-adaptive estimators of these nuisance parameters are used, double-robustness does not readily extend to inference. We present a general theoretical study of the behavior of doubly-robust estimators of an average treatment effect when one of the nuisance parameters is inconsistently estimated. We contrast different approaches for constructing such estimators and investigate the extent to which they may be modified to also allow doubly-robust inference. We find that while targeted maximum likelihood estimation can be used to solve this problem very naturally, common alternative frameworks appear to be inappropriate for this purpose. We provide a theoretical study and a numerical evaluation of the alternatives considered. Our simulations highlight the need and usefulness of these approaches in practice, while our theoretical developments have broad implications for the construction of estimators that permit doubly-robust inference in other problems.

1 Introduction

In recent years, doubly-robust estimators have gained immense popularity in many fields, including causal inference. An estimator is said to be doubly-robust if it is consistent for the target parameter of interest if any of two nuisance parameters is consistently estimated. This property gives doubly-robust estimators a natural appeal: any possible inconsistency in the estimation of one nuisance parameter may be mitigated by the consistent estimation of the other. In many problems, doubly-robust estimators arise spontaneously in pursuit of asymptotic efficiency in that locally efficient estimators often also exhibit double-robustness properties. The form of the efficient influence function of the parameter of interest in the statistical model considered generally determines whether or not this is the case. This turns out to be the case for many problems arising in causal inference, which may explain why doubly-robust estimators arise so frequently in that area. For example, under common causal identification assumptions, the statistical parameter identifying the mean counterfactual response under a point treatment yields a doubly-robust efficient influence function in a nonparametric model (Robins et al., 1994). Thus, locally efficient estimators of this statistical target parameter are naturally doubly-robust.

While the conceptual appeal of doubly-robust estimators is rather clear, questions remain about how these estimators should be constructed in practice. In the literature, it has been long noted that the use of finite-dimensional models is generally overly restrictive to permit consistent estimation of the involved nuisance parameters (Bang and Robins, 2005). Nevertheless, much of the current work on doubly-robust estimation involves parametric working models and estimation via maximum likelihood. Kang and Schafer (2007) showed that doubly-robust estimators can be arbitrarily poorly behaved if both nuisance parameters are inconsistently estimated, leading to recent proposals for estimators that minimize the first-order bias resulting from misspecification (Vermeulen and Vansteelandt, 2014, 2016). While providing a significant improvement over conventional techniques, these estimators nevertheless rely upon at least one nuisance parameter being consistently estimated using a parametric model. An alternative approach argues for aggressively employing flexible, data-adaptive estimation techniques for both nuisance parameters to reduce as much as possible the risk of inconsistency (van der Laan and Rose, 2011).

A general study of the behavior of doubly-robust estimators under inconsistent estimation of a nuisance parameter is needed in order to understand how robust statistical inference can be performed. Surprisingly, this topic has not received very much attention. This could in part be due to the fact that when parametric models are used, the problems arising from model misspecification are well-understood. For example, if nuisance parameters are estimated using maximum likelihood, the resulting estimator of the parameter of interest is asymptotically linear even if one of the nuisance parameter models has been misspecified. Though in this scenario the asymptotic variance of this estimator may not be so easy to calculate explicitly, resampling techniques such as the nonparametric bootstrap may be employed to construct confidence intervals and perform statistical hypothesis tests. When the estimator is the solution of an estimating equation, robust sandwich-type variance estimators may also be available. In contrast, when nuisance parameters are estimated using data-adaptive approaches, including nonparametric smoothing techniques and flexible semiparametric procedures, the complications of inconsistently estimating one nuisance parameter are much more serious. Generally, the resulting estimator is irregular, exhibits large bias and has a convergence rate slower than root-n. As we illustrate in this article, the implications for performing inference are dire: regardless of nominal level, the coverage of naïvely constructed two-sided confidence intervals tends to zero and the type I error rate of two-sided hypothesis tests tends to one as sample size increases. This phenomenon cannot simply be avoided by better variance estimators, and in fact occurs even when the true variance of the estimator is exactly known. Furthermore, the nonparametric bootstrap is no longer a saving grace. Due to the use of data-adaptive procedures and the irregularity of the resulting estimator, this technique is not in general valid for constructing confidence intervals and tests.

In view of these challenges, investigators may believe it simpler to restrict their attention to parametric models. However, this is not an appealing solution since under such a strategy both nuisance parameters, and therefore also the parameter of interest, are likely to be inconsistently estimated. The use of flexible, data-adaptive techniques, such as the Super Learner (van der Laan and Polley, 2007), appears necessary to have any reasonable expectation of consistency for any of the nuisance parameter estimators. The Super Learner is an ensemble learning approach for automatically combining estimators within a library of candidate estimators, possibly including nonparametric, semiparametric or parametric procedures, based on cross-validated risk estimates. Asymptotically, it is guaranteed to perform as well as the best possible such combination (van der Laan and Dudoit, 2003; van der Laan et al., 2006; van der Vaart et al., 2006). As such, in practice, it is a practical and principled way of hedging bets in constructing estimators. However, because it is highly adaptive, research is needed for developing appropriate methods for doubly-robust inference that use flexible estimation tools.

A first theoretical study of the problem of doubly-robust nonparametric inference is reported in van der Laan (2014), which focuses on the counterfactual mean under a single time-point intervention and is based on targeted minimum loss-based estimation. Of course, because the average treatment effect is the difference between two counterfactual means under different treatments, it too is directly addressed in this work. The estimators proposed therein were shown to be doubly-robust not only with respect to consistency but also with respect to asymptotic linearity. Furthermore, under general regularity conditions, the analytic form of their influence function is known, which paves the way for the construction of doubly-robust confidence intervals and p-values. The proposed procedure is quite complex, notably involving an iterative procedure and estimation of additional nuisance parameters. Furthermore, it has never been implemented before. We are thus motivated to study theoretically and numerically the following three questions as pertains to the problem of doubly-robust nonparametric inference on an average treatment effect, or equivalently, on a counterfactual mean:

i) How severe are the effects of inconsistent nuisance parameter estimation on inference using data-adaptive estimators, and how do estimators allowing doubly-robust inference perform?

ii) Can targeted minimum loss-based estimators allowing doubly-robust inference be improved through dimension reduction?

iii) Can simpler alternatives to targeted minimum loss-based estimation be used to construct estimators that are doubly-robust for inference and also easier to implement in practice?

As we shall illustrate via a simulation study, the answer to question i) is that naïvely constructed confidence intervals can have very poor coverage, whereas intervals constructed based on appropriate correction procedures have coverage near their nominal level. This suggests that the methods discussed in this paper are truly needed and that they may indeed be quite useful. For question ii), we demonstrate that it is possible to reduce the dimension of the nuisance parameters introduced in the quest for doubly-robust inference. At the very least, this provides theoretical benefits over the proposal of van der Laan (2014). More importantly, this methodological advance is likely to be critical to any extension of the methods discussed here to the setting of treatments defined by multiple-timepoint interventions. Finally, for question iii), we show that the most popular alternative framework to targeted minimum loss-based estimation, the so-called one-step approach, may not be used to theoretically guarantee doubly-robust inference, though it may still yield an estimator with reasonable performance in practice.

This paper is organized as follows. In Section 2, we review strategies for doubly-robust estimation of a counterfactual mean in a nonparametric model. This sets the stage for the study of doubly-robust inference. In Section 3, we discuss correction procedures using targeted minimum loss-based estimation to recover asymptotic linearity of the parameter estimator under inconsistent estimation of one nuisance parameter. In Section 4, we investigate whether the simpler one-step estimation framework may be used as an alternative to targeted minimum loss-based estimation to perform doubly-robust inference. In Section 5, we provide concluding remarks.

2 Doubly-robust estimation

2.1 Notation and background

Suppose the observed data unit is $O := (W, A, Y) \sim P_0$, where W is a vector of baseline covariates, $A \in \{0, 1\}$ a binary treatment, and Y an outcome, and P_0 is the true data-generating distribution, known only to lie in some model $\mathcal{M}$. We take $\mathcal{M}$ to be a nonparametric model, although arbitrary restrictions on the distribution of A given W are allowed without any impact on the developments herein. We focus on the statistical parameter $\Psi : \mathcal{M} \to \mathbb{R}$ defined as

$$\Psi(P) := \int \bar{Q}(w) dQ_W(w)$$

for each $P \in \mathcal{M}$, where $\bar{Q}(w) = \bar{Q}_P(w) := E_P(Y \mid A = 1, W = w)$ is the so-called outcome regression and $Q_W(w) = Q_{W,P}(w) := P(W \leq w)$ is the distribution function of the covariate vector. The parameter value $\Psi(P)$ represents the treatment-specific, covariate-adjusted mean outcome implied by $P \in \mathcal{M}$. Under additional causal assumptions, it can be interpreted as the mean counterfactual outcome under the treatment corresponding to $A = 1$ (Rubin, 1974). Because all developments below immediately apply to the case $A = 0$, and therefore to the average treatment effect, without loss of generality, we explicitly examine only the case $A = 1$.

As the parameter of interest only depends on P through $Q = Q(P) := (\bar{Q}, Q_W)$, we will at times write $\Psi(Q)$ to denote $\Psi(P)$. We will denote $Q(P_0)$ in shorthand as $Q_0 := (\bar{Q}_0, Q_{W,0})$, where $\bar{Q}_0$ is the true outcome regression and $Q_{W,0}$ the true distribution of W. The propensity score, defined as $g(w) := P(A = 1 \mid W = w)$, plays an important role and throughout, the true propensity score g_0 is assumed to satisfy $g_0(w) > \delta$ for some $\delta > 0$ and all w in the support of $Q_{W,0}$. Below, we make use of empirical process notation, writing Pf to denote $\int f(o)dP(o)$ for each $P \in \mathcal{M}$ and P-integrable function f. We also denote by P_n the empirical distribution function based on $O_1, O_2, \ldots, O_n$, and thus, $P_n f$ is the empirical average $n^{-1} \sum_{i=1}^n f(O_i)$.

We recall that a regular estimator ψ_n of $\psi_0 := \Psi(Q_0)$ is asymptotically linear if and only if it can be written as $\psi_n = \psi_0 + P_n D(P_0) + o_P(n^{-1/2})$, where $D(P_0) \in L_2^0(P_0)$ is a gradient of Ψ at P_0 relative to model $\mathcal{M}$. Here, for each $P \in \mathcal{M}$, we denote by $L_2^0(P)$ the Hilbert space of mean zero finite variance functions endowed with the covariance inner product. The function $D(P) \in L_2^0(P)$ is said to be a gradient of Ψ at P relative to $\mathcal{M}$ if

$$\frac{d}{d\epsilon}\Psi(P_\epsilon)\bigg|_{\epsilon=0} = \int D(P)(o)s(o)dP(o)$$

for any regular one-dimensional parametric submodel $\{P_\epsilon\} \subseteq \mathcal{M}$ with score s for ϵ at $\epsilon = 0$ and such that $P_{\epsilon=0} = P$. Furthermore, such estimators are efficient if and only if their influence function is given by the efficient influence function $D^*(P_0)$.

The efficient influence function is the unique gradient that lies in the tangent space $T_{\mathcal{M}}(P_0) \subseteq L_2^0(P_0)$ of $\mathcal{M}$ at P_0, and it is a critical ingredient in the construction of asymptotically efficient estimators. For an overview of efficiency theory we refer readers to Bickel et al. (1997).

The efficient influence function of Ψ at P relative to $\mathcal{M}$ is

$$D^*(P)(o) = D^*(Q,g)(o) = \frac{a}{g(w)} \left\{y - \bar{Q}(w)\right\} + \bar{Q}(w) - \Psi(Q)$$

with $o := (w, a, y)$ denoting a realized value of O (van der Laan and Robins, 2003).

2.2 Doubly-robust consistency

Suppose that $\bar{Q}_n$ and g_n are estimators of $\bar{Q}_0$ and g_0, respectively, and denote by $\bar{Q}$ and g their respective in-probability limits. We will write $Q_n := (\bar{Q}_n, Q_{W,n})$, where $Q_{W,n}$ is the empirical distribution based on observations $W_1, W_2, \ldots, W_n$. A linearization of the parameter and simple algebraic manipulations allow us to write

$$\begin{aligned}
\Psi(Q_n) - \Psi(Q_0) &= -P_0 D^*(Q_n, g_n) + R(Q_n, Q_0, g_n, g_0) \\
&= (P_n - P_0) D^*(Q_n, g_n) - P_n D^*(Q_n, g_n) + R(Q_n, Q_0, g_n, g_0) \\
&= (P_n - P_0) D^*(Q, g) - P_n D^*(Q_n, g_n) + R(Q_n, Q_0, g_n, g_0) \\
&\qquad + (P_n - P_0)\left\{D^*(Q_n, g_n) - D^*(Q, g)\right\} ,
\end{aligned}$$

where $R(Q_1, Q_2, g_1, g_2) := P_0\{(g_1 - g_2)(\bar{Q}_1 - \bar{Q}_2)/g_1\}$ is the remainder term from the linearization of Ψ using the canonical gradient. As shorthand, we will write $B_n(Q_n, g_n) := P_n D^*(Q_n, g_n)$ and $M_n(Q_n, Q, g_n, g) := \big(P_n - P_0\big)\big\{D^*(Q_n, g_n) - D^*(Q, g)\big\}$. Using this notation, we can write the estimation error $\Psi(Q_n) - \Psi(Q_0)$ as

$$(P_n - P_0) D^*(Q, g) - B_n(Q_n, g_n) + M_n(Q_n, Q, g_n, g) + R(Q_n, Q_0, g_n, g_0) \; . \qquad (1)$$

This representation reduces the analysis of the plug-in estimator $\Psi(Q_n)$ to that of four terms. The first term, $(P_n - P_0)D^*(Q, g)$, is the empirical average of a random variable, $D^*(Q, g)(O)$, with mean zero if either $Q = Q_0$ or $g = g_0$. The latter observation is a simple but fundamental fact underlying much of the doubly-robust estimation literature. Since $Q_{W,n}$ is known to converge to $Q_{W,0}$, we note that $Q = Q_0$ is equivalent to $\bar{Q} = \bar{Q}_0$. The second term, $B_n(Q_n, g_n)$, is a first-order bias term that must be accounted for to allow $\Psi(Q_n)$ to be asymptotically linear. The third term is an empirical process term that is often asymptotically negligible, that is, $M_n := M_n(Q_n, Q, g_n, g) = o_P(n^{-1/2})$. This is true, for example, if $D^*(Q_n, g_n)$ falls in a P_0-Donsker class with probability tending to one and $P_0\{D^*(Q_n, g_n) - D^*(Q, g)\}^2$ converges to zero in probability. For a comprehensive reference on the theory of empirical processes, we encourage readers to consult van der Vaart and Wellner (1996). Finally, the fourth term, $R_n := R(Q_n, Q_0, g_n, g_0)$, is the remainder from the linearization. By inspection, this term tends to zero at a rate determined by how fast the nuisance functions $\bar{Q}_0$ and g_0 are estimated.

To correct for the first-order bias term highlighted above, two general strategies may be used: the one-step Newton-Raphson approach and targeted minimum loss-based estimation. The one-step Newton-Raphson procedure, hereafter referred to as the one-step approach, suggests performing an additive correction for the first-order bias in the parameter space, leading to the estimator

$$\psi_n^+ := \Psi(Q_n) + B_n(Q_n, g_n) \ .$$

This approach appeared early on in the works of Ibragimov et al. (1981) and Pfanzagl (1982), and is the infinite-dimensional extension of the well-known one-step Newton-Raphson technique for efficient estimation in parametric models. In the case considered in this paper, the efficient influence function of the parameter of interest is a linear function of the parameter. As such, the one-step estimator agrees exactly with the solution of the optimal estimating equation for this parameter and is thus equivalent to the so-called augmented inverse probability of treatment estimator (Robins et al., 1994; van der Laan and Robins, 2003). Owing to their simple construction, one-step estimators are generally computationally convenient to implement. However, this convenience comes at a cost. In practice, the one-step correction may produce estimates outside of the parameter space, such as probability estimates either below 0 or above 1. Targeted minimum loss-based estimation, formally developed in van der Laan and Rubin (2006) and comprehensively discussed in van der Laan and Rose (2011), provides a recursive algorithm to convert Q_n into a targeted estimator Q_n^* of Q_0 such that $B_n(Q_n^*, g_n) = 0$, which may then be used to define the targeted plug-in estimator $\psi_n^* := \Psi(Q_n^*)$. The first update of Q_n in this recursive scheme consists of the minimizer of an empirical risk over a least-favorable submodel through Q_n. The process is then repeated using this updated version of Q_n instead of Q_n itself. This updating procedure is iterated until convergence to yield Q_n^*. In the problem considered here, convergence occurs in a single step. In contrast to the one-step approach, targeted minimum-loss based estimation corresponds to performing bias correction in the model space. By virtue of being a plug-in estimator, ψ_n^* may exhibit improved finite-sample behavior relative to its one-step counterpart (Porter et al., 2011).

The large-sample properties of both ψ_n^+ and ψ_n^* can be studied through the representation provided in (1). As discussed above, suppose that the empirical process term M_n is asymptotically negligible. If both Q_0 and g_0 are estimated consistently, so that $Q = Q_0$ and $g = g_0$, and if estimation of these nuisance functions is fast enough to ensure that the remainder term R_n is asymptotically negligible, it follows that ψ_n^+ is asymptotically linear with influence function equal to $D^*(Q_0, g_0)$ and thus asymptotically efficient. The same can be said about ψ_n^* if these same conditions hold replacing Q_n by Q_n^* in both the M_n and R_n terms. If only one of $Q = Q_0$ or $g = g_0$ holds, it is impossible to guarantee the asymptotic negligibility of the remainder term, even when using correctly-specified parametric models. Nevertheless, under very mild conditions, the remainder term R_n based on either of Q_n and Q_n^* tends to zero in probability, and the empirical process term M_n remains asymptotically negligible. Since $D^*(Q,g)(O)$ has mean zero if either $Q = Q_0$ or $g = g_0$ and finite variance, the central limit theorem implies that $(P_n - P_0)D^*(Q,g)$ is $O_P(n^{-1/2})$. It follows then that both ψ_n^+ and ψ_n^* are consistent estimators of ψ_0. This is what is generally referred to as double robustness: consistent estimation of ψ_0 if either at least one of the nuisance functions Q_0 or g_0 is consistently estimated.

2.3 Doubly-robust asymptotic linearity

Doubly-robust asymptotic linearity is a more stringent requirement than doubly-robust consistency. It is also arguably a more important property since without it the construction of valid confidence intervals and tests may be very difficult, if not impossible. A careful study of R_n is required to determine how double-robust inference may be obtained.

When both the outcome regression and propensity score are consistently estimated, R_n is a second-order term consisting of the product of two differences, both tending to zero. Thus, provided $\bar{Q}_0$ and g_0 are estimated sufficiently fast, it is the case that $R_n = o_P(n^{-1/2})$. This holds, for example, if both $\bar{Q}_n - \bar{Q}_0$ and $g_n - g_0$ are $o_P(n^{-1/4})$ with respect to the $L^2(P_0)$ norm. If only one of the outcome regression or propensity score is consistently estimated, one of the differences in R_n does not tend to zero. Consequently, R_n is either of the same order or slower than $(P_n - P_0)D^*(Q,g)$. As such, it at least contributes to the first-order behavior of the estimator, if not determines it entirely. In this case, if a correctly-specified parametric model is used to estimate either $\bar{Q}_0$ or g_0, the delta method generally implies that R_n is asymptotically linear. Then, both ψ_n^+ and ψ_n^* are also asymptotically linear, though their influence function consists of two summands, $D^*(Q,g)$ and the influence function of R_n as an estimator of zero. Correctly specifying a parametric model can seldom be done in realistic settings, however. For this reason, it may be preferable to use flexible, data-adaptive estimators of the nuisance functions to get as close as possible to their true value. In this case, whenever one nuisance is inconsistently estimated, the remainder term R_n tends to zero slowly and dominates the first-order behavior of the estimator of ψ_0. The latter then does not exhibit regular large-sample behavior. Therefore, in this case, the one-step and targeted minimum loss-based estimator are doubly-robust with respect to consistency but not with respect to asymptotic linearity.

To illustrate the deleterious effect of the remainder on inference in these situations, suppose that we construct an asymptotic level $1-\alpha$ two-sided Wald-type confidence interval for ψ_0 based on a consistent estimator ψ_n, say with true standard error s_n. Suppose further that $|\psi_n - \psi_0|/s_n$ tends to $+\infty$ in probability, which often occurs when the bias of ψ_n tends to zero slower than its standard error. Denoting by z_β the β quantile of the standard normal distribution, the coverage of the oracle Wald-type interval $\psi_n \pm z_{1-\alpha/2}s_n$ is given by

$$P_0\left(\psi_n - z_{1-\alpha/2}s_n < \psi_0 < \psi_n + z_{1-\alpha/2}s_n\right) = P_0\left(\left|\frac{\psi_n - \psi_0}{s_n}\right| < z_{1-\alpha/2}\right)$$

and thus tends to zero. This remains true if we replace s_n by any random sequence that converges to zero at the same rate or faster. Were asymptotic linearity preserved under inconsistent estimation of one of the nuisance parameters, $(\psi_n - \psi_0)/s_n$ would instead tend to a standard normal variate. The oracle Wald-type intervals, and in fact any Wald-type interval using a consistent standard error estimator, would have correct asymptotic coverage. This argument therefore stresses the benefit of constructing estimators that are doubly-robust with respect to asymptotic linearity for the sake of obtaining confidence intervals and tests whose validity is doubly-robust.

3 Doubly-robust inference via targeted minimum loss-based estimation

3.1 Existing construction

Recently, van der Laan (2014) proposed a targeted minimum loss-based estimator of ψ_0 that is not only locally efficient and doubly-robust for consistency but also doubly-robust for asymptotic linearity. To do so, he showed that with some additional bias correction R_n may be rendered asymptotically linear with a well-described influence function. This requires approximating the first-order behavior of R_n using additional nuisance parameters. These nuisance parameters consist of a bivariate and univariate regression, defined respectively as

$$\begin{aligned} g_{0,r}(\bar{Q},g)(w) &:= E_{P_0}\left\{A \mid \bar{Q}(W) = \bar{Q}(w), g(W) = g(w)\right\}, \\ \bar{Q}_{0,r}(\bar{Q},g)(w) &:= E_{P_0}\left\{Y - \bar{Q}(W) \mid A = 1, g(W) = g(w)\right\}. \end{aligned}$$

The first nuisance parameter above is the bivariate regression of the true propensity of treatment on an outcome regression and a propensity score, whereas the second is the univariate regression of the residual from an outcome regression on a propensity score in the treated subgroup. The subscript r emphasizes that these nuisance parameters are of reduced dimension relative to g_0 and $\bar{Q}_0$. This dimension reduction is critical since it essentially guarantees that consistent estimators of these parameters can be constructed in practice. For example, we may be unable to consistently estimate g_0, which is a function of the entire vector of potential confounders; however, we can guarantee consistent estimation of $g_{0,r}$, which involves only a bivariate summary of W.

Key to the study of how these additional nuisance parameters may be used to approximate the first-order behavior of the remainder term R_n are the functions

$$D_A(\bar{Q}_{0,r}, g)(o) := -\frac{\bar{Q}_{0,r}(w)}{g(w)} \{a - g(w)\} \ ,$$
$$D_Y(\bar{Q}, g_{0,r})(o) := -\frac{a}{g_{0,r}(w)} \left\{ \frac{g_{0,r}(w) - g(w)}{g(w)} \right\} \{y - \bar{Q}(w)\} \ .$$

In Appendix A, we show that the remainder term $R(Q_n, Q, g_n, g)$ can be represented as

$$R_n^* + I(g = g_0) \left\{ (P_n - P_0) D_A(\bar{Q}_{0,r}, g) - B_{A,n}(\bar{Q}_{n,r}, g_n) + R_{g,n} \right\} \\ + I(\bar{Q} = \bar{Q}_0) \left\{ (P_n - P_0) D_Y(\bar{Q}, g_{0,r}) - B_{Y,n}(\bar{Q}_n, g_{n,r}) + R_{Q,n} \right\} \ , \tag{2}$$

where $B_{A,n}(\bar{Q}_{n,r}, g_n) := P_n D_A(\bar{Q}_{n,r}, g_n)$ and $B_{Y,n}(\bar{Q}_n, g_{n,r}) := P_n D_Y(\bar{Q}_n, g_{n,r})$ are first-order bias terms, and R_n^*, $R_{g,n}$ and $R_{Q,n}$ are second-order terms. The specific form of these terms is provided in Appendix A, where we also discuss sufficient conditions for ensuring their asymptotic negligibility. Importantly, just as the bias term in (1) had to be accounted for to achieve doubly-robust consistency, so too must the bias terms in (2) in order to achieve doubly-robust asymptotic linearity.

An iterative targeted minimum loss-based estimation algorithm was proposed in van der Laan (2014) to produce estimators $\bar{Q}_n^*$, g_n^*, $\bar{Q}_{n,r}^*$ and $g_{n,r}^*$ from initial estimators Q_n and g_n in such a manner as to ensure that

$$B_n(Q_n^*, g_n^*) = B_{A,n}(\bar{Q}_{n,r}^*, g_n^*) = B_{Y,n}(\bar{Q}_n^*, g_{n,r}^*) = o_P(n^{-1/2}) \ .$$

In view of (1) and (2), this implies that $\psi_n^* := \Psi(Q_n^*)$ is asymptotically linear with influence function

$$D^{*,\#}(Q, g) := D^*(Q, g) - I(g = g_0) D_A(\bar{Q}_{0,r}, g) - I(\bar{Q} = \bar{Q}_0) D_Y(\bar{Q}, g_{0,r})$$

provided either $\bar{Q}_0$ or g_0 is estimated consistently. We note that if both $\bar{Q}_0$ and g_0 are estimated consistently, both $D_A(\bar{Q}_0^r, g)$ and $D_Y(\bar{Q}, g_{0,r})$ are identically zero since then $\bar{Q}_{0,r} = 0$ and $g_{0,r} = g_0$. This establishes that asymptotic local efficiency is indeed preserved. We refer readers to Theorem 3 of van der Laan (2014) for a presentation of the corresponding algorithm.

3.2 Novel reduced-dimension construction

In this subsection, we show that it is possible to theoretically improve upon the proposal of van der Laan (2014) through an alternative formulation of a targeted minimum loss-based estimator. In particular, we derive an approximation of the remainder that relies on alternate nuisance parameters of lower dimension than those presented in the previous subsection. This not only renders the involved estimation problem more feasible in practice but it may also pave the way to a tractable generalization of this work to settings wherein the treatment considered is defined longitudinally.

In Appendix B, we argue that the remainder term in (2) can alternatively be represented using $\bar{Q}_{0,r}$ as previously defined and the additional nuisance parameters

$$\begin{aligned} g_{1,0,r}(\bar{Q})(w) &:= E_0\left\{A \mid \bar{Q}(W) = \bar{Q}(w)\right\}, \\ g_{2,0,r}(\bar{Q}, g)(w) &:= E_0\left\{\frac{A - g(W)}{g(W)} \,\middle|\, \bar{Q}(W) = \bar{Q}(w)\right\}. \end{aligned}$$

We note that $g_{1,0,r}$ and $g_{2,0,r}$ consist only of univariate regressions in contrast to the bivariate regression $g_{0,r}(\bar{Q}, g)$ described in the previous subsection and used in van der Laan (2014). As such, nonparametric estimators of these univariate nuisance parameters achieve better rates than those proposed in that work. Use of this alternate representation results in estimators guaranteed to be asymptotically linear under weaker conditions than previously required.

Here, we state the main result dictating the behavior of the novel estimator implied by this parametrization of the remainder term, and we discuss an iterative implementation of this estimator. Redefining

$$D_Y(\bar{Q}, g_{1,0,r}, g_{2,0,r})(o) := \frac{a}{g_{1,0,r}(w)}\; g_{2,0,r}(w)\left\{y - \bar{Q}(w)\right\},$$

we have the following result.

Theorem 1 *Suppose that either $\bar{Q} = \bar{Q}_0$ or $g = g_0$. Provided the nuisance estimators $(\bar{Q}_n^*, \bar{Q}_{n,r}^*, g_n^*, g_{1,n,r}^*, g_{2,n,r}^*)$ satisfy the equations*

$$B_n(Q_n^*, g_n^*) = B_{A,n}(\bar{Q}_{n,r}^*, g_n^*) = B_{Y,n}(\bar{Q}_n^*, g_{1,n,r}^*, g_{2,n,r}^*) = o_P(n^{-1/2})$$

and the second-order terms $R_{Q,n}$ and $R_{g,n}$ described in Appendix B are $o_P(n^{-1/2})$, the plug-in estimator $\psi_n^{*,c} := \Psi(Q_n^*)$ is asymptotically linear with influence function $D^{*,\#}(Q, g)$. Furthermore, $n^{1/2}(\psi_n^{*,c} - \psi_0)$ converges in law to a zero-mean normal random variable with variance estimated consistently by

$$\sigma_n^2 := P_n\left\{D^*(Q_n^*, g_n^*) - D_A(\bar{Q}_{n,r}^*, g_n^*) - D_Y(\bar{Q}_n^*, g_{1,n,r}^*, g_{2,n,r}^*)\right\}^2.$$

An algorithm to construct nuisance estimators that solve the above equations can be devised based on targeted minimum loss-based estimation. Without any loss of generality, suppose that Y is bounded between 0 and 1. Defining $H_1(g)(a, w) := a/g(w)$, $H_2(g_1, g_2)(a, w) := a g_2(w)/g_1(w)$ and $H_3(\bar{Q}, g)(w) := \bar{Q}(w)/g(w)$, we implement the following recursive procedure:

1. construct initial estimates $\bar{Q}_n^0$ and g_n^0 of $\bar{Q}_0$ and g_0, and set $k = 0$;

2. define $H_{1n,k} := H_1(g_n^k)(A,W)$ and $L_{1n,k} := \text{logit}\{\bar{Q}_n^k(W)\}$, fit a logistic regression with outcome Y, covariate $H_{1n,k}$ and offset $L_{1n,k}$ using only data points with $A = 1$, set $\epsilon_{1n,k}$ as the estimated coefficient of $H_{1n,k}$, and define

$$\bar{Q}_n^{k,\circ}(w) := \text{expit}\left[\text{logit}\{\bar{Q}_n^k(w)\} + \epsilon_{1n,k} H_1(g_n^k)(1,w)\right] ;$$

3. construct estimates $g_{1,n,r}^k$ and $g_{2,n,r}^k$ of $g_{1,0,r}$ and $g_{2,0,r}$ based on g_n^k and $\bar{Q}_n^{k,\circ}$;

4. define $H_{2n,k} := H_2(g_{1,n,r}^k, g_{2,n,r}^k)(A,W)$ and $L_{2n,k} := \text{logit}\{\bar{Q}_n^{k,\circ}(W)\}$, fit a logistic regression with outcome Y, covariate $H_{2n,k}$ and offset $L_{2n,k}$ using only data points with $A = 1$, set $\epsilon_{2n,k}$ as the estimated coefficient of $H_{2n,k}$, and define

$$\bar{Q}_n^{k+1}(w) := \text{expit}\left[\text{logit}\{\bar{Q}_n^{k,\circ}(w)\} + \epsilon_{2n,k} H_2(g_{1,n,r}^k, g_{2,n,r}^k)(1,w)\right] ;$$

5. construct estimates $\bar{Q}_{n,r}^k$ of $\bar{Q}_{0,r}$ based on g_n^k and $\bar{Q}_n^{k+1}$;

6. define $H_{3n,k} := H_3(\bar{Q}_{n,r}^k, g_n^k)(W)$ and $L_{3n,k} := \text{logit}\{g_n^k(W)\}$, fit a logistic regression with outcome A, covariate $H_{3n,k}$ and offset $L_{3n,k}$, set $\epsilon_{3n,k}$ as the estimated coefficient of $H_{3n,k}$, and define

$$g_n^{k+1}(w) := \text{expit}\left[\text{logit}\{g_n^k(w)\} + \epsilon_{3n,k} H_3(\bar{Q}_{n,r}^k, g_n^k)(w)\right] ;$$

7. set $k = k+1$ and iterate the above steps until K large enough so that $P_n D^*(Q_n^K, g_n^K) \approx P_n D_A(\bar{Q}_{n,r}^K, g_n^K) \approx P_n D_Y(\bar{Q}_n^K, g_{1,n,r}^K, g_{2,n,r}^K) \approx 0$.

8. set $\bar{Q}_{n,r}^* := \bar{Q}_{n,r}^K$, $g_n^* := g_n^K$, $\bar{Q}_{n,r}^* := \bar{Q}_{n,r}^K$, $g_{1,n,r}^* := g_{1,n,r}^K$ and $g_{2,n,r}^* := g_{2,n,r}^K$.

The important ramification of Theorem 1 is that doubly-robust confidence intervals and tests can readily be crafted. For example, the Wald construction $\psi_n^{*,c} \pm z_{1-\alpha/2}\sigma_n n^{-1/2}$ is a doubly-robust $100 \times (1-\alpha)\%$ asymptotic confidence interval for ψ_0, and prescribing rejection whenever

$$\left| \frac{n^{1/2}(\psi_n^{*,c} - \psi^\circ)}{\sigma_n} \right| > z_{1-\alpha/2}$$

and failure to reject otherwise constitutes a doubly-robust hypothesis test of the null hypothesis $\psi_0 = \psi^\circ$ versus the alternative $\psi_0 \neq \psi^\circ$ with asymptotic level α. Thus, valid statistical inference is preserved when one nuisance parameter is inconsistently estimated, in sharp contrast to conventional doubly-robust estimation, wherein only consistency is preserved.

4 Doubly-robust inference via one-step estimation

In Section 2, we discussed the construction of doubly-robust, locally efficient estimators of ψ_0. We argued that two general strategies, the one-step approach and targeted minimum loss-based estimation, can be used for bias correction. For the sake of constructing asymptotically efficient estimators, these two strategies are generally considered to be alternatives to each other, with targeted minimum loss-based estimation possibly delivering better finite-sample behavior but the one-step approach often simpler to implement. In the previous section, we outlined how the bias-correction feature of the targeted minimum loss-based estimation framework could be leveraged to achieve doubly-robust asymptotic linearity and thus perform doubly-robust inference. Since targeted minimum loss-based estimation can be more complicated to implement than the one-step correction procedure, particularly in view of the iterative nature of the algorithm, it is natural to wonder whether a one-step approach could also be used to account for the additional bias terms that result from the inconsistent estimation of either $\bar{Q}_0$ or g_0 in the problem considered. If so, the resulting one-step estimator could provide a computationally convenient alternative to the complex recursive algorithm involved in the construction of the targeted minimum loss-based estimators.

We recall that the doubly-robust, locally efficient one-step estimator ψ_n^+ was constructed by adding the bias term $B_n^*(Q_n, g_n)$ to the initial plug-in estimator $\Psi(Q_n)$. By extension, it appears sensible to investigate whether the estimator

$$\psi_n^{+,c} := \psi_n^+ + B_{A,n}(\bar{Q}_{n,r}, g_n) + B_{Y,n}(\bar{Q}_n, g_{1,n,r}, g_{2,n,r}) \tag{3}$$

is doubly-robust with respect to asymptotic linearity. By equations (1) and (2), we immediately have that the estimator

$$\psi_n^{oracle} := \psi_n^+ + I(g = g_0)B_{A,n}(\bar{Q}_{n,r}, g_n) + I(\bar{Q} = \bar{Q}_0)B_{Y,n}(\bar{Q}_n, g_{1,n,r}, g_{2,n,r})$$

is asymptotically linear with influence function $D^{\#}(Q, g)$, just as the targeted minimum loss-based estimators in the previous section. Therefore, ψ_n^{oracle} is locally efficient and doubly-robust with respect to asymptotic linearity. Nevertheless, to compute ψ_n^{oracle}, the analyst must know which nuisance parameter, if any, is inconsistently estimated. Such information will generally not be available, except in the case of a randomized trial, where g_0 may be known to the experimenter. To study the properties of $\psi_n^{+,c}$, we note that

$$\psi_n^{+,c} - \psi_n^{oracle} = I(g \neq g_0)B_{A,n}(\bar{Q}_{n,r}, g_n) + I(\bar{Q} \neq \bar{Q}_0)B_{Y,n}(\bar{Q}_n, g_{1,n,r}, g_{2,n,r}) \ . \tag{4}$$

The one-step estimator $\psi_n^{+,c}$ corrects for both inconsistent estimation of $\bar{Q}_0$ and g_0. However, for consistent estimation of ψ_0, no more than one of these two nuisances can in reality be inconsistently estimated. In this case, there is necessarily overcorrection in $\psi_n^{+,c}$ and it is not a priori obvious whether this may be detrimental to the behavior of the estimator. Elucidating this fact requires a careful study of each of the two bias correction terms in settings in which they are not in fact needed. For example, the term $B_{A,n}(\bar{Q}_{n,r}, g_n)$, used to correct for bias resulting from inconsistent estimation of $\bar{Q}_0$, must be analyzed under the scenario wherein it is in fact g_0 that has been inconsistently estimated.

In Appendix C, we show that under reasonable rate conditions, we can represent the first summand on the right-hand side of (4) as

$$B_{A,n}(\bar{Q}_{n,r}, g_n) = P_0 \left\{ \left(\frac{g_0 - g}{g} \right) \bar{Q}_{n,r} \right\} + o_P(n^{-1/2}) \tag{5}$$

when $g \neq g_0$, and we can represent the second summand as

$$B_{Y,n}(\bar{Q}_n, g_{1,n,r}, g_{2,n,r}) = P_0 \left\{ \frac{A}{g_{1,n,r}} (\bar{Q}_0 - \bar{Q}) g_{2,n,r} \right\} + o_P(n^{-1/2})$$

when $\bar{Q} \neq \bar{Q}_0$. This implies that the first-order behavior of $\psi_n^{+,c}$ is driven by these terms. In particular, the rate of convergence of $\psi_n^{+,c}$ is determined by that of the estimators $\bar{Q}_{n,r}$, $g_{1,n,r}$ and $g_{2,n,r}$ of the reduced-dimension nuisance parameters. These terms are unlikely to be estimable at the parametric rate in practice since this would require the correct specification of a parametric model for a complex object. In practice, flexible, data-adaptive techniques are likely to be used to consistently estimate these regression functions. Because the rates achieved by these techniques are generally slower than one over root-n, the estimator $\psi_n^{+,c}$ fails to be root-n-consistent and hence doubly-robust with respect to asymptotic linearity. Using an argument identical to that made in Section 3, we can show that Wald-type confidence intervals for $\psi_n^{+,c}$ have similarly poor asymptotic coverage. Therefore, at least theoretically, the one-step construction does not appear helpful to achieve double-robust inference.

This result warrants further discussion. The above theory shows that the targeted minimum loss-based estimation framework is able to simultaneously account for inconsistent estimation of either the outcome regression or the propensity score without the need to know which is needed. In contrast, the one-step approach is unable to do so: it requires knowledge of which nuisance parameter is possibly inconsistently estimated to retain asymptotically linearity. Without this knowledge, asymptotic linearity cannot be theoretically guaranteed in a doubly-robust fashion. This suggests that for the purpose of performing doubly-robust inference on ψ_0 targeted minimum loss-based estimation may be required. This finding is relevant for future work to derive procedures for doubly-robust inference on other parameters admitting doubly- or multiply-robust estimators.

5 Simulation study

5.1 Data-generating mechanism and setup

In each of the simulations below, we used the following data-generating mechanism. The baseline covariate vector $W := (W_1, W_2)$ has independent components. W_1 is distributed according to a uniform distribution over the interval $(-2, +2)$ and W_2 is a binary random variable with success probability 1/2. The conditional probability of receiving treatment $A = 1$ given $W = (w_1, w_2)$ is given by $g_0(w_1, w_2) := \text{expit}(-w_1 + 2w_1w_2)$. The outcome Y is a binary random variable with conditional probability of occurrence given $A = a$ given by $\bar{Q}_0(a, w) := \text{expit}(0.2a - w_1 + 2w_1w_2)$.

We implemented and compared the performance of the following six distinct estimators:

(1) the standard, uncorrected targeted minimum loss-based estimator;

(2) the corrected targeted minimum loss-based estimator using bivariate nuisance regression, as proposed in van der Laan (2014);

(3) the corrected targeted minimum loss-based estimator using univariate nuisance regressions, as introduced in Theorem 1;

(4) the standard, uncorrected one-step estimator, commonly referred to as the augmented inverse probability weighted estimator;

(5) the corrected one-step estimator using bivariate nuisance regression;

(6) the corrected one-step estimator using univariate nuisance regressions, as displayed in (3).

These estimators were evaluated in each of the three following scenarios:

I. only outcome regression consistently estimated;

II. only propensity score consistently estimated;

III. both outcome regression and propensity score consistently estimated.

The consistently-estimated nuisance parameter, either the outcome regression or the propensity score, was estimated using a bivariate kernel regression estimator with bandwidth selected using cross-validation (Racine and Li, 2004), while the inconsistently-estimated nuisance parameter was estimated using a logistic regression model with main terms only, thus ignoring the interaction between W_1 and W_2. The reduced-dimension nuisance parameters required for the additional correction procedure involved in computing estimators (2), (3), (5) and (6) were estimated using the Nadaraya-Watson estimator with bandwidth selected using cross-validation.

For scenarios I and II, we considered sample sizes $n = 250, 500, 1000, 3000, 5000, 9000$ to study the characteristics of the estimators. For scenario III, theory dictates that all estimators considered are asymptotically equivalent, and so, we only focused on sample sizes $n = 100, 250, 500, 750, 1000$. For each sample size, we randomly generated 5000 data sets. We summarized estimator performance based on four criteria: bias, bias times root-n, nominal coverage of 95% confidence intervals, and accuracy of the standard error estimator. The fourth criterion was studied by comparing the Monte Carlo variance of the estimator and the average value of the estimated variance across simulations. We used these summaries to examine the following hypotheses based on the theoretical developments above:

(a) in scenarios I and II, the bias of estimators (1), (4), (5) and (6) tends to zero slower than one over root-n, whereas that of estimators (2) and (3) does so faster than one over root-n;

(b) in scenarios I and II, the slow convergence of the bias for estimators (1), (4), (5) and (6) adversely affects the nominal confidence interval coverage, while with the corrected targeted minimum loss-based estimators (2) and (3) have asymptotically nominal coverage;

(c) in scenarios I and II, influence function-based variance estimators are accurate for the corrected estimators (2), (3), (5) and (6), but not for the uncorrected estimators (1) and (4);

(d) in scenario III, all estimators have approximately the same behavior.

5.2 Results

We first focus on the results pertaining to scenario I, in which only the outcome regression is consistently estimated. In the top left panel of Figure 1, the bias of each estimator tends to zero, illustrating the conventional double-robustness of these estimators. However, the top right panel supports hypothesis (a) in that the bias of the uncorrected estimators clearly tends to zero slower than one over root-n, while the bias of the corrected targeted minimum loss-based estimators tends to zero faster than this rate. The bias of the corrected one-step estimators is reduced relative to the uncorrected estimators, and for the sample sizes considered, we do not yet see the expected divergence in the bias when inflated by root-n. The bottom left panel indicates strong support for hypothesis (b) in that the coverage of intervals based on the uncorrected estimators is not only far from the nominal level but also U-shaped, suggesting worsening coverage in larger samples, as is expected based on our arguments in Section 3. Intervals based on the corrected estimators have approximately nominal coverage in moderate and large samples. The lower right panel indicates that the bias is not the only factor driving the poor coverage of intervals based on the uncorrected estimators: the variance estimators are also anti-conservative. The

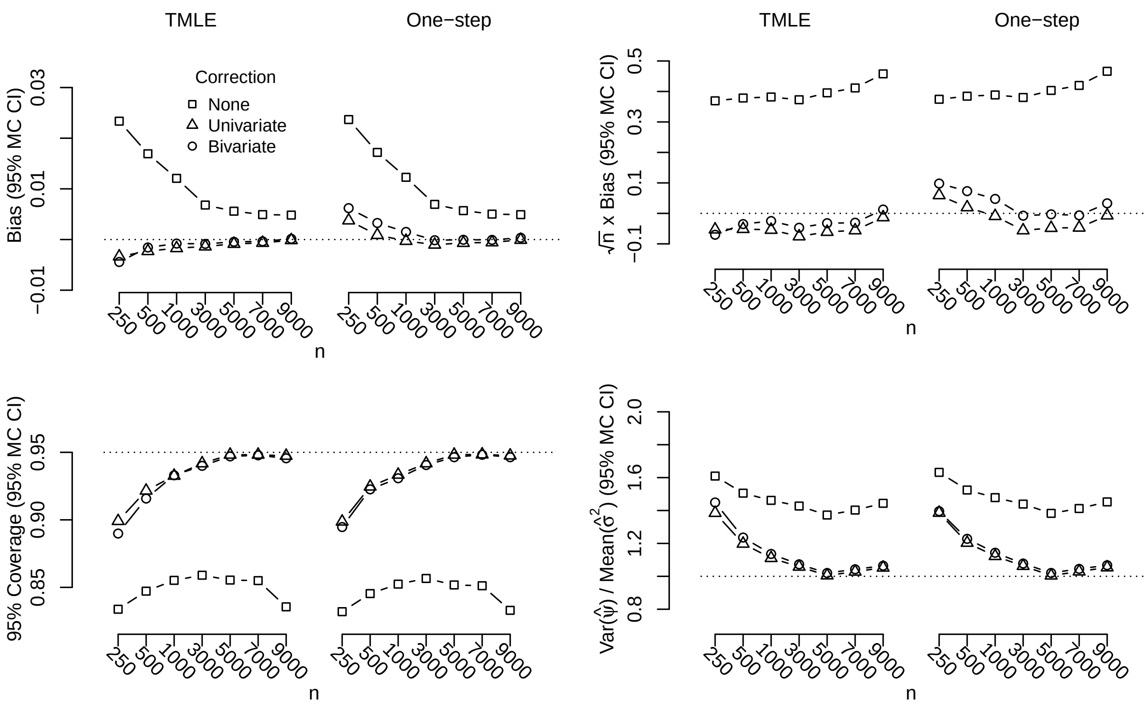

Figure 1: Scenario I: only outcome regression consistently estimated

variance estimators for the corrected estimators are approximately accurate in larger samples, thus supporting hypothesis (c).

We now discuss the results for the scenario in which only the propensity regression is consistently estimated, as summarized in Figure 1. In the top right panel, we see again that the bias of the uncorrected estimators tends to zero slower than one over root-n. In this case, we also find that this is true of the corrected one-step estimators. In contrast, the bias of the corrected targeted minimum loss-based estimators appears to converge to zero faster than one over root-n. The bottom left panel partially supports hypothesis (b): intervals based on the uncorrected estimators achieve near-nominal coverage for moderately large sample sizes in spite of the large bias of these estimators. However, we again find the expected U-shape, with an eventual downturn in coverage as the sample size increases. Intervals based on the corrected targeted minimum loss-based estimators using bivariate nuisance regression have improved coverage throughout, and intervals based on either corrected targeted minimum loss-based estimators have nearly nominal coverage in larger samples. Intervals based on the corrected one-step estimator with the univariate correction achieve approximately nominal coverage, while those based on the one-step estimator with bivariate correction do not, likely due to larger bias, as illustrated in the upper right panel. The bottom right panel shows that the variance estimator for the uncorrected one-step estimator is conservative, while that based on the uncorrected targeted min-

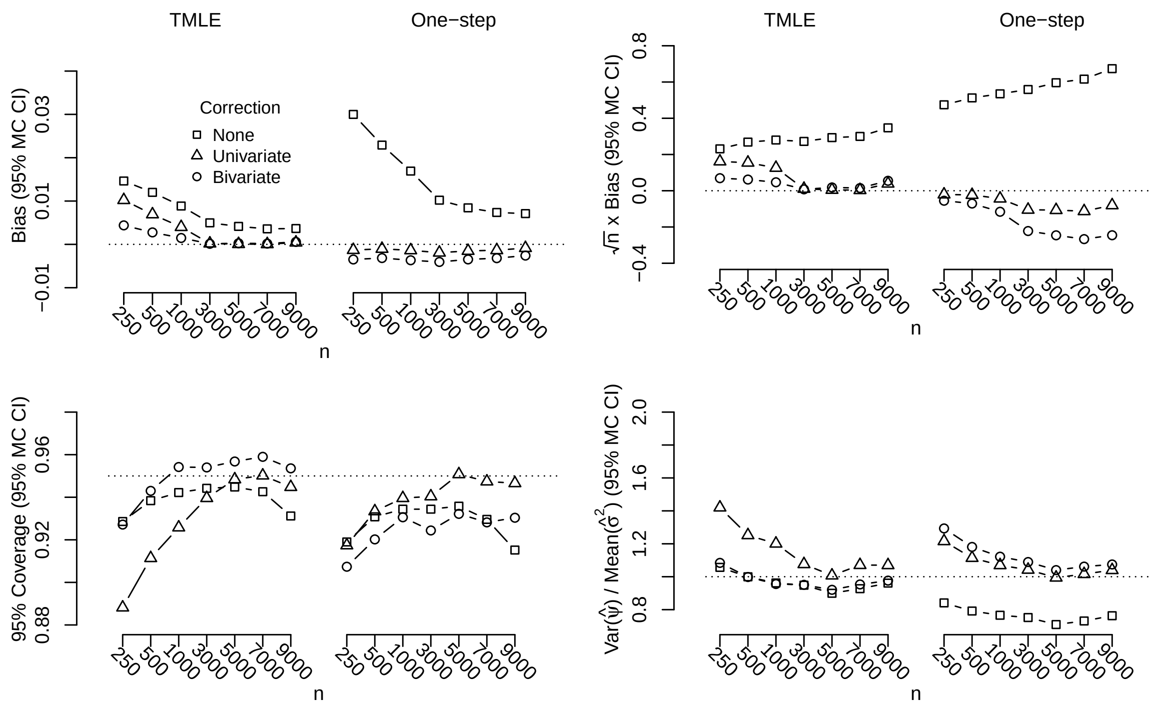

Figure 2: Scenario II: only propensity score consistently estimated

imum loss-based estimator is approximately accurate. The variance estimators based on corrected one-step or targeted minimum loss-based estimators are found to be valid throughout, though the variance estimator based on targeted minimum loss-based using univariate nuisance regressions appears to be substantially anti-conservative in smaller samples.

Finally, Figure 3 supports hypothesis (d): when both the propensity score and outcome regression are consistently estimated, all of the estimators perform approximately equally well, even in smaller samples. This suggests that implementing the correction procedures needed to achieve doubly-robust asymptotic linearity and inference does not come at a cost in terms of estimator performance in situations where the additional corrections are not needed.

6 Concluding remarks

As highlighted earlier, an interesting finding of this work is that it is possible to theoretically guarantee doubly-robust inference under mild conditions using targeted minimum loss-based estimation but not with the more popular one-step approach. While we found the corrected one-step estimators to perform relatively well in simulations, we cannot expect this to hold in all scenarios since theory suggests otherwise.

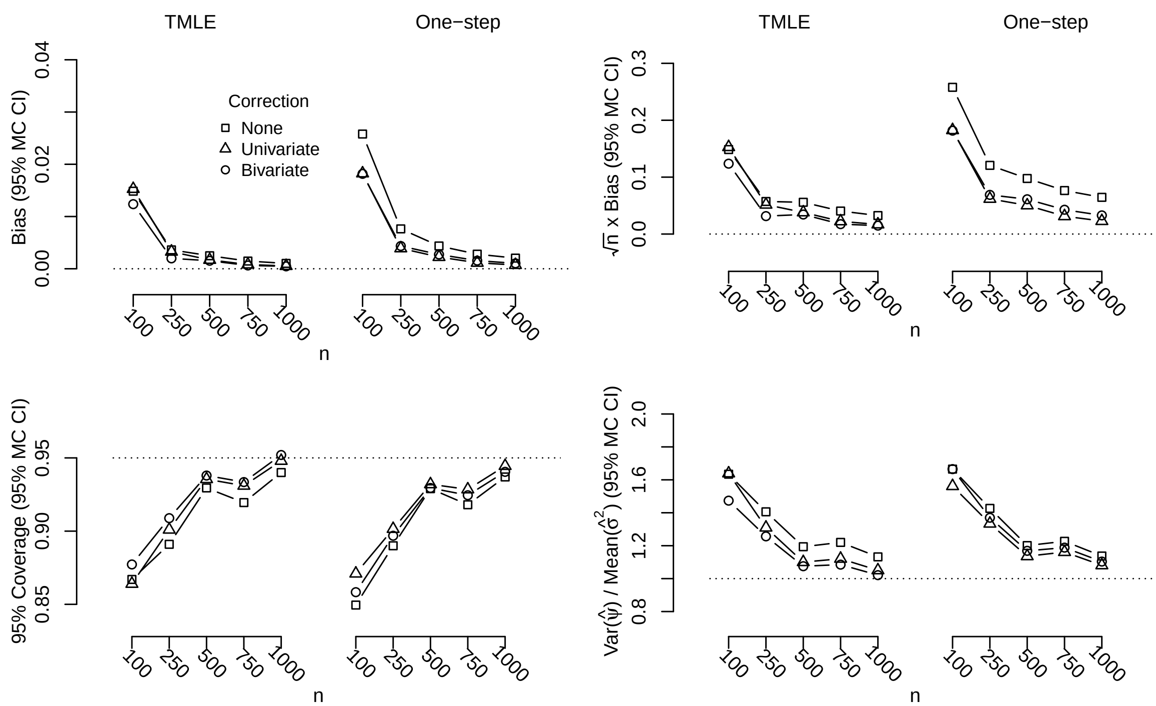

Figure 3: Scenario III: both outcome regression and propensity score consistently estimated

Therefore, in spite of its computational complexity, targeted minimum loss-based estimation may be the preferred approach for providing doubly-robust inference.

It may be fruitful to incorporate universally least favorable parametric submodels, as introduced in van der Laan (2015), into the targeted minimum loss-based estimation algorithms utilized here. Such submodels facilitate the construction of estimators using minimal additional data fitting in the bias-reduction step of the algorithm. Rather than requiring iterations to perform bias reduction, use of these submodels would yield algorithms converging in only a single step. This could result in more expedient computational implementations as well as improved performance in finite samples, particularly in extensions of this work to more complex parameters, including average treatment effects defined by longitudinal interventions.

Appendix A: First-order expansion of remainder

We derive equation (2) and sufficient conditions under which it holds. We note that

$$\begin{aligned}R(\bar{Q}_n, \bar{Q}_0, g_n, g_0) &= P_0\left\{(\bar{Q}_n - \bar{Q}_0)\left(\frac{g_n - g_0}{g_n}\right)\right\}\\ &= P_0\left\{(\bar{Q}_n - \bar{Q}_0)\left(\frac{g_n - g_0}{g}\right)\right\} + R_{1n}\ ,\end{aligned}$$

where we define the second-order remainder term $R_{1,n} := P_0\{(\bar{Q}_n - \bar{Q}_0)(g_n - g_0)(g - g_n)/(g_n g)\}$. Adding and subtracting $\bar{Q}$ and g and simplifying, we find that

$$\begin{aligned}&P_0\left\{(\bar{Q}_n - \bar{Q}_0)\left(\frac{g_n - g_0}{g}\right)\right\}\\ &\quad = P_0\left\{(\bar{Q}_n - \bar{Q})\left(\frac{g - g_0}{g}\right)\right\} + P_0\left\{(\bar{Q} - \bar{Q}_0)\left(\frac{g_n - g}{g}\right)\right\} + R_{2n}\ ,\end{aligned}$$

where we define the second-order term $R_{2,n} := P_0\{(\bar{Q}_n - \bar{Q})(g_n - g)/g\}$. Assuming that either $\bar{Q} = \bar{Q}_0$ or $g = g_0$, we can write

$$P_0\left\{(\bar{Q} - \bar{Q}_0)\left(\frac{g_n - g}{g}\right)\right\} = I(g = g_0)P_0\left\{(\bar{Q} - \bar{Q}_0)\left(\frac{g_n - g_0}{g}\right)\right\}, \tag{6}$$

$$P_0\left\{(\bar{Q}_n - \bar{Q})\left(\frac{g - g_0}{g}\right)\right\} = I(\bar{Q} = \bar{Q}_0)P_0\left\{(\bar{Q}_n - \bar{Q}_0)\left(\frac{g - g_0}{g}\right)\right\}. \tag{7}$$

Examining (6) and with some abuse of notation, we note that

$$P_0\left\{(\bar{Q} - \bar{Q}_0)\left(\frac{g_n - g_0}{g}\right)\right\} = -P_0\left\{\frac{A}{g_0^2}\left(Y - \bar{Q}\right)(g_n - g_0)\right\} = -P_0\left\{\frac{\bar{Q}_{0n,r}}{g_0}(g_n - g_0)\right\},$$

where we set $\bar{Q}_{0n,r}(w) := \mathrm{E}_0\{Y - \bar{Q}(W) \mid g_n(W) = g_n(w), g_0(W) = g_0(w)\}$. Then, we may write

$$
\begin{aligned}
&- P_0\left\{ \frac{\bar{Q}_{0n,r}}{g_0}(g_n - g_0) \right\} \\
&\qquad = -(P_n - P_0)\left\{ \frac{\bar{Q}_{n,r}}{g_0}\ (A - g_n) \right\} + B_{A,n}(\bar{Q}_{n,r}, g_n) + R_{3n} + R_{4n} + M_{1n},
\end{aligned}
$$

where we define

$$
\begin{aligned}
&R_{3n} := P_0\left\{ \left(\frac{\bar{Q}_{0n,r} - \bar{Q}_{0,r}}{g_0} \right)(g_0 - g_n) \right\}, \quad R_{4n} := P_0\left\{ \left(\frac{\bar{Q}_{0,r}}{g_0} - \frac{\bar{Q}_{n,r}}{g_n} \right)(g_0 - g_n) \right\}, \\
&M_{1n} := (P_n - P_0)\left\{ D_A(\bar{Q}_{n,r}, g_n) - D_A(\bar{Q}_{0,r}, g_0) \right\}.
\end{aligned}
$$

If, for example, each of $\bar{Q}_{0n,r} - \bar{Q}_{n,r}$, $\bar{Q}_{n,r} - \bar{Q}_{0,r}$ and $g_n - g_0$ are $o_P(n^{-1/4})$ in $L^2(P_0)$ norm, it generally follows that R_{3n} and R_{4n} are both $o_P(n^{-1/2})$. Furthermore, if $D_A(\bar{Q}_{n,r}, g_n)$ falls in a P_0-Donsker class with probability tending to one and $P_0\{D_A(\bar{Q}_{n,r}, g_n) - D_A(\bar{Q}_{0,r}, g_0)\}^2 = o_P(1)$, it also follows that $M_{1n} = o_P(n^{-1/2})$.

Now, examining (7) and again allowing some abuse of notation, we find that

$$
\begin{aligned}
P_0\left\{ (\bar{Q}_n - \bar{Q}_0)\left(\frac{g - g_0}{g} \right) \right\} &= -P_0\left\{ (\bar{Q}_n - \bar{Q}_0)\left(\frac{A - g}{g} \right) \right\} \\
&= P_0\left\{ \frac{A}{g_{0n,r}} \frac{g_{0n,r} - g}{g}(Y - \bar{Q}_n) \right\},
\end{aligned}
$$

where we define the addition nuisance parameter $g_{0n,r}(\bar{Q}_n, \bar{Q}_0, g) := E_0\{A \mid \bar{Q}_n, \bar{Q}_0, g\}$. Algebraic manipulation allows us to write

$$
\begin{aligned}
&P_0\left\{ \frac{A}{g_{0n,r}} \frac{g_{0n,r} - g}{g}(Y - \bar{Q}_n) \right\} \\
&\qquad = -(P_n - P_0)D_Y(\bar{Q}_0, g_{0,r}) + B_{Y,n}(\bar{Q}_n, g_{n,r}) + R_{5,n} + R_{6,n} + R_{7,n} + M_{2,n}\ ,
\end{aligned}
$$

where we define

$$
\begin{aligned}
&R_{5n} := -P_0\left\{ \frac{g_{0n,r} - g_{0,r}}{g}(\bar{Q}_n - \bar{Q}_0) \right\}, \quad R_{6n} := -P_0\left\{ \frac{g_{n,r} - g_{0,r}}{g_{n,r}}\ (\bar{Q}_n - \bar{Q}_0) \right\}, \\
&R_{7n} := -P_0\left\{ \frac{g_{0,r}}{g_n g}(g_n - g)(\bar{Q}_n - \bar{Q}_0) \right\}, \\
&M_{2n} := (P_n - P_0)\left\{ D_Y(\bar{Q}_n, g_{n,r}) - D_Y(\bar{Q}_0, g_{0,r}) \right\}.
\end{aligned}
$$

If, for example, each of $\bar{Q}_n - \bar{Q}_0$, $g_{0n,r} - g_{0,r}$ and $g_{n,r} - g_{0,r}$ are $o_P(n^{-1/4})$ in $L^2(P_0)$ norm, it generally follows that R_{5n}, R_{6n} and R_{7n} are $o_P(n^{-1/2})$. Furthermore, if $D_Y(\bar{Q}_n, g_{n,r})$ falls in a P_0-Donsker class with probability tending to one and $P_0\{D_Y(\bar{Q}_n, g_{n,r}) - D_Y(\bar{Q}_0, g_{0,r})\}^2 = o_P(1)$, it also follows that $M_{2n} = o_P(n^{-1/2})$.

The above derivations directly imply (2) with $R_n^* := R_{1n} + R_{2n}$, $R_{Q,n} := R_{3n} + R_{4n} + M_{1n}$ and $R_{g,n} := R_{5n} + R_{6n} + R_{7n} + M_{2n}$.

Appendix B: Derivation of reduced-dimension remainder representation

We proceed similarly as above but now with regards to (7). With some abuse of notation, we have

$$\begin{aligned}P_0\left\{(\bar{Q}_n-\bar{Q}_0)\left(\frac{g-g_0}{g}\right)\right\} &= -P_0\left\{(\bar{Q}_n-\bar{Q}_0)\left(\frac{A-g}{g}\right)\right\}\\ &= -P_0\left\{g_{2,0n,r}(\bar{Q}_n-\bar{Q}_0)\right\} = P_0\left\{\frac{A}{g_{1,0n,r}}g_{2,0n,r}(Y-\bar{Q}_n)\right\},\end{aligned}$$

where we define nuisances $g_{1,0n,r}(\bar{Q}_n,\bar{Q}_0)(w) := E_0\{A \mid \bar{Q}_n(W)=\bar{Q}_n(w),\bar{Q}_0(W)=\bar{Q}_0(w)\}$ and $g_{2,0n,r}(\bar{Q}_n,\bar{Q}_0,g)(w) := E_0[\{A-g(W)\}/g(W) \mid \bar{Q}_n(W)=\bar{Q}_n(w),\bar{Q}_0(W)=\bar{Q}_0(w)]$. We can then write

$$\begin{aligned}&P_0\left\{\frac{A}{g_{1,0n,r}}g_{2,0n,r}(Y-\bar{Q}_n)\right\}\\ &\qquad = -(P_n-P_0)D_Y(\bar{Q}_0,g_{1,0,r},g_{2,0,r}) + B_{Y,n}(\bar{Q}_n,g_{1,n,r},g_{2,n,r}) + \tilde{R}_{5n}+\tilde{R}_{6n}+\tilde{M}_{2n},\end{aligned}$$

where we define

$$\begin{aligned}\tilde{R}_{5n} &:= P_0\left\{\left(\frac{A}{g_{1,0n,r}}g_{2,0n,r}-\frac{A}{g_{1,0,r}}g_{2,0,r}\right)(Y-\bar{Q}_n)\right\},\\ \tilde{R}_{6n} &:= P_0\left\{\left(\frac{A}{g_{1,0,r}}g_{2,0,r}-\frac{A}{g_{1,n,r}}g_{2,n,r}\right)(Y-\bar{Q}_n)\right\},\\ \tilde{M}_{2n} &:= (P_n-P_0)\left\{D_Y(\bar{Q}_n,g_{1,n,r},g_{2,n,r})-D_Y(\bar{Q}_0,g_{1,0,r},g_{2,0,r})\right\}.\end{aligned}$$

If, for example, each of $\bar{Q}_n-\bar{Q}_0$, $g_{2,0n,r}-g_{2,0,r}$ and $g_{2,n,r}-g_{2,0,r}$ are $o_P(n^{-1/4})$ in $L^2(P_0)$ norm, it generally follows that $\tilde{R}_{5n}$ and $\tilde{R}_{6n}$ are $o_P(n^{-1/2})$. Furthermore, if $D_Y(\bar{Q}_n,g_{1,n,r},g_{2,n,r})$ falls in a P_0-Donsker class with probability tending to one and $P_0\{D_Y(\bar{Q}_n,g_{1,n,r},g_{2,n,r})-D_Y(\bar{Q}_0,g_{1,0,r},g_{2,0,r})\}^2=o_P(1)$, it also follows that $\tilde{M}_{2n}=o_P(n^{-1/2})$.

This implies that (2) holds with $R_n^* := R_{1,n}+R_{2,n}$, $R_{Q,n} := R_{3n}+R_{4n}+M_{1n}$ and $R_{g,n} := \tilde{R}_{5n}+\tilde{R}_{6n}+\tilde{M}_{2n}$ when the alternative reduced-dimension parametrization of the remainder is used.

Appendix C: Behavior of unnecessary correction terms

We first examine the behavior of $B_{A,n}(\bar{Q}_{n,r},g_n)$ when $\bar{Q}=\bar{Q}_0$. We note that

$$B_{A,n}(\bar{Q}_{n,r},g_n) = P_nD_A(\bar{Q}_{n,r},g_n) = P_0D_A(\bar{Q}_{n,r},g_n)+M_{A,n},$$

where we define the empirical process term

$$M_{A,n} := (P_n - P_0)\{D_A(\bar{Q}_{n,r}, g_n) - D_A(\bar{Q}_{0,r}, g)\} ,$$

which can reasonably be assumed to be $o_P(n^{-1/2})$. The second equality is a consequence of the fact that $D_A(\bar{Q}_{0,r}, g) = 0$ for all g, which follows because $\bar{Q}_{0,r} = 0$. With some abuse of notation, we can write

$$\begin{aligned} P_0 D_A(\bar{Q}_{n,r}, g_n) &= P_0 \left\{ \frac{\bar{Q}_{n,r}}{g_n}(A - g_n) \right\} \\ &= P_0 \left\{ \frac{\bar{Q}_{n,r}}{g_n}(g_0 - g_n) \right\} \\ &= P_0 \left\{ \frac{\bar{Q}_{n,r}}{g}(g_0 - g) \right\} + R_{A,n}, \end{aligned}$$

where we define

$$R_{A,n} := P_0 \left\{ \bar{Q}_{n,r} \left(\frac{g - g_n}{g_n g} \right) (g_0 - g_n) \right\} + P_0 \left\{ \frac{\bar{Q}_{n,r}}{g}(g - g_n) \right\},$$

which is $o_P(n^{-1/2})$ under the rate conditions outlined in Appendices A and B.

We now examine the behavior of $B_{Y,n}(\bar{Q}_n, g_{1,n,r}, g_{2,n,r})$ when $g = g_0$. We have

$$B_{Y,n}(\bar{Q}_n, g_{1,n,r}, g_{2,n,r}) = P_n D_Y(\bar{Q}_n, g_{1,n,r}, g_{2,n,r}) = P_0 D_Y(\bar{Q}_n, g_{1,n,r}, g_{2,n,r}) + M_{Y,n},$$

where we define the empirical process term $M_{Y,n} := (P_n - P_0)\{D_Y(\bar{Q}_n, g_{1,n,r}, g_{2,n,r}) - D_Y(\bar{Q}, g_{1,0,r}, g_{2,0,r})\}$, which can reasonably be assumed to be $o_P(n^{-1/2})$. As above, the second equality is a consequence of the fact that $D_Y(\bar{Q}, g_{1,0,r}, g_{2,0,r}) = 0$ for all $\bar{Q}$, which follows because $g_{2,0,r} = 0$ when $g = g_0$. With some abuse of notation, we can write

$$\begin{aligned} P_0 D_Y(\bar{Q}_n, g_{1,n,r}, g_{2,n,r}) &= P_0 \left\{ \frac{A}{g_{1,n,r}} g_{2,n,r}(Y - \bar{Q}_n) \right\} = P_0 \left\{ \frac{A}{g_{1,n,r}} g_{2,n,r}(\bar{Q}_0 - \bar{Q}_n) \right\} \\ &= P_0 \left\{ \frac{A}{g_{1,0,r}} (\bar{Q}_0 - \bar{Q}) g_{2,n,r} \right\} + R_{Y,n}, \end{aligned}$$

where we define

$$R_{Y,n} := P_0 \left\{ A \left(\frac{g_{1,0,r} - g_{1,n,r}}{g_{1,0,r} g_{1,n,r}} \right) g_{2,n,r}(\bar{Q}_0 - \bar{Q}_n) \right\} + P_0 \left\{ \frac{A}{g_{1,0,r}} g_{2,n,r}(\bar{Q} - \bar{Q}_n) \right\},$$

which is $o_P(n^{-1/2})$ under the rate conditions above.

References

Heejung Bang and James M Robins. Doubly Robust Estimation in Missing Data and Causal Inference Models. *Biometrics*, 61(4):962–973, 2005. ISSN 1541-0420.

P. Bickel, C. Klaassen, Y. Ritov, and J. Wellner. *Efficient and Adaptive Estimation for Semiparametric Models*. Johns Hopkins University Press Baltimore, 1997. ISBN 0801845416.

Ildar Abdulovich Ibragimov, Rafail Z Has' minskii, and Samuel Kotz. *Statistical Estimation: Asymptotic Theory*, volume 2. Springer-Verlag New York, 1981.

Joseph DY Kang and Joseph L Schafer. Demystifying Double Robustness: A Comparison of Alternative Strategies for Estimating a Population Mean from Incomplete Data. *Statistical science*, pages 523–539, 2007.

Johann Pfanzagl. *Contributions to a General Asymptotic Statistical Theory*. Springer, 1982.

Kristin E Porter, Susan Gruber, Mark J van der Laan, and Jasjeet S Sekhon. The relative performance of targeted maximum likelihood estimators. *The International Journal of Biostatistics*, 7(1):1–34, 2011.

Jeff Racine and Qi Li. Nonparametric estimation of regression functions with both categorical and continuous data. *Journal of Econometrics*, 119(1):99–130, 2004.

James M Robins, Andrea Rotnitzky, and Lue Ping Zhao. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association*, 89(427):846–866, 1994. ISSN 0162-1459.

Donald B Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5):688, 1974.

Mark van der Laan and Daniel Rubin. Targeted maximum likelihood learning. *The International Journal of Biostatistics*, 2(1):1–40, 2006.

Mark J van der Laan. Targeted estimation of nuisance parameters to obtain valid statistical inference. *The International Journal of Biostatistics*, 10:29–57, 2014.

Mark J van der Laan. One-step targeted minimum loss-based estimation based on universal least favorable one-dimensional submodels. *UC Berkeley Division of Biostatistics Working Paper Series*, 2015.

Mark J van der Laan and Sandrine Dudoit. Unified cross-validation methodology for selection among estimators and a general cross-validated adaptive epsilon-net estimator: Finite sample oracle inequalities and examples. *UC Berkeley Division of Biostatistics Working Paper Series*, 2003.

Mark J van der Laan and Eric C Polley. Super learner. *Statistical Applications in Genetics and Molecular Biology*, 6(1):1–23, 2007.

Mark J van der Laan and James M Robins. *Unified Methods for Censored Longitudinal Data and Causality*. Springer, 2003. ISBN 0387955569.

Mark J van der Laan and Sherri Rose. *Targeted Learning: Causal Inference for Observational and Experimental Data*. Springer, 2011. ISBN 1441997822.

Mark J van der Laan, Sandrine Dudoit, and Aad W van der Vaart. The cross-validated adaptive epsilon-net estimator. *Statistics & Decisions*, 24(3):373–395, 2006.

Aad W van der Vaart and Jon A Wellner. *Weak Convergence*. Springer, 1996.

Aad W van der Vaart, Sandrine Dudoit, and Mark J van der Laan. Oracle inequalities for multi-fold cross validation. *Statistics & Decisions*, 24(3):351–371, 2006.

Karel Vermeulen and Stijn Vansteelandt. Biased-reduced doubly robust estimation. *Journal of the American Statistical Association*, (forthcoming):00–00, 2014.

Karel Vermeulen and Stijn Vansteelandt. Data-adaptive bias-reduced doubly robust estimation. *The International Journal of Biostatistics*, 12(1):253–282, 2016.