Karel Vermeulen* and Stijn Vansteelandt

Data-Adaptive Bias-Reduced Doubly Robust Estimation

DOI 10.1515/ijb-2015-0029

Abstract: Doubly robust estimators have now been proposed for a variety of target parameters in the causal inference and missing data literature. These consistently estimate the parameter of interest under a semiparametric model when one of two nuisance working models is correctly specified, regardless of which. The recently proposed bias-reduced doubly robust estimation procedure aims to partially retain this robustness in more realistic settings where both working models are misspecified. These so-called *bias-reduced doubly robust estimators* make use of special (finite-dimensional) nuisance parameter estimators that are designed to locally minimize the squared asymptotic bias of the doubly robust estimator in certain directions of these finite-dimensional nuisance parameters under misspecification of both parametric working models. In this article, we extend this idea to incorporate the use of data-adaptive estimators (infinite-dimensional nuisance parameters), by exploiting the bias reduction estimation principle in the direction of only one nuisance parameter. We additionally provide an asymptotic linearity theorem which gives the influence function of the proposed doubly robust estimator under correct specification of a parametric nuisance working model for the missingness mechanism/propensity score but a possibly misspecified (finite- or infinite-dimensional) outcome working model. Simulation studies confirm the desirable finite-sample performance of the proposed estimators relative to a variety of other doubly robust estimators.

Keywords: bias-reduced doubly robust estimation, causal inference, double robustness, missing data, nuisance parameters, super-learning, targeted maximum likelihood estimation (tmle), working models

1 Introduction

The estimation of most statistical parameters demands the postulation of so-called nuisance working models: models that are not of primary interest but needed to obtain a well-behaved estimator. A typical concern is then whether misspecification of these models may induce bias in the resulting estimator of the target parameter. Doubly robust estimators alleviate this concern: they consistently estimate the target parameter when at least one of two nuisance working models is correctly specified, regardless of which [1].

Doubly robust estimators are now available for many statistical parameters and have become very popular over recent years. This is on the one hand so in view of their robustness against model misspecification, but also because of their desirable efficiency properties in many settings (see e.g., Tsiatis et al. [2], Moore and van der Laan [3], Vermeulen et al. [4]) and the fact that they form a compromise between competing estimators that each rely on a single but different working model. In spite of this, doubly robust estimators have also been the subject of recent debate [5]. The reason is that model misspecification is likely to affect all working models in practice, thus making the double-protection property appear of a more academic interest. Initial reports suggested that modest misspecification of both working models might still deliver doubly robust estimators with reasonably small bias [1, 6]. However, later simulation studies signaled that their (finite-sample) bias and variance can sometimes get severely amplified under misspecification of at least one working model and may become especially severe under misspecification of both working models [5, 7].

Vermeulen and Vansteelandt [8] investigated the usefulness of doubly robust estimators from the perspective that all models are wrong. They discovered that, interestingly, some doubly robust estimators

Article note: Karel Vermeulen is a Research Fellow of the Fund for Scientific Research (FWO).

***Corresponding author: Karel Vermeulen**, Department of Applied Mathematics, Computer Science and Statistics, Ghent University, Ghent, Belgium, E-mail: karelb.vermeulen@ugent.be

Stijn Vansteelandt, Department of Applied Mathematics, Computer Science and Statistics, Ghent University, Ghent, Belgium

partially retain their robustness properties even under misspecification of both working models. In particular, they studied the asymptotic bias of doubly robust estimators for a scalar parameter. They found that without knowledge of the true data-generating law, the asymptotic bias of these estimators can typically be locally minimized in certain directions of the finite-dimensional nuisance parameters. This is possible by making use of specific estimators of the nuisance parameters, which target bias reduction. Vermeulen and Vansteelandt [8] referred to their procedure as *bias-reduced doubly robust estimation*.

Vermeulen and Vansteelandt [8] contrasted their bias-reduced doubly robust estimators with a variety of other doubly robust estimators that are primarily aimed at variance reduction under misspecification of one working model. They found their estimators to be highly competitive, although sometimes slightly more biased than so-called targeted maximum likelihood estimators (TMLE, van der Laan and Rubin [9]). TMLE, an estimation principle that may be used for pathwise differentiable target parameters in semiparametric models, reduces bias by making clever use of data-adaptive learning algorithms (infinite-dimensional nuisance working models) such as ensemble learning, and specifically super-learning [10]. In this article, we will investigate how such data-adaptive learning algorithms can be integrated in the theory of bias-reduced doubly robust estimation to aim at further bias reduction. This will also overcome one of the limitations of the proposal by Vermeulen and Vansteelandt [8], that both nuisance working models must be indexed by nuisance parameters of the same dimension.

The article is organized as follows. In Section 2, we introduce some notation and the inferential problem at stake: estimation of a population mean outcome in the presence of outcome missingness that is explainable by measured auxiliary covariates; equivalently, the estimation of a mean counterfactual outcome in the presence of data on all relevant confounders. In Section 3, we briefly review the theory on bias-reduced doubly robust estimation in the context of the example given in Section 2. Next, in Section 4, we outline the proposed extension of the bias-reduced doubly robust estimator to incorporate data-adaptive learning algorithms. The motivation behind this proposal is to fluctuate an initial data-adaptive estimator for the outcome-covariate associations in a way that aims to reduce the bias in the direction of the finite-dimensional parameter indexing the propensity score model. We therefore call this procedure data-adaptive bias-reduced doubly robust estimation. We also show how to perform inference based on the asymptotic linearity of the estimator (under certain regularity conditions to allow for data-adaptive estimation of the outcome working model and the condition that the parametric model for the propensity score is correct). In Section 5, we illustrate the performance of our proposal relative to the original bias-reduced doubly robust estimator and the TMLE procedure [9]. In Section 6, we illustrate the generic nature of the proposal and show how to implement the proposed strategy in a linear instrumental variable analysis. We end with a discussion in Section 7. Finally, these methods are implemented in an R-function, given in Appendix B.

1.1 Some relevant literature overview

The discussions on the Kang and Schafer article [5, 11–15] revealed that many different doubly robust estimators may be obtained for a given target parameter by varying the choice of nuisance working model estimators. Different choices may lead to doubly robust estimators which all have potentially very different behavior and properties under misspecification of at least one working model but are asymptotically equivalent under correct specification of both working models. As a result, many alternatives to maximum likelihood estimation (MLE) have now been proposed in the literature for the estimation of the nuisance parameters/nuisance working models that index certain doubly robust estimators. Rubin and van der Laan [16], Cao et al. [17] and Tsiatis et al. [18] developed doubly robust estimators in specific missing data models which target minimal asymptotic variance when the missingness model is correctly specified. The TMLE (targeted maximum likelihood estimation) procedure [9, 19] and the procedures of Tan [20] and Rotnitzky et al. [21] guarantee doubly robust estimators that fall within the parameter range, with the latter procedures also having desirable efficiency properties. With the exception of TMLE, all these proposals focus on improving the efficiency of doubly robust estimators under misspecification of the working model for the

dependence of the observed outcome on covariates/confounders in missing data/causal inference models. The collaborative TMLE (C-TMLE) procedure of van der Laan and Gruber [22] additionally focuses on the estimation of the missingness/exposure probabilities in a way that aims to attain an optimal bias-variance trade-off. In contrast, Vermeulen and Vansteelandt [8] focus on bias reduction under the constraints of parametric working models. For a review of many of the aforementioned procedures, we refer to Rotnitzky and Vansteelandt [23]; for a comparison of the performance of many of these alternatives, we refer to Tan [20], Porter et al. [24] and Vermeulen and Vansteelandt [8].

2 Doubly robust estimation of a population mean with incomplete data

Consider the i.i.d. sample $\mathcal{O} = (\mathbf{O}_1, \dots, \mathbf{O}_n)$ of size n , where the observed data vector $\mathbf{O} = (RY, R, \mathbf{X})$ is distributed according to a true underlying but unknown probability distribution P_0 . We let Y denote the outcome of interest which is susceptible to missingness, formalized using the missingness indicator R where $R=1$ if Y is observed and $R=0$ if Y is missing. We assume that missingness is explainable by $\mathbf{X}$, a collection of auxiliary covariates, so that the missing at random (MAR, Rubin [25]) assumption holds: $Y \perp\!\!\!\perp R | \mathbf{X}$, where we use $\perp\!\!\!\perp$ to denote statistical independence. Throughout, for some parameter λ , we use λ_0 to denote its (unknown) population value; in particular, $E(Y) = \psi_0$. Furthermore, for a function f of the observed data and a probability distribution P , we let Pf denote the integral $\int f dP$.

Doubly robust estimation of ψ_0 requires specification of two nuisance working models [26]. First, we need a working model $\mathcal{Q} = \{\bar{Q}(\mathbf{X}) | \bar{Q} \text{ in some class of functions}\}$ for the true conditional mean outcome $\bar{Q}_0(\mathbf{X}) = E_0(Y | \mathbf{X})$ (which equals $E_0(Y | R=1, \mathbf{X})$ because of MAR). Let $\mathcal{M}(\mathcal{Q})$ denote the statistical model for the joint distribution of $\mathbf{O}$ implied by the restrictions imposed by $\mathcal{Q}$. Let $\hat{\bar{Q}}_n(\mathbf{X})$ denote an estimator of $\bar{Q}_0(\mathbf{X})$ with probability limit $\bar{Q}^*(\mathbf{X})$; that is, $\hat{\bar{Q}}_n(\mathbf{X}) \xrightarrow{P} \bar{Q}^*(\mathbf{X})$. Under $\mathcal{M}(\mathcal{Q})$, $\bar{Q}^*(\mathbf{X}) = \bar{Q}_0(\mathbf{X})$, but not otherwise. Second, we need a working model $\mathcal{G} = \{g(\mathbf{X}) | g \text{ in some class of functions}\}$ for the true missingness mechanism $g_0(\mathbf{X}) = P_0(R=1 | \mathbf{X})$, for which we assume *positivity*: $g_0(\mathbf{X}) \geq \delta > 0$ with probability one (van der Laan and Rose [19, chap. 10]). Let $\mathcal{M}(\mathcal{G})$ denote the statistical model for the joint distribution of $\mathbf{O}$ implied by the restrictions imposed by the working model $\mathcal{G}$. Let $\hat{g}_n(\mathbf{X})$ denote an estimator of $g_0(\mathbf{X})$ with probability limit $g^*(\mathbf{X})$; that is, $\hat{g}_n(\mathbf{X}) \xrightarrow{P} g^*(\mathbf{X})$. Under $\mathcal{M}(\mathcal{G})$, $g^*(\mathbf{X}) = g_0(\mathbf{X})$. A doubly robust estimator of ψ_0 is then obtained via

$$\hat{\psi}_n(\hat{\bar{Q}}_n, \hat{g}_n) = \frac{1}{n} \sum_{i=1}^n \left[\frac{R_i}{\hat{g}_n(\mathbf{X}_i)} \{Y_i - \hat{\bar{Q}}_n(\mathbf{X}_i)\} + \hat{\bar{Q}}_n(\mathbf{X}_i) \right]. \quad (1)$$

It can easily be verified from the above expression that this estimator is consistent for ψ_0 under the union model $\mathcal{M}(\mathcal{Q}) \cup \mathcal{M}(\mathcal{G})$: as soon as one but not necessarily both working models are correctly specified. Interestingly (provided sufficient regularity for the working models), it is also locally efficient [27] at the intersection model $\mathcal{M}(\mathcal{Q}) \cap \mathcal{M}(\mathcal{G})$ in the following sense: it has smallest asymptotic variance within the class of all estimators that are consistent and asymptotically normal under $\mathcal{M}(\mathcal{G})$, provided that also $\mathcal{M}(\mathcal{Q})$ is correctly specified. At the intersection model, it has the following simple expansion

$$\hat{\psi}_n(\hat{\bar{Q}}_n, \hat{g}_n) - \psi_0 = (P_0 - P_n) \{D^*(\bar{Q}_0, g_0; \psi_0)\} + R_n, \quad (2)$$

with (under sufficient regularity on the working models $\mathcal{Q}$ and $\mathcal{G}$, see for instance Appendix A.1 of van der Laan and Rose [19]) the remainder $R_n = o_p(1/\sqrt{n})$, P_n the empirical distribution which puts mass $1/n$ on each observation $\mathbf{O}_i$, $i = 1, \dots, n$ and $D^*(\bar{Q}_0, g_0; \psi_0)$ the efficient influence function, given by

$$D^*(\bar{Q}_0, g_0; \psi_0)(\mathbf{O}) = \frac{R}{g_0(\mathbf{X})} \{Y - \bar{Q}_0(\mathbf{X})\} + \bar{Q}_0(\mathbf{X}) - \psi_0. \quad (3)$$

The expansion (2) is attractive because it (and thus the asymptotic distribution of the doubly robust estimator) is the same when both nuisance working models are correctly specified, no matter how the nuisance parameters are estimated and no matter whether they are estimated or known. Finally, note that by construction

$$P_n \left[D^* \left\{ \hat{Q}_n, \hat{g}_n; \hat{\psi}_n(\hat{Q}_n, \hat{g}_n) \right\} \right] = 0. \quad (4)$$

3 Bias-reduced doubly robust estimation

Bias-reduced doubly robust estimation is a generic estimation strategy for the nuisance parameters indexing parametric nuisance working models $\mathcal{G}$ and $\mathcal{Q}$, aimed at bias reduction under misspecification of both these models. In Section 3.1, we will introduce such parametric models in the context of the missing data example. In Section 3.2, we will then review the bias-reduced estimation principle.

3.1 Parametric nuisance working models and MLE

Suppose we parameterize the working model for the conditional mean outcome by a q -dimensional parameter γ : $\mathcal{Q}_\gamma = \{\bar{Q}(\gamma)(\mathbf{X}) | \gamma \in \mathbb{R}^q\}$ with $\bar{Q}(\gamma)(\mathbf{X}) = Q(\gamma^T \mathbf{k}(\mathbf{X}))$, Q an appropriate inverse link function and $\mathbf{k} = (1, k_1, \dots, k_{q-1})$; e.g., a linear regression model $\bar{Q}(\gamma)(\mathbf{X}) = \gamma_1 + \gamma_2^T \mathbf{X}$ for a continuous outcome Y . If the model $\mathcal{M}(\mathcal{Q}_\gamma)$ is correctly specified and thus includes P_0 , we let γ_0 be such that $\bar{Q}(\gamma_0) = \bar{Q}_0$. Further, let the p -dimensional parameter α parameterize the working model for the missingness mechanism: $\mathcal{G}_\alpha = \{g(\alpha)(\mathbf{X}) | \alpha \in \mathbb{R}^p\}$ where $g(\alpha)(\mathbf{X}) = G(\alpha^T \mathbf{l}(\mathbf{X}))$, G is an appropriate inverse link function and $\mathbf{l} = (1, l_1, \dots, l_{p-1})$; e.g., a logistic regression model $g(\alpha)(\mathbf{X}) = \text{expit}(\alpha_1 + \alpha_2^T \mathbf{X})$, $\text{expit}(x) = 1/(1 + e^{-x})$. If the model $\mathcal{M}(\mathcal{G}_\alpha)$ is correctly specified and thus includes P_0 , we let α_0 be such that $g(\alpha_0) = g_0$.

Root- n consistent and asymptotically normal estimators $\hat{\gamma}_n$ and $\hat{\alpha}_n$ for the nuisance parameters γ and α can be obtained as solutions to estimating equations $P_n\{D_{\bar{Q}}(\hat{\gamma}_n)\} = \mathbf{0}$ and $P_n\{D_g(\hat{\alpha}_n)\} = \mathbf{0}$ with the estimating functions $D_{\bar{Q}}$ and D_g such that $P_0\{D_{\bar{Q}}(\gamma_0)\} = \mathbf{0}$ if $P_0 \in \mathcal{M}(\mathcal{Q}_\gamma)$ and $P_0\{D_g(\alpha_0)\} = \mathbf{0}$ if $P_0 \in \mathcal{M}(\mathcal{G}_\alpha)$.

Throughout, we will assume that $\hat{\gamma}_n \xrightarrow{p} \gamma^*$ for some γ^* (with $\gamma^* = \gamma_0$ under model $\mathcal{M}(\mathcal{Q}_\gamma)$) and $\hat{\alpha}_n \xrightarrow{p} \alpha^*$ for some α^* (with $\alpha^* = \alpha_0$ under model $\mathcal{M}(\mathcal{G}_\alpha)$) and, moreover, that these estimators are asymptotically linear with influence function $-P_0\{D_{\gamma, \bar{Q}}(\gamma^*)\}^{-1} D_{\bar{Q}}(\gamma^*)$ and $-P_0\{D_{\alpha, g}(\alpha^*)\}^{-1} D_g(\alpha^*)$, where $D_{\gamma, \bar{Q}}(\gamma^*) = \partial D_{\bar{Q}}(\gamma)/\partial \gamma|_{\gamma=\gamma^*}$ and $D_{\alpha, g}(\alpha^*) = \partial D_g(\alpha)/\partial \alpha|_{\alpha=\alpha^*}$.

In practice, maximum likelihood estimation (MLE) (or least squares) is routinely employed to estimate the nuisance parameters. The MLEs $\hat{\gamma}_n^{\text{MLE}}$ and $\hat{\alpha}_n^{\text{MLE}}$ solve the estimating equations $P_n\{D_{\bar{Q}}^{\text{MLE}}(\hat{\gamma}_n^{\text{MLE}})\} = \mathbf{0}$ and $P_n\{D_g^{\text{MLE}}(\hat{\alpha}_n^{\text{MLE}})\} = \mathbf{0}$ with

$$D_{\bar{Q}}^{\text{MLE}}(\gamma)(\mathbf{O}) = R\{Y - \bar{Q}(\gamma)(\mathbf{X})\} Q' \{\gamma^T \mathbf{k}(\mathbf{X})\} \mathbf{k}(\mathbf{X}), \quad (5)$$

$$D_g^{\text{MLE}}(\alpha)(\mathbf{O}) = \frac{R - g(\alpha)(\mathbf{X})}{g(\alpha)(\mathbf{X})\{1 - g(\alpha)(\mathbf{X})\}} G' \{\alpha^T \mathbf{l}(\mathbf{X})\} \mathbf{l}(\mathbf{X}), \quad (6)$$

$Q'(x) = \partial Q(x)/\partial x$ and $G'(x) = \partial G(x)/\partial x$. A doubly robust estimator is then given by $\hat{\psi}_n^{\text{DR, MLE}} \equiv \hat{\psi}_n(\hat{Q}_n^{\text{MLE}}, \hat{g}_n^{\text{MLE}})$, with $\hat{Q}_n^{\text{MLE}} \equiv \bar{Q}(\hat{\gamma}_n^{\text{MLE}})$ and $\hat{g}_n^{\text{MLE}} \equiv g(\hat{\alpha}_n^{\text{MLE}})$. Although MLE is asymptotically efficient and optimal for these nuisance parameters with respect to the corresponding working model, it need not be optimal with respect to the target parameter ψ_0 under misspecification of one of these models. Under such misspecification, the influence function of the doubly robust estimator (and thus its asymptotic distribution) becomes indeed dependent on the choice of root- n consistent estimators of the nuisance parameters (see e.g., Vermeulen and Vansteelandt [8], Proposition 1). This raises the question how to best fit the nuisance working models.

3.2 Bias-reduced doubly robust estimation

3.2.1 Estimation principle

Vermeulen and Vansteelandt [8] aim to reduce the squared asymptotic bias of the doubly robust estimator under misspecification of both working models by making use of particular nuisance parameter estimators. Their starting point is that under possible misspecification of $\mathcal{Q}_\gamma$ and $\mathcal{G}_\alpha$, for arbitrary but fixed nuisance parameter values γ and α , the asymptotic bias of the doubly robust estimator, as a function of P_0 and these nuisance parameters γ and α , is given by $\text{bias}(\gamma, \alpha; \psi_0) = P_0[D^*\{\bar{Q}(\gamma), g(\alpha); \psi_0\}]$. This bias is zero under the union model $\mathcal{M}(\mathcal{Q}_\gamma) \cup \mathcal{M}(\mathcal{G}_\alpha)$, so when we either take $\gamma = \gamma_0$ or $\alpha = \alpha_0$. Away from this model, locally minimizing $\text{bias}^2(\gamma, \alpha; \psi_0)$ – if a local minimum exists – requires finding the values $(\gamma_{\text{BR}}^*, \alpha_{\text{BR}}^*)^T$ that solve the equations $P_0[D_\gamma^*\{\bar{Q}(\gamma_{\text{BR}}^*), g(\alpha_{\text{BR}}^*)\}] = \mathbf{0}$ and $P_0[D_\alpha^*\{\bar{Q}(\gamma_{\text{BR}}^*), g(\alpha_{\text{BR}}^*)\}] = \mathbf{0}$ (Vermeulen and Vansteelandt [8], Theorem 1); here, $D_\gamma^*\{\bar{Q}(\gamma), g(\alpha)\}$ denotes the gradient $\partial D^*\{\bar{Q}(\gamma), g(\alpha); \psi_0\} / \partial \gamma$ and $D_\alpha^*\{\bar{Q}(\gamma), g(\alpha)\}$ denotes the gradient $\partial D^*\{\bar{Q}(\gamma), g(\alpha); \psi_0\} / \partial \alpha$. This is so because these values γ_{BR}^* and α_{BR}^* ensure that $\partial \text{bias}^2(\gamma, \alpha_{\text{BR}}^*; \psi_0) / \partial \gamma|_{\gamma = \gamma_{\text{BR}}^*} = \mathbf{0}$ and $\partial \text{bias}^2(\gamma_{\text{BR}}^*, \alpha; \psi_0) / \partial \alpha|_{\alpha = \alpha_{\text{BR}}^*} = \mathbf{0}$.

The unknown values γ_{BR}^* and α_{BR}^* can be estimated as the solutions $\hat{\gamma}_n^{\text{BR}}$ and $\hat{\alpha}_n^{\text{BR}}$ to the estimating equations

$$P_n[D_\alpha^*\{\bar{Q}(\hat{\gamma}_n^{\text{BR}}), g(\hat{\alpha}_n^{\text{BR}})\}] = \mathbf{0}, \quad (7)$$

$$P_n[D_\gamma^*\{\bar{Q}(\hat{\gamma}_n^{\text{BR}}), g(\hat{\alpha}_n^{\text{BR}})\}] = \mathbf{0}. \quad (8)$$

Here, the function $D_\alpha^*\{\bar{Q}(\gamma), g(\alpha)\}$ is used as an estimating function for γ and the function $D_\gamma^*\{\bar{Q}(\gamma), g(\alpha)\}$ as an estimating function for α . Vermeulen and Vansteelandt [8] (Theorem 2) show that $\hat{\gamma}_n^{\text{BR}} \xrightarrow{P} \gamma_{\text{BR}}^*$ and $\hat{\alpha}_n^{\text{BR}} \xrightarrow{P} \alpha_{\text{BR}}^*$ with $\gamma_{\text{BR}}^* = \gamma_0$ under $\mathcal{M}(\mathcal{Q}_\gamma)$ and $\alpha_{\text{BR}}^* = \alpha_0$ under $\mathcal{M}(\mathcal{G}_\alpha)$. This can be understood upon noting that the double robustness implies that $P_0[D_\gamma^*\{\bar{Q}(\gamma), g(\alpha_0)\}] = \mathbf{0}$ for all γ under $\mathcal{M}(\mathcal{G}_\alpha)$ and $P_0[D_\alpha^*\{\bar{Q}(\gamma_0), g(\alpha)\}] = \mathbf{0}$ for all α under $\mathcal{M}(\mathcal{Q}_\gamma)$. Clearly, computation of these estimators $\hat{\gamma}_n^{\text{BR}}$ and $\hat{\alpha}_n^{\text{BR}}$ of γ_{BR}^* and α_{BR}^* does not demand knowledge of the full data law. The bias-reduced doubly robust estimator is now given by $\hat{\psi}_n^{\text{BR}} \equiv \hat{\psi}_n(\hat{Q}_n^{\text{BR}}, \hat{g}_n^{\text{BR}})$, with $\hat{Q}_n^{\text{BR}} = \bar{Q}(\hat{\gamma}_n^{\text{BR}})$ and $\hat{g}_n^{\text{BR}} = g(\hat{\alpha}_n^{\text{BR}})$.

3.2.2 Illustration: estimation of a population mean with incomplete data

We illustrate the bias-reduced doubly robust estimation principle for the missing data problem discussed in Section 2. Consider the working models $\mathcal{Q}_\gamma$ and $\mathcal{G}_\alpha$ from Section 3.1 but with $\mathbf{k}$ and $\mathbf{l}$ of the same dimension (i.e., such that $q = p$). Estimators for the nuisance parameters $(\hat{\gamma}_n^{\text{BR}, T}, \hat{\alpha}_n^{\text{BR}, T})^T$ are then obtained by solving (7) and (8):

$$n^{-1} \sum_{i=1}^n R_i \{Y_i - \bar{Q}(\hat{\gamma}_n^{\text{BR}})(\mathbf{X}_i)\} \frac{G'\{\hat{\alpha}_n^{\text{MLE}, T} \mathbf{l}(\mathbf{X}_i)\}}{G^2\{\hat{\alpha}_n^{\text{MLE}, T} \mathbf{l}(\mathbf{X}_i)\}} \mathbf{l}(\mathbf{X}_i) = \mathbf{0}, \quad (9)$$

$$n^{-1} \sum_{i=1}^n \left\{ 1 - \frac{R_i}{g(\hat{\alpha}_n^{\text{BR}})(\mathbf{X}_i)} \right\} Q'\{\hat{\gamma}_n^{\text{BR}, T} \mathbf{k}(\mathbf{X}_i)\} \mathbf{k}(\mathbf{X}_i) = \mathbf{0}. \quad (10)$$

Vermeulen and Vansteelandt [8] show how (10) can be solved by optimizing an integrated estimating equation; (9) can be solved via weighted least squares based on the complete cases. They moreover discuss various desirable properties of the doubly robust estimator based on these nuisance parameter estimators. Here, we develop visual insight using two simple examples.

Example 1. Consider a sample $\mathcal{O}$ of size $n = 10^5$, which we take to be large so that we can ignore sampling variability. In this first example, for each individual i , we let $X_i \stackrel{d}{=} N(0, 1)$, $R_i|X_i \stackrel{d}{=} \text{Ber}\{\pi_0(X_i)\}$ with $\pi_0(X_i) = \text{expit}(-1 + X_i^3)$ and $Y_i|X_i \stackrel{d}{=} N(X_i^2, 1)$. The true mean outcome equals $\psi_0 = 1$. One-dimensional working models of the form $g(\alpha)(X) = \text{expit}(\alpha X)$ and $\bar{Q}(\gamma)(X) = \gamma X$ are misspecified. We obtain $\hat{\alpha}_n^{\text{MLE}} = 1.115$ and $\hat{\gamma}_n^{\text{MLE}} = 1.606$, leading to $\hat{\psi}_n^{\text{DR, MLE}} = 0.237$ and $\text{bias}(\hat{\alpha}_n^{\text{MLE}}, \hat{\gamma}_n^{\text{MLE}}; 1) = -0.763$. The bias-reduced doubly robust estimation strategy yields $\hat{\alpha}_n^{\text{BR}} = 2.254$ and $\hat{\gamma}_n^{\text{BR}} = -0.959$, leading to $\hat{\psi}_n^{\text{BR}} = 0.663$ and a smaller bias of $\text{bias}(\hat{\alpha}_n^{\text{BR}}, \hat{\gamma}_n^{\text{BR}}; 1) = -0.337$, as theoretically expected.

Figure 1 shows a contour plot of the log of the squared asymptotic bias as a function of the nuisance parameters α and γ . It thus indicates for each combination of fixed nuisance parameter values γ and α the magnitude of the asymptotic bias of the doubly robust estimator for these choices of γ and α . Darker values indicate bias closer to zero. The cross “ $\times$ ” indicates the point $(\hat{\alpha}_n^{\text{MLE}}, \hat{\gamma}_n^{\text{MLE}}) = (1.115, 1.606)$, which approximates $(\alpha_{\text{MLE}}^*, \gamma_{\text{MLE}}^*)$ and the bullet “ $\bullet$ ” indicates the point $(\hat{\alpha}_n^{\text{BR}}, \hat{\gamma}_n^{\text{BR}}) = (2.254, -0.959)$, which approximates $(\alpha_{\text{BR}}^*, \gamma_{\text{BR}}^*)$. Figure 1 illustrates the defining property of the bias-reduced estimation principle: as one varies the nuisance parameter values allowed by the chosen working models, the squared asymptotic bias of the doubly robust estimator is locally minimized at $(2.254, -0.959) \approx (\alpha_{\text{BR}}^*, \gamma_{\text{BR}}^*)$ in certain directions of γ and α . Bias-reduction is indeed local, as we observe that there are two other regions where even smaller bias near zero is attained. The location of these regions depends strongly on the underlying data generating mechanism P_0 , so that they cannot be obtained by solving an estimating equation without further knowledge of the true data-generating mechanism. $\square$

Example 2. To illustrate that the location of the aforementioned regions depends on the underlying data generating mechanism, we show a similar plot for a second example. With $n = 10^5$, for each individual i , we let $X_i \stackrel{d}{=} N(1, 1)$, $R_i|X_i \stackrel{d}{=} \text{Ber}\{\pi_0(X_i)\}$ with $\pi_0(X_i) = \text{expit}(-1 + X_i^2)$ and $Y_i|X_i \stackrel{d}{=} N(X_i^2, 1)$. Under this data generating mechanism, the true mean outcome equals $\psi_0 = 2$.

In Figure 2, the cross “ $\times$ ” now indicates the point $(\hat{\alpha}_n^{\text{MLE}}, \hat{\gamma}_n^{\text{MLE}}) = (0.737, 2.157)$, which approximates $(\alpha_{\text{MLE}}^*, \gamma_{\text{MLE}}^*)$, leading to $\hat{\psi}_n^{\text{DR, MLE}} = 2.393$ and $\text{bias}(\hat{\alpha}_n^{\text{MLE}}, \hat{\gamma}_n^{\text{MLE}}; 2) = 0.393$. The bullet “ $\bullet$ ” indicates the point $(\hat{\alpha}_n^{\text{BR}}, \hat{\gamma}_n^{\text{BR}}) = (0.609, 1.220)$, which approximates the probability limits $(\alpha_{\text{BR}}^*, \gamma_{\text{BR}}^*)$, leading to $\hat{\psi}_n^{\text{BR}} = 2.316$ and a smaller bias given by $\text{bias}(\hat{\alpha}_n^{\text{BR}}, \hat{\gamma}_n^{\text{BR}}; 2) = 0.316$. The behavior in this second example (see Figure 2) is similar to the behavior in the first example (see Figure 1). However, do note that the regions where smaller bias

Figure 2: Contour plot of the log of the squared asymptotic bias $\log\{\text{bias}^2(\alpha, \gamma; 2)\}$ as a function of the nuisance parameters α and γ for **example 2** with $\times = (0.737, 2.157) \approx (\alpha_{\text{MLE}}^*, \gamma_{\text{MLE}}^*)$, $\bullet = (0.609, 1.220) \approx (\alpha_{\text{BR}}^*, \gamma_{\text{BR}}^*)$ and $\blacktriangle = (0.737, 0.859) \approx (\alpha_{\text{MLE}}^*, \gamma_{\text{MLE-BR}}^*)$.

near zero is attained are located in a different position, not trackable by solving an estimating equation without further knowledge of the true data-generating mechanism. $\square$

As previously noted, it follows from eq. (2) that small perturbations (of the order one over root- n) in the nuisance parameters γ and α do not affect the first-order asymptotic behavior of the doubly robust estimator when the intersection model $\mathcal{M}(\mathcal{Q}_\gamma) \cap \mathcal{M}(\mathcal{G}_\alpha)$ holds. Besides targeting bias reduction, the bias-reduced estimation procedure extends this property to working model misspecification. This is visually suggested by both Figures 1 and 2 which show that the estimators $\hat{\alpha}_n^{\text{BR}}$ and $\hat{\gamma}_n^{\text{BR}}$ are situated in a region where small perturbations of the nuisance parameters do not heavily affect the doubly robust estimator (in contrast to the regions that deliver near-zero bias, which are very unstable). This is unlike the MLE of the nuisance parameters. More formally, the doubly robust estimator $\hat{\psi}_n^{\text{BR}}$ is first-order ancillary [28] under misspecification of the working models [8, Corollary 1). This means that even under possible model misspecification, $\hat{\psi}_n^{\text{BR}}$ still admits the expansion $\hat{\psi}_n^{\text{BR}} - \psi_0 = (P_0 - P_n)[D^*\{\bar{Q}(\gamma_{\text{BR}}^*), g(\alpha_{\text{BR}}^*); \psi_0\}] + o_p(1/\sqrt{n})$ and thus $\hat{\psi}_n^{\text{BR}}$ and $\hat{\psi}_n\{\bar{Q}(\gamma_{\text{BR}}^*), g(\alpha_{\text{BR}}^*)\}$ are asymptotically equivalent. This moreover implies that the asymptotic variance of the doubly robust estimator $\hat{\psi}_n^{\text{BR}}$ can be straightforwardly estimated as one over n times the sample variance of the values $D^*\{\bar{Q}(\hat{\gamma}_n^{\text{BR}}), g(\hat{\alpha}_n^{\text{BR}}); \hat{\psi}_n^{\text{BR}}\}$, without having to adjust for the estimation of the nuisance parameters.

A limitation of the bias-reduced estimation procedure is that it is confined to parametric working models of the same dimension. Equal dimensions are indeed required because the gradient of D^* with respect to α is used as an estimating function for γ and vice versa. This can be overcome by targeting asymptotic bias reduction in the direction of a single nuisance parameter, for instance in the direction of α . The effect of this is best understood from Figures 1 or 2. For a given value $\tilde{\alpha}$ of α , e.g., the MLE $\hat{\alpha}_n^{\text{MLE}}$, and for each value of γ , we evaluate how the bias of the doubly robust estimator changes as we move away from the chosen value $\tilde{\alpha}$. We then choose the outcome regression parameter γ at which no change in the squared asymptotic bias is seen. For example 1, with MLE $\hat{\alpha}_n^{\text{MLE}} = 1.115$, this leads to $\hat{\gamma}_n^{\text{MLE-BR}} = 0.169$, $\hat{\psi}_n^{\text{MLE-BR}} = 0.518$ and $\text{bias}(\hat{\alpha}_n^{\text{MLE}}, \hat{\gamma}_n^{\text{MLE-BR}}; 1) = -0.482$. On Figure 1, this point $(\hat{\alpha}_n^{\text{MLE}}, \hat{\gamma}_n^{\text{MLE-BR}}) = (1.115, 0.169)$ is marked by means of the triangle “ $\blacktriangle$ ”. This is the point where no change of bias is seen in the direction of α . Note that the bias lies between the bias of the MLE and the bias-reduced doubly robust estimator. A qualitatively different result is seen for example 2. With MLE $\hat{\alpha}_n^{\text{MLE}} = 0.737$, we find $\hat{\gamma}_n^{\text{MLE-BR}} = 0.859$, $\hat{\psi}_n^{\text{MLE-BR}} = 2.302$ and $\text{bias}(\hat{\alpha}_n^{\text{MLE}}, \hat{\gamma}_n^{\text{MLE-BR}}; 2) = 0.302$, which is slightly smaller than the bias of the bias-reduced doubly robust estimator. This is an artefact of the true underlying data-generating mechanism P_0 ; we are not aware of theoretical reasons for this bias to be smaller (see example 1 where it is higher). In the next section, we will explain how we can use this idea of targeting asymptotic bias reduction in one direction to be able to make use of data-adaptive learning algorithms.

4 Extending bias-reduced doubly robust estimation beyond parametric working models

In this section, we will relax the restriction to parametric working models by making use of data-adaptive learning algorithms to estimate the conditional mean outcome, as in TMLE. We will continue to work with a parametric working model for the missingness mechanism because of the possible concern that a too flexible, data-adaptive model specification may result in near-positivity violations and thereby distort the performance of the doubly robust estimator; as noted by a referee, this could alternatively be prevented via cross-validation [29]. With concern for bias under such parametric working model, we will apply the bias-reduction principle of Section 3.2 in the direction of the nuisance parameters indexing that working model. A positive side effect of our proposed procedure will be that it no longer constrains the dimensions of both working models to be the same.

4.1 Main idea

Our proposal is to start with a parametric model for the missingness mechanism $g_0(\mathbf{X})$, indexed by α , and an estimator $\hat{\alpha}_n$ of α , e.g., the MLE. Next, an estimator for the conditional mean outcome $\bar{Q}_0(\mathbf{X})$ is obtained, either by maximum likelihood estimation under a parametric model or by using data-adaptive learning algorithms such as super-learning [10]. With concern for misspecification of the missingness model, we next fluctuate this initial estimator of $\bar{Q}_0(\mathbf{X})$ through a parametric fluctuation model, indexed by a finite-dimensional parameter ε of at least the dimension of α . This model includes covariates that are chosen in such a way that the score of the fluctuation parameter ε equals the gradient $D_{\alpha}^*\{\bar{Q}(\gamma), g(\alpha)\}$. By doing so, we aim to reduce the asymptotic bias of the doubly robust estimator in the direction of α , as explained in Section 3.2. In Section 4.2, we will detail how this can be done. Note that we can thus allow for a missingness model of arbitrary dimension by exploiting the bias-reduced estimation principle in only a single dimension.

The idea of extending an initial fit of the conditional mean outcome is not new: this idea was already proposed in Bang and Robins [6] and it also underlies the TMLE procedure of van der Laan and Rubin [9], as well as other improved doubly robust estimation procedures [20, 21]. Our proposal is nonetheless different in that it explicitly targets bias reduction (in the direction of α). In contrast, the TMLE procedure aims at obtaining a doubly robust substitution estimator and the procedures by Tan [20] and Rotnitzky et al. [21] target a bounded doubly robust estimator with desirable efficiency properties.

4.2 Practical implementation of the procedure

The extension of the bias-reduced doubly robust estimation procedure can be implemented using the following four steps.

Step 1: Estimator $\hat{g}_n^{\text{MLE}}$ for the missingness mechanism $g_0(\mathbf{X})$.

Postulate a parametric working model for $g_0(\mathbf{X})$: $\mathcal{G}_{\alpha} = \{g(\alpha)(\mathbf{X}) | \alpha \in \mathbb{R}^p\}$ where $g(\alpha)(\mathbf{X}) = G\{\alpha^T \mathbf{l}(\mathbf{X})\}$ and G is an appropriate inverse link function and $\mathbf{l} = (1, l_1, \dots, l_{p-1})$, e.g., the logistic regression model $G(\cdot) = \text{expit}(\cdot)$. Obtain the MLE $\hat{\alpha}_n^{\text{MLE}}$ solving

$$P_n \left\{ D_g^{\text{MLE}}(\hat{\alpha}_n^{\text{MLE}}) \right\} = \mathbf{0},$$

with $D_g^{\text{MLE}}(\alpha)$ given by eq. (6). Let $\hat{g}_n^{\text{MLE}} = g(\hat{\alpha}_n^{\text{MLE}})$.

Step 2: Initial estimator $\hat{Q}_n^0$ for the conditional mean outcome $\bar{Q}_0$.

The second step of the procedure is to obtain an initial estimator for $\bar{Q}_0$. We describe two possibilities: (1) a parametric model and (2) a super-learner.

1. *Parametric Working Model.* The first option is to postulate a parametric working model for $\bar{Q}_0(\mathbf{X})$: $Q_\gamma = \{\bar{Q}(\gamma)(\mathbf{X}) | \gamma \in \mathbb{R}^q\}$ with $\bar{Q}(\gamma)(\mathbf{X}) = Q\{\gamma^T \mathbf{k}(\mathbf{X})\}$, Q an appropriate inverse link function and $\mathbf{k} = (1, k_1, \dots, k_{q-1})$; e.g., a linear regression model with Q the identity link. Obtain the MLE $\hat{\gamma}_n^{\text{MLE}}$ solving

$$P_n \left\{ D_Q^{\text{MLE}}(\hat{\gamma}_n^{\text{MLE}}) \right\} = \mathbf{0},$$

with $D_Q^{\text{MLE}}(\gamma)$ given by eq. (5). Let $\hat{Q}_n^{\text{MLE}} = \bar{Q}(\hat{\gamma}_n^{\text{MLE}})$.

2. *Super-Learner.* Another option is to obtain an initial estimator based on data-adaptive learning algorithms, such as super-learning [10]. The super-learner is a machine learning algorithm which starts from a library of estimators $\{\hat{Q}_j | j = 1, \dots, J\}$, which may consist of both parametric and nonparametric estimators. It then considers the family of all weighted averages of these estimators: $\hat{Q}_\omega = \sum_{j=1}^J \omega_j \hat{Q}_j$, $\omega = (\omega_1, \dots, \omega_J)^T$, $\sum_{j=1}^J \omega_j = 1$ and $\omega_j \geq 0$ for $j = 1, \dots, J$. Next, the optimal weight vector $\hat{\omega}_n$ is defined to be the choice of ω that minimizes the cross-validated risk with respect to some loss function $\mathcal{L}_{\text{SL}}: (\mathbf{O}, \bar{Q}) \rightarrow \mathcal{L}_{\text{SL}}(\bar{Q})(\mathbf{O}) \in \mathbb{R}$ which satisfies $\bar{Q}_0 = \arg \min_{\bar{Q}} E\{\mathcal{L}_{\text{SL}}(\bar{Q})\}$, e.g., the squared error loss function $\mathcal{L}_2(\bar{Q})(\mathbf{O}) = R\{Y - \bar{Q}(\mathbf{X})\}^2$. The super-learner of $\bar{Q}_0$ is defined as the estimator $\hat{Q}_n^{\text{SL}} = \hat{Q}_{\hat{\omega}_n}$.

For further reference, we let the initial estimator $\hat{Q}_n^0$ denote either $\hat{Q}_n^{\text{MLE}}$ or $\hat{Q}_n^{\text{SL}}$.

Step 3: Fluctuation $\hat{Q}_n^{(c)}$ of the initial estimator $\hat{Q}_n^0$.

To construct an appropriate fluctuation model, we need to choose an appropriate loss-function. In this article, we consider the quasi-log-likelihood loss function with corresponding logistic fluctuation model. This choice is favored over the squared error loss function with linear fluctuation model [30] because it ensures predictions within the admissible range for the outcome. In particular, suppose Y is known to fall in the interval $[a, b]$ with $a < b$. To be able to use the quasi-log-likelihood loss functions and the logistic fluctuation model, we rescale the outcome between zero and one as follows:

$$\tilde{Y} = \frac{Y - a}{b - a}. \quad (11)$$

The procedure we describe below is based on the transformed outcome $\tilde{Y}$ and results in an estimator $\tilde{\psi}_n$ of $\tilde{\psi}_n = E(\tilde{Y})$. Because $E(Y) = (b - a)E(\tilde{Y}) + a$, the final estimator is given by $\hat{\psi}_n = (b - a)\tilde{\psi}_n + a$. For notational convenience, without loss of generality, we will assume that $a = 0$ and $b = 1$ so that $Y = \tilde{Y} \in [0, 1]$ and we can drop the $\sim$ -notation.

Having thus obtained an estimator for the missingness mechanism, such as $\hat{g}_n^{\text{MLE}}$, and an initial estimator $\hat{Q}_n^0$ for the conditional mean outcome taking values in the interval $[0, 1]$, we can now fluctuate $\hat{Q}_n^0$ to target bias reduction in the direction of the finite-dimensional parameter α . Below, we consider three fluctuation models that will accomplish this goal with the second also guaranteeing the final estimator to be a substitution estimator, like the TMLE procedure.

1. *Fluctuation model 1.* Consider the fluctuation model $\{\hat{Q}_n^0(\boldsymbol{\epsilon}^{(1)}) : \boldsymbol{\epsilon}^{(1)} \in \mathbb{R}^p\}$ through the initial estimator $(\hat{Q}_n^0(\mathbf{0}) = \hat{Q}_n^0)$:

$$\text{logit } \hat{Q}_n^0(\boldsymbol{\epsilon}^{(1)})(\mathbf{X}) = \text{logit } \hat{Q}_n^0(\mathbf{X}) + \boldsymbol{\epsilon}^{(1), T} \mathbf{H}^{(1)}(\hat{g}_n^{\text{MLE}})(\mathbf{X}), \quad (12)$$

with $\mathbf{H}^{(1)}(\hat{g}_n^{\text{MLE}})(\mathbf{X}) = [G'\{\hat{\alpha}_n^{\text{MLE}, T} \mathbf{l}(\mathbf{X})\} / G^2\{\hat{\alpha}_n^{\text{MLE}, T} \mathbf{l}(\mathbf{X})\}] \mathbf{l}(\mathbf{X})$ and with $\text{logit}(x) = \log\{x/(1-x)\}$. Then define $\hat{\boldsymbol{\epsilon}}_n^{(1)}$ as the MLE under this parametric submodel, i.e., such that

$$\hat{\boldsymbol{\epsilon}}_n^{(1)} = \arg \min_{\boldsymbol{\epsilon}^{(1)}} P_n \left[\mathcal{L}^{(1)} \left\{ \hat{Q}_n^0(\boldsymbol{\epsilon}^{(1)}) \right\} \right]$$

where $\mathcal{L}^{(1)}$ is the quasi-log-likelihood loss function;

$$\mathcal{L}^{(1)}(\bar{Q})(\mathbf{O}) = -R[Y \log\{\bar{Q}(\mathbf{X})\} + (1 - Y) \log\{1 - \bar{Q}(\mathbf{X})\}]. \quad (13)$$

It is easily verified that $\hat{\boldsymbol{\epsilon}}_n^{(1)}$ solves the estimating equation

$$\sum_{i=1}^n R_i \left\{ Y_i - \hat{Q}_n^0(\hat{\boldsymbol{\epsilon}}_n^{(1)})(\mathbf{X}_i) \right\} \mathbf{H}^{(1)}(\hat{\mathbf{g}}_n^{\text{MLE}})(\mathbf{X}_i) = \mathbf{0}, \quad (14)$$

which can be solved via standard logistic regression of the observed outcomes on the covariates $\mathbf{H}^{(1)}$ using offset logit $\hat{Q}_n^0$. Because the score eq. (14) is like the estimation eq. (9), it targets bias-reduction (in the direction of $\boldsymbol{\alpha}$) in the sense that was previously described. Define the updated estimator $\hat{Q}_n^{(1)} = \hat{Q}_n^0(\hat{\boldsymbol{\epsilon}}_n^{(1)})$. Because the quasi-log-likelihood loss function is a valid loss-function for the conditional mean outcome $\bar{Q}_0$ in the sense that $\bar{Q}_0 = \arg \min_Q \mathcal{L}^{(1)}(\bar{Q})$ (see Gruber and van der Laan [30], Lemma 1), $\hat{Q}_n^{(1)}$ is an improved fit of $\bar{Q}_0$ as compared to $\hat{Q}_n^0$ with respect to the loss function (13).

2. *Fluctuation model 2.* We extend eq. (12) so that the final doubly robust estimator of the target parameter will also be a substitution estimator (also known as a regression doubly robust estimator [12]). For this purpose, define the fluctuation model $\left\{ \hat{Q}_n^0(\boldsymbol{\epsilon}^{(2)}) : \boldsymbol{\epsilon}^{(2)} \in \mathbb{R}^{p+1} \right\}$ through the initial estimator:

$$\begin{aligned} \text{logit } \hat{Q}_n^0(\boldsymbol{\epsilon}^{(2)})(\mathbf{X}) \\ = \text{logit } \hat{Q}_n^0(\mathbf{X}) + \boldsymbol{\epsilon}^{(2,1),T} \mathbf{H}^{(1)}(\hat{\mathbf{g}}_n^{\text{MLE}})(\mathbf{X}) + \boldsymbol{\epsilon}^{(2,2)} \mathbf{H}^{(2)}(\hat{\mathbf{g}}_n^{\text{MLE}})(\mathbf{X}), \end{aligned} \quad (15)$$

$\boldsymbol{\epsilon}^{(2)} = (\boldsymbol{\epsilon}^{(2,1),T}, \boldsymbol{\epsilon}^{(2,2),T})^T$ and $\mathbf{H}^{(2)}(\hat{\mathbf{g}}_n^{\text{MLE}}) = 1/\hat{\mathbf{g}}_n^{\text{MLE}}$. Then define $\hat{\boldsymbol{\epsilon}}_n^{(2)}$ as the MLE under this parametric submodel, i.e., such that

$$\hat{\boldsymbol{\epsilon}}_n^{(2)} = \arg \min_{\boldsymbol{\epsilon}^{(2)}} P_n \left[\mathcal{L}^{(1)} \left\{ \hat{Q}_n^0(\boldsymbol{\epsilon}^{(2)}) \right\} \right]$$

where $\mathcal{L}^{(1)}$ is the quasi-log-likelihood loss function (13). It follows that $\hat{\boldsymbol{\epsilon}}_n^{(2)}$ solves the estimating equation

$$\sum_{i=1}^n R_i \left\{ Y_i - \hat{Q}_n^0(\hat{\boldsymbol{\epsilon}}_n^{(2)})(\mathbf{X}_i) \right\} \times \left(\mathbf{H}^{(1),T}(\hat{\mathbf{g}}_n^{\text{MLE}})(\mathbf{X}_i), \mathbf{H}^{(2)}(\hat{\mathbf{g}}_n^{\text{MLE}})(\mathbf{X}_i) \right)^T = \mathbf{0}, \quad (16)$$

which can be easily solved via standard logistic regression of the observed outcomes on the covariates $\mathbf{H}^{(1)}$ and $\mathbf{H}^{(2)}$ using offset logit $\hat{Q}_n^0$. By adding $\mathbf{H}^{(2)}$ to the fluctuation model, it follows from the second component of eq. (16) that the resulting doubly robust estimator is a substitution estimator (as in the TMLE procedure) because eq. (16) implies that $\sum_{i=1}^n R_i / \hat{\mathbf{g}}_n^{\text{MLE}} \left\{ Y_i - \hat{Q}_n^0(\hat{\boldsymbol{\epsilon}}_n^{(2)})(\mathbf{X}_i) \right\} = 0$ so that $\hat{\psi}_n \{ \hat{Q}_n^0(\hat{\boldsymbol{\epsilon}}_n^{(2)}), \hat{\mathbf{g}}_n^{\text{MLE}} \} = P_n \{ \hat{Q}_n^0(\hat{\boldsymbol{\epsilon}}_n^{(2)}) \}$. Define the updated estimator $\hat{Q}_n^{(2)} = \hat{Q}_n^0(\hat{\boldsymbol{\epsilon}}_n^{(2)})$.

3. *Fluctuation model 3.* There is evidence in the literature that the use of a weighted loss function in the construction of a fluctuation model can improve the stability of the doubly robust estimator of interest; see for example Robins et al. [12], Díaz and Rosenblum [31] for a comparison of an unweighted versus a weighted loss function in the context of TMLE. We therefore propose the following alternative fluctuation model to eq. (12): $\left\{ \hat{Q}_n^0(\boldsymbol{\epsilon}^{(3)}) : \boldsymbol{\epsilon}^{(3)} \in \mathbb{R}^p \right\}$ through the initial estimator:

$$\text{logit } \hat{Q}_n^0(\boldsymbol{\epsilon}^{(3)})(\mathbf{X}) = \text{logit } \hat{Q}_n^0(\mathbf{X}) + \boldsymbol{\epsilon}^{(3),T} \mathbf{I}(\mathbf{X}). \quad (17)$$

Then define $\hat{\boldsymbol{\epsilon}}_n^{(3)}$ as the weighted MLE under this parametric submodel

$$\hat{\boldsymbol{\epsilon}}_n^{(3)} = \arg \min_{\boldsymbol{\epsilon}^{(3)}} P_n \left[\frac{G' \{ \hat{\boldsymbol{\alpha}}_n^{\text{MLE},T} \mathbf{I}(\mathbf{X}) \}}{G^2 \{ \hat{\boldsymbol{\alpha}}_n^{\text{MLE},T} \mathbf{I}(\mathbf{X}) \}} \mathcal{L}^{(1)} \left\{ \hat{Q}_n^0(\boldsymbol{\epsilon}^{(3)}) \right\} \right],$$

with $\mathcal{L}^{(1)}$ the quasi-log-likelihood loss function (13); thus $\hat{\boldsymbol{\epsilon}}_n^{(3)}$ solves the estimating equations

$$0 = \sum_{i=1}^n R_i \left\{ Y_i - \hat{Q}_n^0(\hat{\epsilon}_n^{(2)})(\mathbf{X}_i) \right\} \frac{G' \{ \hat{\alpha}_n^{\text{MLE}, T} \mathbf{l}(\mathbf{X}_i) \}}{G^2 \{ \hat{\alpha}_n^{\text{MLE}, T} \mathbf{l}(\mathbf{X}_i) \}} \mathbf{l}(\mathbf{X}_i), \quad (18)$$

which can be easily solved via standard weighted logistic regression of the observed outcomes on the covariates $\mathbf{l}$ using offset $\logit \hat{Q}_n^0$. As before this targets bias reduction because eq. (18) is like the estimating eq. (9). Define the updated estimator $\hat{Q}_n^{(3)} = \hat{Q}_n^0(\hat{\epsilon}_n^{(3)})$.

We thus obtained three different updated estimators $\hat{Q}_n^{(c)}$, $c=1, 2, 3$. These all aim to reduce bias in the direction of α . As noted, this is because each of the scores for the fluctuation parameters $\epsilon^{(c)}$ ($c=1, 2, 3$) ensure that

$$P_n \left\{ D_{\alpha}^* (\hat{Q}_n^{(c)}, \hat{g}_n^{\text{MLE}}) \right\} = \mathbf{0}.$$

Step 4: Estimating the Target Parameter $\hat{\psi}_n^{(c)}$.

Given the estimators $\hat{g}_n^{\text{MLE}}$ and $\hat{Q}_n^{(c)}$, we obtain the doubly robust estimator $\hat{\psi}_n^{(c)} \equiv \hat{\psi}_n(\hat{Q}_n^{(c)}, \hat{g}_n^{\text{MLE}})$. This is a one-step procedure and does not require iterations. Finally note that by construction (implied by the estimating eq. (16) as a consequence of adding $\mathbf{H}^{(2)}$ to the fluctuation model), $\hat{\psi}_n^{(2)} = \hat{\psi}_n(\hat{Q}_n^{(2)}, \hat{g}_n^{\text{MLE}}) = P_n(\hat{Q}_n^{(2)})$.

Remark: If in the construction of the missingness model, we take $G(x) = \text{expit}(x)$, then $G'(x) = G(x)\{1 - G(x)\}$. For every $c=1, 2, 3$, it then follows from the estimating equations for the fluctuation parameter $\epsilon^{(c)}$ that

$$\sum_{i=1}^n \frac{R_i}{\hat{g}_n^{\text{MLE}}(\mathbf{X}_i)} \left\{ Y_i - \hat{Q}_n^{(c)}(\mathbf{X}_i) \right\} = \sum_{i=1}^n R_i \left\{ Y_i - \hat{Q}_n^{(c)}(\mathbf{X}_i) \right\}. \quad (19)$$

This implies that the doubly robust estimator can be written as

$$\hat{\psi}_n^{(c)} = n^{-1} \sum_{i=1}^n \left\{ R_i Y_i + (1 - R_i) \hat{Q}_n^{(c)}(\mathbf{X}_i) \right\}, \quad (20)$$

which averages the observed outcome for the responders and a predicted outcome for the non-responders, making the doubly robust estimator equal a special type of imputation estimator. This is desirable as this ensures that the doubly robust estimator is sample bounded in the sense that $\hat{\psi}_n^{(c)}$ lies in the observed data range [12, 20].

4.3 Inference

In Appendix A, we present an asymptotic linearity theorem with corresponding influence function for the doubly robust estimators $\hat{\psi}_n^{(c)}$ ($c=1, 2, 3$) under the assumption of a correctly specified working model for the missingness mechanism but a potentially misspecified working model for the conditional mean outcome, where the latter is estimated via either a parametric model or a data-adaptive learning algorithm. This asymptotic linearity of $\hat{\psi}_n^{(c)}$ and the corresponding influence function

$$D^*(\bar{Q}_{(c)}^*, g_0; \psi_0),$$

with $\bar{Q}_{(c)}^*$ the probability limit of $\hat{Q}_n^{(c)}$, provides us with a strategy for inference about the unknown ψ_0 under a correctly specified missingness model. When the outcome model is correctly specified so that $\bar{Q}_{(c)}^* = \bar{Q}_0$, the influence function equals the efficient influence function $D^*(\bar{Q}_0, g_0; \psi_0)$. When the outcome model is misspecified so that $\bar{Q}_{(c)}^* \neq \bar{Q}_0$, the influence function of the doubly robust estimator equals $D^*(\bar{Q}_{(c)}^*, g_0; \psi_0)$, which may differ from the efficient influence function (indeed, efficiency of $\hat{\psi}_n^{(c)}$ is only attained locally at $\bar{Q}_{(c)}^* = \bar{Q}_0$). However, inferences based on this influence function are valid, even though they ignore the

uncertainty in the working model for the missingness mechanism. This is because bias-reduced estimation is used in the direction of α implying that

$$P_0 \left\{ D_{\alpha}^*(\bar{Q}_{(c)}^*, g_0) \right\} = \mathbf{0}.$$

This first-order ancillarity property with respect to α ensures that the resulting doubly robust estimator is insensitive to local changes (of the order one over root- n) in the nuisance parameter α , even under misspecification of the outcome model. Consequently, no correction for the estimation of the missingness mechanism is needed under inconsistency of $\hat{Q}_n^{(c)}$ [8].

Remark: Remember that, when one of the nuisance parameter estimators is not consistent, then the asymptotic distribution of the double-robust estimator changes because of the addition of first-order terms to its expansion. These first-order terms reflect the extent to which the double-robust estimator is affected by the uncertainty in the nuisance parameter estimator(s). A recent proposal by van der Laan [32] targets the estimation of the nuisance parameters so as to make these additional first-order terms asymptotically linear, so that asymptotic 95% confidence intervals can be constructed; in doing so, he assumes that the missingness model is correctly specified (to the extent that it delivers a consistent double-robust estimator of the target parameter in collaboration with the possibly inconsistent estimator of the outcome mean). Our proposal is different in that it makes these additional first-order terms zero, does not assume correct specification of the missingness model regarding the bias reduction property, but also does not allow data-adaptive estimation of the missingness model.

4.3.1 Standard errors

When the missingness model is correctly specified, a standard error for $\hat{\psi}_n^{(c)}$ can be easily calculated as the square root of the sample variance of the estimated influence function divided by n :

$$\widehat{\text{SE}}(\hat{\psi}_n^{(c)}) = \sqrt{\frac{\widehat{\text{Var}}\left\{D^*(\hat{Q}_n^{(c)}, \hat{g}_n^{\text{MLE}}; \hat{\psi}_n^{(c)})\right\}}{n}} \quad (21)$$

with

$$\begin{aligned} \widehat{\text{var}}\left\{D^*(\hat{Q}_n^{(c)}, \hat{g}_n^{\text{MLE}}; \hat{\psi}_n^{(c)})\right\} &= \frac{1}{n} \sum_{i=1}^n \left\{D^*(\hat{Q}_n^{(c)}, \hat{g}_n^{\text{MLE}}; \hat{\psi}_n^{(c)})(\mathbf{O}_i)\right\}^2 \\ &= \frac{1}{n} \sum_{i=1}^n \left[\frac{R_i}{\hat{g}_n(\mathbf{X}_i)} \left\{Y_i - \hat{Q}_n^{(c)}(\mathbf{X}_i)\right\} + \hat{Q}_n^{(c)}(\mathbf{X}_i) - \hat{\psi}_n^{(c)} \right]^2. \end{aligned}$$

An alternative for standard error calculation, which does not demand the assumption that the missingness model is correctly specified, is to use the non-parametric bootstrap. However, there is no theory supporting that the non-parametric bootstrap would produce valid results when the estimators rely on data-adaptive estimation such as super-learning [32].

4.3.2 Confidence intervals and p -values

Given the estimator $\widehat{\text{SE}}(\hat{\psi}_n^{(c)})$ for the standard error of $\hat{\psi}_n^{(c)}$, a confidence interval and p -value can be calculated based on the asymptotic normality of the estimator. A $(1 - \alpha)100\%$ CI is given by

$$\left[\hat{\psi}_n^{(c)} \pm z_{\alpha/2} \widehat{\text{SE}}(\hat{\psi}_n^{(c)}) \right]$$

where $z_{\alpha/2}$ is such that $\Phi(z_{\alpha/2}) = 1 - \alpha/2$ and $\Phi(\cdot)$ the cumulative distribution function of a standard normal random variable. A p -value for the hypothesis test $H_0: \psi_0 = \tilde{\psi}$ versus $H_1: \psi_0 \neq \tilde{\psi}$ can be calculated as

$$p = 2 \left\{ 1 - \Phi \left(\left| \frac{\hat{\psi}_n^{(c)} - \tilde{\psi}}{\widehat{\text{SE}}(\hat{\psi}_n^{(c)})} \right| \right) \right\}.$$

4.3.3 Inference on the original scale

With $\tilde{Y} = (Y - a)/(b - a)$ the transformed data which falls within the interval $[0, 1]$, previously, we let $\hat{\psi}_n^{(c)}$ denote the final estimator for $\tilde{\psi}_0 = E(\tilde{Y})$. Because $E(Y) = (b - a)E(\tilde{Y}) + a$, the final estimator for $\psi_0 = E(Y)$ can be obtained as $\hat{\psi}_n^{(c)} = (b - a)\hat{\psi}_n^{(c)} + a$. When $\sqrt{n}(\hat{\psi}_n^{(c)} - \tilde{\psi}_0)$ converges in distribution to $N(0, \tilde{\sigma}^2)$, then $\sqrt{n}(\hat{\psi}_n^{(c)} - \psi_0)$ will converge in distribution to $N(0, \sigma^2)$ with $\sigma^2 = (b - a)^2 \tilde{\sigma}^2$. Consequently, a standard error for $\hat{\psi}_n^{(c)}$ on the original scale of the outcome, can be obtained as $\widehat{\text{SE}}(\hat{\psi}_n^{(c)}) = (b - a)\widehat{\text{SE}}(\hat{\psi}_n^{(c)})$.

In Appendix B, we provide an R-function that implements these novel methods.

5 Simulation studies

We carried out different simulation studies to compare the performance of the new bias-reduced doubly robust estimators with several alternatives for the estimation of a mean outcome in the presence of incomplete data.

5.1 Estimators

All estimators are based on a parametric working model $\mathcal{G}_\alpha$ for the missingness mechanism $g_0(\mathbf{X})$. Let $\hat{\alpha}_n^{\text{MLE}}$ denote the MLE and let $\hat{g}_n^{\text{MLE}} = g(\hat{\alpha}_n^{\text{MLE}})$. For the conditional mean outcome $\bar{Q}_0(\mathbf{X})$, we consider both estimators based on a parametric working model and based on data-adaptive learning algorithms. First, consider a parametric model $\mathcal{Q}_\gamma$ with $\hat{\gamma}_n^{\text{MLE}}$ the MLE and $\hat{Q}_n^{\text{MLE}} = \bar{Q}(\hat{\gamma}_n^{\text{MLE}})$. Given the parametric working models $\mathcal{G}_\alpha$ and $\mathcal{Q}_\gamma$, let $\hat{\alpha}_n^{\text{BR}}$ and $\hat{\gamma}_n^{\text{BR}}$ denote the nuisance parameter estimators for α and γ obtained via the bias-reduced estimation principle with $\hat{g}_n^{\text{BR}} = g(\hat{\alpha}_n^{\text{BR}})$ and $\hat{Q}_n^{\text{BR}} = \bar{Q}(\hat{\gamma}_n^{\text{BR}})$. Second, we consider the super-learner $\hat{Q}_n^{\text{SL}}$ based on a library consisting of generalized additive and linear models, random forests and adaptive polynomial splines. This can be fitted using the SuperLearner R package [33]. Let $\hat{Q}_n^{\text{MLE}, (c)} = \hat{Q}_n^{\text{MLE}}(\hat{\epsilon}_n^{(c)})$ denote the updated estimators for the conditional mean outcome based on the initial estimator $\hat{Q}_n^{\text{MLE}}$, $c = 1, 2, 3$ and let $\hat{Q}_n^{\text{SL}, (c)} = \hat{Q}_n^{\text{SL}}(\hat{\epsilon}_n^{(c)})$ denote the updated estimators for the conditional mean outcome based on the initial estimator $\hat{Q}_n^{\text{SL}}$, $c = 1, 2, 3$. The estimators under consideration are then the doubly robust estimator based on the MLE for the parametric working models $\hat{\psi}_n^{\text{DR, MLE}} = \hat{\psi}_n(\hat{Q}_n^{\text{MLE}}, \hat{g}_n^{\text{MLE}})$, the bias-reduced doubly robust estimator $\hat{\psi}_n^{\text{BR}} = \hat{\psi}_n(\hat{Q}_n^{\text{BR}}, \hat{g}_n^{\text{BR}})$, the substitution estimator $\hat{\psi}_n^{\text{SL}} = P_n(\hat{Q}_n^{\text{SL}})$ based on standardizing super-learning predictions (which is not doubly robust), the proposed estimators based on both a parametric working model for the missingness mechanism and a parametric conditional mean model for the outcome, $\hat{\psi}_n^{\text{DR, MLE}, (c)} = \hat{\psi}_n(\hat{Q}_n^{\text{MLE}, (c)}, \hat{g}_n^{\text{MLE}})$, and based on a parametric model for the missingness model but a super-learner for the conditional mean outcome, $\hat{\psi}_n^{\text{SL}, (c)} = \hat{\psi}_n(\hat{Q}_n^{\text{SL}, (c)}, \hat{g}_n^{\text{MLE}})$, for $c = 1, 2, 3$. Finally, we consider two versions of TMLE, described in van der Laan and Rubin [9] and Gruber and van der Laan [30]. Both are based on the parametric missingness model $\mathcal{G}_\alpha$ fitted via MLE but the first is based on $\hat{Q}_n^{\text{MLE}}$ and the second is based on $\hat{Q}_n^{\text{SL}}$, leading to the TMLEs $\hat{\psi}_n^{\text{TMLE}}$ and $\hat{\psi}_n^{\text{SL, TMLE}}$, respectively.

Many alternative estimation strategies for these nuisance working models have been considered in the literature (see Section 1.1 for an overview). For a simulation-based comparison of many of these alternatives, we refer to Tan [20], Porter et al. [24] and Vermeulen and Vansteelandt [8].

Remark: In the two simulation settings of the subsequent sections, the outcome Y is not a priori bounded. For those estimators involving a logistic fluctuation model for the initial estimator of the conditional mean outcome, we follow the default setting in the `tmle` R package [34] by taking the bounds to be the observed range of Y but widened by 10% of both the minimum and maximum value. This differs slightly from the approach in Gruber and van der Laan [30]. Specifically, we let $a = \min_{i=1}^n Y_i - 0.1|\min_{i=1}^n Y_i|$ and $b = \max_{i=1}^n Y_i + 0.1|\max_{i=1}^n Y_i|$ and $\tilde{Y} = (Y - a)/(b - a)$. Next, since the initial estimator needs to be represented as a logistic function of its logit transformation, the initial estimator of the conditional mean outcome $\bar{Q}_0(\mathbf{X})$ needs to be bounded away from 0 and 1 because $\text{logit}(x)$ is not defined for $x = 0$ or 1. Therefore, this initial estimator is truncated at $(\zeta, 1 - \zeta)$ for some small $\zeta > 0$. In the simulations, we take $\zeta = 0.005$ as in Gruber and van der Laan [30].

For each of the scenarios considered below, we perform 1,000 Monte Carlo runs at sample sizes of $n = 200$ and 1,000. For each estimator, we calculated the Monte Carlo bias (BIAS), the root mean square error (RMSE) and the Monte Carlo standard deviation (MCSd). For the doubly robust estimators (all but $\hat{\psi}_n^{\text{SL}}$), we also show the average sandwich standard error (ASSE) calculated via eq. (21), and the Monte Carlo coverage of the corresponding 95% Wald confidence intervals (COV) where standard errors are calculated via formula (21).

5.2 Scenario 1: single-covariate setting

5.2.1 Data-generating mechanism

This simulation scenario considers a simple data-generating mechanism where for each i ($i = 1, \dots, n$), $X_i \stackrel{d}{=} N(0, 1)$, $R_i|X_i \stackrel{d}{=} \text{Ber}\{\pi_0(X_i)\}$ and $Y_i|X_i \stackrel{d}{=} N\{m_0(X_i), 1\}$. For each setting, the following parametric working models are used: $\pi(X; \alpha) = \text{expit}(\alpha_1 + \alpha_2 X)$ and $m(X; \gamma) = \gamma_1 + \gamma_2 X$. Simulation experiments with correctly specified parametric working models used $m_0(X) = 1 + X$ and $\pi_0(X) = \text{expit}(\xi X)$ for $\xi = 1, 2$. To allow for misspecification in the outcome model, we additionally generated data using $m_0(X) = X^2$ and $\pi_0(X) = \text{expit}(\xi X)$ for $\xi = 1, 2$. To allow for gross misspecification of the missingness mechanism model, we generated data using $m_0(X) = 1 + X$ and $\pi_0(X) = \text{expit}(-4 + 1.5|X|^{0.5} + 0.75X + 0.5|X|^{1.5})$, as in Vansteelandt et al. [7]. Finally, we also generated data with $m_0(X) = X^2$ and $\pi_0(X) = \text{expit}(-4 + 1.5|X|^{0.5} + 0.75X + 0.5|X|^{1.5})$ to allow for misspecification of both models. In each of the settings, the target parameter $E(Y) = \psi_0$ equals one. Results for the Scenario 1 are given in Table 1 ($n = 200$) and Table 2 ($n = 1000$).

5.2.2 Results

When both working models are correctly specified and weights are not highly variable ($\xi = 1$), all estimators tend to perform similarly in terms of bias and precision. However, when weights become highly variable ($\xi = 2$), estimators $\hat{\psi}_n^{\text{DR, MLE}, (1)}$, $\hat{\psi}_n^{\text{DR, MLE}, (2)}$, $\hat{\psi}_n^{\text{SL}, (1)}$ and $\hat{\psi}_n^{\text{SL}, (2)}$ tend to show some finite-sample bias ($n = 200$), which is resolved when the sample size is increased ($n = 1000$). With highly variable weights, these estimators are also relatively less efficient. When the outcome model is misspecified but the working model for the missingness mechanism is correct, we observe adequate performance for all estimators, but smaller bias and larger precision for the estimators that are based on the super-learner. Furthermore, it appears desirable to use a weighted loss function (as $\hat{\psi}_n^{\text{DR, MLE}, (3)}$ and $\hat{\psi}_n^{\text{SL}, (3)}$ do) as compared to fluctuation models using covariates with the estimated missingness probabilities in the denominator (as for $\hat{\psi}_n^{\text{DR, MLE}, (1)}$, $\hat{\psi}_n^{\text{DR, MLE}, (2)}$, $\hat{\psi}_n^{\text{SL}, (1)}$ and $\hat{\psi}_n^{\text{SL}, (2)}$). In the case where the missingness model is wrong but the outcome model is

Table 1: Simulation results based on 1,000 Monte Carlo replications for Scenario 1, $n = 200$.

EST	BIAS	RMSE	MCS	ASSE	COV	BIAS	RMSE	MCS	ASSE	COV
$n = 200$										
OR correct, MM correct ($\xi = 1$)						OR incorrect, MM correct ($\xi = 1$)				
$\hat{\psi}_n^{\text{DR, MLE}}$	-0.001	0.13	0.13	0.13	0.94	-0.018	0.33	0.33	0.29	0.89
$\hat{\psi}_n^{\text{BR}}$	-0.001	0.13	0.13	0.13	0.94	-0.030	0.21	0.20	0.17	0.88
$\hat{\psi}_n^{\text{SL}}$	0.007	0.13	0.13	—	—	-0.085	0.19	0.17	—	—
$\hat{\psi}_n^{\text{TMLE}}$	-0.000	0.13	0.13	0.13	0.94	-0.036	0.27	0.26	0.24	0.89
$\hat{\psi}_n^{\text{SL, TMLE}}$	0.001	0.13	0.13	0.13	0.94	-0.029	0.17	0.17	0.14	0.89
$\hat{\psi}_n^{\text{DR, MLE, (1)}}$	0.005	0.14	0.14	0.13	0.93	0.028	0.19	0.19	0.16	0.92
$\hat{\psi}_n^{\text{DR, MLE, (2)}}$	0.010	0.14	0.14	0.13	0.93	0.037	0.21	0.21	0.18	0.91
$\hat{\psi}_n^{\text{DR, MLE, (3)}}$	-0.000	0.13	0.13	0.13	0.94	-0.055	0.22	0.22	0.18	0.85
$\hat{\psi}_n^{\text{SL, (1)}}$	0.004	0.14	0.14	0.13	0.93	-0.000	0.16	0.16	0.14	0.92
$\hat{\psi}_n^{\text{SL, (2)}}$	0.010	0.14	0.14	0.13	0.92	0.011	0.17	0.17	0.14	0.92
$\hat{\psi}_n^{\text{SL, (3)}}$	0.001	0.13	0.13	0.13	0.94	-0.028	0.17	0.17	0.14	0.89
OR correct, MM correct ($\xi = 2$)						OR incorrect, MM correct ($\xi = 2$)				
$\hat{\psi}_n^{\text{DR, MLE}}$	-0.002	0.20	0.20	0.17	0.94	-0.060	1.19	1.19	0.58	0.68
$\hat{\psi}_n^{\text{BR}}$	-0.001	0.19	0.19	0.15	0.88	-0.110	0.26	0.24	0.17	0.77
$\hat{\psi}_n^{\text{SL}}$	0.018	0.15	0.15	—	—	-0.249	0.33	0.22	—	—
$\hat{\psi}_n^{\text{TMLE}}$	0.017	0.19	0.19	0.15	0.86	0.013	0.33	0.33	0.24	0.82
$\hat{\psi}_n^{\text{SL, TMLE}}$	0.018	0.19	0.19	0.14	0.85	-0.020	0.23	0.23	0.16	0.82
$\hat{\psi}_n^{\text{DR, MLE, (1)}}$	0.097	0.30	0.28	0.14	0.70	0.304	0.53	0.43	0.22	0.65
$\hat{\psi}_n^{\text{DR, MLE, (2)}}$	0.110	0.31	0.29	0.14	0.67	0.282	0.51	0.43	0.21	0.64
$\hat{\psi}_n^{\text{DR, MLE, (3)}}$	0.003	0.17	0.17	0.14	0.88	-0.187	0.38	0.33	0.19	0.61
$\hat{\psi}_n^{\text{SL, (1)}}$	0.090	0.28	0.27	0.14	0.70	0.128	0.34	0.32	0.17	0.72
$\hat{\psi}_n^{\text{SL, (2)}}$	0.111	0.31	0.29	0.14	0.66	0.100	0.35	0.34	0.17	0.69
$\hat{\psi}_n^{\text{SL, (3)}}$	0.010	0.18	0.18	0.14	0.86	-0.093	0.25	0.23	0.15	0.75
OR correct, MM incorrect						OR incorrect, MM incorrect				
$\hat{\psi}_n^{\text{DR, MLE}}$	-0.008	1.04	1.04	0.76	0.96	5.496	11.46	10.06	4.78	0.90
$\hat{\psi}_n^{\text{BR}}$	-0.003	0.29	0.29	0.22	0.85	1.030	1.23	0.68	0.43	0.34
$\hat{\psi}_n^{\text{SL}}$	0.030	0.29	0.29	—	—	0.155	0.42	0.39	—	—
$\hat{\psi}_n^{\text{TMLE}}$	0.010	0.30	0.30	0.49	0.94	1.199	1.34	0.60	0.94	0.79
$\hat{\psi}_n^{\text{SL, TMLE}}$	-0.002	0.29	0.30	0.46	0.94	0.312	0.50	0.39	0.55	0.91
$\hat{\psi}_n^{\text{DR, MLE, (1)}}$	0.029	0.31	0.31	0.40	0.91	0.411	0.65	0.50	0.70	0.85
$\hat{\psi}_n^{\text{DR, MLE, (2)}}$	0.035	0.33	0.33	0.38	0.89	0.277	0.58	0.51	0.77	0.93
$\hat{\psi}_n^{\text{DR, MLE, (3)}}$	0.014	0.30	0.30	0.49	0.95	0.595	0.77	0.49	0.86	0.87
$\hat{\psi}_n^{\text{SL, (1)}}$	0.032	0.32	0.32	0.38	0.90	0.131	0.41	0.39	0.40	0.87
$\hat{\psi}_n^{\text{SL, (2)}}$	0.037	0.33	0.33	0.37	0.88	0.102	0.41	0.40	0.40	0.87
$\hat{\psi}_n^{\text{SL, (3)}}$	0.001	0.30	0.30	0.46	0.94	0.155	0.39	0.36	0.53	0.92

Note: EST, estimator; BIAS, Monte Carlo Bias; RMSE, root mean square error; MAE, median of absolute errors; MCS, Monte Carlo standard deviation; ASSE, average of sandwich standard errors; COV, Monte Carlo coverage of 95% Wald confidence intervals; OR, outcome regression; MM, missingness model. No convergence for $\hat{\alpha}_n^{\text{BR}}$ was attained in five of the 1,000 runs for the settings OR correct, MM correct ($\xi = 2$) and OR incorrect, MM correct ($\xi = 2$) and in three of the 1,000 runs for the settings OR correct, MM incorrect and OR incorrect, MM incorrect. These cases were deleted in the calculation of the results.

Table 2: Simulation results based on 1,000 Monte Carlo replications for Scenario 1, $n = 1000$.

EST	BIAS	RMSE	MCSD	ASSE	COV	BIAS	RMSE	MCSD	ASSE	COV
$n = 1000$										
OR correct, MM correct ($\xi = 1$)						OR incorrect, MM correct ($\xi = 1$)				
$\hat{\psi}_n^{\text{DR, MLE}}$	0.003	0.06	0.06	0.06	0.96	0.003	0.15	0.15	0.16	0.95
$\hat{\psi}_n^{\text{BR}}$	0.003	0.06	0.06	0.06	0.96	-0.003	0.09	0.09	0.09	0.94
$\hat{\psi}_n^{\text{SL}}$	0.005	0.06	0.06	—	—	-0.024	0.07	0.07	—	—
$\hat{\psi}_n^{\text{TMLE}}$	0.003	0.06	0.06	0.06	0.96	-0.004	0.12	0.12	0.13	0.95
$\hat{\psi}_n^{\text{SL, TMLE}}$	0.003	0.06	0.06	0.06	0.96	-0.003	0.07	0.07	0.07	0.93
$\hat{\psi}_n^{\text{DR, MLE, (1)}}$	0.004	0.06	0.06	0.06	0.95	0.013	0.09	0.09	0.08	0.93
$\hat{\psi}_n^{\text{DR, MLE, (2)}}$	0.005	0.06	0.06	0.06	0.95	0.014	0.10	0.10	0.09	0.93
$\hat{\psi}_n^{\text{DR, MLE, (3)}}$	0.003	0.06	0.06	0.06	0.96	-0.011	0.10	0.10	0.09	0.92
$\hat{\psi}_n^{\text{SL, (1)}}$	0.004	0.06	0.06	0.06	0.95	0.000	0.07	0.07	0.07	0.94
$\hat{\psi}_n^{\text{SL, (2)}}$	0.005	0.06	0.06	0.06	0.94	0.001	0.07	0.07	0.07	0.94
$\hat{\psi}_n^{\text{SL, (3)}}$	0.003	0.06	0.06	0.06	0.95	-0.003	0.07	0.07	0.07	0.93
OR correct, MM correct ($\xi = 2$)						OR incorrect, MM correct ($\xi = 2$)				
$\hat{\psi}_n^{\text{DR, MLE}}$	-0.000	0.09	0.09	0.08	0.94	-0.041	0.50	0.49	0.35	0.78
$\hat{\psi}_n^{\text{BR}}$	0.000	0.08	0.08	0.08	0.91	-0.052	0.14	0.13	0.10	0.80
$\hat{\psi}_n^{\text{SL}}$	0.006	0.07	0.07	—	—	-0.113	0.15	0.10	—	—
$\hat{\psi}_n^{\text{TMLE}}$	0.004	0.09	0.09	0.08	0.90	0.044	0.19	0.19	0.13	0.85
$\hat{\psi}_n^{\text{SL, TMLE}}$	0.004	0.09	0.09	0.07	0.90	-0.012	0.10	0.10	0.08	0.87
$\hat{\psi}_n^{\text{DR, MLE, (1)}}$	0.027	0.12	0.12	0.07	0.78	0.168	0.30	0.24	0.13	0.66
$\hat{\psi}_n^{\text{DR, MLE, (2)}}$	0.030	0.13	0.12	0.07	0.78	0.152	0.27	0.23	0.11	0.64
$\hat{\psi}_n^{\text{DR, MLE, (3)}}$	0.001	0.08	0.08	0.08	0.91	-0.107	0.22	0.19	0.13	0.64
$\hat{\psi}_n^{\text{SL, (1)}}$	0.025	0.12	0.12	0.07	0.78	0.038	0.13	0.12	0.08	0.79
$\hat{\psi}_n^{\text{SL, (2)}}$	0.029	0.13	0.12	0.07	0.77	0.027	0.14	0.13	0.08	0.76
$\hat{\psi}_n^{\text{SL, (3)}}$	0.001	0.08	0.08	0.07	0.91	-0.038	0.11	0.10	0.08	0.82
OR correct, MM incorrect						OR incorrect, MM incorrect				
$\hat{\psi}_n^{\text{DR, MLE}}$	-0.026	0.52	0.52	0.49	0.99	7.239	9.64	6.37	4.10	0.39
$\hat{\psi}_n^{\text{BR}}$	-0.006	0.11	0.11	0.11	0.94	1.237	1.28	0.33	0.28	0.01
$\hat{\psi}_n^{\text{SL}}$	-0.000	0.11	0.12	—	—	0.059	0.16	0.15	—	—
$\hat{\psi}_n^{\text{TMLE}}$	-0.004	0.12	0.12	0.41	1.00	1.222	1.26	0.29	0.66	0.37
$\hat{\psi}_n^{\text{SL, TMLE}}$	-0.005	0.12	0.12	0.39	1.00	0.100	0.18	0.15	0.42	1.00
$\hat{\psi}_n^{\text{DR, MLE, (1)}}$	-0.001	0.12	0.12	0.37	1.00	0.601	0.66	0.26	0.83	0.97
$\hat{\psi}_n^{\text{DR, MLE, (2)}}$	-0.001	0.12	0.12	0.36	1.00	0.573	0.63	0.25	0.80	1.00
$\hat{\psi}_n^{\text{DR, MLE, (3)}}$	0.003	0.16	0.16	0.42	1.00	0.289	0.46	0.36	0.85	0.97
$\hat{\psi}_n^{\text{SL, (1)}}$	0.000	0.13	0.13	0.36	1.00	0.033	0.14	0.14	0.36	1.00
$\hat{\psi}_n^{\text{SL, (2)}}$	0.001	0.12	0.12	0.36	1.00	0.035	0.14	0.14	0.35	1.00
$\hat{\psi}_n^{\text{SL, (3)}}$	0.003	0.15	0.15	0.40	1.00	0.003	0.16	0.16	0.44	1.00

Note: BIAS, Monte Carlo Bias; RMSE, root mean square error; MAE, median of absolute errors; MCSD, Monte Carlo standard deviation; ASSE, average of sandwich standard errors; COV, Monte Carlo coverage of 95% Wald confidence intervals; OR, outcome regression; MM, missingness model.

correct, all estimators show very similar behavior, except for the standard doubly robust estimator $\hat{\psi}_n^{\text{DR, MLE}}$ which performs poorly. Finally, when both working models are misspecified, then $\hat{\psi}_n^{\text{BR}}$ drastically outperforms $\hat{\psi}_n^{\text{DR, MLE}}$ in terms of bias and precision. The proposed estimator $\hat{\psi}_n^{\text{SL}, (3)}$ performs best.

In conclusion, $\hat{\psi}_n^{\text{SL}, (3)}$ performs best in Scenario 1, both in terms of bias and precision, especially when the inverse probability weights become highly variable. Adding the additional covariate $\mathbf{H}^{(2)}$ to the fluctuation model for $\hat{\psi}_n^{\text{DR, MLE}, (2)}$ and $\hat{\psi}_n^{\text{SL}, (2)}$ guarantees a substitution estimator, but does not lead to enhanced performance as compared to $\hat{\psi}_n^{\text{DR, MLE}, (1)}$ and $\hat{\psi}_n^{\text{SL}, (1)}$. This can partly be understood upon noting that $\hat{\psi}_n^{\text{DR, MLE}, (1)}$ and $\hat{\psi}_n^{\text{SL}, (1)}$ already have the form of an imputation estimator (see eq. (20)).

Tables 1 and 2 also show results on the performance of the sandwich standard error calculated using formula (21). The proposed estimator of the sandwich standard error performs well under correct specification of both working models, especially for $\hat{\psi}_n^{\text{SL}, (3)}$. As predicted by the theory, this is also the case when the outcome model is misspecified but the missingness mechanism is correct, but the weights are not highly variable ($\xi = 1$). When the weights become highly variable ($\xi = 2$), the performance is worse. This is not a surprise because convergence to the normal limit distribution then happens more slowly. When the missingness model is misspecified (for both a correctly specified and misspecified outcome model), then the sandwich standard errors overestimate the finite-sample variability of the estimator for all different estimators (except for $\hat{\psi}_n^{\text{BR}}$).

5.3 Scenario 2: Kang and Schafer setting

5.3.1 Data-generating mechanism

This simulation study is taken from Kang and Schafer [5] and often used as a benchmark to evaluate doubly robust estimators for the population mean outcome explainable by measured auxiliary covariates. For each individual $i = 1, \dots, n$, $\mathbf{Z}_i = (Z_{i1}, Z_{i2}, Z_{i3}, Z_{i4})^T \stackrel{d}{=} \mathbf{N}(\mathbf{0}, \mathbf{I})$, $\mathbf{I}$ is the 4×4 identity matrix, $R_i | \mathbf{Z}_i \stackrel{d}{=} \text{Ber}\{\pi_0(\mathbf{Z}_i)\}$ with $\pi_0(\mathbf{Z}) = \text{expit}(-Z_1 + 0.5Z_2 - 0.25Z_3 - 0.1Z_4)$ and $Y_i | \mathbf{Z}_i \stackrel{d}{=} N\{m_0(\mathbf{Z}_i), 1\}$ with $m_0(\mathbf{Z}) = 210 + 27.4Z_1 + 13.7Z_2 + 13.7Z_3 + 13.7Z_4$. Misspecified working models are linear for the outcome model and logistic for the missingness mechanism model, with covariates $\mathbf{X}_i = (X_{i1}, X_{i2}, X_{i3}, X_{i4})^T$ with $X_1 = \exp(Z_1/2)$, $X_2 = Z_2 / \{1 + \exp(Z_1)\} + 10$, $X_3 = (Z_1 Z_3 / 25 + 0.6)^3$ and $X_4 = (Z_2 + Z_4 + 20)^2$. The target parameter $E(Y) = \psi_0$ equals 210. We limit ourselves to the realistic settings where the working models both use either the covariates Z_k or the covariates X_k ($k = 1, \dots, 4$) and thus both working models are correctly specified or both working models are incorrectly specified. We will show simulation results for two scenarios where either $R = 1$ or $R = 0$ denotes the data that are observed and results are shown in Table 3 ($n = 200$) and Table 4 ($n = 1000$).

5.3.2 Results

When both working models are correctly specified, all estimators have comparable performance, especially for $n = 1000$. However, $\hat{\psi}_n^{\text{DR, MLE}, (1)}$, $\hat{\psi}_n^{\text{DR, MLE}, (2)}$, $\hat{\psi}_n^{\text{SL}, (1)}$ and $\hat{\psi}_n^{\text{SL}, (2)}$ tend to be slightly more variable than $\hat{\psi}_n^{\text{DR, MLE}, (3)}$ and $\hat{\psi}_n^{\text{SL}, (3)}$ at $n = 200$, again illustrating the enhanced performance of using a weighted loss function for estimators $\hat{\psi}_n^{\text{DR, MLE}, (3)}$ and $\hat{\psi}_n^{\text{SL}, (3)}$. When both working models are misspecified, as in Kang and Schafer [5], $\hat{\psi}_n^{\text{DR, MLE}}$ shows severe erratic behavior when the observed outcome RY is used; this behavior is partially eliminated when $(1 - R)Y$ is used as the observed outcome [12]. The bias-reduced doubly robust estimator from Section 3, $\hat{\psi}_n^{\text{BR}}$, does not show this severe erratic behavior (for both observed outcomes RY and $(1 - R)Y$). Interestingly, the new estimator $\hat{\psi}_n^{\text{SL}, (3)}$ outperforms all existing estimators $\hat{\psi}_n^{\text{DR, MLE}}$, $\hat{\psi}_n^{\text{BR}}$, $\hat{\psi}_n^{\text{TMLE}}$ and $\hat{\psi}_n^{\text{SL, TMLE}}$ for both observed outcomes RY and $(1 - R)Y$. When $(1 - R)Y$ is used as the observed outcome, $\hat{\psi}_n^{\text{DR, MLE}, (1)}$ and $\hat{\psi}_n^{\text{SL}, (1)}$ show the best performance under double model misspecification.

Tables 3 and 4 also show the performance of the sandwich standard errors (21). Excellent finite-sample behavior is seen when both working models are correctly specified. When both working models are

Table 3: Simulation results based on 1,000 Monte Carlo replications for Scenario 2, $n = 200$.

EST	BIAS	RMSE	MCSD	ASSE	COV	BIAS	RMSE	MCSD	ASSE	COV
observed outcome RY						observed outcome $(1-R)Y$				
$n = 200$										
OR correct, MM correct						OR correct, MM correct				
$\hat{\psi}_n^{\text{DR, MLE}}$	0.099	2.53	2.53	2.57	0.96	0.085	2.53	2.53	2.57	0.96
$\hat{\psi}_n^{\text{BR}}$	0.090	2.54	2.54	2.57	0.96	0.095	2.54	2.55	2.57	0.95
$\hat{\psi}_n^{\text{SL}}$	0.013	2.52	2.52	–	–	0.254	2.55	2.54	–	–
$\hat{\psi}_n^{\text{TMLE}}$	0.028	2.53	2.53	2.56	0.96	0.243	2.55	2.54	2.54	0.95
$\hat{\psi}_n^{\text{SL, TMLE}}$	0.028	2.53	2.53	2.56	0.96	0.243	2.55	2.54	2.54	0.95
$\hat{\psi}_n^{\text{DR, MLE}, (1)}$	–0.112	2.73	2.73	2.56	0.94	0.326	2.72	2.70	2.55	0.94
$\hat{\psi}_n^{\text{DR, MLE}, (2)}$	–0.153	2.80	2.80	2.57	0.94	0.348	2.77	2.75	2.56	0.94
$\hat{\psi}_n^{\text{DR, MLE}, (3)}$	0.029	2.53	2.53	2.56	0.96	0.241	2.55	2.54	2.54	0.95
$\hat{\psi}_n^{\text{SL}, (1)}$	–0.114	2.73	2.72	2.56	0.94	0.324	2.71	2.69	2.55	0.94
$\hat{\psi}_n^{\text{SL}, (2)}$	–0.186	2.80	2.80	2.57	0.94	0.365	2.77	2.75	2.56	0.94
$\hat{\psi}_n^{\text{SL}, (3)}$	0.025	2.53	2.53	2.56	0.96	0.244	2.55	2.54	2.54	0.95
OR incorrect, MM incorrect						OR incorrect, MM incorrect				
$\hat{\psi}_n^{\text{DR, MLE}}$	–15.150	88.60	87.34	15.92	0.93	4.759	6.05	3.74	3.31	0.60
$\hat{\psi}_n^{\text{BR}}$	–2.239	4.45	3.85	2.95	0.82	3.440	4.63	3.10	2.73	0.73
$\hat{\psi}_n^{\text{SL}}$	–1.912	3.68	3.15	–	–	4.742	5.83	3.40	–	–
$\hat{\psi}_n^{\text{TMLE}}$	–6.840	8.44	4.94	4.01	0.57	4.195	5.40	3.40	2.95	0.65
$\hat{\psi}_n^{\text{SL, TMLE}}$	–3.221	4.81	3.57	2.89	0.75	3.706	5.05	3.43	2.68	0.66
$\hat{\psi}_n^{\text{DR, MLE}, (1)}$	–5.962	7.83	5.09	3.05	0.50	1.447	4.02	3.75	2.90	0.84
$\hat{\psi}_n^{\text{DR, MLE}, (2)}$	–6.478	8.41	5.37	3.07	0.46	1.614	4.23	3.91	2.91	0.83
$\hat{\psi}_n^{\text{DR, MLE}, (3)}$	–3.263	5.23	4.09	3.28	0.80	3.656	4.77	3.07	2.66	0.70
$\hat{\psi}_n^{\text{SL}, (1)}$	–3.722	6.08	4.81	2.88	0.65	2.157	4.16	3.56	2.76	0.80
$\hat{\psi}_n^{\text{SL}, (2)}$	–3.944	6.41	5.05	2.89	0.64	2.293	4.38	3.73	2.77	0.79
$\hat{\psi}_n^{\text{SL}, (3)}$	–1.970	3.61	3.02	2.67	0.84	3.350	4.61	3.17	2.57	0.71

Note: EST, estimator; BIAS, Monte Carlo Bias; RMSE, root mean square error; MCSD, Monte Carlo standard deviation; ASSE, average of sandwich standard errors; COV, Monte Carlo coverage of 95% Wald confidence intervals; OR, outcome regression; MM, missingness model. No convergence for $\hat{\alpha}_n^{\text{BR}}$ was attained in 13 of the 1,000 runs for the settings OR correct, MM correct, $n = 200$ for both the observed outcome RY and $(1-R)Y$ and in five of the 1,000 runs for the setting OR incorrect, MM incorrect, $n = 200$ for the observed outcome RY . These cases were deleted in the calculation of the results.

misspecified, the approximation of the sandwich standard errors is better for $n = 1000$ than for $n = 200$, but not sufficient, as is also the case for TMLE. Unlike in Scenario 1, we do observe that they underestimate the true finite-sample variability of the estimators. Overall, the best performance is seen for $\hat{\psi}_n^{\text{SL}, (3)}$.

6 Linear instrumental variable analyses

In this section, we show that the proposed procedure is not restricted to the doubly robust estimator from Section 2 but can be easily extended to other doubly robust estimators. We will do this by illustrating how the principle can be implemented for linear instrumental variable analyses.

Table 4: Simulation results based on 1,000 Monte Carlo replications for Scenario 2, $n = 1000$.

EST	BIAS	RMSE	MCSD	ASSE	COV	BIAS	RMSE	MCSD	ASSE	COV
observed outcome RY						observed outcome $(1 - R)Y$				
$n = 1000$										
OR correct, MM correct						OR correct, MM correct				
$\hat{\psi}_n^{\text{DR, MLE}}$	0.023	1.12	1.12	1.15	0.96	0.022	1.13	1.13	1.15	0.96
$\hat{\psi}_n^{\text{BR}}$	0.022	1.12	1.12	1.15	0.96	0.024	1.13	1.13	1.15	0.96
$\hat{\psi}_n^{\text{SL}}$	-0.006	1.13	1.13	—	—	0.070	1.13	1.12	—	—
$\hat{\psi}_n^{\text{TMLE}}$	0.013	1.12	1.12	1.15	0.96	0.062	1.13	1.13	1.14	0.96
$\hat{\psi}_n^{\text{SL, TMLE}}$	0.011	1.12	1.12	1.15	0.96	0.063	1.13	1.13	1.14	0.96
$\hat{\psi}_n^{\text{DR, MLE}, (1)}$	0.008	1.12	1.12	1.15	0.96	0.070	1.14	1.13	1.14	0.95
$\hat{\psi}_n^{\text{DR, MLE}, (2)}$	0.005	1.12	1.12	1.15	0.96	0.072	1.14	1.14	1.14	0.95
$\hat{\psi}_n^{\text{DR, MLE}, (3)}$	0.013	1.12	1.12	1.15	0.96	0.063	1.13	1.13	1.14	0.95
$\hat{\psi}_n^{\text{SL}, (1)}$	0.001	1.12	1.12	1.15	0.96	0.072	1.13	1.14	1.14	0.95
$\hat{\psi}_n^{\text{SL}, (2)}$	-0.001	1.12	1.12	1.15	0.96	0.078	1.14	1.14	1.14	0.95
$\hat{\psi}_n^{\text{SL}, (3)}$	0.011	1.12	1.12	1.15	0.96	0.064	1.13	1.13	1.14	0.95
OR incorrect, MM incorrect						OR incorrect, MM incorrect				
$\hat{\psi}_n^{\text{DR, MLE}}$	-53.715	469.04	466.19	49.69	0.78	4.509	4.83	1.73	1.60	0.19
$\hat{\psi}_n^{\text{BR}}$	-3.208	3.63	1.69	1.34	0.38	2.970	3.25	1.34	1.25	0.35
$\hat{\psi}_n^{\text{SL}}$	-2.879	3.18	1.52	—	—	3.356	3.74	1.65	—	—
$\hat{\psi}_n^{\text{TMLE}}$	-6.242	6.81	2.73	3.72	0.39	4.075	4.35	1.52	1.40	0.19
$\hat{\psi}_n^{\text{SL, TMLE}}$	-2.281	2.75	1.53	1.53	0.62	2.833	3.20	1.48	1.15	0.35
$\hat{\psi}_n^{\text{DR, MLE}, (1)}$	-3.878	4.55	2.39	1.98	0.51	1.442	2.35	1.86	1.35	0.73
$\hat{\psi}_n^{\text{DR, MLE}, (2)}$	-3.738	4.47	2.45	1.92	0.52	1.570	2.46	1.90	1.34	0.70
$\hat{\psi}_n^{\text{DR, MLE}, (3)}$	-3.351	4.86	1.92	2.82	0.74	3.110	3.39	1.34	1.23	0.30
$\hat{\psi}_n^{\text{SL}, (1)}$	-2.771	3.38	1.94	1.28	0.46	1.919	2.38	1.42	1.17	0.60
$\hat{\psi}_n^{\text{SL}, (2)}$	-3.001	3.60	1.99	1.28	0.41	1.981	2.46	1.46	1.17	0.58
$\hat{\psi}_n^{\text{SL}, (3)}$	-2.188	2.52	1.24	1.21	0.55	2.571	2.91	1.37	1.13	0.41

Note: EST, estimator; BIAS, Monte Carlo Bias; RMSE, root mean square error; MAE, median of absolute errors; MCSD, Monte Carlo standard deviation; ASSE, average of sandwich standard errors; COV, Monte Carlo coverage of 95% Wald confidence intervals; OR, outcome regression; MM, missingness model.

6.1 Doubly robust estimation of linear instrumental variable models

Suppose that we are interested in the causal effect of an exposure A on an outcome Y in the presence of unmeasured confounding U and suppose in particular that for measured covariates $\mathbf{X}$, the vector $(\mathbf{X}^T, U^T)^T$ would be sufficient to control for confounding of the effect of A on Y . Suppose that data is available on an instrumental variable (IV) Z [35–37], which is such that (a) Z is associated with A conditional on $\mathbf{X}$, (b) $Z \perp\!\!\!\perp Y|A, \mathbf{X}, U$ and (c) $Z \perp\!\!\!\perp U|\mathbf{X}$. The observed data is given by the i.i.d. sample $\mathcal{O} = (\mathbf{O}_1, \dots, \mathbf{O}_n)$ of size n with P_0 the unknown underlying data-generating mechanism of the observables $\mathbf{O} = (Y, A, Z, \mathbf{X})$ where we consider both A and Z to be dichotomous. We consider inference for the causal effect ψ_0 indexing the linear instrumental variable model

$$E(Y|A, Z, \mathbf{X}, \mathbf{U}) = E(Y|A=0, Z, \mathbf{X}, \mathbf{U}) + \psi_0 A. \quad (22)$$

Let $\mathcal{M}$ denote the statistical model for P_0 implied by assumptions (a)-(b)-(c) and (22). Okui et al. [38] show that a doubly robust estimator of ψ_0 can be obtained by solving

$$0 = \sum_{i=1}^n \{Z_i - \hat{g}_n(\mathbf{X}_i)\} \{Y_i - \psi A_i - \hat{Q}_n(\mathbf{X}_i)\}. \quad (23)$$

Here, $\hat{g}_n(\mathbf{X})$ is an estimator of the conditional mean of the instrument given the measured confounders $\mathbf{X}$, $g_0(\mathbf{X}) = E(Z|\mathbf{X})$, based on a nuisance working model $\mathcal{G} = \{g(\mathbf{X})|g \text{ in some class of functions}\}$. Let $g^*(\mathbf{X})$ denote the probability limit ($\hat{g}_n(\mathbf{X}) \xrightarrow{P} g^*(\mathbf{X})$) such that $g^*(\mathbf{X}) = g_0(\mathbf{X})$ when $\mathcal{M}(\mathcal{G})$ holds, where $\mathcal{M}(\mathcal{G})$ represents the statistical model for the joint distribution of $\mathbf{O}$ implied by the working model $\mathcal{G}$. Next, $\hat{Q}_n(\mathbf{X})$ is an estimator of the conditional mean of the outcome given the measured confounders $\mathbf{X}$ among the non-exposed, $\bar{Q}_0(\mathbf{X}) = E(Y|A=0, \mathbf{X})$, based on a nuisance working model $\mathcal{Q} = \{\bar{Q}(\mathbf{X})|\bar{Q} \text{ in some class of functions}\}$. Let $\bar{Q}^*(\mathbf{X})$ denote the probability limit ($\hat{Q}_n(\mathbf{X}) \xrightarrow{P} \bar{Q}^*(\mathbf{X})$) such that $\bar{Q}^*(\mathbf{X}) = \bar{Q}_0(\mathbf{X})$ when $\mathcal{M}(\mathcal{Q})$ holds, where $\mathcal{M}(\mathcal{Q})$ represents the statistical model for the joint distribution of $\mathbf{O}$ implied by the working model $\mathcal{Q}$. Consistency of the doubly robust estimator $\hat{\psi}_n(\hat{Q}_n, \hat{g}_n)$ is then attained under the model $\mathcal{M} \cap \{\mathcal{M}(\mathcal{Q}) \cup \mathcal{M}(\mathcal{G})\}$. At the intersection model $\mathcal{M} \cap \mathcal{M}(\mathcal{Q}) \cap \mathcal{M}(\mathcal{G})$, it has the following expansion

$$\hat{\psi}_n(\hat{Q}_n, \hat{g}_n) - \psi_0 = (P_n - P_0)\{D_{IV}^*(\bar{Q}_0, g_0; \psi_0)\} + o_p(1/\sqrt{n}), \quad (24)$$

with influence function $D_{IV}^*(\bar{Q}_0, g_0; \psi_0)$ given by

$$D_{IV}^*(\bar{Q}_0, g_0; \psi_0)(\mathbf{O}) = \frac{\{Z - g_0(\mathbf{X})\} \{Y - \bar{Q}_0(\mathbf{X})\}}{E[A\{Z - g_0(\mathbf{X})\}]} - \psi_0. \quad (25)$$

Note that by construction,

$$P_n \left[D_{IV}^* \left\{ \hat{Q}_n, \hat{g}_n; \hat{\psi}_n(\hat{Q}_n, \hat{g}_n) \right\} \right] = 0.$$

6.2 Practical implementation of the proposed procedure

The proposed data-adaptive bias-reduced doubly robust procedure works as follows:

Step 1: Estimator $\hat{g}_n^{\text{MLE}}$ for the conditional mean $g_0(\mathbf{X})$ of the instrument.

Postulate a parametric working model for $g_0(\mathbf{X})$: $\mathcal{G}_\alpha = \{g(\alpha)(\mathbf{X})|\alpha \in \mathbb{R}^p\}$ where $g(\alpha)(\mathbf{X}) = G\{\alpha^T \mathbf{l}(\mathbf{X})\}$ and G an appropriate inverse link function, e.g., $G(\cdot) = \text{expit}(\cdot)$, and $\mathbf{l} = (1, l_1, \dots, l_{p-1})$. Obtain the MLE $\hat{\alpha}_n^{\text{MLE}}$ using the estimating function $D_g^{\text{MLE}}(\alpha)(\mathbf{O})$ given in eq. (6) but with R replaced by Z . Let $\hat{g}_n^{\text{MLE}} = g(\hat{\alpha}_n^{\text{MLE}})$.

Step 2: Initial estimator $\hat{Q}_n^0$ for the conditional mean outcome among the non-exposed $\bar{Q}_0$.

We again consider two options: (1) Postulate a parametric working model for $\bar{Q}_0(\mathbf{X})$: $\mathcal{Q}_\gamma = \{\bar{Q}(\gamma)(\mathbf{X})|\gamma \in \mathbb{R}^q\}$ where $\bar{Q}(\gamma)(\mathbf{X}) = Q\{\gamma^T \mathbf{k}(\mathbf{X})\}$ and Q an appropriate inverse link function and $\mathbf{k} = (1, k_1, \dots, k_{q-1})$. Obtain the MLE $\hat{\gamma}_n^{\text{MLE}}$ using the estimating function $D_{\bar{Q}}^{\text{MLE}}(\gamma)(\mathbf{O})$ given in eq. (5) but with R replaced by $1-A$. Let $\hat{Q}_n^{\text{MLE}} = \bar{Q}(\hat{\gamma}_n^{\text{MLE}})$; (2) Use a super-learner based on a library $\{\hat{Q}_j|j=1, \dots, J\}$ to obtain a fit $\hat{Q}_n^{\text{SL}}$ among those with $A=0$. Let the initial estimator $\hat{Q}_n^0$ denote either $\hat{Q}_n^{\text{MLE}}$ or $\hat{Q}_n^{\text{SL}}$.

Step 3: Fluctuation $\hat{Q}_n^1$ of the initial estimator $\hat{Q}_n^0$.

Following the same argument as in Step 3 of Section 4.2, we assume that $Y \in [0, 1]$. In this step, we fluctuate the initial estimator $\hat{Q}_n^0$ by means of a parametric fluctuation model $\hat{Q}_n^0(\epsilon)$ such that the score with respect to

$\boldsymbol{\varepsilon}$ targets bias reduction in the direction $\boldsymbol{\alpha}$. For this purpose, define the fluctuation model $\{\hat{Q}_n^0(\boldsymbol{\varepsilon}) : \boldsymbol{\varepsilon} \in \mathbb{R}^p\}$ through the initial estimator ($\hat{Q}_n^0(\mathbf{0}) = \hat{Q}_n^0$):

$$\text{logit } \hat{Q}_n^0(\boldsymbol{\varepsilon})(\mathbf{X}) = \text{logit } \hat{Q}_n^0(\mathbf{X}) + \boldsymbol{\varepsilon}^T \hat{\mathbf{H}}(\hat{g}_n^{\text{MLE}})(Z, \mathbf{X}), \quad (26)$$

with

$$\begin{aligned} \hat{\mathbf{H}}(\hat{g}_n^{\text{MLE}})(Z, \mathbf{X}) &= G' \{ \hat{\boldsymbol{\alpha}}_n^{\text{MLE}, T} \mathbf{I}(\mathbf{X}) \} \mathbf{I}(\mathbf{X}) + [Z - G \{ \hat{\boldsymbol{\alpha}}_n^{\text{MLE}, T} \mathbf{I}(\mathbf{X}) \}] \hat{\mathbf{W}}(\hat{g}_n^{\text{MLE}}), \\ \hat{\mathbf{W}}(\hat{g}_n^{\text{MLE}}) &= \frac{n^{-1} \sum_{j=1}^n A_j G' \{ \hat{\boldsymbol{\alpha}}_n^{\text{MLE}, T} \mathbf{I}(\mathbf{X}_j) \} \mathbf{I}(\mathbf{X}_j)}{n^{-1} \sum_{j=1}^n A_j [Z_j - G \{ \hat{\boldsymbol{\alpha}}_n^{\text{MLE}, T} \mathbf{I}(\mathbf{X}_j) \}]}, \end{aligned}$$

where $\hat{\mathbf{H}}(\hat{g}_n^{\text{MLE}})(Z, \mathbf{X})$ is chosen in a way so that the score for the fluctuation parameter $\boldsymbol{\varepsilon}$ under this fluctuation model is given by the derivative of the influence function $D_{\text{IV}}^* \{ \hat{Q}_n^0(\boldsymbol{\varepsilon}), g(\boldsymbol{\alpha}) \}$ to the parameter $\boldsymbol{\alpha}$ and thus targets bias reduction in the direction of $\boldsymbol{\alpha}$. Define $\hat{\boldsymbol{\varepsilon}}_n = \arg \min_{\boldsymbol{\varepsilon}} P_n \left[\mathcal{L}^{(1)} \{ \hat{Q}_n^0(\boldsymbol{\varepsilon}) \} \right]$, the MLE under this parametric submodel with $\mathcal{L}^{(1)}$ the quasi-log-likelihood function (13). This can be obtained via standard logistic regression of the outcome on the covariates $\hat{\mathbf{H}}$ using offset $\text{logit } \hat{Q}_n^0$ and then defining the updated estimates as $\hat{Q}_n^1 = \hat{Q}_n^0(\hat{\boldsymbol{\varepsilon}}_n)$. As mentioned before, by construction of the fluctuation model, it follows that $\hat{\boldsymbol{\varepsilon}}_n$ solves

$$P_n \left\{ D_{\boldsymbol{\alpha}, \text{IV}}^* (\hat{Q}_n^1, \hat{g}_n^{\text{MLE}}) \right\} = \mathbf{0},$$

with $D_{\boldsymbol{\alpha}, \text{IV}}^* \{ \bar{Q}, g(\boldsymbol{\alpha}) \} = \partial D_{\text{IV}}^* \{ \bar{Q}, g(\boldsymbol{\alpha}) \} / \partial \boldsymbol{\alpha}$, thereby targeting bias reduction in the direction of $\boldsymbol{\alpha}$.

Step 4: Estimating the target parameter $\hat{\psi}_n$.

Given the estimators $\hat{g}_n^{\text{MLE}}$ and $\hat{Q}_n^1$, we obtain the doubly robust estimator $\hat{\psi}_n \equiv \hat{\psi}_n(\hat{Q}_n^1, \hat{g}_n^{\text{MLE}})$.

7 Discussion

In this article, we have proposed an extension to the bias-reduced doubly robust estimation principle of Vermeulen and Vansteelandt [8]. In particular, we relaxed the restriction to parametric nuisance working models by making use of data-adaptive learning algorithms to estimate the conditional mean outcome. We carry on to work with a parametric working model for the missingness mechanism for the inferential problem of Section 2, partly in order to prevent potential positivity violations. However, because the concern of bias is greater for such parametric models, we applied the bias-reduction principle of Section 3.2 in the direction of the nuisance parameters indexing these parametric working models. We furthermore illustrated that the proposed extension is not restricted to doubly robust estimation of the mean outcome susceptible to missingness explainable by measured covariates but also extends to other doubly robust procedures as illustrated in Section 6.

This new procedure follows the spirit of the TMLE procedure of van der Laan and Rubin [9] in the sense it also extends an initial data-adaptive estimator of the relevant part $\bar{Q}_0(\mathbf{X})$ of the conditional outcome distribution, enhancing the performance of the estimator of the target parameter. Fluctuation of initial parametric estimators is also seen in other contexts, e.g., Bang and Robins [6], Tan [20], Rotnitzky et al. [21].

In the implementation of Section 4.2, we could additionally have considered a fluctuation $\hat{g}_n^{\text{MLE}}(\boldsymbol{\varepsilon}_g)$ of the initial estimator $\hat{g}_n^{\text{MLE}}$ for the missingness mechanism so as to target additional bias reduction in the direction of the fluctuation parameters $\boldsymbol{\varepsilon}$ used in the construction of the fluctuation model $\hat{Q}_n^0(\boldsymbol{\varepsilon})$. In preliminary simulation studies, we observed this to be unstable and sometimes non-convergent.

A limitation of our proposal is that the current asymptotic linearity theorem, presented in Appendix A, only guarantees valid inference under a correct working model for the missingness mechanism (but a possibly misspecified finite- or infinite-dimensional outcome working model). This is because misspecification of the working model for the missingness mechanism (but correct specification of the outcome model) would demand acknowledging the uncertainty of the estimator for the outcome working model so as to make the remainder term R_n in the expansion (2) of second-order to obtain valid inference. It is not clear how to accomplish this when using data-adaptive learning algorithms. In Appendix A, we show how this can be done when a parametric working model for the conditional mean outcome is used. An alternative for standard error calculation, which does not demand the assumption of a correctly specified missingness model, is the non-parametric bootstrap, which however lacks supporting theory when the estimators rely on data-adaptive estimation [32] and thus will only provide valid results if one can claim asymptotic linearity of the data-adaptive estimator for the conditional mean outcome model. van der Laan [32] proposes a strategy to obtain valid inference under misspecification of one of both working models in the context of the TMLE procedure. This is accomplished by additionally fluctuating initial estimators of the working models such that the scores with respect to the corresponding fluctuation parameters guarantee the remainder term in the expansion (2) to be asymptotically linear. In the context of the current paper, this would demand adding an additional covariate to the fluctuation models considered in Section 4.2. Further research is needed to see how this can be integrated within the proposed approach.

Appendix

A Asymptotic Linearity Theorem

We present an asymptotic linearity theorem with corresponding influence function for the doubly robust estimators $\hat{\psi}_n^{(c)}$ ($c = 1, 2, 3$) of Section 4 under the assumption of a correctly specified working model for the missingness mechanism but a potentially misspecified working model for the conditional mean outcome, which can be either finite- or infinite-dimensional.

Theorem 1 (Asymptotic linearity of the proposed doubly robust estimator) *With the notation given in the previous section, let $\bar{Q}_{(c)}^*$ denote the probability limit of $\hat{Q}_n^{(c)}$ and let g_{MLE}^* denote the probability limit of $\hat{g}_n^{\text{MLE}}$; thus $\hat{Q}_n^{(c)} \xrightarrow{P} \bar{Q}_{(c)}^*$ and $\hat{g}_n^{\text{MLE}} \xrightarrow{P} g_{\text{MLE}}^*$. Because we assume a correctly specified model for the missingness mechanism, $g_{\text{MLE}}^* = g_0$. This implies that $P_0\{D^*(\bar{Q}_{(c)}^*, g_{\text{MLE}}^*; \psi_0)\} = 0$. By definition of the proposed estimation strategy, we have that*

$$\begin{aligned} P_n\{D^*(\hat{Q}_n^{(c)}, \hat{g}_n^{\text{MLE}}; \hat{\psi}_n^{(c)})\} &= 0, \\ P_n\{D_g^{\text{MLE}}(\hat{\alpha}_n^{\text{MLE}})\} &= \mathbf{0}, \\ P_n\{D_\alpha^*(\hat{Q}_n^{(c)}, \hat{g}_n^{\text{MLE}})\} &= \mathbf{0}. \end{aligned}$$

We make the following regularity conditions;

Donsker class condition

Suppose that the set $\mathcal{D}_1 = \{D^*(\bar{Q}, g; \psi_0) : \bar{Q}, g\}$ is a P_0 -Donsker class where $(\bar{Q}, g)$ varies over a set containing the sequences $(\hat{Q}_n^{(c)}, \hat{g}_n^{\text{MLE}})$ and $(\bar{Q}_{(c)}^*, g_{\text{MLE}}^*)$ with probability tending to one.

Consistency condition of D^*

Assume that

$$P_0 \left[\left\{ D^*(\hat{Q}_n^{(c)}, \hat{g}_n^{\text{MLE}}; \psi_0) - D^*(\bar{Q}_{(c)}^*, g_{\text{MLE}}^*; \psi_0) \right\}^2 \right] \rightarrow 0,$$

$$P_0 \left[\left\{ D^*(\bar{Q}_{(c)}^*, \hat{g}_n^{\text{MLE}}; \psi_0) - D^*(\bar{Q}_{(c)}^*, g_{\text{MLE}}^*; \psi_0) \right\}^2 \right] \rightarrow 0,$$

in probability as $n \rightarrow \infty$.

Glivenko-Cantelli class condition

Suppose that $\mathcal{D}_2 = \{D_{\alpha}^*(\bar{Q}, g) : \bar{Q}, g\}$ is a P_0 -Glivenko-Cantelli class where $(\bar{Q}, g)$ varies over a set containing the sequence $(\hat{Q}_n^{(c)}, \hat{g}_n^{\text{MLE}})$ with probability tending to one.

Consistency condition of D_{α}^*

With $D_{\alpha}^* = (D_{\alpha,1}^*, \dots, D_{\alpha,p}^*)^T$, assume that

$$P_0 \left[\left\{ D_{\alpha,i}^*(\hat{Q}_n^{(c)}, \hat{g}_n^{\text{MLE}}) - D_{\alpha,i}^*(\bar{Q}_{(c)}^*, g_{\text{MLE}}^*) \right\}^2 \right] \rightarrow 0$$

in probability as $n \rightarrow \infty$ for $i = 1, \dots, p$.

Second-order term condition

Define the second order term

$$R_n^* = P_0 \left[R \left\{ \hat{Q}_n^{(c)}(\mathbf{X}) - \bar{Q}_{(c)}^*(\mathbf{X}) \right\} \left\{ \frac{\hat{g}_n^{\text{MLE}}(\mathbf{X}) - g_{\text{MLE}}^*(\mathbf{X})}{\hat{g}_n^{\text{MLE}}(\mathbf{X}) g_{\text{MLE}}^*(\mathbf{X})} \right\} \right]$$

and assume $R_n^* = o_p(1/\sqrt{n})$. Note that this second-order term involves the product of the differences $\hat{Q}_n^{(c)}(\mathbf{X}) - \bar{Q}_{(c)}^*(\mathbf{X})$ and $\hat{g}_n^{\text{MLE}}(\mathbf{X}) - g_{\text{MLE}}^*(\mathbf{X})$.

Then, ψ_n is an asymptotically linear estimator of ψ_0 at P_0 with influence function $D^*(\bar{Q}_{(c)}^*, g_0; \psi_0)$. That is,

$$\hat{\psi}_n - \psi_0 = (P_n - P_0)D^*(\bar{Q}_{(c)}^*, g_0; \psi_0) + o_p(1/\sqrt{n}).$$

In particular, $\sqrt{n}(\hat{\psi}_n - \psi_0)$ converges in distribution to a normal distribution with mean zero and variance $\sigma_0^2 = P_0[\{D^*(\bar{Q}_{(c)}^*, g_0; \psi_0)\}^2]$.

Proof. By definition of the proposed estimator, we know that

$$P_n\{D^*(\hat{Q}_n^{(c)}, \hat{g}_n^{\text{MLE}}; \psi_n^{(c)})\} = 0.$$

Furthermore, because we assume the model for the missingness mechanism to be correctly specified, we have that $P_0\{D^*(\bar{Q}_{(c)}^*, g_{\text{MLE}}^*; \psi_0)\} = 0$. From this, it follows that

$$\begin{aligned}
\hat{\psi}_n - \psi_0 &= \frac{1}{n} \sum_{i=1}^n \left[\frac{R_i}{\hat{g}_n^{\text{MLE}}(\mathbf{X}_i)} \left\{ Y_i - \hat{Q}_n^{(c)}(\mathbf{X}_i) \right\} + \hat{Q}_n^{(c)}(\mathbf{X}_i) - \psi_0 \right] \\
&= P_n \left\{ D^*(\hat{Q}_n^{(c)}, \hat{g}_n^{\text{MLE}}; \psi_0) \right\} \\
&= P_n \left\{ D^*(\hat{Q}_n^{(c)}, \hat{g}_n^{\text{MLE}}; \psi_0) \right\} - P_0 \left\{ D^*(\bar{Q}_{(c)}^*, g_{\text{MLE}}^*; \psi_0) \right\} \\
&= (P_n - P_0) \left\{ D^*(\bar{Q}_{(c)}^*, g_{\text{MLE}}^*; \psi_0) \right\} \\
&\quad + P_n \left\{ D^*(\hat{Q}_n^{(c)}, \hat{g}_n^{\text{MLE}}; \psi_0) - D^*(\bar{Q}_{(c)}^*, g_{\text{MLE}}^*; \psi_0) \right\} \\
&= (P_n - P_0) \left\{ D^*(\bar{Q}_{(c)}^*, g_{\text{MLE}}^*; \psi_0) \right\} \\
&\quad + P_n \left\{ D^*(\hat{Q}_n^{(c)}, \hat{g}_n^{\text{MLE}}; \psi_0) - D^*(\bar{Q}_{(c)}^*, g_{\text{MLE}}^*; \psi_0) \right\} \tag{27}
\end{aligned}$$

$$+ P_n \left\{ D^*(\bar{Q}_{(c)}^*, \hat{g}_n^{\text{MLE}}; \psi_0) - D^*(\bar{Q}_{(c)}^*, g_{\text{MLE}}^*; \psi_0) \right\}. \tag{28}$$

Let us first consider the term (27). We have

$$\begin{aligned}
&P_n \left\{ D^*(\hat{Q}_n^{(c)}, \hat{g}_n^{\text{MLE}}; \psi_0) - D^*(\bar{Q}_{(c)}^*, \hat{g}_n^{\text{MLE}}; \psi_0) \right\} \\
&= (P_n - P_0) \left\{ D^*(\hat{Q}_n^{(c)}, \hat{g}_n^{\text{MLE}}; \psi_0) - D^*(\bar{Q}_{(c)}^*, \hat{g}_n^{\text{MLE}}; \psi_0) \right\} \\
&\quad + P_0 \left\{ D^*(\hat{Q}_n^{(c)}, \hat{g}_n^{\text{MLE}}; \psi_0) - D^*(\bar{Q}_{(c)}^*, \hat{g}_n^{\text{MLE}}; \psi_0) \right\}
\end{aligned}$$

and from the Donsker class condition and the consistency condition of D^* , it follows that

$$(P_n - P_0) \left\{ D^*(\hat{Q}_n^{(c)}, \hat{g}_n^{\text{MLE}}; \psi_0) - D^*(\bar{Q}_{(c)}^*, \hat{g}_n^{\text{MLE}}; \psi_0) \right\} = o_p(1/\sqrt{n}).$$

Next consider the term $P_0 \left\{ D^*(\hat{Q}_n^{(c)}, \hat{g}_n^{\text{MLE}}; \psi_0) - D^*(\bar{Q}_{(c)}^*, \hat{g}_n^{\text{MLE}}; \psi_0) \right\}$. We have that

$$\begin{aligned}
&P_0 \left\{ D^*(\hat{Q}_n^{(c)}, \hat{g}_n^{\text{MLE}}; \psi_0) - D^*(\bar{Q}_{(c)}^*, \hat{g}_n^{\text{MLE}}; \psi_0) \right\} \\
&= P_0 \left\{ D^*(\hat{Q}_n^{(c)}, g_{\text{MLE}}^*; \psi_0) - D^*(\bar{Q}_{(c)}^*, g_{\text{MLE}}^*; \psi_0) \right\} \\
&\quad + P_0 \left\{ D^*(\hat{Q}_n^{(c)}, \hat{g}_n^{\text{MLE}}; \psi_0) - D^*(\bar{Q}_{(c)}^*, \hat{g}_n^{\text{MLE}}; \psi_0) \right. \\
&\quad \left. - D^*(\hat{Q}_n^{(c)}, g_{\text{MLE}}^*; \psi_0) + D^*(\bar{Q}_{(c)}^*, g_{\text{MLE}}^*; \psi_0) \right\}
\end{aligned}$$

We have that

$$\begin{aligned}
&P_0 \left\{ D^*(\hat{Q}_n^{(c)}, \hat{g}_n^{\text{MLE}}; \psi_0) - D^*(\bar{Q}_{(c)}^*, \hat{g}_n^{\text{MLE}}; \psi_0) - D^*(\hat{Q}_n^{(c)}, g_{\text{MLE}}^*; \psi_0) + D^*(\bar{Q}_{(c)}^*, g_{\text{MLE}}^*; \psi_0) \right\} \\
&= P_0 \left(\frac{R}{g_{\text{MLE}}^*(\mathbf{X})} \left[\left\{ Y - \bar{Q}_{(c)}^*(\mathbf{X}) \right\} - \left\{ Y - \hat{Q}_n^{(c)}(\mathbf{X}) \right\} \right] - \frac{R}{\hat{g}_n^{\text{MLE}}(\mathbf{X})} \left[\left\{ Y - \bar{Q}_{(c)}^*(\mathbf{X}) \right\} - \left\{ Y - \hat{Q}_n^{(c)}(\mathbf{X}) \right\} \right] \right) \\
&= P_0 \left[R \left\{ \frac{1}{g_{\text{MLE}}^*(\mathbf{X})} - \frac{1}{\hat{g}_n^{\text{MLE}}(\mathbf{X})} \right\} \left\{ \hat{Q}_n^{(c)}(\mathbf{X}) - \bar{Q}_{(c)}^*(\mathbf{X}) \right\} \right] = R_n^*,
\end{aligned}$$

so that

$$\begin{aligned}
&P_0 \left\{ D^*(\hat{Q}_n^{(c)}, \hat{g}_n^{\text{MLE}}; \psi_0) - D^*(\bar{Q}_{(c)}^*, \hat{g}_n^{\text{MLE}}; \psi_0) \right\} \\
&= P_0 \left\{ D^*(\hat{Q}_n^{(c)}, g_{\text{MLE}}^*; \psi_0) - D^*(\bar{Q}_{(c)}^*, g_{\text{MLE}}^*; \psi_0) \right\} + R_n^*
\end{aligned}$$

and by assumption, the second-order term $R_n^* = o_p(1/\sqrt{n})$. Finally, note that

$$\begin{aligned} & P_0 \left\{ D^* (\hat{Q}_n^{(c)}, g_{\text{MLE}}^*; \psi_0) - D^* (\bar{Q}_{(c)}^*, g_{\text{MLE}}^*; \psi_0) \right\} \\ &= -P_0 \left[\left\{ \hat{Q}_n^{(c)}(\mathbf{X}) - \bar{Q}_{(c)}(\mathbf{X}) \right\} \left\{ \frac{R - g_{\text{MLE}}^*(\mathbf{X})}{g_{\text{MLE}}^*(\mathbf{X})} \right\} \right], \end{aligned}$$

which is $o_p(1/\sqrt{n})$ under the assumption that $g_{\text{MLE}}^*(\mathbf{X})$ equals the true missingness mechanism $g_0(\mathbf{X})$. We conclude that eq. (27) is $o_p(1/\sqrt{n})$.

We next consider the second term (28), which is the contribution of the estimation of the missingness mechanism. From the way we estimated the fluctuation parameter $\boldsymbol{\varepsilon}$, it will follow we can ignore this term; that is, (28) is $o_p(1/\sqrt{n})$, even under misspecification of the estimator of $\bar{Q}_0$. Define the function $\Upsilon_{(c)}(\boldsymbol{\alpha}) = D^* \{ \bar{Q}_{(c)}^*, g(\boldsymbol{\alpha}); \psi_0 \}$. It follows that (28) equals $P_n \{ \Upsilon_{(c)}(\hat{\boldsymbol{\alpha}}_n^{\text{MLE}}) - \Upsilon_{(c)}(\boldsymbol{\alpha}_0) \}$. We know that $\hat{\boldsymbol{\alpha}}_n^{\text{MLE}}$ is an asymptotically linear estimator for $\boldsymbol{\alpha}_0$ with influence function $-P_0 \{ D_{\boldsymbol{\alpha}, g}^{\text{MLE}}(\boldsymbol{\alpha}_0) \} D_g^{\text{MLE}}(\boldsymbol{\alpha}_0)$, that is,

$$\hat{\boldsymbol{\alpha}}_n^{\text{MLE}} - \boldsymbol{\alpha}_0 = (P_n - P_0) \left[-P_0 \{ D_{\boldsymbol{\alpha}, g}^{\text{MLE}}(\boldsymbol{\alpha}_0) \}^{-1} D_g^{\text{MLE}}(\boldsymbol{\alpha}_0) \right] + o_p(1/\sqrt{n})$$

and where $D_{\boldsymbol{\alpha}, g}^{\text{MLE}}(\boldsymbol{\alpha}_0) = \partial D_g^{\text{MLE}}(\boldsymbol{\alpha}) / \partial \boldsymbol{\alpha} |_{\boldsymbol{\alpha}=\boldsymbol{\alpha}_0}$. Consider the Taylor expansion

$$\Upsilon_{(c)}(\hat{\boldsymbol{\alpha}}_n^{\text{MLE}}) - \Upsilon_{(c)}(\boldsymbol{\alpha}_0) = \Upsilon_{(c), \boldsymbol{\alpha}}^T(\boldsymbol{\alpha}_0)(\hat{\boldsymbol{\alpha}}_n^{\text{MLE}} - \boldsymbol{\alpha}_0) + o_p(|\hat{\boldsymbol{\alpha}}_n^{\text{MLE}} - \boldsymbol{\alpha}_0|)$$

with $\Upsilon_{(c), \boldsymbol{\alpha}}(\boldsymbol{\alpha}_0) = \partial \Upsilon_{(c)}(\boldsymbol{\alpha}) / \partial \boldsymbol{\alpha} |_{\boldsymbol{\alpha}=\boldsymbol{\alpha}_0}$. Note that by the asymptotic linearity of $\hat{\boldsymbol{\alpha}}_n^{\text{MLE}}$, $o_p(|\hat{\boldsymbol{\alpha}}_n^{\text{MLE}} - \boldsymbol{\alpha}_0|) = o_p(1/\sqrt{n})$. Consequently,

$$\begin{aligned} & P_n \left\{ \Upsilon_{(c)}(\hat{\boldsymbol{\alpha}}_n^{\text{MLE}}) - \Upsilon_{(c)}(\boldsymbol{\alpha}_0) \right\} \\ &= P_n \left\{ \Upsilon_{(c), \boldsymbol{\alpha}}^T(\boldsymbol{\alpha}_0) \right\} (\hat{\boldsymbol{\alpha}}_n^{\text{MLE}} - \boldsymbol{\alpha}_0) + o_p(1/\sqrt{n}) \\ &= (P_n - P_0) \left[-P_0 \left\{ \Upsilon_{(c), \boldsymbol{\alpha}}^T(\boldsymbol{\alpha}_0) \right\} P_0 \left\{ D_{\boldsymbol{\alpha}, g}^{\text{MLE}}(\boldsymbol{\alpha}_0) \right\}^{-1} D_g^{\text{MLE}}(\boldsymbol{\alpha}_0) \right] + o_p(1/\sqrt{n}) \end{aligned}$$

where the last equality follows from the asymptotic linearity of $\hat{\boldsymbol{\alpha}}_n^{\text{MLE}}$ and the fact that $P_n \left\{ \Upsilon_{(c), \boldsymbol{\alpha}}^T(\boldsymbol{\alpha}_0) \right\} = P_0 \left\{ \Upsilon_{(c), \boldsymbol{\alpha}}^T(\boldsymbol{\alpha}_0) \right\} + o_p(1/\sqrt{n})$. In principle, when the probability limit $\bar{Q}_{(c)}^*$ is different from the truth $\bar{Q}_0$, this term would contribute to the influence function of $\hat{\psi}_n^{(c)}$. However, by definition of $\hat{\boldsymbol{\varepsilon}}_n^{(c)}$ and since that $\hat{Q}_n^{(c)} = \hat{Q}_n^0(\hat{\boldsymbol{\varepsilon}}_n^{(c)})$, we have that $P_n \left\{ D_{\boldsymbol{\alpha}}(\hat{Q}_n^{(c)}, \hat{g}_n^{\text{MLE}}) \right\} = \mathbf{0}$ for every n . From the Glivenko-Cantelli class condition and the consistency condition of $D_{\boldsymbol{\alpha}}^*$, it then follows that $P_n \{ D_{\boldsymbol{\alpha}}(\hat{Q}_n^{(c)}, \hat{g}_n^{\text{MLE}}) \} \xrightarrow{P} P_0 \{ D_{\boldsymbol{\alpha}}(\bar{Q}_{(c)}^*, g_{\text{MLE}}^*) \} = P_0 \{ \Upsilon_{(c), \boldsymbol{\alpha}}(\boldsymbol{\alpha}_0) \}$. Hence, it follows that $P_0 \{ \Upsilon_{(c), \boldsymbol{\alpha}}(\boldsymbol{\alpha}_0) \} = \mathbf{0}$. We conclude that eq. (28) is $o_p(1/\sqrt{n})$.

Putting everything together, we find that

$$\hat{\psi}_n^{(c)} - \psi_0 = (P_n - P_0) \left\{ D^* (\bar{Q}_{(c)}^*, g_0; \psi_0) \right\} + o_p(1/\sqrt{n}),$$

showing asymptotic linearity of $\hat{\psi}_n^{(c)}$. □

Remark. When the updated estimator $\hat{Q}_n^{(c)}$ ($c=1, 2, 3$) is based on a parametric model for the initial estimator, the Donsker class condition will be satisfied. Furthermore, when the updated estimator $\hat{Q}_n^{(c)}$ ($c=1, 2, 3$) is based on the super-learner $\hat{Q}_n^{\text{SL}}$ as an initial estimator and each of the estimators in the library falls in a Donsker class, the Donsker class condition will also be satisfied for $\hat{Q}_n^{\text{SL}}$ because the convex combination of such estimators also falls in that class [39]. In short, the Donsker class condition is satisfied if it holds for each of the estimators in the library of the super-learner. Examples for such estimators are for instance given in van der Laan [32].

Theorem 1 establishes asymptotic linearity of the doubly robust estimators $\hat{\psi}_n^{(c)}$ ($c=1, 2, 3$) under model $\mathcal{M}(\mathcal{G})$. To obtain asymptotic linearity under model $\mathcal{M}(\mathcal{Q})$ (so that $\bar{Q}_{(c)}^* = \bar{Q}_0$), one should assume asymptotic linearity of $\hat{Q}_n^{(c)}$ in the sense that $P_0\{D^*(\hat{Q}_n^{(c)}, g_{\text{MLE}}^*; \psi_0) - D^*(\bar{Q}_{(c)}^*, g_{\text{MLE}}^*; \psi_0)\} = (P_n - P_0)D_{\bar{Q}_{(c)}^*}^* + o_p(1/\sqrt{n})$ (see van der Laan and Rose [19], p. 572 for a discussion concerning this assumption). In this case, the influence function of $\hat{\psi}_n^{(c)}$ under model $\mathcal{M}(\mathcal{Q})$ becomes $D^*(\bar{Q}_{(c)}^*, g_{\text{MLE}}^*; \psi_0) + D_{\bar{Q}_{(c)}^*}^*$. When the initial estimator $\hat{Q}_n^0$ for the conditional mean outcome $\bar{Q}_0$ is parametric such as $\bar{Q}(\gamma)(\mathbf{X}) = Q\{\gamma^T \mathbf{I}(\mathbf{X})\}$, and γ is estimated via a root- n consistent estimator such as the MLE, asymptotic linearity of $\hat{\psi}_n^{(c)}$ ($c=1, 2, 3$) under model $\mathcal{M}(\mathcal{Q})$ can be shown in a straightforward manner, as we argue next. In this case, the updated estimator $\hat{Q}_n^{(c)}$ is based on a parametric model $\bar{Q}(\theta^{(c)})$ for the conditional mean outcome with dimension equal to the dimension of $\theta^{(c)} = (\gamma^T, \varepsilon^{(c), T})^T$. Let $\hat{\theta}_n^{(c)}$ denote the corresponding estimator of $\theta^{(c)}$ with probability limit $\theta_{(c)}^*$. When the parametric model for the conditional mean outcome is correctly specified, $\theta_{(c)}^* = (\gamma_0^T, \mathbf{0}^T)^T$. Under standard regularity conditions, $\hat{\theta}_n^{(c)}$ is asymptotically linear with influence function $D_{\bar{Q}}^*(\theta^{(c)})$; that is, $(\hat{\theta}_n^{(c)} - \theta_{(c)}^*) = (P_n - P_0)D_{\bar{Q}}^*(\theta_{(c)}^*) + o_p(1/\sqrt{n})$. We now show that under $\mathcal{M}(\mathcal{Q})$, equation (27) is not $o_p(1/\sqrt{n})$ because then $g_{\text{MLE}}^* \neq g_0$. For this purpose, define the function $\xi_{(c)}(\theta^{(c)}, \alpha) = D^*\{\bar{Q}(\theta^{(c)}), g(\alpha); \psi_0\}$. It follows from a Taylor expansion and the asymptotic linearity of $\hat{\theta}_n^{(c)}$, that eq. (27) equals

$$\begin{aligned} & P_n \left\{ \xi_{(c)}(\hat{\theta}_n^{(c)}, \hat{\alpha}_n^{\text{MLE}}) - \xi_{(c)}(\theta_{(c)}^*, \hat{\alpha}_n^{\text{MLE}}) \right\} \\ &= P_n \left\{ \xi_{(c), \theta^{(c)}}^T(\theta_{(c)}^*, \hat{\alpha}_n^{\text{MLE}}) \right\} (\hat{\theta}_n^{(c)} - \theta_{(c)}^*) + o_p(|\hat{\theta}_n^{(c)} - \theta_{(c)}^*|) \\ &= P_n \left\{ \xi_{(c), \theta^{(c)}}^T(\theta_{(c)}^*, \hat{\alpha}_n^{\text{MLE}}) \right\} (P_n - P_0)D_{\bar{Q}}^*(\theta_{(c)}^*) + o_p(1/\sqrt{n}), \end{aligned}$$

with $\xi_{(c), \theta^{(c)}}^T(\theta_{(c)}^*, \hat{\alpha}_n^{\text{MLE}}) = \partial \xi_{(c)}(\theta^{(c)}, \hat{\alpha}_n^{\text{MLE}}) / \partial \theta^{(c)}|_{\theta^{(c)} = \theta_{(c)}^*}$. Under suitable regularity conditions (Robins et al. [40], app. B), it follows from the uniform weak law of large numbers that $P_n \left\{ \xi_{(c), \theta^{(c)}}^T(\theta_{(c)}^*, \hat{\alpha}_n^{\text{MLE}}) \right\}$ converges in probability to $P_0 \left\{ \xi_{(c), \theta^{(c)}}^T(\theta_{(c)}^*, \alpha_{\text{MLE}}^*) \right\}$. This shows that eq. (27) can be written as

$$(P_n - P_0) \left[P_0 \left\{ \xi_{(c), \theta^{(c)}}^T(\theta_{(c)}^*, \alpha_{\text{MLE}}^*) \right\} D_{\bar{Q}}^*(\theta_{(c)}^*) \right] + o_p(1/\sqrt{n}).$$

Upon replacing α_0 by α_{MLE}^* in the remainder of the proof of Theorem 1, it follows that the influence function of $\hat{\psi}_n^{(c)}$ under model $\mathcal{M}(\mathcal{Q})$ equals $D^*(\bar{Q}_{(c)}^*, g_{\text{MLE}}^*; \psi_0) + P_0 \left\{ \xi_{(c), \theta^{(c)}}^T(\theta_{(c)}^*, \alpha_{\text{MLE}}^*) \right\} D_{\bar{Q}}^*(\theta_{(c)}^*)$. For this specific setting we thus have that $D_{\bar{Q}_{(c)}^*}^* = P_0 \left\{ \xi_{(c), \theta^{(c)}}^T(\theta_{(c)}^*, \alpha_{\text{MLE}}^*) \right\} D_{\bar{Q}}^*(\theta_{(c)}^*)$. Local-ancillarity with respect to γ can however still be obtained by implementing bias-reduction in the direction of γ as well, enforcing this term to be zero (see Section 3).

Remark. Under model $\mathcal{M}(\mathcal{G})$, $\alpha_{\text{MLE}}^* = \alpha_0$ and hence $P_0 \left\{ \xi_{(c), \theta^{(c)}}^T(\theta_{(c)}^*, \alpha_0) \right\} = \mathbf{0}$. We may hence conclude from the derivation above (but now with the probability limit $\theta_{(c)}^* \neq (\gamma_0^T, \mathbf{0}^T)^T$) that eq. (27) is $o_p(1/\sqrt{n})$ for a parametric initial estimator for $\bar{Q}_0$ which is potentially misspecified, given the parametric model for the missingness mechanism is correctly specified. This argument is more insightful than the Donsker class condition and the second-order term condition.

B R-function

Below, we provide an R-function to obtain the data-adaptive bias-reduced doubly robust estimators $\hat{\psi}_n^{(c)}$ to estimate the mean outcome $E(Y)=\psi_0$ in the presence of incomplete data for both a linear regression working model and a super-learner working model for the initial estimator of the conditional mean outcome and with a logistic regression working model for the propensity score.

As input, the function uses the missingness indicator R , the outcome Y , the auxiliary covariates cov , the estimation method for the initial estimator of the conditional mean outcome `type.initQ=c("par", "SL")` (either "par" for a parametric linear regression model or "SL" for a super-learner), the level ζ at which the initial estimator must be truncated, the type of loss-function for the logistic fluctuation model `fluc=c("unweighted", "weighted")`, and the level of significance α of a hypothesis test of the mean equal to ψ_0 . As output, the function delivers the estimate `est`, which equals $\hat{\psi}_n^{(c)}$, the standard error `se`, which equals $\widehat{SE}(\hat{\psi}_n^{(c)})$ based on eq. (21), a $(1-\alpha)100\%$ Wald confidence interval `ci`, the Wald statistic `Wald.statistic` and the corresponding p -value `p.value` for a test of the null hypothesis $H_0:\psi_0=\tilde{\psi}$. We also provide an example where the procedure is applied to a random dataset, obtained via the Kang and Schafer data-generating mechanism.

```
data.adaptive.biasreduced.DR <-
  function(R,Y,cov,type.initQ=c("par","SL"),
    zeta=0.005,fluc=c("unweighted","weighted"),
    alpha=0.05,psi.tilde=0){
  expit <- function(x){1/(1+exp(-x))}
  logit <- function(x){log(x/(1-x))}
  n <- length(R)
  dat.cov <- data.frame(cov)
  int.cov <- cbind(rep(1,n),cov)
  colnames(dat.cov) <-
    paste("cov.",1:dim(cov)[2],sep="")

  # propensity score
  mler <- glm(R~cov,family="binomial")
  ps.par <- predict(mler,type="response")

  # initial conditional mean outcome
  a <- min(Y[R==1]) - 0.1*abs(min(Y[R==1]))
  b <- max(Y[R==1]) + 0.1*abs(max(Y[R==1]))
  Y.star <- (Y-a)/(b-a)
  {
  if(type.initQ=="par"){
    mley <- lm(Y.star~cov,subset=(R==1))
    initQ <- predict(mley,newdata=dat.cov)
  }
  else if(type.initQ=="SL"){
    Ym.star <- Y.star[R==1]
    dat.cov.m <- dat.cov[R==1,,drop=FALSE]
    SL.library <- c("SL.glm", "SL.randomForest",
      "SL.gam", "SL.polymars", "SL.mean")
    initQ <- SuperLearner(Y=Ym.star,X=dat.cov.m,
      newX=dat.cov,verbose = FALSE,
```

```

        SL.library=SL.library,
        method="method.NNLS")$SL.predict
    }
  }
  initQ.trunc <- ifelse(initQ < zeta,zeta,
    ifelse(initQ > 1-zeta,1-zeta,initQ))

  # fluctuation
  {
    if(fluc=="unweighted"){
      w.cov <- (1-ps.par)/ps.par*int.cov
      fluctuationQ <- glm(Y.star~1+w.cov,
        family=binomial,offset=logit(initQ.trunc),
        subset=(R==1))
      flucQ <- expit(logit(initQ.trunc)+
        as.vector(coef(fluctuationQ)%*%t(w.cov)))
    }
    else if(fluc=="weighted"){
      fluctuationQ <- glm(Y.star~cov,
        family=binomial,offset=logit(initQ.trunc),
        subset=(R==1)) weights=(1-ps.par)/ps.par
      flucQ <- expit(logit(initQ.trunc)+
        as.vector(coef(fluctuationQ)%*%t(int.cov)))
    }
  }
  # doubly robust estimator
  U <- function(R,Y,outcome,ps){
    outcome+R/ps*(Y-outcome)
  }
  est.trunc <- mean(U(R=R,Y=Y.star,
    outcome=flucQ,ps=ps.par))
  psi <- (b-a)*est.trunc+a
  # standard error
  se.psi <- (b-a)*sd(U(R=R,Y=Y.star,
    outcome=flucQ,ps=ps.par))/sqrt(n)

  # 95% confidence interval
  ci.psi <- psi+c(-1,1)*qnorm(1-alpha/2)*se.psi

  # Wald test statistic
  W <- (psi-psi.tilde)/se.psi

  # p-value Wald test
  p.value <- 2*pnorm(abs(W),lower.tail=FALSE)

  return(list(est=psi,se=se.psi,ci=ci.psi,
    Wald.statistic=W,p.value=p.value))
}

```


Example

```
library(SuperLearner)

gen.dataKS <- function(k,n,mech=
  c("normal","reverse"),spec=c("C","I")){
  set.seed(k)
  z1 <- rnorm(n); z2 <- rnorm(n)
  z3 <- rnorm(n); z4 <- rnorm(n)
  z <- cbind(z1,z2,z3,z4)
  colnames(z) <- paste("Z.",1:4,sep="")
  x1 <- exp(0.5*z1)
  x2 <- z2/(1+exp(z1))+10
  x3 <- (0.04*z1*z3+0.6)^3
  x4 <- (z2+z4+20)^2
  y <- rnorm(n,210+27.4*z1+13.7*z2+13.7*z3+13.7*z4,sd=1)
  r <- rbinom(n,1,expit(-1.5*z1+0.75*z2-0.375*z3-0.15*z4))
  if(spec=="C"){
    x1 <- z1;x2 <- z2
    x3 <- z3;x4 <- z4;
    x <- z
  }
  if(mech=="reverse"){
    r <- 1-r
  }
  data <- data.frame(r,y,x1,x2,x3,x4)
  names(data) <- c("r","y","x.1","x.2","x.3","x.4")
  data
}
data <- gen.dataKS(1,n=1000,mech="reverse",spec="I")
R <- data$r
Y <- data$y
Cov <- cbind(data$x.1,data$x.2,data$x.3,data$x.4)

dataadaptive.biasreduced.DR(R=R,Y=Y,cov=Cov,
  type.initQ="SL",zeta=0.005,
  fluc="weighted",alpha=0.05,psi.tilde=210)
```

References

1. Robins JM, Rotnitzky A. Comment on the Bickel and Kwon article, "Inference for Semiparametric Models: Some Questions and an Answer". *Stat Sinica* 2001;11:920–36.
2. Tsiatis AA, Davidian M, Zhang M, Lu X. Covariate adjustment for two-sample treatment comparisons in randomized clinical trials: a principled yet flexible approach. *Stat Med* 2008;27:4658–77.
3. Moore KL, van der Laan MJ. Covariate adjustment in randomized trials with binary outcomes: targeted maximum likelihood estimation. *Stat Med* 2009;28:39–64.
4. Vermeulen K, Thas O, Vansteelandt S. Increasing the power of the Mann-Whitney test in randomized experiments through flexible covariate adjustment. *Stat Med* 2015;35:1012–30.
5. Kang JD, Schafer JL. Demystifying double robustness: a comparison of alternative strategies for estimating a population mean from incomplete data. *Stat Sci* 2007a;22:523–39.

6. Bang H, Robins JM. Doubly robust estimation in missing data and causal inference models. *Biometrics* 2005;61:962–72.
7. Vansteelandt S, Bekaert M, Claeskens G. On model selection and model misspecification in causal inference. *Stat Methods Med Res* 2012;21:7–30.
8. Vermeulen K, Vansteelandt S. Bias-reduced doubly robust estimation. *J Am Stat Assoc* 2015;110:1024–36.
9. van der Laan MJ, Rubin DB. Targeted maximum likelihood learning. *Int J Biostat* 2006;2(1), Article 11.
10. van der Laan MJ, Polley EC, Hubbard AE. “Super Learner,” *Statistical Applications in Genetics and Molecular Biology* 2007;6.
11. Ridgeway G, McCaffrey D. Comment: demystifying double robustness: a comparison of alternative strategies for estimating a population mean for incomplete data. *Stat Sci* 2007;22:540–3.
12. Robins JM, Sued M, Lei-Gomez Q, Rotnitzky A. Comment: performance of double-robust estimators when “Inverse Probability” weights are highly variable. *Stat Sci* 2007;22:544–59.
13. Tan Z. Comment: Understanding OR, PS and DR. *Stat Sci* 2007;22:560–8.
14. Tsiatis AA, Davidian M. Comment: demystifying double robustness: a comparison of alternative strategies for estimating a population mean from incomplete data. *Stat Sci* 2007;22:569–73.
15. Kang JD, Schafer JL. Rejoinder: demystifying double robustness: a comparison of alternative strategies for estimating a population mean from incomplete data. *Stat Sci* 2007b;22:574–80.
16. Rubin DB, van der Laan MJ. Empirical efficiency maximization: improved locally efficient covariate adjustment in randomized experiments and survival analysis. *Int J Biostat* 2008;4:1–40.
17. Cao WH, Tsiatis AA, Davidian M. Improving efficiency and robustness of the doubly robust estimator for a population mean with incomplete data. *Biometrika* 2009;96:723–34.
18. Tsiatis AA, Davidian M, Cao W. Improved doubly robust estimation when data are monotonely coarsened, with application to longitudinal studies with dropout. *Biometrics* 2011;67:536–45.
19. van der Laan MJ, Rose S. Targeted learning, causal inference for observational and experimental data. *Springer series in Statistics*, 2011.
20. Tan Z. Bounded, efficient and doubly robust estimation with inverse weighting. *Biometrika* 2010;97(3):661–82.
21. Rotnitzky A, Lei QH, Sued M, Robins JM. Improved double-robust estimation in missing data and causal inference models. *Biometrika* 2012;99:439–56.
22. van der Laan MJ, Gruber S. Collaborative double robust targeted maximum likelihood estimation. *Int J Biostat* 2010;6.
23. Rotnitzky A, Vansteelandt S. Double-robust methods. In: Molenberghs G, Fitzmaurice G, Kenward MG, Tsiatis AA, Verbeke G, editors. *Handbook of missing data methodology*. Boca Raton: CRC Press, 2015.
24. Porter KE, Gruber S, van der Laan MJ, Sekhon JS. The relative performance of targeted maximum likelihood estimators. *Int J Biostat* 2011;7.
25. Rubin DB. Inference and missing data. *Biometrika* 1976;63:581–90.
26. Scharfstein DO, Rotnitzky A, Robins JM. Adjusting for nonignorable drop-out using semiparametric nonresponse models. *J Am Stat Assoc* 1999a;94:1096–120.
27. Bickel PJ, Klaassen CA, Ritov Y, Wellner JA. Efficient and adaptive estimation for semiparametric models. Baltimore: The Johns Hopkins University Press, 1993.
28. Cox DR. Local ancillarity. *Biometrika* 1980;67:279–86.
29. Petersen ML, Porter KE, Gruber S, Wang Y, van der Laan MJ. Diagnosing and responding to violations in the positivity assumption. *Stat Methods Med Res* 2010;21:31–54.
30. Gruber S, van der Laan M. A targeted maximum likelihood estimator of a causal effect on a bounded continuous outcome. *Int J Biostat* 2010;6(1), Article 26.
31. Díaz I, Rosenblum M. Targeted maximum likelihood estimation using exponential families. *ArXiv e-prints*. *Int J Biostat* 2014;11(2):233–51.
32. van der Laan MJ. Targeted estimation of nuisance parameters to obtain valid statistical inference. *Int J Biostat* 2014;10:29–57.
33. Polley E, van der Laan M. Super Learner Prediction: the SuperLearner package, 2014.
34. Gruber S, van der Laan M. Targeted maximum likelihood estimation: the TMLE package, 2014.
35. Robins J. Correcting for noncompliance in randomized trials using structural nested mean models. *Commun Stat Theor Methods* 1994;23:2379–2412.
36. Wooldridge J. *Econometric analysis of cross section and panel data*. Cambridge, MA: The MIT Press, 2002.
37. Hernán MA, Robins J. Instruments for causal inference. an epidemiologist’s dream? *Epidemiology* 2006;17:360–72.
38. Okui R, Small DS, Tan Z, Robins JM. Doubly robust instrumental variable regression. *Stat Sinica* 2012;22:173–205.
39. van der Vaart AW, Wellner JA. *Weak convergence and empirical processes*. New-York: Springer-Verlag, 1996.
40. Robins JM, Rotnitzky A, Zhao LP. Estimation of regression-coefficients when some regressors are not always observed. *J Am Stat Assoc* 1994;89:846–66.