Taylor & Francis
Taylor & Francis Group

Journal of the American Statistical Association

ISSN: 0162-1459 (Print) 1537-274X (Online) Journal homepage: https://www.tandfonline.com/loi/uasa20

Bias-Reduced Doubly Robust Estimation

Karel Vermeulen & Stijn Vansteelandt

To cite this article: Karel Vermeulen & Stijn Vansteelandt (2015) Bias-Reduced Doubly Robust Estimation, Journal of the American Statistical Association, 110:511, 1024-1036, DOI: 10.1080/01621459.2014.958155

To link to this article: https://doi.org/10.1080/01621459.2014.958155

Bias-Reduced Doubly Robust Estimation

Karel VERMEULEN and Stijn VANSTEELANDT

Over the past decade, doubly robust estimators have been proposed for a variety of target parameters in causal inference and missing data models. These are asymptotically unbiased when at least one of two nuisance working models is correctly specified, regardless of which. While their asymptotic distribution is not affected by the choice of root-n consistent estimators of the nuisance parameters indexing these working models when all working models are correctly specified, this choice of estimators can have a dramatic impact under misspecification of at least one working model. In this article, we will therefore propose a simple and generic estimation principle for the nuisance parameters indexing each of the working models, which is designed to improve the performance of the doubly robust estimator of interest, relative to the default use of maximum likelihood estimators for the nuisance parameters. The proposed approach locally minimizes the squared first-order asymptotic bias of the doubly robust estimator under misspecification of both working models and results in doubly robust estimators with easy-to-calculate asymptotic variance. It moreover improves the stability of the weights in those doubly robust estimators which invoke inverse probability weighting. Simulation studies confirm the desirable finite-sample performance of the proposed estimators. Supplementary materials for this article are available online.

KEY WORDS: Causal inference; Double robustness; Missing data; Nuisance parameters; Semiparametric estimation.

1. INTRODUCTION

Estimation of most statistical parameters requires postulation of so-called nuisance working models: models not of primary scientific interest, but needed to obtain a well-behaved estimator of the target parameter in small to moderate sample sizes. For instance, in studies where outcome data are incomplete in a way that is explainable by measured covariates, estimation of the mean outcome requires modeling the dependence of those covariates on either the outcome or missingness. In typical causal inference problems, estimation of the exposure effect requires modeling the dependence of measured confounders on either the outcome or the exposure. A prevailing concern is that misspecification of those nuisance working models induces bias in the estimator of the target parameter (Robins 1999).

In many missing data and causal inference models, the concern for bias due to model misspecification can be lessened via the use of doubly (or multiply) robust, abbreviated DR, estimators. These consistently estimate the target parameter when at least one of two (or multiple) nuisance working models is correctly specified, regardless of which (Robins and Rotnitzky 2001). Since the seminal work by Scharfstein, Rotnitzky, and Robins (1999a) and Robins and Rotnitzky (2001), a variety of DR estimators have been developed. Bang and Robins (2005) gave an overview of work on DR estimation of the parameters indexing conditional mean models when the outcome data are incomplete, and of marginal treatment effects in causal inference models. DR estimators have also been developed for, for instance, statistical interaction parameters (Vansteelandt et al. 2008), controlled direct effects (Goetgeluk, Vansteelandt, and Goetghebeur 2008), natural direct and indirect effects in mediation analysis (Tchetgen Tchetgen and Shpitser 2012), incomplete covariate problems (Tchetgen Tchetgen and Rotnitzky 2011), and instrumental variables analysis (Okui et al. 2012).

The appeal of DR estimators surpasses the defining property of double protection against model misspecification. Their reliance on multiple nuisance working models, of which only one must be correctly specified, makes them potential "compromise" estimators amidst competing estimators that each rely on a single, but different working model. For instance, in missing data models, DR estimators form a compromise between imputation-based estimators that rely on an imputation model for the incomplete outcome, and inverse probability weighted estimators that rely on a model for the probability of missingness; arguably, they may therefore define the preferred analysis. Additionally, many DR estimators are locally efficient within a broad class of estimators. Because of this, their use has been advocated in randomized trial analyses: by exploiting the known randomization probabilities, they make it possible to increase power via covariate adjustment without risking bias due to model misspecification (Tsiatis et al. 2008; Moore and van der Laan 2009).

Estimation of the nuisance parameters indexing the working models in DR estimators has long been ignored. Theoretical results show that the choice of nuisance parameter estimators has no impact on the asymptotic variance of DR estimators when both working models are correctly specified (Tsiatis 2006). This has led to the default use of maximum likelihood estimators. This standard practice has gradually started to change when simulation studies by Kang and Schafer (2007a) cautioned for potentially disastrous performance of certain DR estimators (relative to simpler estimators) when both working models are misspecified. The discussions of that article (Ridgeway and McCaffrey 2007; Robins et al. 2007; Tan 2007; Tsiatis and Davidian 2007; Kang and Schafer 2007b) reveal that many different DR estimators may exist for a given target parameter, all with potentially very different behavior and properties under misspecification of at least one working model. In particular, when

Karel Vermeulen is Research Fellow of the Fund for Scientific Research (FWO), Department of Applied Mathematics, Computer Science and Statistics, Ghent University, 281 (S9) Krijgslaan, 9000 Ghent, Belgium (E-mail: *karelb.vermeulen@ugent.be*). Stijn Vansteelandt is Professor, Department of Applied Mathematics, Computer Science and Statistics, Ghent University, 9000 Ghent, Belgium (E-mail: *stijn.vansteelandt@ugent.be*). Karel Vermeulen is supported by a Ph.D. grant from the Research Foundation of Flanders (FWO). Stijn Vansteelandt acknowledges support from IAP research network grant P07/06 from the Belgian government (Belgian Science Policy). The authors are grateful to James Robins, Andrea Rotnitzky, Eric Tchetgen Tchetgen, and Zhiqiang Tan for helpful discussions, and to the Associate Editor and three reviewers for comments that helped improve an earlier version of this article.

Journal of the American Statistical Association
September 2015, Vol. 110, No. 511, Theory and Methods
DOI: 10.1080/01621459.2014.958155

a DR estimator exists for a given target parameter, then infinitely many can usually be constructed by varying the choice of nuisance parameter estimators. All resulting DR estimators have the same asymptotic behavior under correct specification of all working models, but a potentially very different behavior under model misspecification.

Rubin and van der Laan (2008), Cao, Tsiatis, and Davidian (2009), and Tsiatis, Davidian, and Cao (2011) developed DR estimators in specific missing data models with desirable efficiency properties when the missingness model is correctly specified. In their development, which generalizes that of Tan (2006), the nuisance parameters indexing the working model for the incomplete outcome are estimated by directly minimizing the variance of the DR estimator. The TMLE (targeted maximum likelihood estimation) procedure (van der Laan and Rose 2011) and the procedures of Tan (2010) and Rotnitzky et al. (2012) guarantee DR estimators of the population mean 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 DR estimators under misspecification of the working model for the full-data distribution (i.e., the dependence of outcome on covariates/confounders in missing data/causal inference models). The collaborative TMLE (C-TMLE) procedure of van der Laan and Gruber (2010), which is a further improvement upon the TMLE procedure, additionally focuses on the estimation of the missingness/exposure probabilities in a way that aims to prevent large variance of the DR estimator.

In this article, as in a recent article by van der Laan (2014), we take a different perspective by focusing on bias reduction rather than variance reduction. This is motivated by the fact that the bias of a DR estimator can get severely amplified under misspecification of at least one working model and may become especially severe under misspecification of both working models (Kang and Schafer 2007a; Vansteelandt, Bekaert, and Claeskens 2012). In particular, in Section 3 we propose a general estimating equation-based strategy, referred to as bias-reduced DR estimation, which locally minimizes the squared first-order asymptotic bias of the DR estimator in the direction of the nuisance parameters under misspecification of both working models. Our proposal differs from van der Laan (2014) in that it avoids bias approximations in view of the difficulty of approximating bias (see later). Unlike most other proposals, it focuses on the estimation of the nuisance parameters in *all* working models, is readily applicable to arbitrary DR estimators, can be adapted to certain multiply robust estimators (see Section 5), and returns DR estimators with easy-to-calculate asymptotic variance. Simulation studies in Section 4 and the analysis of an observational study in Section 6 demonstrate that the proposed estimator is competitive with existing alternatives.

2. DOUBLY ROBUST ESTIMATION

For pedagogic purposes, we first consider the estimation of a population mean outcome in the presence of incomplete data. While seemingly simple, this problem reveals fundamental challenges involving inverse probability weighting. Consider a study design which intends to collect iid data $\{(Y_i, \mathbf{X}_i), i = 1, \dots, n\}$, with Y_i the outcome and $\mathbf{X}_i$ a set of auxiliary covariates for subject i. Estimation of the mean $E(Y)$ is complicated by the fact that Y_i is not available for all individuals. Let R_i denote the missingness indicator, which codes $R_i = 1$ when Y_i is observed and $R_i = 0$ if Y_i is missing. The observed data can then be described as the random sample $\{\mathbf{Z}_i = (R_iY_i, R_i, \mathbf{X}_i), i = 1, \dots, n\}$. Assume that the covariates $\mathbf{X}_i$ contain sufficient information to explain missingness so that the missing at random (MAR) assumption, $Y_i \perp\!\!\!\perp R_i|\mathbf{X}_i$ (Tsiatis 2006), holds. Throughout, for some parameter λ we use λ_0 to denote its (unknown) population value; in particular, $E(Y) = \mu_0$.

Robins, Rotnitzky, and Zhao (1994) showed that when outcome data are missing, consistent estimation of μ_0 requires specification of at least one of the two following working models. The first is a working model for the probability of observing the data—referred to as the propensity score, abbreviated PS, throughout: $P(R = 1|\mathbf{X}) = \pi_0(\mathbf{X}) = \pi(\mathbf{X}; \boldsymbol{\gamma}_0)$, for which we assume $\pi_0(\mathbf{X}) > 0$ with probability one, where $\pi(\mathbf{X}; \boldsymbol{\gamma})$ is a known function, smooth in $\boldsymbol{\gamma}$ and $\boldsymbol{\gamma}_0$ is an unknown p-dimensional parameter; for example, a logistic regression model $\pi(\mathbf{X}; \boldsymbol{\gamma}) = \text{expit}(\gamma_1 + \boldsymbol{\gamma}_2^T\mathbf{X})$. This model is denoted $\mathscr{M}(\boldsymbol{\gamma}) = \{\pi(\mathbf{X}; \boldsymbol{\gamma}) : \boldsymbol{\gamma} \in \mathbb{R}^p\}$. The second is a working model for the conditional mean outcome $E(Y|\mathbf{X}) = m_0(\mathbf{X}) = m(\mathbf{X}; \boldsymbol{\beta}_0)$, where $m(\mathbf{X}; \boldsymbol{\beta})$ is a known function, smooth in $\boldsymbol{\beta}$ and $\boldsymbol{\beta}_0$ is an unknown q-dimensional parameter; for example, a linear model $m(\mathbf{X}; \boldsymbol{\beta}) = \beta_1 + \boldsymbol{\beta}_2^T\mathbf{X}$ for a continuous outcome Y. This model is denoted $\mathscr{M}(\boldsymbol{\beta}) = \{m(\mathbf{X}; \boldsymbol{\beta}) : \boldsymbol{\beta} \in \mathbb{R}^q\}$. Scharfstein, Rotnitzky, and Robins (1999a) showed that a DR estimator of μ_0, with $\tilde{E}_n(V) = n^{-1}\sum_{i=1}^n V_i$, can be obtained as

$$\hat{\mu}_{\text{DR}}(\hat{\boldsymbol{\gamma}}, \hat{\boldsymbol{\beta}}) = \tilde{E}_n\left\{\frac{RY}{\pi(\mathbf{X}, \hat{\boldsymbol{\gamma}})} - \frac{R - \pi(\mathbf{X}, \hat{\boldsymbol{\gamma}})}{\pi(\mathbf{X}, \hat{\boldsymbol{\gamma}})}m(\mathbf{X}, \hat{\boldsymbol{\beta}})\right\}, \quad (1)$$

for root-n consistent and asymptotically normal estimators $\hat{\boldsymbol{\gamma}}$ and $\hat{\boldsymbol{\beta}}$ for the nuisance parameters $\boldsymbol{\gamma}$ and $\boldsymbol{\beta}$ (Tsiatis 2006, chap. 3). This estimator is consistent for μ_0 under the union model $\mathscr{M}(\boldsymbol{\gamma}) \cup \mathscr{M}(\boldsymbol{\beta})$: as soon as one but not necessarily both working models are correctly specified. If the intersection model $\mathscr{M}(\boldsymbol{\gamma}) \cap \mathscr{M}(\boldsymbol{\beta})$ holds, that is, both working models are correctly specified, the DR estimator (1) is locally efficient (Tsiatis 2006) under model $\mathscr{M}(\boldsymbol{\gamma})$: it then has the smallest asymptotic variance within the class of all estimators that are consistent and asymptotically normal under $\mathscr{M}(\boldsymbol{\gamma})$, provided that also $\mathscr{M}(\boldsymbol{\beta})$ is correctly specified. The proposition below, which follows from standard results on M-estimation, gives the asymptotic distribution of the DR estimator (1) in the special case where $\hat{\boldsymbol{\gamma}}$ and $\hat{\boldsymbol{\beta}}$ are solutions to estimating equations $\tilde{E}_n\{\mathbf{U}_\gamma(\mathbf{Z}; \hat{\boldsymbol{\gamma}})\} = \mathbf{0}$ and $\tilde{E}_n\{\mathbf{U}_\beta(\mathbf{Z}; \hat{\boldsymbol{\beta}})\} = \mathbf{0}$. The more general result when the estimating functions $\mathbf{U}_\gamma$ and $\mathbf{U}_\beta$ involve both $\boldsymbol{\gamma}$ and $\boldsymbol{\beta}$ is reported in Appendix A of the online supplemental material.

Proposition 1. Define

$$\phi(\mathbf{Z}; \mu, \boldsymbol{\gamma}, \boldsymbol{\beta}) = \frac{RY}{\pi(\mathbf{X}; \boldsymbol{\gamma})} - \frac{R - \pi(\mathbf{X}; \boldsymbol{\gamma})}{\pi(\mathbf{X}; \boldsymbol{\gamma})}m(\mathbf{X}, \boldsymbol{\beta}) - \mu \quad (2)$$

and denote $\boldsymbol{\gamma}^* = \text{plim}(\hat{\boldsymbol{\gamma}})$, $\boldsymbol{\beta}^* = \text{plim}(\hat{\boldsymbol{\beta}})$, that is, the probability limits of estimators $\hat{\boldsymbol{\gamma}}$ and $\hat{\boldsymbol{\beta}}$ which equal $\boldsymbol{\gamma}_0$ and $\boldsymbol{\beta}_0$, respectively, when the working models are correctly specified, but not necessarily otherwise. Under suitable regularity conditions (Robins, Rotnitzky, and Zhao 1994, app. B), $\hat{\mu}_{\text{DR}}(\hat{\boldsymbol{\gamma}}, \hat{\boldsymbol{\beta}})$ is asymptotically linear with influence function (Tsiatis 2006)

$$\begin{aligned}\tilde{\phi}(\mathbf{Z};\mu_0,\boldsymbol{\gamma}^*,\boldsymbol{\beta}^*) &= \phi(\mathbf{Z};\mu_0,\boldsymbol{\gamma}^*,\boldsymbol{\beta}^*)\\ &-E\left\{\frac{\partial\phi}{\partial\boldsymbol{\gamma}^T}(\mathbf{Z};\mu_0,\boldsymbol{\gamma}^*,\boldsymbol{\beta}^*)\right\}E^{-1}\left\{\frac{\partial\mathbf{U}_\gamma}{\partial\boldsymbol{\gamma}^T}(\mathbf{Z};\boldsymbol{\gamma}^*)\right\}\mathbf{U}_\gamma(\mathbf{Z};\boldsymbol{\gamma}^*)\\ &-E\left\{\frac{\partial\phi}{\partial\boldsymbol{\beta}^T}(\mathbf{Z};\mu_0,\boldsymbol{\gamma}^*,\boldsymbol{\beta}^*)\right\}E^{-1}\left\{\frac{\partial\mathbf{U}_\beta}{\partial\boldsymbol{\beta}^T}(\mathbf{Z};\boldsymbol{\beta}^*)\right\}\mathbf{U}_\beta(\mathbf{Z};\boldsymbol{\beta}^*);\end{aligned}$$

meaning $\hat{\mu}_{\mathrm{DR}}(\hat{\boldsymbol{\gamma}},\hat{\boldsymbol{\beta}})$ can be expanded as $n^{1/2}\{\hat{\mu}_{\mathrm{DR}}(\hat{\boldsymbol{\gamma}},\hat{\boldsymbol{\beta}})-\mu_0\}=n^{1/2}\tilde{E}_n\{\tilde{\phi}(\mathbf{Z};\mu_0,\boldsymbol{\gamma}^*,\boldsymbol{\beta}^*)\}+o_p(1)$, where $o_p(1)$ denotes a term that converges to zero in probability. Consequently, $n^{1/2}[\hat{\mu}_{\mathrm{DR}}(\hat{\boldsymbol{\gamma}},\hat{\boldsymbol{\beta}})-\mu_0-E\{\phi(\mathbf{Z};\mu_0,\boldsymbol{\gamma}^*,\boldsymbol{\beta}^*)\}]$ converges to a normal limit with variance $\mathrm{var}(\tilde{\phi})$.

With $\pi_{\boldsymbol{\gamma}}(\mathbf{X};\boldsymbol{\gamma})=\partial\pi(\mathbf{X};\boldsymbol{\gamma})/\partial\boldsymbol{\gamma}$ and $m_{\boldsymbol{\beta}}(\mathbf{X};\boldsymbol{\beta})=\partial m(\mathbf{X};\boldsymbol{\beta})/\partial\boldsymbol{\beta}$, we obtain from

$$\begin{aligned}&E\left\{\frac{\partial\phi}{\partial\boldsymbol{\gamma}^T}(\mathbf{Z};\mu_0,\boldsymbol{\gamma}^*,\boldsymbol{\beta}^*)\right\}\\ &\quad=E\left[\frac{\pi_0(\mathbf{X})}{\pi^2(\mathbf{X};\boldsymbol{\gamma}^*)}\left\{m(\mathbf{X};\boldsymbol{\beta}^*)-m_0(\mathbf{X})\right\}\pi_{\boldsymbol{\gamma}}^T(\mathbf{X};\boldsymbol{\gamma}^*)\right],\\ &E\left\{\frac{\partial\phi}{\partial\boldsymbol{\beta}^T}(\mathbf{Z};\mu_0,\boldsymbol{\gamma}^*,\boldsymbol{\beta}^*)\right\}\\ &\quad=E\left[\left\{1-\frac{\pi_0(\mathbf{X})}{\pi(\mathbf{X};\boldsymbol{\gamma}^*)}\right\}m_{\boldsymbol{\beta}}^T(\mathbf{X};\boldsymbol{\beta}^*)\right],\end{aligned}$$

that $E\{\partial\phi(\mathbf{Z};\mu_0,\boldsymbol{\gamma}^*,\boldsymbol{\beta}^*)/\partial\boldsymbol{\gamma}^T\}=\mathbf{0}$ and $E\{\partial\phi(\mathbf{Z};\mu_0,\boldsymbol{\gamma}^*,\boldsymbol{\beta}^*)/\partial\boldsymbol{\beta}^T\}=\mathbf{0}$ at $\mathscr{M}(\boldsymbol{\gamma})\cap\mathscr{M}(\boldsymbol{\beta})$ since then $\pi_0(\mathbf{X})=\pi(\mathbf{X};\boldsymbol{\gamma}^*)$ and $m_0(\mathbf{X})=m(\mathbf{X};\boldsymbol{\beta}^*)$. The influence function of the DR estimator then simply becomes $\tilde{\phi}=\phi$, where ϕ is the influence function of $\hat{\mu}_{\mathrm{DR}}(\boldsymbol{\gamma},\boldsymbol{\beta})$; that is, the DR estimator of μ_0 evaluated at fixed nuisance parameter values $\boldsymbol{\gamma}$ and $\boldsymbol{\beta}$. Thus, under correctly specified working models, the choice of root-n consistent estimators of the nuisance parameters does not affect the asymptotic distribution of the DR estimator. This property, which is more generally satisfied for DR estimators (Robins and Rotnitzky 2001), has stimulated the use of standard methods, such as maximum likelihood, to estimate the nuisance parameters (Bang and Robins 2005).

3. BIASED-REDUCED DOUBLY ROBUST ESTIMATION

The property that the choice of root-n consistent estimators of the nuisance parameters does not affect the first-order asymptotic behavior of a DR estimator, is lost as soon as one of both working models is misspecified. Starting from a given DR estimator, infinitely many DR estimators can therefore typically be constructed by varying the choice of nuisance parameter estimators. This calls for estimation strategies for the nuisance parameters that are optimal according to some criterion. In this article, we propose nuisance parameter estimators such that their probability limits locally minimize the squared first-order asymptotic bias of the DR estimator under misspecification of both working models.

To make the presentation as general as possible (with a slight abuse of notation to simplify), let μ_0 denote the (unknown) population value of the scalar target parameter and $\hat{\mu}_{\mathrm{DR}}(\boldsymbol{\gamma},\boldsymbol{\beta})$ a DR estimator for it, based one finite-dimensional working models $\mathscr{M}(\boldsymbol{\gamma})$ and $\mathscr{M}(\boldsymbol{\beta})$ indexed by parameters $\boldsymbol{\gamma}$ and $\boldsymbol{\beta}$ which take on the values $\boldsymbol{\gamma}_0$ and $\boldsymbol{\beta}_0$ when, respectively, $\mathscr{M}(\boldsymbol{\gamma})$ and $\mathscr{M}(\boldsymbol{\beta})$ hold. The observed data is denoted $\{\mathbf{Z}_i:i=1,\ldots,n\}$. Finally, $\phi(\mathbf{Z};\mu,\boldsymbol{\gamma},\boldsymbol{\beta})$ denotes the influence function of the DR estimator $\hat{\mu}_{\mathrm{DR}}(\boldsymbol{\gamma},\boldsymbol{\beta})$.

3.1 Proposal

Consider possibly misspecified working models $\mathscr{M}(\boldsymbol{\gamma})$ and $\mathscr{M}(\boldsymbol{\beta})$ at fixed known values $\boldsymbol{\gamma}$ and $\boldsymbol{\beta}$, respectively. The first-order asymptotic bias of the DR estimator is then given by $\mathrm{bias}(\boldsymbol{\gamma},\boldsymbol{\beta};\mu_0)=E\{\phi(\mathbf{Z};\mu_0,\boldsymbol{\gamma},\boldsymbol{\beta})\}$. In the missing data problem of Section 2, this is also the total finite-sample bias. By the double robustness, $\mathrm{bias}(\boldsymbol{\gamma}_0,\boldsymbol{\beta};\mu_0)=0$ for any $\boldsymbol{\beta}$ under $\mathscr{M}(\boldsymbol{\gamma})$ and $\mathrm{bias}(\boldsymbol{\gamma},\boldsymbol{\beta}_0;\mu_0)=0$ for any $\boldsymbol{\gamma}$ under $\mathscr{M}(\boldsymbol{\beta})$. This property is lost when both nuisance working models are misspecified. Suppose now that there exists a vector $(\boldsymbol{\gamma}_{\mathrm{BR}}^{*T},\boldsymbol{\beta}_{\mathrm{BR}}^{*T})^T$ such that $E\{\partial\phi(\mathbf{Z};\mu_0,\boldsymbol{\gamma}_{\mathrm{BR}}^*,\boldsymbol{\beta}_{\mathrm{BR}}^*)/\partial\boldsymbol{\gamma}\}=\mathbf{0}$ and $E\{\partial\phi(\mathbf{Z};\mu_0,\boldsymbol{\gamma}_{\mathrm{BR}}^*,\boldsymbol{\beta}_{\mathrm{BR}}^*)/\partial\boldsymbol{\beta}\}=\mathbf{0}$, with BR an abbreviation for bias-reduced. The following theorem then shows that $(\boldsymbol{\gamma}_{\mathrm{BR}}^{*T},\boldsymbol{\beta}_{\mathrm{BR}}^{*T})^T$ locally minimizes the squared first-order asymptotic bias in the direction of $\boldsymbol{\gamma}$ and $\boldsymbol{\beta}$, explaining the subscript BR.

Theorem 1. Under suitable regularity conditions (see Appendix C of the online supplemental material), $(\boldsymbol{\gamma}_{\mathrm{BR}}^{*T},\boldsymbol{\beta}_{\mathrm{BR}}^{*T})^T$ locally minimizes the squared first-order bias $\mathrm{bias}^2(\boldsymbol{\gamma},\boldsymbol{\beta};\mu_0)$, with $(\boldsymbol{\gamma}_{\mathrm{BR}}^{*T},\boldsymbol{\beta}_{\mathrm{BR}}^{*T})^T$ a solution to $E\{\partial\phi(\mathbf{Z};\mu_0,\boldsymbol{\gamma}_{\mathrm{BR}}^*,\boldsymbol{\beta}_{\mathrm{BR}}^*)/\partial\boldsymbol{\gamma}\}=\mathbf{0}$ and $E\{\partial\phi(\mathbf{Z};\mu_0,\boldsymbol{\gamma}_{\mathrm{BR}}^*,\boldsymbol{\beta}_{\mathrm{BR}}^*)/\partial\boldsymbol{\beta}\}=\mathbf{0}$.

Proof. Under regularity conditions that allow us to interchange integration and differentiation,

$$\begin{aligned}\frac{\partial}{\partial\boldsymbol{\gamma}}\mathrm{bias}^2(\boldsymbol{\gamma},\boldsymbol{\beta};\mu_0)&=2\mathrm{bias}(\boldsymbol{\gamma},\boldsymbol{\beta};\mu_0)\frac{\partial}{\partial\boldsymbol{\gamma}}\mathrm{bias}(\boldsymbol{\gamma},\boldsymbol{\beta};\mu_0)\\ &=2\mathrm{bias}(\boldsymbol{\gamma},\boldsymbol{\beta};\mu_0)E\left\{\frac{\partial\phi}{\partial\boldsymbol{\gamma}}(\mathbf{Z};\mu_0,\boldsymbol{\gamma},\boldsymbol{\beta})\right\}\end{aligned}$$

and likewise for $\boldsymbol{\beta}$. The result follows since by definition the values $\boldsymbol{\gamma}_{\mathrm{BR}}^*$ and $\boldsymbol{\beta}_{\mathrm{BR}}^*$ satisfy the equations $E\{\partial\phi(\mathbf{Z};\mu_0,\boldsymbol{\gamma}_{\mathrm{BR}}^*,\boldsymbol{\beta}_{\mathrm{BR}}^*)/\partial\boldsymbol{\gamma}\}=\mathbf{0}$ and $E\{\partial\phi(\mathbf{Z};\mu_0,\boldsymbol{\gamma}_{\mathrm{BR}}^*,\boldsymbol{\beta}_{\mathrm{BR}}^*)/\partial\boldsymbol{\beta}\}=\mathbf{0}$. □

In practice, the values $(\boldsymbol{\gamma}_{\mathrm{BR}}^{*T},\boldsymbol{\beta}_{\mathrm{BR}}^{*T})^T$ that solve the mean gradients $E\{\partial\phi(\mathbf{Z};\mu_0,\boldsymbol{\gamma}_{\mathrm{BR}}^*,\boldsymbol{\beta}_{\mathrm{BR}}^*)/\partial\boldsymbol{\gamma}\}=\mathbf{0}$ and $E\{\partial\phi(\mathbf{Z};\mu_0,\boldsymbol{\gamma}_{\mathrm{BR}}^*,\boldsymbol{\beta}_{\mathrm{BR}}^*)/\partial\boldsymbol{\beta}\}=\mathbf{0}$ are unknown and need to be estimated. Therefore, define the estimators $\hat{\boldsymbol{\gamma}}_{\mathrm{BR}}$ and $\hat{\boldsymbol{\beta}}_{\mathrm{BR}}$ as the solutions to the estimating equations

$$\tilde{E}_n\{\partial\phi(\mathbf{Z};\mu_0,\hat{\boldsymbol{\gamma}}_{\mathrm{BR}},\hat{\boldsymbol{\beta}}_{\mathrm{BR}})/\partial\boldsymbol{\beta}\}=\mathbf{0}\tag{3}$$

$$\tilde{E}_n\{\partial\phi(\mathbf{Z};\mu_0,\hat{\boldsymbol{\gamma}}_{\mathrm{BR}},\hat{\boldsymbol{\beta}}_{\mathrm{BR}})/\partial\boldsymbol{\gamma}\}=\mathbf{0}.\tag{4}$$

When the gradient of $\phi(\mathbf{Z};\mu_0,\boldsymbol{\gamma},\boldsymbol{\beta})$ with respect to $\boldsymbol{\gamma}$ or $\boldsymbol{\beta}$ depends on the unknown population value μ_0, a preliminary consistent DR estimator $\hat{\mu}_{\mathrm{DR}}^{\mathrm{prel}}$ is substituted for μ_0 (e.g., the default DR estimator based on MLEs for the nuisance parameters). The following theorem shows that (3) and (4) yield consistent estimators $\hat{\boldsymbol{\gamma}}_{\mathrm{BR}}$ and $\hat{\boldsymbol{\beta}}_{\mathrm{BR}}$ for $\boldsymbol{\gamma}_0$ and $\boldsymbol{\beta}_0$ under $\mathscr{M}(\boldsymbol{\gamma})$ and $\mathscr{M}(\boldsymbol{\beta})$, respectively. The resulting DR estimator $\hat{\mu}_{\mathrm{DR}}(\hat{\boldsymbol{\gamma}}_{\mathrm{BR}},\hat{\boldsymbol{\beta}}_{\mathrm{BR}})$ is referred to as the bias-reduced DR estimator.

Theorem 2. Under suitable regularity conditions (see Appendix C of the online supplemental material), $\hat{\boldsymbol{\gamma}}_{\mathrm{BR}}$ is a consistent estimator for $\boldsymbol{\gamma}_0$ under $\mathscr{M}(\boldsymbol{\gamma})$ and $\hat{\boldsymbol{\beta}}_{\mathrm{BR}}$ is a consistent estimator for $\boldsymbol{\beta}_0$ under $\mathscr{M}(\boldsymbol{\beta})$.

Proof. We give the proof for $\hat{\boldsymbol{\gamma}}_{\mathrm{BR}}$. Under model $\mathscr{M}(\boldsymbol{\gamma})$, the (unknown) population value $\boldsymbol{\gamma}_0$ is well defined. By the double robustness of the estimator $\hat{\mu}_{\mathrm{DR}}(\boldsymbol{\gamma}, \boldsymbol{\beta})$, $\phi(\mathbf{Z}; \mu_0, \boldsymbol{\gamma}_0, \boldsymbol{\beta})$ has mean zero for all $\boldsymbol{\beta}$. Consequently, with $F_0(\mathrm{z})$ denoting the true (unknown) joint distribution function of $\mathbf{Z}$,

$$\mathbf{0} = \frac{\partial}{\partial \boldsymbol{\beta}} E\{\phi(\mathbf{Z}; \mu_0, \boldsymbol{\gamma}_0, \boldsymbol{\beta})\} = \int \frac{\partial}{\partial \boldsymbol{\beta}} \phi(\mathrm{z}; \mu_0, \boldsymbol{\gamma}_0, \boldsymbol{\beta}) \mathrm{d}F_0(\mathrm{z})$$
$$= E\left\{\frac{\partial}{\partial \boldsymbol{\beta}} \phi(\mathbf{Z}; \mu_0, \boldsymbol{\gamma}_0, \boldsymbol{\beta})\right\}$$

for all $\boldsymbol{\beta}$ assuming we can interchange integration and differentiation (see Appendix C of the online supplemental material). Note that we do not take derivatives of $F_0(\mathbf{Z})$ with respect to $\boldsymbol{\beta}$, since the expectation is taken under the true data-generating mechanism, which stays fixed as $\boldsymbol{\beta}$ varies. Hence, at $\mathscr{M}(\boldsymbol{\gamma})$, the gradient $\partial\phi(\mathbf{Z}; \mu_0, \boldsymbol{\gamma}_0, \boldsymbol{\beta})/\partial\boldsymbol{\beta}$ yields an unbiased estimating function for $\boldsymbol{\gamma}$. Under suitable regularity conditions (Robins, Rotnitzky, and Zhao 1994, app. B), it now follows from the uniform WLLN (Newey and McFadden 1994, Lemma 4.3) and the fact that $\mathrm{plim}(\hat{\mu}_{\mathrm{DR}}^{\mathrm{prel}}) = \mu_0$ under $\mathscr{M}(\boldsymbol{\gamma})$ that $\mathbf{0} = \mathrm{plim}[\tilde{E}_n\{\partial\phi(\mathbf{Z}; \hat{\mu}_{\mathrm{DR}}^{\mathrm{prel}}, \hat{\boldsymbol{\gamma}}_{\mathrm{BR}}, \hat{\boldsymbol{\beta}}_{\mathrm{BR}})/\partial\boldsymbol{\beta}\}] = E\{\partial\phi(\mathbf{Z}; \mu_0, \boldsymbol{\gamma}_{\mathrm{BR}}^*, \boldsymbol{\beta}_{\mathrm{BR}}^*)/\partial\boldsymbol{\beta}\}$ with probability limits $\boldsymbol{\gamma}_{\mathrm{BR}}^* = \mathrm{plim}(\hat{\boldsymbol{\gamma}}_{\mathrm{BR}})$ and $\boldsymbol{\beta}_{\mathrm{BR}}^* = \mathrm{plim}(\hat{\boldsymbol{\beta}}_{\mathrm{BR}})$ so that $\boldsymbol{\gamma}_0 = \boldsymbol{\gamma}_{\mathrm{BR}}^*$ and thus $\hat{\boldsymbol{\gamma}}_{\mathrm{BR}}$ is a consistent estimator of $\boldsymbol{\gamma}_0$. □

Theorem 2 shows that $\boldsymbol{\gamma}_{\mathrm{BR}}^* = \boldsymbol{\gamma}_0$ under $\mathscr{M}(\boldsymbol{\gamma})$ and $\boldsymbol{\beta}_{\mathrm{BR}}^* = \boldsymbol{\beta}_0$ under $\mathscr{M}(\boldsymbol{\beta})$. This is not surprising because $\min_{(\boldsymbol{\gamma}^T, \boldsymbol{\beta}^T)^T} \left\{\mathrm{bias}^2(\boldsymbol{\gamma}, \boldsymbol{\beta}; \mu_0)\right\} = 0$ under $\mathscr{M}(\boldsymbol{\gamma}) \cup \mathscr{M}(\boldsymbol{\beta})$. Furthermore, $\mathrm{bias}(\hat{\boldsymbol{\gamma}}_{\mathrm{BR}}, \hat{\boldsymbol{\beta}}_{\mathrm{BR}}; \mu_0) = \mathrm{bias}(\boldsymbol{\gamma}_{\mathrm{BR}}^*, \boldsymbol{\beta}_{\mathrm{BR}}^*; \mu_0) + o(1)$ (see Appendix B of the online supplemental material), and hence, the squared first-order asymptotic bias is also locally minimized when the fixed values $({\boldsymbol{\gamma}_{\mathrm{BR}}^*}^T, {\boldsymbol{\beta}_{\mathrm{BR}}^*}^T)^T$ are replaced by root-n consistent estimators $(\hat{\boldsymbol{\gamma}}_{\mathrm{BR}}^T, \hat{\boldsymbol{\beta}}_{\mathrm{BR}}^T)^T$. However, the bias optimality promised by Theorem 1 may become somewhat illusory when the estimating Equations (3) or (4) depend on the population value μ_0. The reason is that in this case the values $({\boldsymbol{\gamma}_{\mathrm{BR}}^*}^T, {\boldsymbol{\beta}_{\mathrm{BR}}^*}^T)^T$ no longer minimize $\mathrm{bias}^2(\boldsymbol{\gamma}, \boldsymbol{\beta}; \mu_0) = E\{\phi(\mathbf{Z}; \mu_0, \boldsymbol{\gamma}, \boldsymbol{\beta})\}^2$ but instead minimize $\mathrm{bias}^2(\boldsymbol{\gamma}, \boldsymbol{\beta}; \mu^*) = E\{\phi(\mathbf{Z}; \mu^*, \boldsymbol{\gamma}, \boldsymbol{\beta})\}^2$ with $\mu^* = \mathrm{plim}(\hat{\mu}_{\mathrm{DR}}^{\mathrm{prel}})$ which may differ from μ_0 under misspecification of both nuisance working models. The bias optimality of Theorem 1 is therefore limited to DR estimators for which the left-hand sides of (3) and (4) do not depend on the target parameter. Fortunately, many target parameters for which DR estimators exist satisfy this property. For instance, the class given in Robins et al. (2008, sec. 3.1) satisfies this; in particular, this is satisfied for the missing data example in Section 2 (see Equation (2)). When the left-hand sides of (3) and (4) do depend on the target parameter, then the bias optimality of Theorem 1 remains useful for score tests of the null hypothesis that $\mu = \tilde{\mu}$ for some $\tilde{\mu}$; that is, tests of $E\{\phi(\mathbf{Z}; \tilde{\mu}, \boldsymbol{\gamma}_0, \boldsymbol{\beta}_0)\} = 0$. When μ_0 is substituted by $\tilde{\mu}$ in (3) and (4), the resulting (probability limits of the) estimators $\hat{\boldsymbol{\gamma}}_{\mathrm{BR}}$ and $\hat{\boldsymbol{\beta}}_{\mathrm{BR}}$ continue to minimize $E^2\{\phi(\mathbf{Z}; \tilde{\mu}, \boldsymbol{\gamma}, \boldsymbol{\beta})\}$.

A limitation of the proposed approach is that it demands working models of the same dimension because the gradient of ϕ with respect to $\boldsymbol{\beta}$ is used as an estimating function for $\boldsymbol{\gamma}$ and vice versa. Restrictions on the dimension of the working models are also seen in other proposals (Rotnitzky et al. 2012). This can be remedied by enlarging the working models with clever choices of covariates until they reach the same dimension (see Section 7).

Remark 1. The validity of the proposal is predicated on the availability of a DR estimator; it cannot be used for arbitrary estimators. Indeed, reconsider the missing data problem of Section 2 with $\hat{\mu}_{\mathrm{OR}}(\boldsymbol{\beta}) = \tilde{E}_n\{m(\mathbf{X}; \boldsymbol{\beta})\}$ for $m(\mathbf{X}; \boldsymbol{\beta}) = \boldsymbol{\beta}^T(1, \mathbf{X}^T)^T$ an estimator of μ_0. For fixed $\boldsymbol{\beta}$, the influence function of $\hat{\mu}_{\mathrm{OR}}(\boldsymbol{\beta})$ is $\phi_{\mathrm{OR}}(\mathbf{Z}; \mu_0, \boldsymbol{\beta}) = m(\mathbf{X}; \boldsymbol{\beta}) - \mu_0$ and the squared bias is $E\{m(\mathbf{X}; \boldsymbol{\beta}) - \mu_0\}^2$. In this case, $E\{\partial\phi_{\mathrm{OR}}(\mathbf{Z}; \mu_0, \boldsymbol{\beta})/\partial\boldsymbol{\beta}\} = E\{(1, \mathbf{X}^T)^T\}$ does not depend on $\boldsymbol{\beta}$ and hence the gradient $\partial\phi_{\mathrm{OR}}(\mathbf{Z}; \mu_0, \boldsymbol{\beta})/\partial\boldsymbol{\beta}$ does not provide an unbiased estimating function.

3.2 Further Properties

It follows from Proposition 1 (and Appendix A of the online supplemental material) that at $\mathscr{M}(\boldsymbol{\gamma}) \cap \mathscr{M}(\boldsymbol{\beta})$, small perturbations (of the order one over root-n) in $\boldsymbol{\gamma}$ and $\boldsymbol{\beta}$ do not affect the first-order asymptotic behavior of the DR estimator in the sense that $E\{\partial\phi(\mathbf{Z}; \mu_0, \boldsymbol{\gamma}^*, \boldsymbol{\beta})/\partial\boldsymbol{\beta}\} = \mathbf{0}$ for all $\boldsymbol{\beta}$ and $E\{\partial\phi(\mathbf{Z}; \mu_0, \boldsymbol{\gamma}, \boldsymbol{\beta}^*)/\partial\boldsymbol{\gamma}\} = \mathbf{0}$ for all $\boldsymbol{\gamma}$. This local robustness is lost as soon as one of the working models is misspecified. The estimators $\hat{\boldsymbol{\gamma}}_{\mathrm{BR}}$ and $\hat{\boldsymbol{\beta}}_{\mathrm{BR}}$ for the nuisance parameters are designed to restore this local robustness property under model misspecification. It is hence not entirely surprising that (the probability limits of) these estimators locally minimize $\mathrm{bias}^2(\boldsymbol{\gamma}, \boldsymbol{\beta}; \mu_0)$. Specifically, they ensure that the bias-reduced DR estimator $\hat{\mu}_{\mathrm{DR}}(\hat{\boldsymbol{\gamma}}_{\mathrm{BR}}, \hat{\boldsymbol{\beta}}_{\mathrm{BR}})$ is first-order ancillary (Cox 1980) under misspecification of the working models in the sense formalized in the following corollary.

Corollary 1. Under suitable regularity conditions (Robins, Rotnitzky, and Zhao 1994, app. B), $n^{1/2}\{\hat{\mu}_{\mathrm{DR}}(\hat{\boldsymbol{\gamma}}_{\mathrm{BR}}, \hat{\boldsymbol{\beta}}_{\mathrm{BR}}) - \mu_0\} = n^{1/2}\tilde{E}_n\{\phi(\mathbf{Z}; \mu_0, \boldsymbol{\gamma}_{\mathrm{BR}}^*, \boldsymbol{\beta}_{\mathrm{BR}}^*)\} + o_p(1)$ when $\hat{\boldsymbol{\gamma}}_{\mathrm{BR}}$ and $\hat{\boldsymbol{\beta}}_{\mathrm{BR}}$ are the solutions to (3) and (4) with $\boldsymbol{\gamma}_{\mathrm{BR}}^* = \mathrm{plim}(\hat{\boldsymbol{\gamma}}_{\mathrm{BR}})$ and $\boldsymbol{\beta}_{\mathrm{BR}}^* = \mathrm{plim}(\hat{\boldsymbol{\beta}}_{\mathrm{BR}})$.

Proof. This follows from the proof of Proposition 1 because (under standard regularity conditions) $(\hat{\boldsymbol{\gamma}}_{\mathrm{BR}} - \boldsymbol{\gamma}_{\mathrm{BR}}^*)$ and $(\hat{\boldsymbol{\beta}}_{\mathrm{BR}} - \boldsymbol{\beta}_{\mathrm{BR}}^*)$ are $O_p(n^{-1/2})$ (see also Theorem 3.13 in Robins et al. (2008)). □

This first-order ancillarity implies that the first-order asymptotic behavior of $\hat{\mu}_{\mathrm{DR}}(\hat{\boldsymbol{\gamma}}_{\mathrm{BR}}, \hat{\boldsymbol{\beta}}_{\mathrm{BR}})$ is the same as that of $\hat{\mu}_{\mathrm{DR}}(\boldsymbol{\gamma}_{\mathrm{BR}}^*, \boldsymbol{\beta}_{\mathrm{BR}}^*)$, in which $\hat{\boldsymbol{\gamma}}_{\mathrm{BR}}$ and $\hat{\boldsymbol{\beta}}_{\mathrm{BR}}$ are substituted by their probability limits $\boldsymbol{\gamma}_{\mathrm{BR}}^*$ and $\boldsymbol{\beta}_{\mathrm{BR}}^*$, respectively. This has a number of important consequences. First, the asymptotic variance of the DR estimator $\hat{\mu}_{\mathrm{DR}}(\hat{\boldsymbol{\gamma}}_{\mathrm{BR}}, \hat{\boldsymbol{\beta}}_{\mathrm{BR}})$ can be straightforwardly estimated as one over n times the sample variance of the values $\phi\{\mathbf{Z}_i; \hat{\mu}_{\mathrm{DR}}(\hat{\boldsymbol{\gamma}}_{\mathrm{BR}}, \hat{\boldsymbol{\beta}}_{\mathrm{BR}}), \hat{\boldsymbol{\gamma}}_{\mathrm{BR}}, \hat{\boldsymbol{\beta}}_{\mathrm{BR}}\}$, without having to correct for the estimation of the nuisance parameters. Similarly, a score test of the null hypothesis that $\mu = \tilde{\mu}$ for some $\tilde{\mu}$ simplifies to a one-sample t-test of the null hypothesis that $\phi(\mathbf{Z}; \tilde{\mu}, \boldsymbol{\gamma}_0, \boldsymbol{\beta}_0)$ has mean zero. Second, the estimators $\hat{\boldsymbol{\gamma}}_{\mathrm{BR}}$ and $\hat{\boldsymbol{\beta}}_{\mathrm{BR}}$ tend to deliver reasonably efficient DR estimators as will be confirmed in simulation studies in Section 4. This can be intuitively expected because an estimator tends to be less variable when evaluated at fixed nuisance parameter values instead of estimated ones. However, an efficiency benefit relative to the use of maximum likelihood estimation of the nuisance parameters is not theoretically guaranteed. One reason for this is that under model misspecification, different estimators $\hat{\boldsymbol{\gamma}}$ and $\hat{\boldsymbol{\beta}}$ of the nuisance parameters may converge to different probability limits $\boldsymbol{\gamma}^*$ and $\boldsymbol{\beta}^*$ and thereby influence the variance of $\phi(\mathbf{Z}; \mu^*, \boldsymbol{\gamma}^*, \boldsymbol{\beta}^*)$ with μ^* the corresponding probability limit of the DR estimator under model misspecification. A second reason is that an estimator may sometimes vary less when evaluated at estimated rather than known nuisance parameters (Rotnitzky, Li, and Li 2010).

Remark 2. James Robins, Andrea Rotnitzky, Eric Tchetgen Tchetgen, and a referee noted that Robins et al. (2008) use similar estimating equations like (3) and (4) in an intermediate simplifying step in the construction of higher order influence functions, but with a different objective. In their approach, these estimating equations are not used to directly estimate nuisance parameters describing parametric working models. Instead they first obtain (potentially highly data-adaptive) initial estimators of these working models and then use estimating equations similar to (3) and (4) to estimate nuisance parameters describing specific linear extensions of these initial estimators (where the dimension can increase with the sample size); in contrast, we allow for arbitrary but finite-dimensional nuisance working models. As a result, first-order ancillarity (see our Corollary 1 and Theorem 3.13 in Robins et al. (2008)) with respect to their fluctuation parameters is obtained, which simplifies the derivation of higher order influence functions. Our Theorem 2 is also similar to their Lemma 3 but their's allows for infinite-dimensional nuisance parameters.

3.3 Illustration: Missing Data Problem

To illustrate the bias-reduced DR estimation strategy, we return to the missing data problem introduced in Section 2. From the influence function (2) of the DR estimator, it follows that (3) equals $\tilde{E}_n[\{1 - R/\pi(\mathbf{X}, \hat{\boldsymbol{\gamma}}_{\mathrm{BR}})\} m_{\boldsymbol{\beta}}(\mathbf{X}; \hat{\boldsymbol{\beta}}_{\mathrm{BR}})] = \mathbf{0}$. For instance, for $m(\mathbf{X}; \boldsymbol{\beta}) = \beta_1 + \boldsymbol{\beta}_2^T \mathbf{X}$, this becomes

$$\tilde{E}_n\left[\left\{1 - \frac{R}{\pi(\mathbf{X}, \hat{\boldsymbol{\gamma}}_{\mathrm{BR}})}\right\} \left(1, \mathbf{X}^T\right)^T\right] = \mathbf{0}. \tag{5}$$

The first equation, $1 = \tilde{E}_n\{R/\pi(\mathbf{X}; \hat{\boldsymbol{\gamma}}_{\mathrm{BR}})\}$, ensures that the inverse weights sum to the sample size; $\hat{\mu}_{\mathrm{DR}}(\hat{\boldsymbol{\gamma}}_{\mathrm{BR}}, \hat{\boldsymbol{\beta}}_{\mathrm{BR}})$ then equals $\tilde{E}_n\{m(\mathbf{X}; \hat{\boldsymbol{\beta}}_{\mathrm{BR}})\} + \tilde{E}_n[R/\pi(\mathbf{X}; \hat{\boldsymbol{\gamma}}_{\mathrm{BR}})\{Y - m(\mathbf{X}; \hat{\boldsymbol{\beta}}_{\mathrm{BR}})\}]/\tilde{E}_n\{R/\pi(\mathbf{X}; \hat{\boldsymbol{\gamma}}_{\mathrm{BR}})\}$, also considered in Robins et al. (2007). The remaining equations in (5) impose that the sample mean of the covariates, $\tilde{E}_n(\mathbf{X})$, equals the weighted sample mean $\tilde{E}_n\{R\mathbf{X}/\pi(\mathbf{X}; \hat{\boldsymbol{\gamma}}_{\mathrm{BR}})\}$. These restrictions help to ensure stable weights; the bias-reduced DR estimation strategy might therefore alleviate the problem of inefficiency due to highly variable weights (Robins et al. 2007). Restrictions (5) are also known as calibration equations in the survey sampling literature where they are used to improve the simple Horvitz-Thompson estimator by making it unbiased under a linear prediction model (Kott and Liao 2012). For linear outcome models (or more generally whenever $m(\mathbf{X}; \boldsymbol{\beta})$ lies within the span of the gradient $m_{\boldsymbol{\beta}}(\mathbf{X}, \boldsymbol{\beta})$), it now follows from (5) that $\hat{\mu}_{\mathrm{DR}}(\hat{\boldsymbol{\gamma}}_{\mathrm{BR}}, \hat{\boldsymbol{\beta}}_{\mathrm{BR}}) = \tilde{E}_n\{RY/\pi(\mathbf{X}; \hat{\boldsymbol{\gamma}}_{\mathrm{BR}})\} = \tilde{E}_n\{RY/\pi(\mathbf{X}; \hat{\boldsymbol{\gamma}}_{\mathrm{BR}})\}/\tilde{E}_n\{R/\pi(\mathbf{X}; \hat{\boldsymbol{\gamma}}_{\mathrm{BR}})\}$. This demonstrates that the bias-reduced DR estimator remarkably reduces to a simple IPTW estimator, making $\hat{\mu}_{\mathrm{DR}}(\hat{\boldsymbol{\gamma}}_{\mathrm{BR}}, \hat{\boldsymbol{\beta}}_{\mathrm{BR}})$ sample bounded (Robins et al. 2007, sec. 4.1) in the sense that it lies with probability one within the admissible range of observed outcome values (i.e., $[Y_{\min}, Y_{\max}]$, where $Y_{\min} = \min\{Y_i : R_i = 1\}$ and $Y_{\max} = \max\{Y_i : R_i = 1\}$) whenever the outcome Y is continuous with conditional mean linear in $\mathbf{X}$ (or $m(\mathbf{X}; \boldsymbol{\beta})$ lies within the span of the gradient $m_{\boldsymbol{\beta}}(\mathbf{X}, \boldsymbol{\beta})$).

From (2), it follows that the estimating Equation (4) for $\boldsymbol{\beta}$ equals $\tilde{E}_n[R\pi_{\boldsymbol{\gamma}}(\mathbf{X}; \hat{\boldsymbol{\gamma}}_{\mathrm{BR}})/\pi^2(\mathbf{X}; \hat{\boldsymbol{\gamma}}_{\mathrm{BR}})\{Y - m(\mathbf{X}; \hat{\boldsymbol{\beta}}_{\mathrm{BR}})\}] = \mathbf{0}$. For instance, when $\pi(\mathbf{X}; \boldsymbol{\gamma}) = \mathrm{expit}(\gamma_1 + \boldsymbol{\gamma}_2^T \mathbf{X})$, this becomes

$$\tilde{E}_n\left[R\left\{Y - m(\mathbf{X}; \hat{\boldsymbol{\beta}}_{\mathrm{BR}})\right\} \frac{1 - \pi(\mathbf{X}; \hat{\boldsymbol{\gamma}}_{\mathrm{BR}})}{\pi(\mathbf{X}; \hat{\boldsymbol{\gamma}}_{\mathrm{BR}})} \left(1, \mathbf{X}^T\right)^T\right] = \mathbf{0}, \tag{6}$$

which amounts to weighted least squares based on the complete cases with weights $\{1 - \pi(\mathbf{X}_i; \hat{\boldsymbol{\gamma}}_{\mathrm{BR}})\} \times \pi^{-1}(\mathbf{X}_i; \hat{\boldsymbol{\gamma}}_{\mathrm{BR}})$. High (low) weights are thus given to covariate regions with low (high) probability of observed data, thereby forcing the model to fit well in regions with most missing data. The bias-reduced DR estimator can now be equivalently written as a mean imputation estimator $\hat{\mu}_{\mathrm{DR}}(\hat{\boldsymbol{\gamma}}_{\mathrm{BR}}, \hat{\boldsymbol{\beta}}_{\mathrm{BR}}) = \tilde{E}_n\{RY + (1 - R)m(\mathbf{X}; \hat{\boldsymbol{\beta}}_{\mathrm{BR}})\}$, which averages the observed outcome for responders and a predicted outcome for nonresponders. This is desirable as it ensures that the aforementioned boundedness property is also effective whenever the outcome predictions obey the admissible range of the data. For instance, when Y is binary and $m(\mathbf{X}; \boldsymbol{\beta}) = \mathrm{expit}(\beta_1 + \boldsymbol{\beta}_2^T \mathbf{X})$, $m(\mathbf{X}; \hat{\boldsymbol{\beta}}_{\mathrm{BR}})$ falls between 0 and 1 so that $\hat{\mu}_{\mathrm{DR}}(\hat{\boldsymbol{\gamma}}_{\mathrm{BR}}, \hat{\boldsymbol{\beta}}_{\mathrm{BR}})$ is guaranteed to lie between 0 and 1.

While Equation (6) for $\boldsymbol{\beta}$ can be solved using weighted regression in standard statistical software, this is not so for Equation (5) for $\boldsymbol{\gamma}$. To accommodate this, arguing along the lines of Tan (2010), we define the function $\mathscr{F}(\boldsymbol{\gamma}) = \tilde{E}_n[-R\exp\{-\boldsymbol{\gamma}^T(1, \mathbf{X}^T)^T\} - (1 - R)\boldsymbol{\gamma}^T(1, \mathbf{X}^T)^T]$ which is an integrated form of (5) in the sense $\partial\mathscr{F}(\boldsymbol{\gamma})/\partial\boldsymbol{\gamma}$ equals (5). The function $\mathscr{F}(\boldsymbol{\gamma})$ is always concave and bounded on every bounded set for $\boldsymbol{\gamma}$. In Appendix D of the online supplemental material, a condition under which $\mathscr{F}(\boldsymbol{\gamma})$ has a unique maximum is provided and in Appendix J of the supplemental material we provide an R function to obtain $\hat{\mu}_{\mathrm{DR}}(\hat{\boldsymbol{\gamma}}_{\mathrm{BR}}, \hat{\boldsymbol{\beta}}_{\mathrm{BR}})$.

The bias reduction promised by Theorem 1 may be substantial, as we argue next. Let $\delta(\mathbf{X}; \boldsymbol{\gamma}^*) = \pi(\mathbf{X}; \boldsymbol{\gamma}^*) - \pi_0(\mathbf{X})$ denote the degree of model misspecification in the model for the PS at given $\mathbf{X}$ and $\Delta(\mathbf{X}; \boldsymbol{\beta}^*) = m(\mathbf{X}; \boldsymbol{\beta}^*) - m_0(\mathbf{X})$ the degree of model misspecification in the working model for the conditional mean outcome at given $\mathbf{X}$. When both working models are misspecified, the asymptotic bias of the DR estimator can be written as (see e.g., Vansteelandt, Bekaert, and Claeskens 2012)

$$\mathrm{bias}(\boldsymbol{\gamma}^*, \boldsymbol{\beta}^*; \mu_0) = E\left[\frac{\delta(\mathbf{X}; \boldsymbol{\gamma}^*)\Delta(\mathbf{X}; \boldsymbol{\beta}^*)}{\pi(\mathbf{X}; \boldsymbol{\gamma}^*)}\right]. \tag{7}$$

It is thus driven by the degree of misspecification $\delta(\mathbf{X}; \boldsymbol{\gamma}^*)$ and $\Delta(\mathbf{X}; \boldsymbol{\beta}^*)$ but may get inflated in regions with small PS. This inflation is a legitimate concern because in these regions with low PS, the probability of observing Y is low and misspecification in $m(\mathbf{X}; \boldsymbol{\beta})$ is most likely. Bias-reduced DR estimation prevents such inflation. For instance, using the first component of the vector of estimating equations in (6), we obtain that $E[\pi_0(\mathbf{X})\Delta(\mathbf{X}; \boldsymbol{\beta}^*_{\mathrm{BR}})/\pi(\mathbf{X}; \boldsymbol{\gamma}^*_{\mathrm{BR}})] = E\{\Delta(\mathbf{X}; \boldsymbol{\beta}^*_{\mathrm{BR}})\pi_0(\mathbf{X})\}$. This is so whenever the logistic regression model for the PS includes an intercept. The asymptotic bias (7) can then be equivalently written as $E[\Delta(\mathbf{X};\boldsymbol{\beta}^*_{\mathrm{BR}})\{1-\pi_0(\mathbf{X})\}]$, and hence does not get severely inflated in covariate regions with small PS.

Remark 3. Like our proposal, a recent proposal by van der Laan (2014) also focuses on bias reduction. This proposal is different in that it is based on removing an approximation to the first-order bias of the DR estimator by cleverly fitting highly data-adaptive working models. In view of simplicity, the mainstream use of parametric models and the difficulty of obtaining good approximations to the bias, our proposal avoids such approximations and focuses on bias reduction under misspecification of (both) parametric working models.

4. SIMULATION STUDIES

We carried out different simulation studies to compare the performance of the bias-reduced DR estimator $\hat{\mu}_{\mathrm{BR}} \equiv \hat{\mu}_{\mathrm{DR}}(\hat{\boldsymbol{\gamma}}_{\mathrm{BR}}, \hat{\boldsymbol{\beta}}_{\mathrm{BR}})$ with several alternatives for the estimation of a mean outcome in the presence of incomplete data. Nuisance parameters estimated via standard MLE are denoted $\hat{\boldsymbol{\gamma}}_{\mathrm{MLE}}$ and $\hat{\boldsymbol{\beta}}_{\mathrm{MLE}}$. We consider standard estimators $\hat{\mu}_{\mathrm{IPTW}} = \tilde{E}_n\{RY/\pi(\mathbf{X};\hat{\boldsymbol{\gamma}}_{\mathrm{MLE}})\}$, $\hat{\mu}_{\mathrm{OR}} = \tilde{E}_n\{m(\mathbf{X};\hat{\boldsymbol{\beta}}_{\mathrm{MLE}})\}$, and the DR estimator $\hat{\mu}_{\mathrm{MLE}} \equiv \hat{\mu}_{\mathrm{DR}}(\hat{\boldsymbol{\gamma}}_{\mathrm{MLE}}, \hat{\boldsymbol{\beta}}_{\mathrm{MLE}})$. Next we consider the DR estimator $\hat{\mu}_{\mathrm{GBM}}$ which uses a PS estimated under a generalized boosted model (GBM), a multivariate nonparametric regression technique (McCaffrey, Ridgeway, and Morral 2004). We consider the calibrated likelihood estimator $\hat{\mu}_{\mathrm{TAN}}$ of Tan (2010) based on a nonparametric likelihood, and the DR estimator $\hat{\mu}_{\mathrm{PROJ}} \equiv \hat{\mu}_{\mathrm{DR}}(\hat{\boldsymbol{\gamma}}_{\mathrm{MLE}}, \hat{\boldsymbol{\beta}}_{\mathrm{PROJ}})$ of Cao, Tsiatis, and Davidian (2009), which uses an estimator $\hat{\boldsymbol{\beta}}_{\mathrm{PROJ}}$ that minimizes the estimated asymptotic variance of the DR estimator under the assumption of a correctly specified PS model. Finally, we consider two TMLEs $\hat{\mu}_{\mathrm{TMLE}}$ and $\hat{\mu}_{\mathrm{TMLE\text{-}SL}}$ (van der Laan and Rose 2011) based on quasi-log-likelihood loss functions, where $\hat{\mu}_{\mathrm{TMLE}}$ uses ordinary least squares as an initial estimate for the conditional mean outcome, whereas $\hat{\mu}_{\mathrm{TMLE\text{-}SL}}$ uses a super learner (van der Laan and Rose 2011, chap. 3) based on a library consisting of generalized additive and linear models, random forests and adaptive polynomial splines. For each scenario, we perform 1000 Monte Carlo runs at sample sizes of $n = 200$ and 1000. For each estimator, we calculated the Monte Carlo bias (BIAS), the root mean square error (RMSE), the median of absolute errors (MAE) and the Monte Carlo standard deviation (MCSD). Occasionally, no convergence was attained for the estimator $\hat{\boldsymbol{\gamma}}_{\mathrm{BR}}$ at $n = 200$, as indicated below each table.

4.1 Scenario 1

The first simulation scenario considers a simple data-generating mechanism where for each i $(i = 1, \ldots, n)$, $X_i \overset{d}{=} N(0, 1)$, $R_i|X_i \overset{d}{=} \mathrm{Ber}\{\pi_0(X_i)\}$ and $Y_i|X_i \overset{d}{=} N\{m_0(X_i), 1\}$. For each setting, the following working models are used: $\pi(X;\boldsymbol{\gamma}) = \mathrm{expit}(\gamma_1 + \gamma_2 X)$ and $m(X;\boldsymbol{\beta}) = \beta_1 + \beta_2 X$. Simulation experiments with correctly specified working models used $m_0(X) = 1 + X$ and $\pi_0(X) = \mathrm{expit}(\xi X)$ for $\xi = 1, 2$. To allow for misspecification in the outcome model, we additionally generated data using $m_0(X) = X^2$ and $\pi_0(X) = \mathrm{expit}(\xi X)$ for $\xi = 1, 2$. To allow for misspecification of the PS model, we generated data using $m_0(X) = 1 + X$ and $\pi_0(X) = \mathrm{expit}(-4 + 1.5|X|^{0.5} + 0.75X + 0.5|X|^{1.5})$, as in Vansteelandt, Bekaert, and Claeskens (2012). Finally, we also generated data with $m_0(X) = X^2$ and $\pi_0(X) = \mathrm{expit}(-4 + 1.5|X|^{0.5} + 0.75X + 0.5|X|^{1.5})$ to allow for misspecification of both models. In each of the settings, the target parameter $E(Y) = \mu_0$ equals one.

Results for the first simulation scenario are given in Table 1 ($n = 200$) and Table 2 ($n = 1000$). Both tables show similar results. When both working models are correct and weights are not extreme ($\xi = 1$), all estimators perform similarly in terms of bias and precision. When at most one working model is misspecified, $\hat{\mu}_{\mathrm{BR}}$ is competitive with the other DR estimators in terms of RMSE. In these cases, $\hat{\mu}_{\mathrm{BR}}$ shows lower or similar bias than $\hat{\mu}_{\mathrm{GBM}}$, $\hat{\mu}_{\mathrm{TMLE}}$, and $\hat{\mu}_{\mathrm{TMLE\text{-}SL}}$, especially when the outcome model is misspecified. When the PS model is correctly specified, $\hat{\mu}_{\mathrm{BR}}$ outperforms $\hat{\mu}_{\mathrm{PROJ}}$ for low sample size and extreme weights ($\xi = 2$) and performs just slightly worse in other settings. When the PS working model is misspecified, $\hat{\mu}_{\mathrm{BR}}$ drastically outperforms $\hat{\mu}_{\mathrm{PROJ}}$. Finally, when both working models are misspecified, $\hat{\mu}_{\mathrm{BR}}$ partly eliminates the bias amplification of the DR estimator based on standard MLE for the nuisance parameters, although not as much as $\hat{\mu}_{\mathrm{PROJ}}$ and $\hat{\mu}_{\mathrm{TMLE\text{-}SL}}$. Table 1 in the Appendix K of the online supplemental material suggests that this may be an artefact related to the considered data-generating mechanism and sample size: it shows that when both nuisance working models are misspecified, the bias of $\hat{\mu}_{\mathrm{PROJ}}$ keeps on increasing with increasing sample size (and surprisingly also its variance) in contrast to the bias of $\hat{\mu}_{\mathrm{BR}}$ which remains stable; this is in line with the fact that $\hat{\mu}_{\mathrm{BR}}$ minimizes the asymptotic bias. Finally, the smaller bias of $\hat{\mu}_{\mathrm{TMLE\text{-}SL}}$ and $\hat{\mu}_{\mathrm{TAN}}$ under misspecification of both working models is not unexpected because of the richer working models for the conditional mean outcome on which these rely.

Table 3 shows the performance of the sandwich estimator for the standard error for $\hat{\mu}_{\mathrm{BR}}$ computed as the empirical variance of (2) and confirms the asymptotic result of Corollary 1. Not surprisingly, there is undercoverage of the 95% confidence intervals when the inverse PS becomes extreme, especially at low sample size. When weights are extreme, convergence to the normal limit distribution happens more slowly. The coverage is better at $n = 1000$.

4.2 Scenario 2

The second simulation scenario is taken from Kang and Schafer (2007a). For each i $(i = 1, \ldots, n)$, $\mathbf{Z}_i = (Z_{i1}, Z_{i2}, Z_{i3}, Z_{i4})^T \overset{d}{=} \mathbf{N}(\mathbf{0}, \mathbf{I})$ where $\mathbf{I}$ is the 4×4 identity matrix, $R_i|\mathbf{Z}_i \overset{d}{=} \mathrm{Ber}\{\pi_0(\mathbf{Z}_i)\}$ with $\pi_0(\mathbf{Z}) = \mathrm{expit}(-Z_1 + 0.5Z_2 - 0.25Z_3 - 0.1Z_4)$ and $Y_i|\mathbf{Z}_i \overset{d}{=} N\{m_0(\mathbf{Z}_i), 1\}$ 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 PS model, with covariates $\mathbf{X}_i = (X_{i1}, X_{i2}, X_{i3}, X_{i4})^T$ with $X_1 = \exp(Z_1/2)$, $X_2 = Z_2/\{1 + \exp(Z_1)\} + 10$, $X_3 = (Z_1Z_3/25 + 0.6)^3$ and $X_4 = (Z_2 + Z_4 + 20)^2$. The target parameter $E(Y) = \mu_0$ 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 $(k = 1, \ldots, 4)$ and thus both working models are correctly specified or both working models are incorrectly specified. Table 4 shows the simulation results for two

Table 1. Simulation results based on 1000 Monte Carlo replications for Scenario 1, $n = 200$

Estimator	Bias	RMSE	MAE	MCSD	Bias	RMSE	MAE	MCSD
				$n = 200, \mu_0 = 1$				
	OR correct, PS correct ($\xi = 1$)				OR incorrect, PS correct ($\xi = 1$)			
$\hat{\mu}_{\text{OR}}$	−0.0015	0.13	0.09	0.13	−0.349	0.40	0.35	0.19
$\hat{\mu}_{\text{IPTW}}$	−0.0010	0.15	0.09	0.15	−0.004	0.30	0.17	0.30
$\hat{\mu}_{\text{MLE}}$	−0.0014	0.13	0.09	0.13	−0.018	0.33	0.19	0.33
$\hat{\mu}_{\text{BR}}$	−0.0012	0.13	0.09	0.13	−0.030	0.21	0.14	0.20
$\hat{\mu}_{\text{TAN}}$	−0.0011	0.13	0.09	0.13	−0.032	0.18	0.13	0.18
$\hat{\mu}_{\text{PROJ}}$	−0.0023	0.15	0.10	0.15	−0.019	0.17	0.12	0.17
$\hat{\mu}_{\text{GBM}}$	−0.0012	0.13	0.09	0.13	−0.190	0.27	0.20	0.20
$\hat{\mu}_{\text{TMLE}}$	−0.0005	0.13	0.09	0.13	−0.038	0.27	0.18	0.26
$\hat{\mu}_{\text{TMLE-SL}}$	0.0000	0.13	0.09	0.13	−0.027	0.17	0.12	0.17
	OR correct, PS correct ($\xi = 2$)				OR incorrect, PS correct ($\xi = 2$)			
$\hat{\mu}_{\text{OR}}$	−0.0017	0.14	0.09	0.14	−0.81	0.84	0.81	0.22
$\hat{\mu}_{\text{IPTW}}$	0.0030	0.26	0.12	0.26	−0.01	0.92	0.28	0.92
$\hat{\mu}_{\text{MLE}}$	−0.0015	0.20	0.12	0.20	−0.06	1.19	0.44	1.19
$\hat{\mu}_{\text{BR}}$	−0.0010	0.19	0.12	0.19	−0.11	0.26	0.17	0.24
$\hat{\mu}_{\text{TAN}}$	−0.0006	0.21	0.13	0.21	−0.10	0.25	0.16	0.23
$\hat{\mu}_{\text{PROJ}}$	−0.0024	0.33	0.18	0.33	−0.07	0.34	0.19	0.33
$\hat{\mu}_{\text{GBM}}$	−0.0014	0.15	0.10	0.15	−0.54	0.60	0.53	0.26
$\hat{\mu}_{\text{TMLE}}$	0.0029	0.17	0.11	0.17	−0.15	0.36	0.26	0.33
$\hat{\mu}_{\text{TMLE-SL}}$	0.0076	0.18	0.12	0.18	−0.08	0.24	0.15	0.23
	OR correct, PS incorrect				OR incorrect, PS incorrect			
$\hat{\mu}_{\text{OR}}$	−0.0009	0.27	0.17	0.27	0.54	0.96	0.72	0.80
$\hat{\mu}_{\text{IPTW}}$	−1.6122	5.32	0.95	5.07	6.13	16.53	3.07	15.36
$\hat{\mu}_{\text{MLE}}$	−0.0083	1.04	0.31	1.04	5.50	11.46	2.99	10.06
$\hat{\mu}_{\text{BR}}$	−0.0030	0.29	0.18	0.29	1.03	1.23	1.05	0.68
$\hat{\mu}_{\text{TAN}}$	−0.0107	0.30	0.19	0.30	0.43	0.67	0.47	0.51
$\hat{\mu}_{\text{PROJ}}$	−0.0283	0.57	0.23	0.57	−0.09	0.62	0.24	0.62
$\hat{\mu}_{\text{GBM}}$	−0.0040	0.28	0.19	0.28	0.33	0.71	0.48	0.63
$\hat{\mu}_{\text{TMLE}}$	−0.0059	0.28	0.18	0.28	1.10	1.31	1.14	0.70
$\hat{\mu}_{\text{TMLE-SL}}$	−0.0117	0.28	0.19	0.28	0.32	0.52	0.38	0.41

NOTE: Bias, Monte Carlo Bias; RMSE, root mean square error; MAE, median of absolute errors; MCSD, Monte Carlo standard deviation; OR, outcome regression; PS, propensity score. No convergence for $\hat{\boldsymbol{\gamma}}_{\text{BR}}$ was attained in five of the 1000 runs for the settings OR correct, PS correct ($\xi = 2$) and OR incorrect, PS correct ($\xi = 2$) and in three of the 1000 runs for the settings OR correct, PS incorrect and OR incorrect, PS incorrect.

scenarios where either $R = 1$ or $R = 0$ denotes the data that are observed.

As the theory dictates, all DR estimators show similar behavior when both working models are correctly specified. When the observed outcome RY is used, $\hat{\mu}_{\text{MLE}}$ shows severe erratic behavior, corresponding to the results of Kang and Schafer (2007a) but this behavior is partially eliminated when using $(1 - R)Y$ as the observed outcome, in which case $\hat{\mu}_{\text{MLE}}$ now outperforms $\hat{\mu}_{\text{OR}}$ (Robins et al. 2007). The DR estimator $\hat{\mu}_{\text{BR}}$ does not show this severe erratic behavior for both RY and $(1 - R)Y$. In line with Theorem 1, it has smaller bias compared to standard MLE. There is no single alternative outperforming the others for both sample sizes and both settings RY and $(1 - R)Y$. Overall, $\hat{\mu}_{\text{BR}}$ shows competitive performance with the other DR estimators (see also Table 1 in the Appendix K of the online supplemental material for additional results with increasing sample sizes). Table 5 shows the performance of the sandwich estimator for the standard error of $\hat{\mu}_{\text{BR}}$ computed as the empirical variance of (2) and confirms the asymptotic result of Corollary 1.

5. EXTENSION TO OTHER DOUBLY AND MULTIPLY ROBUST ESTIMATORS

5.1 Marginal Treatment Effects

Consider iid data $\{\mathbf{Z}_i = (Y_i, A_i, \mathbf{X}_i), i = 1, \dots, n\}$, where Y_i is the outcome of interest, A_i is a dichotomous treatment taking values zero and one and $\mathbf{X}_i$ is a sufficient set of covariates to control for confounding of the treatment effect, in the sense that $Y(a) \perp\!\!\!\perp A|\mathbf{X}$ for $a \in \{0, 1\}$. Here, $Y(a)$ denotes the counterfactual outcome for treatment level $a \in \{0, 1\}$, which is linked to the observed data through the consistency assumption (i.e., $Y(a) = Y$ iff $A = a$).

To obtain a DR estimator for the marginal treatment effect $\tau = E\{Y(1)\} - E\{Y(0)\} = \mu_0^{(1)} - \mu_0^{(0)}$, we need three working models: a model $\pi(\mathbf{X}; \boldsymbol{\gamma})$ for the PS $P(A = 1|\mathbf{X})$ and models $m^{(a)}(\mathbf{X}; \boldsymbol{\alpha}^{(a)})$ for the conditional mean outcome $E(Y|A = a, \mathbf{X})$ for $a \in \{0, 1\}$. We estimate the treatment effect as $\hat{\tau} = \hat{E}\{Y(1)\} - \hat{E}\{Y(0)\} = \hat{\mu}_{\text{DR}}^{(1)} - \hat{\mu}_{\text{DR}}^{(0)}$ where a DR estimator $\hat{\mu}_{\text{DR}}^{(a)} \equiv \hat{\mu}_{\text{DR}}^{(a)}(\boldsymbol{\gamma}, \boldsymbol{\alpha}^{(a)})$ of $\mu^{(a)}$ is obtained as the solution to the estimating equation $\tilde{E}_n\{\phi^{(a)}(\mathbf{Z}; \hat{\mu}^{(a)}, \boldsymbol{\gamma}, \boldsymbol{\alpha}^{(a)})\} = 0$ for $a \in \{0, 1\}$

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

Estimator	Bias	RMSE	MAE	MCSD	Bias	RMSE	MAE	MCSD
				$n = 1000, \mu_0 = 1$				
	OR correct, PS correct ($\xi = 1$)				OR incorrect, PS correct ($\xi = 1$)			
$\hat{\mu}_{\text{OR}}$	0.0037	0.057	0.039	0.057	−0.349	0.36	0.35	0.09
$\hat{\mu}_{\text{IPTW}}$	0.0022	0.064	0.044	0.064	0.006	0.13	0.08	0.13
$\hat{\mu}_{\text{MLE}}$	0.0029	0.059	0.039	0.058	0.003	0.15	0.10	0.15
$\hat{\mu}_{\text{BR}}$	0.0031	0.059	0.039	0.058	−0.003	0.09	0.07	0.09
$\hat{\mu}_{\text{TAN}}$	0.0033	0.059	0.038	0.059	−0.005	0.08	0.06	0.08
$\hat{\mu}_{\text{PROJ}}$	0.0038	0.060	0.039	0.060	−0.002	0.07	0.05	0.07
$\hat{\mu}_{\text{GBM}}$	0.0032	0.058	0.040	0.058	−0.138	0.16	0.14	0.08
$\hat{\mu}_{\text{TMLE}}$	0.0030	0.058	0.039	0.058	−0.004	0.12	0.09	0.12
$\hat{\mu}_{\text{TMLE-SL}}$	0.0031	0.058	0.039	0.058	−0.002	0.07	0.05	0.07
	OR correct, PS correct ($\xi = 2$)				OR incorrect, PS correct ($\xi = 2$)			
$\hat{\mu}_{\text{OR}}$	0.0033	0.065	0.044	0.065	−0.810	0.82	0.81	0.10
$\hat{\mu}_{\text{IPTW}}$	0.0051	0.120	0.066	0.120	−0.023	0.34	0.17	0.34
$\hat{\mu}_{\text{MLE}}$	−0.0003	0.090	0.057	0.090	−0.041	0.50	0.25	0.49
$\hat{\mu}_{\text{BR}}$	0.0004	0.085	0.057	0.085	−0.052	0.14	0.09	0.13
$\hat{\mu}_{\text{TAN}}$	0.0005	0.091	0.060	0.091	−0.042	0.11	0.07	0.10
$\hat{\mu}_{\text{PROJ}}$	−0.0014	0.109	0.069	0.109	−0.027	0.12	0.07	0.11
$\hat{\mu}_{\text{GBM}}$	0.0026	0.072	0.048	0.071	−0.407	0.43	0.40	0.12
$\hat{\mu}_{\text{TMLE}}$	0.0011	0.080	0.054	0.080	−0.124	0.19	0.14	0.15
$\hat{\mu}_{\text{TMLE-SL}}$	0.0011	0.081	0.055	0.081	−0.041	0.11	0.07	0.10
	OR correct, PS incorrect				OR incorrect, PS incorrect			
$\hat{\mu}_{\text{OR}}$	−0.0056	0.11	0.07	0.11	0.70	0.78	0.70	0.35
$\hat{\mu}_{\text{IPTW}}$	−1.8704	2.50	1.46	1.65	7.34	10.03	5.59	6.84
$\hat{\mu}_{\text{MLE}}$	−0.0264	0.52	0.20	0.52	7.24	9.64	5.57	6.37
$\hat{\mu}_{\text{BR}}$	−0.0057	0.11	0.08	0.11	1.24	1.28	1.22	0.33
$\hat{\mu}_{\text{TAN}}$	−0.0043	0.12	0.08	0.12	0.43	0.47	0.42	0.20
$\hat{\mu}_{\text{PROJ}}$	−0.0007	0.50	0.17	0.50	0.09	0.57	0.19	0.56
$\hat{\mu}_{\text{GBM}}$	−0.0002	0.13	0.09	0.13	0.20	0.27	0.21	0.18
$\hat{\mu}_{\text{TMLE}}$	−0.0052	0.12	0.08	0.12	1.20	1.24	1.20	0.31
$\hat{\mu}_{\text{TMLE-SL}}$	−0.0071	0.12	0.08	0.12	0.10	0.18	0.12	0.15

NOTE: Bias, Monte Carlo Bias; RMSE, root mean square error; MAE, median of absolute errors; MCSD, Monte Carlo standard deviation; OR, outcome regression; PS, propensity score.

(Scharfstein, Rotnitzky, and Robins 1999b), where

$$\phi^{(1)}(\mathbf{Z}; \mu^{(1)}, \boldsymbol{\gamma}, \boldsymbol{\alpha}^{(1)}) = \frac{AY}{\pi(\mathbf{X}; \boldsymbol{\gamma})} - \frac{A - \pi(\mathbf{X}; \boldsymbol{\gamma})}{\pi(\mathbf{X}; \boldsymbol{\gamma})} m^{(1)}(\mathbf{X}; \boldsymbol{\alpha}^{(1)}) - \mu^{(1)}$$

$$\phi^{(0)}(\mathbf{Z}; \mu^{(0)}, \boldsymbol{\gamma}, \boldsymbol{\alpha}^{(0)}) = \frac{(1-A)Y}{1 - \pi(\mathbf{X}; \boldsymbol{\gamma})} + \frac{A - \pi(\mathbf{X}; \boldsymbol{\gamma})}{1 - \pi(\mathbf{X}; \boldsymbol{\gamma})} m^{(0)}(\mathbf{X}; \boldsymbol{\alpha}^{(0)}) - \mu^{(0)}.$$

The proposed estimation strategy proceeds by setting the gradients w.r.t. the nuisance parameters equal to zero, which amounts to solving $(\hat{\boldsymbol{\gamma}}^{(1)}_{\text{BR}}, \hat{\boldsymbol{\alpha}}^{(1)}_{\text{BR}})$ from the system

$$\mathbf{0} = \tilde{E}_n \left[\left\{ 1 - \frac{A}{\pi(\mathbf{X}; \hat{\boldsymbol{\gamma}}^{(1)}_{\text{BR}})} \right\} m^{(1)}_{\boldsymbol{\alpha}^{(1)}}(\mathbf{X}; \hat{\boldsymbol{\alpha}}^{(1)}_{\text{BR}}) \right], \quad (8)$$

$$\mathbf{0} = \tilde{E}_n \left[\left\{ Y - m^{(1)}(\mathbf{X}; \hat{\boldsymbol{\alpha}}^{(1)}_{\text{BR}}) \right\} \frac{A}{\pi^2(\mathbf{X}; \hat{\boldsymbol{\gamma}}^{(1)}_{\text{BR}})} \pi_{\boldsymbol{\gamma}}(\mathbf{X}; \hat{\boldsymbol{\gamma}}^{(1)}_{\text{BR}}) \right] \quad (9)$$

Table 3. Performance of standard error estimates and confidence intervals for the bias-reduced strategy based on 1000 Monte Carlo replications in Scenario 1

	$n = 200$			$n = 1000$		
Setting	MCSD	ASSE	COV	MCSD	ASSE	COV
OR correct, PS correct ($\xi = 1$)	0.13	0.13	0.94	0.06	0.06	0.96
OR correct, PS correct ($\xi = 2$)	0.19	0.15	0.88	0.09	0.08	0.91
OR incorrect, PS correct ($\xi = 1$)	0.20	0.17	0.88	0.09	0.09	0.94
OR incorrect, PS correct ($\xi = 2$)	0.24	0.17	0.77	0.13	0.10	0.80
OR correct, PS incorrect	0.29	0.22	0.85	0.11	0.11	0.94
OR incorrect, PS incorrect	0.68	0.43	0.34	0.33	0.28	0.01

NOTE: MCSD, Monte Carlo standard deviation; ASSE, average of sandwich standard errors; COV, Monte Carlo coverage of 95% Wald confidence intervals; OR, outcome regression; PS, propensity score.

Table 4. Simulation results based on 1000 Monte Carlo replications for Scenario 2, $\mu_0 = 210$

Estimator	Bias	RMSE	MAE	MCSD	Bias	RMSE	MAE	MCSD
	Observed outcome RY				Observed outcome $(1-R)Y$			
	$n = 200$							
	OR correct, PS correct				OR correct, PS correct			
$\hat{\mu}_{\text{OR}}$	0.092	2.52	1.68	2.52	0.088	2.53	1.68	2.53
$\hat{\mu}_{\text{IPTW}}$	−1.761	28.15	13.22	28.10	−0.254	16.64	8.22	16.65
$\hat{\mu}_{\text{MLE}}$	0.099	2.53	1.70	2.53	0.085	2.53	1.73	2.53
$\hat{\mu}_{\text{BR}}$	0.090	2.54	1.71	2.54	0.095	2.54	1.69	2.54
$\hat{\mu}_{\text{TAN}}$	0.094	2.53	1.72	2.53	0.085	2.53	1.69	2.53
$\hat{\mu}_{\text{PROJ}}$	0.090	2.55	1.71	2.55	0.079	2.54	1.72	2.54
$\hat{\mu}_{\text{GBM}}$	0.093	2.53	1.70	2.53	0.088	2.53	1.67	2.53
$\hat{\mu}_{\text{TMLE}}$	0.032	2.53	1.72	2.53	0.238	2.55	1.77	2.54
$\hat{\mu}_{\text{TMLE-SL}}$	0.031	2.53	1.71	2.53	0.241	2.55	1.78	2.54
	OR incorrect, PS incorrect				OR incorrect, PS incorrect			
$\hat{\mu}_{\text{OR}}$	−0.17	3.60	2.51	3.59	7.15	7.76	7.21	3.01
$\hat{\mu}_{\text{IPTW}}$	68.77	453.60	18.43	448.58	−0.80	12.18	6.24	12.16
$\hat{\mu}_{\text{MLE}}$	−15.15	88.60	4.40	87.34	4.76	6.05	5.03	3.74
$\hat{\mu}_{\text{BR}}$	−2.24	4.45	2.78	3.85	3.44	4.63	3.55	3.10
$\hat{\mu}_{\text{TAN}}$	−2.55	4.31	2.96	3.47	4.76	5.79	4.85	3.30
$\hat{\mu}_{\text{PROJ}}$	−0.04	3.93	2.58	3.93	1.00	3.60	2.36	3.46
$\hat{\mu}_{\text{GBM}}$	−0.22	3.46	2.42	3.45	5.84	6.58	5.86	3.03
$\hat{\mu}_{\text{TMLE}}$	−4.38	6.20	4.08	4.39	4.43	5.56	4.47	3.35
$\hat{\mu}_{\text{TMLE-SL}}$	−2.31	4.02	2.72	3.29	3.75	5.07	3.95	3.41
	$n = 1000$							
	OR correct, PS correct				OR correct, PS correct			
$\hat{\mu}_{\text{OR}}$	0.023	1.12	0.76	1.12	0.024	1.13	0.76	1.13
$\hat{\mu}_{\text{IPTW}}$	−0.282	11.27	6.61	11.27	0.052	7.67	3.93	7.67
$\hat{\mu}_{\text{MLE}}$	0.023	1.12	0.76	1.12	0.022	1.13	0.78	1.13
$\hat{\mu}_{\text{BR}}$	0.022	1.12	0.76	1.12	0.024	1.13	0.78	1.13
$\hat{\mu}_{\text{TAN}}$	0.021	1.12	0.76	1.12	0.024	1.13	0.77	1.13
$\hat{\mu}_{\text{PROJ}}$	0.024	1.12	0.76	1.12	0.025	1.13	0.77	1.13
$\hat{\mu}_{\text{GBM}}$	0.022	1.12	0.76	1.12	0.024	1.13	0.77	1.13
$\hat{\mu}_{\text{TMLE}}$	0.013	1.12	0.76	1.12	0.061	1.13	0.76	1.13
$\hat{\mu}_{\text{TMLE-SL}}$	0.013	1.12	0.76	1.12	0.060	1.13	0.77	1.13
	OR incorrect, PS incorrect				OR incorrect, PS incorrect			
$\hat{\mu}_{\text{OR}}$	−0.46	1.64	1.11	1.58	7.18	7.31	7.14	1.35
$\hat{\mu}_{\text{IPTW}}$	161.92	1194.95	36.92	1184.52	−1.00	4.93	3.12	4.83
$\hat{\mu}_{\text{MLE}}$	−53.71	469.04	8.81	466.19	4.51	4.83	4.51	1.73
$\hat{\mu}_{\text{BR}}$	−3.21	3.63	3.10	1.69	2.97	3.25	2.91	1.34
$\hat{\mu}_{\text{TAN}}$	−2.92	3.29	3.02	1.51	4.35	4.62	4.28	1.55
$\hat{\mu}_{\text{PROJ}}$	−1.55	2.03	1.56	1.32	1.39	1.92	1.53	1.33
$\hat{\mu}_{\text{GBM}}$	−0.60	1.56	1.09	1.44	4.83	5.03	4.75	1.39
$\hat{\mu}_{\text{TMLE}}$	−4.73	5.13	4.73	1.99	4.16	4.43	4.11	1.52
$\hat{\mu}_{\text{TMLE-SL}}$	−2.25	2.76	2.34	1.61	2.65	3.12	2.57	1.65

NOTE: Bias, Monte Carlo Bias; RMSE, root mean square error; MAE, median of absolute errors; MCSD, Monte Carlo standard deviation; OR, outcome regression; PS, propensity score. No convergence for $\hat{\gamma}_{\text{BR}}$ was attained in 13 of the 1000 runs for the settings OR correct, PS correct, $n = 200$ for both the observed outcome RY and $(1-R)Y$ and in five of the 1000 runs for the setting OR incorrect, PS incorrect, $n = 200$ for the observed outcome RY.

Table 5. Performance of standard error estimates and confidence intervals for the bias-reduced strategy based on 1000 Monte Carlo replications in Scenario 2

	$n = 200$			$n = 1000$		
Setting	MCSD	ASSE	COV	MCSD	ASSE	COV
OR correct, PS correct (RY)	2.54	2.57	0.96	1.12	1.15	0.96
OR correct, PS correct ($(1-R)Y$)	2.54	2.57	0.95	1.13	1.15	0.96
OR incorrect, PS incorrect (RY)	3.85	2.95	0.82	1.69	1.34	0.38
OR incorrect, PS incorrect ($(1-R)Y$)	3.10	2.73	0.73	1.34	1.25	0.35

NOTE: MCSD, Monte Carlo standard deviation; ASSE, average of sandwich standard errors; COV, Monte Carlo coverage of 95% Wald confidence intervals; OR, outcome regression; PS, propensity score.

and solving $(\hat{\boldsymbol{\gamma}}_{\text{BR}}^{(0)}, \hat{\boldsymbol{\alpha}}_{\text{BR}}^{(0)})$ from a similar system of equations but with A replaced by $1-A$, $\pi(\mathbf{X}; \hat{\boldsymbol{\gamma}}_{\text{BR}}^{(1)})$ replaced by $1-\pi(\mathbf{X}; \hat{\boldsymbol{\gamma}}_{\text{BR}}^{(0)})$, $m^{(1)}(\mathbf{X}; \hat{\boldsymbol{\alpha}}_{\text{BR}}^{(1)})$ replaced by $m^{(0)}(\mathbf{X}; \hat{\boldsymbol{\alpha}}_{\text{BR}}^{(0)})$ and $\pi_{\boldsymbol{\gamma}}(\mathbf{X}; \hat{\boldsymbol{\gamma}}_{\text{BR}}^{(1)})$ by $\pi_{\boldsymbol{\gamma}}(\mathbf{X}; \hat{\boldsymbol{\gamma}}_{\text{BR}}^{(0)})$ and where $\pi_{\boldsymbol{\gamma}}(\mathbf{X}; \boldsymbol{\gamma}) = \partial\pi(\mathbf{X}; \boldsymbol{\gamma})/\partial\boldsymbol{\gamma}$ and $m_{\boldsymbol{\alpha}^{(a)}}^{(a)}(\mathbf{X}; \boldsymbol{\alpha}^{(a)}) = \partial m^{(a)}(\mathbf{X}; \boldsymbol{\alpha}^{(a)})/\partial\boldsymbol{\alpha}^{(a)}$ for $a \in \{0, 1\}$. This results in estimators with similar properties as the estimators in Section 3.3. Note that the DR estimators $\hat{\mu}_{\text{DR}}^{(1)}$ and $\hat{\mu}_{\text{DR}}^{(0)}$, while relying on the same working model $\pi(\mathbf{X}; \boldsymbol{\gamma})$ for the PS, use different estimators for the nuisance parameters indexing that model. In particular, (9) forces the model to fit well in covariate regions with low PS where not being treated is more likely and the equivalent for $\hat{\boldsymbol{\alpha}}_{\text{BR}}^{(0)}$ forces the model to fit well in covariate regions with high PS where being treated is more likely. Additionally, (8) ensures stability of inverse weights equalling one over the PS, while the equivalent for $\hat{\boldsymbol{\gamma}}_{\text{BR}}^{(0)}$ ensures stability of inverse weights equalling one over one minus the PS. This illustrates that the nuisance parameter estimators adapt to the considered estimand. The use of these estimators is illustrated in a reanalysis of the SUPPORT study in Section 6.

A more general development works under the marginal structural model (Robins, Hernán, and Brumback 2000) $E\{Y(a)\} = \beta_0 + \beta_1 a$, where $a \in \mathscr{A}$ may be a continuous exposure level with $\mathscr{A}$ the support of A. In Appendix E of the online supplemental material, we show how a bias-reduced DR estimator for β_1 can be obtained.

5.2 G-Estimation for Semiparametric Regression Models

Consider the semiparametric linear regression model $E(Y|A, \mathbf{X}) = m_0(\mathbf{X}) + \tau_0 A$ (Robins, Mark, and Newey 1992). A DR G-estimator $\hat{\tau}_{\text{G,DR}}(\boldsymbol{\gamma}, \boldsymbol{\alpha})$ of τ_0 is obtained by solving $0 = \tilde{E}_n\{U_{\text{G}}(\mathbf{Z}; \tau, \boldsymbol{\gamma}, \boldsymbol{\alpha})\} = \tilde{E}_n[\{Y - \tau A - m(\mathbf{X}; \boldsymbol{\alpha})\}\{A - \pi(\mathbf{X}; \boldsymbol{\gamma})\}]$, which is unbiased if either the working model $(\mathscr{M}(\boldsymbol{\gamma}))$ $\pi(\mathbf{X}; \gamma)$ for $E(A|\mathbf{X})$ or the working model $(\mathscr{M}(\boldsymbol{\alpha}))$ $m(\mathbf{X}; \boldsymbol{\alpha})$ for $E(Y|A=0, \mathbf{X})$ is correctly specified. For fixed $\boldsymbol{\gamma}$ and $\boldsymbol{\alpha}$, the DR estimator $\hat{\tau}_{\text{G,DR}}(\boldsymbol{\gamma}, \boldsymbol{\alpha})$ admits the expansion $n^{1/2}\{\hat{\tau}_{\text{G,DR}}(\boldsymbol{\gamma}, \boldsymbol{\alpha}) - \tau_0\} = n^{1/2}\tilde{E}_n\{\phi_{\text{G}}(\mathbf{Z}; \tau_0, \boldsymbol{\gamma}, \boldsymbol{\alpha})\} + o_p(1)$ with influence function $\phi_{\text{G}}(\mathbf{Z}; \tau_0, \boldsymbol{\gamma}, \boldsymbol{\alpha}) = E^{-1}[A\{A - \pi(\mathbf{X}; \boldsymbol{\gamma})\}]\{A - \pi(\mathbf{X}; \boldsymbol{\gamma})\}\{Y - m(\mathbf{X}; \boldsymbol{\alpha})\} - \tau_0$. The gradients of $\phi_{\text{G}}(\mathbf{Z}; \tau_0, \boldsymbol{\gamma}, \boldsymbol{\alpha})$ w.r.t. $\boldsymbol{\alpha}$ and $\boldsymbol{\gamma}$ then define the estimating equations for $\boldsymbol{\gamma}$ and $\boldsymbol{\alpha}$; that is, we solve $(\hat{\boldsymbol{\gamma}}_{\text{BR}}, \hat{\boldsymbol{\alpha}}_{\text{BR}})$ from the system $\mathbf{0} = \tilde{E}_n[\{A - \pi(\mathbf{X}; \hat{\boldsymbol{\gamma}}_{\text{BR}})\}m_{\boldsymbol{\alpha}}(\mathbf{X}; \hat{\boldsymbol{\alpha}}_{\text{BR}})]$ and $\mathbf{0} = \tilde{E}_n[\{Y - m(\mathbf{X}; \hat{\boldsymbol{\alpha}}_{\text{BR}})\}\widehat{\text{W}}(A, \mathbf{X}; \hat{\boldsymbol{\gamma}}_{\text{BR}})]$, $\widehat{\text{W}}(A, \mathbf{X}; \hat{\boldsymbol{\gamma}}_{\text{BR}}) = \{A - \pi(\mathbf{X}; \hat{\boldsymbol{\gamma}}_{\text{BR}})\}\tilde{E}_n\{A\pi_{\boldsymbol{\gamma}}(\mathbf{X}; \hat{\boldsymbol{\gamma}}_{\text{BR}})\} - \pi_{\boldsymbol{\gamma}}(\mathbf{X}; \hat{\boldsymbol{\gamma}}_{\text{BR}})\tilde{E}_n[A\{A - \pi(\mathbf{X}; \hat{\boldsymbol{\gamma}}_{\text{BR}})\}]$, $\pi_{\boldsymbol{\gamma}}(\mathbf{X}; \boldsymbol{\gamma}) = \partial\pi(\mathbf{X}; \boldsymbol{\gamma})/\partial\boldsymbol{\gamma}$, $m_{\boldsymbol{\alpha}}(\mathbf{X}; \boldsymbol{\alpha}) = \partial m(\mathbf{X}; \boldsymbol{\alpha})/\partial\boldsymbol{\alpha}$. For a linear outcome model $m(\mathbf{X}; \boldsymbol{\alpha}) = \alpha_0 + \boldsymbol{\alpha}_1^T\mathbf{X}$ and a logistic regression model $\pi(\mathbf{X}; \boldsymbol{\gamma}) = \text{expit}(\gamma_0 + \boldsymbol{\gamma}_1^T\mathbf{X})$ for the PS, the estimating equation for $\boldsymbol{\gamma}$ reduces to standard MLE because $m_{\boldsymbol{\alpha}}(\mathbf{X}; \hat{\boldsymbol{\alpha}}_{\text{BR}}) = (1, \mathbf{X}^T)^T$. The estimating equation for $\boldsymbol{\alpha}$ is obtained by substituting $\pi_{\boldsymbol{\gamma}}(\mathbf{X}; \hat{\boldsymbol{\gamma}}_{\text{BR}}) = \pi(\mathbf{X}; \hat{\boldsymbol{\gamma}}_{\text{BR}})\{1 - \pi(\mathbf{X}; \hat{\boldsymbol{\gamma}}_{\text{BR}})\}(1, \mathbf{X}^T)^T$.

Consider now the semiparametric log-linear model $\log E(Y|A, \mathbf{X}) = m_0(\mathbf{X}) + \tau_0 A$ (Robins, Mark, and Newey 1992). A DR G-estimator $\hat{\tau}'_{\text{G,DR}}(\boldsymbol{\gamma}, \boldsymbol{\alpha})$ of τ_0 is obtained by solving the estimating equation $0 = \tilde{E}_n\{U'_{\text{G}}(\mathbf{Z}; \tau, \boldsymbol{\gamma}, \boldsymbol{\alpha})\} = \tilde{E}_n(\{A - \pi(\mathbf{X}; \boldsymbol{\gamma})\}[Y\exp(-\tau A) - \exp\{m(\mathbf{X}; \boldsymbol{\alpha})\}]$, which is unbiased if either the working model $(\mathscr{M}(\boldsymbol{\gamma}))$ $\pi(\mathbf{X}; \gamma)$ or the working model $(\mathscr{M}(\boldsymbol{\alpha}))$ $m(\mathbf{X}; \boldsymbol{\alpha})$ for $\log E(Y|A = 0, \mathbf{X})$ is correctly specified. Although the gradients of the corresponding influence function $\phi'_{\text{G}}(\mathbf{Z}_i; \tau_0, \boldsymbol{\gamma}, \boldsymbol{\alpha}) = -E^{-1}\{\partial U'_{\text{G}}(\mathbf{Z}; \tau_0, \boldsymbol{\gamma}, \boldsymbol{\alpha})/\partial\tau\}U'_{\text{G}}(\mathbf{Z}_i; \tau_0, \boldsymbol{\gamma}, \boldsymbol{\alpha})$ w.r.t. the nuisance parameters continue to deliver consistent estimators for the nuisance parameters, these estimators no longer ensure minimal squared first-order asymptotic bias under misspecification of both working models because these gradients now depend on the unknown population value τ_0. Nevertheless, Theorem 1 continues to apply for score tests of the null hypothesis that $\tau = \tilde{\tau}$ for some $\tilde{\tau}$. In particular, when the estimators $\hat{\boldsymbol{\gamma}}_{\text{BR}}$ and $\hat{\boldsymbol{\alpha}}_{\text{BR}}$ are defined with the known value $\tilde{\tau}$ substituted for the unknown value of τ, they minimize $E^2\{\phi'_{\text{G}}(\mathbf{Z}; \tilde{\tau}, \hat{\boldsymbol{\gamma}}_{\text{BR}}, \hat{\boldsymbol{\alpha}}_{\text{BR}})\}$.

5.3 Mean Outcome When Missingness is Nonignorable

We reconsider the estimation of a population mean outcome $\mu_0 = E(Y)$ in the presence of incomplete data. Suppose, as in Section 2, that we have iid data $\{\mathbf{Z}_i = (R_iY_i, R_i, \mathbf{X}_i), i = 1, \ldots, n\}$, but that in contrast to Section 2, missingness is not ignorable. Suppose, therefore, that (i) $P(R = 1|\mathbf{X}, Y) > 0$ with probability 1 and (ii) if $P(R = 0|\mathbf{X}) > 0$, $P(R = 0|\mathbf{X}, Y) = \text{expit}\{h_0(\mathbf{X}) + q(\mathbf{X}, Y; \boldsymbol{\kappa})\}$ for some unknown function $h_0(\mathbf{X})$ and a user-specified selection bias function $q(\mathbf{X}, Y; \boldsymbol{\kappa})$ with known $\boldsymbol{\kappa}$ and $q(\mathbf{X}, 0; \boldsymbol{\kappa}) \equiv q(\mathbf{X}, Y; \mathbf{0}) \equiv 0$, for example, $q(\mathbf{X}, Y; \kappa) = \kappa Y$. Let $\mathscr{M}(\boldsymbol{\kappa})$ denote the model for the full data defined by assumptions (i) and (ii). Since $\boldsymbol{\kappa} = \mathbf{0}$ encodes MAR, the selection bias function $q(\mathbf{X}, Y; \boldsymbol{\kappa})$ encodes the degree of deviation from the MAR assumption (Scharfstein, Rotnitzky, and Robins 1999a, b). Scharfstein, Rotnitzky, and Robins (1999a) show that for each choice of $\boldsymbol{\kappa}$, model $\mathscr{M}(\boldsymbol{\kappa})$ places no restriction on the observed data law so that in particular the observed data carry no information of $\boldsymbol{\kappa}$. Model $\mathscr{M}(\boldsymbol{\kappa})$ is hence particularly useful for a sensitivity analysis based upon varying $\boldsymbol{\kappa}$. It can be shown that at $\mathscr{M}(\boldsymbol{\kappa})$, $\mu_0 = E[RY/\pi_0(\mathbf{X}, Y; \boldsymbol{\kappa}) + \{1 - R/\pi_0(\mathbf{X}, Y; \boldsymbol{\kappa})\}m_0(\mathbf{X})]$ with $m_0(\mathbf{X}) = E(Y|R = 0, \mathbf{X}) = E[Y\exp\{q(\mathbf{X}, Y; \boldsymbol{\kappa})\}|R = 1, \mathbf{X}]/E[\exp\{q(\mathbf{X}, Y; \boldsymbol{\kappa})\}|R = 1, \mathbf{X}]$ and $\pi_0(\mathbf{X}, Y; \boldsymbol{\kappa}) = P(R = 1|Y, \mathbf{X}) = [1 + \exp\{h_0(\mathbf{X}) + q(\mathbf{X}, Y; \boldsymbol{\kappa})\}]^{-1}$.

To construct a DR estimator for the target parameter μ_0 (Scharfstein, Rotnitzky, and Robins 1999b), we need two working models: (i) a model $\mathscr{M}(\boldsymbol{\beta})$ for the unknown function $m_0(\mathbf{X})$ given by $m(\mathbf{X}; \boldsymbol{\beta})$, where $m(\mathbf{X}; \boldsymbol{\beta})$ is a known function smooth in a finite-dimensional parameter $\boldsymbol{\beta}$ and (ii) a model $\mathscr{M}(\boldsymbol{\gamma})$ for the unknown function $h_0(\mathbf{X})$ given by $h(\mathbf{X}; \boldsymbol{\gamma})$, where $h(\mathbf{X}; \boldsymbol{\gamma})$ is a known function smooth in a finite-dimensional parameter $\boldsymbol{\gamma}$. The induced working model for $\pi_0(\mathbf{X}, Y; \boldsymbol{\kappa})$ is then $\pi(\mathbf{X}, Y; \boldsymbol{\gamma}, \boldsymbol{\kappa}) = [1 + \exp\{h(\mathbf{X}; \boldsymbol{\gamma}) + q(\mathbf{X}, Y; \boldsymbol{\kappa})\}]^{-1}$. For given $\boldsymbol{\gamma}$ and $\boldsymbol{\beta}$, the target parameter μ_0 can be estimated as $\hat{\mu}_{\text{DR}} \equiv \hat{\mu}_{\text{DR}}(\boldsymbol{\gamma}, \boldsymbol{\beta}; \boldsymbol{\kappa})$, obtained as a solution to the estimating equation $\tilde{E}_n\{\phi_q(\mathbf{Z}; \hat{\mu}_{\text{DR}}, \boldsymbol{\gamma}, \boldsymbol{\beta}, \boldsymbol{\kappa})\} = 0$ where ϕ_q, the influence function of μ_0 (at fixed nuisance parameters), equals ϕ in (2) but with $\pi(\mathbf{X}; \boldsymbol{\gamma})$ replaced by $\pi(\mathbf{X}, Y; \boldsymbol{\gamma}, \boldsymbol{\kappa})$. The bias-reduced DR estimator can be straightforwardly obtained because of the double robustness of $\hat{\mu}_{\text{DR}}(\boldsymbol{\gamma}, \boldsymbol{\beta}; \boldsymbol{\kappa})$ under model $\mathscr{M}(\boldsymbol{\kappa}) \cap \{\mathscr{M}(\boldsymbol{\gamma}) \cup \mathscr{M}(\boldsymbol{\beta})\}$. Upon taking the gradients of ϕ_q with respect to $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$, this amounts to solving $(\hat{\boldsymbol{\gamma}}_{\text{BR}}, \hat{\boldsymbol{\beta}}_{\text{BR}})$ from the system $\mathbf{0} = \tilde{E}_n\{(1 - R[1 + \exp\{h(\mathbf{X}; \hat{\boldsymbol{\gamma}}_{\text{BR}}) + q(\mathbf{X}, Y; \boldsymbol{\kappa})\}])m_{\boldsymbol{\beta}}(\mathbf{X}; \hat{\boldsymbol{\beta}}_{\text{BR}})\}$ and $\mathbf{0} = \tilde{E}_n[R\{Y - m(\mathbf{X}; \hat{\boldsymbol{\beta}}_{\text{BR}})\}\exp\{h(\mathbf{X}; \hat{\boldsymbol{\gamma}}_{\text{BR}}) + q(\mathbf{X}, Y; \boldsymbol{\kappa})\}h_{\boldsymbol{\gamma}}(\mathbf{X};$

$\hat{\boldsymbol{\gamma}}_{\rm BR})]$, with $h_{\boldsymbol{\gamma}}(\mathbf{X};\boldsymbol{\gamma}) = \partial h(\mathbf{X};\boldsymbol{\gamma})/\partial\boldsymbol{\gamma}$ and $m_{\boldsymbol{\beta}}(\mathbf{X};\boldsymbol{\beta}) = \partial m(\mathbf{X};\boldsymbol{\beta})/\partial\boldsymbol{\beta}$. Note that the tilt function $\exp\{q(\mathbf{X}, Y, \boldsymbol{\kappa})\}$ in the latter involves the selection bias function which links the outcome distribution in responders to that in nonresponders, thereby assuring the unbiasedness of this estimating equation under model $\mathscr{M}(\boldsymbol{\kappa}) \cap \mathscr{M}(\boldsymbol{\beta})$. This also results in estimators with similar properties as the estimators in Section 3.3.

Note thus that the results of Section 3.3 immediately extend to nonignorable missingness, unlike certain other alternative nuisance parameter estimation strategies. For instance, in Appendix F of the online supplemental material, we show that strategy of Cao, Tsiatis, and Davidian (2009) does not necessarily lead to an unbiased estimating function for the parameter $\boldsymbol{\beta}$ indexing the working model $m(\mathbf{X};\boldsymbol{\beta})$ for $E(Y|R=0,\mathbf{X})$ when missingness is nonignorable.

5.4 Multiply Robust Estimation in Semiparametric Interaction Models

The principle behind the biased-reduced DR estimation strategy is extensible to certain multiply robust estimators, estimators that are consistent under a union model that assumes that at least one of several working models holds. Consider iid data $\{\mathbf{Z}_i = (Y_i, \mathbf{A}_i, \mathbf{X}_i), i = 1, \ldots, n\}$, where Y_i is the outcome, $\mathbf{A}_i = (A_{i1}, A_{i2})^T$ is a vector of binary exposure variables and $\mathbf{X}_i$ is a vector of extraneous variables. Vansteelandt et al. (2008) developed inference for β_0 in the semiparametric interaction model $\mathscr{M}$ defined by $E(Y|\mathbf{A},\mathbf{X}) = \beta_0 A_1 A_2 + q_1(A_1, \mathbf{X}) + q_2(A_2, \mathbf{X}) + q_0(\mathbf{X})$ where $q_j(A_j, \mathbf{X})$ $(j = 1, 2)$ and $q_0(\mathbf{X})$ are unknown functions satisfying $q_j(0, \mathbf{X}) = 0$. They showed how a multiply robust estimator for β_0 can be obtained under this model. In Appendix G of the online supplemental material, we show how a bias-reduced multiply robust estimator can be obtained.

6. DATA ANALYSIS

We reanalyze data from the Study to Understand Prognoses and Preferences for Outcomes and Risks of Treatments (SUPPORT) conducted in 1989–1994 in $n = 5735$ critically ill patients in five U.S. hospitals to study the effectiveness of right heart catheterization (RHC) in the initial care unit (ICU) of critically ill patients (Connors et al. 1996). RHC is a diagnostic procedure which, at the time of the study by Connors et al. (1996), was thought to lead to better patient outcomes by many physicians. The effectiveness of RHC had not been demonstrated in a randomized clinical trial but based on expert information, a rich set of 72 variables was collected to adjust for potential confounding (see Table 1 in Hirano and Imbens (2002)). The original analysis in Connors et al. (1996) used PS matching and surprisingly found that RHC leads to lower survival as compared to not performing RHC. For each patient, the treatment status A indicates 1 if RHC was applied within 24 h of admission and 0 otherwise. In total, 2184 patients received RHC and 3551 did not. We consider the effect of RHC on 30-day survival Y with $Y = 1$ indicating survival, 0 otherwise. In total, 3817 patients survived and 1918 died within 30 days. Figure 2 in Appendix L of the online supplemental material visualizes the large differences that exist in baseline covariate means between treated and untreated patients (see the x-axis). A detailed description is given in Table 2 of Hirano and Imbens (2002).

To estimate the additive treatment effect $\tau = E\{Y(1)\} - E\{Y(0)\}$, we use the results of Section 5.1. As in Hirano and Imbens (2002), we model the PS $P(A=1|\mathbf{X})$ using a logistic regression including a constant term and all 72 main effects; $\pi(\mathbf{X};\boldsymbol{\gamma}) = \mathrm{expit}\{\boldsymbol{\gamma}^T(1,\mathbf{X}^T)^T\}$. We model the conditional mean outcome $E(Y|A=a,\mathbf{X})$ for $a \in \{0,1\}$ using both a linear and logistic regression model $m^{(a)}_{\rm lin}(\mathbf{X};\boldsymbol{\alpha}^{(a)}) = \boldsymbol{\alpha}^{(a)T}(1,\mathbf{X}^T)^T$ and $m^{(a)}_{\rm logit}(\mathbf{X};\boldsymbol{\alpha}^{(a)}) = \mathrm{expit}(\boldsymbol{\alpha}^{(a)T}(1,\mathbf{X}^T)^T)$ for $a \in \{0,1\}$, including a constant term and all 72 main effects. For the linear outcome model, we obtain estimators $(\hat{\boldsymbol{\gamma}}^{(1)}_{\rm BR,lin}, \hat{\boldsymbol{\alpha}}^{(1)}_{\rm BR,lin}, \hat{\boldsymbol{\gamma}}^{(0)}_{\rm BR,lin}, \hat{\boldsymbol{\alpha}}^{(0)}_{\rm BR,lin})$ solving estimating Equations (8) and (9) for condition $A=1$ and their analogues for condition $A=0$ with $\pi_{\boldsymbol{\gamma}}(\mathbf{X};\boldsymbol{\gamma}) = \{1-\pi(\mathbf{X};\boldsymbol{\gamma})\}\pi(\mathbf{X};\boldsymbol{\gamma})(1,\mathbf{X}^T)^T$ and $m^{(a)}_{\boldsymbol{\alpha}^{(a)},\rm lin}(\mathbf{X};\boldsymbol{\alpha}^{(a)}) = (1,\mathbf{X}^T)^T$. The estimator $\hat{\boldsymbol{\gamma}}^{(a)}_{\rm BR,lin}$ is obtained by maximizing the function $\mathscr{F}^{(a)}_{\rm lin}(\boldsymbol{\gamma}) = \tilde{E}_n[(-1)^a(A-1+a)\exp\left\{(-1)^a\boldsymbol{\gamma}^T(1,\mathbf{X}^T)^T\right\} + (A-a)\boldsymbol{\gamma}^T(1,\mathbf{X}^T)^T]$. For the logistic outcome model, we obtain estimators $(\hat{\boldsymbol{\gamma}}^{(1)}_{\rm BR,logit}, \hat{\boldsymbol{\alpha}}^{(1)}_{\rm BR,logit}, \hat{\boldsymbol{\gamma}}^{(0)}_{\rm BR,logit}, \hat{\boldsymbol{\alpha}}^{(0)}_{\rm BR,logit})$ that solve the same equations but with $m^{(a)}_{\boldsymbol{\alpha}^{(a)},\rm logit}(\mathbf{X};\boldsymbol{\alpha}^{(a)}) = \{1-m^{(a)}_{\rm logit}(\mathbf{X};\boldsymbol{\alpha}^{(a)})\}m^{(a)}_{\rm logit}(\mathbf{X};\boldsymbol{\alpha}^{(a)})(1,\mathbf{X}^T)^T$. The estimator $\hat{\boldsymbol{\gamma}}^{(a)}_{\rm BR,logit}$ is obtained by maximizing the function $\mathscr{F}^{(a)}_{\rm logit}(\boldsymbol{\gamma})$ which is defined like $\mathscr{F}^{(a)}_{\rm lin}(\boldsymbol{\gamma})$ but multiplied by $\{1-m^{(a)}_{\rm logit}(\mathbf{X};\boldsymbol{\alpha}^{(a)})\}m^{(a)}_{\rm logit}(\mathbf{X};\boldsymbol{\alpha}^{(a)})$ within the $\tilde{E}_n$-operator. Because $\mathscr{F}^{(a)}_{\rm logit}(\boldsymbol{\gamma})$ depends on $\boldsymbol{\alpha}^{(a)}$, these functions are maximized using the MLE estimator for $\boldsymbol{\alpha}^{(a)}$. DR estimators for the additive treatment effect τ are then obtained as $\hat{\tau}_{\rm BR,lin} = \hat{\mu}^{(1)}_{\rm DR}(\hat{\boldsymbol{\gamma}}^{(1)}_{\rm BR,lin}, \hat{\boldsymbol{\alpha}}^{(1)}_{\rm BR,lin}) - \hat{\mu}^{(0)}_{\rm DR}(\hat{\boldsymbol{\gamma}}^{(0)}_{\rm BR,lin}, \hat{\boldsymbol{\alpha}}^{(0)}_{\rm BR,lin})$ and $\hat{\tau}_{\rm BR,logit} = \hat{\mu}^{(1)}(\hat{\boldsymbol{\gamma}}^{(1)}_{\rm BR,logit}, \hat{\boldsymbol{\alpha}}^{(1)}_{\rm BR,logit}) - \hat{\mu}^{(0)}(\hat{\boldsymbol{\gamma}}^{(0)}_{\rm BR,logit}, \hat{\boldsymbol{\alpha}}^{(0)}_{\rm BR,logit})$.

The estimators of the PS are different, albeit similar, when estimating $E\{Y(1)\}$ versus $E\{Y(0)\}$. They reveal sufficient overlap of the PS distributions in the RHC group and the no-RHC group (see Figure 1 in Appendix L of the online supplemental material). Figure 2 in Appendix L of the online supplemental material shows that inverse probability of treatment weighting balances the RHC group and the no-RHC group very well. Below we summarize the data analysis results. We obtain an unadjusted effect estimate $\hat{\tau}_{\rm unadj} = -0.0736$ (SE $= 0.0272$, 95% CI -0.1269 to -0.0203) which is prone to potential confounding. The standard DR estimate for the average treatment effect using MLE for all working models equals $\hat{\tau}_{\rm MLE,lin} = -0.0649$ (SE $= 0.0162$, 95% CI -0.0966 to -0.0332) and $\hat{\tau}_{\rm MLE,logit} = -0.0657$ (SE $= 0.0158$, 95% CI -0.0967 to -0.0346). The biased-reduced DR estimation gives more efficient results: we obtain $\hat{\tau}_{\rm BR,lin} = -0.0612$ (SE $= 0.0141$, 95% CI -0.0889 to -0.0335, $\mathrm{avar}(\hat{\tau}_{\rm MLE,lin})/\mathrm{avar}(\hat{\tau}_{\rm BR,lin}) = 1.32$) and $\hat{\tau}_{\rm BR,logit} = -0.0610$ (SE $= 0.0137$, 95% CI -0.0879 to -0.0340, $\mathrm{avar}(\hat{\tau}_{\rm MLE,logit})/\mathrm{avar}(\hat{\tau}_{\rm BR,logit}) = 1.33$). Results for the other improved DR estimators are similar, but less efficient. For example, the calibrated likelihood estimator of Tan gives $\hat{\tau}_{\rm TAN} = -0.0622$ (SE $= 0.0154$, 95% CI -0.0924 to -0.0319) and the TMLE with default super learner gives $\hat{\tau}_{\rm TMLE\text{-}SL} = -0.0586$ (SE $= 0.0149$, 95% CI -0.0877 to -0.0295). Over the different DR methods, the estimates of $E\{Y(1)\}$ range from 0.630 to 0.634 and the estimates of $E\{Y(0)\}$ vary from 0.687 to 0.696.

7. DISCUSSION

In this article, we have proposed a novel strategy for estimating the nuisance parameters indexing the working models in DR estimators. A defining property of the proposed bias-reduced DR estimation strategy is that it locally minimizes the squared first-order asymptotic bias of the DR estimator defined by finite-dimensional nuisance working models. It also makes the DR estimator insensitive to local (one over root n) perturbations of the nuisance parameters. This gets for instance reflected in improved stability of the weights in those DR estimators that invoke inverse weighting. A corresponding efficiency benefit is hence logically anticipated. Formalizing this is, however, complicated by the fact that the choice of root-n estimators for the nuisance parameters affects the asymptotic distribution of the DR estimator not only through their own asymptotic distribution, but also through their probability limits, which can be different for each choice of estimator under model misspecification. In future work, we hope to develop further insight into the theoretical properties of bias-reduced DR estimators as well as confidence intervals obtained by inverting score tests based on this strategy.

The principle of the bias-reduced DR estimator is easy to use and adapts to a wide variety of DR estimators. In that sense, it differs from the various other targeted proposals that have been made over recent years (Cao, Tsiatis, and Davidian 2009; Tan 2010; van der Laan and Gruber 2010; Tsiatis, Davidian, and Cao 2011; van der Laan and Rose 2011; Rotnitzky et al. 2012; van der Laan 2014), some of which are not straightforward or even impossible to adapt to general DR estimators (for instance when the observed data likelihood does not factorize). The simplicity of the proposed approach not only comes through the fact that the estimating functions for the nuisance parameters are readily obtained as gradients of the DR influence function under fixed nuisance parameters, but also through the fact that the asymptotic variance calculation of the resulting DR estimator (and corresponding score tests) can ignore estimation of the nuisance parameters. While the proposed approach is expected to yield estimators with reasonable precision, it does not guarantee minimal variance, unlike other proposals in certain settings (Cao, Tsiatis, and Davidian 2009). In Appendix I of the online supplemental material, we show, however, that the requirement of minimal variance generally leads to complex constrained optimization problems (unless for instance when the PS model is assumed to be correctly specified).

The proposed approach may more generally lend itself better to small-sample inference. For instance, suppose that interest lies in the marginal causal effect $\tau = E\{Y(1) - Y(0)\}$. Because the proposed estimation strategy does not require acknowledging the uncertainty of the estimated nuisance parameters (up to first order), we foresee that it may potentially lend itself better to randomization inference (e.g., permutation tests). How such randomization inference could be accomplished and how it performs in small to large samples will be studied in future work.

A limitation of the bias-reduced DR estimator is that it demands working models of the same dimension. This can in principle be remedied by enlarging the working models with clever choices of covariates until they are of the same dimension. For example, reconsider the DR estimator (1) from Section 2 with working models $m(X;\boldsymbol{\beta}) = \beta_1 + \beta_2 X + \beta_3 X^2$ and $\pi(X;\boldsymbol{\gamma}) = \gamma_1 + \gamma_2 X$ for a one-dimensional covariate X. Taking the gradients of the influence function would lead to two estimating functions for $\boldsymbol{\beta} = (\beta_1, \beta_2, \beta_3)^T$ and three estimating functions for $\boldsymbol{\gamma} = (\gamma_1, \gamma_2)^T$. An additional estimating function for $\boldsymbol{\beta}$ can be obtained by using the extended PS model $\pi(X;\boldsymbol{\gamma}) = \text{expit}\{\gamma_1 + \gamma_2 X + \gamma_3 \zeta(X)\}$ for a cleverly chosen covariate $\zeta(X)$. The proposal then amounts to solving the estimating equations $\tilde{E}_n\{\partial\phi(\mathbf{Z};\mu_0,\boldsymbol{\gamma},\boldsymbol{\beta})/\partial\boldsymbol{\beta}\} = \tilde{E}_n[\{1 - R/\pi(X;\boldsymbol{\gamma})\}(1, X, X^2)^T] = \mathbf{0}$ for $\boldsymbol{\gamma}$ and $\tilde{E}_n\{\partial \phi(\mathbf{Z};\mu_0,\boldsymbol{\gamma},\boldsymbol{\beta})/\partial\boldsymbol{\gamma}\} = \tilde{E}_n(A[\{1 - \pi(X;\boldsymbol{\gamma})\}/\pi(X;\boldsymbol{\gamma})]\{Y - m(X;\boldsymbol{\beta})\}\{1, X, \zeta(X)\}^T) = \mathbf{0}$ for $\boldsymbol{\beta}$. A clever choice for $\zeta(X)$ is $1/\{1 - \pi(X;\hat{\boldsymbol{\gamma}}_{\text{BR}})\}$. Indeed, this choice ensures that $\tilde{E}_n[A\{Y - m(X;\hat{\boldsymbol{\beta}}_{\text{BR}})\}/\pi(X;\hat{\boldsymbol{\gamma}}_{\text{BR}})\}] = 0$, making the bias-reduced DR estimator $\hat{\mu}_{\text{DR}}(\hat{\boldsymbol{\gamma}}_{\text{BR}}, \hat{\boldsymbol{\beta}}_{\text{BR}})$ equal to a substitution estimator $\tilde{E}_n\{m(X;\hat{\boldsymbol{\beta}}_{\text{BR}})\}$. An alternative possibility to cope with nuisance parameters of different dimensions would be to apply the procedure in the direction of just a single nuisance parameter, rather than both.

SUPPLEMENTAL MATERIAL

This file (Appendices A–L) contains proofs, discussions on regularity conditions, extensions, and comments to certain sections, two R-functions, additional simulation results and supporting figures for the data-analysis.

[Received February 2012. Revised July 2014.]

REFERENCES

Bang, H., and Robins, J. M. (2005), "Doubly Robust Estimation in Missing Data and Causal Inference Models," *Biometrics*, 61, 962–972. [1024,1026]

Cao, W. H., Tsiatis, A. A., and Davidian, M. (2009), "Improving Efficiency and Robustness of the Doubly Robust Estimator for a Population Mean With Incomplete Data," *Biometrika*, 96, 723–734. [1025,1029,1034,1035]

Connors, A. F., Speroff, T., Dawson, N. V., Thomas, C., Harrell, F. E., Wagner, D., Desbiens, N., Goldman, L., Wu, A. W., Califf, R. M., William, J. F., Vidaillet, H., Broste, S., Bellamy, P., Joanne, L., and Knaus, W. A. (1996), "The Effectiveness of Right Heart Catheterization in the Initial Care of Critically Ill Patients," *Journal of the American Medical Association*, 276, 889–897. [1034]

Cox, D. R. (1980), "Local Ancillarity," *Biometrika*, 67, 279–286. [1027]

Goetgeluk, S., Vansteelandt, S., and Goetghebeur, E. (2008), "Estimation of Controlled Direct Effects," *Journal of the Royal Statistical Society*, Series B, 70, 1049–1066. [1024]

Hirano, K., and Imbens, G. W. (2002), "Estimation of Causal Effects Using Propensity Score Weighting: An Application of Data on Right Heart Catheterization," *Health Services and Outcomes Research Methodology*, 2, 259–278. [1034]

Kang, J. D. Y., and Schafer, J. L. (2007a), "Demystifying Double Robustness: A Comparison of Alternative Strategies for Estimating a Population Mean From Incomplete Data," *Statistical Science*, 22, 523–539. [1024,1025,1029,1030]

——— (2007b), "Rejoinder: Demystifying Double Robustness: A Comparison of Alternative Strategies for Estimating a Population Mean From Incomplete Data," *Statistical Science*, 22, 574–580. [1024]

Kott, P., and Liao, D. (2012), "Providing Double Protection for Unit Nonresponse With a Nonlinear Calibration-Weighting Routine," *Survey Research Methods*, 6, 105–111. [1028]

McCaffrey, D., Ridgeway, G., and Morral, A. (2004), "Propensity Score Estimation With Boosted Regression for Evaluating Causal Efects in Observational Studies," *Psychological Methods*, 9, 403–425. [1029]

Moore, K. L., and van der Laan, M. J. (2009), "Covariate Adjustment in Randomized Trials with Binary Outcomes: Targeted Maximum Likelihood Estimation," *Statistics in Medicine*, 28, 39–64. [1024]

Newey, W. K., and McFadden, D. (1994), "Large Sample Estimation and Hypothesis Testing," *Handbook of Econometrics*, 4, 2111–2245. [1027]

Okui, R., Small, D. S., Tan, Z., and Robins, J. M. (2012), "Doubly Robust Intrumental Variable Regression," *Statistica Sinica*, 22, 173–205. [1024]

Ridgeway, G., and McCaffrey, D. (2007), "Comment: Demystifying Double Robustness: A Comparison of Alternative Strategies for Estimating a Population Mean for Incomplete Data," *Statistical Science*, 22, 540–543. [1024]

Robins, J. M. (1999), "Association, Causation, and Marginal Structural Models," *Synthese*, 121, 151–179. [1024]

Robins, J. M., Hernán, M. A., and Brumback, B. (2000), "Marginal Structural Models and Causal Inference in Epidemiology," *Epidemiology*, 11, 550–560. [1033]

Robins, J., Li, L., Tchetgen Tchetgen, E., and van der Vaart, A. (2008), "Higher Order Influence Functions and Minimax Estimation of Nonliniear Functionals," *IMS Lecture NotesMonograph Series Probability and Statistics Models: Essays in Honor of David A. Freedman*, 2, 335–421. [1027,1028]

Robins, J. M., Mark, S. D., and Newey, W. K. (1992), "Estimating Exposure Effects by Modeling the Expectation of Exposure Conditional on Confounders," *Biometrics*, 48, 479–495. [1033]

Robins, J. M., and Rotnitzky, A. (2001), Comment on the Bickel and Kwon Article, "Inference for Semiparametric Models: Some Questions and an Answer," *Statistica Sinica*, 11, 920–936. [1024,1026]

Robins, J. M., Rotnitzky, A., and Zhao, L. P. (1994), "Estimation of Regression-Coefficients When Some Regressors are Not Always Observed," *Journal of the American Statistical Association*, 89, 846–866. [1025,1027]

Robins, J. M., Sued, M., Lei-Gomez, Q., and Rotnitzky, A. (2007), "Comment: Performance of Double-Robust Estimators When "Inverse Probability" Weights are Highly Variable," *Statistical Science*, 22, 544–559. [1024,1028,1030]

Rotnitzky, A., Li, L. L., and Li, X. C. (2010), "A Note on Overadjustment in Inverse Probability Weighted Estimation," *Biometrika*, 97, 997–1001. [1028]

Rotnitzky, A., Lei, Q. H., Sued, M., and Robins, J. M. (2012), "Improved Double-Robust Estimation in Missing Data and Causal Inference Models," *Biometrika*, 99, 439–456. [1025,1027,1035]

Rubin, D. B., and van der Laan, M. J. (2008), "Empirical Efficiency Maximization: Improved Locally Efficient Covariate Adjustment in Randomized Experiments and Survival Analysis," *International Journal of Biostatistics*, 4, 1–40. [1025]

Scharfstein, D. O., Rotnitzky, A., and Robins, J. M. (1999a), "Adjusting for Nonignorable Drop-Out Using Semiparametric Nonresponse Models," *Journal of the American Statistical Association*, 94, 1096–1120. [1024,1025,1033]

——— (1999b), "Adjusting For Nonignorable Drop-Out Using Semiparametric Nonresponse Models Rejoinder," *Journal of the American Statistical Association*, 94, 1135–1146. [1031,1033]

Tan, Z. (2006), "A Distributional Approach for Causal Inference Using Propensity Scores," *Journal of the American Statistical Association*, 101, 1619–1637. [1025]

——— (2007), "Comment: Understanding OR, PS and DR," *Statistical Science*, 22, 560–568. [1024]

——— (2010), "Bounded, Efficient and Doubly Robust Estimation With Inverse Weighting," *Biometrika*, 97, 661–682. [1025,1028,1029,1035]

Tchetgen Tchetgen, E. J., and Rotnitzky, A. (2011), "On Protected Estimation of an Odds Ratio Model With Missing Binary Exposure and Confounders," *Biometrika*, 98, 749–754. [1024]

Tchetgen Tchetgen, E. J., and Shpitser, I. (2012), "Semiparametric Theory for Causal Mediation Analysis: Efficiency Bounds, Multiple Robustness, and Sensitivity Analysis," *The Annals of Statistics*, 40, 1816–1845. [1024]

Tsiatis, A. A. (2006), *Semiparametric Theory and Missing Data*, New York: Springer. [1024,1025]

Tsiatis, A. A., and Davidian, M. (2007), "Comment: Demystifying Double Robustness: A Comparison of Alternative Strategies for Estimating a Population Mean From Incomplete Data," *Statistical Science*, 22, 569–573. [1024]

Tsiatis, A. A., Davidian, M., and Cao, W. (2011), "Improved Doubly Robust Estimation When Data are Monotonely Coarsened, With Application to Longitudinal Studies With Dropout," *Biometrics*, 67, 536–545. [1025,1035]

Tsiatis, A. A., Davidian, M., Zhang, M., and Lu, X. (2008), "Covariate Adjustment for Two-Sample Treatment Comparisons in Randomized Clinical Trials: A Principled yet Flexible Approach," *Statistics in Medicine*, 27, 4658–4677. [1024]

van der Laan, M. J. (2014), "Targeted Estimation of Nuisance Parameters to Obtain Valid Statistical Inference," *International Journal of Biostatistics*, 10, 29–57. [1025,1029,1035]

van der Laan, M. J., and Gruber, S. (2010), "Collaborative Double Robust Targeted Maximum Likelihood Estimation," *International Journal of Biostatistics*, 6. doi: 10.2202/1557-4679.1181 [1025,1035]

van der Laan, M. J., and Rose, S. (2011), *Targeted Learning, Causal Inference for Observational and Experimental Data*, New York: Springer. [1025,1029,1035]

Vansteelandt, S., Bekaert, M., and Claeskens, G. (2012), "On Model Selection and Model Misspecification in Causal Inference," *Statistical Methods in Medical Research*, 21, 7–30. [1025,1028,1029]

Vansteelandt, S., Vanderweele, T. J., Tchetgen Tchetgen, E. J., and Robins, J. M. (2008), "Multiply Robust Inference for Statistical Interactions," *Journal of the American Statistical Association*, 103, 1693–1704. [1024,1034]