Statistica Sinica Preprint No: SS-2019-0265

Title	A stable and more efficient doubly robust estimator
Manuscript ID	SS-2019-0265
URL	http://www.stat.sinica.edu.tw/statistica/
DOI	10.5705/ss.202019.0265
Complete List of Authors	Min Zhang and Baqun Zhang
Corresponding Author	Min Zhang
E-mail	mzhangst@umich.edu

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

Statistica Sinica

A stable and more efficient doubly robust estimator

Min Zhang[1] and Baqun Zhang[2]

[1] *University of Michigan, Ann Arbor, and* [2] *Shanghai University of Finance and Economics*

Abstract:

Under the assumption of missing at random, doubly robust (DR) estimators are consistent when either the propensity score or the outcome model is correctly specified. However, despite its appealing theoretic properties, it has been show that the usual augmented inverse probability weighted (IPW) DR estimator may exhibit unsatisfying behavior. We propose an alternative DR method for mean estimation. In this method, we do not directly weight outcomes by the inverse of the estimated propensity scores. Instead we use a nonparametric kernel regression to model the residuals from an outcome regression model as a function of propensity scores. The proposed method does not suffer from the instability of the usual IPW estimator in the event of small estimated propensities. We show that, asymptotically, the new estimator has the double robustness property. Moreover, we show that it is guaranteed to be more efficient than the usual augmented IPW DR estimator when the propensity score model is correct, but the outcome model is incorrect. Our simulation studies show that its finite-sample performance improves upon that of existing DR estimators.

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

Key words and phrases: Causal inference; Comparative effectiveness; Inverse probability weighting; Kernel regression; Missing data; Propensity score.

1 Introduction

Contending with missing data is a common problem in many settings. It is accepted that not properly accounting for missing data can lead to severely biased estimations and invalid inferences. Missing data problems have been an area of active research [Rosenbaum and Rubin, 1984, Scharfstein et al., 1999, Lunceford and Davidian, 2004]. Much of the literature has focused on the situation where the missingness can be assumed to be missing at random (MAR); that is, conditional on the observed variables, the probability of missingness does not depend on the variables that are missing [Little and Rubin, 2019]. In general, methods for dealing with missing data when MAR holds fall into three categories: one models the outcome as a function of covariates, one models the probability of missingness, that is, the propensity score, as a function of covariates [Rosenbaum and Rubin, 1984, Rosenbaum, 1987], or both [Lunceford and Davidian, 2004, Bang and Robins, 2005]. These approaches each have their own advantages and disadvantages in terms of bias, efficiency, robustness, and numerical stability. In general, in methods where only the outcome or the propensity score is modeled, a valid statistical inference depends on the correct specification

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

of the corresponding model, and an incorrect model may lead to inconsistent estimations and invalid inferences. In this sense, these methods are not robust. The so-called doubly robust (DR) methods, where both the outcomes and the propensities are modeled, lead to consistent estimations as long as one of the models, but not necessarily both, is correctly specified [Bang and Robins, 2005], overcoming the issue of nonrobustness to some degree. In addition, DR methods usually have good efficiency properties and achieve the semiparametric efficiency bound if the outcome regression model is correct. DR methods seem to combine the strengths of methods that model either outcomes or propensities, at least theoretically based on asymptotic theory, and are very appealing.

DR estimators have received a lot of attention in the literature, and many different versions have been proposed. In an attempt to demystify double robustness, Kang and Schafer [2007] reviewed several versions of DR estimators and compared them to alternative methods for estimating the population mean when the outcomes are subject to missingness. They found that, although theoretically appealing, DR estimators may have "disastrous" performance when some estimated propensities are small. Kang and Schafer [2007] was followed by several discussion papers, including Robins et al. [2007]. Following the work of Kang and Schafer [2007], there has

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

been considerable interest in improving the usual DR estimators [Tan, 2006, 2007, Cao et al., 2009, Imai and Ratkovic, 2014, Zubizarreta, 2015], particularly, for the setting originally designed by Kang and Schafer [2007]. Most of these efforts have focused on the problem of estimating the population mean when some responses are missing. The methods of Tan [2006, 2007] and Cao et al. [2009] focus on modifying the estimation of the outcome regression model used in the augmentation term. In contrast, the methods of Imai and Ratkovic [2014] and Zubizarreta [2015] attempt to improve performance by seeking a better and more stable way of estimating the weight used in the augmented inverse probability (or propensity) weighted (IPW) DR estimator. These methods can be viewed as modified versions of augmented IPW methods along the line of the original DR estimator.

Approaching the problem from a different perspective, we propose an alternative DR method for mean estimation, where the estimated propensity scores are not used directly for weighting. In contrast to previous methods, we do not change how the outcome regression model or the propensity score model/weights are estimated. We show that, asymptotically, the new estimator has the double robustness property, and it exhibits better efficiency and finite-sample performance than that of existing DR estimators. Interestingly, although the proposed method is not motivated from

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

the perspective of an augmented IPW estimator, we show that the estimator has a close connection with the usual augmented IPW estimator. It has long been noted in the literature that an IPW estimator, including the usual DR estimators, may perform unsatisfactorily in the event of small estimated propensity scores and huge weights [Kang and Schafer, 2007]. Efforts have been made to improve the stability of IPW estimators by trimming or smoothing using Bayesian methods [Elliott and Little, 2000, Elliott, 2008, Austin and Stuart, 2015], although not in the context of DR estimators. We show that the proposed estimator can also be viewed as a principled way to smooth over the inverse of the estimated propensities, therefore reducing the impact of huge weights.

In addition to the references discussed earlier, there have been numerous studies on DR estimators and IPW methods in general. For example, Gruber and van der Laan [2010] studied DR estimators using a targeted maximum likelihood estimation and, Zhou et al. [2019] used a penalized spline method to achieve double robustness. For censored data, Chen et al. [2018] used a kernel-based weighting approach to estimate the survival function of medical cost data subject to censoring. Moreover, efforts have been made to improve IPW-based methods by better estimating the propensity scores using machine learning methods. For example, Pirracchio et al.

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

[2015] used super learner and other machine learning methods to estimate the propensity scores, improving the robustness to a model misspecification of the propensity score.

2. Method

2.1 Notation and background

Consider a study with a random sample of n units from an intended population. Ideally, the full data are $(Y_i, X_i), i = 1, \ldots, n$, independent and identically distributed (i.i.d.) across i, where Y_i and X_i are the response and a vector of auxiliary covariates, respectively, for subject i. Suppose the outcome is subject to missingness. Let R_i be an indicator for observing Y_i, with $R_i = 1$ if Y_i is observed, and $R_i = 0$ if missing. Then, the actually observed data are (R_iY_i, R_i, X_i), i.i.d. across i. Interest focuses on estimating the population mean, $E(Y) \equiv \mu$. We assume the missingness is MAR, denoted by $Y \perp\!\!\!\perp R | X$; that is, the missingness is independent of the outcomes given the observed covariates. When we need to make a causal inference on the treatment effect from the observational data, then even if the outcome variable Y is observed for all subjects, it can still be cast as a missing data problem using the framework of counterfactual outcomes. Hereafter, we only discuss estimating μ, recognizing that the proposed method can be

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

applied directly to compare the treatment effects for observational data.

Various methods are available to adjust for missingness, as reviewed by Kang and Schafer [2007]. In general, these methods involve modeling the outcome or the missingness given covariates, or both. Specifically, the outcome regression-based method builds a model for the outcome Y using the covariates X from the observed data, then estimates μ using the average of the predicted values across all subjects from the fitted model. The consistency of the outcome regression estimator relies on the correct specification of the model for $E(Y|X) = m(X;\beta)$. In contrast, another broad class of methodologies models the probability of the nonmissingness given the covariates, that is, $P(R = 1|X)$, referred to as propensity scores. Propensity scores can be estimated by positing, for example, a logistic regression model that specifies $P(R = 1|X) = \exp(X^T\theta)/\{1 + \exp(X^T\theta)\} \equiv \pi(X;\theta)$. After obtaining the estimated propensity scores, one can weight the contribution of each observed outcome by the inverse of the estimated propensity score, referred to as the IPW estimator. In propensity-score based methods, the consistency of the estimation requires that the model for the propensity score be correctly specified.

In a DR estimator, both the outcome and the propensity scores are modeled, and it remains consistent if either one of the models is correctly

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

specified Bang and Robins [2005]. Hence, the DR estimator affords protection against a misspecification of one of the models, a property referred to as double robustness. The usual augmented IPW DR estimator is given by

$$n^{-1}\sum_{i=1}^{n}\Big\{\frac{R_iY_i}{\widehat{\pi}_i}-\frac{R_i-\widehat{\pi}_i}{\widehat{\pi}_i}m(X_i;\widehat{\beta})\Big\}, \tag{2.1}$$

where $\widehat{\pi}_i = \pi(X_i;\widehat{\theta})$, and $\pi(X_i;\theta)$ and $m(X_i;\beta)$ are the specified models for the propensity and the outcome, respectively. Hereafter, we use $\widehat{\pi}_i$ for brevity, or $\pi(X_i;\widehat{\theta})$ if we would like to emphasize dependence on X_i and $\widehat{\theta}$. The unknown parameters β and θ are usually estimated using the maximum likelihood (ML) method. The first term in (2.1) is an IPW estimator, with the inverse of the propensity serving as the weight, and the second part is an augmentation term. The DR estimator also enjoys good efficiency properties. It is usually more efficient than IPW estimators [Lunceford and Davidian, 2004]. Moreover, if $m(X_i;\beta)$ is correctly specified, it has the smallest asymptotic variance among all estimators that are consistent and asymptotically normal when the propensity model is correct; that is, it is semiparametric efficient.

Despite the appealing theoretical properties based on asymptotics, empirical studies show that the usual DR estimator may exhibit poor performance under some situations in practice [Kang and Schafer, 2007]. They note that the usual DR estimator may be severely biased when both speci-

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

fied models are close to the truth but are not completely correct, and may have "disastrous" performance when some of the estimated propensities are small, even if the propensity score model is correctly specified. Alternative DR estimators have been developed, and some are directly targeted at improving the performance of the usual DR estimator. For example, Kang and Schafer [2007] identified some alternative DR estimators, and Tan [2006] studied a likelihood estimator that possesses the DR property and may alleviate some of the problems associated with the usual DR estimator.

As a follow-up work to Kang and Schafer [2007] and Tan [2007], Cao et al. [2009], abbreviated as CTD below, studied alternative DR estimators that aim to improve the efficiency and robustness of existing DR estimators. The main idea of the CTD projection method is based on the novel observation that when the propensity score model is correct, but the outcome model is incorrect, then the usual DR estimator coupled with the ML estimate of β does not achieve the minimal asymptotic variance among all estimators in (2.1) with an augmentation term $m(X_i; \beta^*)$ for any β^*. They then sought to identify an estimator of β that would lead to an estimator of μ that is both doubly robust and achieves the minimal asymptotic variance when the propensity score model is correct, but the outcome model is

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

incorrect. As a result, the corresponding estimator would be more efficient than the usual DR estimator when propensity score model is correct but outcome model is misspecified. They proposed estimating β by solving

$$\sum_{i=1}^{n}\Big[\frac{R_i}{\widehat{\pi}_i}\frac{1-\widehat{\pi}_i}{\widehat{\pi}_i}\Big\{\begin{array}{c} m_\beta(X_i;\beta) \\ \frac{\pi_\theta(X_i;\widehat{\theta})}{1-\widehat{\pi}_i}\end{array}\Big\}\Big\{Y_i - m(X_i;\beta) - c^T\frac{\pi_\theta(X_i;\widehat{\theta})}{1-\widehat{\pi}_i}\Big\}\Big] = 0, \quad (2.2)$$

where c is a vector that needs to be solved jointly with β , $m_\beta = \partial/\partial\beta\{m(X;\beta)\}$, and $\pi_\theta(X;\theta) = \partial/\partial\theta\{\pi(X;\theta)\}$. Suppose the dimensions of β and θ are p and q, respectively. Then m_β and $\pi_\theta(X;\theta)$ are column vectors with dimensions p and q, respectively, and both sides of (2.2) are of dimension $(p+q)$. Simulation studies using the scenarios designed by Kang and Schafer [2007] and Tan [2007] demonstrated that this method does not suffer the difficulties of the usual DR estimators, as observed by Kang and Schafer [2007]. Furthermore, it achieves comparable or improved performance relative to that of existing methods, including the method of Tan [2006], which is actually closely related to the CTD method, as discussed by Cao et al. [2009]. Nevertheless, these nice properties are again based on large sample theory, which may not necessarily translate into good performance in practice. Taking a closer look at (2.2), the estimation of β in the CTD method is intertwined with the estimated propensities $\pi(X_i;\widehat{\theta})$, and huge weights are given to subjects with small $\pi(X_i;\widehat{\theta})$ (propensities). Therefore, we conjec-

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

ture that the good properties of the CTD estimator when the propensity model is correct or close to correct, but the outcome model is wrong, might be achieved at the expense of worse performance when the outcome model is correct in finite samples.

Other efforts have been made to improve DR estimators by improving the estimation of the weights. Imai and Ratkovic [2014] exploited the dual characteristics of the propensity score as a conditional probability and a covariate balancing score, and proposed estimating the propensity scores using the generalized method-of-moments or the empirical likelihood, which they referred to as the covariate balancing propensity score (CBPS) method. Zubizarreta [2015] proposed a method of directly estimating weights by finding the weights of the minimum variance that balance the empirical distribution of the observed covariates, up to prespecified levels, referring to it as the stable balancing weights (SBW) method. Both methods were applied to the augmented DR estimator and evaluated using the Kang and Schafer [2007] scenarios by their respective authors.

2.2 Proposed method

In contrast to previous methods, based on the usual augmented IPW framework, we propose an alternative and improved DR estimator from a different

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

perspective. In our method, we directly address the issue of sensitivity to small estimated propensity scores that result from the inverse propensity score weighting, and our strategy is to not directly inverse weight by the propensity scores, at least not explicitly. The proposed estimator for μ is given by

$$\widehat{\mu} = n^{-1} \sum_{i=1}^{n} \Big[\frac{\sum_{j=1}^{n} R_j \{Y_j - m(X_j;\widehat{\beta})\} K\big(\frac{\widehat{\pi}_j - \widehat{\pi}_i}{h_n}\big)}{\sum_{j=1}^{n} R_j K\big(\frac{\widehat{\pi}_j - \widehat{\pi}_i}{h_n}\big)} + m(X_i;\widehat{\beta}) \Big], \qquad (2.3)$$

where $K(\cdot)$ is a symmetric kernel function in $\mathcal{R}$, h_n is a bandwidth, and $\widehat{\beta}$ is the usual estimator of β in the outcome regression method, which differs from that used in CTD.

To offer some intuition and motivation for this estimator, we provide a heuristic argument as to why the proposed estimator is expected to possess the double-robustness property. The proposed estimator is motivated from estimating the following quantity:

$$E\big[E\big\{Y - m(X;\beta) | R = 1, \pi(X;\theta)\big\} + m(X;\beta)\big]. \qquad (2.4)$$

If $\pi(X;\theta)$ is the true model for the propensity score, Rosenbaum and Rubin (1983) showed that, conditional on the propensity score, missingness is independent of the confounders and outcomes; that is, $R \perp\!\!\!\perp X | \pi(X;\theta_0)$ and $Y \perp\!\!\!\perp R | \pi(X;\theta_0)$, where θ_0 is the truth such that $P(R = 1|X) = \pi(X;\theta_0)$. It

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

then follows that

$$
\begin{aligned}
(2.4) &= E\big[E\big\{Y - m(X;\beta)|\pi(X;\theta_0)\big\} + m(X;\beta)\big] \\
&= E(Y) - E\{m(X;\beta)\} + E\{m(X;\beta)\} \\
&= E(Y) \equiv \mu,
\end{aligned}
$$

where the first equality follows from $R \perp\!\!\!\perp X|\pi(X;\theta_0)$ and $Y \perp\!\!\!\perp R|\pi(X;\theta_0)$. The above result holds regardless of whether $m(X;\beta)$ is the correct model for Y. This suggests that if the propensity score can be correctly estimated, then estimating the quantity (2.4) may lead to a valid estimator for μ, even if the model for the outcome is wrong.

It is also easy to see that

$$
\begin{aligned}
&E\big\{Y - m(X;\beta)|R = 1, \pi(X;\theta)\big\} \\
&= E\big[E\big\{Y - m(X;\beta)|R = 1, X, \pi(X;\theta)\big\}|R = 1, \pi(X;\theta)\big]. \quad (2.5)
\end{aligned}
$$

If $m(X;\beta)$ is the correct specification of $E(Y|X) = E(Y|R = 1, X)$, then we have

$$
(2.5) = E\big\{m(X;\beta) - m(X;\beta)|R = 1, \pi(X;\theta)\big\} = 0,
$$

and as a result, the targeting quantity satisfies

$$
(2.4) \quad = \quad E\big\{0 + m(X;\beta)\big\} = E(Y) \equiv \mu.
$$

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

Therefore, (2.4) is equal to the target μ if either one of the models is correct and, if we can estimate (2.4), then this estimator is expected to be doubly robust. The proposed estimator (2.3) substitutes unknown parameters β and θ by their estimates, and replaces the outer expectation in (2.4) by the sample average, and $E\big\{Y - m(X;\beta)|R=1,\pi(X;\theta)\big\}$ by the nonparametric Nadaraya–Watson kernel estimator (Fan and Gijbels, 1996). Specifically, the Nadaraya–Watson kernel estimator for $E\big\{Y-m(X;\beta)|R=1,\pi(X;\theta)\big\}$ is $\sum_{j=1}^{n} R_j\{Y_j - m(X_j;\beta)\}K\big(\frac{\pi(X_j;\theta)-\pi(X_i;\theta)}{h_n}\big)/\sum_{j=1}^{n} R_j K\big(\frac{\pi(X_j;\theta)-\pi(X_i;\theta)}{h_n}\big)$ if β and θ are known. Under standard conditions usually assumed for $K(u)$, including $\int K(u)du = 1$, $\int uK(u)du = 0$, $h_n \to 0$, and $nh_n \to \infty$, it can be shown that $n^{-1}\sum_{i=1}^{n} \frac{1}{h_n}K(\frac{x-X_i}{h_n}) \xrightarrow{p} f_X(x)$ and $n^{-1}\sum_{i=1}^{n}\{Y_i\frac{1}{h_n}K(\frac{x-X_i}{h_n})\}$ $\xrightarrow{p} E(Y|x)f_X(x)$, where $f_X(x)$ is the density of X. Therefore, the Nadaraya–Watson kernel estimator $\frac{\sum_{i=1}^{n}\{Y_iK(\frac{x-X_i}{h_n})\}}{\sum_{i=1}^{n}K(\frac{x-X_i}{h_n})}$ estimates $E(Y|x)$ and, similarly, one can obtain the Nadaraya–Watson kernel estimator for $E\big\{Y - m(X;\beta)|R = 1,\pi(X;\theta)\big\}$ detailed above.

In contrast, the usual augmented IPW DR estimator, as well as the estimators of CTD and some other alternatives, is based on or equivalent to directly estimating the following quantity instead:

$$E\Big[\frac{R\{Y - m(X;\beta)\}}{\pi(X;\theta)} + m(X;\beta)\Big], \tag{2.6}$$

because the usual IPW DR estimator in (2.1) can be written equivalently

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

as

$$n^{-1}\sum_{i=1}^{n}\Big[\frac{R_i\{Y_i - m(X_i;\widehat{\beta})\}}{\pi(X_i;\widehat{\theta})} + m(X_i;\widehat{\beta})\Big]. \tag{2.7}$$

Comparing the quantity in (2.6) with that in (2.4), the two quantities differ in their first terms inside the expectation. The first term inside the expectation of (2.6) weights the residual for subjects with observed outcomes by the inverse of his/her propensity $\pi(X;\theta)$. In an estimation, even though $\pi(X_i;\theta)$ is bounded away from zero, the estimated $\pi(X_i;\theta)$ can be close to zero, putting huge weights on those individuals. This leads to numeric instability of the estimators based on this quantity. The proposed estimator may alleviate this issue, because propensities are not used directly as weights. Finally, unlike the CTD method, the estimator does not change the way β is estimated in the outcome regression model and, therefore, we anticipate that it will not suffer from degraded performance in finite samples when the outcome model is correctly specified.

Our discussion above focuses on explaining the difference between the proposed estimator and the various versions of augmented IPW DR estimators; that is, they are motivated from directly estimating (2.4) or (2.7). However, taking a closer look at the proposed estimator, we also see a connection with the augmented IPW DR estimators. By interchanging the order of the summation over i and j, the proposed estimator can be written

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

equivalently as

$$\widehat{\mu} = n^{-1}\sum_{j=1}^{n}\Big[R_j\{Y_j - m(X_j;\widehat{\beta})\}\Big\{\sum_{i=1}^{n}\frac{K\big(\frac{\widehat{\pi}_j-\widehat{\pi}_i}{h_n}\big)}{\sum_{j=1}^{n}R_jK\big(\frac{\widehat{\pi}_j-\widehat{\pi}_i}{h_n}\big)}\Big\} + m(X_j;\widehat{\beta})\Big].$$

Comparing this with the usual augmented IPW DR estimator in (2.7), we see that the only difference between the two estimators is in the weight for $R_j\{Y_j - m(X_j;\widehat{\beta})\}$. Specifically, in the usual IPW DR estimator, the weight is directly the inverse of the estimated propensity score $1/\pi(X_j;\widehat{\theta})$. In contrast, the weight in the proposed estimator is $\sum_{i=1}^{n}\frac{K\big(\frac{\widehat{\pi}_j-\widehat{\pi}_i}{h_n}\big)}{\sum_{j=1}^{n}R_jK\big(\frac{\widehat{\pi}_j-\widehat{\pi}_i}{h_n}\big)}$, which can be shown to converge in probability to $1/E\{R|\pi(X;\theta^*)\}$, where θ^* is the limit of $\widehat{\theta}$. When $\pi(X;\theta)$ is the correct model for $R|X$, then θ^* is equal to the truth, θ_0, and $E\{R|\pi(X;\theta^*)\}$ is equal to $\pi(X;\theta_0)$. In this sense, the proposed estimator resembles an IPW DR estimator, although the motivation for this estimator is quite different. It can be viewed as an IPW estimator in which the propensity is being smoothed to achieve more stability, because the weight for a particular subject now depends on all estimated propensities and their absolute difference (or distance) with the propensity score of the subject, instead of depending only on the propensity for a single subject. It is easier to intuitively see how the stability is achieved by comparing (2.3) with (2.7). For the usual DR estimator in (2.7), each $R_i\{Y_i - m(X_i;\widehat{\beta})\}$ is weighted by $1/\pi(X_i;\widehat{\theta})$, and if the estimated propensity is close to zero, then the huge weight on $R_i\{Y_i - m(X_i;\widehat{\beta})\}$ leads

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

to unstable estimate of μ. Ignoring the second term on $m(X_i;\widehat{\beta})$, we can view this estimator as a summation of many spikes around each observed $Y_i - m(X_i;\widehat{\beta})$. For the proposed estimator in (2.3), however, for each i, the first term is a mountain (all observed $Y_j - m(X_j;\widehat{\beta})$, for $j = 1,\ldots,n$, receive a positive weight) concentrated around the observed $Y_i - m(X_i;\widehat{\beta})$, and the closer $\pi(X_j;\widehat{\theta})$ becomes to $\pi(X_i;\widehat{\theta})$, the larger the weight on $Y_j - m(X_j;\widehat{\beta})$ becomes. Visually, ignoring the second term on $m(X_i;\widehat{\beta})$, the proposed estimator is a summation of many mountains and, as a result, is less sensitive to small estimated propensity scores.

As discussed above, one way to intuitively understand the proposed estimator is that it estimates $E\{Y - m(X;\beta)|R = 1, \pi(X;\theta)\}$ in (2.4) using the nonparametric Nadaraya–Watson kernel estimator. The Nadraya–Watson kernel estimator is a special case of a local polynomial estimator, with the order of the polynomial being zero; that is, the local average kernel estimator. It is well known in the kernel regression literature that a local linear kernel estimator (or high-order local polynomial regression) can reduce the asymptotic bias, especially at the boundary, relative to that of the local average estimator. Then, naturally, one may expect that if we instead estimate $E\{Y - m(X;\beta)|R = 1, \pi(X;\theta)\}$ using more refined estimators, say, the local linear kernel estimator, it may lead to better performance. How-

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

ever, this conjecture is not true. We implemented a similar estimator to the proposed estimator, but replaced the Nadaraya–Watson kernel estimator with a local linear estimator. Our simulations (not shown) show that the performance of the version with the local linear estimator is considerably worse across all scenarios, and the performance is very sensitive to the choice of the bandwidth, in contrast to the proposed estimator. Therefore, we did not pursue this estimator further.

2.3 Asymptotic results

We show in the Supplementary Material that under some mild regularity conditions, if the working model for either the outcome or the propensity score is correctly specified, but not necessarily both, then $\widehat{\mu}$ is consistent for μ, and $\sqrt{n}(\widehat{\mu} - \mu)$ converges in distribution to a normal distribution. We assume the standard regularity conditions required for convergence of $\widehat{\beta}$ and $\widehat{\theta}$ under possibly misspecified models [Tsiatis, 2007] and for consistency and asymptotic normality of the nonparametric kernel estimator [Fan and Gijbels, 1996]. We show that $\widehat{\mu}$ is asymptotically linear, and derive its influence function. Define P and P_n as a probability measure and an empirical measure, respectively, that is, $Pf(X) = \int f(x)P(dx)$ and $P_n f(X) = n^{-1}\sum_{i=1}^{n} f(X_i)$, and we denote $G_n = n^{\frac{1}{2}}(P_n - P)$. When at least one of the working models is correctly specified, we have $n^{\frac{1}{2}}(\widehat{\mu}-\mu)$

$$\begin{aligned}
= & \; G_n\Big\{\frac{E[R\{Y-m(X;\beta^*)\}|\pi(X;\theta^*)]}{E\{R|\pi(X;\theta^*)\}}+\frac{R\{Y-m(X;\beta^*)\}}{E\{R|\pi(X;\theta^*)\}}\\
- & \; \frac{RE[R\{Y-m(X;\beta^*)\}|\pi(X;\theta^*)]}{E^2\{R|\pi(X;\theta^*)\}}+m(X;\beta^*)-\mu\Big\}\\
+ & \; \frac{d}{d\theta}|_{\theta=\theta^*}E\Big[\frac{R\{Y-m(X;\beta^*)\}}{E\{R|\pi(X;\theta)\}}\Big]n^{\frac{1}{2}}(\widehat{\theta}-\theta^*)\\
+ & \; \frac{d}{d\beta}|_{\beta=\beta^*}E\Big[\frac{R\{Y-m(X;\beta)\}}{E\{R|\pi(X;\theta^*)\}}+m(X;\beta)\Big]n^{\frac{1}{2}}(\widehat{\beta}-\beta^*)+o_p(1),
\end{aligned}$$

where β^* and θ^* are the limiting values of $\widehat{\beta}$ and $\widehat{\theta}$ respectively. Under suitable regularity conditions and by standard M-estimation theory, $\widehat{\theta}$ and $\widehat{\beta}$ are asymptotic normal. Therefore, $\widehat{\mu}$ is asymptotically normal with mean zero. Suppose a working logistic regression model and a linear model are specified for the outcome and the propensity score, respectively. Then, $n^{\frac{1}{2}}(\widehat{\beta}-\beta^*) = \frac{1}{\sqrt{n}E(X_iX_i^T)}\sum_{i=1}^n X_i(Y_i - X_i^T\beta^*) + o_p(1)$, and $n^{\frac{1}{2}}(\widehat{\theta}-\theta^*) = \frac{1}{\sqrt{n}E[X_iX_i^T\pi(X_i,\theta^*)\{1-\pi(X_i,\theta^*)\}]}\sum_{i=1}^n X_i\{R_i - \text{expit}(X_i^T\theta^*)\} + o_p(1)$. We obtain the influence function of $\widehat{\mu}$ and have $\sqrt{n}(\widehat{\mu}-\mu) = G_n h(R,X,Y;\beta^*,\theta^*) + o_p(1)$, where the influence function

$$\begin{aligned}
& h(R,X,Y;\beta^*,\theta^*) = \frac{E[R\{Y-m(X;\beta^*)\}|\pi(X;\theta^*)]}{E\{R|\pi(X;\theta^*)\}}+\\
& \frac{R\{Y-m(X;\beta^*)\}}{E\{R|\pi(X;\theta^*)\}}-\frac{RE[R\{Y-m(X;\beta^*)\}|\pi(X;\theta^*)]}{E^2\{R|\pi(X;\theta^*)\}}+m(X;\beta^*)-\mu\\
& +\frac{d}{d\theta}|_{\theta=\theta^*}E\Big[\frac{R\{Y-m(X;\beta^*)\}}{E\{R|\pi(X;\theta)\}}\Big]\frac{X\{R-\pi(X,\theta^*)\}}{E[XX^T\pi(X,\theta^*)\{1-\pi(X,\theta^*)\}]}\\
& +\frac{d}{d\beta}|_{\beta=\beta^*}E\Big[\frac{R\{Y-m(X;\beta)\}}{E\{R|\pi(X;\theta^*)\}}+m(X;\beta)\Big]\frac{X(Y-X^T\beta^*)}{E(XX^T)}.
\end{aligned}$$

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

If one or both working models are correctly specified, the influence function can be further simplified. Specifically, if the working model for $E(Y|R=1,X)=m(X;\beta)$ is correctly specified, the influence function $h(R,X,Y;\beta^*,\theta^*)=$

$$\frac{RY}{E\{R|\pi(X;\theta^*)\}}-\frac{R-E\{R|\pi(X;\theta^*)\}}{E\{R|\pi(X;\theta^*)\}}m(X;\beta_0)-\mu$$
$$+\frac{d}{d\beta}|_{\beta=\beta_0}E\Big[\frac{R\{Y-m(X;\beta)\}}{E\{R|\pi(X;\theta^*)\}}+m(X;\beta)\Big]\frac{X(Y-X^T\beta_0)}{\{E(XX^T)\}}.$$

If the model for $P(R=1|X)=\pi(X;\theta)$ is correctly specified, the influence function $h(R,X,Y;\beta^*,\theta^*)=$

$$\Big\{\frac{RY}{\pi(X;\theta_0)}-\frac{R-\pi(X;\theta_0)}{\pi(X;\theta_0)}\big[E\{Y-m(X;\beta^*)|\pi(X;\theta_0)\}+m(X;\beta^*)\big]-\mu\Big\}$$
$$+\frac{d}{d\theta}|_{\theta=\theta_0}E\Big[\frac{R\{Y-m(X;\beta^*)\}}{E\{R|\pi(X;\theta)\}}\Big]\frac{X\{R-\pi(X,\theta_0)\}}{E[XX^T\pi(X,\theta_0)\{1-\pi(X,\theta_0)\}]}.$$

When both working models are correct, the influence function is

$$\Big\{\frac{RY}{\pi(X;\theta_0)}-\frac{R-\pi(X;\theta_0)}{\pi(X;\theta_0)}m(X;\beta_0)-\mu\Big\},$$

where $\pi(X;\theta_0)=P(R=1|X)$, and it is the semiparametric efficient influence function. Theorems 1 and 2 summarize the results described above.

The variance of $\widehat{\mu}$ can be estimated by n^{-1} times the sample variance of $\widehat{h}(R_i,X_i,Y_i;\widehat{\beta},\widehat{\theta})$, $i=1,\ldots,n$, where $\widehat{h}(R_i,X_i,Y_i;\widehat{\beta},\widehat{\theta})$ is defined as above, except that we replace all marginal expectations by the sample averages, and all conditional expectations by the corresponding Nadaraya–Watson

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

kernel estimator. For example, we replace $E(XX^T)$, $E\{R_i|\pi(X_i;\theta^*)\}$, and $E[R_i\{Y_i - m(X_i;\beta^*)\}|\pi(X_i;\theta^*)]$ by $n^{-1}\sum_{j=1}^n X_j X_j^T$, $\frac{\sum_{j=1}^n R_j K\left(\frac{\widehat{\pi}_j - \widehat{\pi}_i}{h_n}\right)}{\sum_{j=1}^n K\left(\frac{\widehat{\pi}_j - \widehat{\pi}_i}{h_n}\right)}$, and $\frac{\sum_{j=1}^n R_j\{Y_j - m(X_j;\widehat{\beta})\} K\left(\frac{\widehat{\pi}_j - \widehat{\pi}_i}{h_n}\right)}{\sum_{j=1}^n K\left(\frac{\widehat{\pi}_j - \widehat{\pi}_i}{h_n}\right)}$, respectively.

Theorem 1. *When at least one of the working models for the propensity score or the outcome is correctly specified, then $\widehat{\mu}$ is consistent for μ and is asymptotically normal with an influence function defined above.*

Theorem 2. *When both working models for the propensity score and the outcome are correctly specified, then $\widehat{\mu}$ attains the semiparametric efficiency bound.*

Asymptotically, the proposed estimator is equivalent to the usual DR estimator when the outcome regression is correct, regardless of whether the propensity score model is correct. As Cao et al. [2009] commented, when the outcome regression model is correct, it will be fruitless to attempt to further improve efficiency; see Tsiatis and Davidian [2007] for a detailed explanation. Therefore, we focus on the case when the propensity score model is correct, but the outcome regression model is incorrect when discussing efficiency. For simplicity, we first assume the propensity score is known and is not estimated, denoted by $\pi_0(X)$. The asymptotical variance of an estimator is proportional to the variance of its influence function. Following the

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

same argument as in Cao et al. [2009, page 726], the asymptotical variance of the usual DR estimator is proportional to $\mathrm{var}\Big\{\frac{RY}{\pi_0(X)}-\frac{R-\pi_0(X)}{\pi_0(X)}m(X;\beta^*)\Big\}$, which is equal to $E\Big\{\frac{1-\pi_0(X)}{\pi_0(X)}\{Y-m(X;\beta^*)\}^2\Big\}+\mathrm{var}(Y)$, and the asymptotical variance of the proposed estimator is proportional to

$$\begin{aligned}&\mathrm{var}\Big\{\frac{RY}{\pi_0(X)}-\frac{R-\pi_0(X)}{\pi_0(X)}\big[m(X;\beta^*)+E\{Y-m(X;\beta^*)|\pi_0(X)\}\big]\Big\}\\ =\;& E\Big\{\frac{1-\pi_0(X)}{\pi_0(X)}\big[Y-m(X;\beta^*)-E\{Y-m(X;\beta^*)|\pi_0(X)\}\big]^2\Big\}+\mathrm{var}(Y).\end{aligned}$$

The first term of the above expression is equal to

$$\begin{aligned}&E\Big\{\frac{1-\pi_0(X)}{\pi_0(X)}\{Y-m(X;\beta^*)\}^2\Big\}+E\Big\{\frac{1-\pi_0(X)}{\pi_0(X)}E^2\{Y-m(X;\beta^*)|\pi_0(X)\}\Big\}\\ &-2E\Big\{\frac{1-\pi_0(X)}{\pi_0(X)}\{Y-m(X;\beta^*)\}E\{Y-m(X;\beta^*)|\pi_0(X)\}\Big\}.\end{aligned}$$

Using $E(\cdot)=E[E\{\cdot|\pi_0(X)\}]$, the last term is equal to $-2E\Big\{\frac{1-\pi_0(X)}{\pi_0(X)}E^2\{Y-m(X;\beta^*)|\pi_0(X)\}\Big\}$. Therefore, the asymptotic variance of the proposed estimator is proportional to $E\Big\{\frac{1-\pi_0(X)}{\pi_0(X)}\{Y-m(X;\beta^*)\}^2\Big\}+\mathrm{var}(Y)-E\Big\{\frac{1-\pi_0(X)}{\pi_0(X)}E^2\{Y-m(X;\beta^*)|\pi_0(X)\}\Big\}$, which is always less than the variance of the usual DR estimator when the outcome regression model is incorrect. When the propensity scores are not known, but the model is correctly specified, then the influence functions of both the original DR estimator and the proposed one have an additional term representing the effect of estimating θ. It is straightforward to check that these two additional terms are equal. Then, by the same argument as above, we can show that the asymptotic variance

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

of the proposed estimator is that of the original DR estimator, minus a nonnegative term. We summarize the result below.

Theorem 3. *When the model for the propensity score is correct, but the model for the outcome is misspecified, the asymptotic variance of* $\widehat{\mu}$ *is no greater than the asymptotic variance of the usual augmented IPW-based DR estimator.*

4. Simulation Studies

We conducted simulation studies to evaluate the performance of the proposed method and compare it with that of the usual augmented IPW DR estimator, the modified DR method (referred to as the CTD method) proposed by Cao et al. [2009], the CBPS method [Imai and Ratkovic, 2014], and the SBW method [Zubizarreta, 2015]. For comparison, we also included the usual outcome regression method, that is, the average of the predictions from the outcome regression model, fitted using the observed data. For the proposed method, the bandwidth needs to satisfy $h_n \to 0$ and $nh_n \to \infty$. To assess how sensitive the method is to different choices of bandwidth, we implemented it with bandwidth $h_n = n^{-1/3}, n^{-1/4}$, and $n^{-1/5}$, where n is the sample size. Bootstraping with 100 bootstrap samples was used to obtain the standard errors. We chose to use bootstrapping to obtain the

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

standard errors for several reasons. For example, some of the comparison methods (CBPS and SBW) do not provide standard error estimates in the original works. Furthermore, bootstrapping allows one to easily account for variability in estimating propensity scores and fitting outcome regression models, regardless of the methods used to fit the models. The latter point is especially important, because flexible nonparametric and machine learning methods are becoming popular in practice for modeling propensity scores and outcomes; in this case implementing the usual standard error estimates by directly estimating the asymptotic variance becomes rather complicated.

In the first set of simulations, we duplicated the scenario originally designed by Kang and Schafer [2007], which has become a standard scenario under which to compare DR estimators [Cao et al., 2009, Zubizarreta, 2015, Imai and Ratkovic, 2014]. Under this scenario, Kang and Schafer [2007] demonstrated that when both the outcome regression model and the propensity score model are incorrect, but nearly perfect in the sense that they look trustworthy based on model diagnostics, the usual DR estimator may be severely biased and unstable. In this scenario, $Z = (Z_1, \ldots, Z_4)^T$ are generated from independent standard normal distributions, and $X = (X_1, \ldots, X_4)^T$ are defined as $X_1 = \exp(Z_1/2), X_2 = Z_2/\{1 + \exp(Z_1)\} +$

$10, X_3 = (Z_1 Z_3/25 + 0.6)^3, X_4 = (Z_2 + Z_4 + 20)^2$. That is, Z and X can be expressed in terms of each other. Outcomes are generated as $Y = 210 + 27.4Z_1 + 13.7Z_2 + 13.7Z_3 + 13.7Z_4 + \epsilon$, where ϵ is standard normal, and the nonmissingness indicator R is generated according to the true propensity score $P(R = 1|Z) = \text{expit}(-Z_1 + 0.5Z_2 - 0.25Z_3 - 0.1Z_4)$. In real data, the covariates seen by data analysts are X. Naturally, a data analyst that only sees X would fit a linear regression model for Y given X, and a logistic regression model for R given X. As illustrated by Kang and Schafer (2007), although misspecified, these models would appear trustworthy and are nearly as correct. Specifically, the misspecified outcome model is $m(X; \beta) = \beta^T(1, X)$, and the misspecified propensity score model is $\pi(X; \theta) = \text{expit}\{\theta^T(1, X)\}$. As in the previous work, we considered sample sizes $n = 200$ and $n = 1000$.

The results for the first set of simulations are shown in Tables 1 and 2. Under this scenario, when both working models are strictly speaking incorrect, but are nearly perfect, the usual DR estimator is extremely unstable and has huge variability, as demonstrated by the Monte Carlo standard deviation and the root mean squared error. None of the other DR estimators exhibit this type of "disastrous" behavior, with the CTD estimator performing best. The proposed estimators perform comparably with other

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

improved DR estimators, and the performance is not sensitive to different choices of bandwidth. Note that when the propensity score model is correct, but the outcome model is incorrect, the SBW method has relatively larger bias than that of other methods and lower coverage probability. This is not surprising, because this method controls for bias under the condition that the covariates are related to the outcome through a generalized additive form, and this condition is not satisfied under this scenario. For the CTD projection method, when $n = 200$ and when both the outcome regression model and the propensity score model are misspecified, the bootstrap method cannot reliably estimate the uncertainty, and the average standard error is much larger than the Monte Carlo standard deviation. This appears to be a finite-sample problem, and when $n = 1000$, the bootstrap standard error performs as expected. This phenomenon was also observed for the second set of simulations discussed below.

The Kang and Schafer scenario was specifically constructed such that the usual DR estimator may have "disastrous" performance. Thus, the results on this scenario may not generalize to other scenarios. For example, the outcome regression model, when incorrect, is only mildly misspecified and, as a result, the outcome regression method only results in slight bias. To supplement the above simulation study, we compared various methods

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

under a second set of simulations, in which the model misspecifications are due to ignoring some important variables. Four covariates were generated, where X_1 was generated as uniform (0, 1), X_2 as standard normal, X_3 as Bernoulli (0.3), and X_4 as lognormal (0,1). Outcomes were generated according to $Y = 2.5 + X_1/2 + X_2 + X_3 + X_4 + \epsilon$, where ϵ follows a standard normal distribution, and R was generated with propensity score expit$(-1 - X_1/2 + X_2 - X_3 + X_4)$. The proportion of missingness is about 60%. In the methods considered, the misspecified outcome regression and the propensity score models are fitted ignoring X_4.

The results for the second scenario are shown in Tables 3 and 4. Both the usual DR estimator and the proposed estimator are consistent when at least one working model is correctly specified, offering more protection against model misspefication than the outcome regression method. When both working models are incorrect, the usual and proposed DR estimators are all biased, but do not show extreme variability, as observed in the first scenario. Except for the situation in which the propensity score model is correct, but the outcome regression model is incorrect, the proposed method performs very similarly to the usual DR method. Consistent with the results summarized in Theorem 3, when the propensity score model is correct, but the outcome regression model is incorrect, the proposed method is more

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

efficient than the usual DR estimator. The the relative efficiency (the ratio of the mean squared error) of the proposed method relative to that of the usual DR estimator is 1.50 and 1.79 for $n = 200$ and $n = 1000$, respectively. Again, the performance is not sensitive to different choices of bandwidth.

The CTD method exhibits unexpected results. When the sample size is small (n=200), compared with other methods, it has considerably larger variability, especially when the model for the outcome regression is correct, regardless of whether the propensity score model is correct or not. Moreover, the bootstrap method cannot estimate the standard error well for the CTD method, and the average of the standard error estimates is significantly greater than the corresponding Monte Carlo standard deviation. Some of the surprising results are due to the finite-sample performance. When the sample size is 1000, we see that the bootstrap standard error estimates for the projection method work well, and when the outcome regression model is correct, the loss of efficiency of the CTD method relative to that of the usual DR estimator is less. In addition, consistent with the asymptotic results, when n=1000, if the propensity score is correct, but the outcome model is incorrect, the projection estimator is more efficient than the usual estimator, butis still slightly less efficient than the proposed method. Note that weighting the estimating equation for the outcome re-

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

Table 1: Simulation setting 1 (n=200). BIAS: average Monte Carlo bias; MCSD: Monte Carlo standard deviation; RMSE: Root mean squared error; SE: average of standard error; CP: coverage probability of 95% confidence interval. Correct or incorrect refer to corresponding specified model (OR: outcome regression; PS: propensity score). Proposed a, b, and c correspond to bandwidths of $n^{-1/3}, n^{-1/4}$, and $n^{-1/5}$, respectively.

METHODS	BIAS	MCSD	RMSE	SE	CP	BIAS	MCSD	RMSE	SE	CP
	OR correct, PS correct					OR correct, PS incorrect				
OR	-0.054	2.59	2.59	2.55	0.95	-0.054	2.59	2.59	2.55	0.95
Usual DR	-0.055	2.59	2.59	2.55	0.95	-0.067	2.58	2.58	2.62	0.95
CTD pj	-0.049	2.60	2.60	2.64	0.95	-0.053	2.59	2.58	2.57	0.95
CBPS	-0.055	2.59	2.59	2.54	0.94	-0.056	2.59	2.59	2.55	0.94
SBW	-0.054	2.59	2.59	2.54	0.95	-0.055	2.59	2.59	2.55	0.95
Proposed a	-0.055	2.59	2.59	2.56	0.95	-0.055	2.59	2.59	2.56	0.95
Proposed b	-0.055	2.59	2.59	2.56	0.95	-0.056	2.59	2.59	2.56	0.95
Proposed c	-0.055	2.59	2.59	2.56	0.95	-0.056	2.59	2.59	2.56	0.95
	OR incorrect, PS correct					OR incorrect, PS incorrect				
OR	-0.65	3.42	3.48	3.28	0.92	-0.65	3.42	3.48	3.28	0.92
Usual DR	0.27	3.53	3.54	3.48	0.95	-5.69	19.78	20.57	14.92	0.94
CTD pj	-0.028	2.66	2.65	2.78	0.95	-0.85	5.01	5.09	10.71	0.996
CBPS	-0.26	3.29	3.30	3.18	0.94	-2.19	3.58	4.20	3.39	0.88
SBW	1.58	3.10	3.48	3.01	0.91	-0.74	3.39	3.47	3.23	0.92
Proposed a	0.49	3.38	3.41	3.32	0.95	-1.79	3.33	3.78	3.26	0.90
Proposed b	0.56	3.29	3.33	3.23	0.94	-1.68	3.30	3.70	3.22	0.90
Proposed c	0.63	3.23	3.28	3.18	0.95	-1.56	3.28	3.63	3.20	0.90

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

Table 2: Simulation setting 1 (n=1000). Entries are as in Table 1.

METHODS	BIAS	MCSD	RMSE	SE	CP	BIAS	MCSD	RMSE	SE	CP
			OR correct, PS correct					OR correct, PS incorrect		
OR	0.04	1.15	1.15	1.14	0.95	0.04	1.15	1.15	1.14	0.95
Usual DR	0.04	1.15	1.15	1.14	0.95	0.056	1.41	1.41	1.29	0.95
CTD pj	0.04	1.15	1.15	1.15	0.95	0.041	1.15	1.15	1.15	0.95
CBPS	0.04	1.15	1.15	1.14	0.95	0.04	1.15	1.15	0.14	0.95
SBW	0.04	1.15	1.15	1.14	0.95	0.04	1.15	1.15	1.14	0.95
Proposed a	0.042	1.15	1.15	1.14	0.95	0.04	1.15	1.15	1.14	0.95
Proposed b	0.042	1.15	1.15	1.14	0.95	0.04	1.15	1.15	1.14	0.95
Proposed c	0.04	1.15	1.15	1.14	0.95	0.04	1.15	1.15	1.14	0.95
			OR incorrect, PS correct					OR incorrect, PS incorrect		
OR	-0.77	1.50	1.68	1.48	0.91	-0.77	1.50	1.68	1.48	0.91
Usual DR	0.11	1.65	1.65	1.55	0.95	-17.92	166.73	167.6	27.04	0.69
CTD pj	0.08	1.15	1.15	1.16	0.95	-1.36	1.28	1.87	1.35	0.83
CBPS	0.14	1.53	1.54	1.46	0.94	-3.61	2.26	4.25	1.74	0.45
SBW	1.58	1.36	2.09	1.34	0.78	-0.80	1.48	1.69	1.47	0.91
Proposed a	0.27	1.50	1.53	1.49	0.94	-2.13	1.42	2.57	1.39	0.66
Proposed b	0.35	1.44	1.48	1.44	0.94	-2.03	1.41	2.47	1.38	0.68
Proposed c	0.45	1.40	1.47	1.41	0.94	-1.89	1.40	2.35	1.38	0.71

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

Table 3: Simulation setting 2 (n=200). Entries are as in Table 1.

METHODS	BIAS	MCSD	RMSE	SE	CP	BIAS	MCSD	RMSE	SE	CP
	OR correct, PS correct					OR correct, PS incorrect				
OR	0.003	0.18	0.18	0.18	0.94	0.003	0.18	0.18	0.18	0.94
Usual DR	0.003	0.20	0.20	0.20	0.94	0.003	0.19	0.19	0.18	0.94
CTD pj	0.011	0.50	0.50	1.3	0.997	0.006	0.28	0.28	0.45	0.99
CBPS	0.003	0.19	0.19	0.19	0.94	0.003	0.18	0.18	0.18	0.94
SBW	0.003	0.18	0.18	0.18	0.93	0.003	0.18	0.18	0.18	0.94
Proposed a	0.002	0.20	0.20	0.19	0.93	0.003	0.19	0.18	0.18	0.95
Proposed b	0.002	0.19	0.19	0.19	0.93	0.003	0.18	0.18	0.18	0.94
Proposed c	0.002	0.19	0.19	0.18	0.93	0.002	0.18	0.18	0.18	0.94
	OR incorrect, PS correct					OR incorrect, PS incorrect				
OR	0.54	0.23	0.58	0.22	0.33	0.54	0.23	0.58	0.22	0.33
Usual DR	0.023	0.24	0.24	0.27	0.94	0.56	0.25	0.61	0.24	0.36
CTD pj	0.032	0.27	0.27	0.63	0.98	0.56	0.40	0.69	0.69	0.89
CBPS	0.10	0.21	0.23	0.22	0.92	0.55	0.24	0.60	0.24	0.36
SBW	0.004	0.18	0.18	0.18	0.94	0.54	0.23	0.59	0.22	0.35
Proposed a	0.029	0.20	0.20	0.20	0.94	0.55	0.25	0.60	0.24	0.36
Proposed b	0.038	0.20	0.20	0.19	0.94	0.55	0.24	0.60	0.23	0.35
Proposed c	0.048	0.19	0.20	0.19	0.94	0.55	0.24	0.60	0.23	0.33

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

Table 4: Simulation setting 2 (n=1000). Entries are as in Table 1.

METHODS	BIAS	MCSD	RMSE	SE	CP	BIAS	MCSD	RMSE	SE	CP
	OR correct, PS correct					OR correct, PS incorrect				
OR	-0.002	0.080	0.080	0.078	0.95	-0.002	0.080	0.080	0.078	0.95
Usual DR	-0.001	0.09	0.09	0.086	0.93	-0.001	0.083	0.082	0.080	0.94
CTD pj	0.0002	0.10	0.10	0.11	0.96	-0.001	0.10	0.10	0.097	0.93
CBPS	-0.001	0.087	0.087	0.083	0.94	-0.001	0.082	0.082	0.079	0.95
SBW	-0.002	0.080	0.080	0.078	0.94	-0.002	0.080	0.080	0.078	0.95
Proposed a	-0.002	0.087	0.087	0.082	0.94	-0.001	0.083	0.083	0.080	0.94
Proposed b	-0.002	0.085	0.085	0.081	0.94	-0.001	0.082	0.081	0.079	0.94
Proposed c	-0.002	0.084	0.084	0.081	0.94	-0.001	0.081	0.081	0.079	0.94
	OR incorrect, PS correct					OR incorrect, PS incorrect				
OR	0.54	0.10	0.55	0.10	0	0.54	0.10	0.55	0.10	0
Usual DR	-0.0004	0.12	0.12	0.11	0.93	0.56	0.11	0.75	0.11	0.00
CTD pj	0.009	0.096	0.096	0.095	0.94	0.56	0.12	0.57	0.12	0.002
CBPS	0.036	0.103	0.109	0.095	0.91	0.56	0.11	0.57	0.10	0.001
SBW	-0.0003	0.080	0.080	0.078	0.95	0.55	0.10	0.55	0.10	0.001
Proposed a	0.006	0.088	0.089	0.086	0.94	0.56	0.11	0.57	0.11	0
Proposed b	0.013	0.086	0.087	0.084	0.94	0.56	0.11	0.57	0.10	0
Proposed c	0.021	0.085	0.088	0.083	0.93	0.56	0.11	0.57	0.10	0

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

gression model using the inverse of the square of the propensity score, as shown in (2.2), may lead to quite unstable estimates of the parameters in the outcome regression model, which may explain the unsatisfactory finite-sample performance of the CTD method. Unlike scenario 1, where the SBW method shows relatively larger bias, the SBW method seems to have the best performance under scenario 2. This is expected from proposition 4.1 in (Zubizarreta, 2015), because in scenario 2, the covariates have an additive effect on the outcomes. Compared with other methods, CBPS exhibits relatively larger bias and a lower coverage probability when the propensity score model is correct, but the outcome regression model is incorrect. Overall, the proposed methods have comparable or superior performance in all cases.

Finally, we also implemented the usual augmented IPW DR estimator after trimming the estimated propensity score at and smaller than 0.1. The results and a discussion are given in the Supplementary Material.

5. Discussion

Our work follows those of Tan [2006, 2007], Robins et al. [2007], Cao et al. [2009], Imai and Ratkovic [2014], Zubizarreta [2015], and others, in an effort to improve the original DR estimators so that they do not exhibit the

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

"disastrous" behaviors observed by Kang and Schafer [2007], but do enjoy the appealing double-robustness property.

As is clear from (4), the proposed estimator is motivated from the usual outcome regression approach. However, instead of averaging the predictors from the fitted outcome regression model alone, we further model the expectation of the residuals from the outcome regression model, conditional on the propensity scores, and take an average of the predictions from the residual model. In contrast to the usual augmented IPW DR estimator, where the inverses of the propensity scores are used as weights, in the proposed approach, the propensity score is viewed as a predictor and is conditioned on. Because of this, the proposed method does not suffer from the instability problem in the presence of some very small estimated propensity scores. In terms of stability and bias, our simulation studies show that the proposed estimator behaves similarly to the outcome regression method, an estimator that is typically thought to be stable. In the Kang and Schafer [2007] setting, the usual DR estimator exhibits extreme variability when both working models are only mildly misspecified. However, unlike the outcome regression method, the proposed estimator enjoys the double-robustness property, as shown by asymptotic theory and simulation studies. Interestingly, although the proposed estimator is not

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

developed within the framework of augmented inverse propensity weighted estimators, asymptotically, it has an influence function that belongs to the class of augmented IPW estimators. To the best of our knowledge, this is the first time such a connection has been established explicitly. Because of the connection with augmented IPW estimators, an alternative way to understand the proposed method is to view it as an augmented IPW estimator with smoothed weights. Although this perspective offers an intuitive way to understand the proposed estimator, such a connection is not obvious, and it would be difficult to directly come up with ways to smooth over weights. Quite interestingly, although the proposed estimator is not developed from the the perspective of improving efficiency, it enjoys a nice property similar to that of Cao et al. [2009]. Specifically, we show by asymptotic theory and simulations that the proposed estimator is more efficient than the usual DR estimator when the outcome regression model is incorrect, but the propensity score model is correct, and the improvement in efficiency can be considerable, as demonstrated by our simulations. In terms of performance in finite samples, our simulation studies show that, overall, it has quite nice and stable performance under different sample sizes and across scenarios, whereas other existing modified DR estimators may exhibit relatively large bias and/or less satisfactory finite-sample per-

Statistica Sinica: Preprint
doi:10.5705/ss.202019.0265

formance. Finally, we comment that the proposed method is very easy to implement; example code for implementing the method is available at https://github.com/MinZhangUMBiostat/DoubleRobust.

Supplementary Material

In the online Supplementary Material, we provide proofs for the asymptotic results described in Section 2.3, and briefly summarize the simulation results using the trimming method.

References

Peter C Austin and Elizabeth A Stuart. Moving towards best practice when using inverse probability of treatment weighting (iptw) using the propensity score to estimate causal treatment effects in observational studies. *Statistics in medicine*, 34(28):3661–3679, 2015.

Heejung Bang and James M Robins. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61(4):962–973, 2005.

Weihua Cao, Anastasios A Tsiatis, and Marie Davidian. Improving efficiency and robustness of the doubly robust estimator for a population mean with incomplete data. *Biometrika*, 96 (3):723–734, 2009.

Shuai Chen, Wenbin Lu, and Hongwei Zhao. An improved survival estimator for censored medical costs with a kernel approach. *Communications in Statistics-Theory and Methods*, 47(23):5702–5716, 2018.

Michael R Elliott. Model averaging methods for weight trimming. *Journal of official statistics*, 24(4):517, 2008.

Michael R Elliott and Roderick JA Little. Model-based alternatives to trimming survey weights. *Journal of Official Statistics*, 16(3):191, 2000.

Jianqing Fan and Irene Gijbels. *Local polynomial modelling and its applications: monographs on statistics and applied probability 66*, volume 66. CRC Press, 1996.

Susan Gruber and Mark J van der Laan. A targeted maximum likelihood estimator of a causal effect on a bounded continuous outcome. *The International Journal of Biostatistics*, 6(1), 2010.

Kosuke Imai and Marc Ratkovic. Covariate balancing propensity score. *Journal of the Royal Statistical Society: Series B: Statistical Methodology*, pages 243–263, 2014.

Joseph DY Kang and Joseph L Schafer. Demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data. *Statistical science*, 22(4):523–539, 2007.

Roderick JA Little and Donald B Rubin. *Statistical analysis with missing data*, volume 793. John Wiley & Sons, 2019.

Jared K Lunceford and Marie Davidian. Stratification and weighting via the propensity score in estimation of causal treatment effects: a comparative study. *Statistics in medicine*, 23 (19):2937–2960, 2004.

Romain Pirracchio, Maya L Petersen, and Mark Van Der Laan. Improving propensity score estimators' robustness to model misspecification using super learner. *American journal of epidemiology*, 181(2):108–119, 2015.

James Robins, Mariela Sued, Quanhong Lei-Gomez, and Andrea Rotnitzky. Comment: Performance of double-robust estimators when" inverse probability" weights are highly variable. *Statistical Science*, 22(4):544–559, 2007.

Paul R Rosenbaum. Model-based direct adjustment. *Journal of the American Statistical Association*, 82(398):387–394, 1987.

Paul R Rosenbaum and Donald B Rubin. Reducing bias in observational studies using subclassification on the propensity score. *Journal of the American statistical Association*, 79(387): 516–524, 1984.

Daniel O Scharfstein, Andrea Rotnitzky, and James M Robins. Adjusting for nonignorable drop-out using semiparametric nonresponse models. *Journal of the American Statistical Association*, 94(448):1096–1120, 1999.

Zhiqiang Tan. A distributional approach for causal inference using propensity scores. *Journal of the American Statistical Association*, 101(476):1619–1637, 2006.

Zhiqiang Tan. Comment: Understanding or, ps and dr. *Statistical Science*, 22(4):560–568, 2007.

Anastasios Tsiatis. *Semiparametric theory and missing data*. Springer Science & Business Media, 2007.

Anastasios A Tsiatis and Marie Davidian. Comment: Demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data. *Statistical science: a review journal of the Institute of Mathematical Statistics*, 22 (4):569, 2007.

Tingting Zhou, Michael R Elliott, and Roderick JA Little. Penalized spline of propensity methods for treatment comparison. *Journal of the American Statistical Association*, 114 (525):1–19, 2019.

José R Zubizarreta. Stable weights that balance covariates for estimation with incomplete outcome data. *Journal of the American Statistical Association*, 110(511):910–922, 2015.

Department of Biostatistics,

University of Michigan, Ann Arbor, USA.

E-mail: mzhangst@umich.edu

School of Statistics and Management,

Shanghai University of Finance and Economics

E-mail: zhang.baqun@mail.shufe.edu.cn