

ORIGINAL ARTICLE

How should we estimate inverse probability weights with possibly misspecified propensity score models?

Hiroto Katsumata 

Institute of Social Science, The University of Tokyo, Tokyo, Japan
Email: hrt.katsumata@gmail.com

(Received 28 September 2022; revised 22 December 2023; accepted 12 March 2024)

Abstract

Inverse probability weighting is a common remedy for missing data issues, notably in causal inference. Despite its prevalence, practical application is prone to bias from propensity score model misspecification. Recently proposed methods try to rectify this by balancing some moments of covariates between the target and weighted groups. Yet, bias persists without knowledge of the true outcome model. Drawing inspiration from the quasi maximum likelihood estimation with misspecified statistical models, I propose an estimation method minimizing a distance between true and estimated weights with possibly misspecified models. This novel approach mitigates bias and controls mean squared error by minimizing their upper bounds. As an empirical application, it gives new insights into the study of foreign occupation and insurgency in France.

Keywords: causal inference; difference of the convex functions algorithm; inverse probability weighting; Kullback–Leibler divergence; missing data

1. Introduction

Causal inference in social sciences stands out as an increasingly attractive field. Currently, it is acknowledged as a missing data problem, where one of the potential outcomes is missing for each unit (Little and Rubin, 2019). This problem is widely addressed by inverse probability weighting (IPW) based on the propensity score estimation under the conditional ignorability assumption. In a review of papers in the *American Journal of Political Science*, the *American Political Science Review*, and the *Journal of Politics* from 2000 to 2022,¹ I found that 35 articles employ the IPW. Notably, 30 of them appear in the past five years, making it the most widely used method among the popular weighting or matching methods, including the entropy balancing (Hainmueller, 2012), propensity score matching, and genetic matching (Diamond and Sekhon, 2013), in each of the five years. Despite its increasing popularity, however, it is vulnerable to propensity score model misspecification (Kang and Schafer, 2007; Imai and Ratkovic, 2014), which is rather common in practice.

To mitigate the bias due to model misspecification, existing studies have proposed various methods balancing some moments (or kernels) of observed covariates between the target and weighted groups (Hainmueller, 2012; Imai and Ratkovic, 2014; Vermeulen and Vansteelandt, 2015; Zubizarreta, 2015; Chan *et al.*, 2016; Wong and Chan, 2017; Zhao, 2019; Hazlett, 2020; Kallus, 2020; Tan, 2020; Wang and Zubizarreta, 2020). However, the efficacy of these methods is contingent upon stringent assumptions. Specifically, removing bias by balancing specific

¹The details are shown in Supplementary Material A.

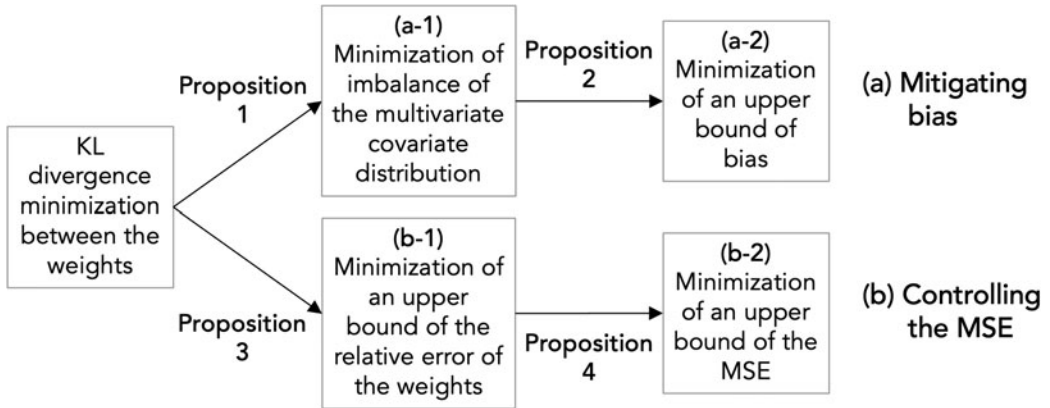


Figure 1. Goals of the proposed method.

moments requires prior knowledge of the true outcome model (Zubizarreta, 2015; Zhao and Percival, 2017; Zhao, 2019; Fan *et al.*, 2021), which is an impractical expectation in the face of real-world scenarios characterized by nonlinearity, interactions, and modifications. Instead, I propose to balance the *multivariate* confounder distribution. This novel perspective, while acknowledged in previous works (Hainmueller, 2012; Zubizarreta, 2015; Li *et al.*, 2018), remains largely unexplored.

In this paper, I propose an estimation method for the IPW with a misspecified propensity score model that mitigates bias and controls the mean squared error (MSE) by minimizing the imbalance in the multivariate covariate distribution rather than only specific moments. The key idea is to use the estimating equations optimized for estimating the (inverse probability) weights rather than the treatment assignment or propensity scores. Specifically, it estimates propensity scores by minimizing the Kullback–Leibler (KL) divergence between the true and estimated weights, which directly quantifies the quality of the IPW estimation. This idea builds upon the quasi maximum likelihood estimation (MLE) with misspecified models, which minimizes the KL divergence between the true and estimated distribution (Akaike, 1973; White, 1982). I demonstrate that the KL divergence between the weights can be calculated up to a constant without requiring knowledge of the true weights or propensity scores and that the proposed method mitigates bias and controls the MSE of parameters of interest by minimizing their upper bounds (Section 3.3 and 3.5).

Figure 1 provides an overview of the proposed method, summarizing its two main goals of the proposed method and two associated intermediate goals for each along with the propositions for the theoretical results. The first main goal is (a) mitigating bias of the parameter of interest. The associated intermediate goals are (a-1) minimizing imbalance in multivariate covariate distribution between the target group and the weighted group and (a-2) minimizing an upper bound of *bias*. The second main goal is (b) controlling the MSE of the parameter of interest. The associated intermediate goals are (b-1) minimizing an upper bound of the relative error of (inverse probability) weights, and (b-2) minimizing an upper bound of the *MSE*. Section 2 presents a numerical example to illustrate that the proposed method achieves these goals better than the MLE.

The proposed method has several additional attractive characteristics. First, the proposed method can incorporate the outcome regression model to substitute for the mean balancing condition (Section 4.3). This allows it to remove bias under the same conditions as the recently proposed covariate balancing methods such as the covariate balancing propensity scores (CBPS) (Imai and Ratkovic, 2014), bias-reduced doubly robust estimator (Vermeulen and Vansteelandt, 2015), calibrated weighting (Tan, 2020), and entropy balancing (Hainmueller, 2012). Second, like the CBPS, it models propensity scores explicitly and estimates them consistently when the model is

correctly specified. While weight calibration methods such as the entropy balancing have an implicit propensity score model (Hainmueller, 2012; Zhao and Percival, 2017; Wang and Zubizarreta, 2020), it is easier for empirical researchers to make a propensity score model by considering the treatment assignment mechanism with their domain knowledge than to incorporate it implicitly through moment conditions. Third, as a direct consequence of these two characteristics, it can become a doubly robust (DR) estimator (Section 4.3) (Seaman and Vansteelandt, 2018). It is consistent either if the propensity score model or the outcome model is correct, and it attains the semiparametric efficiency bound if both the models are correct (Robins *et al.*, 1994).

Despite the theoretical attractiveness, the proposed method has some difficulties in implementation. First, it is not sample-bounded because the estimated weights may not sum to one; second, it is likely to overfit extreme data points; and third, it necessitates the non-convex optimization. I address the first problem by imposing the normalization constraint and deriving the unconstrained optimization problem by the Lagrangian multiplier method (Section 4.1). I address the second problem by utilizing the L2 regularization (Section 4.1). To address the third problem, I develop a majorization-minimization algorithm that iteratively optimizes two convex decompositions of the original function in Section 4.2 (Wu and Lange, 2010).

The proposed method is motivated by an observational study of political devolution and resistance to foreign rule in France during World War II (Ferwerda and Miller, 2014). In Section 6, I apply the proposed method to this study and discover the treatment effect heterogeneity, suggesting that spill-over effects contaminate the treatment effect. I also apply the proposed method to another empirical study in Supplementary Materials H. Section 5 demonstrates the finite-sample performance of the proposed method through simulation studies.

The proposed methodology is implemented via an open-source R package `dbw`, which is available at <https://github.com/hirotokatumata/dbw> and will be soon at CRAN.

2. A numerical example

This section presents a numerical example for illustrating that the proposed method achieves goals (a) and (b), through goals (a-1), and (b-1) better than the MLE, where goals (a-2) and (b-2) are direct consequences of (a-1) and (b-1) as later shown in propositions 2 and 4. It also provides intuition on how the proposed method achieves these goals by showing that it is optimized for estimating inverse probability weights rather than propensity scores.

For simplicity, this example examines a situation where there is only one continuous covariate X_i following the standard normal distribution as well as a binary response variable R_i , indicating that the outcome Y_i is observed when $R_i = 1$ but not observed when $R_i = 0$, for observation $i = 1, 2, \dots, 1000$. The propensity score model is misspecified in the sense that it only includes a linear term for X_i but the true model also includes its squared term X_i^2 . Specifically, the misspecified model is $\Pr(R_i = 1) = \text{expit}(\alpha + \beta X_i)$ whereas the true model is $\Pr(R_i = 1) = \text{expit}(-1 + X + 0.5X^2)$. The inverse probability weights are estimated as $\hat{w}_i = 1/\hat{\Pr}(R_i = 1)$. To consider the possible worst-case scenario for each estimator, I do not specify the data-generating process for the outcome Y_i .

a-1: Distributional imbalance in the covariate. First, I compare the performance of the proposed method with the MLE in terms of distributional imbalance in the covariate X . For this purpose, a commonly used statistic is the Kolmogorov–Smirnov (KS) statistic. The KS statistics are 0.112 for the proposed method and 0.193 for the MLE, indicating that the proposed method has a smaller covariate distributional imbalance than the MLE. To compare the distribution visually, I present a figure showing the cumulative distribution functions of X in Supplementary Material B.

b-1: Relative error of the weights. Next, I compare the proposed method with the MLE in terms of the relative error of the weights. As detailed in Section 3.5, this quantity is defined as $E[(\hat{w}_i/w_i - 1)^2]$, where w_i and \hat{w}_i denote true and estimated weights. The relative error of the

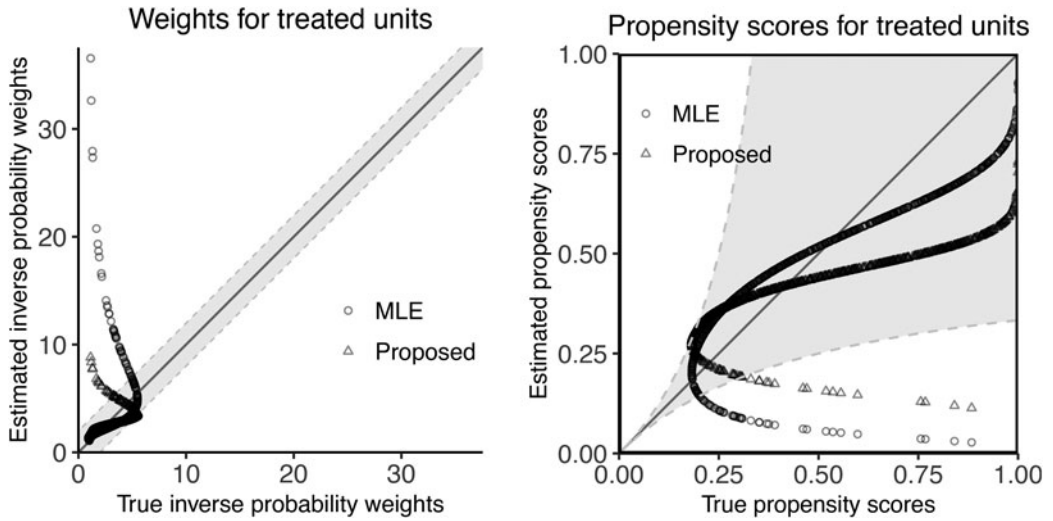


Figure 2. A numerical example: The true and estimated weights and propensity scores with a misspecified propensity score model.

Notes: This figure compares the performance of the proposed estimator (triangles) and MLE estimator (circles) with a misspecified propensity score model in terms of inverse probability weights (left panel) and propensity scores (right panel). In the left (right) panel, the x-axis represents the true weights (propensity scores), and the y-axis represents the estimated ones. The shaded area in the left panel indicates that the difference between the true and estimated weights is small (less than two) whereas the shaded area in the right panel indicates the corresponding area where the difference in weights is small (not the difference in propensity scores).

weights is 0.38 for the proposed method and 4.14 for the MLE, indicating that the relative error of the proposed method is smaller than one-tenth of the MLE.

2.1 Intuition: targeting weights rather than propensity scores

To gain intuition, the left panel of Figure 2 presents the true (the x-axis) and estimated weights (the y-axis) for each observation. The proposed method (triangles) estimates inverse probability weights better than the MLE (circles). Most of the triangles are within or near the shaded narrow corridor around the diagonal line, which indicates that the difference between the true and estimated weights is less than two, whereas many circles lie far from the narrow corridor.

The right panel shows the true (the x-axis) and estimated propensity scores (the y-axis) for each observation. Overall, the MLE (circles) estimates propensity scores better than the proposed method (triangles) but the proposed method restricts the difference more effectively than the MLE when the difference in the resulting inverse probability weights is large shown as the shaded area. This explains why the proposed method approximates the weights better than the MLE.

These results provide intuition when the difference between the two estimators is large: It is when some units have large estimated weights, which is likely to happen when there are unmodeled nonlinear effects of covariates on treatment assignment, there are unmodeled interaction effects of covariates on treatment assignment, or the numbers of the treated and controlled units are unbalanced. In addition, the proposed estimator is especially better than the MLE when the outcomes of these units are extreme (relative to the true average outcome), but not so much when the opposite is true. The difference in the estimated weights between the two estimators is small when no units have large estimated weights.

The root mean squared errors (RMSEs) of the estimated weights are 1.18 for the proposed method and 3.10 for the MLE, and RMSEs of the estimated propensity scores are 0.198 for the proposed method and 0.131 for the MLE, indicating that the proposed method estimates

the weights better than the MLE whereas the MLE estimates the propensity scores better than the proposed method. To estimate the quantity of interest through estimating inverse probability weights, the proposed method is better suited than the MLE.

3. Proposed methodology

3.1 Setup

Suppose that there is a random sample of n units ($i = 1, 2, \dots, n$) from a population. Each unit consists of a triplet (Y_i, R_i, X_i) , where Y_i is an outcome variable, R_i is a binary outcome response indicator variable that takes one when Y_i is observed and zero otherwise, and X_i is a vector of observed covariates whose support is denoted by \mathcal{X} . The propensity score is defined as the conditional response probability given observed covariates and denoted as $\Pr(R_i = 1 | X_i) = \pi(X_i)$. We impose the overlap assumption that the propensity score is bounded away from 0 and 1: $0 < \pi(\mathbf{x}) < 1$ for any $\mathbf{x} \in \mathcal{X}$. We also assume that the outcome is missing at random, i.e., the outcome variable does not account for the response conditional on the observed covariates: $Y_i \perp\!\!\!\perp R_i | X_i$ (Little and Rubin, 2019). These assumptions correspond to the conditional ignorability assumption in causal inference.

For simplicity, the parameter of interest in this study is the average outcome $\mu = \mathbb{E}[Y_i]$ for the target group of $R_i = 1$ and $R_i = 0$, and the proposed method is easily extended to other estimands such as the average treatment effect in causal inference (Kang and Schafer, 2007; Vermeulen and Vansteelandt, 2015; Zubizarreta, 2015; Wong and Chan, 2017; Seaman and Vansteelandt, 2018; Zhao, 2019; Tan, 2020; Wang and Zubizarreta, 2020).

3.2 Problems in the inverse probability weighting with the maximum likelihood estimation

The average outcome can be estimated unbiasedly and consistently by the IPW with true propensity scores: $\sum_{i=1}^n R_i w_i Y_i = \sum_{i=1}^n R_i Y_i / \pi(X_i)$, where $w_i = 1/\pi(X_i)$ is the true inverse probability weight for unit i . In observational studies, however, neither the true propensity scores nor the true inverse probability weights are known. In practice, researchers use a parametric propensity score model $\pi(X_i, \beta)$, where the most common choice is the logistic regression model $\pi(X_i, \beta) = \exp(X_i^T \beta) / (1 + \exp(X_i^T \beta))$. Typically, the nuisance parameters β are estimated using the MLE to obtain the estimated propensity scores for each unit $\hat{\pi}_{MLE}(X_i) = \pi(X_i, \hat{\beta}_{MLE})$. The resulting IPW estimator with the estimated propensity scores is consistent if the propensity score model is correct.

However, if the propensity score model is misspecified, the MLE estimator for β and the resulting IPW estimator is not consistent. Generally, the limiting values of $\hat{\beta}_{MLE}$ are the solution to the following problem (Akaike, 1973; Tan, 2020; White, 1982):

$$\hat{\beta}_{MLE} = \operatorname{argmax}_{\beta} \sum_{i=1}^n \{R_i \log(\pi(X_i, \beta)) + (1 - R_i) \log(1 - \pi(X_i, \beta))\} \tag{1}$$

$$= \operatorname{argmin}_{\beta} \{KL(R, \pi(\mathbf{X}, \beta)) + KL(1 - R, 1 - \pi(\mathbf{X}, \beta))\}. \tag{2}$$

This implies that the MLE estimates $\hat{\beta}_{MLE}$ to minimize the (generalized) KL divergence between the true responses (R and $1 - R$) and their estimates ($\pi(\mathbf{X}, \hat{\beta}_{MLE})$ and $1 - \pi(\mathbf{X}, \hat{\beta}_{MLE})$) for the target group of $R_i = 1$ and $R_i = 0$, where the (generalized) KL divergence is the most common measure of the difference between two points P and Q , which is defined as $KL(P, Q) = \sum_{i=1}^n \{P_i \log(P_i/Q_i) - P_i + Q_i\}$. The problem is that this is not the same as the KL divergence between the true weights and estimated inverse probability weights although the latter determines the performance of the IPW estimator.

3.3 Distribution balancing weighting

This study proposes the distribution balancing weighting (DBW), which estimates such propensity scores $\pi(X_i, \hat{\beta}_{DBW})$ that minimize the KL divergence between the true and the estimated inverse probability weights. The idea of minimizing the KL divergence with possibly misspecified models originates from the quasi MLE, which minimizes the KL divergence between the true and estimated distribution when the true distribution is unknown (Akaike, 1973; White, 1982). Since the propensity scores and their inverse probability weights have a non-linear relationship, minimizing the KL divergence between the true and estimated response assignment like the MLE does not minimize that for the inverse probability weights in general. The DBW applies the idea to the KL divergence between the true and estimated weights and obtains the following KL divergence:

$$\mathbb{E} \left[\text{KL} \left(\frac{R}{\pi(\mathbf{X})}, \frac{R}{\pi(\mathbf{X}, \hat{\beta}_{DBW})} \right) \right] = \mathbb{E} \left[\sum_{i=1}^n \left\{ \frac{R_i}{\pi(X_i, \hat{\beta}_{DBW})} + \log(\pi(X_i, \hat{\beta}_{DBW})) - 1 - \log(\pi(X_i)) \right\} \right]. \tag{3}$$

The DBW estimates coefficients as the minimizers of the following loss function, which is the sample version of the KL divergence above up to a constant:

$$\hat{\beta}_{DBW} = \underset{\beta}{\text{argmin}} \sum_{i=1}^n \left\{ \frac{R_i}{\pi(X_i, \beta)} + \log(\pi(X_i, \beta)) \right\}. \tag{4}$$

As stated formally in proposition 5, the DBW estimates propensity scores and resulting inverse probability weights consistently, like the MLE, when the propensity score model is correctly specified.

3.4 Minimizing an upper bound of bias

The main advantage of the DBW is that it also minimizes the imbalance in the multivariate covariate distribution in the target and weighted groups as stated in proposition 1 below, whose importance has been acknowledged in the existing studies (Hainmueller, 2012; Zubizarreta, 2015; Li *et al.*, 2018; Zhao, 2019). While the propensity scores estimated by the MLE minimize the distributional imbalance in the limit when the propensity score model is correctly specified, it does not when the model is misspecified. The DBW achieves this even when the model is misspecified. The following proposition summarizes this result.

Proposition 1 (Multivariate distribution balancing property): When the propensity scores and resulting inverse probability weights are estimated using (4), the weights minimize the following KL divergence for the multivariate covariate distribution between target and weighted groups in the limit even when the propensity score model is misspecified.

$$\hat{\beta}_{DBW} \xrightarrow{P} \beta_{DBW}^* \tag{5}$$

$$\beta_{DBW}^* = \underset{\beta}{\text{argmin}} \text{KL} \left(f(\mathbf{x}), \frac{f_1(\mathbf{x})}{\pi(\mathbf{X}, \beta)} \right), \tag{6}$$

where β_{DBW}^* is the probability limit of $\hat{\beta}_{DBW}$ and $f(\mathbf{x})$ and $f_1(\mathbf{x})$ is the multivariate covariate distribution for the target group ($\mathcal{S} = \{i \mid R_i \in \{0, 1\}\}$) and the response group ($\mathcal{S}_1 = \{i \mid R_i = 1\}$), respectively.

Proof. For any β , the following equations hold:

$$\sum_{i=1}^n \left\{ \frac{R_i}{\pi(X_i, \beta)} + \log(\pi(X_i, \beta)) \right\} \xrightarrow{P} \int_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) \left\{ \frac{R}{\pi(\mathbf{x}, \beta)} + \log(\pi(\mathbf{x}, \beta)) \right\} d\mathbf{x} \tag{7}$$

$$= \int_{\mathbf{x} \in \mathcal{X}} \left\{ \frac{f_1(\mathbf{x})}{\pi(\mathbf{x}, \beta)} + f(\mathbf{x}) \log(\pi(\mathbf{x}, \beta)) \right\} d\mathbf{x} \tag{8}$$

$$= \text{KL} \left(f(\mathbf{x}), \frac{f_1(\mathbf{x})}{\pi(\mathbf{X}, \beta)} \right) + \text{constant}. \tag{9}$$

By applying the continuous mapping theorem, one can obtain

$$\operatorname{argmin}_{\beta} \sum_{i=1}^n \left\{ \frac{R_i}{\pi(X_i, \beta)} + \log(\pi(X_i, \beta)) \right\} \xrightarrow{P} \operatorname{argmin}_{\beta} \text{KL} \left(f(\mathbf{x}), \frac{f_1(\mathbf{x})}{\pi(\mathbf{X}, \beta)} \right). \tag{10}$$

□

This property is advantageous because it implies that bias is bounded above as stated in the following proposition.

Proposition 2 (Minimizing a sharp upper bound of bias): Let the absolute value of bias be denoted as B and $C = \max(|\mathbb{E}[Y_i - \mu \mid R_i = 1, \mathbf{x}]|)$. A function of imbalance in the multivariate covariate distribution quantified by the KL divergence sharply bounds bias from above as follows:

$$B \leq C \text{KL} \left(f(\mathbf{x}) - \frac{f_1(\mathbf{x})}{\pi(\mathbf{X}, \beta_{\text{DBW}}^*)} \right) \tag{11}$$

Proof.

$$B = \left| \int_{\mathbf{x} \in \mathcal{X}} \mathbb{E}[Y_i - \mu \mid R_i = 1, \mathbf{x}] \left(f(\mathbf{x}) - \frac{f_1(\mathbf{x})}{\pi(\mathbf{X}, \beta_{\text{DBW}}^*)} \right) d\mathbf{x} \right| \tag{12}$$

$$\leq C \int_{\mathbf{x} \in \mathcal{X}} \left| f(\mathbf{x}) - \frac{f_1(\mathbf{x})}{\pi(\mathbf{X}, \beta_{\text{DBW}}^*)} \right| d\mathbf{x} \tag{13}$$

$$\leq C \text{KL} \left(f(\mathbf{x}) - \frac{f_1(\mathbf{x})}{\pi(\mathbf{X}, \beta_{\text{DBW}}^*)} \right), \tag{14}$$

where the last inequality follows from Pinsker’s inequality, and the equalities hold when the propensity score model is correct. □

3.5 Minimizing an upper bound of the mean squared error

This section demonstrates that the proposed method sharply bounds an upper bound of the MSE by showing that the proposed method minimizes a sharp upper bound of the relative error of weights (proposition 3) and the relative error of weights bounds the MSE of the IPW estimator (proposition 4).

First, the KL divergence between the weights, which the proposed method minimizes as a loss function, bounds the relative error of the weights $\xi(\hat{\beta}) = \mathbb{E}[(\pi(X_i)/\pi(X_i, \hat{\beta}) - 1)^2]$ from above as the following proposition states:

Proposition 3 (Minimizing a sharp upper bound of the relative error of weights): When the propensity scores and resulting inverse probability weights are estimated using (4), a sharp upper bound of the relative error of weights is minimized in the limit even when the propensity score model is misspecified. If $\pi(X_i, \hat{\beta}) \geq a\pi(X_i)$ almost surely for some constant $a \in (0, 1]$, the bound is as follows:

$$\xi(\hat{\beta}) \leq 2 \mathbb{E} \left[\text{KL} \left(\frac{R_i}{\pi(X_i)}, \frac{R_i}{\pi(X_i, \hat{\beta})} \right) \right] \quad \text{for } a = 1 \tag{15}$$

$$\xi(\hat{\beta}) \leq \frac{a^2 - 2a + 1}{a(1 - a + a \log(a))} \mathbb{E} \left[\text{KL} \left(\frac{R_i}{\pi(X_i)}, \frac{R_i}{\pi(X_i, \hat{\beta})} \right) \right] \quad \text{for } 0 < a < 1, \tag{16}$$

where the equalities hold when $\pi(X_i, \hat{\beta}) = a\pi(X_i)$, for all i . Proof is available in Supplementary Material C.1.

Note that the MLE does not minimize this upper bound as its coefficients converge to different limits than the DBW unless the propensity score model is correctly specified. Note also that this bound is tighter than Tan (2020)’s bound (Supplementary Material C.1), where the leading term is $5/(3a)$ for a case where $\pi(X_i, \hat{\beta}) \geq a\pi(X_i)$ almost surely for some constant $a \in (0, 1/2]$ instead of $a \in (0, 1]$ in proposition 3. Tan (2020) neither provide an intuitive interpretation of this relative error nor investigate the estimator minimizing it as the loss function.

Second, the MSE of the IPW estimator is bounded above by the relative error of weights as the following proposition states, which is also shown in Proposition 2 of Tan (2020) (a typo in the original proposition corrected):

Proposition 4 (Minimizing a sharp upper bound of the mean squared error): The MSE of the IPW estimator is sharply bounded above by the relative error of weights as:

$$\mathbb{E} \left[\left(\frac{1}{n} \sum_{i=1}^n \frac{R_i Y_i}{\pi(X_i, \hat{\beta})} - \mu \right)^2 \right] \leq \frac{2c}{n\rho} + c \left(1 + \frac{1}{n\rho} \right) \xi(\hat{\beta}), \tag{17}$$

where $\mathbb{E}[Y^2 | X] \leq c$ and $\pi(X_i) \geq \rho$ almost surely for some constants $c > 0$ and $\rho \in (0, 1)$. The equality holds when $Y_i = \mu$ and $\pi(X_i) = \pi(X_i, \hat{\beta}) = 0.5$ for all i . Proof is available in Supplementary Material C.2.

Propositions 3 and 4 together imply that the proposed method minimizes the upper bound of the MSE even when the propensity score model is misspecified.

Table 1. Comparison of the proposed and existing methods for the average outcome estimation

| Estimator | Loss function | Distribution balancing | Mean balancing | PS bounded |
|----------------------|--|------------------------|----------------|------------|
| Proposed estimator | (A) KL divergence between true and estimated weights | ✓ | | ✓ |
| Standard IPW | (B) KL divergence between true and estimated responses | | | ✓ |
| Calibrated weighting | (A) + (B) | | ✓ | ✓ |
| Entropy balancing | (A) KL divergence between true and estimated weights | ✓ | ✓ | |

Notes: The standard IPW is the IPW with the MLE.

3.6 Comparison with the existing methodologies

This section briefly discusses related literature and compares the proposed methods with existing ones. A more detailed discussion including a comparison with the non-parametric methods is provided in Supplementary Material F.

Table 1 summarizes the loss function and three important properties of proposed and three existing methods, the standard IPW, calibrated weighting (Tan, 2020), and entropy balancing (Hainmueller, 2012), for the average outcome estimation, where each of the DBW, calibrated weighting, and entropy balancing has two of those three properties. First, the distribution balancing property minimizes the imbalance in the multivariate covariate distribution between the target and weighted groups (Section 3.3) and an upper bound of the mean squared error (MSE) of the parameter of interest (Section 3.5). The DBW and entropy balancing have this property because they use the same loss function (4) with different link functions, where the DBW uses the logistic function and the entropy balancing uses the exponential function (Wang and Zubizarreta, 2020, See also, Supplementary Material D and F). Second, the calibrated weighting and entropy balancing have the mean balancing property, which is increasingly employed in recently proposed methods (Imai and Ratkovic, 2014; Vermeulen and Vansteelandt, 2015; Zubizarreta, 2015; Zhao, 2019; Wang and Zubizarreta, 2020). This property implies that the estimators are sample-bounded (Robins *et al.*, 2007; Wang and Zubizarreta, 2020). Third, the estimated propensity score is bounded between 0 and 1 except for the entropy balancing, whose implicit propensity scores can take larger values than one (See, Wang and Zubizarreta, 2020). This unboundedness implies that the implicit propensity score model in the entropy balancing is always misspecified, and it cannot be a doubly robust estimator for the average outcome estimation (Zhao, 2019), which contrasts with the double robustness property for the estimation of the average treatment effect for the treated (Zhao and Percival, 2017).

The comparison shows that the drawback of the DBW is the lack of mean balancing and sample-boundedness properties. However, the DBW can incorporate or substitute for these properties. First, it can substitute the mean balancing with outcome models as explained in Section 4.3, which removes bias under the same assumptions as the mean balancing methods such as the calibrated weighting and entropy balancing. I also demonstrate that incorporating the linear outcome model is equivalent to modify the DBW weights to have the mean balancing property. Second, it can incorporate the normalization constraint to gain the sample-boundedness as explained in Section 4.1.

4. Improvement, estimation, and inference

4.1 Normalization, regularization, and estimation

Like the IPW with the MLE, a drawback of the DBW is that estimated weights may not sum to one. To become normalized and thus sample-bounded, the normalized DBW (nDBW) estimator

incorporates the normalization restriction as follows:

$$\hat{\beta}_{\text{nDBW}} = \underset{\beta}{\operatorname{argmin}} \sum_{i=1}^n \log(\pi(X_i, \beta)) \quad \text{s.t.} \quad \sum_{i=1}^n \left(\frac{R_i}{\pi(X_i, \beta)} - 1 \right) = 0. \quad (18)$$

This nDBW estimator minimizes the distribution imbalance and an upper bound of the MSE under the normalization constraint. When the propensity score model is correctly specified, the limiting values of the DBW and nDBW estimators are the same because the (unnormalized) DBW estimator asymptotically satisfies the normalization constraint.

Using the Lagrangian function, the solution to this constrained optimization problem is obtained as the solution to the following unconstrained optimization problem:

$$\hat{\beta}_{\text{nDBW}} = \underset{\beta}{\operatorname{argmin}} \left\{ \sum_{i=1}^n \log(\pi(X_i, \beta)) + \frac{\sum_{i=1}^n (1 - \pi(X_i, \beta))}{n_0} \sum_{i=1}^n \left(\frac{R_i}{\pi(X_i, \beta)} - 1 \right) \right\} \quad (19)$$

$$n_0 = \sum_{i=1}^n (1 - R_i). \quad (20)$$

To estimate coefficients for the nDBW in (19), we can use the M-estimation under the regularity conditions that the solution to (19) is unique and an element of the interior of a compact set. The second condition is satisfied when the objective function in (19) is bounded below, which is similar to the nonseparation condition for the MLE (Vermeulen and Vansteelandt, 2015; Tan, 2020). To satisfy this condition, we can utilize the L2 regularization:

$$\hat{\beta}_{\text{nDBW}} = \underset{\beta}{\operatorname{argmin}} g_{\text{nDBW}}(\beta) \quad (21)$$

$$g_{\text{nDBW}}(\beta) = \sum_{i=1}^n \log(\pi(X_i, \beta)) + \frac{\sum_{i=1}^n (1 - \pi(X_i, \beta))}{n_0} \sum_{i=1}^n \left(\frac{R_i}{\pi(X_i, \beta)} - 1 \right) + \frac{1}{2} \lambda^T \|\beta_{-1}\|^2, \quad (22)$$

where λ is a hyper-parameter controlling for the regularization, β_{-1} is β without the intercept, and $\|\cdot\|$ denotes the Euclidean norm.

As the properties of the M-estimator, the following results regarding the consistency and asymptotic normality of the nDBW estimator are obtained (Stefanski and Boos, 2002).

Proposition 5 (Large sample properties): Under the conditions that the solution to (21) is unique and that (21) is bounded below, $\hat{\beta}_{\text{nDBW}}$ and resulting inverse probability weights $1/\pi(X_i, \hat{\beta}_{\text{nDBW}})$ converge in probability to their limiting values $\hat{\beta}_{\text{nDBW}}^*$ and $1/\pi(X_i, \hat{\beta}_{\text{nDBW}}^*)$, and since (21) is a smooth function, $\hat{\beta}_{\text{nDBW}}$ is asymptotically normally distributed with a certain variance Σ as follows:

$$\sqrt{n}(\hat{\beta}_{\text{nDBW}} - \beta^*) \xrightarrow{d} \mathcal{N}(0, \Sigma) \quad (23)$$

$$\Sigma = g''_{\text{nDBW}}(\beta)^{-1} \mathbb{E}[g'_{\text{nDBW}}(\beta)g'_{\text{nDBW}}(\beta)^T] g''_{\text{nDBW}}(\beta)^{-T}, \quad (24)$$

where β^* is the true values of the coefficients, which is the solution for (3), $g'_{\text{nDBW}}(\beta) = \partial g_{\text{nDBW}}(\beta)/\partial \beta$ and $g''_{\text{nDBW}}(\beta) = \partial g'_{\text{nDBW}}(\beta)/\partial \beta^T$ are the first and second derivatives of

$g_{nDBW}(\beta)$ with respect to β . When the propensity score model is correctly specified, $\hat{\beta}_{nDBW}$ and $1/\pi(X_i, \hat{\beta}_{nDBW})$ are consistently estimated.

To demonstrate the benefits of normalization and regularization, I conduct a simulation in Section 2 for the non-normalized DBW and the nDBW with regularization. First, the RMSEs of the estimated weights are 1.23 for the non-normalized DBW against 1.18 for the nDBW, indicating that the normalization improves the weight estimation.

Second, Figure 3 presents the performance of the proposed estimators with different levels of regularization. The optimal value of the regularization hyper-parameter is 0.22 as depicted by the vertical dotted line. Beyond this value, the RMSE of the weights increases as the regularization becomes stronger, approaching the RMSE of the uniform weights shown in the horizontal dotted line. Although tuning hyper-parameters is still a difficult and unsolved open problem (Zubizarreta, 2015; Wong and Chan, 2017; Zhao, 2019), I recommend using regularization mainly to avoid the separation problem by choosing the smallest one that addresses the separation problem as I do in the simulation studies in Section 5. Another way is choosing the one that minimizes the upper bound of bias estimated by the kernel balance method in Hazlett (2020), which I adopt in the empirical application in Section 6.

4.2 Algorithm

In this section, I explain the algorithm for estimating the DBW without the normalization and regularization for simplicity, and the algorithm for the nDBW with regularization is explained in Supplementary Material E. Since the loss function in (4) is non-convex with respect to β , the convex optimization algorithm such as the Newton’s method is not applicable. For this non-convex optimization, I develop a majorization-minimization algorithm that iteratively optimizes two convex decompositions of the original function (Wu and Lange, 2010). It exploits the characteristics that the minimization problem of the non-convex loss function of (4) is equivalent to the minimization problem of the difference of the two convex functions as follows:

$$g(\beta) \equiv \sum_{i=1}^n \left\{ \frac{R_i}{\pi(X_i, \beta)} + \log(\pi(X_i, \beta)) \right\} = g_1(\beta) - g_2(\beta) \tag{25}$$




Figure 3. The performance of the proposed estimators with different levels of regularization
Notes: The horizontal dotted line indicates the RMSE of the uniform weights.

$$g_1(\beta) = \sum_{i=1}^n \left\{ \frac{R_i}{\pi(X_i, \beta)} + (1 - R_i) \log(\pi(X_i, \beta)) - (1 - R_i) \log(1 - \pi(X_i, \beta)) \right\} \tag{26}$$

$$g_2(\beta) = \sum_{i=1}^n \{ -R_i \log(\pi(X_i, \beta)) - (1 - R_i) \log(1 - \pi(X_i, \beta)) \}. \tag{27}$$

This equivalence implies that the loss function for the DBW (25) is equivalent to the difference of the convex loss functions for the calibrated weighting (26) and the MLE (27). Note that the DBW does *not* minimize the *absolute* difference of the two functions but it minimizes the sum of the loss functions of the calibrated weighting and the negative of the loss function of the MLE. Thus, we should *not* interpret the DBW as the combination of the calibrated weighting and the MLE but rather the calibrated weighting should be regarded as the combination of the DBW and the MLE (Tan, 2020). Since the DBW is better at controlling the mean squared error than the MLE as shown in Section 3.5, the DBW is also better than the calibrated weighting in this regard.

The difference of the convex functions algorithm for the DBW repeats the following two steps until convergence. First, as the majorization step, it constructs a surrogate function $u(\beta, \beta_t)$ for iteration t by the Taylor expansion of the second component of (25) as follows:

$$u(\beta, \beta_t) = g_1(\beta) - \{g_2(\beta_t) + g_2'(\beta_t)^T(\beta - \beta_t)\}, \tag{28}$$

where β_t is the initial values or values obtained in the previous iteration and $g_2'(\beta) = \frac{\partial g_2(\beta)}{\partial \beta}$. Next, as the minimization step, it estimates β_{t+1} such that minimizes (28) with respect to β while keeping β_t fixed. This minimization is easily conducted, e.g., via the Newton’s method, because the objective surrogate function is convex. This algorithm is proved to decrease the loss monotonically in every steps and converge to a stationary point (Wu and Lange, 2010).

4.3 A substitute for the mean balancing condition: doubly robust distribution balancing weighting estimator

Another drawback of the DBW is that it lacks the mean balancing property, which is increasingly employed by recently proposed methods as a safeguard against model misspecification (Hainmueller, 2012; Imai and Ratkovic, 2014; Vermeulen and Vansteelandt, 2015; Zubizarreta, 2015; Zhao, 2019; Tan, 2020; Wang and Zubizarreta, 2020). This removes bias when the true outcome model is a linear regression model with the balanced covariates (Zubizarreta, 2015; Zhao and Percival, 2017; Zhao, 2019; Fan *et al.*, 2021).

To substitute the mean balancing property, the (n)DBW can incorporate the outcome regression model. Specifically, this study considers the augmented IPW where the weights are estimated by the (n)DBW instead of the MLE. The augmented IPW combines the outcome model and propensity score model in the following manner (Robins *et al.*, 1994; Seaman and Vansteelandt, 2018):

$$\hat{\mu} = \frac{1}{n} \sum_{i=1}^n \left(\frac{R_i Y_i}{\pi(X_i, \hat{\beta})} - \frac{R_i - \pi(X_i, \hat{\beta})}{\pi(X_i, \hat{\beta})} \hat{m}(X_i) \right), \tag{29}$$

where $\pi(X_i, \hat{\beta})$ is the estimated propensity scores and $\hat{m}(X_i)$ is the estimated conditional expectation function $m(X_i) = \mathbb{E}[Y_i | X_i]$. The (n)DBW DR estimator uses the WLS to estimate $m(X_i)$ because it has better performance than the OLS (Kang and Schafer, 2007; Imai and Ratkovic, 2014). Under the same condition for the mean balancing methods, the (n)DBW estimator

eliminates bias by using a linear regression outcome model. It can also utilize more flexible methods such as machine learning techniques for estimating $m(X_i)$. This estimator is doubly robust, i.e., consistent when either the propensity score or the outcome model is correctly specified (Vermeulen and Vansteelandt, 2015; Seaman and Vansteelandt, 2018). Moreover, the (n)DBW DR estimator achieves semiparametric efficiency bound if both of the models are correctly specified (Vermeulen and Vansteelandt, 2015; Seaman and Vansteelandt, 2018).

Recently, Chattopadhyay and Zubizarreta (2023) demonstrates that the augmented IPW is expressed as the (non-augmented) IPW with modified weights with the covariate balancing property. Chattopadhyay and Zubizarreta (2023) also proves that the modified weights minimize the Chi-square-type distance from the (inverse probability) weights used in the augmented IPW with the WLS. This implies that the (n)DBW DR estimator can be expressed as the (non-augmented) IPW estimator with weights \hat{w}_{nDBWDR} satisfying the following conditions:

$$\hat{w}_{nDBWDR} = \underset{w}{\operatorname{argmin}} \sum_{i=1}^n R_i \left(\frac{w_i - 1/\pi(X_i, \hat{\beta}_{nDBW})}{1/\pi(X_i, \hat{\beta}_{nDBW})} \right)^2 \quad \text{s.t.} \quad \sum_{i=1}^n (R_i w_i - 1) X_i = 0, \quad (30)$$

which demonstrates that the (n)DBW DR estimator equips the covariate balancing property with the approximate multivariate covariate distribution balancing.

Since the covariate balancing methods satisfy $\frac{1}{n} \sum_{i=1}^n (R_i - \pi(X_i, \hat{\beta})) X_i = 0$ by construction, their DR estimators with the linear outcome model of $\hat{m}(X_i)$ cannot improve their non-DR estimators because the augmented term, the second term in (29), is always 0. I will examine the finite-sample performance of the nDBW DR estimators through simulation studies in Section 5.

5. Simulation studies

This section examines the finite-sample performance of the proposed methods through simulation studies. The standard simulation setting is proposed by Kang and Schafer (2007), which showed that even the DR estimators may suffer from large bias and variance when both the outcome and propensity score models are misspecified. To examine the performance of newly proposed methods, many studies have utilized the same data-generating process (Robins *et al.*, 2007; Imai and Ratkovic, 2014; Vermeulen and Vansteelandt, 2015; Zubizarreta, 2015; Zhao and Percival, 2017; Tan, 2020; Wang and Zubizarreta, 2020).

Following these studies, I use the following data-generating process. There are n units and each unit i has four covariates $X_i = (X_{i1}, X_{i2}, X_{i3}, X_{i4})$, each of which is independently and identically distributed according to the standard normal distribution. Each unit also has an outcome Y_i and a response indicator variable $R_i \in \{0, 1\}$, where outcome is observed only for the response group $R_i = 1$. The response indicator variable is assigned according to the Bernoulli distribution $R_i = \text{Bernoulli}(\pi(X_i))$, where $\pi(X_i)$ is a true propensity score. Following Robins *et al.* (2007); Vermeulen and Vansteelandt (2015), this study considers two types of the true propensity score model, where the signs of the coefficients are reversed each other: the type A model $\pi(X_i) = \text{expit}(-X_{i1} + 0.5 X_{i2} - 0.25 X_{i3} - 0.1 X_{i4})$ and type B model $\pi(X_i) = \text{expit}(X_{i1} - 0.5 X_{i2} + 0.25 X_{i3} + 0.1 X_{i4})$. The target parameter is $\mu = \mathbb{E}[Y_i]$, where the outcome variable Y_i is independently and identically distributed according to the standard normal distribution $Y_i \sim \mathcal{N}(m(X_i), 1)$. Following the extension of the standard simulation study by Tan (2020), this study considers six versions of the true outcome model $m(X_i)$: (Linear 1) $m(X_i) = 210 + 27.4 X_{i1} + 13.7 X_{i2} + 13.7 X_{i3} + 13.7 X_{i4}$ and $\mu = 210$; (Linear 2) $m(X_i) = 210 + 13.7 X_{i1} + 27.4 X_{i2} + 27.4 X_{i3} + 27.4 X_{i4}$ and $\mu = 210$; (Quadratic 1) $m(X_i) = 210 + 27.4 \sum_{j=1}^4 \{\max(X_{ij}, 0)\}^2$ and $\mu = 264.8$; (Quadratic 2) $m(X_i) = 210 + 27.4 \sum_{j=1}^4 \{\max(-X_{ij}, 0)\}^2$ and $\mu = 264.8$; (Exponential 1) $m(X_i) = 210 + 27.4 \sum_{j=1}^4 \exp(X_{ij})$ and $\mu \approx 390.7$; (Exponential 2) $m(X_i) = 210 + 27.4 \sum_{j=1}^4 \exp(-X_{ij})$ and $\mu \approx 390.7$.

To represent misspecified working models, non-linear transforms of the covariates $Z_i = (Z_{i1}, Z_{i2}, Z_{i3}, Z_{i4}) = (\exp(X_{i1}/2), X_{i2}/(1 + \exp(X_{i1})) + 10, (X_{i1}X_{i3}/25 + 0.6)^3, (X_{i1} + X_{i4} + 20)^2)$ are used in some models (Robins *et al.*, 2007; Imai and Ratkovic, 2014; Vermeulen and Vansteelandt, 2015; Zubizarreta, 2015; Zhao and Percival, 2017; Tan, 2020; Wang and Zubizarreta, 2020). Specifically, this study examines the following two scenarios for each simulation setting. In the first scenario, the working propensity score model is correctly specified and uses the logistic regression with X_i as linear predictors whereas in the second scenario it is misspecified and uses the logistic regression with Z_i as linear predictors. Since the focus is on the relative performance of the different propensity score estimation methods, the working outcome model in both the scenarios is misspecified with the linear regression with Z_i as linear predictors.

The propensity scores and/or (inverse probability) weights are estimated by the nDBW and four existing methods (MLE, CBPS, calibrated weighting, and entropy balancing). The unweighted (difference-in-means) and the true propensity score estimators are also used for comparison. For each of these methods, the simple IPW type estimator $\hat{\mu}^{\text{IPW}} = \sum_{i=1}^n R_i Y_i / (n \hat{\pi}(X_i))$ and the augmented IPW type estimator (29) are utilized for the target parameter, where the OLS is used to estimate the outcome regression model. I conduct 2000 Monte Carlo simulations with two different sample sizes ($n = 200$ and 1000) and calculate the bias and root-mean-squared error (RMSE) for each estimator in each of the 48 scenarios (two versions of the true propensity score model, six versions of the true outcome model, two versions of the working propensity score model, and two versions of the sample size). Although the adaptive optimization would improve the estimation, the regularization parameters for the nDBW are simply fixed at $\lambda = 0.03$ for $n = 200$ and $\lambda = 0.007$ for $n = 1000$, which are minimal values to obtain the convergence in all the Monte Carlo draws. Note that the CBPS DR estimator and the bias-reduced doubly robust estimator (Vermeulen and Vansteelandt, 2015) produce the same estimate when the linear outcome regression model uses the same covariates as the propensity score model, i.e., in the misspecified propensity score model cases in this simulation setting.

The results for the augmented IPW type estimators are shown in Table 2–4 and the full results are shown in Table I.1–I.6 in Supplementary Materials I. Note that the unweighted estimator is an IPW type estimator with uniform weights, and the imputation estimator is an augmented IPW type of that.

When the propensity score model is correctly specified, the proposed nDBW DR estimator has at least comparably small bias and RMSE as the calibrated weighting DR, and it has the smallest RMSE in all the cases with a moderate sample size ($n = 1000$) and in most of the cases with a small sample size $n = 200$. Unlike the other estimators, the entropy balancing estimator is not consistent even in the correct propensity score model scenarios except for the linear outcome model cases because it is not doubly robust.

When the propensity score model is misspecified, all the estimators except for the true propensity score DR estimator are no longer consistent, and the MLE DR estimator suffers from large bias and variance when the true propensity score model is type A. Even in these difficult scenarios, the nDBW DR estimator, as well as the calibrated weighting DR estimator, have small biases and RMSEs. To visually compare the relative performance of the estimators with the misspecified propensity score model and moderate sample size ($n = 1000$), Figure 4 shows the absolute bias and RMSE of nDBW DR estimator in the x-axis against those of the calibrated weighting DR (circles), CBPS DR estimator (triangles), and entropy balancing DR (squares) estimators in the y-axis. The symbols above (below) the diagonal line indicate that the nDBW DR estimator performs better (worse) than other estimators. The nDBW estimator outperforms the calibrated weighting DR and CBPS DR estimators both in terms of bias and RMSE in most of the cases. The performance of the nDBW DR estimator is also preferable to the entropy balancing DR estimator, although the latter performs better than the former in some scenarios. Compared with the entropy balancing DR estimator, the nDBW DR estimator has smaller RMSEs by more than 30% in one-third of the cases but larger RMSEs by more than 30% in none of the cases.

Table 2. Simulation results: linear outcome models

| | Type A PS coefficients | | | | Type B PS coefficients | | | |
|--|------------------------|-------|----------|--------|------------------------|-------|----------|-------|
| | n = 200 | | n = 1000 | | n = 200 | | n = 1000 | |
| | Bias | RMSE | Bias | RMSE | Bias | RMSE | Bias | RMSE |
| Linear outcome model 1: correct PS model | | | | | | | | |
| Unweighted | -9.78 | 10.39 | -9.97 | 10.09 | 10.05 | 10.68 | 9.95 | 10.07 |
| nDBW DR | 0.54 | 2.89 | 0.17 | 1.31 | 0.74 | 2.80 | 0.09 | 1.20 |
| MLE DR | 0.36 | 3.65 | 0.01 | 1.74 | 0.44 | 3.74 | 0.00 | 1.76 |
| CBPS DR | 0.49 | 3.25 | 0.07 | 1.59 | 0.38 | 3.08 | 0.02 | 1.46 |
| Calibrated weighting DR | 0.54 | 2.97 | 0.12 | 1.41 | 0.42 | 2.74 | 0.04 | 1.22 |
| Entropy balancing DR | 1.15 | 3.16 | 0.91 | 1.64 | 1.40 | 3.10 | 1.11 | 1.66 |
| True propensity score DR | 0.37 | 3.57 | 0.04 | 1.81 | 0.62 | 4.08 | 0.05 | 1.96 |
| Linear outcome model 1: misspecified PS model | | | | | | | | |
| Unweighted | -10.02 | 10.64 | -9.97 | 10.10 | 9.92 | 10.51 | 10.01 | 10.14 |
| nDBW DR | -1.98 | 3.73 | -2.61 | 2.99 | 2.49 | 3.81 | 1.79 | 2.18 |
| MLE DR | -5.97 | 22.52 | -16.30 | 129.23 | 3.11 | 4.49 | 3.07 | 3.39 |
| CBPS DR/BRDR | -2.73 | 4.36 | -3.57 | 3.97 | 3.07 | 4.39 | 3.33 | 3.63 |
| Calibrated weighting DR | -2.14 | 3.83 | -2.76 | 3.14 | 2.31 | 3.70 | 2.23 | 2.58 |
| Entropy balancing DR | -1.52 | 3.56 | -1.95 | 2.46 | 3.79 | 4.80 | 3.77 | 4.00 |
| True propensity score DR | 0.31 | 3.66 | 0.09 | 1.77 | 0.47 | 4.09 | 0.12 | 1.97 |
| Linear outcome model 2: correct PS model | | | | | | | | |
| Unweighted | -3.59 | 6.18 | -3.76 | 4.37 | 3.76 | 6.29 | 3.73 | 4.34 |
| nDBW DR | 0.07 | 3.64 | 0.03 | 1.65 | 0.62 | 3.66 | 0.08 | 1.63 |
| MLE DR | -0.06 | 4.89 | -0.13 | 2.19 | 0.43 | 5.14 | 0.02 | 2.35 |
| CBPS DR | 0.06 | 4.40 | -0.08 | 2.01 | 0.34 | 4.37 | 0.04 | 2.01 |
| Calibrated weighting DR | 0.10 | 3.77 | -0.05 | 1.72 | 0.37 | 3.73 | 0.06 | 1.67 |
| Entropy balancing DR | 0.02 | 3.85 | -0.08 | 1.74 | 1.18 | 3.91 | 0.95 | 1.91 |
| True propensity score DR | -0.12 | 5.19 | -0.09 | 2.48 | 0.63 | 5.56 | 0.07 | 2.62 |
| Linear outcome model 2: misspecified PS model | | | | | | | | |
| Unweighted | -3.88 | 6.33 | -3.74 | 4.34 | 3.69 | 6.13 | 3.73 | 4.36 |
| nDBW DR | -3.36 | 5.48 | -3.56 | 4.02 | 3.27 | 5.32 | 2.58 | 3.17 |
| MLE DR | -6.05 | 17.36 | -9.29 | 62.15 | 4.34 | 6.43 | 4.25 | 4.73 |
| CBPS DR/BRDR | -4.42 | 6.42 | -4.53 | 4.99 | 4.44 | 6.44 | 4.49 | 4.95 |
| Calibrated weighting DR | -3.55 | 5.59 | -3.73 | 4.18 | 3.27 | 5.33 | 3.15 | 3.66 |
| Entropy balancing DR | -3.49 | 5.70 | -3.51 | 4.03 | 4.54 | 6.33 | 4.45 | 4.86 |
| True propensity score DR | -0.32 | 5.28 | -0.02 | 2.42 | 0.47 | 5.75 | 0.07 | 2.63 |

Notes: This simulation compares the performance of various methods for estimating propensity scores and (inverse probability) weights by investigating combinations of two versions of the true outcome model (Linear 1 and 2) and two versions of coefficients for the propensity score model (type A and B) with the two different numbers of observations ($n=200$ and $n=1000$). For each estimation method, I use two propensity score model specifications (correct and misspecified) and report the bias and RMSE for each in the table.

Finally, the full results are shown in Table I.1–I.6 in Supplementary Materials I, which enable to compare the IPW estimators and their DR variants. In most of the scenarios, the nDBW estimator is improved by incorporating outcome models. In contrast, as explained in Section 4.3, the calibrated weighting and entropy balancing estimators have the same biases and RMSEs as their DR variants in the misspecified propensity score model scenarios, where the linear outcome model uses the same covariates as the propensity score model.

6. Empirical analysis

Can foreign occupiers reduce resistance by devolving political authority to native elites? This question is central in the literature on foreign occupation but also challenging because randomized experiments are almost impossible, and strategic determination of political devolution causes an endogeneity problem in observational studies.

To address this issue, Ferwerda and Miller (2014) utilizes a natural experimental setting in France during World War II, where municipalities were plausibly viewed as randomly assigned into German or Vichy-governed zones in the neighborhood of the demarcation line. After

Table 3. Simulation results: quadratic outcome models

| | Type A PS coefficients | | | | Type B PS coefficients | | | |
|---|------------------------|-------|-----------------|--------|------------------------|-------|-----------------|------|
| | <i>n</i> = 200 | | <i>n</i> = 1000 | | <i>n</i> = 200 | | <i>n</i> = 1000 | |
| | Bias | RMSE | Bias | RMSE | Bias | RMSE | Bias | RMSE |
| Quadratic outcome model 1: correct PS model | | | | | | | | |
| Unweighted | -6.85 | 8.93 | -6.91 | 7.36 | 7.16 | 9.77 | 6.94 | 7.52 |
| nDBW DR | -2.16 | 6.08 | -0.48 | 2.66 | -0.86 | 5.67 | -0.20 | 2.53 |
| MLE DR | -1.43 | 7.45 | -0.31 | 3.53 | -0.53 | 6.95 | -0.14 | 2.96 |
| CBPS DR | -1.58 | 6.63 | -0.39 | 3.11 | -0.55 | 6.18 | -0.14 | 2.79 |
| Calibrated weighting DR | -1.78 | 6.15 | -0.45 | 2.81 | -0.64 | 5.74 | -0.17 | 2.56 |
| Entropy balancing DR | -3.68 | 6.86 | -2.72 | 3.79 | -2.13 | 5.92 | -1.67 | 2.94 |
| True propensity score DR | -1.63 | 7.58 | -0.30 | 3.71 | -0.78 | 6.88 | -0.22 | 3.15 |
| Quadratic outcome model 1: misspecified PS model | | | | | | | | |
| Unweighted | -6.95 | 9.03 | -6.85 | 7.30 | 6.85 | 9.42 | 6.99 | 7.55 |
| nDBW DR | -1.89 | 6.17 | -0.47 | 2.66 | 0.32 | 6.24 | 1.00 | 2.83 |
| MLE DR | 1.71 | 25.10 | 17.40 | 174.96 | 1.02 | 7.76 | 1.34 | 3.51 |
| CBPS DR/BRDR | -2.04 | 6.71 | -0.09 | 3.13 | 1.32 | 7.02 | 1.21 | 3.21 |
| Calibrated weighting DR | -1.83 | 6.22 | -0.43 | 2.73 | 0.65 | 6.30 | 1.07 | 2.87 |
| Entropy balancing DR | -3.63 | 6.87 | -2.40 | 3.59 | -0.95 | 6.30 | -0.62 | 2.70 |
| True propensity score DR | -1.57 | 7.53 | -0.27 | 3.68 | -0.78 | 7.05 | -0.23 | 3.11 |
| Quadratic outcome model 2: correct PS model | | | | | | | | |
| Unweighted | 6.75 | 9.34 | 6.93 | 7.52 | -6.90 | 8.88 | -6.88 | 7.33 |
| nDBW DR | -1.60 | 6.20 | -0.39 | 2.69 | -3.03 | 6.61 | -0.71 | 2.86 |
| MLE DR | -1.11 | 7.63 | -0.12 | 3.43 | -1.53 | 9.63 | -0.17 | 4.80 |
| CBPS DR | -1.33 | 6.87 | -0.23 | 3.16 | -1.75 | 7.36 | -0.33 | 3.74 |
| Calibrated weighting DR | -1.50 | 6.38 | -0.32 | 2.84 | -2.10 | 6.51 | -0.49 | 3.06 |
| Entropy balancing DR | -4.30 | 7.49 | -3.42 | 4.36 | -5.09 | 7.81 | -4.02 | 4.86 |
| True propensity score DR | -1.26 | 7.52 | -0.18 | 3.54 | -1.87 | 10.22 | -0.31 | 5.16 |
| Quadratic outcome model 2: misspecified PS model | | | | | | | | |
| Unweighted | 7.05 | 9.65 | 6.94 | 7.54 | -7.00 | 9.00 | -6.85 | 7.34 |
| nDBW DR | 4.31 | 8.24 | 5.18 | 6.03 | -6.63 | 8.71 | -4.70 | 5.38 |
| MLE DR | 10.72 | 31.50 | 26.53 | 186.73 | -7.23 | 10.21 | -6.63 | 7.37 |
| CBPS DR/BRDR | 5.69 | 9.57 | 6.97 | 7.83 | -7.15 | 9.54 | -7.01 | 7.61 |
| Calibrated weighting DR | 4.67 | 8.46 | 5.46 | 6.30 | -6.40 | 8.63 | -5.40 | 6.05 |
| Entropy balancing DR | 2.30 | 7.29 | 3.15 | 4.41 | -9.16 | 10.71 | -8.50 | 8.88 |
| True propensity score DR | -1.04 | 7.63 | -0.25 | 3.47 | -1.90 | 10.36 | -0.31 | 5.24 |

Notes: This simulation compares the performance of various methods for estimating propensity scores and (inverse probability) weights by investigating combinations of two versions of the true outcome model (Quadratic 1 and 2) and two versions of coefficients for the propensity score model (type A and B) with the two different numbers of observations (*n* = 200 and *n* = 1000). For each estimation method, I use two propensity score model specifications (correct and misspecified) and report the bias and RMSE for each in the table.

Germany defeated France in June 1940, an armistice provided for the division of France into an occupied zone in northern France and the western coast ruled by Germany directly and an unoccupied zone in the south ruled by the Vichy government. In November 1942, in response to the landing of the Allies in North Africa, German forces invaded and occupied the unoccupied southern zone. However, the formal administrative division remained, which kept the extent of political devolution between the two zones different. Ferwerda and Miller (2014) analyzes this period (November 1942 to September 1944), when German military forces were present in both the zones and most of the resistance events occurred (mainly after 1943).

Using the regression discontinuity design, Ferwerda and Miller (2014) finds that political devolution decreases sabotage events in the close neighborhood of the demarcation line, which was drawn arbitrarily at the local level, cutting across the preexisting administrative borders and geographic features such as mountain ranges and rivers. There is, however, a concern about the spill-over effects: Resistance fighters operating from the Vichy zone, not residents in the German-governed zone, may increase sabotage events in the German zone at the demarcation line. To address this concern, it is necessary to compare the municipalities at some distance from the line like the robustness checks in their article. A problem that arises here is that the

Table 4. Simulation results: exponential outcome models

| | Type A PS coefficients | | | | Type B PS coefficients | | | |
|---|------------------------|-------|-----------------|--------|------------------------|-------|-----------------|-------|
| | <i>n</i> = 200 | | <i>n</i> = 1000 | | <i>n</i> = 200 | | <i>n</i> = 1000 | |
| | Bias | RMSE | Bias | RMSE | Bias | RMSE | Bias | RMSE |
| Exponential outcome model 1: correct PS model | | | | | | | | |
| Unweighted | -14.23 | 17.95 | -14.34 | 15.16 | 14.79 | 19.57 | 14.41 | 15.45 |
| nDBW DR | -3.58 | 10.77 | -0.94 | 5.06 | -1.10 | 10.49 | -0.32 | 4.81 |
| MLE DR | -2.44 | 13.31 | -0.58 | 6.69 | -0.58 | 12.93 | -0.22 | 5.52 |
| CBPS DR | -2.64 | 11.91 | -0.75 | 5.85 | -0.67 | 11.60 | -0.22 | 5.21 |
| Calibrated weighting DR | -2.94 | 10.97 | -0.86 | 5.31 | -0.88 | 10.69 | -0.28 | 4.82 |
| Entropy balancing DR | -5.56 | 11.72 | -4.10 | 6.35 | -2.60 | 10.61 | -2.05 | 4.96 |
| True propensity score DR | -2.67 | 13.96 | -0.55 | 6.98 | -0.85 | 13.00 | -0.31 | 5.90 |
| Exponential outcome model 1: misspecified PS model | | | | | | | | |
| Unweighted | -14.52 | 18.28 | -14.28 | 15.09 | 14.28 | 18.98 | 14.47 | 15.48 |
| nDBW DR | -5.13 | 11.96 | -2.63 | 5.88 | 2.05 | 11.98 | 2.78 | 5.88 |
| MLE DR | 0.81 | 61.38 | 45.24 | 601.02 | 3.69 | 14.97 | 4.11 | 7.91 |
| CBPS DR/BRDR | -5.84 | 13.06 | -2.35 | 6.95 | 4.23 | 13.75 | 3.96 | 6.94 |
| Calibrated weighting DR | -5.15 | 12.06 | -2.62 | 6.04 | 2.61 | 12.17 | 3.17 | 6.08 |
| Entropy balancing DR | -7.48 | 13.07 | -5.40 | 7.52 | 0.83 | 11.72 | 1.23 | 5.15 |
| True propensity score DR | -2.69 | 13.86 | -0.47 | 7.26 | -0.99 | 13.08 | -0.34 | 5.91 |
| Exponential outcome model 2: correct PS model | | | | | | | | |
| Unweighted | 14.03 | 18.73 | 14.34 | 15.41 | -14.39 | 18.03 | -14.31 | 15.13 |
| nDBW DR | -2.85 | 11.43 | -0.77 | 5.09 | -5.32 | 12.05 | -1.41 | 5.48 |
| MLE DR | -1.91 | 14.22 | -0.20 | 6.55 | -2.77 | 17.88 | -0.33 | 9.64 |
| CBPS DR | -2.35 | 12.78 | -0.42 | 6.00 | -3.17 | 13.59 | -0.69 | 7.28 |
| Calibrated weighting DR | -2.67 | 11.79 | -0.60 | 5.40 | -3.78 | 12.00 | -1.00 | 5.95 |
| Entropy balancing DR | -7.04 | 13.28 | -5.60 | 7.53 | -8.70 | 13.92 | -6.93 | 8.60 |
| True propensity score DR | -2.06 | 14.14 | -0.31 | 6.79 | -3.27 | 19.39 | -0.55 | 10.54 |
| Exponential outcome model 2: misspecified PS model | | | | | | | | |
| Unweighted | 14.67 | 19.55 | 14.42 | 15.51 | -14.46 | 18.11 | -14.24 | 15.12 |
| nDBW DR | 7.97 | 15.57 | 9.73 | 11.41 | -11.74 | 15.76 | -8.50 | 9.80 |
| MLE DR | 20.04 | 59.54 | 49.55 | 355.21 | -13.17 | 18.47 | -12.22 | 13.54 |
| CBPS DR/BRDR | 10.71 | 18.34 | 13.09 | 14.81 | -13.05 | 17.38 | -12.94 | 14.02 |
| Calibrated weighting DR | 8.62 | 16.01 | 10.26 | 11.93 | -11.35 | 15.56 | -9.81 | 11.04 |
| Entropy balancing DR | 4.84 | 14.05 | 6.36 | 8.68 | -16.36 | 19.28 | -15.34 | 16.07 |
| True propensity score DR | -1.59 | 14.40 | -0.43 | 6.71 | -3.20 | 19.84 | -0.49 | 10.52 |

Notes: This simulation compares the performance of various methods for estimating propensity scores and (inverse probability) weights by investigating combinations of two versions of the true outcome model (Exponential 1 and 2) and two versions of coefficients for the propensity score model (type A and B) with the two different numbers of observations (*n* = 200 and *n* = 1000). For each estimation method, I use two propensity score model specifications (correct and misspecified) and report the bias and RMSE for each in the table.

demarcation line is drawn at random only locally, which leaves some imbalance in the confounders between municipalities with high (the Vichy zone) and low (the German zone) levels of political devolution distant from the demarcation line.

To address this issue, I employ the nDBW DR estimator, and compare the results with the MLE DR, calibrated weighting DR, CBPS DR, entropy balancing DR, and the unweighted difference-in-means estimators. Following Ferwerda and Miller (2014), I construct a propensity score model with the following variables: the distance from the demarcation line, local state capacity (Population from the 1936 census, the distance to the nearest train station, an index variable indicating whether a municipality possessed a telephone bureau, a telegraph station, or a post office in 1939), forest cover and urbanization (the number of actively farmed hectares), the ruggedness of terrain (the mean and standard deviation of the elevation), and the political ideology (leftist and rightist vote shares in the 1936 election).

With these methods, I compare the occurrence of sabotage events, which counts all attacks against infrastructure (largely railroad and communications), between the treatment (German-zone) and control (Vichy-zone) municipalities within some pre-specified bandwidths




Figure 4. Simulation results: Bias and RMSE with misspecified propensity score models

Notes: This figure compares the performance of the DBW DR estimator (x-axis) with the calibrated weighting DR (circles), CBPS DR (triangles), and entropy balancing DR (squares) estimators (y-axis) in terms of the absolute bias (left panel) and RMSE (right panel). Those above the diagonal line indicate that the DBW DR works better than each of these estimators and those below indicate the opposite.

of the distance from the demarcation line (Ferwerda and Miller, 2014). To examine whether the effects are observed only locally near the demarcation line or not, I use rolling windows of 15 km, where municipalities within each of 0–15, 2.5–17.5, ..., 20–35 km from the demarcation line are compared. I denote them by the upper limit: e.g., the results for the municipalities that are 20 km distant from the demarcation line indicate results for those within 5–20 km distant from the line.

6.1 Bias due to imbalance in multivariate covariate distribution

While bias induced by multivariate distributional imbalance is difficult to evaluate, an upper bound of bias can be estimated by the kernel technique (Hazlett, 2020). I use this upper bound to compare the performance of various methods.

Figure 5 presents the upper bounds of bias due to multivariate distributional imbalance for the proposed (shown in red) and other estimators. Those for the calibrated weighting for the distance larger than 30 km are missing because the weights cannot be estimated due to the convergence issue. As the distance from the demarcation line increases, bias due to multivariate distributional imbalance becomes severe, but the proposed estimator mitigates the bias more effectively than the others.

6.2 Results

The results are presented in Figure 6, and the details are shown in Supplementary Material G. The left panel presents point estimates of the ATE for the proposed (shown in red)




Figure 5. Upper bounds of bias due to distributional imbalance for various methods, *Notes:* As the distance from the demarcation line increases, bias due to multivariate distributional imbalance becomes severe, but the proposed estimator mitigates it more effectively than the others.




Figure 6. Political devolution decreases resistance activities only near the demarcation line *Notes:* The left panel presents the ATE estimates by various estimators, diverging as the distance from the demarcation line increases. This divergence is consistent with the estimated upper bounds of bias shown in Figure 5. The right panel presents the ATE estimates and their standard errors estimated by the nDBW DR estimator, which demonstrates that the treatment effects decrease steeply toward zero as the distance from the demarcation line increases.




Figure 7. Distribution balance by various methods.

and other estimators, which diverge when the distance exceeds 25 km. Estimates by the other estimators than the proposed and entropy balancing estimators show unnatural U-shaped effects fluctuating sharply.² These results are consistent with the upper bound of the bias shown in Figure 5.

Why do these estimators provide such different estimates? The key is multivariate distributional imbalance. To examine the distribution balance obtained by these estimators, I conduct the following analysis. First, to represent a possibly complex function of covariates, I prepared the constitutive terms, their squared terms, and all the first-order interactions. Then, to select relevant terms influencing potential outcomes, I estimated outcome models using LASSO and selected terms with non-zero coefficients. This procedure selected sixteen terms out of 54 terms as relevant, of which four were squared terms and nine were interactions, implying that the outcome model consists of a complex function of covariates, highlighting the importance of the distribution balance. Finally, I calculated the absolute mean difference for these selected terms.

Figure 7 presents the results for the treated municipalities with a 20–35 km distance from the demarcation line as an illustration. The proposed estimator (circles) attains better balance than the CBPS estimator (triangles) in all but two terms while avoiding large imbalances unlike the entropy balancing estimator (squares). These results confirm the superiority of the proposed method in the distribution balance.

²The estimates of the calibrated weighting for the distance larger than 30 km are not shown because the weights cannot be estimated due to the convergence issue.

The right panel of [Figure 6](#) presents the ATE estimates estimated by the nDBW DR estimator and their jack-knife standard errors, which demonstrate that the treatment effects decrease steeply toward zero as the distance from the demarcation line increases. The decreasing effect indicates that political devolution reduces resistance activities only near the demarcation line, which raises concern about the spill-over effects that the resistance operating from the Vichy zone increases sabotage events in the German zone. This result suggests the possibility of the contamination of the treatment effects by spill-over effects.

7. Conclusions

The IPW estimators are widely utilized to address missing data problems including causal inference. However, their practical application is susceptible to bias due to propensity score model misspecification. In response, existing studies proposed various methods balancing some moments (or kernels) of observed covariates (Hainmueller, 2012; Imai and Ratkovic, 2014; Vermeulen and Vansteelandt, 2015; Zubizarreta, 2015; Chan *et al.*, 2016; Wong and Chan, 2017; Zhao, 2019; Tan, 2020; Wang and Zubizarreta, 2020), but specifying the moment conditions remains a formidable task because it requires knowledge of the true outcome model (Zubizarreta, 2015; Zhao and Percival, 2017; Zhao, 2019; Fan *et al.*, 2021).

This study proposes the distribution balancing weighting, which mitigates bias and controls the MSE by minimizing their upper bounds. These goals are achieved by balancing covariate distribution through minimizing the KL divergence between the true and estimated weights, inspired by the quasi MLE with misspecified models (Akaike, 1973; White, 1982). The proposed method has several attractive properties as demonstrated in this study.

Finally, I show some future directions for the improvement and the application of the proposed method. First, the key idea of the distribution balancing weighting is to minimize the statistics measuring the discrepancy between the true and estimated weights. This idea may also apply to recently proposed machine learning techniques for propensity score estimation, such as the decision tree, random forest, and generalized random forest. Second, the proposed method may be useful in other missing data problems such as mediation analysis, causal inference with time-series cross-section data, and continuous treatment cases. I am currently exploring these potential directions.

Acknowledgements. I thank Soichiro Yamauchi, the editorial team at *PSRM*, and the anonymous reviewers for their helpful comments. This work was supported by JSPS KAKENHI, Grant-in-Aid for JSPS Fellows (Grant Number 17J03508) and Grant-in-Aid for Early-Career Scientists (Grant Number 21K13225). Competing interests: The author declares none.

Supplementary material. The supplementary material for this article can be found at <https://doi.org/10.1017/psrm.2024.23>.

To obtain replication material for this article, <https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/UAJME4&version=DRAFT&faces-redirect=true>

References

- Akaike H (1973) Information theory and an extension of the likelihood principle. In *Second International Symposium of Information Theory*, Budapest, Hungary, pp. 267–281.
- Chan KCG, Yam SCP and Zhang Z (2016) Globally efficient non-parametric inference of average treatment effects by empirical balancing calibration weighting. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 78, 673–700.
- Chattopadhyay A and Zubizarreta JR (2023) On the implied weights of linear regression for causal inference. *Biometrika* 110, 615–629.
- Diamond A and Sekhon JS (2013) Genetic matching for estimating causal effects: a general multivariate matching method for achieving balance in observational studies. *The Review of Economics and Statistics* 95, 932–45.
- Fan J, Imai K, Lee I, Liu H, Ning Y and Yang X (2021) Optimal covariate balancing conditions in propensity score estimation. *Journal of Business and Economic Statistics* 41, 97–110.
- Ferwerda J and Miller NL (2014) Political devolution and resistance to foreign rule: A natural experiment. *American Political Science Review* 108, 642–60.

- Hainmueller J** (2012) Entropy balancing for causal effects: a multivariate reweighting method to produce balanced samples in observational studies. *Political Analysis* **20**, 25–46.
- Hazlett C** (2020) Kernel balancing: a flexible non-parametric weighting procedure for estimating causal effects. *Statistica Sinica* **30**, 1155–89.
- Imai K and Ratkovic M** (2014) Covariate balancing propensity score. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **76**, 243–63.
- Kallus N** (2020) Generalized optimal matching methods for causal inference. *Journal of Machine Learning Research* **21**, 1–54.
- Kang JDY and Schafer JL** (2007) Demystifying double robustness: a comparison of alternative strategies for estimating a population mean from incomplete data. *Statistical Science* **22**, 523–39.
- Li F, Morgan KL and Zaslavsky AM** (2018) Balancing covariates via propensity score weighting. *Journal of the American Statistical Association* **113**, 390–400.
- Little RJ and Rubin DB** (2019) *Statistical Analysis with Missing Data*. 3 edition. Wiley Series in Probability and Statistics. Hoboken, NJ: John Wiley & Sons.
- Robins JM, Rotnitzky A and Zhao LP** (1994) Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association* **89**, 846–866.
- Robins J, 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.
- Seaman SR and Vansteelandt S** (2018) Introduction to double robust methods for incomplete data. *Statistical Science* **33**, 184–197.
- Stefanski LA and Boos DD** (2002) The calculus of m-estimation. *The American Statistician* **56**, 29–38.
- Tan Z** (2020) Regularized calibrated estimation of propensity scores with model misspecification and high-dimensional data. *Biometrika* **107**, 137–158.
- Vermeulen K and Vansteelandt S** (2015) Bias-reduced doubly robust estimation. *Journal of the American Statistical Association* **110**, 1024–1036.
- Wang Y and Zubizarreta JR** (2020) Minimal dispersion approximately balancing weights: asymptotic properties and practical considerations. *Biometrika* **107**, 93–105.
- White H** (1982) Maximum likelihood estimation of misspecified models. *Econometrica* **50**, 1–25.
- Wong RKW and Chan KCG** (2017) Kernel-based covariate functional balancing for observational studies. *Biometrika* **105**, 199–213.
- Wu TT and Lange K** (2010) The mm alternative to em. *Statistical Science* **25**, 492–505.
- Zhao Q** (2019) Covariate balancing propensity score by tailored loss functions. *The Annals of Statistics* **47**, 965–93.
- Zhao Q and Percival D** (2017) Entropy balancing is doubly robust. *Journal of Causal Inference* **5**, 20160010.
- Zubizarreta JR** (2015) Stable weights that balance covariates for estimation with incomplete outcome data. *Journal of the American Statistical Association* **110**, 910–922.