

Received: 8 February 2019


Revised: 2 November 2019

Accepted: 11 January 2020

DOI: 10.1111/sjos.12457

ORIGINAL ARTICLE**Scandinavian Journal of Statistics**

A framework for covariate balance using Bregman distances

Kevin P. Josey  | **Elizabeth Juarez-Colunga** | **Fan Yang** | **Debashis Ghosh**

Department of Biostatistics and Informatics, Colorado School of Public Health, University of Colorado Anschutz Medical Campus

Correspondence

Debashis Ghosh, Colorado School of Public Health, University of Colorado, Anschutz Medical Campus, 13001 E. 17th Pl., Aurora CO 80045, Mail Stop B119.
Email: Debashis.Ghosh@cuanschutz.edu

Abstract

A common goal in observational research is to estimate marginal causal effects in the presence of confounding variables. One solution to this problem is to use the covariate distribution to weight the outcomes such that the data appear randomized. The propensity score is a natural quantity that arises in this setting. Propensity score weights have desirable asymptotic properties, but they often fail to adequately balance covariate data in finite samples. Empirical covariate balancing methods pose as an appealing alternative by exactly balancing the sample moments of the covariate distribution. With this objective in mind, we propose a framework for estimating balancing weights by solving a constrained convex program, where the criterion function to be optimized is a Bregman distance. We then show that the different distances in this class render identical weights to those of other covariate balancing methods. A series of numerical studies are presented to demonstrate these similarities.

KEYWORDS

average treatment effect, Bregman divergence, causal inference,

This manuscript will be submitted to the Department of Biostatistics and Informatics in the Colorado School of Public Health, University of Colorado Denver, in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Biostatistics for Kevin P. Josey.

© 2020 Board of the Foundation of the Scandinavian Journal of Statistics

Scand J Statist. 2020;1–27.

wileyonlinelibrary.com/journal/sjos

1

1 | INTRODUCTION

Causal inference has been a topic of intense interest in the statistical literature. The focus of causal inference methodology deals with the issue of how to properly evaluate treatment effects in a nonrandomized setting. In many medical and scientific studies, randomization cannot be performed due to logistic, economic, and/or ethical limitations. Under these circumstances, the emergent challenge is the consistent evaluation of treatment effects in the presence of confounding. For example, in Bhagat et al. (2017), a cohort of patients undergoing lung resection surgery are examined to compare the rate of unplanned readmission following thoracoscopic versus open anatomic lung resections. The mode of surgery is not randomized and there are several preoperative characteristics, which inform the method of surgery, that in turn affect the readmission rate. Methods for causal inference seek to minimize any bias induced by these confounding variables.

Two important concepts in causal inference are the potential outcomes model (Rubin, 1974; Splawa-Neyman, Dabrowska, & Speed, 1990) and the propensity score (Rosenbaum & Rubin, 1983). The potential outcomes approach provides a powerful tool for conceptualizing, estimating, and performing inference regarding causal effects. An overview for implementing the potential outcomes model can be found in Imbens and Rubin (2015). They demonstrate that a natural quantity which regularly arises when balancing potential confounders between experimental groups in observational studies is the propensity score (Rosenbaum & Rubin, 1983). The propensity score is defined as the probability of receiving treatment given a set of measured covariates. Based on the assumptions underlying the potential outcomes model and the propensity score, causal inference proceeds in the following stages: (a) a propensity score model is fit using the observed data; (b) diagnostics for covariate balance using the propensity score are evaluated; and (c) estimates of the causal effect are produced by conditioning on the propensity score. Iterating between steps (a) and (b) is often necessary to ensure the homogeneity of the propensity score adjusted covariate distributions.

A key goal for the propensity score model is to achieve covariate balance, which means that the distribution of confounders between the treated and control groups are equivalent. From Rosenbaum and Rubin (1983), the assumptions of strongly ignorable treatment assignment (defined in Section 2.2), in conjunction with the definition of the propensity score, imply that adjustment on the propensity score alone will theoretically achieve balance. However, this result is based on the population propensity score and does not necessarily hold in finite samples. There have been numerous approaches that address the issue of balancing empirical covariate distributions using weighting estimators. We refer to the weights produced by these methods as *balancing weights*. One popular method is to construct propensity scores with covariate balance built into the estimation procedure. Imai and Ratkovic (2013) and Fan, Imai, Liu, Ning, and Yang (2016) introduced the covariate balance propensity score (CBPS) and its subsequent improvement (iCBPS), both of which use generalized methods of moments to fit a logit model with covariate balance serving as

an auxiliary condition. A resulting estimate of the propensity score will automatically achieve

an auxiliary condition. Any resulting estimate of the propensity score will automatically achieve balance by construction.

In the political science literature, Hainmueller (2012) uses maximum entropy density estimation to find balancing weights to estimate the average treatment effect of the treated. The algorithm, termed entropy balancing, finds the vector of balancing weights that minimize the normalized relative entropy from a vector of sampling weights subject to a set of linear equality

constraints about the moments of the covariate distribution. Recent work by Zhao and Percival (2017) shows how this algorithm enjoys a double-robustness property. The general idea of double-robust estimation is to combine covariate information about the treatment assignment and the outcome model into the weighting estimator (Bang & Robins, 2005; Kang & Schafer, 2007). If at least one model is correctly specified, then the resulting causal effect estimate is consistent. When both the outcome and treatment models are correctly specified, then the estimate achieves the semiparametric efficiency bound described by Hahn (1998). Entropy balancing is limited to finding balancing weights to estimate the average treatment effect of the treated, leaving gaps in the procedure for developing doubly robust estimators of other estimands. This issue is related to the choice of Hainmueller (2012) to optimize the normalized relative entropy. By changing the criterion distance function in a convex optimization problem, similar to the one presented by Hainmueller (2012), we can draw parallels to other covariate balancing methods. A similar idea is proposed by Zhao (2019), who shows that CBPS and entropy balancing can be generalized by modifying the score function derived from the respective covariate balance problem. Calibration estimators (Deville & Sarndal, 1992) also produce balancing weights using constrained convex optimization techniques. The proposed methods in Chan, Yam, and Zheng (2015) implicitly extends entropy balancing to include other distance functions. However, they restrict their attention to a non-parametric setting, characterizing their methodology as a departure from the propensity score literature.

Our aim is to extend the work of Hainmueller (2012), Imai and Ratkovic (2013), and Fan et al. (2016) for finding balancing weights that facilitate causal effect estimation when the treatment assignment is not determined by a logit model. We do so by demonstrating how balancing weights can be computed from Bregman distances (Bregman, 1967). Bregman distances have multiple geometric properties that allow for easy estimation of the balancing weights. This geometric interpretation of balancing weights complements the implicit geometry found in classic semiparametric inference. Using the results of our framework, we prove that CBPS (Imai & Ratkovic, 2013) and iCBPS (Fan et al., 2016) are doubly robust estimators of the average treatment effect while assuming the propensity scores follow a logit model. As an extension to CBPS, we propose an estimator for balancing weights akin to the overlap weights discussed by Li, Morgan, and Zaslavsky (2018). We also show how our framework is consistent with the calibration estimator approach of Chan et al. (2015), thereby bridging the empirical covariate balancing methods of entropy balancing, CBPS, iCBPS, and calibration estimators. We are interested in these methods in particular as they do not incorporate a model of the outcome process into their designs in the spirit of Rubin (2008).

The outline of this article is as follows. Section 2 defines the general notation and assumptions that will be applied throughout the manuscript. Section 3 describes the methods for finding balancing weights by solving a constrained optimization problem using Bregman distances as the criterion function. Section 4 describes the similarities between our method and other covariate balancing methods. Section 5 summarizes results from two simulation stud-

ies comparing different covariate balancing methods. This section also contains the results for a replication study of Bhagat et al. (2017) using a variety of different covariate balancing methods. The real dataset illustrates the importance of selecting appropriate covariate balancing methods. Finally, Section 6 concludes with a discussion of the framework and future work.

2 | BACKGROUND AND PRELIMINARIES

2.1 | Notation and definitions

Parameters will be denoted using Greek letters, whereas random variables will be denoted with Roman letters. Boldface letters will denote vectors and matrices, while nonboldface letters represent scalars. For a matrix \mathbf{A} , the transpose is written as \mathbf{A}^T . The symbol ∇f denotes the gradient of a function f . Let $\mathbf{1}_n$ denote the $(n \times 1)$ vector with each entry equal to one. Similarly, let $\mathbf{0}_n$ denote the $(n \times 1)$ vector with each entry equal to zero.

Let \mathbf{X} denote a vector of real-valued covariate measurements, Z denote the random treatment assignment with support $\{0, 1\}$, and Y denote the real-valued outcome variable. The independent sampling units will be indexed by $i = 1, 2, \dots, n$. The $(n \times 1)$ vector of balancing weights will be written as $\mathbf{p} \equiv (p_1, p_2, \dots, p_n)^T$, while the $(n \times 1)$ vector of sampling weights will be written as $\mathbf{q} \equiv (q_1, q_2, \dots, q_n)^T$. We will often write $p_i = p(\mathbf{X}_i)$ for the i th subject to emphasize the fact that the balancing weights are conditioned on the covariates. Define $\{c_j(\mathbf{X}); j = 1, 2, \dots, m\}$, as a set of functions that generate linearly independent features to be balanced between treatment groups. We will refer to these quantities as *balance functions*.

2.2 | Potential outcomes model

Potential outcomes provide a convenient framework for conceptualizing causal effects. This framework was first introduced by Splawa-Neyman et al. (1990) for randomized experiments. The concepts and assumptions necessary to extend this framework to observational data were later formalized by Rubin (1974). The potential outcomes are denoted with a vector $[Y(0), Y(1)]^T$ with $Y(0)$ and $Y(1)$ corresponding to the counterfactual outcome when $Z = 0$ and $Z = 1$, respectively. The conditional expectations for the potential outcomes are denoted with $\mu_0(\mathbf{X}) \equiv \mathbb{E}[Y(0)|\mathbf{X}]$ and $\mu_1(\mathbf{X}) \equiv \mathbb{E}[Y(1)|\mathbf{X}]$. The random outcome is defined by the transformation $Y \equiv ZY(1) + (1 - Z)Y(0)$. Some common causal estimands are the population average treatment effect (ATE), $\tau_{\text{ATE}} \equiv \mathbb{E}[Y(1) - Y(0)]$, and the population average treatment effect of the treated (ATT), $\tau_{\text{ATT}} \equiv \mathbb{E}[Y(1) - Y(0)|Z = 1]$. In any case, the causal effects are nonidentifiable as one of the two required potential outcomes is always missing. This simple observation is the fundamental problem of causal inference. We adopt the setting proposed by Rosenbaum and Rubin (1983) who describe a set of assumptions that will allow us to find consistent estimates of the treatment effect in observational studies. This includes the following assumptions about the data.

Assumption 1 (Strong ignorability). $[Y(0), Y(1)]^T$

The strong ignorability assumption requires that the vector of potential outcomes be independent of the treatment assignment when we condition on the covariates. This assumption further implies that there is no unmeasured confounding. The implication of Assumption 1 along

with the definition of the propensity score as a balance criterion allows us to conclude that $[Y(0), Y(1)]^T \perp\!\!\!\perp Z | \pi(\mathbf{X})$, where $\pi(\mathbf{X}) \equiv \Pr\{Z = 1 | \mathbf{X}\}$ denotes the propensity score (Rosenbaum & Rubin, 1983).

Assumption 2 (Positivity). $0 < \Pr\{Z = 1 | \mathbf{X}\} < 1$ for all \mathbf{X} .

The treatment positivity assumption requires the probability that a subject is assigned to the treatment group as opposed to the control group be bounded away from zero and one. Since

$\Pr\{Z = 1|\mathbf{X}\}$ must be estimated from the covariates, then Assumption 2 equivalently amounts to requiring sufficient overlap between the covariate distributions of the two treatment groups. The feasibility of the convex optimization problems that we will introduce later on are deeply intertwined with Assumption 2. Without sufficient overlap, the estimated balancing weights will either not exist, or be unstable and produce biased estimates of the causal effect.

2.3 | Horvitz–Thompson estimator

The Horvitz–Thompson class of estimators (Hirano, Imbens, & Ridder, 2003; Horvitz & Thompson, 1952) frequently appears in the causal inference literature. For example, the Horvitz–Thompson estimator for the average treatment effect is

$$\hat{\tau}_{\text{ATE}} = \frac{1}{n} \sum_{i=1}^n \left[\frac{Z_i Y_i}{\pi(\mathbf{X}_i)} - \frac{(1 - Z_i) Y_i}{1 - \pi(\mathbf{X}_i)} \right], \quad (1)$$

while the Horvitz–Thompson estimator for τ_{ATT} is

$$\hat{\tau}_{\text{ATT}} = \frac{1}{n_1} \sum_{i=1}^n \left[Z_i Y_i - \frac{\pi(\mathbf{X}_i)(1 - Z_i) Y_i}{1 - \pi(\mathbf{X}_i)} \right], \quad (2)$$

where $n_1 = \sum_{i=1}^n Z_i$. Hahn (1998) was able to show that (1) and (2) have optimal asymptotic properties for estimating τ_{ATE} and τ_{ATT} . Even when we substitute a consistent estimator of the propensity score into (1), the estimator for τ_{ATE} remains consistent and achieves the semiparametric efficiency bound.

A more general form for causal effect estimation is

$$\hat{\tau} = \sum_{i=1}^n \frac{(2Z_i - 1)p(\mathbf{X}_i)Y_i}{\sum_{i=1}^n p(\mathbf{X}_i)Z_i}, \quad (3)$$

which accommodates several different estimands through the choice of $p(\mathbf{X})$. For example, we will see in Section 3.4 that the estimator for τ_{ATT} is similar to estimators of τ_{ATE} with additional constraints placed on the balancing weights so that $p(\mathbf{X}) = q$ whenever $Z = 1$. Equations (1) and (2) provide direction for identifying $p(\mathbf{X})$ within (3) in order to estimate τ_{ATE} and τ_{ATT} , respectively. If the propensity score is known, τ_{ATE} can be estimated by setting $p(\mathbf{X}) = \pi(\mathbf{X})^{-1}$ when $Z = 1$ and $p(\mathbf{X}) = [1 - \pi(\mathbf{X})]^{-1}$ when $Z = 0$. We can also find an estimator for τ_{ATT} by setting $p(\mathbf{X}) = \pi(\mathbf{X})[1 - \pi(\mathbf{X})]^{-1}$ when $Z = 0$ and $p(\mathbf{X}) = 1$ when $Z = 1$. When the propensity score is unknown, finding an estimator for $p(\mathbf{X})$ that produces consistent estimates of τ_{ATE} and τ_{ATT} is relatively straightforward. Estimating balancing weights that also preserve the efficiency of $\hat{\tau}_{\text{ATE}}$ and $\hat{\tau}_{\text{ATT}}$ is a more challenging proposition.

3 | BREGMAN DISTANCES

3.1 | Definition

Let $\Delta^n \subseteq \mathfrak{R}^n$ be a nonempty, convex, and open set with closure $\overline{\Delta}^n$. Define $f : \overline{\Delta}^n \rightarrow \mathfrak{R}$ to be a continuously differentiable, strictly convex function. The Bregman distance generated by the function




FIGURE 1 Example of a Bregman distance for one-dimensional $p', q \in \Delta$. The function $f(p)$ (red line) is strictly convex over $p \in \Delta$. The line tangent to f at q is the blue line. The Bregman function is the distance between the red and blue lines at the point p' [Colour figure can be viewed at wileyonlinelibrary.com]

f is the difference between f evaluated at $\mathbf{p} \in \bar{\Delta}^n$ and the first-order Taylor series approximation of f about $\mathbf{q} \in \Delta^n$, evaluated at \mathbf{p} . In other words, a Bregman distance $D_f : \bar{\Delta}^n \times \Delta^n \rightarrow \mathfrak{R}$ may be defined as

$$D_f(\mathbf{p}||\mathbf{q}) \equiv f(\mathbf{p}) - f(\mathbf{q}) - [\nabla f(\mathbf{q})]^T(\mathbf{p} - \mathbf{q}).$$

Bregman distances are often used to measure the convexity associated with f . Since f is strictly convex over $\bar{\Delta}^n$, it follows that for $\mathbf{p} \in \bar{\Delta}^n$ and $\mathbf{q} \in \Delta^n$, $D_f(\mathbf{p}||\mathbf{q}) \geq 0$ with equality holding if and only if $\mathbf{p} = \mathbf{q}$. This implies that $D_f(\mathbf{p}||\mathbf{q})$ is also strictly convex. A more complete definition of Bregman distances can be found in chapter 2 of Censor and Zenios (1998), which includes additional properties that D_f must satisfy which are not mentioned here. A visual representation of a Bregman distance can be found in Figure 1.

One of the most common examples of a Bregman distance is the unnormalized relative entropy. Let $f(\mathbf{p}) = \sum_{i=1}^n p_i \log(p_i)$ for $\mathbf{p} \in [0, \infty)^n$. We assume $0 \log(0) = 0$ so that the domain of f includes the boundary points contained within the closure of Δ^n . The resulting Bregman distance is written as

$$D_f(\mathbf{p}||\mathbf{q}) = \sum_{i=1}^n p_i \log \frac{p_i}{q_i} - p_i + q_i.$$

The Euclidean distance is another example of a Bregman distance. By selecting $f(\mathbf{p}) = \sum_{i=1}^n p_i^2/2$ for $\mathbf{p} \in \mathfrak{R}^n$, we get

$$D_f(\mathbf{p}||\mathbf{q}) = \sum_{i=1}^n \frac{(p_i - q_i)^2}{2}.$$

$$L(\mathbf{p}||\mathbf{q}) = \sum_{i=1}^n 2$$

In order to simplify the presentation of the methods, we will only consider Bregman distances that are separable. This means $D_f(\mathbf{p}||\mathbf{q}) = \sum_{i=1}^n D_f(p_i ||q_i)$. Note that both the unnormalized relative entropy and the Euclidean distance are separable. Since we require positive weights, we also restrict our focus to convex functions where $\Delta^n \subseteq [0, \infty)^n$ in order to avoid setting additional constraints for $\mathbf{p} \geq \mathbf{0}_n$. Notice that the domain of the unnormalized relative entropy satisfies this

condition while the domain of the Euclidean distance does not. In addition, we will assume throughout that the sampling weights $\mathbf{q} \in \Delta^n$ are fixed by design and known.

3.2 | Constrained optimization and duality

For some $\mathbf{q} \in \Delta^n$, the value $\hat{\mathbf{p}} \in \bar{\Delta}^n$ that minimizes $D_f(\mathbf{p}||\mathbf{q})$ in an unconstrained setting is $\hat{\mathbf{p}} = \mathbf{q}$. In covariate balance problems, we specify a set of linear constraints that the optimal solution must satisfy. Consider the constrained convex optimization problem to

$$\begin{aligned} &\text{minimize} && \sum_{i=1}^n D_f(p_i||q_i) \\ &\text{subject to} && \mathbf{A}^T \mathbf{p} = \mathbf{b}, \end{aligned} \tag{4}$$

where \mathbf{A} is a linearly independent $(n \times m)$ matrix that forms the basis of a linear subspace that defines the constraints of the program and \mathbf{b} is an $(m \times 1)$ vector denoting the margins of those constraints. The entries of \mathbf{A} and \mathbf{b} are denoted with $a_{ij} \in \mathfrak{R}$ and $b_j \in \mathfrak{R}$ ($i = 1, 2, \dots, n$ and $j = 1, 2, \dots, m$), respectively. Equation (4) is often referred to as the primal problem and the corresponding solution is referred to as the primal solution. We denote the set of feasible primal solutions that satisfy the linear constraints in (4) as $\Omega \equiv \{\mathbf{p} : \mathbf{A}^T \mathbf{p} = \mathbf{b}\}$.

Geometrically, the solution to the primal problem is the point

$$\hat{\mathbf{p}} \equiv \arg \min_{\mathbf{p} \in \Omega \cap \bar{\Delta}^n} D_f(\mathbf{p}||\mathbf{q}), \tag{5}$$

which is the generalized projection of $\mathbf{q} \in \Delta^n$ into Ω . Note that $\Omega \cap \bar{\Delta}^n$ is sometimes empty. One solution to avoid this issue is to choose $\mathbf{b} = \mathbf{A}^T \tilde{\mathbf{p}}$, where $\tilde{\mathbf{p}} \in \bar{\Delta}^n$ so that $\tilde{\mathbf{p}} \in \Omega \cap \bar{\Delta}^n$. For example when using the unnormalized relative entropy, we might set $\mathbf{b} = \mathbf{0}_m$ with $\tilde{\mathbf{p}} = \mathbf{0}_n$. However, this condition is not so obvious for some of the covariate balancing problems that we will encounter later on. Instead, we will assume $\Omega \cap \bar{\Delta}^n \neq \emptyset$ throughout. Lemma 1 proves that the generalized projection is unique. The proof appears in section 2.1 of Censor and Zenios (1998). For the sake of completeness, a version of this proof is also found in the online supplement.

Lemma 1. *Suppose $\Omega \cap \bar{\Delta}^n \neq \emptyset$. Then the generalized projection of \mathbf{q} into Ω , defined in (5), is unique.*

When $\hat{\mathbf{p}} \in \Omega \cap \bar{\Delta}^n$, the primal problem can be solved by introducing a vector of Lagrangian multipliers. With Lagrangian multipliers, we can formulate the Lagrangian $L : \bar{\Delta}^n \times \Delta^n \times \mathfrak{R}^m \rightarrow \mathfrak{R}$ for any constrained optimization problem in the form of (4) as

$$L(\mathbf{p}, \mathbf{q}, \boldsymbol{\lambda}) \equiv D_f(\mathbf{p}||\mathbf{q}) + (\mathbf{A}^T \mathbf{p} - \mathbf{b})^T \boldsymbol{\lambda}. \tag{6}$$

Optimizing the Lagrangian with respect to $\mathbf{p} \in \Delta^n$ and $\lambda \in \mathfrak{R}^m$ is an unconstrained problem equivalent to (4). Instead of finding the point $\hat{\mathbf{p}} \in \Omega \cap \Delta^n$ that minimizes $D_f(\mathbf{p}||\mathbf{q})$, we find the vector $\hat{\mathbf{p}}$ that minimizes the Lagrangian with respect to $\mathbf{p} \in \Delta^n$ and $\hat{\lambda}$ that maximizes the Lagrangian with respect to $\lambda \in \mathfrak{R}^m$. In other words, the optimal solution forms a saddle point on (6) over the space $\Delta^n \times \mathfrak{R}^m$ for a fixed $\mathbf{q} \in \Delta^n$.

The following propositions are used in tandem to obtain balancing weights for treatment effect estimation. A necessary condition for these propositions is that the function f , which generates

the Bregman distance D_f be zone consistent with respect to Ω . This means that for any $\mathbf{q} \in \Delta^n$, the Bregman distance produced by f has its generalized projection of \mathbf{q} into Ω contained within the open set Δ^n .

Proposition 1. *Assume that f is zone consistent with respect to Ω . Let $\hat{\mathbf{p}} \in \Omega \cap \Delta^n$ be the generalized projection of \mathbf{q} into Ω , as defined in (5). Then $\hat{\mathbf{p}}$ is uniquely determined by*

$$P_f(\mathbf{q}, \mathbf{A}\hat{\lambda}) \equiv (\nabla f)^{-1}(\nabla f(\mathbf{q}) - \mathbf{A}\hat{\lambda}),$$

where $\hat{\lambda} \in \mathfrak{R}^m$ is also unique.

The proof of Proposition 1 can be found in the online supplement. Given the result of Proposition 1, the convex optimization problem can be solved by estimating $\hat{\lambda} \in \mathfrak{R}^m$ with the dual problem, which is to

$$\begin{aligned} &\text{maximize} && L [P_f(\mathbf{q}, \mathbf{A}\lambda), \mathbf{q}, \lambda] \\ &\text{subject to} && \lambda \in \mathfrak{R}^m. \end{aligned} \tag{7}$$

Propositions 1 and 2 imply that the primal solution to (4) can be obtained by plugging the solution to the dual problem into the generalized projection. The proof of Proposition 2 is adapted from section 3.4 of Bertsekas (1999) and appears in the online supplement.

Proposition 2. *Assume f is zone consistent with respect to Ω . If the primal problem (4) has an optimal solution, then the dual problem (7) also has an optimal solution and the two optimal values are equal.*

3.3 | Balancing weights for the ATE

In this section, we outline the general strategy and guidelines for obtaining balancing weights for estimating τ_{ATE} . The proposed method requires solving the primal problem to

$$\begin{aligned} &\text{minimize} && \sum_{i=1}^n D_f(p_i || q_i) \\ &\text{subject to} && \sum_{i=1}^n p_i(2Z_i - 1)c_j(\mathbf{X}_i) = 0 \quad \text{and} \\ &&& \sum_{i=1}^n p_i Z_i c_j(\mathbf{X}_i) = \sum_{i=1}^n q_i c_j(\mathbf{X}_i) \quad \text{for all } j = 1, 2, \dots, m. \end{aligned} \tag{8}$$

As mentioned in the previous section, (8) can be solved by optimizing the corresponding Lagrangian, which is defined as

$$\begin{aligned}
 LATE(\mathbf{p}, \mathbf{q}, \lambda) \equiv & \sum_{i=1}^n Df(p^i || q^i) + \sum_{j=0}^m \lambda_{j0} \left[\sum_{i=1}^n p^i (2Z_i - 1) c_j(\mathbf{X}_i) \right] \\
 & + \sum_{j=1}^m \lambda_{j1} \left[\sum_{i=1}^n p_i Z_i c_j(\mathbf{X}_i) - \sum_{i=1}^n q_i c_j(\mathbf{X}_i) \right], \tag{9}
 \end{aligned}$$

where $\lambda_0 \equiv (\lambda_{10}, \lambda_{20}, \dots, \lambda_{m0})^T$ and $\lambda_1 \equiv (\lambda_{11}, \lambda_{21}, \dots, \lambda_{m1})^T$ with $\lambda \equiv (\lambda_0^T, \lambda_1^T)^T$. The criterion distance function should be selected so that $P_f \left[q_i, \sum_{j=1}^m c_j(\mathbf{X}_i) \lambda_{j0} \right] = \pi(\mathbf{X}_i)^{-1}$. We can also frame this problem using the notation from Section 3.2. Let \mathbf{A}_0 be an $(n \times m)$ matrix whose elements consist of $a_{ij0} = (2Z_i - 1)c_j(\mathbf{X}_i)$, \mathbf{A}_1 be an $(n \times m)$ with entries $a_{ij1} = Z_i c_j(\mathbf{X}_i)$, $\mathbf{b}_0 = \mathbf{0}_m$, and \mathbf{b}_1 be an $(m \times 1)$ vector with entries $b_{j1} = \sum_{i=1}^n q_i c_j(\mathbf{X}_i)$. We then combine \mathbf{A}_0 and \mathbf{A}_1 to construct $\mathbf{A} = [\mathbf{A}_0, \mathbf{A}_1]$ while \mathbf{b}_0 and \mathbf{b}_1 are concatenated into $\mathbf{b} = [\mathbf{b}_0^T, \mathbf{b}_1^T]^T$.

After differentiating (9) with respect to p_i for some $i = 1, 2, \dots, n$ and setting the resulting derivative to zero, we arrive at the generalized projection evaluated at $\lambda \in \mathfrak{R}^{2m}$. The dual objective function is obtained by substituting the generalized projection for p_i in the Lagrangian. The dual solution solves the dual problem,

$$\hat{\lambda} = \arg \max_{\lambda \in \mathfrak{R}^{2m}} \sum_{i=1}^n L_{\text{ATE}} \left\{ P_f \left[q_i, \sum_{j=1}^m (2Z_i - 1)c_j(\mathbf{X}_i)\lambda_{j0} + \sum_{j=1}^m Z_i c_j(\mathbf{X}_i)\lambda_{j1} \right], q_i, \lambda \right\}. \quad (10)$$

As a result of Propositions 1 and 2, the balancing weights are uniquely determined by the generalized projection evaluated at $\hat{\lambda} \in \mathfrak{R}^{2m}$,

$$\hat{p}(\mathbf{X}_i) = P_f \left[q_i, \sum_{j=1}^m (2Z_i - 1)c_j(\mathbf{X}_i)\hat{\lambda}_{j0} + \sum_{j=1}^m Z_i c_j(\mathbf{X}_i)\hat{\lambda}_{j1} \right], \quad i = 1, 2, \dots, n. \quad (11)$$

It is recommended, and often necessary (see the proofs to Theorems 1 and 2), that one of the balance functions be an intercept—we will assume throughout that $c_1(\mathbf{X}) = 1$ for all \mathbf{X} . This constraint implies $\sum_{\{i:Z_i=1\}} \hat{p}_i = \sum_{\{i:Z_i=0\}} \hat{p}_i$.

3.4 | Balancing weights for the ATT

Next we consider the problem of finding balancing weights to estimate τ_{ATT} . This requires solving the primal problem to

$$\begin{aligned} &\text{minimize} && \sum_{i=1}^n D_f(p_i || q_i) \\ &\text{subject to} && p_i(1 - Z_i)c_j(\mathbf{X}_i) = \sum_{i=1}^n q_i Z_i c_j(\mathbf{X}_i) \quad \text{for all } j = 1, 2, \dots, m. \end{aligned} \quad (12)$$

The criterion Bregman distance should be chosen so that the generalized projection resembles the functional form for the odds of treatment. That is,

m

$$Pf \left[q^i, \sum_{j=1}^m c_j(\mathbf{X}^i) \lambda_j \right] = \frac{\pi(\mathbf{X}^i)}{1 - \pi(\mathbf{X}^i)},$$

where $\lambda \equiv (\lambda_1, \lambda_2, \dots, \lambda_m)^T$. In terms of the notation presented in Section 3.2, (12) is equivalent to (4) by setting $a_{ij} = (1 - Z_i)c_j(\mathbf{X}_i)$, $i = 1, 2, \dots, n$, and $b_j = \sum_{i=1}^n q_i Z_i c_j(\mathbf{X}_i)$, $j = 1, 2, \dots, m$. Similar to the balancing weights for estimating τ_{ATE} , we set $c_1(\mathbf{X}) = 1$ for all \mathbf{X} so that $\sum_{\{i: Z_i=1\}} q_i =$

$\sum_{\{i: Z_i=0\}} \hat{p}_i$. The associated Lagrangian for (12) can then be expressed as

$$L_{\text{ATT}}(\mathbf{p}, \mathbf{q}, \boldsymbol{\lambda}) \equiv \sum_{i=1}^n D_f(p_i || q_i) + \sum_{j=1}^m \lambda_j \left[\sum_{i=1}^n p_i(1 - Z_i)c_j(\mathbf{X}_i) - \sum_{i=1}^n q_i Z_i c_j(\mathbf{X}_i) \right].$$

Notice that when $Z = 1$ and $q \in \Delta$, then $P_f \left[q, (1 - Z) \sum_{j=1}^m c_j(\mathbf{X}) \lambda_j \right] = q$. Moreover, recall from Section 3.1 that $D_f(q || q) = 0$ for some $q \in \Delta$. Therefore, (12) can be reconstructed into the equivalent primal problem to

$$\begin{aligned} &\text{minimize} && \sum_{i=1}^n D_f(p_i || q_i) \\ &\text{subject to} && \sum_{i=1}^n p_i(2Z_i - 1)c_j(\mathbf{X}_i) = 0 \quad \text{for all } j = 1, 2, \dots, m \quad \text{and} \\ &&& p_i = q_i \quad \text{for all } i \in \{i : Z_i = 1\}. \end{aligned}$$

According to Propositions 1 and 2, the balancing weights are evaluated with the resulting generalized projection,

$$\hat{p}(\mathbf{X}_i) = P_f \left[q_i, (1 - Z_i) \sum_{j=1}^m c_j(\mathbf{X}_i) \hat{\lambda}_j \right], \quad i = 1, 2, \dots, n,$$

where the dual vector is estimated by solving for

$$\hat{\boldsymbol{\lambda}} = \arg \max_{\boldsymbol{\lambda} \in \mathfrak{R}^m} \sum_{i=1}^n L_{\text{ATT}} \left\{ P_f \left[q_i, (1 - Z_i) \sum_{j=1}^m c_j(\mathbf{X}_i) \lambda_j \right], q_i, \boldsymbol{\lambda} \right\}.$$

4 | RELATIONSHIP WITH OTHER COVARIATE BALANCE METHODS

4.1 | Entropy balancing

Entropy balancing (Hainmueller, 2012) is a special case of a constrained Bregman distance optimization problem. By setting $f(\mathbf{p}) = \sum_{i=1}^n p_i \log(p_i)$ and $q_i \in (0, \infty)$ for all $i = 1, 2, \dots, n$, we can identify the entropy balancing primal problem, which is to

$$\text{minimize} \quad \sum_{i=1}^n p_i \log \left(\frac{p_i}{q_i} \right) - p_i + q_i$$

$$\text{subject to } \sum_{i=1}^n \left[\begin{array}{c} p_i \\ (q^i) \end{array} \right] \sum_{j=1}^m p_i (1 - Z_i) c_j(\mathbf{X}_i) = \sum_{i=1}^n q_i Z_i c_j(\mathbf{X}_i) \quad \text{for all } j = 1, 2, \dots, m. \quad (13)$$

According to Proposition 2, optimizing (13) is equivalent to maximizing the dual objective function,

$$\hat{\lambda} = \arg \max_{\lambda \in \mathcal{R}^m} \sum_{i=1}^n \left\{ -q_i \exp \left[-(1 - Z_i) \sum_{j=1}^m c_j(\mathbf{X}_i) \lambda_j \right] - q_i Z_i \sum_{j=1}^m c_j(\mathbf{X}_i) \lambda_j \right\}. \tag{14}$$

The vector of balancing weights is obtained by evaluating the generalized projection with the solution to the dual problem, which yields

$$\hat{p}(\mathbf{X}_i) = q_i \exp \left[-(1 - Z_i) \sum_{j=1}^m c_j(\mathbf{X}_i) \hat{\lambda}_j \right], \quad i = 1, 2, \dots, n. \tag{15}$$

In Hainmueller (2012), (13) is written using the normalized relative entropy instead of the unnormalized relative entropy. However, optimizing the normalized relative entropy is simply achieved by (14) and (15) with minor alterations. Let $q'_i = q_i / \sum_{i=1}^n q_i Z_i$ for all $i \in \{i : Z_i = 1\}$ and constrain the intercept so that $\sum_{i=1}^n \hat{p}_i(1 - Z_i) = \sum_{i=1}^n q'_i Z_i = 1$. As previously suggested in Section 3.4, we recommend setting $c_1(\mathbf{X}) = 1$. In doing so, the resulting balancing weights for the control group will sum to one while still satisfying the constraints of the primal problem.

Using the resulting estimating equations for λ and τ_{ATT} in concordance with results from M-estimation theory (Stefanski & Boos, 2002), Zhao and Percival (2017) show that entropy balancing weights produce doubly robust estimates of τ_{ATT} . This means if either $\text{logit}[\pi(\mathbf{X})] \in \text{span}\{c_j(\mathbf{X}) : j = 1, 2, \dots, m\}$ or $\mu_0(\mathbf{X}) \in \text{span}\{c_j(\mathbf{X}) : j = 1, 2, \dots, m\}$, then the balancing weights of (15) applied to (3) is consistent for τ_{ATT} . If both conditions are satisfied, then the estimator achieves the semiparametric efficiency bound derived by Hahn (1998) for estimators of τ_{ATT} . The Horvitz–Thompson estimator for τ_{ATT} that substitutes a consistent estimate of the propensity score for $\pi(\mathbf{X})$ in (2), on the other hand, does not achieve the semiparametric efficiency bound.

4.2 | Covariate balance propensity scores

Another method for covariate balance, developed by Imai and Ratkovic (2013), proposes fitting a logit model for the propensity score,

$$\pi(\mathbf{X}_i) = \frac{\exp \left[\sum_{j=1}^m c_j(\mathbf{X}_i) \lambda_j \right]}{1 + \exp \left[\sum_{j=1}^m c_j(\mathbf{X}_i) \lambda_j \right]}, \quad i = 1, 2, \dots, n, \tag{16}$$

subject to

$$\sum_{i=1}^n \left[\frac{Z_i c_j(\mathbf{X}_i)}{\pi(\mathbf{X}_i)} - \frac{(1 - Z_i) c_j(\mathbf{X}_i)}{1 - \pi(\mathbf{X}_i)} \right] = 0 \quad \text{for all } j = 1, 2, \dots, m. \tag{17}$$

They opted to solve for $\lambda \in \mathfrak{R}^m$ using generalized method of moments (GMM) while at the same time satisfying (17). The estimated propensity scores can be transformed into balancing weights for estimating τ_{ATE} with the inverse probability of treatment weighting estimator. We will refer to the model where the balance functions that appear in (17) are identical to the balance functions within the linear predictor of (16) as the exactly specified CBPS model.

The weights obtained with an exactly specified CBPS model can be expressed in an equivalent manner to a constrained optimization problem following our framework. First, notice that the

fixed effect coefficients of the logit model can double as a vector of dual variables. Next, observe that (17) can be rewritten as

$$\sum_{i=1}^n \left\{ 1 + \exp \left[-(2Z_i - 1) \sum_{j=1}^m c_j(\mathbf{X}_i) \lambda_j \right] \right\} (2Z_i - 1) c_j(\mathbf{X}_i) = 0 \quad \text{for all } j = 1, 2, \dots, m.$$

The CBPS primal problem can then be constructed in order to

$$\begin{aligned} &\text{minimize} && \sum_{i=1}^n \left[(p_i - 1) \log \left(\frac{p_i - 1}{q_i - 1} \right) - p_i + q_i \right] \\ &\text{subject to} && \sum_{i=1}^n p_i (2Z_i - 1) c_j(\mathbf{X}_i) = 0 \quad \text{for all } j = 1, 2, \dots, m. \end{aligned} \tag{18}$$

We call the criterion distance function in (18) the shifted relative entropy, which is generated by setting $f(\mathbf{p}) = \sum_{i=1}^n (p_i - 1) \log(p_i - 1)$, $\mathbf{p} \in [1, \infty)^n$. Notice that (18) also contains fewer constraints than (8). We assume $q_i = 2$ for all $i = 1, 2, \dots, m$. Assuming uniform sampling weights follows the prevailing philosophy of the causal inference literature in which observational data are typically randomly sampled from the population of interest. The solution to the dual problem for (18) finds

$$\hat{\lambda} = \arg \max_{\lambda \in \mathcal{R}^m} \sum_{i=1}^n \left\{ (2Z_i - 1) \sum_{j=1}^m c_j(\mathbf{X}_i) \lambda_j - \exp \left[-(2Z_i - 1) \sum_{j=1}^m c_j(\mathbf{X}_i) \lambda_j \right] \right\}. \tag{19}$$

The principal reason for selecting the shifted relative entropy as the criterion distance function is the resulting balancing weights, which resemble the inverse probability of treatment weights,

$$\hat{p}(\mathbf{X}_i) = 1 + \exp \left[-(2Z_i - 1) \sum_{j=1}^m c_j(\mathbf{X}_i) \hat{\lambda}_j \right], \quad i = 1, 2, \dots, n. \tag{20}$$

A similar derivation of CBPS using the dual function setup was also described by Zhao (2019).

Fan et al. (2016) identifies a condition that the balance functions must satisfy in order for CBPS to be doubly robust for estimating τ_{ATE} . This condition is not obvious from a data analytic context. However, the condition is satisfied if we assume a constant conditional average treatment effect. Under this assumption, we can prove that CBPS is doubly robust using the balancing weights produced by (20).

Assumption 3 (Constant conditional ATE). For all \mathbf{X} , $\mu_1(\mathbf{X}) - \mu_0(\mathbf{X}) = \tau$.

Theorem 1. *Let Assumptions 1 and 2 be given. Suppose $\mathbb{E}[Y(0)]$, $\mathbb{E}[Y(1)]$, and $\mathbb{E}[c_j(\mathbf{X})]$ exist for all $j = 1, 2, \dots, m$. Furthermore, assume $\mathbb{V}[Y(0)] < \infty$ and $\mathbb{V}[Y(1)] < \infty$. Then the balancing weights determined by (19) and (20) applied to (3) is doubly robust in the sense that:*

1. *If $\text{logit}[\pi(\mathbf{X})] = \sum_{j=1}^m c_j(\mathbf{X})\lambda_j$ for some $\lambda_j \in \mathfrak{R}$, $j = 1, 2, \dots, m$, then $\hat{\tau}$ is consistent for τ_{ATE} ;*

2. Under Assumption 3 and if $\mu_0(\mathbf{X}) = \sum_{j=1}^m c_j(\mathbf{X})\beta_j$ for some $\beta_j \in \mathfrak{R}$, $j = 1, 2, \dots, m$, then $\hat{\tau}$ is consistent for τ_{ATE} ;
3. If Conditions 1 and 2 are both satisfied, then

$$\sqrt{n}(\hat{\tau} - \tau_{ATE}) \rightarrow_d \mathcal{N}(0, \Sigma_{\text{semi}})$$

where

$$\Sigma_{\text{semi}} = \mathbb{E} \left\{ \frac{\mathbb{V}[Y(1)|\mathbf{X}]}{\pi(\mathbf{X})} + \frac{\mathbb{V}[Y(0)|\mathbf{X}]}{1 - \pi(\mathbf{X})} \right\}.$$

As an extension to CBPS, consider the primal problem using the Bregman distance generated by setting $f(\mathbf{p}) = \sum_{i=1}^n p_i \log(p_i) + (1 - p_i) \log(1 - p_i)$, $\mathbf{p} \in [0, 1]^n$:

$$\begin{aligned} &\text{minimize} && \sum_{i=1}^n \left[p_i \log \left(\frac{p_i}{q_i} \right) + (1 - p_i) \log \left(\frac{1 - p_i}{1 - q_i} \right) \right] \\ &\text{subject to} && \sum_{i=1}^n p_i (2Z_i - 1) c_j(\mathbf{X}_i) = 0 \quad \text{for all } j = 1, 2, \dots, m. \end{aligned} \tag{21}$$

If we assume $q_i = 1/2$ for all $i = 1, 2, \dots, n$, then according to Propositions 1 and 2, the solution to (21) is

$$\hat{p}(\mathbf{X}_i) = \frac{1}{1 + \exp \left[(2Z_i - 1) \sum_{j=1}^m c_j(\mathbf{X}_i) \hat{\lambda}_j \right]}, \tag{22}$$

where the dual solution is obtained by solving for

$$\begin{aligned} \hat{\lambda} = \arg \max_{\lambda \in \mathfrak{R}^m} & \sum_{i=1}^n \frac{1}{1 + \exp \left[(2Z_i - 1) \sum_{j=1}^m c_j(\mathbf{X}_i) \lambda_j \right]} \log \left[\frac{2}{1 + \exp \left[(2Z_i - 1) \sum_{j=1}^m c_j(\mathbf{X}_i) \lambda_j \right]} \right] \\ & + \sum_{i=1}^n \frac{1}{1 + \exp \left[-(2Z_i - 1) \sum_{j=1}^m c_j(\mathbf{X}_i) \lambda_j \right]} \log \frac{2}{1 + \exp \left[-(2Z_i - 1) \sum_{j=1}^m c_j(\mathbf{X}_i) \lambda_j \right]} \\ & + \sum_{j=1}^m \lambda_j \left\{ \sum_{i=1}^n \left[\frac{(2Z_i - 1) c_j(\mathbf{X}_i)}{1 + \exp \left[(2Z_i - 1) \sum_{j=1}^m c_j(\mathbf{X}_i) \lambda_j \right]} \right] - \sum_{i=1}^n \left[\frac{(2Z_i - 1) c_j(\mathbf{X}_i)}{1 + \exp \left[-(2Z_i - 1) \sum_{j=1}^m c_j(\mathbf{X}_i) \lambda_j \right]} \right] \right\}. \end{aligned} \tag{23}$$

Printed by [Pennsylvania State University - 071.058.108.011 - /doi/pdf/10.1111/sjos.12457] at [14/10/2020].

˘ L — J'

The Bregman distance in this case is referred to as the binary relative entropy. This distance is useful for finding balancing weights that produce estimates for a special case of the weighted average treatment effect called the optimally weighted average treatment effect (OWATE) (Crump, Hotz, Imbens, & Mitnik, 2006),

$$\tau_{\text{OWATE}} \equiv \frac{\mathbb{E}\{\pi(\mathbf{X})[1 - \pi(\mathbf{X})][Y(1) - Y(0)]\}}{\mathbb{E}\{\pi(\mathbf{X})[1 - \pi(\mathbf{X})]\}}.$$

A consistent estimator for τ_{OWATE} is also consistent for τ_{ATE} , with the smallest variance, when we are given Assumption 3 and the potential outcomes have equal variance. Li et al. (2018) further motivates the use of estimators for τ_{OWATE} when there is poor overlap between the treated and the control groups. Equations (23) and (22) provide a dual interpretation of the covariate balance scoring rule for estimating τ_{OWATE} considered by Zhao (2019). By replacing $p(\mathbf{X})$ with (22) in (3), and using arguments similar to the proof of Theorem 1, we derive a doubly robust estimator for τ_{OWATE} with the usual asymptotic properties.

Corollary 1. *Under the same assumptions and conditions as Theorem 1, the balancing weights determined by (22) and (23) applied to (3) is doubly robust for estimating τ_{OWATE} with asymptotic variance*

$$\Sigma_{\text{semi}} = \frac{\mathbb{E}\left(\pi(\mathbf{X})^2[1 - \pi(\mathbf{X})]^2 \left\{ \frac{\mathbb{V}[Y(1)|\mathbf{X}]}{\pi(\mathbf{X})} + \frac{\mathbb{V}[Y(0)|\mathbf{X}]}{1 - \pi(\mathbf{X})} \right\}\right)}{\mathbb{E}\{\pi(\mathbf{X})[1 - \pi(\mathbf{X})]\}^2}. \tag{24}$$

4.3 | Improved CBPSs

The iCBPS approach (Fan et al., 2016) improves upon the CBPS method described in Section 4.2 to better accommodate heterogeneous treatment effects. The objective of this method is to fit a logit model subject to the constraints

$$\begin{aligned} \sum_{i=1}^n \left[\frac{Z_i c_j(\mathbf{X}_i)}{\pi(\mathbf{X}_i)} - \frac{(1 - Z_i) c_j(\mathbf{X}_i)}{1 - \pi(\mathbf{X}_i)} \right] &= 0 \quad \text{and} \\ \sum_{i=1}^n \left[\frac{Z_i}{\pi(\mathbf{X}_i)} - 1 \right] c_j(\mathbf{X}_i) &= 0 \quad \text{for all } j = 1, 2, \dots, m. \end{aligned} \tag{25}$$

Fan et al. (2016) uses GMM to estimate $\hat{\lambda} \in \mathfrak{R}^m$ in (16) subject to (25). This modified approach can be adapted to fit into our proposed framework, with the balancing weights being estimated using dual optimization techniques instead of GMM.

Using the same criterion Bregman distance as the one used in (18), we can obtain balancing weights that satisfy (25) as follows. Assume $q_i = 2$ for all $i = 1, 2, \dots, m$. Define the primal problem for iCBPS as

$$\text{minimize} \quad \sum_{i=1}^n \left[(p_i - 1) \log(p_i - 1) - p_i + 2 \right]$$

$$\begin{aligned}
 \text{subject to } & \sum_{i=1}^n p_i (2Z_i - 1) c_j(\mathbf{X}_i) = 0 \quad \text{and} \\
 & \sum_{i=1}^n p_i Z_i c_j(\mathbf{X}_i) = \sum_{i=1}^n q_i c_j(\mathbf{X}_i) \quad \text{for all } j = 1, 2, \dots, m.
 \end{aligned} \tag{26}$$

As opposed to (18), the iCBPS primal problem follows our guidelines in Section 3.3 more closely. The resulting dual solution solves for

$$\hat{\lambda} = \arg \max_{\lambda \in \mathfrak{R}^{2m}} \sum_{i=1}^n \left\{ (2Z_i - 1) \sum_{j=1}^m c_j(\mathbf{X}_i) \lambda_{j0} + Z_i \sum_{j=1}^m c_j(\mathbf{X}_i) \lambda_{j1} - \sum_{j=1}^m c_j(\mathbf{X}_i) \lambda_{j1} - \exp \left[-(2Z_i - 1) \sum_{j=1}^m c_j(\mathbf{X}_i) \lambda_{j0} - Z_i \sum_{j=1}^m c_j(\mathbf{X}_i) \lambda_{j1} \right] \right\}. \tag{27}$$

The covariate balancing weights differ slightly from (20) due to the additional constraints in (26) with

$$\hat{p}(\mathbf{X}_i) = 1 + \exp \left[-(2Z_i - 1) \sum_{j=1}^m c_j(\mathbf{X}_i) \hat{\lambda}_{j0} - Z_i \sum_{j=1}^m c_j(\mathbf{X}_i) \hat{\lambda}_{j1} \right], \quad i = 1, 2, \dots, n. \tag{28}$$

Note that with the GMM approach, $\lambda \in \mathfrak{R}^m$, whereas with our method, $\lambda \in \mathfrak{R}^{2m}$. This implies that exact balance between covariates is not necessarily achieved with the proposed methods of Fan et al. (2016). A notable deviation from our own recommendations exists within how \mathbf{b} is specified. Even though $q_i = 2$ for all $i = 1, 2, \dots, n$, we set $b_{j1} = \sum_{i=1}^n c_j(\mathbf{X}_i)$ for all $j = 1, 2, \dots, m$. If we were to follow the setup in Section 3.3, we would set $b_{j1} = \sum_{i=1}^n 2c_j(\mathbf{X}_i)$. However, this distinction is minor in the context of uniform sampling weights and should produce similar results. We now show that the weights produced by (27) and (28) applied to (3) is doubly robust given a linear conditional average treatment effect, defined in Assumption 4. Note that Assumption 4 is less stringent than Assumption 3, which was necessary to prove Theorem 1.

Assumption 4 (Linear conditional ATE). For all \mathbf{X} , $\mu_1(\mathbf{X}) - \mu_0(\mathbf{X}) = \sum_{j=1}^m c_j(\mathbf{X}) \alpha_j$ where $\alpha_j \in \mathfrak{R}$ for all $j = 1, 2, \dots, m$.

Theorem 2. Let Assumptions 1 and 2 be given. Suppose $\mathbb{E}[Y_i(0)]$, $\mathbb{E}[Y(1)]$, and $\mathbb{E}[c_j(\mathbf{X})]$ exist for all $j = 1, 2, \dots, m$. Assume $\mathbb{V}[Y(0)] < \infty$ and $\mathbb{V}[Y(1)] < \infty$. Then the balancing weights determined by (27) and (28) applied to (3) is doubly robust in the sense that:

1. If $\text{logit}[\pi(\mathbf{X})] = \sum_{j=1}^m c_j(\mathbf{X}) \lambda_{j0}$ for some $\lambda_{j0} \in \mathfrak{R}$, $j = 1, 2, \dots, m$, then $\hat{\tau}$ is consistent for τ_{ATE} ;
2. Under Assumption 4 and if $\mu_0(\mathbf{X}) = \sum_{j=1}^m c_j(\mathbf{X}) \beta_j$ for some $\beta_j \in \mathfrak{R}$, $j = 1, 2, \dots, m$, then $\hat{\tau}$ is consistent for τ_{ATE} ;
3. If Conditions 1 and 2 are both satisfied, then

$$\sqrt{n}(\hat{\tau} - \tau_{\text{ATE}}) \rightarrow_d \mathcal{N}(0, \Sigma_{\text{semi}}),$$

where

$$\Sigma_{\text{semi}} = \mathbb{E} \left[\frac{1}{\mathbb{V}[Y(1)|\mathbf{X}]} + \frac{1}{\mathbb{V}[Y(0)|\mathbf{X}]} \right].$$

$$\left\{ \pi(\mathbf{X}) \quad 1 - \pi(\mathbf{X}) \right\}$$

4.4 | Calibration estimators

Chan et al. (2015) describes a class of estimators originally introduced by Deville and Sarndal (1992) for survey sampling called calibration estimators. One of the contributions from Chan

et al. (2015) shows how calibration estimators can be applied to covariate balance problems. For some separable generalized distance function $G : \Delta^n \rightarrow \mathfrak{R}$, calibration estimators find balancing weights that solve the primal problem to

$$\begin{aligned} & \text{minimize} && \sum_{i=1}^n G(p_i) \\ & \text{subject to} && \sum_{i=1}^n p_i(1 - Z_i)c_j(\mathbf{X}_i) = \sum_{i=1}^n c_j(\mathbf{X}_i) \quad \text{and} \\ & && \sum_{i=1}^n p_i Z_i c_j(\mathbf{X}_i) = \sum_{i=1}^n c_j(\mathbf{X}_i) \quad \text{for all } j = 1, 2, \dots, m. \end{aligned} \quad (29)$$

Chan et al. (2015) assume uniform sampling weights. Equation (29) is solved by defining the functions $h(p) \equiv G(1 - p)$ and $g(v) \equiv h[(\nabla h)^{-1}(v)] + v - v(\nabla h)^{-1}(v)$ with $p \in \Delta$ and $v \in \mathfrak{R}$ to write the dual objective functions, which we use to solve for

$$\begin{aligned} \hat{\lambda}_0 &= \arg \max_{\lambda \in \mathfrak{R}^m} \sum_{i=1}^n \left\{ g \left[\sum_{j=1}^m (1 - Z_i)c_j(\mathbf{X}_i)\lambda_j \right] - \sum_{j=1}^m c_j(\mathbf{X}_i)\lambda_j \right\} \quad \text{and} \\ \hat{\lambda}_1 &= \arg \max_{\lambda \in \mathfrak{R}^m} \sum_{i=1}^n \left\{ g \left[\sum_{j=1}^m Z_i c_j(\mathbf{X}_i)\lambda_j \right] - \sum_{j=1}^m c_j(\mathbf{X}_i)\lambda_j \right\}. \end{aligned} \quad (30)$$

The resulting balancing weights are obtained by evaluating the first derivative of g at $\hat{\lambda}_0$ and $\hat{\lambda}_1$,

$$\hat{p}(\mathbf{X}_i) = \nabla g \left[\sum_{j=1}^m (1 - Z_i)c_j(\mathbf{X}_i)\hat{\lambda}_{j0} + \sum_{j=1}^m Z_i c_j(\mathbf{X}_i)\hat{\lambda}_{j1} \right]. \quad (31)$$

Similar to our solution for finding balancing weights to estimate τ_{ATE} , the dual variable $\lambda \equiv (\lambda_0, \lambda_1)^T$ has $2m$ entries. Tseng and Bertsekas (1987) and Chan et al. (2015) show that (30) can be solved using any strictly concave $g(v)$, $v \in \mathfrak{R}$, assuming that a feasible solution for (29) exists. Therefore, calibration estimators are not necessarily restricted to Bregman distances. However, if G is a monotone increasing transformation of D_f with respect to $\mathbf{p} \in \Delta^n$, then (29) can be constructed using Bregman distances so that the primal solutions are equivalent. The only difference from the methods we present in Section 3.3 is with the construction of \mathbf{A} and \mathbf{b} . In Theorem 3, we identify the conditions for which our proposed method is equivalent to the calibration estimator approach of Chan et al. (2015).

Theorem 3 Suppose we have a generalized distance $G(\mathbf{p})$ which is a monotone increasing trans

THEOREM 5. Suppose we have a generalized distance $\mathcal{G}(\mathbf{p})$, which is a monotone increasing transformation of some Bregman distance $Df(\mathbf{p}||\mathbf{q})$ with respect to $\mathbf{p} \in \Delta^n$ and $\mathbf{q} \in \Delta^n$ is uniform. Then for

$$\Omega_0 = \left\{ \mathbf{p} : \sum_{i=1}^n p_i Z_i c_j(\mathbf{X}_i) = b_j \text{ and } \sum_{i=1}^n p_i (1 - Z_i) c_j(\mathbf{X}_i) = b_j \text{ for all } j = 1, 2, \dots, m \right\} \text{ and}$$

$$\Omega_1 = \left\{ \mathbf{p} : \sum_{i=1}^n p_i Z_i c_j(\mathbf{X}_i) = b_j \text{ and } \sum_{i=1}^n p_i Z_i c_j(\mathbf{X}_i) = \sum_{i=1}^n p_i (1 - Z_i) c_j(\mathbf{X}_i) \text{ for all } j = 1, 2, \dots, m \right\},$$

$\tilde{\mathbf{p}} = \hat{\mathbf{p}}$ where $\tilde{\mathbf{p}} = \arg \min_{\mathbf{p} \in \Omega_f \cap \Delta^n} G(\mathbf{p})$ and $\hat{\mathbf{p}} = \arg \min_{\mathbf{p} \in \Omega_f \cap \Delta^n} D_f(\mathbf{p} || \mathbf{q})$.

Theorem 1 of Chan et al. (2015) shows that calibration estimators can be used to produce consistent estimates for τ_{ATE} while also attaining the semiparametric efficiency bound described by Hahn (1998). This is accomplished using a nonparametric setup where the balance functions represent a basis for uniformly approximating $\mu_0(\mathbf{X})$, $\mu_1(\mathbf{X})$, and $\pi(\mathbf{X})$. Given this result and Theorem 3 implies that for a sufficiently rich set of balance functions, the Bregman distance weights in conjunction with (3) can produce consistent and efficient estimates of τ_{ATE} . This result is quite useful when the balance functions that determine either the outcome or the treatment assignment are unknown.

Without further defining the distance to be optimized in the primal problem, (30) and (31) are less flexible when considering nonuniform sampling weights. This is especially important when developing iterative estimation algorithms or dealing with more complex balance designs where the data are not sampled uniformly from the population of interest. Furthermore, calibration estimators, as they are described in Chan et al. (2015), achieve a three-way balance between the treated, the controls, and the combined treatment groups for estimating τ_{ATE} . As shown in Theorem 1, this condition is not required when Assumption 3 holds. Zhao (2019) also noted that this condition is not required to achieve global efficiency using covariate balance scoring rules.

5 | NUMERICAL STUDIES

5.1 | Homogeneous treatment effect simulation

In this section, we demonstrate the utility of the proposed methodology using simulated data that assumes a constant conditional average treatment effect (Assumption 3). We generate 1,000 replications of several datasets determined by one of 72 experimental scenarios. For each dataset, we find balancing weights from four different covariate balancing methods to estimate τ_{ATE} . They are:

1. (IPW) Inverse probability of treatment weights where the propensity score follows a logit model fit using maximum likelihood estimation;
2. (CBPS) Inverse probability of treatment weights where the propensity score is fit to an exactly specified logit model subject to (17). The propensity scores are fit using generalized method of moments as implemented in the CBPS package (Fong, Ratkovic, & Imai, 2019);
3. (SENT) Balancing weights that are estimated by minimizing the shifted relative entropy following the results of (27) and (28). Using these balancing weights instead of (19) and (20)

allows us to test the effect of overspecifying the linear constraints when we know Assumption 3 is satisfied;

4. (BENT) Balancing weights that are estimated by minimizing the binary relative entropy subject to the constraints in (21) via the dual and primal solutions of (22) and (23).

We consider an extensive set of experimental scenarios adapted from those examined by Kang and Schafer (2007). These scenarios vary the sample size $n \in \{200, 1,000\}$, the error variance $\sigma^2 \in$

$\{2, 5, 10\}$, the generative process that determines the treatment assignment (indexed by $\{a, b\}$), the outcome process (indexed by $\{a, b\}$), and the correlation between the potential outcomes, $\rho \in \{-0.3, 0, 0.5\}$. The covariates to be balanced (i.e., the balance functions) are distributed as $X_1, X_2, X_3, X_4 \sim \mathcal{N}(0, 1)$. Define the transformations $U_1 = \exp(X_1/2)$, $U_2 = X_2/[1 + \exp(X_1)] + 10$, $U_3 = (X_1 X_3/25 + 0.6)^3$, and $U_4 = (X_2 + X_4 + 20)^2$. The vector $(U_1, U_2, U_3, U_4)^T$ is subsequently standardized to have a mean of zero and marginal variances of one.

The probability that a subject receives the treatment is then determined using the inverse logit link function,

$$\pi_i^{(k)} = \frac{\exp[\eta_i^{(k)}]}{1 + \exp[\eta_i^{(k)}]}, \quad k \in \{a, b\}.$$

Scenarios a and b distinguish whether the log odds of the propensity score is either linear or nonlinear with

$$\begin{aligned} \eta_i^{(a)} &= -X_{i1} + 0.5X_{i2} - 0.25X_{i3} - 0.1X_{i4} \quad \text{and} \\ \eta_i^{(b)} &= -U_{i1} + 0.5U_{i2} - 0.25U_{i3} - 0.1U_{i4}. \end{aligned} \quad (32)$$

The treatment indicators are generated by sampling $Z_i \sim \text{Bin}(1, \pi_i^{(k)})$. For the outcome process, we use the bivariate model

$$\begin{bmatrix} Y_i(0) \\ Y_i(1) \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \mu_i^{(\ell)} \\ \mu_i^{(\ell)} + \tau \end{bmatrix}, \begin{bmatrix} \sigma^2 & \rho\sigma^2 \\ \rho\sigma^2 & \sigma^2 \end{bmatrix} \right),$$

where $\ell \in \{a, b\}$ indexes

$$\begin{aligned} \mu_i^{(a)} &= 210 + 27.4X_{i1} + 13.7X_{i2} + 13.7X_{i3} + 13.7X_{i4} \quad \text{and} \\ \mu_i^{(b)} &= 210 + 27.4U_{i1} + 13.7U_{i2} + 13.7U_{i3} + 13.7U_{i4}. \end{aligned} \quad (33)$$

Once the potential outcomes have been generated, the observed outcome is the potential outcome corresponding to the observed treatment assignment. Each of the covariate balancing methods listed above is provided the design matrix with an intercept and the four original covariates: X_{i1}, X_{i2}, X_{i3} , and X_{i4} for $i = 1, 2, \dots, n$. The causal effect is then estimated using (3) where we substitute $p(\mathbf{X}_i)$ with the balancing weights estimated by each method.

We found that the correlation between the potential outcomes did not affect the resulting estimates of τ_{ATE} . In addition, the effects of altering σ^2 and n had anticipated results. Lower values of σ^2 led to lower standard errors of the causal effect estimate, whereas smaller values of n led to larger standard errors. Therefore, we report the results for $\rho = 0$, $n = 200$, and $\sigma^2 = 10$ in Table 1

and Figure 2. The complete results appear in the online supplement.

For all the methods that we tested, if either the outcome model or the treatment assignment is correctly specified, then the causal effect estimate is unbiased. We see in Table 1 and Figure 2 that the balancing weights obtained with SENT perform as well, or better in some cases, than the exactly specified CBPS model, even though the balancing weights obtained with SENT have twice as many constraints. The Monte Carlo standard error and bias of the estimates for τ_{ATE} are smallest when BENT is used to estimate the balancing weights for every scenario we examined.

TABLE 1 Average estimate, Monte Carlo standard error (MC SE), residual mean squared error, and empirical bias of the constant conditional average treatment effect using the four methods for estimating balancing weights described in Section 5.1

	Treatment		IPW	CBPS	SENT	BENT
	Outcome assignment scenario	scenario				
Avg. estimate (MC SE)	a	a	19.60 (3.28)	20.02 (0.55)	20.02 (0.53)	20.02 (0.51)
	a	b	20.40 (2.70)	20.00 (0.50)	20.01 (0.49)	20.00 (0.49)
	b	a	19.36 (5.27)	19.67 (3.81)	19.75 (2.90)	19.90 (2.21)
	b	b	15.33 (3.57)	15.26 (3.28)	15.83 (2.73)	16.47 (2.31)
Mean squared error (bias)	a	a	10.91 (−0.40)	0.30 (0.02)	0.28 (0.02)	0.26 (0.02)
	a	b	7.42 (0.40)	0.25 (0.00)	0.24 (0.01)	0.24 (0.00)
	b	a	28.21 (−0.64)	14.63 (−0.33)	8.49 (−0.25)	4.87 (−0.10)
	b	b	34.54 (−4.67)	33.14 (−4.74)	24.79 (−4.17)	17.83 (−3.53)

Note: IPW uses inverse probability of treatment weights estimated from a generalized linear model, CBPS uses the covariate balance propensity score weights, SENT uses the constrained optimal solution of the shifted relative entropy, and BENT uses the constrained optimal solution of the binary relative entropy.

This is expected since these weights are used for estimating the τ_{OWATE} (Crump et al., 2006), and because every condition necessary to ensure that an estimator for τ_{OWATE} is also an estimator for τ_{ATE} are met. The Monte Carlo standard errors and mean square error are also uniformly smaller for the average treatment effect estimates when using balancing weights estimated by CBPS and SENT versus IPW. This result indicates that methods which exactly balance the empirical covariate distributions perform better in finite sample settings. When both the outcome and treatment assignment models are misspecified, the four methods for finding balancing weights all produced biased estimates of τ_{ATE} . In these completely misspecified scenarios, the balancing weights estimated with BENT produce the least amount of bias and the lowest standard error for estimating τ_{ATE} .

5.2 | Heterogeneous treatment effect simulation

In this section, we simulate an additional 72 scenarios with a linear conditional average treatment effect to test our proposed methods under Assumption 4. We use the same covariate distributions for $(X_1, X_2, X_3, X_4)^T$ and $(U_1, U_2, U_3, U_4)^T$ as in Section 5.1. We also recycle the conditional mean functions $\mu_i^{(\ell)}$, $\ell \in \{a, b\}$, from (33). To generate the linear conditional average treatment effects, define

(a)

$$\begin{aligned} o_i &= 20 - 13.7X_{i1} + 13.7X_{i4} \quad \text{and} \\ \delta_i^{(b)} &= 20 - 13.7U_{i1} + 13.7U_{i4}. \end{aligned}$$

For outcome scenarios a and b , the bivariate outcome model is defined as

$$\begin{bmatrix} Y_i(0) \\ Y_i(1) \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \mu_i^{(\ell)} \\ \mu_i^{(\ell)} + \delta_i^{(\ell)} \end{bmatrix}, \begin{bmatrix} \sigma^2 & \rho\sigma^2 \\ \rho\sigma^2 & \sigma^2 \end{bmatrix} \right), \quad \ell \in \{a, b\},$$

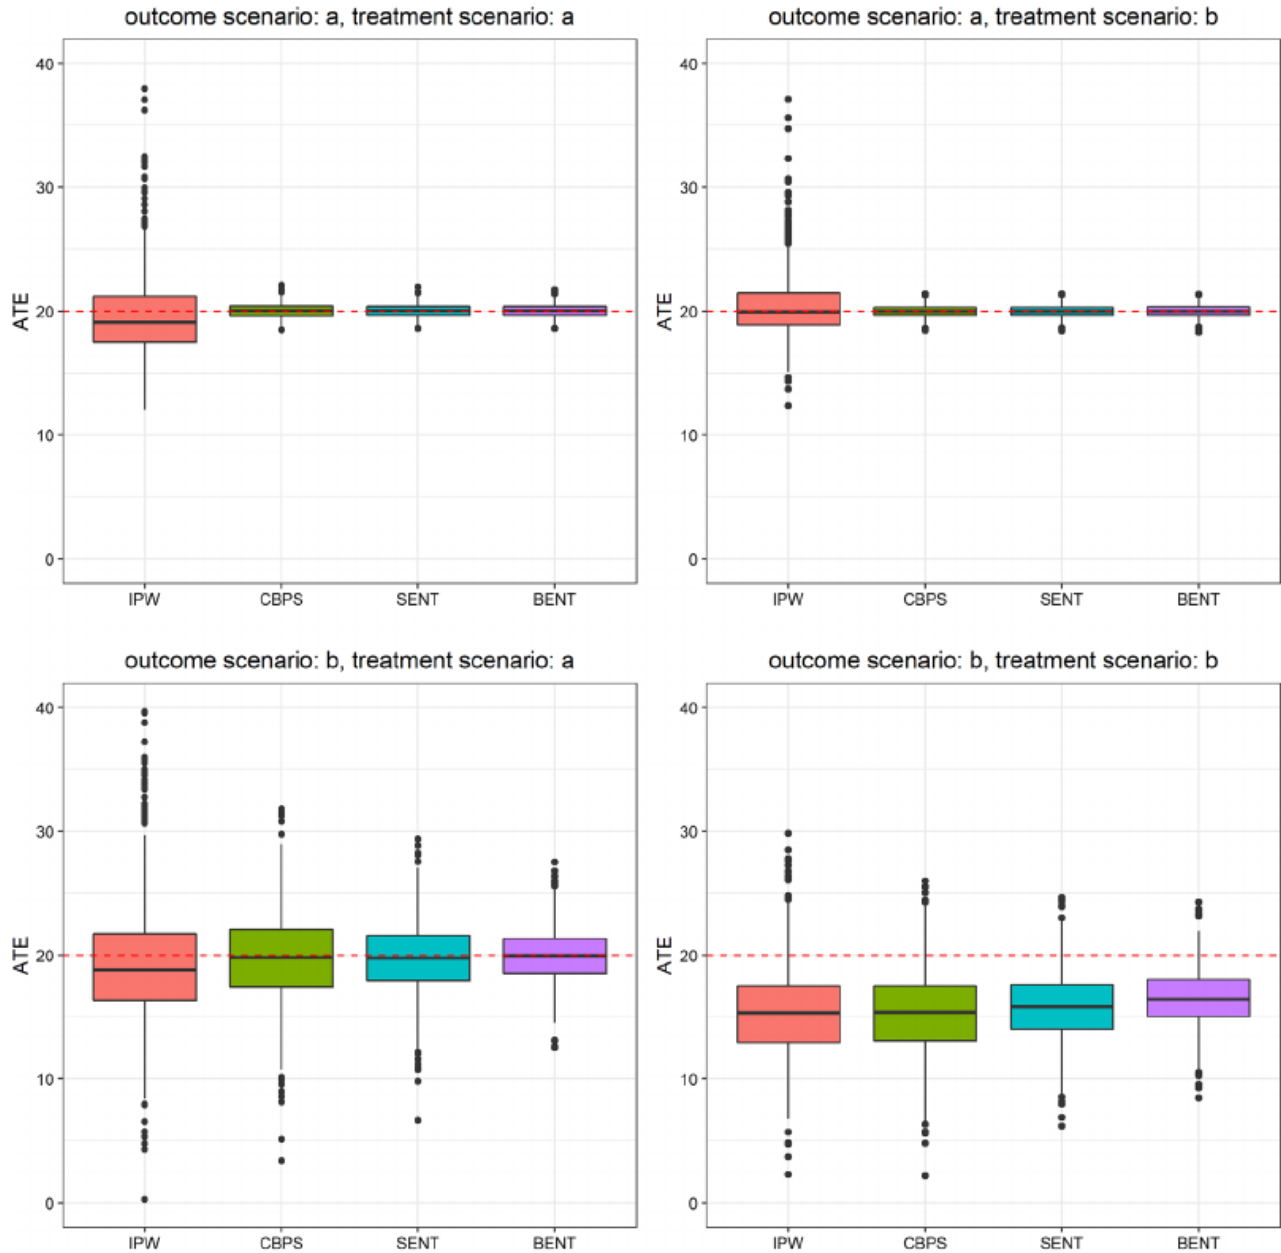


FIGURE 2 A subset of the constant conditional average treatment effect estimates using four different methods for estimating balancing weights. Each boxplot is composed of 1,000 estimates from the replicates that generate the values in Table 1 [Colour figure can be viewed at wileyonlinelibrary.com]

from which we sample $n \in \{200, 1,000\}$ entries. Each unit's treatment assignment is sampled from $\text{Bin}(1, \pi_i^{(k)})$, where $\pi_i^{(k)}$ is determined by $\eta_i^{(k)}$, $k \in \{a, b\}$, which are defined in (32). Similar to the simulations conducted in Section 5.1, we also vary $\sigma^2 \in \{2, 5, 10\}$ and $\rho \in \{-0.3, 0, 0.5\}$. For this set of scenarios, we examine five different covariate balance methods:

Printed by [Pennsylvania State University - 071.058.108.011 - /doi/epdf/10.1111/sjos.12457] at [14/10/2020].

1. (AIPW) Augmented inverse probability weights which uses the estimator

$$\hat{\tau}_{\text{AIPW}} = \frac{1}{n} \sum_{i=1}^n \left\{ \frac{Z_i Y_i}{\hat{\pi}(\mathbf{X}_i)} - \frac{[Z_i - \hat{\pi}(\mathbf{X}_i)] \hat{\mu}_1(\mathbf{X}_i)}{\hat{\pi}(\mathbf{X}_i)} - \frac{(1 - Z_i) Y_i}{1 - \hat{\pi}(\mathbf{X}_i)} - \frac{[Z_i - \hat{\pi}(\mathbf{X}_i)] \hat{\mu}_0(\mathbf{X}_i)}{1 - \hat{\pi}(\mathbf{X}_i)} \right\}. \quad (34)$$

$\hat{\mu}_1(\mathbf{X})$ is fit using linear regression on the treated group and $\hat{\mu}_0(\mathbf{X})$ is fit using linear regression on the controls. $\hat{\pi}(\mathbf{X})$ is fit with logistic regression;

2. (CAL) Calibration weights that solve (30) and (31) for $g(v) = \exp(-v)$, $v \in \mathfrak{R}$. This is equivalent to minimizing the unnormalized relative entropy subject to the linear constraints in (29). The R package ATE developed by Haris and Chan (2015) is used to estimate these balancing weights;
3. (iCBPS) Inverse probability of treatment weights where the propensity scores follow an exactly specified logit model that must satisfy (25) following the results of Fan et al. (2016). The propensity score is estimated using generalized method of moments as implemented in the CBPS package (Fong et al., 2019);
4. (hdCBPS) An augmented version of CBPS that extends (34) by using regularized regression techniques to find $\hat{\mu}_1(\mathbf{X})$ and $\hat{\mu}_0(\mathbf{X})$. The R package CBPS (Fong et al., 2019) is used to implement this method;
5. (SENT) Balancing weights that minimize the shifted relative entropy conditioned on the constraints in (26) where the estimates are obtained from (27) and (28).

For CAL, iCBPS, and SENT, the resulting balancing weights are substituted for $p(\mathbf{X})$ in (3) to estimate τ_{ATE} . The augmented approach of AIPW was first proposed by Robins, Rotnitzky, and Zhao (1994) while hdCBPS uses the augmented estimator proposed by Ning, Peng, and Imai (2018).

As with the previous simulation study, it appears that the correlation between potential outcomes is inconsequential, while the Monte Carlo standard errors predictably decrease when either n increases or σ^2 decreases. A representative selection of results from the experiment where $\rho = 0$, $n = 200$, and $\sigma^2 = 10$ is found in Table 2 and Figure 3. The complete results can be found in the online supplement. This simulation demonstrates that all five methods enjoy the doubly robust property described in Theorem 2. For each scenario, CAL and SENT had similar levels of bias and variation despite using different criterion distance functions. Even though the constraints of iCBPS and SENT are the same, the differences between the optimization techniques of the two methods become quite apparent. The Monte Carlo standard error of the estimates for τ_{ATE} using balancing weights obtained with SENT is smaller than the standard error of the estimates using balancing weights found with iCBPS. AIPW, CAL, and SENT performed similarly whenever the outcome was correctly specified. However, when the outcome model is misspecified and the propensity score is correctly specified, the Monte Carlo standard error and mean squared error were greater with AIPW than SENT and CAL. This suggests, and is further confirmed by hdCBPS, that methods which exactly balance covariate distributions can improve the efficiency of a doubly robust estimator in finite samples. The estimates from hdCBPS performed about as well as SENT and CAL. However, this method was proposed to alleviate issues encountered with high-dimensional covariate data, rendering many of its benefits redundant in this low-dimensional simulation study.

5.3 | Illustrative example of unplanned readmissions after lung

resection

Next, we investigate the results of a real dataset using different weighting and matching methods. In Bhagat et al. (2017), the odds of unplanned, 30-day readmissions are compared between lung cancer patients who receive thoracoscopic versus open lung resections. The study identified 9,510 patients who underwent some form of lung resection from the American College of Surgeons—National Safety and Quality Innovation Program (ACS-NSQIP) database. Of those

TABLE 2 Average estimate, Monte Carlo standard error (MC SE), residual mean squared error, and empirical bias of the linear conditional average treatment effect using the five methods for estimating balancing weights described in Section 5.2

	Treatment		AIPW	CAL	iCBPS	hdCBPS	SENT
	Outcome scenario	assignment scenario					
Avg. estimate (MC SE)	a	a	19.93 (1.51)	19.93 (1.49)	19.06 (3.94)	19.93 (1.51)	19.93 (1.49)
	a	b	20.09 (1.48)	20.09 (1.47)	18.05 (3.50)	20.09 (1.47)	20.09 (1.47)
	b	a	19.80 (7.35)	20.43 (2.76)	19.72 (3.63)	19.59 (3.02)	20.25 (2.78)
	b	b	15.04 (4.52)	16.91 (2.55)	14.21 (4.48)	16.15 (2.76)	16.77 (2.54)
Mean squared error (bias)	a	a	2.28 (−0.07)	2.22 (−0.07)	16.41 (−0.94)	2.28 (−0.07)	2.23 (−0.07)
	a	b	2.18 (0.09)	2.16 (0.09)	16.08 (−1.95)	2.17 (0.09)	2.16 (0.09)
	b	a	53.98 (−0.20)	7.77 (0.43)	13.26 (−0.28)	9.28 (−0.41)	7.78 (0.25)
	b	b	45.04 (−4.96)	16.01 (−3.09)	53.58 (−5.79)	22.45 (−3.85)	16.84 (−3.23)

Note: AIPW uses an augmented inverse probability of treatment approach where the propensity scores are estimated from a generalized linear model, CAL uses the calibration estimated weights, iCBPS uses the covariate balance propensity score weights, hdCBPS is an augmented version of CBPS, and SENT uses the constrained optimal solution of the shifted relative entropy.

9,510 patients, 4,935 (51.9%) received a thoracoscopic resection and 4,575 (48.1%) received an open anatomic resection. The study analysis carried out a greedy one-to-one matching of patients using the estimated propensity score as the criterion matching function (Ho, Imai, King, & Stuart, 2007). The propensity scores were fit with standard logistic regression. This algorithm matched 3,399 thoracoscopic lung resection patients to 3,399 open anatomic lung resection patients, dropping 2,712 patients (28.5%). In doing so, the “treated” group are assumed to be the patients who receive thoracoscopic lung resections and represent a random sample of the target population. Thus, the casual effect being estimated is the average treatment effect of the treated.

We replicated the study conducted in Bhagat et al. (2017) by estimating balancing weights using two different methods. The first method uses entropy balancing (EB) where the estimated balancing weights are obtained with (14) and (15). Recall that these balancing weights applied to (3) is doubly robust (Zhao & Percival, 2017). The second method fits a propensity score model using logistic regression (IPW). With the fitted propensity score, we then use (2) substituting $\hat{\pi}(\mathbf{X})$ for $\pi(\mathbf{X})$ to estimate τ_{ATT} . The causal effect estimates using the propensity score matched (PSM) cohort from the original article are also reported along with the unadjusted (UN) results in Table 3.

Figure 4 shows the amount of imbalance observed for each of the covariates among those

included in the covariate balancing models. We see that across each covariate, entropy balancing perfectly balances the first sample moments of the covariate distribution between the two treatment groups. Logistic regression appears to be less adequate at balancing the covariate moments than matching. However, aside from hospital length of stay, each of the weighted mean differences fell within the conservative 0.05 unit threshold using the inverse probability of treatment weights. The unadjusted differences do not share the same success as their adjusted counterparts, suggesting that some form of balancing should be implemented. After estimating τ_{ATT} , notice in

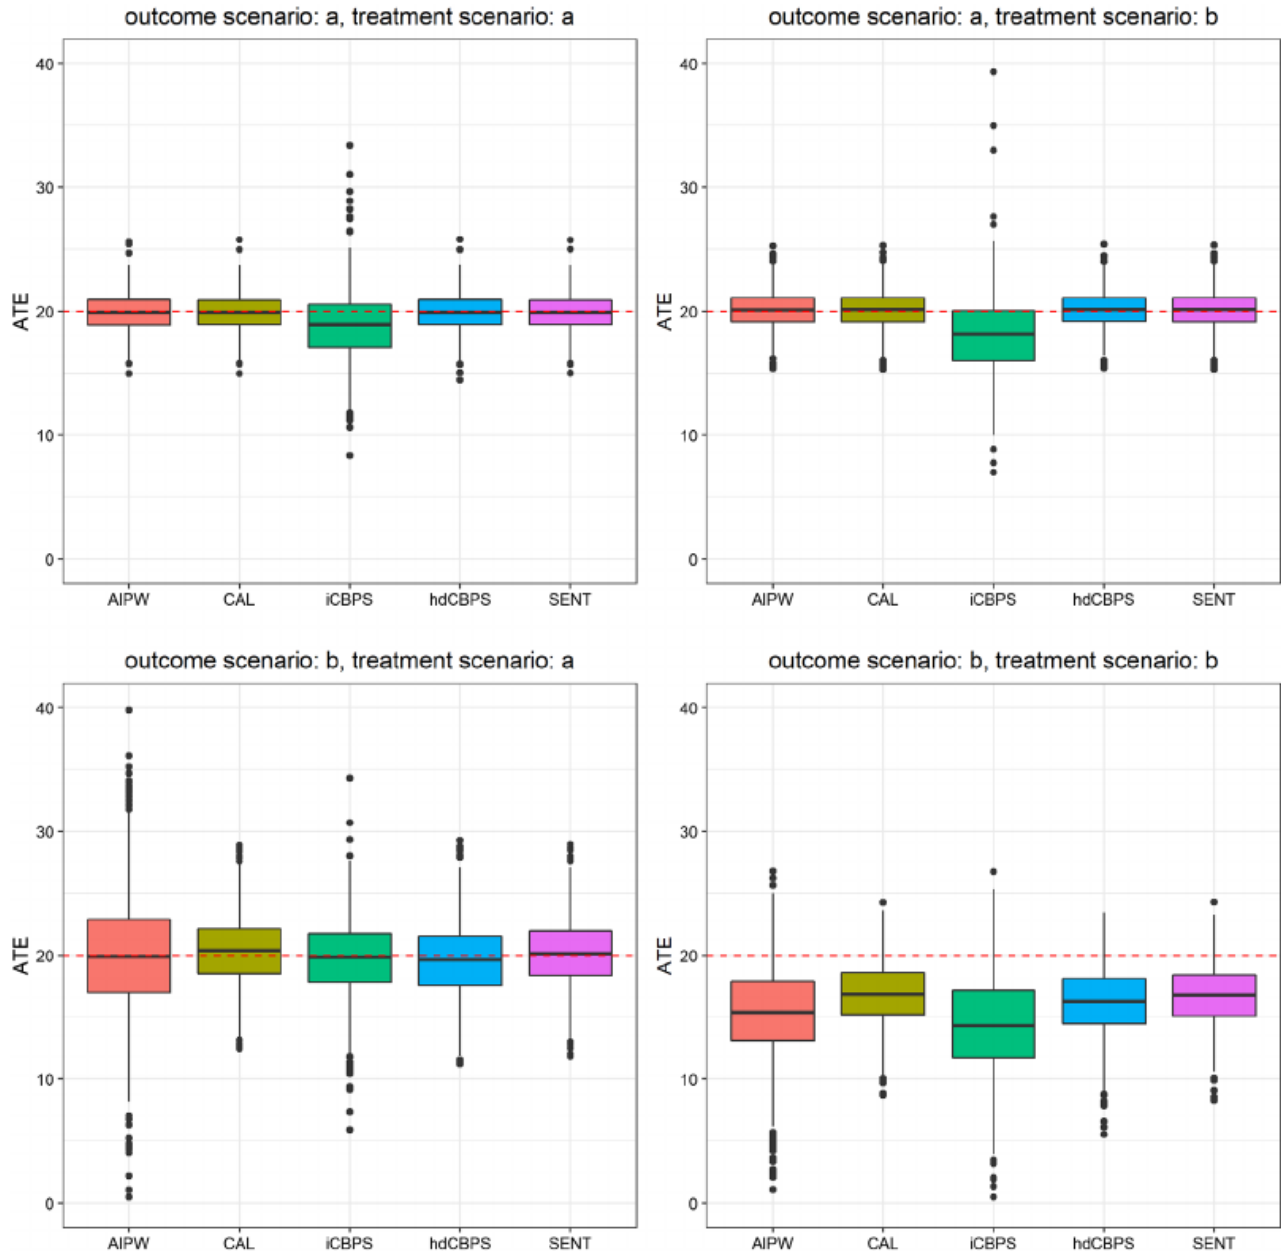


FIGURE 3 A subset of the linear conditional average treatment effect estimates using five different methods for estimating balancing weights. Each boxplot is composed of 1,000 estimates from the replicates that generate the values in Table 2 [Colour figure can be viewed at wileyonlinelibrary.com]

Table 3 that the estimated risk difference is significant when using either the inverse probability of treatment weights or entropy balancing, but is not significant when using propensity score matching. The difference is likely due to the 2,712 patients who were omitted when matching. This discrepancy illuminates and emphasizes the importance of selecting the most appropriate

Printed by [Pennsylvania State University - 071.058.108.011 - /doi/epdf/10.1111/sjos.12457] at [14/10/2020].

method for balancing covariate data, even within a large observational study.


6 | DISCUSSION

The generalized projection of a Bregman distance from a vector of sampling weights onto a set of intersecting hyperplanes is a powerful and flexible tool for normalizing data. In particular,

TABLE 3 ATT estimates for 30-day unplanned readmission in thoracoscopic versus open lung resection patients

Balancing method	Risk difference	SE	95% confidence interval
UN	-0.021	0.006	(-0.032, -0.010)
PSM	-0.010	0.006	(-0.021, 0.001)
IPW	-0.015	0.006	(-0.027, -0.004)
EB	-0.016	0.006	(-0.027, -0.004)

Abbreviations: ATT, average treatment effect of the treated; EB, entropy balancing results; IPW, inverse probability weighted results; PSM, propensity score matched results; UN, unadjusted results.



Printed by [Pennsylvania State University - 071.058.108.011 - /doi/pdf/10.1111/sjos.12457] at [14/10/2020].

0.0 0.1 0.2 0.3 0.4 0.5
Adjusted Absolute Mean Difference

FIGURE 4 Each point represents the adjusted absolute standardized mean difference (x -axis) between thoracoscopic and open lung resection patients. The covariates in the plot (y -axis) are included into each model. The red dotted line marks an absolute standardized mean difference of 0.05 [Colour figure can be viewed at wileyonlinelibrary.com]

this process is quite useful for constructing balancing weights for estimating causal effects. Using properties of dual optimization, we identify a doubly robust estimator for τ_{ATE} and the optimally weighted average treatment effect (Crump et al., 2006) in Theorem 1 and Corollary 1. We also show that the dual interpretation of improved CBPS (Fan et al., 2016) is doubly robust in Theorem 2. In Theorem 3, we present the conditions for which the balancing weights produced by (11) are the same as the weights produced by (31) suggested by Chan et al. (2015). When the true balance functions are unknown, we can use nonparametric methods similar to those suggested by Hirano et al. (2003) and Chan et al. (2015) within our framework to achieve global efficiency.

In the simulation studies we conducted, we observed that the balancing weights that are typically used to estimate τ_{OWATE} had the best performance for estimating τ_{ATE} when Assumption 3 is satisfied. We also observed that including additional constraints as in (8) sometimes had better performance than estimators that require fewer constraints, like CBPS. When we assume a linear conditional average treatment effect, our dual interpretation of iCBPS performed better than the analogous GMM estimator. We then apply our framework to a real dataset of lung resection patients. Here, we demonstrate how the choice of balancing method can have a critical impact on the results of a study.

There are several limitations to our proposed framework. First, each sampling unit's treatment assignment is assumed to be independent from the treatment assignment of the other sampling units. This assumption is sometimes called the no interference assumption. Health outcomes research is rich in observational data from the emergence of the electronic health record. While numerous in size, these datasets are more complex with patients being clustered within regions, hospitals, clinics, and/or practicing physicians. These are all factors that need to be accounted for in some way. How to extend these methods to clustered data settings is currently under investigation. Second, linear equality constraints are often quite stringent. If a particular covariate is difficult to balance, our proposed framework will sometimes fail to find the appropriate balancing weights. Zubizarreta (2015) propose using stable balancing weights which minimize the Euclidean distance subject to linear inequality constraints placed on the weighted sample moments of the covariate distribution. In more recent work, Wang and Zubizarreta (2020) have combined this interval constrained optimization approach with calibration estimators. There is also the issue where the balance functions that generate either the outcome or treatment assignment are high dimensional. This problem is not examined in the presented work. Ning et al. (2018) propose using an augmented approach with the CBPSs of Imai and Ratkovic (2013) and Fan et al. (2016) in the spirit of Robins et al. (1994) and Farrell (2015). Their proposed methodology boasts compelling results as the dimension of covariate distribution increases. It is possible that our approach could be extended using some of the methods proposed by Ning et al. (2018).

In addition to addressing some of the limitations identified in the previous paragraph, in future work we would also like to expand these methods to incorporate multivalued treatment assignments. This would entail modifying the Horvitz–Thompson estimator and also requires extending the constraint matrix \mathbf{A} and target margins \mathbf{b} to facilitate covariate balance between all pairwise combinations of the treatment assignments. Finally, we would like to further investigate methods

for generalizing causal effect estimates to a target population, which would involve estimating \mathbf{q} prior to estimating \mathbf{p} .

ACKNOWLEDGEMENTS

The authors would like to thank Dr. Robert Meguid and the Adult and Child Consortium for Health Outcomes Research and Delivery Science (ACCORDS) program for making the unplanned readmissions data available to us. They would also like to thank Dr. Peter DeWitt for his help

in deploying the associated R package `cba1`. This research is supported by a pilot grant from the Data Science to Patient Value (D2V) initiative from the University of Colorado. The authors acknowledge partial support from NSF award number DMS-1420451.

ORCID

Kevin P. Josey  <https://orcid.org/0000-0003-2490-6272>

REFERENCES

- Bang, H., & Robins, J. M. (2005). Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61(4), 962–973.
- Bertsekas, D. P. (1999). *Nonlinear Programming* (2nd ed.). Belmont, MA: Athena Scientific.
- Bhagat, R., Bronsert, M. R., Ward, A. N., Martin, J., Juarez-Colunga, E., Glebova, N. O., ... Meguid, R. A. (2017). National analysis of unplanned readmissions after thoracoscopic versus open lung cancer resection. *The Annals of Thoracic Surgery*, 104(6), 1782–1790.
- Bregman, L. M. (1967). The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR Computational Mathematics and Mathematical Physics*, 7(3), 200–217.
- Censor, Y., & Zenios, S. A. (1998). *Parallel optimization: Theory, algorithms, and applications* (1st ed.). New York, NY: Oxford University Press.
- Chan, K. C. G., Yam, S. C. P., & Zheng, Z. (2015). 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(3), 673–700.
- Crump, R. K., Hotz, V. J., Imbens, G. W., and Mitnik, O. A. (2006). Moving the goalposts: Addressing limited overlap in the estimation of average treatment effects by changing the estimand. Working Paper 330, National Bureau of Economic Research.
- Deville, J.-C., & Sarndal, C.-E. (1992). Calibration estimators in survey sampling. *Journal of the American Statistical Association*, 87(418), 376–382.
- Fan, J., Imai, K., Liu, H., Ning, Y., and Yang, X. (2016). Improving covariate balancing propensity score: A doubly robust and efficient approach. *Technical Report*.
- Farrell, M. H. (2015). Robust inference on average treatment effects with possibly more covariates than observations. *Journal of Econometrics*, 189(1), 1–23.
- Fong, C., Ratkovic, M., and Imai, K. (2019). CBPS: Covariate balancing propensity score. R package version 0.20.
- Hahn, J. (1998). On the role of the propensity score in efficient semiparametric estimation of average treatment effects. *Econometrica*, 66(2), 315–331.
- Hainmueller, J. (2012). Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies. *Political Analysis*, 20(1), 25–46.
- Haris, A. and Chan, G. (2015). ATE: Inference for average treatment effects using covariate balancing. R package version 0.2.0.
- Hirano, K., Imbens, G. W., & Ridder, G. (2003). Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica*, 71(4), 1161–1189.
- Ho, D. E., Imai, K., King, G., & Stuart, E. A. (2007). Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference. *Political Analysis*, 15(3), 199–236.
- Horvitz, D. G., & Thompson, D. J. (1952). A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association*, 47(260), 663–685.

- Imai, K., & Ratkovic, M. (2013). Covariate balancing propensity score. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *76*(1), 243–263.
- Imbens, G. W., & Rubin, D. B. (2015). *Causal inference for statistics, social, and biomedical sciences: An introduction* (1st ed.). New York, NY: Cambridge University Press.
- Kang, J. D. Y., & Schafer, J. L. (2007). Demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data. *Statistical Science*, *22*(4), 523–539.
- Li, F., Morgan, K. L., & Zaslavsky, A. M. (2018). Balancing covariates via propensity score weighting. *Journal of the American Statistical Association*, *113*(521), 390–400.

- Ning, Y., Peng, S., and Imai, K. (2018). Robust estimation of causal effects via high-dimensional covariate balancing propensity score. arXiv:1812.08683 [stat].
- Robins, J. M., Rotnitzky, A., & Zhao, L. P. (1994). Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association*, 89(427), 846–866.
- Rosenbaum, P. R., & Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1), 41–55.
- Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5), 688–701.
- Rubin, D. B. (2008). For objective causal inference, design trumps analysis. *The Annals of Applied Statistics*, 2(3), 808–840.
- Splawa-Neyman, J., Dabrowska, D. M., & Speed, T. P. (1990). On the application of probability theory to agricultural experiments. essay on principles. section 9. *Statistical Science*, 5(4), 465–472.
- Stefanski, L. A., & Boos, D. D. (2002). The calculus of M-estimation. *The American Statistician*, 56(1), 29–38.
- Tseng, P., & Bertsekas, D. P. (1987). Relaxation methods for linear programs. *Mathematics of Operations Research*, 12(4), 569–596.
- Wang, Y., & Zubizarreta, J. R. (2020). Minimal dispersion approximately balancing weights: Asymptotic properties and practical considerations. *Biometrika*, 107(1), 93–10.
- Zhao, Q. (2019). Covariate balancing propensity score by tailored loss functions. *The Annals of Statistics*, 47(2), 965–993.
- Zhao, Q., & Percival, D. (2017). Entropy balancing is doubly robust. *Journal of Causal Inference*, 5(1).
- Zubizarreta, J. R. (2015). Stable weights that balance covariates for estimation with incomplete outcome data. *Journal of the American Statistical Association*, 110(511), 910–922.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of this article.

How to cite this article: Josey KP, Juarez-Colunga E, Yang F, Ghosh D. A framework for covariate balance using Bregman distances. *Scand J Statist.* 2020;1–27.

<https://doi.org/10.1111/sjos.12457>

APPENDIX TECHNICAL PROOFS

Proofs for all Theorems, Lemmas, Propositions, and Corollaries presented in this article are available online at <http://www.interscience.wiley.com/journal/sjs>.

APPENDIX R PACKAGE AND SIMULATION CODE

The R package used to fit balancing weights as the generalized projection of Bregman distance is still in development with a working version available at <https://github.com/kevjosey/cbal>. The code used to conduct the simulation study in Section 5 is available at the following URL: <https://github.com/kevjosey/cbal-sim>. The code for replicating the study by Bhagat et al. (2017) is available from the authors upon request.