DE GRUYTER J. Causal Infer. 2015; 3(2): 139–155

Samuel David Lendle*, Bruce Fireman and Mark J. van der Laan

Balancing Score Adjusted Targeted Minimum Loss-based Estimation

Abstract: Adjusting for a balancing score is sufficient for bias reduction when estimating causal effects including the average treatment effect and effect among the treated. Estimators that adjust for the propensity score in a nonparametric way, such as matching on an estimate of the propensity score, can be consistent when the estimated propensity score is not consistent for the true propensity score but converges to some other balancing score. We call this property the balancing score property, and discuss a class of estimators that have this property. We introduce a targeted minimum loss-based estimator (TMLE) for a treatment-specific mean with the balancing score property that is additionally locally efficient and doubly robust. We investigate the new estimator's performance relative to other estimators, including another TMLE, a propensity score matching estimator, an inverse probability of treatment weighted estimator, and a regression-based estimator in simulation studies.

Keywords: balancing score, propensity score, causal inference, matching, TMLE

DOI 10.1515/jci-2012-0012

1 Introduction

Estimators based on the propensity score (PS), the probability of receiving a treatment given baseline covariates, are popular for estimation of causal effects such as the average treatment effect (ATE), average treatment effect among the treated (ATT), or the average outcome under treatment. Such methods can be thought of as adjusting for the propensity score in place of baseline covariates, and generally require consistent estimation of the propensity score if it is not known. Common propensity score methods include stratification or subclassification [1–3], inverse probability of treatment weighting (IPTW) [4, 5], and propensity score matching [6–8].

A "balancing score" as defined by Rosenbaum and Rubin [8] is a function of baseline covariates such that treatment and baseline covariates are independent conditional on that function. The propensity score is perhaps the most well-known example of a balancing score, but balancing scores are more general. Typically, propensity score-based methods are said to be consistent when the true propensity score is consistently estimated. Methods that adjust for the propensity score nonparametrically, such as matching or stratification by the propensity score, actually only need that the estimated propensity score converge to some balancing score in order for the parameter of interest to be estimated consistently. However, we are not aware of specific claims in the literature that particular propensity score-based methods are consistent under this weaker condition. We say that an estimator using the propensity score or other balancing score has the balancing score property if it is consistent when the estimated propensity score converges to a balancing score.

Though not guaranteed in general, it is possible for an estimated propensity score based on a misspecified model to converge to a balancing score that is not equal to the true propensity score.

***Corresponding author: Samuel David Lendle,** Group in Biostatistics, University of California, Berkeley, Berkeley, CA, USA, E-mail: lendle@stat.berkeley.edu
Bruce Fireman, Division of Research, Kaiser Permanente, Oakland, CA, USA, E-mail: Bruce.Fireman@kp.org
Mark J. van der Laan, Group in Biostatistics, University of California, Berkeley, Berkeley, CA, USA, E-mail: laan@berkeley.edu

Propensity score-based estimators that have the balancing score property are robust to this sort of estimator misspecification of the PS, while other propensity score-based estimators are not. The balancing score property is desirable because, even though most such estimators were initially developed based on the PS specifically, they inherit this robustness for free. Estimators with the balancing score property are in general not efficient.

An efficient estimator is one that achieves the minimum asymptotic variance of all regular estimators. In many cases, for example when estimating the ATE, ATT, and average outcome under treatment, doubly robust estimators can be constructed. A doubly robust estimator is one that relies on an estimate of both the propensity score and of the outcome regression, the conditional mean of the outcome given baseline covariates and treatment. Doubly robust estimators are consistent if either the estimated propensity score or outcome regression is consistent. Examples include targeted minimum loss-based estimation (TMLE) [9, 10] and augmented inverse probability of treatment weighted estimation (A-IPTW) [11, 12]. In addition to being doubly robust, both TMLE and A-IPTW are efficient when both the propensity score and outcome regression are consistently estimated.

In this article, we discuss a general class of estimators that have the balancing score property. We also construct a TMLE [9, 10] with the balancing score property. This new TMLE not only has the benefit of the robustness provided by the balancing score property, it also is a locally efficient, doubly robust plug-in estimator. This means that our new estimator retains all of the attractive properties of a traditional TMLE while gaining robustness that other estimators with the balancing score property enjoy when the propensity score only converges to a balancing score.

In Section 2, we introduce notation and define the statistical parameter we wish to estimate. In Section 3 we describe a TMLE for the statistical parameter. In Section 4 we discuss the balancing score property and describe the proposed new TMLE. In Section 5 we compare the performance of the new estimator to a traditional TMLE as well as other common estimator and conclude with a discussion in Section 6. A list of notation used throughout the article is provided in Appendix A. Some results and proofs not included in the main text are in Appendix A.2 and two modifications to the TMLE algorithm are presented in Appendix A.3. An example implementation of the proposed new TMLE in R [13] is provided in Appendix A.4.

2 Preliminaries

Consider the random variable $O = (W, A, Y)$ where W is a real-valued vector, A is binary with values in $\{0, 1\}$ and Y is univariate real number. Call the probability distribution of O $P_0 \in \mathcal{M}$ where $\mathcal{M}$ is the statistical model. Assume $P_0(A = 1 \mid W) > 0$ for almost every W. This is sometimes called a positivity assumption. Define the parameter mapping Ψ from $\mathcal{M}$ to $\mathbb{R}$ that maps P to $E_P(E_P(Y \mid A = 1, W))$ where E_P denotes expected value under probability distribution $P \in \mathcal{M}$.

Suppose $A = 1$ indicates some treatment of interest and $A = 0$ represents some control or reference treatment, W represents a vector of baseline covariates measured before treatment, and Y represents some outcome measured after treatment. Then under additional causal assumptions, $\Psi(P_0)$ can be interpreted as a causal quantity. In particular, we may assume that observed treatment A is independent of the counterfactual outcome had each observation received treatment 1 given covariates W. This is known as the randomization assumption or the “no unmeasured confounders” assumption, and the validity depends on the particular application. Under the randomization positivity assumptions, $\Psi(P_0)$ can be interpreted as the average outcome had everyone in the population received treatment 1. In this paper we focus on estimation of the statistical parameter $\Psi(P_0)$, but other similar statistical parameters can, under assumptions, be interpreted as causal parameters such as the ATE or the ATT [14].

For a probability distribution $P \in \mathcal{M}$, $\bar{Q}(a, w) = E_P(Y \mid A = a, W = w)$ is the regression of the outcome on covariates and treatment. Let $Q_W(w) = P(W = w)$ be the distribution of baseline covariates. The conditional distribution of treatment on baseline covariates is called $g(a \mid w) = P(A = a \mid W = w)$, and define the propensity score as $\bar{g}(w) = g(1 \mid w)$, the probability of treatment given covariates w. The parameter mapping Ψ depends on P only through $Q = (\bar{Q}, Q_W)$, so recognizing the abuse of notation, we sometimes write $\Psi(P) = \Psi(Q) = \Psi(\bar{Q}, Q_W)$.

For a distribution $P \in \mathcal{M}$, we make no assumptions on the outcome regression $\bar{Q}$ or on the distribution Q_W of W. We may put some restriction on possible functions g, for example we may know that $P(A \mid W)$ depends only on a subset of W. The model $\mathcal{M}$ is therefore nonparametric or semiparametric.

Let $O_1, \ldots, O_n$ be a data set of n independent and identically distributed random variables drawn from P_0 where $O_i = (W_i, A_i, Y_i)$. We use the subscript 0 to denote the true probability distribution, and n to denote an estimate based on a dataset of size n, so, for example, E_0 denotes expectation with respect to P_0, $\bar{Q}_0(a, w) = E_0(Y \mid A = a, W = w)$, and $\bar{Q}_n$ is an estimate of $\bar{Q}_0$. Let $\psi_0 = \Psi(P_0)$.

3 Targeted minimum loss-based estimation

A plug-in estimator takes an estimate of the distribution P_0, or relevant parts of P_0, and plugs it into the parameter mapping Ψ. In this case, Ψ depends on P through $\bar{Q}$ and Q_W. Using an estimate $\bar{Q}_n$ of $\bar{Q}_0$, and letting Q_{Wn} be the empirical distribution of W, we can calculate the plug-in estimate as

$$\begin{aligned}\Psi(Q_n) &= \int_w \bar{Q}_n(1, W) dQ_{Wn}(w) \\ &= \frac{1}{n}\sum_{i=1}^{n} \bar{Q}_n(1, W_i).\end{aligned}$$

That is, we take the mean of $\bar{Q}_n(1, W)$ with respect to the empirical distribution of W. Plug-in estimators are desirable because they fully utilize known global constraints of Q_0 (by using an estimate Q_n that satisfies these constraints) and guarantee that estimates are in the parameter space, even in small samples. Non-plug-in estimators such as IPTW can produce estimates outside of the parameter space. For instance if our estimand is a probability, a method like IPTW could yield an estimate outside of $[0, 1]$ when the sample size is small.

TMLE is a general framework for constructing a plug-in estimator for ψ_0 with additional properties such as efficiency. TMLE takes an initial estimate of the outcome regression $\bar{Q}_0$, say $\bar{Q}_n^0$, and, using an estimate $\bar{g}_n(W)$ of the propensity score, updates it to $\bar{Q}_n^*$. Using the empirical distribution of W along with the updated $\bar{Q}_n^*$, the final estimate is calculated as $\Psi(\bar{Q}_n^*, Q_{Wn})$. The updated $\bar{Q}_n^*$ is constructed in such a way that the final estimate is efficient or attains other properties. We now review some background and a specific implementation of the TMLE procedure for $\Psi(P_0)$.

An estimator that is asymptotically linear can be written as

$$\sqrt{n}(\psi_n - \psi_0) = \frac{1}{\sqrt{n}}\sum_{i=1}^{n} IC(P_0)(O_i) + o_P(1)$$

for some mean zero function $IC(P_0)$ where $o_P(1)$ is a term that converges in probability to 0. The function $IC(P_0)$ is called the influence curve of the estimator at P_0. For an estimator to be efficient, that is, to have the minimum asymptotic variance among all regular estimators, it must be asymptotically linear with influence curve equal to the so called efficient influence curve [9, 15]. The efficient influence curve for a particular parameter mapping Ψ depends on the model. For our model, regardless of the model for g_0, the efficient influence curve at a $P \in \mathcal{M}$ written in terms of Q and g is

$$D^*(\bar{Q}, Q_W, g)(O) = \frac{A}{g(1 \mid W)}(Y - \bar{Q}(A, W)) + \bar{Q}(1, W) - \Psi(\bar{Q}, Q_W).$$

A derivation of the efficient influence curve is presented in Chapter 4 van der Laan and Rose [9].

Suppose for now Y is binary or bounded by 0 and 1. A modification to the algorithm and a different TMLE are described in Appendix A.3 if this is not the case. The initial estimate $\bar{Q}_n^0$ can be obtained via a parametric model for $E_0(Y \mid A, W)$, such as a generalized linear model [16], or with a data adaptive machine learning algorithm such as the SuperLearner algorithm [9, 17], which combines parametric and data adaptive estimators using cross-validation.

The updating step is defined by a choice of loss function L for Q such that $E_0 L(Q)(O)$ is minimized at Q_0, and a working parametric submodel with finite dimensional real-valued parameter ε, $\{Q(\varepsilon) : \varepsilon\}$ such that $Q(0) = Q$. The submodel is typically chosen so that the efficient influence curve is in the linear span of the components of the "score" $\frac{d}{d\varepsilon} L(Q(\varepsilon)(O))$ at $\varepsilon = 0$. When L is the negative log likelihood, $\frac{d}{d\varepsilon} L(Q(\varepsilon)(O))$ is the score in the usual sense. Starting with $k = 0$, the empirical risk minimizer $\varepsilon_n^k = \arg\min_\varepsilon \sum_{i=1}^n L(Q_n^k(\varepsilon))(O_i)$ is calculated and Q_n^k is updated to $Q_n^{k+1} = Q_n^k(\varepsilon_n^k)$. The process is iterated until $\varepsilon^k \approx 0$, sometimes converging in one step. Details can be found in Refs [9, 10, 18, 19].

Define the loss function $L(Q)(O) = L_Y(\bar{Q})(O) + L_W(Q_W)(O)$ where

$$L_Y(\bar{Q})(O) = -Y \log(\bar{Q}(A, W)) - (1 - Y)\log(1 - \bar{Q}(A, W)).$$

and $L_W(Q_W)(O) = -\log(Q_W(W))$. When Y is binary, $L_Y(\bar{Q})(O)$ is the negative conditional log likelihood of the Bernoulli distribution. Because Y is at least bounded by 0 and 1 if not binary, $L_Y(\bar{Q})(O)$ is a valid loss function for the conditional mean. That is, $\bar{Q}_0 = \arg\min_{\bar{Q}} E_0 L_Y(\bar{Q})(O)$ [20]. The function $L_W(Q_W)(O)$ is the negative log likelihood of the distribution of W, and its true mean is minimized by Q_{W0}. Thus, the sum loss function is a valid loss function for $Q_0 = (\bar{Q}_0, Q_{W0})$.

For a working submodel for $\bar{Q}$, we use

$$\bar{Q}(\varepsilon)(A, W) = \text{logit}^{-1}\left[\text{logit}(\bar{Q}(A, W)) + \varepsilon \frac{A}{g(1 \mid W)}\right]$$

indexed by ε. We call this a logistic working model because it is a logistic regression model with offset logit$\bar{Q}(A, W)$ and single covariate $A/g(1 \mid W)$. The score of this model at $\varepsilon = 0$ is

$$\frac{A}{g(1 \mid W)}(Y - \bar{Q}(A, W)).$$

For Q_W, we can use as working submodel

$$Q_W(\varepsilon')(W) = \{1 + \varepsilon'[\bar{Q}(1, W) - \Psi(Q)]\} Q_W(W)$$

which has score $\bar{Q}(1, W) - \Psi(\bar{Q}, Q_W)$ at $\varepsilon' = 0$. We can see that the efficient influence curve $D^*(P_0)$ can be written as a linear combination of the scores of these submodels when $Q = Q_0$ and $g = g_0$.

The estimate ε_n^0 can be calculated using standard logistic regression software with logit$(\bar{Q}_n^0(A, W))$ as a fixed offset term, and $A/g_n(1 \mid W)$ as a covariate. By using the empirical distribution of W as an initial estimate for Q_{Wn}^0, and negative log likelihood loss function for L_W, the empirical risk is already minimized at Q_{Wn}^0, so $\varepsilon_n'^0 = 0$ and no update is needed. In this case, the algorithm converges in one step, because $\frac{A}{g_n(1 \mid W)}$ is not updated between iterations, so an additional update to $\bar{Q}_n^1$ will yield $\varepsilon_n^1 = 0$. The estimate $\bar{Q}_n^* = \bar{Q}_n^0(\varepsilon_n^0)$ and the TMLE estimate of $\Psi(P_0)$ is calculated as

$$\Psi(\bar{Q}_n^*, Q_{Wn}) = \frac{1}{n} \sum_{i=1}^n \bar{Q}_n^*(1, W_i).$$

Under regularity conditions, the TMLE is asymptotically linear and doubly robust, meaning that if the initial estimate $\bar{Q}_n^0$ is consistent for $\bar{Q}_0$, or $\bar{g}_n$ is consistent for $\bar{g}_0$, then $\Psi(\bar{Q}_n^*, Q_{Wn})$ is consistent for $\Psi(P_0)$. Additionally, when both $\bar{Q}_n^0$ and g_n are consistent, the influence curve of the TMLE is equal to the efficient influence curve, so the estimator achieves the semiparametric efficiency bound. Precise regularity conditions for asymptotic linearity and efficiency are presented in Appendix A.2 in Theorem 3.

4 Balancing score property and proposed estimator

A function b of W is called a balancing score if $A \perp W \mid b(W)$ [8]. Trivially, $b(W) = W$ is a balancing score, and by definition of the propensity score, $\bar{g}_0(W)$, is a balancing score. In general, any function $b(W)$ is a balancing score if and only if there exists some function f such that $\bar{g}_0(W) = f(b(W))$ (Theorem 2 [8]). For example, any monotone transformation of the propensity is a balancing score. Such a function is called a "balancing score" because, conditional on $b(W)$, the distribution of W between the treated and untreated observations is equal or balanced. That is, $P_0(W \mid A = 1, b(W)) = P_0(W \mid A = 0, b(W))$. Rosenbaum and Rubin [8] show that adjusting for a balancing score yields the same estimand as adjusting for the full set of covariates W which we state in Lemma 1 and offer a different proof in Appendix A.2.

Lemma 1 *If $b(W)$ is a balancing score under distribution P, then $E_P(E_P(Y \mid A = 1, b(W))) = \Psi(P)$.*

This result gives rise to methods for estimating $\Psi(P_0)$ based on a balancing score and not on an estimate of $\bar{Q}_0$. The propensity score is the balancing score most commonly used for estimating $\Psi(P_0)$, and frequently used estimators include propensity score matching, stratification, and IPTW. When the propensity score is not known, these estimators rely on an estimated propensity score $\bar{g}_n$, and, under regularity conditions, are consistent when $\bar{g}_n$ is consistent for $\bar{g}_0$. The IPTW estimator, in particular, requires that $\bar{g}_n$ converges to $\bar{g}_0$ for consistency. However, many of these methods, such as propensity score matching and stratification by the propensity score, can be seen as nonparametrically adjusting for the propensity score and only rely on the propensity score being a balancing score. For these estimators, it is sufficient for $\bar{g}_n$ to converge to some balancing score under P_0. We call this property the balancing score property.

In practice, an estimator $\bar{g}_n$ can approximate a balancing score well but not converge to the true propensity score. A parametric logistic regression estimator will estimate some function of the covariates that is a projection of $\bar{g}_0$ onto the model determined by the parametrization of the estimator. If the parametric estimator is correctly specified, this projection will be $\bar{g}_0$. Depending on the true $\bar{g}_0$ and distribution of covariates, it is possible for this projection to be a balancing score or at least approximate some balancing score when the estimator is not correctly specified. For example, suppose the true $\bar{g}_0$ depends on higher order interactions of covariates. Though not the case in general, in some settings a main terms logistic regression may approximate a balancing score well. We explore such a setting via simulation in Section 5. In another example, suppose $\bar{g}_0$ depends on covariates in an additive on the logit scale but not necessarily linear or even smooth way. A logistic regression estimator with linear or possibly higher order polynomial main terms may again approximate some balancing score.

Estimators based only on the propensity score are not doubly robust. We now construct a locally efficient doubly robust estimator with the balancing score property. We start with initial estimators $\bar{Q}_n$ for $\bar{Q}_0$ and $\bar{g}_n$ for $\bar{g}_0$. We then update $\bar{Q}_n$ by nonparametrically regressing Y on A and $\bar{g}_n(W)$ using $\bar{Q}_n(A, W)$ as an offset. Similarly to the TMLE procedure in Section 3, we use this updated estimate of $\bar{Q}_0$ to estimate ψ_0 by plugging it in to the parameter mapping Ψ along with the empirical distribution of W.

To update $\bar{Q}_n$ by further adjusting for A and $\bar{g}_n$, we specify a working model and loss function pair. The working model and loss function pair is somewhat analogous to that in the updating step in the TMLE procedure described in Section 3. The loss function can be the same as that in the TMLE procedure's updating step, but it need not be. Define $\bar{Q}$ and b to be the limits of $\bar{Q}_n$ and $\bar{g}_n$, respectively, as $n \to \infty$. Let Θ be the class of all functions of A and $b(W)$, and let θ be some function in that class. Here $\bar{Q}$ is not necessarily $\bar{Q}_0$ and b is not necessarily $\bar{g}_0$ or even a balancing score. For concreteness, consider two working model and loss function pairs: a logistic working model

$$\bar{Q}^{b,\theta}(A, W) = \text{logit}^{-1}[\text{logit}(\bar{Q}(A, W)) + \theta(A, b(W))] \tag{1}$$

with loss function

$$L'(\bar{Q}^{b,\theta})(O) = -Y\log(\bar{Q}^{b,\theta}(A,W)) - (1-Y)\log(1-\bar{Q}^{b,\theta}(A,W)),$$

which is the negative log likelihood loss when Y is binary, and a linear working model

$$\bar{Q}^{b,\theta}(A,W) = \bar{Q}(A,W) + \theta(A,b(W)) \tag{2}$$

with loss function

$$L'(\bar{Q}^{b,\theta})(O) = (Y - \bar{Q}^{b,\theta}(A,W))^2,$$

the squared error loss. In both working models, we leave the function θ unspecified. We can view a working model used for the updating step in the TMLE procedure as a special case of the working model here by restricting θ to have the form

$$\theta(A,b(W)) = \varepsilon\frac{A}{b(W)}$$

where ε is real, using notation $b(W)$ in place of $g(1 \mid W)$ as used in Section 3.

Define

$$\theta_0 = \arg\min_{\theta\in\Theta} E_0 L'\left(\bar{Q}^{b,\theta}\right)(O).$$

Given $\bar{Q}$, the limit of some estimate for $\bar{Q}_0$, one can think of θ_0, a function of A and $b(W)$, as the residual bias between $E_0(\bar{Q}(A,W) \mid A, b(W))$ and $E_0(Y \mid A, b(W))$ on either the logistic or linear scale. When the initial estimator $\bar{Q}_n$ is consistent, so $\bar{Q} = \bar{Q}_0$, $\theta_0(A,b(W))$ will be 0, because $\bar{Q}$ will already be fully adjusting for A and $b(W)$.

Suppose for now that we have an estimate of θ_0 which we call θ_n. We return to the problem of estimating θ_0 later in this section. Calculate the update of $\bar{Q}_n$ as $\bar{Q}_n^{\bar{g}_n,\theta_n}$ and using this updated regression, a final estimate of ψ_0 is calculated as $\Psi(\bar{Q}_n^{\bar{g}_n,\theta_n}, Q_{Wn})$, which we call a doubly robust balancing score adjusted (DR-BSA) plug-in estimator. In Theorem 1 in Appendix A.2, we show that the DR-BSA estimator doubly robust in the sense that it is consistent when either $\bar{Q} = \bar{Q}_0$ or θ_n consistently estimates θ_0 and b is a balancing score.

When initial estimator $\bar{Q}_n$ does not consistently estimate $\bar{Q}_0$, consistency of the DR-BSA estimate requires that b is a balancing score and θ_0 is consistently estimated. To weaken this requirement, we now construct a TMLE with the balancing score property by using $\bar{Q}_n^0 = \bar{Q}_n^{\bar{g}_n,\theta_n}$ as the initial estimate in the TMLE procedure in Section 3 and updating it to $\bar{Q}_n^*$. The TMLE of $\Psi(P_0)$ is calculated as $\Psi(\bar{Q}_n^*, Q_{Wn})$. We call this a balancing score adjusted TMLE (BSA-TMLE). In Theorem 2 in Appendix A.2, we show that the BSA-TMLE is consistent if any of the three conditions hold: (1) $\bar{Q} = \bar{Q}_0$, (2) $b = \bar{g}_0$, or (3) b is a balancing score and θ_n consistently estimates θ_0. The BSA-TMLE is therefore doubly robust in the usual sense and also has the balancing score property. The BSA-TMLE is a TMLE as described in Section 3 where in addition to attempting to adjust for W, the initial estimator $\bar{Q}_n^0$ is making an extra attempt to adjust for a balancing score. If θ_0 is consistently estimated, then like the standard TMLE, when both the initial estimates of $\bar{Q}_0$ and g_0 are consistent, the influence curve of the BSA-TMLE is the efficient influence curve. Therefore, under regularity conditions, the BSA-TMLE is locally efficient and keeps all of the attractive properties of TMLE while also having the balancing score property.

We now return to the problem of estimating θ_0. The working model in the definition of θ_0 depends is $\bar{Q}^{b,\theta}$ which depends on limits $\bar{Q}$ and b. To estimate θ_n, we use $\bar{Q}_n^{\bar{g}_n,\theta}$ as the working model. If $\bar{g}_n(W)$ is discrete and θ_0 is estimated in a saturated parametric model, $\Psi(\bar{Q}_n^{\bar{g}_n,\theta_n}, Q_{Wn})$ is exactly a TMLE as proved in Lemma 2 in Appendix A.2. When $\bar{g}_n(W)$ is not discrete, it can be discretized into k categories based on quantiles. The parameter θ_0 can be estimated with a saturated parametric model with standard logistic regression software with dummy variables for each stratum and treatment combination, and logit$\bar{Q}_n(A,W)$ as an offset. When $\bar{Q}_n(A,W)$ is unadjusted for W, for example $\bar{Q}_n$ is estimated in a GLM with only an intercept and treatment as a main term, this reduces to usual propensity score stratification. In general, when the number of categories k is fixed and does not grow with sample size, stratification is not consistent, though one hopes that the residual bias is small [2]. If k is too large, there is a possibility of all observations in a particular stratum having the same value for A, in which case $\theta_n(A, W)$ is not well defined. In many applications, the number of strata is often set based on the rule of thumb $k = 5$ recommended by Rosenbaum and Rubin [3]. Though the stratification estimator of ψ_0 is not root-n consistent when k is fixed, the BSA-TMLE removes this remaining bias if g_n consistently estimates the true propensity score while preserving the balancing score property. In practice, the number of strata k can be chosen based on cross-validation in such a way that it can grow with sample size.

Alternatively, when $\bar{g}_n(W)$ is not discrete or has many levels, θ_0 can be estimated in an generalized additive model [21] with $\bar{Q}_n$ as an offset. We can parameterize this model as

$$\bar{Q}_n^{\bar{g}_n,\theta}(A, W) = \operatorname{logit}^{-1}\left[\operatorname{logit}\left(\bar{Q}_n(A, W)\right) + A\theta_1(\bar{g}_n(W)) + (1 - A)\theta_2(\bar{g}_n(W))\right] \tag{3}$$

with $\theta = (\theta_1, \theta_2)$ where θ_1 and θ_2 are unspecified. Other parametric or nonparametric methods can be used and cross-validation based SuperLearning can be used to select the best weighted combination of estimators for θ_0 [9, 17]. When the linear model (2) is used, $\theta_0(A, W) = E_0(Y - \bar{Q}(A, W) \mid A, g_n(1 \mid W))$. In this case, a nearest neighbor or kernel regression can be used where residuals from the initial estimate, $R_i = Y_i - \bar{Q}_n(A_i, W_i)$, are treated as an outcome. This is similar to the bias corrected matching estimator presented by Abadie and Imbens [22].

5 Simulations

We demonstrate properties of the proposed BSA-TMLE in various scenarios, and compare it to other estimators. The estimators compared in simulations include a plug-in estimator based on just the initial estimator of $\bar{Q}_0$ without balancing score adjustment, DR-BSA plug-in estimators without a TMLE update, non-doubly robust BSA plug-in estimators, an inverse probability of treatment weighted estimator, and a TMLE using an initial estimator for $\bar{Q}_0$ not directly adjusted for a balancing score.

The plug-in estimator not adjusted for a balancing score is calculated as $\Psi(\bar{Q}_n, Q_{Wn})$ with $\bar{Q}_n$ as defined in Section 4. We call this the simple plug-in estimator. The DR-BSA plug-in estimator uses the balancing score adjusted $\bar{Q}_n^0$ as in Section 4 and is calculated as $\Psi(\bar{Q}_n^0, Q_{Wn})$. The non-doubly robust BSA plug-in estimator adjusts for the balancing score, but uses as initial $\bar{Q}_n$ an unadjusted estimate that is not a function of W. The non-DR-BSA plug-in estimator can be thought of as only adjusting for $g_n(1 \mid W)$ and not the whole covariate vector W. The IPTW estimator is calculated as

$$n^{-1} \sum_{i=1}^{n} \frac{A_i Y_i}{g_n(1 \mid W_i)}.$$

The estimators we compare are summarized in Table 1.

Table 1 Summary of properties of compared estimators

Estimator	Plug-in	Consistent if			Efficient if
		$\bar{Q}_n \to \bar{Q}_0$	$\bar{g}_n \to \bar{g}_0$	$\bar{g}_n \to$ BS	$\bar{Q}_n \to \bar{Q}_0$ & $\bar{g}_n \to \bar{g}_0$
Simple plug-in	√	√			
BSA	√		√	√	
DR-BSA	√	√	√	√	√†
IPTW			√		
TMLE	√	√	√		√
BSA-TMLE	√	√	√	√	√

†We do now show formally that the DR-BSA estimator is asymptotically linear.

In the simulation studies, we use two methods for adjusting the initial estimator with the propensity score. All simulations were conducted in R [13]. The initial estimator $\bar{Q}_n$ was adjusted with either a generalized additive model (GAM) in eq. (3), or a nearest neighbor approach analogous to propensity score matching. The non-DR-BSA plug-in estimator based on nearest neighbors reduces exactly to a propensity score matching estimator. The GAM was fitted with the mgcv package [21] and the nearest neighbor/propensity score matching type estimator was implement with the Matching package [23].

The initial estimates for $\bar{Q}_0$ and $\bar{g}_0$ are estimated using generalized linear models. Specifically, $\bar{g}_0$ is estimated using logistic regression, and $\bar{Q}_0$ is estimated with least squares when Y is continuous, and logistic regression when Y is binary. To investigate robustness to various kinds of model misspecification, models are either correctly specified, or some relevant covariates are excluded.

The data generating distribution in the simulations was as follows. Baseline covariates W_1, W_2 and W_3 have independent uniform distributions on $[0, 1]$. Treatment A is Bernoulli with mean

$$\mathrm{logit}^{-1}(\beta_0 + \beta_1 W_1 + \beta_2 W_2 + \beta_3 W_3 + \beta_4 W_1 W_2).$$

Outcome Y is either Bernoulli or normal with variance 1 and mean

$$m(\alpha_0 + \alpha_1 W_1 + \alpha_2 W_2 + \alpha_3 W_3 + \alpha_4 A),$$

where m is logit^{-1} if Y is Bernoulli, or the identity if Y is normal. All estimators were evaluated on 1,000 datasets of size $n = 100$ and $n = 1,000$. Bias, variance, and mean squared error (MSE) are calculated for each estimator.

In the first scenario, which we call distribution one, $\alpha = (\alpha_0, \alpha_1, \alpha_2, \alpha_3, \alpha_4) = (-3, 2, 2, 0.5)$ and $\beta = (\beta_0, \beta_1, \beta_2, \beta_3, \beta_4) = (-3, 1, 1, 0, 5)$ so W_1 and W_2 are confounders, and the propensity score depends on the product $W_1 W_2$. The true parameter $\psi_0 \approx 0.0985$ and the variance bound is approximately $1.5691/n$. The variance bound of a parameter in a semiparametric model is the minimum asymptotic variance that a regular estimator can achieve, and depends on the parameter mapping Ψ and the true distribution P_0 [15]. This is analogous with the Cramér-Rao bound in a parametric model. An estimator that asymptotically achieves the variance bound is called efficient.

The first set of results in Table 2 demonstrate the balancing score property. The initial estimate $\bar{Q}_n$ is unadjusted. A correct logistic regression model is specified for $\bar{g}_0$, but predictions are transformed by the Beta cumulative distribution function with both shape parameters equal to 2. Although artificial, this means that $\bar{g}_n$ converges to a monotone transformation of $\bar{g}_0$, which is a balancing score, but does not converge to the true $\bar{g}_0$. We can see that the TMLE not adjusted for the propensity score and the IPTW estimators are not consistent as the bias is not decrease substantially when sample size increase. Conversely, methods where the initially estimate $\bar{Q}_n$ is adjusted with the propensity score, are consistent, as bias is decreasing quickly with sample size.

Table 2 Simulation results for distribution one with $\bar{Q}_n$ unadjusted and $\bar{g}_n$ correctly specified but transformed with Beta CDF

Estimator	$n = 100$			$n = 1{,}000$		
	Bias	Variance	MSE	Bias	Variance	MSE
BSA, NN	0.0276	0.0180	0.0188	0.0026	0.0018	0.0018
BSA, GAM	0.0075	0.0163	0.0163	0.0041	0.0015	0.0015
IPTW	−0.0249	0.0087	0.0093	−0.0246	0.0010	0.0016
TMLE	0.1063	0.0111	0.0224	0.1082	0.0010	0.0127
BSA-TMLE, NN	0.0276	0.0180	0.0188	0.0026	0.0018	0.0018
BSA-TMLE, GAM	0.0070	0.0164	0.0165	0.0037	0.0015	0.0015

Table 3 shows similar performance in a more realistic scenario. In this setting, the initial estimator for $\bar{Q}_n$ is unadjusted, but the logistic regression model for the propensity score is misspecified by excluding the

Table 3 Simulation results for distribution one with $\bar{Q}_n$ unadjusted, and $\bar{g}_n$ misspecified but close to a balancing score

Estimator	n = 100			n = 1,000		
	Bias	Variance	MSE	Bias	Variance	MSE
BSA, NN	0.0311	0.0166	0.0176	0.0027	0.0016	0.0016
BSA, GAM	0.0147	0.0159	0.0161	0.0033	0.0014	0.0014
IPTW	0.0390	0.0410	0.0425	0.0357	0.0025	0.0037
TMLE	0.0096	0.0172	0.0173	0.0098	0.0016	0.0017
BSA-TMLE, NN	0.0311	0.0166	0.0176	0.0027	0.0016	0.0016
BSA-TMLE, GAM	0.0101	0.0189	0.0190	−0.0042	0.0015	0.0016

interaction term W_1W_2. Here predictions are not transformed. Here $\bar{g}_n$ is close to but not exactly a balancing score, but it is close enough that the bias in estimators that nonparametrically adjust for $\bar{g}_n$ is small. The IPTW estimator, however, is still biased at large n because $\bar{g}_n$ is not converging to $\bar{g}_0$. In this case TMLE performs well even with an unadjusted initial estimator but this is not guaranteed when $\bar{g}_n$ is misspecified.

Table 4 examines the performance of estimators when the model for $\bar{g}_0$ is misspecified (only including W_1 in the logistic regression model,) but the initial estimate $\bar{Q}_n$ is a correctly specified model. Here we see that estimates that rely only on estimated propensity score (the non-doubly robust BSA estimators and IPTW,) fail to be consistent, but estimates that use the correctly specified initial estimate of $\bar{Q}_0$, are consistent. Importantly, even when the initial estimate is adjusted with the completely misspecified $\bar{g}_n$, final estimates are still consistent when the initial $\bar{Q}_n$ is correctly specified.

Table 4 Simulation results for distribution one with $\bar{Q}_n$ correctly specified and $\bar{g}_n$ misspecified

Estimator	n = 100			n = 1,000		
	Bias	Variance	MSE	Bias	Variance	MSE
Simple plug-in	0.0071	0.0120	0.0120	0.0011	0.0013	0.0013
BSA, NN	0.1190	0.0126	0.0268	0.1064	0.0014	0.0128
DR-BSA, NN	0.0064	0.0139	0.0140	0.0003	0.0015	0.0015
BSA, GAM	0.1139	0.0116	0.0246	0.1096	0.0012	0.0133
DR-BSA, GAM	0.0152	0.0129	0.0132	0.0015	0.0013	0.0013
IPTW	0.1061	0.0115	0.0228	0.1035	0.0012	0.0119
TMLE	0.0076	0.0129	0.0130	0.0009	0.0013	0.0013
BSA-TMLE, NN	0.0064	0.0139	0.0140	0.0003	0.0015	0.0015
BSA-TMLE, GAM	0.0154	0.0133	0.0136	0.0014	0.0013	0.0013

In a second scenario, called distribution two, Y is conditionally normal with $\alpha = (0, 10, 8, 0, 2)$ and $\beta = (-1, 0, 0, 3, 0)$. Here Y depends on W_1 and W_2 but A does not, so they are not confounders. Additionally, A depends on W_3, but Y does not, so W_3 is an instrumental variable. In this setting, because none of the baseline covariates are confounders, an unadjusted estimator of ψ_0 will be consistent but not efficient, because it will fail to take into account the relationship with the non-confounding baseline covariates W_1 and W_2. Here, the true ψ_0 is 2 and the variance bound is approximately $5.1979/n$.

Table 5 shows results from distribution two where the initial estimate for $\bar{Q}_0$ is the least squares estimate from a linear regression model with A, W_1, W_2, and W_3 are main terms, and the initial estimate for the propensity score is the MLE from a logistic regression model with main terms W_1, W_2, and W_3. Here we see that, although all estimators have low bias, those that only adjust for $\bar{g}_n$, (the non-doubly robust BSA estimators and IPTW,) have much higher variance than those with a correctly specified initial estimate. This demonstrates the importance in terms of efficiency of attempting to estimate $\bar{Q}_0$ well with the initial estimate even when confounding is not a concern.

Table 5 Simulation results from distribution two with $\bar{Q}_n$ correctly specified and $\bar{g}_n$ correctly specified and includes an instrumental variable

Estimator	$n = 100$			$n = 1,000$		
	Bias	Variance	MSE	Bias	Variance	MSE
Simple plug-in	−0.0112	0.0505	0.0506	0.0007	0.0048	0.0048
BSA, NN	0.0080	0.1815	0.1815	0.0020	0.0185	0.0185
DR-BSA, NN	−0.0108	0.0578	0.0579	0.0024	0.0059	0.0060
BSA, GAM	−0.0061	0.3207	0.3208	−0.0008	0.0097	0.0097
DR-BSA, GAM	−0.0112	0.0565	0.0566	0.0010	0.0051	0.0051
IPTW	−0.0072	0.7559	0.7560	−0.0021	0.0231	0.0231
TMLE	−0.0182	0.0575	0.0578	0.0009	0.0052	0.0052
BSA-TMLE, NN	−0.0108	0.0578	0.0579	0.0024	0.0059	0.0060
BSA-TMLE, GAM	−0.0181	0.0587	0.0590	0.0009	0.0053	0.0053

6 Discussion

In this paper, we discuss the balancing score property of estimators that nonparametrically adjust for the propensity score. We see in simulations that, even when the propensity score estimator is not consistent, $\Psi(P_0)$ can be estimated with low bias if the estimate of the propensity score approximates a balancing score well enough. Additionally, we introduce a balancing score adjusted TMLE which has the balancing score property and is also doubly robust and locally efficient, and provide regularity conditions for asymptotic linearity in Appendix A.2.

In order for an estimator to have the balancing score property, we need to estimate some balancing score. We acknowledge that in practice, one does not expect an estimate of the propensity score to converge exactly to a balancing score that is not g_0 in general. However, because the propensity score is a single element of the large class of balancing scores, the condition that an estimated propensity score g_n converges to some balancing score is strictly weaker than requiring g_n to converge to g_0. When g_n fails to converge to g_0, we may still have a chance at approximating a balancing score, and the proposed BSA-TMLE can still reduce bias relative to an estimator that requires that g_n converges to g_0 without sacrificing double robustness or efficiency.

We now discuss some possible generalizations to the work in this paper and areas for further research. The estimators present in this paper are for the statistical parameter $E_0[E_0(Y \mid A = 1, W)]$, which, under assumptions, can be interpreted as the population mean of a variable Y when Y is subject to missingness [24]. The results and similar estimators are immediately applicable to other interesting statistical parameters such as

$$E_0[E_0(Y \mid A = 1, W) - E_0(Y \mid A = 0, W)]$$

and

$$E_0[E_0(Y \mid A = 1, W) - E_0(Y \mid A = 0, W) \mid A = 1]$$

which, under non-testable causal assumptions, can be interpreted as causal parameters called the ATE or ATT, respectively [9, 14]. Additionally, the results are immediately generalizable to the estimation of parameters in marginal structural models [25, 26].

Propensity score-based methods are most often applied in settings where the treatment variable is binary. In settings where the treatment variable is not binary, Imai and Van Dyk [27] generalize the notion of the propensity score to the propensity function, the conditional probability of observed treatment given covariates. Imai and Van Dyk [27] show that the propensity function is a balancing score. When the propensity function can be characterized by a finite dimensional parameter, one can estimate parameters of the distribution of counterfactuals by adjusting for the dimensional characterization of the propensity function in place of all covariates. Using the approach of Imai and Van Dyk [27], the methods in this paper may be extended to develop estimators that are doubly robust and efficient with the balancing score property for more general situations where treatment is categorical or potentially even continuous.

Traditionally, propensity score-based estimators estimate the propensity score based on how well $\bar{g}_n$ approximates the true $\bar{g}_0$. Collaborative targeted minimum loss-based estimation (CTMLE) is a method that chooses an estimator for the propensity score based on how well it helps reduce bias in the estimation of $\Psi(P_0)$ in collaboration with an initial estimate of $\bar{Q}_0$ using cross-validation [9, 28]. In doing so, CTMLE attempts to adjust the propensity score for the most important confounders first and avoid adjustment for instrumental variables. This can lead to improvements in efficiency and robustness to violations of the assumption $P_0(A=a \mid W)>0$. Applying an analogous techniques of estimator selection for balancing score adjusted estimators is an area of further research.

Acknowledgements: Mark van der Laan was supported by NIH grant 5R01AI074345-05. We thank the associate editor and anonymous referees for their insightful comments, which we believe helped greatly improve the quality of the paper.

Funding: U.S. Department of Health and Human Services-National Institutes of Health 5R01AI074345-05.

Appendix

A.1 Notation

- $O=(W,A,Y)$: observed data structure
  - W: vector of covariates
  - A: treatment indicator, 0 or 1
  - Y: univariate outcome
- P: a distribution of O
- $\mathscr{M}$: statistical model, set of possible probability distributions P
- $E_p(\cdot)$: expectation under distribution P
- $Q=(\bar{Q},Q_W)$
  - $\bar{Q}(a,w)=E_P(Y \mid A=a, W=w)$
  - $Q_W(w)=P(W=w)$
- $g(a \mid w)=P(A=a \mid W=w)$
- $\bar{g}(w)=g(1 \mid W)$, also called the propensity score when.
- Ψ: statistical parameter mapping from $\mathscr{M}$ to $\mathbb{R}$.
  - In particular, $\Psi(P)=E_P[E_P(Y \mid A=1,W)]$
  - Also written as $\Psi(Q)$
- $\psi=\Psi(P)$
- Subscript 0: indicates the truth, e.g. $\psi_0=\Psi(P_0)$ is the true parameter value
- Subscript n: indicates an estimate based on n observations, e.g. $\bar{Q}_n$ is an estimate of $\bar{Q}_0$
- $\bar{Q}_n^0$ an initial estimate of $\bar{Q}_0$
- L: loss function
- L_Y: loss function for $\bar{Q}$
- L_W: loss function for Q_W
- $Q(\varepsilon)$ a working submodel through Q
- IC: an influence curve
- D^*: the efficient influence curve

- $\bar{Q}_n^*$ a TMLE updated estimate of some initial $\bar{Q}_n^0$
- $b(w)$: some function of w that is a potential balancing score
- θ: some function of a and $b(w)$
- $\bar{Q}^{b,\theta}$: a working submodel through $\bar{Q}$ for a particular b and θ
- L'a loss function for $\bar{Q}^{b,\theta}$, used in Section 4

A.2 Some results and proofs

Proof of Lemma~1. In this proof, E means expectation with respect to P. First note that $E(Y\,|\,A=1,W,b(W))=E(Y\,|\,A=1,W)$ because b is a function of only W. Next,

$$E[E(Y\,|\,A=1,W)\,|\,A=1,b(W)]=E[E(Y\,|\,A=1,W)\,|\,b(W)]$$

because the inner conditional expectation is a function of only W and $W\perp A\,|\,b(W)$ when b is a balancing score. Thus,

$$\begin{aligned}E[E(Y\,|\,A=1,b(W))]&=E\{E[E(Y\,|\,A=1,W,b(W))\,|\,A=1,b(W)]\}\\&=E\{E[E(Y\,|\,A=1,W)\,|\,A=1,b(W)]\}\\&=E\{E[E(Y\,|\,A=1,W)\,|\,b(W)]\}\\&=E[E(Y\,|\,A=1,W)]\\&=\Psi(P).\end{aligned}$$

□

Theorem 1 *Assume*

$$\Psi((\bar{Q}_n^{g_n,\theta_n},Q_{Wn}))-\Psi((\bar{Q}^{b,\theta_0},Q_{W0}))\to 0,\ \textit{as } n\to\infty.$$

In addition, assume that either $\bar{g}$ is a balancing score or $\bar{Q}=\bar{Q}_0$. Then $\Psi((\bar{Q}_n^{g_n,\theta_n},Q_{Wn}))$ is consistent for ψ_0.

Proof. By definition of θ_0, we have

$$E_0[h(A,b(W))(Y-\bar{Q}^{b,\theta_0}(A,W))]=0$$

for all functions h of A and $b(W)$. Theorem 2 Rosenbaum and Rubin [8] show that b is a balancing score if and only if there exists a function f so that $\bar{g}_0(w)=f(b(w))$ a.e., so we can select the function

$$h(A,b(W))=\frac{A}{f(b(W))}=\frac{A}{\bar{g}_0(W)}.$$

In addition, we also have that $E_0\bar{Q}^{b,\theta_0}(1,W)-\Psi((\bar{Q}^{b,\theta_0},Q_{W,0}))=0$. This proves that

$$P_0D^*(\bar{Q}^{b,\theta_0},Q_{W,0},g_0)=0,$$

where D^* is the efficient influence curve of Ψ at P, and notation

$$P\phi=\int\phi(o)dP(o)$$

for some function ϕ of O and distribution P. Since $P_0D^*(\bar{Q},Q_W,g_0)=\psi_0-\Psi(Q)$, this shows

$$\Psi((\bar{Q}^{b,\theta_0},Q_{W0}))=\Psi((\bar{Q}_0,Q_{W0}))$$

This proves that under the stated consistency condition, we indeed have that $\Psi((\bar{Q}_n^{g_n,\theta_n},Q_{Wn}))$ is consistent for ψ_0. This proves the consistency under the condition that b is a balancing score.

Consider now the case that $\bar{Q}=\bar{Q}_0$. Then $\theta_0=0$ and thus $\bar{Q}^{b,\theta_0}=\bar{Q}_0$. Thus, the limit $\Psi((\bar{Q}^{b,\theta_0},Q_{W0}))=\Psi((\bar{Q}_0,Q_{W0}))$, which proves the second claim of the theorem. □

Theorem 2 *Assume*

$$\Psi((\bar{Q}_n^{g_n,\theta_n}(\varepsilon_n), Q_{Wn})) - \Psi((\bar{Q}^{b,\theta_0}(\varepsilon_0), Q_{W0})) \to 0, \text{ as } n \to \infty,$$

where $\varepsilon_0 = \arg\min_\varepsilon P_0 L(\bar{Q}^{b,\theta_0}(\varepsilon))$.

In addition, assume that b is a balancing score, or $\bar{Q} = \bar{Q}_0$. Then $\varepsilon_0 = 0$ and $\Psi((\bar{Q}_n^{\bar{g}_n,\theta_n}(\varepsilon_n), Q_{Wn}))$ is consistent for ψ_0.

Proof. Firstly, assume b is a balancing score so by Theorem 2 Rosenbaum and Rubin [8] there exists a mapping f so that $g_0(w) = f(b(w))$ a.e.. In the proof of the previous theorem we showed that

$$E_0 \frac{A}{b(W)}(Y - \bar{Q}^{b,\theta_0}(A, W)) = E_0 \frac{A}{g_0(W)}(Y - \bar{Q}^{b,\theta_0}(A, W)) = 0.$$

The left-hand side equals $\frac{d}{d\varepsilon} P_0 L(\bar{Q}^{b,\theta_0}(\varepsilon)) \mid_{\varepsilon=0}$ and this score equation in ε is solved by ε_0. This proves that $\varepsilon_0 = 0$ under the assumption that this score equation $P_0 L(\bar{Q}^{b,\theta_0}(\varepsilon)) = 0$ has a unique solution. The latter follows from the fact that the submodel with single parameter ε has an expected loss that is strictly convex.

This now proves that the limit $\Psi((\bar{Q}^{b,\theta_0}(\varepsilon_0), Q_{W0})) = \Psi((\bar{Q}^{b,\theta_0}, Q_{W,0}))$ so that we can apply the previous theorem which shows that the latter limit equals ψ_0. This proves the consistency of the TMLE when b is a balancing score.

Consider now the case that $\bar{Q} = \bar{Q}_0$. Then $\theta_0 = 0$ and thus $\bar{Q}^{b,\theta_0} = \bar{Q}_0$. Thus, the limit $\Psi((\bar{Q}^{b,\theta_0}, Q_{W0})) = \Psi((\bar{Q}_0, Q_{W0}))$, which proves the consistency under the condition that $\bar{Q} = \bar{Q}_0$. In the latter case, it also follows that $\varepsilon_0 = 0$. □

Lemma 2 *If $\bar{g}_n$ takes only discrete values with support G, then $\Psi((\bar{Q}_n^{\bar{g}_n,\theta_n}, Q_{Wn}))$ is a TMLE if θ_0 is estimated as θ_n using MLE in a saturated parametric model*

$$\text{logit}\bar{Q}_n^{g_n,\theta}(a, w) = \text{logit}(\bar{Q}_n(A, W)) + \sum_{\substack{a \in \{0,1\} \\ c \in G}} \theta_{a,c} I(A = a, \bar{g}_n(W) = c) \tag{4}$$

where $\bar{Q}_n$ is some initial estimator for $\bar{Q}_0$ and I is the indicator function.

Proof of Lemma~2. The MLE θ_n (or empirical risk minimizer for the negative quasi-binomial log likelihood, if Y is not binary), solves the score equations for each parameter $\theta_{a,c}$:

$$0 = \sum_{i=1}^n I(A_i = a, \bar{g}_n(W_i) = c)(Y - \bar{Q}_n^{g_n,\theta_n}(A_i, W_i)).$$

Additionally, any function h of A and $\bar{g}_n(W)$ is in the linear span of basis functions $I(A = a, \bar{g}_n(W) = c)$ for all $a \in \{0, 1\}$, $c \in G$, so

$$0 = \sum_{i=1}^n h(A_i, \bar{g}_n(W_i))(Y - \bar{Q}_n^{g_n,\theta_n}(A_i, W_i)).$$

In particular, the above equation is solved when $h(a, w) = \frac{a}{\bar{g}_n(w)}$, which is the score from the parametric submodel in eq. (4). Thus if the TMLE update is applied to the initial estimate $\bar{Q}_n^0 = \bar{Q}_n^{\bar{g}_n,\theta_n}$, $\varepsilon_n = 0$, and $\bar{Q}_n^* = \bar{Q}_n^0$ so $\Psi((\bar{Q}_n^{\bar{g}_n,\theta_n}, Q_{Wn}))$ is a TMLE. □

Theorem 3 *Define $\Phi_1(Q) = P_0 \bar{Q} \frac{\bar{g} - \bar{g}_0}{\bar{g}}$ and $\Phi_2(g) = P_0(\bar{Q} - \bar{Q}_0)\frac{\bar{g}}{\bar{g}_0}$. Assume $D^*(Q_n^*, g_n)$ falls in a P_0-Donsker class with probability tending to 1; $P_0\{D^*(Q_n^*, g_n) - D^*(Q, g)\}^2 \to 0$ in probability as $n \to \infty$;*

$$P_0(\bar{Q}_0 - \bar{Q}_n^*)(\bar{g}_0 - \bar{g}_n)\frac{(\bar{g} - \bar{g}_n)}{\bar{g}\bar{g}_n} = o_P(1/\sqrt{n});$$

$$P_0(\bar{Q}_n^* - \bar{Q})(\bar{g}_n - \bar{g})/\bar{g} = o_P(1/\sqrt{n});$$

$$P_0(\bar{Q} - \bar{Q}_0)(\bar{g} - \bar{g}_0)/\bar{g} = 0;$$

$\Phi_1(\bar{Q}_n^*)$ and $\Phi_2(\bar{g}_n)$ are asymptotically linear estimators of $\Phi_1(\bar{Q})$ and $\Phi_2(\bar{g})$ with influence curves IC_1 and IC_2, respectively.

Then $\Psi(Q_n^*)$ is asymptotically linear with influence curve $D^*(Q,g) + IC_1 + IC_2$.

Proof. Since $P_0D^*(Q,g) = \psi_0 - \Psi(Q) + P_0(\bar{Q}_0 - \bar{Q})(\bar{g}_0 - \bar{g})/\bar{g}$ (e.g., Zheng and Laan [29]); Zheng and van der Laan [30]), where we use the notation $\bar{Q}(W) = \bar{Q}(1, W)$, this results in the identity:

$$\Psi(Q_n^*) - \psi_0 = (P_n - P_0)D^*(Q_n^*, g_n) + P_0(\bar{Q}_0 - \bar{Q}_n^*)(\bar{g}_0 - \bar{g}_n)/\bar{g}_n.$$

The first term equals $(P_n - P_0)D^*(Q,g) + o_P(1/\sqrt{n})$ if $D^*(Q_n^*, g_n)$ falls in a P_0-Donsker class with probability tending to 1, and $P_0\{D^*(Q_n^*, g_n) - D^*(Q,g)\}^2 \to 0$ in probability as $n \to \infty$ [31, 32]. We write

$$P_0(\bar{Q}_0 - \bar{Q}_n^*)(\bar{g}_0 - \bar{g}_n)/\bar{g}_n = P_0(\bar{Q}_0 - \bar{Q}_n^*)(\bar{g}_0 - \bar{g}_n)/\bar{g} + P_0(\bar{Q}_0 - \bar{Q}_n^*)(\bar{g}_0 - \bar{g}_n)\frac{(\bar{g} - \bar{g}_n)}{\bar{g}\bar{g}_n}.$$

Assume that the last term is $o_P(1/\sqrt{n})$. We now write

$$\begin{aligned} &P_0(\bar{Q}_0 - \bar{Q}_n^*)(\bar{g}_0 - \bar{g}_n)/\bar{g} = P_0(\bar{Q}_n^* - \bar{Q} + \bar{Q} - \bar{Q}_0)(\bar{g}_n - \bar{g} + \bar{g} - \bar{g}_0)/\bar{g} \\ &= P_0(\bar{Q}_n^* - \bar{Q})(\bar{g}_n - \bar{g})/\bar{g} + P_0(\bar{Q}_n^* - \bar{Q})(\bar{g} - \bar{g}_0)/\bar{g} \\ &+ P_0(\bar{Q} - \bar{Q}_0)(\bar{g}_n - \bar{g})/\bar{g} + P_0(\bar{Q} - \bar{Q}_0)(\bar{g} - \bar{g}_0)/\bar{g} \\ &\equiv P_0(\bar{Q}_n^* - \bar{Q})(\bar{g}_n - \bar{g})/\bar{g} + \Phi_1(\bar{Q}_n^*) - \Phi_1(\bar{Q}) \\ &+ \Phi_2(\bar{g}_n) - \Phi_2(\bar{g}) + P_0(\bar{Q} - \bar{Q}_0)(\bar{g} - \bar{g}_0)/\bar{g}, \end{aligned}$$

where $\Phi_1(Q) = P_0\bar{Q}\frac{\bar{g}-\bar{g}_0}{\bar{g}}$ and $\Phi_2(g) = P_0(\bar{Q} - \bar{Q}_0)\frac{\bar{g}}{\bar{g}_0}$. We assume that the first term is $o_P(1/\sqrt{n})$, the last term equals zero (i.e., either $g = g_0$ or $\bar{Q} = \bar{Q}_0$), and $\Phi_1(\bar{Q}_n^*)$ and $\Phi_2(\bar{g}_n)$ are asymptotically linear estimators with influence curves IC_1 and IC_2, respectively. This proves $\Psi(Q_n^*)$ is asymptotically linear with influence curve $D^*(Q,g) + IC_1 + IC_2$. □

A.3 TMLE when *Y* is not bounded by 0 and 1

If Y is not bounded by 0 and 1, but we can assume Y is bounded by l and u with $-\infty < l < u < \infty$, Y can be transformed to $Y^\dagger = \frac{Y-l}{u-l}$. Similarly $\bar{Q}_n^0$ can be transformed to $\bar{Q}_n^{0\dagger} = \frac{\bar{Q}_n^0 - l}{u-l}$. The procedure described in Section 3 can be applied to the data structure $(W, A, Y^\dagger)$ using $\bar{Q}_n^{0\dagger}$ as initial estimator, and the final estimate can be transformed back to the original scale as $\Psi((\bar{Q}_n^*, Q_{Wn})) * (u - l) + l$. When l and u are not known, they can be set to the minimum and maximum of the observed Y as described in [20].

For completeness we can define an alternative TMLE using a linear working model where

$$\bar{Q}_n^0(\varepsilon)(A, W) = \bar{Q}_n^0(A, W) + \varepsilon\frac{A}{g_n(1 \mid W)}$$

with loss function

$$L_Y(\bar{Q})(O) = (Y - \bar{Q}(A, W))^2$$

the squared error loss. Here, $\varepsilon_0 = \arg\min_\varepsilon E_0 L_Y(\bar{Q})(O)$ can be estimated by standard least squares regression software, with $\bar{Q}_n^0(A, W)$ as an offset.

Asymptotically, a TMLE using a linear working model (or linear fluctuation) is the equivalent to a TMLE with a logistic working model, but in practice can perform poorly. This is because if $g_n(1 \mid W_i)$ is very small for some observations, which is more likely in small samples, ε_n^0 can be large in absolute value, having a large effect on $\bar{Q}_n^*$ with a linear fluctuation, which is unbounded. Because of this, if it is reasonable to bound Y by some l and u, it the logistic working model is recommended because $\bar{Q}_n^*$ always respects these bounds, even if ε_n^0 is large.

A.4 Example implementation of a BSA-TMLE estimator in R

```
bsatmle <- function(QnA1, QnA0, gn1, A, Y, family = "binomial") {
   # computes estimates of E(E(Y|A=1, W)) (called ey1 in the
   # output), E(E(Y|A=0, W)) (called ey0), and
   # E(E(Y|A=1, W)) - E(E(Y|A=1, W)) (called ate)
   #
   # Inputs:
   # QnA1, QnA0: vectors, initial estimates of \bar{Q}_n(1, W)
   #             and \bar{Q}_n(O, W)
   # gn1: vector, estimates of g_n(1|W)
   # A: vector, indicator of treatment
   # Y: vector, outcome
   # family: "binomial" for logistic fluctuation, "gaussian"
   #         for linear fluctuation.
   #         if "binomial", Y should be binary or bounded
   #         by 0 and 1

   if (!require(mgcv)) stop("mgcv package is required")
   if (family=="binomial") {
      #use quasibinomial to suppress error messages about
      #non-integer Y
      family <- "quasibinomial"
      link <- qlogis
   } else {
       link <- identity
   }

   QnAA <- ifelse(A==1, QnA1, QnA0)

   # Use a generalized additive model to estimate theta_0
   # using the initial estimate of \bar{Q}
   gamfit <- gam(Y factor(A) + s(gn1, by=factor(A)) + offset(off),
       family, data=data.frame(A=A, gn1=gn1, off=link(QnAA)))

   #Get predictions from gam fit
   QnA1.gam <- predict(gamfit, type="response",
       newdata=data.frame(A=1, gn1=gn1, off=link(QnA1)))
   QnA0.gam <- predict(gamfit, type="response",
       newdata=data.frame(A=0, gn1=gn1, off=link(QnA0)))
   QnAA.gam <- ifelse(A==1, QnA1.gam, QnA0.gam)

   # compute a/g_n(1|W)
   hA1 <- 1/gn1
   hA0 <- -1/(1 - gn1)
   hAA <- ifelse(A==1, hA1, hA0)

   #using glm, fluctuate the gam-updated initial fit of \bar{Q}
   glmfit <- glm(Y -1 + h + offset(off), family,
       data=data.frame(h=hAA, off=link(QnAA.gam)))
```

```
    QnA1.star <- predict(glmfit, type="response",
        newdata=data.frame(h=hA1, off=link(QnA1.gam)))
    QnA0.star <- predict(glmfit, type="response",
        newdata=data.frame(h=hA0, off=link(QnA0.gam)))

    #compute the final estimates
    ey1 <- mean(QnA1.star)
    ey0 <- mean(QnA0.star)

    ate <- ey1-ey0

    list(ey1=ey1, ey0=ey0, ate=ate)
}
```

References

1. Austin PC. The performance of different propensity-score methods for estimating differences in proportions (risk differences or absolute risk reductions) in observational studies. Stat Med 2010;29:2137–48.
2. Lunceford J, Davidian M. Stratification and weighting via the propensity score in estimation of causal treatment effects: a comparative study. Stat Med 2004;23:2937–60.
3. Rosenbaum P, Rubin D. Reducing bias in observational studies using subclassification on the propensity score. J Am Stat Assoc 1984;79:516–24.
4. Robins J, Hernán M, Brumback B. Marginal structural models and causal inference in epidemiology. Epidemiology 2000;11: 550–60.
5. Rosenbaum P. Model-based direct adjustment. J Am Stat Assoc 1987;82:387–94.
6. Caliendo M, Kopeinig S. Some practical guidance for the implementation of propensity score matching. J Econ Surv 2008;22:31–72.
7. Dehejia R, Wahba S. Propensity score-matching methods for nonexperimental causal studies. Rev Econ Stat 2002;84: 151–61.
8. Rosenbaum P, Rubin D. The central role of the propensity score in observational studies for causal effects. Biometrika 1983;70:41.
9. van der Laan MJ, Rose S. Targeted learning: causal inference for observational and experimental data. New York: Springer, 2011.
10. van der Laan MJ, Rubin D. Targeted maximum likelihood learning. Int J Biostat 2006;2. Available at: http://www.degruyter.com/view/j/ijb.2006.2.1/ijb.2006.2.1.1043/ijb.2006.2.1.1043.xml
11. Robins JM, Rotnitzky A, Zhao LP. Estimation of regression coefficients when some regressors are not always observed. J Am Stat Assoc 1994;89:846–66.
12. van der Laan MJ, Robins JM. Unified methods for censored longitudinal data and causality. New York: Springer, 2003.
13. R Core Team. R: a language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing, 2012. Available at: http://www.R-project.org/, ISBN 3-900051-07-0
14. Hahn J. On the role of the propensity score in efficient semiparametric estimation of average treatment effects. Econometrica 1998;66:315–31.
15. Bickel PJ, Klaassen CAJ, Ritov Y, Wellner JA. Efficient and adaptive estimation for semiparametric models. Baltimore, MD: The Johns Hopkins University Press, 1993.
16. McCullagh P, Nelder J. Generalized linear models, Vol. 37. Boca Raton, FL: Chapman & Hall/CRC, 1989.
17. van der Laan MJ, Polley EC, Hubbard AE. Super learner. Stat Appl Genet Mol Biol 2007;6. Available at: http://www.degruyter.com/view/j/sagmb.2007.6.1/sagmb.2007.6.1.1309/sagmb.2007.6.1.1309.xml
18. van der Laan MJ. Targeted maximum likelihood based causal inference: part I. Int J Biostat 2010;6. Available at: http://www.degruyter.com/view/j/ijb.2010.6.2/ijb.2010.6.2.1211/ijb.2010.6.2.1211.xml
19. van der Laan MJ. Targeted maximum likelihood based causal inference: part II. Int J Biostat 2010;6. Available at: http://www.degruyter.com/dg/viewarticle/j$002fijb.2010.6.2$002fijb.2010.6.2.1241$002fijb.2010.6.2.1241.xml
20. Gruber S, Van Der Laan M. A targeted maximum likelihood estimator of a causal effect on a bounded continuous outcome. Int J Biostat 2010;6. Available at: http://www.degruyter.com/view/j/ijb.2010.6.1/ijb.2010.6.1.1260/ijb.2010.6.1.1260.xml
21. Wood SN. Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. J R Stat Soc Ser B (Stat Methodol) 2011;73:3–36.
22. Abadie A, Imbens G. Bias-corrected matching estimators for average treatment effects. J Bus Econ Stat 2011;29:1–11.

23. Sekhon JS. Multivariate and propensity score matching software with automated balance optimization: the matching package for R. J Stat Softw 2011;42:1–52. Available at: http://www.jstatsoft.org/v42/i07/
24. Kang J, Schafer J. Demystifying double robustness: a comparison of alternative strategies for estimating a population mean from incomplete data. Stat Sci 2007;22:523–39.
25. Robins JM. Marginal structural models. In Proceedings of the American Statistical Association. Section on Bayesian Statistical Science, 1–10, 1997.
26. Rosenblum M, van der Laan MJ. Targeted maximum likelihood estimation of the parameter of a marginal structural model. Int J Biostat 2010;6:Article 19.
27. Imai K, Van Dyk DA. Causal inference with general treatment regimes. J Am Stat Assoc 2004;99:854–66.
28. van der Laan MJ, Gruber S. Collaborative double robust targeted maximum likelihood estimation. Int J Biostat 2010;6. Available at: http://www.degruyter.com/view/j/ijb.2010.6.1/ijb.2010.6.1.1181/ijb.2010.6.1.1181.xml
29. Zheng W, Laan MJ Asymptotic theory for cross-validated targeted maximum likelihood estimation. Working Paper 273, U.C. Berkeley Division of Biostatistics Working Paper Series, 2010. Available at: http://www.bepress.com/ucbbiostat/paper273/
30. Zheng W, van der Laan MJ. Targeted maximum likelihood estimation of natural direct effects. Int J Biostat 2012;8. Available at: http://www.degruyter.com/view/j/ijb.2012.8.issue-1/1557-4679.1361/1557-4679.1361.xml
31. van der Vaart AW. Asymptotic statistics. Cambridge and New York: Cambridge University Press, 1998.
32. van der Vaart AW, Wellner JA. Weak convergence and empirical processes. New York: Springer, 1996.