University of California, Berkeley
U.C. Berkeley Division of Biostatistics Working Paper Series

Year 2013 *Paper* 311

Balancing Score Adjusted Targeted Minimum Loss-based Estimation

Samuel D. Lendle* Bruce Fireman† Mark J. van der Laan‡

*University of California, Berkeley, School of Public Health, Division of Biostatistics, lendle@stat.berkeley.edu

†Kaiser Permanente Division of Research, bruce.fireman@kp.org

‡University of California, Berkeley, School of Public Health, Division of Biostatistics, laan@berkeley.edu

Balancing Score Adjusted Targeted Minimum Loss-based Estimation

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

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.

1 Introduction

Estimators based on the propensity score, 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 (Rosenbaum and Rubin, 1984, Lunceford and Davidian, 2004, Austin, 2010), inverse probability of treatment weighting (IPTW) (Rosenbaum, 1987, Robins, Hernán, and Brumback, 2000), and propensity score matching (Rosenbaum and Rubin, 1983, Dehejia and Wahba, 2002, Caliendo and Kopeinig, 2008). Methods that adjust for the propensity score nonparametrically, such as propensity score matching or stratification by the propensity score, can be consistent for the parameter of interest in some cases when the estimated propensity score is not consistent. Specifically, if an estimator of the propensity score converges to a "balancing score" as defined by Rosenbaum and Rubin (1983) then the final estimate can still converge to the true parameter of interest.

We say that an estimator using the propensity score has the balancing score property if it is consistent when the estimated propensity score converges to a balancing score. Such estimators are in general not efficient. In this article, we discuss a general class of estimators that have the balancing score property. We also construct a targeted minimum loss-based estimation (TMLE) (van der Laan and Rubin, 2006, van der Laan and Rose, 2011) that is locally efficient, doubly robust and has the balancing score property.

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 estimator. 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. Some results and proofs not included in the main text are in Appendix A and two modifications to the TMLE algorithm are presented in Appendix B

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 and 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 the average outcome had everyone in the population received treatment $A = 1$. In this paper we focus on estimation of the statistical parameter $\Psi(P_0)$, but other similar statistical parameters can, under assumptions, be interpretable as causal parameters such as the ATE or the ATT (Hahn, 1998).

For a probability distribution $P \in \mathcal{M}$, let $\bar{Q}(a, w) = E_P(Y \mid A = a, W = w)$, $Q_W(w) = P(W = w)$, $Q = (\bar{Q}, Q_W)$, $g(a \mid w) = P(A = a \mid W = w)$, and $\bar{g}(w) = g(1 \mid w)$. The function $\bar{g}$ is called the propensity score. Because the mapping Ψ depends on P only through Q, recognizing the abuse of notation, we sometimes write $\Psi(P) = \Psi(Q) = \Psi((\bar{Q}, Q_W))$. The notation $Pf = \int f(o)dP(o)$. Let $O_1, \ldots, O_n$ be a data set of n independent and identically distributed random variables drawn from P_0 and $O_i = (W_i, A_i, Y_i)$. 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, $\bar{Q}_0(a, w) = E_0(E_0(Y \mid A = 1, W))$, 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 P_0, or relavent parts of P_0, and plugs it into the parameter mapping Ψ. In this case, the 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

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

A targeted minimum loss-based estimator for $\Psi(P_0)$ is a plug-in estimator that takes an estimate of Q_0, say Q_n^0, and, using an estimate $\bar{g}_n(W)$ of the propensity score, updates it to Q_n^*. The final estimate is calculated as $\Psi(Q_n^*)$.

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 (McCullagh and Nelder, 1989), or with a data adaptive machine learning algorithm such as the SuperLearner algorithm (van der Laan, Polley, and Hubbard, 2007, van der Laan and Rose, 2011), which combines parametric and nonparametric models and data adaptive estimators using cross validation. The updating step is defined by a choice of loss function L for Q such that $E_0L(Q)(O)$ 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 such that the efficient influence curve is in the 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 (van der Laan and Rubin, 2006, van der Laan, 2010a,b, van der Laan and Rose, 2011).

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 B if Y is not bounded by 0 and 1. Define the loss function $L(Q)(O) = L_Y(\bar{Q})(O) + L_W(Q_W)(O)$ where $L_W(Q_W)(O) = -\log(Q_W(W))$ and

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

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

$$\bar{Q}_n^0(\varepsilon)(A,W) = \text{logit}^{-1}\left[\text{logit}(\bar{Q}_n^0(A,W)) + \varepsilon\frac{A}{g_n(1 \mid W)}\right]$$

indexed by ε. We call this a logistic working model. The estimate ε_n^0 can be calculated using standard logistic regression software with $\text{logit}(\bar{Q}_n^0(A,W))$

as a fixed offset term, and $\frac{A}{g_n(1|W)}$ as a covariate. By using the empirical distribution of W as an initial estimate for Q^0_{Wn}, and negative log likelihood loss function for L_W, the empirical risk is minimized at Q^0_{Wn}, so no update is needed. In this case, the algorithm converges in one step, because $\frac{A}{g_n(1|W)}$ is not updated between iterations, so an additional update to $\bar{Q}^1_n$ will yield $\varepsilon^1_n = 0$. The estimate $\bar{Q}^*_n$ is calculated as $\bar{Q}^0_n(\varepsilon^0_n)$ 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 $\bar{Q}^0_n$ 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 $\bar{Q}^0_n$ and g_n are consistent, 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 in Theorem 3.

4 Balancing score property and proposed estimator

A function b of W is called a balancing score if $A \perp W \mid b(W)$ (Rosenbaum and Rubin, 1983). Trivially, $b(W) = W$ is a balancing score, and by definition of the propensity score, $\bar{g}_0(W)$, is a balancing score. Another example of a balancing score is any monotone transformation of the propensity 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 (1983) 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.

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

This result gives rise to methods for estimating $\Psi(P_0)$ based only 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 inverse probability of treatment weighting. 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 $\bar{g}_n$ converge 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 converge to a balancing score but not the true propensity score when, for example, the true $\bar{g}_0$ depends on high order interactions between covariates, but a main terms logistic regression does well at approximating a monotone transformation of the balancing score.

Estimators based only on the propensity score are not doubly robust. We wish to construct a locally efficient doubly robust estimator with the balancing score property. Start with an initial estimators $\bar{Q}'_n$ for $\bar{Q}_0$ and $\bar{g}_n$ for $\bar{g}_0$ and call their limits $\bar{Q}'$ and b, respectively, as $n \to \infty$. Define

$$\theta_0 = \arg\min_{\theta \in \Theta} E_0 L'(\bar{Q}'(b, \theta))(O) \tag{1}$$

were L' is a loss function depending on choice of working model for $\bar{Q}'_n(g_n, \theta)$. Consider two working model and loss function pairs: a logistic working model

$$\bar{Q}'_n(\bar{g}_n, \theta)(A,W) = \text{logit}^{-1}[\text{logit}(\bar{Q}'_n(A,W)) + \theta(A, \bar{g}_n(W))] \tag{2}$$

with loss function

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

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

$$\bar{Q}'_n(\bar{g}_n, \theta)(A,W) = \bar{Q}'_n(A,W) + \theta(A, \bar{g}_n(W)) \tag{3}$$

with loss function

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

the squared error loss. For both working models Θ is unrestricted. Let $\bar{Q}_n^0 = \bar{Q}'_n(\theta_n)$ where θ_n is an estimate of θ_0 discussed below. We call the estimate $\Psi((\bar{Q}_n^0, Q_{Wn}))$ a doubly robust balancing score adjusted (DR-BSA) plug-in estimator. In Theorem 1, we show that this estimator is consistent when b is a balancing score or $\bar{Q}' = \bar{Q}_0$ and is therefore doubly robust.

Theorem 1. *Assume*

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

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 $b(W)$. In the Lemma 3 in Appendix A, we prove that b is a balancing score if and only if there exists a function ϕ so that $\bar{g}_0(w) = \phi(b(w))$ a.e., so we can select the function

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

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

$$P_0 D^*(\bar{Q}'(b, \theta_0), Q_{W,0}, g_0) = 0,$$

where

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

is the efficient influence curve of Ψ at P. Since $E_0 D^*(\bar{Q}, Q_W, g_0) = \psi_0 - \Psi(Q)$, 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 $\Psi((\bar{Q}'(b, \theta_0), Q_{W0})) = \Psi((\bar{Q}_0, Q_{W,0}))$, which proves the second claim of the theorem. □

Now, use $\bar{Q}_n^0 = \bar{Q}_n'(g_n, \theta_n)$ as the initial estimator for the TMLE step described in Section 3 to obtain $\bar{Q}_n^*$. 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 we show $\Psi((\bar{Q}_n^*, Q_{Wn}))$ is consistent if $\bar{Q} = \bar{Q}_0$ or b is a balancing score and is therefore doubly robust with the balancing score property.

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 Lemma 3 that there exists a mapping ϕ so that $g_0(w) = \phi(b(w))$ 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 $\frac{d}{d\varepsilon} P_0 L(\bar{Q}'(b, \theta_0)(\varepsilon))\big|_{\varepsilon=0}$ 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 $\Psi((\bar{Q}'(b, \theta_0), Q_{W0})) = \Psi((\bar{Q}_0, Q_{W,0}))$, which proves the consistency under the condition that $\bar{Q}' = \bar{Q}_0$. In the latter case, it also follows that $\varepsilon_0 = 0$. □

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 $\bar{g}_n(W)$ is discrete and θ_0 is estimated in a saturated parametric model, $\Psi((\bar{Q}_n^0, Q_{Wn}))$ is exactly a TMLE as proved in Lemma 2 in Appendix A.

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 maint 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 (Lunceford and Davidian, 2004). 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 (1984). 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. 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, θ_0 can be estimated in an generalized additive model with $\bar{Q}_n^0$ as an offset:

$$\begin{aligned}\bar{Q}_n^0(A,W) = \text{logit}^{-1}[&\text{logit}(\bar{Q}_n'(A,W))+ \\ &A\theta_1(g_n(1 \mid W)) + (1-A)\theta_2(g_n(1 \mid W))]\end{aligned} \tag{4}$$

(Wood, 2011). 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(van der Laan and Rose, 2011, van der Laan et al., 2007). When model (3) is used, 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, estimating $\theta_0(A,W) = E_0(Y - \bar{Q}'(A,W) \mid A, g_n(1 \mid W))$. This is similar to the bias corrected matching estimator presented by Abadie and Imbens (2011).

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 (IPTW), 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

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

In the simulation studies, we use three methods for adjusting the initial estimator with the propensity score. All simulations were conducted in R (R Core Team, 2012). The initial estimator $\bar{Q}_n'$ was adjusted with either a generalized additive model (GAM) in (4), 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 (Wood, 2011) and the nearest neighbor/propensity score matching type estimator was implement with the `Matching` package (Sekhon, 2011).

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

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

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_1W_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 (Bickel, Klaassen, Ritov, and Wellner, 1993). 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 1 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 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 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 3 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

Table 1: 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=1000**		
	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 2: 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=1000**		
	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

Table 3: Simulation results for distribution one with $\bar{Q}'_n$ correctly specified and $\bar{g}_n$ misspecified

Estimator	**n=100**			**n=1000**		
	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

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.

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 4 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 4: 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=1000**		
	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.

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 (Kang and Schafer, 2007). 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 (Hahn, 1998, van der Laan and Rose, 2011). Additionally, the results are immediately generalizable to the estimation of parameters in marginal structural models (Robins, 1997, Rosenblum and van der Laan, 2010).

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 (van der Laan and Gruber, 2010, van der Laan and Rose, 2011). 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 $P_0(A = a|W) > 0$. Applying an analogous techniques of estimator selection for balancing score adjusted estimators is an area of further research.

References

A. Abadie and G.W. Imbens. Bias-corrected matching estimators for average treatment effects. *Journal of Business & Economic Statistics*, 29(1):1–11, 2011.

Peter C Austin. The performance of different propensity-score methods for estimating differences in proportions (risk differences or absolute risk reductions) in observational studies. *Statistics in medicine*, 29(20):2137–48, September 2010. ISSN 1097-0258. doi: 10.1002/sim.3854.

Peter J. Bickel, Chris A. J. Klaassen, Ya'acov Ritov, and Jon A. Wellner. *Efficient and Adaptive Estimation for Semiparametric Models*. The Johns Hopkins University Press, Baltimore, 1993. ISBN 0801845416.

M. Caliendo and S. Kopeinig. Some practical guidance for the implementation of propensity score matching. *Journal of economic surveys*, 22(1): 31–72, 2008.

R.H. Dehejia and S. Wahba. Propensity score-matching methods for nonexperimental causal studies. *Review of Economics and statistics*, 84(1): 151–161, 2002.

S. Gruber and M.J. Van Der Laan. A targeted maximum likelihood estimator of a causal effect on a bounded continuous outcome. *The International Journal of Biostatistics*, 6(1), 2010.

Jinyong Hahn. On the role of the propensity score in efficient semiparametric estimation of average treatment effects. *Econometrica*, 66(2):315–331, March 1998. ISSN 0012-9682. doi: 10.2307/2998560.

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

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

P. McCullagh and J.A. Nelder. *Generalized linear models*, volume 37. Chapman & Hall/CRC, 1989.

R Core Team. *R: A Language and Environment for Statistical Computing.* R Foundation for Statistical Computing, Vienna, Austria, 2012. URL `http://www.R-project.org/`. ISBN 3-900051-07-0.

James M. Robins. Marginal structural models. In *Proceedings of the American Statistical Association. Section on Bayesian Statistical Science*, pages 1–10, 1997.

J.M. Robins, M.Á. Hernán, and B. Brumback. Marginal structural models and causal inference in epidemiology. *Epidemiology*, 11(5):550–560, 2000.

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

P.R. Rosenbaum and D.B. Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41, April 1983.

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

Michael Rosenblum and Mark J van der Laan. Targeted maximum likelihood estimation of the parameter of a marginal structural model. *The international journal of biostatistics*, 6(2):Article 19, January 2010. ISSN 1557-4679. doi: 10.2202/1557-4679.1238.

Jasjeet S. Sekhon. Multivariate and propensity score matching software with automated balance optimization: The Matching package for R. *Journal of Statistical Software*, 42(7):1–52, 2011. URL http://www.jstatsoft.org/v42/i07/.

Mark J. van der Laan. Targeted Maximum Likelihood Based Causal Inference: Part I. *The International Journal of Biostatistics*, 6(2), January 2010a. ISSN 1557-4679. doi: 10.2202/1557-4679.1211.

Mark J van der Laan. Targeted maximum likelihood based causal inference: Part II. *The international journal of biostatistics*, 6(2), January 2010b. ISSN 1557-4679. doi: 10.2202/1557-4679.1241.

Mark J van der Laan and Susan Gruber. Collaborative double robust targeted maximum likelihood estimation. *The international journal of biostatistics*, 6(1), January 2010. ISSN 1557-4679. doi: 10.2202/1557-4679.1181.

Mark J. van der Laan and Sherri Rose. *Targeted Learning: Causal Inference for Observational and Experimental Data*. Springer, New York, 2011. ISBN 1441997814.

Mark J. van der Laan and Daniel Rubin. Targeted Maximum Likelihood Learning. *The International Journal of Biostatistics*, 2(1), January 2006. ISSN 1557-4679. doi: 10.2202/1557-4679.1043.

Mark J van der Laan, Eric C Polley, and Alan E Hubbard. Super learner. *Statistical applications in genetics and molecular biology*, 6(1), January 2007. ISSN 1544-6115. doi: 10.2202/1544-6115.1309.

A. W. van der Vaart. *Asymptotic statistics*. Cambridge University Press, Cambridge, UK New York, NY, USA, 1998.

A. W. van der Vaart and Wellner J. A. *Weak convergence and empirical processes*. Springer, New York, 1996.

Simon N. Wood. Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(1): 3–36, 2011. ISSN 1467-9868. doi: 10.1111/j.1467-9868.2010.00749.x.

Wenjing Zheng and Mark J Van Der Laan. Asymptotic Theory for Cross-validated Targeted Maximum Likelihood Estimation. Working Paper 273, U.C. Berkeley Division of Biostatistics Working Paper Series, 2010. URL `http://www.bepress.com/ucbbiostat/paper273/`.

Wenjing Zheng and Mark J van der Laan. Targeted maximum likelihood estimation of natural direct effects. *The international journal of biostatistics*, 8(1), January 2012.

A Some results and proofs

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

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

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

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

□

Lemma 2. *If $\bar{g}_n$ takes only discrete values with support G, then $\Psi((\bar{Q}_n^0, Q_{Wn}))$ is a TMLE if θ_0 is estimated in a saturated parametric model*

$$logit\bar{Q}_n'(\theta)(a,w) = logit(\bar{Q}_n'(A,W)) + \sum_{\substack{a\in\{0,1\} \\ c\in G}} \theta_{a,c} I(A=a, g_n(1 \mid W)=c) \quad (5)$$

where $\bar{Q}'_n$ is some initial estimator for $\bar{Q}_0$, $\bar{Q}^0_n = \bar{Q}'_n(\theta_n)$ 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 $\theta_{a,c}$:

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

where $\bar{Q}^0_n(a,w) = \bar{Q}'_n(\theta_n)(a,w)$. Additionally, any function h of A and $g_n(1 \mid W)$ is in the linear span of basis functions $I(A = a, g_n(1 \mid W) = c)$ for all $a \in \{0,1\}$, $c \in G$, so

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

In particular, the above equation is solved when $h(a,w) = \frac{a}{g_n(1|w)}$, which is the score from the parametric submodel in (5). Thus if the TMLE update is applied to the estimate $\bar{Q}^0_n$, $\varepsilon_n = 0$, and $\bar{Q}^*_n = \bar{Q}^0_n$ so $\Psi((\bar{Q}^0_n, Q_{Wn}))$ is a TMLE. □

Lemma 3. *The function b is a balancing score if and only if there exists some function ϕ such that $\phi(b(w)) = g_0(w)$ a.e..*

Proof. Suppose b is a balancing score. By definition of the propensity score and the property of the balancing score, we know that

$$g_0(W) = E_0(A \mid W) = E_0(A \mid W, b(W)) = E_0(A \mid b(W)).$$

Thus $\bar{g}_0(W) = \phi(b(W))$, where $\phi(x) = E_0(A \mid b(W) = x)$, which proves that if b is a balancing score, then there exists some function ϕ such that $\phi(b(w)) = g_0(w)$ a.e..

Suppose now that $\bar{g}_0(w) = \phi(b(w))$ a.e. for some ϕ. We have $E_0(A \mid b(W), W) = g_0(W)$, but since $\bar{g}_0(W) = \phi(b(W))$, it follows that $E_0(A \mid b(W), W) = \phi(b(W))$ and thus that $E_0(A \mid b(W), W) = E_0(A \mid b(W))$ so b is a balancing score.

□

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* $n \to \infty$*;*

$$\begin{aligned} P_0(\bar{Q}_0 - \bar{Q}_n^*)(\bar{g}_0 - \bar{g}_n)\frac{(\bar{g}-\bar{g}_n)}{\bar{g}\bar{g}_n} &= o_P(1/\sqrt{n}); \\ P_0(\bar{Q}_n^* - \bar{Q})(\bar{g}_n - \bar{g})/\bar{g} &= o_P(1/\sqrt{n}); \\ P_0(\bar{Q}-\bar{Q}_0)(\bar{g}-\bar{g}_0)/\bar{g} &= 0; \end{aligned}$$

$\Phi_1(\bar{Q}_n^*)$ *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 (2010), Zheng and van der Laan (2012)), where we use the notation $\bar{g}(W) = g(1 \mid W)$ and $\bar{Q}(W) = \bar{Q}(1,W)$, 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 $n \to \infty$ (van der Vaart and A., 1996, van der Vaart, 1998). 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 $D^*(Q,g) + IC_1 + IC_2$. □

B 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 (Gruber and Van Der Laan, 2010).

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 (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.