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

*Year* 2011 *Paper* 280

# Targeted Minimum Loss Based Estimator that Outperforms a given Estimator

Susan Gruber* Mark J. van der Laan†

*University of California, Berkeley, sgruber65@yahoo.com
†University of California - Berkeley, laan@berkeley.edu

# Targeted Minimum Loss Based Estimator that Outperforms a given Estimator

Susan Gruber and Mark J. van der Laan

**Abstract**

Targeted minimum loss based estimation (TMLE) provides a template for the construction of semiparametric locally efficient double robust substitution estimators of the target parameter of the data generating distribution in a semiparametric censored data or causal inference model (van der Laan and Rubin (2006),van der Laan (2008), van der Laan and Rose (2011)). In this article we demonstrate how to construct a TMLE that also satisfies the property that it is at least as efficient as a user supplied asymptotically linear estimator. For the sake of illustration we focus on estimation of the additive average causal effect of a point treatment on an outcome, adjusting for baseline covariates.

# 1 Introduction

Targeted minimum loss based estimation (TMLE) provides a template for the construction of semiparametric locally efficient double robust substitution estimators of the target parameter of the data generating distribution in a semiparametric censored data or causal inference model (van der Laan and Rubin (2006); van der Laan (2008); van der Laan and Rose (2011)). It is assumed that the data set is a realization of $n$ independent and identically distributed random variables, the probability distribution of this random variable is known to be an element of a semiparametric statistical model, and the target parameter (mapping) is defined as a particular function of the possible probability distributions in this semiparametric model. A targeted minimum loss based estimator (TMLE) of the target parameter is defined by an initial estimator of a relevant part of the data generating distribution, a parametric submodel through an initial estimator, a loss function for this relevant part, minimizing the empirical risk of the loss function along the parametric submodel to iteratively update the initial estimator until convergence. This final estimator is the TMLE of the relevant part of the data generating distribution, and the evaluation of its target parameter value is the TMLE of the target parameter. By enforcing that the loss-based score of the submodel (at zero fluctuation of the initial estimator) spans the efficient influence curve of the target parameter (at the initial estimator), it follows that the TMLE of the relevant part of the data generating distribution solves the efficient score estimating equation, making the TMLE locally efficient and double robust, under regularity conditions. By choosing a parametric submodel with extra fluctuation parameters, the TMLE can be arranged to solve additional estimating equations, and thereby satisfy additional properties of interest (e.g., be an imputation estimator, see Gruber and van der Laan (2010a)). One particular example of such an iterative TMLE was presented in the original TMLE article, van der Laan and Rubin (2006), which involved also fluctuating the treatment/censoring mechanism, resulting in a TMLE that also equals an IPTW/IPCW estimator and is guaranteed to outperform the IPTW/IPCW estimator defined by the initial estimator of the treatment/censoring mechanism.

Another property of interest of an estimator is that it is guaranteed to be more efficient than a user supplied class of estimators in the case that the censoring/treatment mechanism is correctly specified. This has been achieved with empirical efficiency maximization (Rubin and van der Laan (2008), Tan (2008),Cao et al. (2009),van der Laan and Rose (2011)). However, in general this technique as presented in Rubin and van der Laan (2008) may come at a cost of losing double robustness (e.g, see Robins and Rotnitzky (1992), van der Laan and Robins (2003). Tan (2008) demonstrates how in the context of estimating equation methodology the double robustness can be preserved. Recently, Rotnitzky et al. (2011) shows how to combine empirical efficiency maximization with double robust locally efficient substitution estimators, by fluctuating the treatment mechanism with a carefully chosen clever covariate derived from the empirical efficiency maximization procedure. Borrowing this fundamental idea, in this article we demonstrate that this enhanced efficiency property can be achieved with the above mentioned TMLE of van der Laan and Rubin (2006) (jointly updating treatment mechanism and outcome regression), by fluctuating the treatment mechanism with this additional clever covariate as suggested by Rotnitzky et al. (2011). For the sake of illustration we focus on estimation of the additive average causal effect of a point treatment on an outcome, adjusting for baseline covariates.

### 1.1 Organization

In Section 2 we present the statistical estimation problem. In Section 3 we present the TMLE, and the enhanced empirically efficient TMLE, and explain its properties. From the presentation in Section 3, for experts familiar with the theory of augmented IPCW estimating equations (Robins and Rotnitzky (1992), van der Laan and Robins (2003)) it will also be clear how this TMLE is generalized to all CAR-censored data and causal inference models. In Section 4 we review the method for empirical efficiency maximization of Rubin and van der Laan (2008), and an adaptive version of it as presented in van der Laan and Rose (2011), used as an ingredient in the enhanced empirically efficient TMLE. In Section 5 we present simulations confirming the enhanced efficiency property of the TMLE presented in Section 3, and comparing it with the (non-double robust) empirical efficiency maximization estimator in Rubin and van der Laan (2008), and a regular TMLE. We end with some concluding remarks. We also provide an appendix with the R-code of the TMLEs implemented in the simulation study.

## 2 The statistical model, target parameter, and estimation problem

Let $O = (W, A, Y) \sim P_0$ be a random variable, where $W$ represents a vector of baseline covariates, $A$ a binary treatment, and $Y$ a continuous or binary outcome with values in $[0, 1]$. Let $g_0(A \mid W)$ be the conditional probability distribution of $A$, given $W$. Consider a statistical model $\mathcal{M}$ that makes no assumptions about the marginal distribution of $W$, and the conditional distribution of $Y$, given $A, W$, but might make assumptions about $g_0$. In particular, it is assumed that $0 < g_0(1 \mid W) < 1$ so that the following target parameter is well defined. The statistical target parameter $\Psi : \mathcal{M} \to \mathbb{R}$ of interest is defined as

$$\Psi(P) = E_P(E_P(Y \mid A = 1, W) - E_P(Y \mid A = 0, W)).$$

If one assumes an underlying nonparametric structural equation model $W = f_W(U_W)$, $A = f_A(W, U_A)$, $Y = f_Y(W, A, U_Y)$ (Pearl (2000)), and the randomization assumption $U_A$ is independent of $U_Y$, given $W$, then $\Psi(P_0)$ identifies the additive causal effect $E_0(Y(1)-Y(0))$, where $Y(a)=f_Y(W,a,U_Y)$ is the treatment-specific counterfactual. For the sake of estimation, we are only concerned with the statistical target parameter.

Let $Q_W(P)$ be the marginal distribution of $W$ under $P$, $\bar{Q}(P)(A,W)=E_P(Y \mid A,W)$, and we will denote corresponding parameter values with $Q_W$ and $\bar{Q}$, respectively. Let $Q(P)=(Q_W(P),\bar{Q}(P))$. Note that $\Psi(P)$ only depends on $P$ through $Q_W(P)$ and $\bar{Q}(P)$. Therefore, we will also use the notation

$$\Psi(Q)=E_{Q_W}\{\bar{Q}(1,W)-\bar{Q}(0,W)\}.$$

Our goal is to estimate $\psi_0=\Psi(Q_0)$ based on observing $n$ i.i.d. copies $O_1,\ldots,O_n$ of $O \sim P_0 \in \mathcal{M}$.

The TMLE requires knowing the canonical gradient/efficient influence curve of the pathwise derivative of $\Psi:\mathcal{M}\rightarrow \mathbb{R}$. The efficient influence curve of $\Psi:\mathcal{M}\rightarrow \mathbb{R}$ at $P$ is given by

$$\begin{aligned} D^*(P)(O) &= \frac{2A-1}{g(A \mid W)}(Y-\bar{Q}(A,W))+\left\{\bar{Q}(1,W)-\bar{Q}(0,W)-\Psi(Q)\right\} \\ &\equiv D^*_Y(P)(O)+D^*_W(P)(W), \end{aligned}$$

where the latter decomposition in a score $D^*_Y(P)$ of the conditional distribution of $Y$, given $A,W$, and score $D^*_W(P)$ of the marginal distribution of $W$ will be utilized in TMLE. In order to establish the enhanced efficiency property of the proposed TMLE we will also utilize the augmented IPCW-representation of the efficient influence curve (Robins and Rotnitzky (1992), van der Laan and Robins (2003)):

$$\begin{aligned} D^*(P)(O) &= \frac{2A-1}{g(A \mid W)}Y-\Psi(Q)-\left\{\frac{\bar{Q}(1,W)}{g(1 \mid W)}+\frac{\bar{Q}(0,W)}{g(0 \mid W)}\right\}(A-g(1 \mid W)) \\ &\equiv D_{IPTW}(Q,g)(O)+D_{CAR}(\bar{Q},g)(O), \end{aligned}$$

where $D_{CAR}(\bar{Q},g)=-H_{CAR}(Q,g)(W)(A-g(1 \mid W))$ with

$$H_{CAR}(Q,g)(W)\equiv\left\{\frac{\bar{Q}(1,W)}{g(1 \mid W)}+\frac{\bar{Q}(0,W)}{g(0 \mid W)}\right\}.$$

In cases where we want to stress the representation of the efficient influence curve $D^*(P)$ as an estimating function in $\psi$ indexed by nuisance parameters $\bar{Q}_0,g_0$, we will also use the notation $D_{IPTW}(\psi_0,g_0)$ for $D_{IPTW}(Q_0,g_0)$, and $D^*(\psi_0,\bar{Q}_0,g_0)$ for $D^*(P_0)$.

Another ingredient of the TMLE presented in the next section is an influence curve $D(P)$ of a competing regular asymptotically linear estimator of $\Psi$ at $P$ in the model $\mathcal{M}$. The TMLE $\psi^*_n$ will be constructed so that it is at least as efficient as this competing estimator at $P_0$ in the case that we estimate $g_0$ consistently. By the representation theorem for the class of gradients in CAR-censored data models (van der Laan and Robins (2003), p. 65), it follows that

$$D(P) = D_{IPTW}(Q, g) + D_{CAR}(\bar{Q}^e, g)$$

for a particular function $\bar{Q}^e = \bar{Q}^e(P)$. Let $\bar{Q}^e_0$ denote the true value of this parameter $P \to \bar{Q}^e(P)$.

The TMLE presented in the next section will use an estimator $\bar{Q}^e_n$ of $\bar{Q}^e_0$ in order to define a clever covariate $H_{CAR}(\bar{Q}^e_n, g^k_n)$ in the definition of the TMLE. As a consequence of this choice of clever covariate, the TMLE $Q^*_n, g^*_n$ will solve

$$0 = P_n D_{IPTW}(\psi^*_n, g^*_n) + D_{CAR}(\bar{Q}^e_n, g^*_n),$$

and thereby have an influence curve at least as efficient as $D(P_0)$, if $g_0$ is estimated consistently.

A particular choice of interest for $\bar{Q}^e_0$ is defined by empirical efficiency maximization over a user-supplied working model as in Rubin and van der Laan (2008). That is, let $\{\bar{Q}_\beta : \beta\}$ be a parametric working model, and define

$$\bar{Q}^e(P_0) = \arg\min_{\bar{Q}_\beta} P_0 \{D_{IPTW}(\psi_0, g_0) + D_{CAR}(\bar{Q}_\beta, g_0)\}^2. \tag{1}$$

Here we used the notation $Pf \equiv \int f(o)dP(o)$. With this choice, $D(P_0)$ represents the influence curve with minimal variance among the class of influence curves $\{D_{IPTW}(Q_0, g_0) + D_{CAR}(\bar{Q}_\beta, g_0) : \beta\}$ indexed by $\beta$.

## 3 The TMLE that is at least as efficient as competing estimator

The TMLE of $\Psi(Q_0)$ as presented in van der Laan and Rubin (2006) is defined by 1) a loss function $L(Q, g) = L(Q) + L(g)$ for $(Q_0, g_0)$ so that $Q_0 = \arg\min_Q P_0 L(Q)$, $g_0 = \arg\min_g P_0 L(g)$, 2) a submodel $\{Q(\epsilon_1) : \epsilon_1\}$ through $Q$ at $\epsilon_1 = 0$, a submodel $\{g(\epsilon_2) : \epsilon_2\}$ through $g$ at $\epsilon_2 = 0$, and 3) an initial estimator $Q^0_n$, $g^0_n$. The TMLE is defined by iterative minimization of the empirical risk, and updating:

$$\begin{aligned} \epsilon_{1n} &= \arg\min_{\epsilon_1} P_n L(Q^0_n(\epsilon_1)) \\ \epsilon_{2n} &= \arg\min_{\epsilon_2} P_n L(g^0_n(\epsilon_2)), \end{aligned}$$

$Q^1_n = Q^0_n(\epsilon_{1n})$, $g^1_n = g^0_n(\epsilon_{2n})$, and this updating process is iterated till $\epsilon_n = (\epsilon_{1n}, \epsilon_{2n}) \approx 0$. The resulting $Q^*_n, g^*_n$ solve the loss-based score equation:

$$P_n \frac{d}{d\epsilon} L(Q^*_n(\epsilon), g^*_n(\epsilon))\Big|_{\epsilon=0} = 0. \tag{2}$$

By defining the loss-function $L$ and submodel through $(Q, g)$, one can control the estimating equation (2) the TMLE solves. In particular, one wants the loss-based scores to span the efficient influence curve $D^*(Q_n^*, g_n^*)$ so that the resulting $\Psi(Q_n^*)$ will be double robust and locally efficient. Below we present a submodel $\{g(\epsilon_2) : \epsilon_2\}$ so that the additional desired enhanced efficiency property is achieved as well.

## 3.1 Initial estimators

Let $Q_{W,n}^0$, $\bar{Q}_n^0$, $g_n^0$, and $\bar{Q}_n^e$ be initial estimators of $Q_{W,0}$, $\bar{Q}_0$, $g_0$, and $\bar{Q}_0^e$, respectively. Let $Q_{W,n}^0 = Q_{W,n}$ be the empirical probability distribution of $W_1, \ldots, W_n$. The estimator of $\bar{Q}_0$ can be based on the least squares or (quasi-)log-likelihood loss function

$$L(\bar{Q})(O) = -\left\{Y \log \bar{Q}(A, W) + (1 - Y) \log\{1 - \bar{Q}(A, W)\}\right\}. \tag{3}$$

This is the log-likelihood loss-function for $\bar{Q}_0$ if $Y$ is binary. We refer to Gruber and van der Laan (2010b) in which this loss function is proposed for TMLE with a continuous bounded outcome $Y \in [0, 1]$. By a simple linear transformation, this also provides a loss function for $Y \in [a, b]$ with bounded $a, b$. In particular, $\bar{Q}_0$ could be estimated with a loss-based super learner using this loss function for the cross-validation selector (van der Laan et al. (2007)).

The estimator of $g_0$ can be based on the log-likelihood loss function $L(g) = -\log g$. The estimation method for $\bar{Q}_0^e$ might depend on the type of parameter it represents. If $\bar{Q}_0^e$ is defined by (1), then one could estimate it as

$$\bar{Q}_n^e = \arg\min_{\bar{Q}_\beta} P_n\{D_{IPTW}(\psi_n^0, g_n^0) + D_{CAR}(\bar{Q}_\beta, g_n^0)\}^2, \tag{4}$$

where $P_n$ denotes the empirical probability distribution of $O_1, \ldots, O_n$, and $\psi_n^0$ represents an estimator of $\psi_0$ that is consistent if $g_n^0$ is consistent. For example, $\psi_n^0$ could be any TMLE that takes $\bar{Q}_n^0$ and $g_n^0$ as initial estimator. We assume that $\bar{Q}_n^e$ is consistent for $\bar{Q}_0^e$ if $g_n^0$ is consistent.

## 3.2 Loss function

We select the log-likelihood loss functions $L(g) = -\log g$, $L(Q_W) = -\log Q_W$ for $g_0$ and $Q_{W,0}$, respectively, and we select $L(\bar{Q})$ (3) as loss function for $\bar{Q}_0$. Let $L(Q, g) \equiv L(\bar{Q}) + L(Q_W) + L(g)$ be the loss function for $(Q_0, g_0)$.

## 3.3 TMLE that is at least as efficient as competing estimator

Let $\bar{g}(W) \equiv g(1 \mid W)$. For a given $\bar{Q}_n^k$, $g_n^k$, define the submodels

$$\begin{aligned}
\text{Logit}\bar{Q}_n^k(\epsilon_1) &= \text{Logit}\bar{Q}_n^k + \epsilon_1 H^*(g_n^k) \\
\text{Logit}\bar{g}_n^k(\epsilon_2) &= \text{Logit}\bar{g}_n + \epsilon_{21} H_{CAR}(\bar{Q}_n^k, g_n^k) + \epsilon_{22} H_{CAR}(\bar{Q}_n^e, g_n^k).
\end{aligned}$$

We also define a submodel $Q_{W,n}(\epsilon_{10}) = (1+\epsilon_{10}D^*_W(Q^k_n))Q_{W,n}$ through the empirical probability distribution $Q_{W,n}$. Let $\epsilon = (\epsilon_{10}, \epsilon_1, \epsilon_2)$. This defines now a submodel $(Q^k_n(\epsilon), g^k_n(\epsilon))$ through $(Q^k_n = (Q^k_{W,n}, \bar{Q}^k_n), g^k_n)$ at $\epsilon = 0$. The scores $\frac{d}{d\epsilon}L(Q^k_n(\epsilon), g^k_n(\epsilon))$ of $(\epsilon_{01}, \epsilon_1)$ at $\epsilon = 0$ spans the efficient influence curve $D^*(Q^k_n, g^k_n)$. The score of $\epsilon_2$ at $\epsilon = 0$ spans any linear combination of $D_{CAR}(\bar{Q}^k_n, g^k_n)$ and $D_{CAR}(\bar{Q}^e_n, g^k_n)$.

Given a current estimator $(Q^k_n, g^k_n)$, we estimate $\epsilon$ with the MLE $\epsilon^k_n$ based on loss function $L(Q, g)$:

$$\begin{aligned}
\epsilon^k_{10n} &= \arg\min_{\epsilon_{10}} -P_n \log Q^k_{W,n}(\epsilon_{10}) \\
\epsilon^k_{1n} &= \arg\min_{\epsilon_1} P_n L(\bar{Q}^k_n(\epsilon_1)) \\
\epsilon^k_{2n} &= \arg\min_{\epsilon_2} -P_n \log g^k_n(\epsilon_2).
\end{aligned}$$

Note that $\epsilon^k_{1n}$ and $\epsilon^k_{2n}$ can be fitted with standard univariate logistic regression using the offset command. We start with $k = 0$. This defines now the first step TMLE update $(Q^1_n = Q^0_n(\epsilon^0_n), g^1_n = g^0_n(\epsilon^0_n))$. We can iterate this updating algorithm till convergence so that $\epsilon_n \approx 0$. Let $(Q^*_n, g^*_n)$ be the final TMLE at convergence. Since $Q^0_{W,n}$ is the empirical probability distribution, we have $\epsilon^k_{0n} = 0$ for all $k$, so that this empirical probability distribution is not updated by the TMLE algorithm, i.e., $Q^*_n = (Q_{W,n}, \bar{Q}^*_n)$. The TMLE of $\psi_0$ is the substitution estimator $\psi^*_n = \Psi(Q^*_n)$.

This particular iterative TMLE algorithm involving updating both $g^0_n$ and $Q^0_n$ was presented and implemented in van der Laan and Rubin (2006)), but without the extra clever covariate $H(\bar{Q}^e_n, g^k_n)$. The important choice of extra clever covariate $H(\bar{Q}^e_n, g^k_n)$ in a model for $g_0$ in order to establish the enhanced efficiency property without losing double robustness was presented in Rotnitzky et al. (2011).

### 3.4 Estimating equations solved by TMLE, and resulting alternative representations of the TMLE

We assume that the algorithm converges. In that case, $(Q^*_n, g^*_n)$ solves the score equations for the sub-model $\{Q^*_n(\epsilon), g^*_n(\epsilon) : \epsilon\}$ at $\epsilon = (\epsilon_{10}, \epsilon_1, \epsilon_2) = 0$. As a consequence, the TMLE solves the following equations:

$$\begin{aligned}
P_n D^*(\psi^*_n, \bar{Q}^*_n, g^*_n) &= 0 \\
P_n D_{IPTW}(\psi^*_n, g^*_n) &= 0 \\
P_n D_{CAR}(\bar{Q}^e_n, g^*_n) &= 0 \\
P_n D^*(\psi^*_n, \bar{Q}^e_n, g^*_n) &= 0.
\end{aligned}$$

This allows a variety of representations of the TMLE. It is a plug in estimator

$$\psi^*_n = \Psi(Q^*_n);$$

it is an IPTW estimator

$$\psi_n^* = \frac{1}{n}\sum_{i=1}^{n}\frac{2A_i - 1}{g_n^*(A_i \mid W_i)}Y_i;$$

it is an augmented IPCW-estimating equation based estimator

$$\psi_n^* = \frac{1}{n}\sum_{i=1}^{n}\frac{2A_i - 1}{g_n^*(A_i \mid W_i)}Y_i - H_{CAR}(Q_n^*, g_n^*)(W_i)(A_i - \bar{g}_n^*(W_i)),$$

corresponding with the implicit estimator $\bar{Q}_n^*, g_n^*$ of the nuisance parameters $(\bar{Q}_0, g_0)$ of the estimating function $D^*(\psi, \bar{Q}_0, g_0)$ in $\psi$; and, finally, it is also an augmented IPCW-estimating equation based estimator

$$\psi_n^* = \frac{1}{n}\sum_{i=1}^{n}\frac{2A_i - 1}{g_n^*(A_i \mid W_i)}Y_i - H_{CAR}(Q_n^e, g_n^*)(W_i)(A_i - \bar{g}_n^*(W_i)),$$

corresponding with estimating $Q_0$ with $Q_n^e$. In the case that $Q_n^e$ is defined by empirical efficiency maximization (1), then the latter estimator is the estimator of Rubin and van der Laan (2008), obtained by maximizing empirical efficiency of the class of estimating functions $D(\psi, \bar{Q}_\beta, g_0)$ (or equivalently, $D(\psi, f_\beta, g_0)$, as reviewed in next section) over the working model $\{\bar{Q}_\beta : \beta\}$ at the (implicit) estimator $g_n^*$ of $g_0$, and defining the estimator of $\psi_0$ as the solution of the corresponding estimating equation.

### 3.5 Properties of TMLE

The TMLE presented above satisfies both the definition of the TMLE as well as the definition of the empirical efficient maximization (estimating equation based) estimator of Rubin and van der Laan (2008), using the implicit estimator $g_n^*$ for $g_0$. As a consequence, it inherits the properties of both the TMLE, as a locally efficient double robust substitution estimator, as well as the empirically efficient maximization estimator of Rubin and van der Laan (2008), as an estimator that is maximally efficient among a user supplied class of asymptotically linear estimators in the case that $g_0$ is estimated consistently. For the sake of being self-contained we present here the rationale resulting in these properties. Formal proofs of these properties would require regularity conditions, and is beyond the scope of this article. Therefore, below we present the general statements, and refer to the general theorems that would have to be applied to formally establish the claimed asymptotic properties. For a completely worked out proof of a TMLE for the additive treatment effect, we refer to Zheng and van der Laan (2010) and van der Laan and Rose (2011).

By the fact that it is a TMLE that solves the efficient influence curve estimating equation $P_n D^*(\psi_n^*, \bar{Q}_n^*, g_n^*) = 0$ it follows that $\psi_n^*$ will be consistent if either $\bar{Q}_n^*$ or $g_n^*$ is consistent. In addition, under regularity conditions (e.g., van der Laan and Robins (2003); van der Laan and Rubin (2006)), $\psi_n^*$ will be an asymptotically linear estimator if either $\bar{Q}_n^*$ or $g_n^*$ is consistent, and it will be efficient if both are consistent. This corresponds with stating that $\psi_n^*$ is a double robust locally efficient estimator.

Before we proceed with explaining the enhanced efficiency property we need to provide the following background on estimating equation based estimators in CAR censored data models (Robins and Rotnitzky (1992), van der Laan and Robins (2003)). Suppose that $\psi_n$ is an estimator that solves the estimating equation $0 = P_n D^*(\psi, \bar{Q}, g_0)$ for some $\bar{Q}$. Then it follows that $\psi_n$ is asymptotically linear with influence curve $D^*(\psi_0, \bar{Q}, g_0)$. In addition, if $\bar{Q}_n$ converges to $\bar{Q}$, then under weak regularity conditions, we have that the solution $\psi_n$ of $P_n D^*(\psi, \bar{Q}_n, g_0)$ is also asymptotically linear with influence curve $D^*(\psi_0, \bar{Q}, g_0)$. By Theorem 2.3 in van der Laan and Robins (2003), if the estimator $g_n^*$ of $g_0$ is such that a particular specified smooth function $\Phi(g_n^*)$ is an efficient estimator of $\Phi(g_0)$ so that its influence curve is an element of the tangent space $T_{CAR}(P_0) = \{S(A \mid W) : E_{g_0}(S \mid W) = 0\}$ of $g$ at $P_0$ under CAR, then, under regularity conditions, the solution $\psi_n$ of $P_n D^*(\psi, \bar{Q}_n, g_n^*) = 0$ is asymptotically linear with an influence curve that has a variance smaller than or equal to the variance of $D^*(\psi, \bar{Q}, g_0)$. That is, consistent (and efficient) estimation of the orthogonal nuisance parameter $g_0$ only improves the efficiency of the estimating equation based estimator of $\psi_0$. Since $g_n^*$ is a pure MLE-based estimator, under regularity conditions, and under the assumption that $g_n^0$ is consistent for $g_0$, one can show that $\Phi(g_n^*)$ is an asymptotically linear estimator of $\Phi(g_0)$ with influence curve in $T_{CAR}(P_0)$.

Given that we know that $P_n D^*(\psi_n^*, \bar{Q}_n^e, g_n^*) = 0$, if $g_n^*$ is consistent for $g_0$ and $\bar{Q}_n^e$ converges to $\bar{Q}_0^e$, it follows that $\psi_n^*$ will be asymptotically linear with an influence curve with variance smaller than or equal to the variance of $D^*(\psi_0, \bar{Q}_0^e, g_0)$. That is, in the case that $g_n^*$ is consistent, the TMLE $\psi_n^*$ is at least as efficient as the competing estimator whose influence curve equals $D^*(\psi_0, \bar{Q}_0^e, g_0)$.

# 4 Empirical efficiency maximization

This section concerns the estimation of $Q_0^e$ that forms an ingredient of the TMLE presented above. We first review empirical efficiency maximization as presented in Rubin and van der Laan (2008), and then we demonstrate how empirical efficiency maximization can be embedded in loss-based learning of $Q_0$ by using as loss function the square of the efficient influence curve (van der Laan and Rose (2011)).

## 4.1 Empirical Efficiency Maximization as in Rubin, van der Laan (2008)

In order to determine a solution that optimizes the variance of the influence curve among a class of influence curves the following method was proposed in Rubin and van der Laan (2008). Firstly, it is noted that

$$\begin{aligned} D^*(g,Q) &= D_{IPTW}(Q,g) - H_{CAR}(Q,g)(A - g(1 \mid W)) \\ &= H_g^*(A,W)(Y - f(Q,g)(W)) - \Psi(Q) \\ &\equiv D^*(g, f(Q,g), \Psi(Q)), \end{aligned}$$

where $H_g^*(A,W) = (2A-1)/g(A \mid W)$, and

$$f(Q,g) = g(1 \mid W)\bar{Q}(0,W) + g(0 \mid W)\bar{Q}(1,W).$$

Note that $D^*(g,f,\psi) = D_{IPTW}(g,\psi) - H_{CAR}(f,g)(A - g(1 \mid W))$, where

$$H_{CAR}(f,g) = \frac{f(W)}{g(1 \mid W)g(0 \mid W)}.$$

Thus, if we find an optimal choice $f^e$ for $f$ among a class of candidates, then that also implies an optimal choice $H_{CAR}(f^e,g)$. Determining the optimal choice for $f$ appears to be convenient, since

$$\mathrm{VAR}_{P_0} D^*(g_0, f, \psi_0) = E_0 \left\{ \{H_{g_0}^*\}^2(A,W)(Y - f(W))^2 \right\}.$$

Given a working model $\{f_\beta : \beta\}$ for $f$, one can now define an optimal choice

$$f^e(P_0) = \arg\min_{f_\beta} P_0 \left\{ \{H_{g_0}^*(A,W)\}^2 (Y - f_\beta(W))^2 \right\}. \tag{5}$$

This choice maps into a corresponding optimal

$$H_{CAR}(f_0^e, g)(W) = \frac{f_0^e(W)}{g(1 \mid W)g(0 \mid W)}.$$

The choice (5) can be estimated with weighted least squares regressing $Y$ on $W$ using weights $H_{g_0}^{*,2}$. An estimator $f_n^e$ of $f_0^e$ results in a clever covariate $H_{CAR}(f_n^e, g_n^k)$ in the $k$-th step of the TMLE-algorithm presented in previous section.

## 4.2 Adaptive empirical efficiency maximization

The choice $\bar{Q}_e(P_0)$ (1) corresponds with minimizing the empirical risk of the loss function $L_{g_0}(\bar{Q}) = \{D^*(\psi_0, \bar{Q}, g_0)\}^2$ over a working model $\{\bar{Q}_\beta : \beta\}$. Note that $L_{g_0}$ is indeed a valid loss function since $\bar{Q}_0 = \arg\min_{\bar{Q}} P_0 L_{g_0}(\bar{Q})$ (van der Laan and Robins (2003)). The strength of this loss function is that its loss-based dissimilarity is given by

$$P_0\{L_{g_0}(\bar{Q}) - L_{g_0}(\bar{Q}_0)\} = P_0 \left\{ D^*(\psi_0, \bar{Q}, g_0) - D^*(\psi_0, \bar{Q}_0, g_0) \right\}^2,$$

which follows by the Theorem of Pythagoras (van der Laan and Rose (2011)). In some cases (as in previous subsection), one can define another (e.g., squared error) loss function that has the same loss-based dissimilarity, but with an empirical risk that might be easier to minimize. The validity of the loss function relies on $g_0$ being known or consistently estimated. At the known $g_0$, this loss-based dissimilarity is targeted towards $\psi_0$ since it concerns approximating the true efficient influence curve, and it also corresponds with minimizing the variance of the influence curves $D^*(\psi_0, \bar{Q}, g_0)$ over $\bar{Q}$.

Instead of working with a single working model, we can use loss-based learning instead, using cross-validation based on this loss function $L_{g_0}$ (van der Laan and Dudoit (2003), van der Laan et al. (2007)). For example, suppose one considers a collection of $K$ working models $\{\bar{Q}^k_{\beta^k} : \beta^k\}$, $k = 1, \ldots, K$. Each working model results in an estimator $\bar{Q}^k_n$ defined by the minimizer of the empirical risk $P_n L_{g_0}(\bar{Q}^k_{\beta^k})$ over the working model indexed by parameter vector $\beta^k$. One can now select the choice $k$ of working model with the $V$-fold cross-validation selector

$$k_n = \arg\min_k \sum_{v=1}^{V} \sum_{i \in \mathrm{Val}_v} L_{g_0}(\bar{Q}^k_{n,v})(O_i),$$

where $\mathrm{Val}_v$ is the validation sample for the $v$-th sample split, and $\bar{Q}_{n,v}$ is the fit of the $k$-th working model based on the training sample $\mathrm{Train}_v$ (i.e., the complement of $\mathrm{VAL}_v$) for the $v$-th sample split, $v = 1, \ldots, V$ The estimator would now be $\bar{Q}^e_n = \bar{Q}^{k_n}_{\beta^{k_n}_n}$, which plays the role of an estimator of $\bar{Q}^e_0$.

As shown in van der Laan and Rose (2011), the general oracle results of the cross-validation selector $k_n$ apply to this loss function $L_{g_0}(\bar{Q})$ (van der Laan and Dudoit (2003)), under the assumption that the efficient influence curve is uniformly bounded in supremum norm (i.e., $\delta < g_0(1 \mid W) < 1 - \delta$ for some $\delta > 0$). As a consequence, under this boundedness condition, if none of the working models are correctly specified, the cross-validation selector will asymptotically make the optimal choice, even if the number $K$ of working models grows polynomial in sample size, while, if one of the working models is correctly specified, then the resulting $\bar{Q}^e_n$ will converge at rate $1/\sqrt{n}$ to $\bar{Q}_0$. These oracle results will also apply if $g_0$ is estimated at a rate faster than at which $\bar{Q}_0$ is estimated (van der Laan and Dudoit (2003)).

As a consequence, if $g_0$ is estimated consistently, the augmented IPTW estimator that uses $\bar{Q}^e_n$ as estimator of $\bar{Q}_0$ will now have an influence curve that is more efficient than $D^*(\psi_0, \bar{Q}^k_{\beta^k}, g_0)$ for any $\beta^k$, and any $k = 1, \ldots, K$. We note that this estimator $\bar{Q}^e_n$ based on using loss-based (super) learning can now be used in the clever covariate $H_{CAR}(\bar{Q}^e_n, g^k_n)$ in the TMLE proposed in the previous section. The TMLE presented in the previous section using this estimator $\bar{Q}^e_n$ as estimator of $\bar{Q}^e_0$ will now not only be a double robust locally efficient substitution estimator, but, if $g_0$ is estimated consistently, it will also be at least as efficient as the augmented IPTW estimator that uses $\bar{Q}^e_n$ as estimator of $\bar{Q}_0$.

If one implements an $L_{g_0}$-based super learner with a library of candidate estimators of $\bar{Q}_0$ that includes nonparametric estimators, so that at least one candidate in the library will be asymptotically consistent for $\bar{Q}_0$, then $\bar{Q}_0^e = \bar{Q}_0$ and $\bar{Q}_n^e$ is now an estimator of the globally optimal $\bar{Q}_0$. However, if one estimates $g_0$ consistently, then $\bar{Q}_n^e$ is also a fully targeted estimator of $\bar{Q}_0$ in the sense that it is tailored to result in a best estimate of the efficient influence curve itself. Again, we can use this estimator $\bar{Q}_n^e$ in the clever covariate $H_{CAR}(\bar{Q}_n^e, g_n^k)$ in the TMLE proposed in the previous section. The resulting TMLE is now not only a double robust locally efficient substitution estimator, but it is also an augmented IPTW estimator that estimates $\bar{Q}_0$ with an estimator $\bar{Q}_n^e$ that is tailored to maximize efficiency in the case that $g_n$ is consistent.

# 5 Simulations illustrating empirical efficiency property of TMLE

We illustrate the additional enhanced efficiency property of the TMLE proposed in the previous section, by comparing the performance of the enhanced TMLE with the "standard" TMLE and the empirical efficiency maximization estimator proposed in Rubin and van der Laan (2008). Two data generating distributions were defined, and the additive treatment effect parameter was estimated for one thousand samples of size $n = 500$ drawn from each. Results in Table 1 below verify that when there are no violations of the positivity assumption, at a misspecified working model for $\bar{Q}_0$ and correctly specified working model for $g_0$, the performance of the enhanced TMLE is on a par with the empirical efficiency maximization estimator, and both outperform the TMLE that does not aim for maximal efficiency.

Data were generated according to the following two mechanisms:

$$\begin{aligned}
W_1, W_2 &\sim N(0,1) \\
g_{0,1}(1 \mid W) &= 0.5 \\
g_{0,2}(1 \mid W) &= \text{Expit}(-0.3 - 0.1W_1 - 0.3W_2) \\
P_{0,1}(Y = 1 \mid A, W) &= \text{Expit}(-1 + A + W_1 + 2.5W_1^2) \\
P_{0,2}(Y = 1 \mid A, W) &= \text{Expit}(-1 + A + W_1 + 2.5W_1^2 - 0.2W_2).
\end{aligned}$$

The first simulation study mimics a randomized controlled trial in which treatment assignment is independent of baseline covariates $W = (W_1, W_2)$. The probability of being assigned to the treatment group is 0.5 for all subjects. In the second study $W_1$ and $W_2$ confound the effect of treatment on the outcome. For this simulation true treatment assignment probabilities ranged between 0.26 and 0.60. The true values of the target parameter for these two simulations are $\psi_{0,1} = 0.1579$ and $\psi_{0,2} = 0.1570$. Theses true values were obtained as an average of the additive effect $(Y_1 - Y_0)$ calculated from the full data for ten samples of size $n = 10^7$. For both studies, (misspecified) logistic linear regression of $Y$ on $(A, W_1)$ was used to obtain

Table 1: Additive treatment effect estimates, 1000 samples ($n = 500$).

| | Simulation 1 | | | Simulation 2 | | |
|---|---|---|---|---|---|---|
| | Bias | Var | MSE | Bias | Var | MSE |
| Unadj | 0.0014 | 0.0017 | 0.0017 | −0.0023 | 0.0017 | 0.0017 |
| TMLE | 0.0015 | 0.0017 | 0.0017 | −0.0006 | 0.0017 | 0.0017 |
| Emp eff | 0.0011 | 0.0015 | 0.0015 | −0.0001 | 0.0015 | 0.0015 |
| TMLE$_{en}$ | 0.0012 | 0.0015 | 0.0015 | −0.0002 | 0.0015 | 0.0015 |

the initial estimate of $\bar{Q}_0$, and the correctly specified logistic regression model was used to obtain the initial estimate of $g_0$. The estimators of $\psi_0$ are of the following form:

$$\psi_n^{TMLE} = \frac{1}{n}\sum_{i=1}^{n}\left\{\bar{Q}_n^{1*}(1, W_i) - \bar{Q}_n^{1*}(0, W_i)\right\},$$

$$\psi_n^{Empeff} = \frac{1}{n}\sum_{i=1}^{n}\frac{2A_i - 1}{g(A_i \mid W_i)}\left(Y_i - f_n^e(W_i)\right),$$

$$\psi_n^{TMLE_{en}} = \frac{1}{n}\sum_{i=1}^{n}\left\{\bar{Q}_n^{k*}(1, W_i) - \bar{Q}_n^{k*}(0, W_i)\right\}.$$

Here $\mathrm{Logit}\bar{Q}_n^{1*} = \mathrm{Logit}\bar{Q}_n^0 + \epsilon_n H_{g_n}^*$, with $\epsilon_n$ fit by maximum likelihood estimation. The target function $f_0^e(W) = f_{c_0,\alpha_0,\beta_0}(W)$, which defines the empirical efficiency maximization estimator, is defined in terms of a working model $f_{c_0,\alpha_0,\beta_0}(W) = c + \mathrm{Expit}(\alpha + \beta W_1)$ (see previous section 4). The true values $(c_0, \alpha_0, \beta_0)$ of the coefficients are fitted with weighted least squares using the *nlm* function in R (Team, 2010) and weights $\{H_{g_n}^*\}^2$. Finally, $\bar{Q}_n^{k^*}(A, W)$ is a targeted estimate of $\bar{Q}_0$ obtained by applying the iterative TMLE procedure described in Section 3 to initial estimates $\bar{Q}_n^0, g_n, f_n^e$, where $k$ denotes the final step. Convergence was defined as $\epsilon_1 < 0.00001$ and $\epsilon_2 < 0.00001$, and typically occurred after two to three iterations (so $k$ typically equals 2 or 3).

Table 1 also reports unadjusted estimates $\psi_n^{unadj} = E_n(Y = 1 \mid A = 1) - E_n(Y = 1 \mid A = 0)$, where $E_0(Y = 1 \mid A)$ is estimated with univariate logistic regression of $Y$ on $A$. The unadjusted estimator is unbiased in simulation 1, but biased in simulation 2.

Results in Table 2 verify the claim made in section 3.4 that in addition to being a double robust locally efficient substitution estimator, the enhanced TMLE is also an IPTW estimator, an augmented IPTW estimating-equation based estimator in which the nuisance parameters $Q_0, g_0$ are estimated with the TMLE $Q_n^*, g_n^*$, and an augmented IPTW estimating-equation based estimating in which the nuisance parameters $Q_0, g_0$ are estimated with $(Q_{W,n}, \bar{Q}_n^e), g_n^*$. Recall that the latter is the empirical efficiency maximization estimator of Rubin and van der Laan (2008), except that $g_0$ is estimated with $g_n^*$ instead of the initial estimator $g_n$.

Table 2: Alternative representations of additive treatment effect estimates.

| | **Simulation 1** | | | **Simulation 2** | | |
|---|---|---|---|---|---|---|
| | Bias | Var | MSE | Bias | Var | MSE |
| TMLE$_{en}$ | 0.0012 | 0.0015 | 0.0015 | $-0.0002$ | 0.0015 | 0.0015 |
| IPTW | 0.0012 | 0.0015 | 0.0015 | $-0.0002$ | 0.0015 | 0.0015 |
| AIPTW$_a$ | 0.0013 | 0.0015 | 0.0015 | $-0.0002$ | 0.0015 | 0.0015 |
| AIPTW$_b$ | 0.0012 | 0.0015 | 0.0015 | $-0.0002$ | 0.0015 | 0.0015 |

We next investigate estimator performance under increasing levels of confounding. In simulation 3a the treatment assignment mechanism is held fixed and confounding is made stronger by increasing the association between $W_2$ and the outcome, $Y$. In simulation 3b the conditional distribution of $Y$ given $(A, W)$ is held fixed while the association between $W_2$ and $A$ increases, leading to violations of the positivity assumption as confounding grows stronger. For each simulation estimates were obtained for 1000 samples of size $n = 500$ with $g_n(1 \mid W)$ bounded away from 0 and 1 at level $(p, 1-p)$, with $p = (10^{-9}, 0.01, 0.025, 0.05, .1)$. Data for simulation 3 were generated as

$$W_1, W_2 \sim N(0,1)$$
$$g_{0,3}(1 \mid W) = \mathrm{Expit}(-0.3 - 0.1W_1 - \gamma_1 W_2)$$
$$P_{0,3}(Y = 1 \mid A, W) = \mathrm{Expit}(-1 + A + W_1 + 2.5W_1^2 - \gamma_2 W_2)$$

with $\gamma_1$ fixed at 0.3 and $\gamma_2$ set to $(0, 0.1, 0.2, \ldots, 2)$ for simulation 3a, and $\gamma_1$ set to $(0, 0.2, \ldots, 1)$ while $\gamma_2$ was fixed at 1, for simulation 3b.

Figure 1 summarizes results for simulation 3a with bounds on $g_n(1 \mid W)$ set to $(10^{-9}, 1 - 10^{-9})$ and (0.1, 0.9). The bias of the unadjusted estimator (1) increases with $\gamma_2$, while the TMLE (2), Emp Eff (3), and TMLE$_{en}$ (4) estimators remain unbiased. When confounding is strong, the unadjusted estimator has the highest variance, followed by TMLE, while as predicted by theory, the variance of the TMLE$_{en}$ estimator closely matches that of the empirical efficiency estimator, designed to minimize variance. Because the treatment assignment mechanism does not lead to a violation of the positivity assumption $(0.14 < g_0(1 \mid W) < 0.77)$, results are the same regardless of the choice of truncation level for $g_n(1 \mid W)$. Estimator performance under increasing practical violations of the positivity assumption is illustrated in Figure 2, which shows results at three truncation levels of $g_n(1 \mid W)$, $(10^{-9}, 1 - 10^{-9})$, (0.025, 0.975), and (0.1, 0.9). Increasing truncation introduces a small amount of bias into TMLE, the empirical efficiency maximization estimator,

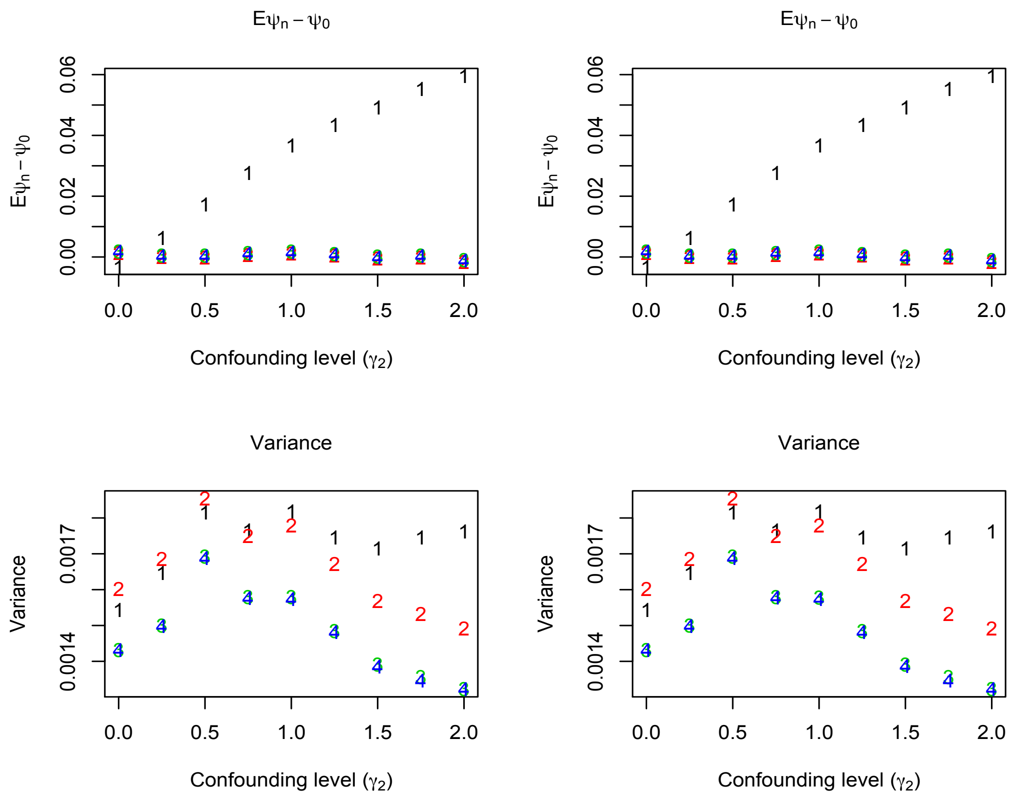


Figure 1: Simulation 3a: Estimator bias and variance at each value of $\gamma_2$, two truncation levels for $g_n(1 \mid W)$, $(10^{-9}, 1-10^{-9})$ (left), and (0.1, 0.9) (right). 1: Unadjusted, 2: TMLE, 3: Emp Eff, 4: TMLE$_{en}$.

and TMLE$_{en}$, but this amount is dwarfed by the bias of the unadjusted estimator. We observe that the variance of all but the unadjusted estimator increases with increased confounding, and is slightly ameliorated by increased truncation of $g_n(1 \mid W)$. At extreme violations of the positivity assumption (Table 3) the variance of TMLE$_{en}$(4) is slightly larger than that of the empirical efficiency maximization estimator (3), but overall these two estimators are very close to one another.

# 6 Discussion

The TMLE represents a template for construction of a loss-based substitution estimator of a target parameter defined on a semiparametric model, defined by a choice of loss function for a relevant part of the data generating distribution, a parametric submodel, and a strategy for iteratively minimizing the empirical risk

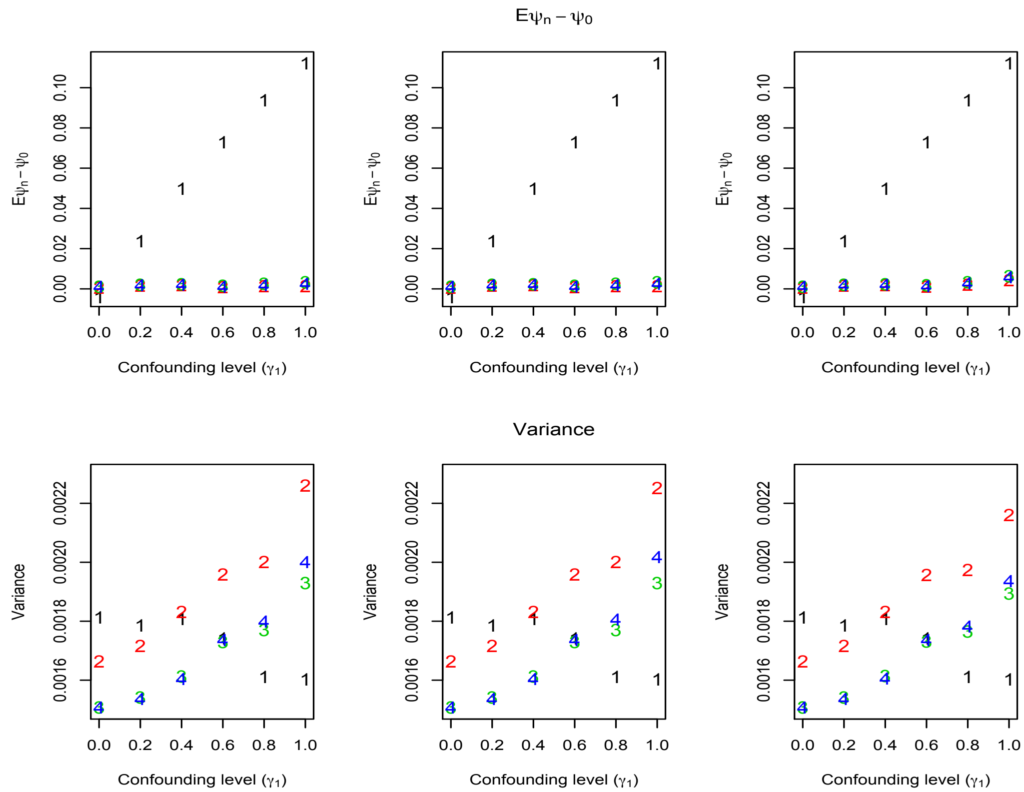


Figure 2: Simulation 3b: Estimator bias and variance at each value of $\gamma_1$. Columns correspond to truncation level for $g_n(1 \mid W)$, $(10^{-9}, 1 - 10^{-9})$ (left), (0.025, 0.05) (middle), and (0.1, 0.9) (right). 1: Unadjusted, 2: TMLE, 3: Emp Eff, 4: TMLE$_{en}$.

Table 3: True conditional treatment assignment probabilities as a function of $\gamma_1$.

| $\gamma_1$ | Range of $g_0(A \mid W)$ | | $\gamma_1$ | Range of $g_0(A \mid W)$ | |
|---|---|---|---|---|---|
| 0 | 0.305 | 0.551 | 0.6 | 0.035 | 0.926 |
| 0.2 | 0.212 | 0.676 | 0.8 | 0.013 | 0.969 |
| 0.4 | 0.090 | 0.837 | 1 | 0.005 | 0.987 |

over the parametric submodel. The choice of submodel and loss function defines the score equations the TMLE will solve. In this manner it can be arranged that the TMLE solves not only the efficient score equation, but also an estimating equation corresponding with the influence curve of a competing estimator. By solving this estimating equation the TMLE is at least as efficient as the competing estimator in the case this competing estimator is asymptotically linear and $g_n^0$ is consistent.

We demonstrated this type of TMLE for the simple point treatment data structure $(W, A, Y)$ and the additive effect parameter. Our presentation is straightforwardly generalized to general CAR-censored data models, and target parameters, since we only relied on a general representation of the efficient influence curve as an augmented IPCW-function as presented in Robins and Rotnitzky (1992); van der Laan and Robins (2003). Suppose now that the target parameter is multivariate. One needs to define the collection of real valued parameters, and one needs to define a competing estimator for each of these real valued parameters. For example, one might define one single real valued parameter as a function of the multivariate parameter, or one might define each component of the target parameter as a real valued parameter. Each of the real valued parameters implies now an influence curve of the corresponding competing estimator. Each of these influence curves implies a clever covariate for the treatment mechanism playing the role of $H(Q_n^e, g_n^k)$ in the above TMLE algorithm. The resulting TMLE will not only be a double robust locally efficient substitution estimator of the target parameter, but it will also estimate each of the real valued parameters in a more efficient way than the competing estimators, in the case that $g_0$ is estimated consistently.

## References

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

S. Gruber and M.J. van der Laan. An application of collaborative targeted maximum likelihood estimation in causal inference and genomics. *The International Journal of Biostatistics*, 2010a.

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(26), 2010b.

J. Pearl. *Causality: Models, Reasoning, and Inference.* Cambridge University Press, Cambridge, 2nd edition, 2000.

J.M. Robins and A. Rotnitzky. Recovery of information and adjustment for dependent censoring using surrogate markers. In *AIDS Epidemiology*, Methodological issues. Bikhäuser, 1992.

A. Rotnitzky, Q. Lei, M. Sued, and J. Robins. Enhanced double-robust locally-efficient estimation in regression models for longitudinal studies. Presentation ENAR 2011 Spring Meeting, article submitted, March 2011.

D.B. Rubin and M.J. van der Laan. Empirical efficiency maximization: Improved locally efficient covariate adjustment in randomized experiments and survival analysis. *The International Journal of Biostatistics*, Vol. 4, Iss. 1, Article 5, 2008.

Z. Tan. Bounded, efficient, and doubly robust estimation with inverse weighting. *Biometrika*, 94:1–22, 2008.

R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2010.

M. J. van der Laan. The construction and analysis of adaptive group sequential designs. Technical Report 232, www.bepress.com/ucbbiostat/paper232, University of California, Berkeley, 2008.

M.J. van der Laan and S. Dudoit. Unified cross-validation methodology for selection among estimators and a general cross-validated adaptive epsilon-net estimator: Finite sample oracle inequalities and examples. Technical report, Division of Biostatistics, University of California, Berkeley, November 2003.

M.J. van der Laan and J.M. Robins. *Unified methods for censored longitudinal data and causality*. Springer, New York, 2003.

M.J. van der Laan and S. Rose. *Targeted Learning: Prediction and Causal Inference for Observational and Experimental Data*. Springer, New York, 2011.

M.J. van der Laan and D. Rubin. Targeted maximum likelihood learning. *The International Journal of Biostatistics*, 2(1), 2006.

M.J. van der Laan, E. Polley, and A. Hubbard. Super learner. *Statistical Applications in Genetics and Molecular Biology*, 6(25), 2007. ISSN 1.

W. Zheng and M. J. van der Laan. Asymptotic theory for cross-validated targeted maximum likelihood estimation. Technical Report 273, www.bepress.com/ucbbiostat/paper273, University of California, Berkeley, 2010.

## Appendix: R Implementation

The R function below calculates the enhanced TMLE for binary outcomes. Required arguments are $Y$ (binary outcome vector), $A$ (binary treatment indicator vector), and initial estimates $\bar{Q}_n^0(A,W)$, $g_n^0(A \mid W)$, and $f_n^e(W)$. $\bar{Q}_n^0(A,W)$ is an $n \times 3$ matrix containing values for $\bar{Q}_n^0(A,W), \bar{Q}_n^0(0,W)$, and $\bar{Q}_n^0(1,W)$ on the logit scale. Predicted values for $g_n(A \mid W)$ are bounded away from 0 and 1.

```
bound <- function(x, bounds){
  x[x<min(bounds)] <- min(bounds)
  x[x>max(bounds)] <- max(bounds)
  return(x)
}

tmle_en <- function(Y, A, g1W, Q, f, gbds = c(10^-9, 1-10^-9)){
  g1w <- bound(g1W, gbds)
  eps1 <- eps2  <- Inf
  epsilon <- .00001
  maxIter <- 30
  iterations <- 0
  while((any(abs(c(eps1, eps2)) > epsilon)) & iterations <= maxIter){
    iterations <- iterations + 1
    h <- cbind(A/g1W - (1-A)/(1-g1W), 1/g1W, -1/(1-g1W))
    m <- glm(Y ~ -1 + offset(Q[,"QAW"]) + h[,1], family=binomial)
    eps1 <- coef(m)
    Q <- Q +  eps1*h

    h2 <- plogis(Q[,"Q1W"])/g1W + plogis(Q[,"Q0W"])/(1-g1W)
    h3 <- f/(g1W * (1-g1W))
    g  <- glm(A ~ -1 + offset(qlogis(g1W)) + h2 + h3, family=binomial)
    g1W  <- bound(predict(g, type = "response"), gbds)
    eps2 <- coef(g)
  }
  Q <- plogis(Q)
  psi.en   <- mean(Q[,"Q1W"] - Q[,"Q0W"])
  psi.IPTW <- mean((A/g1W - (1-A)/(1-g1W)) * Y)
  psi.AIPTWQstargstar <- mean((A/g1W - (1-A)/(1-g1W)) * Y
                      - (Q[,"Q1W"]/g1W - Q[,"Q0W"]/(1-g1W))*(A-g1W))
  psi.AIPTWQegstar    <- mean((A/g1W - (1-A)/(1-g1W)) * Y
                      - f/(g1W * (1-g1W)) * (A-g1W))
  return(c(psi.en, psi.IPTW, psi.AIPTWQstargstar, psi.AIPTWQegstar))
}
```