Benjamin Rich, Erica E. M. Moodie* and David A. Stephens Influence Re-weighted G-Estimation

DOI 10.1515/ijb-2015-0015

Abstract: Individualized medicine is an area that is growing, both in clinical and statistical settings, where in the latter, personalized treatment strategies are often referred to as dynamic treatment regimens. Estimation of the optimal dynamic treatment regime has focused primarily on semi-parametric approaches, some of which are said to be doubly robust in that they give rise to consistent estimators provided at least one of two models is correctly specified. In particular, the locally efficient doubly robust g-estimation is robust to misspecification of the treatment-free outcome model so long as the propensity model is specified correctly, at the cost of an increase in variability. In this paper, we propose data-adaptive weighting schemes that serve to decrease the impact of influential points and thus stabilize the estimator. In doing so, we provide a doubly robust g-estimator that is also robust in the sense of Hampel (15).

Keywords: asymptotic linearity, g-estimation, optimal dynamic treatment regimen structural nested mean models

1 Introduction

The “personalization” of medicine is an area that is receiving ever-growing attention from medical practitioners, statisticians, and computer scientists, to name but some of the interested parties. The U.S. Food and Drug Administration refers to this personalization as “tailoring medical treatment to [...] a patient’s genetic, anatomical, and physiological characteristics”, approaches which are “allowing patients to be treated and monitored more precisely and effectively and in ways that better meet their individual needs” [1]. Within the statistics literature, there have been a great variety of methods proposed for estimating optimal personalized treatment strategies, also known as dynamic treatment regimens (DTRs), from regression-based approaches (Schulte et al. [2] provide an excellent review) to classification-based algorithms (e.g. Zhang et al. [3], Zhao [4]). Although there have been some fully parametric approaches considered [5, 6], the majority of estimation approaches have been semi-parametric (e.g. Murphy [7], Orellana et al. [8], Robins [9], van der Laan and Petersen [10], Zhang et al. [11]).

Semi-parametric methods frequently give rise to estimators that are *locally efficient*, meaning that the estimator’s asymptotic variance achieves the semi-parametric efficiency bound if the true data generating distributing happens to belong to a certain class, yet remain consistent and asymptotically normal (CAN) under broader conditions. This is true, in particular, for any estimator that bears the designation “doubly robust” [12]. In causal inference, doubly robust estimators typically involve the specification of two nuisance models, one modelling the treatment assignment mechanism in terms of confounders (e.g. a propensity score [13]), the other modelling the outcome in terms of confounders and risk factors. The common approach is to specify parametric working models for these nuisance models, though non-parametric alternatives (e.g. trees) could equally be viable. The estimator achieves semi-parametric efficiency if all parametric components, including both nuisance models, are specified correctly.

Depending on the context, one or the other nuisance model may be easier to specify correctly. Randomized trials are one setting where the treatment assignment mechanism is known, for instance. In a pharmacoepidemiological context, it may be possible to construct a realistic propensity score model through consultation with clinicians who understand the factors that are considered in making treatment

*Corresponding author: Erica E. M. Moodie, Department of Epidemiology, Biostatistics, and Occupational Health, McGill University, Montreal, QC, Canada, E-mail: erica.moodie@mcgill.ca

Benjamin Rich, Department of Epidemiology, Biostatistics, and Occupational Health, McGill University, Montreal, QC, Canada
David A. Stephens, Department of Mathematics and Statistics, McGill University, Montreal, QC, Canada

decisions. The effect of the treatment on a clinical outcome may be more difficult to model as it involves complex biological systems, interactions with genes and environment, etc.; it is plausible that partial model misspecification occurs quite commonly in practice. Since standard estimators are not efficient under partial misspecification, it is natural to look for ways to improve on them. In particular, it would be useful to find a means of making the doubly-robust estimators more stable – further “robustifying” such estimators to influential observations.

In this work, we undertake a data-adaptive approach to increasing the stability of the doubly-robust g-estimators of optimal dynamic treatment regimen structural nested mean models (SNMMs) [9]. In the next section, we review g-estimators of DTRs and present results which will be used to develop our proposed approach, in which the g-estimation equation at each interval is modified to incorporate a weight determined by a prediction of the influence of the individual at that interval.

2 Background

2.1 G-estimation of SNMMs for DTRs

A dynamic treatment regimen is a sequence of functions $d = (d_1, \dots, d_K)$ where K is the number of treatment intervals and for each interval $j = 1, \dots, K$ the function $d_j : \mathcal{H}_j \mapsto \mathcal{A}_j$ maps the set of possible observed histories at interval j to the set of possible treatments at interval j . Letting A_j denote treatment received at interval j , L_j denote observations made subsequent to receiving the treatment but prior to receiving the next treatment, and L_0 baseline covariates, an observed history is given by $H_j = (L_0, A_1, L_1, A_2, \dots, L_{j-1})$. In this longitudinal setting, the presence of qualitative interactions dictates that the best treatment strategy involves individualization based on the observed patient history up to the point at which the decision is made. To estimate such a strategy, Robins [9] proposed the SNMM. This model is specified via a *blip function*, which models the interactions between treatment and covariates in the following way. Let $Y = g(O)$ be the final outcome of interest, where $g(\cdot)$ is a suitably defined *reward function* (so that greater values of Y are more desirable). The blip function is defined in terms of potential outcomes (or *counterfactuals*); let $Y(d)$ denote the potential outcome under DTR d . Let $d^{h_j, a, \text{opt}}$ denote the DTR consisting of $a_1, a_2, \dots, a_{j-1}$ (components of h_j) for the first $(j-1)$ treatment intervals, then a at the j^{th} interval, then the optimal DTR for the remaining $(K-j)$ intervals (it is understood that either $(j-1)$ or $(K-j)$ may be zero, in which case either the component before or after a becomes unnecessary). Let 0 denote a reference treatment, such as no treatment or placebo. The blip function for the j^{th} treatment interval is given by

$$\gamma_j(a, h_j) = E[Y(d^{h_j, a, \text{opt}}) - Y(d^{h_j, 0, \text{opt}}) | H_j = h_j, A_j = a]$$

Two key features of the blip function are (1) the blip function evaluated at $a = 0$ is zero (by definition); (2) the blip function determines the optimal DTR through $d^{\text{opt}}(h_j) = \arg \max_a \gamma_j(a, h_j)$. Under standard identifiability and regularity assumptions given in Robins [9], the blip function can be estimated. A SNMM sets $\gamma_j(a, h_j)$ equal to some function of a parameter ψ_j and covariates $x_j = x_j(h_j)$, for instance $\gamma_j(a, h_j; \psi_j) = \psi_j^\top \cdot ax_j$. (Typically, the blip function is assumed correctly specified, but if not, then the parameter ψ_j would correspond to a projection of the truth onto the specified model, resulting in a DTR that is optimal within the class defined by the model.) Let

$$G_j(\psi_j, \underline{\psi}_{j+1}) = Y - \gamma_j(A_j, H_j, \psi_j) + \sum_{m=j+1}^K [\max_{a' \in \mathcal{A}_m} \{\gamma_m(a', H_m; \psi_m)\} - \gamma_m(A_m, H_m; \psi_m)]$$

where $\underline{\psi}_{j+1} = (\psi_{j+1}^\top, \dots, \psi_K^\top)^\top$. Doubly robust estimators of the interval-specific parameters ψ_j are obtained (under regularity conditions) by recursively solving the g-estimating equations

$$0 = \sum_{i=1}^n X_{ij} \{A_{ij} - \pi_j(H_{ij}; \hat{\alpha}_j)\} \{G_{ij}(\psi_j, \hat{\underline{\psi}}_{j+1}) - \eta_j(H_{ij}; \hat{\zeta}_j(\psi_j, \hat{\underline{\psi}}_{j+1}))\}$$

where $\pi_j(h_j; \alpha_j)$ is a posited or working model for $E[A_j|H_j = h_j] = \Pr[A_j = 1|H_j = h_j]$, $\eta_j(h_j; \varsigma_j)$ is a posited or working linear model for $E[G_j(\psi_j, \underline{\psi}_{j+1})|H_j, A_j] = E[Y(d^{h_j, 0, \text{opt}})|H_j, A_j]$, and $\hat{\varsigma}_j(\psi_j, \underline{\psi}_{j+1})$ is the ordinary least-squares estimator of ς_j , determined up to a parameter.

We shall now consider a partially misspecified setting, namely we shall suppose that while the treatment or propensity model $\pi_j(h_j; \alpha_j)$ is correctly specified, the treatment-free outcome model $\eta_j(h_j; \varsigma_j)$ is not. As noted previously, the resulting g-estimator $\hat{\psi}_j$ is CAN but may be susceptible to high variation due to influential observations. This is the point we aim to address. Model diagnostics can help. In a previous paper [14], we proposed residual-based diagnostics for g-estimation of optimal dynamic treatment regimens which can serve as a guide to obtaining a better fitting outcome nuisance model, and also showed in simulations that improved nuisance model fit resulted in improved estimator efficiency for the parameter of interest. In this paper, we will pursue a different approach for obtaining more robust (in the sense of Hampel [15]) g-estimators in the same context through re-weighting. Here we will assume that the model specification is fixed, i.e. cannot be improved through diagnostics, and focus on the properties of the g-estimator for this fixed model under a varying set of weights.

2.2 Weighting in estimating equations: some general results

Suppose that we have observed i.i.d. data $O_1, \dots, O_n$ from a population, and let O denote a generic observation for a single individual from the same population, and o a realization. Let β be a parameter of interest and let $u(o, \beta)$ be a consistent estimating function, i.e. we have $E[u(O, \beta_0)] = 0$, where β_0 denotes the truth (we will suppose that β_0 is the unique solution to this equation, which in the present context is true if models $\gamma_j(a, h_j; \psi_j)$ and $\eta_j(h_j; \varsigma_j)$ are linear in parameters).

Standard M-estimator theory Tsiatis [16] tells us that we can construct an (unweighted) estimator of β that is *regular asymptotically linear* (RAL) and hence consistent asymptotically normal as the solution to $0 = \sum_{i=1}^n u(O_i, \hat{\beta}_n)$. Under regularity conditions, $\hat{\beta}_n$ is consistent, has influence function given by:

$$\phi(o) = - \left[E \left[\frac{\partial u(O, \beta_0)}{\partial \beta} \right] \right]^{-1} u(o, \beta_0),$$

and asymptotic variance given by $M \text{Var}(u(O, \beta_0)) M^T$ where

$$M = E \left[\frac{\partial u(O, \beta_0)}{\partial \beta} \right]^{-1}$$

Now, let $w(O_i)$ be a strictly positive weight function that does not depend on β . We will use the more concise notation $W_i = w(O_i)$, and $U_i(\cdot) = u(O_i, \cdot)$. Note that the weights W_i are data-dependent; this generalized form of M-estimation does require additional technical conditions, which we shall discuss following Lemma 1 below. Then, a weighted estimator $\hat{\beta}_n^w$ is given by the solution to the estimating equation $0 = \sum_{i=1}^n W_i U_i(\hat{\beta}_n^w)$. We will explore the properties of weighted estimators in terms of bias and variance.

First, we observe that the usual Taylor expansion gives

$$0 = \sum W_i U_i(\hat{\beta}_n^w) = \sum W_i U_i(\beta_0) + \sum W_i \frac{\partial U_i(\beta^*)}{\partial \beta} (\hat{\beta}_n^w - \beta_0)$$

where β^* lies between β_0 and $\hat{\beta}_n^w$. So

$$\begin{aligned} \sqrt{n}(\hat{\beta}_n^w - \beta_0) &= - \left[n^{-1} \sum W_i \frac{\partial U_i(\beta^*)}{\partial \beta} \right]^{-1} n^{-1/2} \sum W_i U_i(\beta_0) \\ &\rightarrow^d - E \left[W_i \frac{\partial U_i(\beta_0)}{\partial \beta} \right]^{-1} \sqrt{\text{Var}(W_i U_i(\beta_0))} Z + o(1) \end{aligned}$$

Where Z has the standard normal distribution.

If $\hat{\beta}_n^w$ happens to be $\sqrt{n}$ -consistent, i.e. if $\sqrt{n}(\hat{\beta}_n^w - \beta_0) \rightarrow^d N(0, \text{Var}(\hat{\beta}_n^w))$, then the same M-estimator theory applies to $\hat{\beta}_n^w$ as to $\hat{\beta}_n$, and consequently, under regularity conditions, the asymptotic variance of $\hat{\beta}_n^w$ is given by $M_w \text{Var}(W_i U_i) M_w^T$ where

$$M_w = E \left[W_i \frac{\partial U_i(\beta_0)}{\partial \beta} \right]^{-1}$$

We now give a sufficient condition for the consistency of $\hat{\beta}_n^w$:

Lemma 1 *If the conditional expectation*

$$E[U(\beta_0)|W] = 0,$$

i.e. within levels of constant weight, the estimating function remains consistent, then the weighted estimator $\hat{\beta}_n^w$ is also RAL and CAN.

Proof. If the weights are not data-dependent, then by iterated expectation:

$$E[WU(\beta_0)] = E[E[WU(\beta_0)|W]] = E[WE[U(\beta_0)|W]] = 0,$$

so the usual M-estimator theory also applies to $\hat{\beta}_n^w$. As the weights are, in general, data-dependent, the proof can be made under the limiting values of W_i following the consistency requirements for the estimators used to define W_i using the arguments for adaptive weighted estimators outlined in Section 4.4 of Newey [17].

3 Weighting within the g-estimation framework

For each interval $j = 1, \dots, K$, let $W_{ij} = w_j(O_i)$ be a (scalar) weight assigned to the i^{th} individual in the data set for interval j . The weighted g-estimator based on weight functions $(w_1, \dots, w_K)$ is the solution to the system of estimating equations

$$0 = \sum_{i=1}^n W_{ij} X_{ij} \{A_{ij} - \pi_j(H_{ij}; \hat{\alpha}_j)\} \{G_{ij}(\hat{\psi}_j) - \eta_j(H_{ij}; \hat{\varsigma}_j(\hat{\psi}_j))\},$$

for $j = 1, \dots, K$ where, in the recursive procedure, the j^{th} equation is solved for $\hat{\psi}_j$ with all other parameters substituted by their consistent estimates. The weighted g-estimator is consistent when the conditions of Lemma 1 hold for each estimating function. In particular, under the assumptions considered of sequential randomization and a correctly specified propensity model, this holds if each $w_j(\cdot)$ is functionally independent of $o \setminus h_j$, the set difference between all observed data and the history at interval j , i.e. if the weight at interval j depends only on the history at interval j and we can write $W_{ij} = w_j(H_{ij})$. In that case, it follows by iterated expectation that

$$\begin{aligned} E[X_j \{A_j - \pi_j(H_j; \alpha_j)\} \{G_j(\psi_j) - \eta_j(H_j; \varsigma_j(\psi_j))\} | w_j(H_j)] &= \\ E[E[X_j \{A_j - \pi_j(H_j; \alpha_j)\} \{G_j(\psi_j) - \eta_j(H_j; \varsigma_j(\psi_j))\} | H_j] | w_j(H_j)]. \end{aligned} \quad (1)$$

The inner expectation can be shown to zero under the assumptions of consistency and sequential randomization and a correctly specified propensity model.

3.1 Influence-based weighting schemes

Under correct specification of the treatment-free outcome model, Robins [9] gives an expression for a weight function for g-estimation that is optimal in an efficiency sense (under some further assumptions). The

weight is proportional to the inverse of the conditional variance $\text{Var}(G_j(\underline{\psi}_j)|H_j)$. Under a misspecified treatment-free outcome model, however, this weight function would be unavailable. Joffe and Brensinger [18] discuss weighting in g-estimation, but in the context of an instrumental-variables analysis for estimating treatments effects in randomized trials with non-compliance. Focussing on the single-interval case, their proposal of weighting by the compliance score assigns larger weights to compliers. Some optimality results are given, but again only under correct model specification.

We propose a different approach to constructing weights which seeks to stabilize the standard estimator in a manner which does not depend on correct model specification and can improve its performance as measured by mean squared error in certain cases (as we will see in a simulation study of Section 4).

Data points that are candidates to receive reduced weights in the estimating equation are points with high sample influence. Hinkley [19] and Krasker and Welsch [20] discuss the use of weights constructed as functions of the sample influence in the context of robust regression estimators. Here, the proposal is to identify data points that are highly influential in the estimation of the parameter of interest, and reducing the influence of those individuals through subsequent re-weighting, will stabilize the estimation and make it more robust to misspecification of the nuisance treatment-free outcome model. The sample influence of a data point is, however, constructed from all the available data. Thus, while it is straightforward to describe an algorithm in which weights are constructed such that the weight for individual I is inversely proportional to the sample influence of individual i on the estimate of ψ_j (as we now proceed to do), the asymptotic properties of the resulting estimator are difficult to assess the terms of the estimating equation are not i.i.d. since the weights depend on all of the available data and the treatment mechanism and outcome regression models depend on all data (although the weight functions converge to functions which only rely on the individual's history). As the terms are not independent, the standard weighting theory described above does not apply and the resulting estimator will not in general be consistent for the parameter of interest. We begin, for simplicity, by ignoring these issues, then address them with a simple adaptation of the influence re-weighting g-estimation procedure.

3.2 Influence measures in g-estimation

Influence refers to the stability of statistics (e.g. estimators) to perturbations of the data. The mathematical object used to describe the asymptotic stability of an estimator is the influence function (IF) (Hampel et al. [15], Tsiatis [16], Cook and Weisberg [21]). The idea is to consider the distribution function F from which the observed data $Z_1, \dots, Z_n$ are independently drawn, a statistic $T(F)$, and form the directional derivative

$$\phi_{T,F}(z) = \lim_{\epsilon \rightarrow 0^+} \frac{T((1-\epsilon)F + \epsilon\delta_z) - T(F)}{\epsilon} \quad (2)$$

where δ_z represents a point-mass density at z . Thus, the IF $\phi_{T,F}(z)$ of T with respect to F describes the effect on T of an infinitesimal “contamination” at z [15].

The IF is a useful theoretical tool, but its use in practice for model diagnostics is limited since it depends on the true data generating distribution, which is only available under correct model misspecification and infinite samples. More useful for diagnostics is the finite sample analogue of the IF, obtained through the jackknife (leave-one-out) approach. Consider $\hat{F}_{(i)}$, the empirical distribution function for a data set from which the i^{th} observation has been removed, thus

$$\hat{F}_{(i)} = \sum_{i' \neq i} \left(\frac{1}{n-1} \right) \delta_{z_{i'}}.$$

Then, the sample influence (SI) of observation i is obtained from (2) by setting $F = \hat{F}_{(i)}$, $z = z_i$ and $\epsilon = 1/n$, given

$$SI_i = \frac{T\left(\left(\frac{n-1}{n}\right)\hat{F}_{(i)} + \left(\frac{1}{n}\right)\delta_{z_i}\right) - T(\hat{F}_{(i)})}{\left(\frac{1}{n}\right)} \propto T(\hat{F}) - T(\hat{F}_{(i)}).$$

In standard outcome regression such as $E[Y|X] = m(X; \beta)$, where $Z = (X, Y)$, the usual measures derived from the sample influence (e.g. Seber and Lee [22]) include:

1. DFBETA_i [23], defined by:

$$\hat{\beta} - \hat{\beta}_{(i)};$$

2. DFFIT_i [23], defined by:

$$m(x_i; \hat{\beta}) - m(x_i, \hat{\beta}_{(i)});$$

3. Cook's distance D_i [21], defined by:

$$(\hat{\beta} - \hat{\beta}_{(i)})^T \Sigma^{-1} (\hat{\beta} - \hat{\beta}_{(i)})$$

for a symmetric and positive definite matrix Σ , typically a scaled estimate of the covariance matrix of $\hat{\beta}$.

Preisser and Qaqish [24] extend these definitions to the repeated measures setting. They discuss both the deletion of single observations, whole clusters (individuals in the longitudinal setting), and other subsets.

Consider now the g-estimator $\hat{\psi}_j$. Let $\hat{\psi}_{j(i)}$ be the estimate of ψ_j obtained after deletion of individual i . Then, the sample influence of individual i on the individual components of the estimate $\hat{\psi}_j$ are given by $\hat{\psi}_j - \hat{\psi}_{j(i)}$, which is equivalent to DFBETA above. We will also be interested in the overall or compound influence of individual i , defined in a manner that is analogous to Cook's D_i above, but where we consider the square root:

$$\mathcal{I}_{ij} = \sqrt{(\hat{\psi}_j - \hat{\psi}_{j(i)})^T \hat{\Sigma}_{\psi_j}^{-1} (\hat{\psi}_j - \hat{\psi}_{j(i)})}$$

and $\hat{\Sigma}_{\psi_j}$ is an estimate of the covariance matrix of $\hat{\psi}_j$. There are no computational shortcuts as in the case of linear regression, where elegant mathematics allow the deletion diagnostics to be computed very efficiently [21], here the computation is by brute force. Nonetheless, with closed form g-estimation and fast computers, the computational burden is modest for reasonable sample sizes.

3.3 Influence re-weighted g-estimation

The procedure begins with an initial solution to the weighted g-estimating equations, using some initial set of weights that we denote W_{ij}^0 (note that we are free to set $W_{ij}^0 = 1$ for all i, j). Based on this initial solution, the sample influence of each data point at each interval is computed using the usual leave-one-out approach as described in previous section. Let $\mathcal{I}_{ij}$ denote the influence of the observation made on individual i at interval j . For instance, let

$$\mathcal{I}_{ij} = \sqrt{(\hat{\psi}_{j(i)} - \hat{\psi}_j)^T \hat{\Sigma}_{\psi_j}^{-1} (\hat{\psi}_{j(i)} - \hat{\psi}_j)} \quad (3)$$

where $\hat{\psi}_{j(i)}$ denotes the estimate of ψ_j obtained when the i^{th} individual is deleted from the data set (or equivalently, assigned weight zero), and $\hat{\Sigma}_{\psi_j}$ is an estimate of the covariance matrix of $\hat{\psi}_j$. Define a monotone non-increasing function $\rho: \mathbb{R}^+ \rightarrow \mathbb{R}^+$. For instance, we propose the family of functions

$$\rho_\lambda(x) = \frac{1}{(1 + (\frac{x}{\lambda_1})^2)^{\lambda_2}}$$

where $\lambda = (\lambda_1, \lambda_2) \in \mathbb{R}^+ \times \mathbb{R}^+$ is a tuning parameter. The function $\rho_\lambda(x)$ incorporates a “scale” parameter (λ_1) and a “shape” parameter (λ_2) is smooth and bounded above at 1 so that weights cannot be absurdly large. The procedure assigns to individual i at interval j the weight $W_{ij} = W_{ij}^0 \times \rho(\mathcal{I}_{ij})$. The g-estimating equations are re-solved using this set of weights and a new estimate of ψ is obtained. We refer to an estimator constructed in this manner as an *influence-based weight* (IBW) g-estimator.

In the proposed family ρ_λ , the tuning parameter λ determines the shape of the “down-weighting curve”, i.e. how quickly the weights of points with high influence are shrunk down to zero. While the choice of the λ should be application-specific and is left to the analyst, some guidance can be provided. The consideration is the trade-off between bias and variance. To gain a sense for the bias and variance of the weighted estimator under different choices of λ , bootstrap re-sampling could be used. A less computationally intensive approach is to evaluate the estimator under different choices on a single data set, as Cole and Hernán [25] do in their discussion of the bias-variance trade-off involved in the selection of a truncation point for weights in IPTW estimation. For example, we could study the impact of choosing λ such that the median weight is equal to 0.5, 0.55, 0.6, ..., 0.95 when the weights have been normalized such that the largest weight is 1.

3.4 Re-weighting “inside-the-loop” and “outside-the-loop”

When g-estimation is applied recursively, there are two different variants of the algorithm for computing a re-weighted estimator depending on the ordering of the calculations. We have called them “inside-the-loop” and “outside-the-loop”, making reference to a programming loop with a loop index j that goes down from K to 1. The two variants are sketched in pseudo-code in Figure 1. Essentially, the distinction is whether the initial estimate of ψ_j uses the weighted or unweighted estimate of ψ_m for intervals $j < m$. In the “inside-the-loop” variant, the weighted estimates at later intervals are computed first, and used in the calculation of the initial estimate of ψ_j . In the “outside-the-loop” variant, all initial estimates are computed before any of the weighted estimates. In the simulations, the “inside-the-loop” variant was used because we expect the weighted estimates of ψ_m for $m > j$ to be slightly more stable than the unweighted estimates, and hence lead to a slightly better initial estimate of ψ_j .

(a) “Inside-the-loop”

```

Let  $j$  iterate through treatment intervals from  $K$  to 1:
  Compute initial estimate  $\hat{\psi}_j$  using weighted  $\hat{\psi}_{j+1}^W$ ;
  Let  $i$  iterate over individuals from 1 to  $n$ :
    Compute  $\hat{\psi}_{j(i)}$  and  $\mathcal{I}_{ij}$ ;
    Compute weights  $W_{ij} = W_{ij}^0 \times \rho(\mathcal{I}_{ij})$ ;
  End
  Compute weighted estimate  $\hat{\psi}_j^W$  using weighted  $\hat{\psi}_{j+1}^W$ ;
End

```

(b) “Outside-the-loop”

```

Let  $j$  iterate through treatment intervals from  $K$  to 1:
  Compute initial estimate  $\hat{\psi}_j$  using unweighted  $\hat{\psi}_{j+1}$ ;
End
Let  $i$  iterate over individuals from 1 to  $n$ :
  Let  $j$  iterate through treatment intervals from  $K$  to 1:
    Compute  $\hat{\psi}_{j(i)}$  and  $\mathcal{I}_{ij}$ ;
    Compute weights  $W_{ij} = W_{ij}^0 \times \rho(\mathcal{I}_{ij})$ ;
  End
End
Let  $j$  iterate through treatment intervals from  $K$  to 1:
  Compute weighted estimate  $\hat{\psi}_j^W$  using weighted  $\hat{\psi}_{j+1}^W$ ;
End

```

Figure 1: Algorithms for re-weighting “inside-the-loop” (a) and “outside-the-loop” (b).

3.5 Consistent re-weighting

As stated previously, the estimation procedure described above may have undesirable properties. In particular, because the terms of the estimating equations are not i.i.d. and the conditions of Lemma 1 do not necessarily hold, consistency of the estimator is not assured. We now propose a simple way to adapt the procedure and obtain a different estimator, similarly motivated, that is not subject to these difficulties.

The idea is to construct a set of weights such that the weight for individual i at interval j is a function of the history H_{ij} only. Consistency then follows from Lemma 1, Equation (3), and regularity conditions on the weights. To do so, a predictive model for the sample influence $\mathcal{I}_{ij}$ using H_{ij} as predictors is developed. To compute the new weights, the function ρ is then applied to the predicted sample influence rather than actual sample influence. We will refer to an estimator that uses weights based on a predicted influence as a *predicted influence-based weight* (PIBW) estimator.

The challenge with this approach is in constructing predictive models for the sample influence at each interval. We approached this in two different ways: direct prediction of the compound influence $\mathcal{I}_{ij}$, and prediction of the individual components of $\hat{\psi}_{j(i)}$. For direct prediction of the compound influence $\mathcal{I}_{ij}$, we propose to regress $\log \mathcal{I}_{ij}$ on H_{ij} using some possibly flexible model, and take the exponential of the predicted value as $\hat{\mathcal{I}}_{ij}(h_j)$. The log transformation is appropriate because $\mathcal{I}_{ij}$ is a non-negative quantity. We refer to this approach as PIBW1.

Unlike the compound influence $\mathcal{I}_{ij}$ which is non-negative, the individual components of $\hat{\psi}_{j(i)}$ can take any real value so no transformation is needed and we suppose that the relationship between covariates and components of $\hat{\psi}_{j(i)}$ will be more linear, and that a less flexible model will be required for prediction. However, multiple models are required since each component must be predicted separately. Once these separate predictions have been computed, $\hat{\mathcal{I}}_{ij}(h_j)$ is obtained using (3) but with each component of $\hat{\psi}_{j(i)}$ replaced by its prediction. We refer to this approach as PIBW2.

Yet a third approach that we considered was to orthogonalize the individual components before prediction. This allows for better prediction in the case where the individual components of $\hat{\psi}_{j(i)}$ are highly correlated. Specifically, let $q_{ij} = \hat{\Sigma}_{\psi_j}^{-1/2}(\hat{\psi}_{j(i)} - \hat{\psi}_j)$ where $\hat{\Sigma}_{\psi_j}^{-1/2}$ is the inverse of the Choleski decomposition of the estimate of the covariance matrix of $\hat{\psi}_j$. Then, regress each component of the vectors q_{ij} on H_{ij} to obtain predictions $\hat{q}_{ij}(H_{ij})$, and finally set $\hat{\mathcal{I}}_{ij}(H_{ij}) = \sqrt{\hat{q}_{ij}(H_{ij})^T \hat{q}_{ij}(H_{ij})}$. We refer to the approach of orthogonalizing before predicting as PIBW3.

Once the prediction $\hat{\mathcal{I}}_{ij}(H_{ij})$ has been obtained (using any of the three versions described above), the weights are computed analogously to the IBW case, namely $W_{ij}(H_{ij}) = W_{ij}^0 \times \rho(\hat{\mathcal{I}}_{ij}(H_{ij}))$. For inference, the variance estimator for $\hat{\psi}_j$ can be adjusted to account for the estimation of an additional nuisance model in the usual way (i.e. following the same approach used to adjust for the estimation of the parameter α_j in the propensity score model).

4 Simulations

Simulations were performed using the basic approach described in Moodie, Richardson, and Stephens [26] but with different parameter settings. The basic setup for the simulation is presented in Figure 2. Note that in this data generating model, the outcome Y can depend on treatments A_1 and A_2 and on intermediate outcome L_1 only through the regrets μ_1 and μ_2 if the SNMM specified by the blip functions γ_1 and γ_2 is to hold. Also note that in this setup the true treatment-free outcome model is piecewise linear [26]. In this basic framework, the parameters $\psi, \alpha, \beta, \sigma$ are varied to create the scenarios listed in Table 1. Scenario 1a is a base scenario relative to scenarios 1b–1e. Compared to scenario 1a, scenario 1b almost completely removes the noise in the final outcome model, thus allowing more of the variability in the outcome Y to be explained

Baseline covariate

- $L_0 \sim N(\beta_0, \sigma_0^2)$

Intermediate response

- $L_1|L_0, A_1 \sim N(\beta_{10} + \beta_{11}L_0 + \beta_{12}A_1, \sigma_1^2)$

Treatment propensity

- $\text{logit } P(A_1 = 1|L_0) = \alpha_{10} + \alpha_{11}L_0$
- $\text{logit } P(A_2 = 1|L_0, A_1, L_1) = \alpha_{20} + \alpha_{21}L_0 + \alpha_{22}A_1 + \alpha_{23}L_1$

Blips

- $\gamma_1(a_1, h_1) = a_1(\psi_{10} + \psi_{11}l_0)$
- $\gamma_2(a_2, h_2) = a_2(\psi_{20} + \psi_{21}l_1)$

Regrets

- $\mu_j(a_j, h_j) = \max_{a'_j} \{\gamma_j(a'_j, h_j)\} - \gamma_j(a_j, h_j), \quad j = 1, 2$

Outcome

- $Y|A_2, H_2 \sim N(\beta_{20} + \beta_{21}L_0, \sigma_2^2) - \mu_1(A_1, H_1) - \mu_2(A_2, H_2)$

Figure 2: The basic setup used for simulating data from an SNMM with two treatment intervals.

Table 1: Simulation scenarios.

Scenario	Parameter settings
1a	$\psi_1 = (17, -3.4); \psi_2 = (42, -2.8);$ $\alpha_1 = (0.1, -0.03); \alpha_2 = (0.5, 0, 0, -0.04);$ $\beta_0 = 0; \beta_1 = (-4, 1.25, 6); \beta_2 = (30, 1.6);$ $\sigma_0 = 14; \sigma_1 = 12; \sigma_2 = 30.$
1b	Like scenario 1a but with $\sigma_2 = 1.$
1c	Like scenario 1a but with $\sigma_1 = 3.$
1d	Like scenario 1a but with $\sigma_0 = 28.$
1e	Like scenario 1a but with $\alpha_1 = (0.1, 0.01)$ and $\alpha_2 = (0.5, 0, 0, 0.01).$
1f	Like scenario 1e but with $\sigma_0 = 28.$

by the model. In scenario 1c, the correlation between L_0 and L_1 is increased by decreasing the random variation of L_1 . In scenario 1d, the correlation between L_0 and L_1 is increased by increasing the variability of L_0 . In scenario 1e, the effect of the past observations on treatment assignment is diminished, thus making treatment assignment more random across the levels of L_0 and L_1 . Scenario 1f combines the characteristics of scenarios 1d and 1e.

Sample sizes $n = 200$ and $n = 1,000$ were considered. For each scenario and sample size, each of the estimators considered was evaluated on 1,000 generated data sets. The main criterion used for comparison of estimation methods was the percent root mean squared error (RMSE), which was calculated as the square-root of the mean of the squared difference between parameter estimate and true value across simulation repetitions, expressed as a percentage of the true value of the parameter. More complete simulation results including percent bias (mean absolute deviation between parameter estimate and true value across simulation repetitions, expressed as a percentage of the true value of the parameter) and simulation coverage of estimated 95% confidence intervals (CI) are presented in the Appendix.

First, simple linear models are considered for the specification of the nuisance treatment-free outcome models, specifically:

$$\eta_1(H_1; \varsigma_1) = \varsigma_{10} + \varsigma_{11}L_0,$$

$$\eta_2(H_2; \varsigma_2) = \varsigma_{20} + \varsigma_{21}L_1 + \varsigma_{22}A_1.$$

Because L_0 and L_1 are highly correlated, we do not include L_0 in the treatment-free outcome model at interval 2.

For PIBW estimators, influence was predicted using flexible spline models. Specifically, we regressed $\mathcal{I}_{ij}$ (PIBW1) or each component of $\hat{\psi}_{j(i)}$ (PIBW2) or the components of $\hat{\psi}_{j(i)}$ after orthogonalization (PIBW3) on: $j = 2$ previous treatment A_1 and cubic B-spline function of L_1 with 4 degrees of freedom; and $j = 1$ cubic B-spline function of L_0 , with 4 degrees of freedom.

The knots for the B-splines were chosen automatically by the `bs()` function in R (an internal knot at the median, and boundary knots at the minimum and maximum by default).

We studied the effect of the tuning parameter λ as follows. We set $\lambda_1 = 1$ and varied λ_2 to achieve different median weights between 0.5 and 1. For each choice of median weight, each estimator was applied to 100 simulated data sets from scenarios 1a and 1f of Table 1. The results for scenario 1f with sample sizes $n = 200$ and $n = 1,000$ are presented in Figures 3 and 4 respectively. We note that the choice of λ involves a compromise in the efficiency with which the different components of ψ_j are estimated. The relationship between median weight and RMSE depends on the method of weight construction (IBW vs. PIBW1 vs. PIBW2 vs. PIBW3), and the treatment interval. The results suggest that for scenario 1f and sample size $n = 200$, PIBW1 outperforms PIBW2 and PIBW3 (which are very similar to each other). While there is no choice of λ that is optimal overall, the results suggest that λ should be chosen to produce a lower median weight for PIBW1 than for PIBW2 and PIBW3. For PIBW1, median weights around 0.7 (or lower) for interval $j = 2$, and between 0.7 and 0.9 for interval $j = 1$, produce good results. For PIBW2 and PIBW3 the median weight should be somewhat greater, between 0.8 and 0.9 at interval $j = 2$, and around 0.95 at interval $j = 1$. At sample size $n = 1,000$ (Figure 4), the main difference is that for ψ_{11} the PIBW2 and PIBW3 estimators no longer outperform the standard estimator. Otherwise, the results are qualitatively similar.

Based on these findings, the tuning parameter was set to achieve a median weight of 0.7 and 0.9 at intervals 2 and 1 respectively for PIBW1, 0.9 and 0.95 at intervals 2 and 1 respectively for PIBW2 and PIBW3, and 0.97 and 0.99 at intervals 2 and 1 respectively for IBW. Then for each scenario 1,000 new simulation repetitions were performed (using a different seed for the random number generator). The simulation results for scenarios 1a–1f are summarized in Tables 2 and 3 for the two sample sizes $n = 200$ and $n = 1,000$ respectively. These Tables show the simulation performance of the weighted estimators compared to standard and efficient estimators, measured by percent RMSE. For easier comparison, the RMSE relative to the standard (unweighted) estimator is also shown in parentheses. More complete simulation results, including bias and coverage of 95% confidence intervals estimated using the sandwich estimator of the variance are given in the Appendix. In terms of bias and coverage, the PIBW estimators performed similarly to the standard estimator across all simulation scenarios; minimal bias was observed, and coverage was close to the nominal level (particularly for $n = 1,000$).

We observe that across these simulation scenarios, there was either no real benefit or a modest reduction in variability from using the PIBW estimators versus the standard unweighted estimator. When the PIBW weighting procedure did appear to be beneficial, most of the benefit was observed at the second interval (scenario 1b is an exception), and in the estimation of ψ_{21} in particular. This can be explained by the fact that at interval 2 the model used to predict influence is richer (there are more predictors available). The largest reduction in variance was observed in scenarios 1b and 1f. In very few cases did the PIBW estimators perform worse than the standard estimator, PIBW 1 did not perform well in scenario 1d at interval 2, and PIBW2 and PIBW3 did not perform well in scenarios 1a, 1c and 1e at interval 1. The IBW estimator generally performed worst (though not always) due to incurred bias. Increasing the sample size

Figure 3: Percent RMSE (based on 100 replications) vs. median weight for scenario 1f with sample size $n = 200$ and misspecified linear treatment-free outcome model.

from $n = 200$ to $n = 1,000$ did not substantially change the results. All estimators improved (except IBW due to bias), but their performance relative to one another was similar.

We can gain some insight into why the PIBW estimators performed better than the standard estimator in some scenarios and not in others by looking at the residual plots obtained from fitting the standard estimator. Figure 5 shows plots of residuals at intervals 1 and 2 versus L_0 and L_1 respectively and versus fitted values at the corresponding interval for scenarios 1a, 1c and 1f (a typical data set of size $n = 1,000$ was generated for each scenario). We observe that for these scenarios, the decrease in variability from using a PIBW estimator (scenario 1f shows the most improvement, followed by 1c and then 1a) correlates with how pronounced is the lack of fit of the specified treatment-free outcome model to the true treatment-free outcome model. This is particularly evident in the plots of e_2 vs. L_1 (third column).

So far, simple linear models were considered for the specification of the nuisance treatment-free outcome models. Because the true treatment-free outcome model in this simulation setup is piecewise linear (except in the null scenario 2, where it is linear), the linear treatment-free outcome models may fit quite poorly. In fact, for scenario 1f where the proposed estimators performed best, simple residual plots as

Figure 4: Percent RMSE (based on 100 replications) vs. median weight for scenario 1f with sample size $n = 1,000$ and misspecified linear treatment-free outcome model.

described by Rich et al. [14] would have revealed the poor fit (Figure 5). Thus, we repeated the previous simulations using a more flexible treatment-free outcome model, specifically:

$$\eta_1(H_1; \varsigma_1) = \varsigma_{10} + B_1(L_0; \varsigma_{11})$$

$$\eta_2(H_2; \varsigma_2) = \varsigma_{20} + B_2(L_1; \varsigma_{21}) + \varsigma_{22}A_1,$$

where B_1 and B_2 are cubic B-splines with 4 degrees of freedom, with knots chosen automatically by the `bs()` function in R. Note that this is the same set of predictors used for prediction of influence in the case of the PIBW estimators. The analogues to Tables 2 and 3 when the spline treatment-free outcome model is used are given in the Appendix. Compared to the case of the linear treatment-free outcome model, the variability of the PIBW estimators relative to the standard estimator is greatly increased. With the spline model able to capture more of the features of the true piecewise linear treatment-free outcome model, there is less potential for identifying high influence points that are destabilizing to the standard estimator. It is expected that as the working treatment-free outcome model approaches to the true treatment-free outcome model (and hence the g-estimator approaches semi-parametric efficiency), the benefit of using a weighted estimator diminishes.

Table 2: Performance of weighted estimators compared to standard and efficient estimators for scenarios 1a–1f measured by percent RMSE and RMSE relative to the standard estimator (in parentheses) based on 1,000 data sets of size $n = 200$. The misspecified treatment-free outcome model is linear.

Parameter (truth)	Estimator	Scenario					
		1a	1b	1c	1d	1e	1f
ψ_{10} (17)	Standard	28.43	12.52	29.66	35.93	28.50	37.44
	Efficient	26.55	0.87	26.65	27.00	25.49	25.94
	IBW	28.96	13.11	30.48	37.04	28.71	36.18
	PIBW1	28.52	10.95	29.28	35.32	27.81	34.87
	PIBW2	29.22	10.34	29.67	34.91	28.69	35.15
ψ_{11} (−3.4)	PIBW3	29.22	10.42	29.67	34.95	28.72	35.20
	Standard	10.55	6.07	11.52	8.05	10.96	7.75
	Efficient	9.46	0.32	10.12	5.67	9.34	4.88
	IBW	11.19	6.38	12.07	9.72	10.79	7.28
	PIBW1	10.61	4.89	11.59	7.92	10.34	6.78
ψ_{20} (42)	PIBW2	11.24	4.35	12.09	7.57	11.25	6.98
	PIBW3	11.22	4.36	12.12	7.59	11.28	6.99
	Standard	16.54	12.13	14.52	21.87	18.10	30.95
	Efficient	11.50	0.39	11.16	12.69	10.42	10.95
	IBW	16.74	11.42	14.68	21.48	18.32	31.25
ψ_{21} (−2.8)	PIBW1	16.04	11.00	13.91	23.43	15.91	25.01
	PIBW2	16.33	11.46	14.21	21.42	16.87	25.70
	PIBW3	16.37	11.47	14.24	21.45	16.91	25.80
	Standard	15.68	13.41	17.34	14.16	15.36	15.63
	Efficient	9.95	0.32	11.11	7.07	7.92	4.74
	IBW	20.32	17.86	22.47	22.51	16.00	19.62
	PIBW1	15.43	12.22	16.11	14.88	13.46	13.06
	PIBW2	14.88	10.46	15.58	12.86	13.69	12.88
	PIBW3	14.89	10.48	15.68	12.92	13.69	12.94

Table 3: Performance of weighted estimators compared to standard and efficient estimators for scenarios 1a–1f measured by percent RMSE and RMSE relative to the standard estimator (in parentheses) based on 1,000 data sets of size $n = 1,000$. The misspecified treatment-free outcome model is linear.

Parameter (truth)	Estimator	Scenario					
		1a	1b	1c	1d	1e	1f
ψ_{10} (17)	Standard	12.34	5.19	12.50	16.06	12.68	15.78
	Efficient	11.31	0.39	11.37	12.16	11.62	11.24
	IBW	12.73	5.71	12.87	16.58	12.81	15.26
	PIBW1	12.24	4.51	12.35	15.91	12.55	14.90
	PIBW2	12.42	4.26	12.53	15.63	12.77	14.97
ψ_{11} (−3.4)	PIBW3	12.43	4.30	12.53	15.64	12.78	14.98
	Standard	4.81	2.60	4.86	3.45	4.80	3.08
	Efficient	4.32	0.14	4.32	2.50	4.15	2.08
	IBW	5.72	3.59	5.59	6.59	4.79	2.92
	PIBW1	4.70	1.76	4.72	3.39	4.56	2.68
ψ_{20} (42)	PIBW2	4.95	1.74	5.12	3.23	4.88	2.85
	PIBW3	4.95	1.75	5.12	3.24	4.88	2.85
	Standard	6.86	5.03	6.99	9.72	8.33	13.62
	Efficient	4.98	0.18	5.08	5.51	4.88	4.96
	IBW	6.84	5.24	6.84	10.15	8.67	15.42
ψ_{21} (−2.8)	PIBW1	6.64	4.58	6.51	10.24	7.17	10.88
	PIBW2	6.79	4.72	6.59	9.63	7.63	11.53
	PIBW3	6.79	4.72	6.59	9.65	7.64	11.57
	Standard	7.44	5.60	7.61	6.60	6.52	6.87
	Efficient	4.10	0.13	4.72	2.91	3.28	1.94
	IBW	15.99	15.84	16.34	19.87	8.98	15.42
	PIBW1	6.98	5.11	6.80	7.09	5.76	5.27
	PIBW2	6.68	4.51	6.60	5.62	6.16	5.68
	PIBW3	6.68	4.52	6.60	5.63	6.16	5.67

Figure 5: Residual plots for simulation scenarios 1a, 1c and 1f for the standard (unweighted) g-estimator fit using a misspecified linear treatment-free outcome model. A typical data set of size $n = 1,000$ was generated for each scenario.

5 Discussion

Locally efficient doubly robust g-estimation is robust to misspecification of the treatment-free outcome model so long as the propensity model is specified correctly, but a price for misspecification is paid in efficiency. The weighting schemes proposed in this manuscript are an attempt to counteract this efficiency loss using weights that are adaptively derived to decrease the impact of influential points and thus stabilize the estimator.

The applicability of this methodology depends on how well or how poorly the misspecified treatment-free outcome model approximates the true model, and correspondingly how close the standard estimator is to the efficient estimator. The residual diagnostics given by Rich et al. [14] are a means of assessing the goodness of fit and in some cases may suggest ways of improving the estimator through alternative specifications of the treatment-free outcome model. In some cases, this may obviate the need for the re-weighting approach, but in more complex, high-dimensional settings, the diagnostic techniques may be insufficient to fully address treatment-free outcome model misspecification so the re-weighting approach may prove to be useful. The simulation results presented in this manuscript suggest that in some settings the use of PIBW estimators can result in important decreases in the variability of the estimator (a 15–20% reduction in RMSE or more was observed in some scenarios).

The re-weighting procedure described involves solving the g-estimating equation at interval j , determining the influence of each individual $i = 1, \dots, n$ on the estimate of ψ_j , computing weights based on the (predicted) influence, and re-solving the weighted g-estimating equation. This procedure could be iterated, i.e. determine the influences of each individual on the new weighted estimate, and compute a second set of weights based on these, and so on.

The choice of the tuning parameter λ has an important role to play in determining the efficiency of the weighted estimators and would present a challenge in a real data analysis. The simulation results presented here have demonstrated that the choice is application-specific, and can involve a compromise in the efficiency with which the different components of ψ_j are estimated. These results can provide some guidance for future applications, and further simulations under settings that mimic a specific application may also prove useful. A conservative approach is to err on the side of weights that are too large rather than too small, since in that case the weighted estimators behave more like the standard estimator.

Acknowledgements: This work was carried out when the first author was studying in the Department of Epidemiology, Biostatistics, and Occupational Health of McGill University, and receiving support from the Natural Sciences and Engineering Research Council (NSERC) of Canada. Drs. Moodie and Stephens acknowledge support of NSERC of Canada. Dr. Moodie is supported by a Chercheur-Boursier junior 2 career award from the Fonds de recherche du Quebec-Sante (FRQ-S).

References

1. and Drug Administration, U. S. F. Paving the way for personalized medicine FDA's role in a new era of medical product development, 2013. Available at: <http://www.fda.gov/downloads/scienceresearch/specialtopics/personalizedmedicine/ucm372421.pdf>.
2. Schulte PJ, Tsiatis AA, Laber EB, Davidian M. Q- and A-learning methods for estimating optimal dynamic treatment regimes. *Stat Sci* 2014;29:640–661.
3. Zhang B, Tsiatis A, Davidian M, Zhang M, Laber E. Estimating optimal treatment regimes from a classification perspective. *Stat* 2012a;1:103–14.
4. Zhao Y, Zeng D, Rush A, Kosorok M. Estimating individual treatment rules using outcome weighted learning. *J Am Stat Assoc* 2012;107:1106–18.
5. Arjas E, Saarela O. Optimal dynamic regimes presenting a case for predictive inference. *Int J Biostat* 2010;6: Article 10.
6. Zajonc T. Bayesian inference for dynamic treatment regimes mobility, equity, and efficiency in student tracking. *J Am Stat Assoc* 2012;107:80–92.
7. Murphy SA. Optimal dynamic treatment regimes. *J R Stat Soc Ser B (Stat Methodol)* 2003;65:331–55. Available at: <http://doi.wiley.com/10.1111/1467-9868.00389>.
8. Orellana L, Rotnitzky A, Robins JM. Dynamic regime marginal structural mean models for estimation of optimal dynamic treatment regimes, part I main content. *Int J Biostat* 2010;6. Available at: <http://www.bepress.com/ijb/vol6/iss2/8/>.
9. Robins JM. Optimal structural nested models for optimal sequential decisions. In: DY Lin, PJ Heagerty, editors. *Proceedings of the second Seattle symposium in biostatistics analysis of correlated data*. Springer, 2004:189–326. Available at: <http://www.biostat.harvard.edu/robins/publications/seattlemay04final.pdf>.
10. van der Laan MJ, Petersen ML. Causal effect models for realistic individualized treatment and intention to treat rules. *Int J Biostat* 2007;3. Available at: <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2613338%26tool=pmcentrez%26rendertype=abstract>.
11. Zhang B, Tsiatis A, Laber E, Davidian M. A robust method for estimating optimal treatment regimes. *Biometrics* 2012b;68:1010–18.
12. Tan Z. Bounded, efficient and doubly robust estimation with inverse weighting. *Biometrika* 2010;97:661–82. Available at: <http://biomet.oxfordjournals.org/cgi/doi/10.1093/biomet/asq035>.
13. Rosenbaum PR, Rubin DB. The central role of the propensity score in observational studies for causal effects. *Biometrika* 1983;70:41–55. Available at: <http://biomet.oxfordjournals.org/content/70/1/41.short>.
14. Rich B, Moodie EEM, Stephens DA, Platt RW. Model checking with residuals for g-estimation of optimal dynamic treatment regimes. *Int J Biostat* 2010;6. Available at: <http://www.bepress.com/ijb/vol6/iss2/12/>.
15. Hampel FR, Ronchetti EM, Rousseeuw PJ, Stahel WA. *Robust statistics the approach based on influence functions*. New York: Wiley, 1986. Available at: <http://www.lavoisier.fr/notice/frIWOR363ASRWLKO.html>.
16. Tsiatis AA. *Semiparametric theory and missing data*. New York: Springer, 2006. Available at: http://books.google.ca/books?hl=en&lr=%26id=xqZFt2EMB40C26oi=fnd26pg=PR726dq=tsiatis+semiparametric+missing%26ots=Nce_abJdMv26sig=Bj8z0t1ke6U11qhdlWM5N6OEMRs.
17. Newey WK. Semiparametric efficiency bounds. *J Appl Econ* 1990. Available at: <http://www3.interscience.wiley.com/journal/112749856/abstract>.

18. Joffe MM, Brensinger C. Weighting in instrumental variables and G-estimation. *Stat Med* 2003;22:1285–303. Available at: <http://www.ncbi.nlm.nih.gov/pubmed/12687655>.
19. Hinkley DV. Jackknifing in unbalanced situations. *Technometrics* 1977;19:285–92. Available at: <http://www.jstor.org/stable/1267698>.
20. Krasker WS, Welsch RE. Efficient bounded-influence regression estimation. *J Am Stat Assoc* 1982;77:595–604. Available at: <http://www.jstor.org/stable/2287719>.
21. Cook RD, Weisberg S. Residuals and influence in regression, New York Chapman and Hall, 1982. Available at: <http://purl.umn.edu/37076>.
22. Seber GAF, Lee AJ. Linear regression analysis, 2nd ed. Hoboken, NJ: Wiley, 2003.
23. Belsley DA, Kuh E, Welsch RE. Regression diagnostics identifying influential data and sources of collinearity. New York: Wiley, 1980.
24. Preisser JS, Qaqish BF. Deletion diagnostics for generalised estimating equations. *Biometrika* 1996;83: 551–62.
25. Cole SR, Hernán MA. Constructing inverse probability weights for marginal structural models, *Am J Epidemiol* 2008;168, 656–64, Available at: <http://aje.oxfordjournals.org/cgi/content/abstract/kwn164v1>.
26. Moodie EEM, Richardson TS, Stephens DA. Demystifying optimal dynamic treatment regimes. *Biometrics* 2007; 63:447–55.

Appendix

Table 4: Simulation results for scenario 1a performance of weighted estimators compared to standard and efficient estimators.

Parameter (truth)	Estimator	$n = 200$			$n = 1,000$		
		Percent bias (%)	Percent RMSE (%)	Coverage of 95% CI (%)	Percent bias (%)	Percent RMSE (%)	Coverage of 95% CI (%)
ψ_{11} (17)	Standard	−1.62	28.3	94.9	0.26	12.4	93.7
	Efficient	−0.86	26.3	93.5	0.62	11.3	94.7
	IBW	−1.44	28.8	93.7	0.64	12.8	92.4
	PIBW1	−1.82	28.4	94.6	0.32	12.3	94.3
	PIBW2	−2.10	29.0	93.9	0.24	12.5	94.0
	PIBW3	−2.09	29.0	93.9	0.24	12.5	93.9
ψ_{11} (−3.4)	Standard	1.73	10.9	93.5	0.43	4.9	94.3
	Efficient	0.38	9.6	92.8	0.25	4.3	95.4
	IBW	3.75	11.5	89.5	3.04	5.8	84.8
	PIBW1	2.55	10.9	93.5	0.79	4.7	94.5
	PIBW2	1.67	11.6	92.8	0.45	5.0	94.6
	PIBW3	1.67	11.6	93.1	0.45	5.0	94.6
ψ_{20} (42)	Standard	−0.10	16.5	92.6	−0.28	6.9	96.0
	Efficient	0.14	11.5	93.1	−0.04	5.0	95.3
	IBW	−0.85	16.7	86.7	−0.85	7.0	89.6
	PIBW1	0.17	16.1	92.6	−0.16	6.7	95.8
	PIBW2	−0.63	16.3	93.0	−0.43	6.9	95.2
	PIBW3	−0.70	16.3	92.8	−0.45	6.9	95.2
ψ_{21} (−2.8)	Standard	1.27	15.9	91.6	0.94	7.5	91.8
	Efficient	−0.02	9.9	88.4	0.24	4.2	94.3
	IBW	14.18	20.5	62.0	14.64	16.1	20.2
	PIBW1	3.00	15.6	89.8	1.53	7.1	91.2
	PIBW2	1.70	15.1	91.2	0.75	6.8	93.5
	PIBW3	1.91	15.1	91.4	0.79	6.8	93.5

Table 5: Simulation results for scenario 1b performance of weighted estimators compared to standard and efficient estimators.

Parameter (truth)	Estimator	$n = 200$			$n = 1,000$		
		Percent bias (%)	Percent RMSE (%)	Coverage of 95% CI (%)	Percent bias (%)	Percent RMSE (%)	Coverage of 95% CI (%)
ψ_{11} (17)	Standard	0.40	12.5	90.7	0.07	5.12	93.7
	Efficient	0.03	0.9	93.5	-0.02	0.39	94.7
	IBW	0.67	13.2	90.0	0.45	5.59	92.3
	PIBW1	0.20	11.0	92.0	-0.02	4.43	94.3
	PIBW2	-0.05	10.4	92.9	-0.06	4.17	95.5
	PIBW3	-0.13	10.4	92.8	-0.08	4.21	95.5
ψ_{11} (-3.4)	Standard	1.75	5.9	87.4	0.33	2.55	90.4
	Efficient	-0.01	0.3	92.3	0.00	0.14	94.4
	IBW	3.27	6.3	80.7	2.56	3.55	66.6
	PIBW1	2.29	4.8	86.6	0.35	1.75	92.6
	PIBW2	1.22	4.3	93.5	0.06	1.75	94.3
	PIBW3	1.21	4.3	93.4	0.04	1.76	94.5
ψ_{20} (42)	Standard	-0.25	12.2	91.9	-0.21	5.12	95.5
	Efficient	-0.02	0.4	92.6	-0.01	0.17	93.6
	IBW	-2.24	11.6	85.8	-2.58	5.33	84.8
	PIBW1	0.19	11.1	91.3	-0.03	4.63	94.7
	PIBW2	-0.72	11.6	92.3	-0.30	4.76	94.6
	PIBW3	-0.78	11.6	92.0	-0.31	4.76	94.6
ψ_{21} (-2.8)	Standard	0.83	13.3	88.7	-0.27	5.65	93.8
	Efficient	0.00	0.3	90.5	0.00	0.14	94.9
	IBW	14.89	17.9	45.8	15.18	15.79	2.7
	PIBW1	3.15	12.1	88.1	0.48	5.16	93.6
	PIBW2	1.52	10.4	92.2	-0.09	4.52	95.8
	PIBW3	1.79	10.4	91.8	-0.02	4.52	95.8

Table 6: Simulation results for scenario 1c performance of weighted estimators compared to standard and efficient estimators.

Parameter (truth)	Estimator	$n = 200$			$n = 1,000$		
		Percent bias (%)	Percent RMSE (%)	Coverage of 95% CI (%)	Percent bias (%)	Percent RMSE (%)	Coverage of 95% CI (%)
ψ_{11} (17)	Standard	1.03	29.4	93.2	-0.284	12.5	94.7
	Efficient	0.03	26.4	93.8	-0.144	11.4	94.6
	IBW	1.98	30.2	92.1	0.581	12.8	93.3
	PIBW1	0.58	29.1	93.4	-0.275	12.4	94.5
	PIBW2	0.66	29.5	93.8	-0.261	12.5	94.9
	PIBW3	0.66	29.5	93.9	-0.259	12.5	94.8
ψ_{11} (-3.4)	Standard	1.34	11.6	91.2	0.311	4.8	94.7
	Efficient	-0.37	10.1	91.5	-0.107	4.3	94.8
	IBW	2.84	12.1	87.6	2.450	5.5	86.5
	PIBW1	1.78	11.5	90.7	0.536	4.7	94.9
	PIBW2	1.07	12.1	91.3	0.117	5.1	94.2
	PIBW3	1.08	12.1	91.2	0.122	5.1	94.3
ψ_{20} (42)	Standard	0.05	14.6	94.9	0.410	6.9	93.9
	Efficient	-0.41	11.2	93.9	0.252	5.0	94.5
	IBW	-0.04	14.7	89.0	0.333	6.8	88.1
	PIBW1	0.19	13.9	93.8	0.422	6.4	94.1
	PIBW2	-0.22	14.3	93.2	0.280	6.5	93.7
	PIBW3	-0.23	14.3	93.4	0.276	6.5	93.8
ψ_{21} (-2.8)	Standard	1.13	17.2	92.2	-0.003	7.7	95.1
	Efficient	0.27	11.0	90.7	-0.329	4.7	94.4
	IBW	15.44	22.5	65.3	14.832	16.4	25.8
	PIBW1	3.24	16.1	91.2	0.622	6.8	94.3
	PIBW2	1.80	15.7	92.4	-0.145	6.6	95.9
	PIBW3	2.07	15.8	92.1	-0.097	6.6	96.0

Table 7: Simulation results for scenario 1d performance of weighted estimators compared to standard and efficient estimators.

Parameter (truth)	Estimator	$n = 200$			$n = 1,000$		
		Percent bias (%)	Percent RMSE (%)	Coverage of 95% CI (%)	Percent bias (%)	Percent RMSE (%)	Coverage of 95% CI (%)
ψ_{11} (17)	Standard	-2.19	35.6	94.0	0.44	15.8	94.4
	Efficient	-1.84	27.1	94.1	0.09	12.0	94.6
	IBW	-1.70	36.8	91.8	0.34	16.3	92.4
	PIBW1	-1.80	35.0	93.9	0.48	15.6	94.2
	PIBW2	-2.47	34.9	93.9	0.21	15.4	93.3
	PIBW3	-2.48	35.0	94.1	0.20	15.4	93.3
ψ_{11} (-3.4)	Standard	1.64	8.1	90.6	0.48	3.5	93.8
	Efficient	0.05	5.7	92.1	0.12	2.5	94.4
	IBW	6.08	9.7	76.7	5.55	6.5	49.8
	PIBW1	2.92	7.9	89.5	0.78	3.4	93.9
	PIBW2	1.80	7.6	93.4	0.37	3.2	95.3
	PIBW3	1.88	7.6	93.1	0.38	3.2	95.2
ψ_{20} (42)	Standard	1.50	21.9	93.2	-0.54	9.6	93.9
	Efficient	1.25	12.8	91.5	-0.18	5.6	94.9
	IBW	-2.11	21.7	87.2	-3.74	10.0	86.0
	PIBW1	2.39	23.3	91.7	-0.30	10.1	94.0
	PIBW2	1.14	21.4	92.2	-0.91	9.5	93.7
	PIBW3	0.87	21.4	92.2	-0.99	9.5	93.5
ψ_{21} (-2.8)	Standard	1.11	14.4	90.1	0.23	6.7	94.0
	Efficient	-0.04	7.1	92.3	0.01	3.0	95.7
	IBW	19.05	22.9	40.5	18.94	19.7	2.4
	PIBW1	3.07	15.1	87.6	0.55	7.1	92.5
	PIBW2	2.56	13.2	90.1	0.51	5.7	94.1
	PIBW3	3.03	13.2	89.6	0.64	5.7	94.1

Table 8: Simulation results for scenario 1e performance of weighted estimators compared to standard and efficient estimators.

Parameter (truth)	Estimator	$n = 200$			$n = 1,000$		
		Percent bias (%)	Percent RMSE (%)	Coverage of 95% CI (%)	Percent bias (%)	Percent RMSE (%)	Coverage of 95% CI (%)
ψ_{11} (17)	Standard	1.50	28.9	93.8	0.251	12.8	93.6
	Efficient	1.07	25.8	93.3	0.213	11.8	93.4
	IBW	1.89	29.1	92.2	0.801	12.9	92.4
	PIBW1	1.52	28.3	93.8	0.159	12.7	94.0
	PIBW2	1.43	29.1	94.2	0.133	12.9	93.7
	PIBW3	1.46	29.2	94.0	0.136	12.9	93.7
ψ_{11} (-3.4)	Standard	1.81	10.9	92.6	0.386	4.7	94.0
	Efficient	0.05	9.3	91.6	0.028	4.1	93.3
	IBW	1.89	10.8	90.2	0.446	4.8	91.0
	PIBW1	0.74	10.4	93.3	0.086	4.6	93.1
	PIBW2	1.21	11.3	92.9	0.302	4.9	93.3
	PIBW3	1.21	11.3	93.0	0.299	4.9	93.3
ψ_{20} (42)	Standard	-1.03	18.1	92.7	-0.498	8.4	93.0
	Efficient	-0.06	10.5	94.6	-0.299	4.8	95.0
	IBW	-3.14	18.3	87.0	-2.924	8.7	85.3
	PIBW1	-1.08	15.9	93.7	-0.431	7.2	93.0
	PIBW2	-0.47	16.7	93.7	-0.242	7.6	93.0
	PIBW3	-0.46	16.8	93.8	-0.244	7.6	93.1
ψ_{21} (-2.8)	Standard	-0.90	15.2	91.6	-0.215	6.6	94.9
	Efficient	0.21	7.9	91.0	-0.128	3.3	94.4
	IBW	-6.25	15.9	77.8	-6.363	9.1	66.5
	PIBW1	-1.23	13.4	92.4	-0.479	5.7	95.5
	PIBW2	-0.10	13.8	93.5	-0.118	6.2	94.7
	PIBW3	-0.15	13.8	93.3	-0.127	6.2	94.7

Table 9: Simulation results for scenario 1f performance of weighted estimators compared to standard and efficient estimators.

Parameter (truth)	Estimator	$n = 200$			$n = 1,000$		
		Percent bias (%)	Percent RMSE (%)	Coverage of 95% CI (%)	Percent bias (%)	Percent RMSE (%)	Coverage of 95% CI (%)
ψ_{11} (17)	Standard	0.08	37.2	93.0	-0.33	15.9	94.5
	Efficient	-0.62	25.8	93.6	-0.04	11.4	95.5
	IBW	0.82	35.8	91.2	1.43	15.3	91.7
	PIBW1	0.11	34.8	92.9	-0.12	15.0	94.0
	PIBW2	-0.38	34.9	93.2	-0.21	15.1	94.1
	PIBW3	-0.33	34.9	93.3	-0.18	15.2	94.0
ψ_{11} (-3.4)	Standard	1.64	7.8	89.0	0.20	3.1	92.9
	Efficient	0.24	4.9	91.9	-0.11	2.1	94.5
	IBW	1.49	7.4	84.5	-0.13	3.0	88.5
	PIBW1	0.46	6.8	89.2	-0.19	2.7	94.7
	PIBW2	1.08	7.1	91.2	0.13	2.9	95.1
	PIBW3	1.07	7.1	91.3	0.12	2.9	95.0
ψ_{20} (42)	Standard	-1.31	30.8	92.9	0.45	13.6	94.7
	Efficient	-0.48	11.0	92.4	0.00	4.9	94.1
	IBW	-8.65	31.1	86.5	-8.04	15.5	84.0
	PIBW1	-2.49	25.2	93.2	0.01	10.8	94.9
	PIBW2	-1.04	25.8	93.8	0.35	11.6	95.1
	PIBW3	-1.04	25.9	93.5	0.34	11.6	95.1
ψ_{21} (-2.8)	Standard	-1.25	15.8	90.5	-0.24	6.9	94.5
	Efficient	-0.13	4.7	91.4	-0.03	1.9	95.5
	IBW	-13.77	20.0	64.9	-13.91	15.3	21.8
	PIBW1	-2.81	13.2	90.8	-0.59	5.3	94.7
	PIBW2	-0.59	13.1	94.3	0.11	5.7	96.3
	PIBW3	-0.73	13.2	93.9	0.08	5.7	96.0

Table 10: Simulation results for scenario 2 performance of weighted estimators compared to standard and efficient estimators.

Parameter (truth)	Estimator	$n = 200$			$n = 1,000$		
		Percent bias (%)	Percent RMSE (%)	Coverage of 95% CI (%)	Percent bias (%)	Percent RMSE (%)	Coverage of 95% CI (%)
ψ_{11} (0)	Standard	-0.093	4.34	94.0	-0.102	1.98	94.3
	Efficient	0.127	3.42	98.0	-0.061	1.60	96.9
	IBW	-0.147	4.40	92.4	-0.107	1.99	92.5
	PIBW1	-0.110	4.36	94.2	-0.098	1.98	94.2
	PIBW2	-0.134	4.51	93.6	-0.099	2.02	94.3
	PIBW3	-0.135	4.51	93.6	-0.100	2.02	94.4
ψ_{11} (0)	Standard	0.013	0.33	93.5	0.020	0.14	94.9
	Efficient	0.032	0.41	89.4	0.035	0.17	91.3
	IBW	0.006	0.33	90.8	0.018	0.14	91.1
	PIBW1	0.012	0.34	93.3	0.021	0.15	94.4
	PIBW2	0.013	0.37	93.4	0.020	0.16	93.9
	PIBW3	0.013	0.37	93.3	0.020	0.16	93.9
ψ_{20} (0)	Standard	0.131	5.02	94.9	-0.072	2.19	96.0
	Efficient	0.129	4.06	98.3	-0.007	1.79	97.5
	IBW	0.145	5.21	86.6	-0.086	2.28	88.4
	PIBW1	0.174	5.12	94.7	-0.070	2.21	95.9
	PIBW2	0.109	5.15	94.5	-0.100	2.25	95.7
	PIBW3	0.110	5.15	94.3	-0.101	2.25	95.8
ψ_{21} (0)	Standard	-0.020	0.28	93.3	-0.001	0.12	95.8
	Efficient	-0.057	0.33	86.7	-0.005	0.13	92.6
	IBW	-0.015	0.28	82.3	0.000	0.12	82.4
	PIBW1	-0.018	0.30	92.0	-0.001	0.12	94.6
	PIBW2	-0.019	0.31	92.8	0.000	0.13	94.9
	PIBW3	-0.019	0.31	93.0	0.000	0.13	94.7

Table 11: Simulation results for scenario 3 performance of weighted estimators compared to standard and efficient estimators.

Parameter (truth)	Estimator	$n = 200$			$n = 1,000$		
		Percent bias (%)	Percent RMSE (%)	Coverage of 95% CI (%)	Percent bias (%)	Percent RMSE (%)	Coverage of 95% CI (%)
$\psi_{11} (0)$	Standard	-2.44	45.9	93.3	-0.08	20.4	94.5
	Efficient	2.11	20.3	96.2	0.29	8.6	96.1
	IBW	-3.92	37.5	81.8	-2.34	15.5	81.4
	PIBW1	-0.81	35.1	94.9	-0.37	14.2	96.5
	PIBW2	0.16	34.6	95.5	-0.29	14.4	96.3
	PIBW3	0.39	35.3	95.6	-0.36	14.9	96.0
$\psi_{11} (-1)$	Standard	1.43	25.1	93.5	0.03	11.1	94.5
	Efficient	-0.90	11.9	97.7	-0.14	5.1	97.0
	IBW	9.65	24.7	75.7	7.29	11.4	64.8
	PIBW1	0.44	19.8	93.6	0.17	7.9	95.6
	PIBW2	-0.05	19.7	94.2	0.13	8.1	96.2
	PIBW3	-0.05	20.1	94.3	0.18	8.3	96.3
$\psi_{20} (720)$	Standard	-1.20	21.5	90.5	-0.21	9.8	94.3
	Efficient	-0.77	7.9	93.8	-0.03	3.2	94.4
	IBW	-6.02	19.5	75.1	-5.62	9.9	64.9
	PIBW1	-1.29	15.7	92.4	-0.03	6.5	94.7
	PIBW2	-0.87	15.1	93.4	0.21	6.6	95.3
	PIBW3	-1.03	15.4	93.8	0.25	6.8	95.2
$\psi_{21} (-2)$	Standard	0.71	13.0	90.4	0.14	5.9	94.0
	Efficient	0.39	5.0	93.0	0.01	2.0	94.6
	IBW	0.28	12.1	76.5	0.59	5.1	75.0
	PIBW1	0.80	9.6	92.0	0.03	3.9	94.5
	PIBW2	0.57	9.4	92.8	-0.10	4.0	94.4
	PIBW3	0.63	9.5	92.9	-0.13	4.1	94.5