Romain Neugebauer*, Julie A. Schmittdiel and Mark J. van der Laan

A Case Study of the Impact of Data-Adaptive Versus Model-Based Estimation of the Propensity Scores on Causal Inferences from Three Inverse Probability Weighting Estimators

DOI 10.1515/ijb-2015-0028

Abstract:

Objective: Consistent estimation of causal effects with inverse probability weighting estimators is known to rely on consistent estimation of propensity scores. To alleviate the bias expected from incorrect model specification for these nuisance parameters in observational studies, data-adaptive estimation and in particular an ensemble learning approach known as Super Learning has been proposed as an alternative to the common practice of estimation based on arbitrary model specification. While the theoretical arguments against the use of the latter haphazard estimation strategy are evident, the extent to which data-adaptive estimation can improve inferences in practice is not. Some practitioners may view bias concerns over arbitrary parametric assumptions as academic considerations that are inconsequential in practice. They may also be wary of data-adaptive estimation of the propensity scores for fear of greatly increasing estimation variability due to extreme weight values. With this report, we aim to contribute to the understanding of the potential practical consequences of the choice of estimation strategy for the propensity scores in real-world comparative effectiveness research.

Method: We implement secondary analyses of Electronic Health Record data from a large cohort of type 2 diabetes patients to evaluate the effects of four adaptive treatment intensification strategies for glucose control (dynamic treatment regimens) on subsequent development or progression of urinary albumin excretion. Three Inverse Probability Weighting estimators are implemented using both model-based and data-adaptive estimation strategies for the propensity scores. Their practical performances for proper confounding and selection bias adjustment are compared and evaluated against results from previous randomized experiments.

Conclusion: Results suggest both potential reduction in bias and increase in efficiency at the cost of an increase in computing time when using Super Learning to implement Inverse Probability Weighting estimators to draw causal inferences.

Keywords: inverse probability weighting, super learning, propensity score, data-adaptive estimation, marginal structural model

1 Introduction

Consistent estimation of causal effects with Inverse Probability Weighting (IPW) estimators is known to rely on consistent estimation of propensity scores (PS) [1]. To alleviate the bias expected from incorrect parametric model specification for these nuisance parameters in observational studies, various data-adaptive PS estimation strategies (also known as machine learning algorithms) have been proposed [2–9] as alternatives to the common practice of PS estimation based on an arbitrarily specified parametric model (e. g., logistic model fitted by maximum likelihood) [10–13]. More recently, an ensemble learning approach known as Super Learning (SL) [14] was proposed as a data-adaptive strategy for PS estimation [15–19] to avoid reliance on a single arbitrarily chosen machine learning algorithm.

*Corresponding author: Romain Neugebauer, Division of Research, Kaiser Permanente Northern California, Oakland, CA, USA, E-mail: romain.s.neugebauer@kp.org

Julie A. Schmittdiel, Division of Research, Kaiser Permanente Northern California, Oakland, CA, USA

Mark J. van der Laan, Division of Biostatistics, School of Public Health, University of California, Berkeley, CA, USA

While the theoretical arguments against the use of arbitrarily specified parametric models that do not encode true subject-matter knowledge about the functional form of the PS are evident, the extent to which data-adaptive PS estimation such as SL can improve inferences in practice is not. Despite results from several simulation studies that have addressed this topic previously, some practitioners may view bias concerns over arbitrary parametric assumptions as academic considerations that are inconsequential in practice. They may also be wary of data-adaptive estimation of the PS for fear of greatly increasing estimation variability due to extreme weight values.

With this report, we aim to contribute to the existing literature for understanding the potential practical consequences of the choice of estimation strategy for the PS. This effort is based on secondary analyses from a real-world comparative effectiveness research (CER) study in type 2 diabetes. Three IPW estimators are implemented using both model-based and data-adaptive estimation strategies for the PS. Their practical performances for proper confounding and selection bias adjustment are compared and evaluated against results from previous randomized experiments.

In Section 2, we introduce the CER observational study, its research question, results from prior randomized experiments and we introduce the notation for describing the various estimation approaches detailed in Section 3. Results from three candidate IPW estimators of the same causal estimand based on four candidate estimation approaches for the PS are presented in Section 4. We end with a discussion of these results in Section 5.

2 A case study in comparative effectiveness research in diabetes

In this section, we describe the CER question, answers from previous randomized studies, and the observational study on which is based the evaluation in this report. We also introduce formal notation for representing the data structure and the parameter of interest in this analysis. We note that the presentation in this section appeared in prior work that relied on the same CER study [20, 21].

2.1 Research question and previous trial results

It has long been hypothesized that aggressive glycemic control is an effective strategy to reduce the occurrence of common and devastating microvascular and macrovascular complications of type 2 diabetes (T2DM). A major goal of clinical care of T2DM is minimization of such complications through a variety of pharmacological treatments and interventions to achieve recommended levels of glucose control. The progressive nature of T2DM results in frequent revisiting of treatment decisions for many patients as glycemic control deteriorates. Widely accepted stepwise guidelines start treatment with metformin, then add a secretagogue if control is not reached or deteriorates. Insulin or (less frequently) a third oral agent is the next step. Thus, it is common for T2DM patients to be on multiple glucose-lowering medications.

Previous recommendations specify target hemoglobin A1c of $<7\%$ for most patients [22, 23]. However, evidence supporting the effectiveness of a blanket recommendation is inconsistent across several outcomes [24–31], especially when intensive anti-diabetic therapy is required. In the analyses of this report, we aim to evaluate the impact of progressively more aggressive glucose-lowering strategies on the development or progression of albuminuria, a microvascular complication in T2DM.

In the ACCORD and ADVANCE clinical trials published from 2008 to 2010 [32–34], intensive glucose-lowering strategies using multiple classes of glucose-lowering agents succeeded in reducing A1c levels substantially. In the ADVANCE trial, the more intensive therapy arm aimed to reach an A1c level $<6.5\%$ and achieved a mean A1c level of 6.5% , compared to a mean level of 7.3% in the control arm. In the ACCORD trial, the more intensive arm aimed for an A1c of $<6\%$, and achieved a mean A1c of 6.4% (vs. 7.5% in controls). There is substantial data from both trials [35, 36] to support the hypothesis [37, 38] that, in

general, those with T2DM who are treated to lower A1c levels may have lower rates of onset and progression of albuminuria (e. g., HR: 0.79, 0.66–0.93 in ADVANCE).

2.2 An observational, multi-center, retrospective, cohort study

The effects of intensive treatment remain uncertain, and the optimal target levels of A1c for balancing benefits and risks of therapy are not clearly defined. For these reasons, using data from the electronic health records (EHR) from patients of seven sites of the HMO Research Network [39], a large retrospective cohort study of adults with T2DM was conducted to evaluate the impact of progressively more aggressive glucose-lowering strategies on several clinical outcomes. To properly account for time-dependent confounding and informative selection bias, a parametric dynamic marginal structural model (MSM) [40–45] was fitted using IPW estimation [40, 46, 47] for the purpose of contrasting cumulative risks under the following four treatment intensification (TI) strategies denoted by d_θ : *patient initiates TI at the first time (no grace periods allowed [45]) her first observed A1c level reaches or drifts above $\theta\%$ and patient remains on the intensified therapy thereafter* with $\theta = 7, 7.5, 8$, or 8.5 .

Details of the study design, analytic approach, and results are described elsewhere [48, 49]. In brief, results were consistent with that of ACCORD and ADVANCE and imply that the pattern of results in these trials are applicable to a large population of adults with T2DM treated in routine clinical settings. In particular, findings from the observational study confirmed the benefit of tight glycemic control with respect to the development or progression of albuminuria.

Here, we report on results from secondary analyses of the same observational data using not only alternate PS estimation strategies but also alternate IPW estimators based on a nonparametric dynamic MSM. These secondary analyses aim to contrast the same four counterfactual survival curves indexed by the same TI strategies described above for the purpose of evaluating the impact of four alternate PS estimation strategies on three alternate IPW estimators. We now formally describe the observational data and the causal parameter of interest before detailing the various estimation approaches considered to evaluate it.

2.3 Data, parameter of interest, and assumptions

The observed data on each patient in the cohort consist of measurements on exposure, outcome, and confounding variables made at 90-day intervals between study entry and until each patient's end of follow-up. The time (expressed in units of 90 days) when the patient's follow-up ends is denoted by $\tilde{T}$ and is defined as the earliest of the time to failure, i. e., albuminuria development or progression, denoted by T or the time to a right-censoring event denoted by C . When a patient is right-censored, i. e., $\tilde{T} = C$, the type of right-censoring event experienced by the patient is recorded and denoted by Γ with possible values 1, 2, or 3 to represent end of follow-up by administrative end of study, disenrollment from the health plan, or death respectively. For patients with normoalbuminuria at study entry, i. e., microalbumin-to-creatinine ratio (ACR) < 30 , we defined failure as an ACR measurement indicating either microalbuminuria (ACR 30–300) or macroalbuminuria (ACR > 300). For patients with microalbuminuria at study entry, we defined failure as an ACR measurement indicating macroalbuminuria. We thus excluded patients with a baseline ACR measurement missing (5884) or indicating macroalbuminuria (1608), which yielded the sample size $n = 51,179$. The indicator that the end of follow-up is due to the occurrence of a failure event is denoted by Δ , i. e., $\Delta = 1$ implies that $\tilde{T} = T$ and $\Delta = 0$ implies that $\tilde{T} = C$. At each time point $t = 0, \dots, \tilde{T}$, the patient's exposure to an intensified diabetes treatment is represented by the binary variable $A_1(t)$, and the indicator of the patient's right-censored status at time t is denoted by $A_2(t)$. We thus have $A_2(t) = 0$ for $t = 0, \dots, \tilde{T} - 1$ when $\tilde{T} \geq 1$ and $A_2(\tilde{T}) = 1 - \Delta$. The combination $A(t) = (A_1(t), A_2(t))$ is referred to as the action at time t . At each time point $t = 0, \dots, \tilde{T}$, covariates such as A1c measurements (others are listed in Table 1 of previous published work [21]) are denoted by the multi-dimensional variable $L(t)$ and defined from measurements that occur before the action at time t , $A(t)$, or

are otherwise assumed not to be affected by the actions at time t or thereafter, $(A(t), A(t+1), \dots)$. In addition, the covariates at time t include an outcome measurement denoted by $Y(t)$, i. e., $Y(t) \in L(t)$ for $t=0, \dots, \tilde{T}$. For each time point $t=0, \dots, \tilde{T}+1$, the outcome is the indicator of past failure, i. e., $Y(t)=I(T \leq t-1)$. By definition, the outcome is thus 0 for $t=0, \dots, \tilde{T}$, missing at $t=\tilde{T}+1$ if $\Delta=0$ and, 1 at $t=\tilde{T}+1$ if $\Delta=1$. Indeed, when $\Delta=0$, the patient's end of follow-up is due to occurrence of a right-censoring event during the last follow-up interval $\tilde{T}$ and as a result it is not known to the analyst whether the patient would have experienced failure at that time, i. e., $I(T=\tilde{T})$ and thus $Y(\tilde{T}+1)=I(T \leq \tilde{T})$ are missing. To simplify notation, we use overbars to denote covariate and exposure histories, e. g., a patient's exposure history through time t is denoted by $\bar{A}(t)=(A(0), \dots, A(t))$. Following the MSM framework [40], we approach the observed data in this study as realizations of n independent and identically distributed copies of $O=(\tilde{T}, \Delta, (1-\Delta)\Gamma, \bar{L}(\tilde{T}), \bar{A}(\tilde{T}), \Delta Y(\tilde{T}+1))$ denoted by O_i for $i=1, \dots, n$. The longest observed follow-up time is $\max_{i=1, \dots, n} \tilde{T}_i=36$ (9 years). Details about the approach implemented for mapping EHR data into these coarsened exposure, covariate and outcome data for each patient were described elsewhere [48, Appendix E].

In this study, we aim to evaluate the effect of dynamic treatment interventions on the cumulative risk of failure at a pre-specified time point t_0 , e. g., $t_0=11$ to investigate cumulative risks of failure over three years. The dynamic treatment interventions of interest correspond to treatment decisions made according to given clinical policies for initiation of an intensified therapy based on the patient's evolving A1c level. These policies denoted by d_θ were described above. Formally, these policies are individualized action rules [42] defined as a vector function $d_\theta=(d_\theta(0), \dots, d_\theta(t_0))$ where each function, $d_\theta(t)$ for $t=0, \dots, t_0$, is a decision rule for determining the action regimen (i. e., a treatment and right-censoring intervention) to be experienced by a patient at time t . A decision rule $d_\theta(t)$ maps the action and covariate history measured up to a given time t to an action regimen at time t : $d_\theta(t): (\bar{L}(t), \bar{A}(t-1)) \mapsto (a_1(t), a_2(t))$. In this study, the decision rules of interest are defined such that $d_\theta(t)(\bar{L}(t), \bar{A}(t-1))$ is:

- $(a_1(t), a_2(t))=(0, 0)$ (i. e., no use of an intensified treatment and no right-censoring) if and only if the patient was not previously treated with an intensified therapy (i. e., $A(t-1)=0$) and the A1c level at time t (an element of $L(t)$) was lower than or equal to the threshold θ .
- $(a_1(t), a_2(t))=(1, 0)$ (i. e., use of an intensified treatment and no right-censoring) otherwise.

To simplify notation below, the action regimen

$$(a(0)=d_\theta(0)(L(0)), \quad a(1)=d_\theta(1)(\bar{L}(1), a(0)), \dots, \quad a(t)=d_\theta(t)(\bar{L}(t), \bar{a}(t-1)))$$

through time t is denoted by $d_\theta(\bar{L}(t))$ for any given observed covariate history through time t , $\bar{L}(t)$. The parameter of interest denoted by $\psi^{\theta_1, \theta_2}$ is the difference between the cumulative risks at time t_0 associated with two distinct treatment strategies d_{θ_1} and d_{θ_2} :

$$\psi^{\theta_1, \theta_2}=P(Y_{d_{\theta_1}}(t_0+1)=1)-P(Y_{d_{\theta_2}}(t_0+1)=1),$$

where $Y_{d_\theta}(t_0+1)$ for $\theta=\theta_1, \theta_2$ denotes a patient's potential outcome at time t_0+1 had she been treated between study entry and time t_0 according to the decision rule d_θ . For conciseness, we refer the reader to earlier work [42, 43] [48, Appendices B and D] for a detailed description of the concepts and the counterfactual statistical framework on which relies the definition of this parameter of interest.

Identifiability of this parameter with the observational data above relies on at least three assumptions detailed elsewhere [48, Appendices C]: no unmeasured confounding, positivity, and consistent estimation of the action mechanism (defined in the next section).

If the MSM framework above (missing data framework) is not explicitly resting on the more general structural framework through additional explicit assumptions encoded by a causal diagram [50], then an additional assumption referred to as consistency assumption is made [51, 52].

In addition, a more or less flexible non-saturated MSM may be assumed [10–12, 44, 53]. The assumption encoded by such an MSM typically imposes constraints on the survival curves that underlie the definition of the parameter of interest $\psi^{\theta_1, \theta_2}$. In practice, specification of a non-saturated MSM is essentially an arbitrary

choice that does not encode real knowledge about the true survival curves of interest. The previous CER analysis of these observational data was based on such a MSM although minimal constraints were actually imposed because the MSM chosen was relatively close to saturation. Approaches to hedge against the bias that would arise from MSM misspecification in practice have been proposed [54] and are still being researched [55].

Alternatively, the MSM may be left nonparametric, i. e., no additional assumptions are made (e. g., through specification of a saturated MSM). This is the approach taken here because it reflects the absence of knowledge about the true functional forms of the four survival curves of interest. We note that such an approach may not always be practical due to data sparsity that prevents contrasting the interventions of interest with precision (curse of dimensionality). In such cases, a nonparametric MSM approach based on a working non-saturated model [56] can be adopted to explicitly recognize the limitation of an arbitrarily specified non-saturated MSM in capturing the true parameter of interest $\psi^{\theta_1, \theta_2}$.

3 Inverse probability weighting estimation

In this section, we describe three IPW estimators of the target parameter $\psi^{\theta_1, \theta_2}$ which are each implemented based on one of four PS estimation strategies. More precisely, each IPW estimator consists in, first, estimating each of the two cumulative counterfactual risks $\gamma^{\theta_1} \equiv P(Y_{d_{\theta_1}}(t_0 + 1) = 1)$ and $\gamma^{\theta_2} \equiv P(Y_{d_{\theta_2}}(t_0 + 1) = 1)$, separately, using for both one of the three IPW estimators described below and, second, taking the difference between the resulting two estimates. Inference about the cumulative risk difference (RD) $\psi^{\theta_1, \theta_2}$ from each of these three estimators is derived based on the influence curve of the corresponding IPW estimator and the delta method [57].

3.1 Three inverse probability weighting estimators

From here on, the indicator that a given event $\cdot$ has occurred is denoted by $I(\cdot)$.

3.1.1 Unbounded estimator

The Horvitz-Thompson estimator [58] of γ^θ denoted by $\gamma_{n, HT}^\theta$ is defined as the solution of the estimating equations associated with the following unbiased estimating function:

$$\frac{I(\bar{A}(\tilde{T}(t_0)) = d_\theta(\bar{L}(\tilde{T}(t_0))))}{\prod_{t=0}^{\tilde{T}(t_0)} g(A(t)|\bar{L}(t), \bar{A}(t-1))} Y(\tilde{T}(t_0) + 1) - \gamma^\theta,$$

where $\tilde{T}(t_0) = \min(\tilde{T}, t_0)$ and $g(A(t)|\bar{L}(t), \bar{A}(t-1))$ denotes the conditional probability of the observed action at time t given observed past covariates and actions.

This IPW estimator is a non-convex, linear combination of outcomes as shown by its closed-form expression below:

$$\gamma_{n, HT}^\theta = \frac{1}{n} \sum_{i=1}^n \frac{I(\bar{A}_i(\tilde{T}_i(t_0)) = d_\theta(\bar{L}_i(\tilde{T}_i(t_0))))}{\prod_{t=0}^{\tilde{T}_i(t_0)} g(A_i(t)|\bar{L}_i(t), \bar{A}_i(t-1))} Y_i(\tilde{T}_i(t_0) + 1).$$

Thus, this estimator is unbounded [59, Section 4.1], i. e., the resulting point estimates of the counterfactual cumulative risk is not guaranteed to fall between 0 and 1. We note also that all outcomes from the following patients are completely ignored by this estimator even if these patients followed the treatment strategy d_θ fully or partially before their end of follow-up (because $I(\bar{A}_i(\tilde{T}_i(t_0)) = d_\theta(\bar{L}_i(\tilde{T}_i(t_0)))) = 0$):

- patients who experienced a censoring event before or at time t_0 or,
- patients who experienced the outcome before or at t_0 but who did not follow the treatment strategy d_θ through the failure time or,
- patients who neither experienced the outcome nor a censoring event before or at t_0 and who did not follow the treatment strategy d_θ through t_0 .

Both because of the resulting loss of information and the aforementioned unboundedness property, this IPW estimator is expected [60, Section 2.1] to be less efficient than the other two IPW estimators discussed in the following sections.

The resulting Horvitz-Thompson estimator of $\psi^{\theta_1, \theta_2}$ is denoted by $\psi_{n, HT}^{\theta_1, \theta_2}$ and defined as:

$$\psi_{n, HT}^{\theta_1, \theta_2} = \gamma_{n, HT}^{\theta_1} - \gamma_{n, HT}^{\theta_2}.$$

An estimator of its variance is given by

$$\frac{1}{n^2} \sum_{i=1}^n \left(\frac{I(\bar{A}_i(\tilde{T}_i(t_0)) = d_{\theta_1}(\bar{L}_i(\tilde{T}_i(t_0)))) - I(\bar{A}_i(\tilde{T}_i(t_0)) = d_{\theta_2}(\bar{L}_i(\tilde{T}_i(t_0))))}{\prod_{t=0}^{\tilde{T}_i(t_0)} g(A_i(t)|\bar{L}_i(t), \bar{A}_i(t-1))} Y_i(\tilde{T}_i(t_0) + 1) - \psi_{n, HT}^{\theta_1, \theta_2} \right)^2.$$

This estimator is conservative in the sense that it over-estimates the variance of the unbounded IPW estimator when the probabilities $g(A_i(t)|\bar{L}_i(t), \bar{A}_i(t-1))$ in the previous expressions are replaced with efficient estimates.

3.1.2 Bounded estimator

An alternative [61, 62] [63, Section III. C.] to the previous estimator of γ^θ denoted by $\gamma_{n, bd}^\theta$ is defined as the solution of the estimating equations associated with the following unbiased estimating function:

$$D_{bd}(O|\gamma^\theta, g) = \frac{I(\bar{A}(\tilde{T}(t_0)) = d_\theta(\bar{L}(\tilde{T}(t_0))))}{\prod_{t=0}^{\tilde{T}(t_0)} g(A(t)|\bar{L}(t), \bar{A}(t-1))} (Y(\tilde{T}(t_0) + 1) - \gamma^\theta).$$

Note that the only difference with the previous estimating function is that not only the outcome $Y(\tilde{T}(t_0) + 1)$ but also the target parameter γ^θ is multiplied by the inverse probability weight $\frac{I(\bar{A}(\tilde{T}(t_0)) = d_\theta(\bar{L}(\tilde{T}(t_0))))}{\prod_{t=0}^{\tilde{T}(t_0)} g(A(t)|\bar{L}(t), \bar{A}(t-1))}$. As a result, the IPW estimator is now a convex, linear combination (i. e., a weighted average) of outcomes as shown by its closed-form expression below:

$$\gamma_{n, bd}^\theta = \frac{\sum_{i=1}^n \frac{I(\bar{A}_i(\tilde{T}_i(t_0)) = d_\theta(\bar{L}_i(\tilde{T}_i(t_0))))}{\prod_{t=0}^{\tilde{T}_i(t_0)} g(A_i(t)|\bar{L}_i(t), \bar{A}_i(t-1))} Y_i(\tilde{T}_i(t_0) + 1)}{\sum_{i=1}^n \frac{I(\bar{A}_i(\tilde{T}_i(t_0)) = d_\theta(\bar{L}_i(\tilde{T}_i(t_0))))}{\prod_{t=0}^{\tilde{T}_i(t_0)} g(A_i(t)|\bar{L}_i(t), \bar{A}_i(t-1))}}.$$

This estimator is thus bounded [59, Section 4.1], i. e. the resulting point estimates of the counterfactual cumulative risk is guaranteed to fall between 0 and 1. We note however that implementation of this estimator results in the same loss of information that characterize the Horvitz-Thompson estimator. Because of this loss of information, this IPW estimator is expected to be less efficient than the IPW estimator discussed in the following section.

The resulting bounded estimator of $\psi^{\theta_1, \theta_2}$ is denoted by $\psi_{n, bd}^{\theta_1, \theta_2}$ and defined as:

$$\psi_{n, bd}^{\theta_1, \theta_2} = \gamma_{n, bd}^{\theta_1} - \gamma_{n, bd}^{\theta_2}.$$

An estimator of its variance is given by

$$\frac{1}{n^2} \sum_{i=1}^n \left(C_n^{\theta_1} D_{bd}(O_i | \gamma_{n,bd}^{\theta_1}, g) - C_n^{\theta_2} D_{bd}(O_i | \gamma_{n,bd}^{\theta_2}, g) \right)^2 \text{ using the notation}$$

$$C_n^\theta \equiv \left(\frac{1}{n} \sum_{i=1}^n \frac{I(\bar{A}_i(\tilde{T}_i(t_0)) = d_\theta(\bar{L}_i(\tilde{T}_i(t_0))))}{\prod_{t=0}^{\tilde{T}_i(t_0)} g(A_i(t) | \bar{L}_i(t), \bar{A}_i(t-1))} \right)^{-1} \text{ for } \theta = \theta_1, \theta_2.$$

This estimator is conservative in the sense that it over-estimates the variance of the bounded IPW estimator when the probabilities $g(A_i(t) | \bar{L}_i(t), \bar{A}_i(t-1))$ in the previous expressions are replaced with efficient estimates.

3.1.3 Hazard-based bounded estimator

The previous two IPW estimators are designed to directly evaluate the counterfactual cumulative risk γ^θ . These approaches result in a loss of information because outcomes from patients who only partially follow the dynamic intervention of interest are ignored by these estimators. Instead, an alternate IPW estimator that makes use of all outcomes observed while patients follow the dynamic interventions of interest has commonly been used in practice. This alternate IPW estimator is thus expected to be more efficient than the previous two IPW estimators. Unlike them, it is designed to evaluate the counterfactual hazards associated with the dynamic intervention of interest for $t=0, \dots, t_0$ and the resulting hazard estimates are then mapped into estimates of the counterfactual cumulative risk γ^θ [64] [48, Appendix D]. Details about this indirect IPW estimating approach for γ^θ are described below.

We denote the counterfactual discrete-time hazard under the dynamic intervention d_θ at time t with $\alpha^{t,\theta}$, i. e. $\alpha^{t,\theta} \equiv P(Y_{d_\theta}(t+1) = 1 | Y_{d_\theta}(t) = 0)$. A stabilized bounded IPW estimator of $\alpha^{t,\theta}$ is denoted by $\alpha_{n,bd}^{t,\theta}$ and defined as the solution of the estimating equations associated with the following estimating function:

$$h(O, t, \theta, g, g')(Y(t+1) - \alpha^{t,\theta}) \text{ with } h(O, t, \theta, g, g') \text{ defined as}$$

$$I(Y(t) = 0)I(\bar{A}(t) = d_\theta(\bar{L}(t))) \frac{\prod_{j=0}^t g'(A(j) = d_\theta(\bar{L}(j)) | \bar{A}(j-1) = d_\theta(\bar{L}(j-1)))}{\prod_{j=0}^t g(A(j) | \bar{L}(j), \bar{A}(j-1))} \text{ using the notation } g'(A(j) = d_\theta(\bar{L}(j)) | \bar{A}(j-1) = d_\theta(\bar{L}(j-1)))$$

to denote the probability of a patient following the intervention of interest at time j given that she followed the intervention through time $j-1$.

This IPW estimator is a convex, linear combination of the outcomes at time $t+1$ from all patients who did not experience the event before or at time t and who followed the intervention of interest through time t as shown by the closed-form expression:

$$\alpha_{n,bd}^{t,\theta} = \sum_{i=1}^n \frac{h(O_i, t, \theta, g, g')}{\sum_{i=1}^n h(O_i, t, \theta, g, g')} Y(t+1). \quad (1)$$

As noted in previous work [20, Section 3.3] and when $0 < \alpha^{t,\theta} < 1$, this estimator is equivalent to a stabilized IPW estimator of the coefficient $\beta^{t,\theta}$ denoted by $\beta_{n,bd}^{t,\theta}$ of the following saturated logistic dynamic MSM for the counterfactual hazards under interventions d_θ at time points t :

$$\text{logit}(\alpha^{t,\theta}) = \sum_{j \in \mathcal{T}} \sum_{k \in \Theta} \beta^{j,k} I(t=j, \theta=k),$$

where $\text{logit}(x) = \log(\frac{x}{1-x})$, $t \in \mathcal{T} \equiv \{0, 1, 2, \dots, K\}$ for $K \geq t_0$, and $\Theta = \{7; 7.5; 8; 8.5\}$. More specifically, the estimator $\beta_{n,bd}^{t,\theta}$ is defined by the standard IPW estimation approach to fit a dynamic MSM [42, 43] with the following choice of numerator for the IP weight associated with each person-time outcome $Y_i(t+1)$ contributing to the weighted regression: $\prod_{j=0}^t g'(A(j) = d_\theta(\bar{L}(j)) | \bar{A}(j-1) = d_\theta(\bar{L}(j-1)))$. The equivalence between $\beta_{n,bd}^{t,\theta}$ and $\alpha_{n,bd}^{t,\theta}$ is expressed by the following link: $\beta_{n,bd}^{t,\theta} = \text{logit}(\alpha_{n,bd}^{t,\theta})$.

The collection of estimators $\alpha_{n,bd}^{t,\theta}$ (equivalently $\beta_{n,bd}^{t,\theta}$) for $t=0, \dots, t_0$ can then be mapped into an hazard-based IPW estimator of the counterfactual cumulative risk γ^θ denoted by $\gamma_{n,h}^\theta$ and defined as:

$$\begin{aligned}\gamma_{n,h}^\theta &= 1 - \prod_{t=0}^{t_0} (1 - \alpha_{n,bd}^{t,\theta}) \\ &= 1 - \prod_{t=0}^{t_0} \left(1 - \frac{1}{1 + \exp(-\beta_{n,bd}^{t,\theta})} \right).\end{aligned}$$

In practice, computation of the point estimate of the counterfactual hazards may often be faster using the first equality.

The resulting hazard-based bounded IPW estimator of $\psi^{\theta_1, \theta_2}$ is denoted by $\psi_{n,h}^{\theta_1, \theta_2}$ and defined as:

$$\psi_{n,h}^{\theta_1, \theta_2} = \gamma_{n,h}^{\theta_1} - \gamma_{n,h}^{\theta_2}.$$

An estimator of its variance can be obtained (see proof in the Appendix) using the following closed-form expression derived from the influence curve associated with the IPW estimator of the coefficients of the saturated dynamic MSM:

$$\begin{aligned}\frac{1}{n^2} \sum_{i=1}^n \left[\left((1 - \gamma_{n,h}^{\theta_1}) \sum_{t=0}^{t_0} \alpha_{n,bd}^{t,\theta_1} (I(t=j, \theta_1=k))'_{j \in \mathcal{T}, k \in \Theta} \right. \right. \\ \left. \left. - (1 - \gamma_{n,h}^{\theta_2}) \sum_{t=0}^{t_0} \alpha_{n,bd}^{t,\theta_2} (I(t=j, \theta_2=k))'_{j \in \mathcal{T}, k \in \Theta} \right) IC(O_i | \alpha_{n,bd}, g, g') \right]^2\end{aligned}$$

in which the collection of all estimators $\alpha_{n,bd}^{t,\theta}$ for $t \in \mathcal{T}$ and $\theta \in \Theta$ is denoted by

$$\begin{aligned}\alpha_{n,bd} &\equiv (\alpha_{n,bd}^{0,7}, \alpha_{n,bd}^{0,7.5}, \dots, \alpha_{n,bd}^{0,8.5}, \alpha_{n,bd}^{1,7}, \alpha_{n,bd}^{1,7.5}, \dots), \\ (I(t=j, \theta=k))_{j \in \mathcal{T}, k \in \Theta} &\equiv \begin{pmatrix} I(t=0, \theta=7) \\ I(t=0, \theta=7.5) \\ \vdots \\ I(t=0, \theta=8.5) \\ I(t=1, \theta=7) \\ I(t=1, \theta=7.5) \\ \vdots \end{pmatrix} \text{ for } \theta = \theta_1, \theta_2, \text{ and}\end{aligned}$$

the value of the influence curve for patient i , $IC(O_i | \alpha_{n,bd}, g, g')$, is defined as

$$C_n \sum_{t \in \mathcal{T}} \sum_{\theta \in \Theta} h(O_i, t, \theta, g, g') (I(t=j, \theta=k))_{j \in \mathcal{T}, k \in \Theta} (Y_i(t+1) - \alpha_{n,bd}^{t,\theta})$$

using the square matrix notation

$$C_n \equiv \left(\frac{1}{n} \sum_{i=1}^n \sum_{t \in \mathcal{T}} \sum_{\theta \in \Theta} h(O_i, t, \theta, g, g') \alpha_{n,bd}^{t,\theta} (1 - \alpha_{n,bd}^{t,\theta}) V(t, \theta) \right)^{-1}$$

and

$$V(t, \theta) \equiv (I(t=j, \theta=k))_{j \in \mathcal{T}, k \in \Theta} \times (I(t=j, \theta=k))'_{j \in \mathcal{T}, k \in \Theta}.$$

The estimator above is conservative in the sense that it over-estimates the variance of the hazard-based bounded IPW estimator when the probabilities $g(A_i(t) | \bar{L}_i(t), \bar{A}_i(t-1))$ in the previous expressions are replaced with efficient estimates.

3.2 Four propensity score estimation approaches

In observational studies, the conditional probabilities $g(A(t) = d_\theta(\bar{L}(t)) | \bar{L}(t), \bar{Y}(t) = 0, \bar{A}(t-1) = d_\theta(\bar{L}(t-1)))$ that are required to derive point estimates and inferences with the previous three IPW estimators are unknown and must thus be estimated first. In this section, we start by describing the decomposition of these nuisance parameters based on 5 classes of propensity scores (PS). Next, we describe the four approaches considered for estimating these various PS in this report: two model-based and two data-adaptive estimation approaches. For conciseness, the vector of probabilities $g(A(t) = d_\theta(\bar{L}(t)) | \bar{L}(t), \bar{Y}(t) = 0, \bar{A}(t-1) = d_\theta(\bar{L}(t-1)))$ for $t \in \mathcal{T}$ is denoted by g below.

3.3 Decomposition of the action mechanism

The conditional probability $g(A(t) | \bar{L}(t), \bar{Y}(t) = 0, \bar{A}_1(t-1), \bar{A}_2(t-1) = 0)$ for $t \in \mathcal{T}$ is referred to as the action mechanism at time t and can be factorized based on the following 5 PS:

- PS for TI initiation denoted by $\mu_1(t)$:

$$P(A_1(t) = 1 | \bar{L}(t), \bar{Y}(t) = 0, \bar{A}_1(t-1) = 0, \bar{A}_2(t) = 0)$$

- PS for TI continuation denoted by $\mu_2(t)$:

$$P(A_1(t) = 1 | \bar{L}(t), \bar{Y}(t) = 0, \bar{A}_1(t-2), A_1(t-1) = 1, \bar{A}_2(t) = 0)$$

- PS for right-censoring by administrative end of study denoted by $\mu_3(t)$:

$$P(I(A_2(t) = 1, \Gamma = 1) = 1 | \bar{L}(t), \bar{Y}(t) = 0, \bar{A}_1(t-1), \bar{A}_2(t-1) = 0)$$

- PS for right-censoring by disenrollment from the health plan denoted by $\mu_4(t)$:

$$P(I(A_2(t) = 1, \Gamma = 2) = 1 | \bar{L}(t), \bar{Y}(t) = 0, \bar{A}_1(t-1), \bar{A}_2(t-1) = 0, I(A_2(t) = 1, \Gamma = 1) = 0)$$

- PS for right-censoring by death denoted by $\mu_5(t)$:

$$P(I(A_2(t) = 1, \Gamma = 3) = 1 | \bar{L}(t), \bar{Y}(t) = 0, \bar{A}_1(t-1), \bar{A}_2(t-1) = 0, I(A_2(t) = 1, \Gamma = 1) = 0, I(A_2(t) = 1, \Gamma = 2) = 0).$$

Thus, for patients who did not fail before time t and who followed the decision rule d_θ through time t , an estimate of the nuisance parameter $g(A(t) = d_\theta(\bar{L}(t)) | \bar{L}(t), \bar{Y}(t) = 0, \bar{A}(t-1) = d_\theta(\bar{L}(t-1)))$ can be derived from estimates of the 5 PS above based on the following closed-form expression implied by the factorization of the action mechanism at time t using the chain rule:

$$\left(I(\bar{A}_1(t-1) = 0) \mu_1(t)^{A_1(t)} (1 - \mu(t))^{1-A_1(t)} + I(A_1(t-1) = 1) \mu_2(t) \right) \times (1 - \mu_3(t)) (1 - \mu_4(t)) (1 - \mu_5(t)).$$

3.4 Model-based PS estimation

3.4.1 Logistic models with pooled data over time

A common [10, 11] approach used in practice by analysts to simultaneously estimate each of the 5 PS above for all $t \in \mathcal{T}$ consists in fitting a single model, referred to as a “pooled model”, using data pooled over time t . We also considered such an approach in this report. More specifically, data were pooled over all follow-up times t to fit a separate main-term logistic model for estimating each of the 3 PS for right-censoring ($\mu_3(t)$, $\mu_4(t)$, $\mu_5(t)$) and the PS for TI continuation ($\mu_2(t)$). Data were also pooled for all time points $t > 0$ to fit a single main-term logistic model for estimating the PS for TI initiation after $t = 0$ (i. e., $\mu_1(t)$ for $t > 0$). A separate main-term logistic model was fitted for estimating the PS for TI initiation at $t = 0$ (i. e., $\mu_1(0)$). By “main-term logistic model”, we mean a logistic model with only main terms for each explanatory variable considered

(i. e., no interaction terms between explanatory variables). The explanatory variables considered were all time-independent covariates and the last measurement of time-varying covariates (Markov assumption). In addition, exposure to TI in the last period was included as an explanatory variable for the 3 PS for right-censoring and the latest change in A1c was included as an explanatory variable for estimating all PS. All pooled logistic models also included the variable indexing the 90-day follow-up intervals (i. e., t) as an explanatory variable. We denote the estimator of the nuisance parameter g derived with this approach by g_n .

3.4.2 Logistic models with data stratified by time

In the previous approach, each pooled logistic model encodes the assumption that the associations between the explanatory variables and the PS outcome variable (e. g. death occurrence for PS $\mu_5(t)$) do not change over time. Concern over this assumption [65] motivates instead the use of a different logistic model (referred to as a “stratified model”) to estimate each PS at each time point t separately or, at least, the inclusion of interaction terms between the explanatory variables and functions of t in the pooled models. We note that it is such a concern over time-modified confounding that motivated, in the previous section, the specification of two separate models to estimate the PS for TI initiation: one stratified logistic model to estimate $\mu_1(0)$ and one separate pooled logistic model to estimate all $\mu_1(t)$ for $t > 0$ simultaneously. To fully address concerns over time-modified confounding, we also considered a second PS estimation approach in which, for each time point t separately, 5 main-term logistic models were fitted to estimate each of the 5 PS. The parameterization of these stratified models are the same as that described in the previous section with the difference that the time variable t was omitted from all logistic models. We denote the estimator of the nuisance parameter g derived with this approach by $g_{n,t}$.

3.5 Data-adaptive PS estimation

3.5.1 Logistic models with data-adaptive selection of interaction terms

To lessen the constraints imposed by main-term logistic models, we considered a third PS estimation approach that is based on extending the previous stratified logistic models by data-adaptively including two-way interaction terms between explanatory variables. Due to the large number of explanatory variables in this study and overfitting concerns, we separately implemented the following ad hoc data-adaptive algorithm for selecting the subset of all possible interaction terms to include in each stratified model. For each of the 5 PS and each time $t \in \mathcal{T}$, we first computed 105 two-way interaction terms based on the 15 explanatory variables that were most significantly associated (smallest p-value) with the PS outcome variable in a univariate logistic regression. Second, for each of these 105 terms, we implemented a logistic regression of the PS outcome variable on the interaction term and the two main terms that define the interaction term. Third, we identified the interaction terms with a p-value lower than 0.05. Finally, if more than 50 interaction terms met this criterion, we selected only the 50 terms with the smallest p-value and added them to the stratified logistic model for the PS. We denote the estimator of the nuisance parameter g derived with this approach by $g_{n,t,\times}$.

3.5.2 Super learning

The two model-based estimators g_n and $g_{n,t}$ described earlier do not reflect real subject-matter knowledge about the adequacy of the expit¹ function chosen to properly represent the true values of the 5 PS over time.

¹ $\text{expit}(x) = \frac{1}{1 + \exp(-x)}$.

Indeed, in practice, PS model specification such as choosing a logistic model is typically rooted in tradition, preference, or convenience. To avoid erroneous inference due to such arbitrary model specifications, data-adaptive estimation of the nuisance parameter g may be implemented in practice but consistent estimation then relies on judicious selection of a machine learning algorithm also known as “learner”. We considered such a learner $g_{n,t,\times}$ in the previous section but many other learners have been proposed and can be used as potential candidates for estimating the 5 PS (e. g. [3, 66–74]). Akin to the selection of a parametric model, the selection of a learner does not typically reflect real subject-matter knowledge about the relative suitability of the different learners available, since “in practice it is generally impossible to know a priori which learner will perform best for a given prediction problem and data set” [14]. To hedge against erroneous inference due to arbitrary selection of a learner, Super Learning [14] (SL) may be implemented to combine predicted values from a library of various candidate learners (that includes the arbitrary learner that would have been guessed otherwise) through a weighted average. The selection of the optimal combination of the candidate learners is based on cross-validation [75–78] to protect against overfitting such that the resulting learner (called “super learner”) performs asymptotically as well (in terms of mean error) or better than any of the candidate learners considered. If the arbitrary learners that would have been guessed is based on a parametric model and happens to be correct then using SL instead of the correctly guessed learner only comes at a price of limited increase in prediction variability.

We considered a SL approach to estimate the PS in this report. More specifically, for each time point t separately, 5 super learners were implemented to estimate each of the 5 PS based on the following 10 candidate learners: (i) 5 learners² defined by logistic models with only main terms for the most predictive explanatory variables identified³ by a significant p-value in univariate regressions with 5 significance levels ($\alpha = 1e-30, 1e-10, 1e-5, 0.1$, and 1), and (ii) 5 polychotomous regression learners⁴ based on the most predictive explanatory variables identified by a significant p-value in univariate regressions with the same 5 significance levels. We denote the estimator of the nuisance parameter g derived with this SL approach by $g_{n,t,SL}$.

4 Comparison of practical performances

In this section, we examine and contrast the results from the application of the three IPW estimators based on the four PS estimation approaches just described to address the CER question from the observational study introduced in Section 2. We note that the bounded hazard-based IPW estimators in this report were implemented using nonparametric estimation of the numerator of the stabilized IP weights, i. e.

$$\frac{\prod_{j=0}^t g'(A(j) = d_\theta(\bar{L}(j)) | \bar{A}(j-1) = d_\theta(\bar{L}(j-1)))}{\prod_{j=0}^t g(A(j) | \bar{L}(j), \bar{A}(j-1))}.$$

More specifically, each probability $g'(A(j) = d_\theta(\bar{L}(j)) | \bar{A}(j-1) = d_\theta(\bar{L}(j-1)))$ for $j \in \mathcal{T}$ was estimated by the proportion of patients in the observed data who followed rule d_θ at time t among patients who did not fail before and at time $t-1$ and who followed rule d_θ at time $t-1$. Note that this weight stabilization scheme (i. e., the definition of the numerator chosen for the IP weight) follows guidelines in the literature on dynamic MSM (e. g., Appendix A3 in [45]) in the sense that the numerator of the IP weight is i) only a function of time t and the A1c threshold θ that defines the rules of interest, and ii) is thus not a function of the patient’s exposure history $\bar{A}(t)$ as is typically the case with the numerator of the IP weight commonly used to fit static MSM in practice.

² implemented by the `SL.glm` routine available in the `SuperLearner` R package [79].

³ using the template screening routine `screen.glmP` available in the `SuperLearner` R package.

⁴ implemented by the `SL.polyclass` routine given in [80, Appendix]. This routine implements the `polyclass` learner [72] based on the Bayesian Information Criterion (BIC) as the model selection criterion. To improve computing speed, this learner was favored over the `SL.polymars` routine that is available by default in the `SuperLearner` R package but that relies on cross-validation for model selection.

Figure 1: Crude estimates over 16 quarters of the four survival curves associated with the four TI initiation strategies d_θ with $\theta=7, 7.5, 8, 8.5$.

Results from the 12 IPW estimators considered in this report are also compared to the results from a crude analysis that aims to contrast the survival curves associated with the four treatment intensification strategies of interest d_θ without any adjustment for confounding and selection bias. Such an analysis can be implemented by applying the hazard-based IPW estimator with its stabilized weights set to 1. Figure 1 displays the resulting point estimates of the four survival curves, i.e., crude estimates of $1 - \gamma^\theta$ for $t_0 = 0, \dots, 15$ and $\theta = 7, 7.5, 8, 8.5$. These crude point estimates of the survival curves can be compared to their analogues based on the three IPW estimators implemented with the four PS estimation approaches displayed on Figures 2–4. Each of these figures contains four plots that each displays estimates of the four counterfactual survival curves based on the same IPW estimator but a different PS estimation approach. Table 1 provides details about the comparison of the four survival curves at 3 years using the crude analysis approach. Inferences in Table 1 can be contrasted to that based on the three IPW estimators implemented with the four PS estimation approaches in Table 2.

We note a striking difference between the unbounded IPW estimates of the survival curves based on the PS estimators g_n and $g_{n,t}$ and the unbounded IPW estimates based on $g_{n,t,SL}$ on Figure 2. The first two sets of estimates resemble the crude estimates on Figure 1 with the difference that the crude estimates provide evidence, albeit weak, that is consistent with results from the ACCORD and ADVANCE randomized trials which indicated a potential beneficial effect of more aggressive therapy initiation rules. Indeed, the crude point estimates indicate a consistent ordering of the four survival curves over time with curves indexed by a lower A1c threshold (aggressive rules) above the curves indexed by a higher A1c threshold. While the trend from the crude analysis is restored and more apparent with the IPW estimates based on $g_{n,t,\times}$, it becomes evident with the early and clear separation of the IPW estimates of the survival curves based on $g_{n,t,SL}$. The change in the unbounded IPW results across PS estimation approaches can be explained by examining and contrasting the distributions of the estimated IP weights resulting from each PS estimator. Figure 5 describes the distribution of the non-zero estimated IP weights

Figure 2: Unbounded IPW estimates (no weight truncation) over 16 quarters of the four survival curves associated with the four TI initiation strategies d_θ with $\theta = 7, 7.5, 8, 8.5$. The estimates on the top left, top right, bottom left, and bottom right plots were obtained based on the estimators g_n , $g_{n,t}$, $g_{n,t,\times}$ and $g_{n,t,SL}$, respectively.

$$\frac{I(\bar{A}_i(\tilde{T}_i(t_0)) = d_\theta(\bar{L}_i(\tilde{T}_i(t_0))))}{\prod_{t=0}^{\tilde{T}_i(t_0)} g(A_i(t)|\bar{L}_i(t), \bar{A}_i(t-1))}$$

associated with observed failures $Y_i(\tilde{T}_i(t_0) + 1) = 1$ and pooled over time points $t_0 = 0, \dots, 15$ and rules $\theta = 7, 7.5, 8, 8.5$. We note a clear and increasing shift of the estimated IP weights toward the lower bound 1 as the PS estimation approach becomes more and more nonparametric, i. e. as we go from estimating the PS with g_n , $g_{n,t}$, $g_{n,t,\times}$, and $g_{n,t,SL}$. These results provide supporting evidence that model-based PS estimation can lead to bias which can be corrected with the use of data-adaptive PS estimation such as SL in practice.

The set of results from the hazard-based bounded IPW estimator on Figure 4 follow a very different pattern with little distinctions between the point estimates derived from the four PS estimation approaches and that despite a similar and notable change in the distribution of the stabilized IPW weights. Figure 7 describes the distribution of the non-zero estimated stabilized IP weights $h(O_i, t, \theta, g, g')$ associated with observed failures $Y_i(t+1) = 1$ and pooled over time points $t = 0, \dots, 34$ and rules $\theta = 7, 7.5, 8, 8.5$. We note a clear and increasing shift of the estimated IP weights toward the lower bound 0 and great reduction in the number of larger weights as the PS estimation approach becomes more and more nonparametric.

Figure 3: Bounded IPW estimates (no weight truncation) over 16 quarters of the four survival curves associated with the four TI initiation strategies d_θ with $\theta=7, 7.5, 8, 8.5$. The estimates on the top left, top right, bottom left, and bottom right plots were obtained based on the estimators g_n , $g_{n,t}$, $g_{n,t,x}$ and $g_{n,t,SL}$, respectively.

The set of results from the bounded IPW estimator on Figure 3 follows an intermediate pattern between that of the unbounded and hazard-based bounded estimators in the sense that while the curves become increasingly and more consistently separated in the expected directions as the PS estimation becomes more nonparametric (similar to Figure 2), the distinctions between results from the four PS estimation approaches are less obvious (similar to Figure 4). Figure 6 describes the distribution of the non-zero estimated IP weights

$$\frac{I(\bar{A}_i(\tilde{T}_i(t_0)) = d_\theta(\bar{L}_i(\tilde{T}_i(t_0))))}{\prod_{t=0}^{\tilde{T}_i(t_0)} g(A_i(t)|\bar{L}_i(t), \bar{A}_i(t-1))} \\ \sum_{i=1}^n \frac{I(\bar{A}_i(\tilde{T}_i(t_0)) = d_\theta(\bar{L}_i(\tilde{T}_i(t_0))))}{\prod_{t=0}^{\tilde{T}_i(t_0)} g(A_i(t)|\bar{L}_i(t), \bar{A}_i(t-1))}$$

associated with observed failures $Y_i(\tilde{T}_i(t_0)+1)=1$ and pooled over time points $t_0=0, \dots, 15$ and rules $\theta=7, 7.5, 8.5$. We note a clear and increasing concentration of the estimated IP weights around the median with a reduction in the number of larger weights as the PS estimation approach becomes more and more nonparametric.

Figure 4: Bounded hazard-based IPW estimates (no weight truncation) over 16 quarters of the four survival curves associated with the four TI initiation strategies d_θ with $\theta = 7, 7.5, 8, 8.5$. The estimates on the top left, top right, bottom left, and bottom right plots were obtained based on the estimators g_n , $g_{n,t}$, $g_{n,t,x}$, and $g_{n,t,SL}$, respectively.

We conjecture that the apparent gradual change in the patterns of results across PS estimation approaches as we go from the unbounded, bounded and hazard-based IPW estimators relates to the differences in boundedness and efficiency between the three estimators. More specifically, we conjecture that the successive decrease in sensitivity to bias from errors in PS estimation that result in larger weight values is a consequence of the boundedness property and the improved efficiency of the hazard-based estimator relative to the bounded IPW estimator. This conjecture is also supported by previous work [20] in which we studied an alternate estimator, referred to as Targeted Minimum Loss based Estimator (TMLE), using the same data and based on the same four PS estimation approaches described in this report to evaluate the same causal estimands. TMLE is a substitution estimator and is thus also characterized by the boundedness property. Our previous work showed that TMLE was slightly more efficient than the hazard-based IPW estimator studied here but that TMLE is also not very sensitive to bias from increasing errors in PS estimation across the four PS estimation approaches studied in this report.

In Table 2, we compare the inference we would derive in practice from various estimation choices that led to results consistent with that of previous randomized experiments. It indicates that the

Table 1: Inferences from the crude estimators of six risk differences (RD) at 12 quarters. For each crude RD estimator, the point estimate, standard error, the estimate of the lower and upper bound of the 95 % confidence interval, and the associated p-value are denoted by $\psi_n^{\theta_1, \theta_2}$, $\Gamma_n^{\theta_1, \theta_2}$, $\psi_n^{\theta_1, \theta_2, -}$, $\psi_n^{\theta_1, \theta_2, +}$, and p , respectively.

θ_1	θ_2	$\psi_n^{\theta_1, \theta_2}$	$\Gamma_n^{\theta_1, \theta_2}$	$\psi_n^{\theta_1, \theta_2, -}$	$\psi_n^{\theta_1, \theta_2, +}$	p
8.5	8	1.3e-03	1.5e-03	-1.6e-03	4.2e-03	0.386
8.5	7.5	2.8e-03	3e-03	-3.2e-03	8.8e-03	0.357
8.5	7	0.0335	0.0111	0.0116	0.0553	3e-03
8	7.5	1.5e-03	2.7e-03	-3.7e-03	6.8e-03	0.574
8	7	0.0322	0.0111	0.0105	0.0538	4e-03
7.5	7	0.0307	0.0108	9.6e-03	0.0518	4e-03

Table 2: Comparison of inferences from three IPW estimators of six risk differences (RD) at 12 quarters based on four PS estimators. For each RD estimator, the PS estimator, the point estimate, standard error, its ratio to that of the hazard-based estimator with SL, the estimate of the lower and upper bound of the 95 % confidence interval, and the associated p-value are denoted by g_n , $\psi_n^{\theta_1, \theta_2}$, $\Gamma_n^{\theta_1, \theta_2}$, RE, $\psi_n^{\theta_1, \theta_2, -}$, $\psi_n^{\theta_1, \theta_2, +}$, and p , respectively.

IPW	θ_1	θ_2	g_n	$\psi_n^{\theta_1, \theta_2}$	$\Gamma_n^{\theta_1, \theta_2}$	$\psi_n^{\theta_1, \theta_2, -}$	$\psi_n^{\theta_1, \theta_2, +}$	p	RE
$\psi_{n, HT}^{\theta_1, \theta_2}$	8.5	8	$g_{n,t,SL}$	7.3e-03	4.3e-03	-1.1e-03	0.0156	0.089	1.059
	8.5	7.5	$g_{n,t,SL}$	0.0231	7.5e-03	8.4e-03	0.0378	2e-03	1.054
	8.5	7	$g_{n,t,SL}$	0.0506	0.0129	0.0254	0.0758	0	1.031
	8	7.5	$g_{n,t,SL}$	0.0158	6.8e-03	2.5e-03	0.0292	2e-02	1.055
	8	7	$g_{n,t,SL}$	0.0433	0.0128	0.0182	0.0685	1e-03	1.033
	7.5	7	$g_{n,t,SL}$	0.0275	0.0118	4.4e-03	0.0506	0.019	1.029
$\psi_{n, bd}^{\theta_1, \theta_2}$	8.5	8	$g_{n,t,SL}$	9.8e-03	4.6e-03	8e-04	0.0188	0.032	1.137
	8.5	7.5	$g_{n,t,SL}$	0.0255	7.9e-03	1e-02	0.041	1e-03	1.112
	8.5	7	$g_{n,t,SL}$	0.0572	0.0127	0.0322	0.0822	0	1.021
	8	7.5	$g_{n,t,SL}$	0.0157	7.2e-03	1.5e-03	0.0298	3e-02	1.119
	8	7	$g_{n,t,SL}$	0.0474	0.0127	0.0224	0.0723	0	1.023
	7.5	7	$g_{n,t,SL}$	0.0317	0.0117	8.7e-03	0.0547	7e-03	1.025
$\psi_{n, h}^{\theta_1, \theta_2}$	8.5	8	g_n	0.0109	5.7e-03	-3e-04	0.022	0.056	1.411
	8.5	7.5	g_n	0.0212	8.7e-03	4.1e-03	0.0382	0.015	1.226
	8.5	7	g_n	0.0387	0.0135	0.0122	0.0652	4e-03	1.085
	8	7.5	g_n	0.0103	7.8e-03	-5e-03	0.0256	0.187	1.208
	8	7	g_n	0.0278	0.0138	8e-04	0.0549	0.044	1.11
	7.5	7	g_n	0.0175	0.0133	-8.5e-03	0.0435	0.187	1.161
$\psi_{n, h}^{\theta_1, \theta_2}$	8.5	8	$g_{n,t}$	9.9e-03	5.5e-03	-9e-04	0.0207	0.071	1.361
	8.5	7.5	$g_{n,t}$	0.025	7.8e-03	9.6e-03	0.0404	1e-03	1.103
	8.5	7	$g_{n,t}$	0.0418	0.0126	0.0172	0.0664	1e-03	1.006
	8	7.5	$g_{n,t}$	0.0151	7.1e-03	1.2e-03	0.029	0.033	1.097
	8	7	$g_{n,t}$	0.0319	0.0129	6.7e-03	0.0571	0.013	1.036
	7.5	7	$g_{n,t}$	0.0168	0.012	-6.7e-03	0.0403	0.162	1.05
$\psi_{n, h}^{\theta_1, \theta_2}$	8.5	8	$g_{n,t, \times}$	0.0129	5.5e-03	2.1e-03	0.0237	2e-02	1.37
	8.5	7.5	$g_{n,t, \times}$	0.0337	7.9e-03	0.0183	0.0491	0	1.104
	8.5	7	$g_{n,t, \times}$	0.0483	0.0129	0.0229	0.0737	0	1.037
	8	7.5	$g_{n,t, \times}$	0.0209	6.9e-03	7.3e-03	0.0345	3e-03	1.075
	8	7	$g_{n,t, \times}$	0.0354	0.0129	1e-02	0.0608	6e-03	1.042
	7.5	7	$g_{n,t, \times}$	0.0146	0.0121	-9.2e-03	0.0383	0.23	1.06
$\psi_{n, h}^{\theta_1, \theta_2}$	8.5	8	$g_{n,t,SL}$	8e-03	4e-03	1e-04	0.0159	0.046	1
	8.5	7.5	$g_{n,t,SL}$	0.0216	7.1e-03	7.7e-03	0.0356	2e-03	1
	8.5	7	$g_{n,t,SL}$	0.041	0.0125	0.0166	0.0655	1e-03	1
	8	7.5	$g_{n,t,SL}$	0.0136	6.5e-03	1e-03	0.0263	0.035	1
	8	7	$g_{n,t,SL}$	0.033	0.0124	8.6e-03	0.0574	8e-03	1
	7.5	7	$g_{n,t,SL}$	0.0194	0.0114	-3e-03	0.0418	9e-02	1

Figure 5: Distribution of the IP weights assigned to failures observed under any of the four TI initiation strategies in the unbounded IPW estimation approach. Each plot represents the distribution of the IP weights derived with the estimators g_n , $g_{n,t}$, $g_{n,t,x}$, and $g_{n,t,SL}$ in the interval indicated by the labels on the x axis. These intervals are defined by the 25th, 50th, 75th, 99th, and 99.9th percentiles of the IP weights derived with the estimator g_n .

estimated standard error of the hazard-based bounded IPW estimator generally decreases as the PS estimation approach becomes more nonparametric. This apparent gain in estimation efficiency is explained by the concentration of the IPW weights and the decrease in the proportion of large weights from progressively more flexible PS estimation approaches as shown on Figure 6. We note that these results contradict the position that model-based PS estimation is preferable in practice because data-

Figure 6: Distribution of the IP weights assigned to failures observed under any of the four TI initiation strategies in the bounded IPW estimation approach. Each plot represents the distribution of the IP weights derived with the estimators g_n , $g_{n,t}$, $g_{n,t,x}$, and $g_{n,t,SL}$ in the interval indicated by the labels on the x axis. These intervals are defined by the 25th, 50th, 75th, 99th, and 99.9th percentiles of the IP weights derived with the estimator g_n .

adaptive PS estimation leads to larger weights and thus an increase in IPW estimation variability by revealing practical violations of the positivity assumption. Estimation of the IP weights based on arbitrarily specified parametric models might then be viewed as an implicit weight truncation scheme that restricts the proportion of large weights through smoothing with a misspecified model. On the contrary, the results in this report suggest that model-based estimation of the weights can lead to artificial violation of the positivity assumption in practice (large weights due to model misspecification) when the positivity assumption is in truth not violated. The practical consequence of such model-based estimation is not only biased inference but also increased uncertainty that both could be avoided with data-adaptive PS estimation.

Figure 7: Distribution of the IP weights assigned to failures observed under any of the four TI initiation strategies in the hazard-based bounded IPW estimation approach. Each plot represents the distribution of the IP weights derived with the estimators g_n , $g_{n,t}$, $g_{n,t,x}$, and $g_{n,t,SL}$ in the interval indicated by the labels on the x axis. These intervals are defined by the 25th, 50th, 75th, 99th, and 99.9th percentiles of the IP weights derived with the estimator g_n .

5 Discussion

Results from the secondary analyses in this report illustrate real-world practical consequences of the choice of estimation strategy for the PS in CER based on IPW estimation. We demonstrated that data-adaptive estimation of the PS based on SL can result in substantial bias reduction and gain in estimation efficiency compared to model-based estimation of the PS. These observed gains in efficiency are consistent with known theoretical results that have shown that ignoring information on nuisance parameters (e.g.,

estimating PS even when they are known by study design) can improve the efficiency of the causal effect estimators [81, Section 2.3.7]. The improved estimation performance with SL comes at cost of an increase in computing time. Indeed, the derivation of the PS estimators g_n , $g_{n,t}$, $g_{n,t,\times}$, and $g_{n,t,SL}$ required approximately 10 minutes, 15 minutes, 40 minutes, and 29 hours of computing time in our implementation that is based on an early version of SL [79]. We note that more recent versions of software implementing SL that make use of parallel computing may greatly decrease computing times. Alternate and in particular faster data-adaptive estimation approaches including other ensemble learning methods could be substituted for SL but, to our knowledge, the application of any such alternatives could not be theoretically validated by formal finite sample and asymptotic results such as the ones established for the super learning methodology.

We note that, with data-adaptive PS estimation, the IPW estimators in this report may no longer be asymptotically linear with the influence curves used to derive the conservative variance estimators presented earlier. The application of these variance estimators in practice may thus not be justified theoretically when the PS are estimated data-adaptively. Recent research [82] to derive valid inference when nuisance parameters are estimated with SL led to the development of a new IPW estimation approach that is asymptotically linear with a known influence curve. This approach is currently applicable in point treatment studies only. Additional research is needed to generalize this IPW approach to studies with time-varying exposures and covariates such as the CER study that motivated the present report. Until this generalization is available, one may argue that IPW estimation in real-world CER with time-varying interventions should rely on model-based PS estimation because theoretically valid inferences can then be derived in practice using the conservative variance estimators presented in this report. We counter-argue that these variance estimators are only valid when PS are estimated based on correct models and that the asymptotic coverage of the resulting confidence intervals is 0 if these models are incorrect.

Finally, we note that the variance estimator corresponding with the hazard-based bounded IPW estimator of a RD $\psi_{\theta_1, \theta_2}$ using a saturated MSM presented in Section 3.1.3 generalizes to hazard-based IPW estimation of such RD using a non-saturated (dynamic or static) MSM. To our knowledge, these closed-form expressions for variance estimation were not described previously in published work despite their high practical relevance to avoid bootstrap-based inference.

Acknowledgements: This project was funded under Contract No. HHSA29020050016I from the Agency for Healthcare Research and Quality, US Department of Health and Human Services as part of the Developing Evidence to Inform Decisions about Effectiveness (DEcIDE) program. The authors of this report are responsible for its content. Statements in the report should not be construed as endorsement by the Agency for Healthcare Research and Quality or the US Department of Health and Human Services. This project was also supported by the NIH grant R01 AI074345-07.

Appendix

We denote a working logistic MSM for the counterfactual hazards $\alpha^{t,\theta} \equiv P(Y_{d\theta}(t+1)=1|Y_{d\theta}(t)=0)$ with k coefficients $\beta = (\beta_j)_{j=0,\dots,k}$ by $m(t, \theta|\beta)$, i. e.:

$$m(t, \theta|\beta) = \frac{1}{1 + \exp(-l(t, \theta|\beta))},$$

where the linear part of the working model is denoted by $l(t, \theta|\beta)$ and defined as:

$$l(t, \theta|\beta) \equiv \sum_{j=0}^k \beta_j f_j(t, \theta),$$

for a given set of k functions f_j . For example, in Section 3.1.3, we have $f_j(t, \theta) \equiv I(t = \tau, \theta = k)$ for $(\tau, k) \in \mathcal{T} \times \Theta$ and $j = 1, \dots, k \equiv \text{Card}(\mathcal{T} \times \Theta)$ with $\mathcal{T} \equiv \{0, 1, 2, \dots, K\}$ and $\Theta = \{7; 7.5; 8; 8.5\}$.

By definition of a working MSM [56], the target parameter β_0 (a k -dimensional vector) is defined as follows:

$$\beta_0 \equiv \arg \min_{\beta} \sum_t \sum_{\theta} (E(Y_{d_{\theta}}(t+1) | Y_{d_{\theta}}(t) = 0) - m(t, \theta | \beta))^2 h(t, \theta),$$

for a given user-specified function $h(t, \theta)$. Here, we adopt the choice commonly made with logistic MSM:

$$h(t, \theta) = \frac{\lambda(t, \theta)}{m(t, \theta | \beta)(1 - m(t, \theta | \beta))},$$

where $\lambda(t, \theta)$ represents the user-specified numerator of the stabilized IP weights involved in IPW estimation of β_0 using standard logistic regression software (e.g., $\lambda(t, \theta) \equiv P(\bar{A}(t) = d_{\theta}(\bar{L}(t))) = \prod_{j=0}^t P(A(j) = d_{\theta}(\bar{L}(j)) | \bar{A}(j-1) = d_{\theta}(\bar{L}(j-1)))$). Indeed, with such a choice of function $h(t, \theta)$, an IPW estimator β_n for β_0 can easily be implemented using standard weighted logistic regression software and is defined as the solution of the following unbiased estimating function [42, 43]:

$$\sum_t \sum_{\theta} \frac{I(Y(t) = 0) I(\bar{A}(t) = d_{\theta}(\bar{L}(t)))}{\prod_{j=0}^t P(A(j) | \bar{L}(j), \bar{A}(j-1))} \frac{\lambda(t, \theta)(Y - m(t, \theta | \beta))}{m(t, \theta | \beta)(1 - m(t, \theta | \beta))} \frac{dm(t, \theta | \beta)}{d\beta},$$

which we denote by $D(O|\beta)$ ($a k \times 1$ matrix). We note that the following equalities hold:

$$m(t, \theta | \beta)(1 - m(t, \theta | \beta)) = \frac{\exp(-l(t, \theta | \beta))}{(1 + \exp(-l(t, \theta | \beta)))^2}, \quad (2)$$

$$\frac{dm(t, \theta | \beta)}{d\beta} = \frac{dl(t, \theta | \beta)}{d\beta} \frac{\exp(-l(t, \theta | \beta))}{(1 + \exp(-l(t, \theta | \beta)))^2}. \quad (3)$$

From these equalities, we have:

$$D(O|\beta) = \sum_t \sum_{\theta} \frac{I(Y(t) = 0) I(\bar{A}(t) = d_{\theta}(\bar{L}(t))) \lambda(t, \theta)}{\prod_{j=0}^t P(A(j) | \bar{L}(j), \bar{A}(j-1))} \frac{dl(t, \theta | \beta)}{d\beta} (Y - m(t, \theta | \beta)).$$

The resulting IPW estimator β_n is asymptotically linear with influence curve denoted by $IC(O|\beta)$ and defined as follows [81]:

$$\sqrt{n}(\beta_n - \beta) = \frac{1}{\sqrt{n}} \sum_{i=1}^n IC(O_i | \beta) + o(1)$$

with

$$IC(O|\beta) = -E \left(\frac{dD(O|\beta)}{d\beta'} \right)^{-1} D(O|\beta),$$

where we use the general notation X' for the transpose of any given matrix X . We denote the $k \times k$ matrix $-E \left(\frac{dD(O|\beta)}{d\beta'} \right)^{-1}$ by C and we have:

$$C = E \left(\sum_t \sum_{\theta} \frac{I(Y(t) = 0) I(\bar{A}(t) = d_{\theta}(\bar{L}(t))) \lambda(t, \theta)}{\prod_{j=0}^t P(A(j) | \bar{L}(j), \bar{A}(j-1))} \frac{\exp(-l(t, \theta | \beta))}{(1 + \exp(-l(t, \theta | \beta)))^2} \frac{dl(t, \theta | \beta)}{d\beta} \left(\frac{dl(t, \theta | \beta)}{d\beta} \right)' \right)^{-1}.$$

A substitution estimator $S_{d_{\theta}, n}(t)$ of the counterfactual survival probability $S_{d_{\theta}}(t) = 1 - P(Y_{d_{\theta}}(t+1) = 1)$ can be derived based on the IPW estimator β_n as follows: $S_{d_{\theta}, n}(t) = \prod_{j=0}^t (1 - m(j, \theta | \beta_n))$. This estimator is asymptotically linear with influence curve denoted by $IC_S(O|\beta)$ that can be derived based on the influence curve $IC(O|\beta)$ of the estimator β_n using the delta method as follows [57]:

$$\begin{aligned}
IC_S(O|\beta) &= \frac{d \prod_{j=0}^t (1 - m(j, \theta | \beta))}{d\beta'} IC(O|\beta) \\
&= \left(\prod_{j=0}^t (1 - m(j, \theta | \beta)) \right) \left(\sum_{j=0}^t \frac{\frac{d(1 - m(j, \theta | \beta))}{d\beta'}}{1 - m(j, \theta | \beta)} \right) IC(O|\beta) \\
&\stackrel{(2),(3)}{=} -S_{d_\theta}(t) \left(\sum_{j=0}^t m(j, \theta | \beta) \frac{dl(j, \theta | \beta)}{d\beta'} \right) IC(O|\beta) \\
&= -S_{d_\theta}(t) \left(\sum_{j=0}^t \frac{1}{1 + \exp(-l(j, \theta | \beta))} \frac{dl(j, \theta | \beta)}{d\beta'} \right) IC(O|\beta)
\end{aligned}$$

Similarly, a substitution estimator $\psi_n^{\theta_1, \theta_2}(t)$ of the cumulative risk difference $\psi_n^{\theta_1, \theta_2}(t) = P(Y_{d_{\theta_1}}(t+1)=1) - P(Y_{d_{\theta_2}}(t+1)=1)$ can be derived based on the IPW estimator β_n as follows: $\psi_n^{\theta_1, \theta_2}(t) = \prod_{j=0}^t (1 - m(j, \theta_2 | \beta_n)) - \prod_{j=0}^t (1 - m(j, \theta_1 | \beta_n))$. This estimator is asymptotically linear with influence curve denoted by $IC_\psi(O|\beta)$ that can be derived based on the influence curve $IC(O|\beta)$ of the estimator β_n using the delta method as follows:

$$IC_\psi(O|\beta) = \frac{d \left(\prod_{j=0}^t (1 - m(j, \theta_2 | \beta)) - \prod_{j=0}^t (1 - m(j, \theta_1 | \beta)) \right)}{d\beta'} IC(O|\beta).$$

Using the same derivations as above, we get:

$$IC_\psi(O|\beta) = \left(S_{d_{\theta_1}}(t) \sum_{j=0}^t \frac{1}{1 + \exp(-l(j, \theta_1 | \beta))} \frac{dl(j, \theta_1 | \beta)}{d\beta'} - S_{d_{\theta_2}}(t) \sum_{j=0}^t \frac{1}{1 + \exp(-l(j, \theta_2 | \beta))} \frac{dl(j, \theta_2 | \beta)}{d\beta'} \right) IC(O|\beta).$$

A consistent estimator of the variance of the estimator β_n , $S_{d_{\theta_i}, n}(t)$, and $\psi_n^{\theta_1, \theta_2}$ described above is given by $\frac{1}{n^2} \sum_{i=1}^n (IC(O_i | \beta_n))^2$, $\frac{1}{n^2} \sum_{i=1}^n (IC_S(O_i | \beta_n))^2$ and $\frac{1}{n^2} \sum_{i=1}^n (IC_\psi(O_i | \beta_n))^2$, respectively. When the probabilities $P(A(j) | \bar{L}(j), \bar{A}(j-1))$ involved in these closed-form expressions are unknown and replaced instead with efficient estimates, these variance estimators are conservative in the sense that they over-estimate the estimators' true variance [81].

References

- Robins JM. Association, causation and marginal structural models. *Synthese* 1999;121:151–79.
- McCaffrey DF, Ridgeway G, Morral AR. Propensity score estimation with boosted regression for evaluating causal effects in observational studies. *Psychol Methods* 2004;9:403–25.
- Ridgeway G, McCaffrey DF. Comment: Demystifying double robustness: a comparison of alternative strategies for estimating a population mean from incomplete data. *Stat Sci* 2007;22:540–3.
- Petersen ML, Wang Y, van der Laan MJ, Guzman D, Riley E, Bangsberg DR. Pillbox organizers are associated with improved adherence to HIV antiretroviral therapy and viral suppression: a marginal structural model analysis. *Clin Infect Dis* 2007;45:908–15.
- Moore K, Neugebauer R, Lurmann F, et al. Ambient ozone concentrations cause increased hospitalizations for asthma in children: an 18-year study in Southern California. *Environ Health Perspect* 2008;116:1063–70.
- Moore K, Neugebauer R, Lurmann F, et al. Ambient ozone concentrations and cardiac mortality in Southern California 1983–2000: application of a new marginal structural model approach. *Am J Epidemiol* 2010;171:1233–43.
- Lippman SA, Shade SB, Hubbard AE. Inverse probability weighting in sexually transmitted infection/human immunodeficiency virus prevention research: methods for evaluating social and community interventions. *Sex Transm Dis* 2010;37:512–18.
- Westreich D, Lessler J, Funk MJ. Propensity score estimation: neural networks, support vector machines, decision trees (CART), and meta-classifiers as alternatives to logistic regression. *J Clin Epidemiol* 2010;63:826–33.
- Lee BK, Lessler J, Stuart EA. Improving propensity score weighting using machine learning. *Stat Med* 2010;29:337–46.
- Hernan MA, Brumback B, Robins JM. Marginal structural models to estimate the causal effect of zidovudine on the survival of HIV-positive men. *Epidemiology* 2000;11:561–70.
- Cole SR, Hernan MA, Robins JM, et al. Effect of highly active antiretroviral therapy on time to acquired immunodeficiency syndrome or death using marginal structural models. *Am J Epidemiol* 2003;158:687–94.
- Bodnar LM, Davidian M, Siega-Riz AM, Tsiatis AA. Marginal structural models for analyzing causal effects of time-dependent treatments: an application in perinatal epidemiology. *Am J Epidemiol* 2004;159:926–34.

13. Tager IB, Haight T, Sternfeld B, Yu Z, van Der Laan M. Effects of physical activity and body composition on functional limitation in the elderly: application of the marginal structural model. *Epidemiology* 2004;15:479–93.
14. van der Laan MJ, Polley EC, Hubbard AE. Super learner. *Stat Appl Genet Molecular Biol* 2007;6: Article 25.
15. Neugebauer R, Chandra M, Paredes A, Graham DJ, McCloskey C, Go AS, et al. Modeling approach with super learning for a study on oral bisphosphonate therapy and atrial fibrillation. *J Causal Infer* 2013;1:21–50.
16. Neugebauer R, Fireman B, Roy JA, Raebel Marsha A, Nichols Gregory A, O'Connor PJ. Super learning to hedge against incorrect inference from arbitrary parametric assumptions in marginal structural modeling. *J Clin Epidemiol* 2013;66:S99–109.
17. Gruber S, Logan RW, Jarrin I, Monge S, Hernan MA. Ensemble learning of inverse probability weights for marginal structural modeling in large observational datasets. *Stat Med* 2015;34:106–17.
18. Neugebauer R, Schmittdiel JA, Zhu Z, Rassen JA, Seeger JD, Schneeweiss S. High-dimensional propensity score algorithm in comparative effectiveness research with time-varying interventions. *Stat Med* 2014;34:753–81.
19. Pirracchio R, Petersen ML, van der Laan M. Improving propensity score estimators' robustness to model misspecification using super learner. *Am J Epidemiol* 2015;181:108–19.
20. Neugebauer R, Schmittdiel JA, van der Laan MJ. Targeted learning in real-world comparative effectiveness research with time-varying interventions. *Stat Med* 2014;33:2480–520.
21. Neugebauer R, Schmittdiel JA, Zhu Z, Rassen JA, Seeger JD, Schneeweiss S. High-dimensional propensity score algorithm in comparative effectiveness research with time-varying interventions. *Stat Med* 2014;34:753–81.
22. Nathan DM, Buse JB, Davidson MB, et al. Management of hyperglycemia in Type 2 diabetes: a consensus algorithm for the initiation and adjustment of therapy: a consensus statement from the american diabetes association and the european association for the study of diabetes. *Diab Care* 2006;29:1963–72.
23. Skyler JS, Bergenstal R, Bonow RO, et al. Intensive glycemic control and the prevention of cardiovascular events: implications of the ACCORD, ADVANCE, and VA diabetes trials: a position statement of the american diabetes association and a scientific statement of the american college of cardiology foundation and the american heart association. *Diab Care* 2009;32:187–92.
24. The Diabetes Control and Complications Trial Research Group. The effect of intensive treatment of diabetes on the development and progression of longterm complications in insulin-dependent diabetes mellitus. *N Engl J Med* 1993;329:977–86.
25. Nathan DM, Cleary PA, Backlund JY, et al. Diabetes control and complications trial/epidemiology of diabetes interventions and complications (DCCT/EDIC) study research group. Intensive diabetes treatment and cardiovascular disease in patients with type 1 diabetes. *N Engl J Med* 2005;22:2643–53.
26. Prospective Diabetes UK, Study (UKPDS) Group. Effect of intensive blood-glucose control with metformin on complications in overweight patients with type 2 diabetes (UKPDS 34). *Lancet* 1998;352:854–65.
27. Holman RR, Paul SK, Bethel MA, Matthews DR, Neil HA. 10-year follow-up of intensive glucose control in type 2 diabetes. *N Engl J Med* 2008;359:1577–89.
28. Action to Control Cardiovascular Risk in Diabetes Study Group. Effects of intensive glucose lowering in type 2 diabetes. *N Engl J Med* 2008;358:2545–9.
29. Collaborative Group ADVANCE. Intensive blood glucose control and vascular outcomes in patients with type 2 diabetes. *N Engl J Med* 2008;358:2560–2.
30. William D, Carlos A, Moritz T, et al. Glucose Control and Vascular Complications in Veterans with Type 2 Diabetes. *N Engl J Med* 2009;360:129–39.
31. Ray KK, Seshasai SR, Wijesuriya S, et al. Effect of intensive control of glucose on cardiovascular outcomes and death in patients with diabetes mellitus: a meta-analysis of randomised controlled trials. *Lancet* 2009;373:1765–72.
32. Gerstein HC, Miller M E, Byington RP, et al. Effects of intensive glucose lowering in type 2 diabetes. *N Engl J Med* 2008;358:2545–59.
33. Patel A, MacMahon S, Chalmers J, et al. Intensive blood glucose control and vascular outcomes in patients with type 2 diabetes. *N Engl J Med* 2008;358:2560–72.
34. Duckworth W, Abraira C, Moritz T, et al. Glucose control and vascular complications in veterans with type 2 diabetes. *N Engl J Med* 2009;360:129–39.
35. Ismail-Beigi F, Craven T, Banerji MA, et al. Effect of intensive treatment of hyperglycaemia on microvascular outcomes in type 2 diabetes: an analysis of the ACCORD randomised trial. *Lancet* 2010;376:419–30.
36. O'Connor PJ, Ismail-Beigi F. Near-normalization of glucose and microvascular diabetes complications: data from ACCORD and ADVANCE. *Therap Adv Endocrinol Metabol* 2011;2:17–26.
37. Third Report of the National Cholesterol Education Program (NCEP) Expert Panel on Detection 2002.
38. Nathan DM, Buse JB, Davidson MB, et al. Medical management of hyperglycemia in type 2 diabetes: a consensus algorithm for the initiation and adjustment of therapy: a consensus statement of the American Diabetes Association and the European Association for the Study of Diabetes. *Diabetes Care* 2009;32:193–203.
39. Vogt TM, Elston-Lafata J, Tolsma D, Greene SM. The role of research in integrated healthcare systems: the HMO Research Network. *Am J Manag Care* 2004;10:643–8.

40. Robins JM. Marginal Structural Models in 1997 Proceedings of the American Statistical Association, Section on Bayesian Statistical Science, 1998, 1–10.
41. Murphy SA, van der Laan MJ, Robins JM. Marginal mean models for dynamic treatment regimens. *J Am Stat Assoc* 2001;96:1410–24.
42. van der Laan MJ, Petersen ML. Causal effect models for realistic individualized treatment and intention to treat rules. *Int J Biostat* 2007;3:Article 3.
43. Robins J, Rotnitzky A. Estimation and extrapolation of optimal treatment and testing strategies. *Stat Med* 2008;27:4678–721.
44. Hernan MA, Lanoy E, Costagliola D, Robins JM. Comparison of dynamic treatment regimes via inverse probability weighting. *Basic Clin Pharmacol Toxicol* 2006;98:237–42.
45. Cain LE, Robins JM, Lanoy E, Logan R, Costagliola D, Hernan MA. When to start treatment? A systematic approach to the comparison of dynamic regimes using observational data. *Int J Biostat* 2010;6. Article 18.
46. Hernan MA, Brumback BA, Robins JM. Estimating the causal effect of zidovudine on CD4 count with a marginal structural model for repeated measures. *Stat Med* 2002;21:1689–709.
47. Robins JM, Hernan MA, Brumback B. Marginal structural models and causal inference in epidemiology. *Epidemiology* 2000;11:550–60.
48. Neugebauer R, Fireman B, Roy JA, O'Connor PJ, Selby JV. Dynamic marginal structural modeling to evaluate the comparative effectiveness of more or less aggressive treatment intensification strategies in adults with type 2 diabetes. *Pharmacoepidemiol Drug Saf* 2012;21:99–113.
49. Neugebauer R, Fireman B, Roy JA, O'Connor PJ. Impact of specific glucose-control strategies on microvascular and macrovascular outcomes in 58,000 adults with type 2 diabetes. *Diabetes Care* 2013;36:3510–16.
50. Pearl J. Causal inference in statistics: An overview. *Stat Surv* 2009;3:96–146.
51. VanderWeele TJ. Concerning the consistency assumption in causal inference. *Epidemiology* 2009;20:880–3.
52. Pearl J. On the consistency rule in causal inference: axiom, definition, assumption, or theorem? *Epidemiology* 2010;21:872–5.
53. The HIV-CAUSAL. Collaboration. When to initiate combined antiretroviral therapy to reduce mortality and aids-defining illness in HIV-infected persons in developed countries: an observational study. *Ann Intern Med* 2011;154:509–15.
54. Wang Y, Bembom O, van der Laan MJ. Data-adaptive estimation of the treatment-specific mean Original Research Article. *J Stat Plan Infer* 2007;137:1871–87.
55. Platt RW, Brookhart AM, Cole SR, Westreich D, Schisterman EF. An information criterion for marginal structural models. *Stat Med* 2013;32:1383–93.
56. Neugebauer R, Laan M. Nonparametric causal effects based on marginal structural models. *J Stat Plan Infer* 2006;137:419–34.
57. Vaart AW. Asymptotic statistics. New York NY, USA: Cambridge University Press, 1998.
58. Horvitz DG, Thompson DJ. A generalization of sampling without replacement from a finite universe. *J Am Stat Assoc* 1952;47:663–85.
59. Robins J, Sued M, Lei-Gomez Q, Rotnitzky A. Comment: performance of double-robust estimators when “Inverse Probability” weights are highly variable. *Stat Sci* 2007;22:544–59.
60. Kang JDY, Schafer JL. Demystifying double robustness: a comparison of alternative strategies for estimating a population mean from incomplete data. *Stat Sci* 2007;22:523–39.
61. Robins JM, Rotnitzky A. Recovery of information and adjustment for dependent censoring using surrogate markers. In: Jewell N, Dietz K, Farewell V, editors. *AIDS epidemiology – methodological issues*. Boston, MA: Birkhuser, 1992:297–331.
62. Robins JM, Rotnitzky A, Zhao LP. Estimation of regression coefficients when some regressors are not always observed. *J Am Stat Assoc* 1994;89:846–66.
63. Imbens G. Nonparametric estimation of average treatment effects under exogeneity: A review. *Rev Econ Stat* 2004;6:4–29.
64. Hernan MA. The hazards of hazard ratios. *Epidemiology* 2010;21:13–15.
65. Platt RW, Schisterman EF, Cole SR. Time-modified confounding. *Am J Epidemiol* 2009;170:687–94.
66. Breiman L. Random Forests. *Machine Learning* 2001;45:5–32.
67. Efron B, Hastie T, Johnstone I, Tibshirani R. Least angle regression. *Ann Stat* 2004;32:407–99.
68. Ruczinski I, Kooperberg C, LeBlanc M. Logic regression. *J Comput Graph Stat* 2003;12:475–511.
69. Breiman L, Friedman JH, Olshen RA, Stone CJ. Classification and regression trees. Monterey: Wadsworth & Brooks/Cole, 1984.
70. Hoerl AE, Kennard RW. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics* 1970;12:55–67.
71. Friedman JH. Multivariate adaptive regression splines. *Ann Stat* 1991;19:1–141.
72. Kooperberg C, Bose S, Stone CJ. Polychotomous regression. *J Am Stat Assoc* 1997;92:117–27.
73. Sinisi SE, van der Laan MJ. Deletion/substitution/addition algorithm in learning with applications in genomics. *Stat Appl Genet Mol Biol* 2004;3:Article18.
74. Schneeweiss S, Rassen JA, Glynn RJ, Avorn J, Mogun H, Brookhart MA. High-dimensional propensity score adjustment in studies of treatment effects using health care claims data. *Epidemiology* 2009;20:512–22.
75. Dudoit S, van der Laan MJ. Asymptotics of cross-validated risk estimation in estimator selection and performance assessment. *Stat Methodol* 2005;2:131–54.

76. van der Laan MJ, Dudoit S, Keles S. Asymptotic optimality of likelihood-based cross-validation. *Stat Appl Genet Mol Biol* 2004;3: Article 4.
77. van der Vaart AW, Dudoit S, van der Laan MJ. Oracle Inequalities for multi-fold cross-validation. *Stat Decisions* 2006;24:351–71.
78. van der Laan MJ, Dudoit S, van der Vaart AW. The cross-validated adaptive epsilon-net estimator. *Stat Decisions* 2006;24:373–95.
79. Polley EC. SuperLearner R package (version 1.1-18) <https://github.com/ecpolley/SuperLearner> 2011.
80. Neugebauer R, Fireman B, Roy JA, Raebel MA, Nichols GA, O'Connor PJ. Super learning to hedge against incorrect inference from arbitrary parametric assumptions in marginal structural modeling. *J Clin Epidemiol* 2013;66, Suppl:S99–109.
81. van der Laan MJ, Robins JM. Unified methods for censored longitudinal data and causality. New York: Springer, 2003.
82. van der Laan MJ. Targeted estimation of nuisance parameters to obtain valid statistical inference. *Int J Biostat* 2014;10:29–57.