Research Article

Statistics in Medicine

Received XXXX

(www.interscience.wiley.com) DOI: 10.1002/sim.0000

Targeted Learning in real-world comparative effectiveness research with time-varying interventions

Romain Neugebauer[a]*, Julie A. Schmittdiel[a] and Mark J. van der Laan[b]

In comparative effectiveness research (CER), the aim is often to contrast survival outcomes between exposure groups defined by time-varying interventions. With observational data, standard regressions (e.g., Cox regression) cannot account for time-dependent confounders on causal pathways between exposures and outcome nor for time-dependent selection bias that may arise from informative right-censoring. Inverse Probability Weighting (IPW) estimation to fit Marginal Structural Models (MSMs) has been applied most commonly in practice as the solution to properly adjust for both the time-dependent confounding and selection bias that are expected in longitudinal observational studies. In this report, we describe the application and performance of an alternate estimation approach in such a study. The approach is based on the recently proposed Targeted Learning methodology and consists in Targeted Minimum Loss based Estimation (TMLE) with Super Learning (SL) within a nonparametric MSM. The evaluation is based on the analysis of electronic health records data with both IPW estimation and TMLE to contrast cumulative risks under four more or less aggressive strategies for treatment intensification in adults with type 2 diabetes already on 2+ oral agents or basal insulin. Results from previous randomized experiments provide a surrogate gold standard to validate confounding and selection bias adjustment. Bootstrapping is used to validate analytic estimation of standard errors. This report 1) establishes the feasibility of TMLE in real-world CER based on large healthcare databases, 2) provides evidence of proper confounding and selection bias adjustment with TMLE and SL, and 3) motivates their application for improving estimation efficiency in practice.

Keywords: targeted learning; marginal structural model; inverse probability weighting; efficiency gain; comparative effectiveness; diabetes

1. Introduction

In comparative effectiveness research (CER), the aim is often to contrast survival outcomes between exposure groups defined by time-varying interventions. In observational CER studies, these effects are represented by Marginal Structural Models (MSMs) and their investigation is complicated by the time-dependent confounding and informative right-censoring that are expected with longitudinal data. Standard regression techniques (e.g., Cox regression) are inadequate [1, 2] to account for not only time-dependent confounders on causal pathways between the exposures and outcome but also for time-dependent selection bias that may arise from right-censoring [3]. To date, Inverse Probability Weighting

[a] *Division of Research, Kaiser Permanente Northern California, Oakland, CA*
[b] *Division of Biostatistics, School of Public Health, University of California, Berkeley, CA*
* *Correspondence to: Romain.S.Neugebauer@kp.org*

Contract/grant sponsor: 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.

(IPW) estimation has been the solution of choice to fit MSM in real-world CER studies [4, 5, 6, 7, 8, 9] despite the early development of an alternate estimation approach, augmented-IPW (A-IPW) estimation [10, 11, 12, 13, 14], which is both doubly robust and locally efficient. These two properties may translate in practice into 1) more reliable effect estimates since double robustness provides two chances for proper confounding and selection bias adjustment (i.e., inference can remain valid even if the treatment and action mechanisms on which IPW estimation relies are not estimated consistently), and 2) more precise effect estimates compared to IPW estimates and hence the possibility for earlier detection of differential safety or effectiveness signals. The complexity of implementation of A-IPW estimation with time-varying exposures largely explains its limited use in practice [15, 14].

Recently, Targeted Learning [16] was proposed as an alternative to IPW and A-IPW estimation for drawing causal inferences in problems with both point treatment and time-varying interventions of interest. Targeted Learning encompasses a general doubly robust and efficient estimation methodology that is coupled with Super Learning (SL) [17] to data-adaptively estimate the nuisance parameters on which relies estimation of the estimand of interest. While this approach and its asymptotic properties were derived from a formal theoretical framework and mostly tested with simulated and real-word point treatment data, there remains a need for further practical evaluation of this general methodology for applications in problems with time-varying interventions. Early evaluation of Targeted Learning for CER with time-varying exposures based on real-world data from a large healthcare database [18] revealed the practical complexity of a Targeted Maximum Likelihood Estimation algorithm initially proposed for implementation of Targeted Learning with time-varying interventions. Subsequently, an alternate algorithm was developed which greatly simplifies applications of Targeted Learning in problems with time-varying interventions and a limited number of time-varying covariates [19]. Current experience with these algorithms precludes however their routine applications in real-world CER problems that often require to control for medium to high-dimensional time-varying covariates. More recently, van der Laan and Gruber (2012) [20] derived a Targeted Minimum Loss based Estimation (TMLE) algorithm for evaluating the effect of time-varying interventions based on a general Targeted Learning estimation road map [21, 22] applied with a key identifiability result from Bang and Robins (2005) [14]. Compared to previously proposed Targeted Learning algorithms, TMLE further simplifies implementation of Targeted Learning in CER studies with time-varying interventions and in particular if control for medium to high-dimensional time-varying covariates is needed.

Using a real-world CER study, our first objective is to describe the application of this algorithm, evaluate its computation burden and assess its performance with confounding and selection bias adjustment. Motivation for preferring a TMLE approach over an IPW estimation approach in CER includes i) the mitigation of concerns over violation of the assumption of consistent estimation of the treatment and right-censoring mechanisms (double robustness property), and ii) a possible gain in the precision of the effect estimates with non-rare outcomes (efficiency property). Our second objective here is to evaluate the potential for gain in estimation efficiency with TMLE over IPW estimation in practice.

Our evaluation is based on the analysis of electronic health record data with both IPW estimation and TMLE to contrast cumulative risks under four more or less aggressive strategies for treatment intensification in adults with type 2 diabetes on 2+ oral agents or basal insulin. Results from previous randomized experiments provide a surrogate gold standard to validate confounding and selection bias adjustment. Bootstrapping is used to validate analytic estimation of standard errors on which is based the evaluation of potential efficiency gains with TMLE and SL.

2. Evaluation with a real-world CER study

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.

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.

Current recommendations specify target hemoglobin A1c of $<7\%$ for most patients [23, 24]. However, evidence supporting the effectiveness of a blanket recommendation is inconsistent across several outcomes [25, 26, 27, 28, 29, 30, 31, 32], especially when intensive anti-diabetic therapy is required. In 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 [33, 34, 35], 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 [36, 37] to support the hypothesis [38, 39] 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. In addition, no additional major trials addressing these questions are underway.

For these reasons, using the electronic health records (EHR) from patients of seven sites of the HMO Research Network [40], a large retrospective cohort study of adults with T2DM was assembled 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 dynamic MSM [41, 42, 43, 44] was fitted using IPW estimation [41, 45, 46] 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 her A1c level reaches or drifts above θ% and patient remains on the intensified therapy thereafter' with θ =7, 7.5, 8, or 8.5.

Details of the study design, analytic approach, and results are described elsewhere [47, 48]. 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 to contrast the same four counterfactual survival curves indexed by the TI strategies described above for the purpose of evaluating the performance of TMLE compared to IPW estimation. We now formally describe the observational data and the parameter of interest before describing the TMLE approach and its IPW analog.

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 to 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 follow-up time $\tilde{T}$ is equal to the failure time T is denoted by $\Delta = I(\tilde{T} = T)$. At each time point $t = 0, \ldots, \tilde{T}$, the patient's exposure to an intensified diabetes treatment is represented by the binary variable $A_1(t)$, and the patient's right-censoring status is denoted by the indicator variable $A_2(t) = I(C \leq t)$. The combination $A(t) = (A_1(t), A_2(t))$ is referred to as the action at time t. At each time point $t = 0, \ldots, \tilde{T}$, covariates (e.g., A1c measurements) 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, $(A(t), A(t+1), \ldots)$. In particular, the covariates at time t include an outcome measurement denoted by $Y(t)$, i.e., $Y(t) \in L(t)$ for $t = 0, \ldots, \tilde{T}$. For each time point $t = 0, \ldots, \tilde{T}+1$, 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 $t = 0, \ldots, \tilde{T}$, missing at $t = \tilde{T}+1$ if $\Delta = 0$ and, 1 at $t = \tilde{T}+1$ if $\Delta = 1$. 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 $\bar{A}(t) = (A(0), \ldots, A(t))$. Following the MSM framework [41], 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 $i = 1, \ldots, n$. The longest observed follow-up time is $\max_{i=1,\ldots,n} \tilde{T}_i = 36$ (9 years). Details about the approach implemented for mapping EHR data into the coarsened exposure, covariate and outcome data for each patient was described elsewhere [47, 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 the the 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 [43] defined as a vector function $d_\theta = (d_\theta(0), \ldots, d_\theta(t_0))$ where each function, $d_\theta(t)$ for $t = 0, \ldots, t_0$, 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., $\bar{A}(t-1) = 0$) 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.

The parameter of interest denoted by ψ^{θ_1,θ_2} is the difference between the cumulative risks at time t_0 associated with any two distinct treatment strategies d_{θ_1} and d_{θ_2}:

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

For conciseness, we refer the reader to earlier work [47, Appendices B and D] for a 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 [47, Appendices C]: no unmeasured confounders, positivity, and consistent estimation of the action mechanism. 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 [49], then an additional assumption referred to as consistency assumption is made [50, 51]. Under these identifiability assumptions, the causal parameter ψ^{θ_1,θ_2} can be expressed as a statistical parameter, i.e., a parameter of the *observed* data distribution (as opposed to the counterfactual data distribution).

In addition, a more or less flexible non-saturated MSM may be assumed [4, 52, 53, 54, 5]. The assumption encoded by such an MSM typically imposes constraints on the survival curves that underlie the definition of the parameter of interest ψ^{θ_1,θ_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 since 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 [55] and are still being researched [56].

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.

3. An alternative to Inverse Probability Weighting estimation

Targeted Learning was proposed as a general approach for estimating low dimensional parameters of the unknown probability distribution that underlies the observed data collected in an observational or randomized study. In the Targeted Learning literature, the parameter of interest is referred to as the 'target parameter'. In particular, Targeted Learning can be used to estimate the target parameter ψ^{θ_1,θ_2} under the identifiability assumptions discussed earlier and is then an alternative to IPW estimation of MSM parameters. Similar to other estimation methodologies, estimation with Targeted Learning is based on a statistical model that restricts the set of possible data generating distributions. In particular, such a model often imposes constraints on nuisance parameters (e.g., propensity scores), i.e., parameters that are not of interest but that are used as the building blocks for deriving an estimate of the target parameter. Contrary to standard practice with other estimation approaches, a tenet of Targeted Learning is that estimation should be conducted based on 'honest' statistical models in practice. By 'honest', we mean a model that is specified based on *knowledge* as opposed to models that are specified for *convenience*. In most if not all real-world studies, little subject-matter knowledge is available to honestly restrict the set of possible data generating distributions. Estimation with Targeted Learning is thus typically based on a nonparametric model and machine learning is then used for estimating nuisance parameters instead of an arbitrarily specified model that would likely lead to incorrect study findings. The general Targeted Learning methodology for estimating the target parameter based on machine learning estimates of the nuisance parameters is referred to as

Targeted Minimum Loss based Estimation[†]. More specifically, Targeted Learning is the combination of TMLE with a particular machine learning algorithm discussed later and called Super Learning. TMLE is an estimation methodology based on the 'substitution principle': the target parameter (denoted by ψ) is expressed as a mapping (denoted by Ψ) of a component (denoted by Q) of the data generating distribution, i.e. $\psi = \Psi(Q)$, and a substitution estimator of ψ may then be defined as $\Psi(Q_n)$ for any given estimator Q_n of Q. TMLE involves two steps for estimating the nuisance parameter Q. First, an initial estimator Q_n of Q is implemented. Second, this initial estimator is updated for the purpose of optimizing the bias variance trade-off for the target parameter. The updated estimator of Q is denoted by Q_n^* and is derived based on another component of the data generating distribution (denoted by g and referred to as the action mechanism). TMLE is defined as estimation with the substitution estimator $\Psi(Q_n^*)$. TMLE is mathematically devised such that $\Psi(Q_n^*)$ is a doubly robust and possibly efficient estimator and such that valid statistical inference can be derived with this estimator even with data-adaptive estimation of the nuisance parameters. Double robustness means that $\Psi(Q_n^*)$ is a consistent regular asymptotically linear estimator of the target parameter as long as at least one of two nuisance parameters, Q or g, are consistently estimated. When both nuisance parameters, Q and g, are estimated consistently, TMLE is efficient in the sense that $\Psi(Q_n^*)$ attains the semiparametric efficiency bound in a model that potentially includes constraints on g, i.e., the variance of any regular asymptotically linear estimator of ψ in such a model is greater than or equal to the variance of $\Psi(Q_n^*)$. The mathematical derivation of TMLE for any given observed data and target parameter is based on a well defined technical road map [21, 22] that is comprehensively described in van der Laan and Rose (2011) [16] (e.g. chapter 5). A brief description of this road map is described in van der Laan and Gruber (2012) [20] and applied for estimation of target parameters defined by time-varying interventions such as the parameter ψ^{θ_1,θ_2} of interest in this report. Unlike previous application of the technical road map for TMLE, the application in van der Laan and Gruber (2012) relies on a novel representation of the target parameter with a mapping Ψ defined by an iterative sequence of conditional expectations of the outcome of interest which was originally exploited in Bang and Robins (2005) [14]. This innovation results in a more practicable TMLE algorithm. In sections 3.1 and 3.2, we provide an applied description of this TMLE algorithm for estimating ψ^{θ_1,θ_2} based on the new notation below.

To simplify the formal description of the TMLE algorithm below, we adopt the following new definitions and notation. For any given observed covariate history through time t denoted by $\bar{L}(t)$, the action regimen $(a(0) = d_\theta(0)(L(0)), a(1) = d_\theta(1)(\bar{L}(1), a(0)), \ldots, a(t) = d_\theta(t)(\bar{L}(t), \bar{a}(t-1)))$ through time t is denoted by $d_\theta(\bar{L}(t))$. For a patient who experiences failure before t_0 (i.e., when $\Delta = 1$ and $\tilde{T} < t_0$), we extend the definition of her observed data through $t_0 + 1$ by including the outcome variables $Y(t+1) = I(T \leq t-1) = 1$ for $\tilde{T} < t \leq t_0$. With this extension, the observed data structure becomes:

$$O = (\tilde{T}, \Delta, (1-\Delta)\Gamma, \bar{L}(\tilde{T}), \bar{A}(\tilde{T}), \Delta\bar{Y}(\tilde{T}+1, \max(\tilde{T}, t_0)+1)),$$

where $\Delta\bar{Y}(t, t') = (\Delta Y(t), \ldots, \Delta Y(t'))$ with $t \leq t'$. To simplify expressions below, the outcome $Y(t+1)$ for $\tilde{T} \leq t \leq \max(t_0, \tilde{T})$ when $\Delta = 1$ is also denoted with $L(t+1)$ and the observed data can thus be expressed as:

$$O = (\tilde{T}, \Delta, (1-\Delta)\Gamma, \bar{L}(\tilde{T}), \bar{A}(\tilde{T}), \Delta\bar{L}(\tilde{T}+1, \max(\tilde{T}, t_0)+1)).$$

Finally, we define $\check{T}(t) = \min(\tilde{T}, t)$ for $t = 0, \ldots, t_0$ and the cumulative counterfactual risk $P(Y_{d_\theta}(t_0+1) = 1)$ for any given θ is denoted by γ^θ.

The proposed Targeted Learning approach for estimating ψ^{θ_1,θ_2} consists in estimating each of the two risks γ^{θ_1} and γ^{θ_2}, separately, with the TMLE algorithm described in section 3.1 that was derived based on the following iterative sequence of conditional expectations for representing γ^θ:

$$\gamma^\theta = E\Bigg(E\Big[\ldots E\Big(E\Big[E\Big(Y(t_0+1) \mid \mathcal{F}(t_0)\Big) \mid \mathcal{F}(t_0-1)\Big] \mid \mathcal{F}(t_0-2)\Big) \ldots \mid \mathcal{F}(0)\Big]\Bigg)$$

with $\mathcal{F}(t) = (\bar{A}(\check{T}(t)) = d_\theta(\bar{L}(\check{T}(t))), \bar{L}(t))$. An estimate of the risk difference (RD) of interest, ψ^{θ_1,θ_2}, is then derived by taking the difference between the two resulting estimators denoted by $\gamma_n^{\theta_1,*}$ and $\gamma_n^{\theta_2,*}$, respectively. Inference for the RD is derived based on the influence curve of these two estimators and the delta method [57] as described in section 3.2.

3.1. Point estimation with TMLE

The following TMLE algorithm was adapted from the algorithm proposed by van der Laan and Gruber [20]. Each step below is implemented sequentially for a given θ to estimate the cumulative risk γ^θ:

1. Estimate $P(A(t) = d_\theta(\bar{L}(t)) \mid \bar{L}(t), \bar{Y}(t) = 0, \bar{A}(t-1) = d_\theta(\bar{L}(t-1)))$ denoted by $g^\theta_{A(t)}$ for $t = 0, \ldots, t_0$.

[†]The methodology is also referred to as Targeted Maximum Likelihood Estimation when it is based on the log-likelihood loss function [19].

For each t, $g^{\theta}_{A(t)}$ represents the conditional probability that a patient's exposure and right-censoring status at time t remain concordant with the action implied by the decision rule d_θ given i) that the patient did not fail before t, ii) that her past actions are concordant with action decisions according to rule d_θ, and iii) her past observed covariates $\bar{L}(t)$. For this report, several approaches to estimate $g^{\theta}_{A(t)}$ were implemented but all are based on separate estimation of each element of the factorization of the action mechanism at time t, i.e., $P(A(t) \mid \bar{L}(t), \bar{Y}(t) = 0, \bar{A}_1(t-1), \bar{A}_2(t-1) = 0)$. The same factorization is typically used to estimate the denominator of the weights in IPW estimation. Specifically, the following probabilities were estimated separately with one of several approaches detailed later:

- Propensity score (PS) for TI initiation denoted by μ_1:

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

- PS for TI continuation denoted by μ_2:

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

- PS for right-censoring by administrative end of study denoted by μ_3:

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

  where $I(\cdot)$ denotes an indicator variable
- PS for right-censoring by disenrollment from the health plan denoted by μ_4:

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

- PS for right-censoring by death denoted by μ_5:

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

For patients who followed rule d_θ through t (i.e., for whom $\bar{A}(t) = d_\theta(\bar{L}(t))$), an estimate of the nuisance parameter $g^{\theta}_{A(t)}$ can be derived from estimates of these 5 PS based on the following equality implied by factorizing the action mechanism at time t using the chain rule:

$$g^{\theta}_{A(t)} = \big(I(\bar{A}_1(t-1) = 0)\mu_1^{A_1(t)}(1-\mu_1)^{1-A_1(t)} + I(A_1(t-1) = 1)\mu_2^{A_1(t)}(1-\mu_2)^{1-A_1(t)}\big)(1-\mu_3)(1-\mu_4)(1-\mu_5). \quad (1)$$

The estimate of $g^{\theta}_{A(t)}$ is denoted by $g^{\theta}_{A(t),n}$.

2. Derive an initial estimate of $E(Y(t_0+1) \mid \bar{A}(\check{T}(t_0)) = d_\theta(\bar{L}(\check{T}(t_0))), \bar{L}(t_0))$ denoted by $Q^{\theta}_{L(t_0+1)}(\bar{L}(t_0))$.
   Note that $\bar{L}(t_0+1)$ is always defined in the extended observed data structure when $\bar{A}(\check{T}(t_0)) = d_\theta(\bar{L}(\check{T}(t_0)))$ because $\bar{A}(\check{T}(t_0)) = d_\theta(\bar{L}(\check{T}(t_0)))$ implies either i) $\check{T}(t_0) = t_0$ and $\check{T}(t_0) < \tilde{T}$ or ii) $\check{T}(t_0) = \tilde{T} = T$ and $\check{T}(t_0) \leq t_0$ (since $\check{T}(t_0) = \tilde{T} = C$ is not possible when $\bar{A}(\check{T}(t_0)) = d_\theta(\bar{L}(\check{T}(t_0)))$). The conditional expectation $Q^{\theta}_{L(t_0+1)}(\bar{L}(t_0))$ is thus well defined and we have:

$$Q^{\theta}_{L(t_0+1)}(\bar{L}(t_0)) = 1 + I(\bar{Y}(t_0) = 0)\big(E(Y(t_0+1) \mid \bar{A}(t_0) = d_\theta(\bar{L}(t_0)), \bar{L}(t_0), \bar{Y}(t_0) = 0) - 1\big). \quad (2)$$

   This step thus reduces to the estimation of $E(Y(t_0+1) \mid \bar{A}(t_0) = d_\theta(\bar{L}(t_0)), \bar{L}(t_0), \bar{Y}(t_0) = 0)$, i.e., the conditional probability that a patient experiences the failure event at time t_0 given i) that she experienced no such event previously and no censoring event before and at t_0, ii) that she were continuously treated according to strategy d_θ through t_0, and iii) her covariates through t_0, $\bar{L}(t_0)$. For this report, several approaches to estimate this probability were implemented and are detailed later. All enforce that the estimate lies in the [0,1] interval and all rely solely on data from patients who did not fail before t_0 and who followed rule d_θ through t_0 (i.e., $\bar{Y}(t_0) = 0$ and $\bar{A}(t_0) = d_\theta(\bar{L}(t_0))$). The initial estimate of the nuisance parameter $Q^{\theta}_{L(t_0+1)}(\bar{L}(t_0))$ is denoted by $Q^{\theta}_{L(t_0+1),n}(\bar{L}(t_0))$ and is defined as follows: i) For a patient who did not experience failure before t_0 and who followed rule d_θ through t_0, $Q^{\theta}_{L(t_0+1),n}(\bar{L}(t_0))$ is the estimate of the conditional probability just described, ii) For a patient who did experience failure before t_0 and who followed rule d_θ until failure, $Q^{\theta}_{L(t_0+1),n}(\bar{L}(t_0))$ is set to 1 in accordance with equality (2).
3. Update the initial estimate of $Q^{\theta}_{L(t_0+1)}(\bar{L}(t_0))$.
   This update is implemented by logistic regression for predicting $Y(t_0+1)$ based on an intercept model with an offset variable fitted with weights, and using only data from patients who did not fail before t_0 (i.e., $\bar{Y}(t_0) = 0$) and who

followed rule d_θ through t_0 (i.e., $\bar{A}(t_0) = d_\theta(\bar{L}(t_0))$). The weight and offset associated with the outcome $Y(t_0+1)$ from any patient whose data contribute to this logistic regression are defined as $logit(Q^\theta_{L(t_0+1),n}(\bar{L}(t_0))) = \log\left(\frac{Q^\theta_{L(t_0+1),n}(\bar{L}(t_0))}{1-Q^\theta_{L(t_0+1),n}(\bar{L}(t_0))}\right)$ and $\frac{1}{\prod_{t=0}^{t_0} g^\theta_{A(t),n}}$, respectively. The estimate of the intercept resulting from this weighted logistic regression is denoted by ϵ_n. The updated estimate of $Q^\theta_{L(t_0+1)}(\bar{L}(t_0))$ is denoted by $Q^{\theta,*}_{L(t_0+1),n}(\bar{L}(t_0))$ and is defined as follows: i) For a patient who did not experience failure before t_0 and who followed rule d_θ through t_0, $Q^{\theta,*}_{L(t_0+1),n}(\bar{L}(t_0))$ is $expit[logit(Q^\theta_{L(t_0+1),n}(\bar{L}(t_0))) + \epsilon_n]$ where $expit(t) = \frac{1}{1+\exp(-t)}$. ii) For a patient who did experience failure before t_0 and who followed rule d_θ until failure, $Q^{\theta,*}_{L(t_0+1),n}(\bar{L}(t_0))$ is set to 1 in accordance with equality (2).

4. Repeat the following two steps for $k = t_0 - 1, \ldots, 0$:

   (a) Derive an initial estimate of $E(Q^\theta_{L(k+2)}(\bar{L}(k+1)) \mid \bar{A}(\check{T}(k)) = d_\theta(\bar{L}(\check{T}(k))), \bar{L}(k))$ denoted by $Q^\theta_{L(k+1)}(\bar{L}(k))$.

   Note that $\bar{L}(k+1)$ is always defined in the extended observed data structure when $\bar{A}(\check{T}(k)) = d_\theta(\bar{L}(\check{T}(k)))$ because $\bar{A}(\check{T}(k)) = d_\theta(\bar{L}(\check{T}(k)))$ implies either i) $\check{T}(k) = k$ and $\check{T}(k) < \tilde{T}$ or ii) $\check{T}(k) = \tilde{T} = T$ and $\tilde{T} \leq k$. The conditional expectation $Q^\theta_{L(k+1)}(\bar{L}(k))$ is thus well defined and we have:

   $$Q^\theta_{L(k+1)}(\bar{L}(k)) = 1 + I(\bar{Y}(k) = 0)\big(E(Q^\theta_{L(k+2)}(\bar{L}(k+1)) \mid \bar{A}(k) = d_\theta(\bar{L}(k)), \bar{L}(k), \bar{Y}(k) = 0) - 1\big). \quad (3)$$

   This step thus reduces to the estimation of

   $$E(Q^\theta_{L(k+2)}(\bar{L}(k+1)) \mid \bar{A}(k) = d_\theta(\bar{L}(k)), \bar{L}(k), \bar{Y}(k) = 0), \quad (4)$$

   i.e., the conditional expectation of the continuous measure $Q^\theta_{L(k+2)}(\bar{L}(k+1))$ (itself a conditional expectation between 0 and 1) characterizing a patient at time $k+1$ given i) that she did not experience failure before k and no censoring event before and at k, ii) that she were continuously treated according to strategy d_θ through k, and iii) her baseline and past time-varying covariates through k, $\bar{L}(k)$. For this report, several approaches to estimate this expectation were implemented and are detailed later. All enforce that the estimate lies in the [0,1] interval and all rely solely on data from patients who did not fail before k and who followed rule d_θ through k (i.e., $\bar{Y}(k) = 0$ and $\bar{A}(k) = d_\theta(\bar{L}(k))$). In particular, the datum $Q^{\theta,*}_{L(k+2),n}(\bar{L}(k+1))$ is needed for each of these patients. Among them, some may have followed rule d_θ through $k+1$ and others may only have followed rule d_θ through k. For the first group of patients, we already computed an estimate $Q^{\theta,*}_{L(k+2),n}(\bar{L}(k+1))$ in the latest "update step" while for the second group of patients, such estimates need to be computed here by extrapolation, i.e., using the same protocol employed in the latest "update step" as if these patients also followed rule d_θ at time $k+1$. The initial estimate of the nuisance parameter $Q^\theta_{L(k+1)}(\bar{L}(k))$ is denoted by $Q^\theta_{L(k+1),n}(\bar{L}(k))$ and is defined as follows: i) For a patient who did not experience failure before k and who followed rule d_θ through k, $Q^\theta_{L(k+1),n}(\bar{L}(k))$ is the estimate of conditional expectation (4) above. ii) For a patient who did experience failure before k and who followed rule d_θ until failure, $Q^\theta_{L(k+1),n}(\bar{L}(k))$ is set to 1 in accordance with equality (3).

   (b) Update the initial estimate of $Q^\theta_{L(k+1)}(\bar{L}(k))$.

   This update is implemented by logistic regression for predicting $Q^\theta_{L(k+2)}(\bar{L}(k+1))$ based on an intercept model with an offset variable fitted with weights, and using only data from patients who did not fail before k (i.e., $\bar{Y}(k) = 0$) and who followed rule d_θ through k (i.e., $\bar{A}(k) = d_\theta(\bar{L}(k))$). The weight and offset associated with the outcome $Q^{\theta,*}_{L(k+2),n}(\bar{L}(k+1))$ from any patient whose data contribute to this logistic regression are defined as $logit(Q^\theta_{L(k+1),n}(\bar{L}(k))) = \log\left(\frac{Q^\theta_{L(k+1),n}(\bar{L}(k))}{1-Q^\theta_{L(k+1),n}(\bar{L}(k))}\right)$ and $\frac{1}{\prod_{t=0}^{k} g^\theta_{A(t),n}}$, respectively. The estimate of the intercept resulting from this weighted logistic regression is denoted by ϵ_n. The updated estimate of $Q^\theta_{L(k+1)}(\bar{L}(k))$ is denoted by $Q^{\theta,*}_{L(k+1),n}(\bar{L}(k))$ and is defined as follows: i) For a patient who did not experience failure before k and who followed rule d_θ through k, $Q^{\theta,*}_{L(k+1),n}(\bar{L}(k))$ is $expit[logit(Q^\theta_{L(k+1),n}(\bar{L}(k))) + \epsilon_n]$. ii) For a patient who did experience failure before k and who followed rule d_θ until failure, $Q^{\theta,*}_{L(k+1),n}(\bar{L}(k))$ is set to 1 in accordance with equality (3).

5. Derive the estimate of $E\big(Q^\theta_{L(1)}(L(0))\big)$ denoted by $Q^\theta_{L(0)}$.

   For patients who followed rule d_θ at time 0, we already computed an estimate $Q^{\theta,*}_{L(1),n}(L(0))$ in the latest "update step". For all other patients, an estimate $Q^{\theta,*}_{L(1),n}(L(0))$ is computed here by extrapolation, i.e., using the same

protocol employed in the latest "update step" as if these patients also followed rule d_θ at time 0. Thus, an estimate $Q^{\theta,*}_{L(1),n}(L(0))$ is now available for all n patients in the cohort. The average of these estimates is an estimate of $Q^{\theta}_{L(0)}$ denoted by $Q^{\theta,*}_{L(0),n}$:

$$Q^{\theta,*}_{L(0),n} = \frac{1}{n}\sum_{i=1}^{n} Q^{\theta,*}_{L(1),n}(L_i(0)).$$

This estimate $Q^{\theta,*}_{L(0),n}$ is the TMLE point estimate of the counterfactual cumulative risk of interest γ^θ and we thus also denote it by $\gamma^{\theta,*}_n$.

From the TMLE point estimates $\gamma^{\theta_1,*}_n$ and $\gamma^{\theta_2,*}_n$ obtained by applying twice the 5 steps above with, first, $\theta = \theta_1$ and, second, $\theta = \theta_2$, we derived a point estimate for the parameter of interest ψ^{θ_1,θ_2}. This estimate is denoted by $\psi^{\theta_1,\theta_2,*}_n$ and we have: $\psi^{\theta_1,\theta_2,*}_n = \gamma^{\theta_1,*}_n - \gamma^{\theta_2,*}_n$. Note that the algorithm above can also be repeated for a continuous sequence of time points starting at $t_0 = 0$. The resulting estimates of the counterfactual cumulative risks can then be mapped into an estimate of the corresponding counterfactual survival curve using the link $S_{d_\theta}(t) = 1 - P(Y_{d_\theta}(t+1) = 1)$ where $S_{d_\theta}(t) = P(T_{d_\theta} > t)$ denotes the probability of survival at time t. The estimate of $S_{d_\theta}(t)$ obtained with this approach is denoted by $S^*_{d_\theta,n}(t)$.

To implement the TMLE algorithm above, we need to specify estimation approaches for two vectors of nuisance parameters denoted by g^θ and Q^θ. The nuisance parameter g^θ corresponds with the estimands in step 1, i.e., $g^\theta_{A(t)}$ for $t = 0, \ldots, t_0$. The nuisance parameter Q^θ corresponds with the estimands in all other steps, i.e., $Q^\theta_{L(t+1)}(\bar{L}(t))$ for $t = 0, \ldots, t_0$. The estimator $\gamma^{\theta,*}_n$ is doubly robust in the sense that it is a consistent estimator of the true cumulative risk γ^θ if either the estimator of the nuisance parameter g^θ is a consistent estimator of the true g^θ or if the initial estimator of the nuisance parameter Q^θ is a consistent estimator of the true Q^θ. In addition, $\gamma^{\theta,*}_n$ is efficient if both estimators of the nuisance parameters are consistent.

Given that IPW estimation also relies on an estimate of the nuisance parameter g^θ, we initially evaluated TMLE using the same estimation approach for g^θ that had been implemented in previous work [47]. The approach is based on mapping (1) that links the nuisance parameter g^θ to the 5 PS $\mu_1, \ldots, \mu_5$. Data were pooled for all time points $t = 0, \ldots, 36$ to fit a separate main-term logistic model for estimating each of the 3 PS for right-censoring (μ_3, μ_4, μ_5) and the PS for TI continuation (μ_2). 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., μ_1 for $t > 0$). A separate main-term logistic model was fitted for estimating the PS for TI initiation at $t = 0$ (i.e., μ_1 for $t = 0$). 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. 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 models over time also included the variable indexing the 90-day follow-up intervals (i.e., t) as an explanatory variable.

For deriving the initial estimate of the nuisance parameter Q^θ, we relied on the machine learning algorithm 'DSA' [58, 59]. The DSA implements data-adaptive estimator selection based on cross-validation. The candidate estimators considered were restricted to main-term logistic models of different sizes with the following candidate explanatory variables: all time-independent covariates, the last measurement of time-varying covariates, and the latest change in A1c. To alleviate computing time, the DSA algorithm was implemented with a single 5-fold cross-validation split, without deletion and substitution moves, and with a maximum model size of 10 explanatory variables.

The resulting estimates of the four counterfactual survival curves $S^*_{d_\theta,n}(t)$ for $\theta = 7, 7.5, 8, 8.5$ and $t = 0, \ldots, 15$ (4 years of follow-up) are displayed in Figure 1. The corresponding estimates of the 6 distinct RDs $\psi^{\theta_1,\theta 2,*}_n$ for $t_0 = 11$, i.e., the difference of counterfactual cumulative risks over 3 years, are displayed in Table 1.

Note that while the TMLE algorithm just described may appear complex, it essentially involves the implementation of a sequence of standard regression steps to fit logistic models and derive predicted values from these models. Implementation with standard statistical software is thus relatively trivial as reflected by the computing time needed to derive an estimates $\psi^{\theta_1,\theta_2,*}_n$ in R version 2.13.0 [60]. When we exclude the time needed to derive the estimate of the nuisance parameter g^θ and the initial estimate of the nuisance parameter Q^θ, the completion of the remaining steps of the TMLE algorithm was obtained in about 1 minute. The overall computing time to derive the estimate $\psi^{\theta_1,\theta_2,*}_n$ was however about 40 minutes due to the computing burden imposed by the DSA algorithm. This computing time can thus be greatly shortened with the selection of a faster machine learning algorithm instead of the DSA or through the arbitrary specification of logistic models for the different components of the nuisance parameter Q^θ (as done for estimating the nuisance parameter g^θ). Before evaluating the performance of TMLE by comparison with IPW estimation, we now discuss approaches for deriving inference with TMLE.

3.2. *Inference with TMLE*

The estimator $\gamma_n^{\theta,*}$ defined by the TMLE algorithm above is asymptotically linear with influence curve denoted by $IC_\theta^*(O \mid g^\theta, Q^\theta)$ and defined by:

$$IC_\theta^*(O \mid g^\theta, Q^\theta) = \sum_{t=0}^{t_0+1} D_{\theta,t}^*(O \mid g^\theta, Q^\theta) \text{ with}$$

$$D_{\theta,t}^*(O \mid g^\theta, Q^\theta) = \frac{I\left(\bar{A}(\check{T}(t-1)) = d_\theta(\bar{L}(\check{T}(t-1)))\right)}{\prod_{j=0}^{\check{T}(t-1)} g_{A(j)}^\theta}\left(Q_{L(t+1)}^\theta - Q_{L(t)}^\theta\right), \tag{5}$$

where $Q_{L(t_0+2)}^\theta = Y(t_0+1)$ and $\frac{I\left(\bar{A}(\check{T}(t-1))=d_\theta(\bar{L}(\check{T}(t-1)))\right)}{\prod_{j=0}^{\check{T}(t-1)} g_{A(j)}^\theta}$ is nil at $t=0$ [20].

Note that $D_{\theta,t}^*(O \mid g^\theta, Q^\theta) = 0$ for all t such that either i) $\tilde{T} = C$ and $C+1 \leq t \leq t_0+1$ because we then have $I\left(\bar{A}(\check{T}(t-1)) = d_\theta(\bar{L}(\check{T}(t-1)))\right) = 0$, or ii) $\tilde{T} = T$ and $T+1 < t \leq t_0+1$ because we then have $Q_{L(t+1)}^\theta - Q_{L(t)}^\theta = 0$.

From the delta method, the estimator $\psi_n^{\theta_1,\theta_2,*}$ is thus asymptotically linear with the influence curve $IC_{\theta_1,\theta_2}^*(O \mid g^{\theta_1}, Q^{\theta_1}, g^{\theta_2}, Q^{\theta_2}) = IC_{\theta_1}^*(O \mid g^{\theta_1}, Q^{\theta_1}) - IC_{\theta_2}^*(O \mid g^{\theta_2}, Q^{\theta_2})$, i.e.,

$$\psi_n^{\theta_1,\theta_2,*} - \psi^{\theta_1,\theta_2} = \frac{1}{n}\sum_{i=1}^{n} IC_{\theta_1,\theta_2}^*(O_i \mid g^{\theta_1}, Q^{\theta_1}, g^{\theta_2}, Q^{\theta_2}) + o(\frac{1}{\sqrt{n}}).$$

Under the assumption that $g_n^{\theta_1}$ and $g_n^{\theta_2}$ are consistent estimators, a conservative estimate of the asymptotic standard error of $\psi_n^{\theta_1,\theta_2,*}$ is given by:

$$\sigma_n^{\theta_1,\theta_2,*} = \sqrt{\frac{1}{n^2}\sum_{i=1}^{n}\left[IC_{\theta_1,\theta_2}^*(O_i \mid g_n^{\theta_1}, Q_n^{\theta_1,*}, g_n^{\theta_2}, Q_n^{\theta_2,*})\right]^2}, \tag{6}$$

where $g_n^\theta = (g_{A(0),n}^\theta, \ldots, g_{A(t_0),n}^\theta)$ is the vector of estimates obtained in step 1 of the TMLE algorithm and $Q_n^{\theta,*} = (Q_{L(0),n}^{\theta,*}, \ldots, Q_{L(t_0+1),n}^{\theta,*})$ is the vector of updated estimates obtained at each "update step" of the TMLE algorithm. Thus, computation of the estimated standard error (SE) of the estimator $\psi_n^{\theta_1,\theta_2,*}$ with formula (6) does not add significant computing time to the TMLE approach because it is based on by-products (g_n^θ and $Q_n^{\theta,*}$) of the algorithm for deriving the point estimate.

For each of the six RDs of interest, we can thus compute the lower and upper bounds of the 95% confidence interval (CI) and the p-value associated with the two-sided test of the null hypothesis ($H_0 : \psi^{\theta_1,\theta_2} = 0$) as follows:

$$P(\underbrace{\psi_n^{\theta_1,\theta_2,*} - z_{0.025}\sigma_n^{\theta_1,\theta_2,*}}_{\text{lower bound of the 95\% CI}} \leq \psi^{\theta_1,\theta_2} \leq \underbrace{\psi_n^{\theta_1,\theta_2,*} + z_{0.025}\sigma_n^{\theta_1,\theta_2,*}}_{\text{upper bound of the 95\% CI}}) = 0.95 \text{ and } \underbrace{2\Phi\left(-\left|\frac{\psi_n^{\theta_1,\theta_2,*}}{\sigma_n^{\theta_1,\theta_2,*}}\right|\right)}_{\text{p-value denoted by p}^*} \tag{7}$$

where $z_{0.025} = \Phi^{-1}(0.975)$ and Φ is the cumulative distribution function of the standard normal distribution.

The estimates of the SEs, CIs, and p-values derived analytically from expressions (6) and (7) for each of the 6 distinct RDs ψ^{θ_1,θ_2} are displayed in Table 1. To assess the performance of estimator (6) in evaluating the variability of $\psi_n^{\theta_1,\theta_2,*}$, we also evaluated the variability of $\psi_n^{\theta_1,\theta_2,*}$ based on 10,000 bootstrap samples. For a fair comparison between both estimates of TMLE variability, the point estimates of ψ^{θ_1,θ_2} were derived on each bootstrap sample using the same original estimated nuisance parameters $g_n^{\theta_1}$ and $g_n^{\theta_2}$ (i.e., the 5 PS models in step 1 were not refitted on each bootstrap sample) because estimator (6) is only consistent for estimating the SE of $\psi_n^{\theta_1,\theta_2,*}$ when the nuisance parameter g^θ is known. In Table 1, the ratios (denoted by RSE*) of the estimates of the SE based on the influence curve over that derived from the bootstrap indicate important over-estimation (by up to 61%) of the TMLE SE with formula (6). Given that the performance of the SE estimator (6) relies on correct estimation of the nuisance parameters g^θ, we hypothesized that the poor performance of estimator (6) could be the result of misspecification of the logistic models for the PS on which are based the current estimates of the nuisance parameters g^θ.

Thus, in addition to the first approach described in section 3.1, we also implemented the TMLE algorithm for estimating the 4 survival curves $S_{d_\theta}(t)$ and the 6 RDs ψ^{θ_1,θ_2} with 3 alternate approaches to estimate the nuisance parameter g^θ. These approaches numbered 1 through 4 are progressively more flexible, i.e., nonparametric:

- Approach 1: This is the approach based on the 6 main-term logistic models described in section 3.1. Note that 5 of these models are based on data pooled over time.
- Approach 2: For each time point t separately, 5 main-term logistic models were fitted to estimate each of the 5 PS.
- Approach 3: The logistic models from approach 2 were all modified to include interactions terms between explanatory variables. The interaction terms added to each logistic model were selected using the following algorithm: 105 two-way interaction terms were computed based on the 15 explanatory variables that were the most associated (smallest p-value) with the PS dependent variable (denoted by Z) in a univariate logistic regression. For each of these 105 terms, a logistic regression of Z on the interaction term and the two main terms that define the interaction term was implemented. The p-values p associated with the interaction terms in these 105 regressions were used to identify all interaction terms with $p < 0.05$. If more than 50 interaction terms met this criterion, only the 50 terms with the smallest p were selected and added to the main-term logistic model for the PS.
- Approach 4: The parametric models adopted for estimating g^θ in approaches 1 through 3 above do not reflect true subject-matter knowledge about the 5 PS. To avoid erroneous inference [12, 61] due to arbitrary model specifications, data-adaptive estimation of the nuisance parameter g^θ may be implemented in practice [62, 63] but consistent estimation then relies on judicious selection of a machine learning algorithm also known as 'learner'. Several learners are potential candidates for estimating the 5 PS (e.g., [64, 65, 66, 67, 68, 69, 70, 58, 71, 72]). 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" [17]. To hedge against erroneous inference due to arbitrary selection of a learner, Super Learning [17] 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 [73, 74, 75, 76] to protect against over-fitting 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.

  For each time point t separately, 5 super learners were implemented to estimate each of the 5 PS based on 10 candidate learners: i) 5 learners[‡] defined by logistic models with only main terms for the most predictive explanatory variables identified[§] by a significant p-value in univariate regressions with 5 significance levels (α= 1e-30, 1e-10, 1e-5, 0.1, and 1), and ii) 5 polychotomous regression learners[¶] based on the most predictive explanatory variables identified by a significant p-value in univariate regressions with the same 5 significance levels.

Note that unlike the computing time for the first 3 approaches which is measured in minutes, approach 4 based on SL is computing intensive: approximately 7, 1, 4, 7, and 10 hours were needed to derive the 5 sets of super learners for TI initiation, TI continuation, censoring by administrative end of study, censoring by disenrollment from the health plan, and censoring by death, respectively. To simplify notation, we denote the estimate of g^θ obtained with approaches 1 through 4 above by g^θ_n, $g^\theta_{n,t}$, $g^\theta_{n,t,\times}$, and $g^\theta_{n,t,SL}$, respectively. The TMLE of the 4 survival curves $S_{d_\theta}(t)$ and the 6 RDs ψ^{θ_1,θ_2} based on these four estimates of g^θ are displayed in Table 1 and Figure 1.

In Table 1, it is clear that the ratios (RSE*) of the estimates of the TMLE SE based on the influence curve over that derived from the bootstrap get progressively closer to 1 as the approach used to estimate g^θ becomes more nonparametric. The desired ratios of 1 are approximately reached with SL estimation of the action mechanism, i.e., with $g^\theta_{n,t,SL}$. These results are consistent with our earlier hypothesis for explaining the initial poor performance of the analytic estimator of the SE associated with the TMLE estimator $\psi^{\theta_1,\theta_2,*}_n$ that is based on the estimate g^θ_n: They suggest that 1) the logistic models used in approach 1 to derive the estimate g^θ_n were misspecified, and 2) SL successfully permitted to hedge against such misspecification by correcting for the bias in estimating g^θ with approach 1. Unlike the clear sensitivity of the estimates of TMLE variability to the choice of estimator for g^θ (Table 1), Figure 1 suggests however that the TMLE point estimates are much less sensitive to the approach taken for estimating g^θ since the plots of survival curves associated with approaches 1 through 4 are relatively stable.

To evaluate the performance in time-dependent confounding and selection bias adjustment with the TMLE algorithm, we now compare the inferences obtained with TMLE to that obtained with IPW estimation below.

[‡] implemented by the `SL.glm` routine available in the `SuperLearner` R package [77]

[§] using the template screening routine `screen.glmP` available in the `SuperLearner` R package

[¶] implemented by the `SL.polyclass` routine given in [78, Appendix]. This routine implements the polyclass learner [70] 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.

3.3. Comparison to IPW estimation

In previous work [47], stabilized IPW estimation of the same four survival curves and 6 RDs was implemented based on a parametric dynamic MSM for the four discrete-time hazards. Estimates of the hazards were subsequently mapped into estimates of each of the four corresponding counterfactual survival curves from which inferences about RDs were derived.

Because these IPW estimates were derived in a previous analysis based on additional modeling assumptions encoded by a non-saturated MSM, direct comparison with the TMLE estimates in this report is not optimal. For a fair comparison, we implemented the same general stabilized IPW estimation approach based on the same four approaches for estimating the nuisance parameter g^θ but using the following saturated MSM which is equivalent to a nonparametric MSM:

$$logit(P(Y_{d_\theta}(t+1)=1 \mid Y_{d_\theta}(t)=0)) = \sum_{\theta' \in \{7,7.5,8,8.5\}} \sum_{j=0}^{36} \beta_j^{\theta'} I(\theta=\theta', t=j) \text{ for } \theta \in \{7,7.5,8,8.5\} \text{ and } t=0,\ldots,36.$$

Another difference with the IPW estimation implemented in previous work is that we do not assume that right-censoring due to administrative end of study is uninformative in this report. For simplicity, this assumption had been made in the earlier analyses even though administrative end of study could potentially result in selection bias because of the the cohort being open. When one assumes that some right-censoring events are not informative, IPW estimation can be simplified by ignoring the corresponding conditional probabilities of censoring in the calculation of stabilized weights because weight stabilization results in cancellation of such probabilities in the numerator and denominator of the weights. Because weight stabilization is not possible in TMLE of the cumulative risks γ^θ, estimation of the conditional probability of right-censoring due to administrative end of study is needed for TMLE implementation even if such events can be assumed to be uninformative. For this reason, we relaxed the assumption of uninformative censoring due to administrative end of study when implementing TMLE. We thus also relaxed this assumption when implementing IPW estimation to enable a fair comparison of the two approaches.

The IPW estimator $\beta_{t,n}^\theta$ of each coefficient β_t^θ of the saturated MSM can be derived by a single standard weighted regression as done in previous work. Equivalently, each $\beta_{t,n}^\theta$ can be derived separately by solving the estimating equation associated with the following IPW estimating function for the discrete-time hazard under rule d_θ at time t:

$$D(O \mid \alpha_t^\theta) = I(Y(t)=0) I(\bar{A}(t)=d_\theta(\bar{L}(t))) \frac{\prod_{j=0}^t P(A(j)=d_\theta(\bar{L}(j)) \mid \bar{A}(j-1)=d_\theta(\bar{L}_i(j-1)))}{\prod_{j=0}^t g_{A(j)}^\theta} \left(Y(t+1) - \alpha_t^\theta \right),$$

where the counterfactual discrete-time hazard $P(Y_{d_\theta}(t+1)=1 \mid Y_{d_\theta}(t)=0)$ is denoted by α_t^θ. The resulting estimator denoted by $\alpha_{t,n}^\theta$ is defined as:

$$\alpha_{t,n}^\theta = \frac{\sum_{i=1}^n I(Y_i(t)=0) I(\bar{A}_i(t)=d_\theta(\bar{L}_i(t))) \frac{\prod_{j=0}^t P(A_i(j)=d_\theta(\bar{L}_i(j)) \mid \bar{A}_i(j-1)=d_\theta(\bar{L}_i(j-1)))}{\prod_{j=0}^t g_{A_i(j)}^\theta} Y_i(t+1)}{\sum_{i=1}^n I(Y_i(t)=0) I(\bar{A}_i(t)=d_\theta(\bar{L}_i(t))) \frac{\prod_{j=0}^t P(A_i(j)=d_\theta(\bar{L}_i(j)) \mid \bar{A}_i(j-1)=d_\theta(\bar{L}_i(j-1)))}{\prod_{j=0}^t g_{A_i(j)}^\theta}}$$

and we have: $\beta_{t,n}^\theta = log(\alpha_{t,n}^\theta)$. Computing time is greatly shorten with this second approach for deriving $\beta_{t,n}^\theta$.

Note that unlike TMLE which is doubly robust, IPW estimation relies on consistent estimation of the nuisance parameter g^θ. In this report, however, we aim to evaluate the performance of both approaches under consistent estimation of the nuisance parameter g^θ. Table 1 and Figure 2 display the results of IPW estimation under the four approaches considered for estimating g^θ. In addition, Figure 3 displays the results from a crude analysis that consists in fitting a saturated logistic model for the discrete-time hazards without weights, i.e., without adjustment for time-dependent confounding and selection bias.

A comparison of the crude and IPW estimates of the survival curves on Figures 2 and 3 clearly demonstrates successful adjustment for time-dependent confounding and selection bias with the IPW approach. Whichever the approach adopted for estimating g^θ, the IPW estimates indicate an early separation and consistent ordering of the four survival curves suggesting an increasing beneficial effect of more aggressive therapy initiation rules (i.e., of rules indexed by decreasing A1c thresholds). These results are consistent with that of the ACCORD and ADVANCE randomized trials. Successful performance in bias adjustment with the TMLE approach is illustrated on Figure 1 which indicates the same separation of the four survival curves. The estimates of the survival curves obtained by IPW estimation are visually almost identical to that obtained by TMLE.

The stability of the survival curves on Figures 1 and 2 suggests that estimation bias for g^θ with the first 3 approaches is relatively minor with respect to its impact on confounding and selection bias adjustment with both TMLE and IPW estimation. It is of interest to note that while Table 1 indicates that the analytic estimate of the TMLE variability is largely

affected (RSE$^* > 1$ with g_n^θ, $g_{n,t}^\theta$, and $g_{n,t,\times}^\theta$) by such minor bias in estimation of g^θ, the analytic estimates of IPW estimation variability remain largely unaffected by such bias (RSE≈ 1 whichever the approach for estimating g^θ).

Table 1 also indicates that the SE of the IPW estimators (denoted by $\sigma_n^{\theta_1,\theta_2}$) decreases as the approach taken for estimating g^θ becomes more flexible. This gain in estimation efficiency is explained by the decrease in the proportion of large weights from progressively more flexible approaches as shown in Table 2. In fact, none of the stabilized IPW weights derived from SL are greater than 30. This is quite remarkable given the common intuition that data-adaptive estimation may not be desirable for practical implementation of IPW estimation for fear to reveal violations of the positivity assumption. Intuitively, estimation of the IPW weights based on arbitrarily specified parametric models can then be viewed as an implicit way to restrict the proportion of large weights through smoothing with a misspecified model. This data analysis suggests that parametric estimation of the weights may instead lead to artificial violation of the positivity assumption in practice when the positivity assumption is in truth not violated. These remarks are supported by the examination of data (not shown) from individual patients who have estimated IPW weights above 50 with g_n^θ while being below 50 with $g_{n,t,SL}^\theta$. Such data review shows that SL leads to better prediction of TI initiation and administrative end of study when patients actually initiate TI and remain uncensored. This improved prediction translates into smaller denominators for the IPW weights and thus less variable IPW weights which then results in improving precision in the effect estimates with IPW estimation.

Whether extreme weighting is the result of true or artificial violation of the positivity assumption from misspecification of the model for g^θ, weight truncation [79, 80, 81] is often implemented in practice to improve the precision of IPW estimation. In this analysis, all IPW estimates were derived with and without truncation of the stabilized weights: all stabilized weights above 20 were replaced with value 20. The same rationale for truncating the IPW weights applies to TMLE. Thus, following the same truncation scheme as the one implemented for IPW estimation, the unstabilized weights $\frac{1}{\prod_{t=0}^k g_{A(t)}^\theta}$ used in the 'update steps' of the TMLE algorithm were truncated at 20 divided by the corresponding numerator of the stabilized weights, i.e., replaced by $\min(\frac{1}{\prod_{t=0}^k g_{A(t)}^\theta}, \frac{20}{\prod_{t=0}^k P(A(t)=d_\theta(\bar{L}(t))|\bar{A}(t-1)=d_\theta(\bar{L}(t-1)))})$.

As may be expected from the distributions of the untruncated stabilized weights resulting from the four approaches to estimate g^θ (Table 2), weight truncation at 20 had very little impact on inferences in this study (results not shown). In particular note that with SL, such weight truncation only concerns two rule-person-time observations for which the untruncated weight is below 30. Thus, weight truncation at 20 had essentially no impact on the IPW and TMLE results when the nuisance parameter g^θ is estimated with SL.

4. Targeted Learning for efficiency gains

One motivation for the application of TMLE over IPW estimation is the potential for gain in estimation precision that may arise from the efficiency property of TMLE. Table 1 indicates however little increase in precision with TMLE since the ratios of the IPW SEs to the TMLE SEs range from 2 to 5% only when g^θ is estimated with $g_{n,t,SL}^\theta$. As shown by equality (5) defining the influence curve of the TMLE estimator, gain in precision arises from minimization of the prediction errors from the estimate of the nuisance parameter Q^θ. Estimation precision is thus expected to be potentially improved with TMLE in problems where covariates can predict the outcome well. Lack of efficiency gains from the TMLE results in Table 1 could thus be the result of either the absence of good predictors of the outcome among the measured covariates or inadequate use of the measured covariates for estimating Q^θ (e.g., due to model misspecification). Consistent estimation of the nuisance parameter Q^θ should thus result in improving estimation efficiency with TMLE. This motivates the application of Super Learning for flexible estimation of Q^θ without relying on arbitrary learner choices such as the main-term logistic models considered by the DSA algorithm in the previous implementation of TMLE.

In addition, note that each element of the nuisance parameter Q^θ, i.e., $Q_{\bar{L}(k+1)}^\theta(\bar{L}(k))$ for $k = t_0, \ldots, 0$, is potentially a function of the covariate history $\bar{L}(k)$ in the same way that each PS is potentially a function of past observed covariates. Because treatment decision are typically made based on the latest clinical measurements, the Markov assumption is often justified and only the last measurement of time-varying covariates are then considered as potential explanatory variables to estimate the PS for treatment initiation. So far, we largely relied on such Markov assumptions for estimating not only each PS but also the nuisance parameter Q. The only exception was for the A1c covariate since we not only considered the last measurement of A1c to estimate the 5 PS but also the last change in A1c measurements, i.e., the difference between the last two A1c measurements. While Markov assumptions may be adequate for estimating the PS, it may be argued that clinical outcomes are not only the results of acute effects but also chronic effects and Markov assumptions may thus not permit consistent estimation of the nuisance parameter Q^θ. In addition, some covariates in the observed data were assumed to be unrelated to treatment decisions and right-censoring and where thus excluded from the list of explanatory variables to estimate g^θ. These covariates may however impact the outcome and should then be considered for consistent

estimation of Q^θ.

For these reasons, we implemented TMLE of the RDs ψ^{θ_1,θ_2} based on SL with an expanded list of explanatory variables for estimating the nuisance parameter Q^θ. More specifically, for each k, the list of candidate explanatory variables considered by SL to estimate $Q^\theta_{L(k+1)}(\bar{L}(k))$ was appended with:

- the indicators that A1c, LDL, and ACR were measured in the current period
- the indicator that the ACR is measured in the next period (i.e., that the patient is at risk of possible failure)
- the average of past A1c and LDL measurements
- the average number of past imputed A1c, LDL and ACR measurements
- the standard deviations of the past measurements of A1c and LDL
- the numbers of past A1C measurements above 7, 7.5, 8, 8.5
- the numbers of past LDL measurements above 100,130,160, and 200
- the baseline A1c, LDL and ACR measurements
- the difference between the baseline ACR measurement and the relevant ACR cut-off that would determine subsequent failure
- the difference between the last ACR measurement and the relevant ACR cut-off that would determine subsequent failure

The following 57 candidate learners were considered to estimate $Q^\theta_{L(k+1)}(\bar{L}(k))$ separately with SL for each k and θ:

- 1 learner defined by the intercept logistic model
- 5 learners defined by logistic models with only main terms for the most predictive explanatory variables identified by a significant p-value in univariate regressions with 5 significance levels (α= 0.3, 0.2, 0.1, 0.05, 1)
- 1 learner defined by a logistic model with only main terms for 27 expert-selected covariates (all explanatory variables based on A1c, LDL and ACR measurements)
- 5 learners defined by the `stepAIC` routine in R and based on 5 distinct sets of candidate explanatory variables identified by i) 4 significance levels (α= 0.3, 0.2, 0.1, 0.05) and ii) the same 27 expert-selected covariates
- 5 neural network learners defined based on the same 4 significance levels and 27 expert-selected covariates
- 5 learners defined by Bayes regression and based on the same 4 significance levels and 27 expert-selected covariates
- 5 learners defined by polychotomous regression (`polyclass` with BIC) and based on the same 4 significance levels and 27 expert-selected covariates
- 5 Random Forest learners based on the same 4 significance levels and 27 expert-selected covariates
- 5 learners defined by bagging for classification trees and based on the same 4 significance levels and 27 expert-selected covariates
- 20 learners defined by generalized additive models with smoothing splines of degree 2, 3, 4, or 5 and based on, for each of these degrees, the same 4 significance levels and 27 expert-selected covariates.

In addition to all main-term variables, all these learners also considered two-way interaction terms selected on each training set using the protocol described in section 3.2 (See approach 3 for estimating g^θ).

Note that for all $k < t_0$, the outcome to be predicted when estimating $Q^\theta_{L(k+1)}(\bar{L}(k))$ is continuous between 0 and 1. Because the current R implementation of `polyclass` does not permit the prediction of continuous outcomes between 0 and 1, we modified the 5 learners based on polychotomous regression when $k < t_0$ by dichotomizing the outcomes in the 5-fold training sets based on the 0.5 cut-off, i.e., all outcomes above 0.5 were replaced by 1 and otherwise by 0. To greatly shorten the SL computing time, the same dichotomization was implemented as part of the Random Forest and bagging learners considered for $k < t_0$. In addition for $k = t_0$, SL was only used to predict the outcome of patients who are at risk of failure, i.e., for whom ACR was actually monitored. For all other patients, the estimate of $Q^\theta_{L(t_0+1)}(\bar{L}(t_0))$ was set to 0 by default.

Table 3 displays the results of TMLE based on the SL approach described above for estimating the nuisance parameter Q^θ (g^θ is estimated with $g^\theta_{n,t,SL}$). These results illustrate the potential for gains in efficiency with TMLE compared to IPW estimation since the IPW SEs are now 7 to 11% higher than that of the TMLE estimator.

5. Discussion

Tables 4 and 5 summarize and compare various aspects of the IPW estimation and TMLE algorithms considered in this report. In particular, this work suggests that bias in the estimation of the action mechanism g^θ may differentially impact point and interval estimation with TMLE and IPW estimation. In this one study, while point estimation with both methods and interval estimation with IPW estimation were relatively robust to minor bias in estimation of the action mechanism

g^θ, interval TMLE based on the estimator's influence curve was sensitive to such bias. This observation suggests the use of bootstrapping to evaluate the bias of estimators of the action mechanism and guide the selection of the estimator based on which TMLE inference should be derived in practice, e.g., the data-adaptiveness of the estimation approach for g^θ may be sequentially increased until there is a match between the two estimates of TMLE variance obtained analytically based on the estimator's influence curve and based on bootstrapping.

Super Learning was shown here to be a viable practical data-adaptive estimation approach for the action mechanism despite the common intuition that data-adaptive estimation may not be desirable in practice for fear to reveal violations of the positivity assumption. Instead, estimation based on parametric models - even if misspecificed - may then be preferred in practice because it is viewed as an implicit way to restrict the proportion of large IPW weights through smoothing. The results in this report suggest however that the estimation approaches for the action mechanism that are commonly used in practice based on arbitrary parametric models may lead to artificial violation of the positivity assumption (extreme IPW weights) when this assumption is in truth not violated.

The implementation of TMLE in this work also underscored the need for estimation of all components of the action mechanism even if one or more components are assumed to be uninformative, e.g., conditional probabilities of censoring by administrative end of study. Unlike weight stabilization in IPW estimation which allows the analyst to ignore estimation of 'uninformative components' of the action mechanism because weight stabilization results in cancellation of the corresponding conditional probabilities in the denominator and numerator of the weights, weight stabilization is not possible with the proposed TMLE algorithm and estimates of all components of the action mechanism are thus always needed in practice.

While IPW estimation is grounded in a formal theoretical framework that is opaque to many applied researchers, its implementation is simple enough that some intuitive explanations are available to provide insights into its ability to adjust for confounding and selection bias in practice (e.g., weighting of the observed data results in ghost/pseudo data where confounders are no longer associated with exposures as if patients had actually been randomized to the exposure of interest) [82]. While the Targeted Learning algorithm evaluated in this report is also grounded in theory, there is however no intuitive explanations for justifying its ability to adjust for confounding and selection bias. This report demonstrates that a seemingly convoluted algorithm with no intuitive support is not only computationally feasible in real-world CER based on large healthcare databases but also, most importantly, on a par with IPW estimation in performance with confounding and selection bias adjustment in CER that involve time-varying confounders on causal pathways between the time-varying exposures and outcome. In addition, this application illustrates the algorithm's potential for improving inferences through gains in estimation efficiency compared to IPW estimation.

In practice, however, such gains in efficiency may only reflect over-fitting in the estimation of the nuisance parameter Q^θ (see formula (5)). Estimation methods based on cross-validation for the selection of estimators of the nuisance parameters (such as the DSA algorithm or SL) can then provide protection against such incorrect inferences with TMLE. Unlike estimation of the action mechanism g^θ for which assumptions like the Markov assumption may often be realistically invoked to simplify estimation by only considering the most recent treatment and covariate histories as explanatory variables (e.g., for predicting treatment initiation), results in this report suggest that estimation of Q^θ should generally rely on a richer set of both treatment and covariate histories to realize efficiency gains with TMLE because simplifying assumptions that may apply to components of the the action mechanism are not expected to extend to the nuisance parameter Q^θ in general. If these simplifying assumptions are nevertheless made, suboptimal efficiency gains are expected with TMLE due to inconsistent estimation of Q^θ.

The interpretation of the results from this report are limited by the absence of a true gold standard to compare TMLE and IPW estimation. In addition, while the IPW estimation approach insures that the estimated survival curves are monotone decreasing, the TMLE estimator does not. Isotonization of the estimated survival curves derived by TMLE can be used in practice to enforce monotonic decrease of the estimates of survival probabilities over time. Inference can then be derived based on recent results that have shown that the influence curve of the isotonized estimator of a counterfactual survival curve is identical to the influence curve of the original estimator [83]. Isotonization was not needed here because all estimated survival curves were monotone decreasing. Furthermore, the nonparametric MSM approach adopted in this report may not always be realistic in practice as very little data may be available to contrast the interventions of interest with precision (curse of dimensionality). In such cases, one may revert back to the convenient but elusive approach which consists in arbitrarily choosing a non-saturated MSM which is is expected to be misspecified and thus likely inadequate to provide a consistent estimate of the parameter of interest. Instead, a nonparametric MSM approach based on a *working, non-saturated model* [84] can be adopted to explicitly recognize the limitation of an arbitrarily specified non-saturated MSM in capturing the true parameter of interest ψ^{θ_1,θ_2}. The approach consists in replacing the parameter ψ^{θ_1,θ_2} with a new parameter of interest defined by minimizing the distance between the true survival curves of interest and their estimates from a working model (referred to as a working MSM). Informally, such a nonparametric MSM approach aims to emulate inference from an ideal randomized trial (perfect compliance and no loss to follow-up) that is based on a working (likely misspecified) model for the survival curves in each treatment arm. While IPW estimation is adapted to

estimate the parameters defined by such an MSM approach, the TMLE algorithm studied in this report is not but an alternate algorithm was recently developed for such parameters [85].

Although this work focused on the application of TMLE in CER with observational data, TMLE like IPW estimation is also relevant in randomized experiments with non-adherence and loss to follow-up to properly account for time-dependent confounding and selection bias in 'as treated' or 'per protocol' analyses [86] or to increase estimation efficiency without jeopardizing consistency in intention-to-treat analyses that incorporate covariate information collected prior to randomization [87]. The theoretical properties of the Targeted Learning approach have thus the potential to greatly impact CER through improved confounding and selection bias adjustment (double robustness and data-adaptive estimation) but also more precise effect estimates (efficiency property). Concretely, Targeted Learning could lead to more reliable effect estimates, earlier detection of effectiveness or safety signals, or the ability to detect differential subgroup effects with smaller sample sizes. While these potential advantages are theoretically derived from the large-sample (asymptotic) properties of TMLE, concerns over possible undesirable finite-sample properties that cannot be predicted from theory may be raised. The results in this report do not substantiate such concerns and 1) demonstrate the feasibility of TMLE in real world CER, 2) illustrate the ability of TMLE to properly account for time-dependent confounding and selection bias, and 3) elicit practical evidence of the potential for efficiency gains with TMLE over IPW estimation. While these results motivate further applied investigation of TMLE, they are not meant as a compelling argument for the routine application of TMLE instead of IPW estimation. In this one study, only a relatively modest gain in estimation precision was demonstrated at the cost of an increased computing burden with SL. Additional research is needed to continue evaluation of the translation of TMLE's theoretical properties into practice (finite-sample properties) and, in particular, i) to evaluate the algorithm's double robustness property and, ii) to develop a computationally feasible algorithm to estimate TMLE variability when one does not want to rely solely on consistent estimation of the action mechanism (i.e., consistent estimation of the nuisance parameter g^θ).

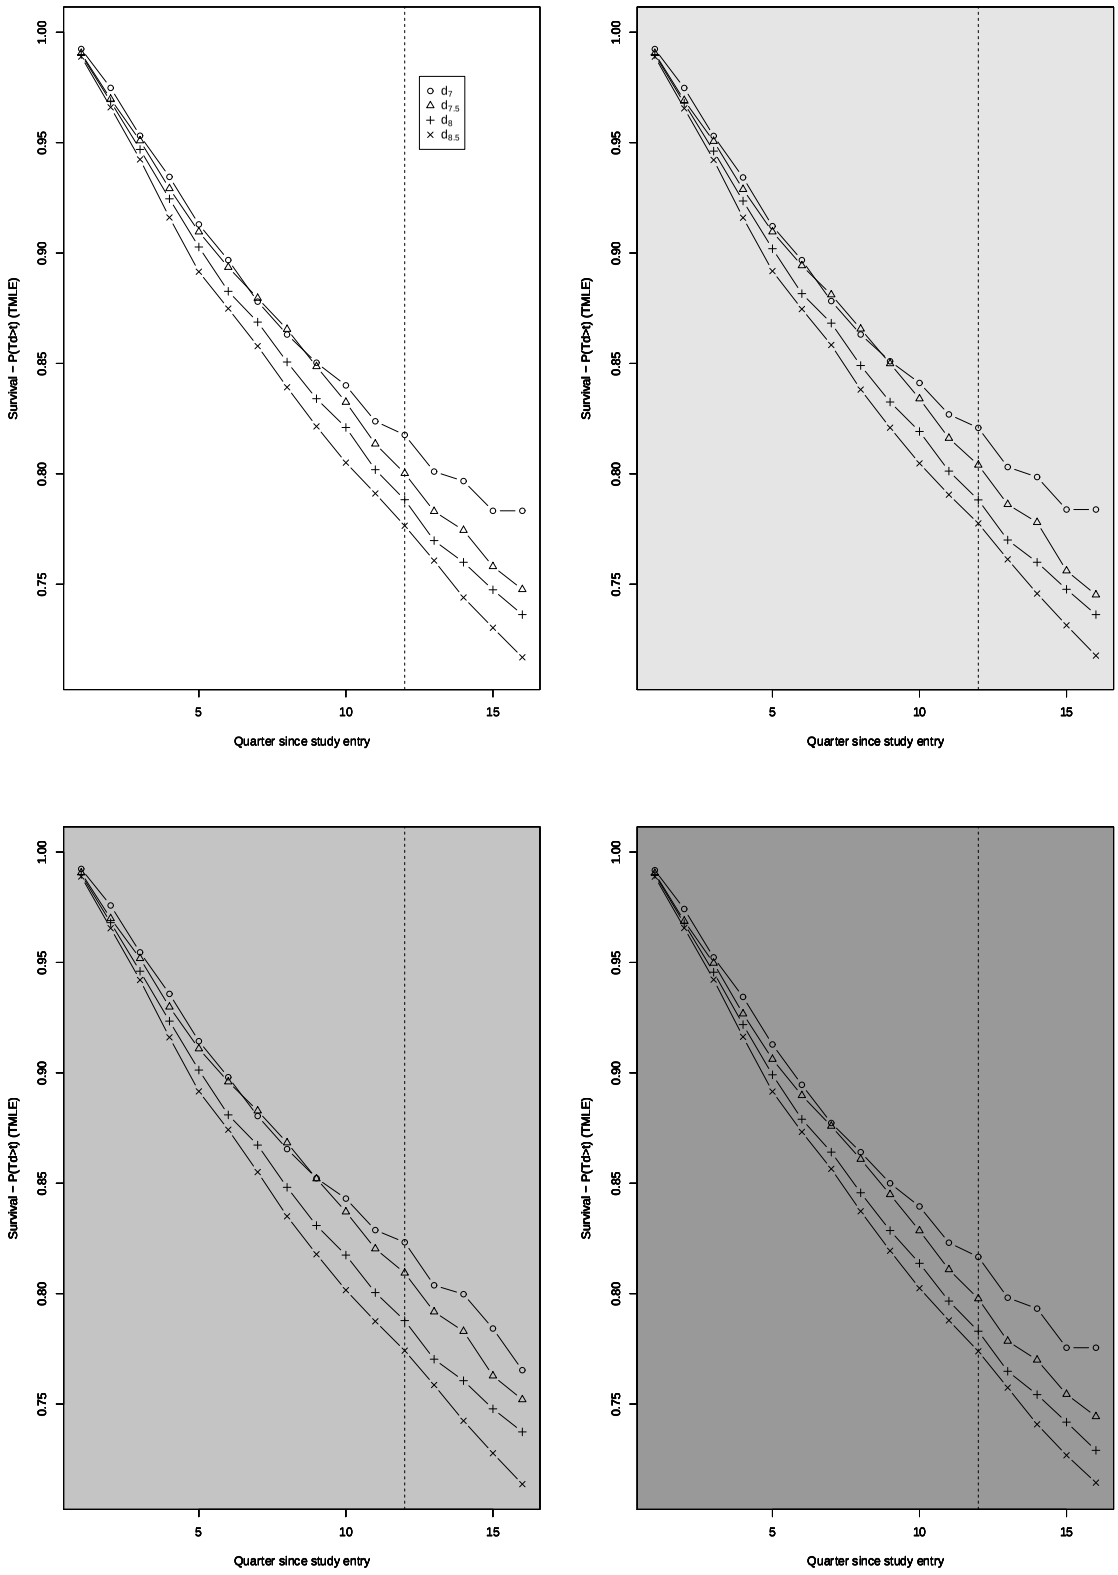

Figure 1. Each plot represents TMLE estimates over 16 quarters of the four counterfactual survival curves corresponding with the four TI initiation strategies d_θ with $\theta = 7, 7.5, 8, 8.5$. The plots located at the top left, top right, bottom left, and bottom right are obtained based on the estimates g_n^θ, $g_{n,t}^\theta$, $g_{n,t,\times}^\theta$, and $g_{n,t,SL}^\theta$ of the nuisance parameter g^θ, respectively.

Table 1. Comparison of inferences from untruncated TMLE and IPW estimation of the 6 RDs at 3 years (12 quarters) based on the four approaches to estimate the nuisance parameter g^θ. Estimates based on TMLE versus IPW estimation are differentiated by the superscript * notation.

			TMLE					IPW estimation					Relative efficiency	
θ_1	θ_2	g^θ	$\psi_n^{\theta_1,\theta_2,*}$	$\psi_n^{\theta_1,\theta_2,*,-}$	$\psi_n^{\theta_1,\theta_2,*,+}$	p*	RSE*	$\psi_n^{\theta_1,\theta_2}$	$\psi_n^{\theta_1,\theta_2,-}$	$\psi_n^{\theta_1,\theta_2,+}$	p	RSE	$\sigma_n^{\theta_1,\theta_2}$	$\frac{\sigma_n^{\theta_1,\theta_2}}{\sigma_n^{\theta_1,\theta_2,*}}$
8.5	8	g_n^θ	0.0118	-5e-03	0.0286	0.167	1.52	0.0109	-3e-04	0.022	0.056	0.99	5.7e-03	
8.5	7.5	g_n^θ	0.0237	-2.3e-03	0.0498	0.074	1.58	0.0212	4.1e-03	0.0382	0.015	0.98	8.7e-03	
8.5	7	g_n^θ	0.0411	9.1e-03	0.0732	0.012	1.23	0.0387	0.0122	0.0652	4e-03	0.98	0.0135	
8	7.5	g_n^θ	0.0119	-0.0117	0.0356	0.323	1.61	0.0103	-5e-03	0.0256	0.187	0.99	7.8e-03	
8	7	g_n^θ	0.0293	-4.8e-03	0.0635	0.092	1.29	0.0278	8e-04	0.0549	0.044	0.98	0.0138	
7.5	7	g_n^θ	0.0174	-0.0165	0.0513	0.315	1.33	0.0175	-8.5e-03	0.0435	0.187	0.97	0.0133	
8.5	8	$g_{n,t}^\theta$	0.0107	-3e-03	0.0243	0.127	1.31	9.9e-03	-9e-04	0.0207	0.071	0.97	5.5e-03	
8.5	7.5	$g_{n,t}^\theta$	0.0264	6.9e-03	0.046	8e-03	1.31	0.025	9.6e-03	0.0404	1e-03	0.99	7.8e-03	
8.5	7	$g_{n,t}^\theta$	0.0433	0.0192	0.0674	0	1	0.0418	0.0172	0.0664	1e-03	1	0.0126	
8	7.5	$g_{n,t}^\theta$	0.0158	-2.1e-03	0.0337	0.084	1.32	0.0151	1.2e-03	0.029	0.033	0.99	7.1e-03	
8	7	$g_{n,t}^\theta$	0.0326	6.7e-03	0.0586	0.014	1.05	0.0319	6.7e-03	0.0571	0.013	0.99	0.0129	
7.5	7	$g_{n,t}^\theta$	0.0168	-7.9e-03	0.0416	0.182	1.08	0.0168	-6.7e-03	0.0403	0.162	0.99	0.012	
8.5	8	$g_{n,t,\times}^\theta$	0.0136	1.2e-03	0.0261	0.032	1.16	0.0129	2.1e-03	0.0237	2e-02	0.98	5.5e-03	
8.5	7.5	$g_{n,t,\times}^\theta$	0.0352	0.0171	0.0533	0	1.2	0.0337	0.0183	0.0491	0	0.98	7.9e-03	
8.5	7	$g_{n,t,\times}^\theta$	0.0491	0.0245	0.0737	0	0.98	0.0483	0.0229	0.0737	0	0.99	0.0129	
8	7.5	$g_{n,t,\times}^\theta$	0.0216	5.2e-03	0.038	1e-02	1.23	0.0209	7.3e-03	0.0345	3e-03	0.98	6.9e-03	
8	7	$g_{n,t,\times}^\theta$	0.0355	0.0102	0.0608	6e-03	1	0.0354	1e-02	0.0608	6e-03	1	0.0129	
7.5	7	$g_{n,t,\times}^\theta$	0.0139	-0.0102	0.038	0.259	1.02	0.0146	-9.2e-03	0.0383	0.23	0.99	0.0121	
8.5	8	$g_{n,t,SL}^\theta$	9.1e-03	1.4e-03	0.0168	0.021	0.99	8e-03	1e-04	0.0159	0.046	0.98	4e-03	1.02
8.5	7.5	$g_{n,t,SL}^\theta$	0.0238	0.0102	0.0374	1e-03	1	0.0216	7.7e-03	0.0356	2e-03	0.98	7.1e-03	1.02
8.5	7	$g_{n,t,SL}^\theta$	0.0427	0.0194	0.0661	0	0.97	0.041	0.0166	0.0655	1e-03	0.98	0.0125	1.05
8	7.5	$g_{n,t,SL}^\theta$	0.0147	2.3e-03	0.0271	2e-02	0.99	0.0136	1e-03	0.0263	0.035	0.98	6.5e-03	1.02
8	7	$g_{n,t,SL}^\theta$	0.0336	0.0104	0.0569	5e-03	0.97	0.033	8.6e-03	0.0574	8e-03	0.98	0.0124	1.05
7.5	7	$g_{n,t,SL}^\theta$	0.0189	-2.5e-03	0.0403	0.083	0.96	0.0194	-3e-03	0.0418	9e-02	0.99	0.0114	1.05

Statist. Med. **2012**, 00 1–25
Prepared using simauth.cls

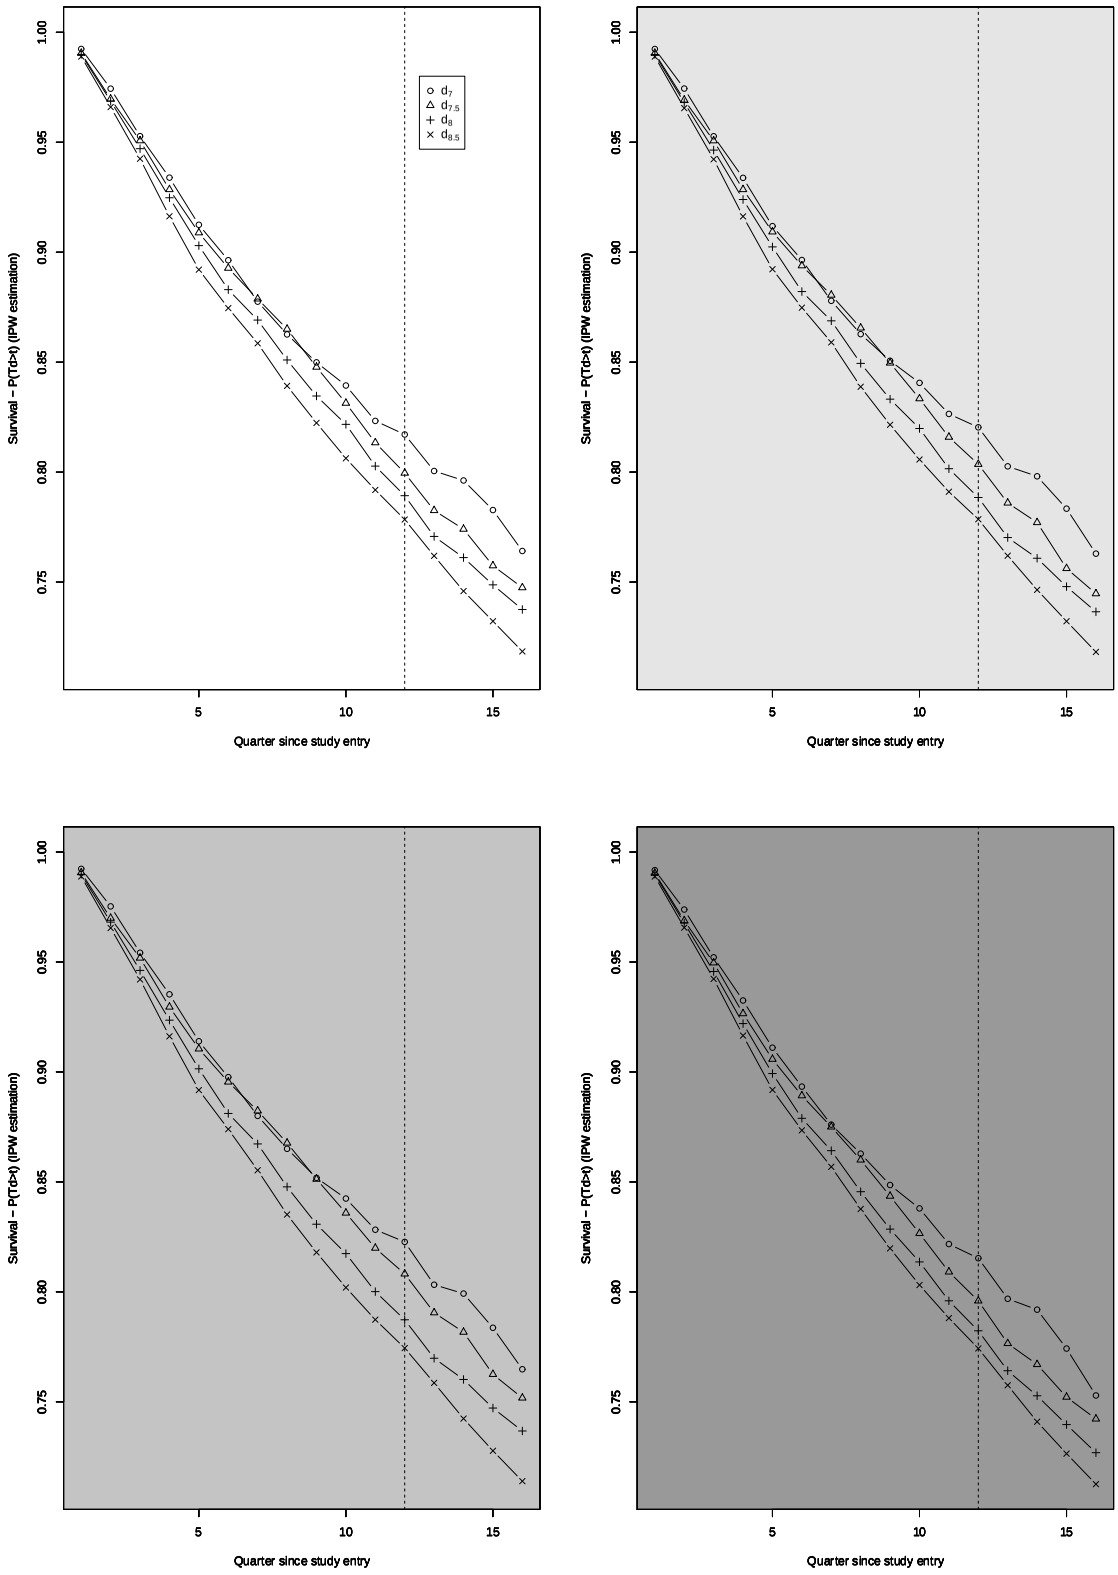

Figure 2. Each plot represents IPW estimates over 16 quarters of the four counterfactual survival curves corresponding with the four TI initiation strategies d_θ with $\theta = 7, 7.5, 8, 8.5$. The plots located at the top left, top right, bottom left, and bottom right are obtained based on the estimates g_n^θ, $g_{n,t}^\theta$, $g_{n,t,\times}^\theta$, and $g_{n,t,SL}^\theta$ of the nuisance parameter g^θ, respectively.

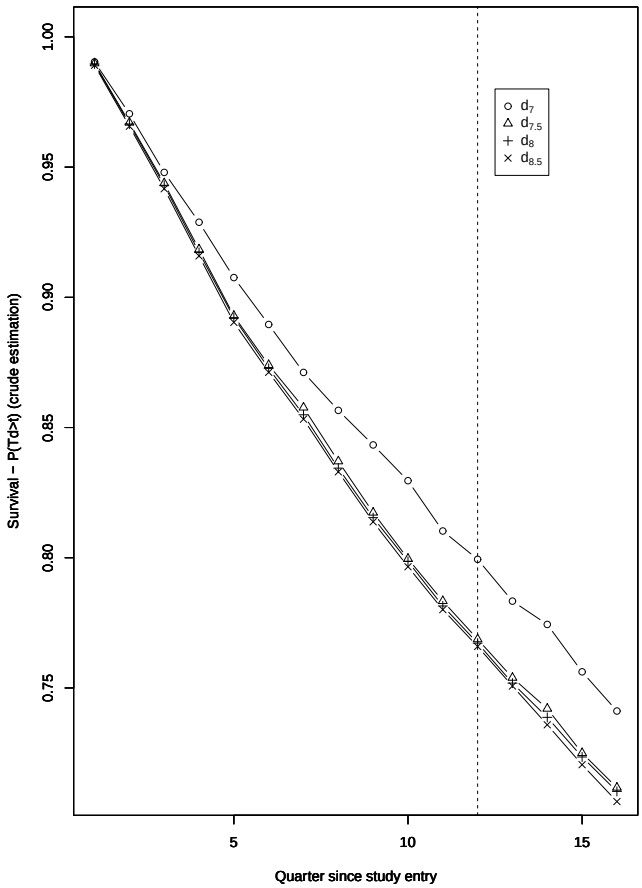

Figure 3. Crude estimates over 16 quarters of the four survival curves associated with the four TI initiation strategies d_θ with $\theta = 7, 7.5, 8, 8.5$.

Table 2. Counts of the number of estimated stabilized IPW weights within specific intervals. For each the four approaches for estimating the weights, these counts describe the distribution of the weight estimates associated with all rule-person-time observations (930956) consistent with a patient following any of the 4 TI decision rules d_θ with $\theta = 7, 7.5, 8, 8.5$. Note that if a patient follows more than one rule at a given time point, her corresponding person-time observation is replicated as many times as the number of rules followed and each replicate is assigned a separate stabilized IPW weight.

Weight range	g^θ_n	$g^\theta_{n,t}$	$g^\theta_{n,t,\times}$	$g^\theta_{n,t,SL}$
<0	0	0	0	0
[0, 0.5[	192037	224426	229283	234103
[0.5, 1[	536883	582320	576979	578316
[1, 10[	189906	116914	119212	117053
[10, 20[	10856	6581	4985	1482
[20, 30[	986	610	410	2
[30, 40[	187	62	40	0
[40, 50[	62	19	19	0
[50, 100[	38	7	24	0
[100, 150[	1	4	3	0
≥ 150	0	13	1	0

Table 3. Comparison of inferences from untruncated TMLE and IPW estimation of the 6 RDs at 3 years (12 quarters) based on SL for estimating the nuisance parameters g^θ and Q^θ. Estimates based on TMLE versus IPW estimation are differentiated by the superscript * notation.

		TMLE					IPW estimation					Relative efficiency
θ_1	θ_2	$\psi_n^{\theta_1,\theta_2,*}$	$\psi_n^{\theta_1,\theta_2,*,-}$	$\psi_n^{\theta_1,\theta_2,*,+}$	p*	$\sigma_n^{\theta_1,\theta_2,*}$	$\psi_n^{\theta_1,\theta_2}$	$\psi_n^{\theta_1,\theta_2,-}$	$\psi_n^{\theta_1,\theta_2,+}$	p	$\sigma_n^{\theta_1,\theta_2}$	$\frac{\sigma_n^{\theta_1,\theta_2}}{\sigma_n^{\theta_1,\theta_2,*}}$
8.5	8	7e-03	-4e-04	0.0143	0.064	3.8e-03	8e-03	1e-04	0.0159	0.046	4e-03	1.07
8.5	7.5	0.0221	9.3e-03	0.0349	1e-03	6.5e-03	0.0216	7.7e-03	0.0356	2e-03	7.1e-03	1.09
8.5	7	0.0386	0.0166	0.0606	1e-03	0.0112	0.041	0.0166	0.0655	1e-03	0.0125	1.11
8	7.5	0.0151	3.5e-03	0.0267	0.011	5.9e-03	0.0136	1e-03	0.0263	0.035	6.5e-03	1.09
8	7	0.0316	9.7e-03	0.0535	5e-03	0.0112	0.033	8.6e-03	0.0574	8e-03	0.0124	1.11
7.5	7	0.0165	-3.8e-03	0.0368	0.11	0.0103	0.0194	-3e-03	0.0418	9e-02	0.0114	1.11

Table 4. Comparison of the properties of IPW estimation and TMLE.The + and - signs indicate relative (potential) advantages/limitations.

	TMLE	**IPW**
Applicable to effects represented by a nonparametric/saturated MSM	✓	✓
Applicable to effects represented by a non-saturated MSM	an alternate algorithm was recently proposed [85]	✓ e.g., [47]
Assumptions required for effect identifiability		
No unmeasured confounders	✓	✓
Positivity	✓	✓
Proper adjustment for time-dependent confounding and selection bias	+ doubly robust	- not doubly robust
Proper analytic variance estimation	- sensitive to bias in g^θ estimation	+ robust to bias in g^θ estimation
Precision	+	-
Feasibility with large, high-dimensional data	+	+
Programming burden		
Point estimation	-	+
Analytic variance estimation	+	+
Computing time	-	+

Table 5. Comparison of the properties of the various estimators of ψ^{θ_1,θ_2} implemented in this report. The two estimators of the nuisance parameter Q^θ based on the DSA algorithm and SL are denoted by $Q^\theta_{n,DSA}$ and $Q^\theta_{n,SL}$, respectively.

	g^θ_n		$g^\theta_{n,t}$		$g^\theta_{n,t,\times}$		$g^\theta_{n,t,SL}$		
	IPW	TMLE $Q^\theta_{n,DSA}$	IPW	TMLE $Q^\theta_{n,DSA}$	IPW	TMLE $Q^\theta_{n,DSA}$	IPW	TMLE $Q^\theta_{n,DSA}$	TMLE $Q^\theta_{n,SL}$
'Sufficient'[†] adjustment for time-dependent confounding and selection bias	✓	✓	✓	✓	✓	✓	✓	✓	✓
Proper[◇] analytic estimation of variance (or ranking by increasing discrepancy with bootstrap estimation)	✓	9	✓	8	✓	7	✓	✓	✓ (not checked[*])
Precision (ranking by increasing variance estimated by bootstrapping or analytically if valid)	9	8[‡]	6.5	4.5[‡]	6.5	4.5[‡]	3	2	1
Robustness to arbitrary parametric assumptions (ranking by decreasing robustness)	9	8	7	6	5	4	3	2	1
Computing burden (ranking by increasing burden)	1	2	3	4	5	6	7	8	9

[†] by 'sufficient', we mean that the inference is concordant with the results from previous randomized trials, i.e., the 6 cumulative RD estimates indicate a clear differentiation and consistent ordering of the effectiveness of the four dynamic treatment interventions to support an increasing beneficial effect of more aggressive therapy initiation rules.

[◇] determined by concordance with bootstrap estimation.

[*] theoretically inferred based on the concordance between the analytic and bootstrap estimates of the variance of TMLE based on $g^\theta_{n,t,SL}$ and $Q^\theta_{n,DSA}$ which suggests consistent estimation of g^θ with $g^\theta_{n,t,SL}$.

[‡] estimated by bootstrapping (data not shown).

References

1. Rosenbaum PR. The consequence of adjustment for a concomitant variable that has been affected by the treatment. *Journal of the Royal Statistical Society, Series A, General* 1984; **147**:656–66.
2. Robins JM. Association, causation and marginal structural models. *Synthese* 1999; **121**:151–179.
3. Hernan MA, Hernandez-Diaz S, Robins JM. A structural approach to selection bias. *Epidemiology* 2004; **15**(5):615–625.
4. Hernan MA, Brumback B, Robins JM. Marginal structural models to estimate the causal effect of zidovudine on the survival of HIV-positive men. *Epidemiology* Sep 2000; **11**(5):561–570.
5. 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**(8):509–515.
6. Toh S, Hernandez-Diaz S, Logan R, Rossouw JE, Hernan MA. Coronary heart disease in postmenopausal recipients of estrogen plus progestin therapy: does the increased risk ever disappear? A randomized trial. *Ann. Intern. Med.* Feb 2010; **152**(4):211–217.
7. Zhang Y, Thamer M, Kaufman JS, Cotter DJ, Hernan MA. High doses of epoetin do not lower mortality and cardiovascular risk among elderly hemodialysis patients with diabetes. *Kidney Int.* Sep 2011; **80**(6):663–669.
8. van der Wal WM, Noordzij M, Dekker FW, Boeschoten EW, Krediet RT, Korevaar JC, Geskus RB. Full loss of residual renal function causes higher mortality in dialysis patients; findings from a marginal structural model. *Nephrol. Dial. Transplant.* Sep 2011; **26**(9):2978–2983.
9. Weintraub WS, Grau-Sepulveda MV, Weiss JM, O'Brien SM, Peterson ED, Kolm P, Zhang Z, Klein LW, Shaw RE, McKay C, *et al.*. Comparative effectiveness of revascularization strategies. *N. Engl. J. Med.* Apr 2012; **366**(16):1467–1476.
10. Robins JM. Robust estimation in sequentially ignorable missing data and causal inference models. *Proceedings of the American Statistical Association*. American Statistical Association: Alexandria, VA, 2000.
11. van der Laan MJ, Robins JM. *Unified methods for censored longitudinal data and causality*. Springer: New York, 2003.
12. Kang JDY, Schafer JL. Demystifying Double Robustness: A Comparison of Alternative Strategies for Estimating a Population Mean from Incomplete Data. *Statistical Science* 2007; **22**(4):523–539.
13. Tsiatis A. *Semiparametric Theory and Missing Data*. Springer: New York, 2006.
14. Bang H, Robins JM. Doubly robust estimation in missing data and causal inference models. *Biometrics* Dec 2005; **61**(4):962–973.
15. Yu Z, van der Laan M. Double robust estimation in longitudinal marginal structural models. *Journal of Statistical Planning and Inference* 2006; **136**(3):1061–1089.
16. van der Laan, Mark J, Rose S. *Targeted Learning: Causal Inference for Observational and Experimental Data*. Springer: New York, 2011.
17. van der Laan M, Polley E, Hubbard A. Super learner. *Statistical Applications in Genetics and Molecular Biology* 2007; **6**(1).
18. Neugebauer R, Silverberg MJ, van der Laan MJ. *Observational study and individualized antiretroviral therapy initiation rules for reducing cancer incidence in HIV-infected patients*, chap. 26. Springer: New York, 2011; 436–456.
19. Stitelman OM, Gruttola VD, van der Laan MJ. A general implementation of TMLE for longitudinal data applied to causal inference in survival analysis. *Technical Report 281*, Division of Biostatistics, UC Berkeley 2011.
20. van der Laan MJ, Gruber S. Targeted Minimum Loss Based Estimation of Causal Effects of Multiple Time Point Interventions. *Int J Biostat* 2012; **8**(1):Article 9.
21. van der Laan M, Rubin D. Targeted Maximum Likelihood Learning. *Int J Biostat* 2006; **2**(1).
22. van der Laan MJ. The Construction and Analysis of Adaptive Group Sequential Designs. *Technical Report 232*, Division of Biostatistics, UC Berkeley 2008.
23. Nathan DM, Buse JB, Davidson MB, Heine RJ, Holman RR, , Sherwin R, Zinman B. 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.
24. Skyler JS, Bergenstal R, Bonow RO, Buse J, Deedwania P, Gale EA, Howard BV, Kirkman MS, Kosiborod M, Reaven P, *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.
25. The Diabetes Control and Complications Trial Research Group. The effect of intensive treatment of diabetes on the development and progression of long-term complications in insulin-dependent diabetes mellitus. *N Engl J Med.* 1993; **329**:977–86.
26. Nathan DM, Cleary PA, Backlund JY, Genuth SM, Lachin JM, Orchard TJ, Raskin P, Zinman B. 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**(353):2643–53.
27. UK Prospective Diabetes 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**:854–865.
28. 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.
29. 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.
30. ADVANCE Collaborative Group. Intensive blood glucose control and vascular outcomes in patients with type 2 diabetes. *N Engl J Med* 2008; **358**:2560–2.
31. Duckworth W, Abraira C, Moritz T, Reda D, Emanuele N, Reaven PD, Zieve FJ, Marks J, Davis SN, Hayward R, *et al.*. Glucose Control and Vascular Complications in Veterans with Type 2 Diabetes. *N Engl J Med* 2009; **360**:129–39.
32. Ray KK, Seshasai SR, Wijesuriya S, Sivakumaran R, Nethercott S, Preiss D, Erqou S, Sattar N. 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.
33. Gerstein HC, Miller ME, Byington RP, Goff DC, Bigger JT, Buse JB, Cushman WC, Genuth S, Ismail-Beigi F, Grimm RH, *et al.*. Effects of intensive glucose lowering in type 2 diabetes. *N. Engl. J. Med.* Jun 2008; **358**(24):2545–2559.
34. Patel A, MacMahon S, Chalmers J, Neal B, Billot L, Woodward M, Marre M, Cooper M, Glasziou P, Grobbee D, *et al.*. Intensive blood glucose control and vascular outcomes in patients with type 2 diabetes. *N. Engl. J. Med.* 2008; **358**(24):2560–2572.
35. Duckworth W, Abraira C, Moritz T, Reda D, Emanuele N, Reaven PD, Zieve FJ, Marks J, Davis SN, Hayward R, *et al.*. Glucose control and vascular complications in veterans with type 2 diabetes. *N. Engl. J. Med.* Jan 2009; **360**(2):129–139.

36. Ismail-Beigi F, Craven T, Banerji MA, Basile J, Calles J, Cohen RM, Cuddihy R, Cushman WC, Genuth S, Grimm RH, *et al.*. Effect of intensive treatment of hyperglycaemia on microvascular outcomes in type 2 diabetes: an analysis of the ACCORD randomised trial. *Lancet* Aug 2010; **376**:419–430.
37. O'Connor PJ, Ismail-Beigi F. Near-normalization of glucose and microvascular diabetes complications: data from ACCORD and ADVANCE. *Therapeutic Advances in Endocrinology and Metabolism* 2011; **2**(1):17–26.
38. *Third Report of the National Cholesterol Education Program (NCEP) Expert Panel on Detection* 2002.
39. Nathan DM, Buse JB, Davidson MB, Ferrannini E, Holman RR, Sherwin R, Zinman B. 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* Jan 2009; **32**(1):193–203.
40. 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**(9):643–8.
41. Robins J. Marginal Structural Models. *1997 Proceedings of the American Statistical Association, Section on Bayesian Statistical Science*, 1998; 1–10.
42. Murphy S, van der Laan M, Robins J. Marginal mean models for dynamic treatment regimens. *Journal of the American Statistical Association* 2001; **96**:1410–1424.
43. van der Laan M, Petersen ML. Causal Effect Models for Realistic Individualized Treatment and Intention to Treat Rules. *Int J Biostat* 2007; **3**(1):Article 3.
44. Robins J RA Orellana L. Estimation and extrapolation of optimal treatment and testing strategies. *Stat Med* 2008; **27**(23):4678–4721.
45. Hernan MA, Brumback BA, Robins JM. Estimating the causal effect of zidovudine on CD4 count with a marginal structural model for repeated measures. *Statistics in Medicine* 2002; **21**:1689–1709.
46. Robins JM, Hernan MA, Brumback B. Marginal structural models and causal inference in epidemiology. *Epidemiology* 2000; **11**(5):550–560.
47. 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* May 2012; **21 Suppl 2**:99–113.
48. Neugebauer R, Fireman B, Roy JA, O'Connor PJ. Dynamic marginal structural modeling to assess impact of more versus less intensive glycemic control strategies on microvascular and macrovascular outcomes in 58,000 adults with type 2 diabetes mellitus. *Submitted to Annals of Internal Med.* 2013; .
49. Pearl J. Causal inference in statistics: An overview. *Statistics Surveys* 2009; **3**:96–146.
50. VanderWeele TJ. Concerning the consistency assumption in causal inference. *Epidemiology* Nov 2009; **20**(6):880–883.
51. Pearl J. On the consistency rule in causal inference: axiom, definition, assumption, or theorem? *Epidemiology* Nov 2010; **21**(6):872–875.
52. Cole SR, Hernan MA, Robins JM, Anastos K, Chmiel J, Detels R, Ervin C, Feldman J, Greenblatt R, Kingsley L, *et al.*. Effect of highly active antiretroviral therapy on time to acquired immunodeficiency syndrome or death using marginal structural models. *Am. J. Epidemiol.* Oct 2003; **158**(7):687–694.
53. 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.* May 2004; **159**(10):926–934.
54. Hernan MA, Lanoy E, Costagliola D, Robins JM. Comparison of dynamic treatment regimes via inverse probability weighting. *Basic Clin. Pharmacol. Toxicol.* Mar 2006; **98**(3):237–242.
55. Wang Y, Bembom O, van der Laan MJ. Data-adaptive estimation of the treatment-specific mean original research article. *Journal of Statistical Planning and Inference* 2007; **137**(6):1871–1887.
56. Platt RW, Brookhart AM, Cole SR, Westreich D, Schisterman EF. An information criterion for marginal structural models. *Stat Med* Sep 2012; .
57. van der Vaart AW. *Asymptotic Statistics*. Cambridge University Press, 1998.
58. Sinisi SE, van der Laan MJ. Deletion/substitution/addition algorithm in learning with applications in genomics. *Stat Appl Genet Mol Biol* 2004; **3**:Article18.
59. Neugebauer R, Bullard J. DSA R package (version 3.1.4). http://www.stat.berkeley.edu/ laan/Software/index.html 2010.
60. R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria 2011. URL `http://www.R-project.org/`, ISBN 3-900051-07-0.
61. Robins J, Sued M, Lei-Gomez Q, Rotnitzky A. Comment: Performance of Double-Robust Estimators When "Inverse Probability" Weights are Highly Variable. *Statistical Science* 2007; **22**(4):544–559.
62. 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.* Oct 2007; **45**:908–915.
63. 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* Aug 2010; **37**:512–518.
64. Breiman L. Random forests. *Machine Learning* 2001; **45**:5–32.
65. Efron B, Hastie T, Johnstone I, Tibshirani R. Least angle regression. *Annals of Statistics* 2004; **32**(2):407–499.
66. Ruczinski I, Kooperberg C, LeBlanc M. Logic regression. *Journal of Computational and Graphical Statistics* 2003; **12**(3):475–511.
67. Breiman L, Friedman JH, Olshen RA, Stone CJ. *Classification and Regression Trees*. Wadsworth & Brooks/Cole: Monterey, 1984.
68. Hoerl A, Kennard R. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics* 1970; **12**(3):55–67.
69. Friedman J. Multivariate adaptive regression splines. *Annals of Statistics* 1991; **19**(1):1–141.
70. Kooperberg C, Bose S, Stone CJ. Polychotomous Regression. *Journal of the American Statistical Association* 1997; **92**(437):117–127.
71. 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* Jul 2009; **20**(4):512–522.
72. Ridgeway G, McCaffrey DF. Comment: Demystifying Double Robustness: A Comparison of Alternative Strategies for Estimating a Population Mean from Incomplete Data. *Statistical Science* 2007; **22**(4):540–543.
73. Dudoit S, van der Laan M. Asymptotics of cross-validated risk estimation in estimator selection and performance assessment. *Statistical Methodology* 2005; **2**:131–154.
74. van der Laan M, Dudoit S, Keles S. Asymptotic optimality of likelihood-based cross-validation. *Statistical Applications in Genetics and Molecular Biology* 2004; **3**(1).
75. van der Vaart A, Dudoit S, van der Laan M. Oracle inequalities for multi-fold cross-validation. *Statistics and Decisions* 2006; **24**(3):351–371.
76. van der Laan M, Dudoit S, van der Vaart A. The cross-validated adaptive epsilon-net estimator. *Statistics and Decisions* 2006; **24**(3):373–395.
77. Polley EC. SuperLearner R package (version 1.1-18). `https://github.com/ecpolley/SuperLearner` 2011.

78. 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. *Journal of Clinical Epidemiology* 2013; **Accepted**.
79. Cole SR, Hernan MA. Constructing inverse probability weights for marginal structural models. *Am. J. Epidemiol.* Sep 2008; **168**:656–664.
80. Bembom O, van der Laan MJ. Data-adaptive selection of the truncation level for Inverse-Probability-of-Treatment-Weighted estimators. *Technical Report 230*, Division of Biostatistics, UC Berkeley 2008.
81. Petersen ML, Porter KE, Gruber S, Wang Y, van der Laan MJ. Diagnosing and responding to violations in the positivity assumption. *Stat Methods Med Res* Feb 2012; **21**(1):31–54.
82. Hernan MA, Robins JM. Estimating causal effects from epidemiological data. *J Epidemiol Community Health* Jul 2006; **60**(7):578–586.
83. Carone M, M J van der Laan. Efficient Substitution Estimation of Counterfactual Survival Function based on Censored Data 2012. Private communication.
84. Neugebauer R, van der Laan M. Nonparametric causal effects based on marginal structural models. *Journal of Statistical Planning and Inference* 2006; **137**(2):419–434.
85. Petersen M, Schwab J, Gruber S, Keiser O, Egger M, van der Laan M. Targeted Maximum Likelihood Estimation for Longitudinal Marginal Structural Working Models. *Journal of Causal Inference* 2013; **Submitted**.
86. Toh S, Hernan MA. Causal inference from longitudinal studies with baseline randomization. *Int J Biostat* 2008; **4**(1):Article 22.
87. Moore KL, Neugebauer R, Valappil T, Laan MJ. Robust extraction of covariate information to improve estimation efficiency in randomized trials. *Stat Med* Aug 2011; **30**(19):2389–2408.