

This is a repository copy of *Exploiting non-systematic covariate monitoring to broaden the scope of evidence about the causal effects of adaptive treatment strategies*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/159869/>

Version: Accepted Version

Article:

Kreif, Noemi, Sofrygin, Oleg, Schmittdiel, Julie A et al. (5 more authors) (2020) Exploiting non-systematic covariate monitoring to broaden the scope of evidence about the causal effects of adaptive treatment strategies. *Biometrics* (Journal of the International Biometric Society). ISSN 1541-0420

<https://doi.org/10.1111/biom.13271>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



Exploiting non-systematic covariate monitoring to broaden the scope of evidence about the causal effects of adaptive treatment strategies

Noémi Kreif¹, Oleg Sofrygin², Julie A. Schmittdiel²,
Alyce S. Adams², Richard W. Grant², Zheng Zhu², Mark J. van der Laan³, and Romain Neugebauer^{2,*}

¹Centre for Health Economics, University of York, York, UK

²Division of Research, Kaiser Permanente Northern California, Oakland, CA, U.S.A.

³Division of Biostatistics, School of Public Health, University of California, Berkeley, CA, U.S.A.

**email*: romain.s.neugebauer@kp.org

SUMMARY: In studies based on electronic health records (EHR), the frequency of covariate monitoring can vary by covariate type, across patients, and over time, which can limit the generalizability of inferences about the effects of adaptive treatment strategies. In addition, monitoring is a health intervention in itself with costs and benefits, and stakeholders may be interested in the effect of monitoring when adopting adaptive treatment strategies. This paper demonstrates how to exploit non-systematic covariate monitoring in EHR-based studies to both improve the generalizability of causal inferences and to evaluate the health impact of monitoring when evaluating adaptive treatment strategies. Using a real world, EHR-based, comparative effectiveness research (CER) study of patients with type II diabetes mellitus, we illustrate how the evaluation of joint dynamic treatment and static monitoring interventions can improve CER evidence and describe two alternate estimation approaches based on inverse probability weighting (IPW). First, we demonstrate the poor performance of the standard estimator of the effects of joint treatment-monitoring interventions, due to a large decrease in data support and concerns over finite-sample bias from near-violations of the positivity assumption (PA) for the monitoring process. Second, we detail an alternate IPW estimator using a no direct effect (NDE) assumption. We demonstrate that this estimator can improve efficiency but at the potential cost of increase in bias from violations of the PA for the treatment process.

KEY WORDS: dynamic treatment regimes, electronic health records, inverse probability weighting, joint interventions, monitoring regimes, no direct effect assumption, time-dependent confounding

This paper has been submitted for consideration for publication in *Biometrics*

This article has been accepted for publication and undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process, which may lead to differences between this version and the Version of Record. Please cite this article as doi: 10.1111/biom.13271

1. Dynamic treatment regimes and non-systematic covariate monitoring

In the management of chronic conditions such as diabetes, adaptive treatment strategies might lead to better outcomes than static treatment regimes. A static treatment regime is a sequence of treatment decisions that are predetermined and the same for all patients that experience it. An adaptive treatment strategy (a.k.a. dynamic treatment regime) is a sequence of treatment decisions that are each updated as a function of the changing circumstances of the patient. The function used to update each treatment decision is referred to as a decision rule. Because dynamic regimes allow for personalization of treatment decisions over time, they better reflect real-world decision making that follows the chronic care model (Bodenheimer et al., 2002). Dynamic treatment regimes can be evaluated using randomized experiments (SMART trials, see e.g. Murphy et al. (2007)) and increasingly in comparative effectiveness research, using observational studies (Cain et al., 2010). Regardless of the study design, evaluating adaptive treatment strategies requires data collection on covariates that enter the decision rules of interest. The pattern and frequency of covariate measurements vary across study designs. Even in randomized trials, the lengths of intervals between monitoring events can vary (Ford et al., 2015). Observational studies make increasing use of routinely collected data sources which were not originally intended for research purposes. In retrospective cohort studies based on electronic health records (EHR) for example, investigators have no control over the data collection process and monitoring is instead a joint decision of the patient and the health care providers, and is not expected to be synchronized between patients. This monitoring variability can lead to the following challenge and opportunity for research.

First, when aiming to establish the optimal dynamic treatment regime, optimality will depend on the frequency of monitoring (Robins et al., 2008). Hence, the difference in monitoring protocols between two populations may invalidate the extrapolation of study results

obtained in one population to the other. Second, monitoring in itself is a health intervention with non-negligible costs to healthcare systems that can also burden patients financially or otherwise. Hence establishing the optimal monitoring regime may be of interest in its own right, and the monitoring variability in retrospective observational studies based on EHR data provides an opportunity to evaluate the health impact of monitoring decisions.

While statistical methods for estimating the causal effects of time-varying interventions (either static or dynamic) are well developed (Robins et al., 2000, 2008; van der Laan and Petersen, 2007) - and we assume basic familiarity with their underlying concepts and application - there is little practical guidance on how to handle the challenge and opportunity just described. A notable exception is Robins et al. (2008) who theoretically developed estimators for the effects of joint dynamic treatment and monitoring strategies which can be used to both improve the generalizability of inferences (see Section 6.2.1 of their work) and evaluate the health impact of monitoring decisions. The standard inverse probability weighting (IPW) estimator for these effects relies on a positivity assumption (PA) about the monitoring process that may be theoretically or practically violated. Robins et al. (2008) propose a new IPW estimator that weakens this PA by introducing a “no direct effect” (NDE) assumption. Neugebauer et al. (2017) build on these results within the nonparametric structural equation modeling framework by deriving new identifiability results, proving the equality of counterfactual outcomes of an original and a modified joint regime using the NDE assumption. This facilitates the derivation of an estimator for the effect of the original joint regime by developing an estimator for the effect of the modified joint regime.

There are few studies that consider the causal effects of monitoring interventions in real-world evaluations. Hernán et al. (2009) considers the effects of joint static treatment and static monitoring interventions where it is assumed that treatment can only change when covariates are monitored. Caniglia et al. (2016) compare dynamic monitoring strategies

(without a joint treatment intervention) where the decision rule for the frequency of monitoring depends on the CD4 cell count. Ford et al. (2015) apply the methods proposed by Robins et al. (2008) to estimate the effect of monitoring with dynamic treatment regimes. These studies focus on applications in HIV research using cohort studies with some level of investigator-control of the data collection process. Additional references are included in the Supporting Information and include reports of statistical methods that are not explicitly linked to a formal causal model but that aim to address related problems with an informative observation process. Here, we apply methods developed in previous theoretical work (Neugebauer et al., 2017) to compare the assumptions and practical performance of standard versus NDE-based IPW estimation of the effects of joint dynamic treatment and static monitoring interventions in a diabetes study based on large EHR data. All analyses were implemented using the `stremr` R package (Sofrygin et al., 2019).

2. Study design

We build on a retrospective cohort study - referred to as the “TI study” - designed to evaluate the effect of various glucose-lowering strategies on clinical outcomes of type II diabetes adults. The available data set includes longitudinal measurements on 51,179 patients from 7 US regions followed for a median time of about 3 years. To be included in the cohort, a patient had to have one or more A1c test $< 7\%$ followed by an A1c test between 7% and 8.5% while taking two or more antidiabetic oral agents or basal insulin. Follow-up started on the date of the A1c $\geq 7\%$ and ended at the earliest of the time to the failure event (i.e., albuminuria development or progression), or the time to a censoring event (i.e., death, disenrollment from the health plan, or administrative end of study). For details on the study context and data, we refer the reader to Neugebauer et al. (2012) (e.g., Appendix E), Neugebauer et al. (2013) (e.g., Table I) and Neugebauer et al. (2015) (e.g., Table I). We extend previous analyses in these articles by evaluating the per-protocol effects of 6 a priori specified joint

dynamic treatment and static monitoring interventions (see Supporting Information for the rationale that led to their selection). The three dynamic treatment interventions considered require that patients initiate treatment intensification (TI) the first time a newly measured A1c reaches a given threshold and that they remain on the intensified therapy thereafter. We consider two monitoring interventions that require periodic A1c monitoring. In the Supporting Information, we consider 10 additional joint interventions.

3. Observed data notation and causal model

For each patient in the cohort, measurements on treatment, monitoring, covariate, censoring and outcome information are updated every 90 days starting at study entry and until the end of follow-up. This analytic unit of time was used in the original analyses of the study data because A1c monitoring was not expected to occur more than once every 90 days. We recently presented sensitivity analyses based on finer units of time (Sofrygin et al., 2019). Follow-up time (denoted by t) is thus expressed in 90-day units, and the first 90 days of follow-up is denoted by $t = 0$. The latest possible follow-up time is denoted by $K + 1 = 36$, corresponding to about 9 years. The binary outcome variable $Y(t)$ represents whether the failure event occurred at the previous time point $t - 1$, and $Y(0) = 0$ by convention. The vector $A(t) = \{A_1(t), A_2(t)\}$ denotes the exposure status at time t and is composed of two binary variables that indicate whether the patient experienced an intensified therapy and a censoring event at time t , respectively. The vector of covariates $Z(t)$ represents patient attributes measured before $A(t)$ (e.g., comorbidity, diagnoses, or vital signs). The covariate $I(t)$ represents the A1c level measured before $A(t)$. $I(t)$ is not monitored at every $t > 0$ for all patients in this study. The binary variable $N(t)$ represents the monitoring decision at time t and indicates whether a measurement for $I(t + 1)$ will be taken. By convention, when A1c is not monitored at time t , $I(t)$ is defined as $I(j)$ where $j < t$ is the latest time point when A1c was monitored. By study design, A1c is always monitored in the first period. We note that,

unlike prior related work (Hernán et al., 2009), treatment $A_1(t)$ can change at any time and not necessarily only when the covariate $I(t)$ is monitored (i.e., not only at t when $N(t-1) = 1$). We jointly denote the outcome and covariates at time t by $L(t) \equiv \{Y(t), Z(t), I(t)\}$ which consists of measurements obtained before $A(t)$. Similar to $I(t)$, each covariate in $Z(t)$ that was not monitored at time t (e.g., blood pressure) is defined using last observed value carried forward (LOVCF) and an indicator of LOVCF is included in the vector of covariates $Z(t)$. Support for this approach to handle partially observed covariates is detailed in Section 2.2 of Kreif et al. (2018) and uses a framework and rationale related to that described in Hernán et al. (2009). To simplify notation, all variables become degenerate after a failure or censoring event occurs. The observed data are realizations of n independent and identically distributed copies O_i of $O = \{L(0), A(0), N(0), \dots, L(K), A(K), N(K), L(K+1)\}$. We use the overbar to denote the history of a variable, e.g. $\bar{Z}(t) = \{Z(0), \dots, Z(t)\}$.

We assume a nonparametric structural equation model (NPSEM) detailed in Neugebauer et al. (2017) that links the observed data distribution to a vector of random disturbances and a fixed vector of functions. We note that the random disturbances of this causal model are not assumed to be mutually independent. In this NPSEM, the observed level of A1c, $I(t)$, is linked to a latent variable $I^0(t)$ that represents the patient's underlying A1c level at time t . More specifically, if a decision to measure $I^0(t)$ was made during follow-up, we observe $I(t)$ to be $I^0(t)$, and otherwise the value of $I^0(t)$ is missing from the observed data O . Because A1c is always monitored in the first period, we have $I(0) = I^0(0)$ for all patients. This NPSEM also defines all counterfactual outcomes referenced in this report (Pearl, 2010).

4. Evaluation of joint dynamic treatment and static monitoring interventions

4.1 Causal estimands

We consider three exposure interventions which each sets the random variable $A(t) = \{A_1(t), A_2(t)\}$ to a realization $a(t) = \{a_1(t), a_2(t)\}$ for each time period, $t = 0, \dots, t_0$ where

$t_0 < K$ is the time point chosen by the analyst when outcomes are compared between intervention groups. Specifically, each intervention is the combination of a dynamic treatment regime and a static right-censoring regime that ensures outcome collection at time t_0 (i.e., $\bar{a}_2(t_0) = 0$ - see Appendix F in the Supporting Information for our selected approach to the competing risk of death). The treatment decisions $a_1(t)$ are adapted to time-varying realizations of patient characteristics based on the following decision rules: “initiate TI the first time a newly measured A1c reaches $x\%$, and remain on intensified therapy thereafter” with $x=7.5; 8; 8.5$. Each intervention requires setting the treatment to $a_1(t) = 1$ for a patient who has not failed yet ($Y(t) = 0$) and who has not yet initiated TI ($A(t-1) = 0$), but for whom a new A1c measurement ($N(t-1) = 1$) has been taken and its level ($I(t)$) has reached threshold x , otherwise, the intervention requires that treatment be not changed (i.e., $a_1(t) = a_1(t-1)$). We formalize the decision rules that define these interventions with vectors of mappings $d_x = (d_x(0), \dots, d_x(K))$ where $d_x(t) : \{A(t-1), N(t-1), Y(t), I(t)\} \mapsto \{a_1(t), a_2(t)\}$ for $t = 0, \dots, K$ and $x \in \mathcal{X} = \{7.5, 8, 8.5\}$. The exposure regime through time t defined by applying the sequence of decision rules d_x with a patient’s observed data O is denoted by $d_x\{\bar{V}(t)\} = [d_x(0)\{V(0)\}, \dots, d_x(t)\{V(t)\}]$, where $V(t) = \{A(t-1), N(t-1), Y(t), I(t)\}$. Figure 6 (top panel) illustrates the dynamic treatment regimes resulting from the application of such decision rules d_x using fake data from two patients (see ‘Step 3A’). We consider two static monitoring regimes defined by regular A1c testing schedules: skip one quarter between A1c tests (i.e., $\bar{n}_1 = (0, 1, 0, 1, \dots)$) and skip 3 quarters between A1c tests (i.e., $\bar{n}_3 = (0, 0, 0, 1, 0, 0, 0, 1, \dots)$). We aim to estimate the causal effects of the six joint interventions (d_x, \bar{n}_y) defined by the three dynamic treatment-censoring regimes d_x with $x \in \mathcal{X}$ and the two static monitoring regimes \bar{n}_y , with $y \in \mathcal{Y} = \{1, 3\}$. The counterfactual outcomes under these joint interventions are denoted by $Y_{d_x, \bar{n}_y}(t+1)$. Their distributions define the causal estimands:

$$\psi^{(d_{x_1}, \bar{n}_{y_1}), (d_{x_2}, \bar{n}_{y_2})}(t_0 + 1) = P\{Y_{d_{x_1}, \bar{n}_{y_1}}(t_0 + 1) = 1\} - P\{Y_{d_{x_2}, \bar{n}_{y_2}}(t_0 + 1) = 1\}, \quad (1)$$

with $(x_1, x_2) \in \mathcal{X}^2$ and $(y_1, y_2) \in \mathcal{Y}^2$. Nine such risk differences were estimated for $t_0 = 0, \dots, 7$ to 1) improve the generalizability of evidence about the comparative effectiveness of the three dynamic treatment regimes by fixing monitoring (i.e., $y_1 = y_2$) and 2) to evaluate the impact of monitoring on the effectiveness of a fixed dynamic treatment regime (i.e., $x_1 = x_2$).

4.2 Identifying assumptions

Identifiability of the above parameters with the observed data O relies on a sequential randomization assumption (SRA) and a PA. Identifiability directly follows from application of the g-formula (Robins, 1986) when the time-varying intervention is expanded to include monitoring. We note that consistency, i.e., the fact that the counterfactual outcomes under the observed treatment-monitoring regime correspond to the observed outcomes follows from the assumed NPSEM (Pearl, 2010). For the time-varying joint interventions that define each of the counterfactual risks in (1), the SRA can be expressed as:

$$Y_{d_x, \bar{n}_y}(t_0 + 1) \perp \{A(t), N(t)\} | \bar{L}(t), \bar{A}(t-1), \bar{N}(t-1), \quad (2)$$

for $t = 0, \dots, t_0$, $x \in \mathcal{X}$ and $y \in \mathcal{Y}$. It states that, for all time points through t_0 , each potential outcome of interest is independent of the current exposure and monitoring decisions conditional on past exposures, monitoring decisions, and confounders.

In EHR-based studies, it is expected that treatment decisions, right-censoring events, or monitoring decisions are affected by factors that also impact the outcome. The SRA holds if all these factors are observed (i.e., included in $\bar{L}(t_0)$). As illustrated below, the SRA can also hold even when some of these factors are unobserved as long as they only affect the exposure or monitoring through observed covariates. Because we defined the exposure as the vector of treatment and censoring status, we note that (2) includes the assumption of no unmeasured sources of selection bias (due to right-censoring).

The causal directed acyclic graph (DAG) depicted by Figure 1 represents commonly

assumed causal relationships between observed treatment, covariates, and outcomes (Robins et al., 2000, fig. 1b), adapted to cohort studies with a time-to-event outcome, and extended to also represent non-systematic monitoring. For simplicity, it is assumed that no covariate other than the A1c measurements and outcomes are collected, follow-up spans only two time points, and right-censoring does not occur. In addition to the observed variables and the latent variable $I^0(t)$, the DAG also includes two potentially unobserved time-varying covariates U_1 and U_2 (e.g., health-seeking behavior such as diet and physical activity) that are risk factors for the outcomes $Y(t)$ and that only impact treatments through observed covariates but that can directly influence the monitoring indicator $N(0)$. The upholding of the SRA (2) can be motivated by ensuring that all backdoor paths from $A_1(t)$, $A_2(t)$, and $N(t)$ to $Y(t_0 + 1)$ are blocked by prior measured variables included in the observed data O . For instance, the SRA will hold if the two gray dashed arrows from U_1 and U_2 to $N(0)$ are deleted and in the absence of an arrow from $I^0(1)$ to $A_1(1)$ which encodes a commonly made assumption that unobserved risk factors for the outcomes can only impact treatment decisions through observed covariates. For instance, a clinician's decision to prescribe a new treatment cannot be influenced by the patient's A1c level if this level was not known to the clinician except through other mediators reported to the clinicians (e.g., symptoms shared by the patient). Such reported information must then be included in the vector of covariate $Z(t)$ for the SRA to hold.

The identifiability of the causal estimand (1) also hinges on a PA:

$$\begin{cases} P[A(t) = d_x(t)\{V(t)\} \mid \bar{L}(t), \bar{Y}(t) = 0, \bar{A}(t-1) = d_x\{\bar{V}(t-1)\}, \bar{N}(t-1) = \bar{n}_y(t-1)] > 0, \\ P[N(t) = n_y(t) \mid \bar{L}(t), \bar{Y}(t) = 0, \bar{N}(t-1) = \bar{n}_y(t-1), \bar{A}(t) = d_x\{\bar{V}(t)\}] > 0, \end{cases} \quad (3)$$

for $t = 0, \dots, t_0$, $x \in \mathcal{X}$ and $y \in \mathcal{Y}$. In particular, this PA requires that for all time periods and any combination of past covariates, there is a positive probability for each patient who previously followed the joint intervention to continue to follow the *monitoring* intervention of interest. Near-violations of this assumption can occur if certain covariates

(e.g., change in A1c control) are strong determinants of monitoring decisions (e.g., the American Diabetes Association recommends that patients who recently changed treatments or whose A1c recently became out of control be monitored more frequently (p. S57 American Diabetes Association (2018))). Near-violations of this assumption due to chance may also be expected if only a small number of patients follow the intervention under evaluation and, in particular, with a large covariate adjustment set (Petersen et al., 2012), and these concerns increase over time as more patients experience a failure or censoring event.

5. Evaluation of the joint interventions under the no direct effect assumption

5.1 Causal estimands

Under an NDE assumption detailed below, the risk differences (1) were shown (Neugebauer et al., 2017) to equal the following causal estimands which are the focus of this Section:

$$\psi^{(d_{x_1}^*, g_{y_1}^*), (d_{x_2}^*, g_{y_2}^*)}(t_0 + 1) = P\{Y_{d_{x_1}^*, g_{y_1}^*}(t_0 + 1) = 1\} - P\{Y_{d_{x_2}^*, g_{y_2}^*}(t_0 + 1) = 1\}, \quad (4)$$

where each $d_{x_j}^*$ with $x_j \in \mathcal{X}$ is a modified version of the dynamic treatment-censoring intervention d_{x_j} and where each $g_{y_j}^*$ with $y_j \in \mathcal{Y}$ is a static monitoring intervention on a subset of the monitoring process. Both interventions $d_{x_j}^*$ and $g_{y_j}^*$ are defined as follows based on the same static monitoring intervention \bar{n}_{y_j} . Instead of requiring a patient to follow the pre-specified monitoring intervention \bar{n}_{y_j} at every time point, the static monitoring intervention $g_{y_j}^*$ requires patients to be monitored at least at the pre-specified set of time points t when $n_{y_j}(t) = 1$, but does not constrain monitoring decisions at other time points, i.e., it allows for the monitoring process to take its natural course for t when $n_{y_j}(t) = 0$. The modified dynamic treatment-censoring intervention $d_{x_j}^*$ differs from the previously defined decision rule d_{x_j} in that it only uses the past A1c measurements which are observed both under the monitoring regime \bar{n}_{y_j} , and under the actual observed monitoring process \bar{N} , i.e., $d_{x_j}^*(t)\{V(t)\} = d_{x_j}(t)\{V^*(t)\}$ where we define $V^*(t) = \{A(t-1), n_{y_j}(t-1)N(t-1), Y(t), n_{y_j}(t-1)I(t)\}$. Thus, $d_{x_j}^*$ corresponds to the original dynamic treatment-censoring intervention d_{x_j} except that it

requires that any additional A1c measurements collected beyond the A1c measurements required under intervention \bar{n}_{y_j} be ignored when applying the decision rule d_{x_j} . Figure 6 (top panel) illustrates the dynamic treatment regimes resulting from the application of such decision rules $d_{x_j}^*$ using fake data from two patients (see 'Step 3B').

While causal estimands (4) can be of interest in their own rights, we focus on their estimation as an indirect approach to the evaluation of causal estimands (1). Because the monitoring regime $g_{y_j}^*$ involves interventions on fewer variables than \bar{n}_{y_j} , it is generally expected that more patients will have an observed monitoring history consistent with the intervention $g_{y_j}^*$. We demonstrate in the Results Section how this potential increase in data support can be exploited to improve inferences in studies with random monitoring to evaluate some monitoring regimes \bar{n}_{y_j} and, in particular, those less frequent in the observed data.

5.2 Identifying assumptions

An SRA and PA are sufficient for identifying causal estimands (4). If causal estimands (1) are indirectly evaluated through causal estimands (4), then an additional identifiability assumption referred to as the NDE assumption is required. It is formalized by the equality of counterfactual covariates under two sets of static interventions (Robins et al., 2008): $L_{a,\bar{n}}^0(t) = L_a^0(t)$ where $L^0(t) = \{Y(t), Z(t), I^0(t)\}$. Using an NPSEM framework, Neugebauer et al. (2017) showed that this assumption is implied by an exclusion restriction assumption that requires that all directed paths from nodes $N(t)$ to subsequent covariates $Y(j)$, $Z(j)$, and $I^0(j)$ be intercepted by treatment or censoring nodes $A(t)$ for $j > t$. With the DAG in Figure 1, this assumption is encoded by the exclusion of the three solid gray arrows from $N(0)$ to $Y(1)$, $I^0(1)$, and $Y(2)$.

The SRA for identification of the causal estimand (4) can be expressed as:

$$\begin{cases} Y_{d_{x_j}^*, g_{y_j}^*}^*(t_0 + 1) \perp \{A(t), N(t)\} | \bar{L}(t), \bar{A}(t-1), \bar{N}(t-1) \text{ for all } t \text{ s.t. } n_{y_j}(t) = 1 \text{ (under } g_{y_j}^*), \\ Y_{d_{x_j}^*, g_{y_j}^*}^*(t_0 + 1) \perp A(t) | \bar{L}(t), \bar{A}(t-1), \bar{N}(t-1) \text{ for all } t \text{ s.t. } n_{y_j}(t) = 0 \text{ (under } g_{y_j}^*), \end{cases} \quad (5)$$

with $t = 0, \dots, t_0$. This assumption is weaker than the SRA (2) as it holds even when some backdoor paths from $N(t)$ to $Y(t_0 + 1)$ are not blocked as long as such open backdoor paths only occur at t when $n_{y_j}(t) = 0$, i.e., for monitoring nodes $N(t)$ that are not intervened upon under intervention $g_{y_j}^*$. This difference between the SRA formulations (2) and (5) is unlikely to result in the selection of different covariates for bias adjustment in practice, as most arguments in favor of or against the upholding of the SRA (2) will typically also apply to the upholding of the SRA (5) and vice versa.

The PA for identification of the causal estimand (4) can be stated as:

$$\left\{ \begin{array}{l} P\left[A(t) = d_{x_j}^*\{V(t)\} \mid \bar{L}(t), \bar{Y}(t) = 0, \bar{A}(t-1) = d_{x_j}^*\{\bar{V}(t-1)\}, \bar{N}(t-1) = g_{y_j}^*\{\bar{N}(t-1)\}\right] > 0 \\ \text{for } t = 0, \dots, t_0 \text{ and,} \\ P\left[N(t) = 1 \mid \bar{L}(t), \bar{Y}(t) = 0, \bar{N}(t-1) = g_{y_j}^*\{\bar{N}(t-1)\}, \bar{A}(t) = d_{x_j}^*\{\bar{V}(t)\}\right] > 0, \\ \text{for } t = 0, \dots, t_0 \text{ such that } n_{y_j}(t) = 1 \text{ (under } g_{y_j}^*), \end{array} \right. \quad (6)$$

where $g_{y_j}^*\{\bar{N}(t)\} = [g_{y_j}^*(0)\{N(0)\}, \dots, g_{y_j}^*(t)\{N(t)\}]$ is defined by the mappings $g_{y_j}^*(k) : N(k) \mapsto N(k)^{1-n_{y_j}(k)}$ as the sequence of monitoring decisions through time t that is consistent with following intervention $g_{y_j}^*$ and compatible with the subset of observed monitoring decisions that are not constrained by intervention $g_{y_j}^*$. We note that the PA (6) requires that the A1c of a patient who previously followed the intervention $(d_{x_j}^*, g_{y_j}^*)$ be possibly monitored at only the time points $t \leq t_0$ when $n_{y_j}(t) = 1$, whichever the patient's covariate and monitoring history. This is in contrast to the PA (3) that requires that a patient who previously followed the intervention (d_{x_j}, \bar{n}_{y_j}) have its A1c monitored at all time points $t \leq t_0$ according to the static intervention $n_{y_j}(t)$, whichever the patient's covariate history. It can thus be expected that the PA (6) will often be more likely to hold in practice than its analog (3). For example for $y_j = 1$, the PA (6) only requires that each patient can possibly be monitored every other quarter and does not place any constraint on the monitoring decisions between these monitoring events, while the PA (3) requires that each patient can possibly experience the exact sequence of monitoring decisions \bar{n}_1 throughout follow-up.

The practical consequences of this weakening of the PA are expected improvements in IPW estimation performance, both in terms of finite sample bias and precision because of the likely more stable IP weights (e.g., fewer extreme values or more compact distribution). However, this potential for improved performance in the estimation of the causal estimand (4) over that of (1) might be offset by an increased risk of practical violation of the PA for the exposure process in (6) compared to that in (3). For example in the TI study, Neugebauer et al. (2012) expected that most patients would not remain unexposed to an intensified therapy if their A1c reached very high levels. Violation of the PA for the exposure process in (6) is then a reasonable concern because the treatment decision at time t according to rule $d_{x_j}^*$ would require that a patient remain unexposed to an intensified therapy - even if a newly measured A1c reached a very high level - if such an A1c was collected at a time when the static monitoring regime $\bar{n}_{y_j}(t)$ would not result in A1c testing (i.e., $n_{y_j}(t-1) = 0$). Because positivity violations when evaluating static treatment regimes often motivates the evaluation of dynamic interventions instead, we expect concerns over violations of the PA for the exposure process in (6) - such as the one just described for the TI study - to apply broadly to other studies. In the Results and Discussion Sections, we examine the trade-off between improved IP weight stability resulting from interventions on fewer monitoring nodes and worsened IP weight stability resulting from poorer adherence to decision rules by comparing the practical performance of the following two IPW estimators for the causal estimand (1) under the NDE assumption.

6. Standard and NDE-based IPW estimators

Both IPW estimators considered for evaluating the risk difference (1) were constructed and implemented using the same general principles summarized below. Details are provided in Sections 1 and 2 of the Supporting Information. Both estimators only differ in the joint interventions that they each evaluate. Whereas the standard IPW estimator directly targets

the interventions of interest (d_{x_j}, \bar{n}_{y_j}) , the NDE-based IPW estimator targets the modified interventions $(d_{x_j}^*, \bar{g}_{y_j}^*)$. In both cases, we first implement an IPW estimator of the discrete-time counterfactual hazards of failure at each time $t + 1$ for $t = 0, \dots, t_0$ for one of the two joint interventions that define the targeted risk difference. Each hazard estimate $\alpha(t)$ is defined by a convex, linear combination of the outcomes at time $t + 1$ from all patients who did not experience the event before or at time t , and who followed the intervention of interest through time t :

$$\alpha(t) = \sum_{i=1}^n \frac{h_i(t)}{\sum_{i=1}^n h_i(t)} Y_i(t+1), \quad (7)$$

where $h_i(t)$ is a stabilized inverse probability weight based on the joint conditional probability that patient i experiences both the exposure and monitoring interventions through time t . To construct the IPW weights $h_i(t)$, we first estimate separate propensity scores for treatment, censoring and monitoring using distinct logistic models. These steps are then replicated for the second joint intervention that defines the targeted risk difference. As detailed elsewhere (Neugebauer et al., 2016), instead of deriving each hazard estimates arithmetically using formula (7), they can all be derived simultaneously by fitting a single *saturated* logistic dynamic marginal structural model (MSM) for the counterfactual hazards (van der Laan and Petersen, 2007; Robins et al., 2008) using a standard weighted logistic regression using an expanded data set (Cain et al., 2010) of person-time observations, replicated for each regime a person follows at each time point, with IP weights $h_i(t)$. IP weights were truncated at 40. The resulting $t_0 + 1$ hazard estimates under each of the targeted interventions are then mapped into estimates of the counterfactual cumulative risk using $\prod_{t=0}^{t_0} (1 - \alpha(t))$. Finally, the difference between any two of these cumulative risk estimates is computed and inference for this risk difference estimate is derived based on the delta method.

We emphasize two major differences in the implementation of the standard versus NDE-based IPW estimators. Whereas only outcomes from patients that have a history of exposure and monitoring concordant with the joint intervention (d_{x_j}, \bar{n}_{y_j}) contribute to the standard

IPW estimator of the hazard $P\{Y_{d_{x_j}, \bar{n}_{y_j}}(t+1) = 1 \mid Y_{d_{x_j}, \bar{n}_{y_j}}(t) = 0\}$, the NDE-based estimator also uses data from patients monitored at least as often as what is required under the static monitoring intervention \bar{n}_y . We note that this potential gain in data support to fit the dynamic MSM in the NDE-based analysis can however be offset by the decrease in the number of patients following rule $d_{x_j}^*$ compared to d_{x_j} . This trade-off is illustrated in Figure 6 (top panel). The stylized patient with id=1 shows a situation where person-time observations can be gained with the NDE assumption (by following $g_{y_j}^*$ when \bar{n}_{y_j} is not followed). In contrast, the patient with id=2 displays a scenario where a person-time observation gained with $g_{y_j}^*$ (at time 2) is lost by not following the modified dynamic treatment intervention $d_{x_j}^*$ (at time 3) because treatment was initiated too early. In addition and as illustrated in the same Figure, the denominator of the IP weights for the NDE-based estimator only includes conditional probabilities of monitoring for time points when monitoring is required under \bar{n}_y whereas the standard estimator includes conditional probabilities of monitoring for all time points.

7. Results

7.1 Data support to evaluate each joint intervention

The two histograms at the bottom of Figure 1 contrast the counts of patients following the two types of static monitoring interventions \bar{n}_y versus g_y^* over the first 8 quarters of follow-up for both $y = 1, 3$. These plots clearly indicate that data support for each intervention g_y^* is significantly larger than for the corresponding \bar{n}_y intervention. This is expected because the intervention g_y^* only requires A1c testing when the static monitoring intervention \bar{n}_y requires it and does not constrain the monitoring process otherwise. As expected, we also find that data support for the joint interventions (d_x, \bar{n}_y) (Figure 2, top) are lower than their intervention analogs without monitoring interventions d_x (Figure 5, top left).

The counts of patients following any of the modified joint interventions (d_x^*, g_y^*) (Figure 3, top) is larger than their analogs (d_x, \bar{n}_y) (Figure 2, top). The discrepancy is particularly

striking with $y = 3$. The histograms in Figure 3 also show that, starting at time 2 and for each $x \in \mathcal{X}$, the number of patients following the joint intervention (d_x^*, g_y^*) defined by \bar{n}_y increases with the decrease in the minimum required frequency of A1c monitoring (i.e., from \bar{n}_1 to \bar{n}_3). The opposite trend is seen for interventions (d_x, \bar{n}_y) in Figure 2. These results address the concern discussed in the prior Section by demonstrating that the expected increase in data support for the joint intervention (d_x^*, g_y^*) compared to (d_x, \bar{n}_y) due to intervening on fewer monitoring variables with g_y^* is not offset - in this study at least - by a decrease in the number of patients who follow the intervention d_x^* instead of d_x .

7.2 Estimates of the causal risk differences

We now contrast the standard and NDE-based estimates of the effects of the 6 joint interventions (d_x, \bar{n}_y) by, first, comparing the three dynamic treatment-censoring interventions when they are all combined with the same monitoring intervention (Figures 2 and 3), and, second, by comparing the two monitoring interventions when they are combined with the same dynamic treatment-censoring intervention (Figure 4). For comparison, we also present the crude and IPW estimates of the same dynamic regimes without monitoring interventions (Figure 5). We note that unadjusted estimates of the cumulative risks provide no evidence of a protective effect of increasingly more aggressive TI strategies while adjusted estimates provide strong evidence that the risk of albuminuria development or progression almost always decrease significantly with the decrease of the A1c threshold at which TI is initiated. This finding is consistent with results from two clinical trials referenced in the Supporting Information.

With the standard analysis (Figure 2), the estimated counterfactual survival curves provide weaker (when $y = 1$) or no evidence (when $y = 3$) of a protective effect of TI initiation at lower A1c thresholds. The standard errors of the IPW risk difference estimates are much larger than those reported in Figure 5 due to reduced data support for joint interventions

compared to their analogs without monitoring interventions. With the NDE analysis (Figure 3), we note the drastic change in the estimated counterfactual survival curves with $y = 3$: the protective effect of TI initiation at lower A1c thresholds which was lost in the standard analysis is recovered. The bottom panels on the same Figure also indicate a large decrease in the standard errors for the risk differences compared to those showed at the bottom of Figure 2. This is explained by the increase in data support for the joint interventions (d_x^*, g_y^*) compared to (d_x, \bar{n}_y) . Overall, the NDE-based estimator provides stronger and more consistent evidence of a protective effect of TI initiation at lower A1c levels.

When contrasting joint interventions defined by different monitoring interventions but the same dynamic treatment-censoring intervention (Figure 4), results from the standard analysis suggest a beneficial effect of more frequent monitoring on the risk of onset or progression of albuminuria. However, the wide confidence intervals for the risk difference fail to provide strong statistical evidence for this effect. Despite the tighter confidence intervals for the risk difference in the NDE-based analysis (bottom right panels) compared to those in the standard analysis (bottom left panels), the general decrease in the point estimates leads to similarly weak statistical evidence of a protective effect of more frequent A1c monitoring.

7.3 Distribution of estimated stabilized IP weights

Figure 6 (bottom panel) summarizes the distribution of the weights $h_i(t)$ in (7) for intervention-person-time observations with non-zero weight values, corresponding with the two sets of joint interventions discussed above (d_x, \bar{n}_y) and (d_x^*, g_y^*) . When comparing these distributions to their analog without monitoring interventions (Figure 5 top right panel), we note that both distributions have slightly shifted right which is expected due to the additional interventions on monitoring. We note a relatively large increase in the proportions of large weights greater than 40 with regimes (d_x, \bar{n}_y) which provides evidence for the theoretical concerns over increased near-violations of the PA (3) discussed in Section 4.2. We conjecture that the

observed near-violations are likely resulting from the large reduction in the number of patients following each joint intervention (i.e., it could be avoided with increased sample sizes) but it may also be indicative of covariates that are strong determinants of monitoring decisions (i.e., structural violations that would persist with increased sample sizes). When comparing the distributions of the IPW weights in Figure 6 between joint interventions types (d_x, \bar{n}_y) and (d_x^*, g_y^*) , we note a relatively large increase in the proportions of large and very large weights greater than 30 with interventions (d_x^*, g_y^*) . This increase in the number of large weights provides evidence of the theoretical concern discussed in Section 5.2 over increased near-violations of the PA requirements for the exposure process in (6) compared to (3).

8. Discussion

With this case study, we aimed to provide detailed practical guidance on how to exploit monitoring variability to evaluate its health impact or to improve the generalizability of CER findings when evaluating dynamic treatment regimes. Although the approaches developed were illustrated with EHR data, they are applicable to any observational study with non-systematic covariate monitoring. We described and compared the implementation of two IPW estimators for evaluating the joint effects of dynamic treatment-censoring and static monitoring interventions. To our knowledge, this report provides the first detailed account of the practical trade-offs between these two IPW estimation approaches. Motivations for their applications also include the joint optimization of treatment and monitoring decisions. With our example, we illustrated the expected poor performance of a standard IPW estimator due to a large decrease in data support to evaluate a given static monitoring regimen which, in turn, can also increase concerns over finite-sample bias from near-violations of the PA for the monitoring process. To alleviate the expected practical limitation of the standard IPW estimator, we demonstrated how an alternative approach that hinges on an NDE assumption can result in much improved estimation efficiency due to increased data support but at the

cost of a potential increase in finite-sample bias due to structural near-violations of the PA for the treatment process.

Compared to an estimator that evaluates the effect of a dynamic treatment without a joint monitoring intervention, the performances of both the standard and NDE-based estimators discussed here hinge on a stronger SRA whose upholding has been questioned in EHR-based studies because “data on the reason for a given visit are often not recorded for data analysis” (Robins et al., 2008). Because we use observational data assembled originally to evaluate dynamic treatment interventions without joint monitoring interventions, it is reasonable to expect residual bias from unmeasured confounding of the effect of A1c monitoring on the outcome. Had the original study aims included the evaluation of the effects of monitoring regimes, additional covariates would have likely been extracted from the EHR to adjust for clinical determinants of A1c lab orders and measures of patients’ prior compliance with such orders. In Web Appendix C of the Supporting Information, we present potential evidence of the violation of the NDE assumption in the TI study. Even if the NDE assumption were violated, we note that inferences from the alternate NDE-based IPW analyses in this paper can remain causally interpretable (using (4) instead of (1)) although results become more difficult to convey and use to inform care. As noted in Robins et al. (2008, p. 4703), we might expect a violation of the NDE assumption in intention-to-treat (ITT) analyses where the dynamic treatment interventions stop but the monitoring interventions continue past TI initiation. Although we implemented a per-protocol analysis, a similar concern arises due to our definition of TI as exposure to any antidiabetic drug not used at study entry. Indeed, monitoring after TI initiation could result in patient switching types of intensified therapy, and if the type of therapy affects patient outcomes, the NDE assumption can be violated. To alleviate this concern, the intervention protocol after TI initiation could be redefined to require continuous exposure to the same drug combination as the one used by

the patient at TI initiation. Such an analysis was not feasible due to the lack of information on drug used after TI initiation in the available data set. In future work, we will describe alternate approaches to mitigate concerns over NDE violations by implementing a locally efficient double robust estimator (TMLE, Neugebauer et al. (2014)) and evaluating stochastic monitoring interventions instead of static ones. Stochastic interventions can define effects that are particularly relevant for patient-centered outcomes research because they can better represent real-world adherence to rigid monitoring schedules such as the ones studied in this report. The weakening of the PA required to identify these effects is expected (Munoz and van der Laan, 2012) to improve estimation performance with both IPW and TMLE.

The nonparametric estimation approach for the discrete-time counterfactual hazards adopted here may not always be practical in many applications if too few subjects follow each of the joint interventions of interest at each time point (curse of dimensionality). Instead, a nonparameteric MSM approach based on a working, non-saturated model (Neugebauer and van der Laan, 2007) can be adopted to explicitly recognize the limitation of an arbitrarily specified non-saturated MSM in capturing the true counterfactual hazard functions while also addressing concerns over poor estimation precision with the nonparametric approach described in this report. Finally, whereas the missing indicator and LOVCF approach we adopted to handle partially observed covariates is warranted for biomarkers whose impact on treatment decisions can be assumed to be entirely mediated by observed covariates, alternate approaches such as multiple imputation might be better suited to handle other sources of missing data such as information known to clinicians but poorly captured by EHR.

9. Acknowledgements

This study was partially supported through a Patient-Centered Outcomes Research Institute (PCORI) Award (ME-1403-12506). All statements in this report, including its findings and conclusions, are solely those of the authors and do not necessarily represent the views of

the Patient-Centered Outcomes Research Institute (PCORI), its Board of Governors or Methodology Committee. The authors thank the following investigators from the HMO research network for making data from their sites available to this study: Denise M. Boudreau, Connie Trinacty, Gregory A. Nichols, Marsha A. Raebel, Kristi Reynolds, and Patrick J. O'Connor. NK was partially funded by the Medical Research Council (MR/L012332/1).

Data Availability Statement

The data that were used in this paper were directly obtained from the electronic health records of participants. Because of privacy issues, the Institutional Review Board cannot allow to make these data publicly available. Metadata can be shared. Parties interested in scientific collaboration using the data may contact Romain Neugebauer.

References

- American Diabetes Association (2018). 6. Glycemic Targets: Standards of Medical Care in Diabetes 2018. *Diabetes Care* **41**, S55–S64.
- Bodenheimer, T., Wagner, E. H., and Grumbach, K. (2002). Improving primary care for patients with chronic illness: the chronic care model, part 2. *JAMA* **288**, 1909–1914.
- Cain, L. E., Robins, J. M., Lanoy, E., Logan, R., Costagliola, D., and Hernán, M. A. (2010). When to start treatment? a systematic approach to the comparison of dynamic regimes using observational data. *The International Journal of Biostatistics* **6**,.
- Caniglia, E. C., Sabin, C., Robins, J. M., Logan, R., Cain, L. E., Abgrall, et al. (2016). When to monitor cd4 cell count and hiv rna to reduce mortality and aids-defining illness in virologically suppressed hiv-positive persons on antiretroviral therapy in high-income countries: a prospective observational study. *Journal of Acquired Immune Deficiency Syndromes* **72**, 214.
- Ford, D., Robins, J. M., Petersen, M. L., Gibb, D. M., Gilks, C. F., Mugenyi, et al. (2015).

The impact of different cd4 cell-count monitoring and switching strategies on mortality in hiv-infected african adults on antiretroviral therapy: an application of dynamic marginal structural models. *American Journal of Epidemiology* **182**, 633–643.

Hernán, M. A., McAdams, M., McGrath, N., Lanoy, E., and Costagliola, D. (2009). Observation plans in longitudinal studies with time-varying treatments. *Statistical Methods in Medical Research* **18**, 27–52.

Kreif, N., Sofrygin, O., Schmittdiel, J., Adams, A., van der Laan, M., Neugebauer, R., et al. (2018). Evaluation of adaptive treatment strategies in an observational study where time-varying covariates are not monitored systematically. *arXiv:1806.11153*.

Munoz, I. D. and van der Laan, M. (2012). Population intervention causal effects based on stochastic interventions. *Biometrics* **68**, 541–549.

Murphy, S. A., Lynch, K. G., Oslin, D., McKay, J. R., and TenHave, T. (2007). Developing adaptive treatment strategies in substance abuse research. *Drug and Alcohol Dependence* **88**, S24–S30.

Neugebauer, R., Fireman, B., Roy, J. A., O’connor, P. J., and Selby, J. V. (2012). Dynamic marginal structural modeling to evaluate the comparative effectiveness of more or less aggressive treatment intensification strategies in adults with type 2 diabetes. *Pharmacoepidemiology and Drug Safety* **21**, 99–113.

Neugebauer, R., Fireman, B., Roy, J. A., and OConnor, P. J. (2013). Impact of specific glucose-control strategies on microvascular and macrovascular outcomes in 58,000 adults with type 2 diabetes. *Diabetes Care* **36**, 3510–3516.

Neugebauer, R., Schmittdiel, J. A., Adams, A. S., Grant, R. W., and van der Laan, M. J. (2017). Identification of the joint effect of a dynamic treatment intervention and a stochastic monitoring intervention under the no direct effect assumption. *Journal of Causal Inference*.

- Neugebauer, R., Schmittdiel, J. A., and Laan, M. J. (2014). Targeted learning in real-world comparative effectiveness research with time-varying interventions. *Statistics in Medicine* **33**, 2480–2520.
- Neugebauer, R., Schmittdiel, J. A., and van der Laan, M. J. (2016). A case study of the impact of data-adaptive versus model-based estimation of the propensity scores on causal inferences from three inverse probability weighting estimators. *The International Journal of Biostatistics* **12**, 131–155.
- Neugebauer, R., Schmittdiel, J. A., Zhu, Z., Rassen, J. A., Seeger, J. D., and Schneeweiss, S. (2015). High-dimensional propensity score algorithm in comparative effectiveness research with time-varying interventions. *Statistics in Medicine* **34**, 753–781.
- Neugebauer, R. and van der Laan, M. (2007). Nonparametric causal effects based on marginal structural models. *Journal of Statistical Planning and Inference* **137**, 419–434.
- Pearl, J. (2010). On the consistency rule in causal inference: axiom, definition, assumption, or theorem? *Epidemiology* **21**, 872–875.
- Petersen, M. L., Porter, K. E., Gruber, S., Wang, Y., and van der Laan, M. J. (2012). Diagnosing and responding to violations in the positivity assumption. *Statistical Methods in Medical Research* **21**, 31–54.
- Robins, J. (1986). A new approach to causal inference in mortality studies with a sustained exposure period application to control of the healthy worker survivor effect. *Mathematical Modelling* **7**, 1393–1512.
- Robins, J., Orellana, L., and Rotnitzky, A. (2008). Estimation and extrapolation of optimal treatment and testing strategies. *Statistics in Medicine* **27**, 4678–4721.
- Robins, J. M., Hernan, M. A., and Brumback, B. (2000). Marginal structural models and causal inference in epidemiology. *Epidemiology* **11**, 550–560.
- Sofrygin, O., Zhu, Z., Schmittdiel, J. A., Adams, A. S., van der Laan, M. J., Neugebauer,

R., et al. (2019). Targeted learning with daily EHR data. *Statistics in Medicine* **38**, 3073–3090.

van der Laan, M. and Petersen, M. L. (2007). Causal Effect Models for Realistic Individualized Treatment and Intention to Treat Rules. *International Journal of Biostatistics* **3**, Article 3.

SUPPORTING INFORMATION

Web Appendices, Tables, and Figures referenced in Sections 1, 4, 6 and 8 are available with this paper at the Biometrics website on Wiley Online Library. Template code and (simulated) data are provided for demonstrating how the analyses presented here can be emulated in other studies.

10. TABLES AND FIGURES

[Figure 1 about here.]

[Figure 2 about here.]

Accepted Article

[Figure 3 about here.]

Accepted Article

[Figure 4 about here.]

Accepted Article

[Figure 5 about here.]

Accepted Article

[Figure 6 about here.]

Received xx 2018. Revised xx 2018. Accepted xx 2018.

Accepted Article




Figure 1. The top diagram is a directed acyclic graph inspired by fig 1b in Robins et al. (2000) and adapted to cohort studies with a time-to-event outcome and non-systematic covariate monitoring. The histogram in the middle panel summarizes the distribution of the length of time between A1c measurements for all patients in the cohort. The histogram in the bottom left panel represents the counts of patients following, over time, one of the two static monitoring interventions \bar{n}_1 and \bar{n}_3 . Each static intervention \bar{n}_y requires that consecutive A1c tests always be separated by y quarter(s). The histogram in the bottom right represents the counts of patients following, over time, one of the two static monitoring interventions g_1^* and g_3^* . Each intervention g_y^* only requires A1c testing when the static monitoring intervention \bar{n}_y requires it (there is no A1c testing restrictions otherwise).




Figure 2. Contrast of treatment regimes under a fixed monitoring regime (without NDE). Each histogram in the top panel represents the counts of patients following, over time, one of the three dynamic treatment-censoring interventions $d_{7.5}$, $d_{8.0}$, or $d_{8.5}$ and each of two static monitoring interventions \bar{n}_1 and \bar{n}_3 . Each static intervention \bar{n}_y requires that consecutive A1c tests always be separated by y quarter(s). Each dynamic interventions d_x requires that therapy be intensified at the first time a new A1c is observed $\geq x\%$. Each plot in the middle panel represents the IPW estimates of the counterfactual survival curves under three joint interventions corresponding with the three dynamic censoring-treatment interventions and one of the two static monitoring interventions. Each plot in the bottom panel represents the IPW estimates of the counterfactual cumulative risk differences over time between two joint interventions that share the same monitoring intervention.




Figure 3. Contrast of treatment regimes under a fixed monitoring regime (with NDE). Each histogram in the top panel represents the counts of patients following, over time, one of the three dynamic treatment-censoring interventions $d_{7.5}^*$, $d_{8.0}^*$, or $d_{8.5}^*$ and each of two static monitoring interventions g_1^* and g_3^* . Each intervention g_y^* only requires A1c testing when the static monitoring intervention \bar{n}_y requires it. The intervention \bar{n}_y requires that consecutive A1c tests always be separated by y quarter(s). The dynamic interventions d_x^* requires that therapy be intensified at the first time a new A1c collected when required by \bar{n}_y is observed $\geq x\%$. Each plot in the middle panel represents the IPW estimates of the counterfactual survival curves under three joint interventions corresponding with the three dynamic censoring-treatment interventions and one of the two static monitoring interventions. Each plot in the bottom panel represents the IPW estimates of the counterfactual cumulative risk differences over time between two joint interventions that share the same monitoring intervention.




Figure 4. Contrast of monitoring regimes under a fixed treatment regime (with and without NDE). Each plot in the top panel represents the IPW estimates of the counterfactual survival curves under two joint interventions corresponding with one of the dynamic censoring-treatment interventions $d_{7.5}$, $d_{8.0}$, or $d_{8.5}$ and two static monitoring interventions \bar{n}_1 and \bar{n}_3 . Each plot in the middle panel represents the IPW estimates of the counterfactual survival curves under two joint interventions corresponding with one of the dynamic censoring-treatment interventions $d_{7.5}^*$, $d_{8.0}^*$, or $d_{8.5}^*$ and two static monitoring interventions g_1^* and g_3^* . Each plot in the bottom panel represents the IPW estimates of the counterfactual cumulative risk differences over time between two joint interventions that share the same treatment-censoring interventions.




Figure 5. Contrast of treatment regimes without monitoring interventions (original analyses). The histogram in the top panel represents the counts of patients following, over time, one of the three dynamic treatment-censoring interventions $d_{7.5}$, $d_{8.0}$, or $d_{8.5}$. The dynamic interventions d_x requires that therapy be intensified at the first time a new A1c is observed $\geq x\%$. Each plot in the middle panel represents the crude and IPW estimates of the counterfactual survival curves under the three interventions. Each plot in the bottom panel represents the IPW estimates of the corresponding counterfactual cumulative risk differences over time. The table in the top panel summarizes the distribution of stabilized inverse probability weights for intervention-person-time observations with a weight value $\neq 0$.

Step 1a. Formulate joint treatment-censoring-monitoring intervention of interest, e.g. (d_8, \bar{n}_1) : "initiate treatment the first time A1c ($I(t)$) drifts above 8% and monitor A1c ($N(t)$) every other follow-up period"															
Step 1b. Define causal estimand, e.g., the counterfactual cumulative risk: $P(Y_{d_8, \bar{n}_1}(t_0 + 1) = 1)$															
Step 2. Structure data into long format respecting temporal ordering and using LVCF to define time-varying covariates					A: Non-NDE Analysis					B: NDE analysis Modified joint treatment-censoring-monitoring intervention, e.g., (d_8^*, g_1^*)					
					Step 3A. Compare observed data to joint intervention			Step 4A. Calculate $h(t)$ with PS		Step 3B. Compare observed data to joint intervention				Step 4B. Calculate $h(t)$ with PS	
id	t	$N(t)$	$I(t)$	$A_1(t)$	$n_1(t)$	$d_8(t)$		$h(t)$	PS	$g_1^*(t)$	$d_8^*(t)$		$h(t)$	PS	
1	0	0	7.9	0	0	$I(0) < 8 \rightarrow a_1(0) = 0$		$\neq 0$	A, N	No intervention	$I(0) < 8 \rightarrow a_1(0) = 0$		$\neq 0$	A	
1	1	1	7.9	0	1	$N(0) = 0 \rightarrow a_1(1) = A_1(0) = 0$		$\neq 0$	A, N	1	$n_1(0) * N(0) = 0 \rightarrow a_1(1) = A_1(0) = 0$		$\neq 0$	A, N	
1	2	1	8.1	1	0	$I(2) > 8 \rightarrow a_1(2) = 1$		0		No intervention	$n_1(1) * I(2) > 8 \rightarrow a_1(2) = 1$		$\neq 0$	A	
1	3	1	8.2	1	1	$A_1(2) = 1 \rightarrow a_1(3) = 1$		0		1	$A_1(2) = 1 \rightarrow a_1(3) = 1$		$\neq 0$	A, N	
2	0	0	7.5	0	0	$I(0) < 8 \rightarrow a_1(0) = 0$		$\neq 0$	A, N	No intervention	$I(0) < 8 \rightarrow a_1(0) = 0$		$\neq 0$	A	
2	1	1	7.5	0	1	$N(0) = 0 \rightarrow a_1(1) = A_1(0) = 0$		$\neq 0$	A, N	1	$n_1(0) * N(0) = 0 \rightarrow a_1(1) = A_1(0) = 0$		$\neq 0$	A, N	
2	2	1	7.9	0	0	$I(2) < 8 \rightarrow a_1(2) = A_1(1) = 0$		0		No intervention	$n_1(1) * I(2) < 8 \rightarrow a_1(2) = A_1(1) = 0$		$\neq 0$	A	
2	3	1	8.1	1	1	$I(3) > 8 \rightarrow a_1(3) = 1$		0		1	$n_1(2) * N(2) = 0 \rightarrow a_1(3) = A_1(2) = 0$		0		
Step 5. Logistic regression on person-time observations, replicated for each joint intervention a patient follows, with $h(t)$ weights															
Step 6. Estimate hazards from the coefficients, transform into estimated counterfactual risks															

IPW	Frequency d_x, \bar{n}_y	%	Cumulative Frequency	Cumulative %	Frequency d_x^*, g_y^*	%	Cumulative Frequency	Cumulative %
<0	0	0.00	0	0.00	0	0.00	0	0.00
[0, 0.5]	11832	3.25	11832	3.25	65094	8.44	65094	8.44
[0.5, 1]	224913	61.81	236745	65.06	601839	78.02	666933	86.46
[1, 10]	125177	34.40	361922	99.47	100577	13.04	767510	99.50
[10, 20]	1714	0.47	363636	99.94	3282	0.43	770792	99.93
[20, 30]	187	0.05	363823	99.99	407	0.05	771199	99.98
[30, 40]	34	0.01	363857	100.00	97	0.01	771296	99.99
[40, 50]	6	0.00	363863	100.00	28	0.00	771324	99.99
[50, 100]	3	0.00	363866	100.00	30	0.00	771354	100.00
[100, 150]	0	0.00	363866	100.00	5	0.00	771359	100.00
≥ 150	0	0.00	363866	100.00	7	0.00	771366	100.00

Figure 6. The figure in the top panel demonstrates the practical steps of the non-NDE and NDE analyses, using illustrative patient data. The table in the bottom panel summarizes the distribution of stabilized inverse probability weights for intervention-person-time observations with a weight value $\neq 0$ corresponding with two sets of joint interventions: 1) the 6 joint interventions corresponding with the three dynamic censoring-treatment interventions $d_{7.5}$, $d_{8.0}$, and $d_{8.5}$, and the two static monitoring interventions \bar{n}_1 and \bar{n}_3 ; 2) the 6 joint interventions corresponding with the three dynamic censoring-treatment interventions $d_{7.5}^*$, $d_{8.0}^*$, and $d_{8.5}^*$, and the two static monitoring interventions g_1^* and g_3^* .