Article:

Zhang, Yingying, Bennett, Alastair, Manca, Andrea orcid.org/0000-0001-8342-8421 et al. (9 more authors) (Accepted: 2025) Estimating the Causal Effect of Realistic Treatment Strategies Using Longitudinal Observational Data. Medical Decision Making. ISSN: 1552-681X (In Press)

Original article

Journal Title
XX(X):1–2

DOI: 10.1177/ToBeAssigned
www.sagepub.com/
SAGE

Estimating the Causal Effect of Realistic Treatment Strategies Using Longitudinal Observational Data

Yingying Zhang[1], Alastair Bennett[1], Andrea Manca[1*], Moshe Mittelman[2,3], Marlijn Hoeks[4], Alex Smith[5], Adele Taylor[5], Reinhard Stauder[6], Theo de Witte[7], Luca Malcovati[8,9], Corine van Marrewijk[4], Noemi Kreif[10]

[1]Centre for Health Economics, University of York, York, United Kingdom
[2]Faculty of Medicine, Tel Aviv University, Tel Aviv-Yafo, Israel
[3]Department of Hematology, Tel Aviv Sourasky Medical Center, Tel Aviv-Yafo, Israel
[4]Department of Hematology, Radboud university medical center, Nijmegen, The Netherlands
[5]Department of Health Sciences, University of York, York, United Kingdom
[6]Department of Public Health, Health Services Research and Health Technology Assessment, UMIT Tirol – The Tyrolean Private University, Austria
[7]Department of Tumor Immunology - Nijmegen Center for Molecular Life Sciences, Radboud university medical center, Nijmegen, The Netherlands
[8]Department of Hematology Oncology, Fondazione IRCCS Policlinico San Matteo, Pavia, Italy
[9]Department of Molecular Medicine, University of Pavia, Pavia, Italy
[10]Department of Pharmacy, University of Washington, Seattle, United States

Corresponding author:
Prof. Andrea Manca, Centre for Health Economics, University of York, Heslington, York, YO10 5DD, United Kingdom.
Email: andrea.manca@york.ac.uk

Acknowledgments

We thank all local investigators, operational team members and patients for their contribution to the collection of data of the European Myelodysplastic Syndromes (EUMDS) Registry as well as their scientific committee and investigators for participating in and contributing to the discussion of realistic dynamic treatment regimes for giving ESA. An earlier version of this manuscript was presented at the The 2024 winter conference for the UK Health Economists Study Group (HESG). We are grateful to those participated in the discussion of the manuscript. All errors and omissions are the authors'.

Competing Interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Financial support for this study was provided entirely by The HTx project. The funding agreement ensured the authors' independence in designing the study, interpreting the data, writing, and publishing the report. The HTx project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement Nº 825162 for 5 years from January 2019. This dissemination reflects only the authors' view and the Commission is not responsible for any use that may be made of the information it contains.

The main aim of HTx is to create a framework for the Next Generation Health Technology Assessment (HTA) to support patient-centered, societally oriented, real-time decision-making on access to and reimbursement for health technologies throughout Europe.

This is also continuing work following the collaboration with the European Unions's Horizon 2020 MDS-RIGHT project (No: 634789), which had initiative for the DTR analyses.

Original article

Journal Title
XX(X):1–12
DOI: 10.1177/ToBeAssigned
www.sagepub.com/
SAGE

Estimating the Causal Effect of Realistic Treatment Strategies Using Longitudinal Observational Data

Abstract

Background Real-world data can inform healthcare decisions by allowing evaluation of nuanced treatment strategies. Longitudinal observational data enable the assessment of dynamic treatment regimes (DTRs), strategies that adapt treatment over time based on patient history, but require causal inference methods to address time-varying confounding. Longitudinal Targeted Minimum Loss-Based Estimation (LTMLE) is a machine learning-based double-robust approach for improved causal effect estimation.

Methods We apply LTMLE to longitudinal registry data to evaluate the impact of erythropoiesis-stimulating agents (ESAs) in the clinical management of low to intermediate-1 risk Myelodysplastic Syndrome (MDS). We define DTRs based on clinically relevant decision rules (e.g. commencing treatment when the haemoglobin level falls below a threshold) and compare them to static treatment regimes (always or never giving ESAs). Outcomes include mortality and health-related quality of life (HRQoL) measured by EQ-5D scores.

Results The static regime of never administering ESAs resulted in declining counterfactual EQ-5D scores and increasing mortality risk over time. In contrast, both the static regime of continuous administration of ESAs and the use of dynamic regimes improved the EQ-5D scores and tended to reduce mortality, although the mortality differences were not statistically significant.

Conclusions The paper provides a case study application of the LTMLE method to evaluate realistic treatment policies under time-varying confounding. The findings support the potential benefits of dynamic treatment strategies for the management of MDS, highlighting the importance of personalised treatment adaptation. The study contributes methodological insights into the applications of LTMLE in small-sample, long-follow-up settings relevant to health technology assessment and policy-making.

Keywords

longitudinal targeted minimum loss-based estimation; Super Learner; time-dependent confounding, EQ-5D, mortality

Highlights

- This study applies the longitudinal targeted minimum loss estimation (LTMLE) method to evaluate the causal effect of static and dynamic treatment strategies using longitudinal observational data.
- We demonstrate the use of the LTMLE method to assess the impact of erythropoiesis stimulating agents (ESAs) on quality-of-life and mortality in patients with low to intermediate-1 risk Myelodysplastic Syndromes (MDS).
- The findings suggest that patients treated under dynamic ESAs treatment regimes show an improved quality-of-life measured by EQ-5D scores and survival compared to those treated under the static treatment regime of never administering ESAs.
- This study contributes to the methodological literature by showcasing the application of the LTMLE method in a small-sample, long-follow-up setting with time-varying confounding, informing health technology assessment and policy decisions.

Introduction

Longitudinal observational data offer important opportunities to generate evidence for comparative effectiveness research. In fact, real-world observational data are increasingly being used to inform policymaking in health care, including regulatory decisions and the evaluation of health technology. When evidence from randomised controlled trials (RCTs) is unavailable, well-designed real-world studies are an acceptable substitute to estimate policy-relevant parameters such as the average treatment effect (ATE). With rich longitudinal data that capture treatment sequences, patient outcomes, and covariates that influence treatment initiation/ switching decisions, researchers can evaluate realistic treatment protocols, the so-called *dynamic treatment regimes* (DTRs). Unlike static treatment regimes, where the sequence of treatments is pre-specified, DTRs allow the decisions to initiate, continue, or switch treatments over time to depend on changing patient characteristics and their treatment responses over time[1–3]. As such, they better reflect clinical decision-making and have greater relevance for practice and policy.

Causal inference from observational data must address the risk of confounding, among other potential sources of bias[4]. Confounding occurs when there are variables that simultaneously affect treatment assignment decisions and health outcomes. Time-varying confounding occurs when the value of certain variables changes over time, influencing both future treatment decisions (e.g. continuation, dose modification, switching) and outcome. This challenge is particularly relevant in longitudinal studies where exposure to treatment and potential confounders are repeatedly measured over time. Traditional approaches such as inverse probability weighting (IPW)[5] and G-estimation[6] address this challenge but rely on correct specification of either the treatment model (IPW) or the outcome model (G-estimation). Double-robust methods model both the treatment mechanism and the outcome mechanism, and can provide unbiased treatment effect estimates if at least one of the two underlying models is correctly specified[7].

Targeted minimum loss-based estimation (TMLE) is a double-robust semiparametric framework that improves flexibility by combining outcome and treatment models[8–10]. TMLE can incorporate Machine Learning (ML) to increase the likelihood of correct model specification of the outcome and the treatment mechanisms, while retaining valid statistical inference, including the estimation of standard errors and confidence intervals[11,12]. TMLE has been used initially to estimate the effects of treatment at a single time point when all potential confounders are baseline variables[13–15]. The approach has been extended to longitudinal data, where time-varying confounding is a primary concern[10], and has been successfully applied to estimate the average causal effects of sustained treatment exposures[12,16–18].

Longitudinal targeted minimum loss-based estimation (LTMLE) is a double-robust method that addresses time-varying confounding; it yields consistent estimates if either the treatment mechanism or the outcome regressions are correctly specified, and achieves greater efficiency than IPW when both models are correctly specified. Despite its potential and the availability of a tutorial that facilitates its practical implementation through the `ltmle` R package[19], there are still relatively few published applications of LTMLE using real-world data (RWD) in contexts directly relevant to health technology assessment (HTA) decision-making.

This paper aims to introduce readers and potential users to LTMLE by illustrating its use in evaluating realistic treatment protocols, a setting relevant to HTA. Using longitudinal data from the European Myelodysplastic Syndromes Registry (EUMDS), we apply LTMLE to evaluate the effects of alternative treatment regimes involving erythropoiesis-stimulating agents (ESAs) in patients with low to intermediate-1 risk myelodysplastic syndromes (LR-MDS).

In clinical practice, the use of ESAs in LR-MDS is often adjusted over time based on the patient's response, measured by haemoglobin levels and transfusion needs, as their rigid use can lead to reduced responsiveness, increased thromboembolic risk, and higher treatment costs[20]. If patients no longer respond to ESAs, their MDS disease status would be reassessed to exclude the possibility of progression. By comparing these strategies to static treatment rules that pre-specify the entire sequence of ESAs administration, e.g. initiate ESAs and continue to administer them - regardless of changing patient characteristics - we aim to identify more efficient, clinically relevant treatment strategies that better reflect real-world decision-making.

As a case study, we estimate the causal effects of static and dynamic treatment strategies on patients' health-related quality of life (HRQOL) measured by the EQ-5D instrument and mortality risks. Rather than providing a technical tutorial, which already exists[19], our goal is to demonstrate the practical relevance and interpretability of the LTMLE method for evaluating treatment strategies using RWD.

The remainder of the paper is structured as follows: the next section introduces the case study, data, and estimation approach; this is followed by the results of our case study. The final section offers a discussion of our findings and the limitations of our study.

Data and Methods

Case Study: Myelodysplastic Syndromes

Myelodysplastic syndromes (MDS) are a family of rare clonal marrow stem-cell disorders, more common in the elderly[21]. At diagnosis, around 75% of patients are classified as having LR-MDS according to the International Prognostic Scoring System (IPSS), which stratifies patients into risk categories: low, intermediate-1, intermediate-2, and high based on bone marrow blast percentage, cytogenetic abnormalities, and number of cytopenias. LR-MDS patients generally have a better prognosis and longer survival than higher-risk groups. The primary goals of treatment in the LR-MDS group are to manage the symptoms of anaemia and improve HRQoL. Anaemia can lead to chronic fatigue and diminished physical, emotional, and cognitive functioning, particularly in older individuals with other comorbidities[22,23]. Red blood cell transfusions (RBCTs) can temporarily reduce anaemia symptoms but may lead to transfusion dependency and iron overload, which can cause organ damage to the liver and heart, with subsequent complications[23–25].

Evidence suggests that early initiation of ESAs in transfusion-independent LR-MDS patients can delay the need for RBCT, maximise their efficacy in terms of response rates and duration, improving HRQoL[22,26]. The current guidelines now recommend ESAs as first-line treatment for LR-MDS patients with symptomatic anaemia[27*], although the evidence base regarding the effectiveness of ESAs in everyday clinical practice and particularly in the older population remains limited. In spite of published treatment protocols, routine practice varies in respect to when clinicians initiate ESAs, the haemoglobin (Hb) levels threshold at which they initiate it, and whether they administer ESAs prior or after RBCT. Importantly, this variability in clinical practice is observed both across countries and within the same healthcare system.

This case study aims to contribute to the evidence base by applying the LTMLE method to assess the impact of different ESAs treatment protocols — both initiation and discontinuation rules — on HRQoL and mortality in LR-MDS patients.

The EUMDS Registry Data

The European Myelodysplastic Syndromes Registry (EUMDS, https://eumds.org/; ID: NCT00600860) is a population-based registry launched in 2008 that prospectively collects detailed patient- and disease-specific information every six months from newly diagnosed MDS patients (within 100 days of diagnosis) recruited from secondary and tertiary care centres across sixteen European countries plus Israel[30]. The registry includes all MDS subtypes classified according to WHO criteria[31,32], although this study focuses on patients with LR-MDS. Patients in EUMDS are followed up until withdrawal (for any reason) or death.

Our study sample includes patients who met the following criteria: diagnosed with LR-MDS, ESAs treatment naïve at baseline (i.e. have not received ESAs but may have received RBCTs prior to diagnosis), and without isolated chromosome 5q deletion (non-del(5q))[†]. We focus on LR-MDS patients as they are the primary candidates for ESAs therapy.

The follow-up period for this analysis spans from 19 March 2008 to 1 September 2019. Follow-up was truncated at the end of 2019 to avoid potential bias from changes in clinical practice, healthcare access, and mortality patterns associated with the COVID-19 pandemic. The large scale and duration of follow-up of EUMDS data makes it possible to estimate the long-term effect of the early introduction of ESAs on HRQoL and mortality in this population.

Notation and the Causal Model

Consider a longitudinal dataset containing n individuals followed from baseline ($t = 0$) at six-month intervals up to time T, where T varies across individuals in our sample. Participants may die or drop out before or at T. At each time point t, we study two outcomes Y_t: (1) HRQoL measured by the EQ-5D instrument, and (2) mortality status.

The treatment indicator $A1_t$ denotes whether a patient receives ESAs at the follow-up visit time t. The censoring indicator $A2_t$ is equal to 1 if the patient is censored at the visit time t and 0 otherwise. Time-varying confounding occurs when a variable L_t affects both treatment $A1_t$ and the outcome of interest Y_{t+1}, and is itself affected by previous treatment $A1_{t-1}$.

We constructed the causal model characterising our setting by reviewing the clinical literature[34] and conducting a focus group discussion with clinicians. An illustrative Directed Acyclic Graph (DAG) is presented in Figure 1. For our research question, baseline confounders L_0 include patient age and the MDS-specific comorbidity index (low, intermediate, or high risk)[‡]. We define two time-varying confounders: haemoglobin (Hb) level L_t^a, and transfusion independence L_t^b, defined as no prior RBCT or fewer than two units of RBCT in the previous six months. We also define a set of time-varying covariates that are not considered confounders as they do not affect treatment decisions but may be helpful in modeling the HRQoL outcomes and mortality risk. They include: (1) a binary indicator for bone marrow blasts $\geq$5%; (2) Karnofsky performance status (0–100 scale, with higher values indicating better function) that efficiently measures geriatric patients' health and functional status; (3) platelet count categories (1 if platelets ≥ 100, 2 if $50 \leq$ platelets < 100, 3 if platelets < 50, with a unit of 10^9/L); and (4) absolute neutrophil count (with a unit of 10^9/L). These variables are excluded from Figure 1 for simplicity.

Figure 1. Directed Acyclic Graph

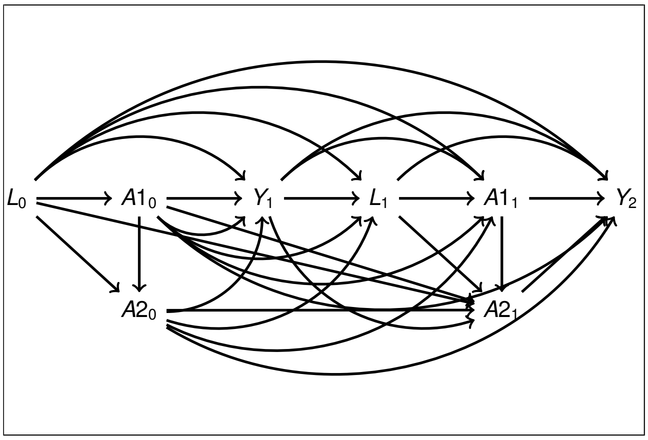

Figure 1 shows the relationships between treatment ($A1_t$), censoring ($A2_t$), outcomes (Y_t), and confounders (L_t) in

*In the U.S., following the COMMANDS trial[28,29], luspatercept is administered as first-line, with ESAs as second-line. However, ESAs remain the first-line in most other countries.

†Patients with isolated 5q deletions were excluded because they are typically managed with lenalidomide as first-line treatment, rather than ESAs, and thus follow a different clinical pathway not comparable to other LR-MDS subtypes[33].

‡Della Porta (2011)[34] found that comorbidity has a significant impact on overall survival and non-leukemic death in patients with very low-, low- and intermediate-risk MDS, underscoring its relevance when deciding on a treatment strategy in MDS patients

time order. The treatment and censoring nodes are preceded by time-varying confounders L_t at time t and followed by the outcome Y_{t+1} at time $t+1$. For simplification, we only show the relationships in the first two time periods ($t = 0, 1$). We also reflect the time ordering in our notation, where we encode the observed data as n independent and identically distributed copies of O:

$$O = (L_0, A1_0, A2_0, Y_1, L_1, A1_1, A2_1, \ldots, \\ L_T, A1_T, A2_T, Y_{T+1}) \quad (1)$$

Baseline confounders L_0 precede the first treatment decision $A1_0$, and no censoring occurs at or before baseline, therefore $A2_0 = 0$. The first post-treatment outcome observed is Y_1.

In the analysis of mortality, the outcome Y_t is a binary variable equal to 1 if death occurs at or before the time period t, and 0 otherwise. As all patients are alive at baseline, we define $Y_0 = 0$. Here, the censoring indicator $A2_t = 0$ means that an individual has not withdrawn from the sample at or before time t.

In the analysis of HRQoL, Y_t is defined as the EQ-5D-3L index, calculated using country-specific tariffs[35,36]. Censoring is defined differently for HRQoL analysis, with $A2_t = 1$ indicating that an individual has left the sample due to either death or withdrawal at or before time t, and 0 otherwise. This differential definition of censoring is necessary because, as described in the following section, selection bias due to censoring is addressed by combining the treatment interventions of interest with a static censoring intervention that counterfactually prevents censoring.

Table 1. Summary statistics at baseline

Baseline Variable	Value
Receive ESAs, n (%)	65 (8%)
EQ-5D scores, mean (SD)	0.71 (0.22)
Age (years), mean (SD)	73.22 (9.78)
MDS comorbidity index: low risk,n (%)	522 (64%)
MDS comorbidity index: intermediate risk,n (%)	293 (36%)
Hb level (g/dL), mean (SD)	9.20 (1.21)
Karnofsky performance status,mean (SD)	81.30 (15.52)
Bone marrow blasts $\geq$ 5%, n (%)	73 (9%)
Platelets $\geq$ 100 (10^9/L), n (%)	636 (78%)
50 $\leq$ Platelets $<$ 100 (10^9/L), n (%)	106 (13%)
Platelets count $<$ 50 (10^9/L), n (%)	82 (10%)
Absolute neutrophil count (10^9/L), n (%)	3.05 (2.88)
RBCT units in current period, mean (SD)	1.64 (3.43)
Transfusion dependent, n (%)	41 (5%)

Notes: Values are reported as mean (standard deviation) for continuous variables, and n (%) for binary variables. Total sample size is 815 at the baseline.

Treatment Protocols Under Evaluation

Our analysis considers joint interventions on treatment $(A1_t)$ and censoring $(A2_t)$, where $A_t = (A1_t, A2_t)$ represents the joint treatment-censoring intervention node at time t. The static censoring component requires $A2_t = 0$ (uncensored) for all $t < t^*$ with $t^* \in 1, ..., T+1$. For the HRQoL analysis (where censoring includes death), this implies counterfactual maintenance of survival and study participation through t; for mortality analysis, it ensures continued study participation without administrative censoring through t. This specification emulates a randomised trial where we intervene to prevent informative censoring, allowing for the estimation of causal effects without selection bias[19].

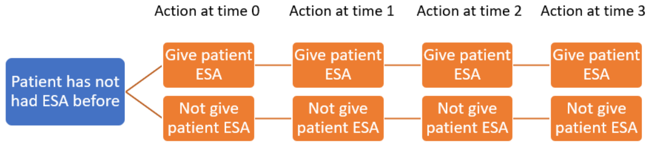

Figure 2. Static Treatment Regimes, Illustrated as a Decision Tree

Notes: Blue shapes represent covariates that are used as inputs in the dynamic treatment regimes, and yellow shapes represent treatment actions.

We define action a_t as an intervention in the treatment-censoring node A_t: $a_t = 1$ corresponds to setting $A1_t$ to 1 (administer ESAs) and $A2_t$ to 0 (keep patient uncensored), while $a_t = 0$ indicates setting $A1_t$ to 0 (not administer ESAs) and as before, $A2_t$ to 0 (keep patient uncensored).

Throughout, the "static treatment regime" refers to fixed sequences of actions, while the "DTR" denotes rules where the action in $A1_t$ depends on time-varying covariates L_t while maintaining always setting $A2_t$ to 0.

Specifically, in our study the static treatment regime "always give ESAs" corresponds to a sequence of treatment-censoring interventions a_t fixed at 1 from the baseline to a selected time period $t^* - 1$ before the end of follow-up ($t^* = 1, .., T+1$): $(a_0 = 1, a_1 = 1, ..., a_{t*-1} = 1)$. Similarly, the static treatment regime "never give ESAs" corresponds to fixing all values of a_t at 0 as $(a_0 = 0, a_1 = 0, ..., a_t^* - 1 = 0)$. Figure 2 illustrates the two static treatment regimes.

In contrast, DTRs aim to capture more realistic, personalised treatment protocols. Such protocols allow the initiation, continuation, or discontinuation of ESAs over time to depend on changing patient characteristics and their previous responses to treatment. We study five DTRs that differ in their strategies for initiating and continuing ESAs, which were developed during focus group discussions with clinical experts.

Figure 3 presents a simplified decision tree that illustrates an example DTR at the first two time points: time 0 (baseline) and time 1 (the first clinical visit after baseline). At baseline, a patient who is transfusion-dependent (TD) is required to start ESAs if their Hb levels fall below a given threshold of 8 g/dL (*DTR1*) or 9 g/dL (*DTR2*), namely if Hb $\leq$ 8 g/dL or $\leq$ 9 g/dL, and should stay off ESAs otherwise. Patients who are non-transfusion-dependent (non-TD) are required to initiate ESAs at a higher Hb threshold of 10 g/dL, namely if Hb $\leq$ 10 g/dL. Then, in the next period, patients who have already started ESAs are required to continue ESAs if they have responded well to treatment. Response is defined based on transfusion dependency and Hb status. A TD patient is considered a responder if they become non-TD after receiving ESAs and their Hb levels do not decline. A non-TD patient is a responder if they remain non-TD after ESAs and the Hb level has not declined. Patients who respond continue receiving ESAs; non-responders discontinue treatment. For each subsequent visit, we apply the same treatment initiation rule for patients who have not yet started ESAs by a given visit, and apply the same treatment continuation/ discontinuation rule for patients who have already started ESAs. The remaining regimes (*DTR3–DTR5*) only consider

Hb levels and do not take transfusion dependency status into account when initiating ESAs, and responses to ESAs are defined based solely on increases in Hb levels after treatment. Specifically, they give ESAs if the patient has an Hb level ≤ 10 (*DTR3*) or 9 (*DTR4*) or 8 (*DTR5*) g/dL, continuing ESAs if the patient's Hb level increases. These rules were developed in consultation with clinical experts to ensure that had relevance to their practices. To save space, we present the results for *DTR1*, which is also the DTR shown in Figure 3, in the main paper, and report results for *DTR2–DTR5* in the Appendix.

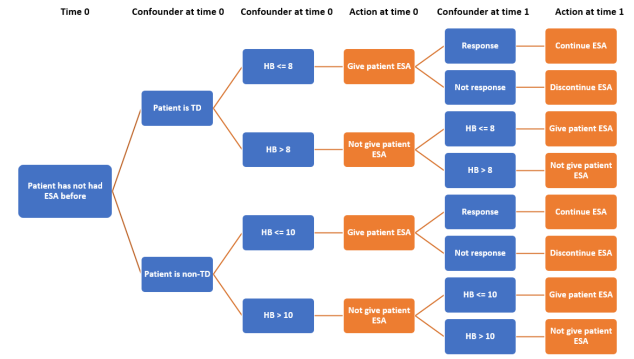

Figure 3. Dynamic Treatment Regime 1, Illustrated as a Decision Tree

Notes: TD stands for transfusion-dependent patients; HB stands for haemoglobin levels; ESA is the treatment of interest. Blue shapes represent covariates that are used as inputs in the dynamic treatment regimes, and yellow shapes represent treatment actions.

Causal Parameters and Identification Assumptions

Our causal parameter of interest is the intervention-specific mean $E[Y^d_{t^*}]$. $Y^d_{t^*}$ is the potential outcome that would be observed in a selected time period before the end of follow-up ($t^* = 1, .., T+1$) if an individual - perhaps contrary to the fact - followed a particular longitudinal intervention d (a static or dynamic treatment regime) up to $t^* - 1$. Intuitively, the expectation of this potential outcome captures the average outcome in the population if everyone followed a given treatment protocol. We evaluate the counterfactual outcomes for the two static treatment regimes and the dynamic regime defined earlier. Due to the static censoring intervention component of all treatment regimes, our estimand in the EQ-5D analysis reflects the mean EQ-5D that would be observed if, perhaps contrary to the fact, patients remained alive and adhered to the treatment regime until $t^* - 1$.

As $E[Y^d_{t^*}]$ is a counterfactual quantity, without further assumptions it cannot be estimated from the observed data. If we simply summarised the observed outcomes for those who actually followed this treatment protocol, our results would be biased due to baseline and time-varying confounding. Two crucial assumptions are necessary to identify $E[Y^d_{t^*}]$: the sequential randomisation assumption[6] and the positivity assumption. Under sequential randomisation, conditional on the observed histories of treatment and confounders, the potential outcome in each time period is independent of the preceding treatment status. This is the longitudinal version of the "no unmeasured confounders" assumption, implying that after controlling for baseline covariates and the observed histories of treatment and confounders, the next treatment decision is "as good as random". The positivity assumption requires that each observation has a positive probability of following the rule d at each time point. For static regimes, this means that each patient in our study must have a positive probability of receiving (or not receiving) ESAs in each time period.

Estimation via LTMLE

General approach Here, we briefly describe the estimation approach using LTMLE. The quantity $E(Y_d(t^*))$ can be written as a sequence of recursively defined conditional expectations, using the longitudinal G-computation formula[10]. This formulation allows for the estimation of the counterfactual mean through a series of sequential regressions. In summary, at each time point, the outcome $Y(t)$ is predicted conditional on the observed past covariates and treatment values, where the treatment is set according to the predefined longitudinal treatment protocol d. This procedure allows for the adjustment of time-varying confounding in a sequential manner. First, only confounding in the last time period is adjusted for by regressing the observed outcome on the treatment variable and confounders in the previous period only, as would be done in a study with only baseline confounding. Then, predictions from this regression are obtained, where the treatment variable is set to the value that would be required by the longitudinal treatment protocol under evaluation. These predictions are subsequently used as the outcome in the next regression, where the treatment variable and confounders in the previous period are controlled for, and predictions are made again. This process is repeated until only the baseline confounders need to be adjusted for, and the expected counterfactual outcome is estimated as the average of the final predictions.

This approach could be subject to misspecification bias if it relied solely on the correct specification of the sequential regressions. A doubly robust and semi-parametric version of this sequential regression estimator can be constructed by including a covariate, which is usually a weight, that uses information from the treatment assignment mechanism[37]. The LTMLE estimator performs this double-robust adjustment in each iterated regression, by updating the predictions with a covariate that is a function of the estimated propensity score. The resulting estimator is doubly robust and consistent if either the treatment mechanism or the sequential regressions are correctly specified[10].

To reduce reliance on parametric assumptions and improve model flexibility, ML algorithms are recommended to estimate both the treatment mechanism and the outcome regressions. The LTMLE accommodates the use of the Super Learner (SL), an ensembling ML algorithm that employs cross-validation to build the best weighted combination of candidate algorithms, instead of selecting only one method[38]. Compared to parametric models such as generalized linear models (GLMs), SL can capture complex, non-linear relationships in the data. Although GLMs may produce lower variance when correctly specified, SL typically reduces bias and improves overall predictive accuracy, highlighting the trade-off between bias and variance in estimator selection.

Implementation We implement the LTMLE method to estimate the counterfactual values of the EQ-5D index and the mortality risks for a series of time periods t^*, under the two static treatment regimes ("never give ESA" and "always give ESAs") and five DTRs developed in consultation with clinicians.

In terms of model specifications, baseline and time-varying confounders are included in both outcome models (sequential regressions) and treatment models (propensity score estimation), as they affect both patient outcomes and treatment decisions. Table 1 reports the summary statistics of variables at baseline. Additional time-varying covariates (see the list of the covariates in the section Notation and Causal Model) are included in the outcome models for predicting HRQoL and mortality risks, as they are assumed to affect the patient outcomes but not the treatment decisions. Time-varying covariates were also included in the censoring model in HRQoL anlaysis, where censoring includes both mortality and withdrawal.

We use both parametric GLMs and the SL algorithms to estimate components of the treatment and censoring mechanisms, as well as the outcome mechanisms. The SL library includes: GLM, Stepwise regression[39], neural networks, generalized additive models[40], Elastic net[41]. In this case study, SL yielded smaller variances around the counterfactual mean parameters; thus SL-based results are reported in the main analysis.

To improve model specifications, reduce model complexity and potential overfitting, we include only one lag of time-varying covariates and confounders (rather than full histories) and adjust for baseline covariates in all models. Missing values are imputed using the method of the last observation carried forward. The 95% confidence intervals and the standard errors are based on the estimated influence curve and are correct asymptotically when both treatment mechanisms and outcome mechanisms are consistently estimated. The LTMLE models are implemented using the `ltmle` package in `R`, version 4.1.1.

Results

We report results separately for static regimes and the DTRs, focusing on estimates obtained using the SL algorithm. In our case study, LTMLE combined with SL algorithm produce lower variance in the estimated counterfactual means and offer greater flexibility than parametric GLMs. Graphical summaries are presented below, with detailed tables available in the Appendix.

After excluding patients with no information on EQ-5D scores across all visits, we have a study sample at baseline of 815 individuals for both the HRQoL and the survival analyses. There are six months between each visit. Table 2 reports the number of patients who followed each of the treatment regimes under investigation at each time point. Most patients follow the static regime "never give ESAs". The number of patients following the *DTR1* is higher than the number of patients following the static rule of always giving ESAs, which is in line with expectations, as a dynamic rule is more realistic than the static rule. The numbers of followers for the other DTRs are not presented here, but *DTR1* has the largest number of patients following the assigned treatment strategy among the five dynamic rules.

Less than 50 patients follow either *DTR1* or the static regime of always administering ESAs after time period 5 (that is, 2.5 years post-baseline). This is consistent with the clinician's feedback that, on average, the effectiveness of ESAs lasts two years, and clinicians cease administering ESAs to patients if they no longer respond to treatment. Therefore, we concentrate on the estimates for the first five time periods. The sample size decreases over time due to patients withdrawing from the registry and due to death.

Table 2. Distribution of patients following different treatment regimes

time	**0**	**1**	**2**	**3**	**4**	**5**
Static1: always give ESAs						
follow	**65**	**54**	**46**	**37**	**28**	**24**
not follow	750	724	631	541	471	396
Static0: Never give ESAs						
follow	**750**	**624**	**525**	**442**	**375**	**316**
not follow	65	154	152	136	124	104
DTR1						
follow	**233**	**170**	**109**	**76**	**55**	**46**
no ESAs	168	126	85	65	51	44
ESAs	65	44	24	11	4	2
not follow	582	608	568	502	444	374
not censor	815	778	677	578	499	420
death	0	11	58	107	148	179
censor	0	26	80	130	168	216
total sample	815	815	815	815	815	815

Note: The bottom panel shows, for each time period, the number of individuals still in the sample (i.e., not censored), the number of deaths, and other censoring events. Among those remaining, we report how many followed or did not follow each treatment rule. For *DTR1*, we also show how many received ESAs among those following the rule.

Health-Related Quality of Life

In this section, we present the LTMLE estimated counterfactual mean EQ-5D index values under the static regimes and DTRs, evaluated at visits 1-5. These results represent the expected EQ-5D outcomes that would have been observed had the patient population, possibly contrary to fact, followed each specified regime.

In Figure 4, we report the estimated counterfactual mean EQ-5D values for patients under the dynamic treatment regime (*DTR1*, subfigure 4a), as well as under the static regimes of "always give ESAs" (*Static1*, subfigure 4c) and "never give ESAs" (*Static0*, subfigure 4f). We estimate the ATEs by contrasting the counterfactual mean EQ-5D scores between *Static1* and *Static0*, and between *DTR1* and either *Static0* or *Static1*. The counterfactual mean EQ-5D index values are the lowest under the *Static0* regime, showing a decreasing trend over time, while the highest values are observed under *DTR1*. The trajectories of HRQoL remain relatively stable over time for the *Static1* and *DTR1* regimes.

We estimate the ATEs by contrasting the counterfactual mean EQ-5D scores between *Static1* and *Static0*, and between *DTR1* and either *Static0* or *Static1*. Subfigure 4e shows that there are significant benefits in terms of HRQoL measured at the 1st, 2nd and 5th time periods for *Static1* compared to *Static0*. Under *DTR1*, patients have significantly higher EQ-5D index values in the 2nd and 5th time points compared to under *Static0* (subfigure 4d). The difference in counterfactual mean EQ-5D scores between *Static1* and *DTR1* is not statistically significant (subfigure 4b). Full point estimates and confidence intervals for these causal contrasts are available in the Appendix.

Mortality Analysis

Figure 5 presents the LTMLE-estimated counterfactual probabilities of death under each treatment regime. It is important to note that we estimate the counterfactual probability of death at each time point, rather than the total survival time (time-to-death) or cumulative survival curves. Patients under the *Static0* regime (subfigure 5f) exhibit the greatest increase in the counterfactual mortality risk over time, followed by those under the *Static1* regime (subfigure 5c). The slowest increase in mortality risk is estimated for patients under the *DTR1* (subfigure 5a).

Pairwise causal contrasts indicate that there is no statistically significant difference in the counterfactual probability of death between the *Static1* and the *Static0* regimes, largely due to the high variance of the estimates. Patients following the *DTR1* are estimated to have lower counterfactual mortality risks than those under either static regime; however, these differences, as measured by the ATEs, do not reach statistical significance. Nonetheless, the observed trends in these counterfactual contrasts support the hypothesis that a more realistic, personalised treatment protocol such as *DTR1* may be associated with better survival outcomes. A full summary of estimated mortality probabilities across all treatment strategies and pairwise comparisons between them is provided in Appendix Table A3.

Discussion and Limitations

The paper applies the LTMLE to longitudinal data from the EUMDS registry to evaluate both static and dynamic ESAs treatment strategies in patients with LR-MDS. We estimate counterfactual mean EQ-5D index values and counterfactual mortality risks, as well as the ATEs comparing different static regimes and DTRs. We utilise SL-based estimation because of its improved flexibility and, in our setting, lower variance compared to parametric models.

The presence of time-varying confounding is a key concern, as LR-MDS has a heterogeneous evolution. Patients may or may not respond; the disease can progress quickly, and treatment decisions depend on evolving Hb levels and transfusion dependence—both influenced by prior ESAs use. After accounting for such time-varying confounding, we find no significant differences in mortality risk across treatment strategies. However, patients following DTRs (where treatment initiation and continuation depend on Hb levels and transfusion status) and those always receiving ESAs show higher EQ-5D index values than patients never treated with ESAs. These findings suggest that DTRs may achieve similar or even better outcomes than the static treatment strategy of continuously administering ESAs, supporting the relevance of personalised care pathways in MDS.

To address potential selection bias from informative right-censoring, particularly due to death, our analyses emulate a hypothetical randomised experiment in which censoring - due to either death or withdrawal from the study - is prevented until the time point selected to evaluate the counterfactual outcomes[42]. While the causal interpretability of results under this conceptual intervention — where deaths are disallowed — has been debated, this approach remains an accepted solution to selection bias in longitudinal studies (see e.g. Neugebauer et al. (2014)[43] and Kreif et al. (2021)[44]). Alternative strategies to handle selection bias include composite outcomes incorporating death, principal stratification, and competing risks frameworks (see Young et al. (2020)[45] for a review).

As with any observational analysis, our estimates may be biased by unmeasured confounding. While the LTMLE framework, combined with ML, aims to adjust for measured confounders flexibly and robustly, it cannot account for variables not captured in the dataset. As patients' prognosis changes rapidly and the EUMDS data is only collected every six months, there may be some residual confounding, such as new comorbidities that may affect both treatment decisions over time and the outcome of interest. In such cases, formal tools for sensitivity analysis can be helpful in evaluating the robustness of results to potential residual confounding. Currently, most available methods for assessing sensitivity to unmeasured confounding are designed for simpler parametric models and are not readily applicable to longitudinal settings with time-varying treatment and censoring[46], especially when using data-adaptive estimation procedures like the LTMLE method. Further methodological development is needed in this area.

Applying LTMLE with limited sample size and prolonged follow-up presents challenges, including attrition and reduced support for complex models. Most previous LTMLE applications have used large samples and fewer time points[11,12,16,17,43,47–49]. Our study contributes to the smaller body of work that shows that LTMLE can still produce informative results in modest samples with extended follow-up[50,51]. Simulation studies and applied analyses have indicated that the precision of LTMLE estimates improves with larger sample sizes and more frequent observations across follow-up time points, but even with small sample sizes, the results remain relatively stable and robust[51].

Another source of bias that should be considered is the potential for measurement error. Clinicians report that dose adjustments of ESAs—rather than abrupt discontinuation—are common in practice for nonresponders, but such granularity is not recorded in the registry.

In this paper, LTMLE is employed to estimate discrete-time mortality risks and we do not construct a full survival curve. However, this is possible with LTMLE, as the method is more broadly applicable to time-to-event outcomes. An illustration is provided by Neugebauer (2014)[43] who estimate the effects of dynamic treatment regimes by contrasting their counterfactual survival curves, constructed from estimates of discrete-time hazards.

Although our results do not identify a clearly optimal ESAs strategy when comparing different DTRs, due to sample uncertainty, our findings suggest that these regimes are at least as effective as static treatment policies where patients receive continuous treatment with ESAs, and may offer improved HRQoL with less overtreatment.

More broadly, our study contributes to the literature by applying a robust longitudinal causal inference method to evaluate the effectiveness of ESA in MDS - a setting where RCT evidence is limited and time-varying confounding is often overlooked. With appropriate methodology, observational data can provide credible and policy-relevant insights.

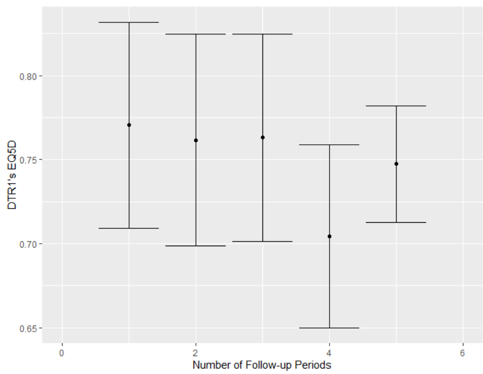

(a) *DTR1*

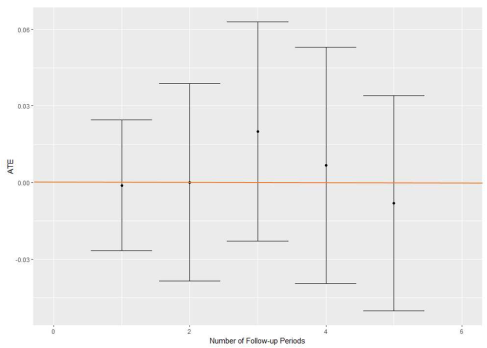

(b) *DTR1* vs. *Static1*

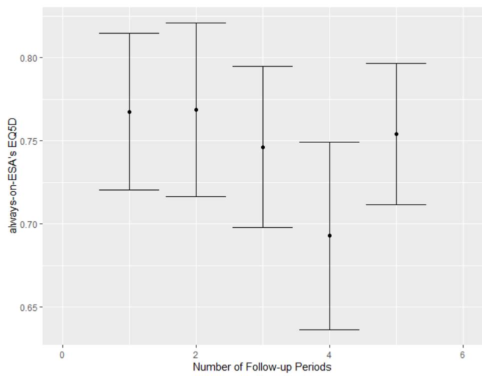

(c) *Static1*

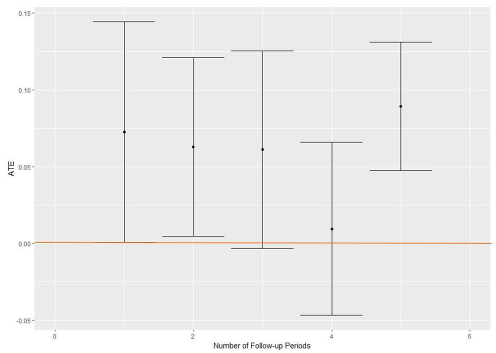

(d) *DTR1* vs. *Static0*

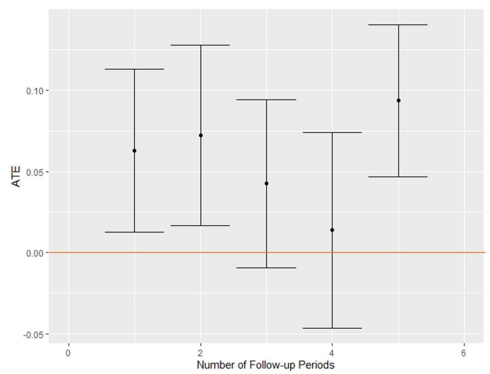

(e) *Static1* vs. *Static0*

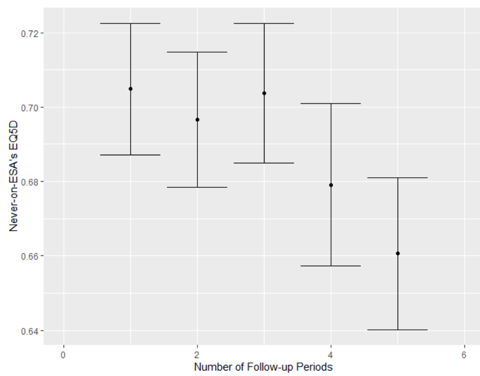

(f) *Static0*

Figure 4. Estimated Counterfactual Mean EQ-5D Scores and ATEs

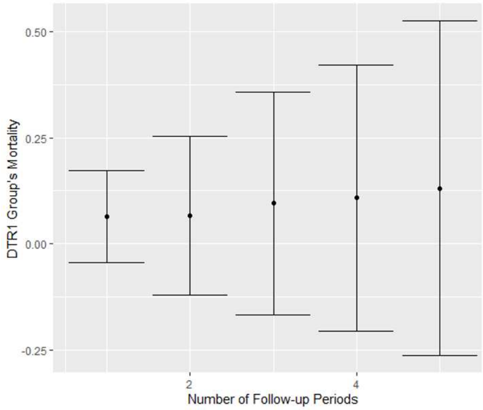

(a) *DTR1*

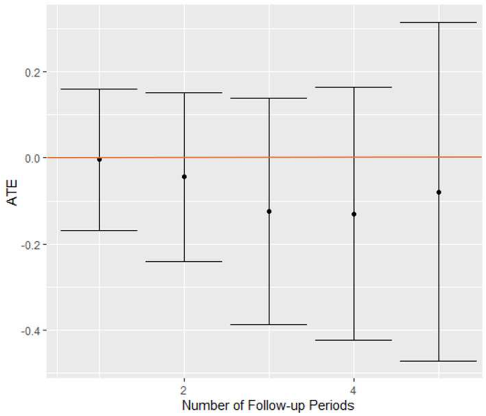

(b) *DTR1* vs. *Static1*

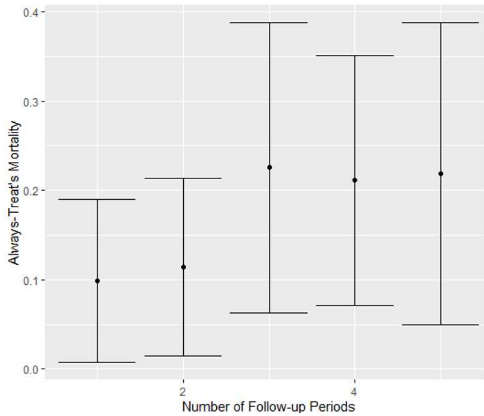

(c) *Static1*

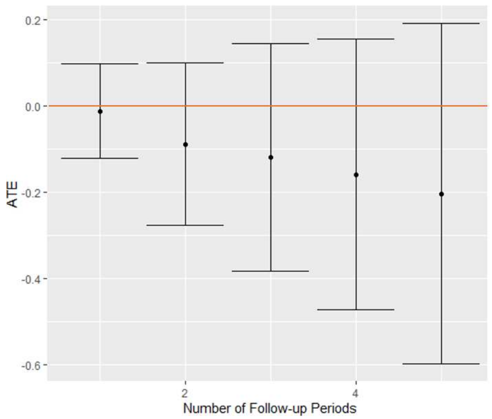

(d) *DTR1* vs. *Static0*

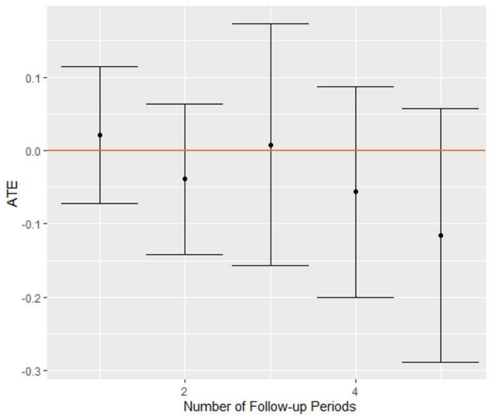

(e) *Static1* vs. *Static0*

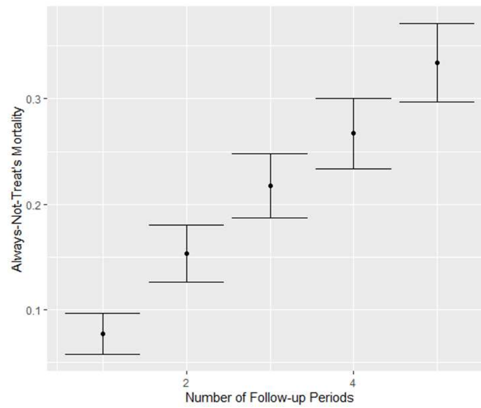

(f) *Static0*

Figure 5. Estimated Counterfactual Mortality Probabilities and ATEs

References

1. Moodie EE, Richardson TS and Stephens DA. Demystifying optimal dynamic treatment regimes. *Biometrics* 2007; 63(2): 447–455.
2. Murphy SA, van der Laan MJ, Robins JM et al. Marginal mean models for dynamic regimes. *Journal of the American Statistical Association* 2001; 96(456): 1410–1423.
3. Chakraborty B and Murphy SA. Dynamic treatment regimes. *Annual review of statistics and its application* 2014; 1: 447–464.
4. Hernán M and Robins J. *Causal Inference: What if.* Chapman & Hill/CRC, 2020.
5. Hernán MÁ, Brumback B and Robins JM. Marginal structural models to estimate the causal effect of zidovudine on the survival of hiv-positive men. *Epidemiology* 2000; : 561–570.
6. Robins J. 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* 1986; 7(9-12): 1393–1512.
7. Robins JM, Rotnitzky A and Zhao LP. Analysis of semiparametric regression models for repeated outcomes in the presence of missing data. *Journal of the american statistical association* 1995; 90(429): 106–121.
8. Van Der Laan MJ and Rubin D. Targeted maximum likelihood learning. *The international journal of biostatistics* 2006; 2(1).
9. Van der Laan MJ, Rose S et al. *Targeted learning: causal inference for observational and experimental data*, volume 10. Springer, 2011.
10. van der Laan MJ and Gruber S. Targeted minimum loss based estimation of causal effects of multiple time point interventions. *The international journal of biostatistics* 2012; 8(1).
11. Schnitzer ME, Moodie EE, van der Laan MJ et al. Modeling the impact of hepatitis c viral clearance on end-stage liver disease in an hiv co-infected cohort with targeted maximum likelihood estimation. *Biometrics* 2014; 70(1): 144–152.
12. Decker AL, Hubbard A, Crespi CM et al. Semiparametric estimation of the impacts of longitudinal interventions on adolescent obesity using targeted maximum-likelihood: accessible estimation with the ltmle package. *Journal of causal inference* 2014; 2(1): 95–108.
13. Gruber S and van der Laan MJ. An application of collaborative targeted maximum likelihood estimation in causal inference and genomics. *The International Journal of Biostatistics* 2010; 6(1).
14. Pearl M, Balzer L and Ahern J. Targeted estimation of marginal absolute and relative associations in case–control data. *Epidemiology* 2016; 27(4): 512–517.
15. Luque-Fernandez MA, Belot A, Valeri L et al. Data-adaptive estimation for double-robust methods in population-based cancer epidemiology: risk differences for lung cancer mortality by emergency presentation. *American journal of epidemiology* 2018; 187(4): 871–878.
16. Schnitzer ME, van der Laan MJ, Moodie EE et al. Effect of breastfeeding on gastrointestinal infection in infants: a targeted maximum likelihood approach for clustered longitudinal data. *The annals of applied statistics* 2014; 8(2): 703.
17. Tran L, Yiannoutsos CT, Musick BS et al. Evaluating the impact of a hiv low-risk express care task-shifting program: a case study of the targeted learning roadmap. *Epidemiologic methods* 2016; 5(1): 69–91.
18. Brown DM, Petersen M, Costello S et al. Occupational exposure to pm2. 5 and incidence of ischemic heart disease: longitudinal targeted minimum loss-based estimation. *Epidemiology (Cambridge, Mass)* 2015; 26(6): 806.
19. Lendle SD, Schwab J, Petersen ML et al. ltmle: an r package implementing targeted minimum loss-based estimation for longitudinal data. *Journal of Statistical Software* 2017; 81: 1–21.
20. Tonia T, Mettler A, Robert N et al. Erythropoietin or darbepoetin for patients with cancer. *Cochrane database of systematic reviews* 2012; (12).
21. Neukirchen J, Schoonen WM, Strupp C et al. Incidence and prevalence of myelodysplastic syndromes: data from the düsseldorf mds-registry. *Leukemia research* 2011; 35(12): 1591–1596.
22. Park S, Kelaidi C, Sapena R et al. Early introduction of esa in low risk mds patients may delay the need for rbc transfusion: a retrospective analysis on 112 patients. *Leukemia research* 2010; 34(11): 1430–1436.
23. Hellström-Lindberg E and van de Loosdrecht A. Erythropoiesis stimulating agents and other growth factors in low-risk mds. *Best practice & research Clinical haematology* 2013; 26(4): 401–410.
24. Jabbour E, Kantarjian HM, Koller C et al. Red blood cell transfusions and iron overload in the treatment of patients with myelodysplastic syndromes. *Cancer: Interdisciplinary International Journal of the American Cancer Society* 2008; 112(5): 1089–1095.
25. Park S, Greenberg P, Yucel A et al. Clinical effectiveness and safety of erythropoietin-stimulating agents for the treatment of low-and intermediate-1- risk myelodysplastic syndrome: a systematic literature review. *British journal of haematology* 2019; 184(2): 134–160.
26. Garelius H, Johnston WT, Smith AG et al. Erythropoiesis-stimulating agents significantly delay the onset of a regular transfusion need in nontransfused patients with lower-risk myelodysplastic syndrome. *Journal of internal medicine* 2017; 281(3): 284–299.
27. Killick SB, Carter C, Culligan D et al. Guidelines for the diagnosis and management of adult myelodysplastic syndromes. *British journal of haematology* 2014; 164(4): 503–525.
28. Platzbecker U, Della Porta MG, Santini V et al. Efficacy and safety of luspatercept versus epoetin alfa in erythropoiesis-stimulating agent-naive, transfusion-dependent, lower-risk myelodysplastic syndromes (commands): interim analysis of a phase 3, open-label, randomised controlled trial. *The Lancet* 2023; 402(10399): 373–385.
29. Della Porta MG, Garcia-Manero G, Santini V et al. Luspatercept versus epoetin alfa in erythropoiesis-stimulating agent-naive, transfusion-dependent, lower-risk myelodysplastic syndromes (commands): primary analysis of a phase 3, open-label, randomised, controlled trial. *The Lancet Haematology* 2024; 11(9): e646–e658.
30. de Swart L, Smith A, Johnston TW et al. Validation of the revised international prognostic scoring system (ipss-r) in patients with lower-risk myelodysplastic syndromes: a report from the prospective european leukaemianet mds (eumds) registry. *British journal of haematology* 2015; 170(3): 372–383.

31. Bennett JM. World health organization classification of the acute leukemias and myelodysplastic syndrome. *International journal of hematology* 2000; 72(2): 131–133.
32. Greenberg P, Cox C, LeBeau MM et al. International scoring system for evaluating prognosis in myelodysplastic syndromes. *Blood, The Journal of the American Society of Hematology* 1997; 89(6): 2079–2088.
33. Duong VH, Komrokji RS and List AF. Efficacy and safety of lenalidomide in patients with myelodysplastic syndrome with chromosome 5q deletion. *Therapeutic advances in hematology* 2012; 3(2): 105–116.
34. Della Porta MG, Malcovati L, Strupp C et al. Risk stratification based on both disease status and extra-hematologic comorbidities in patients with myelodysplastic syndrome. *Haematologica* 2011; 96(3): 441.
35. Foundation ER. *EQ-5D-3L User Guide*, 2018. URL https://euroqol.org/publications/user-guides/. Available from: https://euroqol.org/publications/user-guides/.
36. Stauder R, Yu G, Koinig KA et al. Health-related quality of life in lower-risk mds patients compared with age-and sex-matched reference populations: a european leukemianet study. *Leukemia* 2018; 32(6): 1380–1392.
37. Bang H and Robins JM. Doubly robust estimation in missing data and causal inference models. *Biometrics* 2005; 61(4): 962–973.
38. Van der Laan MJ, Polley EC and Hubbard AE. Super learner. *Statistical applications in genetics and molecular biology* 2007; 6(1).
39. Polley EC and Van Der Laan MJ. chapter Super learner in prediction. bepress, 2010.
40. Hastie TJ. Generalized additive models. In *Statistical models in S*. Routledge, 2017. pp. 249–307.
41. Friedman J, Hastie T and Tibshirani R. Regularization paths for generalized linear models via coordinate descent. *Journal of statistical software* 2010; 33(1): 1.
42. Weuve J, Tchetgen EJT, Glymour MM et al. Accounting for bias due to selective attrition: the example of smoking and cognitive decline. *Epidemiology* 2012; 23(1): 119–128.
43. Neugebauer R, Schmittdiel JA and van der Laan MJ. Targeted learning in real-world comparative effectiveness research with time-varying interventions. *Statistics in medicine* 2014; 33(14): 2480–2520.
44. Kreif N, Sofrygin O, Schmittdiel JA et al. Exploiting nonsystematic covariate monitoring to broaden the scope of evidence about the causal effects of adaptive treatment strategies. *Biometrics* 2021; 77(1): 329–342.
45. Young JG, Stensrud MJ, Tchetgen Tchetgen EJ et al. A causal framework for classical statistical estimands in failure-time settings with competing events. *Statistics in medicine* 2020; 39(8): 1199–1236.
46. D'Agostino McGowan L. Sensitivity analyses for unmeasured confounders. *Current Epidemiology Reports* 2022; 9(4): 361–375.
47. Kreif N, Tran L, Grieve R et al. Estimating the comparative effectiveness of feeding interventions in the pediatric intensive care unit: a demonstration of longitudinal targeted maximum likelihood estimation. *American journal of epidemiology* 2017; 186(12): 1370–1379.
48. Petersen M, Schwab J, Gruber S et al. Targeted maximum likelihood estimation for dynamic and static longitudinal marginal structural working models. *Journal of causal inference* 2014; 2(2): 147–185.
49. Tran L, Yiannoutsos C, Wools-Kaloustian K et al. Double robust efficient estimators of longitudinal treatment effects: comparative performance in simulations and a case study. *The international journal of biostatistics* 2019; 15(2).
50. Schnitzer ME, Lok JJ and Bosch RJ. Double robust and efficient estimation of a prognostic model for events in the presence of dependent censoring. *Biostatistics* 2016; 17(1): 165–177.
51. Schomaker M, Luque-Fernandez MA, Leroy V et al. Using longitudinal targeted maximum likelihood estimation in complex settings with dynamic interventions. *Statistics in medicine* 2019; 38(24): 4888–4911.

Estimating the Causal Effect of Realistic Treatment Strategies Using Longitudinal Observational Data: Supplementary Materials

Table A1: Estimates of Static Treatment Regimes

time	**1**	**2**	**3**	**4**	**5**
Health-Related Quality of Life (HRQoL) Analysis					
ATE (*Static1* vs. *Static0*)	**0.06****	**0.07****	0.04	0.01	**0.09*****
SE	0.03	0.03	0.03	0.03	0.02
Y (*Static1*)	0.77	0.77	0.75	0.69	0.75
SE	0.02	0.03	0.02	0.03	0.02
Y (*Static0*)	0.70	0.70	0.70	0.68	0.66
SE	0.01	0.01	0.01	0.01	0.01
Mortality Analysis					
ATE (*Static1* vs. *Static0*)	0.02	-0.04	0.01	-0.06	-0.12
SE	0.05	0.05	0.08	0.07	0.09
Y (*Static1*)	0.10	0.11	0.23	0.21	0.22
SE	0.05	0.05	0.08	0.07	0.09
Y (*Static0*)	0.08	0.15	0.22	0.27	0.33
SE	0.01	0.01	0.02	0.02	0.02

Note: 1. * 90% significant, ** 95% significant, *** 99% significant. SE = standard error, Y = counterfactual outcome, ATE = average treatment effect. 2. The table presents the LTMLE estimates of counterfactual mean EQ-5D index values (upper panel) and counterfactual mortality risks (lower panel) for patients following the static treatment rules of always giving ESA (*Static1*) or never giving ESA (*Static0*). ATEs were estimated by comparing the counterfactual mean outcomes of patients following the *Static1* to *Static0*.

Table A2: Estimates of Dynamic Treatment Regimes: HRQoL Analysis

time	**1**	**2**	**3**	**4**	**5**
Y (*DTR1*)	0.78	0.76	0.76	0.69	0.75
SE	0.04	0.03	0.03	0.03	0.02
Y (*DTR2*)	0.78	0.76	0.76	0.69	0.77
SE	0.04	0.03	0.03	0.03	0.02
Y (*DTR3*)	0.77	0.77	0.77	0.70	0.76
SE	0.03	0.03	0.03	0.02	0.02
Y (*DTR4*)	0.69	0.69	0.69	0.67	0.67
SE	0.01	0.01	0.01	0.01	0.01
Y (*DTR5*)	0.71	0.73	0.72	0.71	0.68
SE	0.01	0.03	0.02	0.02	0.02
ATE (*DTR1* vs. *Static0*)	**0.07****	**0.06****	**0.06***	0.01	**0.09*****
SE	0.04	0.03	0.03	0.03	0.02
ATE (*DTR1* vs. *Static1*)	0.00	0.00	0.02	0.01	-0.01
SE	0.01	0.02	0.02	0.02	0.02
ATE (*DTR2* vs. *Static0*)	**0.07***	**0.07****	**0.06***	0.02	**0.11*****
SE	0.04	0.03	0.03	0.03	0.02
ATE (*DTR2* vs. *Static1*)	0.01	0.00	0.02	0.00	0.00
SE	0.02	0.02	0.02	0.03	0.02
ATE (*DTR3* vs. *Static0*)	**0.07***	**0.07****	**0.06****	0.02	**0.10*****
SE	0.04	0.03	0.03	0.03	0.02
ATE (*DTR3* vs. *Static1*)	0.00	0.00	0.02	-0.01	0.00
SE	0.03	0.02	0.02	0.03	0.02
ATE (*DTR4* vs. *Static0*)	-0.01	-0.01	-0.01	-0.01	0.01
SE	0.01	0.01	0.01	0.01	0.01
ATE (*DTR4* vs. *Static1*)	**-0.08*****	**-0.08*****	**-0.05***	**-0.02****	**-0.09*****
SE	0.02	0.03	0.03	0.03	0.03
ATE (*DTR5* vs. *Static0*)	0.00	0.03	0.02	0.02	0.02
SE	0.00	0.02	0.02	0.02	0.01
ATE (*DTR5* vs. *Static1*)	**-0.06****	-0.04	-0.03	0.01	**-0.08****
SE	0.03	0.04	0.03	0.03	0.04
ATE (*DTR1* vs. *DTR2*)	0.00	0.00	**-0.01****	0.00	**-0.01*****
SE	0.00	0.00	0.00	0.00	0.00
ATE (*DTR1* vs. *DTR3*)	0.01	0.00	-0.01	0.00	0.00
SE	0.01	0.00	0.00	0.00	0.01
ATE (*DTR1* vs. *DTR4*)	**0.08****	**0.07****	**0.07****	0.02	**0.09*****
SE	0.04	0.03	0.03	0.03	0.02
ATE (*DTR1* vs. *DTR5*)	**0.07***	0.04	0.04	-0.01	**0.06*****
SE	0.04	0.04	0.03	0.03	0.02

Note: 1. * 90% significant, ** 95% significant, *** 99% significant. SE = standard error, Y = counterfactual mean outcome, ATE = average treatment effect. 2. The table presents the LTMLE estimates of counterfactual mean EQ-5D scores for patients following the five dynamic treatment rules (*DTR1-5*). ATEs were estimated by comparing the counterfactual mean EQ-5D scores of patients following a specific DTR to those following the static rule of always giving ESA (*Static1*) or to those following the static rule of never giving ESA (*Static0*), or by comparing the counterfactual mean mortality risks of patients following different DTRs.

Table A3: Estimates of Dynamic Treatment Regimes: Mortality Analysis

time	**1**	**2**	**3**	**4**	**5**
Y (*DTR1*)	0.06	0.07	0.10	0.11	0.13
SE	0.06	0.10	0.13	0.16	0.20
Y (*DTR2*)	0.07	0.06	0.10	0.12	0.13
SE	0.06	0.09	0.13	0.16	0.20
Y (*DTR3*)	0.14	0.12	0.15	0.14	0.14
SE	0.07	0.07	0.10	0.13	0.18
Y (*DTR4*)	0.16	0.16	0.18	0.21	0.25
SE	0.07	0.06	0.07	0.10	0.12
Y (*DTR5*)	0.14	0.15	0.19	0.28	0.31
SE	0.06	0.04	0.05	0.07	0.07
ATE (*DTR1* vs. *Static0*)	-0.01	-0.09	-0.12	-0.16	-0.20
SE	0.06	0.10	0.13	0.16	0.20
ATE (*DTR1* vs. *Static1*)	0.00	-0.04	-0.12	-0.13	-0.08
SE	0.08	0.10	0.13	0.15	0.20
ATE (*DTR2* vs. *Static0*)	0.00	-0.09	-0.12	-0.15	-0.20
SE	0.06	0.09	0.13	0.16	0.20
ATE (*DTR2* vs. *Static1*)	-0.01	-0.05	-0.15	-0.12	-0.08
SE	0.08	0.09	0.13	0.15	0.19
ATE (*DTR3* vs. *Static0*)	0.06	-0.03	-0.07	-0.12	-0.19
SE	0.07	0.07	0.10	0.13	0.18
ATE (*DTR3* vs. *Static1*)	0.02	0.01	-0.08	-0.08	-0.09
SE	0.06	0.06	0.09	0.12	0.17
ATE (*DTR4* vs. *Static0*)	0.08	0.01	-0.03	-0.06	-0.08
SE	0.07	0.06	0.07	0.10	0.12
ATE (*DTR4* vs. *Static1*)	0.05	0.07	-0.03	-0.02	0.02
SE	0.08	0.07	0.10	0.11	0.13
ATE (*DTR5* vs. *Static0*)	0.07	-0.01	-0.02	0.01	-0.03
SE	0.06	0.04	0.05	0.07	0.07
ATE (*DTR5* vs. *Static1*)	0.03	0.05	-0.01	0.04	0.09
SE	0.06	0.06	0.08	0.09	0.11
ATE (*DTR1* vs. *DTR2*)	0.00	0.01	0.01	0.00	0.00
SE	0.02	0.01	0.07	0.05	0.07
ATE (*DTR1* vs. *DTR3*)	-0.06	-0.05	-0.05	-0.05	-0.02
SE	0.08	0.04	0.10	0.04	0.05
ATE (*DTR1* vs. *DTR4*)	-0.06	**-0.10****	-0.10	-0.10	-0.12
SE	0.10	0.02	0.13	0.15	0.19
ATE (*DTR1* vs. *DTR5*)	-0.05	-0.08	-0.11	-0.16	-0.18
SE	0.08	0.10	0.14	0.16	0.18

Note: 1. * 90% significant, ** 95% significant, *** 99% significant. SE = standard error, Y = counterfactual mean outcome, ATE = average treatment effect. 2. The table presents the LTMLE estimates of counterfactual mean mortality risks for patients following the five dynamic treatment rules (*DTR1-5*). ATEs were estimated by comparing the counterfactual mean mortality risks of patients following a specific DTR to those following the static rule of always giving ESA (*Static1*) or to those following the static rule of never giving ESA (*Static0*), or by comparing the counterfactual mean mortality risks of patients following different DTRs.

Figure A1: Comparing counterfactual mean EQ-5D scores of patients following DTR2 and static treatment regimes

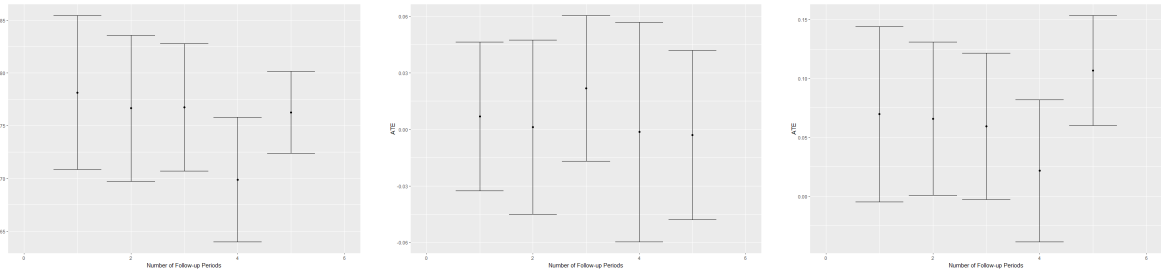

(a) *DTR2* (b) *DTR2* vs. *Static1* (c) *DTR2* vs. *Static0*

Figure A2: Comparing counterfactual mean EQ-5D scores of patients following DTR3 and static treatment regimes

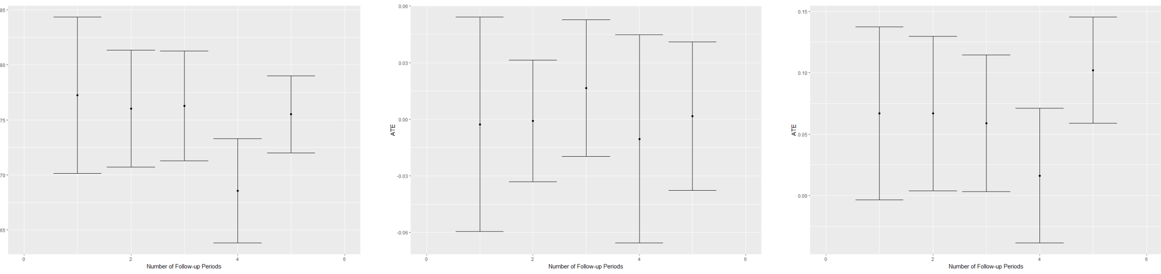

(a) *DTR3* (b) *DTR3* vs. *Static1* (c) *DTR3* vs. *Static0*

Figure A3: Comparing counterfactual mean EQ-5D scores of patients following DTR4 and static treatment regimes

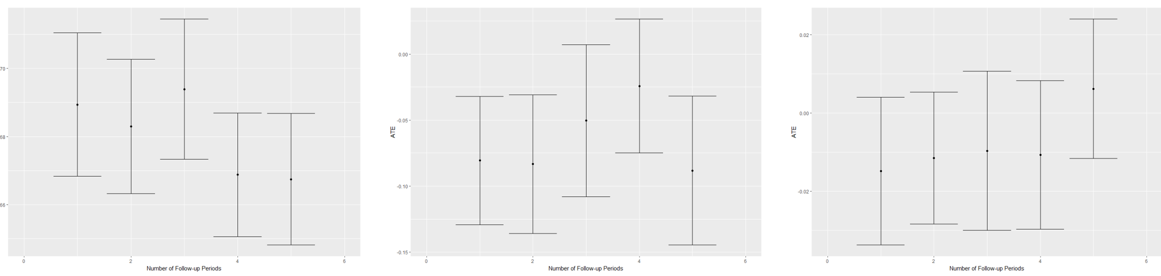

(a) *DTR4* (b) *DTR4* vs. *Static1* (c) *DTR4* vs. *Static0*

Figure A4: Comparing counterfactual mean EQ-5D scores of patients following DTR5 and static treatment regimes

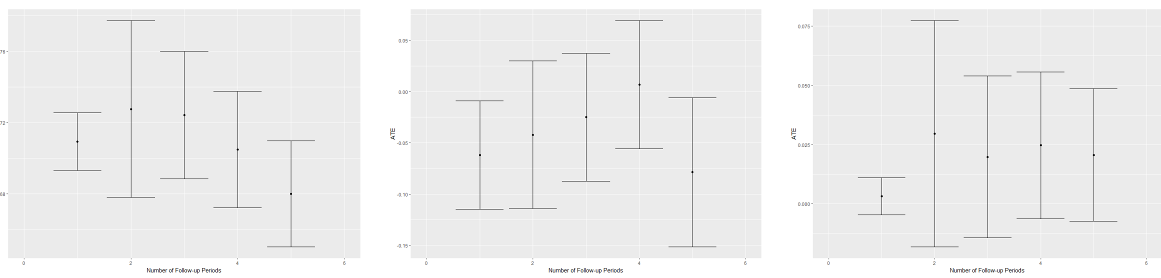

(a) *DTR5* (b) *DTR5* vs. *Static1* (c) *DTR5* vs. *Static0*

Figure A5: Comparing counterfactual mortality probabilities of patients following DTR2 and static treatment regimes

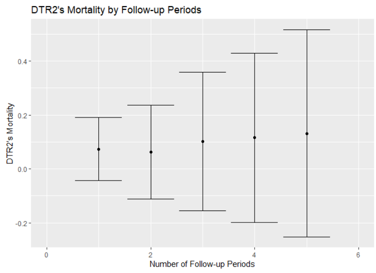

(a) *DTR2*

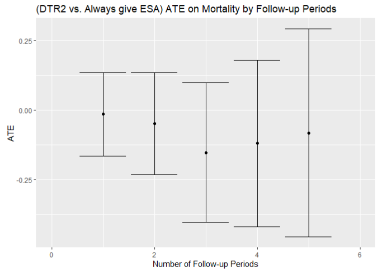

(b) *DTR2* vs. *Static1*

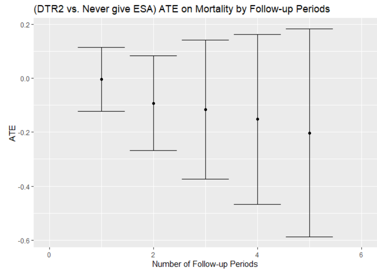

(c) *DTR2* vs. *Static0*

Figure A6: Comparing counterfactual mortality probabilities of patients following DTR3 and static treatment regimes

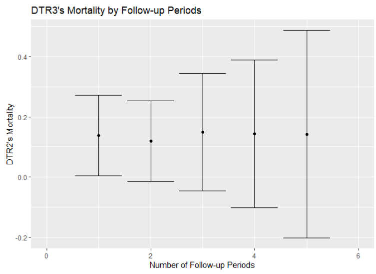

(a) *DTR3*

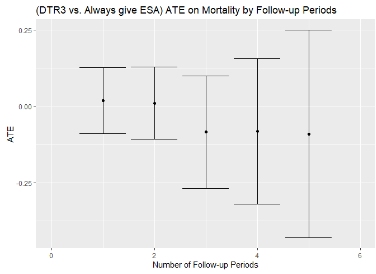

(b) *DTR3* vs. *Static1*

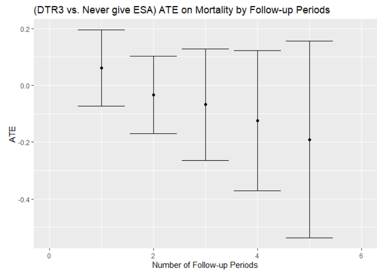

(c) *DTR3* vs. *Static0*

Figure A7: Comparing counterfactual mortality probabilities of patients following DTR4 and static treatment regimes

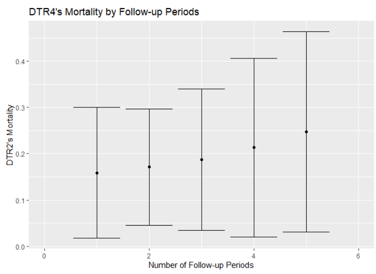

(a) *DTR4*

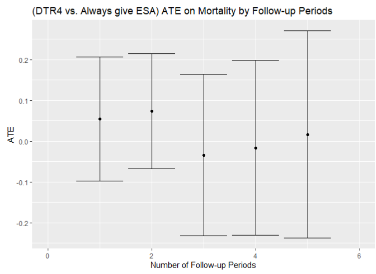

(b) *DTR4* vs. *Static1*

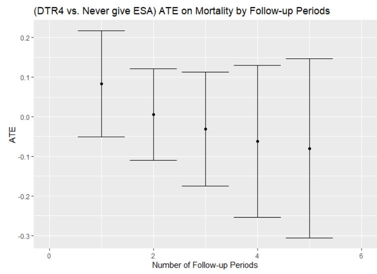

(c) *DTR4* vs. *Static0*

Figure A8: Comparing counterfactual mortality probabilities of patients following DTR5 and static treatment regimes

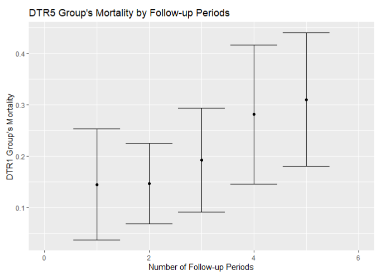

(a) *DTR5*

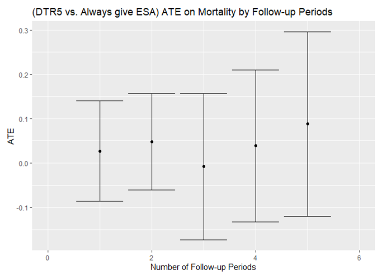

(b) *DTR5* vs. *Static1*

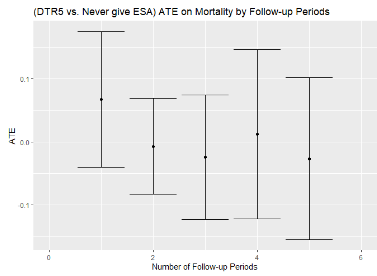

(c) *DTR5* vs. *Static0*