ABSTRACT

Background: Availability of detailed data from electronic health records (EHRs) has increased the potential to examine the comparative effectiveness of dynamic treatment strategies using observational data. Inverse probability (IP) weighting of dynamic marginal structural models can control for time-varying confounders. However, IP weights for continuous treatments may be sensitive to the choice of model.

Methods: We describe a target trial comparing strategies for treating anaemia with darbepoetin, in haemodialysis patients, using EHR data from the UK Renal Registry 2004-2016. Patients received a specified dose (mcg/week), or did not receive darbepoetin. We compare four methods to model time-varying treatment: (A) Logistic regression for zero dose and standard linear regression for log dose; (B) Logistic regression for zero dose and heteroscedastic linear regression for log dose; (C) Logistic regression for zero dose, heteroscedastic linear regression for log dose and multinomial regression for patients who recently received very low or high doses; (D) Ordinal logistic regression.

Results: For this dataset, method C was the only approach that provided a robust estimate of the mortality hazard ratio (HR), with less extreme weights in a fully weighted analysis and no substantial change of the HR point estimate after weight truncation. However, after truncating IP weights at the 95th percentile, estimates were similar across methods.

Conclusions: EHR data can be used to emulate target trials to estimate the comparative effectiveness of dynamic strategies that are sustained over time and adjust treatment to evolving patient characteristics. However, careful model checking, monitoring of large model weights, and adaptation of model strategies to account for these, is essential if an aspect of treatment is continuous.

Key words: Target trial, Continuous treatment, Observational data, Marginal structural model, Anaemia, Erythropoiesis stimulating agents, Haemodialysis

INTRODUCTION

Randomized controlled trials (RCT) are often preferred for establishing and estimating causal effects of health interventions on patient outcomes, but they can be expensive and time-consuming, and the questions of interest are limited by ethical considerations and the need for equipoise. There is increasing interest in using observational data from electronic health records (EHRs) to examine the comparative effectiveness of treatment strategies, but use of this data also poses challenges.[1] In clinical care, treatments received by a patient change over time, based on disease progression and response to prior treatment. Time-varying treatment can lead to the possibility of time-varying confounding, when a risk factor for the outcome also predicts subsequent treatment; and when past treatment predicts current risk factor levels When, in addition, past treatment predicts current risk factors ('treatment-confounder feedback'), effect estimates from conventional methods (e.g., Cox models conditioning on the time-varying confounders) may be biased.[2] Newer methods, including g methods[3] such as inverse probability (IP) weighting of marginal structural models (MSM)[4] and the g-formula,[5] can attempt to avoid this bias.

Erythropoiesis stimulating agents (ESAs) are used to correct and maintain Hb levels in chronic kidney disease (CKD) patients.[6] In UK clinical practice, ESA dosing decisions are based on regular (e.g., monthly) Hb measurements. A clinician reviews the Hb result and decides whether to alter the ESA dose. This is generally done without a written dosing protocol, and may also consider other clinical and laboratory variables, but the optimal Hb target is unknown. Large RCTs in patients with CKD not yet on dialysis found no evidence of benefit of a higher (compared with lower) Hb target for cardiovascular events, or for a composite outcome of death, myocardial infarction, hospitalization for congestive heart failure and stroke, but found increased risk of adverse events.[7-9] Observational studies suggest the best outcomes occur in patients who have high Hb concentrations but require only low doses of ESAs.[10-12] Adverse effects of higher Hb targets seen in RCTs may thus be due to high

ESA doses, particularly in patients with ESA resistance due to other underlying health problems. Hence, results of observational data analyses comparing outcomes when there is a cap on the maximum dose of ESAs might aid clinical decision-making.

'Dynamic' treatment strategies which are sustained over time and adapt treatment to the evolving characteristics of patients, can be assessed using g-methods to control for time-varying treatment-confounder feedback, allowing for valid effect estimation from EHR data. IP weighting, in particular, has become popular in pharmacoepidemiology[13, 14], but there is limited information on the use of these methods when treatment is not a binary decision. In many pharmacoepidemiology applications, patients receive or do not receive a specified drug treatment, with drug dose dependent on patient characteristics. When IP weights are based on models for continuous treatments, estimates of the comparative effectiveness of different treatment strategies may be sensitive to the choice of model.[15] Using a simulation study, Naimi et al.[16] explored different modelling approaches for constructing IP weights for continuous treatments and recommended an ordinal logistic regression approach (with "quantile binning").

In this paper, we describe a target trial comparing dynamic strategies for treatment of anaemia, using the ESA darbepoetin, in haemodialysis patients using EHR data. Patients were untreated or treated with a specified dose (mcg/week) of darbepoetin, depending on haemoglobin (Hb) target levels. The aim of the paper is to describe four models for time-varying treatment, and compare their performance. These models are: (A) Logistic regression models for zero dose and standard linear regression for log dose; (B) Logistic regression models for zero dose and heteroscedastic linear regression for log dose; (C) Logistic regression models for zero dose, heteroscedastic linear regression for log dose, and multinomial regression for patients who recently received very low or very high doses; (D) Ordinal logistic regression. We examine IP weights resulting from each method

and compare resulting estimates of the mortality hazard ratio (HR) of strategies with higher versus lower Hb target levels.

METHODS

Data

The UK Renal Registry (UKRR) collects clinical and biochemical EHR data from all patients receiving renal replacement therapy (RRT) in the UK. Data were extracted quarterly, with the last test result for that quarter recorded.[17] Estimating comparative effectiveness of different treatment strategies in observational studies requires careful measurement of and appropriate adjustment for confounding. The dataset analysed was based on all patients treated in the participating centres during specified periods, and contained information about every ESA prescription decision and the haemoglobin values that led to these, which should mitigate selection bias and lead to the findings being generalizable to other patients receiving haemodialysis. The UKRR obtained bespoke data extractions on haemodialysis patients from 10 centres, including the results of every test (Hb, ferritin, white blood count, albumin, c-reactive protein, urea reduction ratio) along with ESA dose, drug name and treatment date. All these variables are in the UKRR dataset and therefore covered by the Registry's permissions. Further information on the extracted data is in the Supplement.

Data on 8,131 adult (age ≥18 years) haemodialysis patients treated in UK renal centres between 2004 and 2016 were available for analysis. Of these, 6,773 (83.3%) were on darbepoetin at the start of their follow-up and 7,910 (97.3%) were treated with darbepoetin at some time during follow-up. Doses were predominantly in discrete categories (Figure 1, left panel). The likelihood of not receiving darbepoetin increased with patients' measured Hb at the previous visit (Figure 1, middle panel),

while for those on darbepoetin the median dose decreased with increasing values of measured Hb at the previous visit (Figure 1, right panel).

Research Ethics and Informed Consent

The processing of UK Renal Registry data for research has been approved by the NRES Committee North East - Newcastle & North Tyneside 1 Research Ethics Committee, reference 21/NE/0045. A waiver of consent for research purposes has been granted centrally by the Health Research Authority, reference 16/CAG/0064.

Design of the target trial

We designed a target trial to compare the effect on all-cause mortality of ESA treatment strategies based on specified Hb targets (low target range 95-115 g/L, high target range 105-125 g/L) among haemodialysis patients. For the full target trial definition and emulation, see Supplement eTable 1. Each strategy follows a protocol for dose change decisions based on current and previous ESA dose and current and previous Hb (Figure 2) and for acceptable dose changes based on these (Supplement eTable 2). Eligible patients were aged ≥18 years on haemodialysis for at least 3 months at one of 10 renal centres and either on darbepoetin, or not on darbepoetin with a Hb<110 g/L. Those who had a high darbepoetin dose (≥120 darbepoetin mcg/week) and low Hb (<80 g/L) at the time of first eligibility were excluded. Each strategy followed dosing rules based on current Hb, whether darbepoetin dose was changed in the previous month, and whether Hb changed in response to previous dosing (Supplement, Figure 2). Cessation of darbepoetin was permitted while patients' Hb was greater than the upper target for the assigned strategy (eTable 2). We allowed a 'grace period' of up to one month for dose changes to be implemented, when the dosing rule indicated needed changes.[1] Use of grace periods aligns with observed lags in dosing data, makes the strategies more

realistic, and minimizes censoring due to departures from assigned strategies. Randomisation of treatment assignment is emulated via a cloning procedure. Follow-up started after patients completed three months of haemodialysis at a contributing renal centre and ended eight months after baseline, death, or loss to follow-up, whichever happens first. The outcome was all-cause mortality, and the causal contrast of interest was the per protocol effect.

IP weighting of a dynamic marginal structural model

Unlike RCTs, treatment strategies are not explicitly assigned in observational studies, and comparisons rely on treatment received. Inappropriate analysis of observational data, when patient characteristics which vary after the start of follow are used to identify an individual's treatment strategy, can lead to immortal time bias. The "clone, censor and weight" approach proposed by Hernán,[18, 19] attempts to avoid this bias and was used in this study. Briefly, we copied ('cloned') all data for each patient and assigned one clone to each strategy (high versus low Hb target). Clones were censored when the patient's data became inconsistent with the clone's assigned strategy: (i) darbepoetin dose was changed but should have stayed constant; (ii) darbepoetin dose stayed constant but should have changed; or (iii) the darbepoetin dose was changed beyond the range of doses specified by the treatment strategy. The probability of a clone remaining uncensored at each time equals the probability of adhering to the assigned treatment strategy based on past covariate and treatment history. Therefore, models for treatment were used to derive the probability of being censored at each time point and this probability was used to derive IP weights.

Organising darbepoetin data and test results

We included new and established haemodialysis patients during any period between 2004 and 2016 when their treatment centre reported at least 60% of haemodialysis patients being treated with ESAs

(to minimize misclassification). Patients entered the study at the latest of the dates when their centre became eligible and when they first met the target trial eligibility criteria, at which point their follow-up time was set to 0. Data were formatted into discrete time intervals,[20] with one observation per person per 28-day ("month") period for the duration that they remained in the study. This structure allowed lagged variables (e.g., darbepoetin dose during the previous months) to predict subsequent values and ensured lagged values are comparable between patients. Further details of how the daily data for an individual patient was converted to monthly data is in the Supplement.

Notation

Let T_i denote the observed outcome time in months for patient *i* and let $A_i(t)$ denote treatment (darbepoetin dose, mcg/week) received by patient *i* in month *t* with $A_i(t)$=0 if patient *i* was not receiving darbepoetin in month *t*, and where non-zero doses were log-transformed to ensure the distribution was symmetrical. For lagged values of treatment from previous months, we used categories of darbepoetin dose (0, 0.1-20, 20.1-50, 50.1+ mcg/week), in order to allow zero doses in the lagged variable. The vector $L_i(t)$ represents the covariates in month *t* for patient *i*, including cubic splines for Hb (g/L) and lagged Hb from previous month. Further information on covariates is provided in the Supplement. $\bar{A}_i(t)$ denotes treatment history (the vector of darbepoetin dose values from baseline to month *t*) and the matrix $\bar{L}_i(t)$ denotes the history of time-varying covariates for patient *i*. We often suppress the *i* subscript denoting individual patient in the notation because we assume that the random vector for each subject is drawn independently from a distribution common to all subjects.

Models for treatment received

To construct the IP weights, the following 4 models for treatment received were fit to the original data, and then used to estimate treatment censoring weights for the cloned data. At each month, patients not currently treated with darbepoetin could remain off treatment or start darbepoetin, while those receiving treatment could stop treatment, or remain treated with the same or different dose.

For each Method, A-D, parameter estimates from models were used to calculate the probability of adhering to the assigned treatment strategy in each month, among patients (clones) who remained uncensored from treatment strategies. The cumulative probability of adhering to strategy *j* to the end of month *t* is the product of the probabilities of adhering to strategy *j* during each month from 1 to *t*. The IP weights [W(t)] were calculated as 1 / probability of adhering to strategy *j* to the end of month *t*.

Method (A): Logistic regression for zero dose and normal linear regression for log transformed dose

Method A used a two-step modelling process with:

1) logistic regression for the probability of not receiving darbepoetin each month, fit separately by darbepoetin treatment status in the previous month: $Pr(A(t) = 0|\bar{L}(t), \bar{A}(t-2), A(t-1) = 0)$ among those not receiving darbepoetin at *t-1*, and $Pr(A(t) = 0|\bar{L}(t), \bar{A}(t-1), A(t-1) > 0)$ among those receiving darbepoetin at *t-1*, with $A(t)$ coded as one during months on darbepoetin. We used cubic splines for months since baseline to model changes in the probability of not receiving darbepoetin since start of follow up. The models also included current values of covariates at month *t*, and lagged values of Hb (at *t*-1), and treatment (at months *t*-2 and *t*-3).

2) Linear regression model assuming normally-distributed (Gaussian) error terms for the probability density of the log of darbepoetin dose each month, among those receiving darbepoetin at month *t-1* : $f(A(t)|\bar{L}(t), \bar{A}(t-1), A(t-1) > 0)$.

To estimate the IP weights, we defined $R_l(j,t)$ and $R_u(j,t)$ to be the lower and upper limits of the range of acceptable non-zero doses, for strategy *j*=1 or 2 and month *t*, according to the dosing rules (Figure 2 and eTable 2). For example, if a patient was on a dose of 40 mcg/week and the protocol recommended an increase, the acceptable dose range would be 40.1-80 mcg/week. If the protocol specified no change, we extended the acceptable range around the current dose to avoid arbitrarily small probabilities. For example, for a dose of 40 mcg/week, we calculated the probability of having a dose between 35 and 45 mcg/week. The probabilities of zero and non-zero dose were calculated from logistic regression models in step 1 and the probability of prescribed dose being within an acceptable dose range was obtained from the linear model in step 2. We estimated the mean $\mu_{\bar{L}(t),\bar{A}(t-1)} = E[A|\bar{L}(t), \bar{A}(t-1)]$ and estimated the constant variance σ^2 as the root mean square error. The overall probability of adhering to the assigned treatment strategy in the current month was the combined probability of non-zero dose (from step 1) and the probability of dose being in the acceptable range (from step 2). I.e., the probability of adhering to strategy *j* is:

$$\Pr(A(t) = 0|\bar{L}(t), \bar{A}(t-1)) + \Pr(A(t) > 0|\bar{L}(t), \bar{A}(t-1)) \int_{R^l(t,j)}^{R^u(t,j)} f(A(t)|\,\bar{L}(t), \bar{A}(t-1))$$

if a zero dose is acceptable, and

$$\Pr(A(t) > 0|\bar{L}(t), \bar{A}(t-1)) \int_{R^l(t,j)}^{R^u(t,j)} f(A(t)|\,\bar{L}(t), \bar{A}(t-1))$$

if a zero dose is not acceptable.

Method (B): Logistic regression for zero dose and heteroscedastic linear regression for log dose

Method B used the logistic regression models from Method A, but replaced the linear model with a multiplicative heteroscedastic linear regression of log-dose, modelling variance as an exponential function of selected covariates and previous darbepoetin dose. We estimated the mean $\mu_{\bar{L}(t),\bar{A}(t-1)} = E[A|\bar{L}(t), \bar{A}(t-1)]$ and variance $\sigma^2_{\bar{L}(t),\bar{A}(t-1)}$ for all combinations of $\bar{L}(t), \bar{A}(t-1)$. Probabilities were calculated using these predicted mean and standard deviation estimates as per Method A.

Method C: Logistics regression for zero dose, heteroscedastic linear regression for log dose and multinomial regression for coming from very low and very high doses

Method C combined the probabilities estimated from the logistic regression models from Method A and heteroscedastic linear regression model of log-dose from Method B with a multinomial logistic regression model for the change in dose, stratified by extreme dose levels (2.5, 5, 120 and 150 mcg/week). We coded 8 mutually-exclusive groups *G* for an individual's dose change: (1) go off darbepoetin (i.e. move to zero dose), (2) unacceptable decrease in dose for both strategies (i.e. the patient's darbepoetin dose was lowered when the protocol for both strategies said the dose should be constant or increased), (3) acceptable decrease for low Hb strategy only (for when the protocol would only say to lower the darbepoetin dose for the low Hb strategy), (4) acceptable decrease both strategies, (5) keep the darbepoetin dose constant, (6) acceptable increase for both strategies, (7) acceptable increase for high Hb strategy only, and (8) unacceptable increase for both strategies. Further details are shown in eTable 3.

We used multinomial logistic regression to model the probability of a suitable dose for each stratum of non-zero dose at month *t-1*, and calculated the probability of adhering to each treatment strategy in the current month by summing the predicted probabilities of the appropriate acceptable dose

changes: $\sum_1^S \Pr\left(A(t) = g|\bar{L}(t), A(t-1) = D\right)$ where *D* = 2.5, 5, 120 or 150 and *S* is the vector of suitable dose changes (1, ... S) for a given individual in month *t*. For example, from a dose of 2.5 mcg/week, if protocol said a patient on the low Hb strategy should increase their dose, we used the probability of 'Acceptable increase for both strategies'. For a patient following the high Hb strategy, if the protocol recommended dose increase, we summed the probabilities of 'Acceptable increase for both strategies' and 'Accept increase for high Hb strategy only'. The most frequently occurring category, 'Keep the darbepoetin dose the same', was used as the reference group. Because some groups had small numbers, Hb and month (cubic splines) were the only covariates included in the multinomial models.

Method D: Ordinal logistic regression for all dose levels

Finally, in Method D, we transformed dose into an ordinal variable *V*, with 17 levels (coded 0-16) to represent the dosing ladder: 0, 2.5, 5, 10, 15, 20, 25, 30, 40, 50, 60, 70, 80, 100, 120, 150, 180 mcg/week. Doses between ladder rungs were coded as the higher rung e.g., a dose of 35 was coded as 40 mcg/week. We used an ordinal logistic regression model to estimate the probability of each rung, given prior dose and covariates. For each strategy *j*=1 or 2 and month *t*, we defined $O_l(j,t)$ and $O_u(j,t)$ to be the lower and upper limits of the range of acceptable doses from the dosing ladder, including zero dose, according to the dosing rules (Figure 2 and eTable 4). To obtain the probability of prescribed dose being within an acceptable dose range in the current month, we combined the probabilities of each dose within the range:

$$\sum_{O_l}^{O_u} \Pr\left(A(t) = v|\bar{L}(t), A(t-1)\right).$$

Comparative effectiveness estimation

We created IP censoring weights for withdrawal from the target trial if a patient (1) changed from haemodialysis to peritoneal dialysis, (2) had a kidney transplant or (3) was lost to follow-up (Supplement). For all methods above, treatment and withdrawal weights were multiplied together to give a final model weight. Final weights were also truncated at the 90th, 95th and 99th percentile of the weight distribution to mitigate the impact of large weights.

We fit a pooled logistic regression model to the cloned data to estimate the HR for all-cause mortality comparing high versus low Hb targets (reference group). The outcome model included cubic splines for month and used robust standard errors for clustering by patient. We also conducted an unadjusted analysis by fitting an unweighted pooled logistic regression model for mortality after censoring, using only the terms for treatment strategy and time, and sensitivity analyses omitting withdrawal weights.

RESULTS

A total of 8,131 patients met the eligibility criteria and were included in the analyses. There were 355 deaths from 38,337 patient months in the lower Hb strategy and 303 deaths from 37,422 months in the higher Hb strategy. In an unweighted analysis, the estimated mortality HR comparing the higher to the lower Hb strategy was 0.87 (95% CI 0.81, 0.94).

The linear regression model for log-dose in Method A showed evidence of heteroscedasticity: the variance of the residuals varied with predicted dose (eFigure 1, Breusch-Pagan test p<0.001). Method A also resulted in some very large IP weights (>9999) due to high (>50) monthly treatment weights (Table 1). Large weights were most common when patients' prior darbepoetin dose was at an extreme of the dose distribution (Figure 3A), suggesting the model did not predict well for patients with very small or very large prior darbepoetin doses. The 99th, 95th and 90th percentile of the final

weights were: 133.2, 13.0 and 6.7, respectively (Table 1). The fully weighted model was influenced by extreme weights, giving an estimate of the HR of 0.09 (95% CI 0.01, 0.56). After truncating weights at the 99th percentile (113.2), the estimate HR was 0.96 (95% CI 0.81, 1.13).

Method B accounted for heteroscedasticity, and compared to Method A, the 99th, 95th and 90th percentile of the final weights were reduced slightly: 87.1, 11.0 and 5.9, respectively (Table 1). However, final model weights still showed extreme values (>9999) stemming from months with very low or very high prior doses (Figure 3B). The estimated of HR for Method B from the weighted analysis was: 0.05 (95% 0.01, 0.37). After truncating weights at the 99th percentile (87.1), the HR was 0.95 (95% CI 0.81, 1.13).

Method C adapted Method B to deal with large weights by using multinomial models for extreme dosages (Figure 3C). The 99th, 95th and 90th percentile of the overall (cumulative) weights from Method C were: 23.8, 8.6 and 5.3, respectively, and the estimated HR for the weighted analysis was 0.94 (95% CI 0.76, 1.15). This estimated HR was relatively similar after truncating the weights at the 99th percentile (23.8), although CIs were a little narrower: 0.91 (95% CI 0.80, 1.04). The sensitivity analysis, fitting Method C without withdrawal weights found a very similar estimated HR: HR 0.94 (95% CI 0.76, 1.16).

Finally, Method D used an ordinal logistic approach across the levels of dose, but the proportional odds assumption was not met (Brant[21] test p<0.001). The 99th, 95th and 90th percentile of the overall (cumulative) weights using Method D were: 387.7, 33.5 and 14.4, respectively, and the fully weighted analysis appeared influenced by extreme weights (>9999): HR 0.22 (95% CI 0.05, 0.94).

After truncating the model weights at the 95th percentile, estimated HRs were similar across all methods (Table 1). The estimated weighted survival curves for the four methods are presented in eFigure 2. After truncating weights at the 95th percentile, curves were similar across all methods. For completeness, the estimated unweighted survival curve is presented in eFigure 3. Based on the

improved model fit and distribution of weights, we concluded that Method C was the best IP weighting strategy for this dataset and treatment setting.

DISCUSSION

We used observational EHR data obtained from the UKRR to emulate a target trial estimating the effects of higher versus lower target Hb strategies on all-cause mortality in haemodialysis patients, taking account of time-varying confounding by Hb levels. We compared different modelling approaches for deriving IP treatment weights for medication use and dosage (ESA dose, if prescribed) and determined that a flexible modelling approach provided the most robust results in this dataset. Method C had less extreme weights in a fully weighted analysis and was the only approach where there was no substantial change of the HR point estimate after weight truncation. The goal of weight truncation is to improve the variance at the cost of potentially introducing a small amount of residual confounding.[22] Ideally, we expect to see truncation narrow the CI without affecting the point estimate substantially - this was the case for Method C. For the other methods there was a large difference in the HR point estimates between the untruncated and truncated results, and this suggests that the weights were not performing well. Our final weights modelling approach combined logistic regression models for zero dose, heteroscedastic linear regression for log-dose, and multinomial models for extreme doses. However, after truncating the model weights at the 95th percentile, estimates of HRs were similar regardless of method.

Emulating a target trial with detailed observational data has several strengths. Explicitly specifying the protocol of the target trial and describing how to emulate it with observational data ensures synchronization of eligibility and treatment assignment with time zero. This can prevent[23] prevalent user bias,[24] which could occur if prior successful treatment reduces the risk of adverse events following future treatment, and immortal time bias,[18] which can occur if treatment categorization

depends on survival time. The UKRR is a large and highly representative database allowing trends in clinical practice patterns to be captured.[12] The bespoke extract for this project allowed us to carry out an in-depth analysis of observational data. Together, the data and study design allow the estimation of the causal effect of higher versus lower Hb targets outside of an RCT.

Our study also has several limitations. First, some renal centres do not routinely record computerised data on darbepoetin dose or drug type, limiting the number of centres we could include. To minimize inaccurate dose information, we restricted to centres reporting at least 60% of haemodialysis patients being treated with ESAs. Second, the use of a dynamic treatment strategy, while more realistic, meant that we could not use stabilized IP weights that are commonly used for static regimes. Alternative numerators for the weights for dynamic regimes aren't guaranteed to produce estimates that are less variable than those obtained using unstabilized weights.[25] As a result we had some large IP weights in all methods and our variance was high. Confidence intervals were narrower with truncated weights, but truncation can re-introduce some residual confounding. Finally, although we used statistical methods that appropriately control for measured time updated confounders, no observational study can exclude the possibility of unmeasured confounding. However, because Hb results are the main clinical decision factor in ESA dosing, we anticipate that our dataset captured the most important source of confounding. High blood pressure may have led to some ESA doses being delayed or omitted. Because only prescribed doses were recorded in the electronic health record and we did not have data on blood pressure, we were unable to adjust for this.

We encountered some novel methodological challenges when attempting to design a clinically-relevant and computationally-feasible target trial for emulation with the observational data. The main steps we followed in our analysis are shown in Figure 4. Our first attempt at a target trial design aimed to compare restrictive with liberal dosing strategy to achieve a standard Hb target with

darbepoetin doses decreased or increased according to a strict dosing ladder, but this resulted in patients being rapidly censored due to treatments inconsistent with assigned strategies. Whilst an RCT comparing these strategies might be feasible, observational emulation requires that patients actually followed the two dosing strategies being assessed in clinical practice. Our final target trial design reflected more closely treatment policies used by clinicians during the follow up period. A second challenge was the requirement modelling a continuous treatment. Most prior implementations of IP weighting of MSMs have relied on dichotomous treatment strategies, with little guidance available on modelling more complex treatment strategies. Positivity is an important assumption when carrying out IP weighting; if the treatment of interest is binary, e.g., medicine use, the assumption is that there are both treated and untreated individuals at each level of the combination of covariates. When treatment is continuous e.g., dose of a medicine, it must be possible to receive every level of the dose at each level of the combination of covariates. Near violations of this assumption can lead to extreme IP weights. An alternative analytic approach would be to use the g-formula, which is more robust to sparse data and could reduce the difficulties in estimating continuous treatments, but which is less widely known in pharmacoepidemiology. We advise that others who wish to use these methods for similar applications and future directions use large, rich datasets that include detailed information about the treatment and covariates at baseline and throughout follow-up. Careful checking of model assumptions, and examination of the distribution of inverse probability weights, is essential when making causal inferences about the comparative effectiveness of dynamic treatment strategies based on observational data.

Our study demonstrates the emulation of a target trial with data from EHRs to estimate comparative effectiveness of dynamic strategies which are sustained over time and adjust treatment to evolving characteristics of patients. However, careful model checking, monitoring of large model weights, and

adaptation of modelling strategies to account for these, is essential whenever treatment is continuous.

References

1. Hernan MA and Robins JM. Using Big Data to Emulate a Target Trial When a Randomized Trial Is Not Available. *Am J Epidemiol* 2016; 183: 758-764. 2016/03/20. DOI: 10.1093/aje/kwv254.
2. Robins JM. Causal inference from complex longitudinal data. In: Berkane M (ed) *Latent Variable Modeling and Applications to Causality: Lecture Notes in Statistics (120)*. New York: Springer-Verlag, 1997, pp.69-117.
3. Naimi AI, Cole SR and Kennedy EH. An introduction to g methods. *Int J Epidemiol* 2017; 46: 756-762. 2017/01/01. DOI: 10.1093/ije/dyw323.
4. Robins JM, Hernan MA and Brumback B. Marginal structural models and causal inference in epidemiology. *Epidemiology* 2000; 11: 550-560.
5. Keil AP, Edwards JK, Richardson DB, et al. The parametric g-formula for time-to-event data: intuition and a worked example. *Epidemiology* 2014; 25: 889-897. 2014/08/21. DOI: 10.1097/EDE.0000000000000160.
6. Mikhail A, Shrivastava R and Richardson D. Clinical Pracitice Guidelines. *Anaemia of CKD*. 5th 2009-2012 ed. www.renal.org/guidelines: UK Renal Association, 2010.
7. Drueke TB, Locatelli F, Clyne N, et al. Normalization of hemoglobin level in patients with chronic kidney disease and anemia. *N Engl J Med* 2006; 355: 2071-2084. 355/20/2071 pii ;10.1056/NEJMoa062276 doi.
8. Pfeffer MA, Burdmann EA, Chen CY, et al. A trial of darbepoetin alfa in type 2 diabetes and chronic kidney disease. *N Engl J Med* 2009; 361: 2019-2032. NEJMoa0907845 pii ;10.1056/NEJMoa0907845 doi.
9. Singh AK, Szczech L, Tang KL, et al. Correction of anemia with epoetin alfa in chronic kidney disease. *N Engl J Med* 2006; 355: 2085-2098. 355/20/2085 pii ;10.1056/NEJMoa065485 doi.
10. Levin A, Djurdjev O, Duncan J, et al. Haemoglobin at time of referral prior to dialysis predicts survival: an association of haemoglobin with long-term outcomes. *Nephrol Dial Transplant* 2006; 21: 370-377. gfi209 pii ;10.1093/ndt/gfi209 doi.
11. Regidor DL, Kopple JD, Kovesdy CP, et al. Associations between changes in hemoglobin and administered erythropoiesis-stimulating agent and survival in hemodialysis patients. *J Am Soc Nephrol* 2006; 17: 1181-1191. 17/4/1181 pii ;10.1681/ASN.2005090997 doi.
12. Macdougall IC, Tomson CR, Steenkamp M, et al. Relative risk of death in UK haemodialysis patients in relation to achieved haemoglobin from 1999 to 2005: an observational study using UK Renal Registry data incorporating 30,040 patient-years of follow-up. *Nephrol Dial Transplant* 2010; 25: 914-919. gfp550 pii ;10.1093/ndt/gfp550 doi.
13. Xie Y, Bowe B, Gibson AK, et al. Comparative Effectiveness of Sodium-Glucose Cotransporter 2 Inhibitors vs Sulfonylureas in Patients With Type 2 Diabetes. *JAMA Intern Med* 2021; 181: 1043-1053. DOI: 10.1001/jamainternmed.2021.2488.
14. Pawar A, Desai RJ, Gautam N, et al. Risk of admission to hospital for serious infection after initiating tofacitinib versus biologic DMARDs in patients with rheumatoid arthritis: a multidatabase cohort study. *The Lancet Rheumatology* 2020; 2: e84-e98. DOI: 10.1016/S2665-9913(19)30137-7.
15. Hernán MA and Robins JM. *Causal Inference: What If*. Boca Raton: Chapman & Hall/CRC, 2020.
16. Naimi AI, Moodie EE, Auger N, et al. Constructing inverse probability weights for continuous exposures: a comparison of methods. *Epidemiology* 2014; 25: 292-299. DOI: 10.1097/EDE.0000000000000053.
17. Ansell D and Tomson CR. UK Renal Registry 11th Annual Report (December 2008): Chapter 15 The UK Renal Registry, UKRR database, validation and methodology. *Nephron Clin Pract* 2009; 111 Suppl 1: c277-285. DOI: 10.1159/000210004.
18. Hernan MA, Sauer BC, Hernandez-Diaz S, et al. Specifying a target trial prevents immortal time bias and other self-inflicted injuries in observational analyses. *J Clin Epidemiol* 2016; 79: 70-75. 2016/05/31. DOI: 10.1016/j.jclinepi.2016.04.014.

19. Hernan MA. How to estimate the effect of treatment duration on survival outcomes using observational data. *BMJ* 2018; 360: k182. 2018/02/09. DOI: 10.1136/bmj.k182.
20. Fewell Z, Hernán MA, Wolfe F, et al. Controlling for Time-dependent Confounding using Marginal Structural Models. *The Stata Journal* 2004; 4: 402-420. DOI: 10.1177/1536867x0400400403.
21. Brant R. Assessing proportionality in the proportional odds model for ordinal logistic regression. *Biometrics* 1990; 46: 1171-1178.
22. Cole SR and Hernan MA. Constructing inverse probability weights for marginal structural models. *Am J Epidemiol* 2008; 168: 656-664. 20080805. DOI: 10.1093/aje/kwn164.
23. Dickerman BA, Garcia-Albeniz X, Logan RW, et al. Avoidable flaws in observational analyses: an application to statins and cancer. *Nat Med* 2019; 25: 1601-1606. 2019/10/09. DOI: 10.1038/s41591-019-0597-x.
24. Ray WA. Evaluating Medication Effects Outside of Clinical Trials: New-User Designs. *American Journal of Epidemiology* 2003; 158: 915-920. DOI: 10.1093/aje/kwg231.
25. Cain LE, Robins JM, Lanoy E, et al. When to start treatment? A systematic approach to the comparison of dynamic regimes using observational data. *Int J Biostat* 2010; 6: Article 18. DOI: 10.2202/1557-4679.1212.

Table 1. Results from the different modelling approaches estimating the hazard ratio for all-cause mortality comparing high versus low haemoglobin strategies.

Method	Description	Truncation of weights	Distribution of weights					Hazard ratio	95% CI
			Median (IQR)	90th pct	95th pct	99th pct	Max		
A	Logistic regression models for zero dose and **normal linear regression** for log dose	Full	1.6 (1.7)	6.7	13.0	113.2	>9999	0.09	(0.01, 0.56)
		99th pct	1.6 (1.7)	6.7	13.0	113.2	113.2	0.96	(0.81, 1.13)
		95th pct	1.6 (1.7)	6.7	13.0	13.0	13.0	0.92	(0.83, 1.03)
		90th pct	1.6 (1.7)	6.7	6.7	6.7	6.7	0.90	(0.82, 0.99)
B	Logistic regression models for zero dose and **heteroscedastic linear regression** for log dose	Full	1.5 (1.6)	5.9	11.0	87.1	>9999	0.05	(0.01, 0.37)
		99th pct	1.5 (1.6)	5.9	11.0	87.1	87.1	0.96	(0.81, 1.13)
		95th pct	1.5 (1.6)	5.9	11.0	11.0	11.0	0.93	(0.84, 1.04)
		90th pct	1.5 (1.6)	5.9	5.9	5.9	5.9	0.91	(0.83, 1.00)
C	Logistic regression models for zero dose, **heteroscedastic linear regression** for log dose and **multinomial regression** for coming from very low and very high doses	Full	1.5 (1.5)	5.3	8.6	23.8	793.7	0.94	(0.76, 1.15)
		99th pct	1.5 (1.5)	5.3	8.6	23.8	23.8	0.91	(0.80, 1.04)
		95th pct	1.5 (1.5)	5.3	8.6	8.6	8.6	0.91	(0.82, 1.01)
		90th pct	1.5 (1.5)	5.3	5.3	5.3	5.3	0.90	(0.82, 0.99)
D	**Ordinal regression** model for all levels of dose	Full	1.8 (3.7)	14.4	33.5	387.7	>9999	0.22	(0.05, 0.94)
		99th pct	1.8 (3.7)	14.4	33.5	387.7	387.7	1.07	(0.83, 1.37)
		95th pct	1.8 (3.7)	14.4	33.5	33.5	33.5	0.99	(0.87, 1.13)
		90th pct	1.8 (3.7)	14.4	14.4	14.4	14.4	0.93	(0.83, 1.03)

Notes: Full weights means that no truncation has taken place. 99th pct means the weights have been truncated at the 99th percentile. A pooled logistic regression model was fitted to the cloned data to estimate the hazard ratio for all-cause mortality comparing high versus low Hb strategy (reference group).

Figure 1. Distribution of weekly darbepoetin dose (mcg/week) (left panel, A), probability of zero darbepoetin dose by previous Hb level (g/L) (middle panel, B), and darbepoetin dose by previous Hb level (right panel, C).

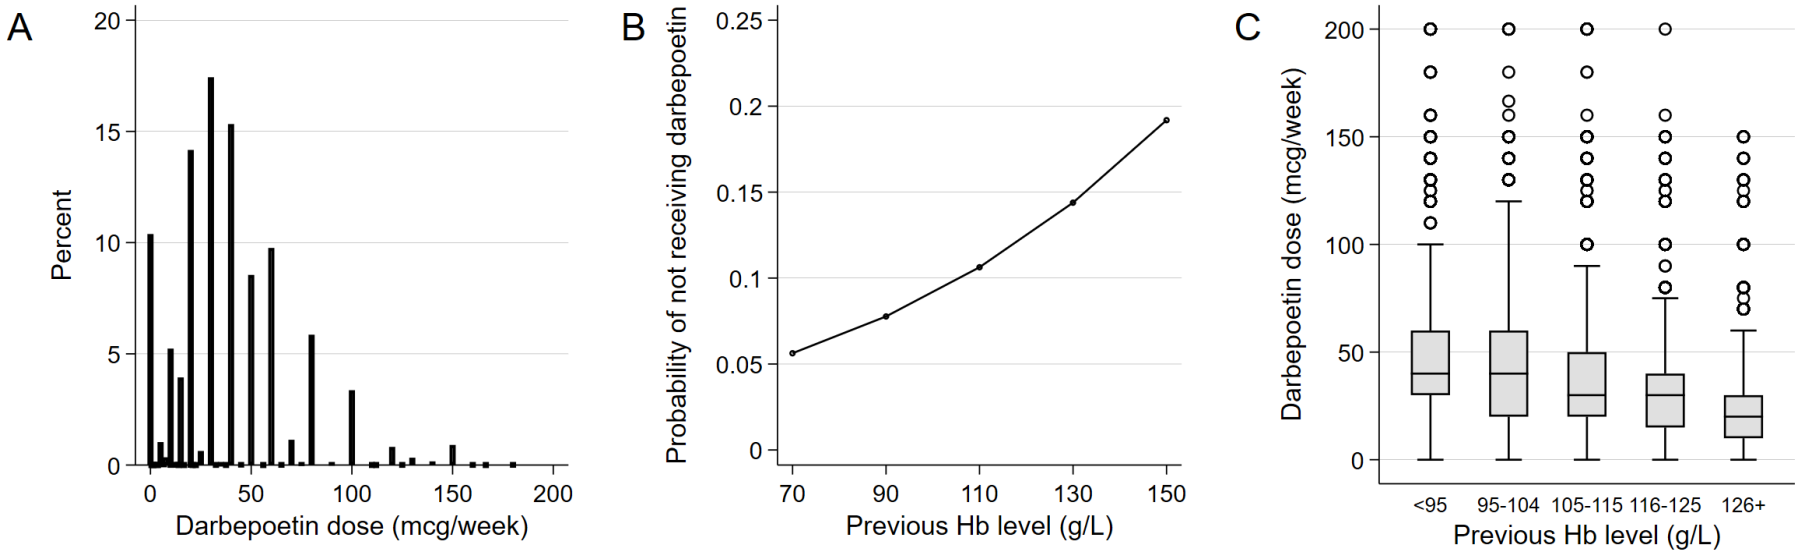

Figure 2. Dose change decisions protocol

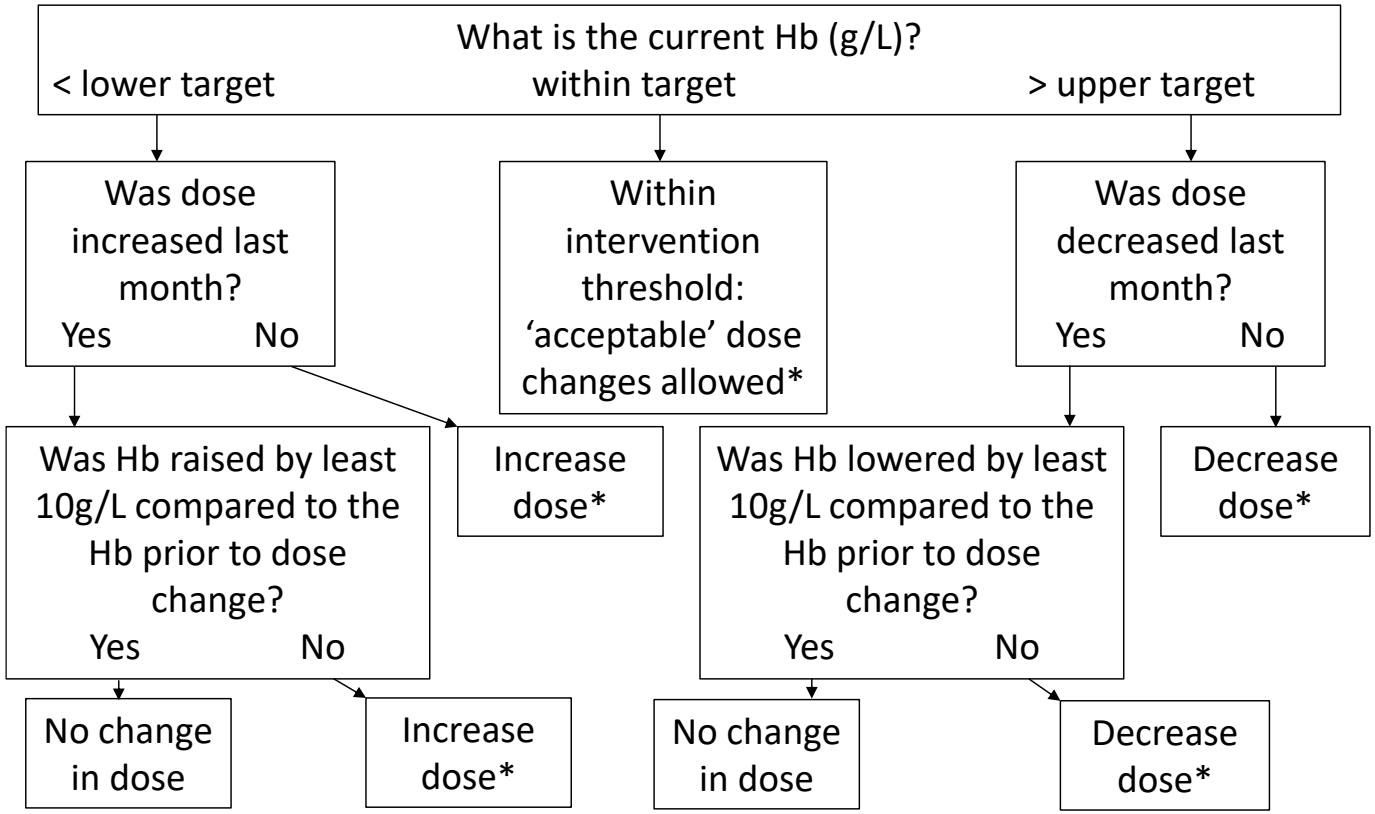

* Dose changes need to be within acceptable level, see Table A2

Figure 3. Monthly treatment weight for the low Hb strategy by previous darbepoetin dose, for Methods A-D

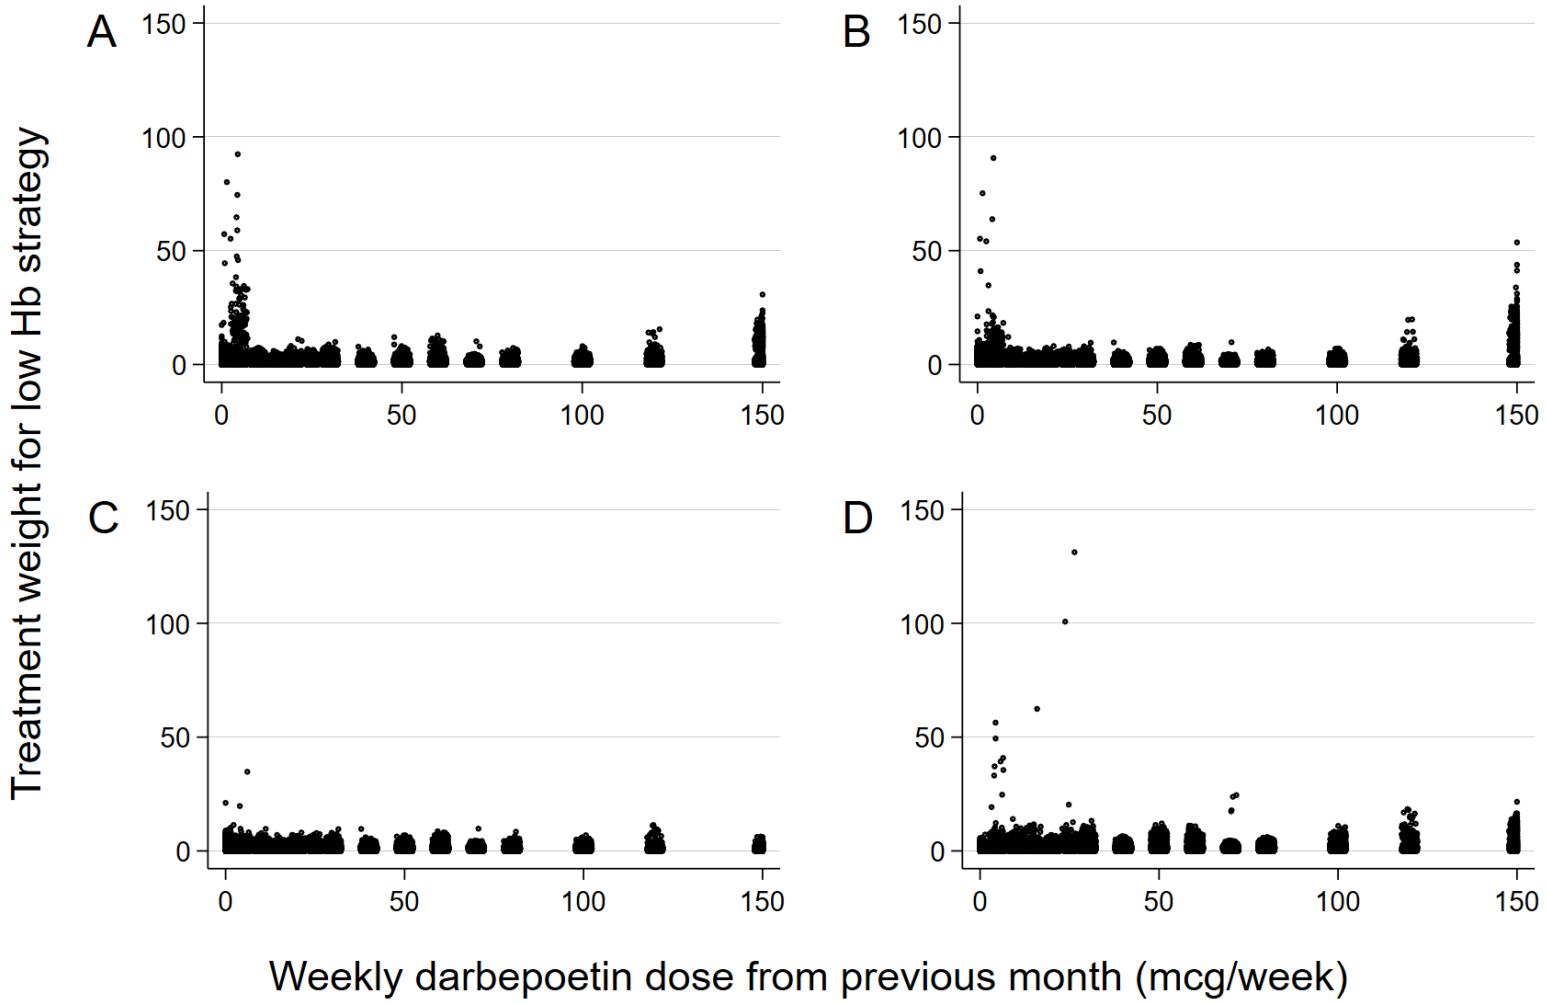

Figure 4. The process followed to arrive at the final estimate of the causal effect of the treatment strategies compared

	Design the target trial.
	Check data contain observations compatible with both arms of the emulated trial.
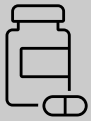	Derive models for treatment and the treatment weights. Check model assumptions.
	Derive models for censoring and the censoring weights, if needed.
	Derive final inverse probability (IP) weights. Assess the reason for any large IPW weights. Calculate truncated weights if some weights are large.
	Fit the final marginal structural model on cloned data.