

HEDG

HEALTH, ECONOMETRICS AND DATA GROUP

THE UNIVERSITY *of York*

WP 14/19

Health econometric evaluation of the effects of a continuous treatment: a machine learning approach

Noémi Kreif ; Iván Díaz; Richard Grieve & David Harrison

August 2014

york.ac.uk/res/herc/hedgwp

Health econometric evaluation of the effects of a continuous treatment: a machine learning approach

Noémi Kreif ^{*1}, Richard Grieve¹, Iván Díaz² and David Harrison³

¹Department of Health Services Research and Policy, London School of Hygiene and Tropical Medicine, London, UK

²Department of Biostatistics, Johns Hopkins Bloomberg School of Public Health, Baltimore, MD

³Intensive Care National Audit & Research Centre (ICNARC), London, UK

July 17, 2014

Abstract

When the treatment under evaluation is continuous rather than binary, the marginal causal effect can be reported from the estimated dose-response function. Here, regression methods can be employed that specify a model for the endpoint, given the treatment and covariates. An alternative is to estimate the generalised propensity score (GPS), which can adjust by the conditional density of the treatment, given the covariates. With either regression or GPS approaches, model misspecification can lead to biased estimates. This paper introduces a machine learning approach, the “Super Learner”, to estimate both the GPS and the dose-response function. The Super Learner selects the convex combination of candidate estimation algorithms, to create new estimators. We take a two stage estimation approach whereby the Super Learner selects a GPS, and then a dose-response function conditional on the GPS. We compare this approach to parametric implementations of the GPS and regression methods.

We contrast the methods in the Risk Adjustment In Neurocritical care (RAIN) cohort study, in which we estimate the marginal causal effects of increasing transfer time from emergency departments to specialised neuroscience centres, for patients with traumatic brain injury. With parametric models for the outcome we find that dose-response curves differ according to choice of parametric specification. With the Super Learner approach to both regression and the GPS, we find that transfer time does not have a statistically significant marginal effect on the outcome.

Keywords: program evaluation; generalised propensity score; machine learning.

JEL Classification Numbers: C1, C5.

*Corresponding author; email: noemi.kreif@lshtm.ac.uk. The authors are grateful to Prof. Kathy Rowan (ICNARC) for data access, Zia Sadique (LSHTM) and Manuel Gomes (LSHTM) for advice on data analysis, to Prof. David Menon (University of Cambridge) and Prof. Martin Smith (UCLH) for clinical advice on the empirical study. This is an independent report commissioned and funded by the Policy Research Programme in the Department of Health to the Policy Innovation Research Unit (PIRU) at LSHTM. The views expressed are not necessarily those of the Department of Health. We thank seminar participants of the Health, Econometrics and Data Group (University of York) for helpful suggestions. We are grateful to Eric Polley for help in using the Super Learner.

1 Introduction

Public policy-makers may be interested in evaluations that estimate the causal effects of treatments measured on a continuous scale. For example, evaluations may attempt to estimate the marginal causal effects of alternative financial incentives for health care providers, different levels of taxation on addictive substances, and increasing doses of a new pharmaceutical. In such settings, the expected outcome at alternative levels of the treatment can be reported from the estimated dose-response function. When estimating such dose-response functions, the key identifying assumption is that all systematic differences between units that received different levels of the treatment, in both observed and unobserved variables that explain the outcome, have been adjusted for. Such adjustment can be performed with regression, which models the outcome as a function of the treatment variable and observed covariates. Regression approaches can be simple, for example ordinary least squares regression, or can take more flexible forms such as fractional polynomials (Royston and Altman, 1994; Royston et al., 2006) or generalised additive models. Here a general concern is whether the regression model is correctly specified (Imai and Van Dyk, 2004).

The generalised propensity score (GPS) approach has been proposed (Imbens, 2000; Hirano and Imbens, 2004) as an alternative to regression for evaluating continuous or multi-valued treatments. Both the standard regression method and the GPS approach assume "unconfoundedness", or that adjusting for observed covariates is sufficient to achieve independence between potential outcomes and the treatment level received. As with the propensity score for binary treatments, rather than adjusting for a vector of covariates, the GPS adjusts for a one-dimensional score, the conditional density of treatment, given baseline covariates. Hirano and Imbens (2004) proposed the GPS for inclusion in a regression model of the outcome as function of the treatment and the GPS. Compared to the multivariable regression approach, this method offers the potential advantage of having to specify the outcome model as a function of only two covariates. This approach has been followed by subsequent empirical work (Bia and Mattei, 2008; Bia et al., 2011; Kluve et al., 2012). Alternative implementations of the GPS include a kernel weighting approach proposed by Flores et al. (2012), and using the GPS for inverse weighting of marginal structural models (Robins et al., 2000). Imai

and Van Dyk (2004) introduced a similar concept, the "propensity function" for estimating the average treatment effect in strata defined by the propensity function. Yang et al. (2014) develop methods of matching and classification on the GPS, to estimate the effects of multi-valued treatments. In this paper we consider further the approach proposed by Hirano and Imbens (2004), for continuous treatments.

A general challenge with the GPS approaches proposed is the specification of the outcome and GPS models. For the outcome model, previous studies have proposed using either fully parametric models with polynomials (Hirano and Imbens, 2004; Bia and Mattei, 2008), or semi-parametric approaches using regression splines (Kluve et al., 2012; Bia et al., 2011). Bia et al. (2011) compared parametric and semiparametric estimators of the outcome model, and found that the estimated dose-response function was robust to the semiparametric approach used, while it was sensitive to parametric specification. For the estimation of the GPS, most studies have assumed a normal or lognormal distribution, and suggested indirect tests of the balancing property (Hirano and Imbens, 2004; Imai and Van Dyk, 2004) to assess model specification. Evidence from the evaluation of binary treatments suggests that misspecification of the propensity score and outcome models can lead to severe bias (Kang et al., 2007). However in the GPS literature there is no research investigating the effects of model misspecification of the GPS on the estimated dose-response function, nor comparing the GPS and regression approaches.

We propose a framework for estimating the marginal causal effects of continuous treatments, which can mitigate the problem of model misspecification. In this approach both the GPS and the outcome model are selected in a data-adaptive way, with the Super Learner. The Super Learner, developed by Van Der Laan and Dudoit (2003), takes the weighted combination of several prediction algorithms with the aim of constructing an improved estimator. The Super Learner has been demonstrated to reduce bias from model misspecification in binary treatment settings (Porter et al., 2011; Kreif et al., 2014), but has not been previously combined with the GPS. The aim of this paper is to consider the Super Learner for estimating dose-response curves, by using it for estimating the GPS, as well as modelling the relationship between the outcome, treatment and the GPS. We compare this method to estimating the outcome model using parametric models. We perform the same contrast for the multivariable

regression approach for estimating the dose-response, relying only on parametric models and then using the Super Learner.

The methods are contrasted in an empirical example, in which we evaluate the marginal causal effect of increasing the transfer time from an emergency department to specialist neuroscience centres for critically ill patients with traumatic brain injury (TBI). We use data from the Risk Adjustment In Neurocritical care (RAIN) cohort study (Harrison et al., 2013).

The paper is organised as follows. Section 2 briefly describes the potential outcomes framework for dose-response functions, and reviews the regression and GPS approaches to estimating dose-response curves. We then introduce the Super Learner for estimating dose-response functions. Section 3 describes the methods used in our empirical application, and presents the results. Section 4 concludes and discusses limitations and further research.

2 Methods

2.1 Dose-response functions

Following Hirano and Imbens (2004), we define dose-response functions in the potential outcomes framework (Rubin, 2005). Let $i = 1$ to n be randomly sampled units. The continuous treatment of interest can take values in $t \in \tau$. For each unit, let $Y_i(t) : t \in \tau$ be a set of potential outcomes, each corresponding to the outcome in a hypothetical world in which $T = t$ is set deterministically. This set of potential outcomes is referred to as the unit-level dose-response function. The main parameter of interest is the average dose-response function, $\mu(t) = E(Y_i(t))$. A further parameter of interest, which can be derived from the previous one, is the marginal treatment effect function (Bia and Mattei, 2008; Bia et al., 2011), capturing the effect of increasing the level of treatment on the expected potential outcome: $(E[Y_i(t)] - E[Y_i(t - \Delta t)])/\Delta t$. For example, with $\Delta t = 1$, the parameter captures the incremental change in the outcome, for a unit change in the level of treatment.

For each unit, we observe a vector of covariates X_i , the level of treatment received, $T_i \in [t_0, t_1]$, and the observed outcome, which corresponds to the potential outcome under the level of treatment received, $Y_i = Y_i(T_i)$. The weak unconfoundedness assumption requires the independence of the potential outcome and the observed level of treatment for each value

of treatment: $Y(t) \perp T | X$ for all $t \in \tau$. Two additional assumptions are required. The consistency assumption requires that the observed outcome corresponds to the potential outcome under the treatment level received, formally, that $T = t$ implies $Y(t) = Y$. The positivity assumption requires that the conditional density of the treatment is non-negative for any covariate values, $P(r(t|X = x) > 0) = 1$. These assumptions imply that after conditioning on the observed covariates, the dose-response curve can be identified using the observed outcomes: $E[Y(t)|X = x] = E[Y|T = t, X = x]$.

These identification results are necessary for both the regression and GPS methods, but these approaches differ in the way that observed covariates are adjusted for.

2.2 Regression methods to estimate the dose-response function

Regression estimators aim to model the observed outcome as a function of the treatment level T and covariates X : $Q(t, x) = E[Y|T = t, X = x]$, and the average dose-response is defined as $\mu(t) = E[Q(t, X)]$. For example, if there is only one covariate, and the outcome is assumed to be a quadratic function of the treatment and the covariate, the conditional expectation of the outcome given covariates and the treatment level actually received could be estimated as $E[Y_i|T_i, X_i] = \alpha_0 + \alpha_1 T_i + \alpha_2 T_i^2 + \alpha_3 X_i + \alpha_4 X_i T_i$. The average dose-response function can be obtained by taking an averaged prediction using the estimated regression coefficients, at each treatment level of interest: $\hat{\mu}(t) = \hat{\alpha}_0 + \hat{\alpha}_1 t + \hat{\alpha}_2 t^2 + \hat{\alpha}_3 \hat{E}(X_i) + \hat{\alpha}_4 \hat{E}(X_i)t$.

2.3 The generalised propensity score method

Hirano and Imbens (2004) define the GPS as follows. Let $r(t, x) \equiv f_{T|X}(t|x)$ be the conditional density of treatment given covariates. Then the generalised propensity score is defined as the random variable $R = r(T, X)$, the conditional density evaluated at the treatment level received and the covariates observed.

The key feature of the GPS is its balancing property, similar to the balancing property of the propensity score of binary treatments. The balancing property states that, within strata of the same value of the GPS evaluated at a given treatment level, $r(t, X)$, the probability that the treatment received equals this treatment level, $T = t$, does not depend on the values of the covariates. This property, combined with the weak unconfoundedness assumption, implies

that the GPS can be used to eliminate any bias associated with differences in the observed covariates among groups of units with different levels of treatment. It also implies that the counterfactual expectation $E(Y_t)$ is identified as $\mu(t) = E[\beta(t, r(t, X))]$, where $\beta(t, r) = E[Y|T = t, R = r]$ is the conditional expectation of the observed outcome given the treatment level and the GPS.

The estimation of dose-response curves using the GPS involves two stages. First, the conditional density of the treatment is estimated, and the GPS is evaluated, at the level of treatment actually received, $\hat{R}_i = \hat{r}(T_i, X_i)$, and for the potential levels of treatment, $\hat{R}_i^t = \hat{r}(t, X_i)$.

In the second stage the conditional expectation of the outcome is estimated, given the treatment level and the GPS, $E[Y|T = t, R = r]$. For example, assuming that the outcome is a quadratic function of the treatment level and the GPS, the following regression will be estimated: $E[Y_i|T_i, \hat{R}_i] = \alpha_0 + \alpha_1 T_i + \alpha_2 T_i^2 + \alpha_3 \hat{R}_i + \alpha_4 T_i \hat{R}_i$. Then the average dose-response function is estimated at each treatment level of interest, by averaging the previously estimated conditional expectation over \hat{R}_i^t . This involves taking the estimated regression coefficients, and obtaining predictions for each unit, by plugging in the treatment level of interest, and the GPS evaluated at the treatment level of interest, for example: $\hat{\mu}(t) = \hat{\alpha}_0 + \hat{\alpha}_1 t + \hat{\alpha}_2 t^2 + \hat{\alpha}_3 \hat{E}(\hat{R}_i^t) + \hat{\alpha}_4 t \hat{E}(\hat{R}_i^t)$.

2.4 The Super Learner in estimating dose-response

The previous sections outlined the general approach for identifying and estimating the dose-response functions using the regression, and the GPS approaches. For both methods, an outcome model needs to be formulated, $Q(t, x)$ for the regression, and $\beta(t, r)$ for the GPS method. For the GPS approach, the conditional density of the treatment received, $r(t, x) \equiv f_{T|X}(t|x)$ also needs to be specified. Models for these quantities are often fitted using the researcher's preferred estimation method (e.g., a parametric model). We propose to use the Super Learner, an automated algorithm for selecting the best-performing estimator among a library of candidate estimators.

The Super Learner (Van Der Laan and Dudoit, 2003; van der Laan et al., 2007; Polley and van der Laan, 2010) is an ensemble prediction algorithm, exploiting machine learning.

In general, machine learning covers a wide range of classification and prediction algorithms. Unlike approaches that assume a fixed statistical model, for example a generalised linear model (GLM) with a gamma error distribution and a log link, machine learning aims to extract the relationship between the endpoint and covariates through a learning algorithm (Lee et al., 2010). Machine learning approaches were demonstrated to reduce bias resulting from model misspecification of the outcome model (Austin, 2012), and the propensity score of a binary treatment (Lee et al., 2010). The Super Learner algorithm uses cross-validation to select a weighted combination of estimates given by different prediction procedures (Polley and van der Laan, 2010). The range of prediction algorithms are pre-selected by the user, potentially including parametric and non-parametric regression models, and more general prediction algorithms. Asymptotically, the Super Learner algorithm performs as well as the best possible combination of the candidate estimators (Van Der Laan and Dudoit, 2003).

2.4.1 The Super Learner to estimate the outcome regression

The objective of the Super Learner is to consider a range of prediction algorithms for the outcome, and construct a new estimator as a convex combination of these estimators, which performs equal or better than any of the candidate estimators asymptotically. Let $\hat{\mu}_j(W) : j = 1, \dots, K$ a list of regression estimators for the conditional expectation of the outcome, $E(Y|W)$. For the regression approach, the vector W includes (x, t) and in the case of the GPS method, it consists of (r, t) .

We partition the sample in V cross-validation splits. Let $\mu_{j,\bar{v}}$ be the j -th candidate trained in the sample excluding split v . We apply each estimator in the training sample, and use the estimated model to predict the outcomes in the validation sample. We compute the cross-validated risk of each estimator, calculated as the squared prediction error,

$$L(\hat{\mu}_j) = \frac{1}{V} \sum_{v=1}^V \frac{1}{n_v} \sum_{i \in v} (Y_i - \hat{\mu}_{j,\bar{v}}(W_i))^2$$

The Super Learner estimator is a convex combinations of the list of estimators, which minimises the cross-validated risk. Formally, we consider the candidate Super Learner esti-

mators

$$\hat{\mu}_\alpha(W) = \sum_j \alpha_j \hat{\mu}_j(W),$$

and choose α as the minimizer of $L(\hat{\mu}_\alpha)$ constrained to $\alpha_j > 0$ and $\sum_j \alpha_j = 1$. The cross validated risk of this estimator, referred to as the "convex super learner" can be obtained by repeating the cross validation procedure from the first step, adding the convex Super Learner to the list of candidate estimators (Polley and van der Laan, 2013).

2.4.2 The Super Learner to estimate the GPS

Here we modify the above algorithm, to select an estimator of the conditional density of the treatment given covariates, by selecting a convex combination of candidate estimators, that minimises the negative log likelihood. Let $\hat{r}_j(t|X) : j = 1, \dots, K$ a list of GPS candidate estimators. For example, the conditional density can be derived from modelling t as a random variable following normal or gamma distribution. Candidate estimators can also include variations of these with different higher order terms in the linear predictor of the mean.

We partition the sample in V cross-validation splits. Let $r_{j,\bar{v}}$ be the j -th candidate trained in the sample excluding split v . We compute the cross-validated risk of each estimator as

$$L(\hat{r}_j) = \frac{1}{V} \sum_{v=1}^V \frac{1}{n_v} \sum_{i \in v} -\log r_{j,\bar{v}}(t_i|W_i)$$

We consider estimator candidates of the form

$$\hat{r}_\alpha(t|W) = \sum_j \alpha_j \hat{r}_j(t|W)$$

We choose α as the minimizer of $L(\hat{r}_\alpha)$ constrained to $\alpha_j > 0$ and $\sum_j \alpha_j = 1$. We use the "rsolnp" package in R to perform the optimisation (Ghalanos and Theussl, 2012).

3 Empirical example

3.1 The RAIN cohort study

Acute Traumatic Brain Injury (TBI) imposes a large burden in terms of cost and mortality (Harrison et al., 2013). An important public policy question is how best to manage critically

ill patients following an acute TBI. In particular, there are large local variations in the time to transferring patients from initial hospital presentation to arrival at a specialised neuroscience centre. The primary aim of the Risk Adjustment In Neurocritical care (RAIN) study was to validate risk prediction models for acute TBI and to use the best models to evaluate the optimum location and comparative costs of neurocritical care in the NHS (Harrison et al., 2013). A total of 67 critical care units participated in the RAIN study, with 3626 admitted patients providing a highly representative sample of patients receiving critical care following acute TBI in the UK.

An important research question the RAIN study aimed to answer was whether adult patients with TBI without an acute neurosurgical lesion benefit from an early decision to transfer to a neuroscience centre. The clinical literature is not conclusive about the benefits of early transfer: while some studies suggest that in patients for whom neurosurgery is not indicated (Bullock et al., 2006), the risks from early transfer and subsequent aggressive protocols of care may be substantial, an alternative view is that an early decision to keep the patient within the non-neuroscience centre can lead to delayed transfers, for example if a critical lesion develops subsequently, with potentially higher risks (Shafi et al., 2008). The RAIN study compared early (within 18 hours of hospital presentation) transfer to a neuroscience centre with no or late (after 24 hours) transfer, for patients who initially present at a non-neuroscience centre and do not require neurosurgery. It was found that at six months, patients in the early transfer group had significantly lower mortality, however higher total costs.

This raised a further research question of how, once the decision to transfer early had been made, variations in transfer time affects expected costs and mortality. Exogenous variation in transfer time is expected due to patient characteristics, local variations in management, and delays due to logistics. By investigating the causal effect of transfer time on outcomes, important insights may be gained for subsequent guideline development, for example on the benefits of reducing logistical barriers to shorter transfer times.

We aim to address this research question by estimating the dose-response relationship between transfer time and 6 month mortality, and transfer time and six months costs.

3.2 Data

We restrict the population of interest to patients with acute TBI who presented outside of a neuroscience centre, did not require neurosurgery, and were transferred to a neuroscience centre within 24 hours. We define transfer time as the time in hours between admission to the emergency department at the presenting hospital and admission to the specialist neuroscience centre. This definition, based on consultation with a panel of clinicians, reflects that a transfer more than 24 hours after hospital presentation implies a decision to delay transfer, rather than an intended early transfer delayed by logistics. Transfer time consists of the time spent at the emergency department, at an intermediate location such as a different ward of the admitting hospital, or a different hospital that is not a neuroscience centre, and time spent in transit between locations.

We aim to control for all variables which are potential confounders in the relationship between transfer time and mortality, and transfer time and costs, i.e. variables which might influence transfer time, and also affect these outcomes. We observe the prognostic variables which might influence the clinicians' decision on when to transfer, measured after the patient has been stabilised, but before the decision has been made to transfer the patient. We use the variables from the IMPACT lab model (Steyerberg et al., 2008), selected based on cross validation in the RAIN study, including clinical factors measured after stabilising the patient (hypoxia, hypotension, motor score, Glasgow Coma Score (GCS), pupil reactivity, Marshall CT classification, presence of traumatic subarachnoid haemorrhage, presence of extradural haematoma, lab measurements). Informed by clinical opinion, we include further important potential confounders: the presence of major extracranial injury, last pre-sedation GCS, variables indicating suspected or confirmed intoxication, age and gender. The descriptive statistics on the key potential confounders, as well as the outcomes, are presented in Table 2.

Mortality is measured as all cause mortality at the time of the six months follow up questionnaire. Six month costs are measured by considering resource use from the first critical care admission following the TBI, and readmissions within six months. The RAIN study prospectively recorded the length of stay (LOS) in critical care for each admission, and whether or not the patient had intracranial neurosurgery for evacuation of a mass lesion.

Each item of resource use was combined with the appropriate unit cost to report a cost per patient for each cost category in 2010–11 prices.

Missing outcome and covariate data has been addressed with multiple imputation using chained equations.

3.3 Implementation

We estimate the dose-response relationship a) between transfer time and mortality and b) between transfer time and costs, using regression and GPS approaches, implementing both methods using parametric models and the Super Learner. We also test the null hypothesis that increasing transfer time (according to units of one hour) has no effect on the expected outcomes.

We consider a range of candidate estimators for the GPS, including normal and gamma models with nonlinear terms and interactions in the linear predictor (see the specifications in Table 1). We use the Super Learner to select the best convex combination of these models, and estimate the GPS. We use these estimated GPSs throughout the analysis, both when the outcome is modelled relying on parametric models, as well as when using the Super Learner.

Several approaches have been proposed to evaluate whether the estimated GPS satisfies the balancing property. Imai and Van Dyk (2004) suggest regressing each covariate on the logarithm of the treatment variable, without and with conditioning on the predicted value of treatment given covariates. We follow the blocking approach proposed by Hirano and Imbens (2004), and categorise the treatment variable, by $k = 3$ quantiles. Let $I_k : k = 1, \dots, K$ be the indicator of each category, and let t_k be the median of each category. For each category, we compute the GPS evaluated at the median, $R_k \equiv r(t_k|X)$. Then for each quantile, we categorise the GPS, R_k into m blocks, again based on quantiles. For each covariate we compare the unadjusted mean differences and corresponding t-statistics between groups of units which belong to a treatment category versus those which do not belong to this treatment category. We then compute the same mean differences and t-statistics, but within each GPS block, defined above, and take the weighted mean of these mean differences, according to number of units in each block.

Kluve et al. (2012) and Flores et al. (2012) emphasise the importance of assessing the

common support after estimating the GPS, which can also be regarded as an indirect test of the positivity assumption. Taking a similar approach to the balance assessment described above, they divide the sample into groups according to the distribution of the treatment variable, and evaluate the GPS at the median of each group. For example, the GPS is evaluated at the median of the first group, for each unit in the sample, and the distribution of the GPS is plotted for those units which actually have treatment levels belonging to this group, versus those which have treatment values outside of these group. For each group the common support can be evaluated by inspecting the overlap of these distributions. Restricting the sample to individuals who are comparable across the three groups simultaneously ensures common support, however this approach can make the estimated treatment effects difficult to interpret. Instead, we report overlap to flag up potential violations of the positivity assumption, but do not drop observations.

For the outcome model, we first consider a range of parametric models, to demonstrate the potential impact of model choice on the estimated dose-response function. For mortality, we consider a range of logistic regression estimators, for the cost endpoint we consider GLMs with gamma distribution and log link. For the regression approaches, we vary the linear predictors of the models to include different degrees of polynomials of the treatment variable and the continuous covariates, and interactions between these. Categorical covariates are introduced as linear terms without interactions. Similarly, we vary the linear predictor in the outcome model for the GPS approach, by including polynomials of the treatment variable, the GPS, and their interactions (see specifications in Table 1).

We include these models in the Super Learner library, as well as a generalised additive model (GAM) with degrees of freedom of two (Hastie, 2013), and a Bayesian GLM approach with linear terms in the linear predictor, and non-informative priors for the coefficients (Gelman and Su, 2013). Each of the candidate algorithms controls for all the pre-specified potential confounders. We specify a 10 folds cross validation for the Super Learner. We use the nonparametric bootstrap to calculate uncertainty around the estimated dose-response function and marginal treatment effect function. Bootstrap standard errors and quantiles have been proved consistent only when the estimator converges in distribution at parametric rates (Van der Vaart, 2000), which prevented us from using more aggressive machine learning

predictors in the Super Learner library.

We draw 1000 bootstrap samples. In each bootstrapped dataset, we re-estimate the GPS and the outcome model, and use these to estimate the points of the dose-response function between $t = 1$ and 24, and the marginal treatment effect between $t = 2$ and 24, using 1 hour increments. We obtain 95% confidence intervals (CIs) of these parameters, based on the quantiles of the bootstrap distribution. We conduct the analysis separately for the 5 multiple imputed datasets. In order to demonstrate the approach, we present results on the first of the imputed datasets. All computation is performed using the R platform (R Core Team, 2013).

< Table 1 around here >

3.4 Results of the empirical example

3.4.1 The estimated GPS, common support, balance

Table 2 describes the sample ($n=488$), while Figure 1a shows the empirical distribution of the transfer time variable. The minimum was at 1.83, and maximum at 23.7 hours. The Super Learner algorithm selected a mixture of the normal and gamma models for the GPS, assigning weights of 18, 32 and 16 and 3% to the normal model candidates, while the gamma models altogether contributed 32%.

The estimated GPS, evaluated at the treatment levels actually received is presented on Figure 1b. Figure 2 presents the assessment of overlap in tertiles of the treatment variable: 1.83 – 5.2, 5.2 – 10.1 and between 10.1 – 23.7 hours. It appears that there is only a small proportion of patients (2.2%) in the comparison of the middle tertile versus the others, for whom there is a lack of overlap at values of the GPS close to zero.

< Table 2 around here >

< Figure 1 around here >

< Figure 2 around here >

Table 3 presents balance statistics as mean differences (t-statistics) before and after adjustment with the GPS. The IMPACT predicted mortality is a composite score expressing the probability of mortality based on the IMPACT lab risk prediction model, and aims to capture the baseline severity of patients. Analogously to the assessment of common support, balance

has been evaluated for the tertiles of the treatment variable, unadjusted and then adjusted for 5 blocks of the GPS. Patients with lower treatment levels were significantly younger, had a lower prognostic score, and were less likely to have a major extracranial injury. After adjustment, mean differences and t-statistics were reduced for those variables where initial imbalance was relatively high (major extracranial injury, age, IMPACT predicted mortality), although for some variables, for example age in the higher treatment categories, imbalance increased.

< Table 3 around here >

3.4.2 Dose-response functions, using parametric models and the Super Learner

We estimated dose response functions for six months mortality and costs, using different parametric specifications of the outcome model for the regression and the GPS approach (upper panel of Figure 3). The plots suggest an increasing relationship in transfer time and expected mortality, however, some of the models estimated a non-monotonic relationship. For example, when the GPS approach used a quadratic term of the treatment variable in the outcome model, as proposed by Hirano and Imbens (2004), the effect of increasing transfer time between 12 and 18 hours appears to reduce rather than increase expected mortality.

The relationship between transfer time and costs seems to be generally U-shaped (Figure 4), however again, different models suggesting divergent curves for certain segments of the distribution of the treatment variable.

Table 4 presents the results from the cross validation performed by the Super Learner algorithm. The table displays the estimated MSEs of the candidate algorithms, alongside the weight each candidate received in the final Super Learner selection, as well as the MSE for the Super Learner estimator.

Among the parametric models, relatively simple outcome models, including linear terms only, provided the best fit in terms of mean squared prediction error. For the mortality endpoint, none of the parametric models from the previous, exploratory analysis received a positive weight in the Super Learner. A combination of the Bayesian GLM prediction algorithm, and the GAM received all weights for the regression, while for the GPS approach, the Bayesian GLM predictor alone was given full weight. For the cost outcome, the pre-

specified parametric models contributed positive weights in the Super Learner estimator, together with the Bayesian GLM approach.

< Figure 3 around here >

< Figure 4 around here >

< Table 4 around here >

The lower panel of Figures 3 and 4 presents the estimated dose-response functions with the Super Learner. For the mortality endpoint, both the regression and GPS approach suggest a monotonically increasing relationship, while for the cost endpoint, the GPS approach suggests that expected costs initially decrease (up to around seven hours) and then increase with transfer time.

In this sample we found the MSE of the best single candidates was slightly lower than that of the convex Super Learner. Because the differences are small, choosing the best candidate estimators would provide similar dose-response functions to the convex Super Learner. In general, van der Laan et al. (2007) suggest that for moderately large samples, such as our study, choosing the convex Super Learner provides more stable estimates than choosing the single best model.

Figures 5 and 6 present point estimates (95% CI) for the marginal treatment effect from the Super Learner estimators. The CIs for the marginal treatment effect function include zero, corresponding to a zero incremental effect of increasing transfer time on the expected outcomes. Hence, the null hypothesis, that increasing transfer time does not have an effect on expected six months costs or mortality, cannot be rejected, at the 5% level of statistical significance. When comparing the width of the CIs between the regression and the GPS approach, we find that the GPS approach reports a larger uncertainty, especially for those areas of the data distribution where the empirical distribution of observed treatment was sparse.

< Figure 5 around here >

< Figure 6 around here >

4 Discussion

This paper provides a framework for estimating marginal causal effects of continuous treatments that does not require the model for the GPS or the outcome regression to be correctly specified. Instead this paper proposes that a data-adaptive method, the Super Learner, can be applied to the GPS approach proposed by Hirano and Imbens (2004). We contrast this approach to parametric implementations of the GPS approach. The paper also compares the GPS approach to regression methods for estimating the dose-response function both with parametric implementation, and using the Super Learner.

The paper illustrates the approach in an empirical example with characteristics typical of program evaluations where the sample size is moderate, and it is necessary to control for many binary and continuous covariates to make a plausible assumption about unconfoundedness. We find that both the regression and the GPS approaches are sensitive to the choice of model specification for the endpoint models, and that the estimated dose-response curves differ by parametric specification. In this example, the Super Learner estimator assigned small weights to nonlinear models, which suggested that the non-monotonic dose-response functions resulting from some of the parametric models were not supported by the data. With the Super Learner, the regression and the GPS approaches led to the same conclusion, namely that the marginal effect of increasing transfer time on mortality and cost was zero (at the 5% level of statistical significance).

In this example, the GPS approach reports wider confidence intervals than the regression approach. This is expected, as estimators using the propensity score are usually less efficient than estimators based on a correctly specified outcome model (Vansteelandt and Daniel, 2014). This may also reflect, that in this setting it was relatively challenging for the Super Learner to select the combination of models that specify the covariate to treatment versus the covariate to outcome relationship. Here, while as part of the RAIN study, an extensive systematic review of previous outcome regression models was undertaken (Harrison et al., 2013), there was little prior information on the form that the GPS may take.

By proposing data-adaptive estimation of the GPS and the outcome regression, this paper builds on the GPS approach proposed by Hirano and Imbens (2004). A related extension

is described in the working paper by Bia et al. (2011) who propose flexible, spline based estimators for the outcome model. Using simulations, they demonstrate that their approach outperforms parametric estimators when the dose-response function is nonlinear. However, even a flexible spline approach requires subjective modelling choices, such as a pre-fixed degree of polynomials and the number of knots in the spline. The Super Learner approach can further increase flexibility, and reduce misspecification, by potentially incorporating these estimators among the candidate predictors.

Our paper highlights a distinguishing feature of Super Learner compared to other model selection approaches, that it combines many estimators, by selecting a combination of predictions from alternative prediction algorithms. That is, the Super Learner aims to provide a better fit to the data than relying on any one prediction algorithm. It is recommended that the analyst requires the Super Learner to consider a rich set of prediction algorithms (van der Laan et al., 2007), to facilitate the consistent estimation of the dose-response function. In selecting the candidates, subject matter knowledge of the data-generating process should be used, for example in this paper we included models with the gamma distribution and log link function to predict costs. Using the Super Learner for estimator selection also encourages transparency in model selection, by reporting the weights and MSE of the candidate predictors.

Our paper adds to the growing literature on the use of the Super Learner for causal inference (Porter et al., 2011; Kreif et al., 2014; Gruber and van der Laan, 2010; Pirracchio et al., 2014), and more generally, to the implementations of machine-learning methods in estimating treatment effects (Lee et al., 2010; Austin, 2012). The setting of continuous treatments posed new challenges for the Super Learner framework, in having to represent the uncertainty in the estimator selection for both the GPS and the outcome regression, which this work has addressed with the non-parametric bootstrap.

This paper has some limitations. First, each of the approaches relies on the assumption of no unmeasured confounders, specifically in the context of continuous treatments, the weak unconfoundedness assumption. This assumption requires that for any level of treatment, the probability of receiving this level is independent of the potential outcomes, conditional on covariates. In our empirical example, this assumption required that factors that cause delays

in transfer, and are also prognostic of the outcomes, are controlled for. We used potential confounders from a previously validated risk prediction model, as well as clinical opinion, to pre-specify a set of variables. However, the possibility for unobserved confounding remains, for example because the covariates are measured at the time of hospital presentation, so subsequent changes in patients' prognosis, which might cause delays in transfer and effect six month outcome, are unmeasured. Kluve et al. (2012) assess the robustness of the dose-response curves to remaining unobserved confounding by also employing instrumental variable estimation. In the absence of appropriate instruments, the effects of unobserved confounding could be examined by extending sensitivity analysis methods to the context of continuous treatments (Rosenbaum, 1987).

Second, in this study covariate balance following adjustment with the GPS did not improve for all variables. An alternative loss function for the Super Learner could explicitly consider a metric that takes into account the balance achieved. For the binary propensity score, Lee et al. (2010) propose a data adaptive algorithm to estimate the GPS, based on balance measures such as the Kolmogorov-Smirnoff test. Such approaches still require subjective choices of the appropriate balance measure, the prioratisation of confounders (Stuart, 2010), and for continuous treatments, the method for categorising the treatment variable. Indeed, the most appropriate balance metric remains a topic of ongoing debate (Kluve et al., 2012).

This work provokes areas of further research. Future simulation studies could examine the sensitivity of the dose-response curve to the misspecification of the GPS, and investigate how this misspecification can be assessed by evaluating the balancing property of the GPS.

While the regression and GPS approach lead to similar conclusions in our empirical example, more generally, a criterion for choosing between these approaches is required. An approach for choosing between dose-response curves was recently proposed by Díaz and van der Laan (2013), building on the principles of Super Learning. This approach requires multiple cross-validation stages, and extensive computational times, but would provide a principled way of choosing between candidate estimation approaches.

References

- Austin, P. C. (2012). Using ensemble-based methods for directly estimating causal effects: an investigation of tree-based g-computation. *Multivariate behavioral research*, 47(1):115–135.
- Bia, M. and Mattei, A. (2008). A stata package for the estimation of the dose-response function through adjustment for the generalized propensity score. *The Stata Journal*, 8(3):354–373.
- Bia, M., Mattei, A., Flores-Lagunes, A., and Flores, C. (2011). Nonparametric estimators of dose-response functions. *CEPS/INSTEAD*.
- Bullock, M., Chesnut, R., Ghajar, J., Gordon, D., Hartl, R., Newell, D. W., Servadei, F., Walters, B. C., and Wilberger, J. (2006). Surgical management of traumatic brain injury. *Neurosurgery*, 58(3):16–24.
- Díaz, I. and van der Laan, M. J. (2013). Targeted data adaptive estimation of the causal dose–response curve. *Journal of Causal Inference*, 1(2):171–192.
- Flores, C. A., Flores-Lagunes, A., Gonzalez, A., and Neumann, T. C. (2012). Estimating the effects of length of exposure to instruction in a training program: the case of job corps. *Review of Economics and Statistics*, 94(1):153–171.
- Gelman, A. and Su, Y.-S. (2013). *arm: Data Analysis Using Regression and Multi-level/Hierarchical Models*. R package version 1.6-09.
- Ghalanos, A. and Theussl, S. (2012). *Rsolnp: General Non-linear Optimization Using Augmented Lagrange Multiplier Method*. R package version 1.14.
- Gruber, S. and van der Laan, M. J. (2010). An application of collaborative targeted maximum likelihood estimation in causal inference and genomics. *The International Journal of Biostatistics*, 6(1).
- Harrison, D., Prabhu, G., Grieve, R., Harvey, S., Sadique, M., Gomes, M., Griggs, K., Walmsley, E., Smith, M., Yeoman, P., et al. (2013). Risk adjustment in neurocritical care (rain)-prospective validation of risk prediction models for adult patients with acute traumatic brain injury to use to evaluate the optimum location and comparative costs of neurocritical care: a cohort study. *Health Technol Assess*, 17(23):1–350.
- Hastie, T. (2013). *gam: Generalized Additive Models*. R package version 1.09.
- Hirano, K. and Imbens, G. W. (2004). The propensity score with continuous treatments. *Applied Bayesian modeling and causal inference from incomplete-data perspectives*, 226164:73–84.
- Imai, K. and Van Dyk, D. A. (2004). Causal inference with general treatment regimes. *Journal of the American Statistical Association*, 99(467).
- Imbens, G. W. (2000). The role of the propensity score in estimating dose-response functions. *Biometrika*, 87(3):706–710.

- Kang, J. D., Schafer, J. L., et al. (2007). Demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data. *Statistical science*, 22(4):523–539.
- Kluve, J., Schneider, H., Uhlendorff, A., and Zhao, Z. (2012). Evaluating continuous training programmes by using the generalized propensity score. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 175(2):587–617.
- Kreif, N., Gruber, S., Radice, R., Grieve, R., and Sekhon, J. S. (2014). Evaluating treatment effectiveness under model misspecification: A comparison of targeted maximum likelihood estimation with bias-corrected matching. *Statistical methods in medical research*, page 0962280214521341.
- Lee, B. K., Lessler, J., and Stuart, E. A. (2010). Improving propensity score weighting using machine learning. *Statistics in medicine*, 29(3):337–346.
- Pirracchio, R., Petersen, M., and van der Laan, M. (2014). Improving propensity score estimators’ robustness to model misspecification using super learner.
- Polley, E. and van der Laan, M. (2013). *SuperLearner: Super Learner Prediction*. R package version 2.0-10.
- Polley, E. C. and van der Laan, M. J. (2010). Super learner in prediction.
- Porter, K. E., Gruber, S., van der Laan, M. J., and Sekhon, J. S. (2011). The relative performance of targeted maximum likelihood estimators. *The international journal of biostatistics*, 7(1):1–34.
- R Core Team (2013). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Robins, J. M., Hernan, M. A., and Brumback, B. (2000). Marginal structural models and causal inference in epidemiology. *Epidemiology*, 11(5):550–560.
- Rosenbaum, P. R. (1987). Sensitivity analysis for certain permutation inferences in matched observational studies. *Biometrika*, 74(1):13–26.
- Royston, P. and Altman, D. G. (1994). Regression using fractional polynomials of continuous covariates: parsimonious parametric modelling. *Applied Statistics*, pages 429–467.
- Royston, P., Altman, D. G., and Sauerbrei, W. (2006). Dichotomizing continuous predictors in multiple regression: a bad idea. *Statistics in medicine*, 25(1):127–141.
- Rubin, D. B. (2005). Causal inference using potential outcomes. *Journal of the American Statistical Association*, 100(469).
- Shafi, S., Diaz-Arrastia, R., Madden, C., and Gentilello, L. (2008). Intracranial pressure monitoring in brain-injured patients is associated with worsening of survival. *Journal of Trauma-Injury, Infection, and Critical Care*, 64(2):335–340.

- Steyerberg, E. W., Mushkudiani, N., Perel, P., Butcher, I., Lu, J., McHugh, G. S., Murray, G. D., Marmarou, A., Roberts, I., Habbema, J. D. F., et al. (2008). Predicting outcome after traumatic brain injury: development and international validation of prognostic scores based on admission characteristics. *PLoS medicine*, 5(8):e165.
- Stuart, E. A. (2010). Matching methods for causal inference: A review and a look forward. *Statistical science: a review journal of the Institute of Mathematical Statistics*, 25(1):1.
- Van Der Laan, M. J. and Dudoit, S. (2003). Unified cross-validation methodology for selection among estimators and a general cross-validated adaptive epsilon-net estimator: Finite sample oracle inequalities and examples.
- van der Laan, M. J., Polley, E. C., and Hubbard, A. E. (2007). Super learner.
- Van der Vaart, A. W. (2000). *Asymptotic statistics*, volume 3. Cambridge university press.
- Vansteelandt, S. and Daniel, R. (2014). On regression adjustment for the propensity score. *Statistics in medicine*.
- Yang, S., Imbens, G., Cui, Z., Faries, D., and Kadziola, Z. (2014). Propensity score matching and subclassification with multi-level treatments. Presented at the 2014 Atlantic Causal Inference Conference.

Table 1: Specifications of candidate algorithms in Super Learner

GPS estimation			
Candidate	Error dist	Link fn	Linear pred
Norm 1	normal	identity	W, A
Norm 2	normal	identity	W, age^2, fg^2, fhb^2
Norm 3	normal	identity	$W, age * fg, age * fhb, fg * fhb, motor3 * fhb$
Norm 4	normal	identity	W, age^2, fg^2, fhb^2
			$age * fg, age * fhb, fg * fhb, motor3 * fhb$
Gam 1	gamma	log	W, A
Gam 2	gamma	log	W, age^2, fg^2, fhb^2
Gam 3	gamma	log	$W, age * fg, age * fhb, fg * fhb, motor3 * fhb$
Gam 4	gamma	log	W, age^2, fg^2, fhb^2
			$age * fg, age * fhb, fg * fhb, motor3 * fhb$
Outcome estimation, regression			
Mortality			
Candidate	Error dist	Link fn	Linear pred
Linear	binomial	logit	W
Linear (costs)			
Linear, int	binomial	logit	$W, A * age, A * fg, A * fhb$
Quadr, int (1)	binomial	logit	$W, A, A^2, A * age$
			$A * fg, A * fhb, A^2 * age$
			$A^2 * fg, A^2 * fhb$
Quadr, int (2)	binomial	logit	$W, A, A^2, A * age, A * fg$
			$A * fhb, A^2 * age, A^2 * fg$
			$A^2 * fhb, A * age^2, A * fg^2$
			$A * fhb^2, age^2, fg^2, fhb^2$
4th	binomial	logit	W, A, A^2, A^3, A^4
GAM	binomial	logit	splines of W, A (df=2)
Bayesian GLM	binomial	logit	W, A
Outcome estimation, GPS			
Mortality			
Linear	binomial	logit	A, r
Linear (costs)			
Linear, int	binomial	logit	$linear, A, r, A * r$
Quadr, int (1)	binomial	logit	$A, r, A^2, r * A, A^2 * r$
Quadr, int (2)	binomial	logit	$A, r, A^2, r * A, A^2 * r, r^2, r^2 * A$
4th	binomial	logit	r, A^2, A^3, A^4
GAM	binomial	logit	splines of r, A (df=2)
Bayesian GLM	binomial	logit	A, r
Costs			
	normal	Identity	A, r
	gamma	log	A, r
	gamma	log	$A, r, A * r$
	gamma	log	$A, r, A^2, r * A, A^2 * r$
	gamma	log	$A, r, A^2, r * A, A^2 * r, r^2, r^2 * A$
	gamma	log	r, A^2, A^3, A^4
	normal	identity	splines of r, A (df=2)
	normal	identity	A, r

W : all covariates, A treatment variable, r : GPS, fg : glucose, fbh : haemoglobin, $motor3$: motor score 3

Table 2: Descriptive statistics of baseline covariates and outcomes

Variable	n=488	
Outcomes		
Dead at six months, n (%)	99	(20.3)
Six months costs (£), mean (SD)	27480	(29741)
Baseline covariates		
IMPACT pred mort, mean (SD)	0.23	(0.17)
Age, mean (SD)	40.33	(17.51)
Major extr inj, n (%)	185	(37.9)
Severe GCS, n (%)	265	(54.3)
Motor score poor, n (%)	226	(46.3)
Any pupil unreactive, n (%)	76	(15.6)

Abbreviations. SD: standard deviation, IMPACT pred mort: predicted mortality from IMPACT risk prediction model, extr inj: extracranial injury, GCS: Glasgow Coma Score.

The following variables had missing values for n observations: dead at six months: n=18, pupil unreactive: n=41, motor score: n=7.

Table 3: Balance unadjusted and adjusted for the GPS: means and t-statistics for the equality of means

Variable	Unadjusted mean (t-stat)			Adjusted mean (t-stat)		
	tx cat [1.83 - 5.2]	tx cat [5.2 - 10.1]	tx cat [10.1-23.7]	tx cat [1.83 - 5.2]	tx cat [5.2 - 10.1]	tx cat [10.1-23.7]
IMPACT pred mort	-0.03 (1.99)	-0.00 (-0.04)	0.03 (1.92)	-0.01 (-0.91)	0.00 (0.18)	0.01 (0.81)
Age	-4.88 (-3.05)	3.87 (2.26)	0.99 (0.58)	-3.09 (-1.81)	4.22 (2.49)	-1.09 (-0.65)
Major extr inj	-0.08 (-1.69)	-0.02 (-0.35)	0.09 (1.99)	-0.02 (-0.32)	0.01 (0.12)	0.02 (0.47)
Severe GCS	0.07 (1.36)	-0.06 (-1.25)	-0.01 (-0.10)	0.03 (0.59)	-0.07 (-1.49)	0.04 (0.73)
Motor score poor	0.018 (0.38)	-0.069 (-1.45)	0.051 (1.06)	0.011 (0.23)	-0.068 (-1.42)	0.050 (1.00)
Any pupil unreactive	-0.01 (-0.33)	-0.03 (-0.92)	0.04 (1.18)	0.01 (0.15)	-0.03 (-0.76)	0.02 (0.61)

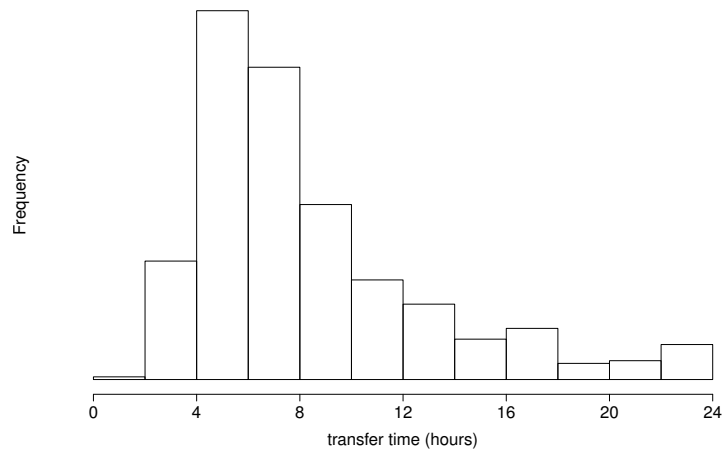
tx cat: treatment category. Mean differences reported, with t-statistics for the equality of the mean in brackets. Each comparison contrast units in a given treatment category vs. the other two treatment categories. There are 122, 123 and 123 patients in each group.

Table 4: Cross validation results and Super Learner weights

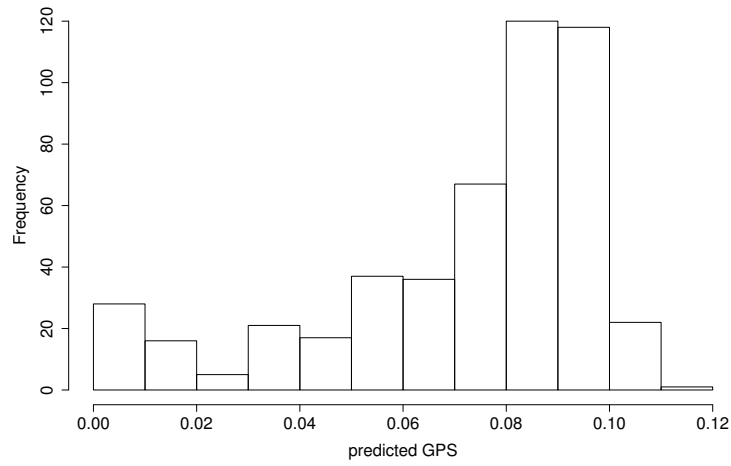
Mortality					
Regression approach			GPS approach		
Candidate predictor	MSE	Weight in Super Learner	MSE	Weight in Super Learner	Weight in Super Learner
Linear	0.147	0.00	0.162		0.00
Linear, int	0.150	0.00	0.162		0.00
Quadratic, int (1)	0.155	0.00	0.162		0.00
Quadratic, int (2)	0.157	0.00	0.164		0.00
4th order	0.148	0.00	0.163		0.00
GAM (df=2)	0.147	0.31	0.162		0.00
Bayesian GLM	0.145	0.69	0.162		1.00
Convex super learner	0.148		0.162		

Costs					
Regression approach			GPS approach		
Candidate predictor	MSE	Weight in Super Learner	MSE	Weight in Super Learner	Weight in Super Learner
Linear (normal)	853058	0.00	883591		0.00
Linear	855597	0.13	883507		0.83
Linear int	864127	0.14	889775		0.00
Quadr, int (1)	941058	0.00	895103		0.17
Quadr, int (2)	897010	0.20	912055		0.00
4th order	891008	0.00	901585		0.00
GAM (df=2)	857542	0.00	887886		0.00
Bayesian GLM	852746	0.52	883579		0.00
Convex super learner	859842		888529		

MSE for costs in 1000 £

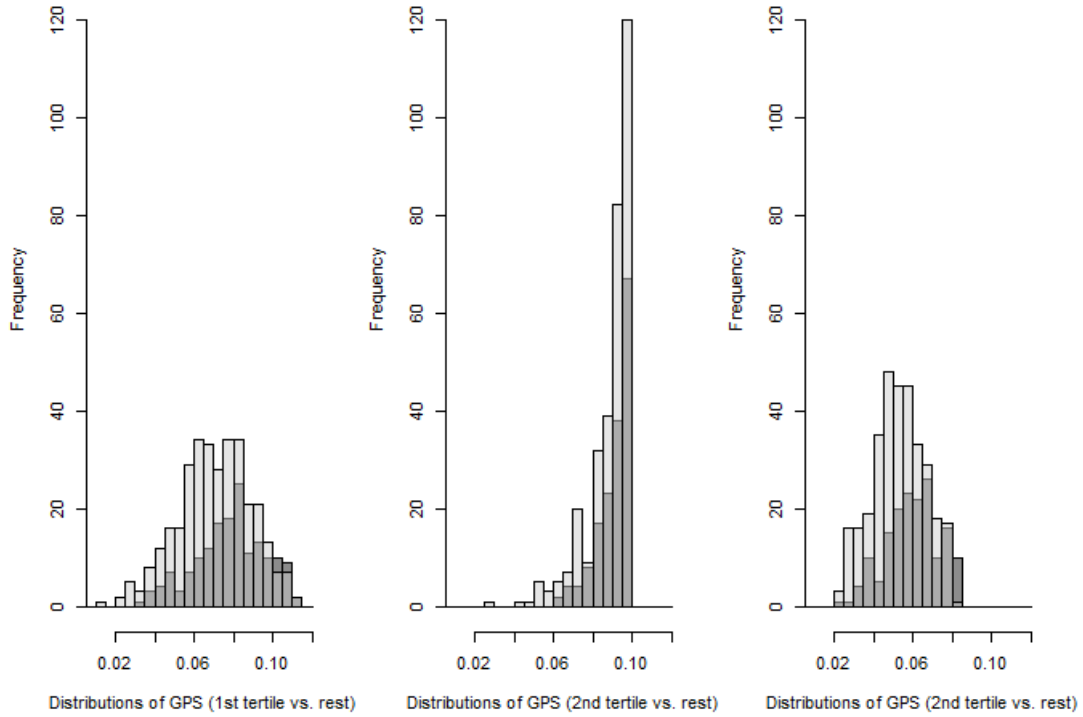


(a) Transfer time (hours)



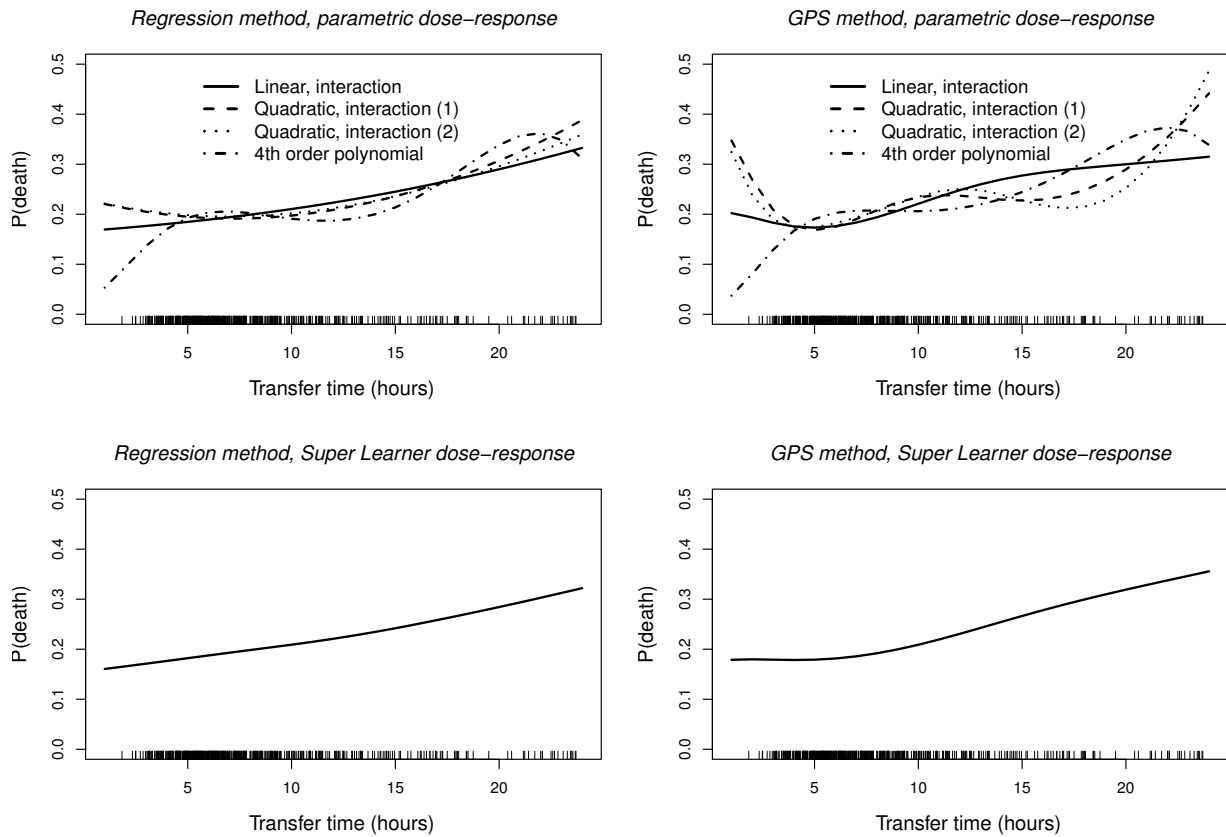
(b) Estimated GPS

Figure 1: The distribution of (a) transfer time (hours) and (b) the estimated GPS



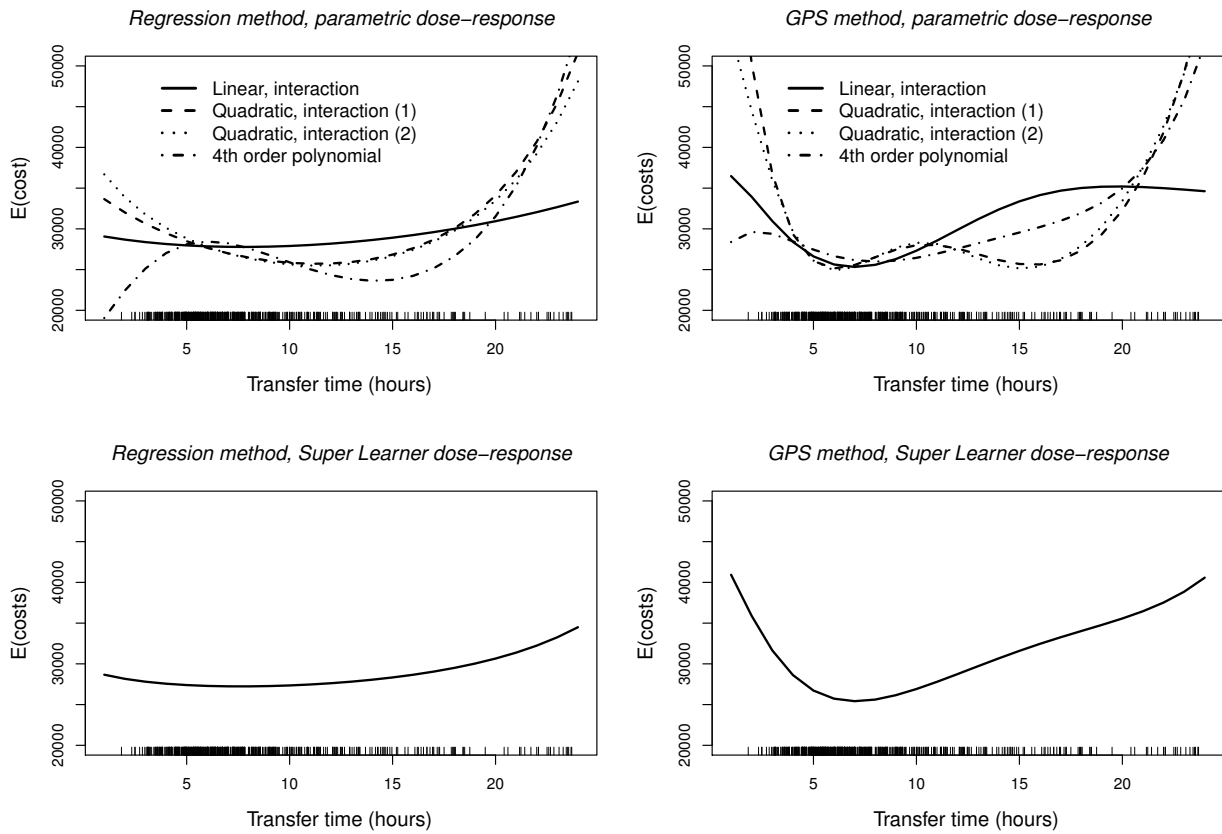
The three graphs compare the distribution of GPS evaluated at the medians of the three treatment groups, 4.66, 6.93 and 12.25. The light grey histogram shows the distribution of the GPS for those who received the treatment level of that category. The dark grey histogram shows the GPS evaluated at the same level, but for those who received treatment of different levels. There are 122, 123 and 123 patients in each group.

Figure 2: Overlap, based on the GPS estimated at medians of tertiles of the transfer time distribution



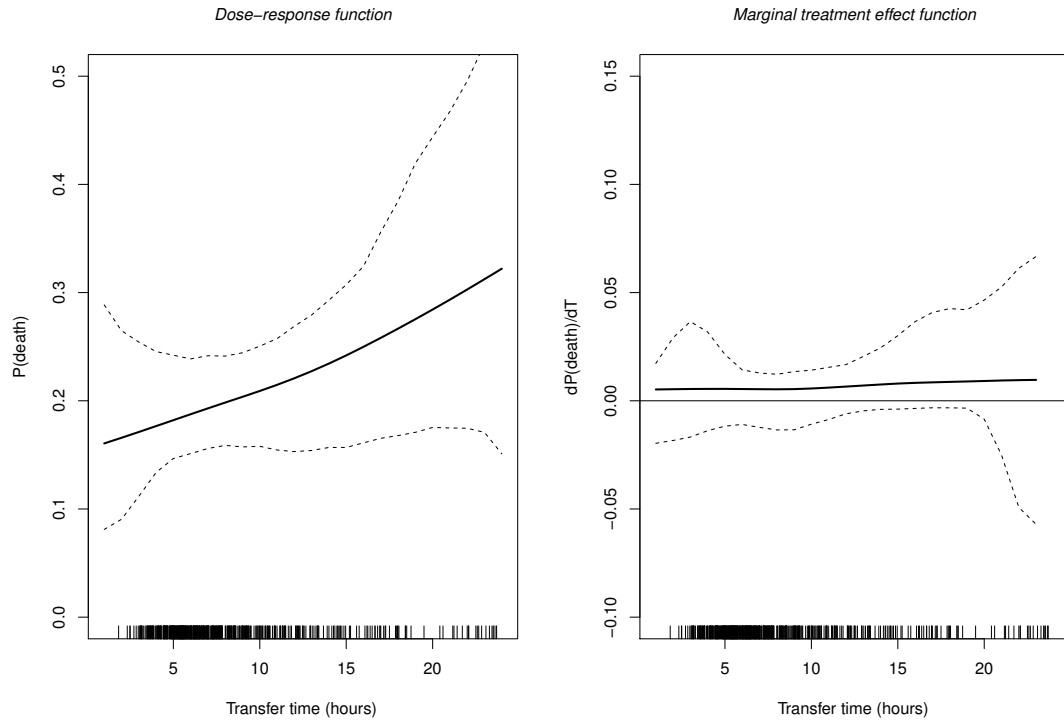
The rug plots demonstrate the distribution of observed transfer times.

Figure 3: Dose-response functions of expected mortality at six months, using regression and GPS, with parametric models and the Super Learner

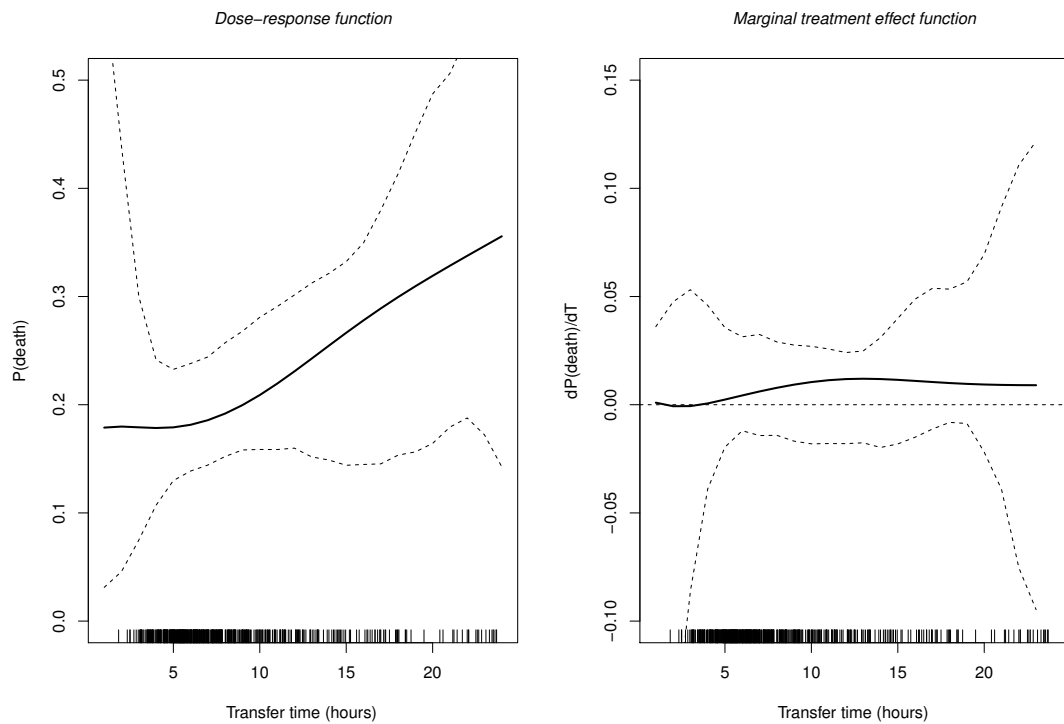


The rug plots demonstrate the distribution of observed transfer times.

Figure 4: Dose-response functions of expected costs at six months, using regression and GPS, with parametric models and the Super Learner



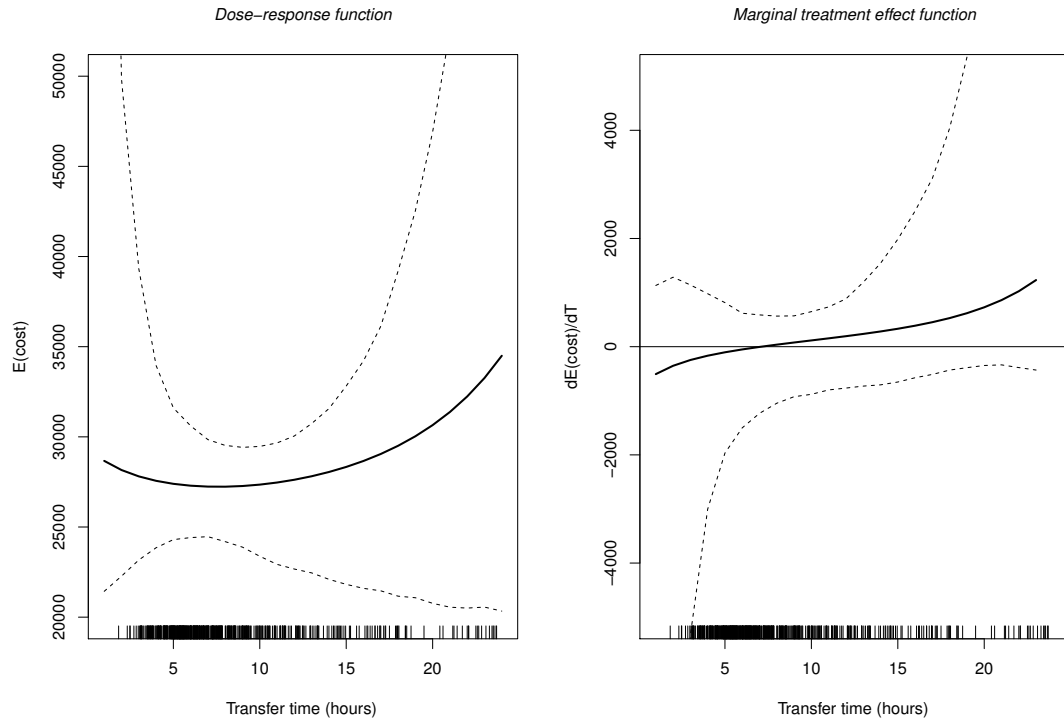
(a) Regression approach, point estimates and 95 % CI



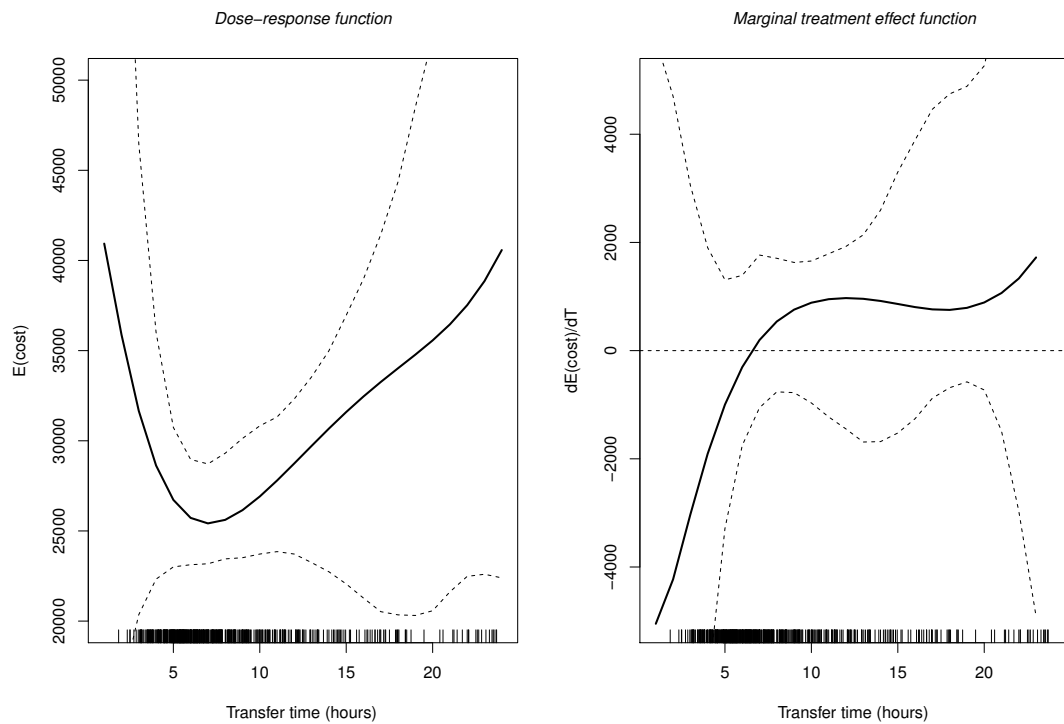
(b) GPS approach, point estimates and 95 % CI

The rug plots demonstrate the distribution of observed transfer times.

Figure 5: Dose-response function and marginal treatment effect function of expected mortality at six months, using the Super Learner



(a) Regression approach, point estimates and 95 % CI



(b) GPS approach, point estimates and 95 % CI

The rug plots demonstrate the distribution of observed transfer times.

Figure 6: Dose-response function and marginal treatment effect function of expected costs at six months, using the Super Learner