

# Data-adaptive doubly robust instrumental variable methods for treatment effect heterogeneity

K. DiazOrdaz, R. Daniel, N. Kreif

Department of Medical Statistics, LSHTM,  
Division of Population Medicine, Cardiff University,  
Centre for Health Economics, University of York.

## Abstract

We consider the estimation of the average treatment effect in the treated as a function of baseline covariates, where there is a valid (conditional) instrument.

We describe two doubly robust (DR) estimators: a locally efficient g-estimator, and a targeted minimum loss-based estimator (TMLE). These two DR estimators can be viewed as generalisations of the two-stage least squares (TSLS) method to semi-parametric models that make weaker assumptions. We exploit recent theoretical results that extend to the g-estimator the use of data-adaptive fits for the nuisance parameters.

A simulation study is used to compare standard TSLS with the two DR estimators' finite-sample performance, (1) when fitted using parametric nuisance models, and (2) using data-adaptive nuisance fits, obtained from the Super Learner, an ensemble machine learning method.

Data-adaptive DR estimators have lower bias and improved coverage, when compared to incorrectly specified parametric DR estimators and TSLS. When the parametric model for the treatment effect curve is correctly specified, the g-estimator outperforms all others, but when this model is misspecified, TMLE performs best, while TSLS can result in large biases and zero coverage.

Finally, we illustrate the methods by reanalysing the COPERS (COPing with persistent Pain, Effectiveness Research in Self-management) trial to make inference about the causal effect of treatment actually received, and the extent to which this is modified by depression at baseline.

## 1 Introduction

There has been an increased interest in estimating the causal effect of treatment actually received (in addition to the intention-to-treat effect) in randomised controlled trials (RCTs) in the presence of treatment non-adherence (Wiles et al., 2014; Dodd et al., 2012; Zhang et al., 2014). An additional challenge is posed by appreciable treatment effect heterogeneity, which is often itself of interest. This is a common issue with psychological interventions (Dunn and Bentall, 2007).

In this work, we consider methods for estimating the dependence of a causal average treatment effect on baseline covariates in RCTs with non-adherence. This is motivated by the COPERS (COping with persistent Pain, Effectiveness Research in Self-management) trial. The intervention introduced cognitive behavioural therapy approaches designed to promote self-efficacy to manage chronic pain, with the primary outcome being pain-related disability. The research team was interested in the causal effect of the received treatment, and whether this effect was modified by a number of baseline variables. Here, we will focus on one possible effect modifier: depression at baseline, measured by the Hospital Anxiety and Depression Scale (HADS).

Instrumental variable (IV) methods are often used to estimate the effect of treatment received in RCTs where randomised treatment is unconfounded by design, but treatment received is not. Assuming that randomised treatment is a valid instrument, and under some additional assumptions reviewed in Section 2, it is possible to identify the average treatment effect in the treated, conditional on baseline covariates  $V$ . In addition to investigating effect modification by a subset of baseline covariates  $V$ , it can in other settings be beneficial to use a larger set  $W$  of baseline covariates for adjustment in the analysis: (i) if the IV assumptions are more plausible conditional on baseline covariates  $W$ , or (ii) to increase the statistical efficiency of the estimators.

A relatively simple method of estimation for this is the so-called two stage least squares (TSLS). In its simplest form, i.e. when  $V$  is null, the first stage predicts the exposure based on an ordinary least squares regression of the exposure on the IV and baseline covariates  $W$ , while the second stage consists of regressing the outcome on the predicted exposure from the first stage and baseline covariates  $W$ , also via ordinary least squares regression. The coefficient corresponding to the predicted exposure in this second model is the TSLS estimator of the desired causal treatment effect. TSLS is robust to the misspecification of the first stage model (Robins, 2000; Wooldridge, 2010) but may be inefficient, especially when the treatment-exposure relationship is non-linear (Vansteelandt and Didelez, 2018). However, as we shall see later, where  $V$  is non-null and the treatment effect varies by baseline covariates, TSLS relies on the outcome model (the second stage) being correctly specified to obtain consistent effect estimates.

Doubly robust (DR) estimators are appealing in such settings, as they estimate consistently the parameter of interest if at least one of the models, for either the exposure or the outcome is correctly specified. In the context of linear IV models with  $V$  null, Okui et al. (2012) proposed a locally-efficient estimating equations DR estimator for the causal effect of treatment in the treated, often called a *g-estimator*. It augments the TSLS estimating equation by adding a model for the instrument given the baseline covariates. This estimator is DR in the sense that it needs to specify correctly either the outcome model or the instrument model. This estimator was generalised to settings where  $V$  is non-null by Vansteelandt and Didelez (2018) and shown to be locally efficient.

Recently, Tóth and van der Laan (2016) proposed a targeted maximum likelihood estimator (TMLE) for the treatment effect in a linear IV model. TMLE is a general approach for causal inference problems yielding semi-parametric substitution estimators that are doubly robust (van der Laan and Rose, 2011).

Although DR estimators offer partial protection in principle against model misspecification,

concerns remain over their performance in practice, when all models are likely to be misspecified (Kang et al., 2007). To further avoid problems of model misspecification, TMLE is usually coupled with machine learning fits of nuisance models, in particular the Super Learner, a cross-validation based estimator selection approach (van der Laan et al., 2007). TMLE and other DR estimators possess a particular orthogonality property that leads to greater suitability with machine learning fits. Estimators based on a single nuisance model can perform poorly when combined with machine learning fits, since the estimator inherits the slow convergence (and hence high finite sample bias) and non-regularity of the machine learning fits themselves, with the latter property making valid statistical inferences extremely complex to obtain, as even the bootstrap is invalid (Benkeser et al., 2017). Some DR estimators, such as TMLE, on the other hand, when combined with machine learning fits, have faster convergence and make (asymptotic) analytic statistical inference tractable via the sampling variance of the corresponding influence functions, under conditions on the convergence of the machine learning algorithms used, which, although not guaranteed to hold by any means, are considerably weaker than what would be required to achieve the same with non-DR estimators (Farrell, 2015).

Chernozhukov et al. (2016) further develop these ideas, and suggest the use of machine learning for estimating the nuisance parameters for a wide class of estimating equation DR estimators, to which the g-estimator presented by Vansteelandt and Didelez (2018) belongs.

We implement these two IV estimators, g-estimation and TMLE, with and without machine learning fits, and compare their performance with that of standard TSLS, in terms of mean bias, root mean squared error (RMSE) and confidence interval (CI) coverage using a simulation study. We also contrast the methods in the case study.


This paper is organised as follows. In the next section, we define the causal parameters of interest and the assumptions for the IV methods. In Section 3.1 we review the standard TSLS, while in Section 3.2, we introduce the g-estimator proposed by Vansteelandt and Didelez (2018). Section 3.3 briefly justifies the use of so-called double machine learning methods (data-adaptive DR estimators), and introduces the Super Learner. The TMLE estimator proposed by Tóth and van der Laan (2016) is described in Section 3.4. In Section 4 we present a simulation study, comparing the performance of these estimators. The proposed methods are then applied to the COPERS RCT in Section 5. We conclude with a discussion.

## 2 Linear instrumental variables models

Let  $W$  be a set of baseline variables,  $Z$  be the randomised treatment indicator,  $A$  be the exposure of interest, the actual treatment received, presumed to be binary. Let  $Y$  be the continuous outcome of interest, and  $U$  be the set of all unobserved common causes of  $A$  and  $Y$ . For simplicity, but without loss of generality, assume that we are only interested in effect modification by a single baseline variable  $V \in W$ .

We use a subscript 0 to denote the true probability distributions, models and parameters. Let the vector of the observed data for the  $i$ -th individual be  $O_i = \{W_i, Z_i, A_i, Y_i\} \sim P_0$ , where  $P_0$  is

Figure 1: DAG depicting a valid conditional instrument  $Z$  for exposure  $A$  in the presence of possible effect modifiers  $W$  and unobserved confounders  $U$ , when the outcome is  $Y$ .



the true underlying distribution from which an independent identically distributed random sample of size  $n$  is drawn. The causal relationships between these variables are encoded by the directed acyclic graph (DAG) shown in Figure 1.

Let the counterfactual or potential outcome  $Y(a)$  be the outcome that would occur if  $A$  were set to  $a \in \{0, 1\}$ . As usual, we assume that no interference and counterfactual consistency hold (Rubin, 1978; VanderWeele, 2009). Following Abadie (2003); Vansteelandt and Didelez (2018), we write the conditional version of the IV assumptions (Angrist et al., 1996), as follows:

- (i) **Conditional unconfoundedness:**  $Z$  is conditionally independent of the unmeasured confounders, conditional on measured covariates  $W$ , i.e.  $Z \perp\!\!\!\perp U|W$ .
- (ii) **Exclusion restriction:** Conditionally on  $W$ ,  $A$  and the confounder  $U$ , the instrument  $Z$  and the response  $Y$  are independent, i.e.  $Z \perp\!\!\!\perp Y|W, U, A$ ,
- (iii) **Instrument relevance:** Also referred to as first stage assumption:  $Z$  is associated with  $A$  conditional on  $W$ , i.e.  $Z \not\perp\!\!\!\perp A|W$ .

A special case is the simple RCT, where randomised treatment is often assumed to be an unconditional valid IV.

In addition, we assume the following conditional mean model for the outcome:

$$E[Y|A, W, Z, U] = \omega_0(W, U) + Am_0(W), \quad (1)$$

where  $\omega_0(W, U)$  and  $m_0(W)$  are unknown functions and  $m_0(W)$  represents the causal treatment effect given covariates  $W$ .

Now by the technical assumption of counterfactual consistency, the conditional mean model (1), together with the conditional IV assumptions, implies (Vansteelandt and Didelez, 2018),

$$E[Y|A, W, Z] = E[Y(0)|A, W, Z] + Am_0(W). \quad (2)$$

While this model can be motivated from model (1), it is often used explicitly as the departure point for causal treatment effect estimation. In fact, Vansteelandt and Didelez (2018) show that these two IV models imply the same restrictions on the observed data distribution. Let  $\mathcal{M}$  denote the statistical model for  $P_0$  implied by the IV assumptions and either model (1) or (2). This is often called the linear IV model. Note that in model  $\mathcal{M}$ , the treatment effect in the treated does not depend on  $Z$ . This is known as the ‘no effect modification’ by  $Z$  assumption (Hernán and Robins,

2006).

The causal effect of interest, the average treatment effect in the treated, conditional on  $V \in W$  taking the value  $v$ , can be written as a function of  $v$  as

$$\text{ATT}(v) = E[Y(1) - Y(0) | A = 1, V = v] = E[m_0(W)|A = 1, V = v]. \quad (3)$$

Since  $\text{ATT}(v)$  is the conditional expectation of  $m_0(W)$  given  $A = 1$  and  $V = v$ , we will first focus on estimating  $m_0(W)$ .

Rearranging equation (2) and marginalising over  $A$ , we can see that

$$\begin{aligned} E[Y - Am_0(W)|Z, W] &= E[Y(0)|Z, W] \\ &= E[Y(0)|W], \end{aligned}$$

the last equality holding by the conditional unconfoundedness assumption. Model  $\mathcal{M}$  thus implies

$$E[Y|Z, W] = m_{y0}(W) + m_0(W)E[A|Z, W], \quad (4)$$

where  $m_{y0}(W) = E[\omega_0(W, U)|W]$  or  $m_{y0}(W) = E[Y(0)|W]$ , depending on whether model (1) or (2) is assumed. Therefore, for a binary IV, we have

$$E[Y|Z = 1, W] - E[Y|Z = 0, W] = m_0(W) (E[A|Z = 1, W] - E[A|Z = 0, W])$$

showing that under  $\mathcal{M}$  and the IV assumptions,  $m_0(W)$  is identified by

$$m_0(W) = \frac{E[Y | Z = 1, W] - E[Y | Z = 0, W]}{E[A | Z = 1, W] - E[A | Z = 0, W]}. \quad (5)$$

Estimation of the conditional expectations in equation (5) would typically involve specifying models for the mean exposure  $E[A|Z, W]$  and the mean outcome  $E[Y|Z, W]$ . Let  $\mu_0(Z, W)$  denote  $E[Y|Z, W]$  under  $P_0$ , which according to equation (4) is a function of  $m_{y0}(W)$ ,  $m_0(W)$  and  $E[A|Z, W]$ . We denote by  $\mathcal{A}_y$  the model for  $m_{y0}(W)$ ,  $\mathcal{A}_m$  the model for  $m_0(W)$ , and  $\mathcal{A}_a$  the model for  $E[A|Z, W]$ .

### 3 Doubly robust estimation for the linear instrumental variable model

The methods presented below make different assumptions about the treatment effect model. As we shall see in Section 3.1 and 3.2, TSLS and the g-estimator assume  $m_0(W)$  is a known parametric function of  $V$ , smooth in  $\psi$ , with the true  $\psi_0$  an unknown parameter of finite dimension  $d$ .

To illustrate the methods, we consider throughout a situation where interest lies in the main effect modification by a single variable  $V \in W$ , i.e.  $E[m_0(W)|A = 1, V] = \psi_c + \psi_v V$ . Under this working model for the effect modification by  $V$ , the parameter of interest is  $\psi = (\psi_c, \psi_v)$ , where  $\psi_c$

represents the main causal treatment effect, and  $\psi_v$  is the effect modification by  $V$ . In this case,  $d = 2$ .

The TSLS and g-estimation require the specification of a parametric model for  $m_0(W)$ , so in what follows, these methods presume that this working model is the true model, i.e  $m_0(W)$  takes the parametric form

$$m(W, \psi) = \psi_c + \psi_v V. \quad (6)$$

In contrast, the TMLE method does not assume a parametric model for  $m_0(W)$ , but instead obtains an initial estimate for  $m_0(W)$  using machine learning, which is then updated according to the TMLE algorithm (see details in Section 3.5). This updated estimate is then projected onto a user-supplied parametric working model for the treatment effect curve on  $V$ ,  $m_\psi(V; \psi)$ . Importantly, the working parametric model is not assumed to be the true model for  $m_0(W)$ . Using the working model  $m_\psi(V; \psi) = \psi_c + \psi_v V$ , TMLE also estimates  $\psi = (\psi_c, \psi_v)$ .

### 3.1 TSLS

Estimation of the expectations in equation (5) is often done via an approach known as two-stage least squares (TSLS). The first stage fits a linear treatment selection model, that is a model for  $A$  conditional on the instrument and the baseline covariates of interest, and then, in a second stage, a regression model of the outcome on the predicted treatment received and baseline covariates is fitted. We write

$$E[A|Z, W] = m_{a0}(Z, W) \quad (7)$$

$$E[Y|Z, W] = m_{y0}(W) + m_0(W)m_{a0}(Z, W), \quad (8)$$

In principle, there are many parametric choices for the second stage models,  $\mathcal{A}_y$  and  $\mathcal{A}_m$ . For TSLS to be consistent, the first stage model  $\mathcal{A}_a$  must be the parametric linear regression implied by the second stage, that is, it must include all the covariates and interactions that appear in the second stage model.

For example, if we assume that  $\mathcal{A}_m : m(W; \psi) = \psi_c + \psi_v V$ , and  $\mathcal{A}_y : m_y(W; \beta) = \beta^\top W$ , where abusing notation we assume the vector of ones is the first column of  $W$ , then the first-stage would involve two equations, as follows

$$\begin{aligned} E[A|Z, W] &= \alpha_z Z + \alpha_{zv} ZV + \alpha_v V + \sum_{W_i \in W \setminus V} \alpha_{wi} W_i \\ E[AV|Z, W] &= \lambda_z Z + \lambda_{zv} ZV + \lambda_v V + \sum_{W_i \in W \setminus V} \lambda_{wi} W_i \end{aligned} \quad (9)$$

where again,  $W$  includes 1 to allow for an intercept. Because estimation of these two first-stage models is done separately without acknowledging that the model for  $A$  should imply the model for  $AV$ , the resulting TSLS estimator may be inefficient (Vansteelandt and Didelez, 2018).

Vansteelandt and Didelez (2018) show that the second stage of the TSLS, equation (8), can be

written as the solution to an estimating equation

$$0 = \sum_{i=1}^n e_y(Z_i, W_i) \{Y_i - m_{y0}(W_i; \beta) - m_0(W_i; \psi)m_{a0}(Z_i, W_i; \alpha_0)\}, \quad (10)$$

with  $e_y(Z_i, W_i)$  any conformable index vector function of dimension  $\dim(\beta) + d$ . The estimators  $\hat{\beta}$  and  $\hat{\psi}$  obtained solving equation (10), after substituting  $\hat{\alpha}$  for  $\alpha_0$  (the estimator from the first stage), are consistent asymptotically normal (CAN), under  $\mathcal{A}_y \cap \mathcal{A}_m$ , i.e. even when  $\mathcal{A}_a$  is misspecified (Robins, 2000; Wooldridge, 2010). Moreover, in the specific settings where the treatment effect model  $\mathcal{A}_m$  is linear in the covariates and the instrument is independent of  $W$ , TSLS is also robust to misspecification of  $\mathcal{A}_y$ . We refer the interested reader to Vansteelandt and Didelez (2018), Appendix B Proposition 5, for the proof.

This means that for estimators which are DR in the more general settings, with either a non-linear  $\mathcal{A}_m$  or where  $Z$  depends on  $W$ , methods beyond TSLS need to be considered.

### 3.2 G-estimation

Okui et al. (2012) propose an estimator for  $\psi = (\psi_c, \psi_v)$ , constructed by augmenting the estimating equation (10), by multiplying the summand by  $Z_i - E[Z|W_i]$ . The model  $\mathcal{A}_z$  for the conditional distribution of the IV  $E[Z|W] := P_0(Z = 1|W)$  is often called the instrument propensity score, and it is assumed to be a known mass function  $g_0(W; \gamma)$ , smooth on a finite dimensional parameter  $\gamma$ .

More specifically, the IV g-estimator of the causal parameter  $\psi_0$  is obtained as a solution to the following estimating equation (Okui et al., 2012)

$$0 = \sum_n (e(Z_i, W_i; \gamma_0) - E[e(Z_i, W_i; \gamma_0)|W_i]) \{Y_i - m_{y0}(W_i; \beta_0) - m_0(W_i; \psi)A_i\}, \quad (11)$$

where  $e(Z, W)$  is any conformable vector function (i.e. of the appropriate dimension).

This can be made (locally) efficient by choosing (Vansteelandt and Didelez, 2018)

$$e(Z, W; \gamma_0) = \sigma_0^{-2}(Z, W) \begin{pmatrix} 1 \\ V \end{pmatrix} \left[ m_{a0}(Z, W; \alpha_0) - \frac{E_{g(W; \gamma_0)}[\sigma_0^{-2}(Z, W)m_{a0}(Z, W; \alpha_0)|W]}{E_{g(W; \gamma_0)}[\sigma_0^{-2}(Z, W)|W]} \right] \quad (12)$$

and  $\sigma_0^2(Z, W) = \text{Var}\{Y - Am(W; \psi)|Z, W\}$ .

This estimator (denoted by IV-g) can be made feasible by replacing  $\alpha_0$ ,  $\beta_0$  and  $\gamma_0$  by their corresponding consistent estimators  $\hat{\alpha}$ ,  $\hat{\beta}$  and  $\hat{\gamma}$ , and setting  $\sigma_0^2$  equal to 1 (as it is just a proportionality constant). It has been shown to be CAN under  $\mathcal{A}_m \cap \{\mathcal{A}_z \cup \mathcal{A}_y\}$ , and hence consistent whenever the model for  $m_0(W)$  is correctly specified, and either the instrument propensity score or outcome model are correctly specified (Okui et al., 2012). The addition of the instrument propensity score model is particularly helpful when the dependence between  $Z$  and the baseline covariates is known, as would be the case in a randomised trial, thus guaranteeing robustness against misspecification of the outcome model.

Moreover, the IV-g estimator is efficient when  $\mathcal{A}_y$ ,  $\mathcal{A}_a$  and  $\mathcal{A}_z$  are correctly specified (Vansteelandt and Didelez, 2018).

landt and Didelez, 2018), a property referred to as ‘locally efficiency’.

Since the IV g-estimator is CAN, its asymptotic variance is the variance of its influence function, denoted  $D(\psi)$  (Newey, 1990), i.e.  $\text{Var}(\psi) = E[D(\psi)^\top D(\psi)]$ .  $D(\psi)$  can be written as

$$D_i(\psi) = M(Z_i, W_i, A_i)^{-1} K(Z_i, W_i) \begin{pmatrix} 1 \\ V_i \end{pmatrix} \{Y_i - m_y(W_i; \beta)\} - \begin{pmatrix} \psi_c \\ \psi_v \end{pmatrix} \quad (13)$$

with

$$K(Z, W) = m_{a0}(Z, W; \alpha_0) - E_{g(W; \gamma_0)}[m_{a0}(Z, W; \alpha_0) | W] \quad (14)$$

and

$$M(Z, W, A) = AK(Z, W) \begin{pmatrix} 1 & V \\ V & V^2 \end{pmatrix} \quad (15)$$

Finally, the variance can be estimated by the sample variance of the estimated influence function, obtained by replacing  $\alpha_0$ ,  $\beta_0$  and  $\gamma_0$  by their corresponding consistent estimators  $\hat{\alpha}$ ,  $\hat{\beta}$  and  $\hat{\gamma}$ ,

$$\hat{V}(\hat{\psi}) = n^{-1} \text{Var}_n(\hat{D}(\hat{\psi})),$$

where we have used the subscript  $n$  to denote the sample variance on a sample of size  $n$ .

This variance estimator ignores the nuisance parameter estimation. Robust standard errors can be obtained via the bootstrap or a sandwich estimator.

### 3.3 Data-adaptive estimation for the IV-g estimator

The IV-g estimator presented thus far is restricted to using parametric working models for the nuisance parameters. In this section, we will relax this restriction by using instead machine learning methods to obtain data-adaptive estimates of the nuisance functionals and plugging these fits into the estimating equation (11). Unlike fixed parametric models where the functional form is chosen by the analyst, these methods use an algorithm, often called a learner, to find the best fitting estimator. In fact, we need not restrict ourselves to a single learner, instead we use meta-learning or ensemble approaches, in particular the Super Learner (SL) (van der Laan et al., 2007). The SL uses cross validation to find the optimal weighted convex combination of multiple candidate estimators specified by the user in the SL *library*. The library can include parametric and non-parametric estimators. The SL has been shown to perform asymptotically as well as the best learner included in its library, so that adding additional algorithms improves the performance of the SL. The finite sample performance of the SL has been demonstrated extensively in simulations (van der Laan et al., 2007; Porter et al., 2011; Pirracchio et al., 2015).

Machine learning estimation approaches for DR estimators for the IV model have been studied recently. Chernozhukov et al. (2016) give sufficient conditions to guarantee that using data-adaptive fits for the nuisance functionals in a DR estimator results in valid inferences. Briefly, the score function  $\mathcal{S}$  of the DR estimator needs to be Neyman orthogonal to the nuisance parameters i.e. the path-wise (or Gateaux) derivative of the score function exists and vanishes at the true value of



the nuisance parameters. Then, as long as the data-adaptive estimators for all nuisance functionals converge to their respective truths, and the product of their convergence rates is faster than  $n^{-\frac{1}{2}}$ , the DR estimator will be CAN. Convergence rates for these data-adaptive estimators depend on the smoothness and number of covariates included (Györfi et al., 2006).

Under sufficient regularity conditions, and if  $\mathcal{A}_a$  is estimated using parametric models, the arguments used in Chernozhukov et al. (2016) Section 4.2 can be applied to show that the IV-g estimator is still CAN when using data-adaptive fits for  $\mathcal{A}_y$  and  $\mathcal{A}_z$  as long as:

$$\|\hat{g}(W) - g_0(W)\|_{P_{0,2}} \times \left( \|\hat{m}_a(W) - m_{a0}(W)\|_{P_{0,2}} + \|\hat{m}_y(W) - m_{y0}(W)\|_{P_{0,2}} \right) < O(n^{-\frac{1}{2}}), \quad (16)$$

where  $\|\cdot\|_{P_{0,2}}$  denotes the  $L^2$  norm defined by  $P_0$ . We refer the interested reader to Chernozhukov et al. (2016) for the technical details.

The IV-g estimator implemented here estimates  $\mathcal{A}_m$  and  $\mathcal{A}_y$  jointly using parametric models that include only main terms for both, but uses data-adaptive fits for  $\mathcal{A}_a$  and the instrument propensity score. Thus the rates in condition (16) are not quite sufficient. We conjecture that in addition we need that  $\|\hat{m}_a(Z, W) - m_{a0}(Z, W)\|_{P_{0,2}} < O(n^{-\frac{1}{4}})$ .

The use of data-adaptive fits for nuisance functionals has been extensively exploited within the TMLE literature which we review next.

### 3.4 Targeted minimum loss estimation (TMLE)

Targeted minimum loss estimation (TMLE) is a general approach for causal inference, which has been adopted on a wide range of causal problem (Gruber and van der Laan, 2010; van der Laan and Rose, 2011; Zheng and van der Laan, 2012; van der Laan and Gruber, 2012; Petersen et al., 2014).

TMLE is a semi-parametric, double-robust estimation approach, that combines estimates of the nuisance functional and outcome mechanism in a way that guarantees a consistent estimator of the target parameter if either the nuisance parameters or the outcome mechanism are estimated consistently. If both are estimated consistently, TMLE is efficient (van der Laan and Rose, 2011). These initial estimates can be obtained by specifying parametric models or, most commonly via machine learning, in particular the Super Learner (van der Laan et al., 2007). The use of data-adaptive estimates reduces bias, and improves the chances of achieving efficiency, and accurate statistical inference.

TMLE estimators for the linear IV model have been recently proposed by Tóth and van der Laan (2016). We describe below in more detail the linear non-iterative TMLE, which we denote by IV-TMLE.

### 3.5 IV-TMLE

Let  $\Psi : \mathcal{P} \mapsto \mathbb{R}^2$  be the target parameter mapping, from the space of all possible models for the true distribution of the data  $P_0$  to  $\mathbb{R}^2$ , defined by the treatment effect modification working parametric

model  $m_\psi = \psi_c + \psi_v V$ ,  $\Psi(P_0) = (\psi_c, \psi_v) = \psi_0$  as the solution to

$$E \left[ \begin{pmatrix} 1 \\ V \end{pmatrix} \{m_0(W) - (\psi_c + \psi_v V)\} \right] = 0.$$

We note that  $\Psi$  only depends on  $P_0$  through  $m_0$  and the distribution of the covariates  $P_{W,0}$ . We denote this relevant part by  $Q_0 = Q_0(P_0) = (m_0, Q_W)$ , with  $Q_W = P_{W,0}$ , while under  $\mathcal{M}$ ,  $m_0(W)$  depends on  $\mu_0(Z, W)$  and  $m_{a0}(Z, W)$ , so that  $Q_0 = (\mu_0(Z, W), m_{a0}(Z, W), Q_W)$ .

The TMLE algorithm starts by obtaining initial estimates of  $\mu_0(Z, W)$ ,  $m_{a0}(Z, W)$ , and the instrument propensity score  $g_0(W)$ . We denote these initial estimates by writing a hat over the object.

The next step in the construction of a TMLE requires the specification of a loss function  $L(P)$ , such that  $E[L(P_0)] = \min_{P \in \mathcal{P}} P_0(L(P))$ , i.e. the expectation of the loss function is minimised at the true probability distribution. For the linear-fluctuation IV-TMLE, we use the square error loss function  $L(P) = E[D^*(\Psi(P))^\top D^*(\Psi(P))]$ , where  $D^*(\Psi(P_0))$  is the efficient influence function (EIF) under the linear IV model  $\mathcal{M}$  and the working model  $m_\psi = \psi_c + \psi_v V$ , which can be written as (see Tóth and van der Laan (2016) for the details):

$$\begin{aligned} D^*(\Psi(P_0)) &= c \zeta^{-2}(W) \{m_{a0}(Z, W) - E_0(m_{a0}(Z, W)|W)\} \{Y - m_{a0}(Z, W)m_0(W) - m_{y0}(W)\} \\ &\quad - c \zeta^{-2} \{(m_{a0}(Z, W) - E_0[m_{a0}(Z, W)|W])m_0(W)\} (A - m_{a0}(Z, W)) + D_W(Q_W), \end{aligned}$$

where  $D_W(Q_W) = c \{m_0(W) - m_\psi(V)\}$  and

$$c = \text{Var}(V)^{-1} \begin{pmatrix} E[V^2] - E[V]V \\ V - E[V] \end{pmatrix}.$$

The term  $\zeta^2(W)$  is associated with instrument strength, and can be written as

$$\zeta^2(W) = \text{Var}_{g_0(W)}[m_{a0}(Z, W)|W] = m_{a0}(1, W)^2 \text{Var}(Z|W) = m_{a0}(1, W)^2 g_0(W)(1 - g_0(W)). \quad (17)$$

Now, we proceed to the targeting step, using a linear fluctuating model that updates the initial fit for  $m_0$ . For this, initial fits for  $m_y(W)$ ,  $m_a(W)$  and  $\mu_0(Z, W)$  are used in model (4) to define the initial estimate for  $m_0$ ,  $\widehat{m}(W)$ . The linear fluctuating model is now:

$$m^*(\epsilon)(W) = \widehat{m}(W) + h(W)^T \epsilon, \quad (18)$$

where the so-called clever covariate is defined as

$$h(W) = \text{Var}(V)^{-1} \begin{pmatrix} E[V^2] - E[V]V \\ V - E[V] \end{pmatrix} \zeta^{-2}. \quad (19)$$

With these choices for  $L$  and the fluctuating model,  $\epsilon$  can be found by solving the EIF equation,

$$\frac{1}{n} \sum_{i=1}^n D^* \{m^*(\epsilon)(W)\} = 0.$$

This is equivalent to solving for  $\epsilon$  a system of  $d$  linear equations:

$$\frac{1}{n} \sum_{i=1}^n \hat{h}(W) \{\hat{m}_a(Z, W) - E_{\hat{g}(W)}[\hat{m}_a(Z, W)|W]\} \left\{ Y - A \left( \hat{m}(W) + \hat{h}(W)^\top \epsilon \right) - \hat{m}_y(W) \right\} = 0, \quad (20)$$

where  $\hat{h}(W)$  is obtained by plugging in the initial fits  $\hat{m}_a(Z, W)$  and  $\hat{g}(W)$  into equation (19). The TMLE estimator of  $m_0(W)$ , denoted by  $m^*$ , is now obtained by substituting  $\epsilon^*$ , the solution to equation (20) into equation (18). Finally, we project  $m^*$  onto the working model  $m_\psi$  by OLS, obtaining  $(\psi_c^*, \psi_v^*)$ , the TMLE estimator of the causal parameter of interest.

Tóth and van der Laan (2016) showed that this approach yields an estimator which is consistent when (i) the initial estimators of  $m_{a0}$  and  $g_0(W)$  are consistent, (ii) the initial estimators of  $m_0$  and  $g_0(W)$  are consistent, or (iii) the initial estimators of  $m_0$  and  $m_{y0}$  are consistent.

### 3.6 Comparison of the two estimators

Recall that the IV-g estimator assumes a parametric model for  $m_0(W)$ , e.g.  $m_0(W; \psi) = \psi_c + \psi_v V$  (equation (6)), while the IV-TMLE obtains a first estimate of  $m_0(W)$  which can include all other covariates in  $W$  not only  $V$ , and then after updating this through the TMLE procedure projects the targeted estimator onto a user-supplied working parametric model, e.g.  $m_\psi(W) = m(W; \psi) = \psi_c + \psi_v V$ , which is not assumed to be correctly specified.

The IV-g estimator gains efficiency from assuming equation (6) when equation (6) holds. Suppose however that the  $\mathcal{A}_m$  model (6) is misspecified (e.g. that the treatment effect curve depends on more covariates, not just  $V$ , or that the relationship is not linear). Interestingly, the IV-g estimator will continue to be a consistent estimator of the parameters  $\psi_c, \psi_v$ , whenever the mean exposure model  $m_{a0}(Z, W)$  is correctly specified and  $\text{Cov}(\{m_{a0}(Z, W) - E(m_{a0}(Z, W)|W)\}, A|W)$  is constant in  $W$ . This is because the IV-g estimator solves

$$E[\{m_{a0}(Z, W) - E(m_{a0}(Z, W)|W)\} \{Y - m_{y0}(W) - Am(W; \psi)\}] = 0 \quad (21)$$

This is equivalent to

$$\begin{aligned} 0 &= E[\{m_{a0}(Z, W) - E(m_{a0}(Z, W)|W)\} \{Y - m_{y0}(W) - Am(W; \psi)\}] \\ &= E[\{m_{a0}(Z, W) - E(m_{a0}(Z, W)|W)\} \{m_{y0}(W) + Am_0(W) - m_{y0}(W) - Am(W; \psi)\}] \\ &= E[\{m_{a0}(Z, W) - E(m_{a0}(Z, W)|W)\} \{A(m_0(W) - m(W; \psi))\}] \end{aligned}$$

so that,

$$m(W; \tilde{\psi}) = \text{plim } m(W; \hat{\psi}) = \frac{E[(m_{a0}(Z, W) - E(m_{a0}(Z, W)|W)) Am_0(W)]}{E[(m_{a0}(Z, W) - E(m_{a0}(Z, W)|W)) A]} \quad (22)$$

This converges to  $E[m_0(W)]$  if  $\text{Cov}(\{m_{a0}(Z, W) - E(m_{a0}(Z, W)|W)\}, A|W)$  is constant, and continues to hold in large samples, when  $m_{a0}(Z, W)$ ,  $m_{y0}(W)$  and the instrument propensity score  $g_0(W)$  are replaced by their consistent estimators respectively.

If however  $\text{Cov}(\{m_{a0}(Z, W) - E(m_{a0}(Z, W)|W)\}, A|W)$  is not constant in  $W$ , then the IV-g estimator under a misspecified  $\mathcal{A}_m$  converges to a weighted average of the true estimand of interest.

### 3.6.1 Weak instrument instability for IV-TMLE

The other difference between these estimators is the reliance (or not) on the exposure model. In particular, the clever covariate  $h(W)$  used to define the IV-TMLE estimator depends on the inverse of  $\text{Var}_{g_0(W)}(m_{a0}(Z, W)|W)$ . This captures the strength of the instrument in predicting the exposure given  $W$ , and the variance of the IV-TMLE estimators become very large when this is very small. To stabilise the estimators, we choose the maximum of the estimated value of  $\text{Var}_{\hat{g}(W)}(\hat{m}_a(Z, W)|W)$ , and 0.025 when constructing the clever covariate for a given data set. Such corrections are not necessary for the IV-g estimator.

## 4 Simulation Study

We perform a small factorial simulation study to assess the finite sample performance of the alternative methods, under the different combinations of  $\mathcal{A}_y$ ,  $\mathcal{A}_a$  or  $\mathcal{A}_m$  being in turn correctly specified or not, while the instrument model is always correct. We write  $\mathbb{1}(\mathcal{A}_k \neq m_{k0})$  as an indicator function, for scenarios where the assumed model indexed by  $k$  is misspecified.

We generate the RCT data, with two-sided non-adherence, i.e. both randomly allocated groups have a non-zero probability of receiving the opposite treatment. We begin by generating five independent standard normal variables  $W_1, \dots, W_4$  and  $V$ . These are the observed baseline covariates, of which one is the effect modifier  $V$ . We also generate a standard normal unobserved confounder  $U$ . We generate randomised treatment also independently of the other variables,  $Z \sim \text{Bern}(0.6)$ , and then simulate the binary treatment received  $A$  as a function the baseline variables, the unobserved confounder, and the instrument:

$$\text{logit}(\text{Pr}(A = 1|Z, W, U)) = (1.5Z + 0.03V + 0.01W_1 + 0.01W_2 + 0.01W_3 + 0.01W_4 - \mathbb{1}(\mathcal{A}_a \neq m_{a0})(5ZW_1 + 0.03U)).$$

Notice that we are generating  $A$  in such a way that the condition necessary for the IV-g estimator to converge to the parameter of interest when  $\mathcal{A}_m$  is wrong is no longer satisfied.

The continuous outcome  $Y$  is then simulated as follows:

$$\begin{aligned}
m_y(W) &= (1 - \mathbb{1}(\mathcal{A}_y \neq m_{y0}))(0.5 + 0.5V + 0.01W_1 + 0.01W_2 + 0.01W_3 + 0.01W_4) \\
&+ \mathbb{1}(\mathcal{A}_y \neq m_{y0})\exp\{0.05 + 0.05V + 0.001W_1 + 0.001W_2 + 0.001W_3 + 0.001W_4 \\
&- 0.2V(W_1 + W_2 + W_3 + W_4)\} \\
m(W) &= 0.5 + 0.5V + \mathbb{1}(\mathcal{A}_m \neq m_0)3(W_1 + W_2 + W_3 + W_4) \\
Y &\sim N(m_y(W) + m(W)A + U, 1),
\end{aligned}$$

We simulate RCTs with two different sample sizes,  $n = 500$  and  $n = 10,000$ , and generate 1,000 replicates for each scenario. We perform analyses with TSLS, IV-g and IV-TMLE. The latter two are implemented with and without the use of SL.

Parametric TSLS, IV-g and TMLE are fitted using the main terms only. For TSLS the first stage are as per equation (9), while for IV-g and TMLE, these are logistic models. The IV propensity score is also a logistic model. The SEs of the parametric IV-g and TMLE are obtained by bootstrapping (percentile 95% CI using 1999 bootstrap samples).

For SL fits in IV-g, since  $A$  and  $Z$  are binary, we use a library including glm (generalised linear models), step (stepwise model selection using AIC), svm (support vector machines), gam (generalised additive models). For the glm, step and gam learners, linear and second-order terms are used.

For the IV-TMLE, we use the same library as above for  $g(W)$  and  $m_a(W)$ , while for  $E[Y|Z, W]$  and  $m_y(W)$ , we include only glm, step, svm and polymars (multivariate adaptive polynomial spline regression), in order to preserve the linear structure of model (4).

The methods described in Chernozhukov et al. (2016) require the use of sample splitting in order to guarantee valid inference when using machine learning for fitting the nuisance models. We do not perform such sample splitting for the IV-g or the TMLE estimators, other than that already being used for the cross validation embedded in the Super Learner. A limitation with this pragmatic approach, in that the variance estimator based on the (E)IF ignores the fact that these factors are not independent from each other, as we have used the data to both obtain the fits and then evaluate the (E)IF. van der Laan and Luedtke (2015) suggest using cross-validation. For example, if using 10-fold cross-validation we can estimate the variance,

$$\hat{V} = \frac{1}{10} \sum_{k=1}^{10} \frac{1}{np} \sum_{i \in N(k)} D(\psi)^\top D(\psi)(O_i),$$

where  $p$  is the proportion of observations in each validation sample  $N(k)$ , and  $D(\psi)$  is the IF (or the EIF) of the parameter of interest.

We compute the mean bias of the estimates, coverage of 95% confidence intervals (CI), and root mean square error (RMSE).

## 4.1 Results from the simulation

Figure 2 shows the mean bias (top) and CI coverage rate (bottom) corresponding to scenarios with sample size of  $n = 500$ . The corresponding results when  $n = 10,000$  appear in Figure 3.

As expected, when all models are correctly specified, all methods show close to zero bias for both of the target parameters, already with small ( $n = 500$ ) sample size, with coverage levels close to the nominal (between 92.5 and 97.5%) for TSLS and IV-g. While the bootstrap CI corresponding to parametric TMLE resulted in over-coverage (99%), the SL IV-TMLE shows some under-coverage, which improves at large ( $n = 10,000$ ) sample. This is a consequence of the variance estimator used here ignoring the fact that the data have been used to both obtain the fits and then evaluate the EIF (instead of using sample splitting). The RMSE were indistinguishable across the methods. This pattern held both when the methods were implemented using parametric models, and when with the SL.

TSLS performs well when  $\mathcal{A}_m$  is correctly specified regardless of the other models, but performs very poorly when  $\mathcal{A}_a$  and  $\mathcal{A}_m$  are misspecified (bias 200 times the size of the effect and 0 coverage).

Under those misspecified scenarios when the double robust properties were expected to provide protection, these small levels of bias, and good coverage has generally held. Adding a misspecified treatment effect model  $\mathcal{A}_m$  has doubled RMSE across the methods, but has not affected bias or coverage, and this held true when a misspecified  $\mathcal{A}_y$  model was used, as long as the exposure model  $\mathcal{A}_a$  was still correctly specified. However, when the exposure model was misspecified (the last two columns of the Figures), the estimators started showing sensitivity. When the treatment effect model  $\mathcal{A}_m$  and the exposure model  $\mathcal{A}_a$  were both misspecified, all of the parametric estimators considered showed high levels of bias. When implementing the IV-g and IV-TMLE methods using the Super Learner, the bias and coverage, and consequently the RMSE, returned to the levels reported under correct specification. An exception is the coverage for the IV-g estimator of  $\psi_c$ , which remained very low (<50%). This is mostly due to the substantial remaining bias, stemming from the fact that the estimator converges to a weighted average of the true  $m_0(W)$  when the parametric  $\mathcal{A}_m$  is mis-specified (equation (22)). When  $\mathcal{A}_y$  is also misspecified, we observe essentially the same levels of bias, RMSE and coverage, as under the previously described scenario. Across the scenarios, the effect of sample size was not significant in the reported bias, RMSE and coverage levels, except for the scenarios when  $\mathcal{A}_m$  and  $\mathcal{A}_a$  were misspecified. Here, the IV-g estimator with parametric implementation reported extremely high bias, RMSE and Monte Carlo errors. With large sample size, the bias was smaller, although still substantial, and the Monte Carlo errors reduced to close to zero.

We summarise the results of the simulation study by comparing estimated RMSEs, Figure 4. Under  $n = 10,000$ , we conclude that both DR estimators have reported performance according to their theoretical double-robust properties, and the TSLS method showed similar performance to the parametric implementation of the IV-g method. Both DR methods have benefitted from the data-adaptive estimation of the nuisance parameters: the performance of the estimators have not been harmed in the correctly specified scenarios, and RMSE has been greatly reduced in the

Figure 2: Performance (Bias and Coverage) of TSLS, TMLE and IV-g estimators, when the sample size is ( $n = 500$ ). Scenarios with correct or misspecified  $\mathcal{A}_a$  and  $\mathcal{A}_y$  vary by column. Settings where the parametric  $\mathcal{A}_m$  is correct are plotted in blue, while  $\mathcal{A}_m$  mis-specified is plotted in red. The hollow shapes correspond to estimates obtained with parametric models, while the solid shapes correspond to SL. The true values of  $(\psi_c, \psi_v)$  are  $(0.5, 0.5)$ , the bias is presented with its Monte Carlo error CIs. Dotted line in the bias plot is the 0 line, the dashed lines in the coverage plot are the 92.5 and 97.5 % coverage rates

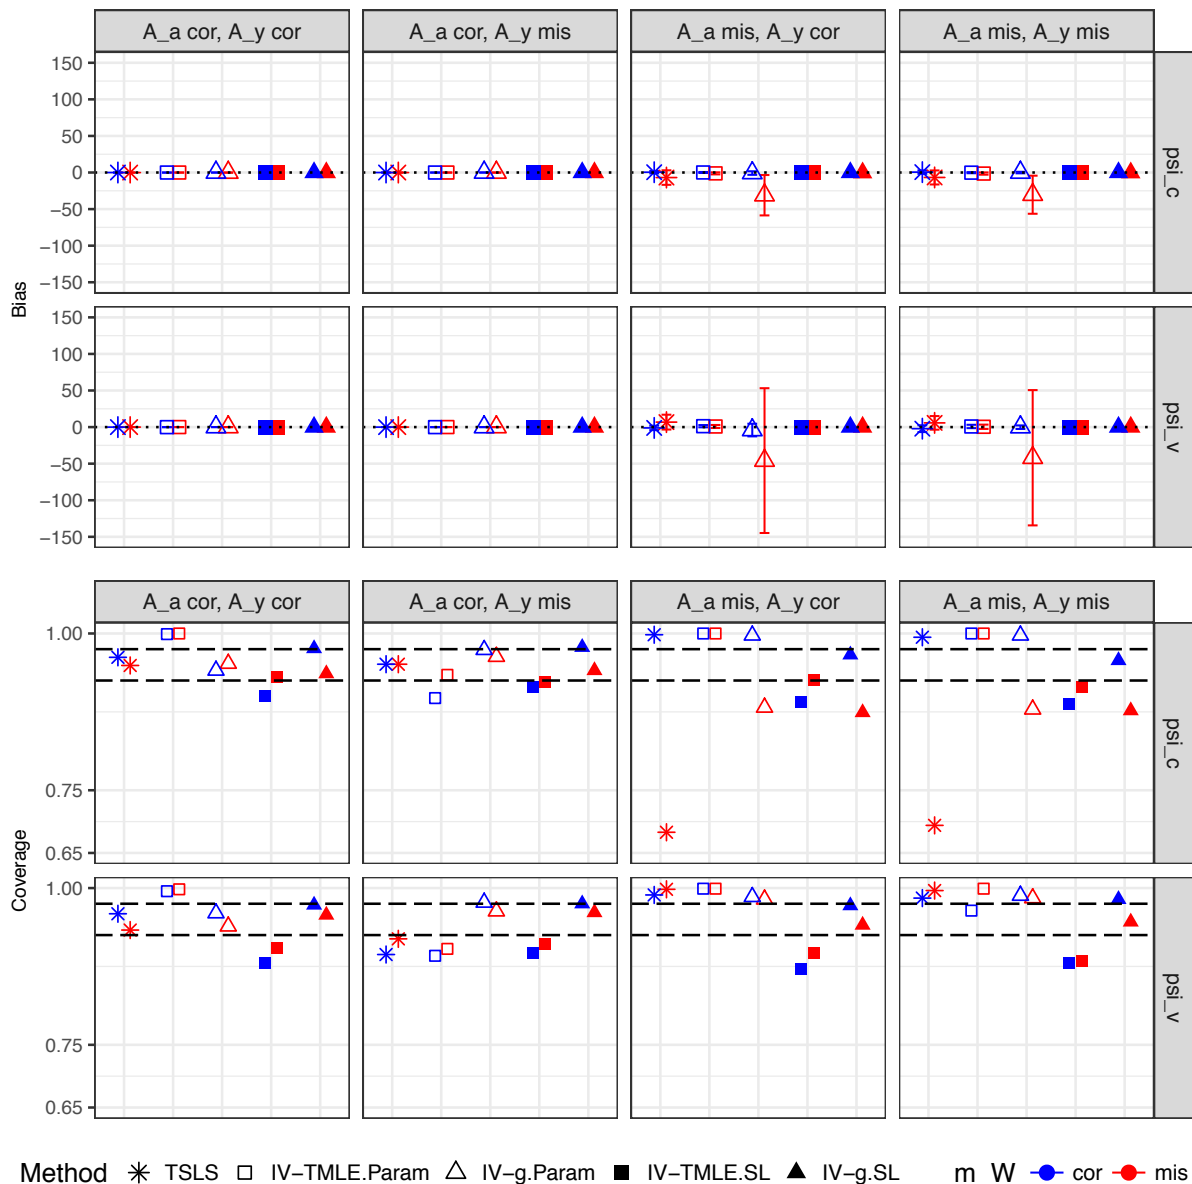


Figure 3: Performance (Bias and Coverage) of TSLS, TMLE and IV-g estimators, when the sample size is  $(n = 10,000)$ . Scenarios with correct or misspecified  $\mathcal{A}_a$  and  $\mathcal{A}_y$  vary by column. Settings where the parametric  $\mathcal{A}_m$  is correct are plotted in blue, while  $\mathcal{A}_m$  mis-specified is plotted in red. The hollow shapes correspond to estimates obtained with parametric models, while the solid shapes correspond to SL. The true values of  $(\psi_c, \psi_v)$  are  $(0.5, 0.5)$ , the bias is presented with its Monte Carlo error CIs. Dotted line in the bias plot is the 0 line, the dashed lines in the coverage plot are the 92.5 and 97.5 % coverage rates.

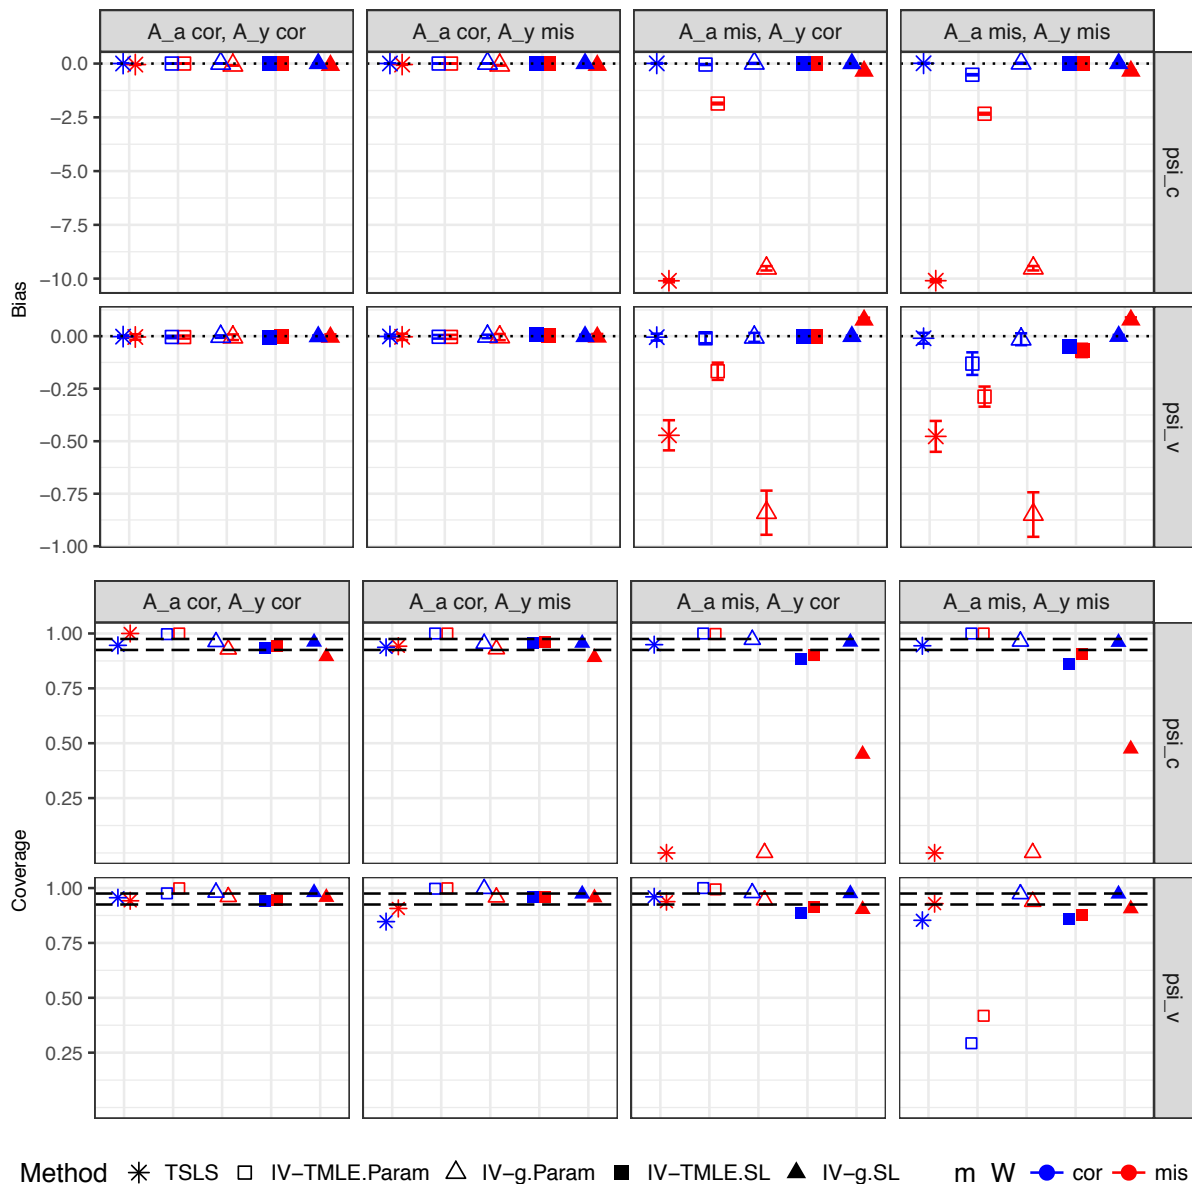
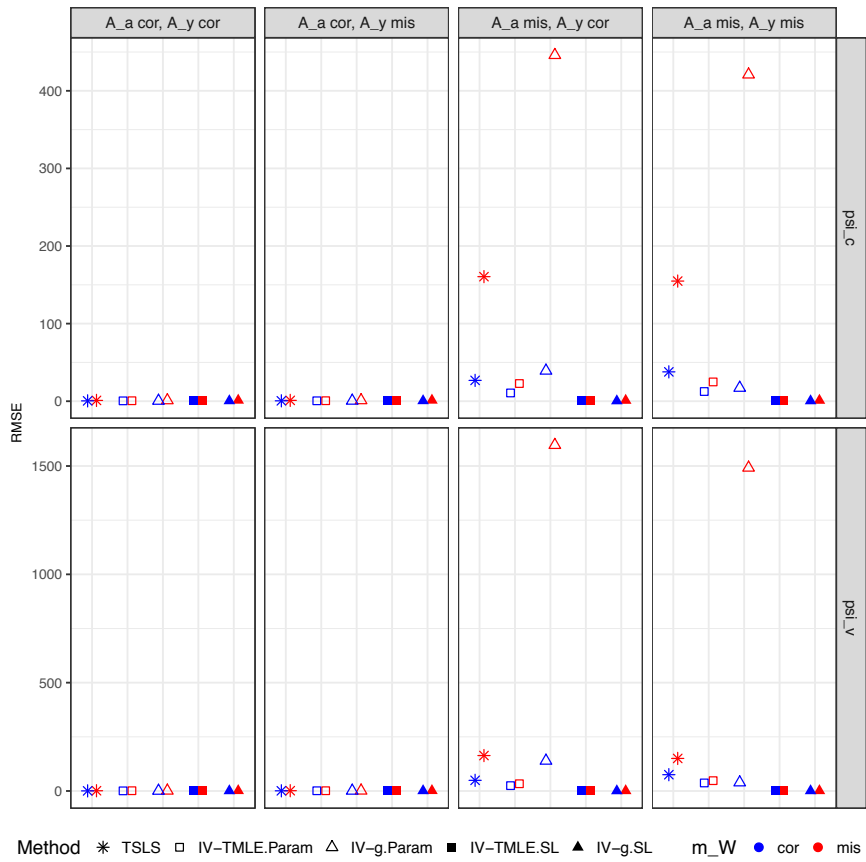





Figure 4: RMSE of TSLS, TMLE and IV-g estimators, when the sample size is (a)  $n = 500$  (top) and (b)  $n = 10,000$ . Scenarios with correct or misspecified  $\mathcal{A}_a$  and  $\mathcal{A}_y$  vary by column. Settings where the parametric  $\mathcal{A}_m$  is correct are plotted in blue, while  $\mathcal{A}_m$  mis-specified is plotted in red. The hollow shapes correspond to estimates obtained with parametric models, while the solid shapes correspond to SL.



(a)  $n=500$



(b)  $n=10,000$

scenarios when the DR properties did not provide protection against misspecification. When the assumed model  $\mathcal{A}_m$  is correctly specified, IV-g outperforms all other methods, with the smallest RMSE. In contrast, when  $\mathcal{A}_m$  is misspecified, IV-TMLE performs best.

## 5 Motivating example: the COPERS trial

We now illustrate the methods in practice by applying each in turn to the motivating example. The COping with persistent Pain, Effectiveness Research in Self-management trial (COPERS) was a randomised controlled trial across 27 general practices and community services in the UK. It recruited 703 adults with musculoskeletal pain of at least 3 months duration, and randomised 403 participants to the active intervention and a further 300 to the control arm. The mean age of participants was 59.9 years, with 81% white, 67% female, 23% employed, 85% with pain for at least 3 years, and 23% on strong opioids.

Intervention participants were offered 24 sessions introducing them to cognitive behavioural (CB) approaches designed to promote self-management of chronic back pain. The sessions were delivered over three days within the same week with a follow-up session 2 weeks later. At the end of the 3-day course participants received a relaxation CD and self-help booklet. Controls received usual care and the same relaxation CD and self-help booklet.

The primary outcome was pain-related disability at 12 months, using the Chronic Pain Grade (CPG) disability sub-scale. This is a continuous measure on a scale from 0 to 100, with higher scores indicating worse pain-related disability.

In the active treatment, only 179 (45%) attended all 24 sessions, and 322 (86.1%) received at least one session. The control arm participants had no access to the active intervention sessions. Participants and group facilitators were not masked to the study arm they belonged to.

The intention-to-treat analysis found no evidence that the COPERS intervention had an effect on improving pain-related disability at 12 months in people with long-established, chronic musculoskeletal pain ( $-1.0$ , 95% CI  $-4.8$  to  $2.7$ ).

Poor attendance to the sessions was anticipated, and so obtaining causal treatment effect estimates was a pre-defined objective of the study. The original report included a causal treatment effect analysis using TSLS, using a binary indicator for treatment received (attending at least half of the sessions), and assuming that randomisation was a valid instrument for treatment received (Taylor et al., 2016). The IV model adjusted for the following baseline covariates: site of recruitment, age, gender and HADS score and the CPG score at baseline. This IV analysis found no evidence of a treatment effect on CPG at 12 months amongst the compliers ( $-1.0$ , 95% CI  $-5.9$  to  $3.9$ ).

The COPERS study also performed a number of subgroup analyses to investigate treatment effect heterogeneity, but did not carry out IV analysis with effect modification. However, treatment heterogeneity in the causal effect is still of interest.

For our re-analyses, the data set consists of 652 participants followed up for 12 months, 374 allocated to active treatment, and 278 in the control (93% of those recruited). Thirty-five individuals

(5%) have missing primary outcome data, and a further 4 (<1%) have missing baseline depression score, leaving a sample size of 613.

We focus on the causal effect of receiving at least one treatment session as a function of depression at baseline measured using the Hospital Anxiety and Depression Scale (HADS).

We argue that random allocation is a valid IV: the assumptions concerning unconfoundedness and instrument relevance are justified by design. The exclusion restriction assumption seems plausible with our choices for  $A$ , as only those participants receiving at least one training sessions would know how to use the CB coping mechanisms and potentially to improve their disability. It is unlikely that that random allocation has a direct effect, though since participants were not blinded to their allocation, we cannot completely rule out some psychological effects of knowing one belongs to the control or active group on pain and disability.

We perform each of the methods in turn, TSLS, IV-g and IV-TMLE to estimate  $ATT(v)$ . As Table 1 summarises, the use of DR methods, even after using Super Learner does not result in a material change in the point estimates or SEs, compared to standard TSLS. All 5 estimators give the same inference, namely, there is no evidence of an average treatment effect in the treated, and there is no evidence of effect modification by baseline depression. The lack of statistical significance may be due to small numbers of participants in the trial. While not statistically significant, the direction of the treatment effect modification is interesting: the point estimates indicate the treatment may benefit more those with higher depression symptoms at baseline (the ATT resulting in a reduction in the disability score).

Table 1: ATT of the COPERS intervention on CPG, with all-or-nothing binary exposure  $A$ , main effect  $\psi_c$  and effect modification by depression)  $\psi_v$ .

	$\psi_c$	SE	$\psi_v$	SE
TSLS	2.94	4.67	-0.58	0.57
IV-g	2.78	4.66	-0.53	0.54
IV-g SL	2.10	4.75	-0.45	0.54
IV-TMLE	3.16	4.74	-0.64	0.56
IV-TMLE SL	2.22	4.88	-0.51	0.58

## 6 Discussion

This paper compared the performance of two doubly robust estimators for the ATT dependent on baseline covariates, in the presence of unmeasured confounding, but where a valid (conditional) IV is available. These estimators were implemented with and without the use of data-adaptive estimates of the nuisance parameters. We have demonstrated empirically through simulations that the IV-g estimator has good finite sample performance when using data-adaptive fits for the nuisance parameters, provided the parametric model assumed for the treatment effect curve is correctly specified. The IV-TMLE offers greater flexibility, not relying on a correctly specified parametric model, and performs very well under most scenarios. The price for this greater flexibility is less efficiency, compared with that of the IV-g, when the (parametric) model for the treatment

effect curve is correctly specified. As the simulations show, the use of data-adaptive fits for the nuisance models greatly reduces bias, and improves coverage, resulting in much smaller RMSEs, when compared with using parametric nuisance models, and thus data-adaptive fits should be used.

The methods were motivated and tested in the context of estimating the ATT with effect modification in RCTs with non-adherence to randomised treatment with binary exposure and a continuous outcome. However, the methods presented here are applicable to other settings. One situation may be where the IV assumptions are believed to be satisfied only after conditioning on baseline covariates, making this applicable to certain observational settings. Extensions to situations with continuous exposure are also straight-forward if one is prepared to assume linearity of the treatment effect curve (Tóth and van der Laan, 2016; Vansteelandt and Didelez, 2018).

We have focused on the  $ATT(v)$  as the estimand of interest, but Ogburn et al. (2015) have shown that the same functional of the observed data can be used to identify under monotonicity the local average treatment effects conditional on baseline covariates,  $LATE(v)$ . In fact, much of the previous literature regarding estimation of instrumental variable models with covariates has assumed monotonicity. In particular, for the special case where  $V = W$ , previous methods include full parametric specifications suitable when both the IV and exposure are binary by Little and Yau (1998) and Hirano et al. (2000) and a semi-parametric model proposed by Abadie (2003). In the case where  $V$  is null, Frölich (2007) characterised two distinct non-parametric estimation methods, while Tan (2006) proposed a DR estimator which is consistent when either the instrument propensity score and either the outcome or the exposure parametric models are correctly specified.

For the  $ATT(v)$ , Robins (1994) and Tan (2010) proposed DR estimators in the case where for  $V = W$ , and  $V$  is a strict subset of  $W$  respectively. The DR estimator presented by Okui et al. (2012) and Vansteelandt and Didelez (2018) builds on the work of Tan (2010). For the special case when  $V$  is null, Vansteelandt and Didelez (2018) proposed other DR estimators which are locally efficient, and also constructed a bias-reduced DR IV estimator. Several authors have proposed data-adaptive estimators for the linear IV model with no effect modification, beginning with a TSLS where the first stage is fitted using LASSO with a data-adaptive penalty (Belloni et al., 2012). The bias-reduced DR IV estimator has also been implemented when  $V$  is null using data-adaptive fits for the conditional mean outcome in the unexposed  $\mathcal{A}_y$  (Vermeulen and Vansteelandt, 2016). Chernozhukov et al. (2016) proposed two other IV DR data-adaptive estimators and gave conditions under which data-adaptive fits can be used for  $\mathcal{A}_z$ ,  $\mathcal{A}_a$  and  $\mathcal{A}_y$ . Comparing these DR estimators to the those presented here would be a promising avenue for future research.

The present study has some limitations. Formal proof of the regularity conditions and rates of convergence needed to obtain uniformly valid inferences when using data-adaptive techniques for the IV g-estimator presented here are still lacking, and should be the subject of further study. Similarly, the required convergence conditions for the IV-TMLE estimator have not been formally reported. Another limitation is that we did not seek to quantify the rates of convergence attained by algorithms included in the SL library. This is because in general the rates of convergence of the individual machine learning algorithms depend on the number of included variables, and other tuning parameters, making the assessment of rates of convergence complex. A potential promising

solution for this could be to include the highly adaptive lasso (HAL) (Benkeser and Laan, 2016) in the SL library, as this has been proven under sufficient regularity conditions to converge at rates faster than  $n^{-\frac{1}{4}}$ .

A number of extensions to the work presented here are of interest. The IV-g method implemented here jointly estimates  $m_0(W)$  and  $m_y(W)$ , and thus used parametric models for both. This is not necessary, and an alternative strategy where  $m_y(W)$  is estimated beforehand and the fitted values plugged in into the estimating equation (11) is possible, thus allowing the use of data-adaptive fits for the model  $\mathcal{A}_y$ . Future work could extend the bias-reduced DR estimator to the linear IV model with effect modification, and compare this with IV-TMLE and a fully data-adaptive version of the IV-g estimator.

## Acknowledgements

We thank Prof. Stephanie Taylor and the COPERS trial team for access to the data. We also thank Prof. Stijn Vansteelandt for commenting on an earlier draft of this paper, and Boriska Tóth for sharing her code implementing IV-TMLE.

Karla DiazOrdaz is supported by UK Medical Research Council Career development award in Biostatistics MR/L011964/1.

## References

- Abadie, A. (2003). Semiparametric instrumental variable estimation of treatment response models. *Journal of econometrics*, 113(2):231–263.
- Angrist, J. D., Imbens, G. W., and Rubin, D. B. (1996). Identification of causal effects using instrumental variables. *Journal of the American Statistical Association*, 91(434):pp. 444–455.
- Belloni, A., Chen, D., Chernozhukov, V., and Hansen, C. (2012). Sparse models and methods for optimal instruments with an application to eminent domain. *Econometrica*, 80(6):2369–2429.
- Benkeser, D., Carone, M., Laan, M. V. D., and Gilbert, P. (2017). Doubly robust nonparametric inference on the average treatment effect. *Biometrika*, 104(4):863–880.
- Benkeser, D. and Laan, M. V. D. (2016). The highly adaptive lasso estimator. In *2016 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*, pages 689–696.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., et al. (2016). Double machine learning for treatment and causal parameters. *arXiv preprint arXiv:1608.00060*.
- Dodd, S., White, I., and Williamson, P. (2012). Nonadherence to treatment protocol in published randomised controlled trials: a review. *Trials*, 13(1):84.

- Dunn, G. and Bentall, R. (2007). Modelling treatment-effect heterogeneity in randomized controlled trials of complex interventions (psychological treatments). *Statistics in Medicine*, 26(26):4719–4745.
- Farrell, M. H. (2015). Robust inference on average treatment effects with possibly more covariates than observations. *Journal of Econometrics*, 189(1):1–23.
- Frölich, M. (2007). Nonparametric IV estimation of local average treatment effects with covariates. *Journal of Econometrics*, 139(1):35 – 75. Endogeneity, instruments and identification.
- Gruber, S. and van der Laan, M. J. (2010). A targeted maximum likelihood estimator of a causal effect on a bounded continuous outcome. *The International Journal of Biostatistics*, 6(1).
- Györfi, L., Kohler, M., Krzyzak, A., and Walk, H. (2006). *A distribution-free theory of nonparametric regression*. Springer Science & Business Media.
- Hernán, M. A. and Robins, J. M. (2006). Instruments for causal inference: an epidemiologist’s dream? *Epidemiology*, 17(4):360–372.
- Hirano, K., Imbens, G. W., Rubin, D. B., and Zhou, X.-H. (2000). Assessing the effect of an influenza vaccine in an encouragement design. *Biostatistics*, 1(1):69–88.
- 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.
- Little, R. J. and Yau, L. H. (1998). Statistical techniques for analyzing data from prevention trials: Treatment of no-shows using rubin’s causal model. *Psychological Methods*, 3(2):147.
- Newey, W. K. (1990). Semiparametric efficiency bounds. *Journal of Applied Econometrics*, 5(2):99–135.
- Ogburn, E. L., Rotnitzky, A., and Robins, J. M. (2015). Doubly robust estimation of the local average treatment effect curve. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 77(2):373–396.
- Okui, R., Small, D. S., Tan, Z., and Robins, J. M. (2012). Doubly robust instrumental variable regression. *Statistica Sinica*, 22(1):173–205.
- Petersen, M., Schwab, J., Gruber, S., Blaser, N., Schomaker, M., and van der Laan, M. (2014). Targeted maximum likelihood estimation for dynamic and static longitudinal marginal structural working models. *Journal of causal inference*, 2(2):147–185.
- Pirracchio, R., Petersen, M. L., and van der Laan, M. (2015). Improving propensity score estimators’ robustness to model misspecification using super learner. *American journal of epidemiology*, 181(2):108–119.

- 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.
- Robins, J. M. (1994). Correcting for non-compliance in randomized trials using structural nested mean models. *Communications in Statistics - Theory and Methods*, 23(8):2379–2412.
- Robins, J. M. (2000). Robust estimation in sequentially ignorable missing data and causal inference models. In *Proceedings of the American Statistical Association, Section on Bayesian Statistical Science*, pages pp. 6–10.
- Rubin, D. B. (1978). Bayesian inference for causal effects: The role of randomization. *The Annals of statistics*, pages 34–58.
- Tan, Z. (2006). Regression and weighting methods for causal inference using instrumental variables. *Journal of the American Statistical Association*, 101(476):1607–1618.
- Tan, Z. (2010). Marginal and nested structural models using instrumental variables. *Journal of the American Statistical Association*, 105(489):157–169.
- Taylor, S. J., Carnes, D., and Homer, Kate, e. a. (2016). Improving the self-management of chronic pain: Coping with persistent pain, effectiveness research in self-management (copers). *Programme Grants for Applied Research*, 4(14).
- Tóth, B. and van der Laan, M. J. (2016). TMLE for marginal structural models based on an instrument. *U.C. Berkeley Division of Biostatistics Working Paper Series. Working Paper 350*.
- van der Laan, M. and Rose, S. (2011). *Targeted Learning: Causal Inference for Observational and Experimental Data*. Springer Series in Statistics.
- van der Laan, M. J. and Gruber, S. (2012). Targeted minimum loss based estimation of causal effects of multiple time point interventions. *The international journal of biostatistics*, 8(1).
- van der Laan, M. J. and Luedtke, A. R. (2015). Targeted learning of the mean outcome under an optimal dynamic treatment rule. *Journal of causal inference*, 3(1):61–95.
- van der Laan, M. J., Polley, E. C., and Hubbard, A. E. (2007). Super learner. *Statistical applications in genetics and molecular biology*, 6(1):1–21.
- VanderWeele, T. J. (2009). Concerning the consistency assumption in causal inference. *Epidemiology*, 20(6):880–883.
- Vansteelandt, S. and Didelez, V. (2018). Improving the robustness and efficiency of covariate adjusted linear instrumental variable estimators. *to appear in Scandinavian Journal of Statistics*.
- Vermeulen, K. and Vansteelandt, S. (2016). Data-adaptive bias-reduced doubly robust estimation. *The international journal of biostatistics*, 12(1):253–282.

- Wiles, N. J., Fischer, K., Cowen, P., Nutt, D., Peters, T. J., Lewis, G., and White, I. R. (2014). Allowing for non-adherence to treatment in a randomized controlled trial of two antidepressants (citalopram versus reboxetine): an example from the genpod trial. *Psychological Medicine*, 44(13):2855–2866.
- Wooldridge, J. M. (2010). *Econometric Analysis of Cross Section and Panel Data*. MIT Press.
- Zhang, Z., Peluso, M. J., Gross, C. P., Viscoli, C. M., and Kernan, W. N. (2014). Adherence reporting in randomized controlled trials. *Clinical Trials*, 11(2):195–204.
- Zheng, W. and van der Laan, M. J. (2012). Targeted maximum likelihood estimation of natural direct effects. *The international journal of biostatistics*, 8(1):1–40.