



Deposited via The University of York.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/166368/>

Version: Published Version

Article:

DiazOrdaz, Karla, Rhian, Daniel and Kreif, Noemi (2020) Data-adaptive doubly robust instrumental variable methods for treatment effect heterogeneity. *Journal de la Société Française de Statistique*. pp. 135-163. ISSN: 2102-6238

Reuse

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

Takedown

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

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

Titre: Estimateurs doublement robustes avec apprentissage automatique pour l'estimation de l'hétérogénéité de l'effet traitement dans les modèles à variables instrumentales

Karla DiazOrdaz¹, Rhian Daniel² and Noemi Kreif³

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 g-estimator and a targeted minimum loss-based estimator (TMLE). These 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 and use data-adaptive estimation of the nuisance parameters for the g-estimator.

A simulation study is used to compare standard TSLS with the two DR estimators' finite-sample performance when using (1) parametric or (2) data-adaptive estimation of the nuisance parameters.

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.

The methods are also applied to the COPERS (COping with persistent Pain, Effectiveness Research in Self-management) trial to make inferences about the causal effect of treatment actually received, and the extent to which this is modified by depression at baseline.

Keywords: Instrumental variables, doubly robustness, machine learning estimation, heterogeneous treatment effects,, g-estimation, TMLE

AMS 2000 subject classifications: 62F35, 62G05, 46N30

1. Introduction

There has been an increased interest in estimating the causal effect of treatment actually received in randomised controlled trials (RCTs) in the presence of treatment non-adherence, in addition to the intention-to-treat effect, as highlighted by the International Council for Harmonisation addendum to guideline E9 (Statistical Principles for Clinical Trials, addendum on Estimands). 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](#)).

¹ Department of Medical Statistics, London School of Hygiene and Tropical Medicine, London United Kingdom.
E-mail: karla.diaz-ordaz@lshtm.ac.uk

² Division of Population Medicine, Cardiff University, Wales, United Kingdom.
E-mail: danielr8@cardiff.ac.uk

³ Centre for Health Economics, University of York, United Kingdom.
E-mail: noemi.kreif@york.ac.uk

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 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 regresses 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, 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 DR 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 (van der Laan and Rose, 2011).

Although DR estimators offer in principle partial protection against model misspecification, concerns remain over their performance in practice, when all models are likely to be misspecified (Kang et al., 2007). To alleviate biases due to model misspecification, TMLE is usually coupled with machine learning estimation of the nuisance parameters, using in particular the Su-

per 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 estimation. 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 estimators, with the latter phenomenon making valid statistical inferences complex to obtain (van der Vaart, 2014). In addition, since the resulting estimators can be irregular, the nonparametric bootstrap is in general not valid (Bickel et al., 1997). Some DR estimators, such as TMLE, on the other hand, when combined with machine learning estimation of the nuisance functionals, have faster convergence and make (asymptotic) analytic statistical inference tractable via the sampling variance of the corresponding influence functions, under empirical processes conditions, assuming that the convergence rates of the machine learning estimators (to their respective truths) used are fast enough (van der Laan and Rubin, 2006; Farrell, 2015).


Building on previous literature that establishes conditions for one-step and estimating equations estimators to be (asymptotically) Neyman orthogonal (Newey and McFadden, 1994; van der Laan and Robins, 2003) as well as previous work that used sample splitting to avoid empirical processes conditions (Bickel, 1982), Chernozhukov et al. (2018) proposed the use of sample splitting when using machine learning for estimating the nuisance parameters, thus widening the class of estimating equations DR estimators that can be estimated data-adaptively. In particular, Chernozhukov et al. (2018) give regularity conditions for estimating equations estimators of the linear IV model, which can be adapted for the g-estimator introduced by Vansteelandt and Didelez (2018). Thus, we implement here the g-estimator with and without machine learning estimation for nuisance parameters. We compare its performance with that of a TMLE (Tóth and van der Laan, 2016), again implemented either parametrically or data-adaptively, and standard parametric 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 by applying them to the illustrative RCT.

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 machine learning estimation for the nuisance models of the 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 in Section 6.

2. Linear instrumental variables models

Let W be a set of baseline variables, Z be the randomised treatment indicator and A be the exposure of interest, the actual treatment received, assumed to be binary. Denote by Y the continuous outcome of interest, and by U the set of all unobserved common causes of A and Y . Further, assume that (U, W) would be a sufficient set to control for the confounding in the effect of A on Y , were U observed. For simplicity, but without loss of generality, we assume that interest lies in estimating effect modification by a single baseline variable $V \in W$.

FIGURE 1. DAG depicting a valid conditional instrument Z for exposure A in the presence of observed and unobserved confounders W and U respectively, where the outcome is Y .



Let a subscript 0 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 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 potential outcome $Y(a)$ be the outcome that would occur if A were set to $a \in \{0, 1\}$. As usual, we assume *no interference*, i.e. the potential outcomes of the i -th individual are unrelated to the treatment status of all other individuals, and *counterfactual consistency*, for all individuals $Y = Y(z)$ and $A = A(z)$ if $Z = z$, and $Y = Y(z, a)$ if $(Z, A) = (z, a)$, for all z and all a (Rubin, 1978; VanderWeele, 2009).

Following Abadie (2003) and 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. $Z \not\perp\!\!\!\perp A|W$.

Assumptions (i) and (ii) can be shown to imply (ii') $Y(a) \perp\!\!\!\perp Z|W$, for all a , which is an alternative assumption often invoked independently (Robins, 1994; Vansteelandt and Didelez, 2018; Swanson et al., 2018).

In addition to these IV assumptions, we assume the following partially linear conditional mean model for the outcome:

$$E[Y|A, W, Z, U] = \bar{\omega}_0(W, U) + Am_0(W), \quad (1)$$

where $\bar{\omega}_0(W, U)$ and $m_0(W)$ are unknown functions, with $m_0(W)$ representing the causal treatment effect curve given covariates W . The assumption of linearity in A is necessary to identify $m_0(W)$ using an instrument. With binary exposure A , this assumption always holds.

Under these assumptions, the conditional mean model (1) implies the so-called linear structural mean model (Robins, 1994):

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

Proof: We begin by observing that eq. (1) implies $E[Y|A = 0, W, Z, U] = \omega_0(W, U)$. Thus, we can re-write eq. (1) as

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

where we use the fact that $Y = Y(z, a) = Y(a)$ by counterfactual consistency and exclusion restriction. Now, since $Y(a) \perp\!\!\!\perp A|(U, W)$, since U and W are sufficient to control for confounding between A and Y , we have $E[Y(0)|A = 0, W, Z, U] = E[Y(0)|A, W, Z, U]$, and thus

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

where the last step uses the fact that the right hand side, $Am_0(W)$ is independent of U (Vansteelandt and Didelez, 2018) ■.

While the linear structural mean model eq. (2) 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, namely $E[Y - Am_0(W)|Z, W] = E[Y - Am_0(W)|W]$. Therefore, we denote by \mathcal{M} 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 model \mathcal{M} assumes the treatment effect curve $m_0(W)$ 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

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

Since $ATT(v)$ is the conditional expectation of $m_0(W)$ given $A = 1$ and $V = v$, we focus on identifying $m_0(W)$.

Rearranging equation (2), we have

$$\begin{aligned} E[Y|A, W, Z] - Am_0(W) &= E[Y(0)|A, W, Z], \\ E\{E[Y|A, W, Z]|W, Z\} - E[Am_0(W)|W, Z] &= E\{E[Y(0)|A, W, Z]|W, Z\}, \\ E[Y - Am_0(W)|Z, W] &= E[Y(0)|Z, W], \\ E[Y - Am_0(W)|W] &= E[Y(0)|W], \end{aligned} \quad (4)$$

where in the second step we marginalise over A , and the last equality holding since $Y(0) \perp\!\!\!\perp Z|W$ (Assumption ii’).

Model \mathcal{M} thus implies

$$E[Y|Z, W] = \omega_0(W) + m_0(W)E[A|Z, W], \quad (5)$$

where $\omega_0(W) = E[Y - Am_0(W)|W]$ being equal to $E[\omega_0(W, U)|W]$ or $E[Y(0)|W]$, depending on whether model (1) or (2) is assumed.

Equation (5) implies $E[Y|Z = 1, W] - E[Y|Z = 0, W] = m_0(W)(E[A|Z = 1, W] - E[A|Z = 0, W])$, therefore for a binary IV, under \mathcal{M} and the (conditional) IV assumptions, $m_0(W)$ is identified by

$$m_0(W) = \frac{E[Y | Z = 1, W] - E[Y | Z = 0, W]}{E[A | Z = 1, W] - E[A | Z = 0, W]}, \quad (6)$$

Estimation of the conditional expectations in equation (6) would typically involve specifying models for the mean exposure $E[A|Z, W]$ and the mean outcome $E[Y|Z, W]$.

Denote by $\omega(W)$ the model for $E[Y - Am_0(W)|W]$ and $\pi(Z, W)$ the model for $E[A|Z, W]$. Finally, let $\mu(Z, W)$ denote the implied model for $E[Y|Z, W]$.

3. Doubly robust estimation for the linear instrumental variable model

To illustrate the methods, we consider throughout a situation where interest lies in the main effect modification by a single variable $V \in W$, with a working parametric model for the treatment effect curve as a function of this single variable being:

$$m(W; \psi) = \psi_c + \psi_v V. \quad (7)$$

The statistical parameter of interest is therefore $\psi = (\psi_c, \psi_v)$, where ψ_c represents the main causal treatment effect, and ψ_v is the effect modification by V . The function $m(W; \psi)$ can be interpreted as a working model for $E[m_0(W)|A = 1, V]$. Importantly, the working parametric model is not assumed to be the true model for $m_0(W)$.

3.1. TSLS

Estimation of the expectations in equation (6) 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, the second stage is a linear model for eq. (5), that is a linear regression for the outcome on the predicted treatment received and baseline covariates. We write

$$E[A|Z, W] = \pi(Z, W), \quad (8)$$

$$E[Y|Z, W] = \omega(W) + m(W)\pi(Z, W). \quad (9)$$

In principle, there are many parametric choices for the second stage models, $\omega(W)$ and $m(W)$. For TSLS to be consistent, the first stage model $\pi(Z, W)$ must be the parametric linear regression implied by the second stage, i.e. it must include all the covariates and interactions appearing in the second stage model.

For example, if we assume working models $m(W; \psi) = \psi_c + \psi_v V$, and $\omega(W; \beta) = \beta^\top W$, 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

$$\begin{aligned} E[A|Z, W] &= \alpha_z Z + \alpha_{zV} ZV + \alpha_v V + \sum_{W_i \in W \setminus V} \alpha_{wi} W_i, \\ E[AV|Z, W] &= \lambda_z Z + \lambda_{zV} ZV + \lambda_v V + \sum_{W_i \in W \setminus V} \lambda_{wi} W_i, \end{aligned} \quad (10)$$

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 AZ , the resulting TSLS estimator may be inefficient (Vansteelandt and Didelez, 2018).

Vansteelandt and Didelez (2018) show that standard TSLS estimation of β and ψ in equation (9) is equivalent to solving an estimating equation of the form

$$0 = \sum_{i=1}^n e_y(Z_i, W_i) \{Y_i - \omega(W_i; \beta) - m(W_i; \psi) \pi(Z_i, W_i; \alpha_0)\}, \quad (11)$$

for a given α_0 , where $e_y(Z_i, W_i)$ is any conformable index vector function of dimension $\dim(\beta) + d$.

The estimators $\hat{\beta}$ and $\hat{\psi}$ obtained solving equation (11), after substituting $\hat{\alpha}$ for α_0 (the estimator from the first stage), are consistent asymptotically normal (CAN), when both models $\omega(W; \beta)$ and $m(W; \psi)$ are correctly specified, i.e. even when $\pi(W; \alpha)$, the first stage model for the exposure, is misspecified (Robins, 2000; Wooldridge, 2010). Moreover, in the specific settings where the treatment effect curve $m(W; \psi)$ is linear in the covariates and the instrument is independent of W , TSLS is also robust to misspecification of $\omega(W; \beta)$. We refer the interested reader to Vansteelandt and Didelez (2018), Appendix B Proposition 5, for the proof.

This means that for estimators which are doubly robust in the more general settings, with either a treatment effect model $m(W; \psi)$ that depends on the covariates (treatment effect heterogeneity), or where the instrument Z depends on W , methods beyond TSLS need to be considered.

3.2. G-estimation

Thus far, we have shown that in treatment effect modification settings with a binary conditional IV, the TSLS IV estimator is consistent if the treatment-free outcome model $\omega(W)$ is correctly specified. An approach to obtaining a doubly robust estimator involves modelling $E[Z|W_i]$ in structural nested mean models (Robins, 1994). Then, the parameter of interest ψ can be estimated using G-estimation.

G-estimation exploits the idea that, on average, there is no residual association between Z and $E[Y - Am_0(W)|W]$. This suggests an estimation strategy for finding the parameters that make the empirical conditional covariance between Z and the treatment-free potential outcome $Y(0)$ equal to 0. The resulting estimator is consistent if either the model for the conditional expectation $E(Z|W)$ or the treatment-free outcome model $\omega(W)$ or both are correctly specified, and the assumption that partially linear IV model \mathcal{M} for the conditional mean of Y given W and Z is correct. The model for the conditional distribution of the binary IV $g_0(W; \gamma_0) = E[Z|W] = P_0(Z = 1|W)$ is often called the instrument propensity score, and it is assumed to be a known function of W , smooth in a finite dimensional parameter γ_0 .

Okui et al. (2012) showed that this g-estimator for $\psi = (\psi_c, \psi_v)$ can be obtained as a solution to the following estimating equation (Okui et al., 2012)

$$0 = \sum_n (e(Z_i, W_i; \gamma_0) - E[e(Z_i, W_i; \gamma_0)|W_i]) \{Y_i - \omega_0(W_i; \beta_0) - m_0(W_i; \psi) A_i\}, \quad (12)$$

where $e(Z, W)$ is any conformable vector function, i.e. of the appropriate dimension $\dim(\beta) + d$, with $d = \dim(\psi)$.

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

$$e(Z, W; \gamma_0) = \sigma_0^{-2}(Z, W) \begin{pmatrix} 1 \\ V \end{pmatrix} \left[\pi_0(Z, W; \alpha_0) - \frac{E[\sigma_0^{-2}(Z, W)\pi_0(Z, W; \alpha_0)|W]}{E[\sigma_0^{-2}(Z, W)|W]} \right] \quad (13)$$

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

This estimator requires the user to specify working parametric models for $E(Y - Am(W; \psi)|W)$ and $E(Z|W)$, i.e. to specify working models for $\omega(W; \beta)$ and $g(W; \gamma)$. The estimator (denoted by IV-g) considered here estimates both the parameter of interest ψ and the nuisance parameter β jointly, though approaches that estimate β consistently first have also been proposed (Okui et al., 2012). This can be made feasible by replacing α_0 , β_0 and γ_0 by their corresponding consistent estimators $\hat{\alpha}$, $\hat{\beta}$ and $\hat{\gamma}$, and setting σ_0^2 equal to 1 (as it is just a proportionality constant). It has been shown to be CAN if either the model for $E(Y - Am(W; \psi)|A, W)$ or the model for $E(Z|W)$ is correct, and hence consistent whenever the model for $m_0(W)$ is correctly specified (Okui et al., 2012). The addition of the instrument propensity score model to the TSLS estimating equations (11) 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 all three models are correctly specified (Vansteelandt and Didelez, 2018).

The influence function of the IV-g estimator can be written as

$$D_i(\psi)(O_i) = M(Z_i, W_i, A_i)^{-1} K(Z_i, W_i) \begin{pmatrix} 1 \\ V_i \end{pmatrix} \{Y_i - \omega(W_i; \beta)\} - \begin{pmatrix} \psi_c \\ \psi_v \end{pmatrix} \quad (14)$$

with

$$K(Z, W) = \pi(Z, W; \alpha) - E_{g(W; \gamma_0)}[\pi(Z, W; \alpha)|W] \quad (15)$$

and

$$M(Z, W, A) = AK(Z, W) \begin{pmatrix} 1 & V \\ V & V^2 \end{pmatrix}. \quad (16)$$

Since the IV g-estimator is CAN, the asymptotic variance is the variance of its influence function, i.e. $\text{Var}(\psi) = E[D(\psi)^\top D(\psi)]$ (Newey, 1990). Therefore, we obtain an estimate of the variance by the sample variance of the estimated influence function, obtained by plugging-in consistent estimators for α , β and γ ,

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

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.

The IV-g estimator gains efficiency from assuming that the working model for the treatment effect curve, equation (7), $m_\psi(W) = m(W; \psi) = \psi_c + \psi_v V$, holds when this is correct. However when model (7) is misspecified (e.g. that the treatment effect curve depends on more covariates, not just V , or that the relationship is not linear), the IV-g estimator will behave as a projection onto the working parametric model, so long as the mean exposure model $\pi_0(Z, W)$ is correctly specified and $\text{Cov}(\{\pi_0(Z, W) - E(\pi_0(Z, W)|W)\}, A|W)$ is constant in W .

3.3. Data-adaptive estimation

The IV-g estimator presented thus far is restricted to using parametric working models for the nuisance parameters. Since all working models are likely to be misspecified in practice, the resulting estimator is unlikely to be consistent.

An increasingly popular strategy to avoid the bias introduced by such model misspecification and have valid inferences is to use machine learning estimators for the nuisance parameters. This is made possible since DR estimators can converge at fast rates (\sqrt{n}) to the true parameter, and are therefore CAN, even when the nuisance functionals have been estimated via machine learning, under either empirical process conditions (e.g. Donsker class) restricting the complexity of the nuisance functionals, or using sample splitting (van der Laan and Rose, 2011; Farrell, 2015; Chernozhukov et al., 2017; Kennedy, 2016; Athey et al., 2018).

Briefly, if the score function \mathcal{S} of the DR estimator is *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 is CAN and inference can be based on the IF. Convergence rates for these data-adaptive estimators depend on the smoothness and number of covariates included (Györfi et al., 2006).

Machine learning estimation of the nuisance parameters of DR estimators for the partially linear IV model has been studied recently. Chernozhukov et al. (2018) give sufficient conditions to guarantee that using data-adaptive fits for the nuisance functionals in DR estimators constructed from estimating equations based on Neyman-orthogonal scores results in valid inferences. In particular, consider the score function

$$\mathcal{S}_i = \{Z_i - g(W_i)\} \{Y_i - \omega(W_i) - m(W_i; \psi)A_i\}, \quad (17)$$

where $g(W_i)$ and $\omega(W_i)$ are L^2 -functions with respect to P_0 , mapping $W \mapsto \mathbb{R}$. Assuming Y, A and Z are bounded and with finite variance bounded away from zero, the estimator obtained as a solution to the estimating equation with score (17) is CAN even after plugging in data-adaptive nuisance estimators, as long as these satisfy:

$$\|\hat{g}(W) - g_0(W)\| \times \|\hat{\omega}(W) - \omega_0(W)\| < o_P(n^{-\frac{1}{2}}), \quad (18)$$

where $\|\cdot\| = \|\cdot\|_{P,2}$ i.e. the L^2 -functions with respect to P_0 .

We refer the interested reader to Chernozhukov et al. (2018) for the technical details.

Since the g-estimator for the IV model is Neyman orthogonal, data-adaptive IV-g estimators can be obtained by solving equation (12) after data-adaptive estimates for $\omega(W)$ (the treatment-free outcome model), $\pi(Z, W)$ (the exposure model) and/or $g(W)$ (the instrument propensity score) have been plugged in. Under sufficient regularity conditions, and provided the data-adaptive models used converge sufficiently fast to their respective true parameter, the resulting IV g-estimator is CAN.

For example, solving equation (12) where we have plugged in fits from a parametric model for $\pi(Z, W)$ and data-adaptive estimates for $E[Y - Am_0(W)|W]$ and $E[Z|W]$, the arguments used in Chernozhukov et al. (2018) can be applied directly to show that the IV-g estimator is CAN

when eq. (18) holds. To see why, consider a parametric model for $\pi(Z, W) = \alpha_0 + \alpha_1 Z + \alpha_2 V$, so that the score function for eq. (12) is $\mathcal{S} = (1 \quad V)^\top \{\alpha_1 (Z - g(W)) (Y - \omega(W) - Am(W; \psi))\}$, which is of the form eq. (17).

The data-adaptive IV-g estimator implemented here uses data-adaptive estimates for $E[A|Z, W]$ and $E[Z|W]$ but estimates jointly the parametric $m(W; \psi)$ and $\omega(W, \beta)$ as before. The resulting estimator can be shown to be CAN if the nuisance models converge to their respective truths at the rates of convergence in equation (18), under sufficient regularity conditions. See the Appendix for a sketch of the proof.

To obtain the data-adaptive estimates, we use 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). 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 influence-function based estimation approach, which incorporates a “targeting” step that guarantees the resulting estimator has a well behaved higher-order residual term. Most commonly, it combines estimates of nuisance functionals and an initial estimate of the target parameter. These initial estimates can be obtained by specifying parametric models or, under empirical processes conditions (e.g. Donsker class) which can be relaxed using sample splitting (Zheng and van der Laan, 2011), via machine learning. Typically, the TMLE literature uses the Super Learner with cross-validation (van der Laan et al., 2007). Assuming the data-adaptive estimates converge to their respective truths sufficiently fast, the resulting TMLE is CAN. We refer the interested reader to van der Laan and Rose (2011) and van der Laan and Rose (2018).

Tóth and van der Laan (2016) proposed three TMLE estimators for the (partially) linear IV model. In the next section, we describe in more detail the non-iterative linear TMLE, which we denote by IV-TMLE.

3.4.1. 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 projecting the treatment effect curve onto the

working parametric model $m_\psi = \psi_c + \psi_v V$, i.e. $\Psi(P_0) = (\psi_c, \psi_v) = \psi_0$ is the solution to

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

We note that Ψ only depends on P_0 through m_0 and the distribution of Z and the covariates P_{0W} . We denote this relevant part by $Q_0 = (m_0, g_0, Q_{0W})$ with $Q_{0W} = P_{0W}$.

Under the IV model $\mathcal{M} : E[Y|Z, W] = \omega_0(W) + m_0(W)\pi_0(Z, W)$, the treatment effect curve $m_0(W)$ depends on $\mu_0(Z, W) = E[Y|W, Z]$ and $\pi_0(Z, W)$, and thus construction of a TMLE for the IV model starts by obtaining initial estimates of $\mu(Z, W)$, $\pi(Z, W)$, and the instrument propensity score $g(W)$. We denote these initial estimates by a 0 superscript. From these, and model (5), we calculate an initial estimate for $m(W)$, denoted $m^0(W)$.

The next step in the construction of a TMLE requires the specification of a loss function $L(P)$, such that the expectation of the loss function is minimised at the true probability distribution, $E_0[L(P_0)(O)] = \min_{P \in \mathcal{P}} E_0[L(P)(O)]$. Here, we use the square error loss function. Under the IV model \mathcal{M} and the working model for the treatment effect curve $m_\psi(V) = \psi_c + \psi_v V$, the efficient influence function (EIF) can be written as:

$$D^*(m, g, Q_W)(O) = h(W) \{ \pi_0(Z, W) - E_0[\pi_0(Z, W)|W] \} \{ Y - \pi_0(Z, W)m_0(W) - \omega_0(W) \} - h(W) \{ (\pi_0(Z, W) - E_0[\pi_0(Z, W)|W])m_0(W) \} (A - \pi_0(Z, W)) + D_W(Q_W), \quad (19)$$

where $h(W)$ is the so-called *clever covariate*, defined as

$$h(W) = \text{Var}(V)^{-1} \begin{pmatrix} E[V^2] - E[V]V \\ V - E[V] \end{pmatrix} \zeta^{-2}(W) \quad (20)$$

with the term $\zeta^2(W)$, which is associated with instrument strength, being

$$\begin{aligned} \zeta^2(W) &= \text{Var}_{Z|W}(\pi(Z, W)|W), \\ &= E[\{ \pi(Z, W) - \sum_{z \in \{0,1\}} \pi(z, W)g(Z = z, W) \}^2 | W], \\ &= \{ \pi(1, W) - \pi(0, W) \}^2 g(W)(1 - g(W)), \\ &= \{ \pi(1, W) - \pi(0, W) \}^2 \text{Var}(Z|W). \end{aligned} \quad (21)$$

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

The *targeting* step involves fitting a linear model for $m(W)$ on the single ‘‘clever’’ covariate $h(W)$ with the initial estimate $m^0(W)$ as an offset,

$$m^*(\epsilon)(W) = m^0(W) + h(W)^T \epsilon. \quad (22)$$

Estimation of the coefficient in equation (22) involves solving the empirical EIF equation,

$$\frac{1}{n} \sum_{i=1}^n D^*(m^*(\epsilon), g^0, Q_W)(O_i) = 0, \quad (23)$$

or equivalently, solving for ϵ a system of d linear equations:

$$\frac{1}{n} \sum_{i=1}^n h^0(W_i) \{ \pi^0(Z_i, W_i) - E_{g^0(W_i)}[\pi^0(Z_i, W_i)|W_i] \} \left\{ Y_i - A_i \left(m^0(W_i) + h^0(W_i)^T \epsilon \right) - \omega^0(W_i) \right\} = 0,$$

where $h^0(W_i)$ is obtained by plugging in $\pi^0(Z_i, W_i)$ and $g^0(W_i)$ into the equations defining the clever covariate (21) and (20).

Denote by ε^* the solution to equation (23). Then, the non-iterative linear TMLE estimator of $m_0(W)$ is obtained by substituting ε^* into equation (22). Finally, we project the resulting function $m^*(\varepsilon^*)(W)$ onto the working model m_ψ by OLS, obtaining (ψ_c^*, ψ_v^*) , the TMLE estimator of the statistical parameters of interest.

Tóth and van der Laan (2016) showed that this approach results in an estimator which is double-robust, i.e. consistent when (i) the initial estimators of $\pi_0(Z, W)$ and g_0 are consistent, (ii) the initial estimators of m_0 and g_0 are consistent, or (iii) the initial estimators of m_0 and ω_0 are consistent. However, using a linear fluctuation model has the drawback that the resulting estimates are not guaranteed to be constrained within the bounds implied by the data.

We remark that the variance of the IV-TMLE estimators becomes very large when the term $\zeta^2(W)$ is very small. Since $\zeta^2(W)$ captures the strength of the instrument in predicting the exposure given W , the IV-TMLE estimators become unstable with large variance when the instrument is weak. To stabilise the estimators, we choose the maximum of the estimated value of $\zeta^2(W)$ and 0.025 when constructing the clever covariate for a given data set.

4. Simulation Study

We perform a factorial simulation study to assess the finite sample performance of the alternative methods to estimate the statistical parameter of interest, under the different combinations of ω , π or m being in turn correctly specified or not, while the instrument model is always correct. We write $\mathbb{1}(k \neq k_0)$ as an indicator function for scenarios where the assumed model for $k \in \{\omega, \pi, m\}$ is misspecified.

We generate data to mimic a randomised controlled trial with two-sided non-adherence, i.e. both randomly allocated groups have a non-zero probability of receiving the opposite treatment. There are two different sample sizes, small $n = 500$ and large $n = 10,000$. 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 $A \sim \text{Bern}(\pi_0(W, V, U, Z))$, i.e. the probability of getting the active treatment is a function of the baseline variables, the unobserved confounder, and the instrument, namely

$$\text{logit}(\pi_0) = 1.5Z + 0.03V + 0.01W_1 + 0.01W_2 + 0.01W_3 + 0.01W_4 + 0.03U - \mathbb{1}(\pi \neq \pi_0)(5ZW_1).$$

Notice that we are generating the exposure A in such a way that the condition necessary for the IV-g estimator to converge to the parameter of interest when \mathcal{M} is wrong is no longer satisfied, for settings where the true $\pi_0 \neq \pi$.

The continuous outcome Y is then simulated from $Y \sim N(\mu_0, 1)$, with μ_0 given by:

$$\begin{aligned}\mu_0 &= \omega_0(W) + m_0(W)A + U, \\ \omega_0(W) &= \{1 - \mathbb{1}(\omega \neq \omega_0)\}\{0.5 + 0.5V + 0.01W_1 + 0.01W_2 + 0.01W_3 + 0.01W_4\} \\ &\quad + \mathbb{1}(\omega \neq \omega_0)\exp\{0.05 + 0.05V + 0.001W_1 + 0.001W_2 + 0.001W_3 \\ &\quad\quad + 0.001W_4 - 0.2V(W_1 + W_2 + W_3 + W_4)\}, \\ m_0(W) &= 0.5 + 0.5V + \mathbb{1}(m \neq m_0)\{3(W_1 + W_2 + W_3 + W_4)\},\end{aligned}$$

which means that the true $E[m_0(W)|A = 1, V = v] = 0.5 + 0.5V$, i.e. $\psi_0 = (\psi_{c0}, \psi_{v0}) = (0.5, 0.5)$.

We 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 data-adaptive estimation of nuisance models. For the TSLS, the first stage is as per equation (10). Parametric IV-g and TMLE use main terms logistic models for the instrument propensity score and the treatment model, namely

$$\text{logit}(g(W; \gamma)) = \text{logit}\{P(Z = 1|W)\} = \gamma_0 + \sum_{i=1}^4 \gamma_i W_i + \gamma_5 V,$$

and

$$\text{logit}(\pi(W, Z; \alpha)) = \text{logit}\{P(A = 1|Z, W)\} = \alpha_2 Z + \sum_{i=1}^4 \alpha_i W_i + \alpha_5 V$$

For the data-adaptive estimation of $\pi(Z, W)$ and $g(W)$, we use the Super Learner. Since A and Z are binary, the library used includes glm (generalised linear models), step (stepwise model selection using AIC), svm (support vector machines, with radial basis functions) and gam (generalised additive models), with linear and second-order terms used for the glm, step and gam learners.

In addition, for the IV-TMLE, we use need data-adaptive estimates of the continuous outcome. The library of learners used for $\mu(Z, W)$ and $\omega(W)$ includes glm, step, svm and polymars (multivariate adaptive polynomial spline regression), chosen in order to preserve the linear structure of the partially linear IV model (5).

The SEs of the parametric IV-g and TMLE are obtained by bootstrapping (percentile 95% confidence intervals (CI) using 1999 bootstrap samples), while for the data-adaptive estimators the SEs are based on the empirical variance of the estimated (E)IF.

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

4.1. Results from the simulation

Figures 2 and 3 show the mean bias (top) and CI coverage rate (bottom) corresponding to scenarios with sample size of $n = 500$ and $n = 10,000$ respectively. For clarity, the figures show only the methods resulting in absolute bias less than 2 are plotted, corresponding to those having absolute bias $< 400\%$ of the true parameter value. The excluded results are reported in Table 4 in the Appendix.

When all models are correctly specified (first column, plotted in black), all methods show close to zero bias for both of the target parameters. At large sample sizes ($n = 10,000$), the coverage

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 π and ω vary by column, m correctly specified is plotted in black while m misspecified is plotted in grey. The hollow shapes correspond to parametric nuisance models estimation, and the solid shapes correspond to estimators using data-adaptive nuisance model estimates. The bias is presented with Monte Carlo Error CIs. Results corresponding to bias ≥ 2 in absolute value are not plotted, but can be found in Table 4. 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.

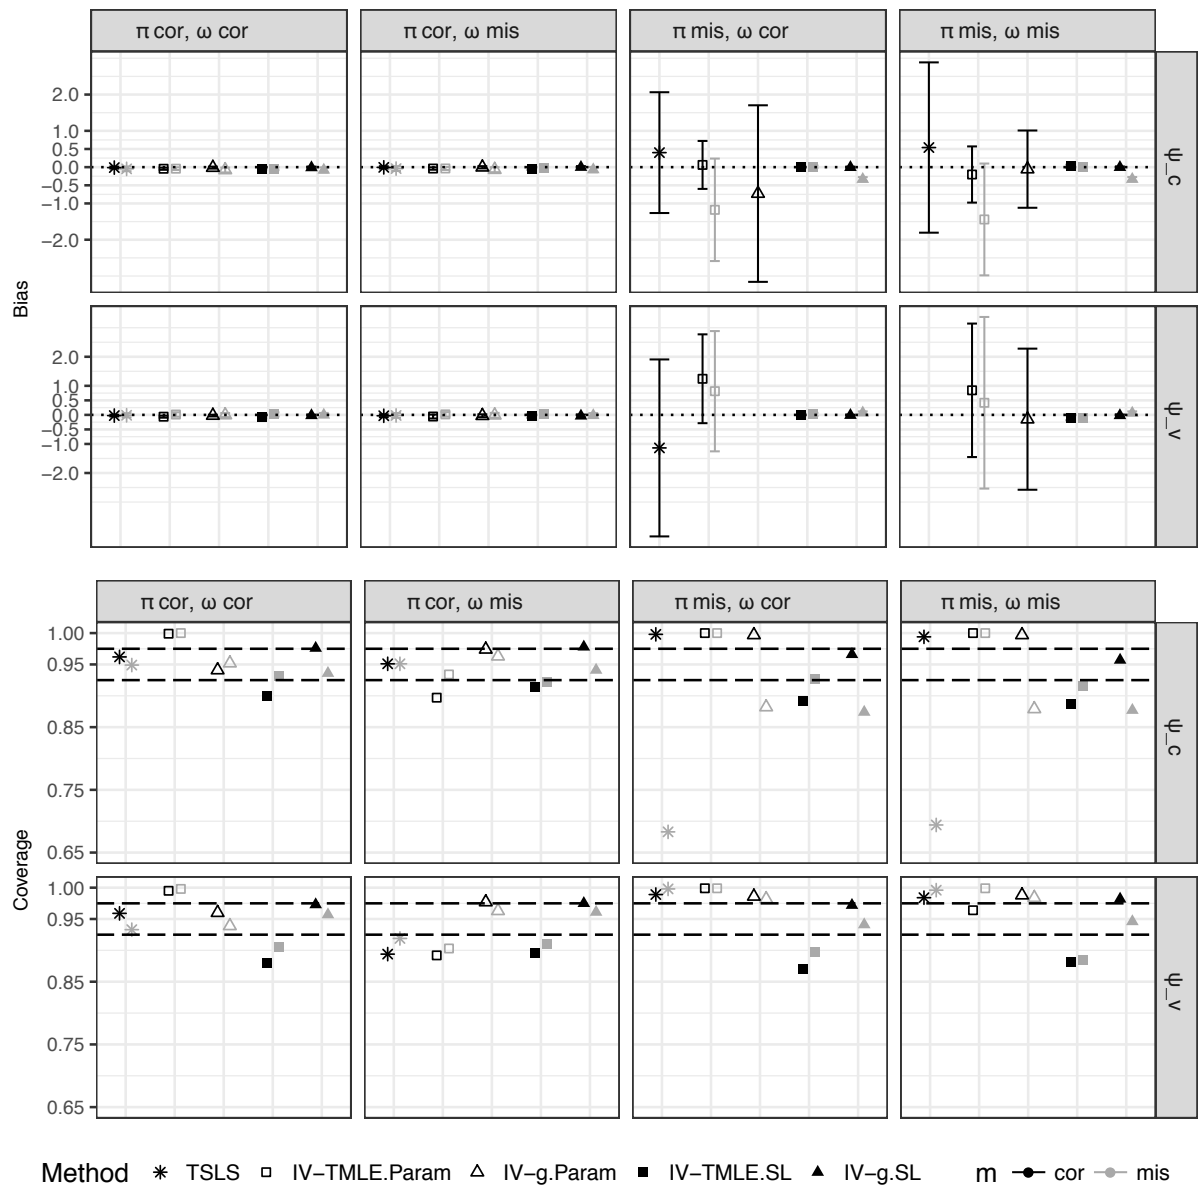



FIGURE 3. Performance (Bias and Coverage) of TSLs, TMLE and IV-g estimators, when the sample size is $n = 10,000$. Scenarios with correct or misspecified π and ω vary by column, m correctly specified is plotted in black while m misspecified is plotted in grey. The hollow shapes correspond to parametric nuisance models estimation, and the solid shapes correspond to estimators using data-adaptive nuisance model estimates. The bias is presented with Monte Carlo Error CIs. Results corresponding to bias ≥ 2 in absolute value are not plotted, but can be found in Table 4. 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.



levels are close to the nominal value (between 92.5 and 97.5%) for TSLS and IV-g estimator. In contrast, the bootstrapped CIs corresponding to parametric TMLE result in over-coverage (99%), while the EIF-based CIs for the data-adaptive TMLE shows under-coverage, which is especially in the small sample size scenarios ($n = 500$), dropping below 90% for ψ_v . This low-coverage phenomenon of the EIF-based CIs for TMLE estimators has been noted before by [van der Laan and Gruber \(2011\)](#) and [Petersen et al. \(2014\)](#).

TSLS performs well when m and π are correctly specified (second column), but when the exposure model π is misspecified (3rd and 4th columns), it performs poorly, even in scenarios where m and the IV model are correctly specified (plotted in black), demonstrating numerically the lack of double robustness. When m is misspecified (plotted in grey), TSLS results in bias $\geq 200\%$ of the true effect (not plotted in the Figures, see Table 4). Consequently, the coverage of the CIs is poor, being close to 0 in the larger sample size settings.

Both parametric TMLE and g-estimator result in small levels of bias and good coverage under those misspecified scenarios when the double robust properties are expected to provide protection. For example, with m correctly specified, the g-estimator has small bias and good coverage even when the exposure model π and the outcome model ω are misspecified (the last column of the Figures). TMLE on the other hand shows some significant bias, even at large samples $n = 10,000$. However, implementing the IV-g and IV-TMLE methods using the Super Learner returns the bias and coverage to the levels reported under correct specification, with TMLE still showing coverage under 92.5%.

Where m is misspecified, we would expect the g-estimator to behave as a projection of the true treatment effect curve onto the working parametric model $m_\psi(W)$, as long as the model for the exposure $\pi(Z, W)$ is correctly specified and it is such that $\text{Cov}(\{\pi_0(Z, W) - E(\pi_0(Z, W)|W)\}, A|W)$ is constant in W . Since the data generating models are such that the true $\pi_0(W, Z)$ has constant covariance with the received exposure given W , we can see in the first two columns of Figures 2 and 3, that the g-estimator performance is adequate when the parametric model $\pi(W, Z)$ is correctly specified (empty triangles plotted in grey). In contrast, for scenarios where the true $\pi_0(W, Z)$ does not have conditional constant covariance with A given W (third and fourth columns), there is substantial remaining bias even after using data-adaptive fits for the nuisance models, especially for the intercept ψ_c (see for example, in the last column of Figure 3 plotted in grey).

Tables 1 and 2 report the RMSE results. When m is correctly specified, IV-g outperforms all other methods, with the smallest RMSE. Where the working parametric model for the treatment effect curve $m(W)$ is misspecified, TMLE has smaller RMSE in most settings. Looking at the larger sample $n = 10,000$, we can 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 scenarios when the DR properties do not provide protection against misspecification.

TABLE 1. RMSE of the TSLS, TMLE and IV-g estimators, when the sample size is $n = 500$.

Scenario	Nuisance models estimation	Parameter	Method	RMSE			
				$m(W)$ correct	$m(W)$ mis		
π cor, ω cor	Parametric	ψ_c	TSLS	0.446	1.030		
			IV-g	0.443	1.084		
			IV-TMLE	0.473	0.606		
		ψ_v	TSLS	0.480	1.131		
			IV-g	0.479	1.132		
			IV-TMLE	0.580	1.234		
	SL	ψ_c	IV-g	0.439	1.159		
			IV-TMLE	0.475	0.614		
			ψ_v	IV-g	0.468	1.117	
			IV-TMLE	0.586	1.160		
		π cor, ω mis	Parametric	ψ_c	TSLS	0.520	1.065
					IV-g	0.517	1.119
IV-TMLE	0.548				0.655		
ψ_v	TSLS			0.782	1.314		
	IV-g			0.788	1.338		
	IV-TMLE			1.073	1.262		
SL	ψ_c		IV-g	0.495	1.183		
			IV-TMLE	0.616	0.685		
			ψ_v	IV-g	0.753	1.295	
			IV-TMLE	1.111	1.368		
	π mis, ω cor		Parametric	ψ_c	TSLS	26.835	160.523
					IV-g	39.241	446.003
IV-TMLE		10.649			22.791		
ψ_v		TSLS		49.141	163.553		
		IV-g		139.285	1596.750		
		IV-TMLE		24.685	33.392		
SL		ψ_c	IV-g	0.316	0.990		
			IV-TMLE	0.472	0.557		
			ψ_v	IV-g	0.309	0.822	
			IV-TMLE	0.756	0.858		
		π mis, ω mis	Parametric	ψ_c	TSLS	37.825	154.812
					IV-g	17.157	420.882
IV-TMLE	12.496				24.872		
ψ_v	TSLS			75.209	150.203		
	IV-g			39.172	1491.963		
	IV-TMLE			37.065	47.649		
SL	ψ_c		IV-g	0.367	1.011		
			IV-TMLE	0.743	0.788		
			ψ_v	IV-g	0.557	0.925	
			IV-TMLE	1.446	1.437		

TABLE 2. RMSE of TSLS, TMLE and IV-g estimators, when the sample size is $n = 10,000$.

Scenario	Nuisance models estimation	Parameter	Method	RMSE			
				$m(W)$ correct	$m(W)$ mis		
π cor, ω cor	Parametric	ψ_c	TSLS	0.092	0.207		
			IV-g	0.092	0.228		
			IV-TMLE	0.092	0.112		
		ψ_v	TSLS	0.090	0.213		
			IV-g	0.090	0.214		
			IV-TMLE	0.091	0.113		
	SL	ψ_c	IV-g	0.092	0.255		
			IV-TMLE	0.093	0.112		
			ψ_v	IV-g	0.090	0.215	
			IV-TMLE	0.093	0.114		
		π cor, ω mis	Parametric	ψ_c	TSLS	0.107	0.213
					IV-g	0.107	0.234
IV-TMLE	0.107				0.125		
ψ_v	TSLS			0.140	0.240		
	IV-g			0.134	0.236		
	IV-TMLE			0.140	0.156		
SL	ψ_c		IV-g	0.104	0.260		
			IV-TMLE	0.136	0.143		
			ψ_v	IV-g	0.133	0.237	
			IV-TMLE	0.206	0.163		
	π mis, ω cor		Parametric	ψ_c	TSLS	0.270	10.160
					IV-g	0.333	9.660
IV-TMLE		0.349			1.957		
ψ_v		TSLS		0.269	1.250		
		IV-g		0.345	1.890		
		IV-TMLE		0.382	0.682		
SL		ψ_c	IV-g	0.069	0.408		
			IV-TMLE	0.075	0.098		
			ψ_v	IV-g	0.066	0.196	
			IV-TMLE	0.076	0.100		
		π mis, ω mis	Parametric	ψ_c	TSLS	0.317	10.158
					IV-g	0.378	9.658
IV-TMLE	0.686				2.427		
ψ_v	TSLS			0.416	1.277		
	IV-g			0.456	1.908		
	IV-TMLE			0.875	0.822		
SL	ψ_c		IV-g	0.078	0.409		
			IV-TMLE	0.094	0.128		
			ψ_v	IV-g	0.107	0.212	
			IV-TMLE	0.153	0.516		

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 3 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 five estimators result in the same conclusions, namely that there is no evidence of an average treatment effect in the treated, and also that there is no evidence of effect modification by baseline depression. This result could be due to small numbers of participants in the trial, or indeed our definition of being exposed to treatment (attending at least one session). Nevertheless, the direction of the treatment effect modification is interesting, indicating that the treatment may benefit more those with higher depression symptoms at baseline, suggesting a reduction in the disability score.

TABLE 3. *ATT of the COPERS intervention on CPG, with all-or-nothing binary exposure A, main effect ψ_c and effect modification by depression ψ_v .*

	ψ_c	SE	95% CI	ψ_v	SE	95% CI
TSLS	2.94	4.67	(-6.21, 12.09)	-0.58	0.57	(-1.70, 0.54)
IV-g	2.78	4.66	(-6.35, 11.91)	-0.53	0.54	(-1.59, 0.53)
IV-g SL	2.10	4.75	(-7.21, 11.41)	-0.45	0.54	(-1.51, 0.61)
IV-TMLE	3.16	4.74	(-6.13, 12.45)	-0.64	0.56	(-1.74, 0.46)
IV-TMLE SL	2.22	4.88	(-7.34, 11.78)	-0.51	0.58	(-1.65, 0.63)

6. Discussion

This paper compared the performance of two doubly robust estimators for the ATT conditional on a baseline covariate, i.e. $ATT(v)$, 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 does not rely on a correctly specified parametric working model, and instead models the whole treatment effect curve, projecting the final estimates onto the working model of interest. This allows us to define the parameters of interest even under a misspecified treatment effect curve. However, it is less efficient compared with the IV g-estimator when the parametric working model for the treatment effect curve is correctly specified. The g-estimator on the other hand can suffer large biases when the assumed treatment effect curve is misspecified. As the simulations show, the use of data-adaptive fits for the nuisance models greatly reduces bias, and improves coverage for both estimators, 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 (Little and Yau, 1998; Hirano et al., 2000) as well as semi-parametric models (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 the instrument propensity score and either the outcome or the exposure parametric models are correctly specified.

For the $ATT(v)$, Robins (1994) proposed DR estimators in settings where $V = W$, while Tan (2010) did so in settings where 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 $\omega(W)$ (Vermeulen and Vansteelandt, 2016). Chernozhukov et al. (2018) proposed two other IV DR data-adaptive estimators and gave conditions under which data-adaptive fits can be used for the law of the instrument Z given W , $g(W)$, the treatment model $\pi(Z, W)$ and $\omega(W)$. Comparing these DR estimators to the those presented here would be a promising avenue for future research.

The present study has some limitations. Firstly, we did not use sample-splitting in our estimators. Evaluating the effect of doing so in point estimation and variance estimation is a promising extension. In addition, 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 $\omega(W)$ and $m(W)$, and thus used parametric models for both. This is not necessary, and an alternative strategy where $\omega(W)$ is estimated beforehand and the fitted values are plugged into the estimating equation (12) is possible, thus allowing the use of data-adaptive fits for the model $\omega(W)$. 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 was supported by UK Medical Research Council Career development award in Biostatistics MR/L011964/1 and UK Wellcome Trust Institutional Strategic Support Fund-LSHTM Fellowship 204928/Z/16/Z.

Appendix

6.1. Consistent data-adaptive g-estimation

Here, we give here a sketch of the proof of consistency of the data-adaptive g-estimator. Similarly to Chernozhukov et al. (2018) and Farrell (2015), we let \mathcal{P} be the class of probability distributions for O that obey the partially linear IV model, such that for each $P \in \mathcal{P}$, the restrictions $E[Y - Am_0(W)|W] = \omega_0(W)$, $E[A|Z, W] = \pi_0(Z, W)$, and $E[Z|W] = g_0(W)$ hold. Let η_0 denote the nuisance functional describing $g_0(W)$, $\omega_0(W)$, and $\pi_0(Z, W)$. For simplicity, we sketch the arguments under the null, that is $m_0(W; \psi_0) = 0$. In addition, unlike Chernozhukov et al. (2018), we do not use sample splitting, and proceed instead under empirical processes conditions which are from now on assumed to hold.

Denote by $\varepsilon_Y = Y - \omega_0(W)$ and $\varepsilon_A = \pi_0(Z, W) - E(\pi_0(Z, W)|W)$. Since the instrument Z is binary, we can write $E(\pi_0(Z, W)|W) = \sum_Z \pi_0(Z, W)g_0(W)$, where we use the shorthand $g_0(W) = g_0(Z = z, W) = Pr(Z = z|W)$.

We want to find conditions guaranteeing that $\sqrt{n}(\hat{\psi} - \psi_0) = o_p(1)$, where $\hat{\psi}$ has been estimated with the IV g-estimator which used data-adaptive estimates for the nuisance parameters η .

We begin by writing

$$\begin{aligned} \sqrt{n}(\hat{\psi} - \psi_0) = & \\ & \frac{1}{\sqrt{n}} \sum_i \mathcal{S}(O_i; \psi_0, \hat{\eta}_0) - E_P[\mathcal{S}(O_i; \psi_0, \hat{\eta}_0)] - \mathcal{S}(O_i, \psi_0, \eta_0) - E_P[\mathcal{S}(O_i; \psi_0, \eta_0)] + \sqrt{n}E_P[\mathcal{S}(O_i, \psi_0, \hat{\eta}_0)], \end{aligned} \quad (24)$$

where \mathcal{S} is the score corresponding to the estimating equation (12), with $\sigma_0^2 = 1$, i.e.:

$$\mathcal{S}(O_i; \psi_0, \eta_0) = \{\pi_0(Z_i, W_i) - E_P[\pi_0(Z_i, W_i)g_0(W_i)]\} \{Y_i - \omega_0(W_i) - m_0(W_i; \psi_0)A_i\}, \quad (25)$$

The first part can be shown to be $o_p(1)$ where $\|\mathcal{S}(O_i, \psi_0, \hat{\eta}_0) - \mathcal{S}(O_i, \psi_0, \eta_0)\| = o_p(1)$, by Chebyshev's inequality.

Therefore, we want to give sufficient conditions for

$$\|\mathcal{S}(O_i, \psi_0, \hat{\eta}_0) - \mathcal{S}(O_i, \psi_0, \eta_0)\|_p = o_p(1). \quad (26)$$

Using recursive expansion around each true nuisance functional, $\mathcal{S}(O_i, \psi_0, \widehat{\eta}_0)$ can be written as

$$\begin{aligned}
&= \{ \pi_0(Z, W) - \sum_z \pi_0(Z, W) g_0(W) \} \{ Y - \omega_0(W) \} + \{ \pi_0(Z, W) - \sum_z \pi_0(Z, W) g_0(W) \} (\omega_0(W) - \widehat{\omega}_0(W)) \\
&+ \sum_z \pi_0(Z, W) (g_0(W) - \widehat{g}_0(W)) \{ Y - \omega_0(W) \} + \sum_z \pi_0(Z, W) (g_0(W) - \widehat{g}_0(W)) (\omega_0(W) - \widehat{\omega}_0(W)) \\
&+ \{ (\pi_0(Z, W) - \widehat{\pi}_0(Z, W)) - \sum_z (\pi_0(Z, W) - \widehat{\pi}_0(Z, W)) g_0(W) \} \{ Y - \omega_0(W) \} \\
&+ \{ (\pi_0(Z, W) - \widehat{\pi}_0(Z, W)) - \sum_z (\pi_0(Z, W) - \widehat{\pi}_0(Z, W)) g_0(W) \} (\omega_0(W) - \widehat{\omega}_0(W)) \\
&+ \sum_z (\widehat{\pi}_0(Z, W) - \pi_0(Z, W)) (g_0(W) - \widehat{g}_0(W)) \{ Y - \omega_0(W) \} \\
&+ \sum_z (\widehat{\pi}_0(Z, W) - \pi_0(Z, W)) (g_0(W) - \widehat{g}_0(W)) (\omega_0(W) - \widehat{\omega}_0(W)),
\end{aligned}$$

which can be further simplified to:

$$\begin{aligned}
&= \varepsilon_A \varepsilon_y + \varepsilon_A \{ \omega_0(W) - \widehat{\omega}_0(W) \} - \sum_z \pi_0(Z, W) \{ g_0(W) - \widehat{g}_0(W) \} \varepsilon_y \\
&- \sum_z \pi_0(Z, W) (g_0(W) - \widehat{g}_0(W)) (\omega_0(W) - \widehat{\omega}_0(W)) \\
&+ \varepsilon_y \left\{ (\widehat{\pi}_0(Z, W) - \pi_0(Z, W)) - \sum_z (\widehat{\pi}_0(Z, W) - \pi_0(Z, W)) g_0(W) \right\} \\
&+ (\widehat{\pi}_0(Z, W) - \pi_0(Z, W)) (\omega_0(W) - \widehat{\omega}_0(W)) \\
&- \sum_z (\widehat{\pi}_0(Z, W) - \pi_0(Z, W)) g_0(W) (\omega_0(W) - \widehat{\omega}_0(W)) \\
&+ \varepsilon_y \sum_z (\widehat{\pi}_0(Z, W) - \pi_0(Z, W)) (g_0(W) - \widehat{g}_0(W)) \\
&- \sum_z (\widehat{\pi}_0(Z, W) - \pi_0(Z, W)) (g_0(W) - \widehat{g}_0(W)) (\omega_0(W) - \widehat{\omega}_0(W)).
\end{aligned}$$

Therefore

$$\begin{aligned}
& \|\mathcal{S}(O_i, \psi_0, \widehat{\eta}_0) - \mathcal{S}(O_i, \psi_0, \eta_0)\|_P \\
\leq & \|\varepsilon_A \{\omega_0(W) - \widehat{\omega}_0(W)\}\|_P + \|\varepsilon_y \left\{ (\widehat{\pi}_0(Z, W) - \pi_0(Z, W)) - \sum_z (\widehat{\pi}_0(Z, W) - \pi_0(Z, W)) g_0(W) \right\}\|_P \\
& + \|\varepsilon_y \sum_z \pi_0(Z, W) (\widehat{g}_0(W) - g_0(W))\|_P \\
& + \|(\omega_0(W) - \widehat{\omega}_0(W)) \left\{ (\widehat{\pi}_0(Z, W) - \pi_0(Z, W)) - \sum_z (\widehat{\pi}_0(Z, W) - \pi_0(Z, W)) g_0(W) \right\}\|_P \\
& + \|(\omega_0(W) - \widehat{\omega}_0(W)) \sum_z \pi_0(Z, W) (\widehat{g}_0(W) - g_0(W))\|_P \\
& + \left\| \sum_z (\widehat{\pi}_0(Z, W) - \pi_0(Z, W)) (g_0(W) - \widehat{g}_0(W)) (\omega_0(W) - \widehat{\omega}_0(W)) \right\|_P \\
\leq & \sqrt{C} \|\omega_0(W) - \widehat{\omega}_0(W)\|_P + \sqrt{C} \left\| \left\{ (\widehat{\pi}_0(Z, W) - \pi_0(Z, W)) - \sum_z (\widehat{\pi}_0(Z, W) - \pi_0(Z, W)) g_0(W) \right\} \right\|_P \\
& + \sqrt{C} \|(\pi_0(1, W) - \pi_0(0, W)) (\widehat{g}_0(W) - g_0(W))\|_P \\
& + \|(\omega_0(W) - \widehat{\omega}_0(W)) \left\{ (\widehat{\pi}_0(Z, W) - \pi_0(Z, W)) - \sum_z (\widehat{\pi}_0(Z, W) - \pi_0(Z, W)) g_0(W) \right\}\|_P \\
& + \|(\omega_0(W) - \widehat{\omega}_0(W)) (\pi_0(1, W) - \pi_0(0, W)) (\widehat{g}_0(W) - g_0(W))\|_P \\
& + \left\| \sum_z (\widehat{\pi}_0(Z, W) - \pi_0(Z, W)) (g_0(W) - \widehat{g}_0(W)) (\omega_0(W) - \widehat{\omega}_0(W)) \right\|_P,
\end{aligned}$$

where we assume that there exists a constant $C > 0$, such that

$$\begin{aligned}
P(E[\{\pi_0(Z, W) - \sum_z \pi_0(Z, W) g_0(W)\}^2] \leq C) &= 1, \text{ and} \\
P(E[\{Y - \omega_0(W)\}^2] \leq C) &= 1.
\end{aligned} \tag{27}$$

Finally, since $\|(\pi_0(1, W) - \pi_0(0, W))\|$ is bounded and

$$\|\widehat{\pi}_0(Z, W) - \pi_0(Z, W)\| = o_{P_0}(1), \tag{28}$$

$$\|\widehat{\omega}_0(W) - \omega_0(W)\| = o_{P_0}(1), \tag{29}$$

$$\|\widehat{g}_0(W) - g_0(W)\| = o_{P_0}(1), \tag{30}$$

by definition of $\omega_0(W)$ and $g_0(W)$, we conclude that the assumption (26) needed for the first part of eq. (24) to be $o_P(1)$ holds.

Now, for the second term in eq. (24), we want conditions such that

$$\sqrt{n} E_P[\mathcal{S}(\psi_0, \widehat{\eta}_0)] = o_P(1), \tag{31}$$

we have

$$\begin{aligned}
\sqrt{n}E_P[\mathcal{L}(\psi_0, \widehat{\eta}_0)] &= \sqrt{n}E_P[\widehat{\pi}_0(Z, W) - E_P[\widehat{\pi}_0(Z, W)](Y - \widehat{\omega}_0(W))] \\
&= \sqrt{n}E_P[\widehat{\pi}_0(Z, W) - \sum_z \{\widehat{\pi}_0(Z, W)\widehat{g}_0(W)\}(Y - \widehat{\omega}_0(W))] \\
&= \sqrt{n}E_P[\widehat{\pi}_0(Z, W) - \sum_z \{\widehat{\pi}_0(Z, W)\widehat{g}_0(W)\}(\omega_0 - \widehat{\omega}_0(W))] \\
&= \sqrt{n}E_P \left[\widehat{\pi}_0(Z, W) - \sum_z \{\widehat{\pi}_0(Z, W)(g_0(W) - \widehat{g}_0(W))\}(\omega_0 - \widehat{\omega}_0(W)) \right] \\
&= \sum_z \{\widehat{\pi}_0(Z, W)(g_0(W) - \widehat{g}_0(W))\} \{\omega_0(W) - \widehat{\omega}_0(W)\}.
\end{aligned}$$

Now the norm of the first term of this expression is such that

$$\begin{aligned}
\left\| \sum_z \widehat{\pi}_0(Z, W)(g_0(W) - \widehat{g}_0(W)) \right\| &= \left\| \widehat{\pi}_0(1, W)(g_0(1, W) - \widehat{g}_0(1, W)) + \widehat{\pi}_0(0, W)(g_0(0, W) - \widehat{g}_0(0, W)) \right\| \\
&= \left\| (\widehat{\pi}_0(1, W) - \widehat{\pi}_0(0, W))(g_0(1, W) - \widehat{g}_0(1, W)) \right\| \\
&\leq \left\| (\widehat{\pi}_0(1, W) - \widehat{\pi}_0(0, W)) \right\| \left\| (g_0(1, W) - \widehat{g}_0(1, W)) \right\|,
\end{aligned}$$

where we have used Cauchy-Schwarz inequality in the last step. Now, since $\left\| (\widehat{\pi}_0(1, W) - \widehat{\pi}_0(0, W)) \right\|$ is bounded, assuming

$$\left\| g_0(W) - \widehat{g}_0(W) \right\| \left\| \omega_0(W) - \widehat{\omega}_0(W) \right\| = o_P(n^{-\frac{1}{2}}), \quad (32)$$

is sufficient to guarantee eq. (31) holds.

In summary, to guarantee the data-adaptive IV g-estimator is CAN assumptions (27), (28) and (32) need to hold. These conditions are essentially the same found by Chernozhukov et al. (2018).

6.2. *Extra results*TABLE 4. *Scenarios excluded from the Bias and Coverage Figures.*

Scenario	m model	Parameter	Method	bias	MCE CI		coverage
<i>n = 500</i>							
π mis, ω cor	cor	ψ_v	IV-g ^a	-4.057	-12.691	4.577	0.986
	mis	ψ_c	TSLs	-6.881	-16.826	3.064	0.683
	mis	ψ_c	IV-g ^a	-31.033	-58.624	-3.442	0.882
	mis	ψ_v	TSLs	6.785	-3.348	16.918	0.998
	mis	ψ_v	IV-g ^a	-45.811	-144.787	53.165	0.982
π mis, ω mis	cor	ψ_v	TSLs	-2.171	-6.834	2.492	0.984
	mis	ψ_c	TSLs	-6.740	-16.330	2.850	0.694
	mis	ψ_c	IV-g ^a	-30.361	-56.392	-4.330	0.879
	mis	ψ_v	TSLs	5.751	-3.557	15.059	0.996
	mis	ψ_v	IV-g ^a	-41.902	-134.385	50.581	0.984
<i>n = 10,000</i>							
π mis, ω cor	mis	ψ_c	TSLs	-10.101	-10.169	-10.032	0.000
	mis	ψ_c	IV-g ^a	-9.521	-9.622	-9.420	0.001
π mis, ω mis	mis	ψ_c	TSLs	-10.097	-10.166	-10.028	0.000
	mis	ψ_c	IV-g ^a	-9.518	-9.620	-9.416	0.001
	mis	ψ_c	IV-TMLE ^a	-2.333	-2.374	-2.291	1.000

^a All nuisance models fitted parametrically.

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.
- Athey, S., G. W. Imbens, and S. Wager (2018). Approximate residual balancing: debiased inference of average treatment effects in high dimensions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 80(4), 597–623.
- 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.
- Bickel, P., F. (1982). On adaptive estimation. *Annals of Statistics* 10, 647–71
- Bickel, P., F. Götzte, and W. van Zwet (1997). Resampling fewer than n observations: Gains, losses, and remedies for losses. *Statistica Sinica* 7, 1–31.
- Chernozhukov, V., D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, and W. Newey (2017). Double/debiased/neyman machine learning of treatment effects. *American Economic Review* 107(5), 261–65.
- Chernozhukov, V., D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, and J. Robins (2018). Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal* 21(1), C1–C68.
- 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.
- Kennedy, E. H. (2016). *Semiparametric Theory and Empirical Processes in Causal Inference*, Chapter Statistical Causal Inferences and Their Applications in Public Health Research, pp. 141–167. Cham: Springer International Publishing.
- 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.
- Newey, W. K. and McFadden, D. (1994) Large sample estimation and hypothesis testing, in *Handbook of Econometrics*, Elsevier B.V., 2111–2245.
- 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*, 6–10.
- Rubin, D. B. (1978). Bayesian inference for causal effects: The role of randomization. *The Annals of statistics*, 34–58.
- Swanson, S. A., Hernan, M. A., Miller, M., Robins, J. M. and Richardson T. S. (2018) Partial Identification of the Average Treatment Effect Using Instrumental Variables: Review of Methods for Binary Instruments, Treatments, and Outcomes. *Journal of the American Statistical Association* **113**(522):933–947.
- 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. J. and D. Rubin (2006). Targeted maximum likelihood learning. *The International Journal of Biostatistics* 2(1).
- 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. and Rose, S. (2018). *Targeted Learning in Data Science: Causal Inference for Complex Longitudinal Studies*. Springer International Publishing
- van der Laan, M. J. and Robins, J.M., (2003). *Unified Methods for Censored Longitudinal Data and Causality*, Springer Series in Statistics.
- van der Laan, M. J. and Gruber, S. (2011). Targeted minimum loss based estimation of an intervention specific mean outcome. *newblock U.C. Berkeley Division of Biostatistics Working Paper Series. Working Paper 290*.
- 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.
- van der Vaart, A. (2014). Higher Order Tangent Spaces and Influence Functions. *Statistical Science* 29(4), 679–686.
- 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.
- Zheng, W. and van der Laan, M. J. (2011). Cross-validated targeted minimum-loss-based estimation. *Targeted*

Learning. Springer, New York, USA. pp. 459–474.