

RESEARCH ARTICLE**Tutorial: Introduction to computational causal inference using reproducible Stata, R and Python code**Matthew J. Smith¹ | Camille Maringe¹ | Bernard Rachet¹ | Mohammad A. Mansournia² | Paul N. Zivich³ | Stephen R. Cole³ | Miguel Angel Luque-Fernandez^{*1,4,5}¹Inequalities in Cancer Outcomes Network, Department of Non-communicable Disease Epidemiology. London School of Hygiene and Tropical Medicine, London, U.K.²Department of Epidemiology and Biostatistics. Tehran University of Medical Sciences, Tehran, Iran.³Department of Epidemiology. University of North Carolina at Chapel Hill, North Carolina, U.S.⁴Non-communicable Disease and Cancer Epidemiology Group, Instituto de investigacion Biosanitaria de Granada (ibs.GRANADA), University of Granada, Granada, Spain⁵Biomedical Network Research Centers of Epidemiology and Public Health (CIBERESP), Madrid, Spain.**Correspondence**

* Miguel Angel Luque-Fernandez, Email: miguel-angel.luque@lshtm.ac.uk

Present Address

Keppel St, Bloomsbury, London WC1E 7HT, United Kingdom

Abstract The purpose of many health studies is to estimate the effect of an exposure on an outcome. It is not always ethical to assign an exposure to individuals in randomised controlled trials, instead observational data and appropriate study design must be used. There are major challenges with observational studies, one of which is confounding that can lead to biased estimates of the causal effects. Controlling for confounding is commonly performed by simple adjustment for measured confounders; although, often this is not enough. Recent advances in the field of causal inference have dealt with confounding by building on classical standardisation methods. However, these recent advances have progressed quickly with a relative paucity of computational-oriented applied tutorials contributing to some confusion in the use of these methods among applied researchers. In this tutorial, we show the computational implementation of different causal inference estimators from a historical perspective where different estimators were developed to overcome the limitations of the previous one. Furthermore, we also briefly introduce the potential outcomes framework, illustrate the use of different methods using an illustration from the health care setting, and most importantly, we provide reproducible and commented code in Stata, R and Python for researchers to apply in their own observational study. The code can be accessed at <https://github.com/migariane/TutorialCausalInferenceEstimators>

KEYWORDS:

Causal Inference; Regression adjustment; G-methods; G-formula; Propensity score; Inverse probability weighting; Double-robust methods; Machine learning; Targeted maximum likelihood estimation; Epidemiology; Statistics; Tutorial

1 | INTRODUCTION

The questions that motivate most studies in the health, social and behavioral sciences are causal in nature. For example, what is the mortality risk difference amongst patients offered a surgical procedure versus those who did not receive it in a given population?¹ Causal inference methods may be used to answer this scientific question where a clinical trial is unfeasible or unethical (e.g. a surgical procedure).² Causal inference methods are based on the counterfactual framework introduced by Neyman in the randomized experiment setting, then extended to time-fixed and time-varying observational studies by Rubin

and Robins, respectively.^{3,4}

Over the years, rapid on-going advances in the field causal inference have created a set of different approaches to estimate the causal effect of a treatment (or exposure) on an outcome (i.e., methods that incorporate propensity scores, the G-computation, or a combination of both, namely double-robust methods). Overall, double-robust methods now are preferred over naive regression approaches because the latter are biased under misspecification of a parametric outcome model when the research question is causal in nature.^{5,6}

In this tutorial we introduce the most commonly used estimators in causal inference from a practical computational perspective: allowing readers to learn by using the replicable code. We use the Stata statistical software (*StataCorp. 2020. Stata Statistical Software: Release 16. College Station, TX: StataCorp LLC, USA*). We also provide readers with access to code in R statistical software (*R Development Core Team (2020). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria*) and Python statistical software (*Python Software Foundation (2020)*) and a link to a GitHub repository where readers can access the data, along with all software code, to replicate our empirical example accessed at <https://github.com/migariane/TutorialCausalInferenceEstimators>.

To illustrate the computation of the most common causal inference estimators we use a classical empirical data set from the prospective cohort study of Connors *et al* (1996).⁷ The data is open source and available at: <http://biostat.mc.vanderbilt.edu/wiki/Main/DataSets>. The study evaluated the effectiveness of right heart catheterisation (RHC) on short-term mortality (30 days) of 5,735 critically ill adult patients, 2,184 treated and 3,551 controls in five United States teaching hospitals receiving care in an ICU for 1 of 9 prespecified disease categories between 1989 and 1994. In our illustration, the outcome is short-term mortality (Y) defined as 30 days after ICU admission, RHC as the main treatment or exposure (A), and in (W) we include the set of confounders.

In the following sections, we will illustrate the computational implementation of different estimators for our estimand of interest: the average treatment effect (ATE), also known as the risk difference. In section two we introduce the causal language, followed by section three that details the computation of the G-formula non-parametrically and parametrically using the G-computation method. Then, in section four we introduce the methods based on the inverse probability-of-treatment weights (IPTW). Afterwards, in section five, we shortly introduce the computation of double-robust methods, and finally in section six we present the targeted maximum likelihood estimation (TMLE) framework. We illustrate the computation of the different estimators using Stata in a set of boxes containing commented code.

2 | CAUSAL LANGUAGE: POTENTIAL OUTCOMES

Commonly, the aim of an observational study is to answer a scientific question that characterises the effect of an exposure or treatment on an outcome. This question is translated to an *estimand* whose estimate will provide a relevant answer to the characterisation of the effect of the exposure on the outcome (i.e., by how many times is the risk of the outcome increased, given a particular level of the exposure). On the other hand, *estimators* are mathematical functions using the observed values of the observations in the sample (that is, a function of the random variables) that generates the quantitative value for the estimand. The estimators are represented by algebraic equations that explicitly describe a function of the realized observations.

The most common estimand in causal inference is the ATE and is the focus of this tutorial. The ATE is a function of the underlying distribution of the counterfactual outcomes, which can be estimated non-parametrically or parametrically.⁴ The ATE is defined by an average of the difference of two random variables (i.e., the potential outcomes $Y(1)$ and $Y(0)$).^{2,8} Thus, to describe the ATE we need to introduce the language of the Potential Outcomes Framework also known as (a.k.a) Neyman-Rubin Potential Outcomes framework.²

To illustrate the framework we use the study of Connors based on intensive care medicine *et al* (1996).⁷ Let Y , the outcome, denote the vital status of the patient in an intensive care unit (ICU) at 30 days after admission. Let A , the exposure, denote the

treatment variable of whether or not the patient received right heart catheterisation (RHC) during his stay at the ICU. Let C , a confounder, denote the gender of the patient.

In this RHC study, each patient has two *potentially observed* outcomes (i.e., $Y(a)$), where the first is $Y(1)$ if they received RHC, and the second is $Y(0)$ if they did not receive RHC.³ We say *potentially observed* because only one of these two outcomes can ever be observed as each patient only receives one of the treatments. As an example from table 1, Ronald has two potential outcomes: firstly, $Y(0) = 1$ says that if Ronald were not to receive RHC then he would die within 30 days, and secondly, $Y(1) = 0$ says that if Ronald were to receive RHC then he would not die within 30 days. The two potential outcomes mentioned thus far can also be stated as $Y(0)$ and $Y(1)$, respectively.

Patient	Y	A	C	Y(0)	Y(1)
Ronald	1	0	0	1	0
Malala	1	1	1	1	1
Florence	1	1	1	0	1
Karl	0	1	0	0	0
Marie	1	0	1	1	1
Greta	0	1	1	0	0
Miguel	1	0	0	1	0
Alice	0	1	1	0	0
Matthew	1	1	0	1	1
Ada	0	0	1	0	0

TABLE 1 Potential outcomes framework: C = Binary confounder, A = Binary exposure, Y = Binary outcome

Causal effects of a treatment on an outcome are defined in terms of a contrast between the *potential outcomes* under different treatment levels (i.e. $Y(0)$ and $Y(1)$), as formalised by Rubin.⁹ An individual causal effect exists for patient i if $Y_i(0) \neq Y_i(1)$. For example, there is an individual causal effect for Ronald as $Y_i(0) \neq Y_i(1)$, but not for Karl as $Y(0) = Y(1) = 0$. However, only one potential outcome can ever be observed, corresponding to the treatment that was actually taken. Therefore, individual causal effects are not possible to estimate. To elaborate, in this example, the causal effect of RHC on risk of death is a comparison of the potential outcomes for patient i , such that the ATE is defined as

$$\text{ATE} = E[\underbrace{P(Y(1)|A=1) + P(Y(1)|A=0)}_{\text{Unobserved}}] - E[\underbrace{P(Y(0)|A=0) + P(Y(0)|A=1)}_{\text{Unobserved}}],$$

where \mathbf{E} refers to expectation (i.e., average) and \mathbf{P} refers to probability. The potential outcome, $Y(1)|A=0$, is counterfactual to $Y(1)|A=1$ because $Y(1)|A=0$ states what the outcome for patient i (who did not receive the treatment, i.e., $A=0$) would have been if the patient had counterfactually received it, i.e., $A=1$. Ronald's observed outcome was death before 30 days ($Y(0)=1$) given he did not receive RHC, and his (unobserved) counterfactual outcome was alive at 30 days ($Y(1)=0$) given he did receive RHC.

Since the patient receives only $A=1$ or $A=0$, only one outcome for each patient can be observed, the counterfactual is not observed. In experimental designs, the randomisation makes the ATE estimate, i.e., $E(Y|A=1) - E(Y|A=0)$, unbiased for the ATE estimand:

$$\text{ATE} = E[Y(1) - Y(0)].$$

Since randomisation makes confounding random and the expected value of the distribution of confounders averaged over

allocations are the same:¹⁰

$$E[P(Y(1)|A = 1)] = E[P(Y(1)|A = 0)] \text{ and } E[P(Y(0)|A = 0)] = E[P(Y(0)|A = 1)].$$

However, randomisation is often unethical in clinical trials or unfeasible in observational studies. We must make certain (untestable) assumptions to identify potential outcomes from the observed data and estimate the causal estimand.⁴ Given that the potential outcomes are not necessarily directly observed from the data, to identify the ATE from observable data (i.e., from Table 1) the following three assumptions are made:

1. Counterfactual consistency

Counterfactual consistency holds if the observed outcome for all exposed individuals equals their outcome if they had been exposed, and likewise for unexposed individuals. This means that the definition of the exposure is consistent for all individuals. For example, in table 1, Ronald's observed outcome equals his potential outcome if he had not been exposed ($Y(0) = 1$). This means that the definition of the exposure is consistent for Ronald (the same applies for all the other patients). Analytically, consistency is represented by:

$$Y = AY(1) + (1 - A)Y(0)$$

2. Conditional exchangeability

In randomised studies, conditional exchangeability, and in fact, marginal exchangeability and full exchangeability, hold because the exposed individuals, had they not been exposed, would have had the same average outcome as the unexposed, and vice versa. This cannot be guaranteed in observational studies but it can be assumed to hold if the unmeasured risk factors of the outcome are equally distributed between the exposed and the unexposed groups conditional on the measured confounders. Thus, using the language of the potential outcomes the conditional exchangeability assumption (a.k.a conditional independence, unconfoundedness or ignorability) is given:

$$Y(a) \perp\!\!\!\perp A \mid \mathbf{C} \quad \forall a \in \{0, 1\}$$

Hence, the conditional mean independence is given

$$E[Y(a) \mid A = 1, C = c] = E[Y(a) \mid A = 0, C = c] = E[Y(a) \mid C = c] \quad \forall a \in \{0, 1\}$$

3. Positivity

Positivity holds if the conditional probability of being exposed (and similarly for being unexposed) is greater than zero. Therefore, if $P(C = c) > 0$ then $P(A = a \mid C = c) > 0 \quad \forall C \in \mathbf{c}, a \in \{0, 1\}$. We must also assume noninterference, meaning that the potential outcome of one individual was not influenced by the exposure of any other individual (i.e., independence of observations).

Transitioning from unobserved potential outcomes to a setting where we estimate the causal effect can be done by applying these assumptions. Suppose we are interested in estimating the ATE, and for simplicity, under the assumption there is only one confounder, namely C .

$$ATE = E[Y(1)] - E[Y(0)]$$

Note that we do have the possibility to estimate the average treatment among the treated (ATT) represented by:

$$ATT = E[Y(1) \mid A = 1] - E[Y(0) \mid A = 1]$$

By the law of total probability

$$P[Y(a) = 1] = \sum_{\mathbf{c}} P[Y(a) = 1 \mid C = c] P(C = c)$$

By conditional exchangeability the right hand side is

$$\sum_c P[Y(a) = 1 \mid A = a, C = c] P(C = c)$$

This is possible since we are assuming that within levels of C the predictors of the outcome are equally distributed between the treatment (RHC) and non-treatment (non-RHC) groups: that is we have achieved what would happen if patients were randomised to each treatment group. If we assume consistency the right hand side is

$$\sum_c P[Y = 1 \mid A = a, C = c] P(C = c)$$

Thus, under these assumptions, the ATE is defined as

$$P(Y(1) = 1) - P(Y(0) = 1) \\ \sum_c P[Y = 1 \mid A = 1, C = c] \Pr(C = c) - \sum_c P[Y = 1 \mid A = 0, C = c] P(C = c), \quad (1)$$

and the average treatment effect among the treated (ATT) is defined as

$$\sum_c P[Y = 1 \mid A = 1, C = c] \Pr(C = c \mid A = 1) - \sum_c P[Y = 1 \mid A = 0, C = c] P(C = c \mid A = 1) \quad (2)$$

We have transitioned from the unobserved potential outcomes to a setting where we can estimate our causal estimand using the distribution of the observed data using formula (1), namely the G-formula.¹¹ To do it we are going to compute different estimators of the ATE using the prospective cohort study of Connors *et al* (1996).⁷ Figure (1) represents the association of the vector of predefined confounders (i.e., \mathbf{W}) on the treatment (A) and the outcome (Y).




FIGURE 1 Y: mortality within 30 days ; A: right heart catheterisation; \mathbf{W} : full set of confounders outlined in Connors *et al.* (1996)

We assume that: all the variables included in \mathbf{W} are confounders of the effect of A on Y; there are no intermediate variables; and there is no residual confounding. Therefore, the set of covariates included in \mathbf{W} suffices to assume the conditional mean independence to estimate the ATE. In our example the outcome (Y) is mortality within 30 days, the treatment (A) is the right heart catheterisation, and we use patients' sex, age, education, race and the presence of a carcinoma as confounders represented by the vector \mathbf{W} in Figure 1. Note, that the original data include a highly dimensional set of confounders; however, for illustrative purposes, we include five of confounders. With Connor's study as an example, and using a set of different causal inference estimators for the ATE and the ATT, we estimate the standardised short term mortality risk difference of death for those patients who were offered RCH versus those who were not offered RHC.

3 | G-FORMULA: FORMAL DEFINITION AND APPLICATION

3.1 | Non-parametric G-formula

Classical epidemiological methods use regression adjustment to explain the main effect of a risk factor on a disease or outcome. Regression adjustment, controlling for confounding, requires making the assumption that the effect is constant across levels of confounders (\mathbf{W}) included in the model.⁶ However, in non-randomised observational studies, the effect measure is confounded given the different distribution of individual characteristics at baseline. In the instance of differential age distributions between two groups, age-standardisation to one population's age distribution, or to an external standard, is performed. The G-formula is a generalisation of the classical standardisation procedure, and it allows obtaining an unconfounded marginal estimation of the ATE that is given for a binary treatment by:⁴

$$\text{ATE} = \sum_{\mathbf{w}} [P(Y = 1 | A = 1, \mathbf{W} = \mathbf{w}) - P(Y = 1 | A = 0, \mathbf{W} = \mathbf{w})] P(\mathbf{W} = \mathbf{w}). \quad (3)$$

where,

$$P(Y = 1 | A = a, \mathbf{W} = \mathbf{w}) = \frac{P(\mathbf{W} = \mathbf{w}, A = a, Y = 1)}{\sum_y P(\mathbf{W} = \mathbf{w}, A = a, Y = y)}$$

is the conditional probability distribution of the outcome $Y = y$, given the exposure or treatment $A = a$, and the set of confounders $\mathbf{W} = \mathbf{w}$. Note that the implementation of the G-formula only requires the use of the total law of probability and the expectation of random variables. In probability theory, the total probability is a fundamental rule relating marginal probabilities to conditional probabilities. Likewise, marginal and conditional probabilities are used to estimate the expectation of random variables.

In the following set of 7 boxes we show how to compute the estimate of the effect of RHC on short term mortality using the non-parametric G-formula using Stata (the same computing approach is provided in the appendix for R and Python users). We will use gender as the unique confounder. It is an unrealistic scenario but it is made intentionally to allow readers to readily appreciate the computing implementation of the non-parametric G-formula. In contrast to the naive approach, the way we adjust for confounding based on the generalisation of standardisation (G-formula) is more consistent: we assume that the RHC effect on mortality can differ by sex while using the naive approach assumes that the effect is the same for both sexes.⁶ It is a subtle difference but provides a better adjustment for confounding. We do the same for the average treatment among the treated (ATT). Sometimes, the ATT is of interest (i.e., under the presence of near positivity violations). We use the bootstrap for statistical inference and provide 95% confidence intervals (CI) for the ATE. Finally, we also show readers how to implement the non-parametric estimation of the ATE and ATT using a regression modelling approach based on a fully saturated regression model.

Box 1: Setting the data

Here we declare the global variables Y, A, C, and W to match the presented algebraic nomenclature. i.e., Y: outcome, A: treatment, C: one unique confounder, and W: a set of confounders. We will use these global variables throughout the implementation of the different methods.

```

1  clear
2  set more off
3  use "rhc.dta"
4  global Y death_d30 // Outcome
5  global A rhc // Exposure or treatment
6  global C gender // One unique confounder of the set of W confounders
7  global W gender age edu race carcinoma // A set of five confounders

```

Box 2: Regression - a naive approach

We first introduce a naive approach to estimating the ATE; we regress the outcome over the treatment and adjust for the confounder (i.e., gender). In naive regression, it is assumed that there is an equal probability of an individual being a certain level of the confounder (i.e., gender). Naively, we say there is strong evidence ($p < 0.001$) that the risk of death within 30 days is 7.4%

higher amongst those with RHC and 95% CI (4.84-9.86) (Table 3).

```
regress $Y $A $C // Risk differences = 7.4%; 95% CI (4.84-9.86); p<0.001
```

Box 3: Non-parametric G-Formula for the ATE

First, we compute the marginal probability of the confounder, saving it in a matrix "m", and generate two new variables named *genderf* for females and *genderm* for males (i.e., the proportion of females was 44%, thus 56% are males, which shows unequal proportions). Afterwards, we compute, and save in a matrix, the expected conditional probabilities of the outcome by levels of the exposure and the confounder. We substitute the results of the matrix into the G-formula, given in (2), and compute the ATE. The risk of death within 30 days is 7.4% higher amongst those with RHC compared to those patients who did not get the RCH.

```
1  proportion $C          // Marginal probability of C
2  matrix m=e(b)
3  gen genderf = m[1,1]
4  sum genderf
5  gen genderm = m[1,2]
6  sum genderm
7  ssc install sumup
8  sumup $Y, by($A $C)    // Expected conditional probabilities of the outcome by levels of A and C
9  // from sumup command extract the conditional means by the given levels of A and C i.e. zero and one
10 matrix y00 = r(Stat1) // [6,1] matrix for E(Y|A=0,C=0)
11 matrix y01 = r(Stat2) // [6,1] matrix for E(Y|A=0,C=1)
12 matrix y10 = r(Stat3) // [6,1] matrix for E(Y|A=1,C=0)
13 matrix y11 = r(Stat4) // [6,1] matrix for E(Y|A=1,C=1)
14 // see "matrix list y00": position subscript [3,1] is the one of interest
15 // Applying the G-formula
16 gen EY1 = ((y11[3,1]-y01[3,1]))*genderm
17 gen EY0 = ((y10[3,1]-y00[3,1]))*genderf
18 qui: mean EY1 EY0
19 matrix ATE = r(table)
20 display "The ATE is: " ATE[1,1] + ATE[1,2] // Applying the G-formula
21 drop EY1 EY0
22 // Also one can try
23 gen ATE = ((y11[3,1]-y01[3,1]))*genderm + ((y10[3,1]-y00[3,1]))*genderf
24 qui sum ATE
25 drop ATE
```

Note that the results from the naive and G-formula approaches are the same (Table 3). However, the interpretation varies because the naive approach is a conditional estimation interpreted as the individual risk among those who actually had the exposure holding the levels of the confounder constant at their mean value (i.e., gender) at each combination of the levels of the exposure and outcome (i.e., instead of using the marginal probabilities of males and females as in row 15 from box 3, the mean value of gender is used). Note that it is difficult to conceptualize the mean value of a categorical variable such as gender. On the other hand, the G-formula is the contrast of the conditional mean outcome given treatment A for both male and females standardised using the marginal probability of gender (i.e., males and females) and averaged across the levels of males and females in the sample (i.e., row 15 in box 3) and therefore must be interpreted at a population level. The interpretation of the conditional treatment effect, adjusted for gender, from the naive approach becomes difficult to interpret after standardisation using the mean value of gender and requires the assumption that the effect of the treatment of interest is constant across the levels of the confounders included in the model (i.e., is the same for males and females).^{6,12} However, in observational studies the effect of risk factors or treatments might differ across individuals with different susceptibilities or characteristics. Therefore, the G-Formula has become a powerful alternative to the multivariable regression adjustment to better control for confounding and evaluate the effects of treatment and exposures in observational studies.⁴

So far we have calculated the ATE amongst those who had RHC or not, but often there are identifiability issues where the probability of the exposure or treatment is nearly zero or zero for some groups. Therefore, the average treatment effect amongst the treated (ATT) might be of interest as it is the marginal contrast of the potential outcomes amongst only those who had RHC. Therefore, this comes with a slightly weaker assumption of conditional mean exchangeability such that the probability of the potential outcome among those who died if, possibly contrary to fact, they had not received the treatment (i.e., RHC in our example) is independent of the exposure given the confounder.

Box 4: Non-parametric G-Formula for the ATT

In box 5 we show how to compute the ATT non-parametrically based on the implementation of the G-formula. The estimate of the ATT is 7.3% and is interpreted as the increased risk of death within 30 days amongst patients with RHC, had they not been treated.

```

1 // Expected conditional probability of the outcome amongst those with RHC
2 proportion $C if $A==1
3 matrix m=e(b)
4 gen genderfatet = m[1,1]
5 gen gendermatet = m[1,2]
6 // G-formula
7 gen EY1 = ((y11[3,1]-y01[3,1])*gendermatet
8 gen EY0 = ((y10[3,1]-y00[3,1])*genderfatet
9 qui: mean EY1 EY0
10 matrix ATT = r(table)
11 display "The ATT is: " ATT[1,1] + ATT[1,2] // Applying the G-formula
12 drop EY1 EY0
13 // Also one can try
14 gen ATT = ((y11[3,1]-y01[3,1])*gendermatet + ((y10[3,1]-y00[3,1])*genderfatet
15 qui sum ATT
16 drop ATT

```

3.1.1 | Statistical inference: The bootstrap

Note that classical statistical inference methods do not apply for the causal inference estimators. This is because the standard errors (SE) of the causal inference estimators do not account for the different steps we need to make when we balance the confounders between exposure groups. We use the bootstrap procedure for inference implemented in Stata with the command *bootstrap*.¹³ Usually, the bootstrap is used to approximate the variance of the estimator (i.e., G-computation or IPTW for the ATE).^{13,14} When estimating the variance using the bootstrap method, the observed data is thought of as representing the entire target population, and each draw (with replacement) from data mimics the sampling variability. Under certain assumptions, this set of draws will return estimates of the sampling distribution that are equivalent to having actually repeated the sampling from the original target population.¹³ Typically, for procedures that use parametric models, the bootstrap would give reliable estimation for the variance of the estimator. However, note, it does not account for the bias engendered by model misspecification, so it just provides sampling variability for whatever the estimator is estimating.¹³ The accuracy with which the bootstrap distribution estimates the sampling distribution depends on the number of observations in the original sample and the number of replications in the bootstrap. To implement the bootstrap procedure in Stata we need to define a program that estimates the non-parametric G-formula and then samples (with replacement) the ATE to derive the confidence intervals for the ATE and the ATT. In box 6 we provide the code to compute the SE for the ATE, and the ATT, using Stata.

Box 5: Bootstrap 95% Confidence Intervals (CI) for the ATE/ATT estimated using the Non-parametric G-Formula

```

1 // ATE
2 capture drop program ATE
3 program define ATE, rclass // As before but now define a program to estimate the ATE
4     capture drop y1
5     capture drop y0
6     capture drop ATE
7     sumup $Y, by($A $C)
8     matrix y00 = r(Stat1)
9     matrix y01 = r(Stat2)
10    matrix y10 = r(Stat3)
11    matrix y11 = r(Stat4)
12    gen ATE = ((y11[3,1]-y01[3,1])*genderm + ((y10[3,1]-y00[3,1])*genderf
13    qui sum ATE
14    return scalar ate = `r(mean)'
15 end
16 qui bootstrap r(ate), reps(1000): ATE // Bootstrap 1000 estimates of the ATE
17 estat boot, all
18 drop ATE
19
20 // ATT
21 capture program drop ATT
22 program define ATT, rclass

```



```

23     capture drop y1
24     capture drop y0
25     capture drop ATT
26     sumup $Y, by($A $C)
27     matrix y00 = r(Stat1)
28     matrix y01 = r(Stat2)
29     matrix y10 = r(Stat3)
30     matrix y11 = r(Stat4)
31     gen ATT = ((y11[3,1]-y01[3,1])*gendermatet + ((y10[3,1]-y00[3,1])*genderfatet
32     qui sum ATT
33     return scalar att = `r(mean)'
34 end
35 qui bootstrap r(att), reps(1000): ATT
36 estat boot, all
37 drop ATT

```

After bootstrapping, the estimate of the ATE and the ATT was 7.4% and 7.3%, respectively. We obtain three sets of confidence intervals for each of ATE and ATT. For the ATE, the first (N) is an approximation based on the normal distribution (95% CI 4.79-9.95), the second (P) is based on the percentile of the bootstrap distribution (4.74-9.84), and the third (BC) is based on the bias-corrected (4.91 - 10.01) (Table 3). We now do the same for the ATT. The 95% CIs for the ATT are: normal-approximation (4.74-9.91), percentile (4.71-10.00), and bias-corrected (4.70-10.00). Note that the percentile interval is a simple "first-order" interval that is formed from quantiles of the bootstrap distribution. However, it is based only on bootstrap samples and does not adjust for skewness in the bootstrap distribution. Thus, we will report the BC 95% CI.¹⁵

Remaining with the non-parametric implementation of the G-formula, we turn our attention to computing the ATE and ATT using a fully saturated regression model (still with only one confounder) using the full information of the sample (including the interactions). In Stata there are two different approaches that we illustrate in boxes 7 and 8 using the commands *predictnl* and *margins*.

Box 6: Non-parametric G-Formula using a fully saturated regression model in Stata (A)

To estimate the ATE using a fully saturated regression model we need to include all the possible interactions between the treatment and the different levels of the confounders (if categorical otherwise with a continuous confounder). We use the # facility of Stata to include the interaction between A and C. The *coeflegend* option asks Stata to provide the list of the labels of the variables in the analysis. The labels are then used for the *predictnl* command, which allows the computation of the non-parametric predictions based on the combination of the conditional probabilities from the regression coefficients. Finally, we average over the predictions to get the non-parametric estimate for the ATE but biased 95% CI.

```

1 regress $Y ibn.$A ibn.$A#c.($C), noconstant vce(robust) coeflegend
2 predictnl ATE = (_b[1bn.rhc] + _b[1bn.rhc#c.gender]*gender) - (_b[0bn.rhc] + _b[0bn.rhc#c.gender]*gender)
3 mean ATE

```

Box 7: Non-parametric G-Formula using a fully saturated regression model in Stata (B)

A simpler option would be to use the *margins* command to estimate the marginal probabilities using the option *vce(unconditional)*. Then the difference in marginal probabilities between the treated versus non-treated is implemented using the contrast option from the *margins* command. Note, that here we obtain the same estimate of the ATE as 7.4% (95%CI 4.83-9.91) but the correct 95%CI has been calculated using the Delta method (Table 3). The Delta method is a statistical approach to derive standard error of an asymptotically normally distributed estimator. It uses a first order Taylor approximation which is how we approximate the distribution of a function using a tangent line (i.e., the first derivative).¹⁶ Therefore, using the Delta method here we assume that the ATE estimate from the G-computation is normally distributed and asymptotically linear.¹⁷

```

1 regress $Y ibn.$A ibn.$A#c.($C), noconstant vce(robust) // Fully saturated model specification
2 margins $A , vce(unconditional) // Marginal probability for A
3 margins r.$A , contrast(nowald) // Difference in marginal probability between treatment groups

```

3.2 | Parametric G-formula

In contrast to the non-parametric methods (i.e., probability distribution free), to compute the ATE parametrically we assume there is a particular probability distribution that fits the distribution of our data. To compute the ATE, we first regress (using a simple linear regression model) the outcome over the confounder separately for each treatment group, and then contrast the difference in the expected probabilities between the two treatment groups. The algebraic form of the ATE under the G-computation is given by

$$\text{ATE} = \frac{1}{n} \sum_{i=1}^n (\hat{E}(Y_i | A_i = 1, \mathbf{W}_i) - \hat{E}(Y_i | A_i = 0, \mathbf{W}_i)) \quad (4)$$

In box 8 we provide the code to compute, by hand, the ATE based on the parametric G-formula for one confounder using parametric regression adjustment.

Box 8: Parametric regression adjustment implementation of the G-Formula

```

1 regress $Y $C if $A==1 // Expected probability amongst those with RHC
2 predict double y1hat
3 regress $Y $C if $A==0 // Expected probability amongst those without RHC
4 predict double y0hat
5 mean y1hat y0hat // Difference between the expected conditional probabilities
6 lincom _b[y1hat] - _b[y0hat] // ATE and biased confidence interval

```

The risk of mortality amongst those with RHC is 7.4% higher compared to those without RHC. Note that if we want to naively compute a 95% for the linear contrast between the marginal potential outcomes we will get a biased CI that does not account for the two-step procedure to get the marginal probabilities.

Box 9: Parametric regression adjustment using Stata's *teffects*

In box 10, we just confirm the result we got for the ATE estimate using the STATA's *teffects* command including the 'ra' option to perform the regression adjustment. Note the difference between the naive and the *teffects* 95% CIs. The *teffects* uses the Delta method to correct for the two step estimation procedure and provide appropriate statistical inference.

```

1 teffects ra ($Y $C) ($A) //Parametric G-Formula implementation in Stata

```

With the *teffects* command in Stata the ATE is 7.4% which is the same as we obtained by hand. However, note that the 95%CI for the ATE using the command from Stata is more conservative than using the naive approach without accounting for the two-step estimation procedure (i.e., 95%CI: 4.83 - 9.91 and 95%CI: 7.35 - 7.39, respectively for the *teffects* and the naive approaches) (Table 3).

Box 10: Bootstrap for the parametric regression adjustment

Again, if we want to compute the 95% CI by hand using Stata we must use the bootstrap procedure (refer to Box 6 for an explanation).

```

1 capture program drop ATE
2 program define ATE, rclass
3     capture drop y1
4     capture drop y0
5     reg $Y $C if $A==1
6     predict double y1, xb
7     quiet sum y1
8     reg $Y $C if $A==0
9     predict double y0, xb
10    quiet sum y0
11    mean y1 y0
12    lincom _b[y1]-_b[y0]
13    return scalar ace =`r(estimate)´
14 end
15 qui bootstrap r(ace), reps(1000): ATE
16 estat boot, all

```

```
17 | drop ATE
```

After bootstrapping, the estimate of the ATE is 7.4% and the bias-corrected 95%CI: (4.73-10.23) (Table 3).

Box 11: Parametric multivariable regression adjustment implementation of the G-Formula

As is often the case, there is almost always more than one confounder. The parametric computation of the G-formula can easily be extended to include more than one confounder: remember that **W** includes all of the confounders in our study.

```
1 regress $Y $W if $A==1 // Saturated regression model with all confounders for those with RHC
2 predict double y1hat
3 regress $Y $W if $A==0 // Saturated regression model with all confounders for those without RHC
4 predict double y0hat
5 mean y1hat y0hat // ATE is the difference in expectations
6 lincom _b[y1hat] - _b[y0hat] // The estimation from this command gives a biased confidence interval for
 the ATE
```

The ATE of those with RHC (i.e. risk of mortality amongst those with RHC) is 8.3%, (95% CI: 8.22-8.43) higher compared to those without RHC. Note that the 95%CI provided by the *lincom* Stata command is biased as it is not accounting for the two-step estimation procedure to derive the ATE.

Box 12: Parametric multivariable regression adjustment using Stata's *teffects* command

In box 13 we use the Stata's *teffects* command to confirm our results: note that we now include **W** instead of the single confounder **C**.

```
1 teffects ra ($Y $W) ($A)
```

We obtain the same results with Stata's command as with our calculations by hand (ATE 8.3%; 95% CI: 5.80-10.85). However, note again the difference between the 95% CI estimated naively and using the *teffects* command (Table 3).

Box 13: Parametric multivariate regression adjustment using Stata's *margins* command

In box 14 we show another way of obtaining the ATE under the parametric G-formula approach using the Stata *margins* command after fitting a fully saturated regression model. First, we regress the dependent variable (Y) over the treatment (A) and the interaction of A with all of the other independent variables (W). We do it using the Stata functionalities *ibn* to force the estimation of all the levels of a categorical variable and *#* to indicate the interaction between all the levels of A with all of the other variables included in the model (W). Then, the *margins* command calculates the predicted value of the expectation of the outcome given the treatment and the confounders i.e. $E(Y|A, \mathbf{W})$ and then reports the mean value of those predictions for each level of the treatment (A). The *vce* option specified whether the standard errors should account for heteroscedasticity: *robust* accounts for this but *unconditional* does not. Finally, to compute the ATE and provide corrected 95% CI based on the Delta method, we use the *contrast* option to contrast the outcomes between those who received RHC and those who did not. Note the results are the same as before using the *teffects* command (i.e., ATE 8.3%; 95% CI: 5.80-10.85) (Table 3).

```
1 regress $Y ibn.$A ibn.$A#c.($W), noconstant vce(robust) // Fully saturated model
2 margins $A, vce(unconditional) // E(Y|A=1,W), E(Y|A=0,W) and Delta method for the standard errors (i.e., vce
 unconditional) and 95%CI
3 margins r.$A, contrast(nowald) // ATE and Delta method for the standard error and 95%CI
```

Box 14: Bootstrap for the multivariable parametric regression adjustment

Finally, in box 15 we show how to compute the bootstrap 95% CIs for the G-computation implementation of the G-formula using regression adjustment in Stata.

```
1 capture program drop ATE
2 program define ATE, rclass
3 capture drop y1
4 capture drop y0
5 reg $Y $W if $A==1
```

```

6   predict double y1, xb
7   quiet sum y1
8   reg $Y $W if $A==0
9   predict double y0, xb
10  quiet sum y0
11  mean y1 y0
12  lincom _b[y1]-_b[y0]
13  return scalar ace =`r(estimate)'
14  end
15  qui bootstrap r(ace), reps(1000): ATE dots
16  estat boot, all
17  drop ATE

```

After bootstrapping, the estimate of the ATE is 8.3%. The bootstrapped bias-corrected 95%CI are (5.77-10.83) (Table 3).

Box 15: Computing the parametric marginal risk ratio after regression adjustment

Before moving onto other methods, we show that *teffects* can be used for different estimands such as the marginal Risk Ratio (RR), which is the ratio between the $E(E(Y|A=1, \mathbf{W}))$ and $E(E(Y|A=0, \mathbf{W}))$. To do this in Stata we use the *aequations* option to display the regression model coefficients used to predict the outcomes as well as the coefficients used to predict the exposure in order to manipulate them and compute the marginal RR estimate as follows:

$$100 \times \left(\frac{E_w(\hat{E}(Y|A = 1, \mathbf{W}) - \hat{E}(Y|A = 0, \mathbf{W}))}{E_w(\hat{E}(Y|A = 0, \mathbf{W}))} - 1 \right) = \frac{E_w(\hat{E}(Y|A = 1, \mathbf{W}))}{E_w(\hat{E}(Y|A = 0, \mathbf{W}))}$$

Note that again we use the Delta method to derive the 95%CI for the marginal RR using the *nlcom* command in Stata.

```

1 teffects ra ($Y $W) ($A), aequations // aequations display the estimation auxiliary equations
2 teffects ra ($Y $W) ($A), coeflegend // coeflegend displays the name of the coefficients
3 nlcom 100*_b[ATE:r1vs0.$A]/_b[POmean:0.$A] // nlcom uses the Delta method to compute the standard error
4 // 27.4% increase in risk
5 teffects ra ($Y $W) ($A), pom coeflegend // pom displays E(Y|A=1,W) and E(Y|A=0,W), respectively
6 nlcom _b[POmeans:1.rhc]/_b[POmeans:0bn.rhc]
7 // 27.4% increase in relative risk

```

The relative risk of mortality is 27% (95% CI 18.25 - 36.61) higher amongst those with RHC compared to those without RHC. Note that using the Stata command *nlcom* we estimate the SE for the RR based on the Delta method accounting for the two-step estimating procedure.

4 | INVERSE PROBABILITY OF TREATMENT WEIGHTING

4.1 | Inverse probability weighting based on the propensity score

In observational settings, some individuals will be more likely than others to be exposed (A) due to their characteristics. Suppose some individuals who were exposed were unlikely to be exposed based on a specific set of features encapsulated in a particular vector of confounders (W). To balance the exposure risk we increase the weight (upweight) these individuals on the outcome variable (Y) by the inverse of their probability of treatment or exposure (A) (i.e., propensity score), so that they represent themselves but also the other individuals with similar characteristics, who were unexposed. Concurrently, we upweight those patients that were unlikely to be unexposed. The resulting dataset is unchanged apart from now A and W are independent. Therefore, a comparison of Y(1) to Y(0) gives a marginal causal effect under the three identification assumptions and further assuming the propensity score model is correctly specified. Therefore, the inverse probability of treatment weighting (IPTW) estimators also compute the G-formula parametrically but based on the inverse probability of treatment or exposure.^{4,18} Originally, they were motivated from the classical Horvitz and Thompson survey estimator used to upweight on the outcome variable by the inverse probability that it is observed to account for the missingness process.¹⁹ Therefore, regression adjustment and IPTW estimators are both computational implementations of the G-formula, i.e. generalisation of the standardisation procedure. Thus, these

methods are also known as G-computation methods as they are based on the computational implementation of the G-Formula.

Analytically departing from the identification assumptions for the ATE for the regression adjustment G-computation estimand ($ATE = E_{\mathbf{W}}(E(Y|A=1, \mathbf{W}) - E_{\mathbf{W}}(Y|A=0, \mathbf{W}))$), we can rewrite the same estimand as a function of the distribution of A given W i.e., $P(A|W)$ a.k.a propensity score or treatment mechanism.

Therefore, the IPTW estimand of the ATE is given by

$$ATE = \hat{E} \left(\frac{\mathbb{I}(A = a)}{g(A = 1 | \mathbf{W})} Y \right), \quad (5)$$

where $g(A = a | \mathbf{W}) = P(A = a | \mathbf{W})$ and \mathbb{I} refers to the binary indicator $A = a$ set equal to 0 or 1.

There is a modified version of the IPTW estimator in (5) that is more efficient and is used in practice, consisting on a stabilised version of the weights (i.e., ranging between zero and one). The stabilised version of the IPTW estimator is given by

$$ATE = \frac{\hat{E} \left(\frac{\mathbb{I}(A=a)}{g(A=1|\mathbf{W})} Y \right)}{\hat{E} \left(\frac{\mathbb{I}(A=a)}{g(A=1|\mathbf{W})} \right)}, \quad (6)$$

where $g(A = a | \mathbf{W}) = P(A = a | \mathbf{W})$ and \mathbb{I} refers to the binary indicator $A = a$ set equal to 0 or 1.

By repeated use of the law of total expectation, the IPTW and the G-computation regression adjustment estimands for a binary treatment are equivalent as given by

$$\begin{aligned} E \left(\frac{\mathbb{I}(A = 1)}{P(A = 1 | \mathbf{W})} Y \right) &= E \left(E \left(\frac{\mathbb{I}(A = 1) Y}{P(A = 1 | \mathbf{W})} \middle| A, \mathbf{W} \right) \right) \\ &= E \left(\frac{\mathbb{I}(A = 1)}{P(A = 1 | \mathbf{W})} E(Y | A, \mathbf{W}) \right) \\ &= E \left(\frac{\mathbb{I}(A = 1)}{P(A = 1 | \mathbf{W})} E(Y | A = 1, \mathbf{W}) \right) \\ &= E \left(E \left(\frac{\mathbb{I}(A = 1)}{P(A = 1 | \mathbf{W})} E(Y | A = 1, \mathbf{W}) \middle| \mathbf{W} \right) \right) \\ &= E \left(\frac{E(Y | A, \mathbf{W})}{P(A = 1 | \mathbf{W})} P(A = 1 | \mathbf{W}) \right) = E[E(Y | A = 1, \mathbf{W})]. \end{aligned}$$

In box 16 we show how to compute the IPTW by hand in two steps:

- First, the propensity score model is fitted (i.e., a logistic regression model for a binary treatment; however, data-adaptive and machine learning techniques can be used too)
- Then the sampling weights are generated based on the inverse probability of treatment. Note that the weights are just the implementation of the classical Horvitz-Thompson survey estimator,¹⁹ (see rows 3 and 4) also known as unstabilised weights.

Note that when there are near violations of the positivity assumption the unstabilised weights can have large values increasing the variance and hence the uncertainty of the ATE estimation. Therefore, it is advisable to explore the distribution of the weights in order to evaluate the extent to which they balance the distribution of confounders across the levels of the treatment (i.e., equally distributed). It is common to provide a table with the standardised treatment differences between unweighted and weighted confounders. The visualization of the overlap of the propensity scores by level of the treatment is commonly used to identify and visualise positivity or near positivity violations.

Box 16: Computation of the IPTW estimator for the ATE

```

1  logit $A $W, vce(robust) nolog           // Propensity score model for the exposure
2  predict double ps                       // Propensity score prediction
3  generate double ipw1 = ($A==1)/ps       // Sampling weights for the treated group
4  generate double ipw0 = ($A==0)/(1-ps)   // Sampling weights for the non-treated group
5  regress $Y [pw=ipw1]                   // Weighted outcome probability among treated
6  scalar Y1 = _b[_cons]
7  regress $Y [pw=ipw0]                   // Weighted outcome probability among non treated
8  scalar Y0 = _b[_cons]
9  display "ATE =" Y1 - Y0

```

You will notice that the risk difference between those with RHC and those without is 8.3% (confidence intervals are estimated below in box 18).

Box 17: Bootstrap computation for the IPTW estimator

As before, we can obtain confidence intervals using the bootstrap procedure implemented in box 18.

```

1  program drop ATE
2  program define ATE, rclass
3  capture drop y1
4  capture drop y0
5  regress $Y [pw=ipw1]
6  matrix y1 = e(b)
7  gen double y1 = y1[1,1]
8  regress $Y [pw=ipw0]
9  matrix y0 = e(b)
10 gen double y0 = y0[1,1]
11 mean y1 y0
12 lincom _b[y1]-_b[y0]
13 return scalar ace = `r(estimate)'
14 end
15 qui bootstrap r(ace), reps(1000): ATE
16 estat boot, all
17 drop ATE

```

After bootstrapping, the estimate of the ATE is 8.3%. The bootstrapped bias-corrected confidence are: (5.71-10.90).

Box 18: Computation of the IPTW estimator for the ATE using Stata's *teffects* command

We now confirm this result in box 19 using Stata's *teffects* command. Note that the Horvitz-Thompson estimator is implemented using the *ipw* option. We obtain the same point estimate for the ATE and slightly different, but consistent, 95%CI based on the robust SE derived from the functional Delta method (i.e., ATE 8.3%; 95% CI 5.78 - 10.83) (Table 2).

```
teffects ipw ($Y) ($A $W, logit), nolog vsquish
```

Box 19: Assessing IPTW balance

In box 19 we show how to explore "balance" i.e., that the distribution of the confounders are balanced between those with RHC and those without, after re-weighting the contributions of patients using IPTW. When applying weights we must be careful as we are assuming that the exposure has been balanced across the levels of the confounders. In Stata we use the *tebalance* option after using the *teffects* command but the balance can be assessed by hand as well.

```

1  qui teffects ipw ($Y) ($A $W)
2  tebalance summarize // Stata's tebalance
3
4  // tebalance by hand (gender)
5  egen genderst = std(gender) // Standardization
6  logistic $A $W // Propensity score
7  predict double ps
8  gen ipw = .
9  replace ipw=($A==1)/ps if $A==1
10 replace ipw=($A==0)/(1-ps) if $A==0
11 regress genderst $A // Raw difference
12 regress genderst $A [pw=ipw] // Standardized difference

```

After weighting, the two exposure groups appear to be well balanced. Prior to weighting, there was some imbalance between genders, education level and those who had a carcinoma (absolute values of the standardised differences close to or beyond 10%). A variance ratio equal to 1 before and after weighting informs us that the distribution of the exposure across the levels of the confounder is the same (i.e., perfectly balanced). Note that the variance ratio for the continuous variable age is 0.79, which is a bit further from 1 than the variance ratio for the original (unweighted) sample (Table (2)).

TABLE 2 Distribution of the confounders before and after applying weights

Confounder	Standardised differences		Variance ratio	
	Raw	Weighted	Raw	Weighted
Gender	0.0931272	0.0004124	0.9771947	0.9999057
Age	-0.0613524	-0.0038196	0.8174922	0.7899075
Education	0.0913642	-0.0025822	1.0147230	1.0250380
Race	-0.0022396	0.0023428	1.0295870	1.0254230
Carcinoma	-0.1051837	0.0012232	0.8386081	1.0226510

There is no definitive value at which the exposure is considered unbalanced; however, as a guideline, a variance ratio less than 0.5 must clearly indicate that the data is not balanced and the potential for the positivity violation must be explored (i.e., when $P(A = a | C = c)$ is near to zero or one). An additional strategy is to check the distribution of the weights: if very large weights are identified this confirms the violation of the positivity assumption. Again there is no clear consensus but when there are very large weights researchers often set the weights to a less extreme value. This is done by truncating the distribution of the confounder to the 5th and 95th percentiles: removing the data at the extremes of the distribution. Note that "truncating" or "trimming" the weights reduces non positivity and variance but at the expense of introducing bias.²⁰ In extreme cases, changing the estimand to the ATT could be another solution.

Box 20: Assessing IPTW overlap by hand

Therefore, independently of the balance, it is also important to check for any violations of the identifiability conditions. Starting with the positivity assumption, overlapping the weights gives a visual impression of whether the weights are balanced. In box 21 we show how to visualise the "overlap" using a kernel density estimate by the levels of the exposure. (Figure (2)). shows that there appears to be a suitable amount of overlap.

```

1  sort $A
2  by $A: summarize ps
3  kdensity ps if $A==1, generate(x1pointsa d1A) nograph n(10000) // Non-parametric kernel density estimate of
   the distribution of the propensity score among treated patients
4  kdensity ps if $A==0, generate(x0pointsa d0A) nograph n(10000) // Non-parametric kernel density estimate of
   the distribution of the propensity score among non-treated patients
5  label variable d1A "density for RHC=1"
6  label variable d0A "density for RHC=0"
7  twoway (line d0A x0pointsa , yaxis(1))(line d1A x1pointsa , yaxis(2))

```

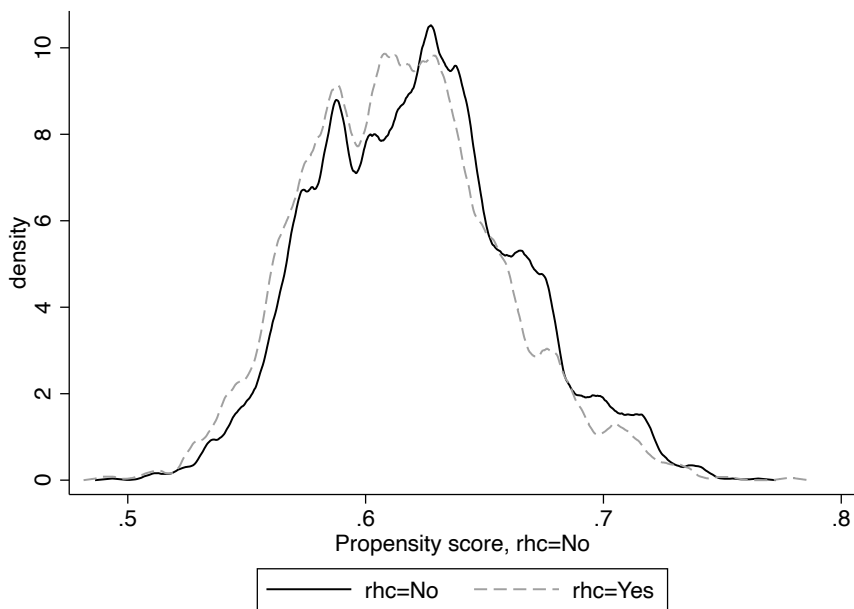
Box 21: Assessing overlap using teffects overlap

Note that the overlap plots can be obtained with Stata's command *overlap* after calling the *teffects* command. Note that we are using data where the balance and overlap are nearly perfect but in real-world observational data the balance and overlap will differ.

```

1  qui: teffects ipw ($Y) ($A $W, logit), nolog vsquish
2  teffects overlap

```

FIGURE 2 Propensity score overlap by exposure status

4.2 | Marginal structural model with stabilised weights

A marginal structural model (MSM) is a weighted regression model for the outcome (Y) as dependent variable and treatment or exposure (A) as independent variable. In box 23 we show how to compute a MSM. Note that the coefficient of the MSM for the treatment A is the ATE estimate. To fit the MSM we use as sampling weights (i.e., Horvitz and Thompson) the IPTW computed as the inverse of the probability of treatment or propensity score (see box 17). However, the variance of the ATE estimate can be inflated in a non-saturated MSM due to the presence of large weights. Larger weights are the result of potential near-violation of the positivity assumption. Thus, to counter it, one may use a stabilised version of the weights.²¹ Stabilised weights only differ from the weights described above in that they do not take a simple inverse but instead divide the baseline probability of selecting a treatment (estimated from a model with no covariates) by the probability of selecting treatment given the covariates. The mean of the stabilised weights is equal to one. Therefore, the stabilised weights produce ATE estimates that have smaller variance being more robust against near-positivity violations. Finally, note that for statistical inference we use the *vce(robust)* to estimate the correct SE for the ATE using the Delta method.²²

Box 22: Computation of the IPTW estimator for the ATE using a MSM

```

1 // baseline treatment probabilities
2 logit $A, vce(robust) nolog
3 predict double nps, pr
4 // propensity score model
5 logit $A $W, vce(robust) nolog
6 predict double dps, pr
7 // Unstabilised weight
8 gen ipw = .
9 replace ipw=($A==1)/dps if $A==1
10 replace ipw=($A==0)/(1-dps) if $A==0
11 sum ipw
12
13 // Stabilized weight
14 gen sws = .
15 replace sws = nps/dps if $A==1
16 replace sws = (1-nps)/(1-dps) if $A==0
17 sum sws
18
19 // MSM

```



```

20 reg $Y $A [pw=ipw], vce(robust) // MSM unstabilized weight
21 reg $Y $A [pw=sws], vce(robust) // MSM stabilized weight

```

4.3 | Inverse probability weighting plus regression adjustment

The IPTW-RA is an estimator using a G-computation regression adjustment that incorporates the estimated stabilised IPTW. It has been shown that the IPTW-RA helps to correct the estimator when the regression function is misspecified provided that the propensity score model for the treatment is correctly specified. When the regression function is correctly specified, the weights do not affect the consistency of the estimator even if the model from which they are derived is misspecified.²³ Note that by blending both approaches we only increase our chances to get a more consistent estimate of the ATE and this is why estimators that combine both modelling approaches are named double robust. When one uses G-computation methods only, they simply rely on extrapolation of the treatment effects when there are identifiability issues due to sparsity and near-positivity violations. Adding the IPTW to the regression adjustment allows evaluation of the balance of the treatment and of possible positivity violations, increasing the researcher's awareness of the limitations of causal inference modelling. We also encourage researchers when possible to explore the implementation of the non-parametric G-formula (just using some of the most important confounders) to really understand their data at hand and identify potential problems relative to the curse of dimensionality from finite samples (i.e., zero empty cells for a given combination of conditional probabilities from the different variables included in analysis needed to implement the G-formula).

Although IPTW with regression adjustment (IPTW-RA) is usually more efficient than IPTW, it also relies on different parametric modelling assumptions: i) a parametric G-computation regression adjustment model, and ii) a parametric logit model for the propensity score of binary exposures. The G-computation weighted model uses the weights calculated from the predictions of the propensity score logit model. An estimated propensity score that is close to 0 or 1 is problematic, since it implies that some individuals will receive a very large weight leading to imprecise and unstable estimates. Here, the positivity assumption is required (i.e., if $P(C = c) > 0$ then $P(A = a | C = c) > 0 \forall C \in c, a \in \{0, 1\}$) and stabilised weights are used (see code from Box 23). After reweighting the individuals, we generate a new pseudo-population (weighted population) from which the data generation does not follow a theoretical distribution and individuals are no longer independent. Therefore, we must use robust standard errors based on the Delta method applied to a function such as the Sandwich estimator of the SE.²⁴

Box 23: Computation of the IPTW-RA estimator for the ATE and bootstrap for statistical inference

```

1 capture program drop ATE
2 program define ATE, rclass
3     capture drop y1
4     capture drop y0
5     reg $Y $W if $A==1 [pw=sws]
6     predict double y1, xb
7     quiet sum y1
8     return scalar y1=`r(mean)'
9     reg $Y $W if $A==0 [pw=sws]
10    predict double y0, xb
11    quiet sum y0
12    return scalar y0=`r(mean)'
13    mean y1 y0
14    lincom _b[y1]-_b[y0]
15    return scalar ate =`r(estimate)'
16 end
17 qui bootstrap r(ate), reps(1000): ATE
18 estat boot, all
19 drop ATE

```

After bootstrapping, the estimate of the IPTW-RA ATE is 8.3%, bias-corrected 95%CI (5.76-10.87). The results are the same as those obtained using the Stata's *teffects* command with the option *ipwra* presented in box 25 (i.e., ATE: 8.3% and 95% CI: 5.79-10.84) (Table 3).

Box 24: Computation of the IPTW-RA estimator for the ATE using Stata's *teffects*

Note that using *ipwra* we specify two models (i.e., the model for the outcome and the model for the exposure). As with the

parametric G-formula approach, we can also estimate the RR: 27.4% higher mortality among treated than non-treated (95% CI: 18.20-36.56), which is similar to the estimate obtained by the parametric G-formula.

```

1 teffects ipwra ($Y $W) ($A $W), pom coeflegend
2 nlcom _b[P0means:1.rhc]/_b[P0means:0bn.rhc] // marginal RR and 95%CI (Delta method)
3 // 27.4% increased risk

```

5 | AUGMENTED INVERSE PROBABILITY WEIGHTING

The AIPTW estimator is an IPTW estimator that includes an augmentation term, which corrects the estimator when the treatment model is misspecified. When the treatment model is correctly specified, the augmentation term vanishes as the sample size becomes large. Thus, the AIPTW estimator is more efficient than the IPTW. However, like the IPTW, the AIPTW does not perform well when the predicted treatment probabilities are too close to zero or one (i.e., near positivity violations). The augmentation term has expectation zero and includes the expectation of the propensity score and the regression adjustment outcome. Thus, the AIPTW integrates two parametric models i.e., a model for the outcome and a model for the treatment.^{25,26} The AIPTW estimator produces a consistent doubly-robust estimate of the parameter if either of the two models has been correctly specified.^{23,26}

Focusing on the IPTW estimator for the ATE in (3), let $\hat{\mu}$ be the expectation of the ATE using IPTW, more formally this is

$$\hat{\mu}_a = \hat{E} \left(\frac{\mathbb{1}(A = a)}{g(A | \mathbf{W})} Y \right),$$

we can rewrite the equation in the form of an estimating equation (see glossary) as

$$\sum_i \left(\frac{A_i Y_i}{g(A_i | \mathbf{W}_i)} - \mu_a \right) = 0,$$

As long as the estimating function has mean zero then $\hat{\mu}$ is a consistent estimator of μ , where $\hat{\mu}_a = \hat{E}(Y_i | A_i = a, \mathbf{W}_i)$. If we augment the estimating function using a mean-zero term, $\frac{A_i - g(A_i | \mathbf{W}_i)}{g(A_i | \mathbf{W}_i)}$, including the propensity score expectation ($g(A_i | \mathbf{W}_i)$), we have integrated both the estimation of the treatment mechanism and the mean outcome ($E(Y_i | A_i, \mathbf{W}_i)$), then

$$\sum_i \left(\frac{A_i Y_i}{g(A_i | \mathbf{W}_i)} - \left(\frac{A_i - g(A_i | \mathbf{W}_i)}{g(A_i | \mathbf{W}_i)} \right) E(Y_i | A_i = a, \mathbf{W}_i) \right) = 0,$$

rearranging the equation we can see that the AIPTW estimator is a combination of inverse weighting and outcome regression defined for a binary treatment as

$$\underbrace{\frac{1}{n} \sum_{i=1}^n ((\hat{Y}_i | A_i = 1, \mathbf{W}_i) - (\hat{Y}_i | A_i = 0, \mathbf{W}_i))}_{\text{G-computation-Regression-Adjustment}} + \underbrace{\frac{1}{n} \sum_{i=1}^n \left(\frac{A_i [Y_i - E(Y_i | A_i = 1, \mathbf{W}_i)]}{g(A_i = 1 | \mathbf{W}_i)} - \frac{(1 - A_i) [Y_i - E(Y_i | A_i = 1, \mathbf{W}_i)]}{g(A_i = 0 | \mathbf{W}_i)} \right)}_{\text{Zero-expectation}}, \quad (7)$$

where the ATE from the AIPTW estimator is defined as

$$\text{AIPTW-ATE} = \mu_1 - \mu_0$$

$$\mu_1 = \frac{1}{n} \sum_{i=1}^n \left(\hat{Y}_i | A_i = 1, \mathbf{W}_i \right) + \frac{A_i [Y_i - E(Y_i | A_i = 1, \mathbf{W}_i)]}{g(A_i = 1 | \mathbf{W}_i)}, \quad (8)$$

$$\mu_0 = \frac{1}{n} \sum_{i=1}^n \left(\hat{Y}_i | A_i = 0, \mathbf{W}_i \right) + \frac{(1 - A_i) [Y_i - E(Y_i | A_i = 1, \mathbf{W}_i)]}{g(A_i = 0 | \mathbf{W}_i)}.$$

Note that the second term in equation (7) can be interpreted as playing the role of two nuisance parameters of the AIPTW estimating function. The nuisance parameters are represented as a weighted sum of the residuals for the conditional mean of the outcome.¹⁷ The AIPTW equation (7) shows that the AIPTW estimator equals the G-formula estimator if the outcome model is correctly specified irrespective of the exposure model. Likewise, it will be equal to the IPTW estimator if the exposure model is correctly specified, irrespective of the outcome model.

In box 25 we show how to compute the AIPTW estimator for the ATE using Stata.

Box 25: Computation of the AIPTW estimator for the ATE and bootstrap for statistical inference

```

1 // Step (i) prediction model for the outcome using G-computation regression adjustment
2 qui glm $Y $A $W, fam(bin)
3 predict double QAW, mu
4 qui glm $Y $W if $A==1, fam(bin)
5 predict double Q1W, mu
6 qui glm $Y $W if $A==0, fam(bin)
7 predict double Q0W, mu
8
9 // Step (ii): prediction model for the treatment
10 qui logit $A $W
11 predict double dps, pr
12 qui logit $A
13 predict double nps, pr
14 gen sws = .
15 replace sws = nps/dps if $A==1
16 replace sws = (1-nps)/(1-dps) if $A==0
17
18 // Step (iii): Estimation equation based on analytical formula 5
19 gen double y1 = (sws*($Y-QAW) + (Q1W))
20 quiet sum y1
21 return scalar y1=`r(mean)'
22 gen double y0 = (sws*($Y-QAW) + (Q0W))
23 quiet sum y0
24 return scalar y0=`r(mean)'
25 mean y1 y0
26 lincom _b[y1] - _b[y0]
27
28 // step (iv) Bootstrap confidence intervals
29 capture program drop ATE
30 program define ATE, rclass
31 capture drop y1
32 capture drop y0
33 capture drop Q*
34 qui glm $Y $A $W, fam(bin)
35 predict double QAW, mu
36 qui glm $Y $W if $A==1, fam(bin)
37 predict double Q1W, mu
38 qui glm $Y $W if $A==0, fam(bin)
39 predict double Q0W, mu
40 gen double y1 = (sws*($Y-QAW) + (Q1W))
41 quiet sum y1
42 return scalar y1=`r(mean)'
43 gen double y0 = (sws*($Y-QAW) + (Q0W))
44 quiet sum y0
45 return scalar y0=`r(mean)'
46 mean y1 y0
47 lincom _b[y1] - _b[y0]
48 return scalar ate =`r(estimate)'

```

```

49 end
50 qui bootstrap r(ate), reps(1000): ATE
51 estat boot, all
52 drop ATE

```

After bootstrapping the ATE is 8.3% and the bias-corrected 95%CI confidence intervals are: (5.91-10.86) (Table 3). In box 27, we show the same results using Stata's *teffects* command with the *aipw* option. Note that we have to specify the model for the treatment and the model for the outcome. Likewise, we can compute the marginal RR as described above.

Box 26: Computation of the AIPTW estimator for the ATE and marginal risk ratio using Stata's *teffects*

```

1 teffects aipw ($Y $W) ($A $W, logit)
2 // marginal Relative Risk
3 nlcom 100*_b[r1vs0.$A]/_b[P0mean:0.$A]
4 //another way to compute it
5 teffects aipw ($Y $W) ($A $W, logit), pom coeflegend
6 nlcom _b[P0means:1.rhc]/_b[P0means:0bn.rhc]

```

The ATE is 8.3% (95% CI: 5.79 - 10.84) (Table 3) and the marginal RR is 27.4% (95%CI: 18.2-36.6).

6 | DATA-ADAPTIVE ESTIMATION: ENSEMBLE LEARNING TARGETED MAXIMUM LIKELIHOOD ESTIMATION

Targeted Maximum Likelihood Estimation (TMLE) is a semi-parametric doubly-robust method that reduces the bias of an initial estimate by targeting an estimate closer to the true model specification by allowing for flexible estimation using non-parametric data-adaptive machine-learning methods.⁵ It therefore requires weaker assumptions than its competitors. There are several tutorials published elsewhere but we recommend our tutorial as it provides a step-by-step guided implementation of TMLE illustrated in a realistic scenario based on cancer epidemiology where assumptions about correct model specification and positivity (i.e., when a study participant had zero probability of receiving the treatment) are nearly violated.²⁷ The advantages of TMLE have repeatedly been demonstrated in both simulation studies and applied analyses.⁵ Evidence shows that TMLE consistently provides the least biased ATE estimate compared with other double-robust estimators such as the IPTW-RA and AIPTW.⁵

In box 27, we provide the computational implementation of TMLE by hand (without data-adaptive estimation) to guide and interpret the different steps involved in the TMLE. A description of the theory behind these steps can be found elsewhere.²⁷ In Step 1, we estimate the expected outcome (given treatment and confounders): this is called the plug in initial estimate of the estimator obtained via G-computation, namely Q_0 . In Step 2, we define the expected exposure given the confounders as we did previously for the estimation of the propensity score in box 17, namely g_0 . Steps 1 and 2 are similar to the double-robust methods of AIPTW; however, we now come to the advantage of TMLE. In Step 3, we regress the predicted treatment values and predicted outcome introduced in the model as an offset on the observed outcome. The parameter estimates (epsilon) for that regression are used to correct the initial estimations of Q_0 . In other words, we reduce the residual bias and optimized the bias-variance trade-off for the estimate of the ATE so that we can obtain valid statistical inference. Note that the TMLE framework adds the possibility to estimate the Q_0 and g_0 models using data adaptive machine learning algorithms and selecting the best model or an ensemble of the models.⁵ It has been shown that using machine learning algorithms reduces misspecification estimation bias.²⁸

Box 27: Computational implementation of TMLE by hand

```

1 * Step 1: prediction model for the outcome Q0 (g-computation)
2 glm $Y $A $W, fam(binomial)
3 predict double QAW_0, mu
4 gen aa=$A
5 replace $A = 0
6 predict double QOW_0, mu
7 replace $A= 1
8 predict double Q1W_0, mu
9 replace $A = aa
10 drop aa

```

```

11 // Q to logit scale
12 gen logQAW = log(QAW / (1 - QAW))
13 gen logQ1W = log(Q1W / (1 - Q1W))
14 gen logQ0W = log(Q0W / (1 - Q0W))
15
16
17 * Step 2: prediction model for the treatment g0 (IPTW)
18 glm $A $W, fam(binomial)
19 predict gw, mu
20 gen double H1W = $A / gw
21 gen double H0W = (1 - $A) / (1 - gw)
22
23 * Step 3: Computing the clever covariate H(A,W) and estimating the parameter (epsilon) (MLE)
24 glm $Y H1W H0W, fam(binomial) offset(logQAW) noconstant
25 mat a = e(b)
26 gen eps1 = a[1,1]
27 gen eps2 = a[1,2]
28
29 * Step 4: update from Q0W and Q1W to Q0W_1 and Q1W_1
30 gen double Q1W_1 = exp(eps1 / gw + logQ1W) / (1 + exp(eps1 / gw + logQ1W))
31 gen double Q0W_1 = exp(eps2 / (1 - gw) + logQ0W) / (1 + exp(eps2 / (1 - gw) + logQ0W))
32
33 * Step 5: Targeted estimate of the ATE
34 gen ATE = (Q1W_1 - Q0W_1)
35 summ ATE
36 global ATE = r(mean)
37
38 * Step 6: Statistical inference (functional Delta method)
39 qui sum(Q1W_1)
40 gen EY1tmle = r(mean)
41 qui sum(Q0W_1)
42 gen EY0tmle = r(mean)
43
44 gen d1 = (($A * ($Y - Q1W_1)/gw)) + Q1W_1 - EY1tmle
45 gen d0 = ((1 - $A) * ($Y - Q0W_1)/(1 - gw)) + Q0W_1 - EY0tmle
46
47 gen IC = d1 - d0
48 qui sum IC
49 gen varIC = r(Var) / r(N)
50
51 global LCI = $ATE - 1.96*sqrt(varIC)
52 global UCI = $ATE + 1.96*sqrt(varIC)
53 display "ATE:" %05.4f $ATE _col(15) "95%CI: " %05.4f $LCI ", " %05.4f $UCI

```

Note, that in box 27, the residual bias is reduced by solving an equation that calculates how much to update, or fluctuate, our initial outcome estimates i.e.,

$$\text{logit}(E[Y|A, \mathbf{W}]) = \text{logit}(\hat{E}[Y|A, \mathbf{W}]) + \epsilon H(A, \mathbf{W}),$$

where $E(Y|A, \mathbf{W})$ represents the expectation of the outcome (Y) given the treatment status (A) and the set of confounders (\mathbf{W}). To solve this equation, we fit an intercept-free logistic regression (using the clever covariate H as the only predictor of the observed outcome) and the initially predicted outcome (under the observed treatment) as an offset (step 3 rows 24-27). Fitting the logistic regression, using maximum likelihood procedures, solves an estimating equation which yields many useful statistical properties of TMLE, such as: 1) the final estimate is consistent as long as either the outcome or treatment model are estimated correctly (consistently); 2) if both of these models are estimated consistently, the final estimate achieves its smallest possible variance as the sample size approaches infinity (namely, efficiency).

Next, we added the coefficient ϵ of the clever covariate H in the previous step to the expected outcome for all observations from the model fitted in Step 1 using (step 4: rows 30-31), updating the Q_0 and g_0 models predictions.

$$\hat{E}[Y|A = 1, \mathbf{W}] = \text{expit}(\text{logit}(E[Y|A = 1, \mathbf{W}]) + \epsilon H(1, \mathbf{W})), \text{ and}$$

$$\hat{E}[Y|A = 0, \mathbf{W}] = \text{expit}(\text{logit}(E[Y|A = 0, \mathbf{W}]) + \epsilon H(0, \mathbf{W})).$$

Finally, we compute the ATE as the difference between expectations of the updated Q_0 and g_0 predictions in the previous step (i.e., $\hat{E}[Y|A = 1, \mathbf{W}] - \hat{E}[Y|A = 0, \mathbf{W}]$) and provide statistical inference using the functional Delta method.^{5,29}

It is worth noting that Steps 3 and 4, which are improvements to AIPTW and IPTW-RA estimators, are the very concepts that make TMLE more robust against near positivity violations and force the estimator to respect the boundaries of the limits of the parameter space (ie., the probabilities stay between 0 and 1). For example, to estimate the ATE using the AIPTW estimator the researcher sets the estimation equation equal to zero. However, solving the estimating equation when there are near violations of the positivity assumption can cause the estimator to fall outside the boundaries of the parameter space (i.e., 0 and 1).

The ATE estimate is 8.3% (5.82 - 10.87) (Table 3), which is consistent with all the previous estimates using different estimators. Note that statistical inference is performed using the functional Delta method that is based on the Influence Function (IF).^{30,17,5}

6.1 | Statistical inference for data-adaptive estimators: Functional Delta Method

The IF is a fundamental object of semi-parametric theory that allows us to characterise a wide range of estimators and their efficiency.^{30,17,5} The IF of a regular asymptotic and linear estimator $\hat{\psi}$ of $\psi(\theta)$, where θ , is a random variable based on independent and identically distributed samples O_i which captures the first order asymptotic behaviour of $\hat{\psi}$, such that

$$n^{\frac{1}{2}} \hat{\psi} - \psi(\theta) = n^{-\frac{1}{2}} \sum_{i=1}^n IF(O_i; \theta) + o_p(1).$$

Where $o_p(1)$ represents the remainder term from the first order approximation that converges to zero in probability when the sample size converges to infinity. Mathematically we can identify the IF as being the second term of a first degree Taylor approximation.²⁹ From the variance of the IF we derive the SE of the ATE from the TMLE estimator. Therefore, the functional Delta method based on the IF readily allows the application of the Central Limit Theorem and, therefore, to compute Wald-type confidence intervals.⁵ However, using the IF for statistical inference may require larger sample sizes to avoid finite-sample issues. Recent research and theoretical developments support the use of doubly-robust cross fit estimators (i.e., cross-validated) to improve statistical inference robustness based on the functional Delta method.³¹

In box 28, we outline how to compute the ATE using data-adaptive procedures implemented in the *eltmle* user-written Stata command.³² *eltmle* implements the TMLE framework for the ATE, the marginal RR and odds ratio for a binary, or continuous, outcome with a binary treatment. *eltmle* includes the use of data-adaptive estimation of the propensity score g_0 and regression outcome Q_0 models via ensemble learning,³³ which is implemented calling the *SuperLearner* package v.2.0-21 from R.^{34,33} The super-learner uses V-fold cross-validation (10-fold by default) to assess the performance of prediction regarding the potential outcomes and the propensity score as weighted averages of a set of machine learning algorithms. The *SuperLearner* has default algorithms implemented in the base installation of the *tmle-R* package v.1.2.0-5.³⁵ The default algorithms include the following: i) stepwise selection, ii) generalized linear modeling (GLM), iii) a GLM variant that includes second order polynomials and two-by-two interactions of the main terms included in the model. Additionally, *eltmle* has an option to include Bayes Generalized Linear Models and Generalized Additive Models as additional Super-Learner algorithms. Future implementations will offer more advanced machine learning algorithms.

Box 28: TMLE with data-adaptive estimation using the Stata's user written *eltmle*

```

1  preserve
2  eltmle $Y $A $W, tmle // install via "ssc install eltmle" or "github install migariane/eltmle"
3  restore

```

The ATE is 8.3% (5.82-10.87). We also obtain estimates of the causal risk ratio (CRR 1.28; 95% CI 1.19-1.37) and the marginal odds ratio (MOR 1.45; 95% CI 1.29-1.61). These, and previous, results are presented in Table 3.

7 | SIMULATION

The motivation of this section is to share with the reader the template needed to run a simple Monte Carlo simulated experiment to compare all the different methods provided in the tutorial. For simplicity we run one sample estimate. However, we provide the results and code in R of a Monte Carlo experiment with 1,000 samples based on the same template as the one presented here and available at <https://github.com/migariane/TutorialCausalInferenceEstimators>.

In Box 29 we show the data generation process to create random variables including the confounders, the treatment and the outcome. Afterward, we estimate the simulated value for the ATE. Then, we compute the ATE using all the different estimators described above, and finally we compare their performance based on the relative bias with respect to the value of the simulated ATE. The variables generated in this simulation reflect the context of cancer survival epidemiology. This simulation includes a binary outcome (Y), potential outcomes i.e., Y(1) and Y(0) and a binary treatment (A). The confounders W reflect the commonly analysed cancer patient characteristics: deprivation level (w1, five categories), age at diagnosis (w2), cancer stage (w3, four categories) and comorbidity (w4, four categories).

Box 29: Data generation for the Monte Carlo experiment

```

1 // Data generation
2 clear
3 set obs 1000
4 set seed 777
5
6 // Confounders
7 gen w1 = round(runiform(1, 5)) //Quintiles of Socioeconomic Deprivation
8 gen w2 = rbinomial(1, 0.45) //Binary: probability age >65 = 0.45
9 gen w3 = round(runiform(0, 1) + 0.75*(w2) + 0.8*(w1)) //Stage
10 recode w3 (5/6=1) //Stage (TNM): categorical 4 levels
11 gen w4 = round(runiform(0, 1) + 0.75*(w2) + 0.2*(w1)) //Comorbidities: categorical four levels
12
13 // Binary treatment
14 gen A = (rbinomial(1, invlogit(-1 - 0.15*(w4) + 1.5*(w2) + 0.75*(w3) + 0.25*(w1) + 0.8*(w2)*(w4))))
15 // Potential outcomes and simulated ATE (one run)
16 gen Y1 = (invlogit(-3 + 1 + 0.25*(w4) + 0.75*(w3) + 0.8*(w2)*(w4) + 0.05*(w1))) // Potential outcome 1
17 gen Y0 = (invlogit(-3 + 0 + 0.25*(w4) + 0.75*(w3) + 0.8*(w2)*(w4) + 0.05*(w1))) // Potential outcome 2
18 gen psi = Y1-Y0 // Simulated ATE
19 gen Y = A*(Y1) + (1 - A)*Y0 // Outcome
20
21 // ATE from different estimators
22 teffects ra (Y w1 w2 w3 w4) (A)
23 estimates store ra
24 teffects ipw (Y) (A w1 w2 w3 w4)
25 estimates store ipw
26 teffects ipwra (Y w1 w2 w3 w4) (A w1 w2 w3 w4)
27 estimates store ipwra
28 teffects aipw (Y w1 w2 w3 w4) (A w1 w2 w3 w4)
29 estimates store aipw
30 qui reg psi
31 estimates store psi
32 estout psi ra ipw ipwra aipw
33
34 // ATE from Ensemble learning maximum likelihood estimation (ELTMLE)
35 preserve
36 eltmle Y A w1 w2 w3 w4, tml
37 restore
38
39 // Relative bias for the ATE
40 * Regression adjustment
41 display abs(0.1787 - 0.203419)/0.1787
42 0.1383268 // 13.8% bias
43 * IPTW
44 display abs(0.1787 - 0.2776)/0.1787
45 0.55344152 // 55.3% bias
46 * IPTW-RA
47 display abs(0.1787 - .2052088)/0.1787
48 0.1483424 // 14.8% bias
49 * AIPTW

```

```

50 display abs(0.1787 - 0.2030)/0.1787
51 0.13598209 // 13.6% bias
52 * ELTMLE
53 display abs(0.1787 - 0.1784)/0.1787
54 0.00167879 // 0% bias

```

It is clear to see that, compared to the true ATE of 17.9%, all of the methods (RA [13.8%], IPTW [55.3%], IPTW-RA [14.8%], AIPTW [13.6%]) produce a biased estimate, but ELTMLE produces an estimate that is unbiased relative to the true ATE. Note that the relative bias of IPTW is very large; IPTW relies on the positivity assumption, which in this simulation is violated because there was a low number of individuals with a higher comorbidity value. Without correcting for this imbalance in the data the methods that rely on this assumption will be vulnerable to bias.

8 | CONCLUSION

Overall, all the methods introduced here include the estimation of the G-formula (non-parametrically or parametrically), which is a generalization of standardization, and the inverse probability of treatment weighting (IPTW).⁴ However, there are other estimators based on matching strategies that we did not discuss in our tutorial.¹⁸

Table 3 shows all the ATE results for all the different estimators we introduced in the tutorial. Overall, all the methods showed a consistent result for the ATE. In a well-balanced data set, all of the methods used to address confounding consistently provide the same result (3). The RHC data (demonstrated in this paper) is used to teach causal inference methods because of its extremely well balanced distribution of confounders across levels of the treatment (RHC). However, in most observational studies, data is not usually well balanced and there are potentially near violations of the positivity assumption that must always be checked.

We introduce the different estimators in regards to their chronological development; the methods were developed to answer the limitations of the previous approach. For example, parametric estimators were developed to answer the question of the curse of dimensionality of the non-parametric implementation of the G-formula. Then, issues related to extrapolation for the G-computation, and the instability of the estimations due to large weights for the IPTW estimators, encouraged the development of double-robust methods. AIPTW was a strong candidate to answer this issue by incorporating semi-parametric theory and methods to causal inference. However, researchers soon realised that it did not solve the estimation equation equal to zero due to the fact that it is not a substitution estimator or plug-in estimator (see glossary). Thus, to overcome this limitation of the AIPTW estimator, data-adaptive estimation using machine learning algorithms and ensemble learning to estimate the nuisance parameters from the regression and propensity score models, combined to a smart way to solve the estimation equation were developed later in time. Evidence shows that the double robust estimators (particularly TMLE) obtain less biased estimates of the true causal effect in comparison to naive estimators such as multivariate regression.

Evidence shows that comparing the underlying properties of each method based on Monte Carlo experiments, only TMLE provides the numerous properties of estimating the probability distribution that enable it to out-perform the others. The properties of the estimator are: loss-based, well-defined, unbiased, efficient and can be used as a substitution estimator. Maximum likelihood estimation (MLE) based methods (stratification, propensity score and parametric regression) and estimating equations (IPTW and AIPTW) do not have all of the properties of TMLE and will underperform in comparison. For more detailed comparisons between the different methods, we refer the reader to Chapter 6 of TMLE.⁵ It is important to highlight that in contrast to the AIPTW estimator, TMLE respects the global constraints of the statistical model (i.e. $P_0(0 < Y < 1) = 1$) and solves the estimation equations being equal to zero.

However, even if TMLE is less prone to errors due to misspecification than alternative methods like inverse probability weighting, there is some question regarding the validity of the robustness of inference produced by TMLE in non-parametric settings. This is an area of ongoing work. Furthermore, TMLE is currently limited to analysis conducted in R, and the only current implementation of TMLE outside of R is the user written program *elmtree* for Stata,³². *elmtree* is not completely native to Stata but rather calls the *SuperLearner* R package to calculate the predictions of the exposure and outcome models. More work is

required to implement TMLE in other statistical packages.³⁵

Causal inference is a growing field with many developments during the recent years. Modern causal inference methods allow machine learning to be used when strong assumptions for parametric models are not entirely valid or reasonable. Overall, due to the difficulty of correctly specifying parametric models in high dimensional data, we advocate for the use of doubly-robust estimators with ensemble learning. However, using these approaches may require larger sample sizes to avoid finite-sample issues. More importantly, the development of renewed statistical inference for estimators using data adaptive procedures given that there was not theoretical support for the use of the Bootstrap is now well established.^{17,36} However, recent developments support the use of doubly-robust cross fit estimators to improve statistical inference robustness based on the functional Delta method.³¹ Tutorials introducing the use and derivation of the functional Delta method and Influence Curve for applied researchers are needed. The tutorial presented here may help applied researchers to gain a better understanding of computational implementation of different causal inference estimators from an applied perspective.

TABLE 3 Results of the Average Treatment Effect (ATE) from the methods.

Method	<i>teffects</i>		<i>Bootstrap</i>	
	ATE	Naive 95% CI	ATE	BS 95% CI
Univariable				
Regression	7.4	4.84 – 9.86	n/a	n/a
NPG - 1C	n/a	n/a	7.4	4.91 – 10.01
NPG - FS	7.4	4.83 – 9.91	n/a	n/a
PG - 1C	7.4	4.83 – 9.91	7.4	4.73 – 10.23
PG - FS	8.3	5.80 – 10.85	8.3	5.77 – 10.83
Multivariable				
Regression	8.3	5.74 – 10.72	n/a	n/a
IPW - PS	8.3	5.78 – 10.83	8.3	5.71 – 10.90
IPW - RA	8.3	5.79 – 10.84	8.3	5.76 – 10.87
AIPW	8.3	5.79 – 10.84	8.3	5.91 – 10.86
TMLE	8.3	5.82 – 10.87	n/a	n/a
ELTMLE	8.3	5.82 – 10.87	n/a	n/a

NPG = Non-Parametric G-formula, **FS** = Fully saturated, **PG** = Parametric G-formula, **IPW** = Inverse Probability Weighting, **PS** = Propensity Score, **RA** = Regression Adjustment, **AIPW** = Augmented Inverse Probability Weighting, **TMLE** = Targeted Maximum Likelihood Estimation by hand, **ELTMLE** = Ensemble Learning Targeted Maximum Likelihood Estimation using Stata *elmtle* package.

FUNDING

MALF is supported by a Miguel Servet I Investigator Award (grant CP17/00206 EU-FEDER) from the National Institute of Health, Carlos III (ISCIII), Madrid, Spain. His funders had no role in the study design, data collection, data analysis, data interpretation, or writing of the report.

AUTHORS CONTRIBUTIONS

The article arises from the motivation to disseminate the principles of modern epidemiology amongst clinicians and applied researchers. MALF developed the concept, designed the first draft of the article and the computing code. All authors interpreted and reviewed the code and the data, drafted and revised the manuscript. All authors read and approved the final version of the manuscript. MALF is the guarantor of the article.

ACKNOWLEDGEMENTS

We thank Fiona Ingleby and Cristina Renzi for their comments and testing the code included in the boxes of the article.

GLOSSARY

The glossary is adapted from the book *"Targeted learning: causal inference for observational and experimental data"* and a recent publication introducing the TMLE framework.^{37,5}

- **Data-generating process (DGP)**
The true mechanism that generated the observed data, with the corresponding data-generating probability distribution which produces the observed samples that were collected.
- **Estimator Estimate**
A function of the sample of observations (that is, a function of the random variables) that generates estimates. The realized value of an estimator, or a function of the realized observations.
- **Counterfactual**
A contrary-to-fact value said to arise from hypothetically imposing an intervention on a system represented by a structural causal model. For example, the potential outcome $Y(a)$ is a counterfactual that arises from a hypothetical intervention that sets the treatment A to level a .
- **Statistical model**
A set (family) of probability distributions that could describe the data-generating process. Note that, outside simulation exercises, the true data-generating process is unknown.
- **Fully-saturated model**
The data one would have observed in the ideal (impossible) experiment, and the set of possible probability distributions of the full-data random variable. In a causal model, the full data includes the counterfactual values of the outcome (i.e., potential outcomes) under all treatment/exposure conditions.
- **Model misspecification**
A scenario in which the statistical model, which is postulated to contain the distribution describing the data-generating process, fails to actually contain the corresponding true data-generating distribution.
- **Parametric statistical model**
A family of probability distributions indexed by a finite set of model parameters. For example, a linear model traditionally assumes the outcome is a linear function of covariates plus a normally distributed error term with constant variance. Its parameters are the coefficients on the covariates and the variance of the error term.
- **Non-parametric statistical model**
A family of probability distributions that cannot be indexed by a finite set of parameters. That is, the set of parameters indexing this family of distributions is infinite-dimensional. Most often, when making minimal assumptions, the data-generating process cannot be defined by a finite set of parameters, making the set of parameters infinite-dimensional. For example, if all we know about the data-generating process is that we have access to n independent and identically distributed (i.i.d.) samples, then the statistical model for the data-generating process is a non-parametric statistical model.

- Target estimand or target parameter
A function of the true (unknown) data-generating process that one is interested in estimating, and represents the mathematical formulation of the motivating question of interest.
- Maximum likelihood estimation
The most common method for estimating parameters in a finite-dimensional model (i.e., parametric statistical model). As the name implies, such estimates are generated by finding a set of parameter estimates that maximize the likelihood function of the observed data.
- Score equation
The gradient (i.e., multi-variable generalization of the derivative) of the log-likelihood function of the data with respect to the parameter(s). This equation provides information on the degree of change resulting from very small perturbations of the parameter values.
- Regular estimator
A class of estimators that converge in distribution to some limit distribution even if one samples from a slightly perturbed data distribution. Such estimators, if also asymptotically linear, accommodate inference by way of their asymptotic convergence to a Normal distribution.
- Plug-in (substitution) estimator
An estimator that generates an estimate of the true parameter value by “plugging in” estimates of relevant parts of the data-generating distribution into the parameter mapping. This method is commonly referred to as the plug-in principle. For example, “plugging in” targeted Super Learner fit of the conditional mean under $A = 1$ and $A = 0$ generates an estimate of the average treatment effect.

References

1. Pearl Judea. Causal Inference in Statistics: An Overview. *Statistics Surveys*. 2009;3:96–146.
2. Rubin Donald B. The design versus the analysis of observational studies for causal effects: parallels with the design of randomized trials. *Stat Med*. 2007;26(1):20–36.
3. Rubin Donald B. Estimating causal effects of treatments in randomized and nonrandomized studies.. *Journal of educational Psychology*. 1974;66(5):688.
4. Robins James. A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Mathematical Modelling*. 1986;7(9):1393–1512.
5. Laan Mark, Rose Sherri. *Targeted learning: causal inference for observational and experimental data*. Springer Series in Statistics; 2011.
6. Luque-Fernandez Miguel Angel, Redondo-Sanchez Daniel, Schomaker Michael. Effect Modification and Collapsibility in Evaluations of Public Health Interventions. *American Journal of Public Health*. 2019;109(3):e12-e13.
7. Connors Alfred F., Speroff Theodore, Dawson Neal V., et al. The effectiveness of right heart catheterization in the initial care of critically ill patients. *Journal of the American Medical Association*. 1996;276(11):889–897.
8. Gutman Roe, Rubin Donald B. Estimation of causal effects of binary treatments in unconfounded studies. *Stat Med*. 2015;34(26):3381–3398.
9. Rubin Donald B. Causal inference using potential outcomes. *Journal of the American Statistical Association*. 2011;.
10. Greenland Sander, Mansournia Mohammad Ali. Limitations of individual causal models, causal graphs, and ignorability assumptions, as illustrated by random confounding and design unfaithfulness. *European Journal of Epidemiology*. 2015;30(10):1101–1110.

11. Robins James M. Association, Causation, and Marginal Structural Models. *Synthese*. 1999;121(1/2):151–179.
12. Keil Alexander P., Edwards Jessie K., Richardson David B., Naimi Ashley I., Cole Stephen R.. The parametric g-formula for time-to-event data: Intuition and a worked example. *Epidemiology*. 2014;.
13. Efron Bradley, Tibshirani Robert. *An introduction to the bootstrap*. New York: Chapman & Hall; 1993.
14. Efron Bradley, Efron B. *The jackknife, the bootstrap and other resampling plans*. SIAM; 1982.
15. Jung Kwanghee, Lee Jaehoon, Gupta Vibhuti, Cho Gyeongcheol. Comparison of Bootstrap Confidence Interval Methods for GSCA Using a Monte Carlo Simulation. *Frontiers in Psychology*. 2019;10:2215.
16. Oehlert Gary W.. A note on the delta method. *American Statistician*. 1992;46(1):27–29.
17. Kennedy Edward H. *Semiparametric theory and empirical processes in causal inference*. 2016.
18. Rosenbaum Paul R, Rubin Donald B. The central role of the propensity score in observational studies for causal effects. *Biometrika*. 1983;70(1):41–55.
19. Horvitz D. G., Thompson D. J.. A Generalization of Sampling Without Replacement From a Finite Universe. *Journal of the American Statistical Association*. 1952;47(260):663.
20. Cole Stephen R., Hernán Miguel A.. *Constructing inverse probability weights for marginal structural models*. 2008.
21. Robins James M, Hernán Miguel Ángel, Brumback Babette. Marginal Structural Models and Causal Inference in Epidemiology. *Epidemiology*. 2000;11(5).
22. Boos Dennis D, Stefanski L A. *Essential statistical inference: theory and methods*. Springer; 2013.
23. Kang Joseph D Y, Schafer Joseph L. Demystifying Double Robustness: A Comparison of Alternative Strategies for Estimating a Population Mean from Incomplete Data. *Statistical Science*. 2007;22(4):523–539.
24. Tsiatis . *Semiparametric Theory and Missing Data Springer Series in Statistics: Amazon.es: Tsiatis, Anastasios A.: Libros en idiomas extranjeros* .
25. Robins James M, Rotnitzky Andrea, Zhao Lue Ping. Estimation of Regression Coefficients When Some Regressors Are Not Always Observed. *Journal of the American Statistical Association*. 1994;89(427):846–866.
26. Bang Heejung, Robins James M. Doubly robust estimation in missing data and causal inference models. *Biometrics*. 2005;61(4):962–973.
27. Luque-Fernandez Miguel Angel, Schomaker Michael, Racht Bernard, Schnitzer Mireille E.. Targeted maximum likelihood estimation for a binary treatment: A tutorial. *Statistics in Medicine*. 2018;37(16):2530–2546.
28. Naimi Ashley I, Mishler Alan E, Kennedy Edward H. Challenges in Obtaining Valid Causal Effect Estimates with Machine Learning Algorithms. 2017;.
29. Luque-Fernandez M A. Delta Method in Epidemiology: An Applied and Reproducible Tutorial. *GitHub repository*. 2020;.
30. Vaart A W van der. *Asymptotic statistics*. Cambridge, UK: Cambridge University Press; 1998.
31. Zivich Paul N, Breskin Alexander. Machine learning for causal inference: on the use of cross-fit estimators. 2020;.
32. migariane/meltml: Ensemble Learning Targeted Maximum Likelihood Estimation for Stata users | Zenodo .
33. Van Der Laan Mark J., Polley Eric C., Hubbard Alan E.. Super learner. *Statistical Applications in Genetics and Molecular Biology*. 2007;6(1).
34. Polley Eric, Laan Mark van der. Super Learner In Prediction. *U.C. Berkeley Division of Biostatistics Working Paper Series*. 2010;.

35. Gruber Susan, Laan Mark. tml: An R Package for Targeted Maximum Likelihood Estimation. *U.C. Berkeley Division of Biostatistics Working Paper Series*. 2011;.
36. Non- and Semi-Parametric Maximum Likelihood Estimators and the Von Mises Method (Part 1) [with Discussion and Reply] on JSTOR .
37. Coyle Jeremy R., Hejazi Nima S., Malenica Ivana, et al. Targeting Learning: Robust Statistics for Reproducible Research. 2020;.

