



HAL
open science

Causal effect on a target population: a sensitivity analysis to handle missing covariates

Bénédicte Colnet, Julie Josse, Gaël Varoquaux, Erwan Scornet

► **To cite this version:**

Bénédicte Colnet, Julie Josse, Gaël Varoquaux, Erwan Scornet. Causal effect on a target population: a sensitivity analysis to handle missing covariates. *Journal of Causal Inference*, 2022, 10 (1), pp.372-414. 10.1515/jci-2021-0059 . hal-03473691v1

HAL Id: hal-03473691

<https://hal.science/hal-03473691v1>

Submitted on 9 Dec 2021 (v1), last revised 15 Mar 2023 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Causal effect on a target population: a sensitivity analysis to handle missing covariates

Bénédicte Colnet^{*} Julie Josse[†] Erwan Scornet[‡] Gaël Varoquaux[§]

November 16, 2021

Abstract

Randomized Controlled Trials (RCTs) are often considered as the gold standard to conclude on the causal effect of a given intervention on an outcome, but they may lack of external validity when the population eligible to the RCT is substantially different from the target population. Having at hand a sample of the target population of interest allows to generalize the causal effect. Identifying this target population treatment effect needs covariates in both sets to capture all treatment effect modifiers that are shifted between the two sets. However such covariates are often not available in both sets. Standard estimators then use either weighting (IPSW), outcome modeling (G-formula), or combine the two in doubly robust approaches (AIPSW). In this paper, after completing existing proofs on the complete case consistency of those three estimators, we compute the expected bias induced by a missing covariate, assuming a Gaussian distribution and a semi-parametric linear model. This enables sensitivity analysis for each missing covariate pattern, giving the sign of the expected bias. We also show that there is no gain in imputing a partially-unobserved covariate. Finally we study the replacement of a missing covariate by a proxy. We illustrate all these results on simulations, as well as semi-synthetic benchmarks using data from the Tennessee Student/Teacher Achievement Ratio (STAR), and with a real-world example from critical care medicine.

Keywords: Average treatment effect (ATE); distributional shift; external validity; generalizability; transportability.

1 Introduction

Context Randomized Controlled Trials (RCTs) are often considered as the gold standard to conclude on the causal effect of a given intervention on an outcome, as it is applied randomly (Imbens and Rubin, 2015). Yet, they may face a lack of external validity, when the population eligible to the RCT is substantially different from the target population of the intervention policy (Rothwell, 2005). Indeed, if there are treatment effect modifiers with a different distribution in the target population than that in the trial, generalizing –or transporting– the causal effects measured on the RCT is necessary to estimate the target population causal effect. Using covariates present in both RCT and an observational sample of the target population, this target population average treatment effect (ATE) can be identified and estimated with a variety of methods (Imbens et al., 2005; Cole and Stuart, 2010; Stuart et al., 2011; Pearl and Bareinboim, 2011; Bareinboim and Pearl, 2013; Tipton, 2013; Bareinboim et al., 2014; Pearl and Bareinboim, 2014; Kern et al., 2016; Bareinboim and Pearl, 2016; Buchanan et al., 2018; Stuart et al., 2018; Dong et al., 2020), reviewed in Colnet et al. (2020) and Degtiar and Rose (2021).

The conundrum is the following: the RCT yields an unbiased estimate of the treatment effect, but subject to a population shift, while the observational data at hand gives a unbiased sample of the target population of interest, but with confounding effects. In this context, two main approaches exist to estimate the target population ATE from a RCT. The *Inverse Probability of Sampling Weighting* (IPSW) reweights the RCT sample so that it resembles the target population with respect to the necessary covariates for generalization, while the G-formula models the outcome, using the RCT sample, with and without treatment conditionally to the same covariates, and then to extend the model to the target population of interest. These two methods can be combined in a doubly-robust

^{*}Parietal project-team, INRIA Saclay, Centre de Mathématiques Appliquées, Institut Polytechnique de Paris, Palaiseau, France. (email: benedicte.colnet@inria.fr).

[†]Inria Sophia-Antipolis, Montpellier, France.

[‡]Centre de Mathématiques Appliquées, UMR 7641, École polytechnique, CNRS, Institut Polytechnique de Paris, Palaiseau, France.

[§]Parietal project-team, INRIA Saclay, France.

approach – *Augmented Inverse Probability of Sampling Weighting* (AIPSW) – that enjoys better statistical properties. These methods rely on covariates to capture the heterogeneity of the treatment and the population distributional shift. But the two datasets are seldom acquired as part of a homogeneous effort and as a result they come with different covariates (Susukida et al., 2016; Lesko et al., 2016; Stuart and Rhodes, 2017; Egami and Hartman, 2021; Li et al., 2021). Restricting the analysis to the covariates in common raises the risk of omitting an important one leading to identifiability issues. Controlling biases due to unobserved covariates is of crucial importance to causal inference, were it is known as *sensitivity analysis* (Cornfield et al., 1959; Imbens, 2003; Rosenbaum, 2005). Here, we study this problem for estimators of the target population causal effect, accounting for the different covariates in the two datasets.

Prior art The problem of the missing covariate is central in causal inference as, in an observational study, one can never prove that there is no hidden confounding (Cochrane et al., 1972; Pearl, 2009; Imbens and Rubin, 2015; Hernan, 2020). In that setting, sensitivity analysis strives to assess how far confounding would affect the conclusion of a study (for example, would the ATE be of a different sign with such a hidden confounder). Such approaches date back to a study on the effect of smoking on lung cancer (Cornfield et al., 1959), and have been further developed for both parametric (Imbens, 2003; Rosenbaum, 2005; Dorie et al., 2016; Ichino et al., 2008; Cinelli and Hazlett, 2020) and semi-parametric situations (Franks et al., 2019; Veitch and Zaveri, 2020). Typically, the analysis translates expert judgment into mathematical expression of how much the confounding affects treatment assignment and the outcome, and finally how much the estimated treatment effect is biased (giving for example bounds). In practice the expert must usually provide *sensitivity parameters* that reflect plausible properties of the missing confounder. Classic sensitivity analysis, dedicated to ATE estimation from observational data, use as sensitivity parameters the *impact of the missing covariate on treatment assignment probability* along with the *strength on the outcome* of the missing confounder. However, given that these quantities are hardly directly transposable when it comes to generalization, these approaches cannot be directly applied to estimating the population treatment effect. These parameters have to be respectively replaced by the *sampling bias* and the *strength as treatment effect modifier*, in adequate sensitivity analysis method.

For sensitivity analysis of the generalization shift¹, Andrews and Oster (2019) study the case of a totally unobserved covariate and model the strength of the missing covariate as a treatment effect modifier and on the sampling mechanism in a linear generative model. Dahabreh et al. (2019) propose a sensitivity analysis that models the bias of the conditional average treatment effect without the missing covariate. As the analysis starts from two data sets, the missing covariate can also be partially-observed in one of the two data set, which opens the door to new dedicated methods, in addition to sensitivity methods for totally-missing covariates. Following this observation, Nguyen et al. (2017, 2018) proposes a sensitivity analysis dedicated to this specific purpose. In particular they propose a method to handle the case where a covariate is present in the RCT but not in the observational data set, and proposes a sensitivity analysis under the hypothesis of a linear generative model for the outcome. These three works are to our knowledge the only sensitivity analysis designed for the generalisation purpose. As the missing covariate is partially observed, practitioners sometimes impute this covariate relying on other observed covariates, though this approach is poorly documented so far. For example Lesko et al. (2016) proposes imputation of the CD4 count of the target population using a range of plausible distribution. Imputation is also detailed in the context of individual participant data in meta-analysis (Resche-Rigon et al., 2013; Jolani et al., 2015). A related setting of interest is that where the missing covariate can be approached with a proxy in the two data sets, which has been introduced in the case of transportability by Pearl and Bareinboim (2011) using linguistic skills is a proxy for the age. Can it solve, or at least mitigate, the unobserved variables bias?

Contributions In this work we investigate the problem of a missing covariate that affects the identifiability of the target population average treatment effect (ATE), a common situation when combining different data sources. This work comes after the identifiability assessment, that is we consider that the necessary set of covariates to generalize is known, but a necessary covariate is totally or partially missing. This work focuses on three main estimators IPSW, G-formula, and AIPSW. In Section 2 we recall the definitions of these three estimators in the complete case – that is when all the necessary covariates to ensure identifiability are observed. This section also gives new consistency proofs for IPSW, G-formula, and AIPSW estimators. This same section ends with specific notations to describe all the missing covariate patterns. In Section 3, we propose a method to handle a semi-parametric generative process considering a linear conditional average treatment effect (CATE). We prove that under this semi-parametric model, the bias induced when generalizing the effect with the completely observed set

¹Note that the term *sampling selection bias* or *selection bias* is also used in the literature, but for this work we prefer the term generalization shift as the term bias is used later on for estimators’ bias.

of covariates of the IPSW estimator is the same as that for the G-formula, and we quantify it. We also prove that a linear imputation of a partially missing covariate does not reduce this bias. As mentioned in the introduction, and unlike classic sensitivity analysis, several missing data patterns can be observed: either totally missing or missing in one of the two sets. Therefore Section 3 provides sensitivity analysis frameworks for all the possible missing data patterns, including the case of a proxy variable that would replace the missing one. These results can be useful for users as they may be tempted to consider the intersection of common covariates between the RCT and the observational data. We detail how the different patterns imply either one or two sensitivity parameters. To give users an interpretable analysis, and due to the specificity of the sensitivity parameters at hands, we propose an alternative to Austen plots (Imbens, 2003; Veitch and Zaveri, 2020) that are used to represent and communicate on the sensitivity analysis results. In particular our representation includes the sign of the bias along with a bias’ landscape rather than a threshold. Section 4 presents an extensive synthetic simulation analysis that illustrates theoretical results along with a semi-synthetic data simulation using the Tennessee Student/Teacher Achievement Ratio (STAR) experiment evaluating the effect of class size on children performance in elementary schools (Krueger, 1999) in Section 4.2. Finally, Section 5 provides a real-world analysis to assess the effect of acid tranexomic on the Disability Rating Score (DRS) for trauma patients when a covariate is totally missing.

2 Problem setting: generalizing a causal effect

Subsection 2.1 details the notation used, and recalls RCT’s consistency and identifiability condition necessary to generalize the treatment effect. Subsection 2.2 contains our contribution on consistency of the G-formula, IPSW, and AIPSW estimators. Finally subsection 2.3 introduces the notation when a covariate from the adjustment set is totally or partially unobserved.

2.1 Context and notations

2.1.1 Notations

Notations are grounded on the potential outcome framework (Imbens and Rubin, 2015). We model each observations i in the RCT or observational population as described by a random tuple $(X_i, Y_i(0), Y_i(1), A_i, S_i)$ drawn from a distribution $(X, Y(0), Y(1), A, S) \in \mathcal{X} \times \mathcal{Y}^2 \times \{0, 1\}^2$, such that the observations are iid. For each observations, X_i is a p -dimensional vector of covariates, A_i denotes the binary treatment assignment (with $A_i = 0$ if no treatment and $A_i = 1$ if treated), $Y_i(a)$ is the continuous outcome had the subject been given treatment a (for $a \in \{0, 1\}$), and S_i is a binary indicator for RCT eligibility (i.e., meet the RCT inclusion and exclusion criteria) and willingness to participate if being invited to the trial ($S_i = 1$ if eligible and $S_i = 0$ if not). Assuming consistency of potential outcomes, and no interaction between treated and non-treated subject (SUTVA assumption), we denote by $Y_i = A_i Y_i(1) + (1 - A_i) Y_i(0)$ the outcome realized under treatment assignment A_i .

We define the conditional average treatment effect (CATE):

$$\forall x \in \mathcal{X}, \quad \tau(x) = \mathbb{E}[Y(1) - Y(0) \mid X = x],$$

and the population average treatment effect (ATE):

$$\tau = \mathbb{E}[Y(1) - Y(0)] = \mathbb{E}[\tau(X)].$$

We model the patients belonging to an RCT sample of size n and in an observational data sample of size m by $n + m$ independent random tuples: $\{X_i, Y_i(0), Y_i(1), A_i, S_i\}_{i=1}^{n+m}$, where the RCT samples $i = 1, \dots, n$ are identically distributed according to $\mathcal{P}(X, Y(0), Y(1), A, S \mid S = 1)$, and the observational data samples $i = n + 1, \dots, n + m$ are identically distributed according to $\mathcal{P}(X, Y(0), Y(1), A, S)$. We also denote $\mathcal{R} = \{1, \dots, n\}$ the index set of units observed in the RCT study, and $\mathcal{O} = \{n + 1, \dots, n + m\}$ the index set of units observed in the observational study. For each RCT sample $i \in \mathcal{R}$, we observe $(X_i, A_i, Y_i, S_i = 1)$, while for observational data $i \in \mathcal{O}$, we consider the setting where *we only observe the covariates* X_i , which is a common case in practice. Typical data set is presented on Table 2. A selection diagram - being a causal diagram enriched with a selection variable denoted S where each member of S corresponds to a mechanism by which the populations differ (Pearl and Bareinboim, 2011; Bareinboim and Pearl, 2016) - representative of the whole situation is presented on Figure 1. This DAG illustrates that we are in a non-nested case. Associated encoded assumptions for generalizability are recalled later on (see subsection 2.1.3).




Figure 1: Selection diagram illustrating the general framework: note the specific usage of a square for S , called a "selection variable", either denoting differences in covariates' distributions between data sets (also called the sources) when pointing to $\rightarrow X$, or denoting intervention or not when pointing to $\rightarrow A$. Note that we consider the conditional potential outcomes to be the same in both sources as no arrow points from S to Y . (Pearl and Bareinboim, 2011).

Because the samples in the RCT and observational data do not follow the same covariate distribution, the ATE τ is different from the RCT (or sample²) average treatment effect τ_1 which can be expressed as:

$$\tau \neq \tau_1 = \mathbb{E}[Y(1) - Y(0) \mid S = 1] .$$

This inequality is the core of the *lack of external validity* introduced in the beginning of the work, but formalized with a mathematic expression.

For the rest of the work we denote $\mu_a(x)$ the conditional mean outcome under treatment $a \in \{0, 1\}$ in the observational data: $\mu_a(x) = \mathbb{E}[Y(a) \mid X = x]$. $\mu_{(a)}(\cdot)$ are also called responses surfaces and can be estimated by fitting them on the controls and treated individuals in the RCT respectively to obtain $\hat{\mu}_0(\cdot)$ and $\hat{\mu}_1(\cdot)$. We denote respectively by $e_1(x)$ the propensity score in the RCT population $e_1(x) = \mathbb{P}(A = 1 \mid X = x, S = 1)$. Note that, the function $e_1(x)$ is imposed by the trial characteristics and usually is a constant e_1 , for example $e_1 = 0.5$ (other cases include stratified RCT trials).

All along this document estimators are denoted with the number of observations used in index, for example $\hat{\tau}_n$ if the estimator depends only on the RCT individuals (for example the difference-in-means estimator introduced later on in Definition 1), or $\hat{\tau}_{n,m}$ if it depends on both datasets.

2.1.2 RCT internal validity and estimator

Assumptions already introduced in the Section 2 are the *consistency of treatment assignment* assumption ($Y = AY(1) + (1 - A)Y(0)$) and randomization within the RCT ($Y(a) \perp\!\!\!\perp A \mid S = 1$). Note that to ensure the internal validity of the RCT a last assumption has to be verified, that is, all units from the RCT have a non-zero probability to be in the treated or in the control group.

Assumption 1 (Positivity of trial treatment assignment - Imbens and Rubin (2015)). $\exists \eta_1 > 0, \forall x \in \mathcal{X}, \eta_1 \leq e_1(x) \leq 1 - \eta_1$

These three assumptions guarantee the RCT's internal validity. Under these assumptions the difference-in-means estimator, denoted $\hat{\tau}_{DM,n}$, is a consistent estimator of τ_1 .

Definition 1 (Difference-in-means estimator - Wager (2020)). *When the RCT is not stratified, that is e_1 is constant, then*

$$\hat{\tau}_{DM,n} = \frac{1}{n_1} \sum_{A_i=1} Y_i - \frac{1}{n_0} \sum_{A_i=0} Y_i \quad , \text{ where } n_a = |\{i : A_i = a\}| , \quad (1)$$

and $n_0 + n_1 = n$.

The estimator of a stratified RCT, when $e_1(x)$ is not a constant, is recalled in appendix, along with consistency proofs (see Section C). Note that τ_1 can be estimated with other estimators, using regression adjustments, which enjoys better statistical properties (Wager et al., 2016; Wager, 2020).

²Usually τ_1 is also called the Sample Average Treatment Effect (SATE), when τ is named the Population Average Treatment Effect (PATE) (Stuart et al., 2011; Miratrix et al., 2017; Egami and Hartman, 2021; Degtiar and Rose, 2021).

2.1.3 Assumptions to generalize towards the target population

In order to generalize the RCT estimate to the target population, three additional assumptions are required for identification of the target population ATE τ .

Assumption 2 (Representativity of the observational data). *For all $i \in \mathcal{O}, X_i \sim \mathcal{P}(X)$ where \mathcal{P} is the target population distribution.*

Then, a key assumption concerns the set of covariates that allows the identification of the target population treatment effect. This implies a conditional independence relation being called the *ignorability assumption on trial participation* or *S-ignorability* (Imbens et al., 2005; Stuart et al., 2011; Tipton, 2013; Hartman et al., 2015; Pearl, 2015; Kern et al., 2016; Stuart and Rhodes, 2017; Nguyen et al., 2018; Egami and Hartman, 2021).

Assumption 3 (Ignorability assumption on trial participation - Stuart et al. (2011)). $Y(1) - Y(0) \perp\!\!\!\perp S \mid X$.

Assumption 3 ensures that the ATE is non parametrically identifiable. The covariates X needed to generalize correspond to covariates being *both* treatment effect modifiers *and* subject to a distributional shift between the RCT sample and the target population. Such a set is sometimes called a *separating set* (Egami and Hartman, 2021). Note that Assumption 3 implies a functional relationship between the covariates and the potential outcomes in the RCT and in the target population, that is, for all $x \in \mathcal{X}, \tau_1(x) = \mathbb{E}[Y(1) - Y(0) \mid X = x, S = 1] = \mathbb{E}[Y(1) - Y(0) \mid X = x] = \tau(x)$. Such assumption, named the *transportability* assumption, is encoded in the DAG on Figure 1, where the absence of an S variable pointing to Y , encodes the assumption that X and A specific effects on Y are invariant across the two populations (Pearl and Bareinboim, 2011). This last assumption is effectively used in the estimator identification, as detailed in appendix (see Section C). But in this work we keep considering the widely used Assumption 3.

Assumption 4 (Positivity of trial participation - Stuart et al. (2011)). *There exists a constant c such that for all x with probability 1, $\mathbb{P}(S = 1 \mid X = x) \geq c > 0$*

In this work we assume that a separating set X is already identified, and we focus on the consistency of procedures, along with a sensitivity analysis (detailed in Section 3).

2.2 Estimators of the target population ATE and their consistency

This subsection assumes that previously introduced assumptions to generalize hold, and focuses on finite sample estimation and estimators' consistency. To properly transport the ATE, we successively recall three methods being the G-formula (Lesko et al., 2017; Pearl and Bareinboim, 2011; Dahabreh et al., 2019), the Inverse Propensity Weighting Score (IPSW) (Cole and Stuart, 2010; Lesko et al., 2017; Buchanan et al., 2018), and the Augmented IPSW (AIPSW) estimators (Dahabreh et al., 2020). Note that other methods exist, such as calibration weighting (Dong et al., 2020).

There is an abundant literature studying identification: whether or not it is theoretically possible to recover the true value of the causal effect given assumptions on a data-generating mechanism. For example when a causal quantity can be re-expressed in do-free term, the quantity is said to be identifiable (Pearl, 2009). In identifiable settings, a consistent estimator is desirable: it is a procedure that, given sufficient data, yields estimates of the true causal effect. Most of the existing works on estimating the target population causal effect focus on identification or establish consistency for parametric models or oracle estimators which are not bona fide estimation procedures as they require knowledge of the population data-generation mechanism (Cole and Stuart, 2010; Stuart et al., 2011; Lunceford and Davidian, 2004; Buchanan et al., 2018; Correa et al., 2018; Dahabreh et al., 2019; Egami and Hartman, 2021). To our knowledge, we found no explicit consistency results for the G-formula, IPSW, and AIPSW procedures for a totally non-parametric case, when either the CATE or the weights are estimated from the data without prior knowledge. Here we give the corresponding explicit consistency results (Section A gives the proofs).

2.2.1 G-formula

First, the G-formula estimator consists in modeling the expected values for each potential outcome conditional to the covariates.

Definition 2 (G-formula - Dahabreh et al. (2019)). *The G-formula is denoted $\hat{\tau}_{G,n,m}$, and defined as*

$$\hat{\tau}_{G,n,m} = \frac{1}{m} \sum_{i=n+1}^{n+m} (\hat{\mu}_{1,n}(X_i) - \hat{\mu}_{0,n}(X_i)), \quad (2)$$

where $\hat{\mu}_{a,n}(X_i)$ is an estimator of $\mu_a(X_i)$ obtained on the RCT sample. These intermediary estimates are called nuisance components.

One can observe that the surface response estimators condition the consistency of the estimator. We introduce an assumption on the convergence rate of those nuisance parameters.

Assumption 5 (Consistency of surface responses' estimation). *We denote $\hat{\mu}_{0,n}$ and $\hat{\mu}_{1,n}$ estimators of μ_0 and μ_1 respectively, and denote \mathcal{D}_n the RCT sample, so that*

- (H1-G) For $a \in \{0, 1\}$, almost surely, $\mathbb{E}[\hat{\mu}_{a,n}(X) \mid \mathcal{D}_n] \rightarrow \mathbb{E}[\mu_a(X)]$ when $n \rightarrow \infty$,
- (H2-G) For $a \in \{0, 1\}$, there exist C_1, N_1 so that for all $n \geq N_1$, almost surely, $\mathbb{E}[\hat{\mu}_{a,n}^2(X) \mid \mathcal{D}_n] \leq C_1$.

Theorem 1 (G-formula consistency). *Consider the G-formula estimator in Definition 2 along with Assumptions 1, 2, 3, 4 (identifiability), and Assumption 5 (consistency). Then,*

$$|\hat{\tau}_{G,n,m} - \tau| \xrightarrow{a.s.} 0, \text{ when } n, m \rightarrow \infty.$$

The proof of Theorem 1 is in appendix (see section A).

2.2.2 IPSW

Another approach, called Inverse Propensity Weighting Score (IPSW), consists in weighting the RCT sample so that it resembles the target population distribution.

Definition 3 (Inverse Propensity Weighting Score - IPSW - Stuart et al. (2011); Buchanan et al. (2018)). *The IPSW estimator is denoted $\hat{\tau}_{\text{IPSW},n,m}$, and defined as*

$$\hat{\tau}_{\text{IPSW},n,m} = \frac{1}{m} \sum_{i=1}^n \frac{Y_i}{\hat{\alpha}_{n,m}(X_i)} \left(\frac{A_i}{e_1(X_i)} - \frac{1 - A_i}{1 - e_1(X_i)} \right), \quad (3)$$

where $\hat{\alpha}_{n,m}$ is an estimate of the odd ratio of the indicator of being in the RCT:

$$\alpha(x) = \frac{P(i \in \mathcal{R} \mid \exists i \in \mathcal{R} \cup \mathcal{O}, X_i = x)}{P(i \in \mathcal{O} \mid \exists i \in \mathcal{R} \cup \mathcal{O}, X_i = x)}.$$

This intermediary quantity to estimate, $\alpha(\cdot)$, is called a nuisance component.

Similarly to the G-formula we introduce consistency assumption on the nuisance components of the IPSW, so that the asymptotic consistency of the IPSW can be proven.

Assumption 6 (Consistency assumptions - IPSW). *Denoting $\frac{n}{m\hat{\alpha}_{n,m}(x)}$, estimated weights on the set of observed covariates X , and considering Y , the following conditions hold,*

- (H1-IPSW) $\sup_{x \in \mathcal{X}} \left| \frac{n}{m\hat{\alpha}_{n,m}(x)} - \frac{f_X(x)}{f_{X|S=1}(x)} \right| = \epsilon_{n,m} \xrightarrow{a.s.} 0$, when $n, m \rightarrow \infty$,
- (H2-IPSW) Y is square-integrable.

Theorem 2 (IPSW consistency). *Consider the IPSW estimator in Definition 3 along with Assumptions 1, 2, 3, and 4 (identifiability) and Assumption 6 (consistency). Then,*

$$|\hat{\tau}_{\text{IPSW},n,m} - \tau| \xrightarrow{a.s.} 0, \text{ when } n, m \rightarrow \infty$$

Theorem 2 completes the work from Cole and Stuart (2010); Stuart et al. (2011); Bareinboim and Pearl (2013); Buchanan et al. (2018); Egami and Hartman (2021), assuming neither oracle estimator, nor parametric assumptions on $\alpha(\cdot)$. Proof is available in appendix (see Section A).

2.2.3 AIPSW

The model for the expectation of the outcomes among randomized individuals (used in the G-estimator in Definition 2) and the model for the probability of trial participation (used in IPSW estimator in Definition 3) can be combined to form an Augmented IPSW estimator (AIPSW) that has a doubly robust statistical property. Doubly robust estimators allow to have two chances to specify the nuisance models correctly, either the model for the probability of participation or the expectation of the outcome.

Definition 4 (Augmented IPSW - AIPSW - Dahabreh et al. (2020)). *The AIPSW estimator is denoted $\hat{\tau}_{\text{AIPSW},n,m}$, and defined as*

$$\hat{\tau}_{\text{AIPSW},n,m} = \frac{1}{m} \sum_{i=1}^n \frac{n}{m \hat{\alpha}_{n,m}(X_i)} \left[\frac{A_i (Y_i - \hat{\mu}_{1,n}(X_i))}{e_1(X_i)} - \frac{(1 - A_i) (Y_i - \hat{\mu}_{0,n}(X_i))}{1 - e_1(X_i)} \right] + \frac{1}{m} \sum_{i=n+1}^{m+n} (\hat{\mu}_{1,n}(X_i) - \hat{\mu}_{0,n}(X_i)).$$

To ensure AIPSW consistency, additional assumptions are required on the procedure and the estimated nuisance parameters.

Assumption 7 (Consistency assumptions - AIPSW). *The nuisance parameters are bounded, so that*

- (H1-AIPSW) *The estimated density ratio is bounded, so that $\exists \alpha_0 > 0, \sup_{x \in \mathcal{X}} \left| \frac{n}{m \hat{\alpha}_{n,m}(x)} \right| = \frac{1}{\alpha_0}$,*
- (H2-AIPSW) *The estimated surface responses $\hat{\mu}_{a,n}$ are obtained following a cross-fitting estimation,*
- (H3-AIPSW) *The estimated surface responses $\hat{\mu}_{a,n}$ where $a \in \{0, 1\}$ converge towards functions, even if these functions are not the true surface responses. More precisely, there exist two bounded functions $\xi_1, \xi_0 : \mathcal{X} \rightarrow \mathbb{R}$, such that $\forall a \in \{0, 1\}, \lim_{n \rightarrow +\infty} \sup_{x \in \mathcal{X}} |\xi_a(x) - \hat{\mu}_{a,n}(x)| = 0$.*

Theorem 3 (AIPSW consistency). *Consider the AIPSW estimator in Definition 4, along with Assumptions 1, 2, 3, 4 hold (identifiability), and Assumption 7 (consistency). If Assumption 5 or Assumption 6 also holds then,*

$$|\hat{\tau}_{\text{AIPSW},n,m} - \tau| \xrightarrow{a.s.} 0, \text{ when } n, m \rightarrow \infty$$

Theorem 3 completes Dahabreh et al. (2020)'s work providing a consistency result. Proof is available in appendix (see Section A). The doubly-robustness to model misspecification is illustrated with simulations in Dahabreh et al. (2020); Dong et al. (2020); Colnet et al. (2020).

Which variables to use when estimating the generalization shift Theorems 1, 2, and 3 show that the G-formula, IPSW, and AIPSW estimators give consistent estimation of the target population causal effect when using as covariates treatment-effect modifiers that are shifted between the RCT and the target population. Adding more variables can be counter-productive, as it increases the risk of violating assumption 4, the positivity of trial participation, or add colliders (Pearl, 2015).

2.3 Situation of interest: Missing covariate on one dataset

We study the common situation where both data sets (RCT and observational) contain a different subset of the total covariates X . X can be decomposed as $X = X_{\text{mis}} \cup X_{\text{obs}}$ where X_{obs} denotes the covariates that are presents in both data sets, the RCT and the observational study. X_{mis} denotes the covariates that are either partially-observed in one of the two data sets or totally unobserved in both data sets. We denote by *obs* (resp. *mis*) the index set of observed (resp. missing) covariates. An illustration of the typical data considered are presented in Table 2, with an example of two missing data patterns.

As a consequence, estimators of the target population ATE may be used on X_{obs} only. For clear notation we add a subscript *obs* on any estimator applied on the set X_{obs} rather than X .

Definition 5 (G-formula estimator on the set $X_{\text{obs}} = X \setminus X_{\text{mis}}$). *We denote $\hat{\tau}_{G,n,m,\text{obs}}$ the G-formula estimator applied on the covariates set X_{obs} rather than X . This estimator uses the same observations, but is different from $\hat{\tau}_{G,n,m}$ (Definition 2) as it requires estimation of nuisance components $\hat{\mu}_{a,n,\text{obs}}(x_{\text{obs}})$ for $a \in \{0, 1\}$, which corresponds to the estimated surface response on the set X_{obs} instead of X .*

	Set	Covariates			A	Y
		X ₁	X ₂	X ₃		
1	\mathcal{R}	1.1	20	5.4	1	10.1
	\mathcal{O}	-6	45	8.3	0	8.4
n	\mathcal{R}	0	15	6.2	1	14.5
n + 1	\mathcal{O}
	\mathcal{O}	-2	52	NA	NA	NA
	\mathcal{O}	-1	35	NA	NA	NA
n + m	\mathcal{O}	-2	22	NA	NA	NA

	Set	Covariates			A	Y
		X ₁	X ₂	X ₃		
1	\mathcal{R}	1.1	20	NA	1	10.1
	\mathcal{O}	-6	45	NA	0	8.4
n	\mathcal{R}	0	15	NA	1	14.5
n + 1	\mathcal{O}
	\mathcal{O}	-2	52	3.4	NA	NA
	\mathcal{O}	-1	35	3.1	NA	NA
n + m	\mathcal{O}	-2	22	5.7	NA	NA

Figure 2: Typical structure of the data considered, where a covariate - quantitative or categorical - would be available in the RCT, but not in the observational data set (left) or the reverse situation (right). In this specific example $obs = \{1, 2\}$ ($mis = \{3\}$), corresponds to common (resp. different) covariates in the two datasets.

Definition 6 (IPSW estimator on the set $X_{obs} = X \setminus X_{mis}$). We denote $\hat{\tau}_{IPSW, n, m, obs}$ the IPSW estimator applied on the covariates set X_{obs} rather than X . This estimator uses the same observations, but is different from $\hat{\tau}_{IPSW, n, m}$ (Definition 3) as it requires estimation of nuisance components $\hat{\alpha}_{n, m, obs}(x_{obs})$, which corresponds to the estimated odd ratio of the indicatrix of being in the RCT on the set X_{obs} instead of X .

As detailed later on, such estimators may suffer from bias due to the omitted variable(s) which violates Assumption 3, that is:

$$Y(1) - Y(0) \not\perp S \mid X_{obs}$$

As mentioned earlier, we consider that the user is aware of which variables are treatment effect modifiers and subject to a distributional shift, and we only consider the case where the ignorability on trial participation (Assumption 3) is broken due to totally or partially unobserved set of covariates X_{mis} . Such covariate are denoted *key covariates*.

3 Impact of a missing key covariate for a linear CATE

For the rest of the work, we assume $X, Y(0), Y(1) \in \mathbb{R}^{p+2}$, while previous results did not require such assumption.

3.1 Toward a semi-parametric model

So far, we have not introduced a generative model for Y , as it was not needed for the purpose of generalization. To analyze the effect of a missing covariate, we introduce a nonparametric generative model. We focus on zero-mean additive-error representations, considering that the potential outcomes are generated according to:

$$Y(A) = \mu(A, X) + \varepsilon_A, \tag{4}$$

for any function $\mu \in \mathbf{L}^2(\{0, 1\} \times \mathcal{X} \rightarrow \mathbb{R})$ and such that $\mathbb{E}[\varepsilon_A \mid X] = 0$.

Lemma 1. Assume that the nonparametric generative model of Equation (4) holds, then there exist a function $g : \mathcal{X} \rightarrow \mathbb{R}$ such that

$$Y(A) = g(X) + A\tau(X) + \varepsilon_A, \quad \text{where } \tau(X) := \mathbb{E}[Y(1) - Y(0) \mid X] \tag{5}$$

Lemma 1 simply follows from rewriting Equation (4) accounting for the fact that A is binary and $Y \in \mathbb{R}$ (derivations are detailed in appendix C.3). Such a decomposition is often used in the literature (Chernozhukov et al., 2017; Hahn et al., 2019; Nie and Wager, 2020). This model allows to have a simpler expression of the treatment effect without any additional assumptions, due to the discrete nature of A . This model enables placing independent functional form on the CATE $\tau(X)$, with the interest to highlight treatment effect modifiers variables, such as variables that intervenes in the CATE $\tau(X)$.

Remark on the binary outcome case Note that model (5) is only valid in the case of a continuous outcome Y . In the case of a binary outcome, the model from (4) is not equivalent to (5). The challenge is that A is not additively separable, which Chernozhukov et al. (2017) denotes as *treatment effects that are fully heterogeneous or interactive*. For example, if we consider $y \in \{0, 1\}$ generated with the following logistic regression model, $\text{logit}(\mathbb{P}[y = 1 | X]) = x_1 + x_2 + A$, then the treatment effect would not be homogeneous, and all the covariates would be treatment effect modifiers introducing a coupling between A and X through the logit. In the following work we use Lemma 1 to have a semi-parametric assumption on the CATE $\tau(X)$ and propose a sensitivity analysis.

3.2 Expression of the missing-covariate bias

In this part, we study a special case of the previous nonparametric model in Lemma 1, where the CATE depends linearly on X . We admit there exist $\delta \in \mathbb{R}^p$, and $\sigma \in \mathbb{R}^+$ such that:

$$Y = g(X) + A\langle X, \delta \rangle + \varepsilon, \quad \text{where } \varepsilon \sim \mathcal{N}(0, \sigma^2), \quad (6)$$

where $\tau(X) := \langle X, \delta \rangle$ and $\varepsilon_0 = \varepsilon_1 = \varepsilon$. Under this model, the Average Treatment Effect (ATE) takes the following form:

$$\begin{aligned} \tau &= \mathbb{E}[Y(1)] - \mathbb{E}[Y(0)] \\ &= \mathbb{E}[g(X) + A\langle X, \delta \rangle | A = 1] - \mathbb{E}[g(X) + A\langle X, \delta \rangle | A = 0] \\ &= \langle \mathbb{E}[X], \delta \rangle \end{aligned}$$

Only variables being both treatment effect modifier ($\delta_j \neq 0$) and subject to a distributional change between the RCT and the target population are necessary to generalize the ATE. If such a key covariate is missing, it leads to a biased estimation of the target population ATE. We want to express the bias of such an omitted variable on the transported ATE. But first, we have to specify a context in which a certain permanence of the relationship between X_{obs} and X_{mis} in the two data sets holds. Therefore, we introduce the *Transportability of covariates relationship* assumption.

Assumption 8 (Transportability of covariates relationship). *The distribution of X is Gaussian, that is, $X \sim \mathcal{N}(\mu, \Sigma)$, and transportability of Σ is true, that is, $X | S = 1 \sim \mathcal{N}(\mu_{RCT}, \Sigma)$.*

The plausibility of assumption 8 can be partially-assessed through a statistical test on $\Sigma_{obs,obs}$ for example Box's M test (Box, 1949), supported with visualizations (Friendly and Sigal, 2020) (see a discussion in appendix, Section G). Prior consistency assumptions 5 and 6 formalize the assumptions on the rate of convergence of the nuisance components needed to control the asymptotic bias of the G-formula and IPSW procedures applied on the observed set of covariates X_{obs} . Note that Assumption 6 needs to be slightly extended.

Assumption 9 (Extension of Assumption 6). *Assume Assumption 6, along with,*

- (*H4-IPSW*) For all n, m large enough, $\mathbb{E}[\epsilon_{n,m,obs}]$ exists,

where $\epsilon_{n,m,obs}$ was introduced in Assumption 6.

Theorem 4. *Assume that Model (6) holds, along with Assumptions 2, 3, 4 (identifiability), and Assumption 8. Let B be the following quantity:*

$$B = \sum_{j \in mis} \delta_j \left(\mathbb{E}[X_j] - \mathbb{E}[X_j | S = 1] - \Sigma_{j,obs} \Sigma_{obs,obs}^{-1} (\mathbb{E}[X_{obs}] - \mathbb{E}[X_{obs} | S = 1]) \right),$$

where $\Sigma_{obs,obs}$ is the sub matrix of Σ corresponding to observed index rows and columns, and $\Sigma_{j,obs}$ is the row j with column corresponding to observed index of Σ . Then, the two following statements hold.

- Granting Assumption 5, the asymptotic bias of $\hat{\tau}_{G,n,m,obs}$ is equal to B , that is

$$\tau - \lim_{n,m \rightarrow \infty} \mathbb{E}[\hat{\tau}_{G,n,m,obs}] = B. \quad (7)$$

- Similarly, granting Assumption 9, the asymptotic bias of $\hat{\tau}_{IPSW,n,m,obs}$ is equal to B , that is

$$\tau - \lim_{n,m \rightarrow \infty} \mathbb{E}[\hat{\tau}_{IPSW,n,m,obs}] = B. \quad (8)$$

Proof is given in appendix (see Section B). Note that this bias is a function of the covariates distributional shift (and more precisely, the expected values shift), and also of the CATE due the coefficients of the missing covariates δ_j for $j \in mis$. In simple words, the higher the distributional shift, and the higher the treatment modifying strength, the higher the bias. Simulations are performed to illustrate this theoretical bias, see Figure 5.

From the bias to sensitivity analysis The above theoretical bias can be used to translate expert judgments about the strength of the missing covariates, which corresponds to sensitivity analysis. Section 3.3 details the case of a totally unobserved covariate, Section 3.4 the case of a missing covariate in RCT, Section 3.4 the case of a missing covariate in the observational sample. Section 3.5 completes the previous sections presenting an extension to Austen plots. Finally Section 3.6 details the imputation case, and Section 3.7 the case of a proxy variable.

3.3 Sensitivity analysis when a key covariate is totally unobserved

In this part, X_{mis} is considered totally unobserved.

The approach we propose relies on Imbens (2003)'s prototypical framework. Suppose we have a totally unobserved key covariate, from which the association with observed covariates is known and similar in the trial and the observational study³: $X_{mis} \perp\!\!\!\perp X_{obs}$ and $X_{obs} \perp\!\!\!\perp X_{mis} \mid S = 1$. We also suppose that the complete parametric model is:

$$Y = g(X) + A(\delta, X_{obs}) + A\delta_{mis} X_{mis} + \varepsilon \quad (9)$$

Those assumptions help defining a situation where the asymptotic bias can be quantified, as presented in Lemma 1.

Corollary 1 (Sensitivity model). *Assume that Model (6) holds, along with Assumptions 2, 3, 4 (identifiability), Assumption 8, and $X_{mis} \perp\!\!\!\perp X_{obs}$ and $X_{obs} \perp\!\!\!\perp X_{mis} \mid S = 1$, along with the sensitivity model (9). Then the following statement holds,*

$$\tau - \lim_{n, m \rightarrow \infty} \mathbb{E}[\hat{\tau}_{G, n, m, obs}] = \delta_{mis} \Delta_m$$

where $\Delta_m = \mathbb{E}[X_{mis}] - \mathbb{E}[X_{mis} \mid S = 1]$.

Corollary 1 is a direct consequence of Theorem 4, particularized for the case where $X_{obs} \perp\!\!\!\perp X_{mis}$ and $X_{obs} \perp\!\!\!\perp X_{mis} \mid S = 1$. In this expression, Δ_m and δ_{mis} are called the sensitivity parameters. To estimate the bias implied by an unobserved covariate, we have to determine how strongly X_{mis} is a treatment effect modifier (through δ_{mis}), and how strongly it is linked to the trial inclusion (through the shift between the trial sample and the target population $\Delta_m = \mathbb{E}[X_{mis}] - \mathbb{E}[X_{mis} \mid S = 1]$). Note that this setting is also proposed by Andrews and Oster (2019), but has small differences that we review at the end of this subsection. Nguyen et al. (2018) also mention a similar method in their appendix, but do not perform simulation or particular implementation. Table 1 summarizes the similarities and differences with Imbens (2003) and Andrews and Oster (2019)'s approaches.

	Imbens (2003)	Andrews and Oster (2019)	Sensitivity model
Assumption on covariates	$X_{mis} \perp\!\!\!\perp X_{obs}$	$X_{mis} \perp\!\!\!\perp X_{obs}$	$X_{mis} \perp\!\!\!\perp X_{obs}$
Model on Y	Linear model	Linear model	Linear CATE (9)
Other assumption	Model on A (logit)	Model on S (logit)	-
First sensitivity parameter	Strength on Y, using δ_{mis}	Strength on Y, using δ_{mis}	Strength on Y, using δ_{mis}
Second sensitivity parameter	Strength on A (logit's coefficient)	Strength on S (logit's coefficient)	Δ_m : shift of X_{mis}

Table 1: Summary of the differences in between Imbens (2003)'s method, being a prototypical method for sensitivity analysis for observational data and hidden confounding, and Andrews and Oster (2019)'s method and our method.

Our approach has different assumptions than Andrews and Oster (2019)'s model, with no assumption needed on the trial selection process. In addition, the sensitivity parameter accounting for the distributional change in between the RCT and the observational study is the shift Δ_m on the missing covariate X_{mis} , which is a more interpretable quantity compared to a parametric assumption on the sampling mechanism (for example a logistic regression). For example, medical experts or economists can postulate an *a priori* on the shift that can be translated on the bias.

³If S was a child of covariates in a causal diagram, the independence assumption could be challenged. Indeed, in such a situation the association of X_{obs} and X_{mis} may be different between the trial sample and the target population due to collider bias when conditioning on sample membership. A nested-trial design would fall in this situation. In our set up, we consider that the selection variable is a parent of the covariates (see causal diagram on Figure 1), which exclude being in a collider bias situation.

In practice in this setting of a completely missing covariate sensitivity analysis can be proceed using procedure 1. To represent the amplitude of bias depending on the sensitivity parameters values, a graphical aid derived from Austen plots (Imbens, 2003; Veitch and Zaveri, 2020) is developed and adapted toward a heatmap, see Section 3.5.

Procedure 1: Totally unobserved covariate

```

init :  $\delta_{mis} := [\dots]$ ; // Define range for plausible  $\delta_{mis}$  values
init :  $\Delta_m := [\dots]$ ; // Define range for plausible  $\Delta_m$  values
Compute all possible bias  $\delta_{mis}\Delta_m$  (see Lemma 1)
return Austen plot

```

A partially-observed covariate could always be removed so that this sensitivity analysis could be conducted for every missing data patterns (the variable being missing in the RCT or in the observational data). But this simple method has drawbacks. In particular, the independence assumption ($X_{mis} \perp\!\!\!\perp X_{obs}$) needed here is a strong one, and dropping a partially-observed covariate seems inefficient as it discards available information. Therefore, in the following subsections we propose methods that use the partially-observed covariate – when available – to improve the bias estimation and to remove this independence assumption.

3.4 Sensitivity analysis when a key covariate is partially unobserved

When partially available, we propose to use X_{mis} to have a better estimate of the bias. Unlike the above, this approach does not need the partially observed covariate to be independent of all other covariates, but rather captures the dependencies from the data.

Observed in observational study *In this part, X_{mis} is considered unobserved in the RCT.*

Suppose one key covariate X_{mis} is observed in the observational study, but not in the RCT. Under Assumption 8, the bias of $\hat{\tau}_{G,n,m,obs}$ is derived in Theorem 4. The quantitative bias is informative as it depends only on the regression coefficients δ , and on the shift in expectation between covariates. Furthermore, parts of the bias term can be estimated from the data, helping to reveal sensitivity parameters:

$$B = \underbrace{\delta_{mis}}_{X_{mis}'s\ strength} \left(\underbrace{\mathbb{E}[X_{mis}] - \mathbb{E}[X_{mis} | S = 1]}_{\text{Shift of } X_{mis}: \Delta_m} - \underbrace{\Sigma_{mis,obs} \Sigma_{obs,obs}^{-1} (\mathbb{E}[X_{obs}] - \mathbb{E}[X_{obs} | S = 1])}_{\text{Can be estimated from the data}} \right) \quad (10)$$

The covariance term $\Sigma_{mis,obs} \Sigma_{obs,obs}^{-1}$ can be estimated from the observed data, using the observational study where the necessary covariates are all observed and assuming (8). The shift for the observed set of covariates can be estimated. The remaining parameters, δ_{mis} corresponding to the missing covariates in the complete linear model and $\Delta_m = \mathbb{E}[X_{mis}] - \mathbb{E}[X_{mis} | S = 1]$ are not identifiable from the observed data. These two parameters correspond respectively to the treatment effect modifier strength and the distributional shift importance of the missing covariates. They can be turned into an estimate of the bias using this explicit formulation. In practice, when facing the case of a variable observed only in the observational study, the procedure to follow to perform a sensitivity analysis is summarized in procedure 2. Simulations illustrate how these sensitivity parameters can be used, along with graphical visualization derived from Austen plots (see section 4).

Procedure 2: Observed in observational

```

init :  $\delta_{mis} := [\dots]$ ; // Define range for plausible  $\delta_{mis}$  values
init :  $\Delta_{mis} := [\dots]$ ; // Define range for plausible  $\Delta_{mis}$  values
Estimate  $\Sigma_{obs,obs}$ ,  $\Sigma_{mis,obs}$ , and  $\mathbb{E}[X_{obs}]$  on the observational dataset;
Estimate  $\mathbb{E}[X_{obs} | S = 1]$  on the RCT dataset;
Compute all possible bias for range of  $\delta_{mis}$  and  $\Delta_{mis}$  according to Theorem 4.
return Austen plot

```

Data-driven approach to determine sensitivity parameter Note that giving a range for the shift is probably easier, than giving a range of the coefficients. We propose an empirical method to have a practical data-driven estimation of δ_{mis} . First, learn a linear model of X_{mis} from observed covariates X_{obs} on the observational data, then impute the missing covariate in the trial, and finally obtain $\hat{\delta}_{mis}$ with a Robinson procedure on the imputed trial data. This can give an idea of a possible value for δ_{mis} . This method is used in the semi-synthetic simulation (see Section 4.2).

Observed in RCT *In this part, X_{mis} is considered unobserved in the Observational study.*

This method is already developed by [Nguyen et al. \(2017, 2018\)](#), and we briefly recall its principle in this part. Note that we extend this method by considering a semi-parametric model (6), while they considered a completely linear model. For this missing covariate pattern, only one sensitivity parameter is necessary. As the RCT is the complete data set, the regression coefficients δ of (6) can be estimated for all the key covariates, leading to an estimate $\hat{\delta}_{mis}$ for the partially unobserved covariate. [Nguyen et al. \(2017, 2018\)](#) showed that:

$$\tau = \langle \delta_{obs}, \mathbb{E}[X_{obs}] \rangle + \langle \delta_{mis}, \underbrace{\mathbb{E}[X_{mis}]}_{\text{Unknown}} \rangle. \quad (11)$$

In this case, and as the influence of X_{mis} as a treatment effect moderator can be estimated from the data through $\hat{\delta}_{mis}$, only one sensitivity parameter is needed. Here, it consists of a range of plausible $\mathbb{E}[X_{mis}]$, for example according to a domain expert prior, which is also interpretable.

Note that δ_{mis} can be estimated following a R-learner procedure ([Robinson, 1988](#); [Wager, 2020](#); [Nie and Wager, 2020](#)), in particular in the case of a semi-parametric model with a linear CATE as in (6). The Robinson procedure, also called R-learner, is recalled in Appendix (see Section D). This allows to extend [Nguyen et al. \(2018\)](#)'s work to the semi-parametric case.

Procedure 3: Observed in RCT

```

init :  $\mathbb{E}[X_{mis}] := [\dots]$ ; // Define range for plausible  $\mathbb{E}[X_{mis}]$  values
Estimate  $\delta$  with the Robinson procedure, that is:
Run a non-parametric regression  $Y \sim X$  on the RCT, and denote  $\hat{m}_n(x) = \mathbb{E}[Y | X = x, S = 1]$  the
obtained estimator;
Then define transformed features  $\tilde{Y} = Y - \hat{m}_n(X)$  and  $\tilde{Z} = (A - e_1(X))X$ . Finally, estimate  $\hat{\delta}$  running
the OLS regression on  $\tilde{Y} \sim \tilde{Z}$ ;
Estimate  $\mathbb{E}[X_{obs}]$  on the observational dataset;
Compute all possible bias for range of  $\mathbb{E}[X_{mis}]$  according to (11).
return Austen plot

```

3.5 Vizualization: Austen plots

In this part, the Austen plot proposed concerns the case where two sensitivity parameters are needed, which corresponds to the covariate is either completely missing, or missing in the RCT. When the covariate is missing in the observational data, a simple univariate plot can be drawn.

For now on, each of the sensitivity method suppose to translate sensitivity parameter(s) and to compute the range of bias associated. A last step is to communicate or visualize the range of bias, which is slightly more complicated when there are two sensitivity parameters. Austen plot is a way to aid such judgement ([Imbens, 2003](#); [Veitch and Zaveri, 2020](#)). It consists in having a two-dimensional plot, each of the axis representing the sensitivity parameter, and the solid curve is the set of sensitivity parameters that leads to an estimate that induces a certain bias' threshold. Here, we adapt this method to our settings with several changes. Because coefficients interpretation is hard, a typical practice is to translate a regression coefficient into a partial R^2 . For example, [Imbens \(2003\)](#) prototypical example proposes to interpret the two parameters with partial R^2 . In our case, a close quantity can be used:

$$R^2 \sim \frac{\mathbb{V}[\delta_{mis} X_{mis}]}{\mathbb{V}[\sum_{j \in obs} \hat{\delta}_j X_j]} \quad (12)$$

where the denominator term is obtained when regressing Y on X_{obs} . If this R^2 coefficient is close to 1, then the missing covariate has a similar influence on Y compared to other covariates. On the contrary, if R^2 is close to 0, then the impact of X_{mis} on Y as a treatment effect modifier is small compared to other covariates. But in our case one of the sensitivity parameter is really palpable as it is the covariate shift Δ_m . We advocate keeping the regression coefficient and shift as sensitivity parameter rather than a R^2 to help practitioners as it allows to keep the sign of the bias, than can be in favor of the treatment or not and help interpreting the sensitivity analysis.

On Figure 3 we present a glimpse of the simulation result, to introduce the principle of the Austen plot, with on the left the representation using R^2 and on the right a representation keeping the raw sensitivity parameters. In this plot, we consider the covariate X_3 to be missing, so that we represent *what would be the bias if we missed X_3 ?* The associated sensitivity parameters are represented on each axis. In other word, the Austen plot shows how strong an unobserved key covariate would need to be to induce a bias that would force to reconsider the conclusion




Figure 3: Austen plots: On this figure X_3 is supposed to be a missing covariate. (Left) Regular Austen plot showing how strong an key covariate would need to be to induce a bias of ~ 6 in function of the two sensitivity parameters Δ_m and partial R^2 when a covariate is totally unobserved. (Right) The exact same simulation data are represented, while using rather δ_{mis} than the partial R^2 , and superimposing the heatmap of the bias which allows to reveal the general landscape along with the sign of the bias.

of the study because the bias is above a certain threshold, that is represented by the blue line. For example in our simulation set-up, X_3 is below the threshold as illustrated on Figure 3. The threshold can be proposed by expert, and here we proposed the absolute difference between $\hat{\tau}_{n,m,obs}$ and the RCT estimate $\hat{\tau}_1$ as a natural quantity. In particular, we observe that keeping the sign of the sensitivity parameter allows to be even more confident on the direction of the bias.

3.6 Partially observed covariates: imputation

In this part, X_{mis} is considered unobserved in the RCT or in the observational study.

Another practically appealing solution is to impute the partially-observed covariate, based on the complete data set (whether it is the RCT or the observational one) following Procedure 4. We analyse theoretically in this section the bias of such procedure in Lemma 2.

To ease the mathematical analysis, we focus on a G-formula estimator based on oracles quantities: the best imputation function and the surface responses are assumed to be known. While these are not available in practice, they can be approached with consistent estimates of the imputation functions and the surface responses. The precise formulation of our oracle estimates are given in Definition 7 and Definition 8.

Definition 7 (Oracle estimator when covariate is missing in the observational data set). *Assume that the RCT is complete and that the observational sample contains one missing covariate X_{mis} . We assume that we know*

- (I) *the true response surface μ_1 and μ_0*
- (II) *the true linear relation between X_{mis} as a function of X_{obs} and.*

Our oracle estimate $\hat{\tau}_{G,\infty,m,imp}$ consists in applying the G-formula with the true response surfaces μ_1 and μ_0 (I) on the observational sample, in which the missing covariate has been imputed by the best (linear) function (II).

Definition 8 (Oracle estimator when covariate is missing in the RCT data set). *Assume that the observational sample is complete and that the RCT contains one missing covariate X_{mis} . We assume that we know*

- (I) *the true linear relation between X_{mis} as a function of X_{obs} , which leads to the optimal imputation \hat{X}_{mis} and*
- (II) *the true response surfaces, $\mathbb{E} [Y(a)|X_{obs}, \hat{X}_{mis}, S = 1]$, for $a \in \{0, 1\}$.*

Our oracle estimate $\hat{\tau}_{G,\infty,\infty,imp}$ consists in optimally imputing the missing variable X_{mis} in the RCT (I). Then, the G-formula is applied to the observational sample, with the surface responses that have been perfectly fitted on the completed RCT sample.

Lemma 2 (Oracle bias of imputation in a Gaussian setting). Assume that the CATE is linear (6) and that Assumption 8 holds. Let B be the following quantity:

$$B = \delta_{mis} \left(\mathbb{E}[X_{mis}] - \mathbb{E}[X_{mis} | S = 1] - \Sigma_{j,obs} \Sigma_{obs,obs}^{-1} (\mathbb{E}[X_{obs}] - \mathbb{E}[X_{obs} | S = 1]) \right).$$

- **Complete RCT.** Assume that the RCT is complete and that the observational data set contains a missing covariate X_{mis} . Consider the oracle estimator $\hat{\tau}_{G,\infty,m,imp}$ defined in Definition 7. Then,

$$\tau - \lim_{m \rightarrow \infty} \mathbb{E}[\hat{\tau}_{G,\infty,m,imp}] = B$$

- **Complete Observational.** Assume that the observational data set is complete and that the RCT contains a missing covariate X_{mis} . Consider the oracle estimator $\hat{\tau}_{G,\infty,\infty,imp}$ defined in Definition 8. Then,

$$\tau - \mathbb{E}[\hat{\tau}_{G,\infty,\infty,imp}] = B$$

Derivations are detailed in appendix (see Subsection B.2). Lemma 2 proves that, when the data are Gaussian, learning the imputation on one data set and applying it on the other one leads to a biased estimate, where the asymptotic value of the bias is quantified. Simulations in Section 4 show that the average bias of a finite-sample imputation procedure is similar to the bias of $\hat{\tau}_{G,\infty,\infty,obs}$. This suggests that considering the oracle quantities instead of (finite-sample) estimators yields the correct value for the bias, which results mostly from the distributional shift between the two data sets and not from an intrinsic bias of the imputation or surface response estimates. Yet, simulations suggest that beyond the oracle bias, imputing the missing in the observational study reduces estimation variance when using the G-formula or IPSW estimators in a finite sample size (Figure 11).

Procedure 4: Linear imputation

Model X_{mis} a linear combination of X_{obs} on the complete data set;
 Impute the missing covariate with \hat{X}_{mis} with the previous fitted model;
 Compute $\hat{\tau}$ with the G-formula using the imputed data set $X_{obs} \cup \hat{X}_{mis}$;
return $\hat{\tau}$

3.7 Using a proxy variable for the missing covariate

In this part, X_{mis} is considered totally or partially unobserved, but could be replaced with a proxy variable in both datasets.

Another solution is to use a so-called proxy variable, as illustrated by the structural diagram in Figure 4. Such a variable is a variable associated with the missing one, but not a treatment effect modifier. Could a proxy covariate, available in both data set, overcome this issue? Impact of a proxy in the case of a linear model is documented in econometrics (Chen et al., 2005, 2007; Angrist and Pischke, 2008; Wooldridge, 2016). An example of a proxy variable is the height of children as a proxy for their age. Note that in this case, even if the age is present in one of the two datasets, then only the children’s height is kept in for this method.




Figure 4: Selection diagram illustrating the proxy variable case: Note that there is no arrow from S to the proxy variable, indicating no change in the relation between the proxy and the missing covariate in the trial sample and target population (Pearl and Bareinboim, 2011).

Here, we propose a framework to handle a missing key covariate, with a proxy variable and estimate the bias reduction accounting for the additional noise brought by the proxy.

Assumption 10 (Proxy framework). *Assume that $X_{mis} \perp X_{obs}$, and that there exist a proxy variable X_{prox} such that,*

$$X_{prox} = X_{mis} + \eta$$

where $\mathbb{E}[\eta] = 0$, $\text{Var}[\eta] = \sigma_{prox}^2$, and $\text{Cov}(\eta, X_{mis}) = 0$. In addition we suppose that $\text{Var}[X_{mis}] = \text{Var}[X_{mis} | S = 1] = \sigma_{mis}^2$.

Definition 9. We introduce $\hat{\tau}_{G,n,m,prox}$ the G-formula estimator where X_{mis} is substituted by X_{prox} as detailed in assumption 10.

Lemma 3. *Assume that the generative linear model (6) holds, along with assumption (8) and the proxy framework (10), therefore the bias of $\hat{\tau}_{G,n,m,prox}$ is:*

$$\tau - \mathbb{E}[\hat{\tau}_{G,n,m,prox}] = \delta_{mis} \Delta_m \left(1 - \frac{\sigma_{mis}^2}{\sigma_{mis}^2 + \sigma_{prox}^2} \right)$$

We denote $\hat{\delta}_{prox}$, the estimated coefficient for X_{prox} that can be obtained using a R-learner or T-learner procedure when regressing Y on the set $X_{obs} \cup X_{prox}$.

Corollary 2. *The bias in lemma 3 can be estimated using the following expression:*

$$\tau - \mathbb{E}[\hat{\tau}_{G,n,m,prox}] = \hat{\delta}_{prox} (\mathbb{E}[X_{prox}] - \mathbb{E}[X_{prox} | S = 1]) \frac{\sigma_{prox}^2}{\sigma_{mis}^2}$$

Proof of Lemma 3 and Corollary 2 is detailed in appendix (Proof B.3). Note that, as expected, the average bias reduction strongly depends on the quality of the proxy. In the limit case, if $\sigma_{prox} \sim 0$ so that the correlation between the proxy and the missing variable is one, then the bias is null. In general, if $\sigma_{prox} \gg \sigma_{mis}$ then the proxy variable does not diminish the bias.

Finally, the practical approach is detailed in procedure 5. Note that it requires to have a range of possible σ_{prox} values. We recommend to use the data set on which the proxy along with the partially-unobserved covariate are present, and to obtain an estimation of this quantity on this subset.

Procedure 5: Proxy variable

```

init :  $\sigma_{prox} := [\dots]$ ; // Define range for plausible  $\sigma_{prox}$  values
if  $X_{mis}$  is in RCT then
  init :  $\Delta_{mis} := [\dots]$ ; // Define range for plausible  $\Delta_{mis}$  values
  Estimate  $\delta_{mis}$  with the Robinson procedure (see Procedure 3 for details);
  Compute all possible bias for range of  $\sigma_{prox}$  according to Lemma B.3.
else
  Estimate  $\delta_{prox}$  with the Robinson procedure (see Procedure 3 for details);
  Estimate  $\mathbb{E}[X_{prox}]$  and  $\mathbb{E}[X_{prox} | S = 1]$ ;
  Compute all possible bias for range of  $\sigma_{prox}$  according to Corollary 2.
return Bias's range

```

3.8 Conclusion on the linear CATE case

In the above, we extended the work of Nguyen et al. (2017, 2018); Andrews and Oster (2019), by studying i) a method for each missing covariate pattern, ii) a strategy that proposes to impute the partially-observed covariate, iii) replacing a missing variable with a noisy proxy. All these methods rely on different assumptions recalled in Table 2. Note that one of the key assumption to use the partially observed covariate is to suppose a somewhat constant link between the observed covariate and the partially missing one in the RCT sample and in the target population. This is why we supposed that the variance-covariance matrix is the same in RCT sample and in the observational study (Assumption 8).

4 Synthetic and semi-synthetic simulations

In this section we illustrate each results and methods from Section 3. In particular the synthetic simulations results are summarized, while a detailed presentation of the method along with comments is available in appendix for the interested reader (see Section E).

Missing covariate pattern	Assumption(s) required	Procedure's label
Totally unobserved covariate	$X_{mis} \perp\!\!\!\perp X_{obs}$	1
Partially observed in observational study	Assumption 8	2
Partially observed in RCT	No assumption	3
Proxy variable	Assumptions 8 and 10	5

Table 2: Summary of the assumptions and results pointer for all the sensitivity methods according to the missing covariate pattern when the generative outcome is semi-parametric with a linear CATE (6).

4.1 Synthetic simulations

While results from Section 3 apply to any choice of function for $g(x)$ in (6), here we illustrate the case of a linear function. All the simulations are available on [github](#)⁴, and include non-linear forms for $g(x)$ where results hold.

Simulations parameters We use a similar simulation framework as in [Dong et al. \(2020\)](#) and [Colnet et al. \(2020\)](#), where 5 covariates are generated – independently except for X_1 and X_5 which correlation is set at 0.8 except when explicitly mentioned – with $X_j \sim \mathcal{N}(1, 1)$ for each $j = 1, \dots, 5$. The trial selection process is defined using a logistic regression model, such that:

$$\logit \{P(S = 1 | X)\} = \beta_{s,0} + \beta_{s,1} X_1 + \dots + \beta_{s,5} X_5. \quad (13)$$

This selection process implies that the variance-covariance matrix in the RCT sample and in the target population *is not exactly* the same, so that assumption 8 does not hold all along this simulation scheme. This is interesting to illustrate the resilience to small violations on assumption 8.

The outcome is generated according to a linear model, following (6):

$$Y(a) = \beta_0 + \beta_1 X_1 + \dots + \beta_5 X_5 + a(\delta_1 X_1 + \dots + \delta_5 X_5) + \varepsilon \text{ with } \varepsilon \sim \mathcal{N}(0, 1). \quad (14)$$

In this simulation, $\beta = (5, 5, 5, 5, 5)$, and other parameters are further detailed in Table 3.

Covariates	X_1	X_2	X_3	X_4	X_5
Treatment effect modifier	Yes	Yes	Yes	No	No
Linked to trial inclusion	Yes	No	Yes	Yes	No
δ	$\delta_1 = 30$	$\delta_2 = 30$	$\delta_3 = -10$	$\delta_4 = 0$	$\delta_5 = 0$
β_s	$\beta_{s,1} = -0.4$	$\beta_{s,2} = 0$	$\beta_{s,3} = -0.3$	$\beta_{s,4} = -0.3$	$\beta_{s,5} = 0$
$\cdot \perp\!\!\!\perp X_1$	-	$X_2 \perp\!\!\!\perp X_1$	$X_3 \perp\!\!\!\perp X_1$	$X_4 \perp\!\!\!\perp X_1$	$X_5 \not\perp\!\!\!\perp X_1$

Table 3: Simulations parameters.


First a sample of size 10,000 is drawn from the covariate distribution. From this sample, the selection model (13) is applied which leads to an RCT sample of size $n \sim 1000$. Then, the treatment is generated according to a Bernoulli distribution with probability equal to 0.5, $e_1(x) = e_1 = 0.5$. Finally the outcome is generated according to (14). The observational sample is obtained by drawing a new sample of size $m = 10,000$ from the distribution of the covariates. These parameters imply a target population ATE of $\tau = \sum_{j=1}^5 \delta_j \mathbb{E}[X_j] = \sum_{j=1}^5 \delta_j = 50$. Also, the sample selection ($S = 1$) in (13) with these parameters is biased toward lower values of X_1 (and indirectly X_5), and higher values of X_3 . This situation illustrate a case where $\tau_1 \neq \tau$. Empirically, we obtained $\tau_1 \sim 44$.

Illustration of Theorem 4 Figure 5 presents results of a simulation with 100 repetitions with the full set of covariates (on the Figure see none), and the impact of missing covariate(s) when using the G-formula or the IPSW to generalize. The theoretical bias from Theorem 4 are also represented.

The figure shows that, despite the slight violation of assumption 8, Theorem 4 describes well the bias. The absence of covariates X_2, X_4 and/or X_5 does do not affect the ATE generalization because they are not simultaneously treatment effect modifiers and shifted between the RCT sample and the target population. In addition, the sign of the bias depends on the sign of the coefficient as highlighted by X_1 and X_3 covariate. We also added the case

⁴[BenedicteColnet/unobserved-covariate](#)

Figure 5: Illustration of Theorem 4: Simulation results for the linear model with one, or more, missing covariate(s) when generalizing the treatment effect using the G-formula (Definition 2) or IPSW (Definition 3) estimator on the set of observed covariates. The missing covariate(s) are indicated on the bottom of the plots. The continue blue line refers to τ_1 , the ATE on the RCT population, and the red dashed line to the ATE in the target population. The theoretical bias is obtained from Theorem 4. Simulations are repeated 100 times.



where the two correlated variables X_1 and X_5 are missing, we observe that the bias is higher than when only X_1 is missing. This result is also predicted by the Theorem 4 and illustrates a case where variables acting on Y without being treatment effect modifiers and linked to trial inclusion can be of interest when a key covariate is unobserved or partially-unobserved, through correlation, to diminish the bias.

Sensitivity analysis for a totally unobserved covariate To illustrate this case, the missing covariate has to be supposed independent of all the others, therefore for this paragraph the correlation between X_1 and X_5 is exceptionally set to 0 (see Table 3). Then, according to lemma 1, the two sensitivity parameters δ_{mis} and the shift Δ_m can be used to produce a Austen plot for the bias on the transported ATE. The procedure 1 summarizes the different steps, and the Austen plot’s output result was presented in Figure 3.

Sensitivity analysis when missing a covariate in the RCT In this case the two sensitivity parameters are needed: δ_{mis} and Δ_m . Here we present a situation where all covariates are successively partially missing in the RCT. The sensitivity analysis is represented similarly as in Figure 3. Because each missing variable implies a different landscape due to the dependence relation to other covariates (as stated in Theorem 4), each variable requires a different heatmap (except if covariates are all independent). Figure 6 presents results from the simulation.

Figure 6 illustrates the benefit of this method accounting for other correlated covariates. Sensitivity parameters that invalidate the study are larger when X_1 is missing than when X_2 is missing, as X_2 is independent of the other covariates and X_1 is correlated to X_5 that is included in the study. With the same δ_{mis} value, the bias if X_1 is missing with $\Delta_{mis} = 0.25$ is below 6 (the blue curve) where it would not be the case for X_2 .

Sensitivity analysis when missing a covariate in the observational data Only one sensitivity parameter is needed, being $\mathbb{E}[X_{mis}]$ as recalled in (11). This simulation is used, considering that X_1 is missing, and the procedure 3 is applied. Results are presented in Table 4.

Sensitivity parameter $\mathbb{E}[X_{mis}]$	0.8	0.9	1.0	1.1	1.2
Empirical average $\hat{\tau}_{G,n,m,obs}$	44.0	47.0	50.0	53.0	56.0
Empirical standard deviation $\hat{\tau}_{G,n,m,obs}$	0.3	0.3	0.3	0.3	0.4

Table 4: Simulations results when applying procedure 3: Results of the simulation considering X_1 being partially observed in the RCT, and using the sensitivity method of Nguyen et al. (2017), but with a Robinson procedure to handle semi-parametric generative functions. When varying the sensitivity parameters, the estimated ATE is close to the true ATE ($\tau = 50$) when the sensitivity parameter is closer to the true one ($\mathbb{E}[X_{mis}] = 1$). The results are presented for 100 repetitions.




Figure 6: Simulations results when applying procedure 2: Heatmaps with a blue curve showing how strong an unobserved key covariate would need to be to induce a bias of $\tau - \tau_1 \sim 6$ in function of the two sensitivity parameters Δ_m and δ_{mis} when a covariate is totally unobserved. Each heatmap illustrates a case where the covariate would be missing (indicated on the top of the map), given all other covariates. The cross indicate the coordinate of true sensitivity parameters, in adequation with the bias empirically observed in Figure 5. The bias landscape depends on the dependence of the covariate with other observed covariates, as illustrated with an asymmetric heatmap when X_1 is partially observed, due to the presence of X_5 .

Simulations illustrating imputation (Lemma 2) and usage of a proxy (Lemma 3) are available in appendix, in Section E.

Conclusion Simulations illustrate and support Theorem 4 and all the sensitivity analysis patterns. Additional simulations also illustrate the case of imputation and the proxy in appendix (see Section E). Interestingly the simulation set up violates Assumption 8 as the variance-covariance matrix in the two sources are not the same due to the logistic model to emulate RCT eligibility (13). This empirically supports that reality deviation from the framework of Assumption 8 can be supported, at least when deviation are small.

4.2 Semi-synthetic simulation: the STAR experiment

This semi-synthetic benchmark illustrates the sensitivity analysis in case of a partially observed covariate.


4.2.1 Simulation details

To illustrate the previous methods on data for which the data generation process is not under our control, we use the data from a randomized controlled trial, the Tennessee Student/Teacher Achievement Ratio (STAR) study. This RCT is a pioneering randomized study from the domain of education (Angrist and Pischke, 2008), started in 1985, and designed to estimate the effects of smaller classes in primary school, on the results. This experiment showed a strong payoff to smaller classes (Finn and Achilles, 1990). In addition, the effect has been shown to be heterogeneous (Krueger, 1999), where class size has a larger effect for minority students and those on subsidized lunch. For our purposes, we focus on the same subgroup of children, same treatment (small versus regular classes), and same outcome (average of all grades at the end) as in Kallus et al. (2018).

4509 students are concerned by the treatment randomization, with treatment assignment at first grade only. On the whole data, we estimated an average treatment effect of 12.80 additional points on the grades (95% CI [10.41-15.2]) with the difference-in-means estimator (1). We consider this estimate as the ground truth τ as it is the global RCT. Then, we generate a random sample of 500 children to serve as the observational study. From the rest of the data, we sample a biased RCT according to a logistic regression that defines probability for each class to be selected in the RCT, and using only the variable `g1surban` informing on the neighborhood of the school. The final selection is performed using a Bernoulli procedure, which leads to 563 children in the RCT. The resulting RCT is such that τ_1 is 4.85 (95% CI [-2.07-11.78]) which is underestimated. This is due to the fact that the selection is performed toward children that benefit less from the class size reduction according to previous studies (Finn and Achilles, 1990; Krueger, 1999; Kallus et al., 2018). When generalizing the ATE with the G-formula on the full set of covariates, estimating the nuisance components with a linear model (resp. random forest), and estimating the confidence intervals with a stratified bootstrap (500 repetitions), the target population ATE is recovered with an

estimate of 13.05 (95% CI [5.13-20.88]) (resp. 13.02 (95% CI [4.40-20.04])). When not including the covariate on which the selection is performed (`g1surban`) leads to a biased generalized ATE of 5.87 (95% CI [-2.34-13.06]) (resp. 6.41 (95% CI [-3.02-12.77])). These results are represented on Figure 7.

Figure 7: Simulated STAR data: True target population ATE estimation using all the STAR's RCT data is represented (difference-in-means (1)), along with the ATE estimate of a biased simulated RCT (difference-in-means (1)), and the generalized estimate of the ATE is called Generalized ATE, and is estimated with all covariates or without `g1surban` with the G-formula (Definition 2) using a linear model for the nuisance components. In the case of the G-formula, the confidence intervals are estimated with a stratified bootstrap (500 repetitions). The true population ATE is the red dashed line. Similar results are obtained when nuisance components are estimated with random forest.



4.2.2 Application of the sensitivity methods


We now successively consider two different missing covariate patterns to apply the methods developed in Section 3.

Considering `g1surban` is missing in observational study Then Nguyen et al. (2017)'s method can be applied, and when considering a range of plausible $\mathbb{E}[\text{g1surban}]$, one can estimate the target population ATE. Applying such a method and specifying the following range]2.1, 2.7[(containing the true value for $\mathbb{E}[\text{g1surban}]$) then a range for the generalized ATE is]9.5, 16.7[. Recalling that the *ground truth* value is 12.80 (95% CI[10.41-15.2]), the estimated range is in good overlap with the ground truth value. In other words, a user with this specification of the range would correctly conclude that without this key variable, the generalized ATE is probably underestimated.

Considering `g1surban` is missing in the RCT Figure 8 illustrates the method when the missing covariate is in the RCT data set (see Procedure 2). This method relies on Assumption 8, which we test with a Box M-test on Σ (though in practice such a test could only be performed on $\Sigma_{obs,obs}$). Including only numerical covariates would reject the null hypothesis (p -value = 0.034). Note that beyond violating Assumption 8, some variables are categorical (eg `race` and `gender`). Further discussions about violation of this assumption are available in appendix (Section G).

In this application, applying recommendations in Section 3.4⁵ allowed us to get $\delta_{\text{g1surban}} \sim 11$. We consider that the shift is correctly given by domain expert, and so the true shift is taken with uncertainty corresponding to the 95% confidence interval of a difference in mean. Finally, Figure 8 allows to conclude on a positive bias, that is $\mathbb{E}[\hat{\tau}_{G,n,m,obs}] \leq \tau$. Note that our method underestimate a bit the true bias, with an estimated bias of 6.4 when the true bias is 7.08, delimited with the continue red curve on the top right.

Figure 8: Sensitivity analysis of STAR data: considering the covariate `g1surban` is missing in the RCT. The black cross indicates the point estimate value for the bias would an expert have the true sensitivity values (6.4) and the true bias value is represented with the red line (7.08).



⁵See paragraph entitled *Data-driven approach to determine sensitivity parameter*.

5 Application on critical care data

A motivating application for this work is the generalization to a French target population –represented by the **Traumabase** registry– of the CRASH-3 trial (**CRASH-3, 2019**), evaluating Tranexamic Acide (TXA) to prevent death from traumatic brain injury.

CRASH-3 A total of 175 hospitals in 29 different countries participated to the randomized and placebo-controlled trial, called CRASH-3 (**Dewan et al., 2012**), where adults with Traumatic Brain Injury (TBI) suffering from intracranial bleeding were randomly administrated TXA (**CRASH-3, 2019**). The inclusion criteria of the trial are patients with a Glasgow Coma Scale (GCS)⁶ score of 12 or lower or any intracranial bleeding on CT scan, and no major extracranial bleeding, leading to 9202 patients. The outcome studied is Disability Rating Scale (DRS) after 28 days of injury in patients treated within 3 hours of injury. Such an index is a composite ordinal indicator ranging from 0 to 29, the higher the value, the stronger the disability. This outcome has some drawbacks in the sense that TXA diminishes the probability to die from TBI, and therefore may increase the number of high DRS values (**Brenner et al., 2018**). Therefore, to avoid a censoring or truncation due to death, we consider two approaches: either all individuals are kept with deceased individuals imputed a DRS score of 30, *or* only individuals with a mild or moderate TBI and who survived (according to **CRASH-3 (2019)**) are kept. Both approaches provide similar results. The first approach is presented in this section, and the second approach in appendix (see Section **F**). Six covariates are present at baseline, with age, sex, time since injury, systolic blood pressure, Glasgow Coma Scale score (GCS), and pupil reaction. The difference-in-means estimators (**1**) applied on CRASH3 gives an ATE of -0.3 with [95% CI -0.8 0.2]), therefore not giving evidence of an effect of TXA on DRS.


Traumabase To improve decisions and patient care in emergency departments, the **Traumabase** group, comprising 23 French Trauma centers, collects detailed clinical data from the scene of the accident to the release from the hospital. The resulting database, called the Traumabase, comprises 23,000 trauma admissions to date, and is continually updated. The Traumabase currently comprises around 8,270 patients suffering from TBI.

Predicting the treatment effect on the Traumabase data We want to generalize the treatment effect to the French patients - represented by the Traumabase data base. However an important treatment effect modifier is missing, that is the time between treatment and the trauma. For example, **Mansukhani et al. (2020)** reveal a 10% reduction in treatment effectiveness for every 20-min increase in time to treatment (TTT). In addition TTT is probably shifted between the two populations. Therefore this covariate breaks assumption **3** (ignorability on trial participation), and we propose to apply the methods developed in Section **3**.

Sensitivity analysis The concatenated data set with the RCT and observational data contains 11 667 observations (with $n = 2 690$ and $m = 8 977$). Considering a totally missing covariate, we apply procedure **1**. We assume that time-to-treatment (TTT) is independent of all other variables, for example the ones related to the patient baseline characteristics (e.g. age) or to the severity of the trauma (e.g. the Glasgow score). Clinicians support this assumption as the time to receive the treatment depends on the time for the rescuers to come to the accident area, and not on the other patient characteristics. We first estimated the target population treatment effect with the set of observed variables and the G-formula estimator (Definition **2**), leading to an estimated target population ATE of -0.035 (95% CI [-0.38 - 0.28]). The ATE is estimated using random forests, and the confidence interval with non-parametric stratified bootstrap. As the omission of the TTT variable could affect this conclusion, the sensitivity analysis gives insights on the potential bias. Figure **9** represents the Austen plot of the sensitivity analysis. To compare treatment effect modifier influence and shifts, all covariate are centered and scaled. Therefore the x-axis of the Austen plot now corresponds to the shift in standard deviation compared to the RCT sample. The more a covariate is on the right, the higher the expected value of this quantity is in the observational data set, here the Traumabase. To give a concrete example, the `Glasgow.initial` score has a higher expected value in the Traumabase. In particular here the difference is of 1.47 point higher in the Traumabase, roughly 0.5 standard deviation of the values taken in the RCT. The y-axis represents the coefficient of a normalized linear regression. Bootstrap is used to compute the 95% confidence interval of $\hat{\delta}_{mis}$ and $\hat{\Delta}_{mis}$. As the TTT variable is not supposed to be a stronger treatment effect modifier than the Glasgow score, should the shift be strong this would not lead to a strong bias of the estimated population treatment effect.

⁶The Glasgow Coma Scale (GCS) is a neurological scale which aims to assess a person’s consciousness. The lower the score, the higher the gravity of the trauma.

Figure 9: CRASH-3 and Traumabase: Sensitivity analysis performed with the joint data CRASH-3 and Traumabase, apply procedure 1 with all the observed covariate. Variables have been scaled to interpret coefficient, and center on the RCT’s expected mean so that the shift Δ_{mis} corresponds to the number of standard deviation from CRASH-3 expected values for each covariate.



Conclusion

We have studied sensitivity analyses for causal-effect generalization: the impact of an unobserved confounder, a missing covariate required for identifiability of the causal effect. Generalization settings deal with two datasets, one RCT and on observational data. Hence, we propose procedures suited for when the covariate is missing in one or the other data, or both. We also investigate solving the issue with imputation or a proxy variable. In particular:

1. To go beyond the common requirement that the unobserved confounder is independent from the observed covariates, we instead assume that their covariance is transported (Assumption 8). Our simulation study (4) shows that even with a slightly deformed covariance, the proposed sensitivity analysis procedure give useful estimates of the bias.
2. Our procedures use a sensitivity parameter with a direct interpretation: the shift of the missing covariate, denoted Δ_m . We hope that this will facilitate specifying the sensitivity analysis by domain experts.
3. Leveraging the high interpretability of our sensitivity parameter, our framework concludes on the sign of the estimated bias. This sign is important as accepting a treatment effect highly depends on the direction of the generalization shift. We integrate the above methods into the existing Austen plot visualization, using a heatmap to represent the sign of the estimated bias.

Our proposal inherits limitations from the more standard sensitivity analysis methods with observational data, namely the need of a semi-parametric assumption on the outcome model to translate expert judgments on the bias, along with a Gaussian covariate distribution. Therefore, future extensions of this work could explore ways to relax either the parametric assumption or the distributional assumption to support more robust sensitivity analysis.

Acknowledgments

We would like to acknowledge helpful discussions with Dr Marine LE MORVAN. We also would like to acknowledge the insights, discussions, and medical expertise from the Traumabase group and physicians, in particular Drs François-Xavier AGERON, and Tobias GAUSS. In addition, none of the data analysis part could have been done without the help of Dr. Ian Roberts and the CRASH-3 group, who shared with us the clinical trial data.

A Consistency of G-formula, IPSW, and AIPSW

This appendix contains the proofs of consistency theorems given in Section 2.2.

A.1 Consistency of G-formula

Proof of Theorem 1. First, we derive the expectation of $\hat{\tau}_{G,n,m}$, which satisfies:

$$\begin{aligned}
\mathbb{E}[\hat{\tau}_{G,n,m}] &= \mathbb{E}[\mathbb{E}[\hat{\tau}_{G,n,m} \mid \mathcal{D}_n]] && \text{Conditional expectation property} \\
&= \mathbb{E}\left[\mathbb{E}\left[\frac{1}{m} \sum_{i=n+1}^{n+m} \hat{\mu}_{1,n}(X_i) - \mu_{0,n}(X_i) \mid \mathcal{D}_n\right]\right] \\
&= \mathbb{E}\left[\frac{1}{m} \sum_{i=n+1}^{n+m} \mathbb{E}[\hat{\mu}_{1,n}(X_i) - \hat{\mu}_{0,n}(X_i) \mid \mathcal{D}_n]\right] && \text{Linearity} \\
&= \mathbb{E}[\hat{\mu}_{1,n}(X) - \hat{\mu}_{0,n}(X)],
\end{aligned}$$

since, for all $n+1 \leq i, j \leq n+m$

$$\mathbb{E}[\hat{\mu}_{1,n}(X_i) - \hat{\mu}_{0,n}(X_i) \mid \mathcal{D}_n] = \mathbb{E}[\hat{\mu}_{1,n}(X_j) - \hat{\mu}_{0,n}(X_j) \mid \mathcal{D}_n].$$

Consequently, according to (H1-G) and (H2-G),

$$\lim_{n \rightarrow \infty} \mathbb{E}[\hat{\tau}_{G,n,m}] = \mathbb{E}[\mu_1(X) - \mu_0(X)] = \tau. \quad (15)$$

Therefore $\hat{\tau}_{G,n,m}$ is asymptotically unbiased as long as the surface responses are asymptotically unbiased estimators. Now, we prove that the variance of $\hat{\tau}_{G,n,m}$ is asymptotically null. According to the law of total variance,

$$\text{Var}[\hat{\tau}_{G,n,m}] = \mathbb{E}[\text{Var}[\hat{\tau}_{G,n,m} \mid \mathcal{D}_n]] + \text{Var}[\mathbb{E}[\hat{\tau}_{G,n,m} \mid \mathcal{D}_n]]. \quad (16)$$

Regarding the first term in (16), we have, almost surely,

$$\begin{aligned}
\text{Var}[\hat{\tau}_{G,n,m} \mid \mathcal{D}_n] &= \text{Var}\left[\frac{1}{m} \sum_{i=n+1}^{n+m} \hat{\mu}_{1,n}(X_i) - \hat{\mu}_{0,n}(X_i) \mid \mathcal{D}_n\right] \\
&= \frac{1}{m} \text{Var}[\hat{\mu}_{1,n}(X) - \hat{\mu}_{0,n}(X) \mid \mathcal{D}_n] && \text{i.i.d.} \\
&\leq \frac{4}{m} (\text{Var}[\hat{\mu}_{1,n}(X) \mid \mathcal{D}_n] + \text{Var}[\hat{\mu}_{0,n}(X) \mid \mathcal{D}_n]) && \text{Cauchy-Schwarz} \\
&\leq \frac{8C_1}{m}. && \text{(H2-G)}
\end{aligned}$$

To rewrite the second term in (16), note that, almost surely

$$\begin{aligned}
(\mathbb{E}[\hat{\tau}_{G,n,m} \mid \mathcal{D}_n])^2 &\leq \mathbb{E}[(\hat{\mu}_{1,n}(X) - \hat{\mu}_{0,n}(X))^2 \mid \mathcal{D}_n] && \text{Jensen} \\
&\leq 4 (\mathbb{E}[\hat{\mu}_{1,n}(X)^2 \mid \mathcal{D}_n] + \mathbb{E}[\hat{\mu}_{0,n}(X)^2 \mid \mathcal{D}_n]) \\
&\leq 8C_1. && \text{(H2-G)} \quad (17)
\end{aligned}$$

Note also that, almost surely,

$$\mathbb{E}[\hat{\tau}_{G,n,m} \mid \mathcal{D}_n] = \mathbb{E}[\hat{\mu}_{1,n}(X) - \hat{\mu}_{0,n}(X) \mid \mathcal{D}_n], \quad (18)$$

which converges to τ almost surely, as n tends to infinity, according to (H1-G). Since $\mathbb{E}[\hat{\tau}_{G,n,m} \mid \mathcal{D}_n]$ is uniformly upper bounded by inequality (17) and converge in distribution, we have

$$\text{Var}[\mathbb{E}[\hat{\tau}_{G,n,m} \mid \mathcal{D}_n]] \rightarrow \text{Var}[\tau] = 0. \quad (19)$$

Finally, we have,

$$\begin{aligned}
\text{Var}[\hat{\tau}_{G,n,m}] &= \mathbb{E}[\text{Var}[\hat{\tau}_{G,n,m} \mid \mathcal{D}_n]] + \text{Var}[\mathbb{E}[\hat{\tau}_{G,n,m} \mid \mathcal{D}_n]] \\
&\leq \frac{8C_1}{m} + \text{Var}[\mathbb{E}[\hat{\tau}_{G,n,m} \mid \mathcal{D}_n]],
\end{aligned}$$

which tends to zero as m, n tend to infinity. This concludes the proof. \square

A.2 Consistency of IPSW

Proof of Theorem 2. First, we consider an oracle estimator $\hat{\tau}_{\text{IPSW},n}^*$ that does know the weights $\frac{f_X(x)}{f_{X|S=1}(x)}$, that is

$$\tau_{\text{IPSW},n}^* = \frac{1}{n} \sum_{i=1}^n Y_i \frac{f_X(x_i)}{f_{X|S=1}(x_i)} \left(\frac{A_i}{e_1(x)} - \frac{1-A_i}{1-e_1(x)} \right).$$

Because of the finite variance of Y , the strong law of large numbers (also called Kolmogorov's law) allows to state that:

$$\tau_{\text{IPSW},n}^* \xrightarrow{a.s.} \mathbb{E} \left[Y \frac{f_X(x)}{f_{X|S=1}(x)} \left(\frac{A}{e_1(x)} - \frac{1-A}{1-e_1(x)} \right) \mid S=1 \right] = \tau. \quad (20)$$

This result is largely detailed in [Egami and Hartman \(2021\)](#) (see their appendix). Now, we need to prove that this result also holds for the true estimate $\hat{\tau}_{\text{IPSW},n,m}$. To this aim, we first use the triangle inequality:

$$\begin{aligned} |\hat{\tau}_{\text{IPSW},n,m} - \tau_{\text{IPSW},n}^*| &= \left| \frac{1}{n} \sum_{i=1}^n \left(\frac{A_i Y_i}{e_1(x_i)} - \frac{(1-A_i) Y_i}{1-e_1(x_i)} \right) \left(\frac{n}{\hat{\alpha}_{n,m}(x_i)m} - \frac{f_X(x_i)}{f_{X|S=1}(x_i)} \right) \right| \\ &\leq \frac{1}{n} \sum_{i=1}^n \left| \frac{A_i Y_i}{e_1(x_i)} \left(\frac{n}{\hat{\alpha}_{n,m}(x_i)m} - \frac{f_X(x_i)}{f_{X|S=1}(x_i)} \right) \right| \\ &\quad + \frac{1}{n} \sum_{i=1}^n \left| \frac{(1-A_i) Y_i}{1-e_1(x_i)} \left(\frac{n}{\hat{\alpha}_{n,m}(x_i)m} - \frac{f_X(x_i)}{f_{X|S=1}(x_i)} \right) \right| \\ &\leq \frac{1}{n} \sqrt{\sum_{i=1}^n \left(\frac{A_i Y_i}{e_1(x_i)} \right)^2} \sqrt{\sum_{i=1}^n \left(\frac{n}{\hat{\alpha}_{n,m}(x_i)m} - \frac{f_X(x_i)}{f_{X|S=1}(x_i)} \right)^2} \\ &\quad + \frac{1}{n} \sqrt{\sum_{i=1}^n \left(\frac{(1-A_i) Y_i}{1-e_1(x_i)} \right)^2} \sqrt{\sum_{i=1}^n \left(\frac{n}{\hat{\alpha}_{n,m}(x_i)m} - \frac{f_X(x_i)}{f_{X|S=1}(x_i)} \right)^2} \quad \text{Cauchy-Schwarz} \\ &= \frac{\epsilon_{n,m}}{\sqrt{n}} \left(\sqrt{\sum_{i=1}^n \left(\frac{A_i Y_i}{e_1(x_i)} \right)^2} + \sqrt{\sum_{i=1}^n \left(\frac{(1-A_i) Y_i}{1-e_1(x_i)} \right)^2} \right) \quad \text{(H1-IPSW)} \\ &\leq \frac{2\epsilon_{n,m}}{\eta_1 \sqrt{n}} \sqrt{\sum_{i=1}^n Y_i^2} \quad A \leq 1 \text{ and Assumption 1.} \end{aligned}$$

which tends to zero according to (H1-IPSW) as $n, m \rightarrow \infty$ and because Y is square integrable (H2-IPSW). Finally, combining the last inequality with equation (20), we have

$$\begin{aligned} |\hat{\tau}_{\text{IPSW},n,m} - \tau| &\leq |\hat{\tau}_{\text{IPSW},n,m} - \hat{\tau}_{\text{IPSW},n}^*| + |\hat{\tau}_{\text{IPSW},n}^* - \tau| \\ &\leq \frac{2\epsilon_{n,m}}{\eta_1 \sqrt{n}} \sqrt{\sum_{i=1}^n Y_i^2} + |\hat{\tau}_{\text{IPSW},n}^* - \tau|, \end{aligned}$$

which tends to zero almost surely, according to (20), (H1-IPSW) and (H2-IPSW). \square

A.3 Consistency of AIPSW

The proof of Theorem 3 is based on Assumption 7 and either Assumption 5 or Assumption 6. Therefore the proof contains two parts.

Proof of Theorem 3. Note that because of (H2-AIPSW), the estimated response surfaces are obtained with cross-fitting. The procedure supposes to divide up the data into K evenly sized folds, where K is typically set to 5 or 10. Let $k(\cdot)$ be a mapping from the sample indices $i = 1, \dots, n$ to the K evenly sized data folds, and fit $\hat{\mu}_{0,n}$ and $\hat{\mu}_{1,n}$

with cross-fitting over the K folds using methods tuned for optimal predictive accuracy. For $i \in \{1, n\}$, $\hat{\mu}_{0,n}^{-k(i)}$ and $\hat{\mu}_{1,n}^{-k(i)}$ denote predictions made without using the data fold that the i^{th} training example belongs to. When nothing is mentioned, that is the notation $\hat{\mu}_{0,n}$ and $\hat{\mu}_{1,n}$ is adopted, then it supposes that the estimated surface responses uses the whole data set.

First, grant Assumption 5.

We recall that the AIPSW estimator $\hat{\tau}_{\text{AIPSW},n,m}$ is defined as

$$\begin{aligned} \hat{\tau}_{\text{AIPSW},n,m} &= \frac{1}{n} \sum_{i=1}^n \frac{n}{m\hat{\alpha}_{n,m}(X_i)} \frac{A_i \left(Y_i - \hat{\mu}_{1,n}^{-k(i)}(X_i) \right)}{e_1(X_i)} & A_{m,n} \\ &\quad - \frac{1}{n} \sum_{i=1}^n \frac{n}{m\hat{\alpha}_{n,m}(X_i)} \frac{(1 - A_i) \left(Y_i - \hat{\mu}_{0,n}^{-k(i)}(X_i) \right)}{1 - e_1(X_i)} & B_{m,n} \\ &\quad + \frac{1}{m} \sum_{i=n+1}^{m+n} (\hat{\mu}_{1,n}(X_i) - \hat{\mu}_{0,n}(X_i)) & C_{m,n} \end{aligned}$$

To facilitate the rest of the derivations, each term of $\hat{\tau}_{\text{AIPSW},n,m}$ is denoted with a letter $A_{m,n}$, $B_{m,n}$, or $C_{m,n}$. Because Assumption 5 holds and according to Theorem 1, we have

$$|C_{n,m} - \tau| \xrightarrow{a.s.} 0, \text{ when } n, m \rightarrow \infty.$$

Now, consider the term $A_{n,m}$, so that,

$$\begin{aligned} |A_{n,m}| &\leq \frac{1}{n} \sqrt{\sum_{i=1}^n \left(\frac{n}{m\hat{\alpha}_{n,m}(X_i)} \frac{1}{e_1(X_i)} \right)^2} \sqrt{\sum_{i=1}^n \left(A_i Y_i - \hat{\mu}_{1,n}^{-k(i)}(X_i) \right)^2} & \text{Cauchy-Schwarz} \\ &\leq \frac{1}{n} \frac{1}{\eta_1} \sqrt{\sum_{i=1}^n \left(\frac{n}{m\hat{\alpha}_{n,m}(X_i)} \right)^2} \sqrt{\sum_{i=1}^n \left(A_i Y_i - \hat{\mu}_{1,n}^{-k(i)}(X_i) \right)^2} & \text{Assumption 1} \\ &\leq \frac{1}{\eta_1} \frac{1}{\alpha_0} \sqrt{\sum_{i=1}^n \frac{1}{n} \left(A_i Y_i - \hat{\mu}_{1,n}^{-k(i)}(X_i) \right)^2} & \text{H1-AIPSW} \\ &\rightarrow 0, \text{ when } n, m \rightarrow \infty. & \text{Assumption 5} \end{aligned}$$

The last step assumes the dominated convergence theorem. The reasoning is the same for the term $B_{n,m}$. Therefore,

$$|\hat{\tau}_{\text{AIPSW},n,m} - \tau| \xrightarrow{a.s.} 0, \text{ when } n, m \rightarrow \infty.$$

Now, grant Assumption 6 instead of Assumption 5. Note that the AIPSW estimate can be rewritten as

$$\begin{aligned} \hat{\tau}_{\text{AIPSW},n,m} &= \frac{1}{n} \sum_{i=1}^n \frac{n}{m\hat{\alpha}_{n,m}(X_i)} \left(\frac{A_i Y_i}{e_1(X_i)} - \frac{(1 - A_i) Y_i}{1 - e_1(X_i)} \right) & D_{n,m} \\ &\quad - \frac{1}{n} \sum_{i=1}^n \left(\frac{n}{m\hat{\alpha}_{n,m}(X_i)} - \frac{f_X(X_i)}{f_{X|S=1}(X_i)} \right) \left(\frac{A_i \hat{\mu}_{1,n}^{-k(i)}(X_i)}{e_1(X_i)} \right) & E_{n,m} \\ &\quad + \frac{1}{n} \sum_{i=1}^n \left(\frac{n}{m\hat{\alpha}_{n,m}(X_i)} - \frac{f_X(X_i)}{f_{X|S=1}(X_i)} \right) \left(\frac{(1 - A_i) \hat{\mu}_{0,n}^{-k(i)}(X_i)}{1 - e_1(X_i)} \right) & F_{n,m} \\ &\quad - \frac{1}{n} \sum_{i=1}^n \frac{f_X(X_i)}{f_{X|S=1}(X_i)} \left(\frac{A_i \hat{\mu}_{1,n}^{-k(i)}(X_i)}{e_1(X_i)} - \frac{(1 - A_i) \hat{\mu}_{0,n}^{-k(i)}(X_i)}{1 - e_1(X_i)} \right) & G_n \\ &\quad + \frac{1}{m} \sum_{i=n+1}^{m+n} (\hat{\mu}_{1,n}(X_i) - \hat{\mu}_{0,n}(X_i)). & C_{n,m} \end{aligned}$$

Now, we study the convergence of each of these terms separately. First, we observe that the term $D_{n,m}$ corresponds to the IPSW estimator (Definition 3). According to Assumption 6, this term converges uniformly to τ . Let us now consider the term $E_{n,m}$. First, note that, according to Assumption 7 (H3-AIPSW), the estimated surface responses are uniformly bounded for n large enough, that is, there exist $N \in \mathbb{N}, \mu_M > 0$ such that, for all $a \in \{0, 1\}$,

$$\sup_{x \in \mathcal{X}} |\hat{\mu}_{a,n}(x)| \leq \mu_M.$$

It follows that, for all n large enough,

$$\begin{aligned} |E_{n,m}| &\leq \frac{1}{n} \sqrt{\sum_{i=1}^n \left(\frac{n}{m\hat{\alpha}_{n,m}(X_i)} - \frac{f_X(X_i)}{f_{X|S=1}(X_i)} \right)^2} \sqrt{\sum_{i=1}^n \left(\frac{A_i \hat{\mu}_{1,n}^{-k(i)}(X_i)}{e_1(X_i)} \right)^2} && \text{Cauchy-Schwarz} \\ &\leq \frac{1}{n} \sqrt{\sum_{i=1}^n \left(\frac{n}{m\hat{\alpha}_{n,m}(X_i)} - \frac{f_X(X_i)}{f_{X|S=1}(X_i)} \right)^2} \frac{1}{\eta_1} \sqrt{\sum_{i=1}^n \left(\hat{\mu}_{1,n}^{-k(i)}(X_i) \right)^2} && \text{Assumption 1} \\ &\leq \frac{1}{\sqrt{n}} \sqrt{\sum_{i=1}^n \left(\frac{n}{m\hat{\alpha}_{n,m}(X_i)} - \frac{f_X(X_i)}{f_{X|S=1}(X_i)} \right)^2} \frac{\mu_M}{\eta_1} && \text{Assumption 7} \\ &\rightarrow 0, \text{ when } n, m \rightarrow \infty. && \text{Assumption 6} \end{aligned}$$

The reasoning is the same for the term $F_{n,m}$, which also converges uniformly toward 0 when $n, m \rightarrow \infty$. Now, let us consider the terms G_n and $C_{n,m}$. By Assumption (H3-AIPSW), for all $\varepsilon > 0$, for all n large enough, for all $x \in \mathcal{X}$,

$$\hat{\mu}_{1,n}(x) \in [\xi_1(x) - \varepsilon, \xi_1(x) + \varepsilon].$$

Therefore, for all n large enough, and for all m ,

$$\frac{1}{m} \sum_{i=n+1}^{m+n} (\xi_1(X_i) - \varepsilon) \leq \frac{1}{m} \sum_{i=n+1}^{m+n} \hat{\mu}_{1,n}(X_i) \leq \frac{1}{m} \sum_{i=n+1}^{m+n} (\xi_1(X_i) + \varepsilon).$$

By the law of large numbers, for all m, n large enough,

$$\mathbb{E}[\xi_1(X)] - 2\varepsilon \leq \frac{1}{m} \sum_{i=n+1}^{m+n} \hat{\mu}_{1,n}(X_i) \leq \mathbb{E}[\xi_1(X)] + 2\varepsilon,$$

which proves that $\lim_{m,n \rightarrow \infty} \frac{1}{m} \sum_{i=n+1}^{m+n} \hat{\mu}_{1,n}(X_i) = \mathbb{E}[\xi_1(X)]$. Similarly, $\lim_{m,n \rightarrow \infty} \frac{1}{m} \sum_{i=n+1}^{m+n} \hat{\mu}_{0,n}(X_i) = \mathbb{E}[\xi_0(X)]$, and consequently,

$$\lim_{m,n \rightarrow \infty} C_{n,m} = \mathbb{E}[\xi_1(X)] - \mathbb{E}[\xi_0(X)].$$

We can apply the same reasoning for the term G_n , by taking into account the fact that it uses a cross-fitting strategy. By Assumption 7 (H3-AIPSW), for all $\varepsilon > 0$, for all n large enough, for all $x \in \mathcal{X}$, for all $i \in \{1, \dots, n\}$,

$$\hat{\mu}_{1,n}^{-k(i)}(x) \in [\xi_1(x) - \varepsilon, \xi_1(x) + \varepsilon].$$

Using this inequality, one can obtain

$$\frac{1}{n} \sum_{i=1}^n \frac{f_X(X_i)}{f_{X|S=1}(X_i)} \frac{A_i}{e_1(X_i)} (\xi_1(x) - \varepsilon) \leq \frac{1}{n} \sum_{i=1}^n \frac{f_X(X_i)}{f_{X|S=1}(X_i)} \frac{A_i}{e_1(X_i)} \hat{\mu}_{1,n}^{-k(i)}(X_i) \leq \frac{1}{n} \sum_{i=1}^n \frac{f_X(X_i)}{f_{X|S=1}(X_i)} \frac{A_i}{e_1(X_i)} (\xi_1(x) + \varepsilon).$$

The same reasoning as above shows that, for all $a \in \{0, 1\}$,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \frac{f_X(X_i)}{f_{X|S=1}(X_i)} \frac{A_i}{e_1(X_i)} \hat{\mu}_{a,n}^{-k(i)}(X_i) = \mathbb{E} \left[\frac{f_X(X_i)}{f_{X|S=1}(X_i)} \frac{A_i}{e_1(X_i)} \xi_a(X) \right] = \mathbb{E}[\xi_a(X)].$$

Therefore,

$$\lim_{n \rightarrow \infty} G_n = -(\mathbb{E}[\xi_1(X)] - \mathbb{E}[\xi_0(X)]),$$

hence $G_n + C_{n,m} \xrightarrow{a.s.} 0$, when $n, m \rightarrow \infty$, which concludes the proof. \square

B Proofs for the missing covariate setting

This section gathers proofs related to the case where key covariates (treatment effect modifiers with distributional shift) are missing. In particular this appendix contains the proofs of results presented in Section 3.

B.1 Proof of Theorem 4

This proof contains the proof for the G-formula and the IPSW. Each of the proof relies on an oracle estimator that knows the nuisance components (could it be the surface response or the weights), and then the consistency of the estimators is derived using assumptions on the convergence rate of the nuisance components, so that estimators have a behavior close to that of the oracles.

Proof of Theorem 4 - G-formula. We first introduce an oracle G-formula estimator $\tau_{G,obs}^*$ on the set of observed covariates X_{obs} , defined as

$$\begin{aligned}\tau_{G,obs}^* &= \mathbb{E}[\mathbb{E}[Y(1) - Y(0) \mid X_{obs} = x_{obs}, S = 1]] \\ &= \mathbb{E}[\mathbb{E}[\langle \delta, X \rangle \mid X_{obs} = x_{obs}, S = 1]] && \text{Semi-linear generative model (6)} \\ &= \mathbb{E}[\mathbb{E}[\langle \delta, X_{obs} \rangle + \langle \delta, X_{mis} \rangle \mid X_{obs} = x_{obs}, S = 1]] \\ &= \mathbb{E}[\langle \delta, X_{obs} \rangle] + \mathbb{E}[\mathbb{E}[\langle \delta, X_{mis} \rangle \mid X_{obs} = x_{obs}, S = 1]] && \text{Assumption 3.}\end{aligned}$$

To complete the derivations, it suffices to observe that if Assumption 8 holds, then X is a Gaussian vector distributed as $\mathcal{N}(\mu, \Sigma)$. Hence, according to (Ross, 1998), the conditional expectation $\mathbb{E}[X_{mis} \mid X_{obs}]$ satisfies

$$\mathbb{E}[X_{mis} \mid X_{obs} = x_{obs}] = \mathbb{E}[X_{mis}] + \Sigma_{mis,obs} (\Sigma_{obs,obs})^{-1} (x_{obs} - \mathbb{E}[X_{obs}]). \quad (21)$$

Therefore,

$$\tau_{G,obs}^* = \mathbb{E}[\langle \delta, X_{obs} \rangle] + \langle \delta, \mathbb{E}[X_{mis} \mid S = 1] + \Sigma_{mis,obs} \Sigma_{obs,obs}^{-1} (\mathbb{E}[X_{obs}] - \mathbb{E}[X_{obs} \mid S = 1]) \rangle.$$

Using this result, and deriving the difference with τ leads to:

$$\begin{aligned}\tau - \tau_{G,obs}^* &= \mathbb{E}[\langle \delta, X_{obs} \rangle] + \mathbb{E}[\langle \delta, X_{mis} \rangle] \\ &\quad - \left(\mathbb{E}[\langle \delta, X_{obs} \rangle] + \langle \delta, \mathbb{E}[X_{mis} \mid S = 1] + \Sigma_{mis,obs} \Sigma_{obs,obs}^{-1} (\mathbb{E}[X_{obs}] - \mathbb{E}[X_{obs} \mid S = 1]) \rangle \right) \\ &= \langle \delta, \mathbb{E}[X_{mis}] - \mathbb{E}[X_{mis} \mid S = 1] - \Sigma_{mis,obs} \Sigma_{obs,obs}^{-1} (\mathbb{E}[X_{obs}] - \mathbb{E}[X_{obs} \mid S = 1]) \rangle \\ &= \sum_{j \in mis} \delta_j \left(\mathbb{E}[X_j] - \mathbb{E}[X_j \mid S = 1] - \Sigma_{j,obs} \Sigma_{obs,obs}^{-1} (\mathbb{E}[X_{obs}] - \mathbb{E}[X_{obs} \mid S = 1]) \right) && \text{Re-written}\end{aligned}$$

Note that this formula holds for any number of missing covariates. The last step of the proof is to show that under (H1-G) and (H2-G), then the asymptotic bias of $\hat{\tau}_{G,n,m,obs}$ corresponds to $\tau - \tau_{G,obs}^*$. This derivation is the same as in the proof of the G-formula consistency, same arguments and assumptions are used, except that we consider the surface responses on the set of observed covariates X_{obs} . \square

Proof of Theorem 4 - IPSW. First, we consider an oracle $\tau_{IPSW,obs}^*$ that does know the weights $\frac{f_{X_{obs}}(x_{obs})}{f_{X_{obs} \mid S=1}(x_{obs})}$,

defined as

$$\begin{aligned}
\tau_{\text{IPSW},obs}^* &= \mathbb{E}\left[\tau(X) \frac{f_{X_{obs}}(x_{obs})}{f_{X_{obs}|S=1}(x_{obs})} \mid S = 1\right] \\
&= \int_{\mathcal{X}} \tau(X) \frac{f_{X_{obs}}(x_{obs})}{f_{X_{obs}|S=1}(x_{obs})} f_{X|S=1}(x) dx \\
&= \int_{\mathcal{X}} \tau(X) \frac{f_{X_{obs}}(x_{obs})}{f_{X_{obs}|S=1}(x_{obs})} f_{(X_{mis}|X_{obs}=x_{obs},S=1)}(x_{mis}) f_{X_{obs}|S=1}(x_{obs}) dx \\
&\hspace{15em} \text{Bayes rules} \\
&= \int_{\mathcal{X}_{obs}} \left(\int_{\mathcal{X}_{mis}} \tau(X) f_{(X_{mis}|X_{obs}=x_{obs},S=1)}(x_{mis}) dx_{mis} \right) f_{X_{obs}}(x_{obs}) dx_{obs} \\
&\hspace{15em} \text{Fubini's theorem} \\
&= \int_{\mathcal{X}_{obs}} \left(\int_{\mathcal{X}_{mis}} \sum_{j=1}^p x_j \delta_j f_{(X_{mis}|X_{obs}=x_{obs},S=1)}(x_{mis}) dx_{mis} \right) f_{X_{obs}}(x_{obs}) dx_{obs} \\
&\hspace{15em} \text{Because of a linear CATE} \\
&= \int_{\mathcal{X}_{obs}} \left(\int_{\mathcal{X}_{mis}} \left(\sum_{j \in obs} x_j \delta_j + \sum_{i \in mis} x_i \delta_i \right) f_{(X_{mis}|X_{obs}=x_{obs},S=1)}(x_{mis}) dx_{mis} \right) f_{X_{obs}}(x_{obs}) dx_{obs} \\
&\hspace{15em} \text{Because } X = (X_{mis}, X_{obs}) \\
&= \int_{\mathcal{X}_{obs}} \sum_{j \in obs} x_j \delta_j f_{X_{obs}}(x_{obs}) dx_{obs} \\
&\quad + \int_{\mathcal{X}_{obs}} \left(\int_{\mathcal{X}_{mis}} \sum_{i \in mis} x_i \delta_i f_{(X_{mis}|X_{obs}=x_{obs},S=1)}(x_{mis}) dx_{mis} \right) f_{X_{obs}}(x_{obs}) dx_{obs}.
\end{aligned}$$

Because conditional distribution in case of a multivariate Gaussian distribution has an explicit form (Ross, 1998), one can derive the density of $f_{(X_{mis}|X_{obs}=x_{obs},S=1)}(x_{mis})$, and in particular under Assumption 8, and $\forall i \in mis$ and $\forall x_{obs} \in \mathcal{X}_{obs}$, we obtain

$$\mathbb{E}[X_i \mid X_{obs} = x_{obs} \mid S = 1] = \mathbb{E}[X_i \mid S = 1] + \Sigma_{mis,obs} (\Sigma_{obs,obs})^{-1} (x_{obs} - \mathbb{E}[X_{obs} \mid S = 1]).$$

Therefore, $\forall x_{obs} \in \mathcal{X}_{obs}$:

$$\begin{aligned}
\int_{\mathcal{X}_{mis}} \sum_{i \in mis} x_i \delta_i f_{(X_{mis}|X_{obs}=x_{obs},S=1)}(x_{mis}) dx_{mis} &= \sum_{i \in mis} \int_{\mathcal{X}_{mis}} x_i \delta_i f_{(X_{mis}|X_{obs}=x_{obs},S=1)}(x_{mis}) dx_{mis} \\
&= \sum_{i \in mis} \delta_i \mathbb{E}[X_i \mid X_{obs} = x_{obs} \mid S = 1] \\
&= \sum_{i \in mis} \delta_i (\mathbb{E}[X_i \mid S = 1] \\
&\quad + \Sigma_{mis,obs} (\Sigma_{obs,obs})^{-1} (x_{obs} - \mathbb{E}[X_{obs} \mid S = 1])).
\end{aligned}$$

Finally,

$$\begin{aligned}
\tau_{\text{IPSW},obs}^* &= \int_{\mathcal{X}_{obs}} \left(\sum_{j \in obs} x_j \delta_j \right) f_{X_{obs}}(x_{obs}) dx_{obs} \\
&\quad + \int_{\mathcal{X}_{obs}} \left(\sum_{i \in mis} \delta_i \left(\mathbb{E}[X_i | S = 1] + \Sigma_{mis,obs} (\Sigma_{obs,obs})^{-1} (x_{obs} - \mathbb{E}[X_{obs} | S = 1]) \right) \right) f_{X_{obs}}(x_{obs}) dx_{obs} \\
&= \sum_{j \in obs} \delta_j \mathbb{E}[X_j] + \sum_{i \in mis} \delta_i \left(\mathbb{E}[X_i | S = 1] + \Sigma_{mis,obs} (\Sigma_{obs,obs})^{-1} (\mathbb{E}[X_{obs}] - \mathbb{E}[X_{obs} | S = 1]) \right) \\
&= \tau - \sum_{i \in mis} \delta_i \mathbb{E}[X_i] + \sum_{i \in mis} \delta_i \left(\mathbb{E}[X_i | S = 1] + \Sigma_{mis,obs} (\Sigma_{obs,obs})^{-1} (\mathbb{E}[X_{obs}] - \mathbb{E}[X_{obs} | S = 1]) \right) \\
&= \tau - \sum_{i \in mis} \delta_i \left(\mathbb{E}[X_i] - \mathbb{E}[X_i | S = 1] - \Sigma_{mis,obs} (\Sigma_{obs,obs})^{-1} (\mathbb{E}[X_{obs}] - \mathbb{E}[X_{obs} | S = 1]) \right),
\end{aligned}$$

because under Model 6,

$$\tau = \sum_{j=1}^p \delta_j \mathbb{E}[X_j] = \sum_{i \in mis} \delta_i \mathbb{E}[X_i] + \sum_{j \in obs} \delta_j \mathbb{E}[X_j].$$

Consequently, the oracle estimator $\tau_{\text{IPSW},obs}^*$ is biased, and satisfies

$$\tau - \tau_{\text{IPSW},obs}^* = \sum_{i \in mis} \delta_i \left(\mathbb{E}[X_i] - \mathbb{E}[X_i | S = 1] - \Sigma_{mis,obs} (\Sigma_{obs,obs})^{-1} (\mathbb{E}[X_{obs}] - \mathbb{E}[X_{obs} | S = 1]) \right).$$

The last step of the proof is to show that the asymptotic bias of $\hat{\tau}_{\text{IPSW},n,m,obs}$ converges toward the bias of the oracle estimator $\tau_{\text{IPSW},obs}^*$. This last part is similar to the consistency proof for $\hat{\tau}_{\text{IPSW},n,m}$ (see Section A), but as it contains small adaptations, as highlighted with the additional assumptions (H4-IPSW), the derivations are detailed.

$$\begin{aligned}
\tau_{\text{IPSW},n,obs}^* &= \frac{1}{n} \sum_{i=1}^n Y_i \frac{f_{X_{obs}}(x_{i,obs})}{f_{X_{obs}|S=1}(x_{i,obs})} \left(\frac{A_i}{e_1(x_i)} - \frac{1 - A_i}{1 - e_1(x_i)} \right) \\
&\xrightarrow{a.s.} \mathbb{E} \left[Y \frac{f_{X_{obs}}(x_{obs})}{f_{X_{obs}|S=1}(x_{obs})} \left(\frac{A}{e_1(x)} - \frac{1 - A}{1 - e_1(x)} \right) \mid S = 1 \right] = \tau_{\text{IPSW},obs}^*.
\end{aligned}$$

Then, we consider the IPSW estimator $\hat{\tau}_{\text{IPSW},n,m,obs}$ satisfying

$$\begin{aligned}
|\mathbb{E}[\hat{\tau}_{\text{IPSW},n,m,obs} - \tau_{\text{IPSW},n,obs}^*]| &\leq \mathbb{E} [|\hat{\tau}_{\text{IPSW},n,m,obs} - \tau_{\text{IPSW},n,obs}^*|] \\
&= \mathbb{E} \left[\left| \frac{1}{n} \sum_{i=1}^n \left(\frac{A_i Y_i}{e_1(x_i)} - \frac{(1 - A_i) Y_i}{1 - e_1(x_i)} \right) \left(\frac{n}{\hat{\alpha}_{n,m}(x_{i,obs}) m} - \frac{f_{X_{obs}}(x_{i,obs})}{f_{X_{obs}|S=1}(x_{i,obs})} \right) \right| \right] \\
&\leq \mathbb{E} \left[\left| \frac{1}{n} \sum_{i=1}^n \left(\frac{A_i Y_i}{e_1(x_i)} - \frac{(1 - A_i) Y_i}{1 - e_1(x_i)} \right) \right| \epsilon_{n,m,obs} \right] && \text{Using H2-IPSW} \\
&\leq \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n \left(\frac{2M}{\eta_1} \right) \epsilon_{n,m} \right] && \text{Using H1-IPSW and H3-IPSW} \\
&= \frac{2M}{\eta_1} \mathbb{E} [\epsilon_{n,m,obs}], && \text{Using H4-IPSW}
\end{aligned}$$

which tends to zero according to (H2-IPSW) as $n, m \rightarrow \infty$.

Combining the last inequality with the first result on $\tau_{\text{IPSW},n,obs}^*$, gives

$$\begin{aligned} |\mathbb{E}[\hat{\tau}_{\text{IPSW},n,m,obs} - \tau_{\text{IPSW},obs}^*]| &\leq |\mathbb{E}[\hat{\tau}_{\text{IPSW},n,m,obs} - \hat{\tau}_{\text{IPSW},n,obs}^*]| + |\mathbb{E}[\hat{\tau}_{\text{IPSW},n,obs}^* - \tau_{\text{IPSW},obs}^*]| \\ &\leq \frac{2M}{\eta_1} \mathbb{E}[\epsilon_{n,m,obs}] + |\hat{\tau}_{\text{IPSW},n}^* - \tau| \\ &\rightarrow 0. \end{aligned}$$

Finally, $\mathbb{E}[\hat{\tau}_{\text{IPSW},n,m,obs} - \tau_{\text{IPSW},obs}^*] \rightarrow 0$ as $n, m \rightarrow \infty$. Hence, the asymptotic bias of $\hat{\tau}_{\text{IPSW},n,m,obs}$, satisfies

$$\begin{aligned} \tau - \mathbb{E}[\hat{\tau}_{\text{IPSW},n,m,obs}] &= \tau - \tau_{\text{IPSW},obs}^* \\ &= \sum_{i \in mis} \delta_i \left(\mathbb{E}[X_i] - \mathbb{E}[X_i | S = 1] - \Sigma_{mis,obs} (\Sigma_{obs,obs})^{-1} (\mathbb{E}[X_{obs}] - \mathbb{E}[X_{obs} | S = 1]) \right). \end{aligned}$$

□

B.2 Imputation

This part contains the proof of Lemma 2.

Proof. This proof is divided into two parts, depending on the missing covariate pattern.

Consider the RCT as the complete dataset We assume that the linear link between the missing covariate X_{mis} and the observed one X_{obs} in the trial population is known, so is the true response surfaces $\mu_1(\cdot)$ and $\mu_0(\cdot)$. We consider the estimator $\hat{\tau}_{G,\infty,m,imp}$ based on the two previous oracles quantities. We denote by $c_0, \dots, c_{\#obs}$ the coefficients linking X_{obs} and X_{mis} in the trial, so that, on the event $S = 1$,

$$X_{mis} = c_0 + \sum_{j \in obs} c_j X_j + \varepsilon, \quad (22)$$

where ε is a Gaussian noise satisfying $\mathbb{E}[\varepsilon | X_{obs}] = 0$ almost surely. Since we assume that the true link between X_{mis} and X_{obs} is known (that is we know the coefficients c_0, \dots, c_d), the imputation of the missing covariate on the observational sample writes

$$\hat{X}_{mis} := c_0 + \sum_{j \in obs} c_j X_j. \quad (23)$$

We denote \tilde{X} the imputed data set composed of the observed covariates and the imputed one in the observational sample. The expectation of the oracle estimator $\hat{\tau}_{G,\infty,m,imp}$ is defined as,

$$\begin{aligned} \mathbb{E}[\hat{\tau}_{G,\infty,m,imp}] &= \mathbb{E} \left[\frac{1}{m} \sum_{i=n+1}^{n+m} \mu_1(\tilde{X}_i) - \mu_0(\tilde{X}_i) \right] && \text{By definition of } \hat{\tau}_{G,\infty,m,imp} \\ &= \mathbb{E} \left[\frac{1}{m} \sum_{i=n+1}^{n+m} \langle \delta, \tilde{X}_i \rangle \right] && \text{Linear CATE (6)} \\ &= \mathbb{E} \left[\frac{1}{m} \sum_{i=n+1}^{n+m} \left(\left(\sum_{j \in obs} \delta_j X_{j,i} \right) + \delta_{mis} \hat{X}_{mis,i} \right) \right] \end{aligned}$$

Because of the finite variance of X_{obs} and \hat{X}_{mis} the law of large numbers allows to state that:

$$\lim_{m \rightarrow \infty} \mathbb{E}[\hat{\tau}_{G,\infty,m,imp}] = \left(\sum_{j \in obs} \delta_j \mathbb{E}[X_j] \right) + \delta_{mis} \mathbb{E}[\hat{X}_{mis}].$$

Due to Assumption 8, the distribution of the vector X is Gaussian in both populations, and one can use (21) to write the conditional expectation in the trial population, that is

$$\mathbb{E}[X_{mis} | X_{obs}, S = 1] = \mathbb{E}[X_{mis} | S = 1] + \Sigma_{mis,obs} \Sigma_{obs,obs}^{-1} (X_{obs} - \mathbb{E}[X_{obs} | S = 1]). \quad (24)$$

Combining (22) and (24), one can obtain:

$$c_0 + \sum_{j \in obs} c_j X_j = \mathbb{E}[X_{mis} | S = 1] + \Sigma_{mis,obs} \Sigma_{obs,obs}^{-1} (X_{obs} - \mathbb{E}[X_{obs} | S = 1]). \quad (25)$$

Now, we can compute,

$$\begin{aligned} \mathbb{E}[\hat{X}_{mis}] &= \mathbb{E}[c_0 + \sum_{j \in obs} c_j X_j] \\ &= \mathbb{E} \left[\mathbb{E}[X_{mis} | S = 1] + \Sigma_{mis,obs} \Sigma_{obs,obs}^{-1} (X_{obs} - \mathbb{E}[X_{obs} | S = 1]) \right] \\ &= \mathbb{E}[X_{mis} | S = 1] + \Sigma_{mis,obs} \Sigma_{obs,obs}^{-1} (\mathbb{E}[X_{obs}] - \mathbb{E}[X_{obs} | S = 1]). \end{aligned} \quad (25)$$

This last result allows to conclude that,

$$\lim_{m \rightarrow \infty} \mathbb{E}[\hat{\tau}_{G,\infty,m,imp}] = \left(\sum_{j \in obs} \delta_j \mathbb{E}[X_j] \right) + \delta_{mis} \left(\mathbb{E}[X_{mis} | S = 1] + \Sigma_{mis,obs} \Sigma_{obs,obs}^{-1} (\mathbb{E}[X_{obs}] - \mathbb{E}[X_{obs} | S = 1]) \right).$$

Finally, as $\tau = \sum_{j=1}^p \delta_j \mathbb{E}[X_j]$,

$$\tau - \lim_{m \rightarrow \infty} \mathbb{E}[\hat{\tau}_{G,\infty,m,imp}] = \delta_{mis} \left(\mathbb{E}[X_{mis}] - \mathbb{E}[X_{mis} | S = 1] - \Sigma_{mis,obs} \Sigma_{obs,obs}^{-1} (\mathbb{E}[X_{obs}] - \mathbb{E}[X_{obs} | S = 1]) \right),$$

which concludes this part of the proof.

Consider the observational data as the complete data set We assume here that the true relations between X_{mis} and X_{obs} is known and the true response model is also known. We denote by $\tau_{G,\infty,\infty,imp}$ the estimator based on these two quantities.

More precisely, we denote by $c_0, \dots, c_{\#obs}$ the coefficients linking X_{obs} and X_{mis} in the observational population, so that

$$X_{mis} = c_0 + \sum_{j \in obs} c_j X_j + \varepsilon, \quad (26)$$

where ε is a Gaussian noise satisfying $\mathbb{E}[\varepsilon | X_{obs}] = 0$ almost surely.

As the estimator is an oracle, the relation in (26) is used to impute the missing covariate in the observational sample, so that

$$\hat{X}_{mis} := c_0 + \sum_{j \in obs} c_j X_j. \quad (27)$$

We denote \tilde{X} the imputed data set composed of the observed covariates and the imputed one in the trial population. Note that the \hat{X}_{mis} is a linear combination of X_{obs} in the trial population, and thus a measurable function of X_{obs} .

This property is used below and labelled as (27). As $\tau_{G,\infty,\infty,imp}$ is an oracle, one have:

$$\begin{aligned}\mathbb{E}[\tau_{G,\infty,\infty,imp}] &= \mathbb{E}\left[\mathbb{E}[Y(1) - Y(0) \mid \tilde{X}, S = 1]\right] \\ &= \mathbb{E}\left[\mathbb{E}[Y(1) - Y(0) \mid \hat{X}_{mis}, X_{obs}, S = 1]\right] \\ &= \mathbb{E}\left[\mathbb{E}[Y(1) - Y(0) \mid X_{obs}, S = 1]\right]\end{aligned}\tag{27}$$

$$= \mathbb{E}\left[\left(\sum_{j \in obs} \delta_j X_j\right) + \delta_{mis} \mathbb{E}[X_{mis} \mid X_{obs}, S = 1]\right].\tag{6}$$

$$\begin{aligned}&= \left(\sum_{j \in obs} \delta_j \mathbb{E}[X_j]\right) + \delta_{mis} \mathbb{E}\left[\mathbb{E}[X_{mis} \mid X_{obs}, S = 1]\right] \\ &= \left(\sum_{j \in obs} \delta_j \mathbb{E}[X_j]\right) \\ &\quad + \delta_{mis} \left(\mathbb{E}[X_{mis} \mid S = 1] + \Sigma_{mis,obs} \Sigma_{obs,obs}^{-1} (\mathbb{E}[X_{obs}] - \mathbb{E}[X_{obs} \mid S = 1])\right).\end{aligned}\tag{24}$$

Finally, as $\tau = \sum_{j=1}^p \delta_j \mathbb{E}[X_j]$,

$$\tau - \mathbb{E}[\tau_{G,\infty,\infty,imp}] = \delta_{mis} \left(\mathbb{E}[X_{mis}] - \mathbb{E}[X_{mis} \mid S = 1] - \Sigma_{mis,obs} \Sigma_{obs,obs}^{-1} (\mathbb{E}[X_{obs}] - \mathbb{E}[X_{obs} \mid S = 1])\right),$$

which concludes this part of the proof. \square

B.3 Proxy variable

Proof of Lemma 3. Recall that we denote $\hat{\tau}_{G,n,m,prox}$ the G-formula estimator using X_{prox} instead of X_{mis} in the G-formula. The derivations of $\hat{\tau}_{G,n,m,prox}$ give:

$$\begin{aligned}\mathbb{E}[\hat{\tau}_{G,n,m,prox}] &= \mathbb{E}\left[\mathbb{E}[Y \mid X_{obs}, X_{prox}, S = 1, A = 1] - \mathbb{E}[Y \mid X_{obs}, X_{prox}, S = 1, A = 0]\right] \\ &\quad \text{Definition of } \hat{\tau}_{G,n,m,prox} \\ &= \mathbb{E}\left[\mathbb{E}[g(X) + \langle \delta, X \rangle \mid X_{obs}, X_{prox}, S = 1] - \mathbb{E}[g(X) \mid X_{obs}, X_{prox}, S = 1]\right] \\ &= \mathbb{E}\left[\mathbb{E}[\langle \delta, X \rangle \mid X_{obs}, X_{prox}, S = 1]\right] \\ &= \sum_{j \in obs} \delta_j \mathbb{E}[X_j] + \delta_{mis} \mathbb{E}\left[\mathbb{E}[X_{mis} \mid X_{obs}, X_{prox}, S = 1, A = 1]\right] \\ &\quad \text{Linearity of } Y \text{ (6)} \\ &= \sum_{j \in obs} \delta_j \mathbb{E}[X_j] + \delta_{mis} \mathbb{E}\left[\mathbb{E}[X_{mis} \mid X_{prox}, S = 1]\right] \\ &\quad X_{mis} \perp\!\!\!\perp X_{obs} \text{ (10) and } X \perp\!\!\!\perp A \mid S = 1\end{aligned}$$

The framework of the proxy variable (10) allows to have an expression of the conditional expectation of X_{mis} (Ross, 1998):

$$\mathbb{E}\left[\mathbb{E}[X_{mis} \mid X_{prox}, S = 1]\right] = \mathbb{E}[X_{mis} \mid S = 1] + \frac{\text{Cov}(X_{mis}, X_{prox})}{\mathbb{V}[X_{prox}]} (X_{prox} - \mathbb{E}[X_{prox} \mid S = 1]),$$

where

$$\begin{aligned}\mathbb{V}[X_{prox}] &= \mathbb{V}[X_{mis} + \eta] \\ &= \mathbb{V}[X_{mis}] + \mathbb{V}[\eta] + 2 \underbrace{\text{Cov}(\eta, X_{mis})}_{=0(10)} \\ &= \sigma_{mis}^2 + \sigma_{prox}^2\end{aligned}$$

and

$$\begin{aligned}
\text{Cov}(X_{mis}, X_{prox}) &= \mathbb{E}[X_{mis}X_{prox}] - \mathbb{E}[X_{prox}]^2 & \mathbb{E}[X_{mis}] &= \mathbb{E}[X_{prox}] \\
&= \mathbb{E}[X_{prox}^2 - \eta X_{prox}] - \mathbb{E}[X_{prox}]^2 \\
&= \mathbb{E}[X_{prox}^2] - \mathbb{E}[X_{prox}]^2 - \mathbb{E}[\eta X_{prox}] \\
&= \mathbb{V}[X_{prox}] - \mathbb{E}[\eta X_{mis}] - \mathbb{E}[\eta^2] \\
&= \sigma_{mis}^2 + \sigma_{prox}^2 - 0 - \sigma_{prox}^2 \\
&= \sigma_{mis}^2
\end{aligned}$$

Therefore, we have

$$\mathbb{E}[\mathbb{E}[X_{mis} | X_{prox}, S = 1]] = \mathbb{E}[X_{mis} | S = 1] + \frac{\sigma_{mis}^2}{\sigma_{mis}^2 + \sigma_{prox}^2} (X_{prox} - \mathbb{E}[X_{prox} | S = 1]),$$

which allows us to complete the first derivation:

$$\begin{aligned}
\mathbb{E}[\hat{\tau}_{G,n,m,prox}] &= \sum_{j \in \text{obs}} \delta_j \mathbb{E}[X_j] + \delta_{mis} \mathbb{E} \left[\mathbb{E}[X_{mis} | S = 1] + \frac{\sigma_{mis}^2}{\sigma_{mis}^2 + \sigma_{prox}^2} (X_{prox} - \mathbb{E}[X_{prox} | S = 1]) \right] \\
&= \sum_{j \in \text{obs}} \delta_j \mathbb{E}[X_j] + \delta_{mis} \left(\mathbb{E}[X_{mis} | S = 1] + \frac{\sigma_{mis}^2}{\sigma_{mis}^2 + \sigma_{prox}^2} (\mathbb{E}[X_{prox}] - \mathbb{E}[X_{prox} | S = 1]) \right) \\
&= \sum_{j \in \text{obs}} \delta_j \mathbb{E}[X_j] + \delta_{mis} \left(\mathbb{E}[X_{mis} | S = 1] + \frac{\sigma_{mis}^2}{\sigma_{mis}^2 + \sigma_{prox}^2} (\mathbb{E}[X_{mis}] - \mathbb{E}[X_{mis} | S = 1]) \right),
\end{aligned}$$

since $\mathbb{E}[X_{prox} | S = 1] = \mathbb{E}[X_{mis} | S = 1]$ and $\mathbb{E}[X_{prox}] = \mathbb{E}[X_{mis}]$. Recalling that $\tau = \sum \delta_j \mathbb{E}[X_j]$, the final form of the bias of $\hat{\tau}_{G,n,m,prox}$ can be obtained as

$$\tau - \mathbb{E}[\hat{\tau}_{G,n,m,prox}] = \delta_{mis} (\mathbb{E}[X_{mis}] - \mathbb{E}[X_{mis} | S = 1]) \left(1 - \frac{\sigma_{mis}^2}{\sigma_{mis}^2 + \sigma_{prox}^2} \right).$$

□

Proof of Corollary 2. Note that the final expression of the bias obtained in the previous proof can not be estimated in all missing covariate patterns. For example, if X_{mis} is partially observed in the RCT, then an estimate of δ_{mis} can be computed, and therefore the bias can be estimated. But in all other missing covariate pattern, a temptation is to estimate δ_{prox} from the regression of Y against $X = (X_{obs}, X_{prox})$ with an OLS procedure. [Wooldridge \(2016\)](#) details the infinite sample estimate of such a coefficient:

$$\lim_{n,m \rightarrow \infty} \mathbb{E}[\hat{\delta}_{prox}] = \delta_{mis} \frac{\sigma_{mis}^2}{\sigma_{mis}^2 + \sigma_{prox}^2}$$

Note that the quantity $\frac{\sigma_{mis}^2}{\sigma_{mis}^2 + \sigma_{prox}^2}$ is always lower than 1, therefore if $\delta_{mis} \geq 1$, then $\hat{\delta}_{prox}$ underestimates δ_{mis} . This phenomenon is called the attenuation bias. This point is documented by [Wooldridge \(2016\)](#), and is due to heteroscedasticity in the plug-in regression:

$$\text{Cov}[X_{prox}, \varepsilon] = \text{Cov}[X_{mis} + \eta, \varepsilon - \delta_{mis}\eta] = -\delta_{mis}\sigma_{\eta}^2 \neq 0$$

This asymptotic estimate can be plugged-in into the previous bias estimation:

$$\tau - \mathbb{E}[\hat{\tau}_{G,n,m,prox}] = \hat{\delta}_{prox} (\mathbb{E}[X_{prox}] - \mathbb{E}[X_{prox} | S = 1]) \frac{\sigma_{prox}^2}{\sigma_{mis}^2}$$

□

C Proof of existing results

This section contains only existing derivations, but gathered here with same notations as in Section 2 for better readability. It contains consistency proofs for the difference-in-means estimator (1), and identifiability derivations for the G-formula (Definition 2) and IPSW (Definition 3). We also recall the proof of Lemma 1.

C.1 Consistency within the RCT

In this subsection we recall the identification of τ_1 and how the consistency of the the difference-in-means estimator for τ_1 can be demonstrated. This proof is also detailed in [Wager \(2020\)](#).

C.1.1 Identifiability

Identification of τ_1 . Recall that, by definition, $e_1(x) = P(A = 1 \mid X = x, S = 1)$. Thus, we have,

$$\begin{aligned}
\tau_1(x) &= \mathbb{E}[Y(1) - Y(0) \mid X = x, S = 1] && \text{By definition} \\
&= \mathbb{E}[Y(1) \mid X = x, S = 1] - \mathbb{E}[Y(0) \mid X = x, S = 1] && \text{Linearity of } \mathbb{E}[\cdot] \\
&= \frac{\mathbb{E}[A \mid X = x, S = 1] \mathbb{E}[Y(1) \mid X = x, S = 1]}{e_1(x)} \\
&\quad - \frac{\mathbb{E}[1 - A \mid X = x, S = 1] \mathbb{E}[Y(0) \mid X = x, S = 1]}{1 - e_1(x)} && \text{Multiplication by } \frac{e_1(x)}{e_1(x)} \text{ and } \frac{1 - e_1(x)}{1 - e_1(x)} \\
&= \frac{\mathbb{E}[AY(1) \mid X = x, S = 1]}{e_1(x)} - \frac{\mathbb{E}[(1 - A)Y(0) \mid X = x, S = 1]}{1 - e_1(x)} && \text{RCT ignorability: } Y(0), Y(1) \perp\!\!\!\perp A \mid X = x, S = 1 \\
&= \frac{\mathbb{E}[AY \mid X = x, S = 1]}{e_1(x)} - \frac{\mathbb{E}[(1 - A)Y \mid X = x, S = 1]}{1 - e_1(x)} && \text{SUTVA} \\
&= \mathbb{E}\left[\frac{A}{e_1(x)}Y - \frac{1 - A}{1 - e_1(x)}Y \mid X = x, S = 1\right].
\end{aligned}$$

From this expression, we obtain

$$\tau_1 = \mathbb{E}[\tau_1(X) \mid S = 1] = \mathbb{E}\left[\mathbb{E}\left[\frac{A}{e_1(x)}Y - \frac{1 - A}{1 - e_1(x)}Y \mid X = x, S = 1\right]\right],$$

which allows τ_1 identifiability. □

C.1.2 RCT's consistency

In a RCT, the propensity score is known, and can depend on x . This dependence allows so-called stratified RCT where the treatment assignment proportion can vary among individuals, even if the assignment remains random. To demonstrate the consistency of the difference-in-means estimator, we suppose that $e_1(x)$ is a constant for simplicity.

Lemma 4. *Consistency and asymptotic variance of $\hat{\tau}_{DM,n}$ Under the SUTVA assumption ($Y = Y(A)$) and the random treatment assignment assumption ($A \perp\!\!\!\perp \{Y(0), Y(1)\}$), the difference-in-means estimator is unbiased and consistent for the average treatment effect τ_1 , so that:*

$$\hat{\tau}_{DM,n} \xrightarrow{p} \tau_1$$

In addition, the finite sample variance of $\hat{\tau}_{DM,n}$ is given by:

$$\text{Var}[\hat{\tau}_{DM,n}] = \frac{1}{n_0} \text{Var}[Y(0)] + \frac{1}{n_1} \text{Var}[Y(1)]$$

where n_0 is the number of individuals that are not treated and n_1 the number of treated individuals, such that $n = n_0 + n_1$

Proof of Lemma 4. This proof recalls [Wager \(2020\)](#) approach, with notations adaptations.

We first demonstrate the unbiasedness. Consider $a \in 0, 1$, the treatment assignment, so that we can derive:

$$\begin{aligned}
\mathbb{E}\left[\frac{1}{n_a} \sum_{A=a} Y\right] &= \mathbb{E}[Y \mid A = a] && \text{IID} \\
&= \mathbb{E}[Y(a) \mid A = a] && \text{SUTVA} \\
&= \mathbb{E}[Y(a)] && \text{Random treatment assignment}
\end{aligned}$$

Therefore, $\hat{\tau}_{DM,n}$ is an unbiased estimator:

$$\mathbb{E}[\hat{\tau}_{DM,n}] = \mathbb{E}[Y(1)] - \mathbb{E}[Y(0)] = \tau.$$

To demonstrate the consistency, the second step is to express the variance of $\hat{\tau}_{DM,n}$. We obtain:

$$\text{Var}[\hat{\tau}_{DM,n} \mid n_0, n_1] = \frac{1}{n_0} \text{Var}[Y(0)] + \frac{1}{n_1} \text{Var}[Y(1)].$$

Therefore, one can write:

$$\mathbb{E}\left[(\hat{\tau}_{DM,n} - \tau)^2\right] = \text{Var}(\hat{\tau}_{DM,n}) \xrightarrow{n \rightarrow +\infty} 0.$$

Therefore the difference-in-means estimator is a consistent estimator of τ (in addition to being unbiased). The last part of the proof concerns the asymptotic variance. According to the central limit theorem, we have:

$$\sqrt{n}(\hat{\tau}_{DM,n} - \tau) \xrightarrow[n \rightarrow \infty]{d} \mathcal{N}(0, V_{DM,n}), \text{ where } V_{DM,n} = \text{Var}[Y(0)]/(1 - e_1) + \text{Var}[Y(1)]/e_1$$

The empirical variance of the difference-in-means estimator is such as:

$$\hat{V}_{DM,n} = \frac{1}{n_1 - 1} \sum_{A_i=1} \left(Y_i - \frac{1}{n_1} \sum_{A_i=1} Y_i \right)^2 + \frac{1}{n_0 - 1} \sum_{A_i=0} \left(Y_i - \frac{1}{n_0} \sum_{A_i=0} Y_i \right)^2.$$

A plug-in estimator can be used to obtain a confidence interval, such that:

$$\lim_{n \rightarrow \infty} \mathbb{P} \left[\tau \in \left(\hat{\tau}_{DM,n} \pm \Phi^{-1}(1 - \alpha/2) \sqrt{\hat{V}_{DM,n}/n} \right) \right] = 1 - \alpha.$$

where Φ denotes the standard Gaussian cumulative distribution function. □

C.2 Identification of complete case estimators

In this part we recall the identification derivations of the G-formula (Definition 2) and IPSW (Definition 3) when all the necessary covariates are observed (that is Assumptions 3 and 4 hold).

C.2.1 G-formula

We recall identification derivations for the G-formula:

$$\begin{aligned} \tau &= \mathbb{E}[\tau(X)] && \text{By definition of the CATE and ATE} \\ &= \mathbb{E}[\tau_1(X)] && \text{Transportability or the CATE, Assumption 3} \\ &= \mathbb{E}[\mathbb{E}[Y(1) \mid X, S = 1] - \mathbb{E}[Y(0) \mid X, S = 1]] && \text{Definition of } \tau_1(X) \text{ and Assumption 4} \\ &= \mathbb{E}[\mathbb{E}[Y \mid X, A = 1, S = 1] - \mathbb{E}[Y \mid X, A = 0, S = 1]]. && \text{Consistency assumption: } Y = AY + (1 - A)Y \end{aligned}$$

C.2.2 Inverse-propensity sampling weighting estimator

We recall identification derivations for the IPSW:

$$\begin{aligned} \tau &= \mathbb{E}[\tau(X)] && \text{By definition of the CATE and ATE} \\ &= \mathbb{E}[\tau_1(X)] && \text{Assumption 3} \\ &= \mathbb{E} \left[\frac{f_X(x)}{f_{X|S=1}(x)} \tau_1(X) \mid S = 1 \right], && \text{Assumption 4} \end{aligned}$$

where f_X (resp. $f_{X|S=1}$) denotes the density of X (resp. $X \mid S = 1$). We used the following derivations,

$$\mathbb{E}[\tau_1(X)] = \int_{\mathcal{X}} \tau_1(x) f_X(x) = \int_{\mathcal{X}} \tau_1(x) \frac{f_X(x)}{f_{X|S=1}(x)} f_{X|S=1}(x) = \mathbb{E} \left[\frac{f_X(x)}{f_{X|S=1}(x)} \tau_1(X) \mid S = 1 \right].$$

Now, recall that we denote $\alpha(x)$ the conditional odds ratio of the indicator of being or not in the RCT, defined as

$$\begin{aligned}\alpha(x) &= \frac{P(i \in \mathcal{R} \mid \exists i \in \mathcal{R} \cup \mathcal{O}, X_i = x)}{P(i \in \mathcal{O} \mid \exists i \in \mathcal{R} \cup \mathcal{O}, X_i = x)} \\ &= \frac{P(i \in \mathcal{R})}{P(i \in \mathcal{O})} \times \frac{P(X_i = x \mid i \in \mathcal{R})}{P(X_i = x \mid i \in \mathcal{O})} && \text{Using Bayes rules and Assumption 4} \\ &= \frac{n}{m} \times \frac{f_{X|S=1}(x)}{f_X(x)},\end{aligned}$$

where we used, for the last equality, the fact that $\frac{P(i \in \mathcal{R})}{P(i \in \mathcal{O})} \sim \frac{n}{m}$ and the assumption that we consider the observational study as an unbiased sample of the target population (Assumption 2). Therefore,

$$\tau = \mathbb{E} \left[\frac{n}{m\alpha(X)} \tau_1(X) \mid S = 1 \right].$$

This quantity can be further developed, underlying $\tau_1(X)$ as derived in Proof C.1.1, which gives

$$\tau = \mathbb{E} \left[\frac{nY}{m\alpha(X)} \left(\frac{A}{e_1(x)} - \frac{1-A}{1-e_1(x)} \right) \mid S = 1 \right].$$

This last quantity can be estimated from the data, and the expression in a finite sample is recalled in Definition 3.

C.3 Proof of Lemma 1

The decomposition is well-known, and for example is used by Chernozhukov et al. (2017); Hahn et al. (2019); Nie and Wager (2020).

Proof. Recall that the outcome model is given by (4), that is

$$Y(A) = \mu(A, X) + \varepsilon_A,$$

where Y takes values in \mathbb{R} . The CATE can then be written as

$$\tau(x) = \mathbb{E}[Y(1) - Y(0) \mid X = x] = \mu(1, x) - \mu(0, x) = \mu_1(x) - \mu_0(x).$$

Letting $g(x) = \mu(0, x)$, we obtain

$$\mathbb{E}[Y(A) \mid X = x] = g(X) + A\tau(X).$$

□

D Robinson procedure

This appendix recall the so-called Robinson procedure that aims at estimating the CATE coefficients δ in a semi-parametric equation such as (6). This method was developed by Robinson (1988) and has been further extended (Chernozhukov et al., 2017; Wager, 2020; Nie and Wager, 2020). Such a procedure is called a R-learner, where the R denotes *Robinson* or *Residuals*. We recall the procedure,

1. Run a non-parametric regressions $Y \sim X$ using a parametric or non parametric method. The best method can be chosen with a cross-validation procedure. We denote $\hat{m}_n(x) = \mathbb{E}[Y \mid X = x]$ the estimator obtained.
2. Define the transformed features $\tilde{Y} = Y - \hat{m}_n(X)$ and $\tilde{Z} = (A - 0.5)X$, using the previous procedure \hat{m}_n .
3. Estimate $\hat{\delta}_n$ running the OLS regression on the transformed features $\tilde{Y} \sim \tilde{Z}$.

If the non-parametric regressions of $m(x)$ satisfies $\mathbb{E}[(\hat{m}(X) - m(X))^2]^{\frac{1}{2}} = o_P\left(\frac{1}{n^{1/4}}\right)$, then the procedure to estimate δ is \sqrt{n} -consistent and asymptotically normal,

$$\sqrt{n}(\hat{\delta} - \delta) \Rightarrow \mathcal{N}(0, V_R), \quad V_R = \text{Var}[\tilde{Z}]^{-1} \text{Var}[\tilde{Z}\tilde{Y}] \text{Var}[\tilde{Z}]^{-1}$$


See Chernozhukov et al. (2017); Wager (2020) for details.

E Synthetic simulation - Extension

This section completes the synthetic simulation presented in Section 4.

Simulation parameters Parameters chosen highlight different covariate roles and strength importance. In this setting, covariates X_1, X_2, X_3 are the so-called treatment effect modifiers due to a non-zero δ coefficients, and X_1, X_3, X_4 are shifted from the RCT sample and the target population distribution due to a non-zero β_s coefficient. This situation is in evidence on Figure 10. Therefore covariate X_1 and X_3 are necessary to generalize the treatment effect, because in both groups. Because in the simulation X_2 and X_4 are independent, the set X_1 and X_3 is also sufficient to generalize. Only X_2 has the same marginal distribution in the RCT sample and in the observational study. Note that the amplitude and sign of different coefficients used, along with dependence between variables allows to illustrate several phenomena. For example X_3 is less shifted in between the two samples compared to X_1 because $|\beta_{s,1}| \leq |\beta_{s,3}|$.

Figure 10: DAG associated with the simulation for the linear CATE: Colorized covariates $X_1, X_2,$ and X_3 correspond to treatment effect modifiers due to a non-zero δ coefficients, and the association between X_1 and X_5 is represented with a dashed arrow (unobserved confounding).



Additional comments on Figure 5 Note that depending on the correlation strength between X_1 and X_5 , the missingness of X_1 can lead to different coefficients estimations when using the G-formula estimation, and different bias on the ATE. Table 5 illustrates this situation, where the higher the correlation, the higher the error on the coefficients estimations, but the lower the bias on the ATE when only X_1 is missing.

Table 5: Coefficients estimated in the simulation: Simulation with X_1 as the missing covariate repeated 100 times, means of estimated coefficients for X_5 and bias on ATE using the Robinson procedure.

ρ_{X_1, X_5}	$\delta_5 - \hat{\delta}_5$	$\tau - \hat{\tau}_{obs}^g$
0.05	-6.34	8.24
0.5	-16.78	6.35
0.95	-28.56	0.93

Imputation When a covariate is partially observed, at temptation is to imputed the missing part with a model learned on the complete part as detailed in procedure 4. Section 3 illustrates Lemma 2, as it shows that linear imputation does not diminish the bias compared to a case where the generalization is performed using only the restricted set of observed covariates. On Figure 11 we simulated all the missing covariate patterns (in RCT or in observational) considering X_1 is partially missing, with varying correlation strength between X_5 and X_1 , and fitting a linear imputation model. Imputation does not lead to a lower bias than totally removing the partially observed covariate. Therefore, in case of a partially missing covariate we advocate running a sensitivity analysis rather than a linear imputation.

Proxy variable Finally and to illustrate Lemma 3, the simulation is extended to replace X_1 by a proxy variable, generated following (10) with a varying σ_{prox} . The generalized ATE is estimated with the G-formula. The experiments is repeated 20 times per σ_{prox} values. Results are presented on Figure 12. Whenever σ_{prox} is small compared to σ_{mis} (which is equal to one in this simulation), therefore the bias is small.




Figure 11: Simulations results when imputing (procedure 4): Results when imputing X_1 with a linear model fitted on the complete data set (either the RCT or the observational). All the missing covariate pattern are simulated using either the G-formula or the IPSW estimators. The impact of the correlation between X_1 and X_5 is investigated. Each simulation is repeated 100 times. All procedures have a similar bias as the procedure ignoring the partially-missing covariate (totally.missing), so that a linear imputation (procedure 4) improves neither the bias nor the variance.




Figure 12: Simulation results for proxy variable (procedure 5) Simulation when a key covariate is replaced by a proxy following the proxy-framework (see Assumption 10). The theoretical bias (3) is represented along with the empirical values obtained when generalizing the ATE with the plugged-in G-formula estimator.

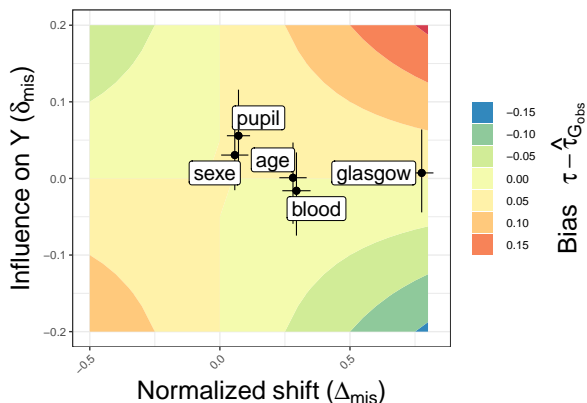
F Application on critical care data - Extension

As explained in Section 5 considering the DRS as the outcome can lead to truncation by death. This part presents alternative analysis if only mild and moderately injured patients (CRASH-3, 2019), with a Glasgow score above 8, are kept along with individuals who survived. When running the exact same analysis leads to keeping 7536 individuals with $m = 1975$ and $n = 5561$, $\hat{\tau}_1$ of -0.20 with [%95 CI $-0.67 - 0.26$], and $\hat{\tau}_{G,obs}$ of -0.09 with [%95 CI $-0.49 - 0.39$]. The corresponding austen plot is presented on Figure 13. Results are similar to Section 5 with no effect of TXA on the outcome, and small heterogeneous treatment effect.

G Homogeneity of the variance-covariance matrix

Recall that Assumption 8 states that the covariance matrices in both data sets are identical. This assumption, which may appear to be very restrictive, can be partially tested on the set of observed covariates. In this section, we present such a test (Box’s M-test Box, 1949), which illustrates the validity of Assumption 8 on some particular data set. Taking one step further, we study the impact of Assumption 8 violation on the resulting estimate.

Figure 13: CRASH-3 and Traumabase: Sensitivity analysis performed with the joint data CRASH-3 and Traumabase keeping only mild and moderate individuals who survived to TBI, apply procedure 1 with all the observed covariates. Variables have been scaled to interpret coefficient, and center on the RCT’s expected mean so that the shift Δ_m corresponds to the number of standard deviation from CRASH-3 expected values for each covariate.



G.1 Statistical test and visualizations

Friendly and Sigal (2020) detail available tests to assess if covariance matrices from two data sample are equal. Despite its sensitivity to violation, Box’s M-test (Box, 1949) can be used test the equality. In particular the package `heplots` contains the tests and visualizations in R. The command line to perform the test is detailed below.

```
library(heplots)
boxM(data[, c("X1", "X2", "X3", "X4")], group = data$S)
```

Even if we cannot bring a general rule to know if the covariance matrices are equal, we can display some examples in which Assumption 8 holds. For instance, Friendly and Sigal (2020) report that the `skull` data is an example of a real data set with multiple sources where there are substantial differences among the means of groups, but little evidence for heterogeneity of their covariance matrices.

G.1.1 Semi-synthetic experiment: STAR

While doing the semi-synthetic experiment on the `Star` data set, the Box M-test rejects the null hypothesis when considering only numerical covariates (`age`, `g1freelunch`, `gkfreelunch`, and `g1surban`) with a p-value of 0.022. This indicates that the preservation of the variance-covariance structure between the two simulated sources does not hold. To help support conclusions, one can visualize how the variance covariance matrix vary in between the two sources, as presented on Figure 14, supporting that the changes in the variance-covariance are not very strong.

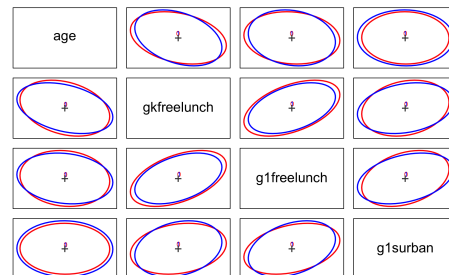


Figure 14: Pairwise data ellipses for the STAR data, centered at the origin. This view allows to compare the variances and covariances for all pairs of variables.

G.1.2 Traumabase and CRASH-3

Note that this part’s purpose is only to illustrate the principle as the application performed in Section 5 relies on the independence between the time to treatment and all other covariates, and not Assumption 8.

Similarly, one can further inspect how far the variance and covariance change in between the two sources. Pairwise data ellipses are presented on Figure 15 for CRASH-3 and Traumabase patients, suggesting rather strong difference in the variance-covariance matrix. As expected Box M-test largely rejects the null hypothesis.

It is interesting to note that in some cases the variance covariance matrix is identical in between two populations. For example we tested whether the two major trauma centers in France present heterogeneity in the variance-covariance matrix, and the Box M test does not reject the null hypothesis.

G.2 Extension of the simulations

Simulations presented in Section 4 can be extended to illustrate empirically the consequences of a poorly specified Assumption 8. Suppose X_1 is the unobserved covariate, and that the variance-covariance matrix is not the same in

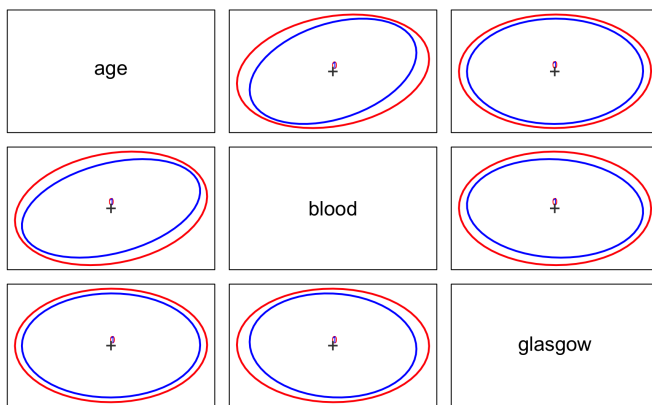


Figure 15: Pairwise data ellipses for the CRASH-3 and Traumabase data, centered at the origin. CRASH-3 data are in blue and Traumabase data in red. This view allows to compare the variances and covariances for all pairs of variables. While the mean are really different in the two sources, the variances and covariances are not so different.


the randomized population ($S = 1$) as in the target population. But the heterogeneities in between the two sources can be different in their nature, affecting covariates depending or not from X_1 . We can imagine two situations, a situation (A) where the link in between X_1 and X_5 is different in the two sources, and another situation (B) where the link in between X_2 and X_3 is not the same. The situation is illustrated on Figures 16a and 16b with pairwise data ellipses. Note that with $n = 1000$ and $m = 10000$ a Box-M test largely rejects the null-hypothesis with a similar statistic value for both situations. When computing the bias according to Theorem 4 and repeating the experiment 50 times, empirical evidence is made that the localization of the heterogeneity impacts or not the bias computation. As presented on Figure 16c, situation A affects the bias computation, when situation B keeps the bias estimation valid.

G.3 Recommendations


Our current recommendations when considering the Assumption 8 is, first, to visualize the heterogeneity of variance-covariance matrix with pairwise data ellipses on $\Sigma_{obs,obs}$. A statistical test such as a Box-M test can be applied on $\Sigma_{obs,obs}$. We also want to emphasize that a statistical test depends on the size of the data sample, when what really matters in this assumption for the sensitivity analysis to be valid is the permanence of covariance structure of the missing covariates with the strongly correlated observed covariates. Simulations presented on Figure 16c is somehow an empirical pathological case where the variance-covariance matrix are equivalently different when considering a statistical test, but leads to different consequences on the validity of Theorem 4, and therefore the sensitivity analysis.

References


- Andrews, I. and E. Oster (2019, 2018). A simple approximation for evaluating external validity bias. *Economics Letters* 178, 58–62. Working Paper.
- Angrist, J. D. and J.-S. Pischke (2008, December). *Mostly Harmless Econometrics: An Empiricist’s Companion*. Princeton University Press.
- Bareinboim, E. and J. Pearl (2013). A general algorithm for deciding transportability of experimental results. *Journal of Causal Inference* 1(1), 107–134.
- Bareinboim, E. and J. Pearl (2016). Causal inference and the data-fusion problem. *Proceedings of the National Academy of Sciences* 113(27), 7345–7352.
- Bareinboim, E., J. Tian, and J. Pearl (2014). Recovering from selection bias in causal and statistical inference. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, AAAI’14, pp. 2410–2416. AAAI Press.
- Box, G. E. P. (1949, 12). A general distribution theory for a class of likelihood criteria. *Biometrika* 36(3-4), 317–346.



(a) Situation A - Pairwise data ellipses



(b) Situation B - Pairwise data ellipses



(c) ATE estimation in the two situations where $\hat{\tau}_{G,obs}$ is estimated considering X_1 is missing and denoted ATE.uncomplete, while the bias B is estimated following Theorem 4 giving ATE.corrected ($\hat{\tau}_{G,obs} + \hat{B}$).

Figure 16: Effect of a different variance-covariance matrix on the ATE estimation, where heterogeneity between the two variance-covariance matrix is introduced as presented in (a) and (b), and on (c) the impact on the estimated average treatment effect (ATE). Situations A and B result in a similar statistics when using a Box-M test, but leads to very different impact on the bias estimation as visible on (c). The simulation are repeated 50 times, with a similar outcome generative model as in (14), and $n = 1000$ and $m = 10000$.

Brenner, A., M. Arribas, J. Cuzick, V. Jairath, S. Stanworth, K. Ker, H. Shakur-Still, and I. Roberts (2018). Outcome measures in clinical trials of treatments for acute severe haemorrhage. *Trials* 19, 533.

Buchanan, A. L., M. G. Hudgens, S. R. Cole, K. R. Mollan, P. E. Sax, E. S. Daar, A. A. Adimora, J. J. Eron, and M. J. Mugavero (2018). Generalizing evidence from randomized trials using inverse probability of sampling weights. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 181, 1193–1209.

Chen, X., H. Hong, and D. Nekipelov (2007). Measurement error models.

Chen, X., H. Hong, and E. Tamer (2005, 02). Measurement error models with auxiliary data. *Review of Economic Studies* 72, 343–366.

Chernozhukov, V., D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. Newey, and J. Robins (2017, 06). Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal* 21.

Cinelli, C. and C. Hazlett (2020, February). Making sense of sensitivity: extending omitted variable bias. *Journal of the Royal Statistical Society Series B* 82(1), 39–67.

- Cochrane, A. L. et al. (1972). Effectiveness and efficiency: random reflections on health services, Volume 900574178. Nuffield Provincial Hospitals Trust London.
- Cole, S. R. and E. A. Stuart (2010). Generalizing evidence from randomized clinical trials to target populations: The actg 320 trial. American Journal of Epidemiology 172, 107–115.
- Colnet, B., I. Mayer, G. Chen, A. Dieng, R. Li, G. Varoquaux, J.-P. Vert, J. Josse, and S. Yang (2020). Causal inference methods for combining randomized trials and observational studies: a review.
- Cornfield, J., W. Haenszel, E. C. Hammond, A. M. Lilienfeld, M. B. Shimkin, and E. L. Wynder (1959, 01). Smoking and Lung Cancer: Recent Evidence and a Discussion of Some Questions. JNCI: Journal of the National Cancer Institute 22(1), 173–203.
- Correa, J., J. Tian, and E. Bareinboim (2018). Generalized adjustment under confounding and selection biases.
- CRASH-3 (2019). Effects of tranexamic acid on death, disability, vascular occlusive events and other morbidities in patients with acute traumatic brain injury (CRASH-3): a randomised, placebo-controlled trial. The Lancet 394(10210), 1713–1723.
- Dahabreh, I. J., S. E. Robertson, J. A. Steingrimsson, E. A. Stuart, and M. A. Hernán (2020). Extending inferences from a randomized trial to a new target population. Statistics in Medicine 39(14), 1999–2014.
- Dahabreh, I. J., S. E. Robertson, E. J. T. Tchetgen, E. A. Stuart, and M. A. Hernán (2019). Generalizing causal inferences from individuals in randomized trials to all trial-eligible individuals. Biometrics 75, 685–694.
- Dahabreh, I. J., J. M. Robins, S. J. Haneuse, and M. A. Hernán (2019). Generalizing causal inferences from randomized trials: counterfactual and graphical identification. arXiv preprint arXiv:1906.10792.
- Dahabreh, I. J., J. M. Robins, S. J.-P. A. Haneuse, I. Saeed, S. E. Robertson, E. A. Stuart, and M. A. Hernán (2019). Sensitivity analysis using bias functions for studies extending inferences from a randomized trial to a target population.
- Degtiar, I. and S. Rose (2021). A review of generalizability and transportability.
- Dewan, Y., E. Komolafe, J. Mejía-Mantilla, P. Perel, I. Roberts, and H. Shakur-Still (2012, 06). CRASH-3: Tranexamic acid for the treatment of significant traumatic brain injury: study protocol for an international randomized, double-blind, placebo-controlled trial. Trials 13, 87.
- Dong, L., S. Yang, X. Wang, D. Zeng, and J. Cai (2020). Integrative analysis of randomized clinical trials with real world evidence studies. arXiv preprint arXiv:2003.01242.
- Dorie, V., M. Harada, N. Carnegie, and J. Hill (2016, September). A flexible, interpretable framework for assessing sensitivity to unmeasured confounding. Statistics in Medicine 35(20), 3453–3470.
- Egami, N. and E. Hartman (2021). Covariate selection for generalizing experimental results: Application to large-scale development program in uganda.
- Finn, J. D. and C. M. Achilles (1990). Answers and questions about class size: A statewide experiment. American Educational Research Journal 27(3), 557–577.
- Franks, A., A. D’Amour, and A. Feller (2019, 04). Flexible sensitivity analysis for observational studies without observable implications. Journal of the American Statistical Association, 1–38.
- Friendly, M. and M. Sigal (2020). Visualizing tests for equality of covariance matrices. The American Statistician 74(2), 144–155.
- Hahn, P. R., J. S. Murray, and C. Carvalho (2019). Bayesian regression tree models for causal inference: regularization, confounding, and heterogeneous effects.
- Hartman, E., R. Grieve, R. Ramsahai, and J. S. Sekhon (2015). From sample average treatment effect to population average treatment effect on the treated: combining experimental with observational studies to estimate population treatment effects. Journal of the Royal Statistical Society: Series A (Statistics in Society) 178(3), 757–778.
- Hernan, MA Robins, J. (2020). Causal Inference: What If.

- Ichino, A., T. Nannicini, and F. Mealli (2008, 04). From temporary help jobs to permanent employment: What can we learn from matching estimators and their sensitivity? Journal of Applied Econometrics 23, 305–327.
- Imbens, G. (2003). Sensitivity to exogeneity assumptions in program evaluation. The American Economic Review.
- Imbens, G., J. Hotz, and J. Mortimer (2005). Predicting the efficacy of future training programs using past. Journal of Econometrics 125(1-2), 241–270.
- Imbens, G. W. and D. B. Rubin (2015). Causal Inference in Statistics, Social, and Biomedical Sciences. Cambridge UK: Cambridge University Press.
- Jolani, S., T. Debray, H. Koffijberg, S. van Buuren, and K. Moons (2015). Imputation of systematically missing predictors in an individual participant data meta-analysis: a generalized approach using mice. Statistics in medicine 34(11), 1841–1863.
- Kallus, N., A. M. Puli, and U. Shalit (2018). Removing hidden confounding by experimental grounding. In Advances in neural information processing systems, pp. 10888–10897.
- Kern, H., E. Stuart, J. Hill, and D. Green (2016, 01). Assessing methods for generalizing experimental impact estimates to target populations. Journal of Research on Educational Effectiveness 9, 1–25.
- Krueger, A. B. (1999). Experimental Estimates of Education Production Functions. The Quarterly Journal of Economics 114(2), 497–532.
- Lesko, C. R., A. L. Buchanan, D. Westreich, J. K. Edwards, M. G. Hudgens, and S. R. Cole (2017). Generalizing study results: a potential outcomes perspective. Epidemiology 28, 553–561.
- Lesko, C. R., S. R. Cole, H. I. Hall, D. Westreich, W. C. Miller, J. J. Eron, J. Li, M. J. Mugavero, and for the CNICS Investigators (2016, 01). The effect of antiretroviral therapy on all-cause mortality, generalized to persons diagnosed with HIV in the USA, 2009–11. International Journal of Epidemiology 45(1), 140–150.
- Li, F., A. L. Buchanan, and S. R. Cole (2021). Generalizing trial evidence to target populations in non-nested designs: Applications to aids clinical trials.
- Lunceford, J. K. and M. Davidian (2004). Stratification and weighting via the propensity score in estimation of causal treatment effects: A comparative study. In Statistics in Medicine, pp. 2937–2960.
- Mansukhani, R., L. Frimley, H. Shakur-Still, L. Sharples, and I. Roberts (2020, 07). Accuracy of time to treatment estimates in the crash-3 clinical trial: impact on the trial results. Trials 21.
- Miratrix, L. W., J. S. Sekhon, A. G. Theodoridis, and L. F. Campos (2017). Worth weighting? how to think about and use weights in survey experiments.
- Nguyen, T., B. Ackerman, I. Schmid, S. Cole, and E. Stuart (2018, 12). Sensitivity analyses for effect modifiers not observed in the target population when generalizing treatment effects from a randomized controlled trial: Assumptions, models, effect scales, data scenarios, and implementation details. PLOS ONE 13, e0208795.
- Nguyen, T. Q., C. Ebnesajjad, S. R. Cole, E. A. Stuart, et al. (2017). Sensitivity analysis for an unobserved moderator in rct-to-target-population generalization of treatment effects. The Annals of Applied Statistics 11(1), 225–247.
- Nie, X. and S. Wager (2020, 09). Quasi-oracle estimation of heterogeneous treatment effects. Biometrika 108.
- Pearl, J. (2009). Causality (2 ed.). Cambridge, UK: Cambridge University Press.
- Pearl, J. (2015). Generalizing experimental findings. Journal of Causal Inference 3(2), 259–266.
- Pearl, J. and E. Bareinboim (2011, Aug.). Transportability of causal and statistical relations: A formal approach. Proceedings of the AAAI Conference on Artificial Intelligence 25(1).
- Pearl, J. and E. Bareinboim (2014). External Validity: From Do-Calculus to Transportability Across Populations. Statistical Science 29(4), 579 – 595.

- Resche-Rigon, M., I. White, J. Bartlett, S. A. Peters, and S. Thompson (2013, 07). Multiple imputation for handling systematically missing confounders in meta-analysis of individual participant data. Statistics in medicine 32.
- Robinson, P. (1988). Root- n -consistent semiparametric regression. Econometrica 56(4), 931–54.
- Rosenbaum, P. (2005, 10). Sensitivity Analysis in Observational Studies, Volume 4.
- Ross, S. M. (1998). A First Course in Probability (Fifth ed.). Upper Saddle River, N.J.: Prentice Hall.
- Rothwell, P. M. (2005). External validity of randomised controlled trials: “to whom do the results of this trial apply?”. The Lancet 365, 82–93.
- Stuart, E. A., B. Ackerman, and D. Westreich (2018). Generalizability of randomized trial results to target populations: Design and analysis possibilities. Research on social work practice 28(5), 532–537.
- Stuart, E. A., S. R. Cole, C. P. Bradshaw, and P. J. Leaf (2011). The use of propensity scores to assess the generalizability of results from randomized trials. Journal of the Royal Statistical Society: Series A (Statistics in Society) 174, 369–386.
- Stuart, E. A. and A. Rhodes (2017). Generalizing treatment effect estimates from sample to population: A case study in the difficulties of finding sufficient data. Evaluation review 41(4), 357–388.
- Susukida, R., R. Crum, E. Stuart, C. Ebnesajjad, and R. Mojtabai (2016, 01). Assessing sample representativeness in randomized control trials: Application to the national institute of drug abuse clinical trials network. Addiction 111, n/a–n/a.
- Tipton, E. (2013). Improving generalizations from experiments using propensity score subclassification: Assumptions, properties, and contexts. Journal of Educational and Behavioral Statistics 38, 239–266.
- Veitch, V. and A. Zaveri (2020). Sense and sensitivity analysis: Simple post-hoc analysis of bias due to unobserved confounding.
- Wager, S. (2020). Stats 361: Causal inference.
- Wager, S., W. Du, J. Taylor, and R. J. Tibshirani (2016). High-dimensional regression adjustments in randomized experiments. Proceedings of the National Academy of Sciences 113(45), 12673–12678.
- Wooldridge, J. M. (2016). Introductory econometrics: A modern approach. Nelson Education.