



HAL
open science

Comparison of meta-learners for estimating multi-valued treatment heterogeneous effects

Naoufal Acharki, Ramiro Lugo, Antoine Bertoncello, Josselin Garnier

► To cite this version:

Naoufal Acharki, Ramiro Lugo, Antoine Bertoncello, Josselin Garnier. Comparison of meta-learners for estimating multi-valued treatment heterogeneous effects. 40th International Conference on Machine Learning, Jul 2023, Honolulu, United States. pp.91-132. hal-04392279

HAL Id: hal-04392279

<https://hal.science/hal-04392279>

Submitted on 13 Jan 2024

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.

Comparison of meta-learners for estimating multi-valued treatment heterogeneous effects

Naoufal Acharki^{1,2} Ramiro Lugo² Antoine Bertoncello² Josselin Garnier¹

Abstract

Conditional Average Treatment Effects (CATE) estimation is one of the main challenges in causal inference with observational data. In addition to Machine Learning based-models, nonparametric estimators called *meta-learners* have been developed to estimate the CATE with the main advantage of not restraining the estimation to a specific supervised learning method. This task becomes, however, more complicated when the treatment is not binary as some limitations of the naive extensions emerge. This paper looks into meta-learners for estimating the heterogeneous effects of multi-valued treatments. We consider different meta-learners, and we carry out a theoretical analysis of their error upper bounds as functions of important parameters such as the number of treatment levels, showing that the naive extensions do not always provide satisfactory results. We introduce and discuss meta-learners that perform well as the number of treatments increases. We empirically confirm the strengths and weaknesses of those methods with synthetic and semi-synthetic datasets.

1. Introduction

With the rapid development of Machine Learning (ML) and its efficiency in predicting outcomes, the question of counterfactual prediction “*what would happen if?*” arises. Engineers and specialists want to know how the outcome would be affected after an intervention on a parameter. It will help them personalize the parameter at efficient levels and optimize the outcome. Recently, much effort has been devoted to supervised ML to find the optimal intervention strategy. Yet, the results are not always convincing. These models cannot distinguish between correlations and causal

relationships in the data (Pearl, 2019).

Based on the Potential Outcomes theory (Neyman, 1923; Rubin, 1974), epidemiologists and statisticians have developed a set of tools that reduce the inference of causal effects to statistical inference under certain assumptions about, for example, the data-generating process. They have been successfully applied in many fields such as medicine (Alaa & van der Schaar, 2017), economics (Knaus et al., 2020b), public policy (Imai & Strauss, 2011) and marketing (Diemert et al., 2018) to infer causal effects. However, most of these studies are limited to a binary treatment setting whereas many causal questions in real-world cases are not binary. It would be helpful to give an in-depth analysis of the impact of the treatment across its possible levels instead of just considering a binary scenario where the treatment is either assigned or not. In addition, the heterogeneity of effects may provide valuable information regarding this treatment’s effectiveness and help users personalize their intervention policies and strategies.

The problem of Causal Inference beyond binary treatment settings is gaining attention from the Causal Inference community. There are two major challenges: Firstly, the lack of the ground truth due to the fundamental problem of causal inference (Holland, 1986) makes Heterogeneous Treatment Effects estimation more challenging (Alaa & van der Schaar, 2018) as standard metrics can not be used to assess performances. Secondly, binarizing the multi-valued treatment setting leads to a violation of the Stable Unit Treatment Value Assumption (SUTVA) as it violates the principle of “no hidden variations of treatment”. It may yield a bias known as position bias in recommendation systems (Chen et al., 2023; Wu et al., 2022). It happens when some units tend to have/select high treatment values and, therefore, there is a hidden variation of the treatment that is not taken into consideration when one attempt to binarize the treatment. From a statistical point of view, this bias was established by Heiler & Knaus (2022) who show that the binarization of multi-valued treatments does not disassociate the heterogeneity of the treatment from the heterogeneity of the effects of each value.

The extension of the binary setting does not seem trivial as several versions are possible and turn out to have different

¹CMAP, Ecole polytechnique, Institut Polytechnique de Paris, Palaiseau, France ²TotalEnergies One Tech, Palaiseau, France. Correspondence to: Naoufal Acharki <naoufal.acharki@gmail.com>.

behaviours. Moreover, – to the best of our knowledge – there is no result so far on the impact of the number of possible treatments on the performances of heterogeneous treatment effect (nonparametric or ML-based) estimators.

In this paper, we study the problem of nonparametric estimation of Heterogeneous Treatment Effects, also known as Conditional Average Treatment Effects (CATEs), for multi-valued treatments. We consider nonparametric estimators, also referred to as *meta-learners* (Künzel et al., 2019) or *model-agnostic algorithms* (Curth & van der Schaar, 2021a). We put our main focus on discussing the theoretical properties of meta-learners for estimating CATEs. Finally, along the lines of Curth & van der Schaar (2021a), we consider it significant to draw strengths and weaknesses theoretically and compare scenarios in which some methods would perform better than others.

Contributions. The paper considers meta-learners for multi-valued treatments. First, we generalize *meta-learners* to the multi-treatment setting for CATE estimation. We overview, in particular, Debiased Machine Learning (DML) estimators in observational studies and we establish a new version for the X-learner based on regression adjustment for multi-valued treatments. We also highlight the multi-treatments R-learner’s main drawbacks. Second, and this is the major contribution of the paper, we conduct a theoretical analysis of the proposed meta-learners for multiple treatments based on an asymptotic bias-variance analysis (see Silverman (1986) for an example of this analysis for kernel density estimation). We analyze the biases and upper bounds on errors of the M-, DR-, X-learners as well as the T- and naive X-learners. Thanks to this analysis, we can identify the effect of the number of possible treatment levels, in addition to other parameters such as the propensity score lower bound and the outcome model estimation error. This approach is different from what has been conducted in binary treatment with the minimax approach (Künzel et al., 2019; Curth & van der Schaar, 2021a; Kennedy, 2020) as it allows a direct analysis of the roles of the important parameters (e.g. the impact of the number of possible treatments K) instead of relying on the smoothness of treatment effects. Following this analysis, we present some key points about the expected performances of each meta-learner then we present a summarizing table of our findings. We note also that our analysis sheds new light on binary meta-learners’ performances as it also clarifies the influence of the sampling probability for both T- and naive X-learners.

2. Related work

Meta-learners for CATEs estimation. The recent interest in the CATE’s estimation has motivated the Causal Inference community to develop numerous algorithms and

methods (see Caron et al. (2022b) for a review). This includes a wide variety of statistical and ensemble methods (Hill, 2011; Athey & Imbens, 2016; Alaa & van der Schaar, 2017; Wager & Athey, 2018; Powers et al., 2018; Hahn et al., 2020; Caron et al., 2022a) as well as neural networks (Johansson et al., 2016; Shalit et al., 2017; Yoon et al., 2018; Shi et al., 2019) (see Dorie et al. (2019) for a review of -hybrid- ML models for causal inference). In contrast, some methods, known as *meta-learners*, are nonparametric and do not require a specific ML method. The theory of *meta-learners* was initially introduced and discussed by Künzel et al. (2019) for the CATE estimation in the binary setting with three meta-learners: the S-learner, the T-learner (which use either a *Single* or *Two* models) and the X-learner. Later, Kennedy (2020) proposes the DR-learner (Doubly-Robust) to overcome the problem of model misspecifications when estimating nuisance functions (e.g. the propensity score and outcome models). Nie & Wager (2020) present the R-learner that estimates the CATE by minimizing an orthogonalized loss function. Curth & van der Schaar (2021a) consider the PW-learner (Propensity Weighting, also known as the M-learner) and the RA-learner (Regression-Adjustment), which is an improved version of the X-learner. They show that, under some conditions, the RA-, PW-, and DR-learners can attain the oracle convergence rate.

Multiple and continuous treatments. Recently, there has also been an increased interest in causal inference with multi-valued and continuous treatments. The theoretical work of Imbens (2000); Lechner (2001); Frölich (2002); Imai & Dyk (2004) extended the potential outcome framework and the propensity score to the non-binary treatment setting, including also continuous treatments. The average dose-response estimation was considered and successfully applied in many domains in medicine and economics (Dominici et al., 2002; Flores, 2007; Kallus, 2017; Saini et al., 2019; Lin et al., 2019; Hu et al., 2020; Knaus, 2022). Additionally, Colangelo & Lee (2020) apply doubly debiased machine learning methods to dose-response modelling with continuous treatment. The CATE’s estimation, however, remains less prominently studied in the literature. Hill (2011) (briefly) and Hu et al. (2020) consider Bayesian additive regression trees for the estimation of counterfactual response and causal effects. An extension of GRF to multi-valued treatments is developed by Tibshirani et al. (2020) (which can be seen as an M-learner with multi-valued treatments). Schwab et al. (2020); Harada & Kashima (2021); Nie et al. (2021) and Kaddour et al. (2021) applied neural networks and representations learning to estimate counterfactual response curves for multiple continuous treatments (more precisely for graph-structured treatments). Zhao & Harinen (2019) naively extended binary meta-learners (X- and R-learners) without any theoretical analysis of their behaviour.

The remainder of the paper is structured as follows. Section

3 presents the CATE estimation for multi-treatments. In Section 4 we introduce CATE nonparametric estimators (meta-learners) and we discuss their consistency. In Section 5, we establish the theoretical analysis of meta-learners' error bounds and provide some discussions. We present numerical experiments and results in Section 6. Finally, we present our conclusion in Section 7.

3. Problem setting

To address the problem of causal inference under multiple treatments, we follow the Rubin-Neyman model as extended by Imbens (2000); Lechner (2001); Imai & Dyk (2004) and we consider the following statistical problem.

We suppose the existence of $Y(t)$, the real-valued counterfactual outcome that would have been observed under treatment level $t \in \mathcal{T} = \{t_0, t_1, \dots, t_K\}$. We consider $(\mathbf{X}, T, Y(t)_{t \in \mathcal{T}}) \sim \mathbb{P}$ where $\mathbf{X} = (X^{(1)}, \dots, X^{(d)}) \in \mathcal{D} \subseteq \mathbb{R}^d$ denotes a random vector of covariates and T denotes the treatment assignment random variable. We suppose finally that we observe data that has the form of an independent and identically distributed sample of n units $\mathbf{D}_{\text{obs}} = (D_{\text{obs},i})_{i=1}^n$ where $D_{\text{obs},i} = (\mathbf{X}_i, T_i, Y_{\text{obs},i})$ is distributed as $(\mathbf{X}, T, Y_{\text{obs}})$ and $Y_{\text{obs}} = Y(T)$ (consistency assumption). We define the Generalized Propensity Score (GPS) $r(t, \mathbf{x}) := \mathbb{P}(T = t | \mathbf{X} = \mathbf{x})$ (Imbens, 2000) as the generalization of the classical Propensity Score with the same balancing propriety (Rosenbaum & Rubin, 1983) to remove *selection bias* in observational studies.

We aim to infer the effect of the treatment T on the outcome Y . More precisely, we want to estimate the CATE between two levels of treatment t_k and t_0 , defined as

$$\tau_k(\mathbf{x}) = \mathbb{E}[Y(t_k) - Y(t_0) | \mathbf{X} = \mathbf{x}], \quad (1)$$

which can be interpreted as the expected treatment effect between levels t_0 (defined as the baseline treatment value) and t_k given covariates $\mathbf{X} = \mathbf{x}$. Note that other definitions and alternatives of the CATE are possible (Kaddour et al., 2021).

Unfortunately, it is impossible to infer this effect directly. We observe only one potential outcome corresponding to the treatment T (i.e. the real outcome) for every unit. All other potential outcomes are missing (inherently unobservable). Consequently, to identify the causal effects from the observed sample data, we shall consider the following assumptions (Assumption 3.1 is unfortunately untestable).

Assumption 3.1 (Unconfoundedness). The treatment mechanism is unconfounded given the observed covariates $Y(t) \perp\!\!\!\perp \mathbf{1}\{T = t\} \mid \mathbf{X}$ for all $t \in \mathcal{T}$.

Assumption 3.2 (Overlap). The probability of receiving the treatment given the observed covariates is positive i.e. there

exists $r_{\min} > 0$ such that $r_{\min} \leq \mathbb{P}(T = t | \mathbf{X} = \mathbf{x})$ for all $t \in \mathcal{T}$ and $\mathbf{x} \in \mathcal{D}$.

With the previous assumptions, the expected potential outcome satisfies $\mathbb{E}(Y(t) | \mathbf{X} = \mathbf{x}) = \mathbb{E}(Y_{\text{obs}} | T = t, \mathbf{X} = \mathbf{x})$ and the CATE can be estimated.

The problem of the CATE estimation can be seen as a non-parametric estimation problem. We tackle it by generalizing the notion of meta-learners to derive consistent estimators. This task can be achieved either by modelling the CATE directly in one or two steps: by decomposing it into regularized regression problems or by addressing a minimization problem with respect to an appropriate loss function. Moreover, all considered meta-learners below, except the R-learner, have the advantage of supporting any supervised regression method (e.g. random forest, gradient boosting methods, neural networks).

4. Generalizing meta-learners to multi-valued treatments

In the following, we adopt a similar taxonomy of CATEs estimators as Curth & van der Schaar (2021a) and Knaus et al. (2020a). Namely, direct plug-in (one-step) meta-learners, pseudo-outcome (two-step) meta-learners and Neyman-Orthogonality-based learners (R-learner).

4.1. Direct plug-in meta-learners

This subsection presents direct plug-in meta-learners, also known as *one-step* learners. They estimate the CATE in (1) by targeting it directly from \mathbf{D}_{obs} . They are the naive extension of the T- and S-learners in the binary case.

T-learner with multiple treatments. T-learner is a naive approach to estimating CATEs. It consists on estimating the *two* conditional response surfaces $\mu_t(\mathbf{x}) = \mathbb{E}(Y(t) | \mathbf{X} = \mathbf{x})$ using $\mathbf{S}_t = \{i, T_i = t\}$ for $t \in \{t_k, t_0\}$, then estimates the CATE as $\hat{\tau}_k^{(T)}(\mathbf{x}) := \hat{\mu}_{t_k}(\mathbf{x}) - \hat{\mu}_{t_0}(\mathbf{x})$.

The T-learning approach does not account for the interaction between T and Y and creates different models for different treatments. Still, it may suffer from selection bias (Curth & van der Schaar, 2021b), i.e. when the outcome models μ_t are estimated with respect to the wrong distribution when sampling $(D_{\text{obs},i})_{i \in \mathbf{S}_t}$. Therefore, $\hat{\mu}_t$ should be estimated by minimizing the expected squared error on the nominal *weighted* distribution using Importance Sampling (Hassanpour & Greiner, 2019); see Appendix A for details.

S-learner with multiple treatments. The S-learner in multi-valued treatments uses the identification of the CATE and considers a *single* model $\mu(t, \mathbf{x}) = \mathbb{E}(Y_{\text{obs}} | T = t, \mathbf{X} = \mathbf{x})$. μ is estimated using the whole dataset

\mathbf{D}_{obs} . The CATE can be computed therefore as $\hat{\tau}_k^{(S)}(\mathbf{x}) := \hat{\mu}(\mathbf{x}, t_k) - \hat{\mu}(\mathbf{x}, t_0)$.

Including the treatment T as a feature and sharing information between covariates \mathbf{X} and T may provide better predictions. However, this advantage is conditioned by the ability of the base learner to capture and distinguish contributions of both \mathbf{X} and T on Y_{obs} as we will see in Section 5. Note that the S-learner may also suffer from confounding and regularization biases (Chernozhukov et al., 2018; Hahn et al., 2020) when estimating $\hat{\mu}$.

4.2. Pseudo-outcome meta-learners

To overcome *selection bias*, a usual alternative is to consider specific representations of the observed outcome Y_{obs} , called *pseudo-outcome* Z_k . They incorporate *nuisance components* that generally include valuable information such as the dependence between covariates \mathbf{X} and T (i.e. the GPS r) and the occurrence of a particular treatment assignment. Under the *well-specification* of nuisance components, regressing Z_k on \mathbf{X} produces a *consistent* estimator i.e. $\mathbb{E}(Z_k | \mathbf{X} = \mathbf{x}) = \tau_k(\mathbf{x})$ while keeping the same sample size as \mathbf{D}_{obs} . In the following, we say that an estimator ($\hat{\mu}$ or \hat{r}) is *well-specified* if it is based on a well-specified statistical model, that is, the class of distributions assumed for modelling contains the unknown probability distribution from which the sample used for estimation is drawn.

M-learner with multiple treatments. The *M-learner* (Horvitz & Thompson, 1952), where M refers to the *modified* learned pseudo-outcome in the algorithm, is inspired from the Inverse Propensity Weighting (IPW). It is defined in the multi-valued setting, for $k = 1, \dots, K$, as the regression of Z_k^M such that

$$Z_k^M = \frac{\mathbf{1}\{T = t_k\}}{\hat{r}(t_k, \mathbf{X})} Y_{\text{obs}} - \frac{\mathbf{1}\{T = t_0\}}{\hat{r}(t_0, \mathbf{X})} Y_{\text{obs}},$$

where \hat{r} is an estimator of the GPS r .

DR-learner with multiple treatments. The *Doubly Robust* (DR) method (Robins et al., 1994; Kennedy et al., 2017; Kennedy, 2020) is helpful in overcoming the problem of the model's misspecification by estimating two components, the outcome model μ . and the GPS r , instead of relying on the correctness of one (and the only) parameter. If $\hat{\mu}$ and \hat{r} denote arbitrary estimators of the outcome μ and the GPS r (we assume that \hat{r} satisfies Assumption 3.2), then the DR-learner regresses $Z_{\hat{\mu}, \hat{r}, k}^{DR}$ such that:

$$\begin{aligned} Z_{\hat{\mu}, \hat{r}, k}^{DR} &= \frac{Y_{\text{obs}} - \hat{\mu}_T(\mathbf{X})}{\hat{r}(t_k, \mathbf{X})} \mathbf{1}\{T = t_k\} + \hat{\mu}_{t_k}(\mathbf{X}) \\ &\quad - \frac{Y_{\text{obs}} - \hat{\mu}_T(\mathbf{X})}{\hat{r}(t_0, \mathbf{X})} \mathbf{1}\{T = t_0\} - \hat{\mu}_{t_0}(\mathbf{X}). \end{aligned}$$

X-learner with multiple treatments. The X-learner (Künzel et al., 2019), also known as *Regression-Adjustment* (RA)-learning in a version developed by Curth & van der Schaar (2021a), has been proposed as an alternative to T-learning in the case where one treatment group is over-represented. The idea consists of a *cross* procedure of estimation between observations Y_{obs} and outcome models when one of the treatments occurs. For $k = 1, \dots, K$, we define the *Regression-Adjustment* pseudo-outcome Z_k^X as

$$\begin{aligned} Z_k^X &= \mathbf{1}\{T = t_k\} (Y_{\text{obs}} - \hat{\mu}_{t_0}(\mathbf{X})) + \sum_{l \neq k} \mathbf{1}\{T = t_l\} \times \\ &\quad (\hat{\mu}_{t_k}(\mathbf{X}) - Y_{\text{obs}}) + \sum_{l \neq k} \mathbf{1}\{T = t_l\} (\hat{\mu}_{t_l}(\mathbf{X}) - \hat{\mu}_{t_0}(\mathbf{X})). \end{aligned}$$

For comparison purposes, we consider also the naive extension of the binary X-learner (Künzel et al., 2019) to multi-treatments as proposed by Zhao & Harinen (2019). This extension considers two random variables $D^{(k)} := Y(t_k) - \hat{\mu}_{t_0}(\mathbf{X})$ and $D^{(0)} := \hat{\mu}_{t_k}(\mathbf{X}) - Y(t_0)$ where $\hat{\mu}_{t_k}$ and $\hat{\mu}_{t_0}$ are trained on the samples \mathbf{S}_{t_k} and \mathbf{S}_{t_0} . Then it regresses $(D_i^{(k)})_{i \in \mathbf{S}_{t_k}}$ and $(D_i^{(0)})_{i \in \mathbf{S}_{t_0}}$ on \mathbf{X} to obtain $\hat{\tau}^{(k)}$ and $\hat{\tau}^{(0)}$, and estimates the CATE as:

$$\begin{aligned} \hat{\tau}_k^{(X, \text{nv})}(\mathbf{x}) &:= \frac{\hat{r}(t_k, \mathbf{x})}{\hat{r}(t_k, \mathbf{x}) + \hat{r}(t_0, \mathbf{x})} \hat{\tau}^{(k)}(\mathbf{x}) + \\ &\quad \frac{\hat{r}(t_0, \mathbf{x})}{\hat{r}(t_k, \mathbf{x}) + \hat{r}(t_0, \mathbf{x})} \hat{\tau}^{(0)}(\mathbf{x}). \end{aligned}$$

We note that the consistency of the M- and DR-learners is already established in the literature (Knaus, 2022). We show it also for the extended X-learner in Appendix A. We note also that there are three main approaches possible to learn nuisance components (r and μ) and then estimate the τ_k , namely, Full-Sample, Sample-Split and Cross-Fit methods (Okasa, 2022). This paper does not discuss estimation procedures and adopts the Full-Sample strategy.

4.3. Neyman-Orthogonality based learner: R-learner

The R-learner is based mainly on the Robinson (1988) decomposition and was proposed by Nie & Wager (2020) to provide a flexible CATE estimator avoiding regularization bias. It states that the potential outcome error $\epsilon = Y_{\text{obs}} - \mu_T(\mathbf{X})$ satisfies $\mathbb{E}(\epsilon | T, \mathbf{X}) = 0$ and

$$\epsilon = Y_{\text{obs}} - m(\mathbf{X}) - \sum_{k=1}^K (\mathbf{1}\{T = t_k\} - r(t_k, \mathbf{X})) \tau_k(\mathbf{X}),$$

where $m(\mathbf{x}) = \mathbb{E}(Y_{\text{obs}} | \mathbf{X} = \mathbf{x})$ is the observed outcome model. Therefore, considering the mean squared error of ϵ (the generalized R-loss function) and minimizing it estimate

K models $\{\widehat{\tau}_k^{(R)}\}_{k=1}^K$ simultaneously such that

$$\{\widehat{\tau}_k^{(R)}\}_{k=1}^K := \operatorname{argmin}_{\{\bar{\tau}_k\} \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n \left[(Y_{\text{obs},i} - \widehat{m}(\mathbf{X}_i)) - \sum_{k=1}^K (\mathbf{1}\{T_i = t_k\} - \widehat{r}(t_k, \mathbf{X}_i)) \bar{\tau}_k(\mathbf{X}_i) \right]^2,$$

where \widehat{m} (respectively, \widehat{r}) is an estimator of m (respectively, r) and \mathcal{F} is the space of candidate models $[\{\bar{\tau}_k\}_{k=1}^K]$.

The generalized R-learner suffers from two major limitations: First, it cannot be written as *weighted* supervised learning problem with a specific pseudo-outcome. Only parametric families \mathcal{F} can be considered in the multi-treatment regime. The second drawback is the non-identifiability of the generalized R-loss described previously in (4.3) without a regularization. Indeed, this problem does not have a unique solution (see Appendix A for details) and leads thus to poor estimation performance. This point is also shown recently by Zhang et al. (2022) for continuous treatments and our numerical results in Appendix D also confirm that the R-learner fails to estimate CATEs $(\tau_k)_{k=1}^K$.

5. Theoretical analysis of the error upper bound

In this section, we analyze the error's upper bounds for different meta-learners. The theoretical analysis will be carried out under the following framework:

Assumption 5.1. We assume that (T, \mathbf{X}) satisfies the overlap assumption 3.2 and that, for $t \in \mathcal{T}$, the outcome $Y(t)$ is generated from a function $f : \mathbb{R} \times \mathbb{R}^d \rightarrow \mathbb{R}$ such that

$$Y(t) = f(t, \mathbf{X}) + \varepsilon(t), \quad (2)$$

where $\varepsilon(t)$ are i.i.d. Gaussian $\mathcal{N}(0, \sigma^2)$ and independent of (T, \mathbf{X}) .

Assumption 5.2. We assume the existence of $\beta_t \in \mathbb{R}^p$ such that, for all $t \in \mathcal{T}$ and $\mathbf{x} \in \mathcal{D}$

$$f(t, \mathbf{x}) = \sum_{j=0}^{p-1} \beta_{t,j} f_j(\mathbf{x}) = \mathbf{f}(\mathbf{x})^\top \beta_t,$$

where f_j are some bounded predefined basis functions.

The assumption of a product effect is reasonable. One can show the universality of this representation in the Reproducing Kernel Hilbert Space (RKHS) (Proposition 1 of Kaddour et al. (2021)) if we allow the dimension p to be large enough.

Under these two assumptions, the CATE τ_k can be written as:

$$\tau_k(\mathbf{x}) = \sum_{j=0}^{p-1} \beta_{k,j}^* f_j(\mathbf{x}) = \mathbf{f}(\mathbf{x})^\top \beta_k^*,$$

where $\beta_k^* = (\beta_{k,j}^*)_{j=0}^{p-1} = \beta_{t_k} - \beta_{t_0} \in \mathbb{R}^p$.

From a theoretical point of view, the S-learner corresponds to the naive Ordinary Least Square (OLS) $\widehat{\beta}_k$ of β_k^* . The statistical task of the CATE's estimation holds immediately. However, we cannot properly analyse the base-learner's ability to learn β_k^* under confounding effects.

Theorem 5.3. Under Assumptions (5.1-5.2), the OLS estimators $\widehat{\beta}_k^*$ of the T-learner and the naive X-learner are unbiased and have an asymptotic covariance matrix $\mathbb{V}(\widehat{\beta}_k^*) = \mathbf{C}/n$, whose terms \mathbf{C}_{ij} are bounded by:

$$\mathcal{E}^T = \mathcal{E}^{X,nv} = \mathcal{O}\left(\frac{1}{\rho(t_k)} + \frac{1}{\rho(t_0)}\right),$$

where $\mathbb{P}(T = t) = \rho(t) > 0$ for all $t \in \mathcal{T}$.

We consider now pseudo-outcome meta-learners (M-, DR- and X-learners). When investigating the pseudo-outcomes Z_k , one can see that, for $k = 1, \dots, K$, these pseudo-outcomes are linear with respect to Y_{obs} i.e.

$$Z_k = A_{t_k}(T, \mathbf{X}) Y_{\text{obs}} + B_{t_k}(T, \mathbf{X}),$$

where $A_{t_k}(T, \mathbf{X})$ and $B_{t_k}(T, \mathbf{X})$ are given for each pseudo-outcome meta-learner.

Theorem 5.4. Under Assumptions (5.1-5.2), the OLS estimator $\widehat{\beta}_k^*$ is unbiased if the nuisance parameters ($\widehat{\mu}$ and \widehat{r}) are well-specified, and has an asymptotic covariance matrix $\mathbb{V}(\widehat{\beta}_k^*) = \mathbf{C}/n$, whose terms, for all $\epsilon > 0$, are bounded by:

$$\mathcal{E}^M = \mathcal{O}\left(\frac{1}{r_{\min}^{1+\epsilon}}\right) \text{ for the M-learner,}$$

$$\mathcal{E}^{DR} = \mathcal{O}\left(\frac{\operatorname{err}(\widehat{\mu}_{t_k}) + \operatorname{err}(\widehat{\mu}_{t_0})}{r_{\min}^{1+\epsilon}}\right) \text{ for the DR-learner,}$$

$$\mathcal{E}^X = \mathcal{O}\left(K^2 \sum_{l \neq k} \operatorname{err}(\widehat{\mu}_{t_l})\right) \text{ for the X-learner,}$$

where $\operatorname{err}(\widehat{\mu}_t) = \mathbb{E}_{\mathbf{X}}[(f(t, \mathbf{X}) - \widehat{\mu}_t(\mathbf{X}))^2]$ is the expected mean squared error of $\widehat{\mu}_t$.

Sketch of the proofs of Theorems 5.3 and 5.4. Both proofs are similar and are structured in 3 steps: 1) Express the OLS estimator $\widehat{\beta}_k^*$ as a function of the true β_k^* ; 2) Apply the multivariate Central Limit Theorem and the Slutsky theorem (or the Delta method in the general case with a biased $\widehat{\beta}_k^*$); 3) Bound the asymptotic covariance matrix terms. The full proofs can be found in Appendix B.

Remark 5.5. For $K = 2$, one recovers the error's upper bounds of pseudo-outcome meta-learners and the important parameters as shown by Curth & van der Schaar (2021a). However, the influence of the sampling probability ρ for the T- and naive X-learners is original.

Finally, regarding the binarized (Kaddour et al., 2021) or generalized R-learners, they cannot be expressed as a *model-agnostic* regression problem but rather as a minimization problem of a loss function (generalized R-loss or binarized R-loss). Therefore, the actual theoretical analysis of the error’s upper bound cannot be conducted similarly.

Using Theorems 5.3 and 5.4, we establish the following discussions (see Table 1 for a summary).

5.1. General comments and insights

About Theorem 5.3. The significant implication of Theorem 5.3 is the ability to anticipate the performances of the T-learner with respect to the treatment distribution when the probability $\rho(t)$ of sampling the treatment value t is small. The performances of the T- and naive X-learners become poor when the number K of treatments increases because these learners imply learning in small samples \mathbf{S}_t . Moreover, the naive X-learner also does not seem to really differ from the T-learner in any meaningful way (see Appendix B.2). We note that this result is original and cannot be obtained by the minimax approach.

About Theorem 5.4. From this theorem, one can establish the relationship between the number of observations n and the number of treatment values K for a given error bound. The same theorem is also useful for an RCT setting (r_{\min} is known) and when the error of the outcome model is small for all treatment levels except for one treatment value t .

Moreover, although pseudo-outcome meta-learners have the advantage of learning on the full sample, they, unfortunately, may lead to high error and poor performance in how nuisance components intervene. On the one hand, the GPS is in the M- and DR-learners denominators. The error bound is likely to be high when there is a lack of overlap or when the number of treatments K increases (we have necessarily $r_{\min} \leq 1/K$). On the other hand, the upper bounds of the X- and DR-learners depend on the quality of the estimated $\hat{\mu}$. One can expect that the more precise the outcome models are, the lower the error is.

5.2. Specific comments for each meta-learner

M-learner. Without surprise, the M-learner is very sensitive to the estimated GPS \hat{r} and suffers from high variance. This is even more critical as the number K of treatments increases.

DR-learner. The error’s estimation term $\text{err}(\hat{\mu}_t)$ of $\hat{\mu}_t$ and $\hat{\mu}_{t_0}$ in the numerator dramatically improves the DR-learner’s performances.

X-learner. The X-learner incorporates only $\hat{\mu}$, and does not imply the GPS r . Thus, the X-learner is likely to have the smallest error compared to other meta-learners when the

overlap assumption is not sufficiently respected. However, the consistency of $(\hat{\mu}_t)_{t \neq t_k}$ is required to estimate CATEs correctly.

M-learner vs DR-learner. If the outcome models $\hat{\mu}_t$ and $\hat{\mu}_{t_0}$ are well-specified, the error’s upper bound is expected to be smaller for the DR-learner than for the M-learner. However, if they are misspecified (but the propensity score is well-specified), then there is no guarantee that the DR-learner would perform better than M-learner. It may perform even worse, as we will see in Appendix D (Table 8).

M-learner vs X-learner. The X-learner is likely to have a lower error upper bound if the expected squared error $\text{err}(\hat{\mu}_t)$ is small and if some conditions on K and r_{\min} hold.

X-learner vs DR-learner. Analytically, it is difficult to anticipate which meta-learner would perform better. This depends mainly on $\text{err}(\hat{\mu}_t)$, K and r_{\min} , whom, in some cases, make the X-learner have less error than the DR-learner and the opposite in other cases. Still, our numerical results in Appendix D show that the X-learner outperforms the DR-learner in most cases.

X-learner vs Naive X-learner. We cannot theoretically compare these meta-learners without knowledge about the distribution of T . However, we can see numerically that the X-learner clearly outperforms the naive X-learner.

We conclude this subsection with a small discussion about the generalized and the binarized R-learners (see Appendix C). Indeed, the binary R-loss function may be solved separately in a low sample (instead of \mathbf{D}_{obs}) and one can obtain a unique solution, unlike the generalized R-loss. However, the optimization procedure is two-stage iterative and computationally heavy. We do not consider it in Section 6.

Table 1. Summary table of multi-treatments meta-learners.

Meta-learner	Advantages	Disadvantages
T-learner (nv X-learner)	Simple approach	Selection bias Low samples
S-learner	Simple approach	Confounding effects Regularization bias
M-learner	Consistency	High variance
DR-learner	Consistency Doubly Robust	Possibly high variance
X-learner	Consistency Low variance	Non-intuitive
R-learner	Interaction effects	Non-identifiability
Bin R-learner	Identifiability	Computational cost

In Table 1, *Bin R-learner* refers to the binarized R-learner (Kaddour et al., 2021) and *nv X-learner* refers to the naive extension of the X-learner described in subsection 4.2. *Possibly high variance* refers to the case where the variance can be significantly high due to the lack of overlap caused by the inverse propensity weighting in some DGPs. *Selection bias* refers to the bias that occurs when sampling \mathbf{S}_t and comparing units directly, as we describe in Proposition 4.1 in Appendix A. The confounding effects represent the statistical and intrinsic dependence between the treatment T and the covariates \mathbf{X} , which prevent some base-learner (e.g. random forest) from distinguishing and disassociating the treatment T and the covariates \mathbf{X} . Finally, *low samples* refer to cases where the samples $\mathbf{S}_t = \{i, T_i = t\}$ become small for some treatment levels and the model has then few observations to learn from.

5.3. Practical recommendations for selecting a meta-learner

In this subsection, based on previous insights and our numerical findings (see Section 6 for more detail), we provide some instructions and recommendations for selecting meta-learners given a dataset \mathbf{D}_{obs} :

- Pseudo-outcomes meta-learners and the S-learner are preferred in low sample regimes.
- It is recommended to examine the distribution of T to understand at which values of t the T-learner may fail to learn CATEs.
- The S-learner remains a simple and reasonable choice, especially when $K \geq 10$.
- The X-learner is useful to learn the CATE τ_t when there is not enough information about a specific treatment value t .
- If one is interested in quantifying uncertainties, then the DR-learner is recommended as it allows having a consistent estimation of the CATE.
- After estimating the GPS r , one needs to check if there is a possible lack of overlap. If a lack of overlap is detected, one should avoid considering the M- and DR-learners.

6. Numerical experiments

In this section, we assess the performances of the different meta-learners, additional numerical results are shown in Appendix D.

In synthetic or semi-synthetic examples where the CATEs are known, the error in estimation is given by **mPEHE** (respectively, **sdPEHE**), the mean (respectively, the standard deviation) of the Precision in Estimation of Heterogeneous

Effect (PEHE) (Hill, 2011; Shalit et al., 2017) defined as the mean squared error in the estimation of the treatment effect $\hat{\tau}_k$, over all possible treatment levels t_k for $k = 1, \dots, K$:

$$\mathbf{mPEHE} = \sqrt{\frac{1}{K} \sum_{k=1}^K PEHE(\hat{\tau}_k)^2},$$

where $PEHE(\hat{\tau}_k)^2 = \frac{1}{n} \sum_{i=1}^n (\hat{\tau}_k(\mathbf{X}_i) - \tau_k(\mathbf{X}_i))^2$, and,

$$\mathbf{sdPEHE} = \sqrt[4]{\frac{1}{K-1} \sum_{k=1}^K \left(PEHE(\hat{\tau}_k)^2 - \mathbf{mPEHE}^2 \right)^2}$$

Those metrics will be used to compare meta-learners under different scenarios (sample size n , number of possible treatments K , the correctness of nuisance parameters). We do not consider here model-fitting of base-learners. All hyperparameters (e.g. the number of trees, depth etc.) are fixed to their default values during all experiments. In addition, we do not consider Neural Networks because it would require choosing between at least five possible architectures (Curth & van der Schaar, 2021a) to define tasks of learning nuisance components and estimating CATEs τ_k while the main focus of the paper is on the choice of the meta-learner.

6.1. Synthetic datasets: analytical functions in randomized and non-randomized studies

In this subsection, we begin by empirically evaluating the performances of meta-learners when the treatment T is taking $K+1 = 10$ possible values in $[0, 1]$ on a linear outcome:

$$Y(t) \text{ of the form (2) with } f(t, X) = (1+t)X, \text{ and } X \sim \mathcal{U}[0, 1],$$

in Randomized Controlled Trials (RCT) setting where the T and X are independent. Second, we evaluate meta-learners on the hazard rate outcome:

$$Y(t) \text{ of the form (2) with } f(t, \mathbf{X}) = t + \|\mathbf{X}\| \exp(-t\|\mathbf{X}\|) \text{ and } \mathbf{X} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_5)$$

in a non-randomized setting as will be described below.

To simulate observational data, instead of removing some rows, we create a selection bias in the data by selecting preferentially only observations with specific characteristics (see subsection D.1 in Appendix D). This strategy comes in line with the findings and recommendations of Curth et al. (2021) about creating a biased sub-sample and evaluating CATEs' estimators.

The GPS is estimated using the XGBoost model, and the outcome models μ_t are either estimated by the T- or S-learning approaches. In Tables 2 and 3 and Appendix D,

Table 2. **mPEHE** and **sdPEHE** for XGBoost and RandomForest; linear model in a RCT setting with $n = 2000$ units.

Meta-learner	XGBoost	RandomForest
T-Learner	0.065 (0.019)	0.041 (0.016)
S-Learner	0.033 (0.018)	0.032 (0.028)
NvX-Learner	0.060 (0.019)	0.037 (0.016)
M-Learner	1.25 (0.610)	1.22 (0.621)
DR-Learner	0.068 (0.019) – 0.063 (0.020)	0.068 (0.018) – 0.068 (0.018)
X-Learner	0.063 (0.020) – 0.033 (0.017)	0.045 (0.016) – 0.061 (0.040)
RLin-Learner	0.135 (0.130)	0.137 (0.128)

For the DR- and X-learners: μ_t are estimated by T- (left value) or S- (right value). The bold font indicates the best meta-learner (row) per base-learner (column).

 Table 3. **mPEHE** and **sdPEHE** for XGBoost and RandomForest. Hazard rate model in an observational setting with $n = 10000$ units.

Meta-learner	XGboost	RandomForest
T-Learner	0.183 (0.039)	0.286 (0.155)
RegT-Learner	0.176 (0.044)	0.286 (0.155)
S-Learner	0.176 (0.056)	0.306 (0.153)
NvX-Learner	0.190 (0.096)	0.336 (0.200)
M-Learner	1.61 (0.505)	1.58 (0.472)
DR-Learner	0.168 (0.045) - 0.178 (0.048)	0.304 (0.158) – 0.322 (0.162)
X-Learner	0.167 (0.053) – 0.172 (0.057)	0.302 (0.169) – 0.332 (0.167)
RLin-Learner	0.231 (0.081)	0.186 (0.123)

For the DR- and X-learners: μ_t are estimated by T- (left value) or S- (right value). The bold font indicates the best meta-learner (row) per base-learner (column).

RLin-learner denotes the generalized R-learner with linear regression models in (4.3) with $p = 2$. For each meta-learner (row) and base-learner (column), we indicate the **mPEHE** followed by the **sdPEHE** between brackets.

In Tables 2 and 3, we find that, as expected, the M-learner predicts poorly. The T- and naive X-learners give better predictions for Random Forest, whereas the S-learner gives better results for XGBoost. Regularizing T-learner (RegT-Learner) against selection bias increases its performance. The X- and DR-learners improve the predictions of the S-learner for XGBoost, but this improvement is not always observable for Random Forests. Unfortunately, the actual results (and also additional numerical experiments in Appendix D) confirm the claim: The RLin-learner generally fails to identify CATEs correctly.

Despite these satisfying results, we highlight the problem of over-fitted gradient boosting models and Random Forest by comparing them with the linear model in Appendix D. This problem should be taken further while estimating CATEs. We think that using out-sample prediction supervised models might solve this problem.

We consider now the effect of increasing K on the hazard rate function with XGBoost. The results are shown in Figure

1 in Appendix D.3. On the one hand, the performances of the T- and the naive X-learners become compromised. The regularized T-learner suffers also from the same issue (with a linear effect with respect to K), which can be also quantified on the DR- and X-learners when applying the regularized T-learning. For $K \geq 20$, the T- and the naive X-learners perform better than the previous meta-learners. Regarding the M-learner, it has poor performances in all cases as can be expected with Theorem 5.4. On the other hand, the S-learner stabilises once K is large enough and, consequently, the DR- and X-learners also stabilize when applying the S-learning for large K . Therefore, we recommend the S-learner’s estimated potential outcome model when $K \geq 10$ for pseudo-outcome meta-learners. To conclude, two-step meta-learners are robust. In particular, the X-learner improves the quality of plug-in meta-learners; when it does not, the differences are very small.

6.2. Semi-synthetic dataset: estimating heterogeneous treatment effects on a non-randomized dataset.

In this subsection, we consider a multistage fracturing Enhanced Geothermal System (EGS) (Olasolo et al., 2016). We assume that the heat extraction performance satisfies the physical model: $Q_{well}(\ell_L) = Q_{fracture} \times \ell_L / d \times \eta_d$, where

Table 4. **mPEHE** and **sdPEHE** for XGBoost and RandomForest. Heat Extraction model in an observational setting.

Meta-learner	XGBoost	RandomForest
T-learner	0.172 (0.052)	0.157 (0.067)
RegT-Learner	0.156 (0.042)	0.154 (0.067)
S-learner	0.101 (0.040)	0.218 (0.129)
NvX-Learner	0.102 (0.042)	0.143 (0.067)
M-learner	1.04 (0.423)	0.898 (0.417)
DR-learner	0.148 (0.042) – 0.097 (0.029)	0.164 (0.068) – 0.203 (0.108)
X-learner	0.142 (0.041) – 0.094 (0.034)	0.173 (0.077) – 0.211 (0.120)
RLin-learner	0.357 (0.274)	0.362 (0.278)

For the DR- and X-learners: μ_t are estimated by T- (left value) or S- (right value). The bold font indicates the best meta-learner (row) per base-learner (column).

$Q_{fracture}$ is the *unknown* heat extraction performance from a single fracture, that can be generated using a numerical model with eight input parameters including reservoir characteristics and fracture design. ℓ_L is the Lateral Length of the well, d is the average spacing between two fractures and η_d is the stage efficiency penalizing the individual contribution when fractures are close to each other. We refer to Appendix E for a detailed description of the model and the semi-synthetic dataset.

We consider the Lateral Length as treatment T with $K+1 = 13$ possible values and the covariates $\mathbf{X} \in \mathbb{R}^{11}$ are the remaining variables. We also consider a logarithmic transformation of the heat performance for a meaningful **mPEHE**, and we normalize the treatment T . Following the *preferential selection*, we sample $n = 10000$ units such that wells with high lateral length are likely to have larger fractures and vice versa. The GPS is estimated using gradient boosting models. Table 4 resumes the **mPEHE** and **sdPEHE** in brackets for different meta-learners. Most findings of subsection 6.1 remain valid: XGBoost model is generally a better choice than Random Forests (except for T-learning); the X-learner, followed by the DR-learner, outperforms all other learners.

7. Conclusion

We have investigated heterogeneous treatment effects estimation with multi-valued treatments. In addition to standard plug-in meta-learners, we have considered representations to build pseudo-outcome meta-learners, and we have proposed the generalized Robinson decomposition to build the R-learner. Using the bias-variance analysis, we have conducted an in-depth analysis of the error’s upper bounds of pseudo-outcomes meta-learners. Thanks to this analysis, we could discuss the advantages and limitations of each pseudo-outcome meta-learner. In particular, we have identified the impacts of the number of treatment levels and the lower bound r_{\min} on the M-, DR and X-learners. Through

synthetic and semi-synthetic industrial datasets, we have illustrated the performances of different meta-learners in a non-randomized case where some covariates are confounded with the treatment. We have demonstrated the ability of the X-learner to reconstruct the ground truth model. We have also highlighted how the choice of base-learner can affect the quality of CATEs estimation.

8. Software and Data

The code and the semi-synthetic dataset in subsection 6.2 are available at <https://github.com/nacharki/multipleT-MetaLearners>.

9. Acknowledgements

The authors thank Marianne Clausel, Alessandro Leite, Audrey Poinsot, Georges Oppenheim and the anonymous reviewers for helpful feedback and discussions. This work was supported by TotalEnergies and the French National Agency for Research and Technology (ANRT) (Grant n° 2019/0714).

References

- Alaa, A. and van der Schaar, M. Limits of estimating heterogeneous treatment effects: Guidelines for practical algorithm design. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *PMLR*, pp. 129–138. PMLR, 10–15 Jul 2018.
- Alaa, A. M. and van der Schaar, M. Bayesian inference of individualized treatment effects using multi-task gaussian processes. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS’17, pp. 3427–3435. Curran Associates Inc., 2017. ISBN 9781510860964.

- Athey, S. and Imbens, G. Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, 113(27):7353–7360, 2016. doi: 10.1073/pnas.1510489113.
- Caron, A., Baio, G., and Manolopoulou, I. Shrinkage bayesian causal forests for heterogeneous treatment effects estimation. *Journal of Computational and Graphical Statistics*, pp. 1–13, 2022a.
- Caron, A., Baio, G., and Manolopoulou, I. Estimating individual treatment effects using non-parametric regression models: A review. *Journal of the Royal Statistical Society Series A*, 185(3):1115–1149, July 2022b. doi: 10.1111/rssa.12824.
- Chen, J., Dong, H., Wang, X., Feng, F., Wang, M., and He, X. Bias and debias in recommender system: A survey and future directions. *ACM Transactions on Information Systems*, 41(3):1–39, 2023.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., and Robins, J. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 01 2018. ISSN 1368-4221. doi: 10.1111/ectj.12097.
- Colangelo, K. and Lee, Y.-Y. Double debiased machine learning nonparametric inference with continuous treatments. *arXiv preprint arXiv:2004.03036*, 2020.
- Curth, A. and van der Schaar, M. Nonparametric estimation of heterogeneous treatment effects: From theory to learning algorithms. In Banerjee, A. and Fukumizu, K. (eds.), *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *PMLR*, pp. 1810–1818. PMLR, 13–15 Apr 2021a.
- Curth, A. and van der Schaar, M. On inductive biases for heterogeneous treatment effect estimation. In Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems*. PMLR, 2021b.
- Curth, A., Svensson, D., Weatherall, J., and van der Schaar, M. Really doing great at estimating CATE? a critical look at ML benchmarking practices in treatment effect estimation. In *35th Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- Diemert, E., Betlei, A., Renaudin, C., and Amini, M.-R. A large scale benchmark for uplift modeling. In *Proceedings of the AdKDD and TargetAd Workshop, KDD, London, United Kingdom, August, 20, 2018*. ACM, 2018.
- Dominici, F., Daniels, M., Zeger, S. L., and Samet, J. M. Air pollution and mortality: estimating regional and national dose-response relationships. *Journal of the American Statistical Association*, 97(457):100–111, 2002.
- Dorie, V., Hill, J., Shalit, U., Scott, M., and Cervone, D. Automated versus do-it-yourself methods for causal inference: Lessons learned from a data analysis competition. *Statistical Science*, 34(1):43–68, 2019.
- Flores, C. A. Estimation of Dose-Response Functions and Optimal Doses with a Continuous Treatment. Technical Report 0707, University of Miami, Department of Economics, November 2007.
- Frölich, M. Programme evaluation with multiple treatments. *Wiley-Blackwell: Journal of Economic Surveys*, 2002.
- Hahn, P., Murray, J., and Carvalho, C. Bayesian regression tree models for causal inference: Regularization, confounding, and heterogeneous effects (with discussion). *Bayesian Analysis*, 15(3):965–1056, 2020. ISSN 1936-0975. doi: 10.1214/19-BA1195.
- Harada, S. and Kashima, H. Graphite: Estimating individual effects of graph-structured treatments. In *Proceedings of the 30th ACM International Conference on Information and Knowledge Management*, pp. 659–668. Association for Computing Machinery, 2021. ISBN 9781450384469.
- Hassanpour, N. and Greiner, R. Counterfactual regression with importance sampling weights. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence, IJCAI-19*, pp. 5880–5887. International Joint Conferences on Artificial Intelligence Organization, 7 2019. doi: 10.24963/ijcai.2019/815.
- Heiler, P. and Knaus, M. C. Effect or Treatment Heterogeneity? Policy Evaluation with Aggregated and Disaggregated Treatments. Iza discussion papers, Institute of Labor Economics (IZA), September 2022.
- Hill, J. L. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240, 2011. doi: 10.1198/jcgs.2010.08162.
- Holland, P. W. Statistics and causal inference. *Journal of the American Statistical Association*, 81(396):945–960, 1986.
- Horvitz, D. G. and Thompson, D. J. A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association*, 47(260):663–685, 1952. ISSN 01621459.
- Hu, L., Gu, C., Lopez, M., Ji, J., and Wisnivesky, J. Estimation of causal effects of multiple treatments in observational studies with a binary outcome. *Statistical methods in medical research*, 29(11):3218–3234, 2020.

- Imai, K. and Dyk, D. A. V. Causal inference with general treatment regimes: Generalizing the propensity score. *Journal of the American Statistical Association*, 99(467): 854–866, 2004.
- Imai, K. and Strauss, A. Estimation of heterogeneous treatment effects from randomized experiments, with application to the optimal planning of the get-out-the-vote campaign. *Political Analysis*, 19(1):1–19, 2011. doi: 10.1093/pan/mpq035.
- Imbens, G. W. The role of the propensity score in estimating dose-response functions. *Biometrika*, 87(3):706–710, 2000.
- Johansson, F., Shalit, U., and Sontag, D. Learning representations for counterfactual inference. In Balcan, M. F. and Weinberger, K. Q. (eds.), *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *PMLR*, pp. 3020–3029. PMLR, 20–22 Jun 2016.
- Kaddour, J., Zhu, Y., Liu, Q., Kusner, M. J., and Silva, R. Causal effect inference for structured treatments. In Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp. 24841–24854. Curran Associates, Inc., 2021.
- Kallus, N. Recursive partitioning for personalization using observational data. In *International conference on machine learning*, pp. 1789–1798. PMLR, 2017.
- Kennedy, E. H. Optimal doubly robust estimation of heterogeneous causal effects, 2020.
- Kennedy, E. H., Ma, Z., McHugh, M. D., and Small, D. S. Non-parametric methods for doubly robust estimation of continuous treatment effects. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79(4):1229–1245, 2017.
- Knaus, M. C. Double machine learning-based programme evaluation under unconfoundedness. *The Econometrics Journal*, 25(3):602–627, 2022.
- Knaus, M. C., Lechner, M., and Strittmatter, A. Machine learning estimation of heterogeneous causal effects: Empirical Monte Carlo evidence. *The Econometrics Journal*, 24(1):134–161, 06 2020a.
- Knaus, M. C., Lechner, M., and Strittmatter, A. Heterogeneous employment effects of job search programmes: A machine learning approach. *Journal of Human Resources*, pp. 0718–9615R1, Mar 2020b. ISSN 1548-8004. doi: 10.3368/jhr.57.2.0718-9615r1.
- Künzel, S. R., Sekhon, J. S., Bickel, P. J., and Yu, B. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences*, 116(10):4156–4165, Feb 2019. ISSN 1091-6490. doi: 10.1073/pnas.1804597116.
- Lechner, M. Identification and estimation of causal effects of multiple treatments under the conditional independence assumption. In Lechner, M. and Pfeiffer, F. (eds.), *Econometric Evaluation of Labour Market Policies*, pp. 43–58. Physica-Verlag HD, 2001.
- Lin, L., Zhu, Y., and Chen, L. Causal inference for multi-level treatments with machine-learned propensity scores. *Health Services and Outcomes Research Methodology*, 19(2):106–126, 2019.
- Neyman, J. On the Application of Probability Theory to Agricultural Experiments. Essay on Principles. Section 9. *Statistical Science*, 5:465 – 472, 1923. Translated and published in English in 1990.
- Nie, L., Ye, M., qiang liu, and Nicolae, D. {VCN}et and functional targeted regularization for learning causal effects of continuous treatments. In *International Conference on Learning Representations*, 2021.
- Nie, X. and Wager, S. Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika*, 09 2020. doi: 10.1093/biomet/asaa076.
- Okasa, G. Meta-learners for estimation of causal effects: Finite sample cross-fit performance. *arXiv preprint arXiv:2201.12692*, 2022.
- Olasolo, P., Juárez, M., Morales, M., Liarte, I., et al. Enhanced geothermal systems (egs): A review. *Renewable and Sustainable Energy Reviews*, 56:133–144, 2016.
- Pearl, J. The seven tools of causal inference, with reflections on machine learning. *Commun. ACM*, 62(3):54–60, feb 2019. ISSN 0001-0782. doi: 10.1145/3241036.
- Powers, S., Qian, J., Jung, K., Schuler, A., Shah, N. H., Hastie, T., and Tibshirani, R. Some methods for heterogeneous treatment effect estimation in high dimensions. *Statistics in medicine*, 37(11):1767–1787, 2018.
- Robins, J. M., Rotnitzky, A., and Zhao, L. P. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association*, 89(427):846–866, 1994.
- Robinson, P. M. Root-n-consistent semiparametric regression. *Econometrica*, 56(4):931–954, 1988.
- Rosenbaum, P. R. and Rubin, D. B. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 04 1983.

- Rubin, D. Estimating causal effects if treatment in randomized and nonrandomized studies. *J. Educ. Psychol.*, 66, 01 1974.
- Saini, S. K., Dhamnani, S., Aakash, Ibrahim, A. A., and Chavan, P. Multiple treatment effect estimation using deep generative model with task embedding. In *The World Wide Web Conference, WWW '19*, pp. 1601–1611. Association for Computing Machinery, 2019. ISBN 9781450366748. doi: 10.1145/3308558.3313744.
- Schwab, P., Linhardt, L., Bauer, S., Buhmann, J., and Karlen, W. Learning counterfactual representations for estimating individual dose-response curves. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34: 5612–5619, 04 2020. doi: 10.1609/aaai.v34i04.6014.
- Shalit, U., Johansson, F. D., and Sontag, D. Estimating individual treatment effect: Generalization bounds and algorithms. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML'17*, pp. 3076–3085. JMLR.org, 2017.
- Shi, C., Blei, D., and Veitch, V. Adapting neural networks for the estimation of treatment effects. In *Advances in neural information processing systems*, volume 32. PMLR, 2019.
- Silverman, B. W. *Density estimation for statistics and data analysis*, volume 26. CRC press, 1986.
- Sudret, B. Global sensitivity analysis using polynomial chaos expansions. *Reliability engineering & system safety*, 93(7):964–979, 2008.
- Tibshirani, J., Athey, S., and Wager, S. *grf: Generalized Random Forests*, 2020. URL <https://CRAN.R-project.org/package=grf>. R package version 1.1.0.
- Wager, S. and Athey, S. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242, 2018.
- Wilks, S. Multidimensional statistical scatter. *Collected Papers, Contributions to Mathematical Statistics*, 01 1967.
- Wilks, S. S. Certain generalizations in the analysis of variance. *Biometrika*, 24(3/4):471–494, 1932. ISSN 00063444.
- Wu, P., Li, H., Deng, Y., Hu, W., Dai, Q., Dong, Z., Sun, J., Zhang, R., and Zhou, X.-H. On the opportunity of causal learning in recommendation systems: Foundation, estimation, prediction and challenges. In Raedt, L. D. (ed.), *Proceedings of the 31st International Joint Conference on Artificial Intelligence, IJCAI-22*, pp. 5646–5653. International Joint Conferences on Artificial Intelligence Organization, 7 2022. doi: 10.24963/ijcai.2022/787.
- Yoon, J., Jordon, J., and van der Schaar, M. GANITE: estimation of individualized treatment effects using generative adversarial nets. In *6th International Conference on Learning Representations, ICLR 2018*, 2018.
- Zhang, Y., Kong, D., and Yang, S. Towards R-learner of conditional average treatment effects with a continuous treatment: T-identification, estimation, and inference. *arXiv preprint arXiv:2208.00872*, 2022.
- Zhao, Z. and Harinen, T. Uplift modeling for multiple treatments with cost optimization. In *2019 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*, pp. 422–431. IEEE, 2019.

The Appendix of the paper is divided into the following sections:

- Appendix A contains the proofs of the main propositions in Sections 3: Regularization of the T-learner, the consistency of the naive and extended X-learners, and the generalization of the Robinson decomposition. It also includes a solution for the generalized R-learner with linear regression models.
- Appendix B is divided into two parts: The first part B.1 establishes the bias-variance analysis of pseudo-outcome meta-learners (proof of Theorem 5.4). It provides the main framework for the proof based on the Delta Method and Slutsky’s theorem, which will then be applied to the M-, DR-, and X-learners to show their error’s upper bound. The second part B.2 establishes the bias-variance of the T- and the naive X-learners (proof of Theorem 5.3). This appendix follows the same logic as Appendix B.1’s main proof.
- Appendix C discusses further the comparison of the generalized R-learner and the binarized R-learner.
- Appendix D concerns Section 6.1 with more insights about the numerical experiments and results.
- Appendix E describes the Semi-synthetic dataset used in subsection 6.2

A. Proofs of propositions of Section 4

Proposition A.1 (Regularizing the T-learner against selection bias). *For a treatment level $t \in \mathcal{T}$, the expected squared error of the estimator $\hat{\mu}_t$ on the outcome surface μ_t satisfies:*

$$\begin{aligned} \mathbb{E}_{\mathbf{X} \sim \mathbb{P}(\cdot)} [(\hat{\mu}_t(\mathbf{X}) - \mu_t(\mathbf{X}))^2] = \\ \mathbb{E}_{\mathbf{X} \sim \mathbb{P}(\cdot | T=t)} \left[\frac{\mathbb{P}(T=t)}{r(t, \mathbf{X})} (\hat{\mu}_t(\mathbf{X}) - \mu_t(\mathbf{X}))^2 \right]. \end{aligned} \tag{3}$$

where $\mathbb{P}(\cdot)$ is the marginal distribution of \mathbf{X} and $\mathbb{P}(\cdot | T=t)$ is the conditional distribution of \mathbf{X} given $T=t$.

A.1. Proof of Proposition A.1

This proof is similar to the proof of equation (5) in supplementary of [Curth & van der Schaar \(2021a\)](#). For simplicity, we assume that the distribution of \mathbf{X} and the conditional distribution of \mathbf{X} given $T=t$ are absolutely continuous with respect to the Lebesgue measure over \mathbb{R}^d . Let $p_{\mathbf{X}}(\mathbf{x})$ denote the probability distribution function of \mathbf{X} , let $p(\mathbf{x} | T=t)$ denote the probability distribution function of \mathbf{X} given $T=t$ and let $R_t = \int (\hat{\mu}_t(\mathbf{x}) - \mu_t(\mathbf{x}))^2 p(\mathbf{x} | T=t) d\mathbf{x}$.

$$\begin{aligned}
 \mathbb{E}_{\mathbf{X} \sim \mathbb{P}(\cdot)} [(\hat{\mu}_t(\mathbf{X}) - \mu_t(\mathbf{X}))^2] &= \int (\hat{\mu}_t(\mathbf{x}) - \mu_t(\mathbf{x}))^2 p(\mathbf{x}) d\mathbf{x} \\
 &= \mathbb{P}(T = t) \int (\hat{\mu}_t(\mathbf{x}) - \mu_t(\mathbf{x}))^2 p(\mathbf{x} | T = t) d\mathbf{x} + \sum_{t' \neq t} \mathbb{P}(T = t') \int (\hat{\mu}_t(\mathbf{x}) - \mu_t(\mathbf{x}))^2 p(\mathbf{x} | T = t') d\mathbf{x} \\
 &= \mathbb{P}(T = t) R_t + \sum_{t' \neq t} \mathbb{P}(T = t') \int (\hat{\mu}_t(\mathbf{x}) - \mu_t(\mathbf{x}))^2 \frac{p(\mathbf{x} | T = t')}{p(\mathbf{x} | T = t)} p(\mathbf{x} | T = t) d\mathbf{x} \\
 &= \mathbb{P}(T = t) R_t + \sum_{t' \neq t} \mathbb{P}(T = t') \int (\hat{\mu}_t(\mathbf{x}) - \mu_t(\mathbf{x}))^2 \frac{\frac{\mathbb{P}(T=t'|\mathbf{x})p(\mathbf{x})}{\mathbb{P}(T=t')}}{\frac{\mathbb{P}(T=t|\mathbf{x})p(\mathbf{x})}{\mathbb{P}(T=t)}} p(\mathbf{x} | T = t) d\mathbf{x} \quad (\text{Bayes rule}) \\
 &= \mathbb{P}(T = t) R_t + \mathbb{P}(T = t) \sum_{t' \neq t} \int (\hat{\mu}_t(\mathbf{x}) - \mu_t(\mathbf{x}))^2 \frac{\mathbb{P}(T = t' | \mathbf{x})}{\mathbb{P}(T = t | \mathbf{x})} p(\mathbf{x} | T = t) d\mathbf{x} \quad (4) \\
 &= \mathbb{P}(T = t) R_t + \mathbb{P}(T = t) \int (\hat{\mu}_t(\mathbf{x}) - \mu_t(\mathbf{x}))^2 \frac{\sum_{t' \neq t} \mathbb{P}(T = t' | \mathbf{x})}{\mathbb{P}(T = t | \mathbf{x})} p(\mathbf{x} | T = t) d\mathbf{x} \\
 &= \mathbb{P}(T = t) R_t + \mathbb{P}(T = t) \int \frac{1 - r(t, \mathbf{x})}{r(t, \mathbf{x})} (\hat{\mu}_t(\mathbf{x}) - \mu_t(\mathbf{x}))^2 p(\mathbf{x} | T = t) d\mathbf{x} \\
 &= \mathbb{P}(T = t) \int \left(1 + \frac{1 - r(t, \mathbf{x})}{r(t, \mathbf{x})} \right) (\hat{\mu}_t(\mathbf{x}) - \mu_t(\mathbf{x}))^2 p(\mathbf{x} | T = t) d\mathbf{x} \\
 &= \mathbb{E}_{\mathbf{X} \sim p(\cdot | T=t)} \left[\frac{\mathbb{P}(T = t)}{r(t, \mathbf{X})} (\hat{\mu}_t(\mathbf{X}) - \mu_t(\mathbf{X}))^2 \right].
 \end{aligned}$$

A.2. Consistency of the X-learner

By direct calculations, we show that

$$\mathbb{E}(Z_k^X | \mathbf{X} = \mathbf{x}) = \mathbb{E}[\mathbf{1}\{T = t_k\} Y(t_k) | \mathbf{X} = \mathbf{x}] - r(t_k, \mathbf{x}) \mu_{t_0}(\mathbf{x}) + \sum_{l \neq k} r(t_l, \mathbf{x}) (\mu_{t_k}(\mathbf{x}) \quad (5)$$

$$- \mathbb{E}[\mathbf{1}\{T = t_l\} Y(t_l) | \mathbf{X} = \mathbf{x}]) + \sum_{l \neq k} r(t_l, \mathbf{x}) (\mu_{t_l}(\mathbf{x}) - \mu_{t_0}(\mathbf{x})) \quad (6)$$

$$= r(t_k, \mathbf{x}) \mu_{t_k}(\mathbf{x}) - r(t, \mathbf{x}) \mu_{t_0}(\mathbf{x}) + \sum_{l \neq k} (r(t_l, \mathbf{x}) \mu_t(\mathbf{x}) - r(t_l, \mathbf{x}) \mu_{t'}(\mathbf{x})) \quad (7)$$

$$+ \sum_{l \neq k} r(t_l, \mathbf{x}) (\mu_{t_l}(\mathbf{x}) - \mu_{t_0}(\mathbf{x})) \quad (\text{by Assumption 3.1}) \quad (8)$$

$$= r(t_k, \mathbf{x}) \mu_{t_k}(\mathbf{x}) - r(t_k, \mathbf{x}) \mu_{t_0}(\mathbf{x}) + \sum_{l \neq k} r(t_l, \mathbf{x}) \mu_{t_k}(\mathbf{x}) - \sum_{l \neq k} r(t_l, \mathbf{x}) \mu_{t_0}(\mathbf{x}) \quad (9)$$

$$= (\mu_{t_k}(\mathbf{x}) - \mu_{t_0}(\mathbf{x})) \left(r(t_k, \mathbf{x}) + \sum_{l \neq k} r(t_l, \mathbf{x}) \right) \quad (10)$$

$$= \mu_{t_k}(\mathbf{x}) - \mu_{t_0}(\mathbf{x}) = \tau_k(\mathbf{x}). \quad (11)$$

A.3. Consistency of the naive extension of the X-learner

Let us consider the two random variables $D^{(k)} := Y(t_k) - \mu_{t_0}(\mathbf{X})$ and $D^{(0)} := \mu_{t_k}(\mathbf{X}) - Y(t_0)$. We have

$$\begin{aligned}
 \tau^{(k)}(\mathbf{x}) &= \mathbb{E}(D^{(k)} | \mathbf{X} = \mathbf{x}) = \mathbb{E}(Y(t_k) - \mu_{t_0}(\mathbf{X}) | \mathbf{X} = \mathbf{x}) \\
 &= \mathbb{E}[Y(t_k) | \mathbf{X} = \mathbf{x}] - \mu_{t_0}(\mathbf{x}) \\
 &= \mu_{t_k}(\mathbf{x}) - \mu_{t_0}(\mathbf{x}) = \tau_k(\mathbf{x}),
 \end{aligned} \quad (12)$$

and

$$\begin{aligned}
 \tau^{(0)}(\mathbf{x}) &= \mathbb{E}(D^{(0)} | \mathbf{X} = \mathbf{x}) = \mathbb{E}(\mu_{t_k}(\mathbf{X}) - Y(t_k) | \mathbf{X} = \mathbf{x}) \\
 &= \mu_{t_k}(\mathbf{x}) - \mathbb{E}[\mathbb{E}(Y(t_k) | \mathbf{X} = \mathbf{x})] \\
 &= \mu_{t_k}(\mathbf{x}) - \mu_{t_0}(\mathbf{x}) = \tau_k(\mathbf{x}).
 \end{aligned} \quad (13)$$

Therefore,

$$\tau_k^{(X, \text{nv})}(\mathbf{x}) = \frac{r(t_k, \mathbf{x})}{r(t_k, \mathbf{x}) + r(t_0, \mathbf{x})} \tau^{(k)}(\mathbf{x}) + \frac{r(t_0, \mathbf{x})}{r(t_k, \mathbf{x}) + r(t_0, \mathbf{x})} \tau^{(0)}(\mathbf{x}) = \tau_k(\mathbf{x}). \quad (14)$$

A.4. Generalizing the Robinson decomposition

We show first the Neyman-Orthogonality propriety, i.e. $\mathbb{E}(\epsilon \mid T, \mathbf{X}) = 0$. Indeed, for $t \in \mathcal{T}$ and $\mathbf{x} \in \mathcal{D}$, we have

$$\begin{aligned} \mathbb{E}[\epsilon \mid T = t, \mathbf{X} = \mathbf{x}] &= \mathbb{E}[Y_{\text{obs}} - \mu_T(\mathbf{X}) \mid T = t, \mathbf{X} = \mathbf{x}] \\ &= \mathbb{E}[Y(t) - \mu_T(\mathbf{X}) \mid T = t, \mathbf{X} = \mathbf{x}] \\ &= \mu_t(\mathbf{x}) - \mu_t(\mathbf{x}) = 0. \end{aligned} \quad (15)$$

Thus, the observed outcome model satisfies:

$$\begin{aligned} \mathbb{E}(Y_{\text{obs}} \mid \mathbf{X} = \mathbf{x}) &= \mathbb{E}\left[\epsilon + \sum_{k=0}^K \mathbf{1}\{T = t_k\} \mu_{t_k}(\mathbf{X}) \mid \mathbf{X} = \mathbf{x}\right] \\ &= \mathbb{E}[\mathbb{E}[\epsilon \mid T, \mathbf{X}] \mid \mathbf{X} = \mathbf{x}] + \sum_{k=0}^K \mathbb{E}[\mathbf{1}\{T = t_k\} \mid \mathbf{X} = \mathbf{x}] \mu_{t_k}(\mathbf{x}) \\ &= \sum_{k=0}^K \mu_{t_k}(\mathbf{x}) r(t_k, \mathbf{x}) = \mu_{t_0}(\mathbf{x}) r(t_0, \mathbf{x}) + \sum_{k=1}^K \mu_{t_k}(\mathbf{x}) r(t_k, \mathbf{x}) \\ &= \mu_{t_0}(\mathbf{x}) \left[1 - \sum_{k=1}^K r(t_k, \mathbf{x})\right] + \sum_{k=1}^K \mu_{t_k}(\mathbf{x}) r(t_k, \mathbf{x}) \\ &= \mu_{t_0}(\mathbf{x}) + \sum_{k=1}^K r(t_k, \mathbf{x}) [\mu_{t_k}(\mathbf{x}) - \mu_{t_0}(\mathbf{x})] \\ &= \mu_{t_0}(\mathbf{x}) + \sum_{k=1}^K r(t_k, \mathbf{x}) \tau_k(\mathbf{x}) = m(\mathbf{x}). \end{aligned} \quad (16)$$

By gathering both quantities :

$$\begin{aligned} Y_{\text{obs}} - m(\mathbf{X}) &= \sum_{k=0}^K \mathbf{1}\{T = t_k\} \mu_{t_k}(\mathbf{X}) - \mu_{t_0}(\mathbf{X}) - \sum_{k=1}^K r(t_k, \mathbf{X}) \tau_k(\mathbf{X}) + \epsilon \\ &= \mathbf{1}\{T = t_0\} \mu_{t_0}(\mathbf{X}) + \sum_{k=1}^K \mathbf{1}\{T = t_k\} \mu_{t_k}(\mathbf{X}) - \mu_{t_0}(\mathbf{X}) - \sum_{k=1}^K r(t_k, \mathbf{X}) \tau_k(\mathbf{X}) + \epsilon \\ &= (\mathbf{1}\{T = t_0\} - 1) \mu_{t_0}(\mathbf{X}) + \sum_{k=1}^K (\mathbf{1}\{T = t_k\} \mu_{t_k}(\mathbf{X}) - r(t_k, \mathbf{X}) \tau_k(\mathbf{X})) + \epsilon \\ &= \sum_{k=1}^K (\mathbf{1}\{T = t_k\} \mu_{t_k}(\mathbf{X}) - r(t_k, \mathbf{X}) \tau_k(\mathbf{X})) - \sum_{k=1}^K \mathbf{1}\{T = t_k\} \mu_{t_0}(\mathbf{X}) + \epsilon \\ &= \sum_{k=1}^K (\mathbf{1}\{T = t_k\} \mu_{t_k}(\mathbf{X}) - \mathbf{1}\{T = t_k\} \mu_{t_0}(\mathbf{X}) - r(t_k, \mathbf{X}) \tau_k(\mathbf{X})) + \epsilon \\ &= \sum_{k=1}^K [\mathbf{1}\{T = t_k\} - r(t_k, \mathbf{X})] \tau_k(\mathbf{X}) + \epsilon. \end{aligned} \quad (17)$$

Therefore, we obtain the generalized Robinson decomposition for the multi-treatment regime.

A.5. Solving the generalized R-learner for linear models

For $k = 1, \dots, K$, we assume that $\bar{\tau}_k$ belongs to the family of linear regression models such that:

$$\mathcal{F} = \left\{ \left\{ \bar{\tau}_k(\mathbf{x}) := \beta_{k,0} + \sum_{j=1}^{p-1} \beta_{k,j} f_j(\mathbf{x}) \right\}_{k=1}^K / \boldsymbol{\beta}_k = (\beta_{k,0}, \dots, \beta_{k,p-1})^\top \in \mathbb{R}^p \right\}. \quad (18)$$

f_j are predefined functions (e.g. polynomial functions). It is also possible to use a matrix notation and write $\bar{\tau}_k(\mathbf{X}) = \mathbf{H}\boldsymbol{\beta}_k$ where $\mathbf{H} = (f_j(\mathbf{X}_i)) \in \mathbb{R}^{n \times p}$ assumed to be full rank matrix $\text{rank}(\mathbf{H}) = p \leq n$.

Let $\bar{Y} = (\bar{Y}_i)_{i=1}^n$ and $\bar{T}_k = (\bar{T}_{i,k})_{i=1}^n$ such that $\bar{Y}_i = Y_{\text{obs},i} - \widehat{m}(\mathbf{X}_i)$ and $\bar{T}_{i,k} = \mathbf{1}\{t_i = t_k\} - \widehat{r}(t_k, \mathbf{X}_i)$. Let $\epsilon = (\epsilon_i)_{i=1}^n$ denote the vector of errors obtained for the generalized Robinson (1988) decomposition in Proposition 3.3.

We show immediately that \mathcal{L} , the generalized R-loss function associated with the mean squared error of ϵ in (2) in the paper, is quadratic with respect to $\boldsymbol{\beta}$. Indeed,

$$\begin{aligned} \mathcal{L}(\{\bar{\tau}_k\}_{t \neq t^{(0)}}) &= \frac{1}{n} \epsilon^\top \epsilon = \frac{1}{n} \left(\bar{Y} - \sum_{k=1}^K \bar{T}_k \odot (\mathbf{H}\boldsymbol{\beta}_k) \right)^\top \left(\bar{Y} - \sum_{k=1}^K \bar{T}_k \odot (\mathbf{H}\boldsymbol{\beta}_k) \right) \\ &= \frac{1}{n} \left[\bar{Y}^\top \bar{Y} - 2 \sum_{k=1}^K \bar{Y}^\top (\bar{T}_k \odot (\mathbf{H}\boldsymbol{\beta}_k)) + \sum_{k,k'=1}^K (\bar{T}_k \odot (\mathbf{H}\boldsymbol{\beta}_k))^\top (\bar{T}_{k'} \odot (\mathbf{H}\boldsymbol{\beta}_{k'})) \right] \\ &= \frac{1}{n} \left(\bar{Y}^\top \bar{Y} - 2 \sum_{k=1}^K \bar{Y}^\top \mathbf{D}_{\bar{T}_k} \mathbf{H} \boldsymbol{\beta}_k + \sum_{k,k'=1}^K \boldsymbol{\beta}_k^\top \mathbf{H}^\top \mathbf{D}_{\bar{T}_k} \mathbf{D}_{\bar{T}_{k'}} \mathbf{H} \boldsymbol{\beta}_{k'} \right), \end{aligned} \quad (19)$$

where \odot is the Hadamard product (element-wise product). The last line holds because $\bar{T}_k \odot (\mathbf{H}\boldsymbol{\beta}_k) = \mathbf{D}_{\bar{T}_k} \mathbf{H} \boldsymbol{\beta}_k$ with $\mathbf{D}_{\bar{T}_k}$ is the diagonal matrix of the vector $\bar{T}_k = (\bar{T}_{i,k})_{i=1}^n$.

By differentiating $\partial \mathcal{L} / \partial \boldsymbol{\beta}_k = 0$ for $k = 1, \dots, K$:

$$\begin{cases} -\mathbf{a}_1 + \mathbf{B}_1 \widehat{\boldsymbol{\beta}}_1 + \sum_{k=2}^K \mathbf{C}_{1k} \widehat{\boldsymbol{\beta}}_k = 0 \\ \vdots \\ -\mathbf{a}_K + \sum_{k=1}^{K-1} \mathbf{C}_{Kk} \widehat{\boldsymbol{\beta}}_k + \mathbf{B}_K \widehat{\boldsymbol{\beta}}_K = 0 \end{cases} \quad (20)$$

$$\Leftrightarrow \begin{bmatrix} \mathbf{B}_1 & \mathbf{C}_{12} & \cdots & \mathbf{C}_{1K} \\ \mathbf{C}_{21} & \mathbf{B}_2 & \cdots & \mathbf{C}_{2K} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{C}_{K1} & \mathbf{C}_{K2} & \cdots & \mathbf{B}_K \end{bmatrix} \begin{bmatrix} \widehat{\boldsymbol{\beta}}_1 \\ \widehat{\boldsymbol{\beta}}_2 \\ \vdots \\ \widehat{\boldsymbol{\beta}}_K \end{bmatrix} = \begin{bmatrix} \mathbf{a}_1 \\ \mathbf{a}_2 \\ \vdots \\ \mathbf{a}_K \end{bmatrix}, \quad (21)$$

where

$$\mathbf{a}_j = \frac{1}{n} \mathbf{H}^\top \mathbf{D}_{\bar{T}_j} \bar{Y} \in \mathbb{R}^p, \quad (22)$$

$$\mathbf{B}_j = \frac{1}{n} \mathbf{H}^\top \mathbf{D}_{\bar{T}_j}^2 \mathbf{H} \in \mathbb{R}^{p \times p}, \quad (23)$$

$$\mathbf{C}_{ij} = \frac{1}{n} \mathbf{H}^\top \mathbf{D}_{\bar{T}_i} \mathbf{D}_{\bar{T}_j} \mathbf{H} \in \mathbb{R}^{p \times p}. \quad (24)$$

Let $\boldsymbol{\beta} = (\boldsymbol{\beta}_1^\top, \dots, \boldsymbol{\beta}_K^\top)^\top \in \mathbb{R}^{K \times p}$ and consider the block matrix \mathbf{A} defined as

$$\mathbf{A} = \begin{bmatrix} \mathbf{B}_1 & \mathbf{C}_{12} & \cdots & \mathbf{C}_{1K} \\ \mathbf{C}_{21} & \mathbf{B}_2 & \cdots & \mathbf{C}_{2K} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{C}_{K1} & \mathbf{C}_{K2} & \cdots & \mathbf{B}_K \end{bmatrix}. \quad (25)$$

The matrix \mathbf{A} is real symmetric and satisfies:

$$\begin{aligned}\boldsymbol{\beta}^\top \mathbf{A} \boldsymbol{\beta} &= \sum_{1 \leq k, l \leq K} \boldsymbol{\beta}_k^\top \mathbf{H}^\top \mathbf{D}_{\bar{T}_k} \mathbf{D}_{\bar{T}_l} \mathbf{H} \boldsymbol{\beta}_l \\ &= \left\| \sum_{k=1}^K \mathbf{D}_{\bar{T}_k} \mathbf{H} \boldsymbol{\beta}_k \right\|^2 \geq 0.\end{aligned}\tag{26}$$

This result shows that \mathbf{A} is positive semi-definite, all its eigenvalues are nonnegative and also proves the existence of a minimizer $\hat{\boldsymbol{\beta}}$ to the loss function \mathcal{L} . However, this is not sufficient to prove the uniqueness of the solution as one cannot prove all eigenvalues are positive.

The solution $\hat{\boldsymbol{\beta}}$ to Problem (4.3) in the main paper with the minimal norm is given by

$$\hat{\boldsymbol{\beta}} = \mathbf{A}^+ \mathbf{a},\tag{27}$$

where \mathbf{A}^+ is the Moore–Penrose inverse of \mathbf{A} and $\mathbf{a} = (\mathbf{a}_1^\top, \dots, \mathbf{a}_K^\top)^\top$.

B. Theoretical analysis of the error bounds.

B.1. Error estimation of pseudo-outcome meta-learners.

Step 0. Set-up of the theorem

In the following subsection, we will analyze the error estimation of each two-step meta-learner. Given Assumption 5.1 stating that the observations are generated from a function f respecting the two causal assumptions (3.1-3.2), each unit i has the following observed and potential outcomes

$$\begin{aligned}Y_i(t_k) &= f(t_k, \mathbf{X}_i) + \varepsilon_i(t_k), \\ Y_i(t_0) &= f(t_0, \mathbf{X}_i) + \varepsilon_i(t_0).\end{aligned}\tag{28}$$

where $\varepsilon_i(t)$ are i.i.d. Gaussian $\mathcal{N}(0, \sigma^2)$ and independent of $(T_i, \mathbf{X}_i)_{i=1}^n$. As a consequence, the noise $(\varepsilon_i)_{i=1}^n = (\varepsilon_i(T_i))_{i=1}^n$ is also Gaussian $\mathcal{N}(0, \sigma^2)$ and is independent of $(T_i, \mathbf{X}_i)_{i=1}^n$.

The CATE model τ_k for each $k = 1, \dots, K$ can be written as:

$$\begin{aligned}\tau_k(\mathbf{x}) &= \mathbb{E}(Y(t_k) - Y(t_0) \mid \mathbf{X} = \mathbf{x}) \\ &= \mathbb{E}(f(t_k, \mathbf{X}) - f(t_0, \mathbf{X}) + \epsilon^* \mid \mathbf{X} = \mathbf{x}) \\ &= f(t_k, \mathbf{x}) - f(t_0, \mathbf{x})\end{aligned}\tag{29}$$

where ϵ^* is a noise independent of \mathbf{X} (and T) and satisfying $\mathbb{E}(\epsilon^*) = 0$.

Under the assumption 5.2, we write $\tau_k(\mathbf{X}) = f(t_k, \mathbf{X}) - f(t_0, \mathbf{X}) = \mathbf{H} \boldsymbol{\beta}_k^*$ where $\boldsymbol{\beta}_k^* = \boldsymbol{\beta}_{t_k} - \boldsymbol{\beta}_{t_0}$ and $\mathbf{H} = (\mathbf{H}_{ij}) \in \mathbb{R}^{n \times p}$ is the regression matrix, assumed to be full rank matrix, such that $\mathbf{H}_{ij} = f_j(\mathbf{X}_i)$ for $i = 1, \dots, n$ and $j = 0, \dots, p - 1$. With pseudo-outcome meta-learners, we consider a random variable Z_k for a fixed t_k such that

$$Z_{k,i} = A_{t_k}(T_i, \mathbf{X}_i) Y_{\text{obs},i} + B_{t_k}(T_i, \mathbf{X}_i), \quad i = 1, \dots, n,$$

where the functions $A_{t_k}(T, \mathbf{X})$ and $B_{t_k}(T, \mathbf{X})$ are given for each pseudo-outcome meta-learner.

Step 1. Identification of $\hat{\boldsymbol{\beta}}_k$ and $\boldsymbol{\beta}_k^*$

The regression coefficients $\hat{\boldsymbol{\beta}}_k$ are given by the Ordinary Least Squares (OLS) method

$$\hat{\boldsymbol{\beta}}_k = (\mathbf{H}^\top \mathbf{H})^{-1} \mathbf{H}^\top \mathbf{z}_k,\tag{30}$$

where $\mathbf{z}_k = (Z_{k,i})_{1 \leq i \leq n}$. Thus,

$$\begin{aligned}
 \hat{\boldsymbol{\beta}}_k &= (\mathbf{H}^\top \mathbf{H})^{-1} \mathbf{H}^\top \mathbf{z}_k \\
 &= (\mathbf{H}^\top \mathbf{H})^{-1} \mathbf{H}^\top (A_{t_k}(T_i, \mathbf{X}_i) Y_{\text{obs},i} + B_{t_k}(T_i, \mathbf{X}_i))_{i=1}^n \\
 &= (\mathbf{H}^\top \mathbf{H})^{-1} \mathbf{H}^\top (A_{t_k}(T_i, \mathbf{X}_i) f(T_i, \mathbf{X}_i) + B_{t_k}(T_i, \mathbf{X}_i) + A_{t_k}(T_i, \mathbf{X}_i) \epsilon_i)_{i=1}^n \\
 &= (\mathbf{H}^\top \mathbf{H})^{-1} \mathbf{H}^\top (\tau_k(\mathbf{x}) + A_{t_k}(T_i, \mathbf{X}_i) f(T_i, \mathbf{X}_i) - \tau_k(\mathbf{x}) + B_{t_k}(T_i, \mathbf{X}_i) + A_{t_k}(T_i, \mathbf{X}_i) \epsilon_i)_{i=1}^n \\
 &= (\mathbf{H}^\top \mathbf{H})^{-1} \mathbf{H}^\top (\mathbf{H} \boldsymbol{\beta}_k^* + A_{t_k}(T_i, \mathbf{X}_i) f(T_i, \mathbf{X}_i) - \tau_k(\mathbf{x}) + B_{t_k}(T_i, \mathbf{X}_i) + A_{t_k}(T_i, \mathbf{X}_i) \epsilon_i)_{i=1}^n \\
 &= \boldsymbol{\beta}_k^* + (\mathbf{H}^\top \mathbf{H})^{-1} \mathbf{H}^\top (A_{t_k}(T_i, \mathbf{X}_i) f(T_i, \mathbf{X}_i) - \tau_k(\mathbf{x}) + B_{t_k}(T_i, \mathbf{X}_i) + A_{t_k}(T_i, \mathbf{X}_i) \epsilon_i)_{i=1}^n \\
 &= \boldsymbol{\beta}_k^* + (\mathbf{H}^\top \mathbf{H})^{-1} \mathbf{H}^\top \tilde{\boldsymbol{\epsilon}}_k
 \end{aligned}$$

where $\tilde{\boldsymbol{\epsilon}}_{k,i} = \psi_k(T_i, \mathbf{X}_i) + A_{t_k}(T_i, \mathbf{X}_i) \epsilon_i$ and $\psi_k(T_i, \mathbf{X}_i) = A_{t_k}(T_i, \mathbf{X}_i) f(T_i, \mathbf{X}_i) - \tau_k(\mathbf{X}_i) + B_{t_k}(T_i, \mathbf{X}_i)$ to simplify notations.

Let us consider the random vector $\mathbf{Z}_k^{(n)}$ such that

$$\mathbf{Z}_k^{(n)} = \left(\frac{1}{n} (\mathbf{H}^\top \tilde{\boldsymbol{\epsilon}}_k)_1, \dots, \frac{1}{n} (\mathbf{H}^\top \tilde{\boldsymbol{\epsilon}}_k)_p, \frac{1}{n} (\mathbf{H}^\top \mathbf{H})_{11}, \dots, \frac{1}{n} (\mathbf{H}^\top \mathbf{H})_{pp} \right)^\top \in \mathbb{R}^{p+p^2}, \quad (31)$$

that allows us to write $\hat{\boldsymbol{\beta}}_k$ as

$$\begin{aligned}
 \hat{\boldsymbol{\beta}}_k &= \boldsymbol{\beta}_k^* + (\mathbf{H}^\top \mathbf{H})^{-1} \mathbf{H}^\top \tilde{\boldsymbol{\epsilon}}_k \\
 &= \boldsymbol{\beta}_k^* + \left(\frac{1}{n} \mathbf{H}^\top \mathbf{H} \right)^{-1} \left(\frac{1}{n} \mathbf{H}^\top \tilde{\boldsymbol{\epsilon}}_k \right) \\
 &= \boldsymbol{\beta}_k^* + \phi(\mathbf{Z}_k^{(n)}),
 \end{aligned} \quad (32)$$

where $\phi : \mathbb{R}^{p+p^2} \rightarrow \mathbb{R}^p$ is a \mathcal{C}^1 -function.

Step 2. The asymptotic behaviour of the OLS estimator's mean and covariance

In order to apply the Central Limit Theorem (CLT) later, we show that the vector $\mathbf{Z}_k^{(n)}$ can be written as sum of *i.i.d.* random vectors $\mathbf{Z}_{k,i}$.

$$\begin{aligned}
 \mathbf{Z}_k^{(n)} &= \left(\frac{1}{n} (\mathbf{H}^\top \tilde{\boldsymbol{\epsilon}}_k)_1, \dots, \frac{1}{n} (\mathbf{H}^\top \tilde{\boldsymbol{\epsilon}}_k)_p, \frac{1}{n} (\mathbf{H}^\top \mathbf{H})_{11}, \dots, \frac{1}{n} (\mathbf{H}^\top \mathbf{H})_{pp} \right)^\top \in \mathbb{R}^{p+p^2} \\
 &= \left(\frac{1}{n} \sum_{i=1}^n \mathbf{H}_{i1} \tilde{\boldsymbol{\epsilon}}_{k,i}, \dots, \frac{1}{n} \sum_{i=1}^n \mathbf{H}_{ip} \tilde{\boldsymbol{\epsilon}}_{k,i}, \frac{1}{n} \sum_{i=1}^n \mathbf{H}_{i1} \mathbf{H}_{i1}, \dots, \frac{1}{n} \sum_{i=1}^n \mathbf{H}_{ip} \mathbf{H}_{ip} \right)^\top \\
 &= \frac{1}{n} \sum_{i=1}^n (\mathbf{H}_{i1} \tilde{\boldsymbol{\epsilon}}_{k,i}, \dots, \mathbf{H}_{ip} \tilde{\boldsymbol{\epsilon}}_{k,i}, \mathbf{H}_{i1} \mathbf{H}_{i1}, \dots, \mathbf{H}_{ip} \mathbf{H}_{ip})^\top = \frac{1}{n} \sum_{i=1}^n \mathbf{Z}_{k,i}.
 \end{aligned} \quad (33)$$

The mean \mathbf{m} of the vector $\mathbf{Z}_k^{(n)}$ satisfies

$$\begin{aligned}
 \mathbf{m} &= \mathbb{E}(\mathbf{Z}_k^{(n)}) = \frac{1}{n} \sum_{i=1}^n \mathbb{E}(\mathbf{Z}_{k,i}) = \mathbb{E}(\mathbf{Z}_{k,i}) \\
 &= \left(h_1, \dots, h_p, F_{11}, \dots, F_{pp} \right)^\top,
 \end{aligned} \quad (34)$$

where, for $j, j' = 1, \dots, p$,

$$\begin{aligned}
 h_j &= \mathbb{E}[f_{j-1}(\mathbf{X})(\psi_k(T, \mathbf{X}) + A_{t_k}(T, \mathbf{X})\epsilon)] = \mathbb{E}(f_{j-1}(\mathbf{X})\psi_k(T, \mathbf{X})) \\
 F_{jj'} &= \mathbb{E}(f_{j-1}(\mathbf{X})f_{j'-1}(\mathbf{X})).
 \end{aligned} \quad (35)$$

The covariance matrix \mathbf{C} of $\mathbf{Z}_k^{(n)}$ has entries

$$\begin{aligned} \mathbf{C}_{jj'} &= \text{Cov} \left(\mathbf{Z}_j^{(k)}, \mathbf{Z}_{j'}^{(k)} \right) = \mathbb{E}(\mathbf{Z}_j^{(k)}, \mathbf{Z}_{j'}^{(k)}) - \mathbb{E}(\mathbf{Z}_j^{(k)})\mathbb{E}(\mathbf{Z}_{j'}^{(k)}) \\ &= \begin{cases} \mathbb{E}(f_{j-1}(\mathbf{X})f_{j'-1}(\mathbf{X})(\psi_k(T, \mathbf{X}) + A_{t_k}(T, \mathbf{X})\epsilon)^2) - h_j h_{j'} & \text{if } j, j' \in \{1, \dots, p\} \\ \mathbb{E}(f_{\tilde{k}-1}(\mathbf{X})f_{\tilde{k}'-1}(\mathbf{X})f_{l-1}(\mathbf{X})f_{l'-1}(\mathbf{X})) - F_{\tilde{k}\tilde{k}'}F_{ll'} & \text{if } j, j' \in \{p+1, \dots, p^2\} \\ \mathbb{E}(f_{\tilde{k}-1}(\mathbf{X})f_{\tilde{k}'-1}(\mathbf{X})f_{j-1}(\mathbf{X})(\psi_k(T, \mathbf{X}) + A_{t_k}(T, \mathbf{X})\epsilon)) - h_j F_{\tilde{k}\tilde{k}'} & \text{otherwise.} \end{cases} \\ &= \begin{cases} \mathbb{E}(f_{j-1}(\mathbf{X})f_{j'-1}(\mathbf{X})\psi_k^2(T, \mathbf{X})) + \sigma^2\mathbb{E}(f_{j-1}(\mathbf{X})f_{j'-1}(\mathbf{X})A_{t_k}^2(T, \mathbf{X})) - h_j h_{j'} & \text{if } j, j' \in \{1, \dots, p\} \\ \mathbb{E}(f_{\tilde{k}-1}(\mathbf{X})f_{\tilde{k}'-1}(\mathbf{X})f_{l-1}(\mathbf{X})f_{l'-1}(\mathbf{X})) - F_{\tilde{k}\tilde{k}'}F_{ll'} & \text{if } j, j' \in \{p+1, \dots, p^2\} \\ \mathbb{E}(f_{\tilde{k}-1}(\mathbf{X})f_{\tilde{k}'-1}(\mathbf{X})f_{j-1}(\mathbf{X})\psi_k(T, \mathbf{X})) - h_j F_{\tilde{k}\tilde{k}'} & \text{otherwise,} \end{cases} \end{aligned} \quad (36)$$

where $\tilde{k}, \tilde{k}' = \eta^{-1}(j)$ (respectively, $l, l' = \eta^{-1}(j')$) such that η is the correspondence indexes map between \mathbf{m} and F in $\mathbf{m}_j = F_{\tilde{k}\tilde{k}'}$ (respectively, $\mathbf{m}_{j'} = F_{ll'}$) when $j \geq p+1$ (respectively, when $j' \geq p+1$).

By considering now the vector

$$\mathbf{S}^{(n)} = \sqrt{n}(\mathbf{Z}_k^{(n)} - \mathbf{m}) = \frac{1}{\sqrt{n}} \sum_{i=1}^n (\mathbf{Z}_{k,i} - \mathbf{m}), \quad (37)$$

one can show by the multivariate CLT that

$$\mathbf{S}^{(n)} = \sqrt{n}(\mathbf{Z}_k^{(n)} - \mathbf{m}) \xrightarrow{\mathcal{L}} \mathcal{N}(\mathbf{0}, \mathbf{C}). \quad (38)$$

This allows us to write $\hat{\beta}_k$ as function of $\mathbf{S}^{(n)}$ and \mathbf{m} . Indeed,

$$\begin{aligned} \hat{\beta}_k &= \beta_k^* + (\mathbf{H}^\top \mathbf{H})^{-1} \mathbf{H}^\top \tilde{\epsilon} \\ &= \beta_k^* + \phi(\mathbf{Z}^{(n)}) \\ &= \beta_k^* + \phi(\mathbf{m} + \mathbf{S}^{(n)}/\sqrt{n}) \\ &= \beta_k^* + \Phi(\mathbf{S}^{(n)}, \mathbf{m}), \end{aligned} \quad (39)$$

where $\Phi : \mathbb{R}^{p+p^2} \times \mathbb{R}^{p+p^2} \rightarrow \mathbb{R}^p$ is also \mathcal{C}^1 -function.

Since $\sqrt{n}(\mathbf{S}^{(n)} - \mathbf{0}) \xrightarrow{\mathcal{L}} \mathcal{N}(\mathbf{0}, \mathbf{C})$, one obtains by the Delta method

$$\sqrt{n}[\Phi(\mathbf{S}^{(n)}, \mathbf{m}) - \Phi(\mathbf{0}, \mathbf{m})] \xrightarrow{\mathcal{L}} \mathcal{N} \left(\mathbf{0}, J_{\Phi}^{(1)}(\mathbf{0}, \mathbf{m})^\top \mathbf{C} J_{\Phi}^{(1)}(\mathbf{0}, \mathbf{m}) \right), \quad (40)$$

where $J_{\Phi}^{(1)}(\mathbf{0}, \mathbf{m})$ is the Jacobian matrix at the first $p+p^2$ coordinates of Φ at $(\mathbf{0}, \mathbf{m})$.

By denoting \mathbf{g}_n , a Gaussian noise with zero-mean and covariance matrix $\mathbf{C}' = J_{\Phi}^{(1)}(\mathbf{0}, \mathbf{m})^\top \mathbf{C} J_{\Phi}^{(1)}(\mathbf{0}, \mathbf{m})$, the previous equation is equivalent to

$$\hat{\beta}_k = \beta_k^* + \Phi(\mathbf{S}_n, \mathbf{m}) \approx \beta_k^* + \Phi(\mathbf{0}, \mathbf{m}) + \mathbf{g}_n/\sqrt{n}. \quad (41)$$

For n large, the expansions of the first moments is of the form:

$$\mathbb{E}(\hat{\beta}_k) \approx \beta_k^* + \Phi(\mathbf{0}, \mathbf{m}). \quad (42)$$

and, the asymptotic variance is also of the form:

$$\mathbb{V}(\hat{\beta}_k) \approx \frac{1}{n} J_{\Phi}^{(1)}(\mathbf{0}, \mathbf{m})^\top \mathbf{C} J_{\Phi}^{(1)}(\mathbf{0}, \mathbf{m}). \quad (43)$$

This result holds whether the nuisance parameters in A_t and B_t are well-specified or not, so there is no guarantee that $\Phi(\mathbf{0}, \mathbf{m}) = \mathbf{0}$ and the estimator $\hat{\beta}_k$ may be biased.

In the following, we assume that the nuisance parameters in A_t and B_t are well-specified i.e. $\mathbb{E}(\psi_k(T, \mathbf{X}) \mid \mathbf{X} = \mathbf{x}) = 0$ in such way that $\mathbb{E}(Z_k \mid \mathbf{X} = \mathbf{x}) = \tau_k(\mathbf{x})$, or equivalently, $\mathbb{E}(\mathbf{H}^\top \tilde{\epsilon}_k) = \mathbf{0}$. Consequently, the estimator of $\hat{\beta}_k$ is unbiased. In this case, computing the variance $\mathbb{V}(\hat{\beta}_k)$ becomes much easier and more explicit.

On the one hand, by the multivariate Central Theorem Limit (CTL)

$$\frac{1}{\sqrt{n}} \mathbf{H}^\top \tilde{\epsilon}_k \xrightarrow{\mathcal{L}} \mathcal{N}(\mathbf{0}, \Sigma) \quad (44)$$

which is equivalent to

$$\frac{1}{\sqrt{n}} \mathbf{H}^\top \tilde{\epsilon}_k \approx \mathbf{g}_n, \quad (45)$$

where \mathbf{g}_n is a Gaussian noise with zero-mean and covariance matrix of Σ with entries

$$\begin{aligned} \Sigma_{jj'} &= \mathbb{E}[f_j(\mathbf{X})f_{j'}(\mathbf{X})(\psi_k(T, \mathbf{X}) + A_{t_k}(T, \mathbf{X})\epsilon)^2] \\ &= \mathbb{E}(f_j(\mathbf{X})f_{j'}(\mathbf{X})\psi_k^2(T, \mathbf{X})) + \sigma^2 \mathbb{E}(f_j(\mathbf{X})f_{j'}(\mathbf{X})A_{t_k}^2(T, \mathbf{X})). \end{aligned} \quad (46)$$

On the other hand, by the law of large numbers (LLN), we have $1/n(\mathbf{H}^\top \mathbf{H}) \xrightarrow{a.s.} \mathbf{F}$, thus $1/n(\mathbf{H}^\top \mathbf{H}) \xrightarrow{P} \mathbf{F}$. Since \mathbf{F} is invertible, then

$$n(\mathbf{H}^\top \mathbf{H})^{-1} \xrightarrow{P} \mathbf{F}^{-1}, \quad (47)$$

where $\mathbf{F} = (F_{jj'})_{0 \leq j, j' \leq p-1}$ and $F_{jj'} = \mathbb{E}(f_j(\mathbf{X})f_{j'}(\mathbf{X}))$.

By Slutsky's theorem,

$$\begin{aligned} \sqrt{n}(\hat{\beta}_k - \beta_k^*) &= n(\mathbf{H}^\top \mathbf{H})^{-1} \cdot 1/\sqrt{n} \mathbf{H}^\top \tilde{\epsilon} \\ &\xrightarrow{\mathcal{L}} \mathcal{N}(\mathbf{0}, \mathbf{F}^{-1} \Sigma \mathbf{F}^{-1}). \end{aligned} \quad (48)$$

We can deduce that the asymptotic mean and variance are of the form

$$\begin{aligned} \mathbb{E}(\hat{\beta}_k) &= \beta_k^*, \\ \mathbb{V}(\hat{\beta}_k) &\approx \frac{1}{n} \mathbf{F}^{-1} \Sigma \mathbf{F}^{-1}. \end{aligned} \quad (49)$$

Step 3. Obtaining the error upper bound

The determinant of the variance matrix, also known as the generalized variance by Wilks (1932; 1967) is usually used as a scalar measure of overall multidimensional scatter and can be useful to compare the variance of each meta-learner.

In our case, comparing the generalized variance is equivalent to comparing $\det\left(\frac{1}{n}\Sigma\right)$ of each pseudo-outcome meta-learner since

$$\det(\mathbb{V}(\hat{\beta}_k)) = (\det \mathbf{F}^{-1})^2 \det\left(\frac{1}{n}\Sigma\right) = \frac{1}{(\det \mathbf{F})^2} \det\left(\frac{1}{n}\Sigma\right), \quad (50)$$

with, obviously, $\det(\Sigma) > 0$ because Σ is symmetric positive definite.

In some cases, the polynomials f_j are chosen to be orthonormal with respect to the distribution of \mathbf{X} (e.g. Polynomials Chaos (Sudret, 2008)). A consequence of this choice implies that \mathbf{F} is the identity matrix. Therefore, in the following, we focus on computing and bounding Σ terms. The assumptions (3.1-5.2) and the following lemma will be used for this purpose.

Lemma B.1. *If X_1, \dots, X_m is a sequence of random variables and $b > 1$, then*

$$\begin{aligned} \left| \mathbb{E}\left[\left(\sum_{i=1}^m X_i\right)^2\right] \right| &\leq m \sum_{i=1}^m \mathbb{E}[|X_i^2|], \\ \left| \mathbb{E}\left[\left(\sum_{i=1}^m X_i\right)^b\right] \right| &\leq m^{(b-1)} \sum_{i=1}^m \mathbb{E}[|X_i^b|]. \end{aligned} \quad (51)$$

Proof. The first inequality is obtained by Cauchy-Schwarz, whereas the second inequality can be proved by Jensen inequality. Indeed, for $b > 1$, the function $x \mapsto x^b$ is convex for $x > 0$ and

$$\left| \frac{\sum_{i=1}^m X_i}{m} \right|^b \leq \frac{\sum_{i=1}^m |X_i|^b}{m}. \quad (52)$$

Therefore,

$$\left| \mathbb{E} \left[\left(\sum_{i=1}^m X_i \right)^b \right] \right| \leq \mathbb{E} \left[\left| \sum_{i=1}^m X_i \right|^b \right] \leq m^{(b-1)} \sum_{i=1}^m \mathbb{E} \left[|X_i|^b \right]. \quad (53)$$

□

In the following and by Assumption 5.1, $f_j(\mathbf{X}) \in L^a$ i.e. $f_j(\mathbf{X})$ has all possible finite moments for all $j \in \{0, \dots, p-1\}$. Moreover, there exists $C > 0$ such that:

$$\forall t \in \mathcal{T}, \forall \mathbf{x} \in \mathcal{D} : |f(t, \mathbf{x})| \leq C. \quad (54)$$

B.1.1. ERROR ESTIMATION OF THE M-LEARNER

Let $a, b > 1$ such that $1/a + 1/b = 1$. We denote $\delta_{jj'}^{(a)} = |\mathbb{E}(f_j^a(\mathbf{X})f_{j'}^a(\mathbf{X}))|^{1/a}$. By Hölder inequality we show that for the M-learner:

$$\begin{aligned} |\mathbb{E}(f_j(\mathbf{X})f_{j'}(\mathbf{X})\psi_k^2(T, \mathbf{X}))| &\leq |\mathbb{E}(f_j^a(\mathbf{X})f_{j'}^a(\mathbf{X}))|^{1/a} \cdot |\mathbb{E}(\psi_k^{2b}(T, \mathbf{X}))|^{1/b} \quad (\text{Hölder}) \\ &\leq \delta_{jj'}^{(a)} \left(2^{2b-1} \mathbb{E} \left[\left(\frac{\mathbf{1}\{T=t_k\}}{r(t_k, \mathbf{X})} - 1 \right)^{2b} f^{2b}(t_k, \mathbf{X}) \right. \right. \\ &\quad \left. \left. + \left(\frac{\mathbf{1}\{T=t_k\}}{r(t_k, \mathbf{X})} - 1 \right)^{2b} f^{2b}(t_k, \mathbf{X}) \right] \right)^{1/b} \\ &\quad (\text{Lemma B.1 with } m=2) \\ &\leq 2^{(2b-1)/b} \delta_{jj'}^{(a)} \left(\mathbb{E} \left[2^{2b-1} \left(\frac{\mathbf{1}\{T=t_k\}}{r^{2b}(t, \mathbf{X})} + 1 \right) f^{2b}(t_k, \mathbf{X}) \right] \right. \\ &\quad \left. + \mathbb{E} \left[2^{2b-1} \left(\frac{\mathbf{1}\{T=t_k\}}{r^{2b}(t_k, \mathbf{X})} + 1 \right) f^{2b}(t_k, \mathbf{X}) \right] \right)^{1/b} \quad (\text{Lemma B.1}) \\ &\leq 2^{2(2b-1)/b} \delta_{jj'}^{(a)} \left(\mathbb{E} \left[\left(\frac{\mathbf{1}\{T=t_k\}}{r^{2b}(t, \mathbf{X})} + 1 \right) f^{2b}(t_k, \mathbf{X}) \right] \right. \\ &\quad \left. + \mathbb{E} \left[\left(\frac{\mathbf{1}\{T=t_k\}}{r^{2b}(t_k, \mathbf{X})} + 1 \right) f^{2b}(t_k, \mathbf{X}) \right] \right)^{1/b} \quad (55) \\ &\leq 2^{2(2b-1)/b} \delta_{jj'}^{(a)} \left(\mathbb{E} \left[\left(\frac{1}{r^{2b-1}(t, \mathbf{X})} + 1 \right) f^{2b}(t_k, \mathbf{X}) \right] \right. \\ &\quad \left. + \mathbb{E} \left[\left(\frac{1}{r^{2b-1}(t_k, \mathbf{X})} + 1 \right) f^{2b}(t_k, \mathbf{X}) \right] \right)^{1/b} \\ &\leq 2^{2(2b-1)/b} \delta_{jj'}^{(a)} \left(\frac{1}{r_{\min}^{2b-1}} + 1 \right)^{1/b} (C^{2b} + C^{2b})^{1/b} \quad (\text{Bounding } r \text{ and } f) \\ &\leq 2^{2(2b-1)/b} \delta_{jj'}^{(a)} \left(\frac{1}{r_{\min}^{2b-1}} + \frac{1}{r_{\min}^{2b-1}} \right)^{1/b} 2^{1/b} C^b \\ &\leq 2^{2(2b-1)/b} \delta_{jj'}^{(a)} \frac{2^{1/b}}{r_{\min}^{(2b-1)/b}} 2^{1/b} C^b \\ &\leq 2^4 \delta_{jj'}^{(a)} \frac{1}{r_{\min}^{(2b-1)/b}} C^b = \frac{16}{r_{\min}^{(2b-1)/b}} \delta_{jj'}^{(a)} C^b. \end{aligned}$$

On the other term, one obtains similarly:

$$\begin{aligned}
 |\mathbb{E}(f_j(\mathbf{X})f_{j'}(\mathbf{X})A_{t_k}^2(T, \mathbf{X}))| &\leq |\mathbb{E}(f_j^a(\mathbf{X})f_{j'}^a(\mathbf{X}))|^{1/a} \cdot |\mathbb{E}(A_{t_k}^{2b}(T, \mathbf{X}))|^{1/b} \quad (\text{H\"older}) \\
 &\leq \delta_{jj'}^{(a)} |\mathbb{E}(A_{t_k}^{2b}(T, \mathbf{X}))|^{1/b} \\
 &\leq \delta_{jj'}^{(a)} \left(2^{2b-1} \mathbb{E}\left(\frac{\mathbf{1}\{T=t_k\}}{r(t_k, \mathbf{X})}\right)^{2b} + \mathbb{E}\left(\frac{\mathbf{1}\{T=t_k\}}{r(t_k, \mathbf{X})}\right)^{2b} \right)^{1/b} \quad (\text{Lemma B.1}) \\
 &\leq 2^{(2b-1)/b} \sigma^2 \delta_{jj'}^{(a)} \left(\mathbb{E}\left(\frac{\mathbf{1}\{T=t_k\}}{r^{2b}(t_k, \mathbf{X})}\right) + \mathbb{E}\left(\frac{\mathbf{1}\{T=t_k\}}{r^{2b}(t_k, \mathbf{X})}\right) \right)^{1/b} \\
 &\leq 2^{(2b-1)/b} \sigma^2 \delta_{jj'}^{(a)} \left(\frac{2}{r_{\min}^{2b-1}} \right)^{1/b} = \frac{4}{r_{\min}^{(2b-1)/b}} \sigma^2 \delta_{jj'}^{(a)}.
 \end{aligned} \tag{56}$$

Thus, by combining the two terms, one gets:

$$\begin{aligned}
 \left| \Sigma_{jj'}^{(M)} \right| &\leq |\mathbb{E}(f_j(\mathbf{X})f_{j'}(\mathbf{X})\psi_k^2(T, \mathbf{X}))| + \sigma^2 |\mathbb{E}(f_j(\mathbf{X})f_{j'}(\mathbf{X})A_{t_k}^2(T, \mathbf{X}))| \\
 &\leq \frac{16}{r_{\min}^{(2b-1)/b}} \delta_{jj'}^{(a)} C^b + \frac{4}{r_{\min}^{(2b-1)/b}} \sigma^2 \delta_{jj'}^{(a)} \\
 &\leq \frac{1}{r_{\min}^{(2b-1)/b}} (16 C^b + 4\sigma^2) \delta_*^{(b)},
 \end{aligned} \tag{57}$$

where $\delta_*^{(b)} = \max_{j,j'} \left| \mathbb{E}(f_j^{b/(b-1)}(\mathbf{X})f_{j'}^{b/(b-1)}(\mathbf{X})) \right|^{(b-1)/b} = \max_{j,j'} \delta_{jj'}^{(a)}$.

Therefore, for all $\epsilon = b - 1 > 0$, there exists $C_M = 4C + \sigma^2$ such that

$$\left| \Sigma_{jj'}^{(M)} \right| \leq 4r_{\min}^{1/(1+\epsilon)-2} \delta_*^{(1+\epsilon)} C_M. \tag{58}$$

In particular, if $\epsilon \ll 1$ then $1/(1+\epsilon) - 2 \approx -(1+\epsilon)$ and

$$\left| \Sigma_{jj'}^{(M)} \right| \leq \frac{4}{r_{\min}^{1+\epsilon}} \delta_*^{(1+\epsilon)} C_M \tag{59}$$

B.1.2. ERROR ESTIMATION OF THE DR-LEARNER.

In this case, we have

$$A_{t_k}(T, \mathbf{X}) = \frac{\mathbf{1}\{T=t_k\}}{r(t_k, \mathbf{X})} - \frac{\mathbf{1}\{T=t_k\}}{r(t_k, \mathbf{X})}, \tag{60}$$

$$B_{t_k}(T, \mathbf{X}) = \mu_{t_k}(\mathbf{X}) - \mu_{t_k}(\mathbf{X}) - \left(\frac{\mathbf{1}\{T=t_k\}}{r(t_k, \mathbf{X})} - \frac{\mathbf{1}\{T=t_k\}}{r(t_k, \mathbf{X})} \right) \mu_T(\mathbf{X}). \tag{61}$$

We need just to compute the upper bound of $\mathbb{E}(f_j(\mathbf{X})f_{j'}(\mathbf{X})\psi_k^2(T, \mathbf{X}))$ such that

$$\begin{aligned}
 \psi_k(T, \mathbf{X}) &= A_{t_k}(T, \mathbf{X})f(T, \mathbf{X}) - \tau_k(\mathbf{x}) + B_{t_k}(T, \mathbf{X}) \\
 &= \left(\frac{\mathbf{1}\{T=t_k\}}{r(t_k, \mathbf{X})} - 1 \right) f(t_k, \mathbf{X}) - \left(\frac{\mathbf{1}\{T=t_k\}}{r(t_k, \mathbf{X})} - 1 \right) f(t_k, \mathbf{X}) + \mu_{t_k}(\mathbf{X}) \left(1 - \frac{\mathbf{1}\{T=t_k\}}{r(t_k, \mathbf{X})} \right) \\
 &\quad - \mu_{t_k}(\mathbf{X}) \left(1 - \frac{\mathbf{1}\{T=t_k\}}{r(t_k, \mathbf{X})} \right) \\
 &= \left(\frac{\mathbf{1}\{T=t_k\}}{r(t_k, \mathbf{X})} - 1 \right) (f(t_k, \mathbf{X}) - \mu_{t_k}(\mathbf{X})) - \left(\frac{\mathbf{1}\{T=t_k\}}{r(t_k, \mathbf{X})} - 1 \right) (f(t_k, \mathbf{X}) - \mu_{t_k}(\mathbf{X}))
 \end{aligned} \tag{62}$$

Similarly to the previous calculus, we show that for the DR-learner

$$\begin{aligned}
 |\mathbb{E}(f_j(\mathbf{X})f_{j'}(\mathbf{X})\psi_k^2(T, \mathbf{X}))| &\leq |\mathbb{E}(f_j^a(\mathbf{X})f_{j'}^a(\mathbf{X}))|^{1/a} \cdot |\mathbb{E}(\psi_k^{2b}(T, \mathbf{X}))|^{1/b} \quad (\text{Hölder}) \\
 &\leq \delta_{jj'}^{(a)} \left(2^{2b-1} \mathbb{E} \left[\left(\frac{\mathbf{1}\{T=t_k\}}{r(t_k, \mathbf{X})} - 1 \right)^{2b} (f(t_k, \mathbf{X}) - \mu_{t_k}(\mathbf{X}))^{2b} \right. \right. \\
 &\quad \left. \left. + \left(\frac{\mathbf{1}\{T=t_k\}}{r(t_k, \mathbf{X})} - 1 \right)^{2b} (f(t_k, \mathbf{X}) - \mu_{t_k}(\mathbf{X}))^{2b} \right] \right)^{1/b} \quad (\text{Lemma B.1}) \\
 &\leq 2^{(2b-1)/b} \delta_{jj'}^{(a)} \left(\mathbb{E} \left[\left(\frac{\mathbf{1}\{T=t_k\}}{r(t_k, \mathbf{X})} - 1 \right)^{2b} (f(t_k, \mathbf{X}) - \mu_{t_k}(\mathbf{X}))^{2b} \right] \right. \\
 &\quad \left. + \mathbb{E} \left[\left(\frac{\mathbf{1}\{T=t_k\}}{r(t_k, \mathbf{X})} - 1 \right)^{2b} (f(t_k, \mathbf{X}) - \mu_{t_k}(\mathbf{X}))^{2b} \right] \right)^{1/b} \\
 &\leq 2^{(2b-1)/b} \delta_{jj'}^{(a)} \left(\mathbb{E} \left[2^{2b-1} \left(\frac{\mathbf{1}\{T=t_k\}}{r^{2b}(t_k, \mathbf{X})} + 1 \right) (f(t_k, \mathbf{X}) - \mu_{t_k}(\mathbf{X}))^{2b} \right] \right. \\
 &\quad \left. + \mathbb{E} \left[2^{2b-1} \left(\frac{\mathbf{1}\{T=t_k\}}{r^{2b}(t_k, \mathbf{X})} + 1 \right) (f(t_k, \mathbf{X}) - \mu_{t_k}(\mathbf{X}))^{2b} \right] \right)^{1/b} \quad (\text{Lemma B.1}) \quad (63) \\
 &\leq 2^{2(2b-1)/b} \delta_{jj'}^{(a)} \left(\mathbb{E} \left[\left(\frac{1}{r^{2b-1}(t_k, \mathbf{X})} + 1 \right) (f(t_k, \mathbf{X}) - \mu_{t_k}(\mathbf{X}))^{2b} \right] \right. \\
 &\quad \left. + \mathbb{E} \left[\left(\frac{1}{r^{2b-1}(t_k, \mathbf{X})} + 1 \right) (f(t_k, \mathbf{X}) - \mu_{t_k}(\mathbf{X}))^{2b} \right] \right)^{1/b} \\
 &\leq 2^{2(2b-1)/b} \delta_{jj'}^{(a)} \left(\frac{1}{r_{\min}^{(2b-1)/b}} + 1 \right) \left(\mathbb{E} \left[(f(t_k, \mathbf{X}) - \mu_{t_k}(\mathbf{X}))^{2b} \right] \right. \\
 &\quad \left. + \mathbb{E} \left[(f(t_k, \mathbf{X}) - \mu_{t_k}(\mathbf{X}))^{2b} \right] \right)^{1/b} \\
 &\leq 2^{2(2b-1)/b} \delta_{jj'}^{(a)} \left(\frac{1}{r_{\min}^{(2b-1)/b}} + 1 \right) \left[\left(\mathbb{E}(f(t_k, \mathbf{X}) - \mu_{t_k}(\mathbf{X}))^{2b} \right)^{1/b} \right. \\
 &\quad \left. + \mathbb{E}(f(t_k, \mathbf{X}) - \mu_{t_k}(\mathbf{X}))^{2b} \right]^{1/b} \quad (\text{Subadditivity of } |\mathbf{X}|^{1/b})
 \end{aligned}$$

Hence,

$$\begin{aligned}
 |\Sigma_{jj'}^{(\text{DR})}| &\leq 2^{2(2b-1)/b} \delta_{jj'}^{(a)} \left(\frac{1}{r_{\min}^{(2b-1)/b}} + 1 \right) \left[\left(\mathbb{E}(f(t_k, \mathbf{X}) - \mu_{t_k}(\mathbf{X}))^{2b} \right)^{1/b} \right. \\
 &\quad \left. + \left(\mathbb{E}(f(t_k, \mathbf{X}) - \mu_{t_k}(\mathbf{X}))^{2b} \right)^{1/b} \right] + \frac{4}{r_{\min}^{(2b-1)/b}} \sigma^2 \delta_{jj'}^{(a)} \\
 &\leq 2^{2(2b-1)/b} \delta_*^{(b)} \left(\frac{1}{r_{\min}^{(2b-1)/b}} + 1 \right) \left[\left(\mathbb{E}(f(t_k, \mathbf{X}) - \mu_{t_k}(\mathbf{X}))^{2b} \right)^{1/b} \right. \\
 &\quad \left. + \left(\mathbb{E}(f(t_k, \mathbf{X}) - \mu_{t_k}(\mathbf{X}))^{2b} \right)^{1/b} \right] + \frac{4}{r_{\min}^{(2b-1)/b}} \sigma^2 \delta_*^{(b)} \quad (64)
 \end{aligned}$$

We consider now $\epsilon = b - 1 > 0$, and we assume that $\epsilon \ll 1$, then

$$\begin{aligned}
 &2^{2(2b-1)/b} \delta_*^{(b)} \left(\frac{1}{r_{\min}^{(2b-1)/b}} + 1 \right) \left[\left(\mathbb{E}(f(t_k, \mathbf{X}) - \mu_{t_k}(\mathbf{X}))^{2b} \right)^{1/b} + \left(\mathbb{E}(f(t_k, \mathbf{X}) - \mu_{t_k}(\mathbf{X}))^{2b} \right)^{1/b} \right] \\
 &+ \frac{4}{r_{\min}^{(2b-1)/b}} \sigma^2 \delta_*^{(b)} \approx 4 \delta_*^{(1+\epsilon)} \left(\frac{1}{r_{\min}^{1+\epsilon}} + 1 \right) \left(\mathbb{E}(f(t_k, \mathbf{X}) - \mu_{t_k}(\mathbf{X}))^{2b} + \mathbb{E}(f(t_k, \mathbf{X}) - \mu_{t_k}(\mathbf{X}))^{2b} \right) \\
 &\quad + 4 \sigma^2 \delta_*^{(1+\epsilon)} \frac{1}{r_{\min}^{1+\epsilon}}. \quad (65)
 \end{aligned}$$

Consequently,

$$\left| \sum_{jj'}^{(\text{DR})} \right| \leq 4 \left(\frac{C_{DR}^* + \sigma^2}{r_{\min}^{1+\epsilon}} + C_{DR}^* \right) \delta_*^{(1+\epsilon)}, \quad (66)$$

where $C_{DR}^* = \mathbb{E}(f(t_k, \mathbf{X}) - \mu_{t_k}(\mathbf{X}))^2 + \mathbb{E}(f(t_k, \mathbf{X}) - \mu_{t_k}(\mathbf{X}))^2 = \text{err}(\mu_{t_k}) + \text{err}(\mu_{t_k})$.

B.1.3. ERROR ESTIMATION OF THE X-LEARNER.

In this case, we have

$$A_{t_k}(T, \mathbf{X}) = 2 \times \mathbf{1}\{T = t_k\} - 1, \quad (67)$$

$$B_{t_k}(T, \mathbf{X}) = (1 - \mathbf{1}\{T = t_k\})\mu_{t_k}(\mathbf{X}) - \mu_{t_k}(\mathbf{X}) + \sum_{l \neq k} \mathbf{1}\{T = t_l\}\mu_{t_l}(\mathbf{X}). \quad (68)$$

One can write ψ_k as

$$\begin{aligned} \psi_k(T, \mathbf{X}) &= A_{t_k}(T, \mathbf{X})f(T, \mathbf{X}) - \tau_k(\mathbf{x}) + B_{t_k}(T, \mathbf{X}) \\ &= (2 \mathbf{1}\{T = t_k\} - 1)f(T, \mathbf{X}) - (f(t_k, \mathbf{X}) - f(t_k, \mathbf{X})) + (1 - \mathbf{1}\{T = t_k\}) \\ &\quad \mu_{t_k}(\mathbf{X}) - \mu_{t_k}(\mathbf{X}) + \sum_{l \neq k} \mathbf{1}\{T = t_l\}\mu_{t_l}(\mathbf{X}) \\ &= (1 - \mathbf{1}\{T = t_k\})(\mu_{t_k}(\mathbf{X}) - f(t_k, \mathbf{X})) - (\mu_{t_k}(\mathbf{X}) - f(t_k, \mathbf{X})) \\ &\quad + \sum_{l \neq k} \mathbf{1}\{T = t_l\}(\mu_{t_l}(\mathbf{X}) - f(t_l, \mathbf{X})) = a_k + \sum_{l \neq k} b_l. \end{aligned} \quad (69)$$

where

$$a_k = (1 - \mathbf{1}\{T = t_k\})(\mu_{t_k}(\mathbf{X}) - f(t_k, \mathbf{X})) - (\mu_{t_k}(\mathbf{X}) - f(t_k, \mathbf{X})), \quad (70)$$

$$b_l = \mathbf{1}\{T = t_l\}(\mu_{t_l}(\mathbf{X}) - f(t_l, \mathbf{X})). \quad (71)$$

Similarly to the M- and DR-learners calculus, and using lemma B.1:

$$\begin{aligned}
 |\mathbb{E}(f_j(\mathbf{X})f_{j'}(\mathbf{X})\psi_k^2(T, \mathbf{X}))| &\leq |\mathbb{E}(f_j^a(\mathbf{X})f_{j'}^a(\mathbf{X}))|^{1/a} \cdot |\mathbb{E}(\psi_k^{2b}(T, \mathbf{X}))|^{1/b} \\
 &\leq \delta_{jj'}^{(a)} \left| \mathbb{E}(a_t + \sum_{l \neq k} b_l)^{2b} \right|^{1/b} \quad (\text{Hölder}) \\
 &\leq \delta_{jj'}^{(a)} \left(2^{2b-1} \left(\mathbb{E}(a_t^{2b}) + \mathbb{E}(\sum_{l \neq k} b_l)^{2b} \right) \right)^{1/b} \quad (\text{Lemma B.1 with } m = 2) \\
 &\leq 2^{(2b-1)/b} \delta_{jj'}^{(a)} \left(\mathbb{E}(a_t^{2b}) + \mathbb{E}(\sum_{l \neq k} b_l)^{2b} \right)^{1/b} \\
 &\leq 2^{(2b-1)/b} \delta_{jj'}^{(a)} \left[2^{2b-1} \left(\mathbb{E}((1 - \mathbf{1}\{T = t_k\})^{2b} (\mu_{t_k}(\mathbf{X}) - f(t_k, \mathbf{X}))^{2b}) \right. \right. \\
 &\quad \left. \left. + \mathbb{E}(\mu_{t_k}(\mathbf{X}) - f(t_k, \mathbf{X}))^{2b} \right) + (K - 1)^{2b-1} \right. \\
 &\quad \left. \times \sum_{l \neq k} \mathbb{E}(\mathbf{1}\{T = t_l\} (\mu_{t_l}(\mathbf{X}) - f(t_l, \mathbf{X}))^{2b}) \right]^{1/b} \\
 & \quad (\text{Lemma B.1 with } m = 2 \text{ on the 1st term, and } m = (K - 1) \text{ on the 2nd term}) \\
 &\leq 2^{(2b-1)/b} \delta_{jj'}^{(a)} \left[2^{2b-1} \left(\mathbb{E}(\mu_{t_k}(\mathbf{X}) - f(t_k, \mathbf{X}))^{2b} + \mathbb{E}(\mu_{t_k}(\mathbf{X}) - f(t_k, \mathbf{X}))^{2b} \right) \right. \\
 &\quad \left. + (K - 1)^{2b-1} \sum_{l \neq k} \mathbb{E}(\mu_{t_l}(\mathbf{X}) - f(t_l, \mathbf{X}))^{2b} \right]^{1/b} \\
 &\leq 2^{(2b-1)/b} \delta_{jj'}^{(a)} \left[2^{(2b-1)/b} \left(\mathbb{E}(\mu_{t_k}(\mathbf{X}) - f(t_k, \mathbf{X}))^{2b} \right)^{1/b} + 2^{(2b-1)/b} \left(\mathbb{E}(\mu_{t_k}(\mathbf{X}) \right. \right. \\
 &\quad \left. \left. - f(t_k, \mathbf{X}))^{2b} \right)^{1/b} + (K - 1)^{(2b-1)/b} \sum_{l \neq k} \left(\mathbb{E}(\mu_{t_l}(\mathbf{X}) - f(t_l, \mathbf{X}))^{2b} \right)^{1/b} \right] \\
 &\leq 2^{2(2b-1)/b} \delta_{jj'}^{(a)} \left[\left(\mathbb{E}(\mu_{t_k}(\mathbf{X}) - f(t_k, \mathbf{X}))^{2b} \right)^{1/b} + \left(\mathbb{E}(\mu_{t_k}(\mathbf{X}) - f(t_k, \mathbf{X}))^{2b} \right)^{1/b} \right. \\
 &\quad \left. + \left(\frac{K - 1}{2} \right)^{(2b-1)/b} \sum_{l \neq k} \left(\mathbb{E}(\mu_{t_l}(\mathbf{X}) - f(t_l, \mathbf{X}))^{2b} \right)^{1/b} \right].
 \end{aligned} \tag{72}$$

Given that $\mathbb{E}(f_j(\mathbf{X})f_{j'}(\mathbf{X})A_{t_k}^2(T, \mathbf{X})) = \mathbb{E}(f_j(\mathbf{X})f_{j'}(\mathbf{X})) = \delta_{jj'}^{(1)}$, we deduce finally

$$\begin{aligned}
 \left| \Sigma_{jj'}^{(X)} \right| &\leq |\mathbb{E}(f_j(\mathbf{X})f_{j'}(\mathbf{X})\psi_k^2(T, \mathbf{X}))| + \sigma^2 |\mathbb{E}(f_j(\mathbf{X})f_{j'}(\mathbf{X})A_{t_k}^2(T, \mathbf{X}))| \\
 &\leq 2^{2(2b-1)/b} \delta_{jj'}^{(a)} \left[\left(\mathbb{E}(\mu_{t_k}(\mathbf{X}) - f(t_k, \mathbf{X}))^{2b} \right)^{1/b} + \left(\mathbb{E}(\mu_{t_k}(\mathbf{X}) - f(t_k, \mathbf{X}))^{2b} \right)^{1/b} \right. \\
 &\quad \left. + \left(\frac{K - 1}{2} \right)^{(2b-1)/b} \sum_{l \neq k} \left(\mathbb{E}(\mu_{t_l}(\mathbf{X}) - f(t_l, \mathbf{X}))^{2b} \right)^{1/b} \right] + \sigma^2 \delta_{jj'}^{(1)} \\
 &\leq 2^{2(2b-1)/b} \delta_*^{(b)} \left[\left(\mathbb{E}(\mu_{t_k}(\mathbf{X}) - f(t_k, \mathbf{X}))^{2b} \right)^{1/b} + \left(\mathbb{E}(\mu_{t_k}(\mathbf{X}) - f(t_k, \mathbf{X}))^{2b} \right)^{1/b} \right. \\
 &\quad \left. + \left(\frac{K - 1}{2} \right)^{(2b-1)/b} \sum_{l \neq k} \left(\mathbb{E}(\mu_{t_l}(\mathbf{X}) - f(t_l, \mathbf{X}))^{2b} \right)^{1/b} \right] + \sigma^2 \delta_*^{(1)}
 \end{aligned} \tag{73}$$

where $\delta_*^{(1)} = \max_{j, j'} \mathbb{E}(f_j(\mathbf{X})f_{j'}(\mathbf{X}))$.

As in the previous cases, we consider now $\epsilon = b - 1 > 0$ with $\epsilon \ll 1$, then

$$\begin{aligned}
 & 2^{2(2b-1)/b} \delta_*^{(b)} \left[\left(\mathbb{E}(\mu_{t_k}(\mathbf{X}) - f(t_k, \mathbf{X}))^{2b} \right)^{1/b} + \left(\mathbb{E}(\mu_{t_k}(\mathbf{X}) - f(t_k, \mathbf{X}))^{2b} \right)^{1/b} \right. \\
 & \quad \left. + \left(\frac{K-1}{2} \right)^{(2b-1)/b} \sum_{l \neq k} \left(\mathbb{E}(\mu_{t_l}(\mathbf{X}) - f(t_l, \mathbf{X}))^{2b} \right)^{1/b} \right] + \sigma^2 \delta_*^{(1)} \\
 & \approx 4 \delta_*^{(1+\epsilon)} \left(\mathbb{E}(f(t_k, \mathbf{X}) - \mu_{t_k}(\mathbf{X}))^2 + \mathbb{E}(f(t_k, \mathbf{X}) - \mu_{t_k}(\mathbf{X}))^2 \right) \\
 & \quad + \frac{(K-1)^2}{4} \sum_{l \neq k} \mathbb{E}(\mu_{t_l}(\mathbf{X}) - f(t_l, \mathbf{X}))^2 + \sigma^2 \delta_*^{(1)}.
 \end{aligned} \tag{74}$$

Therefore,

$$\left| \Sigma_{jj'}^{(X)} \right| \leq 4 \delta_*^{(1+\epsilon)} C_X + \sigma^2 \delta_*^{(1)}. \tag{75}$$

where $C_X = \text{err}(\mu_{t_k}) + \text{err}(\mu_{t_k}) + \frac{(K-1)^2}{4} \sum_{l \neq k} \text{err}(\mu_{t_l})$.

B.2. Error estimation of the T- and naive X-learners.

In this subsection, we propose to conduct the bias-variance analysis of the T-learner and the naive extension of the X-learner. Some steps of this proof are quite similar to the proof of Appendix B.1.

B.2.1. ERROR ESTIMATION OF THE T-LEARNER.

STEP 0. SET-UP

For all $t \in \mathcal{T}$, we define the set $\mathbf{S}_t = \{i, T_i = t\}$ with $n_t = |\mathbf{S}_t|$. Under Assumptions (3.1-5.2), the T-learner of the CATE can be defined as

$$\hat{\tau}_k^{(T)}(\mathbf{x}) = \mathbf{f}(\mathbf{x})^\top \hat{\beta}_{t_k} - \mathbf{f}(\mathbf{x})^\top \hat{\beta}_{t_0} = \mathbf{f}(\mathbf{x})^\top (\hat{\beta}_{t_k} - \hat{\beta}_{t_0}), \tag{76}$$

where $\hat{\beta}_{t_k}$ and $\hat{\beta}_{t_0}$ are the OLS estimators of β_{t_k} and β_{t_0} such that:

$$\hat{\beta}_{t_k} = (\mathbf{H}_k^\top \mathbf{H}_k)^{-1} \mathbf{H}_k^\top \mathbf{y}_k, \tag{77}$$

$$\hat{\beta}_{t_0} = (\mathbf{H}_0^\top \mathbf{H}_0)^{-1} \mathbf{H}_0^\top \mathbf{y}_0, \tag{78}$$

where $\mathbf{H}_k = (f_j(\mathbf{X}_i))_{i \in \mathbf{S}_{t_k}, j} \in \mathbb{R}^{n_k \times p}$ (respectively, $\mathbf{H}_0 = (f_j(\mathbf{X}_i))_{i \in \mathbf{S}_{t_0}, j} \in \mathbb{R}^{n_0 \times p}$) is the regression matrix and $\mathbf{y}_k = (Y_{\text{obs}, i})_{i \in \mathbf{S}_{t_k}}$ (respectively, $\mathbf{y}_0 = (Y_{\text{obs}, i})_{i \in \mathbf{S}_{t_0}}$).

STEP 1. IDENTIFICATION OF $\hat{\beta}_k$

By similar calculus, we show that:

$$\begin{aligned}
 \hat{\beta}_{t_k} &= (\mathbf{H}_k^\top \mathbf{H}_k)^{-1} \mathbf{H}_k^\top \mathbf{y}_k \\
 &= (\mathbf{H}_k^\top \mathbf{H}_k)^{-1} \mathbf{H}_k^\top (f(t_k, \mathbf{X}_i) + \varepsilon_i(t_k))_{i \in \mathbf{S}_{t_k}} \\
 &= \beta_{t_k} + (\mathbf{H}_k^\top \mathbf{H}_k)^{-1} \mathbf{H}_k^\top \boldsymbol{\epsilon}_k \\
 &= \beta_{t_k} + \frac{1}{\sqrt{n_k}} \left(\frac{1}{n_k} \mathbf{H}_k^\top \mathbf{H}_k \right)^{-1} \left(\frac{1}{\sqrt{n_k}} \mathbf{H}_k^\top \boldsymbol{\epsilon}_k \right).
 \end{aligned} \tag{79}$$

where $\boldsymbol{\epsilon}_k = (\varepsilon_i(t_k))_{i=1}^n$ are i.i.d. Gaussian $\mathcal{N}(0, \sigma^2)$ and independent of $(T_i, \mathbf{X}_i)_{i=1}^n$.

Thus,

$$\sqrt{n}(\hat{\beta}_{t_k} - \beta_{t_k}) = \sqrt{\frac{n}{n_k}} \left(\frac{1}{n_k} \mathbf{H}_k^\top \mathbf{H}_k \right)^{-1} \left(\frac{1}{\sqrt{n_k}} \mathbf{H}_k^\top \boldsymbol{\epsilon}_k \right). \tag{80}$$

By similar calculus

$$\sqrt{n}(\hat{\beta}_{t_0} - \beta_{t_0}) = \sqrt{\frac{n}{n_0}} \left(\frac{1}{n_0} \mathbf{H}_0^\top \mathbf{H}_0 \right)^{-1} \left(\frac{1}{\sqrt{n_0}} \mathbf{H}_0^\top \boldsymbol{\epsilon}_0 \right). \tag{81}$$

Therefore, by considering $\widehat{\boldsymbol{\beta}}_k = \widehat{\boldsymbol{\beta}}_{t_k} - \widehat{\boldsymbol{\beta}}_{t_0}$,

$$\begin{aligned}\sqrt{n}(\widehat{\boldsymbol{\beta}}_k - \boldsymbol{\beta}_k^*) &= \sqrt{n}(\widehat{\boldsymbol{\beta}}_{t_k} - \boldsymbol{\beta}_{t_k}) + \sqrt{n}(\widehat{\boldsymbol{\beta}}_{t_0} - \boldsymbol{\beta}_{t_0}) \\ &= \sqrt{\frac{n}{n_k}} \left(\frac{1}{n_k} \mathbf{H}_k^\top \mathbf{H}_k \right)^{-1} \left(\frac{1}{\sqrt{n_k}} \mathbf{H}_k^\top \boldsymbol{\epsilon}_k \right) + \sqrt{\frac{n}{n_0}} \left(\frac{1}{n_0} \mathbf{H}_0^\top \mathbf{H}_0 \right)^{-1} \left(\frac{1}{\sqrt{n_0}} \mathbf{H}_0^\top \boldsymbol{\epsilon}_0 \right).\end{aligned}\quad (82)$$

STEP 2. THE ASYMPTOTIC BEHAVIOUR OF THE OLS ESTIMATOR'S MEAN AND COVARIANCE

Let $\mathbf{a} = (\mathbf{a}_k, \mathbf{a}_0) \in \mathbb{R}^{2p}$ and let ϕ_n denote the characteristic function of the vector $(\frac{1}{\sqrt{n_k}} \mathbf{H}_k^\top \boldsymbol{\epsilon}_k, \frac{1}{\sqrt{n_0}} \mathbf{H}_0^\top \boldsymbol{\epsilon}_0)$. We have

$$\begin{aligned}\phi_n(\mathbf{a}) &= \mathbb{E} \left[\exp i \mathbf{a}^\top \left(\frac{1}{\sqrt{n_k}} \mathbf{H}_k^\top \boldsymbol{\epsilon}_k, \frac{1}{\sqrt{n_0}} \mathbf{H}_0^\top \boldsymbol{\epsilon}_0 \right) \right] \\ &= \mathbb{E} \left[\exp i \left(\mathbf{a}_k^\top \frac{1}{\sqrt{n_k}} \mathbf{H}_k^\top \boldsymbol{\epsilon}_k + \mathbf{a}_0^\top \frac{1}{\sqrt{n_0}} \mathbf{H}_0^\top \boldsymbol{\epsilon}_0 \right) \right] \\ &= \mathbb{E} \left[\exp i \left(\frac{1}{\sqrt{n_k}} \sum_{m=1}^{n_k} \mathbf{a}_k^\top (H_{mj} \varepsilon_m(t_k))_{j=0}^{p-1} + \frac{1}{\sqrt{n_0}} \sum_{m=1}^{n_0} \mathbf{a}_0^\top (H_{mj} \varepsilon_m(t_0))_{j=0}^{p-1} \right) \right] \\ &= \mathbb{E} \left[\exp i \left(\frac{1}{\sqrt{n}} \sum_{m=1}^n \mathbf{a}_k^\top (H_{mj})_{j=0}^{p-1} \varepsilon_m(t_k) \mathbf{1}\{T_m = t_k\} \times \frac{\sqrt{n}}{\sqrt{n_k}} \right) \right. \\ &\quad \left. + \frac{1}{\sqrt{n}} \sum_{m=1}^n \mathbf{a}_0^\top (H_{mj})_{j=0}^{p-1} \varepsilon_m(t_0) \mathbf{1}\{T_m = t_0\} \times \frac{\sqrt{n}}{\sqrt{n_0}} \right].\end{aligned}\quad (83)$$

Now, let us consider the vector $\mathbf{Z}^{(n)} = (\mathbf{Z}_k^{(n)}, \mathbf{Z}_0^{(n)}) \in \mathbb{R}^{2p}$ such that

$$\begin{aligned}\mathbf{Z}^{(n)} &= \left(\frac{1}{n} \sum_{m=1}^n H_{m1} \varepsilon_m(t_k) \mathbf{1}\{T_m = t_k\}, \dots, \frac{1}{n} \sum_{m=1}^n H_{mp} \varepsilon_m(t_k) \mathbf{1}\{T_m = t_k\}, \right. \\ &\quad \left. \frac{1}{n} \sum_{m=1}^n (H_{m1} \varepsilon_m(t_0) \mathbf{1}\{T_m = t_0\}), \dots, \frac{1}{n} \sum_{m=1}^n H_{mp} \varepsilon_m(t_0) \mathbf{1}\{T_m = t_0\} \right) \\ &= \frac{1}{n} \sum_{m=1}^n \left(H_{m1} \varepsilon_m(t_k) \mathbf{1}\{T_m = t_k\}, \dots, H_{mp} \varepsilon_m(t_k) \mathbf{1}\{T_m = t_k\}, \right. \\ &\quad \left. H_{m1} \varepsilon_m(t_0) \mathbf{1}\{T_m = t_0\}, \dots, H_{mp} \varepsilon_m(t_0) \mathbf{1}\{T_m = t_0\} \right) \\ &= \frac{1}{n} \sum_{m=1}^n \mathbf{Z}_m.\end{aligned}\quad (84)$$

The mean $\mathbf{m} = (\mathbf{m}_k, \mathbf{m}_0)$ of the vector \mathbf{Z}_m satisfies, for $j = 0, \dots, p-1$,

$$m_{k,j} = \mathbb{E}[f_j(\mathbf{X}) \varepsilon(t_k) \mathbf{1}\{T = t_k\}] = 0, \quad (85)$$

$$m_{0,j} = \mathbb{E}[f_j(\mathbf{X}) \varepsilon(t_0) \mathbf{1}\{T = t_0\}] = 0, \quad (86)$$

and its covariance matrix that satisfies, for $j, j' = 1, \dots, 2p$,

$$\begin{aligned} \text{Cov}(\mathbf{Z}_{m,j}, \mathbf{Z}_{m,j'}) &= \begin{cases} \mathbb{E}[f_{j-1}(\mathbf{X})f_{j'-1}(\mathbf{X})\varepsilon^2(t_k)\mathbf{1}\{T = t_k\}] & \text{if } j, j' \in \{1, \dots, p\} \\ \mathbb{E}[f_{j-1}(\mathbf{X})f_{j'-1}(\mathbf{X})\varepsilon^2(t_0)\mathbf{1}\{T = t_0\}] & \text{if } j, j' \in \{p+1, \dots, 2p\} \\ \mathbb{E}[f_{j-1}(\mathbf{X})f_{j'-1}(\mathbf{X})\varepsilon(t_k)\varepsilon(t_0)\mathbf{1}\{T = t_k\}\mathbf{1}\{T = t_0\}] & \text{otherwise.} \end{cases} \\ &= \begin{cases} \sigma^2\mathbb{E}[f_{j-1}(\mathbf{X})f_{j'-1}(\mathbf{X})\mathbf{1}\{T = t_k\}] & \text{if } j, j' \in \{1, \dots, p\} \\ \sigma^2\mathbb{E}[f_{j-1}(\mathbf{X})f_{j'-1}(\mathbf{X})\mathbf{1}\{T = t_0\}] & \text{if } j, j' \in \{p+1, \dots, 2p\} \\ 0 & \text{otherwise,} \end{cases} \\ &= \begin{cases} \sigma^2\rho(t_k)\mathbf{F}_{k,jj'} & \text{if } j, j' \in \{1, \dots, p\}, \\ \sigma^2\rho(t_0)\mathbf{F}_{0,jj'} & \text{if } j, j' \in \{p+1, \dots, 2p\}, \\ 0 & \text{otherwise,} \end{cases} \end{aligned} \quad (87)$$

where the matrices $\mathbf{F}_k = (\mathbb{E}[f_{j-1}(\mathbf{X})f_{j'-1}(\mathbf{X}) | T = t_k])_{j,j'} \in \mathbb{R}^p$ and $\mathbf{F}_0 = (\mathbb{E}[f_{j-1}(\mathbf{X})f_{j'-1}(\mathbf{X}) | T = t_0])_{j,j'} \in \mathbb{R}^p$ are supposed to be invertible. Note that, for an integrable function h , $\mathbb{E}[h(\mathbf{X})\mathbf{1}\{T = t\}] = \mathbb{P}(T = t)\mathbb{E}[h(\mathbf{X}) | T = t]$.

Therefore, using the multivariate CLT on $\mathbf{Z}^{(n)}$, we get

$$\begin{aligned} \begin{pmatrix} \sqrt{n} \mathbf{a}_k^\top \mathbf{Z}_k^{(n)} \\ \sqrt{n} \mathbf{a}_0^\top \mathbf{Z}_0^{(n)} \end{pmatrix} &= \begin{pmatrix} \frac{1}{\sqrt{n}} \sum_{m=1}^n \mathbf{a}_k^\top (H_{mj})_{j=0}^{p-1} \varepsilon_m(t_k) \mathbf{1}\{T_m = t_k\} \\ \frac{1}{\sqrt{n}} \sum_{m=1}^n \mathbf{a}_0^\top (H_{mj})_{j=0}^{p-1} \varepsilon_m(t_0) \mathbf{1}\{T_m = t_0\} \end{pmatrix} \\ &\xrightarrow{\mathcal{L}} \mathcal{N} \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma^2\rho(t_k)\mathbf{a}_k^\top \mathbf{F}_k \mathbf{a}_k & 0 \\ 0 & \sigma^2\rho(t_0)\mathbf{a}_0^\top \mathbf{F}_0 \mathbf{a}_0 \end{pmatrix} \right). \end{aligned} \quad (88)$$

On the other hand,

$$\left(\frac{\sqrt{n}}{\sqrt{n_k}}, \frac{\sqrt{n}}{\sqrt{n_0}} \right) = \left(\frac{\sqrt{n}}{\sqrt{\sum_{m=1}^n \mathbf{1}\{T_m = t_k\}}}, \frac{\sqrt{n}}{\sqrt{\sum_{m=1}^n \mathbf{1}\{T_m = t_0\}}} \right) \xrightarrow{a.s.} \left(\frac{1}{\sqrt{\rho(t_k)}}, \frac{1}{\sqrt{\rho(t_0)}} \right), \quad (89)$$

where $\rho(t) = \mathbb{P}(T = t)$.

Thus, by the Slutsky theorem:

$$\begin{aligned} \frac{1}{\sqrt{n}} \sum_{m=1}^n (\mathbf{a}_k^\top (H_{mj})_{j=0}^{p-1} \varepsilon_m(t_k) \mathbf{1}\{T_m = t_k\}) \frac{\sqrt{n}}{\sqrt{n_k}} + \mathbf{a}_0^\top (H_{mj})_{j=0}^{p-1} \varepsilon_m(t_0) \mathbf{1}\{T_m = t_0\} \frac{\sqrt{n}}{\sqrt{n_0}} \\ \xrightarrow{\mathcal{L}} \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{a}_k^\top \mathbf{F}_k \mathbf{a}_k + \sigma^2 \mathbf{a}_0^\top \mathbf{F}_0 \mathbf{a}_0). \end{aligned} \quad (90)$$

Therefore,

$$\phi_n(\mathbf{a}) \xrightarrow{n \rightarrow +\infty} \mathbb{E} \left[\exp i \left(\mathbf{a}_k^\top \sigma^2 \rho(t_k) \mathbf{F}_k \mathbf{a}_k + \mathbf{a}_0^\top \sigma^2 \rho(t_0) \mathbf{F}_0 \mathbf{a}_0 \right) \right] = \phi(\mathbf{z}_k, \mathbf{z}_0)(\mathbf{a}), \quad (91)$$

where \mathbf{z}_k and \mathbf{z}_0 are two independent zero-mean random vectors with covariance matrices $\sigma^2 \mathbf{F}_k$ and $\sigma^2 \mathbf{F}_0$ respectively.

As shown previously in Appendix B.1, we can prove immediately that $(1/n_k \mathbf{H}_k^\top \mathbf{H}_k)^{-1} \xrightarrow{P} \mathbf{F}_k^{-1}$. Moreover, $n_k/n \xrightarrow{a.s.} \rho(t_k)$ so $n_k/n \xrightarrow{P} \rho(t_k)$. Thus

$$\sqrt{\frac{n}{n_k}} \left(\frac{1}{n_k} \mathbf{H}_k^\top \mathbf{H}_k \right)^{-1} \xrightarrow{P} \frac{1}{\sqrt{\rho(t_k)}} \mathbf{F}_k^{-1}. \quad (92)$$

Finally, given Equation (82) and using the Slutsky theorem, we get

$$\begin{aligned} \sqrt{n}(\hat{\beta}_k - \beta_k^*) &\xrightarrow{\mathcal{L}} \mathcal{N} \left(\mathbf{0}, \frac{1}{\rho(t_k)} \mathbf{F}_k^{-1} \sigma^2 \mathbf{F}_k \mathbf{F}_k^{-1} + \frac{1}{\rho(t_0)} \mathbf{F}_0^{-1} \sigma^2 \mathbf{F}_0 \mathbf{F}_0^{-1} \right) \\ &= \mathcal{N} \left(\mathbf{0}, \frac{\sigma^2}{\rho(t_k)} \mathbf{F}_k^{-1} + \frac{\sigma^2}{\rho(t_0)} \mathbf{F}_0^{-1} \right). \end{aligned} \quad (93)$$

Here also, we can deduce that the asymptotic mean and covariance matrix are of the form

$$\begin{aligned}\mathbb{E}(\widehat{\boldsymbol{\beta}}_k) &= \boldsymbol{\beta}_{t_k} - \boldsymbol{\beta}_{t_0} = \boldsymbol{\beta}_k^*, \\ \mathbb{V}(\widehat{\boldsymbol{\beta}}_k) &\approx \frac{1}{n} \left(\frac{1}{\rho(t_k)} \mathbf{F}_k^{-1} + \frac{1}{\rho(t_0)} \mathbf{F}_0^{-1} \right) \sigma^2.\end{aligned}\quad (94)$$

STEP 3. OBTAINING THE ERROR UPPER BOUND

The asymptotic covariance matrix is given by the matrices \mathbf{F}_k^{-1} and \mathbf{F}_0^{-1} . We assume that the polynomials f_j are chosen to be orthonormal, and that, conditionally to T , their distribution is not significantly different. One can anticipate, therefore, that $\mathbf{F}_k, \mathbf{F}_0 \approx \mathbf{F}$ and easily identify the error's upper bound of the T-learner as:

$$\frac{1}{\rho(t_k)} + \frac{1}{\rho(t_0)}.\quad (95)$$

B.3. Error estimation of the naive X-learner.

Let \bar{r} denote a fixed arbitrary estimator of the GPS (see remark 3.3) and respecting the assumption 3.2, that is, $r_{\min} \leq \bar{r}(t, \mathbf{x})$. Let $\widehat{\mu}_{t_k}$ denote the estimator of μ_{t_k} . The model $\widehat{\mu}_{t_k}$ is trained using the sample \mathbf{S}_{t_k} , the OLS estimator $\widehat{\boldsymbol{\beta}}_{t_k}$ satisfies

$$\begin{aligned}\widehat{\boldsymbol{\beta}}_{t_k} &= (\mathbf{H}_k^\top \mathbf{H}_k)^{-1} \mathbf{H}_k^\top \mathbf{y}_k \\ &= \boldsymbol{\beta}_{t_k} + (\mathbf{H}_k^\top \mathbf{H}_k)^{-1} \mathbf{H}_k^\top \boldsymbol{\epsilon}_k\end{aligned}\quad (96)$$

where $\mathbf{y}_k = (Y_{\text{obs},i})_{i \in \mathbf{S}_{t_k}}$ and $\boldsymbol{\epsilon}_k = (\varepsilon_i(t_k))_{i \in \mathbf{S}_{t_k}}$.

Similarly, the OLS estimator of μ_{t_0} satisfies also

$$\widehat{\boldsymbol{\beta}}_{t_0} = \boldsymbol{\beta}_{t_0} + (\mathbf{H}_0^\top \mathbf{H}_0)^{-1} \mathbf{H}_0^\top \boldsymbol{\epsilon}_0,\quad (97)$$

where $\mathbf{y}_0 = (Y_{\text{obs},i})_{i \in \mathbf{S}_{t_0}}$ and $\boldsymbol{\epsilon}_0 = (\varepsilon_i(t_0))_{i \in \mathbf{S}_{t_0}}$.

We recall now the definition of the naive extension of the X-learner:

$$\widehat{r}_k^{(\text{X}, \text{nv})}(\mathbf{x}) = \frac{\bar{r}(t_k, \mathbf{x})}{\bar{r}(t_k, \mathbf{x}) + \bar{r}(t_0, \mathbf{x})} \widehat{r}^{(k)}(\mathbf{x}) + \frac{\bar{r}(t_0, \mathbf{x})}{\bar{r}(t_k, \mathbf{x}) + \bar{r}(t_0, \mathbf{x})} \widehat{r}^{(0)}(\mathbf{x}).\quad (98)$$

where the estimators $\widehat{r}^{(k)}$ and $\widehat{r}^{(0)}$ are built respectively on \mathbf{S}_{t_k} and \mathbf{S}_{t_0} by regressing $(D_i^{(k)})_{i \in \mathbf{S}_{t_k}} = (Y_i(t_k) - \widehat{\mu}_{t_0}(\mathbf{X}_i))_{i \in \mathbf{S}_{t_k}}$ and $(D_i^{(0)})_{i \in \mathbf{S}_{t_0}} = (\widehat{\mu}_{t_k}(\mathbf{X}_i) - Y_i(t_0))_{i \in \mathbf{S}_{t_0}}$ on \mathbf{X} . In the following, we denote $\widehat{r}^{(k)}(\mathbf{x}) = \mathbf{f}(\mathbf{x})^\top \widehat{\boldsymbol{\beta}}^{(k)}$ and $\widehat{r}^{(0)}(\mathbf{x}) = \mathbf{f}(\mathbf{x})^\top \widehat{\boldsymbol{\beta}}^{(0)}$. Here, $\widehat{\boldsymbol{\beta}}^{(k)}$ denotes the OLS estimator of $\widehat{r}^{(k)}$ and is given by:

$$\begin{aligned}\widehat{\boldsymbol{\beta}}^{(k)} &= (\mathbf{H}_k^\top \mathbf{H}_k)^{-1} \mathbf{H}_k^\top (Y_{\text{obs},i} - \widehat{\mu}_{t_0}(\mathbf{X}_i))_{i \in \mathbf{S}_{t_k}} \\ &= (\mathbf{H}_k^\top \mathbf{H}_k)^{-1} \mathbf{H}_k^\top (Y_{\text{obs},i} - \mathbf{f}(\mathbf{X}_i)^\top \widehat{\boldsymbol{\beta}}_{t_0})_{i \in \mathbf{S}_{t_k}} \\ &= (\mathbf{H}_k^\top \mathbf{H}_k)^{-1} \mathbf{H}_k^\top \mathbf{y}_k - (\mathbf{H}_k^\top \mathbf{H}_k)^{-1} \mathbf{H}_k^\top \mathbf{H}_k \widehat{\boldsymbol{\beta}}_{t_0} \\ &= \widehat{\boldsymbol{\beta}}_{t_k} - \widehat{\boldsymbol{\beta}}_{t_0} = \widehat{\boldsymbol{\beta}}_k,\end{aligned}\quad (99)$$

where $\widehat{\boldsymbol{\beta}}_k = \widehat{\boldsymbol{\beta}}_{t_k} - \widehat{\boldsymbol{\beta}}_{t_0}$ is the T-learner OLS estimator as given in (93).

By similar calculus, we show that

$$\begin{aligned}\widehat{\boldsymbol{\beta}}^{(0)} &= (\mathbf{H}_0^\top \mathbf{H}_0)^{-1} \mathbf{H}_0^\top (\widehat{\mu}_{t_k}(\mathbf{X}_i) - Y_{\text{obs},i})_{i \in \mathbf{S}_{t_0}} \\ &= (\mathbf{H}_0^\top \mathbf{H}_0)^{-1} \mathbf{H}_0^\top (\mathbf{f}(\mathbf{X}_i)^\top \widehat{\boldsymbol{\beta}}_{t_k} - Y_{\text{obs},i})_{i \in \mathbf{S}_{t_0}} \\ &= (\mathbf{H}_0^\top \mathbf{H}_0)^{-1} \mathbf{H}_0^\top \mathbf{H}_0 \widehat{\boldsymbol{\beta}}_{t_k} - (\mathbf{H}_0^\top \mathbf{H}_0)^{-1} \mathbf{H}_0^\top \mathbf{y}_0 \\ &= \widehat{\boldsymbol{\beta}}_{t_k} - \widehat{\boldsymbol{\beta}}_{t_0} = \widehat{\boldsymbol{\beta}}_k.\end{aligned}\quad (100)$$

It results that

$$\begin{aligned}\widehat{\tau}_k^{(X, \text{nv})}(\mathbf{x}) &= \frac{\bar{r}(t_k, \mathbf{x})}{\bar{r}(t_k, \mathbf{x}) + \bar{r}(t_0, \mathbf{x})} \mathbf{f}(\mathbf{x})^\top \widehat{\boldsymbol{\beta}}^{(k)} + \frac{\bar{r}(t_0, \mathbf{x})}{\bar{r}(t_k, \mathbf{x}) + \bar{r}(t_0, \mathbf{x})} \mathbf{f}(\mathbf{x})^\top \widehat{\boldsymbol{\beta}}^{(0)} \\ &= \left(\frac{\bar{r}(t_k, \mathbf{x})}{\bar{r}(t_k, \mathbf{x}) + \bar{r}(t_0, \mathbf{x})} + \frac{\bar{r}(t_0, \mathbf{x})}{\bar{r}(t_k, \mathbf{x}) + \bar{r}(t_0, \mathbf{x})} \right) \mathbf{f}(\mathbf{x})^\top \widehat{\boldsymbol{\beta}}_k = \mathbf{f}(\mathbf{x})^\top \widehat{\boldsymbol{\beta}}_k\end{aligned}\quad (101)$$

In the end, the naive X-learner is no more than a simple T-learner, the error's upper bound of the naive X-learner is given therefore by:

$$\sigma^2 \left(\frac{1}{\rho(t_k)} + \frac{1}{\rho(t_0)} \right). \quad (102)$$

C. Discussion about the binarized R-learner.

Another alternative to R-learning to continuous treatments is proposed by (Kaddour et al., 2021). The approach considers both Assumptions 5.1 and 5.2 on the outcome $Y(t) = \mathbf{f}(\mathbf{X})^\top \boldsymbol{\beta}_t + \varepsilon(t)$, then establishes the binarized (Robinson, 1988) decomposition such that

$$Y_{\text{obs}} - m(\mathbf{X}) = \mathbf{f}(\mathbf{X})^\top (\boldsymbol{\beta}_T - e^\beta(\mathbf{X})) + \varepsilon, \quad (103)$$

where $\varepsilon = \varepsilon(T)$, $m(\mathbf{x}) = \mathbb{E}(Y_{\text{obs}} | \mathbf{X} = \mathbf{x})$ and $e^\beta(\mathbf{x}) = \mathbb{E}(\boldsymbol{\beta}_T | \mathbf{X} = \mathbf{x})$.

Considering the mean squared error of ε as a loss function and minimizing it allows us to identify the optimal $\widehat{\boldsymbol{\beta}}$ and therefore CATEs. Given two nuisance estimators \widehat{m} and \widehat{e}^β of m and e^β , one needs to solve the following problem:

$$\widehat{\boldsymbol{\beta}} = \operatorname{argmin}_{\boldsymbol{\beta} \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n \left[(Y_{\text{obs},i} - \widehat{m}(\mathbf{X}_i)) - \mathbf{f}(\mathbf{X}_i)^\top (\boldsymbol{\beta}_{T_i} - \widehat{e}^\beta(\mathbf{X}_i)) \right]^2, \quad (104)$$

where \mathcal{F} is the space of candidate models $\boldsymbol{\beta}$. The previous problem corresponds to a classical OLS estimator and has, therefore, a unique solution.

If the space of candidate models \mathcal{F} is separable, then the optimization problem can be divided into the following sub-problems:

$$\begin{aligned}\widehat{\boldsymbol{\beta}}_{t_0} &= \operatorname{argmin} \frac{1}{n_k} \sum_{i \in \mathbf{S}_{t_0}} \left[(Y_{\text{obs},i} - \widehat{m}(\mathbf{X}_i)) - \mathbf{f}(\mathbf{X}_i)^\top (\boldsymbol{\beta}_{t_0} - \widehat{e}^\beta(\mathbf{X}_i)) \right]^2 \\ &\quad \vdots \\ \widehat{\boldsymbol{\beta}}_{t_K} &= \operatorname{argmin} \frac{1}{n_K} \sum_{i \in \mathbf{S}_{t_K}} \left[(Y_{\text{obs},i} - \widehat{m}(\mathbf{X}_i)) - \mathbf{f}(\mathbf{X}_i)^\top (\boldsymbol{\beta}_{t_K} - \widehat{e}^\beta(\mathbf{X}_i)) \right]^2.\end{aligned}$$

However, this approach does not consider the interactions between different $\widehat{\boldsymbol{\beta}}_t$ and is computationally heavy when the number of possible treatments K becomes larger. It also requires specifying the family of models \mathcal{F} and precise the dimension p for Assumption 5.2.

There are two main differences between the generalized R-learner and the binarized: 1) In the binarized R-learner, $(\widehat{\boldsymbol{\beta}}_{t_k})_{k=1}^K$ may be solved separately but using a small sample (\mathbf{S}_{t_k} instead of \mathbf{D}_{obs}); 2) The solution $(\widehat{\boldsymbol{\beta}}_{t_k})_{k=1}^K$ of the binarized R-learner is unique and is given by the OLS estimator of the binarized R-loss function.

D. Additional details about simulated analytical functions in section 6.1

In this section, we consider a treatment T with $K + 1 = 10$ possible values in $\mathcal{T} = \{t_k := \frac{k}{K}, k \in \{0, \dots, K\}\}$, drawn from a uniform distribution, and the following outcome functions.

The linear model outcome for $X \in \mathbb{R}$:

$$Y(t) | X \sim \mathcal{N}((1+t)X, \sigma^2). \quad (105)$$

The multivariate hazard rate (Imbens, 2000) outcome satisfies for $\mathbf{X} \in \mathbb{R}^5$:

$$Y(t) | \mathbf{X} \sim \mathcal{N}(t + \|\mathbf{X}\| \exp(-t\|\mathbf{X}\|), \sigma^2). \quad (106)$$

We compute the exact components of each model in the following subsections: the GPS r , the potential outcome models μ_t and the observed outcome model m .

D.1. Computing nuisance components

THE GENERALIZED PROPENSITY SCORE (GPS).

In the first design (RCT), we sample n units such that T and \mathbf{X} are independent. The true propensity score is known

$$r(t, \mathbf{X}) = \mathbb{P}(T = t) = 1/(K + 1) \text{ for } t \in \mathcal{T}. \quad (107)$$

In the second design (observational studies), we combine $K + 2$ samples in a single sample of n units. The first sample \mathbf{D}_{K+1} contains $n_{K+1} = n/2$ units where the treatment is assigned randomly: \mathbf{X} and T are independent, $\mathbb{P}(T = t) = 1/(K + 1)$, $\mathbf{X} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_5)$ when the hazard rate model is applied and $X \sim \mathcal{U}(0, 1)$ when the linear model is applied. For $k = 0, \dots, K$, the sample \mathbf{D}_k contains $n_k = n/(2(K + 1))$ units and the distribution of (\mathbf{X}, T) does not respect a RCT setting. For the linear model, the joint distribution of (X, T) is given by:

$$T = \frac{k}{K} \text{ and } X \text{ follows a uniform distribution } \mathcal{U}(I_k) \text{ with } I_k = \left[\frac{k}{K+1}, \frac{k+1}{K+1} \right). \quad (108)$$

For the hazard rate model, the joint distribution of (\mathbf{X}, T) is given by:

$$T = \frac{k}{K}, X_1 \text{ follows a truncated standardized normal distribution on } I_k = \left[q_{\frac{k}{K+1}}, q_{\frac{k+1}{K+1}} \right) \quad (109)$$

and X_j follow a standardized normal distribution $\mathcal{N}(0, 1)$ for $2 \leq j \leq 5$,

where q_α is the α -quantile of the standardized normal distribution with $q_0 = -\infty$ and $q_1 = +\infty$. This strategy of selecting preferentially only observations with certain characteristics is called *preferential selection* sampling and creates thus a selection bias on observed data.

For all $k \in \{0, \dots, K\}$, the true propensity score satisfies for the linear model:

$$r(t_k, x) = \begin{cases} \frac{K+2}{2(K+1)} & \text{if } x \in I_k, \\ \frac{1}{2(K+1)} & \text{otherwise.} \end{cases} \quad (110)$$

and, for the hazard rate model, it satisfies for $\mathbf{x} \in \mathbb{R}^5$:

$$r(t_k, \mathbf{x}) = \begin{cases} \frac{K+2}{2(K+1)} & \text{if } x_1 \in I_k, \\ \frac{1}{2(K+1)} & \text{otherwise.} \end{cases} \quad (111)$$

Proof. We show proof for the hazard rate model with normal distribution. The proof remains the same for the linear model in a non-randomized setting.

Let A be a given event, and then

$$\mathbb{P}(A) = \sum_{k=0}^{K+1} \frac{n_k}{n} \mathbb{P}_k(A), \quad (112)$$

where \mathbb{P} is the observed probability distribution of the combined sample and \mathbb{P}_k denotes the probability measure induced by (107), (109) and the unconfoundedness assumption 3.1.

Given the treatment $T = t_j$ and covariate vector $\mathbf{x} = (x, x_2, \dots, x_5)$, we have

$$\begin{aligned} r(t_j, \mathbf{x}) &= \mathbb{P}(T = t_j \mid X_1 = x) \\ &= \lim_{\delta \rightarrow 0} \mathbb{P}(T = t_j \mid X_1 \in [x, x + \delta]) \\ &= \lim_{\delta \rightarrow 0} \frac{\mathbb{P}(T = t_j, X_1 \in [x, x + \delta])}{\mathbb{P}(X_1 \in [x, x + \delta])}. \end{aligned} \quad (113)$$

On the one hand,

$$\begin{aligned} \mathbb{P}(T = t_j, X_1 \in [x, x + \delta]) &= \sum_{k=0}^{K+1} \frac{n_k}{n} \mathbb{P}_k(T = t_j, X_1 \in [x, x + \delta]) \\ &= \frac{n_j}{n} \mathbb{P}_j(T = t_j, X_1 \in [x, x + \delta]) + \frac{n_{K+1}}{n} \mathbb{P}_{K+1}(T = t_j, X_1 \in [x, x + \delta]) \\ &= \frac{n_j}{n} \mathbb{P}_j(X_1 \in [x, x + \delta]) + \frac{n_{K+1}}{n} \mathbb{P}_{K+1}(T = t_j) \mathbb{P}_{K+1}(X_1 \in [x, x + \delta]) \\ &= \frac{1}{2(K+1)} \mathbb{P}_j(X_1 \in [x, x + \delta]) + \frac{1}{2(K+1)} \mathbb{P}_{K+1}(X_1 \in [x, x + \delta]). \end{aligned} \quad (114)$$

For $x \in \mathbb{R}$, there exists a unique j_0 such that $x \in I_{j_0}$. For δ small enough, we have $[x, x + \delta] \subset I_{j_0}$ and, consequently, $[x, x + \delta] \cap I_j = \emptyset$ for all $j \neq j_0$. This implies:

$$\mathbb{P}_j(X_1 \in [x, x + \delta]) = \frac{\mathbb{P}_{K+1}(X_1 \in [x, x + \delta], X_1 \in I_j)}{\mathbb{P}_{K+1}(X_1 \in I_j)} = \frac{\mathbb{P}_{K+1}(X_1 \in [x, x + \delta])}{\mathbb{P}_{K+1}(X_1 \in I_j)} \mathbf{1}\{j = j_0\}. \quad (115)$$

Therefore,

$$\begin{aligned} \mathbb{P}(T = t_j, X_1 \in [x, x + \delta]) &= \frac{1}{2(K+1)} \mathbb{P}_{K+1}(X_1 \in [x, x + \delta]) \left(\frac{\mathbf{1}\{j = j_0\}}{\mathbb{P}_{K+1}(X_1 \in I_{j_0})} + 1 \right) \\ &= \left(\frac{1}{2} \mathbf{1}\{j = j_0\} + \frac{1}{2(K+1)} \right) \mathbb{P}_{K+1}(X_1 \in [x, x + \delta]). \end{aligned} \quad (116)$$

On the other hand,

$$\begin{aligned} \mathbb{P}(X_1 \in [x, x + \delta]) &= \sum_{k=0}^{K+1} \frac{n_k}{n} \mathbb{P}_k(X_1 \in [x, x + \delta]) \\ &= \frac{1}{2(K+1)} \sum_{k=0}^K \frac{\mathbb{P}_{K+1}(X_1 \in [x, x + \delta], X_1 \in I_k)}{\mathbb{P}_{K+1}(X_1 \in I_k)} + \frac{1}{2} \mathbb{P}_{K+1}(X_1 \in [x, x + \delta]) \\ &= \frac{1}{2(K+1)} \frac{\mathbb{P}_{K+1}(X_1 \in [x, x + \delta])}{\mathbb{P}_{K+1}(X_1 \in I_{j_0})} + \frac{1}{2} \mathbb{P}_{K+1}(X_1 \in [x, x + \delta]) \\ &= \frac{1}{2} \mathbb{P}_{K+1}(X_1 \in [x, x + \delta]) + \frac{1}{2} \mathbb{P}_{K+1}(X_1 \in [x, x + \delta]) \\ &= \mathbb{P}_{K+1}(X_1 \in [x, x + \delta]) \end{aligned} \quad (117)$$

Finally,

$$\begin{aligned}
 r(t_j, \mathbf{x}) &= \lim_{\delta \rightarrow 0} \frac{\mathbb{P}(T = t_j, X_1 \in [x, x + \delta])}{\mathbb{P}(X_1 \in [x, x + \delta])} \\
 &= \lim_{\delta \rightarrow 0} \frac{\left(\frac{1}{2}\mathbf{1}\{j = j_0\} + \frac{1}{2(K+1)}\right)\mathbb{P}_{K+1}(X_1 \in [x, x + \delta])}{\mathbb{P}_{K+1}(X_1 \in [x, x + \delta])} \\
 &= \frac{1}{2}\mathbf{1}\{j = j_0\} + \frac{1}{2(K+1)} \\
 &= \frac{(K+1)\mathbf{1}\{j = j_0\} + 1}{2(K+1)} \\
 &= \begin{cases} \frac{K+2}{2(K+1)} & \text{if } x \in I_j, \\ \frac{1}{2(K+1)} & \text{otherwise.} \end{cases}
 \end{aligned} \tag{118}$$

□

THE POTENTIAL OUTCOME MODELS.

The potential outcome models are given directly by the conditional mean. For the linear model, μ_t satisfies for all $t \in \mathcal{T}$ and $x \in [0, 1]$:

$$\mu_t(x) = (1 + t)x. \tag{119}$$

For the hazard rate model, μ_t is given by:

$$\mu_t(\mathbf{x}) = t + \|\mathbf{x}\| \exp(-t\|\mathbf{x}\|). \tag{120}$$

THE OBSERVED OUTCOME MODELS.

For the linear model, the observed outcome model m can be computed as:

$$\begin{aligned}
 m(x) &= \mathbb{E}(Y_{\text{obs}} \mid X = x) \\
 &= \mathbb{E}((1 + T)\mathbf{X} \mid X = x) \\
 &= (1 + \mathbb{E}(T \mid X = x))x \\
 &= \left(1 + \sum_{k=0}^K r(t_k, x)t_k\right)x,
 \end{aligned} \tag{121}$$

where r is given by (110).

and, for the hazard rate model, m can be computed as:

$$\begin{aligned}
 m(\mathbf{x}) &= \mathbb{E}(\mathbb{E}(Y_{\text{obs}} \mid \mathbf{X}, T) \mid \mathbf{X} = \mathbf{x}) \\
 &= \mathbb{E}(T + \|\mathbf{X}\| \exp(-T\|\mathbf{X}\|) \mid \mathbf{X} = \mathbf{x}) \\
 &= \mathbb{E}(T \mid \mathbf{X} = \mathbf{x}) + \|\mathbf{x}\| \mathbb{E}(\exp(-T\|\mathbf{X}\|) \mid \mathbf{X} = \mathbf{x}) \\
 &= \sum_{k=0}^K r(t_k, \mathbf{x})t_k + \sum_{k=0}^K \|\mathbf{x}\| r(t_k, \mathbf{x}) \exp(-t_k\|\mathbf{x}\|),
 \end{aligned} \tag{122}$$

where r is given by (111).

D.2. Additional numerical results and plots.

In this section, we present the results of different simulations and scenarios for linear and hazard rate models with $K+1 = 10$, $n = 2000$ for the linear model, and $n = 10000$ for the Hazard rate model. In the randomized setting, the sample \mathbf{D}_{obs} is sampled randomly, and the propensity score is given by (107). In a non-randomized setting, the sample \mathbf{D}_{obs} is given by preferential selection as described in Section D and the GPS is given by (110). When we say that "the models' nuisance components are exact", we replace the expression of the estimators $\hat{\mu}_t$, \hat{m} or \hat{r} with the expressions obtained in Section D.

Linear model in a randomized setting.

Table 5. **mPEHE** and **sdPEHE** for three different ML base-learners; Case where nuisance components are exact.

Meta-learner	XGBoost	RandomForest	Linear Model
M-Learner	2.23 (1.20)	2.09 (1.08)	0.087 (0.096)
DR-Learner	0.165 (0.034)	0.140 (0.033)	$9.65 (7.84) 10^{-3}$
X-Learner	0.022 (0.004)	0.029 (0.004)	$1.42 (1.46) 10^{-3}$
RLin-Learner	$10.2 (8.42) 10^{-3}$		

Table 6. **mPEHE** and **sdPEHE** for three different ML base-learners; Case where nuisance components are well-specified.

Meta-learner	XGBoost	RandomForest	Linear Model
T-Learner	0.065 (0.019)	0.041 (0.016)	$10.0 (8.37) 10^{-3}$
S-Learner	0.033 (0.018)	0.032 (0.028)	$3.03 (2.42) 10^{-3}$
NvX-Learner	0.060 (0.019)	0.037 (0.016)	$10.0 (8.37) 10^{-3}$
M-Learner	1.25 (0.610)	1.22 (0.621)	0.201 (0.191)
DR-Learner	0.068 (0.019) – 0.063 (0.020)	0.068 (0.018) – 0.068 (0.018)	$10.0 (9.14) - \mathbf{5.27 (4.36) 10^{-3}}$
X-Learner	0.063 (0.020) – 0.033 (0.017)	0.045 (0.016) – 0.061 (0.040)	$10.0 (8.37) - 3.28 (2.98) 10^{-3}$
RLin-Learner	0.135 (0.130)	0.137 (0.128)	0.073 (0.063)

For the DR- and X-learners: μ_t are estimated by T-learning (left value) or S-learning (right value).

Table 7. **mPEHE** and **sdPEHE** for three different ML base-learners; Case where the propensity score is misspecified.

Meta-learner	XGBoost	RandomForest	Linear Model
M-Learner	3.86 (2.95)	3.68 (2.80)	1.45 (0.99)
DR-Learner	0.145 (0.108)	0.245 (0.179)	0.014 (0.015)
X-Learner	0.033 (0.017)	0.061 (0.040)	$3.28 (2.98) 10^{-3}$
RLin-Learner	0.336 (0.272)	0.338 (0.277)	0.338 (0.215)

Table 8. **mPEHE** and **sdPEHE** for three different ML base-learners; Case where the outcome models are misspecified.

Meta-learner	XGBoost	RandomForest	Linear Model
M-Learner	1.25 (0.610)	1.22 (0.621)	0.201 (0.191)
DR-Learner	0.811 (0.386)	0.888 (0.378)	0.308 (0.362)
X-Learner	0.304 (0.330)	0.303 (0.330)	0.275 (0.328)
RLin-Learner	0.073 (0.062)		

Table 9. **mPEHE** and **sdPEHE** for three different ML base-learners; Case where nuisance components are misspecified.

Meta-learner	XGBoost	RandomForest	Linear Model
M-Learner	3.86 (2.95)	3.68 (2.80)	1.45 (0.99)
DR-Learner	1.87 (1.31)	2.09 (1.48)	0.828 (0.496)
X-Learner	0.304 (0.330)	0.303 (0.330)	0.275 (0.328)
RLin-Learner	0.277 (0.178)		

Linear model in non-randomized setting

Table 10. **mPEHE** and **sdPEHE** for three different ML base-learners; Case where nuisance components are exact.

Meta-learner	XGBoost	RandomForest	Linear Model
M-Learner	3.69 (1.80)	2.96 (1.47)	0.153 (0.177)
DR-Learner	0.276 (0.081)	0.206 (0.056)	10.9 (8.47) 10^{-3}
X-Learner	0.022 (0.004)	0.028 (0.004)	1.69 (1.11) 10^{-3}
RLin-Learner	11.0 (11.1) 10^{-3}		

Table 11. **mPEHE** and **sdPEHE** for three different ML base-learners; Case where nuisance components are well-specified.

Meta-learner	XGBoost	RandomForest	Linear Model
T-Learner	0.067 (0.023)	0.043 (0.016)	10.5 (10.0) 10^{-3}
RegT-Learner	0.059 (0.021)	0.042 (0.016)	13.0 (11.1) 10^{-3}
S-Learner	0.033 (0.018)	0.060 (0.055)	6.46 (5.11) 10^{-3}
NvX-Learner	0.062 (0.023)	0.039 (0.017)	9.56 (10.0) 10^{-3}
M-Learner	1.35 (0.82)	1.14 (0.72)	0.196 (0.153)
DR-Learner	0.065 (0.022) – 0.065 (0.027)	0.069 (0.026) – 0.096 (0.056)	13.0 (11.1) – 8.17 (6.30) 10^{-3}
X-Learner	0.059 (0.021) – 0.034 (0.017)	0.046 (0.016) – 0.084 (0.058)	14.7 (11.6) – 6.49 (5.23) 10^{-3}
RLin-Learner	0.155 (0.137)	0.124 (0.114)	0.108 (0.097)

For the DR- and X-learners: μ_t are estimated by T-learning (left value) or S-learning (right value).

Hazard rate model in randomized setting

Table 12. **mPEHE** and **sdPEHE** for three different ML base-learners; Case where nuisance components are exact.

Meta-learner	XGBoost	RandomForest	Linear Model
M-Learner	4.27 (1.45)	4.21 (1.28)	0.529 (0.188)
DR-Learner	0.127 (0.022)	0.144 (0.044)	0.106 (0.094)
X-Learner	0.045 (0.025)	0.087 (0.049)	0.106 (0.094)
RLin-Learner	0.107 (0.094)		

Table 13. **mPEHE** and **sdPEHE** for three different ML base-learners; Case where nuisance components are well-specified.

Meta-learner	XGBoost	RandomForest	Linear Model
T-Learner	0.175 (0.046)	0.263 (0.144)	0.113 (0.091)
S-Learner	0.159 (0.048)	0.260 (0.130)	0.662 (0.421)
NvX-Learner	0.176 (0.091)	0.313 (0.188)	0.113 (0.092)
M-Learner	1.57 (0.471)	1.79 (0.453)	0.824 (0.522)
DR-Learner	0.165 (0.049) – 0.159 (0.047)	0.281 (0.144) – 0.275 (0.137)	0.114 (0.094) – 0.464 (0.286)
X-Learner	0.163 (0.057) – 0.154 (0.051)	0.279 (0.157) – 0.279 (0.146)	0.113 (0.092) – 0.644 (0.380)
RLin-Learner	0.245 (0.136)	0.241 (0.136)	0.717 (0.450)

For the DR- and X-learners: μ_t are estimated by T-learning (left value) or S-learning (right value).

Hazard rate model in non-randomized setting

Table 14. **mPEHE** and **sdPEHE** for three different ML base-learners; Case where nuisance components are exact.

Meta-learner	XGBoost	RandomForest	Linear Model
M-Learner	6.28 (1.88)	5.74 (1.60)	3.72 (1.42)
DR-Learner	0.138 (0.029)	0.139 (0.044)	0.110 (0.097)
X-Learner	0.044 (0.025)	0.087 (0.050)	0.110 (0.097)
RLin-Learner		0.299 (0.176)	

Table 15. **mPEHE** and **sdPEHE** for three different ML base-learners; Case where nuisance components are well-specified.

Meta-learner	XGboost	RandomForest	Linear Model
T-Learner	0.183 (0.039)	0.286 (0.155)	0.129 (0.094)
RegT-Learner	0.176 (0.044)	0.286 (0.155)	0.121 (0.098)
S-Learner	0.176 (0.056)	0.306 (0.153)	0.671 (0.428)
NvX-Learner	0.190 (0.096)	0.336 (0.200)	0.129 (0.094)
M-Learner	1.61 (0.505)	1.58 (0.472)	0.906 (0.557)
DR-Learner	0.168 (0.045) - 0.178 (0.048)	0.304 (0.158) – 0.322 (0.162)	0.121 (0.098) – 0.518 (0.327)
X-Learner	0.167 (0.053) – 0.172 (0.057)	0.302 (0.169) – 0.332 (0.167)	0.120 (0.094) – 0.652 (0.388)
RLin-Learner	0.231 (0.081)	0.186 (0.123)	1.05 (0.651)

For the DR- and X-learners: μ_t are estimated by T-learning (left value) or S-learning (right value).

D.3. Asymptotic performances when K increases.

In this subsection, we consider the effect of increasing K on the hazard rate function with XGBoost. For each value K , we sample $J = 10$ different non-randomized samples following the *preferential selection* as defined previously in Appendix D.1. The **mPEHE** is then computed by averaging the **mPEHE** over the $J = 10$ samples. The results are drawn in the figure below.




Figure 1. Variation of meta-learners' performances as functions of the number of possible treatment values K for the hazard rate function in an observational design setting. (a): All meta-learners; (b): Without the M-learner

When we consider the effect of increasing K on the hazard rate function with a linear model (with $p = 2$), we notice the expected effect of K on the M-learner: The error bound is increasing with K . This means that the behaviour of the M-learner with XGBoost can be explained by the over-fitting of the base-learner.




Figure 2. Variation of meta-learner’s performances when the number of possible treatment values K for the hazard rate function in observational design setting with a linear model. (a): All meta-learners; (b): Focus on the T-learning methods.

E. Description of the semi-synthetic dataset.

Motivation

The difficulty in evaluating a causal model’s performance in real-world applications motivates the need to create a semi-synthetic dataset. In this subsection, we consider a multistage fracturing Enhanced Geothermal System (EGS).

Enhanced Geothermal Systems (EGS) are geothermal wells that generate geothermal energy by creating fluid connectivity in low-permeability conductive rocks through hydraulic, thermal, or chemical stimulation. The EGS concept involves extracting heat by constructing a subsurface fracture system to which water can be added via injection wells. Indeed, rocks are permeable due to slight fractures and pore spaces between mineral grains, and the injected water is heated by contact with the rock and returns to the surface through production wells. Moreover, Enhanced geothermal systems (EGS) have a high potential for developing and supplying renewable energy sources that are more efficient and cheaper than traditional hydrocarbon resources.

For energy companies, the goal is to optimize the design of the geothermal well (fracture spacing, Lateral Length etc.) to generate the maximum geothermal energy. However, some economic and operational problems present challenges: On the one hand, if the fractures are too small or too few, rocks will not be exploited sufficiently. On the other hand, if the number of fractures in a given rock is too high, the fractures may cool down faster. We would have a costly design that will not

maximize the extracted heat.

We assume that the heat extraction performance of the EGS satisfies the following physical model:

$$Q_{well} = Q_{fracture} \times \ell_L/d \times \eta_d, \tag{123}$$

where Q_{well} is the heat extraction performance delivered by the well (output), $Q_{fracture}$ is the *unknown* heat extraction performance from a single fracture that can be generated using a complex seven-parameter model, including reservoir characteristics and fracture design, ℓ_L is the Lateral Length of the well, d is the average spacing between two fractures and η_d is the stage efficiency penalizing the individual contribution when fractures are close to each other. We refer to Figure 3 for a graphical description of the EGS and its inputs/output.

Finally, the model in (123) respects the unconfoundedness assumption 3.1, and we can control all its variables in the simulations. We note that, in practice, all inputs are continuous with a given density. However, we discretize these variables in their input space to create a full factorial design.

Description of the data-set

This section describes the data-generating process of our semi-synthetic dataset simulating the heat delivered by a multistage fracturing EGS. The process involved the creation of a conceptual reservoir model and the modelling of multiple wells’ completion scenarios. The output (heat extraction performance) obtained from physics-based simulation experiments was tabulated with inputs in the semi-synthetic dataset.

The input data for the model were fabricated to ease confidentiality and non-disclosure information issues. However, data has been selected from reliable sources such as field observations, journals and books to be within the range of interest. Doing so allowed the building of a plain but representative reservoir model that would provide realistic results of an EGS.

The heat extraction performance from a single fracture ($Q_{fracture}$) is determined using fracture length, fracture height, fracture width, fracture permeability, reservoir porosity, reservoir permeability and pore pressure. Modelling and simulation work were done using preprocessor and reservoir simulation tools PETREL and ECLIPSE.

The four physical parameters of the fracture were investigated, and the list of values used for each parameter can be observed in Table 16. In the end, $10 \times 10 \times 2 \times 3 = 600$ fracture’s simulation cases have been realized.

Table 16. Fracture parameters and their range of variation for simulations.

Variable	Range of variation
Fracture length (ft)	[100, 1000] by a step of 100 ft
Fracture height (ft)	[50, 500] by a step of 50 ft
Fracture width (in)	{0.1, 0.2}
Fracture Permeability (md)	{30000, 85000, 19000}

To emulate distinct reservoir schemes, it was decided to vary three main parameters; porosity, permeability and pore pressure. For porosity and permeability, the simulator takes the minimum and maximum values and estimates the physical properties across the reservoir. Three different multipliers were applied to define three (Low, Base and High) scenarios. Concerning pore pressure, three specific values were defined to simulate under-normal, normal (base) and overpressure (high) gradient conditions. Therefore, $3 \times 3 \times 3 = 27$ possible scenarios were defined. Table 17 displays the range of minimum and maximum values for the three reservoir parameters to be varied.

Table 17. Reservoir parameters and their range of variation for simulations.

Variable	Range of variation
(K_{min} , K_{max}) (md)	{(0.0054, 0.0157), (0.054, 0.157), (0.109, 0.314)}
(Por_{min} , Por_{max}) (dec)	{(0.0054, 0.0157), (0.054, 0.157), (0.109, 0.314)}
Pore pressure (psi)	{5000, 7000, 9000}




Figure 3. The Causal DAG associated with the multistage EGS. Nodes in yellowish brown represent the reservoir characteristics, they can only be simulated, but in reality, we cannot intervene in these variables. Nodes in Dark green represent the fracture design. Engineers control them, and intervening in them is possible whenever there is a need to make a new fracture in the well. Nodes in blue represent a well’s design and can be chosen arbitrarily by engineers or statisticians. Nodes in black denote the outputs. $Q_{fracture}$ is only given by the simulator, whereas Q_{well} is given by the physical model in (123). Note that this graph contains nine nodes, but both K_{min} and K_{max} represent the same physical parameter K , and the same remark is valid for POR_{min} and POR_{max} .

By combining different reservoir scenarios with single fracture simulations, we obtained a single dataset with 16,200 possible cases for a fracture in a reservoir then we simulated the heat extraction performance for each experiment. Simulation’s results were tabulated in the dataset "Single_Fracture_Simulation_Cases_16200.csv".

The next step is to define well characteristics (lateral lengths and fracture spacing) to evaluate the heat extraction performance of the well when reservoir and fracture properties are not changed.

Table 18. Well parameters and their range of variation.

Variable	Range of variation
Lateral length (ft)	[2000, 14000] by a step of 1000 ft
Fracture spacing (ft)	[100, 500] by a step of 100 ft

Regarding the spacing efficiency coefficient, this coefficient was used to model interactions between fractures and penalize the heat extraction performance of a single fracture in the presence of other close fractures, that is, when the spacing between two fractures is small. Indeed, if the fractures are spaced too close, there may not be enough thermal energy in the rock to heat the water, decreasing the heat extraction efficiency. Modelling this efficiency led to the efficiency table "Fracture_Efficiency.csv" that describes what would be the well’s heat performance behaviour with respect to the fracture spacing selected. Based on this table, one can interpolate the efficiency to draw the curve (see Figure 4) and thus obtain the spacing efficiency coefficient for any desired value fracture spacing.

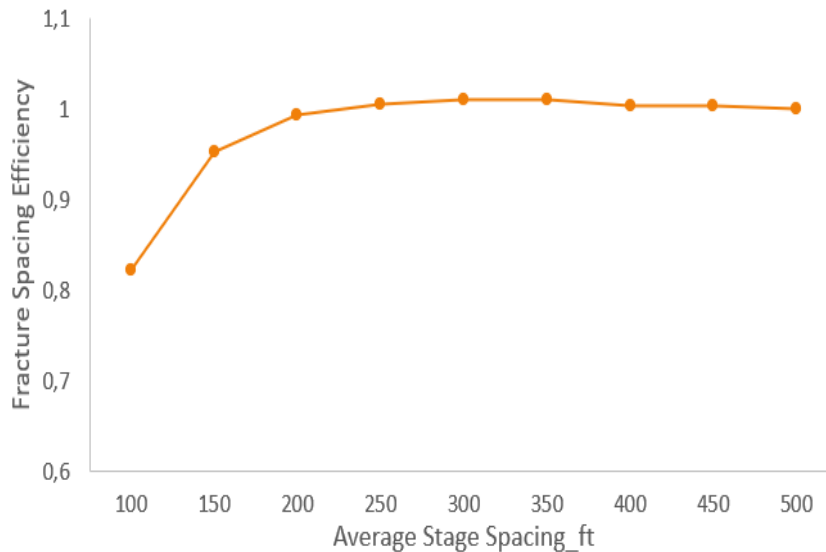


Figure 4. Cross plot between fracture spacing efficiency and average stage spacing.

The final generation of the semi-synthetic dataset "Main_Dataset.csv" was achieved by combining two main tables created using the R programming language. This table allows calculating the heat performance of a well for any lateral length and fracture spacing between 500 ft and 100 ft with the associated spacing efficiency coefficient defined in the efficiency table, following the physical model in (123).

The three datasets are available at <https://github.com/nacharki/multipleT-MetaLearners>.

Finally, we emphasize that this study's design methodology focused on generating a semi-synthetic dataset using reservoir numerical simulation and creating a new benchmarking dataset for comparing and validating causal inference methods. Indeed, following the last step of forming the final dataset "Main_Dataset.csv", any user can define different distributions (with different values) on lateral lengths in the range [2000, 14000] and fracture spacing in range [100, 500], pick-up the corresponding spacing efficiency coefficients using the curve drawn in Figure 4 and generate a new semi-synthetic dataset by extrapolating them with "Single_Fracture_Simulation_Cases_16200.csv" dataset.

The creation of a non-randomized biased dataset.

The idea of this step was to create a collection of biased data from the main semi-synthetic dataset to emulate observational data found in real-world situations. For example, geothermal wells with larger lateral lengths are likely to have more fractures (expensive wells are located in better geological areas). The opposite is seen for smaller wells that tend to be associated with fewer fractures. This situation creates a discrepancy between what engineers expect from physical models and what they observe in the field data. The biased data, with 9,992 observations, was generated by following the *preferential selection* strategy from the main dataset. Figure 5 shows the difference between the *real* heat extraction performance of the EGS and the observed heat extraction performance on the field: low (under-estimated) heat performance for small wells and high (over-estimated) heat performance for large wells.

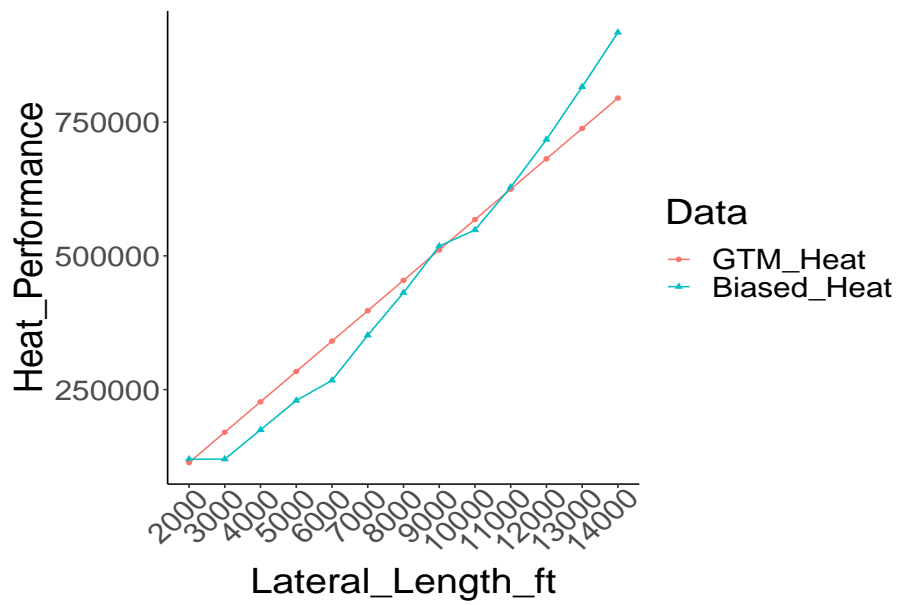


Figure 5. An illustration of selection bias on the heat performance. Red line: The heat extraction performance on the main dataset (i.e. Ground Truth Model). Blue line: The heat performance on the biased dataset (i.e. observed response).