



HAL
open science

How to select predictive models for causal inference?

Matthieu Doutreligne, Gaël Varoquaux

► **To cite this version:**

Matthieu Doutreligne, Gaël Varoquaux. How to select predictive models for causal inference?. 2023. hal-03946902v2

HAL Id: hal-03946902

<https://hal.science/hal-03946902v2>

Preprint submitted on 24 May 2023 (v2), last revised 8 Jan 2024 (v4)

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.



HOW TO SELECT PREDICTIVE MODELS FOR CAUSAL INFERENCE?

Matthieu Doutréline^{1,2,*} and Gaël Varoquaux¹

¹Inria, Soda, Saclay, France

²Mission Data, Haute Autorité de Santé, Saint-Denis, France

*Corresponding author: matthieu.doutreligne@inria.fr

May 24, 2023

ABSTRACT

As predictive models –eg from machine learning– give likely outcomes, they may be used to reason on the effect of an intervention, a causal-inference task. The increasing complexity of health data has opened the door to a plethora of models, but also the Pandora box of model selection: which of these models yield the most valid causal estimates? Here we highlight that classic machine-learning model selection does not select the best outcome models for causal inference. Indeed, causal model selection should control both outcome errors for each individual, treated or not treated, whereas only one outcome is observed. Theoretically, simple risks used in machine learning do not control causal effects when treated and non-treated population differ too much. More elaborate risks build proxies of the causal error using “nuisance” re-weighting to compute it on the observed data. But does computing these nuisance adds noise to model selection? Drawing from an extensive empirical study, we outline a good causal model-selection procedure: using the so-called R -risk; using flexible estimators to compute the nuisance models on the train set; and splitting out 10% of the data to compute risks.

Keywords Model Selection, Treatment Effect, G-formula, Observational Study, Machine Learning,

1 Introduction

1.1 Extending prediction to prescription requires causal model selection

Increasingly rich data drives new predictive models. In health, new risks or prognostic models leverage routinely-collected data, sometimes with machine learning [50]: predicting morbidity-related outcomes from administrative data [86], heart failure from claims [19], sepsis from clinical records [32], suicide attempts from patient records and questionnaires [78]... Data may be difficult to control and model, but claims of accurate prediction can be established on left-out data [3, 60, 83]. Given a model predicting of an outcome of interest, it is tempting to use it to guide decisions: will an individual benefit or not from an intervention such as surgery [22]? This is a causal-inference task and principled approaches can build on contrasting the prediction of the outcome with and without the treatment [79, 11].

Outcome modeling is an integral part of the causal modeling toolkit, under the names of G-computation, G-formula [65], Q-model [79], or conditional mean regression [85]. On observational data, causal inference of treatment effects is brittle to un-accounted for confounding, but it can assess effectiveness and safety in real-world practice [10, 29] or off-label drug usage [62] for potential repurposing [34, 21]. Epidemiology has historically focused on methods that model treatment assignment [8, 25], based on the propensity score [68] but recent empirical results [85, 20] show benefits of outcome modeling to estimate average treatment effect (ATE). A major benefit of using outcome modeling for causal inference is that these methods naturally go beyond average effects, estimating individualized or conditional average treatment effects (CATE), important for precision medicine. For this purpose, such methods are also invaluable on randomized trials [81, 45, 31].

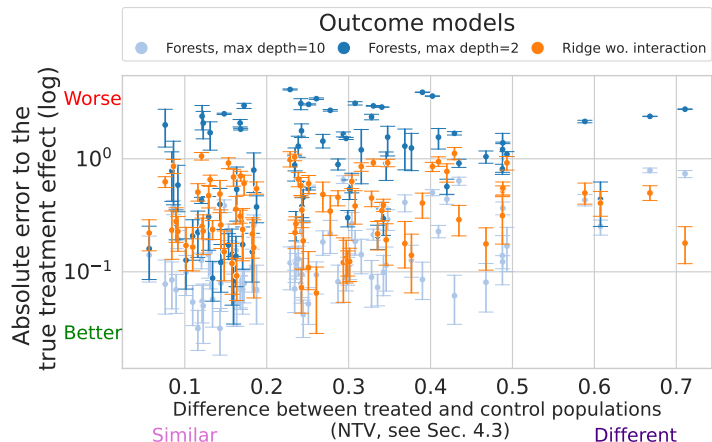
Recent developments have seen the multiplication of predictive modeling methods. Leaving aside the, overwhelming, machine-learning literature, even methods specifically designed for causal inference are numerous: Bayesian Additive Regression Trees [30], Targeted Maximum Likelihood Estimation [44, 74], causal boosting [61], causal multivariate adaptive regression splines (MARS) [61], random forests [84, 5], Meta-learners [41], R-learners [54], Doubly robust estimation [14]... The wide variety of methods leaves the applied researcher with the difficult choice of selecting between different estimators based on the data at hand. Indeed, estimates may vary markedly when using different models. For instance, Figure 1 shows large variations obtained across four different outcome estimators on 2016 semi-synthetic datasets [20]. Flexible models such as random forests are doing well in most settings except for setups where treated and untreated populations differ markedly in which case a simple linear model (ridge) is to be preferred. However a different choice of hyper-parameters (max depth= 2) for random forests yield the poorest performances. A simple rule of thumb such as preferring more flexible models does not work in general; model selection is necessary.

Standard practices to select models in predictive settings rely on cross-validation on the error on the outcome [60, 83]. However, as we will see, these model-selection practices may not pick the best models for causal inference, as they can be misled by inhomogeneities due to treatment allocation. Given complex, potentially noisy, data, which model is to be most trusted to yield valid causal estimates? Because there is no single learner that performs best on all data sets, there is a pressing need for clear guidelines to select outcome models for causal inference.

Objectives and structure of the paper In this paper, we study *model selection procedures* with a focus on practical settings: *finite samples* settings and without *well-specification* assumption. One question is whether model-selection

Figure 1: **Different outcome models lead to different estimation error on the Average Treatment Effects**, with six different outcome models on 77 classic simulations where the true causal effect is known [20]. The models are random forests, ridge regression with and without interaction with the treatment, (hyper-parameters detailed in Appendix A). The different configurations are plotted as a function of increasing difference between treated and untreated population –detailed in subsection 4.3.

There is no systematic best performer: choosing the best model among a family of candidate estimators is important.



procedures, that rely on data split, can estimate reliably enough the complex risks, theoretically motivated for causal inference. Indeed, these risks force of departure from standard model-selection procedures as they come with more quantities to estimate, which may bring additional variance, leading to worse model selection.

We first give a simple illustration of the problem of causal model selection and quickly review the important prior art. Then, in Section 2, we set causal model selection in the *potential outcome* framework and detail the causal risks and model-selection procedure. Section 3 gives our theoretical result. In section 4, we run a thorough empirical study, with many different settings covered. Finally, we comment our findings in Section 5. Results outline how to best select outcome models for causal inference with an adapted cross-validation that estimate the so-called R -risk. The R -risk modulates observed prediction error to compensate for systematic differences between treated and non-treated individuals. It relies on the two *nuisance* models, themselves estimated from data and thus imperfect; yet these imperfections do not undermine the benefit of the R -risk.

1.2 Illustration: the best predictor may not estimate best causal effects

Using a predictor to reason on causal effects relies on contrasting the prediction of the outcome for a given individual with and without the treatment –as detailed in subsection 2.1. Given various predictors of the outcome, we are interested in selecting those that estimate best the treatment effect. Standard predictive modeling or machine-learning practice selects the best predictor, *ie* the one that minimizes the expected error. However, the best predictor may not be the best model to reason about causal effects of an intervention, as we illustrate below.

Figure 2 gives a toy example: the outcome $Y \in [0, 1]$, the probability of death, a binary treatment $A \in \{0, 1\}$ and a covariate $X \in \mathbb{R}$ which summarizes the patient health status (eg. the Charlson co-morbidity index [13]). We simulate a situation for which the treatment is beneficial (decreases mortality) for patients with high Charlson scores (bad health status). On the contrary, the treatment has little effect for patients in good condition (small Charlson scores).

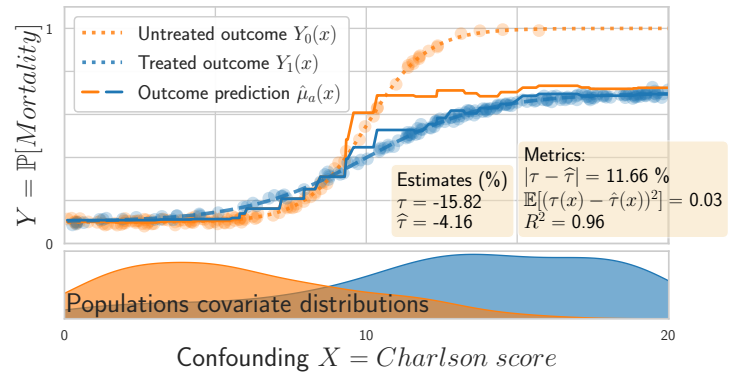
Figure 2a shows a random forest predictor with a counter-intuitive behavior: it predicts well on average the outcome (as measured by a regression R^2 score) but perform poorly to estimate causal quantities: the average treatment effect τ (as visible via the error $|\tau - \hat{\tau}|$) or the individual treatment effect (the error $\mathbb{E}[(\tau(x) - \hat{\tau}(x))^2]$). On the contrary, Figure 2b shows a linear model with smaller R^2 score but better causal inference.

Figure 2: **Illustration:** a) a random-forest estimator with high performance for standard prediction (high R^2) but that yields poor causal estimates (large error between true effect τ and estimated $\hat{\tau}$), b) a linear estimator with smaller prediction performance leading to better causal estimation.

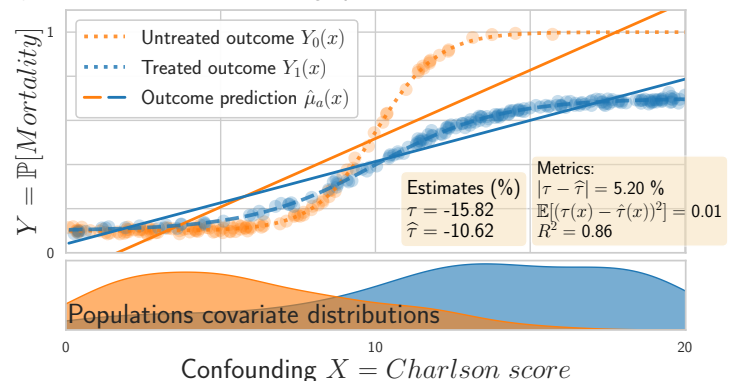
Selecting the estimator with the smallest error to the individual treatment effect $\mathbb{E}[(\tau(x) - \hat{\tau}(x))^2]$ –the τ -risk, def. 1 – would lead to the best causal estimates; however computing this error is not feasible: computing it requires access to unknown quantities: $\tau(x)$.

While the random forest fits the data better than the linear model, it gives worse causal inference because its error is very inhomogeneous between the treated and untreated. The R^2 score does not capture this inhomogeneity.

a) Random forest, good average prediction but bad causal inference



b) Linear model, worse average prediction but better causal inference



Intuitively, the problem is that causal estimation requires controlling an error on both treated and non-treated outcome for the same individual: the observed outcome, and the non-observed *counterfactual* one. The linear model is misspecified –the outcome functions are not linear–, leading to poor R^2 ; but it interpolates better to regions where there are few untreated individuals –high Charlson score– and thus gives better causal estimates. Conversely, the random forest puts weaker assumptions on the data, thus has higher R^2 score but is biased by the treated population in the region with poor overlap region, leading to bad causal estimates.

This toy example illustrates that the classic minimum Mean Square Error criterion is not suited to choosing a model among a family of candidate estimators for causal inference.

1.3 Prior work: model selection for outcome modeling (g-computation)

A natural way to select a predictive model for causal inference would be an error measure between a causal quantity such as the conditional average treatment effects (CATE) and models’ estimate. But such error is not a “feasible” risk: it cannot be computed solely from observed data and requires oracle knowledge.

Simulation studies of causal model selection Using eight simulations setups from [61], where the oracle CATE is known, Schuler et al. [73] compare four causal risks, concluding that for CATE estimation the best model-selection risk is the so-called R -risk [54] –def. 6, below. Their empirical results are clear for randomized treatment allocation but less convincing for observational settings where both simple Mean Squared Error –MSE, μ -risk(f) def. 2– and reweighted MSE – μ -risk $_{IPW}$ def. 3– appear to perform better than R -risk on half of the simulations. Another work [1] studied empirically both MSE and reweighted MSE risks on the semi-synthetic ACIC 2016 datasets [20], but did not include the R -risk and looked only at the agreement of the best selected model with the true CATE risk – τ -risk(f) def. 1–, not on the full ranking of methods compared to the true CATE. We complete these prior empirical work by studying a wider variety of data generative processes and varying the influence of overlap, an important parameter of the data generation process which makes a given causal metric appropriate [17]. We also study how to best adapt cross-validation procedures to causal metrics which themselves come with models to estimate.

Theoretical studies of causal model selection Several theoretical works have proposed causal model selection procedures that are *consistent*: select the best model in a family given asymptotically large data. These work typically rely on introducing a CATE estimator in the testing procedure. For instance matching [67], an IPW estimate [27], a doubly robust estimator [71], or debiasing the error with influence functions [1]. However, for theoretical guarantees to hold, the test-set correction needs to converge to the oracle: it needs to be flexible enough –or well-posed– and asymptotic data. From a practical perspective, knowing that such requirements are met implies having a good CATE estimate, which amounts to having solved the original problem of causal model selection. We study how causal model-selection procedures behave outside of these settings.

Statistical guarantees on causal estimation procedures Much work in causal inference has focused on building procedures that guarantee asymptotically consistent estimators. Targeted Machine Learning Estimation (TMLE) [44, 74] and Double Machine Learning [14] both provide estimators for Average Treatment Effect combining flexible treatment and outcome models. Here also, theories require asymptotic regimes and models to be *well-specified*.

By contrast, Johansson et al. [38] studies causal estimation without assuming that estimators are well specified. They derive an upper bound on the oracle error to the CATE (τ -risk) that involves the error on the outcome and the similarity of the distributions between the features of treated and control patients. However, they focus on using this upper bound for estimation, and do not give insights on model selection. In addition, for hyperparameter selection, they rely on a plugin estimate of the τ -risk built with counterfactual nearest neighbors, which has been shown ineffective [73].

2 Formal setting of causal inference and model selection

2.1 The Neyman-Rubin Potential Outcomes framework

Settings The Neyman-Rubin Potential Outcomes framework [52, 35] enables statistical reasoning on causal treatment effects: Given an outcome $Y \in \mathbb{R}$ (eg. mortality risk or hospitalization length), function of a binary treatment $A \in \mathcal{A} = \{0, 1\}$ (eg. a medical act, a drug administration), and baseline covariates $X \in \mathcal{X} \subset \mathbb{R}^d$, we observe the factual distribution, $O = (Y(A), X, A) \sim \mathcal{D} = \mathbb{P}(y, x, a)$. However, we want to model the existence of potential observations (unobserved ie. counterfactual) that correspond to a different treatment. Thus we want quantities on the counterfactual distribution $O^* = (Y(1), Y(0), X, A) \sim \mathcal{D}^* = \mathbb{P}(y(1), y(0), x, a)$.

Popular quantities of interest (estimands) are: at the population level, the Average Treatment Effect

$$\text{(ATE)} \quad \tau \stackrel{\text{def}}{=} \mathbb{E}_{Y(1), Y(0) \sim \mathcal{D}^*} [Y(1) - Y(0)];$$

to model heterogeneity, the Conditional Average Treatment Effect

$$\text{(CATE)} \quad \tau(x) \stackrel{\text{def}}{=} \mathbb{E}_{Y(1), Y(0) \sim \mathcal{D}^*} [Y(1) - Y(0) | X = x].$$

Causal assumptions Some assumptions are necessary to assure identifiability of the causal estimands in observational settings [70]. We assume the usual strong ignorability assumptions, composed of 1) *unconfoundedness* $\{Y(0), Y(1)\} \perp\!\!\!\perp A | X$, 2) *strong overlap* ie. every patient has a strictly positive probability to receive each treatment, 3) *consistency*, and 4) *generalization* (detailed in Appendix B). In this work, we investigate the role of the overlap [17], which is testable with data.

Estimating treatment effects with outcome models Should we know the two expected outcomes for a given X , we could compute the difference between them, which gives the causal effect of the treatment. These two expected outcomes can be computed from the observed data: the consistency 3 and unconfoundedness 1 assumptions imply the equality of two different expectations:

$$\mathbb{E}_{Y(a) \sim \mathcal{D}^*} [Y(a) | X = x] = \mathbb{E}_{Y \sim \mathcal{D}} [Y | X = x, A = a] \quad (1)$$

On the left, the expectation is taken on the counterfactual unobserved distribution. On the right, the expectation is taken on the factual observed distribution conditionally on the treatment. This equality is referred as the g-formula identification [64]. For the rest of the paper, the expectations will always be taken on the factual observed distribution \mathcal{D} , and we will omit to explicitly specify the distribution. This identification leads to outcome based estimators (ie. g-computation estimators[79]), targeting the ATE τ with outcome modeling:

$$\tau = \mathbb{E}_{Y \sim \mathcal{D}^*} [Y(1) - Y(0) | X = x] = \mathbb{E}_{Y \sim \mathcal{D}} [Y | A = 1] - \mathbb{E}_{Y \sim \mathcal{D}} [Y | A = 0] \quad (2)$$

This equation has two central quantities: the conditional expectancy function of the outcome associated to specific covariates and treatment or not is, often called the *response function*:

$$\text{(Response function)} \quad \mu_a(x) \stackrel{\text{def}}{=} \mathbb{E}_{Y \sim \mathcal{D}} [Y | X = x, A = a]. \quad (3)$$

Given a sample of data and the oracle response functions μ_0, μ_1 , the finite sum version of Equation 2 leads to an estimator of the ATE written:

$$\hat{\tau} = \frac{1}{n} \left(\sum_{i=1}^n \mu_1(x_i) - \mu_0(x_i) \right) \quad (4)$$

This estimator is an oracle **finite sum estimator** by opposition to the population expression of τ , $\mathbb{E}[\mu_1(x_i) - \mu_0(x_i)]$, which involves an expectation taken on the full distribution \mathcal{D} , which is observable but requires infinite data. For each estimator ℓ taking an expectation over \mathcal{D} , we use the symbol $\hat{\ell}$ to note its finite sum version.

Similarly to the ATE, for the CATE, at the individual level:

$$\tau(x) = \mu_1(x) - \mu_0(x) \quad (5)$$

Robinson decomposition Another decomposition of the outcome model plays an important role, the R-decomposition [66]: introducing two quantities, the conditional mean outcome and the probability to be treated (known as propensity score [68]):

$$\text{(Conditional mean outcome)} \quad m(x) \stackrel{\text{def}}{=} \mathbb{E}_{Y \sim \mathcal{D}} [Y | X = x]. \quad (6)$$

$$\text{(Propensity score)} \quad e(x) \stackrel{\text{def}}{=} \mathbb{P}[A = 1 | X = x], \quad (7)$$

the outcome can be written

$$\text{(R-decomposition)} \quad y(a) = m(x) + (a - e(x))\tau(x) + \varepsilon(x; a) \quad \text{with} \quad \mathbb{E}[\varepsilon(X; A) | X, A] = 0 \quad (8)$$

m and e are often called *nuisances* [14]; they are unknown in general.

2.2 Model-selection risks, oracle and feasible

Causal model selection We formalize model selection for causal estimation. Thanks to the g-formula identification (Equation 1), a given outcome model $f : \mathcal{X} \times \mathcal{A} \rightarrow \mathcal{Y}$ –learned from data or built from domain knowledge– induces feasible estimates of the ATE and CATE (eqs 4 and 5), $\hat{\tau}_f$ and $\hat{\tau}_f(x)$. Let $\mathcal{F} = \{f : \mathcal{X} \times \mathcal{A} \rightarrow \mathcal{Y}\}$ be a family of such estimators. Our goal is to select the best candidate in this family for the observed dataset O using a risk of interest ℓ :

$$f_\ell^* = \operatorname{argmin}_{f \in \mathcal{F}} \ell(f, O) \quad (9)$$

We now detail possible risks ℓ , risks useful for causal model selection, and how to compute them.

The τ -risk: an oracle error risk As we would like to target the CATE, the following evaluation risk is natural:

Definition 1 (τ -risk(f)) *also called PEHE [72, 30]:*

$$\tau\text{-risk}(f) = \mathbb{E}_{X \sim p(X)} [(\tau(X) - \hat{\tau}_f(X))^2]$$

its finite-sum version over the observed data:

$$\widehat{\tau\text{-risk}}(f) = \sum_{x \in O} (\tau(x) - \hat{\tau}_f(x))^2$$

However these risks are not feasible because the oracles $\tau(x)$ are not accessible, with the observed data $(Y, X, A) \sim \mathcal{D}$.

Feasible error risks Feasible risks are based on the prediction error of the outcome model and *observable* quantities.

Definition 2 (Factual μ -risk) [75] *This is the usual Mean Squared Error on the target y . It is what is typically meant by “generalization error” in supervised learning and estimated with cross-validation:*

$$\mu\text{-risk}(f) = \mathbb{E}_{(Y, X, A) \sim \mathcal{D}} [(Y - f(X; A))^2]$$

The following risks use the nuisances e –propensity score, def 7– and m –conditional mean outcome, def 6. We give the definitions as *semi-oracles*, function of the true unknown nuisances, but later instantiate them with estimated nuisances, noted (\check{e}, \check{m}) . Semi-oracles risks are superscripted with the $*$ symbol.

Definition 3 (μ -risk $_{IPW}^*$) [42] *Let the inverse propensity weighting function $w(x, a) = \frac{a}{e(x)} + \frac{1-a}{1-e(x)}$, we define the semi-oracle Inverse Propensity Weighting risk,*

$$\mu\text{-risk}_{IPW}^*(f) = \mathbb{E}_{(Y, X, A) \sim \mathcal{D}} \left[\left(\frac{A}{e(X)} + \frac{1-A}{1-e(X)} \right) (Y - f(X; A))^2 \right]$$

Definition 4 (τ -risk $_{IPW}^*$) [84] *The CATE $\tau(x)$ can be estimated propensity score, with a regression against inverse propensity weighted outcomes [4, 27, 84]. From this objective, we can derive the τ -risk $_{IPW}$.*

$$\tau\text{-risk}_{IPW}^*(f) = \mathbb{E}_{(Y, X, A) \sim \mathcal{D}} \left[\left(Y \frac{A - e(X)}{e(X)(1 - e(X))} - \tau_f(X) \right)^2 \right] = \mathbb{E}_{(Y, X, A) \sim \mathcal{D}} \left[\left(Y \left(\frac{A}{e(X)} - \frac{1 - A}{1 - e(X)} \right) - \tau_f(X) \right)^2 \right]$$

Definition 5 (U -risk *) [41, 54] *Based on the Robinson decomposition –eq. 8 [66], the U-learner uses the $A - e(X)$ term in the denominator. The derived risk is:*

$$U\text{-risk}^*(f) = \mathbb{E}_{(Y, X, A) \sim \mathcal{D}} \left[\left(\frac{Y - m(X)}{A - e(X)} - \tau_f(X) \right)^2 \right]$$

Note that extreme propensity weights in the denominator term might inflate errors in the numerator due to imperfect estimation of the mean outcome m the numerator errors, leading to a highly biased metric.

Definition 6 (R -risk *) [54, 73] *The R-risk also uses two nuisance m and e :*

$$R\text{-risk}^*(f) = \mathbb{E}_{(Y, X, A) \sim \mathcal{D}} [(Y - m(X)) - (A - e(X)) \tau_f(X)]^2$$

It is also motivated by the Robinson decomposition –eq. 8 [66]. It performs well in various simulations where using lasso, boosting or kernel ridge regressions for both the nuisances (\check{e}, \check{m}) and the target $\hat{\tau}(x)$ [54].

These risks are summarized in Table 1.

Table 1: Review of causal risks

Risk	Equation	Reference
$mse(\tau(X), \tau_f(X)) = \tau\text{-risk}$	$\mathbb{E}_{X \sim p(X)} [(\tau(X) - \hat{\tau}_f(X))^2]$	Eq. 1 [30]
$mse(Y, f(X)) = \mu\text{-risk}$	$\mathbb{E}_{(Y, X, A) \sim \mathcal{D}} [(Y - f(X; A))^2]$	Def. 2 [73]
$\mu\text{-risk}_{IPW}^*$	$\mathbb{E}_{(Y, X, A) \sim \mathcal{D}} \left[\left(\frac{A}{e(X)} + \frac{1-A}{1-e(X)} \right) (Y - f(X; A))^2 \right]$	Def. 3 [42]
$\tau\text{-risk}_{IPW}^*$	$\mathbb{E}_{(Y, X, A) \sim \mathcal{D}} \left[\left(Y \left(\frac{A}{e(X)} - \frac{1-A}{1-e(X)} \right) - \hat{\tau}_f(X) \right)^2 \right]$	Def. 4 [84]
$U\text{-risk}^*$	$\mathbb{E}_{(Y, X, A) \sim \mathcal{D}} \left[\left(\frac{Y - m(X)}{A - e(X)} - \hat{\tau}_f(X) \right)^2 \right]$	Def. 5 [54]
$R\text{-risk}^*{}^1$	$\mathbb{E}_{(Y, X, A) \sim \mathcal{D}} \left[\left((Y - m(X)) - (A - e(X)) \hat{\tau}_f(X) \right)^2 \right]$	Def. 6 [54]

¹ Called $\tau\text{-risk}_R$ in Schuler et al. [73].

2.3 Estimation and model selection procedure

Causal model selection (as in Equation 9) may involve estimating various quantities from the observed data: the outcome model f , its induced risk as introduced in the previous section, and possibly nuisances required by the risk. Given a dataset with N samples, we split out a train and a test sets (\mathcal{T}, \mathcal{S}). We fit each candidate estimator $f \in \mathcal{F}$ on \mathcal{T} . We also fit the nuisance models (\check{e}, \check{m}) on the train set \mathcal{T} , setting hyperparameters by a nested cross-validation before fitting the nuisance estimators with these parameters on the full train set. Causal quantities are then computed by applying the fitted candidate estimators $f \in \mathcal{F}$ on the test set \mathcal{S} . Finally, we compute the model-selection metrics for each candidate model on the test set. This procedure is described in Algorithm 1 and Figure 3.

As extreme inverse propensity weights induce high variance, clipping can be useful for numerical stability [82, 36].

Algorithm 1 Evaluation of selection procedures for one simulation

Given a train and a test sets $(\mathcal{T}, \mathcal{S}) \sim \mathcal{D}$, a family of candidate estimators $\{f \in \mathcal{F}\}$, a set of causal metrics $\ell \in \mathcal{L}$:

1. Prefit: Learn estimators for unknown nuisance quantities (\check{e}, \check{m}) on the training set \mathcal{T}
 2. Fit: $\forall f \in \mathcal{F}$ learn $\hat{f}(\cdot, a)$ on \mathcal{T}
 3. Model selection: $\forall x \in \mathcal{S}$ predict $(\hat{f}(x, 1), \hat{f}(x, 0))$ and evaluate each candidate estimator with each causal metric $\mathcal{M}(f, \mathcal{S})$. For each causal metric $\ell \in \mathcal{L}$ and each candidate estimator $f \in \mathcal{F}$, store the metric value: $\ell(f, \mathcal{S})$ – possibly function of \check{e} and \check{m}
-

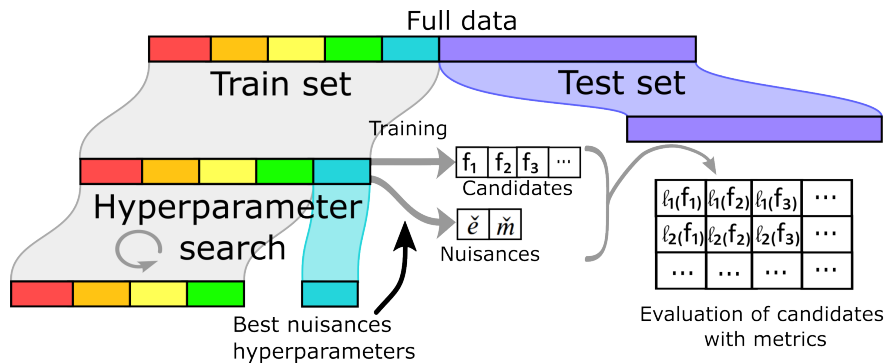


Figure 3: Estimation procedure for causal model selection.

3 Theory: Links between feasible and oracle risks

We now relate two feasible risks, $\mu\text{-risk}_{IPW}$ and the R -risk to the oracle τ -risk. Both results make explicit the role of overlap for the performances of causal risks.

These bounds depend on a specific form of residual that we now define: for each potential outcome, $a \in \{0; 1\}$, the variance conditionally on x is [75]:

$$\sigma_y^2(x; a) \stackrel{\text{def}}{=} \int_y (y - \mu_a(x))^2 p(y | x = x; A = a) dy$$

Integrating over the population, we get the Bayes squared error: $\sigma_B^2(a) = \int_{\mathcal{X}} \sigma_y^2(x; a) p(x) dx$ and its propensity weighted version: $\tilde{\sigma}_B^2(a) = \int_{\mathcal{X}} \sigma_y^2(x; a) p(x; a) dx$. In case of a purely deterministic link between the covariates, the treatment, and the outcome, these residual terms are null.

3.1 Upper bound of τ -risk with μ -risk_{IPW}

Proposition 1 (Upper bound with μ -risk_{IPW}) [38] *Given an outcome model f , let a weighting function $w(x; a) = \frac{a}{e(x)} + \frac{1-a}{1-e(x)}$ as the Inverse Propensity Weight. Then, under overlap (assumption 2), we have:*

$$\tau\text{-risk}(f) \leq 2 \mu\text{-risk}_{IPW}(w, f) - 2 (\sigma_B^2(1) + \sigma_B^2(0))$$

This result has been derived in previous work [38]. It links μ -risk_{IPW} to the squared residuals of each population thanks to a reweighted mean-variance decomposition. For completeness, we provide the proof in .

The upper-bound comes from the triangular inequality applied to the residuals of both populations. Interestingly, the two quantities are equal when the absolute residuals on treated and untreated populations are equal on the whole covariate space, ie for all $x \in \mathcal{X}$, $|\mu_1(x) - f(x, 1)| = |\mu_0(x) - f(x, 0)|$. The main source of difference between the oracle τ -risk and the reweighted mean squared error, μ -risk_{IPW}, comes from heterogeneous residuals between populations. These quantities are difficult to characterize as they are linked both to the estimator and to the data distribution. This bound indicates that minimizing the μ -risk_{IPW} helps to minimize the τ -risk, which leads to interesting optimization procedures [38]. However, there is no guarantee that this bound is tight, which makes it less useful for model selection.

Assuming strict overlap (probability of all individuals being treated or not bounded away from 0 and 1 by η , appendix B), the above bound simplifies into a looser one involving the usual mean squared error: $\tau\text{-risk}(f) \leq \frac{2}{\eta} \mu\text{-risk}(f) - 2 (\sigma_B^2(1) + \sigma_B^2(0))$. For weak overlap (propensity scores not bounded far from 0 or 1), this bound is very loose (as shown in Figure 2) and is not appropriate to discriminate between models with close performances.

3.2 Reformulation of the R -risk as reweighted τ -risk

We now derive a novel rewriting of the R -risk, making explicit its link with the oracle τ -risk.

Proposition 2 (R -risk as reweighted τ -risk) *Given an outcome model f , its R -risk appears as weighted version of its τ -risk (Proof in Appendix C.2):*

$$R\text{-risk}^*(f) = \int_x e(x)(1 - e(x)) (\tau(x) - \tau_f(x))^2 p(x) dx + \tilde{\sigma}_B^2(1) + \tilde{\sigma}_B^2(0) \quad (10)$$

The R -risk targets the oracle at the cost of an overlap re-weighting and the addition of the reweighted Bayes residuals, which are independent of f . In good overlap regions the weights $e(x)(1 - e(x))$ are close to $\frac{1}{4}$, hence the R -risk is close to the desired gold-standard τ -risk. On the contrary, for units with extreme overlap violation, these weights goes down to zero with the propensity score.

3.3 Interesting special cases

Randomization special case If the treatment is randomized as in RCTs, $p(A = 1 | X = x) = p(A = 1) = p_A$, thus μ -risk_{IPW} takes a simpler form:

$$\mu\text{-risk}_{IPW} = \mathbb{E}_{(Y, X, A) \sim \mathcal{D}} \left[\left(\frac{A}{p_A} + \frac{1-A}{1-p_A} \right) (Y - f(X; A))^2 \right]$$

However, even if we have randomization, we still can have large differences between τ -risk and μ -risk_{IPW} coming from heterogeneous errors between populations as noted in Section 3.1 and shown experimentally in simulations [73].

Concerning the R -risk, replacing $e(x)$ by its randomized value p_A in Proposition 2 yields the oracle τ -risk up to multiplicative and additive constants:

$$R\text{-risk} = p_A (1 - p_A) \tau\text{-risk} + (1 - p_A) \sigma_B^2(0) + p_A \sigma_B^2(1) \quad (11)$$

Therefore, optimizing estimators for CATE with R -risk* in the randomized setting is optimal if we target the τ -risk. This explains the strong performances of R -risk in randomized setups [73] and is a strong argument in favor of this risk for heterogeneity estimation in RCTs.

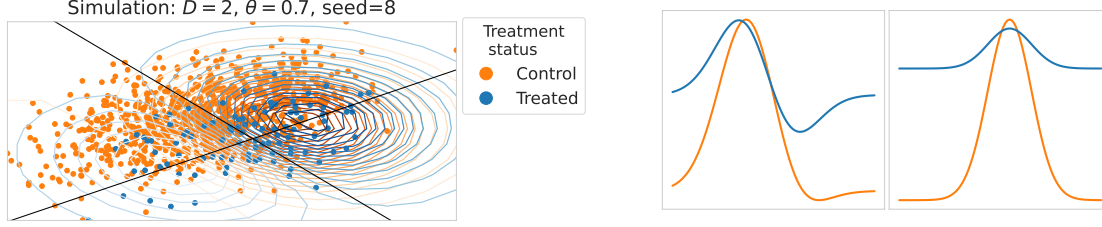


Figure 4: Example of the simulation setup in the input space with two knots *–ie.* basis functions. The left panel gives views of the observations in feature space, while the right panel displays the two response surfaces on a 1D cut along the black lines drawn on the left panel.

Oracle Bayes predictor Consider the case where we have access to the oracle Bayes predictor for the outcome *ie.* $f(x, a) = \mu(x, a)$, then all risks are equivalent up to the residual variance:

$$\tau\text{-risk}(\mu) = \mathbb{E}_{X \sim p(X)}[(\tau(X) - \tau_\mu(X))^2] = 0 \quad (12)$$

$$\mu\text{-risk}(\mu) = \mathbb{E}_{(Y, X, A) \sim p(Y; X; A)}[(Y - \mu_A(X))^2] = \int_{\mathcal{X}, \mathcal{A}} \varepsilon(x, a)^2 p(a | x) p(x) dx da \leq \sigma_B^2(0) + \sigma_B^2(1) \quad (13)$$

$$\mu\text{-risk}_{IPW}(\mu) = \sigma_B^2(0) + \sigma_B^2(1) \quad \text{follows from Lemma 1} \quad (14)$$

$$R\text{-risk}(\mu) = \tilde{\sigma}_B^2(0) + \tilde{\sigma}_B^2(1) \leq \sigma_B^2(0) + \sigma_B^2(1) \quad \text{follows directly from Proposition 2} \quad (15)$$

Thus, differences between causal risks only matter in finite sample regimes. Universally consistent learners converge to the Bayes risk in asymptotic regimes, making all model selection risks equivalent. However, in practice choices must be made in non-asymptotic regimes.

4 Empirical Study

We evaluate the following causal metrics, oracle and feasible versions, presented in Table 1: $\widehat{\mu\text{-risk}}_{IPW}^*$, $\widehat{R\text{-risk}}^*$, $\widehat{U\text{-risk}}^*$, $\widehat{\tau\text{-risk}}_{IPW}^*$, $\widehat{\mu\text{-risk}}$, $\widehat{\mu\text{-risk}}_{IPW}$, $\widehat{R\text{-risk}}$, $\widehat{U\text{-risk}}$, $\widehat{\tau\text{-risk}}_{IPW}$. We benchmark the metrics in a variety of settings: many different simulated data generation processes and three semi-simulated datasets¹.

4.1 Caussim: Extensive simulation settings

Data Generation Process We use simulated data, on which the ground-truth causal effect is known. Going further than prior empirical studies of causal model selection [73, 1], we use multiple generative processes, to reach more general conclusions (as discussed in Appendix 20).

We generate the response functions using random bases. Basis extension methods are common in biostatistics, *eg* functional regression with splines [33, 58]. By allowing the function to vary at specific knots, they give flexible *–non-linear–* models of the studied mechanisms. Taking inspiration from splines, we use random approximation of Radial Basis Function (RBF) kernels [63] to generate the response surfaces. RBF use the same process as polynomial splines but replace polynomial by Gaussian kernels. Unlike polynomials, Gaussian kernels have decreasing influences in the input space. This avoids unrealistic divergences of the population response surfaces at the ends of the feature space.

The number of basis functions *–ie. knots–*, controls the complexity of the ground-truth response surfaces and treatment. We first use this process to draw the non-treated response surface μ_0 and the causal-effect τ . We then draw the observations from a mixture two Gaussians, for the treated and non treated. We vary the separation between the two Gaussians to control the amount of overlap between treated and control populations, as it an important parameter for causal inference (related to η which appears in section 3.1). Finally, we generate the observed outcomes adding Gaussian noise. We generated such datasets 1000 times, with uniformly random overlap parameters $\theta \in [0, 2.5]$. Appendix E.1 gives more details on the data generation.

¹Scripts for the simulations and the selection procedure are available at <https://github.com/strayMat/caussim>.

Family of candidate estimators We test model selection on a family of candidate estimators that approximate imperfectly the data-generating process. To build such an estimator in two steps, we first use a RBF expansion similar as the one used for the data-generation generation process. Concretely, we choose two random knots and apply a transformation of the raw data features with the same Gaussian kernel used for the data-generation mechanism. This step is referred as the featurization. Then, we fit a linear regression on these transformed features. We consider two ways of combining these steps for outcome mode; using common nomenclature [41], we refer to these regression structures as different meta-learners which differ on how they model, jointly or not, the treated and the non treated:

- S-Learner: A single learner for both population, taking the treatment as a supplementary covariate.
- SftLearner: A single set of basis functions is sampled at random for both populations, leading to a given feature space used to model both the treat and the non treated, then two separate different regressors are fitted on this representation.
- T-Learner: Two completely different learners for each population, hence separate featurization and separate regressors.

We do not include more elaborated meta-learners such as R-learner [54] or X-learner [41]. Our goal is not to have the best possible learner but to have a variety of sub-optimal learners in order to compare the different causal metrics. For the same reason, we did not include more powerful outcome models such as random forests or boosting trees.

For the regression step, we fit a Ridge regression on the transformed features with 6 different choices of the regularization parameter $\lambda \in [10^{-3}, 10^{-2}, 10^{-1}, 1, 10^1, 10^2]$, coupled with a T-Learner or a SftLearner. We sample 10 different random basis for the learning procedure and the featurization yielding a family \mathcal{F} of 120 candidate estimators.

4.2 Semi-simulated datasets

Datasets We also use semi-simulated datasets, where a known synthetic causal effect is added to real –non synthetic– covariate. We study datasets used in previous work to evaluate causal inference:

ACIC 2016 [20]: The dataset is based on the Collaborative Perinatal Project [55], a RCT conducted on a cohort of pregnant women to identify causes of infants’ developmental disorders. The initial intervention was a child’s birth weight ($A = 1$ if weight $< 2.5kg$), and outcome was the child’s IQ after a given follow-up period. The study contained $N = 4802$ data points with $D = 55$ features (5 binary, 27 count data, and 23 continuous). They simulated 77 different setups with varying parameters for treatment and response generation models, treatment assignment probabilities, overlap, and interactions between treatment and covariates². We used 10 different seeds for every setup, totalizing 770 dataset instances.

ACIC 2018 [77]: The raw covariates data comes from the Linked Births and Infant Deaths Database (LBIDD) [48] with $D = 177$ covariates. Treatment and outcome models have been simulated with complex models to reflect different scenarii. The data do not provide the true propensity scores, so we evaluate only feasible metrics, which do not require this nuisance parameter. We used all 432 datasets³ of size $N = 5000$.

Twins [47]: It is an augmentation of the real data on twin births and mortality rates in the USA from 1989-1991 [2]. There are $N = 11984$ samples (pairs of twins), and $D = 50$ covariates⁴. The outcome is the mortality and the treatment is the weight of the heavier twin at birth. This is a "true" counterfactual dataset –as remarked in [16]– in the sense that we have both potential outcomes with each twin. They simulate the treatment with a sigmoid model based on GESTAT10 (number of gestation weeks before birth) and x the 45 other covariates:

$$\mathbf{t}_i | \mathbf{x}_i, \mathbf{z}_i \sim \text{Bern}(\sigma(w_o^\top \mathbf{x} + w_h(\mathbf{z}/10 - 0.1))) \quad \text{with } w_o \sim \mathcal{N}(0, 0.1 \cdot I), \quad w_h \sim \mathcal{N}(5, 0.1) \quad (16)$$

We built upon this equation, adding a non-constant slope in the treatment sigmoid, allowing us to control the amount of overlap between treated and control populations. We sampled uniformly 1000 different overlap parameters between 0 and 2.5, totalizing 1000 dataset instances. Unlike the previous datasets, only the overlap varies for these instances. The response surfaces are fixed by the original twin outcomes.

Family of candidate estimators For these three datasets, the family of candidate estimators are gradient boosting trees for both the response surfaces and the treatment⁵ with S-learner, learning rate in $\{0.01, 0.1, 1\}$, and maximum number of leaf nodes in $\{25, 27, 30, 32, 35, 40\}$ resulting in a family of size 18.

²Original R code available at <https://github.com/vdorie/aciccomp/tree/master/2016> to generate 77 simulations settings.

³Using the scaling part of the data, from github.com/IBM-HRL-MLHLS/IBM-Causal-Inference-Benchmarking-Framework

⁴We obtained the dataset from <https://github.com/AMLab-Amsterdam/CEVAE/tree/master/datasets/TWINS>

⁵Scikit-learn regressor, `HistGradientBoostingRegressor`, and classifier, `HistGradientBoostingClassifier`.

Nuisance estimators Drawing inspiration from the TMLE literature that uses combination of flexible machine learning methods [74], we use as models for the nuisances \check{e} (respectively \check{m}) a form of meta-learner: a stacked estimator of ridge and boosting classifiers (respectively regressions). We select hyper-parameters with randomized search on a validation set \mathcal{V} and keep them fix for model selection (detailed of the hyper parameters in Appendix E.2). As extreme inverse propensity weights induce high variance, we use clipping [82, 36] to bound $\min(\check{e}, 1 - \check{e})$ away from 0 with a fixed $\eta = 10^{-10}$, ensuring strict overlap for numerical stability.

4.3 Measuring overlap between treated and non treated

Good overlap, or “positivity” between treated and control population is crucial for causal inference as it is required by the positivity assumption 2 for causal identification. It is typically assessed by qualitative methods using population histograms (as in Figure 2) or side-by-side box plots, or quantitative approaches such as Standardized Mean Difference [7, 8]. While these methods are useful to decide if positivity holds, they do not summarize a dataset’s overlap in a single measure. Rather, we compute the divergence between the population covariate distributions $\mathbb{P}(X|A = 0)$ and $\mathbb{P}(X|A = 1)$ to characterize the behavior of causal risk [17, 38]. We introduce the Normalized Total Variation (NTV), a divergence based on the sole propensity score. Details are given in Appendix D.

4.4 Empirical results: factors driving good model selection across datasets

The R -risk is the best metric Each metric ranks differently the candidate models. Figure 5 shows the agreement between the ideal ranking of methods given the oracle τ -risk and the different causal metrics under evaluation. We measure this agreement with a relative⁶ Kendall tau κ (eq. 20) [39]. Given the importance of overlap in how well metrics approximate the oracle τ -risk (subsection C.1), we separate strong and weak overlap –defined as first and last tertile of the Normalized Total Variation, eq. 18.

⁶To remove the variance across datasets (some datasets lead to easier model selection than others), we report values for one metric relative to the mean of all metrics for a given dataset instance: $\text{Relative } \kappa(\ell, \tau\text{-risk}) = \kappa(\ell, \tau\text{-risk}) - \text{mean}_\ell(\kappa(\ell, \tau\text{-risk}))$

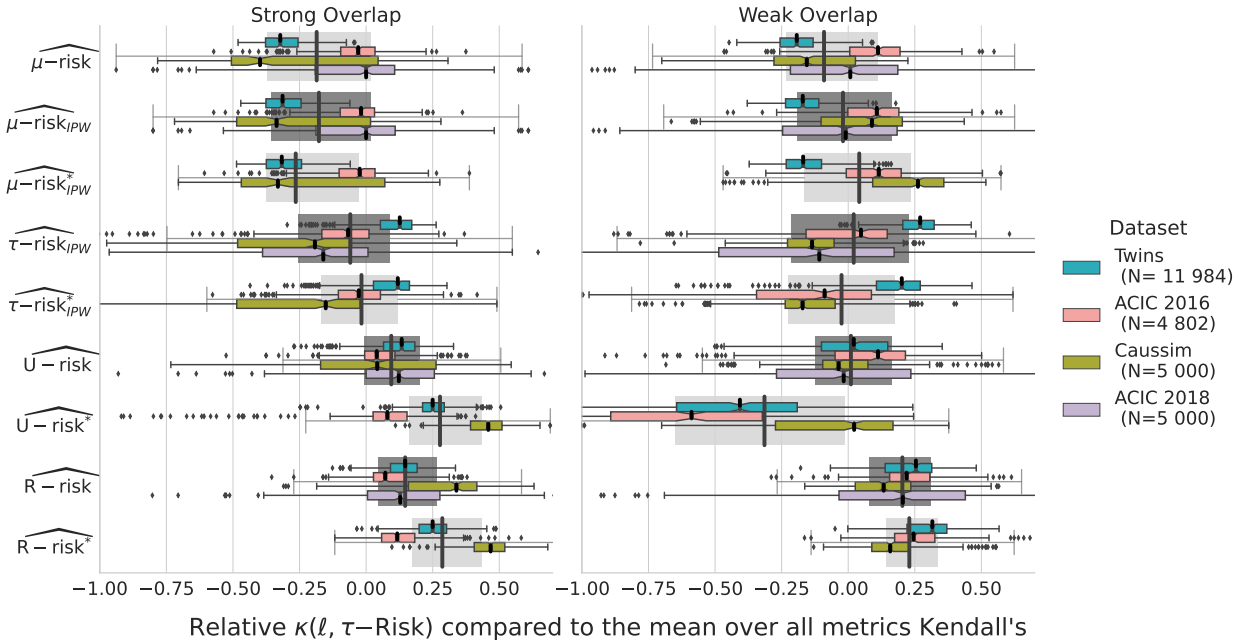


Figure 5: **The R -risk is the best metric:** Relative Kendall’s τ agreement with τ -risk, measured as the difference between each metric Kendall’s τ and the mean kendall’s τ over all metric: $\kappa(\ell, \tau\text{-risk}) - \text{mean}_\ell(\kappa(\ell, \tau\text{-rRisk}))$. Strong and Weak overlap correspond to the first and last tertiles of the overlap distribution measured with Normalized Total Variation eq. 18. Appendix E.3 presents the same results measured with absolute Kendall’s in Figure 14 and with τ -risk gains in Figure 15. Table 4 displays the median and IQR for the relative Kendall’s results.

Among all metrics the classical mean squared error (ie. factual μ -risk) is worse and reweighting it with propensity score (μ -risk_{IPW}) does not bring much improvements. The R -risk, which includes a model of mean outcome and propensity scores, leads to the best performances. Interestingly, the U -risk, which uses the same nuisances, has good performances for strong overlap but deteriorates in weak overlap, probably due variance inflation when dividing by extreme propensity scores.

Beyond the rankings, the differences in terms of absolute ability to select the best model are large: The model selected by the R -risk achieves a τ -risk only 1% higher than that of the best possible candidate for strong overlap on Caussim, but selecting with the μ -risk or μ -risk_{IPW} –as per machine-learning practice– leads to 10% excess risk and using τ -risk_{IPW} –as in some causal-inference methods [4, 27]– leads to 100% excess risk (Figure 15). Across other datasets, the R -risk consistently decreases the risk compared to the μ -risk: 0.1% compared to 1% on ACIC2016, 1% compared to 20% on ACIC 2018, and 0.05% compared to 1% on Twins.

Model selection is harder in settings of low population overlap Model selection for causal inference becomes more and more difficult with increasingly different treated and control populations (Figure 6). The absolute Kendall’s coefficient correlation with τ -risk drops from values around 0.9 (excellent agreement with oracle selection) to 0.6 on both Caussim and ACIC 2018 (Appendix E.3, Figure 14).

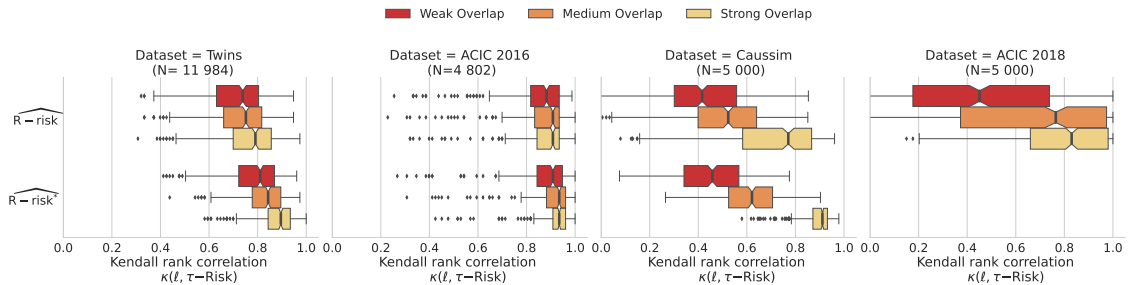


Figure 6: **Model selection is harder in settings of low population overlap:** Kendall’s τ agreement with τ -risk. Strong, medium and Weak overlap are the tertiles of the overlap distribution measured with NTV eq. 18. Appendix E.3 presents results for all metrics in Figure 17 and in absolute Kendall’s and continuous overlap values in Figure 14.

Nuisances can be estimated on the same data as outcome models Using the train set \mathcal{T} both to fit the candidate estimator and the nuisance estimates is a form of double dipping which can lead errors in nuisances to be correlated to that of outcome models [54]. In theory, these correlations can bias model selection and, strictly speaking, there are theoretical arguments to split out a third separated data set –a “nuisance set”– to fit the nuisance models. The drawback is that it depletes the data available for model estimation and selection. However, Figure 7 shows that there is no substantial differences between a procedure with a separated nuisance set and the simpler shared nuisance-candidate set procedure (results for all metrics in Figure 16).

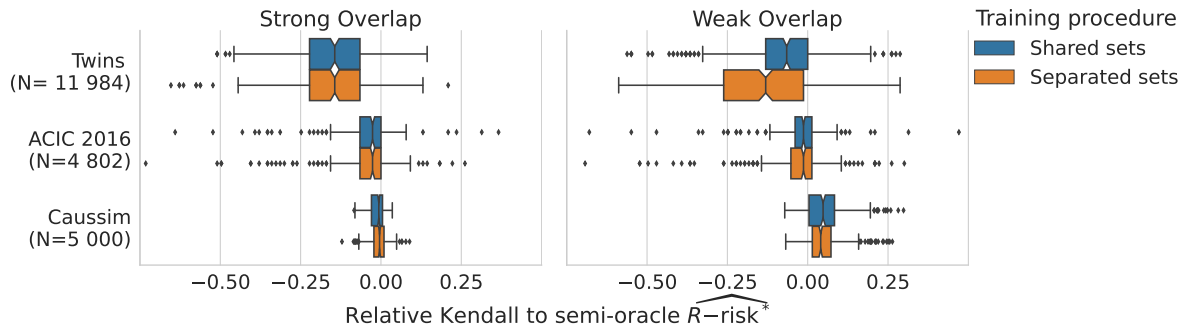


Figure 7: **Nuisances can be estimated on the same data as outcome models:** Results for the R -risk are similar between the [shared nuisances/candidate set](#) and the [separated nuisances set](#) procedures. Figure 16 details results for all metrics.

Stacked models are good overall estimators of nuisances Oracle versions of every risks recover more often the best estimator. However, in configurations that we investigated, stacked nuisances estimators (boosting and linear)

lead to feasible metrics with close performances to the oracles ones. This suggests that the corresponding estimators recover well-enough the true nuisances. One may wonder if it may be useful to use simpler models for the nuisances, in particular in settings where there are less data or where the true models are linear. Figure 8 compares causal model selection estimating nuisances with stacked estimators or linear model. It comprises the Twins data, where the true propensity model is linear, and a downsampled version of this data, to study a situation favorable to linear models. In these settings, stacked and linear estimations of the nuisances performs equivalently. Overall, Figure 19 suggests that to estimate nuisance it suffices to use adaptive models as built by stacking linear models and gradient-boosted trees. Figure 18 in the appendices, details other causal metrics.

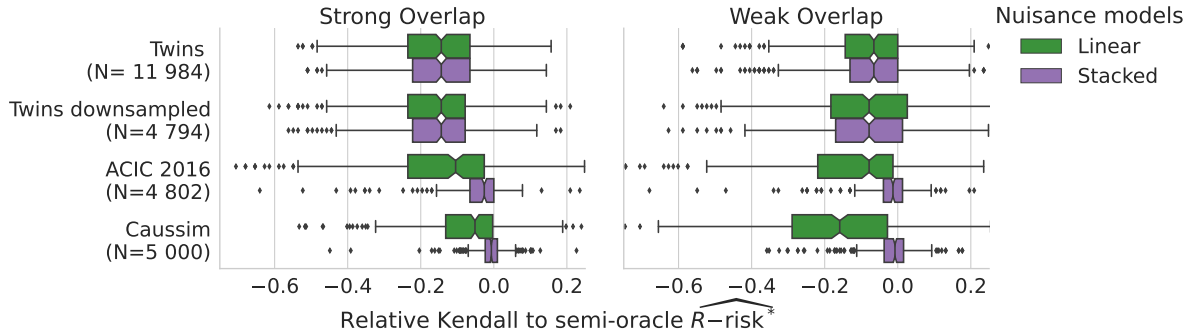


Figure 8: **Stacked models are good overall estimators of the nuisances**: Results are shown only for the R-risk. Details for every metrics are provided in Figure 18. For Twins, where the true propensity model is linear, **stacked** and **linear** estimations of the nuisances performs equivalently, even for a downsampled version (N=4794).

Use 90% of the data to estimate outcome models, 10% to select them For best causal modeling, the analyst often faces a compromise: given a finite data sample, should she allocate most of the data to estimate the outcome model, thus maximizing chances of achieving a high-quality outcome model but leaving little data for model selection. Or, she could choose a bigger test set for model selection and effect estimation. For causal model selection, there is no established practice (as reviewed in Appendix F).

We investigate such tradeoff varying the ratio between train and test data size. For this, we first split out 30% of the data as a holdout set \mathcal{V} on which we use the oracle response functions to derive silver-standard estimates of the causal quantities of interest. We then use the standard estimation procedure on the remaining 70% of the data, splitting it into train \mathcal{T} and test \mathcal{S} of varying sizes. We finally measure the error between this estimate and the silver-standard one.

We consider two different analytic goals: estimating a average treatment effect –a single number used for policy making– and a CATE –A full model of the treatment effect as a function of covariates X . Given that the latter is a much more complex object than the former, the optimal train/test ratio might vary. To measure errors, we use for the ATE the relative absolute ATE bias between the ATE computed with the selected outcome model on the test set, and the true ATE as evaluated on the holdout set \mathcal{V} . For the CATE, we compare the τ -risk of the best selected model applied on the holdout set \mathcal{V} . We explore this trade-off for the ACIC 2016 dataset and the R-risk.

Figure 9 shows that a train/test ratio of 0.9/0.1 (K=10) appears as a good trade-off for CATE and ATE estimation, though there is little difference with a split of 0.8/0.2 (K=5).

5 Discussion and conclusion

Predictive models are increasingly used to reason about causal effects. Our results highlight that they should be selected, validated, and tuned using different procedures and error measures than those classically used to assess prediction (estimating the so-called μ -risk). Rather, selecting the best outcome model according to the R -risk (eq. 6) leads to more valid causal estimates. Estimating this risk requires a markedly more complex procedure than standard cross-validation used *e.g.* in machine learning: it involves fitting nuisance models necessary for model evaluation, though our empirical results show that these can be learned on the same set of data as the outcome model evaluated. A poor estimation of the nuisance models may compromise the benefits of the more complex R -risk (as shown in Figure 8). However controlling and selecting these latter models is easier because they are associated to errors on observed distributions. Our empirical results show that when selecting these models in a flexible family of models the R -risk dominates simpler risks for model selection. Results show that going from an oracle R -risk –where the nuisances are known– to a feasible R -risk –where the nuisances are estimated– decreases only very slightly the model-selection

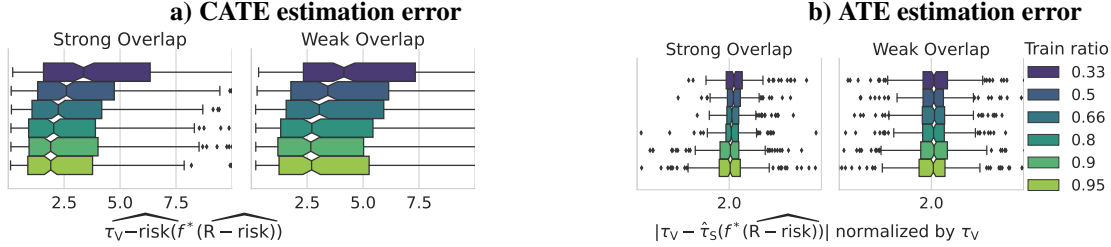


Figure 9: **a) For CATE, a train/test ratio of 0.9/0.1 appears as a good trade-off.** b) For ATE, there is a small signal pointing also to 0.9/0.1 ($K=10$). for ATE. Experiences on 10 replications of all 78 instances of the ACIC 2016 data.

performance of the R -risk. This may be explained by theoretical results that suggest that estimation errors on both nuisances partly compensate out in the R -risk [18, 51, 40, 54, 14].

For strong overlap, the conventional procedure (μ -risk) appears theoretically motivated (subsection 3.1), however empirical results show that even in this regime the R -risk brings a sizeable benefit, in agreement with Schuler et al. [73].

Our results points to $K=5$ or 10 for the choice of the number of folds used in cross-validation for CATE and ATE, aligned with previous empirical evidence for ATE [14].

Extension to binary outcome While we focused on continuous outcomes, in medicine, the target outcome is often a categorical variable such as mortality status or diagnosis. In this case, it may be interesting to focus on other estimands than the Average Treatment Effect $\mathbb{E}[Y(1)] - \mathbb{E}[Y(0)]$, for instance the relative risk $\frac{\mathbb{P}(Y(1)=1)}{\mathbb{P}(Y(0)=1)}$ or the odd ratio, $\frac{\mathbb{P}(Y(1)=1)/[1-\mathbb{P}(Y(1)=1)]}{\mathbb{P}(Y(0)=1)/[1-\mathbb{P}(Y(0)=1)]}$ are often used [6]. While the odds ratio is natural for case-control studies [69], a good choice of measure can reduce heterogeneity [15]. Using as an estimand the log of these values is suitable to additive models (for reasoning or noise assumptions). In the log domain, the relative risk or the odds ratio are written as a difference, as the ATE: $\log \mathbb{P}(Y(1) = 1) - \log \mathbb{P}(Y(0) = 1)$ or $\log(\mathbb{P}(Y(1) = 1)/[1 - \mathbb{P}(Y(1) = 1)]) - \log \mathbb{P}(Y(0) = 1)/[1 - \mathbb{P}(Y(0) = 1)]$. Hence, the framework studied here (subsection 2.1) can directly apply. It is particularly easy for the log odds ratio, as it is the output of a logistic regression or any model with a cross-entropy loss.

Going further The R -risk needs good estimation of nuisance models. The propensity score e calls for a control on the estimation of the individual posterior probability. We have used the Brier score to select these models, as it is minimized by the true individual probability. Regarding model-selection for propensity score, an easy mistake is to use expected calibration errors popular in machine learning [59, 87, 53, 49] as these select not for the individual posterior probability but for an aggregate error rate [57]. An open question is whether a better metric than the brier score can be designed that controls for $e(1 - e)$, the quantity used in the R -risk, rather than e .

The quality of model selection varies substantially from one data-generating mechanism to another. The overlap appears as an important parameter: when the treated and untreated, causal model selection is very hard. However, remaining variance in the empirical results suggests that other parameters of the data generation processes come into play. Intuitively, the complexity of the response surfaces and the treatment heterogeneity interact with overlap violations: when extrapolations to weak-overlap regions is hard, causal model selection is hard.

Nevertheless, from a practical perspective, our study establishes that the R -risk is the best option to select predictive models for causal inference, without requiring assumptions on the data-generating mechanism, the amount of data at hand, or the specific estimators used to build predictive models.

Acknowledgments

We acknowledge fruitful discussions with Bénédicte Colnet.

Financial disclosure

None reported.

Conflict of interest

The authors declare no potential conflict of interests.

References

- [1] Alaa, Ahmed and Schaar, Mihaela Van Der. “Validating Causal Inference Models via Influence Functions”. In: *International Conference on Machine Learning* (May 24, 2019), pp. 191–201.
- [2] Almond, Douglas, Chay, Kenneth Y., and Lee, David S. “The Costs of Low Birth Weight”. In: *The Quarterly Journal of Economics* 120.3 (2005). Publisher: Oxford University Press, pp. 1031–1083.
- [3] Altman, Douglas G, Vergouwe, Yvonne, Royston, Patrick, and Moons, Karel GM. “Prognosis and prognostic research: validating a prognostic model”. In: *Bmj* 338 (2009).
- [4] Athey, Susan and Imbens, Guido. “Recursive partitioning for heterogeneous causal effects”. In: *Proceedings of the National Academy of Sciences* 113.27 (2016), pp. 7353–7360.
- [5] Athey, Susan, Tibshirani, Julie, and Wager, Stefan. “Generalized random forests”. In: *Annals of Statistics* 47.2 (Apr. 2019). Publisher: Institute of Mathematical Statistics, pp. 1148–1178.
- [6] Austin, Peter C and Stuart, Elizabeth A. “Estimating the effect of treatment on binary outcomes using full matching on the propensity score”. In: *Statistical methods in medical research* 26.6 (2017), pp. 2505–2525.
- [7] Austin, Peter C. “An Introduction to Propensity Score Methods for Reducing the Effects of Confounding in Observational Studies”. In: *Multivariate Behavioral Research* 46.3 (May 2011), pp. 399–424.
- [8] Austin, Peter C. and Stuart, Elizabeth A. “Moving towards best practice when using inverse probability of treatment weighting (IPTW) using the propensity score to estimate causal treatment effects in observational studies”. In: *Statistics in Medicine* 34.28 (Dec. 10, 2015), pp. 3661–3679.
- [9] Battocchi, Keith, Dillon, Eleanor, Hei, Maggie, Lewis, Greg, Oka, Paul, Oprescu, Miruna, and Syrgkanis, Vasilis. *EconML: A Python Package for ML-Based Heterogeneous Treatment Effects Estimation*. <https://github.com/microsoft/EconML>. Version 0.14.0. 2019.
- [10] Black, Nick. “Why we need observational studies to evaluate the effectiveness of health care”. In: *Bmj* 312.7040 (1996), pp. 1215–1218.
- [11] Blakely, Tony, Lynch, John, Simons, Koen, Bentley, Rebecca, and Rose, Sherri. “Reflection on modern methods: when worlds collide—prediction, machine learning and causal inference”. In: *International journal of epidemiology* 49.6 (2020), pp. 2058–2064.
- [12] Bouthillier, Xavier, Delaunay, Pierre, Bronzi, Mirko, Trofimov, Assya, Nichyporuk, Brennan, Szeto, Justin, Mohammadi Sepahvand, Nazanin, Raff, Edward, Madan, Kanika, Voleti, Vikram, Ebrahimi Kahou, Samira, Michalski, Vincent, Arbel, Tal, Pal, Chris, Varoquaux, Gael, and Vincent, Pascal. “Accounting for Variance in Machine Learning Benchmarks”. In: *Proceedings of Machine Learning and Systems* 3 (Mar. 15, 2021), pp. 747–769.
- [13] Charlson, Mary E., Pompei, Peter, Ales, Kathy L., and MacKenzie, C.Ronald. “A new method of classifying prognostic comorbidity in longitudinal studies: Development and validation”. In: *Journal of Chronic Diseases* 40.5 (Jan. 1987), pp. 373–383.
- [14] Chernozhukov, Victor, Chetverikov, Denis, Demirer, Mert, Duflo, Esther, Hansen, Christian, Newey, Whitney, and Robins, James. “Double/Debiased Machine Learning for Treatment and Structural Parameters”. In: *The Econometrics Journal* (2018), p. 71.
- [15] Colnet, Bénédicte, Josse, Julie, Varoquaux, Gaël, and Scornet, Erwan. “Risk ratio, odds ratio, risk difference... Which causal measure is easier to generalize?” In: *arXiv preprint arXiv:2303.16008* (2023).
- [16] Curth, Alicia, Svensson, David, and Weatherall, James. “Really Doing Great at Estimating CATE? A Critical Look at ML Benchmarking Practices in Treatment Effect Estimation”. In: *Neurips Process 2021* (2021), p. 14.
- [17] D’Amour, Alexander, Ding, Peng, Feller, Avi, Lei, Lihua, and Sekhon, Jasjeet. “Overlap in observational studies with high-dimensional covariates”. In: *Journal of Econometrics* 221.2 (2021), pp. 644–654.
- [18] Daniel, Rhian M. “Double Robustness”. In: *Wiley StatsRef: Statistics Reference Online*. John Wiley & Sons, Ltd, 2018, pp. 1–14.
- [19] Desai, Rishi J, Wang, Shirley V, Vaduganathan, Muthiah, Evers, Thomas, and Schneeweiss, Sebastian. “Comparison of machine learning methods with traditional models for use of administrative claims with electronic medical records to predict heart failure outcomes”. In: *JAMA network open* 3.1 (2020), e1918962–e1918962.

- [20] Dorie, Vincent, Hill, Jennifer, Shalit, Uri, Scott, Marc, and Cervone, Dan. “Automated versus do-it-yourself methods for causal inference: Lessons learned from a data analysis competition”. In: *Statistical Science* 34.1 (2019), pp. 43–68.
- [21] Dudley, Joel T, Deshpande, Tarangini, and Butte, Atul J. “Exploiting drug–disease relationships for computational drug repositioning”. In: *Briefings in bioinformatics* 12.4 (2011), pp. 303–311.
- [22] Fontana, Mark Alan, Lyman, Stephen, Sarker, Gourab K, Padgett, Douglas E, and MacLean, Catherine H. “Can machine learning algorithms predict which patients will achieve minimally clinically important differences from total joint arthroplasty?” In: *Clinical orthopaedics and related research* 477.6 (2019), p. 1267.
- [23] Gao, Zijun, Hastie, Trevor, and Tibshirani, Robert. “Assessment of heterogeneous treatment effect estimation accuracy via matching”. In: *Statistics in Medicine* 40.17 (2021).
- [24] Gretton, Arthur, Borgwardt, Karsten M, Rasch, Malte J, Schölkopf, Bernhard, and Smola, Alexander. “A kernel two-sample test”. In: *The Journal of Machine Learning Research* 13.1 (2012), pp. 723–773.
- [25] Grose, Elysia, Wilson, Samuel, Barkun, Jeffrey, Bertens, Kimberly, Martel, Guillaume, Balaa, Fady, and Khalil, Jad Abou. “Use of Propensity Score Methodology in Contemporary High-Impact Surgical Literature”. In: *Journal of the American College of Surgeons* 230.1 (Jan. 1, 2020). Publisher: Elsevier, 101–112.e2.
- [26] Gruber, Susan and van der Laan, Mark J. “tmle: An R Package for Targeted Maximum Likelihood Estimation”. In: *Journal of Statistical Software* 51.13 (2012). doi:10.18637/jss.v051.i13, pp. 1–35.
- [27] Gutierrez, Pierre and Gerardy, Jean-Yves. “Causal Inference and Uplift Modeling A review of the literature”. In: *Proceedings of The 3rd International Conference on Predictive Applications and APIs*. Proceedings of Machine Learning Research 67 (2016), p. 14.
- [28] Hernán, MA and Robins, JM. *Causal Inference: What If*. 2020.
- [29] Hernán, Miguel A. “Methods of Public Health Research — Strengthening Causal Inference from Observational Data”. In: *New England Journal of Medicine* 385.15 (Oct. 7, 2021), pp. 1345–1348.
- [30] Hill, Jennifer L. “Bayesian Nonparametric Modeling for Causal Inference”. In: *Journal of Computational and Graphical Statistics* 20.1 (Jan. 1, 2011), pp. 217–240.
- [31] Hoogland, Jeroen, IntHout, Joanna, Belias, Michail, Rovers, Maroeska M, Riley, Richard D, E. Harrell Jr, Frank, Moons, Karel GM, Debray, Thomas PA, and Reitsma, Johannes B. “A tutorial on individualized treatment effect prediction from randomized trials with a binary endpoint”. In: *Statistics in medicine* 40.26 (2021), pp. 5961–5981.
- [32] Horng, Steven, Sontag, David A, Halpern, Yoni, Jernite, Yacine, Shapiro, Nathan I, and Nathanson, Larry A. “Creating an automated trigger for sepsis clinical decision support at emergency department triage using machine learning”. In: *PLoS one* 12.4 (2017), e0174708.
- [33] Howe, Channele J., Cole, Stephen R., Westreich, Daniel J., Greenland, Sander, Napravnik, Sonia, and Eron, Joseph J. “Splines for trend analysis and continuous confounder control”. In: *Epidemiology (Cambridge, Mass.)* 22.6 (Nov. 2011), pp. 874–875.
- [34] Hurler, Mark R, Yang, Lun, Xie, Qing, Rajpal, Deepak K, Sanseau, Philippe, and Agarwal, Pankaj. “Computational drug repositioning: from data to therapeutics”. In: *Clinical Pharmacology & Therapeutics* 93.4 (2013), pp. 335–341.
- [35] Imbens, Guido W. and Rubin, Donald B. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press, 2015.
- [36] Ionides, Edward L. “Truncated Importance Sampling”. In: *Journal of Computational and Graphical Statistics* 17.2 (June 2008), pp. 295–311.
- [37] Jesson, Andrew, Mindermann, Sören, Shalit, Uri, and Gal, Yarin. “Identifying Causal-Effect Inference Failure with Uncertainty-Aware Models”. In: *Advances in Neural Information Processing Systems* 33 (Oct. 22, 2020), pp. 11637–11649.
- [38] Johansson, Fredrik D., Shalit, Uri, Kallus, Nathan, and Sontag, David. “Generalization Bounds and Representation Learning for Estimation of Potential Outcomes and Causal Effects”. In: *arXiv:2001.07426 [cs, stat]* (Mar. 17, 2021).
- [39] Kendall, M. G. “A new measure of rank correlation”. In: *Biometrika* 30.1-2 (June 1938), pp. 81–93.
- [40] Kennedy, Edward H. “Optimal doubly robust estimation of heterogeneous causal effects”. In: *arXiv preprint arXiv:2004.14497* (2020).

- [41] Künzel, Sören R., Sekhon, Jasjeet S., Bickel, Peter J., and Yu, Bin. “Metalearners for estimating heterogeneous treatment effects using machine learning”. In: *Proceedings of the National Academy of Sciences* 116.10 (Mar. 5, 2019). Publisher: National Academy of Sciences Section: PNAS Plus, pp. 4156–4165.
- [42] Laan, Mark J van der, Laan, MJ, and Robins, JM. *Unified methods for censored longitudinal data and causality*. Springer Science & Business Media, 2003.
- [43] Laan, Mark J. van der, Polley, Eric C., and Hubbard, Alan E. “Super Learner”. In: *Statistical Applications in Genetics and Molecular Biology* 6 (Sept. 16, 2007).
- [44] Laan, Mark J. van der and Rose, Sherri. *Targeted Learning*. Springer Series in Statistics. 2011.
- [45] Lamont, Andrea, Lyons, Michael D, Jaki, Thomas, Stuart, Elizabeth, Feaster, Daniel J, Tharmaratnam, Kukatharmini, Oberski, Daniel, Ishwaran, Hemant, Wilson, Dawn K, and Van Horn, M Lee. “Identification of predicted individual treatment effects in randomized clinical trials”. In: *Statistical methods in medical research* 27.1 (2018), pp. 142–157.
- [46] Loiseau, Nicolas, Trichelair, Paul, He, Maxime, Andreux, Mathieu, Zaslavskiy, Mikhail, Wainrib, Gilles, and Blum, Michael G. B. “External control arm analysis: an evaluation of propensity score approaches, G-computation, and doubly debiased machine learning”. In: *BMC Medical Research Methodology* 22 (2022).
- [47] Louizos, Christos, Shalit, Uri, Mooij, Joris, Sontag, David, Zemel, Richard, and Welling, Max. “Causal Effect Inference with Deep Latent-Variable Models”. In: *Advances in neural information processing systems* (Nov. 6, 2017).
- [48] MacDorman, M. F. and Atkinson, J. O. “Infant mortality statistics from the linked birth/infant death data set–1995 period data”. eng. In: *Monthly Vital Statistics Report* 46.6 Suppl 2 (Feb. 1998), pp. 1–22.
- [49] Minderer, Matthias, Djolonga, Josip, Romijnders, Rob, Hubis, Frances, Zhai, Xiaohua, Houlsby, Neil, Tran, Dustin, and Lucic, Mario. “Revisiting the Calibration of Modern Neural Networks”. In: *Advances in Neural Information Processing Systems* 34 (2021), pp. 15682–15694.
- [50] Mooney, Stephen J and Pejaver, Vikas. “Big data in public health: terminology, machine learning, and privacy”. In: *Annual review of public health* 39 (2018), p. 95.
- [51] Naimi, Ashley I, Mishler, Alan E, and Kennedy, Edward H. “Challenges in Obtaining Valid Causal Effect Estimates with Machine Learning Algorithms”. In: *American Journal of Epidemiology* (2021).
- [52] Naimi, Ashley I and Whitcomb, Brian W. “Defining and Identifying Average Treatment Effects”. In: *American Journal of Epidemiology* (2023).
- [53] Niculescu-Mizil, Alexandru and Caruana, Rich. “Predicting good probabilities with supervised learning”. en. In: *Proceedings of the 22nd international conference on Machine learning - ICML '05*. ACM Press, 2005, pp. 625–632.
- [54] Nie, Xinkun and Wager, Stefan. “Quasi-Oracle Estimation of Heterogeneous Treatment Effects”. In: *Biometrika* 108.2 (Dec. 13, 2017), pp. 299–319.
- [55] Niswander, Kenneth R. and Stroke, United States National Institute of Neurological Diseases and. *The Women and Their Pregnancies: The Collaborative Perinatal Study of the National Institute of Neurological Diseases and Stroke*. Google-Books-ID: A0bdVhldQkC. National Institute of Health, 1972. 562 pp.
- [56] Pedregosa, Fabian, Varoquaux, Gaël, Gramfort, Alexandre, Michel, Vincent, Thirion, Bertrand, Grisel, Olivier, Blondel, Mathieu, Prettenhofer, Peter, Weiss, Ron, Dubourg, Vincent, Vanderplas, Jake, Passos, Alexandre, Cournapeau, David, Brucher, Matthieu, Perrot, Matthieu, and Duchesnay, Édouard. “Scikit-learn: Machine Learning in Python”. In: *Journal of Machine Learning Research* 12.85 (2011), pp. 2825–2830.
- [57] Perez-Lebel, Alexandre, Morvan, Marine Le, and Varoquaux, Gaël. “Beyond calibration: estimating the grouping loss of modern neural networks”. In: *arXiv preprint arXiv:2210.16315* (2022).
- [58] Perperoglou, Aris, Sauerbrei, Willi, Abrahamowicz, Michal, and Schmid, Matthias. “A review of spline function procedures in R”. In: *BMC Medical Research Methodology* 19.1 (Mar. 6, 2019), p. 46.
- [59] Platt, John C. and Platt, John C. “Probabilistic Outputs for Support Vector Machines and Comparisons to Regularized Likelihood Methods”. In: *Advances in Large Margin Classifiers* (1999), pp. 61–74.
- [60] Poldrack, Russell A, Huckins, Grace, and Varoquaux, Gael. “Establishment of best practices for evidence for prediction: a review”. In: *JAMA psychiatry* 77.5 (2020), pp. 534–540.
- [61] Powers, Scott, Qian, Junyang, Jung, Kenneth, Schuler, Alejandro, Shah, Nigam H., Hastie, Trevor, and Tibshirani, Robert. “Some methods for heterogeneous treatment effect estimation in high dimensions”. In: *Statistics in Medicine* 37.11 (2018), pp. 1767–1787.

- [62] Radley, David C, Finkelstein, Stan N, and Stafford, Randall S. “Off-label prescribing among office-based physicians”. In: *Archives of internal medicine* 166.9 (2006), pp. 1021–1026.
- [63] Rahimi, Ali and Recht, Benjamin. “Random Features for Large-Scale Kernel Machines”. In: *Advances in Neural Information Processing Systems*. Vol. 20. 2008.
- [64] Robins, James. “A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect”. In: *Mathematical Modelling* 7.9 (Jan. 1, 1986), pp. 1393–1512.
- [65] Robins, James M. and Greenland, Sander. “The role of model selection in causal inference from non experimental data”. In: *American Journal of Epidemiology* 123.3 (Mar. 1986), pp. 392–402.
- [66] Robinson, P. M. “Root-N-Consistent Semiparametric Regression”. In: *Econometrica* 56.4 (1988). Publisher: [Wiley, Econometric Society], pp. 931–954.
- [67] Rolling, Craig A. and Yang, Yuhong. “Model selection for estimating treatment effects”. In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 76.4 (Sept. 2014), pp. 749–769.
- [68] Rosenbaum, Paul R and Rubin, Donald B. “The central role of the propensity score in observational studies for causal effects”. In: *Biometrika* 70 (1983), pp. 41–55.
- [69] Rothman, KJ, Greenland, S, and Lash, TL. “Case-control studies, chapter 8”. In: *Modern epidemiology* (2008), pp. 111–127.
- [70] Rubin, Donald B. “Causal Inference Using Potential Outcomes”. In: *Journal of the American Statistical Association* 100.469 (Mar. 1, 2005). Publisher: Taylor & Francis, pp. 322–331.
- [71] Saito, Yuta and Yasui, Shota. “Counterfactual Cross-Validation: Stable Model Selection Procedure for Causal Inference Models”. In: *International Conference on Machine Learning*. PMLR. 2020, pp. 8398–8407.
- [72] Schulam, Peter and Saria, Suchi. “Reliable Decision Support using Counterfactual Models”. In: *Advances in neural information processing systems* 30 (Mar. 30, 2017).
- [73] Schuler, Alejandro, Baiocchi, Michael, Tibshirani, Robert, and Shah, Nigam. “A comparison of methods for model selection when estimating individual treatment effects”. In: *arXiv:1804.05146 [cs, stat]* (June 13, 2018).
- [74] Schuler, Megan S. and Rose, Sherri. “Targeted Maximum Likelihood Estimation for Causal Inference in Observational Studies”. In: *American Journal of Epidemiology* 185.1 (Jan. 1, 2017), pp. 65–73.
- [75] Shalit, Uri, Johansson, Fredrik D, and Sontag, David. “Estimating individual treatment effect: generalization bounds and algorithms”. In: *International Conference on Machine Learning*. PMLR. 2017, pp. 3076–3085.
- [76] Shimoni, Yishai, Karavani, Ehud, Ravid, Sivan, Bak, Peter, Ng, Tan Hung, Alford, Sharon Hensley, Meade, Denise, and Goldschmidt, Yaara. “An Evaluation Toolkit to Guide Model Selection and Cohort Definition in Causal Inference”. In: *arXiv preprint arXiv:1906.00442* (2019).
- [77] Shimoni, Yishai, Yanover, Chen, Karavani, Ehud, and Goldschmidt, Yaara. “Benchmarking Framework for Performance-Evaluation of Causal Inference Analysis”. In: *arXiv:1802.05046 [cs, stat]* (Mar. 20, 2018).
- [78] Simon, Gregory E, Johnson, Eric, Lawrence, Jean M, Rossom, Rebecca C, Ahmedani, Brian, Lynch, Frances L, Beck, Arne, Waitzfelder, Beth, Ziebell, Rebecca, Penfold, Robert B, et al. “Predicting suicide attempts and suicide deaths following outpatient visits using electronic health records”. In: *American Journal of Psychiatry* 175.10 (2018), pp. 951–960.
- [79] Snowden, Jonathan M., Rose, Sherri, and Mortimer, Kathleen M. “Implementation of G-computation on a simulated data set: demonstration of a causal inference technique”. In: *American Journal of Epidemiology* 173.7 (Apr. 2011), pp. 731–738.
- [80] Sriperumbudur, Bharath K., Fukumizu, Kenji, Gretton, Arthur, Schölkopf, Bernhard, and Lanckriet, Gert R. G. “On integral probability metrics, ϕ -divergences and binary classification”. In: *arXiv:0901.2698 [cs, math]* (Oct. 12, 2009).
- [81] Su, Xiaogang, Peña, Annette T, Liu, Lei, and Levine, Richard A. “Random forests of interaction trees for estimating individualized treatment effects in randomized trials”. In: *Statistics in medicine* 37.17 (2018), pp. 2547–2560.
- [82] Swaminathan, Adith and Joachims, Thorsten. “Counterfactual risk minimization: Learning from logged bandit feedback”. In: *International Conference on Machine Learning*. PMLR. 2015, pp. 814–823.
- [83] Varoquaux, Gaël and Colliot, Olivier. *Evaluating machine learning models and their diagnostic value*. 2022.
- [84] Wager, Stefan and Athey, Susan. “Estimation and Inference of Heterogeneous Treatment Effects using Random Forests”. In: *Journal of the American Statistical Association* 113.523 (July 3, 2018), pp. 1228–1242.

-
- [85] Wendling, T., Jung, K., Callahan, A., Schuler, A., Shah, N. H., and Gallego, B. “Comparing methods for estimation of heterogeneous treatment effects using observational data from health care databases”. In: *Statistics in Medicine* 37.23 (2018), pp. 3309–3324.
 - [86] Yurkovich, Marko, Avina-Zubieta, J Antonio, Thomas, Jamie, Gorenchtein, Mike, and Lacaille, Diane. “A systematic review identifies valid comorbidity indices derived from administrative health data”. In: *Journal of clinical epidemiology* 68.1 (2015), pp. 3–14.
 - [87] Zadrozny, Bianca and Elkan, Charles. “Obtaining calibrated probability estimates from decision trees and naive Bayesian classifiers”. en. In: (2001), p. 8.

A Variability of ATE estimation on ACIC 2016

Figure 1 shows ATE estimations for six different models used in g-computation estimators on the 76 configurations of the ACIC 2016 dataset. Outcome models are fitted on half of the data and inference is done on the other half –ie. train/test with a split ratio of 0.5. For each configuration, and each model, this train test split was repeated ten times, yielding non parametric variance estimates [12].

Outcome models are implemented with `scikit-learn` [56] and the following hyper-parameters:

Outcome Model	Hyper-parameters grid
Random Forests	Max depth: [2, 10]
Ridge regression without treatment interaction	Ridge regularization: [0.1]
Ridge regression with treatment interaction	Ridge regularization: [0.1]

Table 2: Hyper-parameters grid used for ACIC 2016 ATE variability

B Causal assumptions

We assume the following four assumptions, referred as strong ignorability and necessary to assure identifiability of the causal estimands with observational data [70]:

Assumption 1 (Unconfoundedness)

$$\{Y(0), Y(1)\} \perp\!\!\!\perp A | X$$

This condition –also called ignorability– is equivalent to the conditional independence on $e(X)$ [68]: $\{Y(0), Y(1)\} \perp\!\!\!\perp A | e(X)$.

Assumption 2 (Overlap, also known as Positivity)

$$\eta < e(x) < 1 - \eta \quad \forall x \in \mathcal{X} \text{ and some } \eta > 0$$

The treatment is not perfectly predictable. Or with different words, every patient has a chance to be treated and not to be treated. For a given set of covariates, we need examples of both to recover the ATE.

As noted by [17], the choice of covariates X can be viewed as a trade-off between these two central assumptions. A bigger covariates set generally reinforces the ignorability assumption. In the contrary, overlap can be weakened by large \mathcal{X} because of the potential inclusion of instruments: variables only linked to the treatment which could lead to arbitrarily small propensity scores.

Assumption 3 (Consistency) *The observed outcome is the potential outcome of the assigned treatment:*

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

Here, we assume that the intervention A has been well defined. This assumption focuses on the design of the experiment. It clearly states the link between the observed outcome and the potential outcomes through the intervention [28].

Assumption 4 (Generalization) *The training data on which we build the estimator and the test data on which we make the estimation are drawn from the same distribution \mathcal{D}^* , also known as the “no covariate shift” assumption [37].*

C Proofs: Links between feasible and oracle risks

C.1 Upper bound of τ -risk with μ -risk_{IPW}

For the bound with the μ -risk_{IPW}, we will decompose the CATE risk on each factual population risks:

Definition 7 (Population Factual μ -risk) [75]

$$\mu\text{-risk}_a(f) = \int_{\mathcal{Y} \times \mathcal{X}} (y - f(x; A = a))^2 p(y; x = x | A = a) dy dx$$

Applying Bayes rule, we can decompose the μ -risk on each intervention:

$$\mu\text{-risk}(f) = p_A \mu\text{-risk}_1(f) + (1 - p_A) \mu\text{-risk}_0(f) \text{ with } p_A = \mathbb{P}(A = 1)$$

These definitions allows to state a intermediary result on each population:

Lemma 1 (Mean-variance decomposition) *We need a reweighted version of the classical mean-variance decomposition.*

For an outcome model $f : x \times A \rightarrow \mathcal{X}$. Let the inverse propensity weighting function $w(a; x) = ae(x)^{-1} + (1 - a)(1 - e(x))^{-1}$.

$$\int_{\mathcal{X}} (\mu_1(x) - f(x; 1))^2 p(x) dx = p_A \mu\text{-risk}_{IPW,1}(w, f) - \sigma_{Bayes}^2(1)$$

And

$$\int_{\mathcal{X}} (\mu_0(x) - f(x; 0))^2 p(x) dx = (1 - p_A) \mu\text{-risk}_{IPW,0}(w, f) - \sigma_{Bayes}^2(0)$$

Proof 1

$$\begin{aligned} p_A \mu\text{-risk}_{IPW,1}(w, f) &= \int_{\mathcal{X} \times \mathcal{Y}} \frac{1}{e(x)} (y - f(x; 1))^2 p(y | x; A = 1) p(x; A = 1) dy dx \\ &= \int_{\mathcal{X} \times \mathcal{Y}} (y - f(x; 1))^2 p(y | x; A = 1) \frac{p(x; A = 1)}{p(x; A = 1)} p(x) dy dx \\ &= \int_{\mathcal{X} \times \mathcal{Y}} [(y - \mu_1(x))^2 + (\mu_1(x) - f(x; 1))^2 + 2(y - \mu_1(x))(\mu_1(x) - f(x; 1))] p(y | x; A = 1) p(x) dy dx \\ &= \int_{\mathcal{X}} \left[\int_{\mathcal{Y}} (y - \mu_1(x))^2 p(y | x; A = 1) dy \right] p(x) dx + \int_{\mathcal{X} \times \mathcal{Y}} (\mu_1(x) - f(x; 1))^2 p(x) p(y | x; A = 1) dx dy \\ &+ 2 \int_{\mathcal{X}} \left[\int_{\mathcal{Y}} (y - \mu_1(x)) p(y | x; A = 1) dy \right] (\mu_1(x) - f(x; 1)) p(x) dx \\ &= \int_{\mathcal{X}} \sigma_y^2(x, 1) p(x) dx + \int_{\mathcal{X}} (\mu_1(x) - f(x; 1))^2 p(x) dx + 0 \end{aligned}$$

Proposition 1 (Upper bound with mu-IPW) *Let f be a given outcome model, let the weighting function w be the Inverse Propensity Weight $w(x; a) = \frac{a}{e(x)} + \frac{1-a}{1-e(x)}$. Then, under overlap (assumption 2),*

$$\tau\text{-risk}(f) \leq 2 \mu\text{-risk}_{IPW}(w, f) - 2(\sigma_{Bayes}^2(1) + \sigma_{Bayes}^2(0))$$

Proof 2

$$\tau\text{-risk}(f) = \int_{\mathcal{X}} (\mu_1(x) - \mu_0(x) - (f(x; 1) - f(x; 0)))^2 p(x) dx$$

By the triangle inequality $(u + v)^2 \leq 2(u^2 + v^2)$:

$$\tau\text{-risk}(f) \leq 2 \int_{\mathcal{X}} [(\mu_1(x) - f(x; 1))^2 + (\mu_0(x) - f(x; 0))^2] p(x) dx$$

Applying Lemma 1,

$$\begin{aligned} \tau\text{-risk}(f) &\leq 2[p_A \mu\text{-risk}_{IPW,1}(w, f) + (1 - p_A) \mu\text{-risk}_{IPW,0}(w, f)] - 2(\sigma_{Bayes}^2(0) + \sigma_{Bayes}^2(1)) \\ &= 2 \mu\text{-risk}_{IPW}(w, f) - 2(\sigma_{Bayes}^2(0) + \sigma_{Bayes}^2(1)) \end{aligned}$$

C.2 Reformulation of the R -risk as reweighted τ -risk

Proposition 2 (R -risk as reweighted τ -risk) **Proof 3** *We consider the R -decomposition: [66],*

$$y(a) = m(x) + (a - e(x))\tau(x) + \varepsilon(x; a) \tag{17}$$

Where $\mathbb{E}[\varepsilon(X; A)|X, A] = 0$ We can use it as plug in the R-risk formula:

$$\begin{aligned}
R\text{-risk}(f) &= \int_{\mathcal{Y} \times \mathcal{X} \times \mathcal{A}} [(y - m(x)) - (a - e(x))\tau_f(x)]^2 p(y; x; a) dy dx da \\
&= \int_{\mathcal{Y} \times \mathcal{X} \times \mathcal{A}} [(a - e(x))\tau(x) + \varepsilon(x; a) - (a - e(x))\tau_f(x)]^2 p(y; x; a) dy dx da \\
&= \int_{\mathcal{X} \times \mathcal{A}} (a - e(x))^2 (\tau(x) - \tau_f(x))^2 p(x; a) dx da \\
&+ 2 \int_{\mathcal{Y} \times \mathcal{X} \times \mathcal{A}} (a - e(x)) (\tau(x) - \tau_f(x)) \int_{\mathcal{Y}} \varepsilon(x; a) p(y | x; a) dy p(x; a) dx da \\
&+ \int_{\mathcal{X} \times \mathcal{A}} \int_{\mathcal{Y}} \varepsilon^2(x; a) p(y | x; a) dy p(x; a) dx da
\end{aligned}$$

The first term can be decomposed on control and treated populations to force $e(x)$ to appear:

$$\begin{aligned}
&\int_{\mathcal{X}} (\tau(x) - \tau_f(x))^2 [e(x)^2 p(x; 0) + (1 - e(x))^2 p(x; 1)] dx \\
&= \int_{\mathcal{X}} (\tau(x) - \tau_f(x))^2 [e(x)^2 (1 - e(x)) p(x) + (1 - e(x))^2 e(x) p(x)] dx \\
&= \int_{\mathcal{X}} (\tau(x) - \tau_f(x))^2 (1 - e(x)) e(x) [1 - e(x) + e(x)] p(x) dx \\
&= \int_{\mathcal{X}} (\tau(x) - \tau_f(x))^2 (1 - e(x)) e(x) p(x) dx.
\end{aligned}$$

The second term is null since, $\mathbb{E}[\varepsilon(x, a)|X, A] = 0$.

The third term corresponds to the modulated residuals 3 : $\tilde{\sigma}_B^2(0) + \tilde{\sigma}_B^2(1)$

D Measuring overlap

Motivation of the Normalized Total Variation Computing overlap when working only on samples of the observed distribution, outside of simulation, requires a sophisticated estimator of discrepancy between distributions, as two data points never have the same exact set of features. Maximum Mean Discrepancy [24] is typically used in the context of causal inference [75, 38]. However it needs a kernel, typically Gaussian, to extrapolate across neighboring observations. We prefer avoiding the need to specify such a kernel, as it must be adapted to the data which is tricky with categorical or non-Gaussian features, a common situation for medical data.

For simulated and some semi-simulated data, we have access to the probability of treatment for each data point, which sample both densities in the same data point. Thus, we can directly use distribution discrepancy measures and rely on the Normalized Total Variation (NTV) distance to measure the overlap between the treated and control propensities. This is the empirical measure of the total variation distance [80] between the distributions, $TV(\mathbb{P}(X|A = 1), \mathbb{P}(X|A = 0))$. As we have both distribution sampled on the same points, we can rewrite it a sole function of the propensity score, a low dimensional score more tractable than the full distribution $\mathbb{P}(X|A)$:

$$\widehat{NTV}(e, 1 - e) = \frac{1}{2N} \sum_{i=1}^N \left| \frac{e(x_i)}{p_A} - \frac{1 - e(x_i)}{1 - p_A} \right| \quad (18)$$

Formally, we can rewrite NTV as the Total Variation distance between the two population distributions. For a population $O = (Y(A), X, A) \sim \mathcal{D}$:

$$\begin{aligned}
NTV(O) &= \frac{1}{2N} \sum_{i=1}^N \left| \frac{e(x_i)}{p_A} - \frac{1 - e(x_i)}{1 - p_A} \right| \\
&= \frac{1}{2N} \sum_{i=1}^N \left| \frac{P(A = 1|X = x_i)}{p_A} - \frac{P(A = 0|X = x_i)}{1 - p_A} \right|
\end{aligned}$$

Thus NTV approximates the following quantity in expectation over the data distribution \mathcal{D} :

$$\begin{aligned}
NTV(\mathcal{D}) &= \int_{\mathcal{X}} \left| \frac{p(A = 1|X = x)}{p_A} - \frac{p(A = 0|X = x)}{1 - p_A} \right| p(x) dx \\
&= \int_{\mathcal{X}} \left| \frac{p(A = 1, X = x)}{p_A} - \frac{p(A = 0, X = x)}{1 - p_A} \right| dx \\
&= \int_{\mathcal{X}} |p(X = x|A = 1) - p(X = x|A = 0)| dx
\end{aligned}$$

For countable sets, this expression corresponds to the Total Variation distance between treated and control populations covariate distributions : $TV(p_0(x), p_1(x))$.

Measuring overlap without the oracle propensity scores: For ACIC 2018, or for non-simulated data, the true propensity scores are not known. To measure overlap, we rely on flexible estimations of the Normalized Total Variation, using gradient boosting trees to approximate the propensity score. Empirical arguments for this plug-in approach is given in Figure 10.

Empirical arguments We show empirically that NTV is an appropriate measure of overlap by :

- Comparing the NTV distance with the MMD for Caussim which is gaussian distributed in Figure 12,
- Verifying that setups with penalized overlap from ACIC 2016 have a higher total variation distance than unpenalized setups in Figure 11.
- Verifying that the Inverse Propensity Weights extrema (the inverse of the ν overlap constant appearing in the overlap Assumption 2) positively correlates with NTV for Caussim, ACIC 2016 and Twins in Figure 13. Even if the same value of the maximum IPW could lead to different values of NTV, we expect both measures to be correlated : the higher the extrem propensity weights, the higher the NTV.

Estimating NTV in practice Finally, we verify that approximating the NTV distance with a learned plug-in estimates of $e(x)$ is reasonable. We used either a logistic regression or a gradient boosting classifier to learn the propensity models for the three datasets where we have access to the ground truth propensity scores: Caussim, Twins and ACIC 2016. We respectively sampled 1000, 1000 and 770 instances of these datasets with different seeds and overlap settings. We first run a hyperparameter search with cross-validation on the train set, then select the best estimator. We refit on the train set this estimator with or without calibration by cross validation and finally estimate the normalized TV with the obtained model. This training procedure reflects the one described in Algorithm 1 where nuisance models are fitted only on the train set.

The hyper parameters are : learning rate $\in [1e - 3, 1e - 2, 1e - 1, 1]$, minimum samples leaf $\in [2, 10, 50, 100, 200]$ for boosting and L2 regularization $\in [1e - 3, 1e - 2, 1e - 1, 1]$ for logistic regression.

Results in Figure 10 comparing bias to the true normalized Total Variation of each dataset instances versus growing true NTV indicate that calibration of the propensity model is crucial to recover a good approximation of the NTV.

E Experiments

E.1 Details on the data generation process

We use Gaussian-distributed covariates and random basis expansion based on Radial Basis Function kernels. A random basis of RBF kernel enables modeling non-linear and complex relationships between covariates in a similar way to

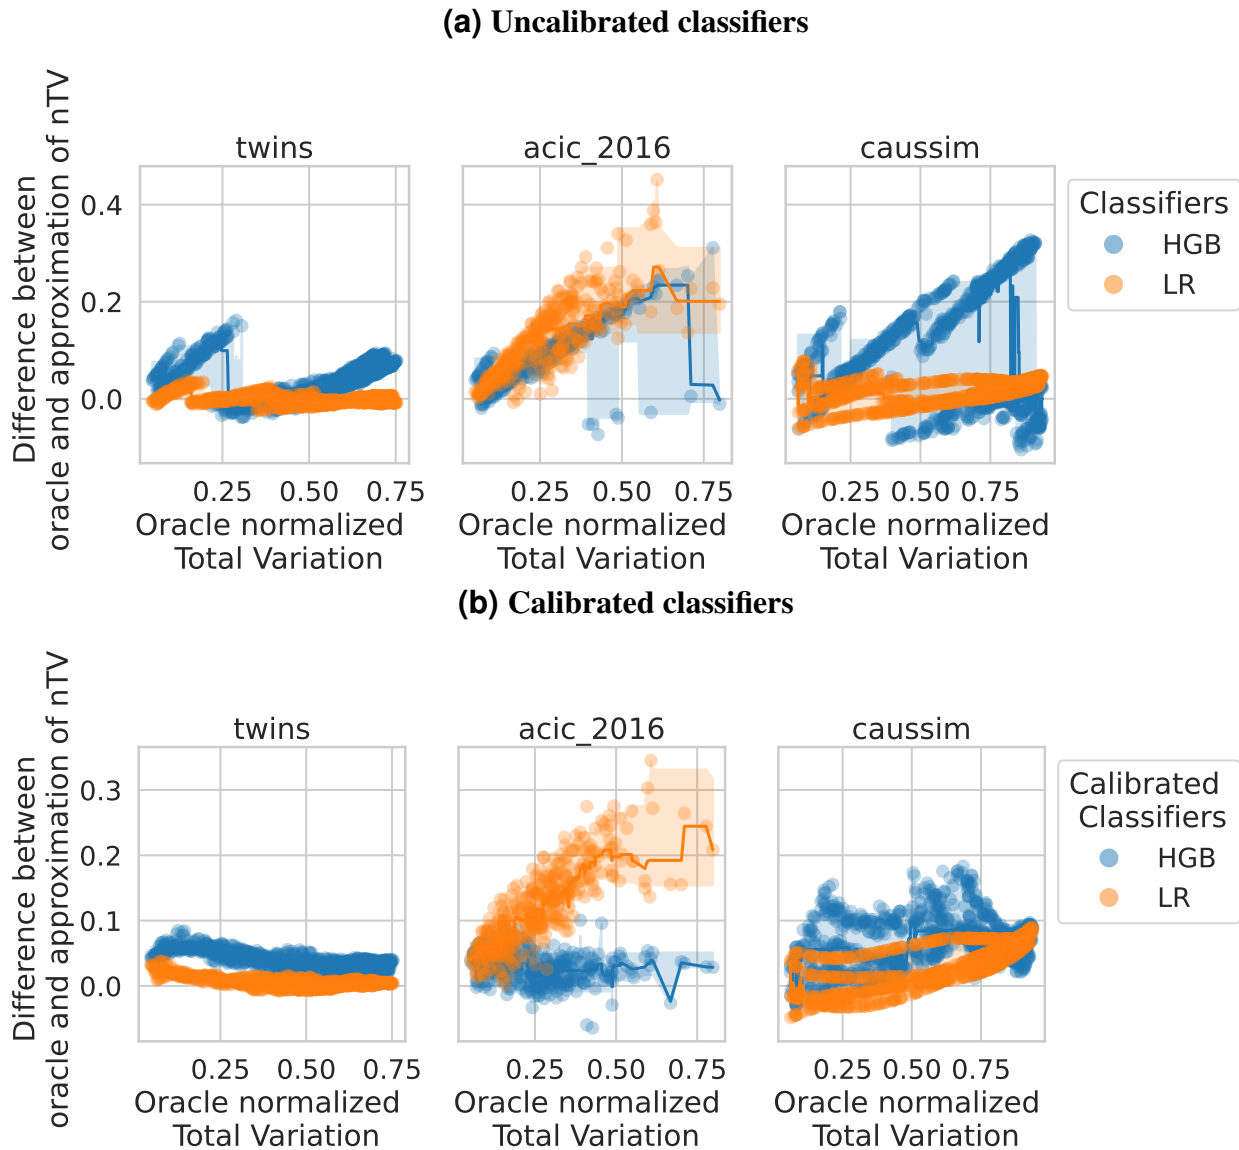


Figure 10: a) Without calibration, estimation of NTV is not trivial even for boosting models. b) Calibrated classifiers are able to recover the true Normalized Total Variation for all datasets where it is available.

Figure 11: NTV recovers well the overlap settings described in the ACIC paper [20]

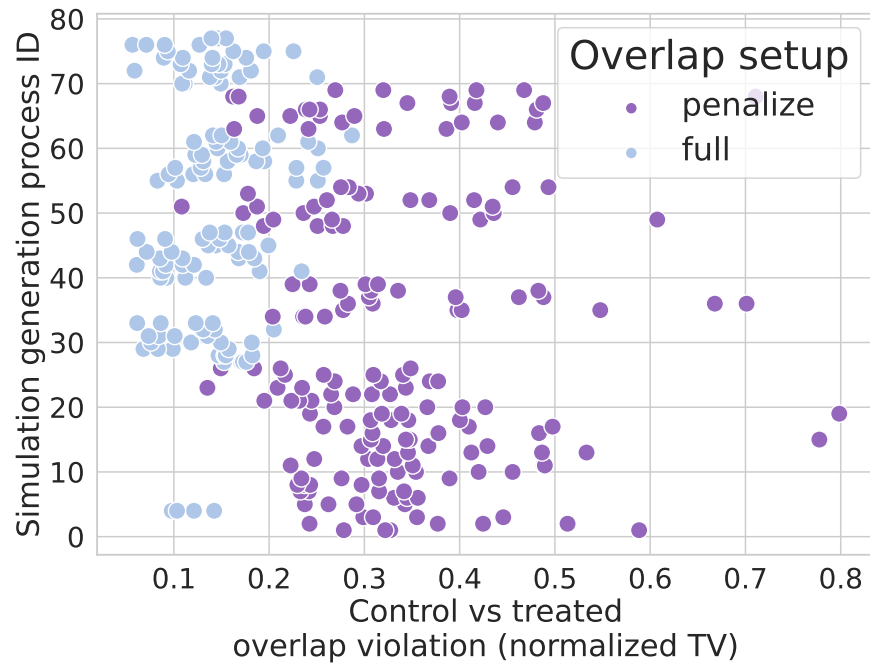
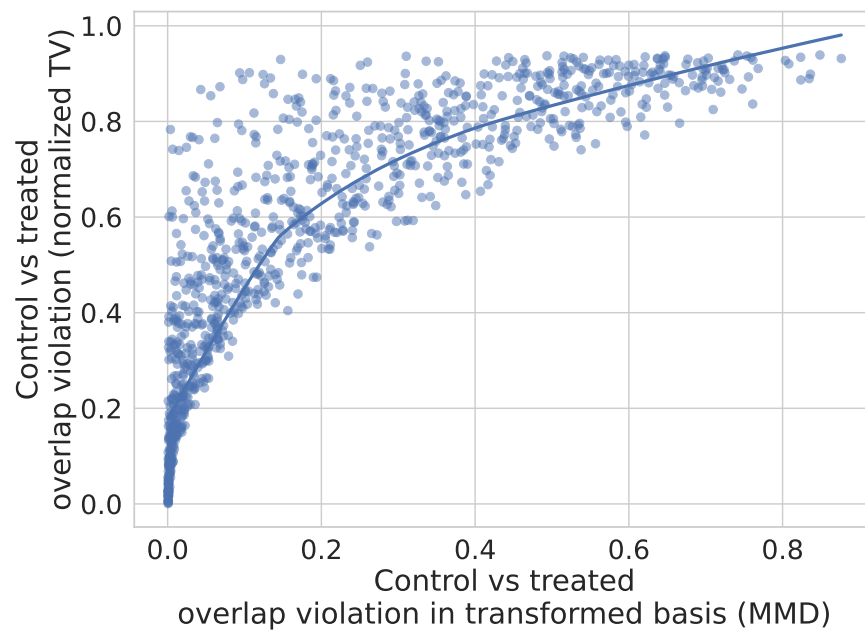


Figure 12: Good correlation between overlap measured as normalized Total Variation and Maximum Mean Discrepancy (200 sampled Caussim datasets)



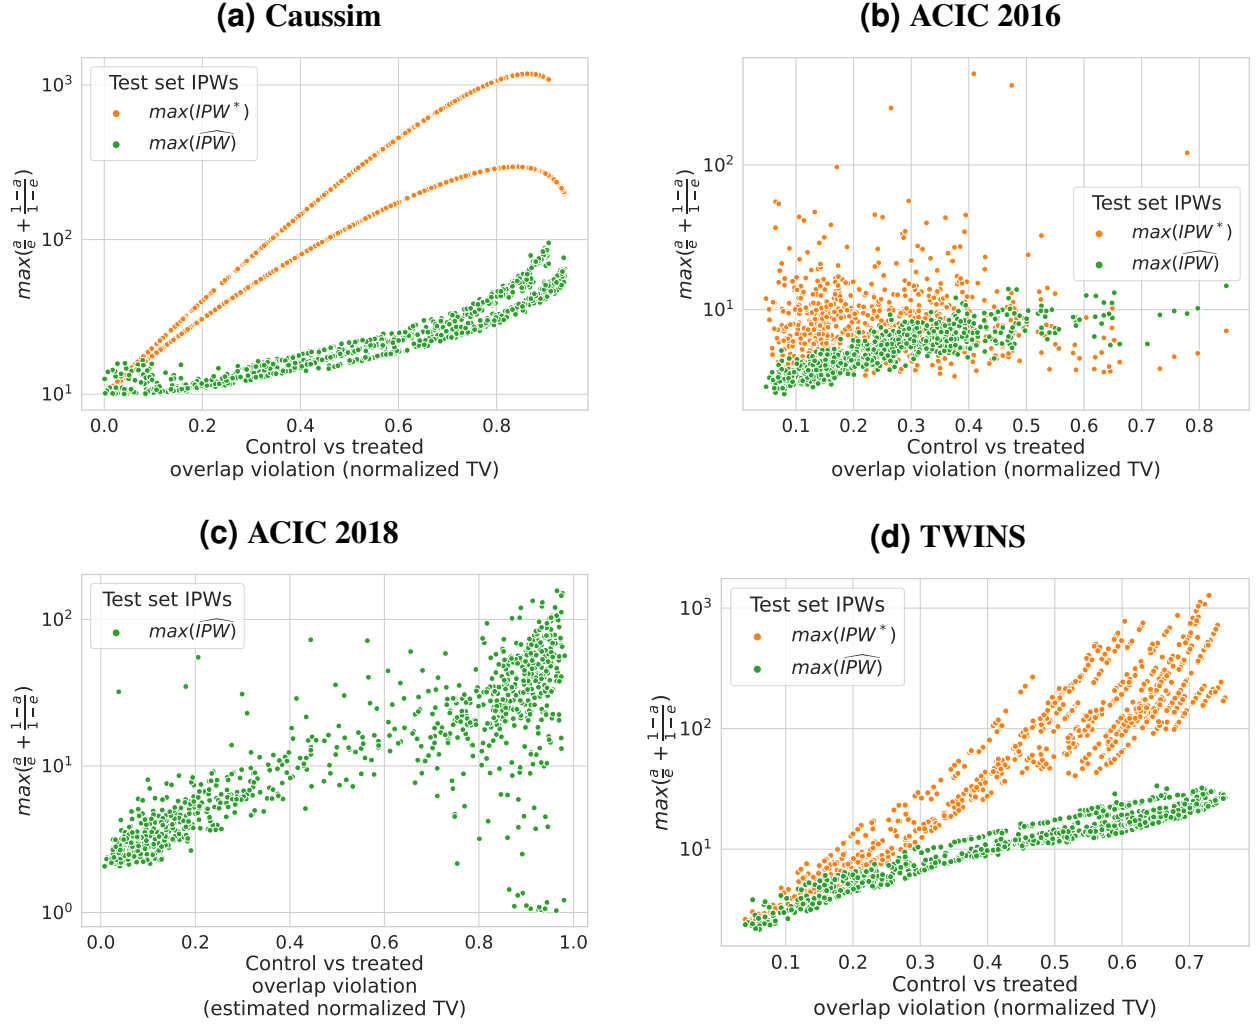


Figure 13: Maximal value of Inverse Propensity Weights increases exponentially with the overlap as measure by Normalized Total Variation.

the well known spline expansion. The estimators of the response function are learned with a linear model on another random basis (which can be seen as a stochastic approximation of the full data kernel [63]). We carefully control the amount of overlap between treated and control populations, a crucial assumption for causal inference.

- The raw features for both populations are drawn from a mixture of Gaussians: $\mathbb{P}(X) = p_A \mathbb{P}(X|A = 1) + (1 - p_A) \mathbb{P}(X|A = 0)$ where $\mathbb{P}(x|A = a)$ is a rotated Gaussian:

$$\mathbb{P}(x|A = a) = W \cdot \mathcal{N}\left(\begin{bmatrix} (1 - 2a)\theta \\ 0 \end{bmatrix}; \begin{bmatrix} \sigma_0 & 0 \\ 0 & \sigma_1 \end{bmatrix}\right) \quad (19)$$

with θ a parameter controlling overlap (bigger yields poorer overlap), W a random rotation matrix and $\sigma_0^2 = 2$; $\sigma_1^2 = 5$.

This generation process allows to analytically compute the oracle propensity scores $e(x)$, to simply control for overlap with the parameter θ , the distance between the two Gaussian main axes and to visualize response surfaces.

- A basis expansion of the raw features increases the problem dimension. Using Radial Basis Function (RBF) Nystroem transformation ⁷, we expand the raw features into a transformed space. The basis expansion sam-

⁷We use the [Sklearn implementation](#), [56]

ples randomly a small number of representers in the raw data. Then, it computes an approximation of the full N -dimensional kernel with these basis components, yielding the transformed features $z(x)$.

We generate the basis following the original data distribution, $[b_1..b_D] \sim \mathbb{P}(x)$, with $D=2$ in our simulations. Then, we compute an approximation of the full kernel of the data generation process $RBF(x, \cdot)$ with $x \sim \mathbb{P}(x)$ with these representers: $z(x) = [RBF_\gamma(x, b_d)]_{d=1..D} \cdot Z^T \in \mathbb{R}^D$ with RBF_γ being the Gaussian kernel $K(x, y) = \exp(-\gamma\|x - y\|^2)$ and Z the normalization constant of the kernel basis, computed as the root inverse of the basis kernel $Z = [K(b_i, b_j)]_{i,j \in 1..D}^{-1/2}$

- Functions μ_0, τ are distinct linear functions of the transformed features:

$$\mu_0(x) = [z(x); 1] \cdot \beta_\mu^T$$

$$\tau(x) = [z(x); 1] \cdot \beta_\tau^T$$

- Adding a Gaussian noise, $\varepsilon \sim \mathcal{N}(0, \sigma(x; a))$, we construct the potential outcomes: $y(a) = \mu_0(x) + a\tau(x) + \varepsilon(x, a)$

We generated 1000 instances of this dataset with uniformly random overlap parameters $\theta \in [0, 2.5]$.

E.2 Model selection procedures

Nuisances estimation The nuisances are estimated with a stacked regressor inspired by the Super Learner framework, [43]). The hyper-parameters are optimized with a random search with following search grid detailed in Table 3. All implementations come from `scikit-learn` [56].

Model	Estimator	Hyper-parameters grid
Outcome, m	StackedRegressor (HistGradientBoostingRegressor, ridge)	ridge regularization: [0.0001, 0.001, 0.01, 0.1, 1, 10, 100] HistGradientBoostingRegressor learning rate: [0.01, 0.1, 1] HistGradientBoostingRegressor max leaf nodes: [10, 20, 30, 50]
Treatment, e	StackedClassifier (HistgradientBoostingClassifier, LogisticRegression)	LogisticRegression C: [0.0001, 0.001, 0.01, 0.1, 1, 10, 100] HistGradientBoostingClassifier learning rate: [0.01, 0.1, 1] HistGradientBoostingClassifier max leaf nodes: [10, 20, 30, 50]

Table 3: Hyper-parameters grid used for nuisance models

E.3 Additional Results

Definition of the Kendall’s tau, κ The Kendall’s tau is a widely used statistics to measure the rank correlation between two set of observations. It measures the number of concordant pairs minus the discordant pairs normalized by the total number of pairs. It takes values in the $[-1, 1]$ range.

$$\kappa = \frac{(\text{number of concordant pairs}) - (\text{number of discordant pairs})}{(\text{number of pairs})} \quad (20)$$

Values of relative $\kappa(\ell, \tau\text{-risk})$ compared to the mean over all metrics Kendall’s as shown in the boxplots of Figure 5

Figure 14 - Results measured in absolute Kendall’s

Figure 15 - Results measured as distance to the oracle tau-risk To see practical gain in term of τ -risk, we plot the results as the normalized distance between the estimator selected by the oracle τ -risk and the estimator selected by each causal metric.

Then, $\widehat{R\text{-risk}}^*$ is more efficient than all other metrics. The gain are substantial for every datasets.

Figure 16 - Stacked models for the nuisances is more efficient For each metrics the benefit of using a stacked model of linear and boosting estimators for nuisances compared to a linear model. The evaluation measure is Kendall’s tau relative to the oracle $R\text{-risk}^*$ to have a stable reference between experiments. Thus, we do not include in this analysis the ACIC 2018 dataset since $R\text{-risk}^*$ is not available due to the lack of the true propensity score.

Metric	Dataset	Strong Overlap		Weak Overlap	
		Median	IQR	Median	IQR
$\widehat{\mu\text{-risk}}$	Twins (N=11 984)	-0.32	0.12	-0.19	0.12
	ACIC 2016 (N=4 802)	-0.03	0.13	0.11	0.19
	Caussim (N=5 000)	-0.40	0.55	-0.16	0.31
	ACIC 2018 (N=5 000)	0.00	0.30	0.01	0.40
$\widehat{\mu\text{-risk}}_{IPW}$	Twins (N= 11 984)	-0.31	0.13	-0.17	0.12
	ACIC 2016 (N=4 802)	-0.02	0.13	0.11	0.19
	Caussim (N=5 000)	-0.34	0.50	0.09	0.31
	ACIC 2018 (N=5 000)	0.00	0.30	-0.01	0.43
$\widehat{\mu\text{-risk}}_{IPW}^*$	Twins (N= 11 984)	-0.32	0.13	-0.17	0.13
	ACIC 2016 (N=4 802)	-0.02	0.13	0.11	0.21
	Caussim (N=5 000)	-0.33	0.54	0.26	0.27
	Twins (N= 11 984)	0.13	0.12	0.27	0.12
$\widehat{\tau\text{-risk}}_{IPW}$	ACIC 2016 (N=4 802)	-0.07	0.18	0.05	0.31
	Caussim (N=5 000)	-0.19	0.43	-0.14	0.18
	ACIC 2018 (N=5 000)	-0.16	0.40	-0.11	0.66
	Twins (N= 11 984)	0.12	0.14	0.20	0.16
$\widehat{\tau\text{-risk}}_{IPW}^*$	ACIC 2016 (N=4 802)	-0.03	0.16	-0.09	0.43
	Caussim (N=5 000)	-0.15	0.46	-0.17	0.19
	Twins (N= 11 984)	0.13	0.12	0.02	0.25
	ACIC 2016 (N=4 802)	0.04	0.11	0.11	0.26
$\widehat{U\text{-risk}}$	Caussim (N=5 000)	0.04	0.43	-0.04	0.17
	ACIC 2018 (N=5 000)	0.12	0.26	-0.02	0.50
	Twins (N= 11 984)	0.25	0.08	-0.41	0.45
	ACIC 2016 (N=4 802)	0.08	0.13	-0.59	0.57
$\widehat{U\text{-risk}}^*$	Caussim (N=5 000)	0.46	0.12	0.02	0.44
	Twins (N= 11 984)	0.15	0.10	0.25	0.18
	ACIC 2016 (N=4 802)	0.07	0.12	0.22	0.15
	Caussim (N=5 000)	0.34	0.26	0.13	0.21
$\widehat{R\text{-risk}}$	ACIC 2018 (N=5 000)	0.13	0.27	0.21	0.47
	Twins (N= 11 984)	0.25	0.10	0.32	0.15
	ACIC 2016 (N=4 802)	0.12	0.12	0.25	0.15
	Caussim (N=5 000)	0.47	0.11	0.16	0.14

Table 4: Values of relative $\kappa(\ell, \tau\text{-risk})$ compared to the mean over all metrics Kendall’s as shown in the boxplots of Figure 5

Figure 17 Low population overlap hinders model selection for all metrics

Figure 18 - Stacked models for the nuisances is more efficient For each metrics the benefit of using a stacked model of linear and boosting estimators for nuisances compared to a linear model. The evaluation measure is Kendall’s tau relative to the oracle $R\text{-risk}^*$ to have a stable reference between experiments. Thus, we do not include in this analysis the ACIC 2018 dataset since $R\text{-risk}^*$ is not available due to the lack of the true propensity score.

Figure 19 - Flexible models are performant in recovering nuisances even in linear setups

Selecting different seeds and parameters is crucial to draw conclusions One strength of our study is the various number of different simulated and semi-simulated datasets. We are convinced that the usual practice of using only a small number of generation processes does not allow to draw statistically significant conclusions.

Figure 20 illustrate the dependence of the results on the generation process for caussim simulations. We highlighted the different trajectories induced by three different seeds for data generation and three different treatment ratio instead of 1000 different seeds. The result curves are relatively stable from one setup to another for $R\text{-risk}$, but vary strongly for $\mu\text{-risk}$ and $\mu\text{-risk}_{IPW}$.

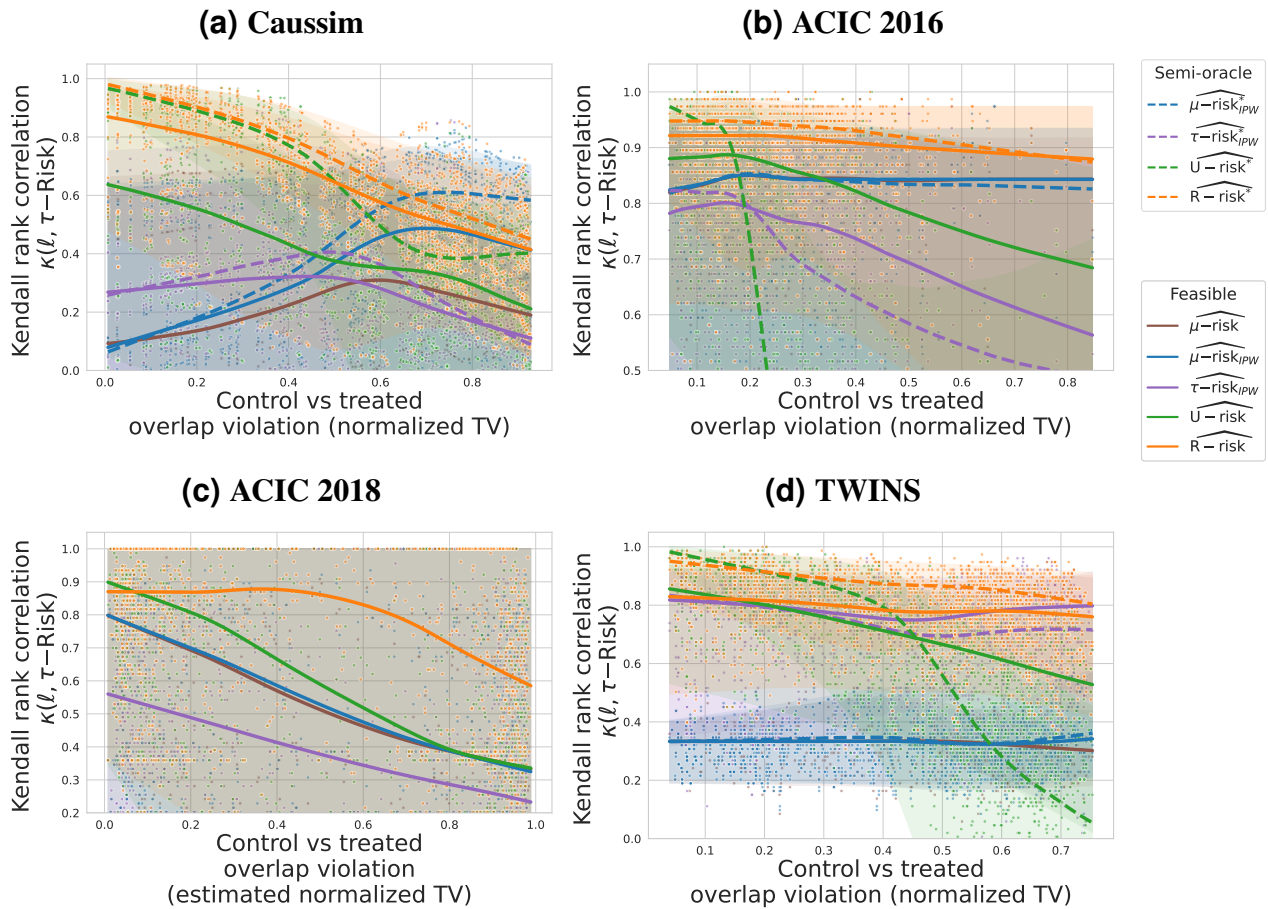


Figure 14: Agreement with τ -risk ranking of methods function of overlap violation. The lines represent medians, estimated with a lowess. The transparent bands denote the 5% and 95% confidence intervals.

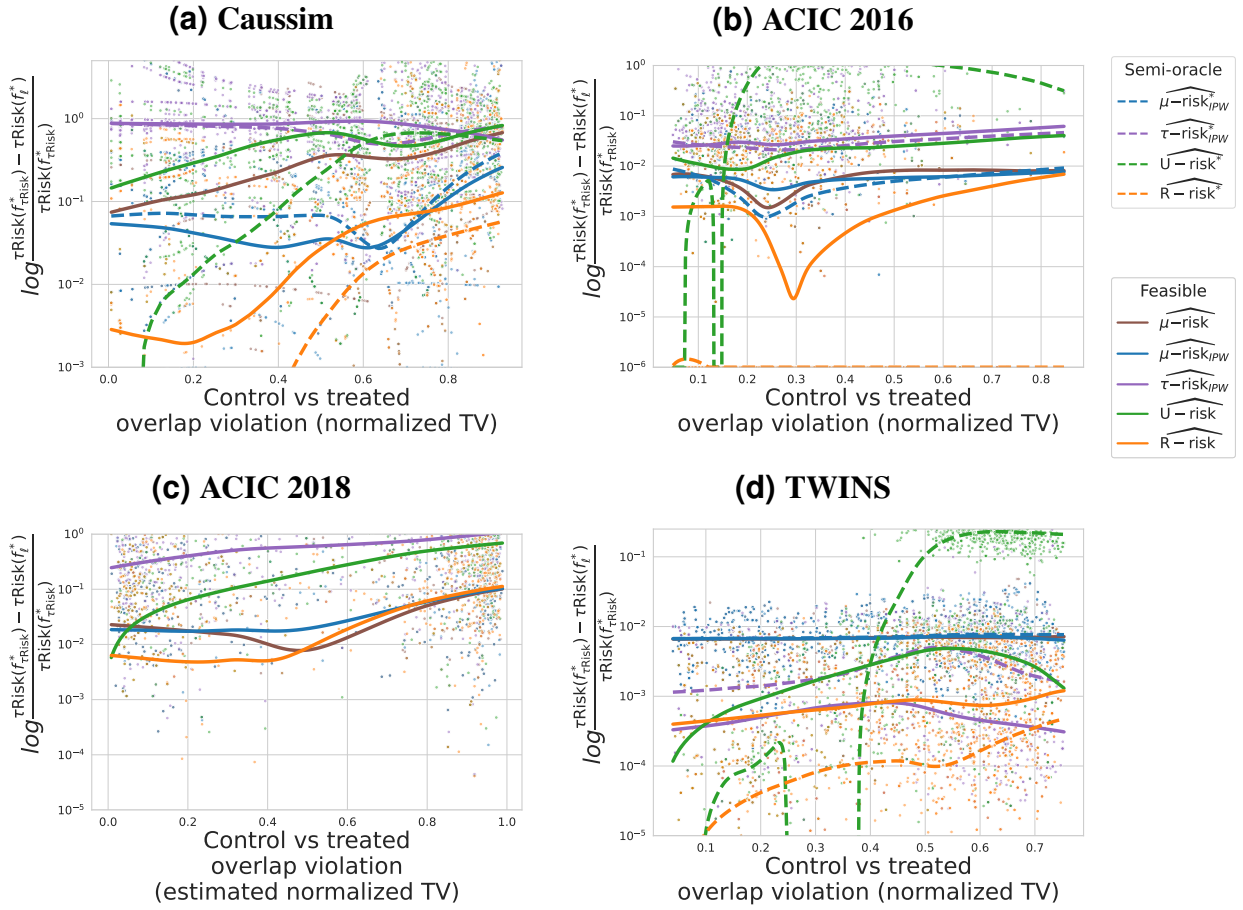
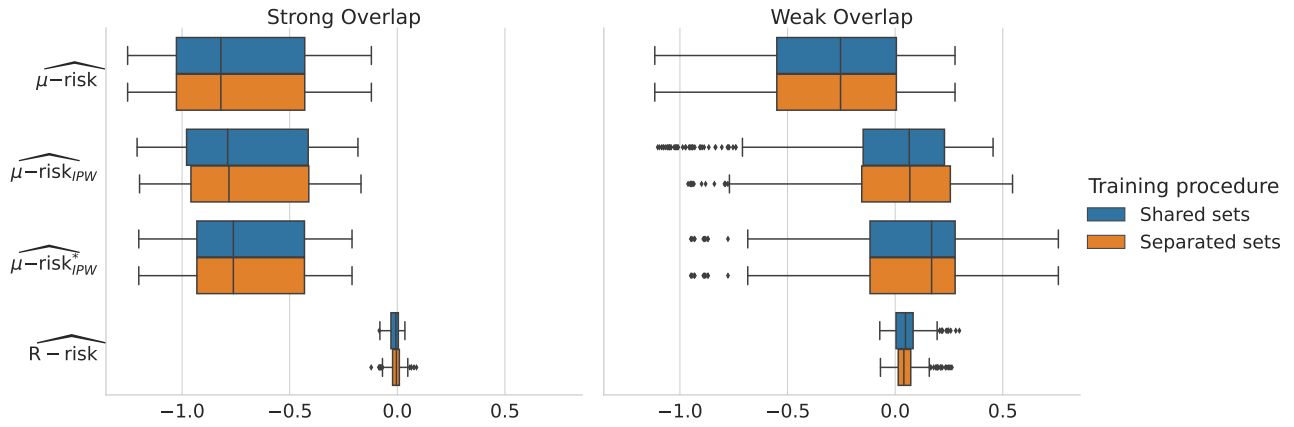
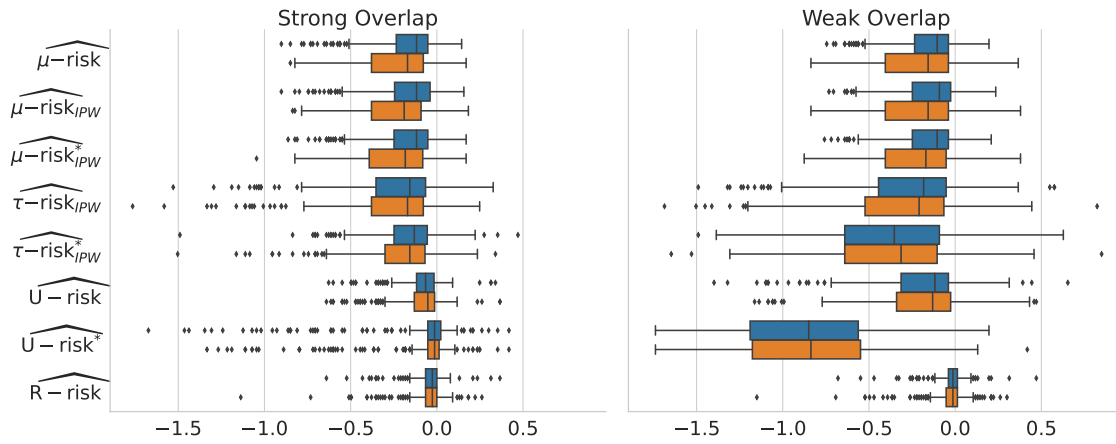


Figure 15: Metric performances by normalized tau-risk distance to the best method selected with τ -risk. All nuisances are learned with the same estimator stacking gradient boosting and ridge regression. Dotted and plain lines corresponds to 60% lowest quantile estimates. This choice of quantile allows to see better the oracle metrics lines for which outliers with a value of 0 distort the curves.

(a) Caussim



(b) ACIC 2016



(c) Twins

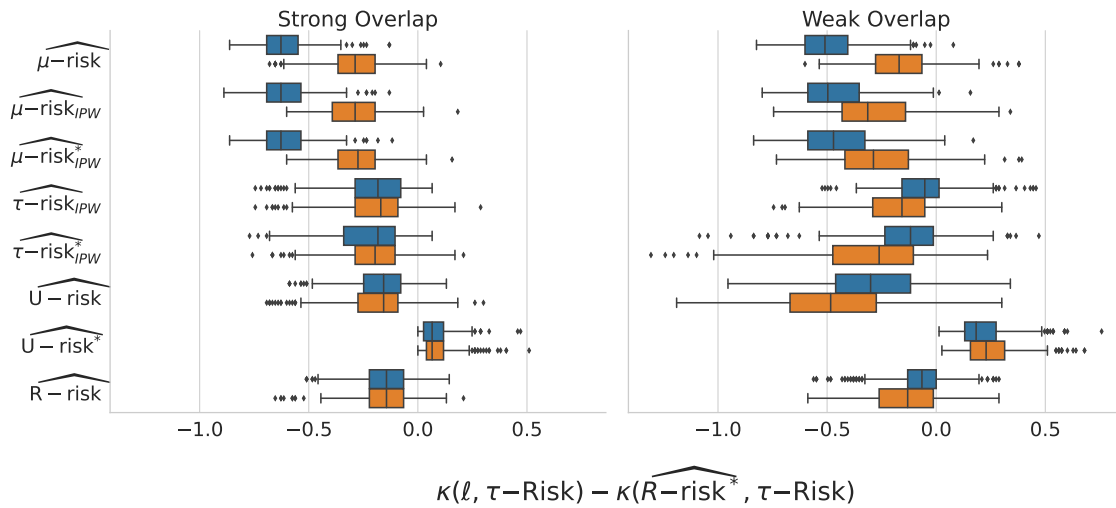


Figure 16: Results are similar between the Shared nuisances/candidate set and the Separated nuisances set procedure. The experience has not been run on the full metrics for Caussim due to computation costs.

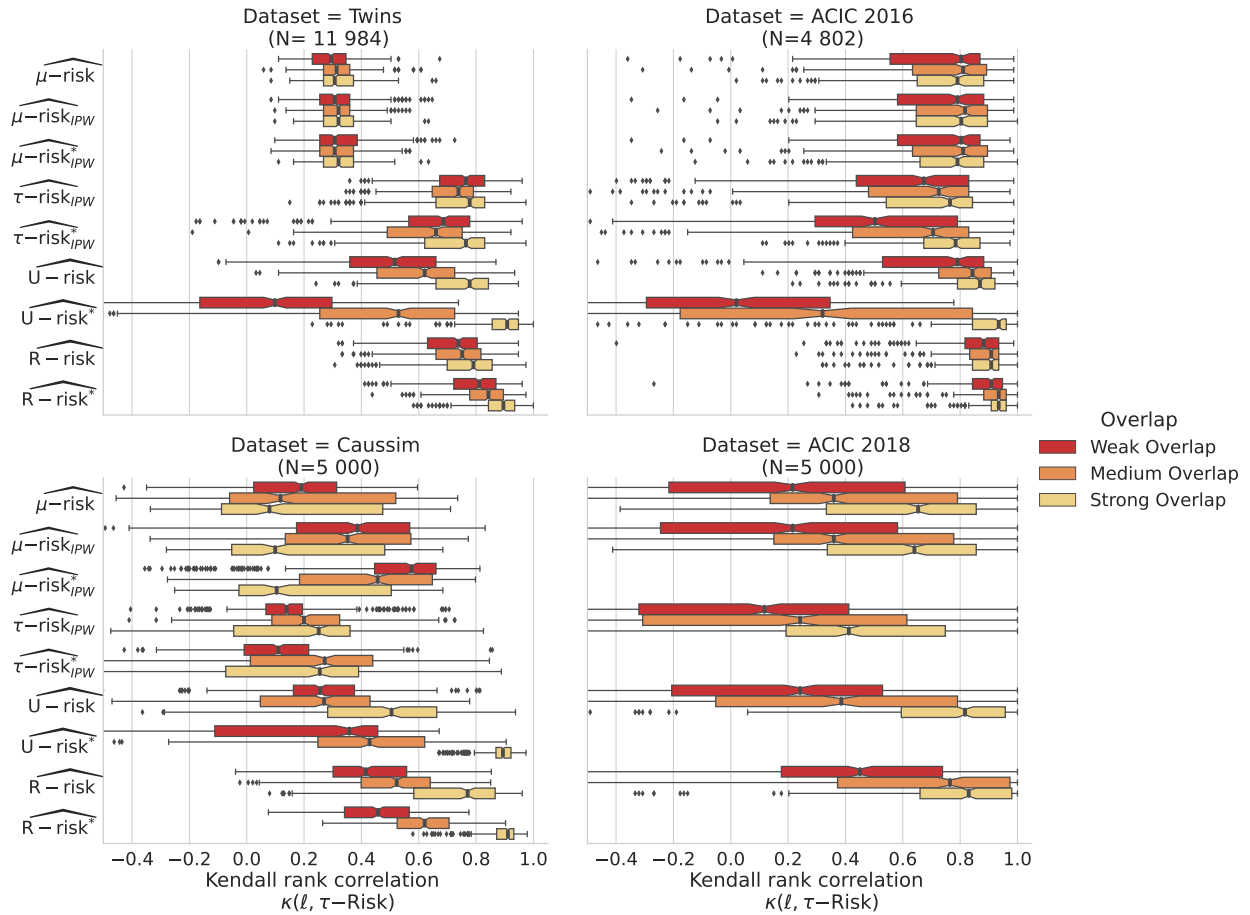


Figure 17: **Low population overlap hinders causal model selection for all metrics:** Kendall’s τ agreement with τ -risk. Strong, medium and Weak overlap correspond to the tertiles of the overlap distribution measured with Normalized Total Variation eq. 18.

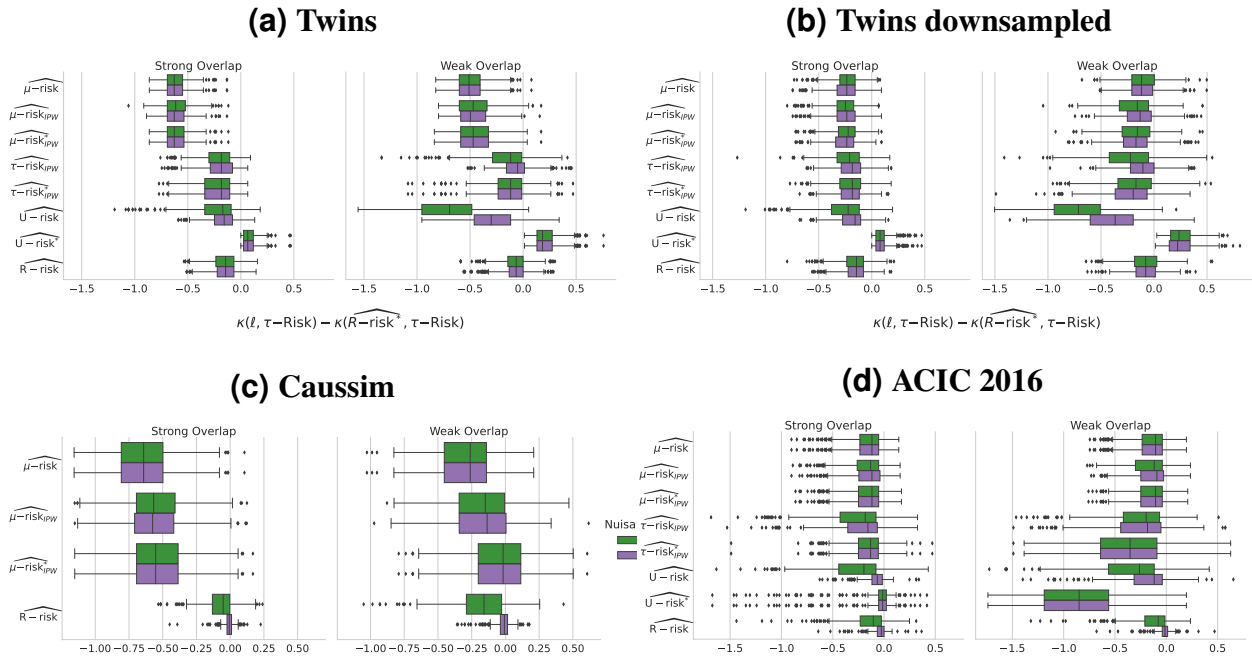


Figure 18: Learning the nuisances with **stacked models** (linear and gradient boosting) is important for successful model selection with R-risk. For Twins dataset, there is no improvement for **stacked models** compared to **linear models** because of the linearity of the propensity model.

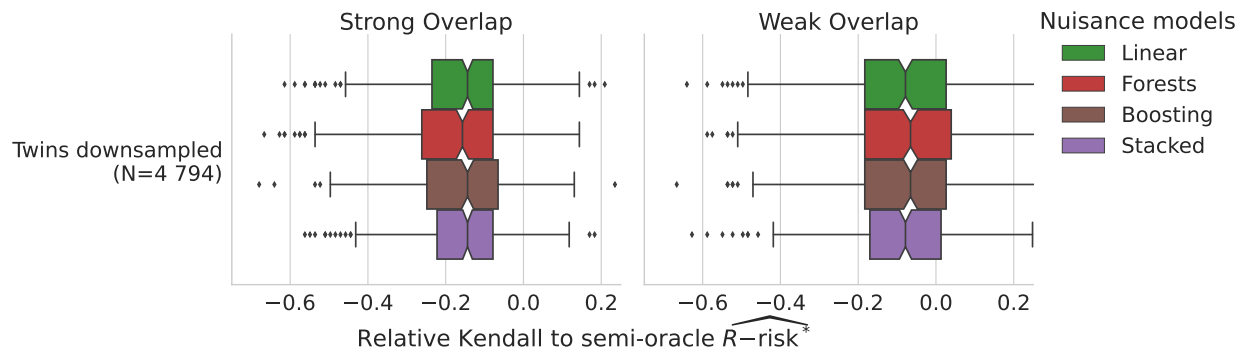
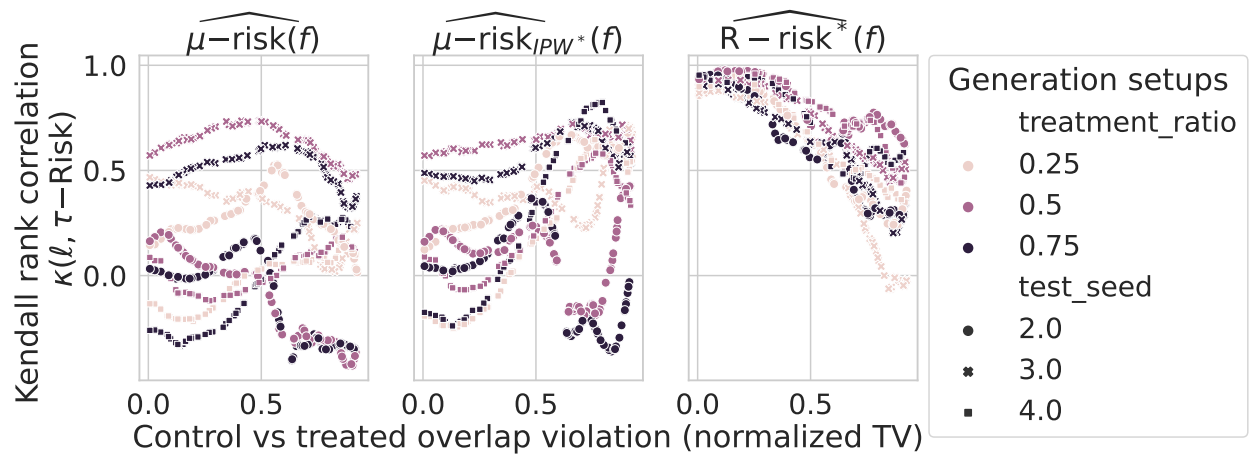


Figure 19: **Flexible models are performant in recovering nuisances in the downsampled Twins dataset.** The propensity score is linear in this setup, making it particularly challenging for flexible models compared to linear methods.

Figure 20: Kendall correlation coefficients for each causal metric. Each (color, shape) pair indicates a different (treatment ratio, seed) of the generation process.



F Heterogeneity in practices for data split

Splitting the data is common when using machine learning for causal inference, but practices vary widely in terms of the fraction of data to allocate to train models, outcomes and nuisances, and to evaluate them.

Before even model selection, data splitting is often required for estimation of the treatment effect, ATE or CATE, for instance to compute the nuisances required to optimize the outcome model (as the R -risk, definition 6). The most frequent choice is use 80% of the data to fit the models, and 20% to evaluate them. For instance, for CATE estimation, the R-learner has been introduced using K-folds with $K = 5$ and $K = 10$: 80% of the data (4 folds) to train the nuisances and the remaining fold to minimize the corresponding R-loss [54]. Yet, it has been implemented with $K=5$ in causallib [76] or $K=3$ in econML [9]. Likewise, for ATE estimation, Chernozhukov et al. [14] introduce doubly-robust machine learning, recommending $K=5$ based on an empirical comparison $K=2$. However, subsequent works use doubly robust ML with varying choices of K : Loiseau et al. [46] use $K=3$, Gao et al. [23] use $K=2$. In the econML implementation, K is set to 3 [9]. Naimi et al. [51] evaluate various machine-learning approaches –including R-learners– using $K=5$ and 10, drawing inspiration from the TMLE literature which sets $K=5$ in the TMLE package [26].

Causal model selection has been much less discussed. The only study that we are aware of, Schuler et al. [73], use a different data split: a 2-folds train/test procedure, training the nuisances on the first half of the data, and using the second half to estimate the R -risk and select the best treatment effect model.