

# Rethinking the Win Ratio: A Causal Framework for Hierarchical Outcome Analysis

Mathieu Even

Theremia and PreMeDICAL Inria-Inserm, University of Montpellier, France  
and

Julie Josse

PreMeDICAL Inria-Inserm, University of Montpellier, France

March 18, 2025

## Abstract

For *hierarchical* multivariate outcomes, the FDA recommends the Win Ratio and Generalized Pairwise Comparisons approaches Pocock et al. [2011], Buyse [2010]. However, as far as we know, these empirical methods lack causal or statistical foundations to justify their broader use in recent studies. To address this gap, we establish causal foundations for hierarchical comparison methods. We define related causal effect measures, and highlight that depending on the methodology used to compute Win Ratio, the causal estimand targeted can be different, as proved by our consistency results, which may then lead to reversed and incorrect treatment recommendations in heterogeneous populations, as we illustrate through striking examples.

In order to compensate for this fallacy, we introduce a novel, individual-level yet identifiable causal effect measure that better approximates the ideal, non-identifiable individual-level estimand. We prove that computing Win Ratio or Net Benefits using a Nearest Neighbor pairing approach between treated and controlled patients, an approach that can be seen as an extreme form of stratification, leads to estimating this new causal estimand measure. We extend our methods to observational settings via propensity weighting, distributional regression to address the curse of dimensionality, and a doubly robust framework. We prove the consistency of our methods, and the double robustness of our augmented estimator. These methods are straightforward to implement, making them accessible to practitioners. Finally, we validate our approach using synthetic data and on CRASH-3 [CRASH et al., 2019], a major clinical trial focused on assessing the effects of tranexamic acid in patients with traumatic brain injury.

*Keywords:* Hierarchical Outcomes Analysis, Multiple Outcomes, Randomized Control Trials, Observational data, Distributional Regression.

# 1 Introduction

Quantifying the benefit of a treatment in clinical research can be challenging, especially when outcomes are complex, multidimensional, or involve competing risks. Traditional statistical approaches often struggle to capture the nuanced relationships between such outcomes, limiting their ability to provide clinically meaningful insights. In these cases, innovative methods are required to address the inherent complexity of the data, to go beyond considering a single composite summary outcome. The Win Ratio [Pocock et al., 2011] and Generalized Pairwise Comparisons [Buyse, 2010] have emerged as powerful tools to evaluate treatment effects by comparing groups through hierarchical and multidimensional assessments of outcomes, to the point where they appear in the recent Food and Drugs Administration (FDA) guidances for handling multiple outcomes (see the FDA report *Multiple Endpoints in Clinical Trials Guidance for Industry, 2022*).

Clinical trials often involve competing risks and multiple outcomes, where prioritizing one type of event over another can lead to more clinically meaningful interpretations. For example, in cardiovascular (CV) studies, time to death may be prioritized over hospitalizations, reflecting the relative importance of these events to patients and clinicians. The Win Ratio and General Pairwise Comparisons provide a natural mechanism to account for this prioritization by structuring comparisons hierarchically, allowing for multidimensional comparisons and nuanced definitions of “wins” and “losses.” Pocock et al. [2011], Buyse [2010] form pairs of patients, each pair consisting of a patient in the control group and a treated patient. Each pair is then considered as a “Win” if the outcome of the treated patient is considered as more favorable than the control one, as a “Loss” if it is considered as less favorable, and as a “Null” if the two patient outcomes are comparable. We here illustrate the hierarchical comparison process, drawing inspiration from examples in cardiovascular trials [Pocock et al., 2011, Redfors et al.,

2020]. These studies evaluate multiple endpoints consisting of death, stroke, and heart failure hospitalizations (HFH), prioritizing events based on their clinical severity. In this hierarchy, death is considered the most severe outcome, followed by stroke, and finally HFH. For each patient pair, comparisons proceed as follows. (i) Determine which patient died first during their shared follow-up period. If one patient died earlier, the other patient is deemed the "winner" for this pair. (ii) If neither patient died, assess who experienced a stroke first. The patient with the later or no stroke is considered the "winner." (iii) If neither patient died nor had a stroke, compare the number of HFHs during follow-up. The patient with fewer HFHs is declared the "winner." This hierarchical process stops at the first event that distinguishes between a win or loss for the pair. If no event differentiates the pair, the outcome is recorded as a tie: depending on the methodology used, a tie can then be counted as a loss, as 1/2 instead of 1 or 0 (for respectively win or loss), or simply discarded. This approach ensures that clinically meaningful priorities are respected while maximizing the utility of the available data. For  $Y_i$  and  $Y_j$  the (multidimensional) outcomes of two patients  $i$  and  $j$ , respectively treated and controlled, we write

$$Y_i > Y_j,$$

for a win of  $i$  over  $j$  and  $Y_i \sim Y_j$  for a tie. The Win-Proportion is then defined as

$$\text{Win Proportion} \stackrel{\text{def}}{=} \frac{\#\text{Wins}}{\#\text{Pairs}},$$

the Win Ratio [Pocock et al., 2011] as:

$$\text{Win Ratio} \stackrel{\text{def}}{=} \frac{\#\text{Wins}}{\#\text{Losses}},$$

and the Net-Benefit of the treatment [Buyse, 2010] as:

$$\text{Net Benefit} \stackrel{\text{def}}{=} \frac{\#\text{Wins} - \#\text{Losses}}{\#\text{Pairs}}.$$

Treatment recommendations are then made using these computed values. If the win proportion is (significantly) above 0.5, treatment should be preferred over non-treatment, while for the Win Ratio and the Net Benefit the threshold values are respectively 1 and 0.

## 1.1 Contributions and outline of the paper

The way pairs of treated and control patients are formed is determined by the methodology chosen for the trial. *Complete pairings* are the historical prominent approach [Pocock et al., 2011, Buyse, 2010], and consist of choosing all possible pairs of treated and controlled individuals. Formally, if  $\mathcal{N}_1$  and  $\mathcal{N}_0$  are respectively the set of treated and controlled patients, the set of pairs used for complete pairings is  $\mathcal{C}_{\text{Tot}} = \mathcal{N}_1 \times \mathcal{N}_0$ . Dong et al. [2018] then introduced the *stratified* Win Ratio, an approach that consists in stratifying patients according to risks. In terms of pairings, the stratified approach translates into choosing pairs of treated and controlled patients that should have similar responses to treatment. Stratifying thus requires to evaluate risks, based on available covariates denoted as  $X_i$  for some patient  $i$ . In this paper, we introduce a third pairing approach: *Nearest Neighbor* pairings, that consist in choosing pairs  $\mathcal{C}_{\text{NN}} \subset \mathcal{N}_0 \times \mathcal{N}_1$  such that for all  $(i, j) \in \mathcal{C}_{\text{NN}}$ ,  $i$  is the treated patient with features  $X_i$  that are the closest features amongst all controlled patients to the features  $X_j$  of the controlled patient  $j$ . Nearest Neighbor pairings are in fact reminiscent of stratified pairings, since they can be seen as the extreme limit of stratification.

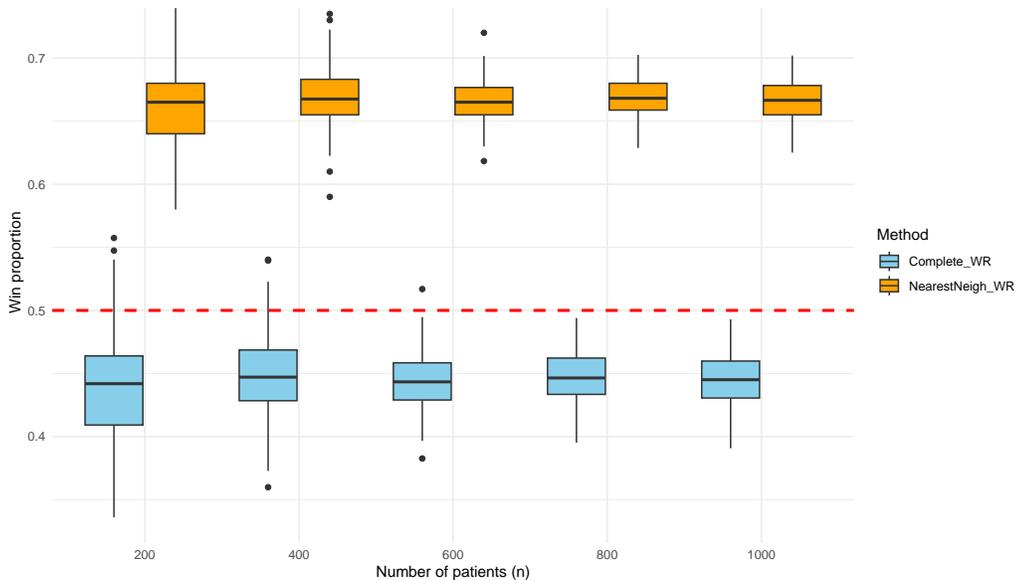


Figure 1: Comparison of the win proportion  $p_W$  computed with complete pairings and Nearest Neighbor pairings. Setting of Example 1. Boxplots over 100 runs. The two approaches lead to different treatment recommendations (above and below 0.5).

Our first contribution in this paper is to expose the fact that considering different pairings between treated and controlled patients (complete pairings or Nearest Neighbor pairings) may lead to different treatment recommendations. We illustrate this in Figure 1 on a simple synthetic example. This has major consequences: choosing different methodologies leads to different treatment recommendations. We argue that this is because these methods are typically limited to descriptive settings, relying on empirical procedures without a formal causal foundation. For instance, there is no well-defined estimand for what the Win Proportion, Win Ratio, and Net Benefit estimators seek at estimating. Despite their utility, integrating these methods into a cohesive causal inference framework remains underexplored. This paper formalizes Win Ratio and General Pairwise Comparisons approaches within a counterfactual paradigm [Splawa-Neyman et al., 1990], enabling a more robust interpretation of treatment effects in terms of potential outcomes. By doing so, we aim at addressing the limitations of existing methodologies

to bridge the gap between descriptive statistical techniques and causal reasoning, facilitating their application in complex clinical trials and broadening their use to observational studies.

In particular, studying hierarchical comparisons in a causal inference frameworks leads to introducing estimands for our estimators. We show that depending on the choice of pairings considered, the estimand approached may change and have dramatically different properties, thus explaining Figure 1. With complete pairings and in randomized controlled trials, the Win Proportion is a consistent estimator of the following causal estimand: <sup>1</sup>

$$\mathbb{P}(Y_i > Y_j) ,$$

where  $i$  and  $j$  are respectively two treated and controlled patients. This estimand is a *population-level* estimand. It as such does not capture the individual-level treatment effects: it is the probability that an individual sampled uniformly with treatment fares better than another individual sampled uniformly without treatment. An *individual-level* estimand would instead be the probability that a given individual fares better with than without treatment. Given  $Y_i(1)$  and  $Y_i(0)$  the potential outcomes of a given patient with and without treatment [Splawa-Neyman et al., 1990], this ideal individual-level estimand formally writes as:

$$\mathbb{P}(Y_i(1) > Y_i(0)) .$$

However, this causal estimand is *non-identifiable*. It depends on the joint distribution of the potential outcomes, which is *never* observed. These two estimands were previously mentioned in a series of works [Mao, 2017, Guo and Ni, 2022, Chen et al., 2024, Yin et al., 2022, Zhang et al., 2022, Chiaruttini et al., 2024].

Our next contribution is thus to introduce a new identifiable *and* individual-level estimand, that will serve as *proxy* for the non-identifiable individual-level causal estimand previously

---

<sup>1</sup>Note that our framework in the paper goes beyond estimating probabilities, and that for generality purposes we instead formalize this as  $\mathbb{E}[w(Y_i, Y_j)]$  for some *contrast* or *win* function  $w$ . The particular example presented in the introduction for readability purposes is  $w(y, y') = \mathbb{1}_{\{y > y'\}}$ .

defined. Our proxy is a better proxy than the population-level estimand, as it compares the outcome of an individual being treated, with the outcome of a controlled patient *that has the exact same features*, rather than any random controlled patient. Informally, our newly introduced estimand writes as

$$\mathbb{P}(Y_i > Y_{i'}) ,$$

where  $i$  is a treated patient, and  $i'$  an independent controlled patient that satisfies  $X_{i'} = X_i$  ( $i'$  may not exist, it is the result of a mind experiment). We show that in randomized controlled trials, the Win Proportion with Nearest Neighbors pairings is a consistent estimator for the causal estimand we introduce, thereby proving its identifiability. Using a simple example, we show that our newly introduced estimand should be preferred over the population-level estimand, making appear a paradox where changing the causal measure changes treatment effects. We provide discussions on the properties of our estimand in Section 2.3.

The last methodology part of the paper is devoted to estimating the estimand we introduce. In Section 4, we extend the Nearest Neighbors approach to the observational setting by incorporating propensity weights. In order to face the fact that Nearest Neighbors may be slow to converge in the presence of high dimensional input features and to face their lack of robustness in the presence of missing data in the covariates, we introduce a distributional regression approach in Section 4.2 to estimate our new causal estimand. The distributional regression estimators we propose benefit from very simple implementations, and as such we believe they can be easily used by practitioners. We propose a doubly-robust estimator, that combines the strengths of the distributional regression approach and of the weighted Nearest Neighbors approach.

We finally illustrate our methodologies on synthetically generated observational data and on a real world RCT. The CRASH-3 trial (Clinical Randomisation of an Antifibrinolytic in Significant Head Injury) was a large-scale, multi-center randomized controlled trial (RCT) that included over 12,000 patients in 175 hospitals across 29 countries. In particular, our Win Ratio

study of the CRASH-3 dataset in Section 6 [CRASH et al., 2019] compares the existing Win Ratio methodologies with our proposed methods. We offer guidances to compute Win Ratio and confidence intervals for the estimators, and our illustration on the CRASH-3 dataset shows that while traditional approaches fail to offer statistically significant results (1 is in the computed confidence intervals), the methods we propose in this paper offer significant results (in favor of placebo treatment here), due to the fact that an individual-level estimand is being targeted, rather than a population-level one.

## 1.2 Related works

**Hierarchical Outcome Analysis and practical advancements.** Pocock et al. [2011], Buyse [2010] introduced the Win Ratio and Generalized Pairwise Comparisons with the Net Benefit of a treatment. These methodologies are inspired by *Wilcoxon-Mann-Whitney tests* [Wilcoxon, 1945, Mann and Whitney, 1947], that consist in testing if a real-valued random variable  $Y$  is *stochastically* larger than another random variable  $Z$ . The quantity that arises in such tests is  $\mathbb{1}_{\{Y \geq Z\}}$  and its mean,  $\mathbb{P}(Y \geq Z)$ . Several studies have employed the Win Ratio or Generalized Pairwise Comparisons methodologies to analyze hierarchical outcomes. For instance, Pocock et al. [2023] used a stratified Win Ratio approach in the EMPULSE trial, which included 530 patients evenly split between treatment and placebo groups. Their analysis demonstrated how comparisons of all patient pairs contributed to "wins" for empagliflozin and placebo across four levels of the outcome hierarchy, resulting in an unstratified Win Ratio of 1.38, with accompanying confidence intervals and related metrics. They also discussed appropriate and inappropriate interpretations of this ratio. Similarly, Backer et al. [2024] applied hierarchical analysis to study reductions in treatment dosage and intensity for acute promyelocytic leukemia using Generalized Pairwise Comparisons, prioritizing efficacy outcomes such as event-free survival at two years over tolerability outcomes, including four pre-specified toxicities

common in this context. In another example, [Boentert et al. \[2024\]](#) applied the Win Ratio methodology to the COMET trial [[Diaz-Manera et al., 2021](#)], a Phase 3 study on late-onset Pompe disease. Their analysis showcased how the Win Ratio approach could be used to assess multiple endpoints in the orphan drug context, providing a more comprehensive evaluation of treatment benefits compared to previous analyses of the COMET trial.

[Redfors et al. \[2020\]](#) provide a comprehensive overview of the Win Ratio methodology, offering insights into its design and reporting for clinical studies. [Ajufu et al. \[2023\]](#) highlight several fallacies associated with the Win Ratio and Generalized Pairwise Comparisons. These include the observation that a "win" does not always equate to a clinically meaningful benefit, emphasizing the importance of accounting for "ties" or "nulls" rather than ignoring them. They also caution that using patient-reported outcomes in comparisons may introduce biases and that baseline risk stratification may not fully balance the risk profiles of paired subjects. Additionally, [Mao \[2024\]](#) identify further limitations of hierarchical comparison methods, such as challenges related to outcome censoring. A major issue with current Win Ratio and Generalized Pairwise Comparisons lies in the fact that these methodologies are only restricted to Randomized Controlled Trials (RCTs) and cannot be applied to observational studies, and another unaddressed limit is missing data in covariates. Finally, [Dong et al. \[2018\]](#) proposed a *stratified Win Ratio* approach, inspired by stratified odds ratio approaches [[Cochran, 1954](#), [Mantel and Haenszel, 1959](#)]. Our approach and results in this paper tend towards recommending the use of as much stratification as possible when dealing with Win Ratio and more generally comparison-based estimators.

**Formalization of hierarchical outcome analyses.** All previous cited works above provide guidances and recommendations based on empirical observations and trial examples. There has been advances towards formalizing estimands for hierarchical comparisons and adapting

these methods to observational data, based on Wilcoxon-Mann-Whitney testing [Wilcoxon, 1945, Mann and Whitney, 1947]. Mao [2017] formalized pairwise comparisons in a  $U$ -statistics setting, and provide (augmented) inverse propensity weighting for estimating a population-level estimand in observational studies. The estimand Mao [2017] introduced writes as the probability that a given randomly selected treated individual fares better than (*i.e.* wins against) another randomly selected controled individual, and is also studied by Chen et al. [2024], Yin et al. [2022] for rank-sum-tests or more generally for learning with contrast functions [Guo and Ni, 2022]. This is a population-level estimand in the sense that it compares the distribution of the outcome of two randomly selected patient. Several subsequent works then drew inspiration from Mao [2017] to apply this in cancer studies [Chiaruttini et al., 2024] or refine the proposed method in the presence of dependent subjects [Zhang et al., 2022] or clusters [Zhang and Jeong, 2021].

**Causal effect measures to assess treatment effects.** One of the key contribution of our paper is to design an adequate estimand for pairwise comparisons, that Win Ratio and Generalized Pairwise Comparisons methologies seek at approaching. The causal estimand we define is derived from a new causal measure. A causal measure is a functional of the joint distribution of the potential outcomes, and can vary from the Average Treatment Effect (ATE) with the Risk Difference (RD) to the Risk Ratio, the Odds Ratio, etc. Reporting results using different causal measures may lead to different conclusions or interpretations, and choosing the right is never straightforward [Colnet et al., 2024]. Previous estimands related to the Win Ratio consist of an *individual-level* estimand that cannot be estimated due to non-identifiability issues, and of a *population-level* estimand that is identifiable, but that compares the distribution of the treated population, with the distribution of the control population [Mao, 2017, Guo and Ni, 2022]. See Fay and Li [2024] for an extended discussion on individual and population-level

estimands and causal measures, and Gao et al. [2024] for approaches to handle non-identifiable causal estimands. The estimand we introduce lies in-between the two previously cited estimands: it is individual-level *and* identifiable. As opposed to some causal measures like the risk-ratio that are not directly collapsible [Fay and Li, 2024, Groenwold et al., 2011, Didelez and Stensrud, 2021] where taking the population-level estimand ( $\mathbb{E}[Y_i(1)]/\mathbb{E}[Y_i(0)]$  instead of  $\mathbb{E}[Y_i(1)/Y_i(0)]$  for the risk ratio) makes sense and can be interpreted in terms of treatment recommendations ( $\mathbb{E}[Y_i(1)]/\mathbb{E}[Y_i(0)] = 2$  means that treating everyone leads to an averaged outcome twice larger compared to treating no one), in comparison-based estimands this is not the case, as will be highlighted in our paper. As such, one must be very careful at what one is estimating, and defining the true quantity of interest becomes even more important.

Finally, another example of application of our framework beyond hierarchical outcomes analysis is *causal inference on distribution functions* [Lin et al., 2023]. If outcomes are general objects in a metric space (for instance, histograms), a contrast function  $w(y, y') = d(y, y')$  can be used. The metric can for instance be the Wasserstein metric if we work with continuous histogram), quantifying how much two outcomes  $y$  and  $y'$  are different. Here, the intuitive population-level causal measure would be to compute  $\bar{y}_1, \bar{y}_0$  the *Wasserstein barycenters* of outcomes in respectively treated and control groups. Our framework enables to go beyond such population-level estimands, while keeping identification possible.

## 2 Causal inference framework

### 2.1 Definitions and assumptions

We classically assume that we have access to  $n$  independent and identically distributed (*i.i.d.*) patients. Each patient  $i \in [n]$  is characterized by his *features* vector  $X_i$ , that lies in some feature space  $\mathcal{X}$ , a treatment assignment  $T_i \in \{0, 1\}$ , and an observed response (or outcome)

$Y_i \in \mathcal{Y}$ .  $\mathcal{Y}$  is the outcome set, and outcomes might be multivariate (for instance,  $\mathcal{Y} \subset \mathbb{R}^d$  with  $d \geq 2$ ).  $T_i = 0$  and  $T_i = 1$  respectively correspond to patient  $i$  being in the control group (non-treated) or in the test group (treated). Equivalently, the control and test groups can be replaced by two different treatment options. We use the *potential outcome* framework [Splawa-Neyman et al., 1990], that formalizes the concept of an intervention by positing the existence of two values  $Y_i(0)$  and  $Y_i(1)$  for the outcomes of interest, for the two situations where the patient has been exposed to treatment or not. These values are called *potential outcomes*, and they lie in some outcomes space  $\mathcal{Y}$ . The following assumption, often stated as the Stable Unit Treatment Values Assumption, is made throughout the paper, and states that the outcome is equal to the potential outcome given treatment.

**Assumption 1** (SUTVA). *We have  $Y_i = Y_i(T_i)$  for all  $i \in \{1, \dots, n\}$ .*

We will also make the two following assumptions, often referred to as *unconfoundedness* and *overlap/positivity*, respectively. Both these assumptions will be directly verified in the randomized controlled trial setting (RCT setting, for which  $T_i \perp\!\!\!\perp X_i$ ) that we will first consider in Section 3. The *observational* setting considered after in Sections 4 and 4.2 will rely on Assumptions 2 and 3.

**Assumption 2** (Unconfoundedness). *We have  $\{Y_i(0), Y_i(1)\} \perp\!\!\!\perp T_i | X_i$ .*

**Assumption 3** (Positivity). *There exists  $\eta \in (0, 1)$  such that for all  $x \in \mathcal{X}$ , we have*

$$\eta \leq \pi(x) \leq 1 - \eta,$$

where  $\pi(x) = \mathbb{P}(T_i = 1 | X_i = x)$  is the probability of being treated given  $X_i = x$ .

We want to know if a given patient would fare better under treatment than without it. In the potential outcomes framework, there exists several causal measures to quantify this, amongst

which the *Risk Difference* (RD) if  $\mathcal{Y} \subset \mathbb{R}$ , for which the Average Treatment Effect writes as:

$$\tau_{\text{RD}} \stackrel{\text{def}}{=} \mathbb{E} [Y_i(1) - Y_i(0)] .$$

However, we would like to handle more general outcomes, and Section 2.3 defines more general causal measures.

We introduce the “win” function on  $\mathcal{Y} \times \mathcal{Y}$ , that quantifies how a given outcome  $y \in \mathcal{Y}$  fares when compared to another outcome  $y' \in \mathcal{Y}$ . Our framework generalizes the *lexicographic order*<sup>2</sup> beyond the setting introduced by Pocock et al. [2011].

**Definition 1** (Win function). *Let*

$$w : (y, y') \in \mathcal{Y}^2 \mapsto w(y, y') \in [0, 1] ,$$

be the win function, taking two outputs  $y, y'$  to compare, and outputting a value between 0 and 1.

A typical example is:

$$w(y|y') = \begin{cases} 1 & \text{if } y > y' \\ \frac{1}{2} & \text{if } y \sim y' \\ 0 & \text{if } y < y' \end{cases} , \quad (1)$$

where  $>$  is an order on  $\mathcal{Y}$ , that we refer to as *clinical order*, and  $\sim$  means that the outcomes are similar or cannot be compared. If “ties” are discarded, the win function writes as:

$$w(y|y') = \begin{cases} 1 & \text{if } y > y' \\ 0 & \text{if } y \leq y' \end{cases} , \quad (2)$$

Pocock et al. [2011] first introduced the Win Ratio to handle composite endpoints in clinical trials based on clinical priorities, without formalizing it using the potential outcomes framework.

---

<sup>2</sup>The *lexicographic order* is a total order on  $\mathbb{R}^d$ , defined as  $y > y'$  if and only if  $\inf \{k \in [d], y_k > y'_k\} < \inf \{k \in [d], y_k < y'_k\}$ , with the convention that the infimum of an empty set is  $+\infty$ . In plain words, the *lexicographic order* amounts to order vectors as words in the dictionary.

In their example, outcomes  $Y_i$  are of dimension 2 and potential outcomes  $Y_1(t), Y_2(t) \in \mathbb{R} \cup \{\infty\}$  would respectively correspond to the time after treatment (or non treatment) before an eventual cardiovascular death event, and the time after (non-)treatment before an eventual heart-failure hospitalization.  $Y_i(t) = \infty$  means that no such event occurred. The win function here writes as:  $w(y|y') = \mathbb{1}_{\{y_1 > y'_1\}} + \mathbb{1}_{\{y_1 = y'_1, y_2 > y'_2\}}$ , *i.e.*  $w(y|y') = 1$  if and only if  $y$  is strictly larger than  $y'$  for the lexicographic order. Note that in this example, ties are counted as losses, which may not always be the case, as this may cause problems when ties should not be discarded nor treated as losses, as pointed out by [Ajufo et al. \[2023\]](#).

**Notations and terminology.** For a sequence  $(Z_n)_{n \geq 0}$  of random variables with values in a metric space  $(\mathcal{Z}, d)$  and  $Z$  a random variable in  $\mathcal{Z}$ , we say that  $Z_n$  converges in probability towards  $Z$  if for any  $\varepsilon > 0$  we have that  $\mathbb{P}(d(Z_n, Z) > \varepsilon) \rightarrow 0$ . We say that a measurable event  $\mathcal{E}$  is almost sure if its probability is 1. We say that  $Z_n$  converges almost surely towards some value  $\ell$  if the event  $\{Z_n \rightarrow \ell\}$  is almost sure.  $Z_n$  is a consistent estimator of some quantity  $\ell$  if  $Z_n$  converges in probability towards the constant random variable  $\ell$ .

## 2.2 Traditional Win Ratio, Net-Benefit and Mann-Whitney-Wilcoxon comparison estimators

Let  $\mathcal{N}_t = \{i : T_i = t\}$  for  $t = 0, 1$  be respectively the control and treated groups. Below, we formalize the estimators used by [Pocock et al. \[2011\]](#), [Buyse \[2010\]](#) for Win Ratio and Generalized Comparisons. We refer to these estimators as *traditional* or *historical* Win Ratio and Net-Benefit. [Pocock et al. \[2011\]](#), [Buyse \[2010\]](#) form pairs  $\mathcal{C} \subset \mathcal{N}_1 \times \mathcal{N}_0$  and define

$$n_W \stackrel{\text{def}}{=} \sum_{(y, y') \in \mathcal{C}} w(y|y'), \quad (3)$$

as the number of wins and  $n_L = |\mathcal{C}| - n_W$  as the number of losses. The *Win Proportion* is defined as:

$$\hat{p}_W \stackrel{\text{def}}{=} \frac{n_W}{n_W + n_L}, \quad (4)$$

and the *Win Ratio* [Pocock et al., 2011] is then defined as:

$$\hat{R}_{WR} \stackrel{\text{def}}{=} \frac{\hat{p}_W}{1 - \hat{p}_W} = \frac{n_W}{n_L}. \quad (5)$$

Another related quantity is the *Net Benefit* [Buyse, 2010]:

$$\hat{\Delta}_{NB} \stackrel{\text{def}}{=} 2\hat{p}_W - 1 = \frac{n_W - n_L}{n_W + n_L}. \quad (6)$$

Computing  $\hat{p}_W$  leads to values for both  $\hat{R}_{WR}$  and  $\hat{\Delta}_{NB}$ . Based on  $\hat{R}_{WR}$ ,  $\hat{\Delta}_{NB}$  or  $\hat{p}_W$ , the treatment  $T = 1$  can be judged favorable compared to non-treatment (or equivalently, to the treatment option corresponding to  $T = 0$ ) if

$$\hat{p}_W > \frac{1}{2},$$

which is equivalent to  $\hat{R}_{WR} > 1$  or  $\hat{\Delta}_{NB} > 0$ . Uncertainties and variance needs however to be taken into account to rule out a treatment option for another.

The choice of the set of pairs  $\mathcal{C} \subset \mathcal{N}_0 \times \mathcal{N}_1$  of control and test individuals used to compute the number of wins in Equation (3) has a crucial impact on the quantity being computed. The pair set  $\mathcal{C}$  might vary from the two natural following extremes:

1. *Complete pairings*, for which we have:

$$\mathcal{C}_{\text{Tot}} \stackrel{\text{def}}{=} \mathcal{N}_0 \times \mathcal{N}_1.$$

Complete pairings are the prevalent strategy in Win Ratio or Generalized Pairwise Comparisons analyses.

2. *Nearest-neighbor pairings*, for which

$$\mathcal{C}_{\text{NN}} \stackrel{\text{def}}{=} \{(i, \sigma^*(i)) \mid i \in \mathcal{N}_0\},$$

where  $\sigma^*(i) : \mathcal{N}_0 \rightarrow \mathcal{N}_1$  matches  $i \in \mathcal{N}_0$  to the treated individual  $j \in \mathcal{N}_1$  that has closest features in  $\mathcal{X}$  *i.e.*,  $\sigma^*(i) \in \operatorname{argmin}_{j \in \mathcal{N}_1} \|X_i - X_j\|^2$ .  $\mathcal{C}_{\text{NN}}$  can also be generalized to  $k$ -Nearest Neighbors.

**Handling mixed factor and numerical features.** For continuous features, Nearest Neighbors can be naturally applied after an eventual reweighting of the different features, to prevent from scale effects. For categorical features, Nearest Neighbors can be applied with one-hot encodings for instance, leading to Hamming-distances. In the presence of different categorical features of different importance, these can be weighted according to their importance. In the CRASH-3 dataset that we study in Section 6, we handle mixed categorical and numerical features in two different ways. The first one is to consider the Mahalanobis distance, that naturally balances variabilities and scale of the different features. The second approach, that we believe to be even more robust and significant, uses a Factor Analysis of Mixed Data (FAMD). Since nearest neighbors algorithm relies on distance metrics (like Euclidean or Mahalanobis distances), it struggles with mixed data types. FAMD transforms both numerical and categorical variables into a common latent space, ensuring a more meaningful distance calculation. Furthermore, nearest neighbors algorithm suffers from high-dimensional data because distances become less meaningful in higher dimensions. FAMD captures the most important variations in fewer dimensions, improving NN's effectiveness. This approach is detailed in Section 6 when studying the CRASH-3 data.

As highlighted in Example 1, using different pairings leads to different Win Proportions (and thus to different Win Ratios and Net Benefits), to the point where treatment recommendations may even differ.

**Example 1.** Suppose that we have  $n = 6$  individuals with univariate and real outcomes ( $\mathcal{Y} \subset \mathbb{R}$ ). Assume that for  $i = 1, 2, 3, 4$ , individuals are men (for which  $X_i = 0$ ) and we have  $Y_i(1) =$

$y_1 > y_0 = Y_i(0)$  while for  $i = 5, 6$  individuals are women (for which  $X_i = 1$ ) and we have  $Y_i(1) = y_1 < y_0 = Y_i(0)$ , and that for  $i \in \{1, \dots, 6\}$  we have  $T_i = 1$  if  $i$  is an odd number. Assume then that  $y'_0 > y_1 > y_0 > y'_1$ . Then, we have:

$$\hat{p}_W = \begin{cases} \frac{2}{3} & \text{if } \mathcal{C} = \mathcal{C}_{\text{NN}} \\ \frac{4}{9} & \text{if } \mathcal{C} = \mathcal{C}_{\text{Tot}} \end{cases},$$

leading to  $\hat{p}_W > 1/2$  or  $\hat{p}_W < 1/2$  and thus to different treatment decisions depending on the coupling pairs chosen. Here, treatment favors 4 out of the 6 patients, while no-treatment only favors 2 out of the 6: complete pairings thus favors the treatment option that benefits to only a minority of patients. This is further illustrated in Figure 1 in the setting of Example 1, with  $n/3$  women and  $2n/3$  men, in a RCT setting with treatment probability of  $1/2$ .

### 2.3 From estimators to causal measures and estimands

The quantities introduced so far —  $\hat{p}_W, \hat{R}_{\text{WR}}, \hat{\Delta}_{\text{NB}}$  —, are data-dependent estimators (hence the  $\hat{\cdot}$  notation). To efficiently capture treatment effects and treatment comparisons, we need to first answer the following crucial question: *what is the estimand that these estimators seek at estimating?* In this subsection, we review the two causal estimands previously defined and studied for generalized comparisons and Win Ratio analyses, before introducing our new estimand. The first estimand is a natural individual-level estimand, but cannot be estimated as it is not identifiable. Most works therefore introduced a population-level estimand, as a proxy for the non-identifiable individual-level causal estimand. This causal estimand may however not capture treatment effects correctly, hence our new causal estimand that is both individual-level and identifiable.

**A natural but non-identifiable causal estimand.** In our causal inference framework, we wish to determine if individuals would fare better if treated or non-treated. Here,  $w$  is a contrast

function that compares two outcomes  $y, y' \in \mathcal{Y}$ :  $w(y|y')$  quantifies the relative favorability of  $y$  compared to  $y'$ . If we worked with the Risk Difference, we would have  $w(y|y') = y - y'$ , and the ATE would write  $\tau_{\text{ATE}} = \mathbb{E}[w(Y_i(1)|Y_i(0))] = \mathbb{E}[Y_i(1) - Y_i(0)]$ . In the general case, with general contrast functions (our “win” function), this leads to consider the following quantity:

$$\tau_{\text{indiv}} \stackrel{\text{def}}{=} \mathbb{E}[w(Y_i(1)|Y_i(0))] . \quad (7)$$

The quantity  $\tau_{\text{indiv}}$  is a *causal effect measure* (or causal measure, for short) and a *causal estimand*, that compares the two potential outcomes of a given individual using the contrasts/win function  $w$ . If  $w$  is as in Equation (1), we have:

$$\tau_{\text{indiv}} = \mathbb{P}(Y_i(1) > Y_i(0)) + \frac{1}{2}\mathbb{P}(Y_i(1) \sim Y_i(0)) ,$$

while if  $w$  is as in Equation (2), we have:

$$\tau_{\text{indiv}} = \mathbb{P}(Y_i(1) > Y_i(0)) .$$

However, as highlighted by several works [Mao, 2017, Guo and Ni, 2022, Chen et al., 2024, Yin et al., 2022, Zhang et al., 2022, Chiaruttini et al., 2024], this causal measure is *non-identifiable* since estimating it in general requires the knowledge of the *joint distribution* of the potential outcomes, which is never observed <sup>3</sup>.  $\tau_{\text{indiv}}$  is indeed an *individual-level* causal estimand [Fay and Li, 2024], as it is directly a function of the joint probability distribution  $\mathcal{P}(\{Y_i(0), Y_i(1)\})$ , as opposed to *population-level* causal estimands that are functions of  $\{\mathcal{P}(Y_i(0)), \mathcal{P}(Y_i(1))\}$ . What makes it possible to estimate  $\mathbb{E}[w(Y_i(1)|Y_i(0))]$  with the Risk Difference is the fact that thanks

---

<sup>3</sup>If  $Y_i(0), Y_i(1) \sim \text{Bernoulli}(1/2)$ , the ATE with the risk difference  $\mathbb{E}[Y_i(1) - Y_i(0)]$  is identifiable (via e.g. taking the mean on test and control groups, in a RCT setting) and equal to 0, while the ATE with  $w(y|y') = \mathbb{1}_{\{Y_i(1) > Y_i(0)\}}$  writes as  $\mathbb{P}(Y_i(1) > Y_i(0))$  is not identifiable. Indeed, in that latter case, coupling  $(Y_i(1), Y_i(0))$  as  $Y_i(1) = Y_i(0)$  gives  $\mathbb{P}(Y_i(1) > Y_i(0)) = 0$ , while taking independent potential outcomes leads to  $\mathbb{P}(Y_i(1) > Y_i(0)) = \frac{1}{4}$ . Since the distribution of the observations  $(X_i, T_i, Y_i)$  does not change by taking either coupling but the value of  $\mathbb{E}[w(Y_i(1)|Y_i(0))]$  changes, we can say that we have non-identifiability.

to the linearity of the contrast function  $w$ ,  $\mathbb{E}[w(Y_i(1)|Y_i(0))]$  is both a population and an individual-level causal measure.

**A population-level causal estimand.** To circumvent this non-identifiability issue of individual-level causal measures, Mao [2017], Guo and Ni [2022], Chen et al. [2024], Yin et al. [2022], Zhang et al. [2022], Chiaruttini et al. [2024] consider a population-level causal measure  $\tau_{\text{pop}}$  instead of the individual-level one that defines  $\tau_{\text{indiv}}$ . Their causal measure writes as:

$$\tau_{\text{pop}} \stackrel{\text{def}}{=} \mathbb{E}[w(Y_i(1)|Y_j(0))] , \quad (8)$$

where  $i$  and  $j$  are two different and independent individuals. Considering two independent individuals leads to a different measure, that can now be estimated. We will show in Theorem 1 that with total pairings,  $\hat{p}_W$  in Equation (9) is a consistent estimator of  $\tau_{\text{pop}}$ .

**Our individual-level and identifiable causal estimand.** Using  $\tau_{\text{pop}}$  leads to comparing individuals that may not be comparable, hence the following causal measure we introduce. It is an identifiable relaxation of  $\tau_{\text{indiv}}$ , defined by comparing patient  $i$  with features  $X_i$  to an independent copy *that has the same features*.  $\tau_\star$  is an individual-level causal estimand.

**Definition 2.** For  $x \in \mathcal{X}$ , let  $\{Y^{(x)}(0), Y^{(x)}(1)\}$  be an independent copy of  $\{Y_i(0), Y_i(1)\}|X_i = x$ .

Let

$$\tau_\star(x) \stackrel{\text{def}}{=} \mathbb{E}[w(Y^{(x)}(1)|Y_i(0))|X_i = x] ,$$

and define

$$\tau_\star \stackrel{\text{def}}{=} \mathbb{E}[w(Y^{(X_i)}(1)|Y_i(0))] . \quad (9)$$

The quantities defined  $\tau_\star(x)$  and  $\tau_\star$  are respectively a conditional effect measure and a causal effect measure [Pearl, 2009]. By construction, the causal measure built satisfies *direct collapsibility* [Colnet et al., 2024, Definition 4] and is *logic-respecting* [Colnet et al., 2024, Definition 6].

We will show in Theorem 1 that the causal estimand  $\tau_\star$  is identifiable as opposed to  $\tau_{\text{indiv}}$ , since for  $\mathcal{C} = \mathcal{C}_{\text{NN}}$  the estimator  $\hat{p}_W$  (Equation (4)) is consistent for  $\tau_\star$ .

**Remark 1.** Under additional assumptions such as *potential independence* (i.e.,  $Y_i(1) \perp\!\!\!\perp Y_i(0) \mid X_i$ ), we have that  $\tau_\star = \tau_{\text{indiv}}$ . This assumption of potential independence is however quite strong and may be considered unlikely to hold true. It means that all that accounts for treatment effects is included in  $X_i$ , precluding e.g. any unmeasured factor. Some examples (such as for instance comparing 2 different doses of a same treatment) make it impossible to assume conditional independence in general, hence the appeal of  $\tau_\star$ , since Definition 2 does not need to make any assumption for  $\tau_\star$  to be well-defined.

We now define the *statistical estimands* related to our causal measures and to our causal estimands. Recall that statistical estimands are functions of measurable quantities; as such, they cannot make appear counterfactual quantities such as potential outcomes, which is not the case for causal measures and causal estimands. For instance, the statistical estimand related to the causal estimand  $\tau_{\text{RD}}$  is  $\mathbb{E}[Y_i | T_i = 1] - \mathbb{E}[Y_i | T_i = 0]$ . For the population-level causal measure  $\tau_{\text{pop}}$ , the related statistical estimand writes as:

$$\mathbb{E}[w(Y_i | Y_j) | T_i = 1, T_j = 0] ,$$

while the statistical estimand related to  $\tau_\star$  writes as

$$\mathbb{E}[w(Y^{(X_i)}(1), Y_i) | T_i = 0] ,$$

or, equivalently:

$$\mathbb{E}[w(Y_i, Y^{(X_i)}(0)) | T_i = 1] .$$

Under Assumption 1 (SUTVA), we have that these statistical estimands are equal to their associated causal estimands.

From the estimands  $\tau = \tau_\star$  or  $\tau = \tau_{\text{pop}}$ , the Win Ratio estimand related to the estimator  $\hat{R}_{\text{WR}}$  defined in Equation (5) writes as:

$$R_{\text{WR}} \stackrel{\text{def}}{=} \frac{\tau}{1 - \tau},$$

while the Net Benefit estimand related to the estimator  $\hat{\Delta}_{\text{NB}}$  defined in Equation (6) writes as:

$$\Delta_{\text{NB}} \stackrel{\text{def}}{=} 2\tau - 1.$$

**Remark 2.** Having estimators and confidence intervals for estimating  $\tau$  with some estimator  $\hat{\tau}$  is equivalent to having estimates and confidence intervals for  $R_{\text{WR}}$  or  $N_{\text{NB}}$  for estimators  $\hat{\Delta} = 2\hat{\tau} - 1$  and  $\hat{R} = \frac{\hat{\tau}}{1 - \hat{\tau}}$ . Indeed, if we have a confidence interval of the form  $\mathbb{P}(\hat{\tau} \in [\tau - \varepsilon_1, \tau + \varepsilon_2]) \geq 1 - \alpha$ , we also have  $\mathbb{P}(\hat{\Delta} \in [\Delta_{\text{NB}} - 2\varepsilon_1, \Delta_{\text{NB}} + 2\varepsilon_2]) \geq 1 - \alpha$  and  $\mathbb{P}(\hat{R} \in [R_{\text{WR}} - \varepsilon'_1, R_{\text{WR}} + \varepsilon'_2]) \geq 1 - \alpha$ , where  $\varepsilon'_1 = \frac{\tau}{1 - \tau} - \frac{\tau - \varepsilon_1}{1 + \varepsilon_1 - \tau}$  and  $\varepsilon'_2 = \frac{\tau + \varepsilon_2}{1 - \varepsilon_2 - \tau} - \frac{\tau}{1 - \tau}$  if  $0 < \tau - \varepsilon_1$  and  $\tau + \varepsilon_2 < 1$ .

**Remark 3** (On the well-posedness of Definition 2). Let  $(E, \mathcal{E})$  and  $(F, \mathcal{F})$  be two probability spaces.  $\nu : E \times \mathcal{F} \rightarrow [0, 1]$  is a transition kernel if it satisfies  $\forall x \in E$ ,  $\nu(x, \cdot)$  is a probability measure on  $(F, \mathcal{F})$  and  $\forall B \in \mathcal{F}$ ,  $\nu(\cdot, B)$  is  $\mathcal{E}$ -measurable. If  $X \in \mathbb{R}^p$  and  $Y \in \mathbb{R}^d$ , the conditional law of  $Y$  given  $X$  is a kernel  $\nu$  on  $(\mathbb{R}^d \times \mathcal{B}(\mathbb{R}^p))$  that satisfies:

$$\mathbb{P}_{(X,Y)} = \mathbb{P}_{X,\nu}, \quad \text{i.e.} \quad \forall (A, B) \in \mathcal{B}(\mathbb{R}^p) \times \mathcal{B}(\mathbb{R}^d), \quad \mathbb{P}(X \in A, Y \in B) = \int_{x \in A} \nu(x, B) \mathbb{P}_X(dx).$$

Such a transition kernel always exists: this is Miloslav Jiřina's Theorem [Jiřina, 1959] for Borelian measures and random variables. We write

$$\mathbb{P}(Y \in B | X = x) \stackrel{\text{def}}{=} \nu(x, B).$$

To build two independent copies of  $Y$  given  $X = x$ , we thus draw  $Y_x, Y'_x$  with:

$$\forall (B, B') \in \mathcal{B}(\mathbb{R}^d)^2, \quad \mathbb{P}(Y_x \in B, Y'_x \in B') = \nu(x, B)\nu(x, B').$$

We thus have created a copy of  $Y$  that satisfies the two following properties: (i) it is independent of  $Y$  conditionally on  $X$ ; (ii) conditionally on  $X$ , the distributions of  $Y'$  and  $Y$  are the same.

Their joint distribution with  $X$  writes as:

$$\mathbb{P}(X \in A, Y \in B, Y^{(X)} \in B') = \mathbb{P}(X \in A) \mathbb{E}[\nu(X, B)\nu(X, B')|X \in A],$$

for all  $(A, B, B') \in \mathcal{B}(\mathbb{R}^p) \times \mathcal{B}(\mathbb{R}^d)^2$ .

Example 2 below is the estimand-wise version of Example 1. It shows that the causal measures  $\tau_\star$  and  $\tau_{\text{pop}}$  are not equivalent and may lead to different treatment recommendations if populations are heterogeneous. In particular, Example 2 justifies our preference and recommendations towards using  $\tau_\star$  over  $\tau_{\text{pop}}$ .

**Example 2.** Assume that outcomes are univariate  $(Y_i(0), Y_i(1) \in \mathbb{R})$ , and that  $\mathcal{X} = \mathcal{X}_1 \cup \mathcal{X}_2$  with  $\mathbb{P}(X_i \in \mathcal{X}_1) = 1 - \alpha$ ,  $\mathbb{P}(X_i \in \mathcal{X}_2) = \alpha$  such that:

$$Y_i(1) = y_1 > y_0 = Y_i(0)|X_i \in \mathcal{X}_1, \quad Y_i(1) = y'_1 < y'_0 = Y_i(0)|X_i \in \mathcal{X}_2,$$

almost surely. We then have, if  $y'_0 > y_1 > y_0 > y'_1$ :

$$\mathbb{E}[w(Y^{(X_i)}(1)|Y_i(0))] = 1 - \alpha, \quad \mathbb{E}[w(Y_i(1)|Y_j(0))] = (1 - \alpha)^2.$$

Thus, if  $\frac{1}{2} > \alpha > 1 - \frac{1}{\sqrt{2}}$ , we have:

$$\mathbb{E}[w(Y^{(X_i)}(1)|Y_i(0))] > \frac{1}{2} > \mathbb{E}[w(Y_i(1)|Y_j(0))],$$

leading to different conclusions in terms of treatment efficacy (see Figure 1).

We highlight the fact that the phenomenon appearing in Example 2 is *not* reminiscent of Simpson's paradox [Simpson, 1951, Wagner, 1982], as understood in the popular sense. Simpson's paradox states a trend might appear in *all* subgroups of a population, but still reverse when considering the average over all population. Here, the paradox in Example 2 is of a very

different nature, since in both cases the average over the whole population is considered. The difference lies in the way the average is taken: different causal measures might lead to different treatment effects. Taking the average of individual effects over the global population (as done when considering  $\tau_\star$ ) or comparing the whole treated group distribution with the control distribution (as done when considering  $\tau_{\text{pop}}$ ) can lead to opposite trends.

Similarly to  $\tau_{\text{pop}}$ ,  $\tau_\star$  it serves as a computable proxy to approximate the ideal value  $\tau_{\text{indiv}}$  that cannot be approximated in general. We however argue that, thanks to simple examples such as the one just above,  $\tau_\star$  is a *better* proxy than  $\tau_{\text{pop}}$ , since it captures more information. This is intuitively the case: both  $(Y^{(X_i)}(1), Y_i(0))$  and  $(Y_j(1), Y_i(0))$  are couplings of the random variables  $Y_i(1)$  and  $Y_i(0)$ . The coupling  $(Y^{(X_i)}(1), Y_i(0))$  is however naturally closer to the coupling  $(Y_i(1), Y_i(0))$  than the coupling  $(Y_j(1), Y_i(0))$ , since  $(Y^{(X_i)}(1), Y_i(0))$  takes into account covariate effects, leading to:

$$d_{\ell^2}((Y^{(X_i)}(1), Y_i(0)), (Y_i(1), Y_i(0))) \leq d_{\ell^2}((Y_j(1), Y_i(0)), (Y_i(1), Y_i(0))),$$

if the marginals are absolutely continuous with respect to the Lebesgue measure, and where  $d_{\ell^2}$  is the  $\ell^2$  distance between densities. Finally, next proposition formalizes the excess risk when using  $\tau_\star$  or  $\tau_{\text{pop}}$  as proxis for  $\tau_{\text{indiv}}$ .

**Proposition 1.** *We have:*

$$|\tau_\star - \tau_{\text{indiv}}| \leq d_{\text{TV}}(P_{Y^{(X_i)}(1), Y_i(0)}, P_{Y_i(1), Y_i(0)}),$$

and

$$|\tau_{\text{pop}} - \tau_{\text{indiv}}| \leq d_{\text{TV}}(P_{Y_j(1), Y_i(0)}, P_{Y_i(1), Y_i(0)}),$$

where  $P_{Y^{(X_i)}(1), Y_i(0)}$ ,  $P_{Y_i(1), Y_i(0)}$ ,  $P_{Y_j(1), Y_i(0)}$  are respectively the joint distributions of  $(Y^{(X_i)}(1), Y_i(0))$ ,  $(Y_i(1), Y_i(0))$  and  $(Y_j(1), Y_i(0))$ , and  $d_{\text{TV}}$  is the total-variation distance between distributions.

Furthermore, if the win function  $w$  is 1-Lipschitz, we have that:

$$|\tau_\star - \tau_{\text{indiv}}| \leq \mathcal{W}_1(P_{Y^{(X_i)}(1), Y_i(0)}, P_{Y_i(1), Y_i(0)}),$$

and

$$|\tau_{\text{pop}} - \tau_{\text{indiv}}| \leq \mathcal{W}_1(P_{Y_j(1), Y_i(0)}, P_{Y_i(1), Y_i(0)}),$$

where  $\mathcal{W}_1$  is the 1–Wasserstein distance between distributions.

### 3 Consistency of traditional Win Ratio, Net-Benefit and Win Proportions in the RCT setting

In this section, we study  $\hat{p}_W$  in light of the causal measures we previously defined. The quantity  $\hat{p}_W$  is the Win Proportion (Equation (4)), that is used to compute  $\hat{R}_{\text{WR}}$  and  $\hat{\Delta}_{\text{NB}}$ , referred to as *traditional Win Ratio and Net-Benefit estimators*. As expected from Examples 1 and 2, the behavior of  $\hat{p}_W$  crucially depends on the pairings considered. The following theorem shows that, in a Randomized Controlled Trial (RCT) setting, the Win Proportion is the most natural estimator for the estimands  $\tau = \tau_\star$  and  $\tau = \tau_{\text{pop}}$ . Indeed, for a Nearest Neighbor pairing choice, the Win Proportion is a consistent estimator of  $\tau_\star$ , while for a complete pairing choice the Win Proportion is a consistent estimator of  $\tau_{\text{pop}}$ .

**Theorem 1** (Consistency of Win Ratio). *Let*

$$\hat{p}_W^{(n_0, n_1, \mathcal{C})} \stackrel{\text{def}}{=} \frac{1}{|\mathcal{C}|} \sum_{(i, j) \in \mathcal{C}} w(Y_j | Y_i),$$

for  $\mathcal{C} \subset \mathcal{N}_0 \times \mathcal{N}_1$ . Assume that Assumptions 1 to 3 hold, and assume further that we are in the RCT setting:  $T_i \perp\!\!\!\perp X_i$ .

1. Complete pairing. We have:

$$\hat{p}_W^{(n_0, n_1, \mathcal{C}_{\text{Tot}})} \xrightarrow{\mathbb{P}} \tau_{\text{pop}} \stackrel{\text{def}}{=} \mathbb{E}[w(Y_i(1) | Y_j(0))], \quad i \neq j,$$

where the limit in probability is taken as  $n_0, n_1 \rightarrow \infty$  and  $\mathcal{C}_{\text{Tot}} = \mathcal{N}_0 \times \mathcal{N}_1$ .

2. Nearest Neighbor. Assume that  $(x, y) \mapsto \mathbb{E}[w(Y_i(1)|y)|X_i = x]$  is continuous in its first variable  $x$  and that  $\mathcal{X}$  is compact. Then, let  $\sigma^* : \mathcal{N}_0 \rightarrow \mathcal{N}_1$  be defined as:

$$\forall i \in \mathcal{N}_1, \quad \sigma^*(i) \in \operatorname{argmin}_{j \in \mathcal{N}_1} \|X_i - X_j\|,$$

where if there are two possible choices or more for  $\sigma^*(i)$ , we choose uniformly at random.

Let  $\mathcal{C}_{\text{NN}} = \{(i, \sigma^*(i)), i \in \mathcal{N}_0\}$ . We have:

$$\hat{p}_W^{(n_0, n_1, \mathcal{C}_{\text{NN}})} \xrightarrow[\mathbb{P}]{} \tau_\star \stackrel{\text{def}}{=} \mathbb{E} \left[ \mathbb{E} \left[ w(Y^{(X_i)}(1) | Y_i(0)) | X_i \right] \right],$$

where the limit in probability is taken as  $n_0, n_1 \rightarrow \infty$ .

The consistency results of Theorem 1 formalize the intuition brought by Examples 1 and 2 and illustrated in Figure 1. The choice of pairing  $\mathcal{C}$  is crucial, and leads to the estimation of very different quantities. For complete pairings  $\mathcal{C}_{\text{Tot}}$ , the causal estimand that is estimated is the population-level one,  $\tau_{\text{pop}}$ . For Nearest Neighbor pairings  $\mathcal{C}_{\text{NN}}$ , the causal estimand that is estimated is the individual-level one that we introduced,  $\tau_\star$ , defined in Definition 2. As highlighted in Examples 1 and 2 and in Figure 1, complete pairings and the population-level causal estimand have undesired behaviors and may even lead to different treatment recommendations. When the features  $X_i$  are expressive enough, this fallacy of  $\tau_{\text{pop}}$  and of complete pairings is solved by considering our causal measure  $\tau_\star$  and Nearest Neighbor pairings instead.

However, it is worth mentioning that, while  $\tau_\star$  seems like the causal estimand one would seek at approaching, using Nearest Neighbor pairings to approach  $\tau_\star$  has some downsides too. The first weakness of this approach is the fact that it is restricted to the *RCT* setting: the consistency result above in Theorem 1.2 crucially relies on  $T_i$  being independent of  $X_i$ . A first direction will thus be to relax the RCT setting to an observational one: this is what we do in Section 4, where we combine an Inverse Propensity Weighting approach with Nearest Neighbor pairings. Then, the second weakness of this Nearest Neighbor approach is that it

suffers from the curse of high dimensions: if the convergence speed in terms of samples  $n$  required will drop exponentially as the dimension of the features  $X_i$  increases. This is a known weakness of Nearest Neighbors approach in general [Biau and Devroye, 2015]. Although they are one the most studied non-parametric learning methods for which theoretical results can be derived, Nearest Neighbors suffer from the fact that in high dimensions, points are likely to be separated if  $n$  is not exponentially large in the dimension. This motivates our aim for a more general and systemic method to estimate our causal estimand  $\tau_\star$ : in Section 4.2, we provide a distributional regression point of view on the problem and develop new methodologies for efficiently estimating  $\tau_\star$ . This new approach also has the benefit of being able to handle missing values in the covariates, unlike Nearest Neighbors approaches.

Finally, we conclude this section with Remark 4: forming risk strata may be an interesting strategy, that lies in-between the two extreme choices  $\mathcal{C}_{\text{Tot}}$  and  $\mathcal{C}_{\text{NN}}$  for pairing sets. The causal estimand related to a strata function can also be defined, and we expect results similar to those of Theorem 1 to hold. We argue that the individual-level causal estimand  $\tau_\star$  can in fact be recovered using infinitesimal stratas, further justifying the strength of Definition 2.

**Remark 4** (Win Ratio and comparisons with strata). The following notion of strata could be developed further, as an in-between the extreme cases presented in Theorem 1. Let  $r : \mathcal{X} \rightarrow \mathcal{R}$  be a risk strata function, on some metric space  $(\mathcal{R}, d)$ . Patients with similar risks ( $d(r(X_i), r(X_j))$  small) are expected to have similar control and treated behaviors. Using this strata, several studies build the following pairs of treated and control patients:

$$\mathcal{C}_{\hat{r}, \varepsilon} \stackrel{\text{def}}{=} \{(i, j) \in \mathcal{N}_0 \times \mathcal{N}_1 \mid d(r(X_i), r(X_j)) \leq \varepsilon\},$$

to obtain the estimator  $\hat{p}_W$  (Equation (4)) using these pairs. Now, define the  $\varepsilon$ -strata causal estimand as:

$$\tau_\star^{r, \varepsilon} \stackrel{\text{def}}{=} \mathbb{P}(Y_i(1) > Y_j(0) \mid d(r(X_i), r(X_j)) \leq \varepsilon),$$

where  $i$  and  $j$  are two different and independent indices. Under adequate, if  $\mathcal{C} = \mathcal{C}_{\hat{r},\varepsilon}$  we expect to have that:

$$\hat{p}_W \xrightarrow{\mathbb{P}} \tau_{\star}^{r,\varepsilon}.$$

Furthermore, our causal measure  $\tau_{\star}$  can in fact be seen as the limit of  $\tau_{\star}^{r,\varepsilon}$  for  $\varepsilon \rightarrow 0$ . However, this limit may not always be well-defined, hence the strength of Definition 2.

## 4 Observational setting

In this section, we introduce new methods to approximate  $\tau_{\star}$  beyond the RCT setting previously considered. In Section 4.1, we generalize the Nearest Neighbor pairing approach using Inverse Propensity Weights (IPW), to account for non-constant probability of treatment. Then, to address both missing values issues in the covariates (that can hardly be handled by Nearest Neighbors) and the slow convergence in large dimensions, we propose a Distributional Regression approach in Section 4.2, with a direct regression estimator, and a doubly-robust estimator, that both leverage recent advances in distributional regression (Distributional Random Forests, [Cevid et al. \[2022\]](#)).

### 4.1 An Inverse Propensity Weighting approach

We here generalize the traditional estimator  $\hat{p}_W$  with Nearest Neighbors (Theorem 1), that we used to compute Win Ratios and Net-Benefit in the previous section, to the observational setting in order to estimate  $\tau_{\star}$  beyond randomized controlled trials. We will use approximated propensity scores. Recall that for  $x \in \mathcal{X}$ ,  $\pi(x)$  is the probability of being treated, conditionally on  $X_i = x$  (as defined in Assumption 3):

$$\forall x \in \mathcal{X}, \quad \pi(x) \stackrel{\text{def}}{=} \mathbb{P}(T_i = 1 | X_i = x),$$

and

$$1 - \pi(x) \stackrel{\text{def}}{=} \mathbb{P}(T_i = 0 | X_i = x) .$$

$\pi(X_i)$  is then called the *propensity* score of patient  $i$ . In this section, we assume that we have access to approximated propensity scores, via  $\hat{\pi}$  an approximation of  $\pi$ .

We assume that  $\hat{\pi}$  is independent from the samples  $(X_i, T_i, Y_i)_{i \in [n]}$ : it has been computed using independent samples, via for instance a sample splitting approach. Note that this independence assumption could be removed using cross-fitting techniques, as for instance done in [Athey and Wager \[2021\]](#). In the observational setting, we only assume that [Assumption 2](#) holds (uncounfoundeness), instead of assuming that  $T_i \perp\!\!\!\perp X_i$ . Thus, estimators used in RCTs might be biased when used on observational data due to counfounders between treatment and outcomes. To unbiased RCT estimates, one thus resorts to *Inverse Propensity Weighting* (IPW) [[Robins et al., 1994](#), [Horvitz and Thompson, 1952](#)], as classically done for instance estimating the Average Treatment Effect with the Risk Difference in observational settings.

We thus adapt the Nearest Neighbor estimator (that is a consistent estimator for  $\tau_\star$  in the RCT setting), defined as:

$$\frac{1}{n_0} \sum_{i \in \mathcal{N}_0} w(Y_{\sigma_1^*(i)} | Y_i), \tag{10}$$

and used in [Theorem 1.2](#), where

$$\forall i \in \mathcal{N}_1, \quad \sigma_1^*(i) \in \operatorname{argmin}_{j \in \mathcal{N}_1} \|X_i - X_j\| ,$$

with uniform sampling if there are equalities. We replace this estimator by:

$$\hat{\tau}_{\text{IPW}} \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i \in \mathcal{N}_0} w(Y_{\sigma_1^*(i)} | Y_i) (1 - \hat{\pi}(X_i))^{-1}. \tag{11}$$

Note that in the context of a RCT, propensities are known and are constant (for all  $x \in \mathcal{X}$ ,  $\pi(x) = \pi$ ), and that  $n_0/n$  is an unbiased and consistent estimate of  $\pi$ . As such, Equation (10) is simply a specific case of Equation (11). We next show that  $\hat{\tau}_{\text{IPW}}$  is indeed a generalization of  $p_W$  with  $\mathcal{C}_{\text{NN}}$  to observational data, since it is a consistent estimator of the same causal estimand.

**Theorem 2.** Assume that Assumptions 1 to 3 hold. Assume that  $\hat{\pi}$  satisfies:

1. Pointwise consistency almost surely:  $\forall x \in \mathcal{X}$ , we have  $\mathbb{P}(\hat{\pi}(x) \rightarrow \pi(x)) = 1$ ;
2. Mean consistency:  $\mathbb{E}[|\hat{\pi}(X_i) - \pi(X_i)|] \rightarrow 0$ ;
3. Finite and bounded second moment of propensity scores:

$$\limsup \mathbb{E} \left[ \frac{1}{(1 - \hat{\pi}(X_i))^2} \right] \quad \text{and} \quad \limsup \mathbb{E} \left[ \frac{1}{\hat{\pi}(X_i)^2} \right] < \infty.$$

Assume finally that  $(x, y) \mapsto \mathbb{E}[w(Y_i(1)|y)|X_i = x]$  is continuous in its first variable  $x$  and that  $\mathcal{X}$  is compact. Then,  $\hat{\tau}_{\text{IPW}}$  (Equation (11)) is a consistent estimator of  $\tau_\star$  (Definition 2):

$$\frac{1}{n} \sum_{i \in \mathcal{N}_0} w(Y_{\sigma_1^*(i)}|Y_i)(1 - \hat{\pi}(X_i))^{-1} \xrightarrow{\mathbb{P}} \tau_\star,$$

as  $n_0, n_1 \rightarrow \infty$ .

As opposed to the inverse propensity weighting estimators provided by Mao [2017] and by many subsequent works [Chen et al., 2024, Chiaruttini et al., 2024, Guo and Ni, 2022, Yin et al., 2022, Zhang et al., 2022], our estimator  $\hat{\tau}_{\text{IPW}}$  is consistent for  $\tau_\star$  rather than for  $\tau_{\text{pop}}$ .

## 4.2 Distributional Regression Approach

We now introduce a distributional regression approach for estimating  $\tau_\star$ , to address the fact that the Nearest Neighbor approach cannot handle missing values in the covariates and may suffer from slow convergence if the dimension is too large. Our distributional regression approach is the counterpart of the *Two-learners* or *Plug-in G-formula* approaches, used for (C)ATE estimation with the risk difference. When estimating

$$\tau_{\text{RD}} = \mathbb{E}[Y_i|T_i = 1] - \mathbb{E}[Y_i|T_i = 0],$$

with a regression approach, the idea is to learn two (non)parametric estimates  $\hat{\mu}_t : \mathcal{X} \rightarrow \mathbb{R}$  of  $\mu_t : x \mapsto \mathbb{E}[Y_i|T_i = t, X_i = x]$ , to approximate  $\tau_{\text{RD}}$  with:

$$\frac{1}{n} \sum_{i=1}^n \hat{\mu}_1(X_i) - \hat{\mu}_0(X_i).$$

A naive adaptation of this regression approach to our problem of estimating  $\tau_*$  would thus be to first learn  $x \mapsto \hat{\mu}_t(x)$  (non)parametric estimates of  $\mu_t : x \mapsto \mathbb{E}[Y_i|X_i = x, T_i = t]$  as before. Then, let  $\hat{p}(X_i) = w(\hat{y}_1(X_i)|\hat{y}_0(X_i))$ , that aims at estimating  $\mathbb{E}[w(Y_i(1)|Y_i(0))|X_i]$ , to obtain the estimator  $\frac{1}{n} \sum_{i=1}^n \hat{p}(X_i)$ . However, as opposed to (C)ATE estimation, we don't have linearity of  $w$  here, so that even if the conditional expectations are perfectly estimated ( $\hat{\mu}_t = \mu_t$ ), we won't even have consistency. Regressing the conditional expectations makes us loose information on the way.

Hence the *distributional regression approach*, since we need to go beyond learning conditional expectations. First, notice that we have, using the independence between  $Y^{(X_i)}$  and  $Y_i$  conditionally on  $X_i$ :

$$\mathbb{E}[w(Y^{(X_i)}(1)|Y_i(0))|X_i] = \int_{\mathcal{Y}} \mathbb{E}[w(Y_i(1)|y)|X_i] d\mathbb{P}(Y_i(0) = y|X_i),$$

and equivalently:

$$\mathbb{E}[w(Y^{(X_i)}(1)|Y_i(0))|X_i] = \int_{\mathcal{Y}} \mathbb{E}[w(y|Y_i(0))|X_i] d\mathbb{P}(Y_i(1) = y|X_i).$$

Define:

$$q_t : (x, y) \in \mathcal{X} \times \mathcal{Y} \mapsto \mathbb{E}[w(tY_i(t) + (1-t)y|ty + (1-t)Y_i(t))|X_i = x],$$

so that

$$q_1(x) = \mathbb{E}[w(Y_i(1)|y)|X_i = x],$$

and

$$q_0(x) = \mathbb{E}[w(y|Y_i(0))|X_i = x].$$

Thus, if we learn a (non)parametric estimate  $\hat{q}_t$  of  $q_t$  for  $t \in \{0, 1\}$ , we have a candidate estimator for  $\tau_\star$  if we are able to sample from  $d\mathbb{P}(Y_i(t)|X_i)$

In the RCT setting, noting  $\pi = \mathbb{P}(T_i = 1)$ , two candidate estimators would be:

$$\hat{S}_t = \frac{1}{n} \sum_{i \in [n]} \left\{ t \hat{q}_1(X_i, Y_i) \frac{1 - T_i}{1 - \pi} + (1 - t) \hat{q}_0(X_i, Y_i) \frac{T_i}{\pi} \right\},$$

for  $t \in \{0, 1\}$ . Both  $\hat{S}_0$  and  $\hat{S}_1$  are unbiased estimators of  $\tau_\star$  if we have  $\hat{q}_t = q_t$  (perfect estimation).

In the observational setting, this can be generalized to, using IPW weights:

$$\hat{S}_1 = \frac{1}{n} \sum_{i=1}^n \hat{q}_1(X_i, Y_i) \frac{1 - T_i}{1 - \hat{\pi}(X_i)}, \quad \text{or} \quad \hat{S}_0 = \frac{1}{2n} \sum_{i=1}^n \hat{q}_0(X_i, Y_i) \frac{T_i}{\hat{\pi}(X_i)},$$

where  $\hat{\pi}$  are propensity scores estimated on an independent dataset. Indeed, we then have if the exact propensity scores are known ( $\hat{\pi} = \pi$ ):

$$\begin{aligned} \mathbb{E} \left[ \hat{q}(X_i, Y_i) \frac{1 - T_i}{1 - \pi(X_i)} \right] &= \mathbb{E} [\mathbb{E} [\hat{q}(X_i, Y_i) \mid X_i, T_i = 0]] \\ &= \mathbb{E} \left[ \int \hat{q}_i(X_i, y) dP(Y_i(0) = y \mid X_i) \right]. \end{aligned}$$

If the exact conditional expectations are known (*i.e.*, if  $\hat{q}_t = q_t$ ), the above expression is then equal to  $\tau_\star$ . This approach suffers from both the nuisance factors of the estimated propensity scores  $\hat{\pi}$  and of the distributional regressions  $\hat{q}_t$ . It however appears that we can go beyond this dependency on propensity scores, by carefully combining distributional regressions performed on control and test groups. The approach we introduce next only relies on distributional regression, and is as such closer to more traditional regression approaches such as the *Two learner* or *Plug-in G-formula* methods.

#### 4.2.1 The direct distributional regression estimator

Let:

$$\hat{\tau}_{\text{reg}} \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n (1 - T_i) \hat{q}_1(X_i, Y_i) + T_i \hat{q}_0(X_i, Y_i), \quad (12)$$

be the *distributional regression estimator*. We have the following first result, that justifies the use of  $\hat{\tau}_{\text{reg}}$ . If distributional regression is perfect (no estimation error, i.e.  $\hat{q}_t = q_t$ ), then  $\hat{\tau}_{\text{reg}}$  is an unbiased estimator of  $\tau_\star$ .

**Proposition 2.** *Assume that Assumptions 1 and 2 hold (SUTVA and unconfoundedness) and that  $\hat{q}_t = q_t$  for  $t \in \{0, 1\}$ . Then,  $\hat{\tau}_{\text{reg}}$  is an unbiased estimator of  $\tau_\star$ .*

*Proof.* If there's no error in estimation and using unconfoundedness:

$$\begin{aligned}
& \mathbb{E} [(1 - T_i)\hat{q}_1(X_i, Y_i) + T_i\hat{q}_0(X_i, Y_i)] \\
&= \mathbb{E} [\mathbb{P}(T_i = 0|X_i) \mathbb{E} [w(Y^{(X_i)}(1)|Y_i)|X_i, Y_i, T_i = 0]] \\
&\quad + \mathbb{E} [\mathbb{P}(T_i = 1|X_i) \mathbb{E} [w(Y_i|Y^{(X_i)}(0))|X_i, Y_i, T_i = 1]] \\
&= \mathbb{E} [(1 - \pi(X_i))\mathbb{E} [w(Y^{(X_i)}(1)|Y_i(0))|X_i] + \pi(X_i)\mathbb{E} [w(Y_i(1)|Y^{(X_i)}(0))|X_i]] \\
&= \mathbb{P}(Y^{(X_i)}(1) \geq Y_i(0)) ,
\end{aligned}$$

using that conditionally on  $X_i$ , the random variables  $(Y_i(0), Y_i(1))$  and  $(Y^{(X_i)}(0), Y^{(X_i)}(1))$  are independent and identically distributed.  $\square$

Moving beyond simple unbiasedness, provided that the estimation errors tend to zero (in mean over the population), we prove in the next Theorem that we have consistency of  $\hat{\tau}_{\text{reg}}$ , and even asymptotic normality if the estimation error is sufficiently small.

**Theorem 3** (Consistency and asymptotic normality). *Assume that Assumptions 1 to 3 hold.*

*We have the followings.*

1. *Assume that for  $t \in \{0, 1\}$ , we have:*

$$\mathbb{E} [|\hat{q}_t(X_i, Y_i) - q_t(X_i, Y_i)| | T_i = 1 - t] \rightarrow 0 .$$

*Then, the estimator  $\hat{\tau}_{\text{reg}}$  defined in Equation (12) converges almost surely to  $\tau_\star$  (defined in Definition 2).*

2. Assume that for  $t \in \{0, 1\}$ , we have:

$$\mathbb{E} [|\hat{q}_t(X_i, Y_i) - q_t(X_i, Y_i)| | T_i = 1 - t] = o(n^{-1/2}).$$

Then, the estimator  $\hat{\tau}_{\text{reg}}$  defined in Equation (12) satisfies:

$$\sqrt{n}(\hat{\tau}_{\text{reg}} - \tau_{\star}) \xrightarrow{\mathbb{P}} \mathcal{N}(0, \sigma_{\infty}^2),$$

where

$$\begin{aligned} \sigma_{\infty}^2 &= \mathbb{P}(T_i = 1) \text{var}(\mathbb{E}[w(Y_i(1), Y_i(0)) | X_i, Y_i(0), T_i = 1]) \\ &\quad + \mathbb{P}(T_i = 0) \text{var}(\mathbb{E}[w(Y_i(1), Y_i(0)) | X_i, Y_i(1), T_i = 0]) \end{aligned}$$

is the variance of the probability of a win conditioned on covariates, treatment, and counterfactual.

The assumption of Theorem 3.1 will hold, as long as the distributional regressors are consistent. This will for instance be the case of Distributional Random Forests [Cevid et al., 2022] that we use (defined and explained further in Section 4.2.3), under very mild assumptions. The assumption of Theorem 3.2 is much stronger, and requires a fast parametric rate of convergence. It will hold if for instance we perform logistic regression (Section 4.2.3) for a well-specified distributional regression problem. Under such assumptions, the asymptotic normality yields asymptotically valid confidence intervals.

#### 4.2.2 The doubly robust estimator

We can finally build the following doubly robust estimator, that combines both the distribution regression estimator  $\hat{\tau}_{\text{reg}}$  (Equation (12)) and the inverse propensity weighting estimator  $\hat{\tau}_{\text{IPW}}$  (Equation (11)), by using distributional regression estimates  $\hat{q}_t$  and approximated propen-

sity scores  $\hat{\pi}$ :

$$\begin{aligned} \hat{\tau}_{\text{AIPW}} \stackrel{\text{def}}{=} & \frac{1}{n} \sum_{i=1}^n \left\{ (1 - T_i) \hat{q}_1(X_i, Y_i) - \lambda (\hat{q}_1(X_i, Y_i) - w(Y_{\sigma_0(i)} | Y_i)) \frac{1 - T_i}{(1 - \hat{\pi}(X_i))} \right\} \\ & + \frac{1}{n} \sum_{i=1}^n \left\{ T_i \hat{q}_0(X_i, Y_i) - (1 - \lambda) (\hat{q}_0(X_i, Y_i) - w(Y_i | Y_{\sigma_1(i)})) \frac{T_i}{\hat{\pi}(X_i)} \right\}, \end{aligned} \quad (13)$$

where  $\lambda \in (0, 1)$  is a parameter independent of the  $n$  datasamples,  $\sigma_0, \sigma_1$  are Nearest Neighbor pairings on control and test groups, in the sense that

$$\sigma_0 \in \operatorname{argmin}_{\sigma: \mathcal{N}_0 \rightarrow \mathcal{N}_1} \sum_i \|X_i - X_{\sigma(i)}\|^2,$$

and

$$\sigma_1 \in \operatorname{argmin}_{\sigma: \mathcal{N}_1 \rightarrow \mathcal{N}_0} \sum_j \|X_j - X_{\sigma(j)}\|,$$

where in case of equality in the argmin, a uniform sampling over the minimizers is performed.

We have the following (weak) double robustness property for  $\tau_{\text{AIPW}}$ : if either  $\hat{\pi}$  or both  $\hat{q}_t$  are good estimators of the respective quantities they seek at estimating, then  $\hat{\tau}_{\text{AIPW}}$  is a consistent estimator of  $\tau_\star$ .

**Theorem 4.** *Assume that either (i) the assumptions of Theorem 2 or (ii)  $\eta' \leq \hat{\pi} \leq 1 - \eta'$  a.s., and the assumptions of Theorem 3 holds. Assume that  $\lambda \rightarrow \mathbb{P}(T_i = 0)$  almost surely. Then, we have:*

$$\hat{\tau}_{\text{AIPW}} \xrightarrow[\mathbb{P}]{} \tau_\star.$$

Classically, our augmented estimator combines both non-parametric estimators  $\hat{q}_t$ , that here come from distribution regressions, and estimated propensity weights  $\hat{\pi}$ . There is here however an additional and less conventional parameter  $\lambda$ . The role of  $\lambda$  is to balance between treated and control groups. Having  $\lambda$  that converges almost surely to  $\mathbb{P}(T_i = 0)$  is a necessary assumption for our doubly robust estimator to be a consistent estimator of  $\tau_\star$ . This assumption can easily be imposed by setting  $\lambda = \frac{1}{n'} \sum_{j=1}^{n'} \mathbb{1}_{\{T'_j=0\}}$  for some independent samples  $\{(X'_j, T'_j, Y'_j), j \in [n']\}$ , obtained by randomly splitting our dataset.

### 4.2.3 Distributional random forests and logistic regression for estimating $q_1(x, y)$

So far, we have not specified the methods used to perform distributional regression to learn the (non)parametric estimates  $\hat{q}_t$  of  $q_t$ . We describe in this section a non-parametric regression approach using *distributional random forests* [Cevid et al., 2022, Bénard et al., 2024], and a parametric approach using logistic regression.

**Distributional random forests (DRF).** Our goal is to estimate

$$q_1(x, y) = \mathbb{P}(w(Y_i|y)|X_i = x, T_i = 1) ,$$

with distributional random forests [Cevid et al., 2022, Bénard et al., 2024]. Let  $\mathcal{H}$  a Hilbert space with a kernel  $k(\cdot, \cdot)$  defined on  $\mathcal{Y} \times \mathcal{Y}$  (usually the Gaussian kernel). For any probability measure  $\mathbb{P}$  on  $\mathcal{Y}$ , let  $\phi(\mathbb{P}) = \mathbb{E}[k(Z, \cdot)] \in \mathcal{H}$ , for  $Z \sim \mathbb{P}$ . Distributional random forests approximate  $\mu(x, t) = \phi(\mathbb{P}_{Y|X=x, T=t})$  using splitting rules in the Hilbert space  $\mathcal{H}$ , and as such are simply generalization of vanilla random forests to Hilbert spaces. Then, Cevid et al. [2022] use the fact that for kernels such as the Gaussian kernel, learning kernel representations amounts to learning probability distributions. Thus, learning with samples  $(X_1, Y_1), \dots, (X_m, Y_m)$ ,  $\mu(x, t)$  is approximated via  $\hat{\mu}_m(x, t)$  that takes the form:

$$\hat{\mu}_m(x, t) = \sum_{i=1}^m \omega_i(x, t)k(Y_i, \cdot) \in \mathcal{H} ,$$

for weights  $\omega_i(\cdot, \cdot)$  learnt by the forest. The formula on the right hand side can be written as  $\sum_{i=1}^m \omega_i(x, t)k(Y_i, \cdot) = \phi(\sum_{i=1}^m \omega_i(x, t)\delta_{Y_i})$ , where for  $y \in \mathcal{Y}$ ,  $\delta_y$  is a Dirac of mass 1 at point  $y$ .

The distribution  $\mathbb{P}_{Y|X=x, T=t}$  is thus approximated by:

$$\hat{\mathbb{P}}_{Y|X=x, T=t}^{(m)} = \sum_{i=1}^m \omega_i(x, t)\delta_{Y_i} .$$

Hence, the probability  $\mathbb{P}(Y \in \mathcal{S}|X = x, T = t)$  can be estimated by:

$$\hat{\mathbb{P}}^{(m)}(Y \in \mathcal{S}|X = x, T = t) \stackrel{\text{def}}{=} \sum_{i=1}^m \omega_i(x, t)\delta_{\{Y_i \in \mathcal{S}\}} .$$

More generally, any conditional expectation  $\mathbb{E}[h(Y)|X = x, T = t]$  for some measurable  $h : \mathcal{Y} \rightarrow \mathbb{R}$  can be approximated by:

$$\hat{\mathbb{E}}^{(m)}[h(Y)|X = x, T = t] \stackrel{\text{def}}{=} \sum_{i=1}^n \omega_i(x, t) h(Y_i).$$

Now, we remark that the quantity that we wish to estimate writes as  $q_1(x, y) = \mathbb{P}(h_y(Y_i)|X_i = x, T_i = 1)$ , where  $h_y(Y_i) = w(Y_i|y)$ . Thus, our estimate  $\hat{q}_1(x, y)$  of  $q_1(x, y) = \mathbb{P}(w(Y_i|y)|X_i = x, T_i = 1)$  is:

$$\hat{q}_1(x, y) \stackrel{\text{def}}{=} \sum_{i=1}^n \omega_i(x, 1) w(Y_i|y),$$

while our estimate  $\hat{q}_0(x, y)$  of  $q_0(x, y) = \mathbb{P}(w(y|Y_i)|X_i = x, T_i = 0)$  is:

$$\hat{q}_0(x, y) \stackrel{\text{def}}{=} \sum_{i=1}^n \omega_i(x, 0) w(y|Y_i).$$

In practice, these steps are implemented in a very concise way, using the following pseudo code to implement the distributional regression estimator (Equation (12)).

1. Dataset  $\mathcal{D} = \{(X_i, T_i, Y_i), i \in [n]\}$  is split between a train set  $\mathcal{D}_{\text{train}} = \{(X_i, T_i, Y_i), i \in [m]\}$  and an inference set  $\mathcal{D}_{\text{inference}} = \{(X_i, T_i, Y_i), m + 1 \leq i \leq n\}$ .
2. A Distributional Random Forest (DRF) is trained on  $\mathcal{D}_{\text{train}}$  to predict  $Y_i$  from  $(X_i, T_i)$ , using the R package `drf` [2020], an implementation of DRFs as introduced by Cevid et al. [2022].
3. Apply the DRF to predict on the inference set  $\mathcal{D}_{\text{inference}}$ , to obtain the weights  $\omega_i(X_j, T_j)$  for  $i$  a train point and  $j$  in the inference set.
4. Compute and output:

$$\frac{1}{n - m} \sum_{j=m+1}^n \sum_{i=1}^m \left[ \mathbb{1}_{\{T_j=0\}} \omega_i(X_j, 0) w(Y_i|Y_j) + \mathbb{1}_{\{T_j=1\}} \omega_i(X_j, 1) w(Y_j|Y_i) \right].$$

**Logistic regression for outcomes**  $\mathcal{Y} = \{0, 1\}^d$ . We now present a parametric approach to estimate  $q_t$ ,  $t \in \{0, 1\}$ , using linear logistic regression. This approach is here described in the case of multivariate binary outcomes  $\mathcal{Y} = \{0, 1\}^d$  with the win function  $w(y|y') = \mathbb{1}_{\{y > y'\}}$ , but can be generalized to any categorical outcomes. The idea is to fit a generalized linear regression model to learn  $x \in \mathcal{X} \mapsto \mathbb{P}(Y_i = y | X_i = x, T_i = t)$  for  $y \in \mathcal{Y} = \{0, 1\}^d$  and  $t \in \{0, 1\}$ : this is a multiclass classification problem. We parameterize our classifier using  $(u_k^{(t)})_{k \in [d]}$  where  $u_k^{(t)} \in \mathbb{R}^p$  for  $k \in [d]$  and  $t \in \{0, 1\}$  (for  $\mathcal{X} \subset \mathbb{R}^p$ ), and learn a function of the form:

$$(x, t) \in \mathcal{X} \times \{0, 1\} \quad \mapsto \quad \prod_{k=1}^d \text{expit}(\langle u_k^{(t)}, x \rangle),$$

where  $\text{expit}(s) = \frac{e^s}{1+e^s}$  for  $s \in \mathbb{R}$ . Interactions can then be imposed, by setting some constraints on the  $u_k^{(t)}$ , such as  $u_k^{(t)} = u_\ell^{(t)}$  for all  $k, \ell$  and fixed  $t$ , for full interactions. The weights are learnt by minimizing an empirical loss of the form:

$$\mathcal{L} \left( \left\{ u_k^{(t)}, k \in [d], t \in \{0, 1\} \right\} \right) \stackrel{\text{def}}{=} \frac{1}{m} \sum_{i=1}^m \sum_{k=1}^d \log \left( \text{expit}(-2(Y_i - 1/2) \langle u_k^{(T_i)}, X_i \rangle) \right).$$

If this model is well-specified (in the sense that the data is indeed generated by Bernoulli random variables of the form  $\mathbb{P}(Y_{i,k}(t) = 1) = \text{expit}(\langle u_k^{(t), \star}, X_i \rangle)$ ), we expect a fast parametric statistical rate and the assumption of Theorem 3.2 to hold.

## 5 Experiments on randomly generated data

We first start with experiments on random synthetic data.

### 5.1 The impact of the dimension on data with correlated and non-correlated outcomes

We generate synthetic observational data as follows. Inputs  $X_i$  and outputs  $Y_i$  respectively lie in  $\mathbb{R}^p$  and  $\{0, 1\}^d$ . The *win* function on  $\{0, 1\}^d$  we then use is  $w(y, y') = \mathbb{1}_{\{y > y'\}}$ , for  $>$  the

lexicographic order as in Pocock et al. [2011].

1. Covariates are generated as standard multivariate Gaussian random variables  $X_i \sim \mathcal{N}(0, I_p)$ .
2. Treatment assignments follow a Bernoulli law of mean  $\sigma(\langle X_i, v \rangle)$ , where  $v \in \mathbb{R}^p$  is unitary and  $\sigma : \mathbb{R} \rightarrow [0, 1]$ .
3. Potential outcomes  $Y_i(0), Y_i(1) \in \{0, 1\}^d$  follow multidimensional Bernoulli laws of parameters  $\langle u_i^{(t)}, X_i \rangle, i \in [d]$  for  $u_i^{(t)} \in \mathbb{R}^p$ :

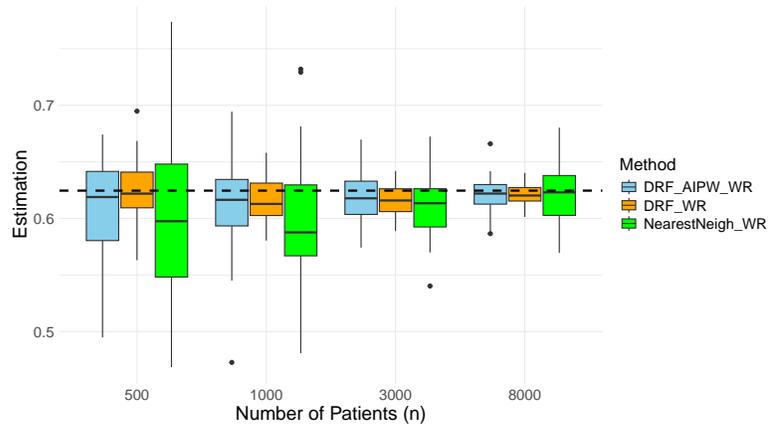
$$\mathbb{P}(Y_i(t) = (e_1, \dots, e_d)) = \prod_{i=1}^d \left( e_i \sigma(\langle u_i^{(t)}, X_i \rangle) + (1 - e_i)(1 - \sigma(\langle u_i^{(t)}, X_i \rangle)) \right),$$

for any  $(e_1, \dots, e_d) \in \{0, 1\}^d$ .

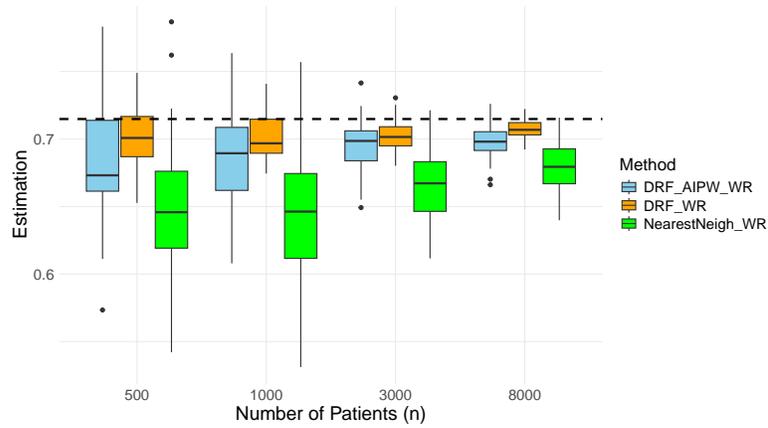
We then distinguish two scenarios in terms of outcomes: correlated and uncorrelated ones, uncorrelated outcomes being harder for distributional regression (more parameters to learn).

1. The first one is when the different outcomes have very strong correlations: vectors  $(u_i^{(t)})_{i \in [d]}$  are highly correlated. We model this using  $u^{(0)}, u^{(1)}$  two unitary vectors, and setting  $u_i^{(t)} = u^{(t)}$  for all  $i \in [d]$ . This is the **correlated outcomes setting**.
2. The second scenario is the **uncorrelated outcomes setting**, where we instead take  $u_i^{(t)}$  as random unitary vectors: there is no correlation between the different multiple outcomes.

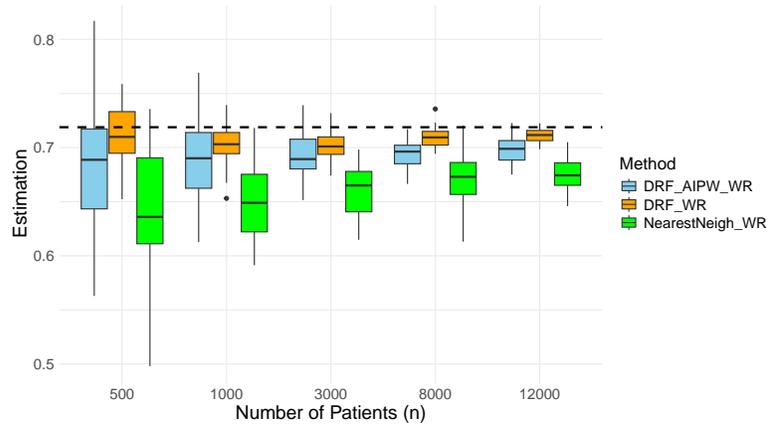
Figures 2 and 3 study the impact of the dimension ( $p, d$  increasing) in the correlated outcomes setting, while Figures 4 and 5 study the impact of the dimension in the uncorrelated outcomes setting. We also provide experiments where AIPW and IPW estimators use oracle propensity weights (*‘cheating’* estimators, as referred to in the plots), to show that the bias is due to the dimension and nearest neighbors rather than propensities that are not well estimated. These illustrate the shortcomings of the nearest neighbor approach when the dimension becomes larger, and the strength of our distributional regression approach.



(a)  $d = p = 3$

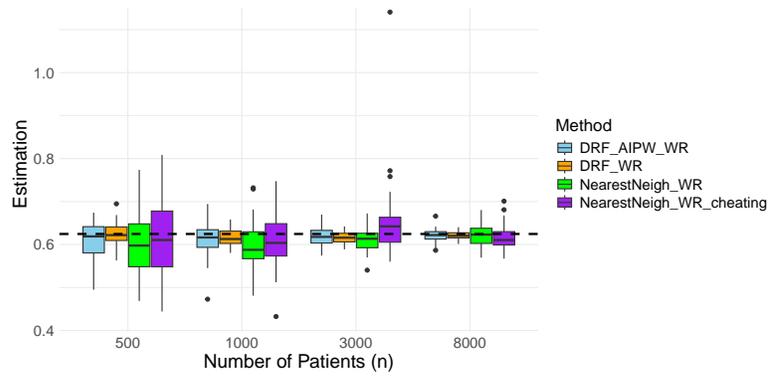


(b)  $d = p = 8$

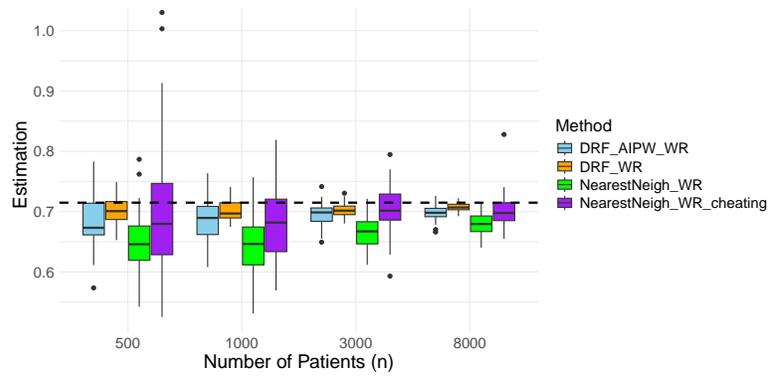


(c)  $d = p = 20$

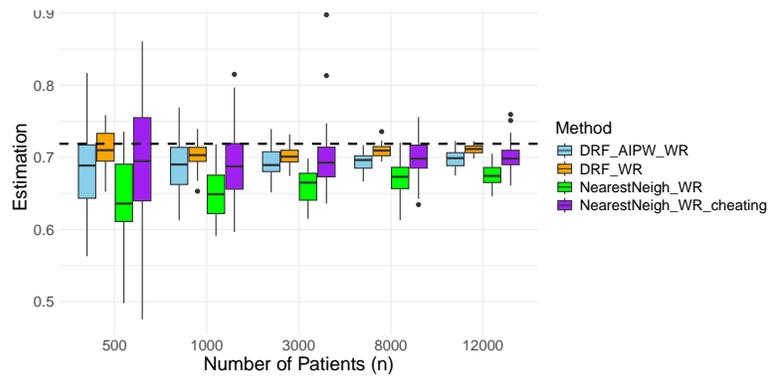
Figure 2: Testing for the impact of the dimension, correlated outcomes setting. Boxplots over 100 runs. DRF AIPW WR, DRF WR and NearestNeigh WR respectively correspond to the AIPW method in Equation (13), the direct distributional approach in Equation (12) and the weighted Nearest Neighbor approach in Equation (11). For the AIPW and distributional regression approach, Distributional Random Forests (DRFs, Bénard et al. [2024]) are used to



(a)  $d = p = 3$

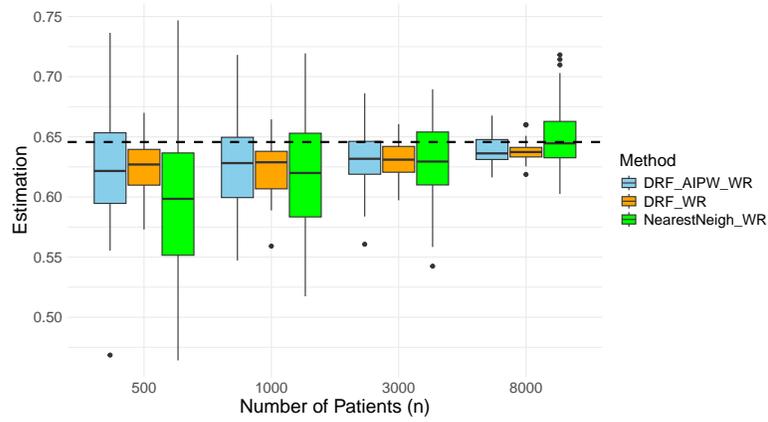


(b)  $d = p = 8$

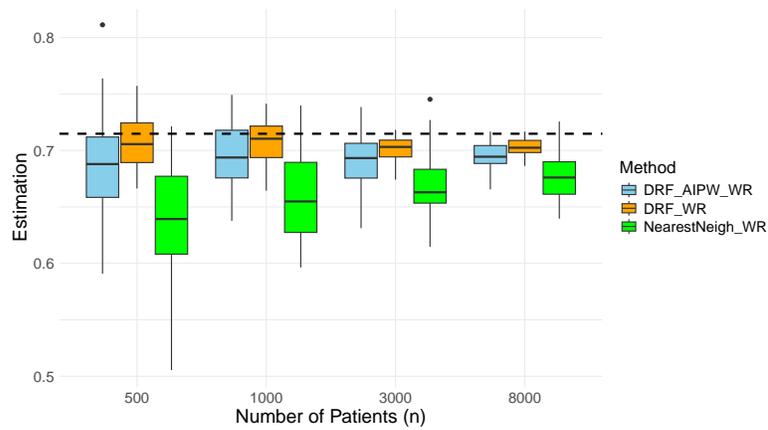


(c)  $d = p = 20$

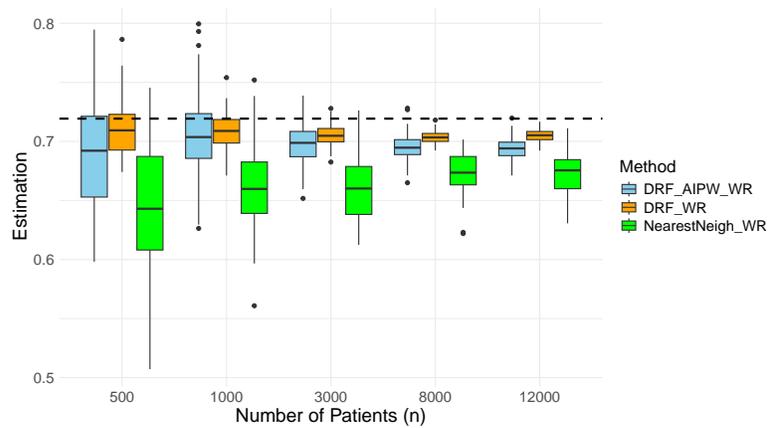
Figure 3: Same setting as in Figure 2, with an added method: Nearest Neighbor with a ‘cheating’ option, that corresponds to exactly plugging in the propensity scores instead of estimating them. Boxplots over 100 runs.



(a)  $d = p = 3$

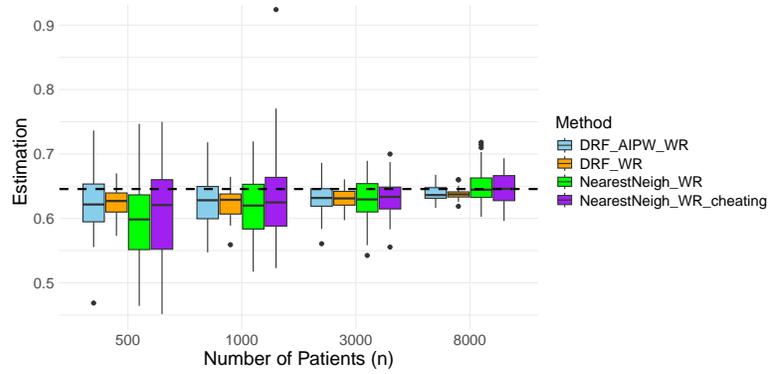


(b)  $d = p = 8$

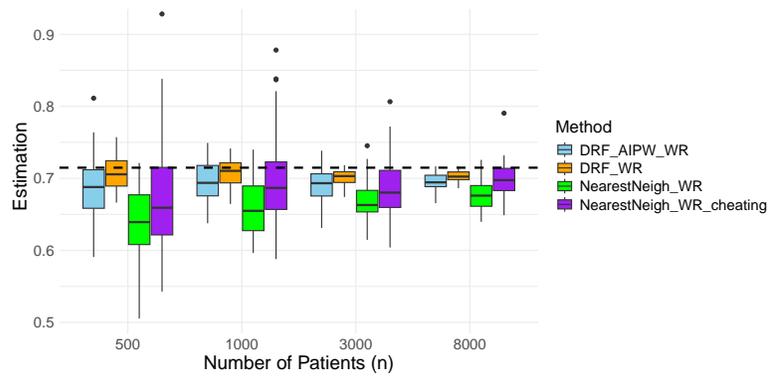


(c)  $d = p = 20$

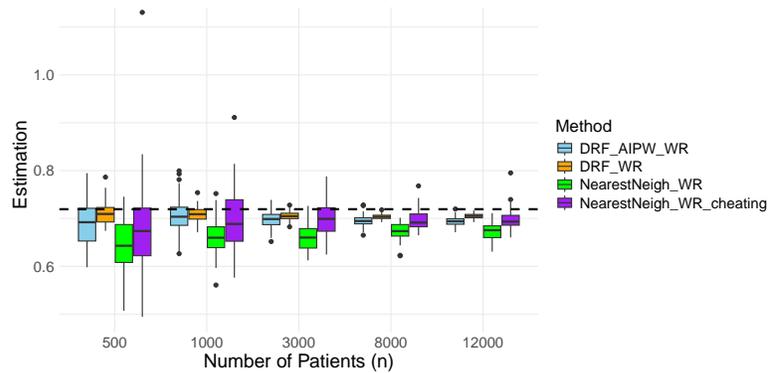
Figure 4: Testing for the impact of the dimension, uncorrelated outcomes setting. Boxplots over 100 runs. DRF AIPW WR, DRF WR and NearestNeigh WR as in Figure 2



(a)  $d = p = 3$



(b)  $d = p = 8$



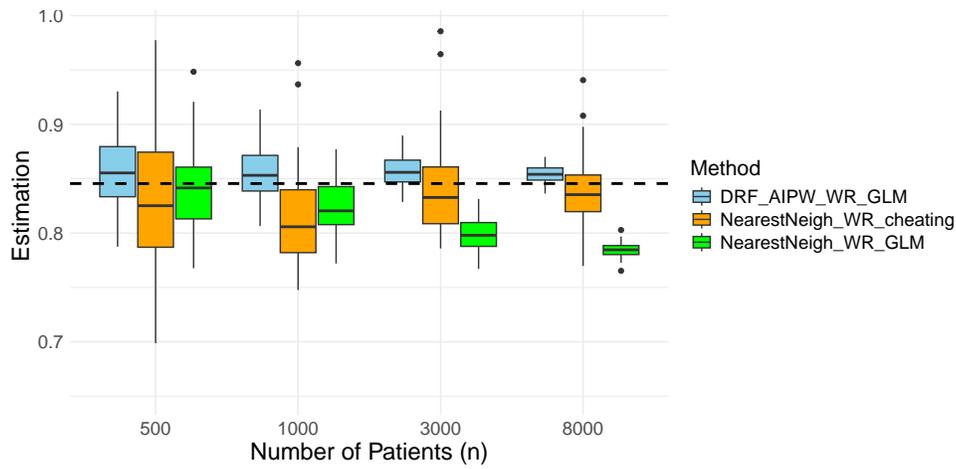
(c)  $d = p = 20$

Figure 5: Testing for the impact of the dimension, uncorrelated outcomes setting. Boxplots over 100 runs. DRF AIPW WR, DRF WR, NearestNeigh WR with and without ‘CHEATING’ option as in Figure 3

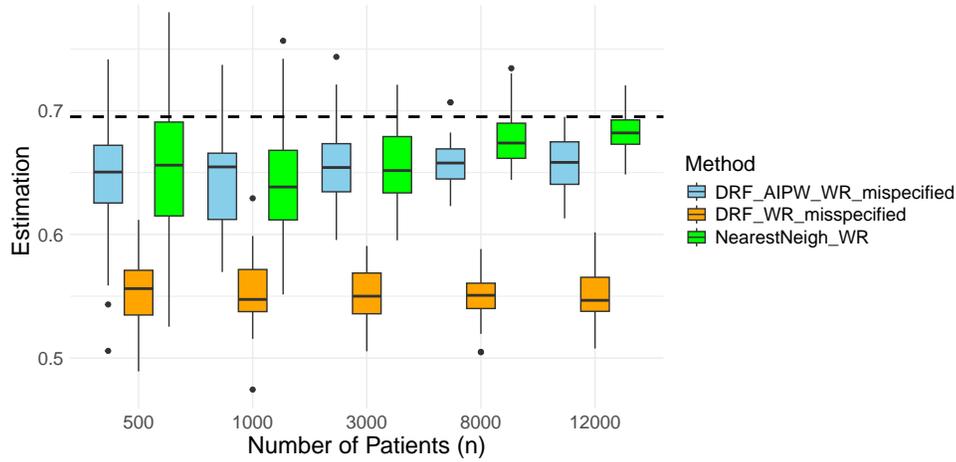
## 5.2 Misspecification and double robustness

We next provide two synthesized experiments, to illustrate the double-robustness of our augmented estimator (Equation (13)), that respectively correspond to Figure 6a and Figure 6b.

1. In the first experiment, we test for double robustness by misspecifying propensities. Instead of using probability forests to learn propensities, we use a linear classifier that we train using logistic regression. We generate probability of treatment *non-linearly*, of the form  $\sigma(X_{i,1}X_{i,2})$  for some  $\sigma : \mathbb{R} \rightarrow [0, 1]$  (Cdf of a Gaussian, in our case). Treatment responses are then generated as multi dimensional Bernoulli random variables, of means  $\sigma((X_{i,1} - X_{i,2})^2)$  and  $\sigma((X_{i,1} + X_{i,2})^2)$  for respectively treated and non-treated individuals.
2. The second experiment tests for double robustness by misspecifying in the distributional regression. We perform distributional regression as in Section 4.2.3 with logistic regression on the outcomes, by training as in the homogeneous setting, *i.e.* by imposing  $u_1^{(t)} = \dots = u_p^{(t)}$  for  $t = 0, 1$ . We generate treatment assignments and responses exactly as in Figure 5 (heterogeneous setting).



(a) Testing for robustness to misspecified propensities,  $d = p = 8$



(b) Testing for robustness to misspecified distributional regression,  $d = p = 5$

Figure 6: Testing for double robustness, by misspecifying either propensities or distributional regression.. Boxplots over 100 runs. DRF AIPW WR, DRF WR, NearestNeigh WR with as in Figure 3, ‘misspecified’ refers to learning a linear logistic regression for propensities (Figure 6a), or doing logistic distribution regression as in Section 4.2.3 and imposing a correlated outcomes for distributional regression (fig. 6b) while the outcomes are generated uncorrelated.

## 6 Application to the CRASH-3 RCT

We finally illustrate our methodologies to perform a Win Ratio analysis of the CRASH-3 trial dataset, that comes from the CRASH-3 clinical trial [CRASH et al. \[2019\]](#), which studied the effects of tranexamic acid (TXA) in traumatic brain injury (TBI). In this section, we illustrate the different ways our methodologies can be used to derive Win Ratio estimates with confidence intervals. Overall, our conclusions are that different methods should be used for increased robustness of the results, since the properties of the estimands that are being targeted may differ.

### 6.1 Presentation of the dataset

**Data description and preprocessing.** The CRASH-3 RCT contains information on 12, 737 patients. In order to have lighter computations, we chose to use only a random sample (without replacement) of 6, 000 patients for all our analysis. Missing data is imputed using mice [[Van Buuren and Groothuis-Oudshoorn, 2011](#)].

1. **Patients covariates** include: siteId (hospital identifier), sex (male/female), age (years), timeSinceInjury (hours since injury), sbpStatus (systolic blood pressure category), systolicBloodPressure (mmHg), gcsEyeOpening, gcsMotorResponse, gcsVerbalResponse (Glasgow Coma Scale scores), gcsTiming (time of GCS assessment), pupilReact (pupil reactivity), majorExtracranial (major extracranial injury), intraCranialBleeding (intracranial bleeding), epidural, subdural, subarachnoid, parenchymal, intraventricular (types of brain injuries), eligible (study eligibility), consent (study consent), eyeOpening, communicationAbility, motorResponse, feeding, toileting, grooming, levelOfFunctioning, employability, walking, washingDressing (functional outcomes), painDiscomfort, anxietyDepression, agitationAggression, fatigue (quality of life indicators), daysIcu (ICU stay duration),

neuroHaemEvac (hematoma evacuation surgery), neuroOther (other neurosurgery), and estBloodLoss (estimated blood loss).

2. Patients are either assigned **Placebo** (6,321) or **TXA** (6,416 patients).
3. **Primary outcomes** are *death events* in the 28 days following trauma, that we encode as 1 or 0.
4. **Secondary outcomes** are vascular risks. We encode them as 1 (if there is at least one stroke, heart attack, pulmonary embolism or deep vein thrombosis) or 0 (if there are no such events).
5. **Tertiary outcomes** are the number of days the patient stayed in the hospital (censored at 28 days).

Computing the average treatment effect for each of these 3 outcomes lead respectively to the confidence intervals, where  $Y_1, Y_2, Y_3$  are respectively our death, secondary effects and hospitalization duration outcomes:

$$\mathbb{E}[Y_1(1) - Y_1(0)] \in [-0.0032, 0.0033],$$

$$\mathbb{E}[Y_2(1) - Y_2(0)] \in [-0.0021, 0.0079],$$

and

$$\mathbb{E}[Y_3(1) - Y_3(0)] \in [-0.2305, 0.4512],$$

computed using the difference of means estimator. None of these intervals are significant. Here are some key conclusions from the study made by [CRASH et al. \[2019\]](#), to further contextualize the dataset:

1. TXA reduces the risk of death due to TBI, but only if given within 3 hours of injury. The earlier TXA is given, the greater the benefit. No significant benefit if given after 3 hours.

2. TXA did not increase the risk of Stroke, Heart attack, Pulmonary embolism, Deep vein thrombosis (DVT).
3. Greatest Benefit for Mild-to-Moderate TBI. TXA had the most impact on patients with mild-to-moderate TBI (Glasgow Coma Scale 9-15). No significant survival benefit in severe TBI ( $GCS \leq 8$ ), possibly because of the high fatality rate.
4. No Increase in Disability. TXA did not increase the number of survivors with severe disability. Patients who survived had similar functional outcomes to those in the placebo group.
5. Safe and Cost-Effective. TXA is cheap and widely available, making it a practical treatment for emergency trauma care. Safe for use in pre-hospital settings and emergency departments.

The final takeaway is that TXA is an effective, safe, and low-cost intervention that can save lives in TBI when given early (within 3 hours). However, it does not help much in severe TBI and must be administered as soon as possible after injury.

## 6.2 The different methodologies used

We compared the following methodologies in Figure 7:

1. Traditional win ratio, computed using all pairs of the dataset using the WINS package [Cui and Huang, 2021].
2. Stratified Win Ratio [Dong et al., 2018]: we stratify according to the time since injury the patient received the treatment (TXA or placebo). We computed the median of time since injury, and made 2 stratas: patients below and above this median.

3. Nearest neighbor approach for win ratio, as introduced in this paper. We either: **(i)** use the Manahalahobis distance on the covariates, to naturally balance between different scales, or **(ii)** perform a **Factor Analysis of Mixed Data (FAMD)** of the dataset (the equivalent of a PCA for combined categorical and numerical covariates) from the FactoMineR package [Lê et al., 2008]. We perform the FAMD analysis of the covariates, keep 95% of the variance explained, and perform nearest neighbors on these dimensions. Nearest neighbors are computed using the MatchIt package. Since nearest neighbors algorithm relies on distance metrics (like Euclidean or Manahalahobis distances), it struggles with mixed data types. FAMD transforms both numerical and categorical variables into a common latent space, ensuring a more meaningful distance calculation. Furthermore, nearest neighbors algorithm suffers from high-dimensional data because distances become less meaningful in higher dimensions. FAMD captures the most important variations in fewer dimensions, improving NN’s effectiveness.

4. Another version of nearest neighbors: *Optimal Matchings*, that solve (in the case where  $|\mathcal{N}_1| \leq |\mathcal{N}_0|$ )

$$\min_{\sigma: \mathcal{N}_1 \rightarrow \mathcal{N}_0 \text{ injective}} \sum_{i \in \mathcal{N}_1} \|X_{\sigma(i)} - X_i\|^2.$$

The difference with Nearest Neighbors as studied in our paper is that two treated units cannot be assigned to the same control unit. We use this on both the full covariates and on the features in the FAMD latent space.

5. Distributional Random Forests, as described in Section 4.2. We estimate the *Win Proportion* (obtained with  $w(y, y') = \mathbb{1}_{\{y > y'\}}$ ) and the *Loss Proportion* (obtained with  $w(y, y') = \mathbb{1}_{\{y \leq y'\}}$ ), and estimate the Win Ratio as the ratio between win et loss proportions.

6. Our doubly robust approach, as described in Section 4.2. For the DRF and AIPW DRF, forests are taken with 1000 trees.

For confidence intervals, we either use those given by the WINS package for computing win ratios (that use asymptotic gaussian approximations), or use bootstrapping with 1,000 bootstraps for distributional regression approaches. Gaussian approximations for Win Ratio consist of outputting confidence intervals that are of the form:

$$\text{Win Ratio} \in \frac{\#Wins \pm 1,96\sqrt{\#Wins}}{\#Losses \pm 1,96\sqrt{\#Losses}}.$$

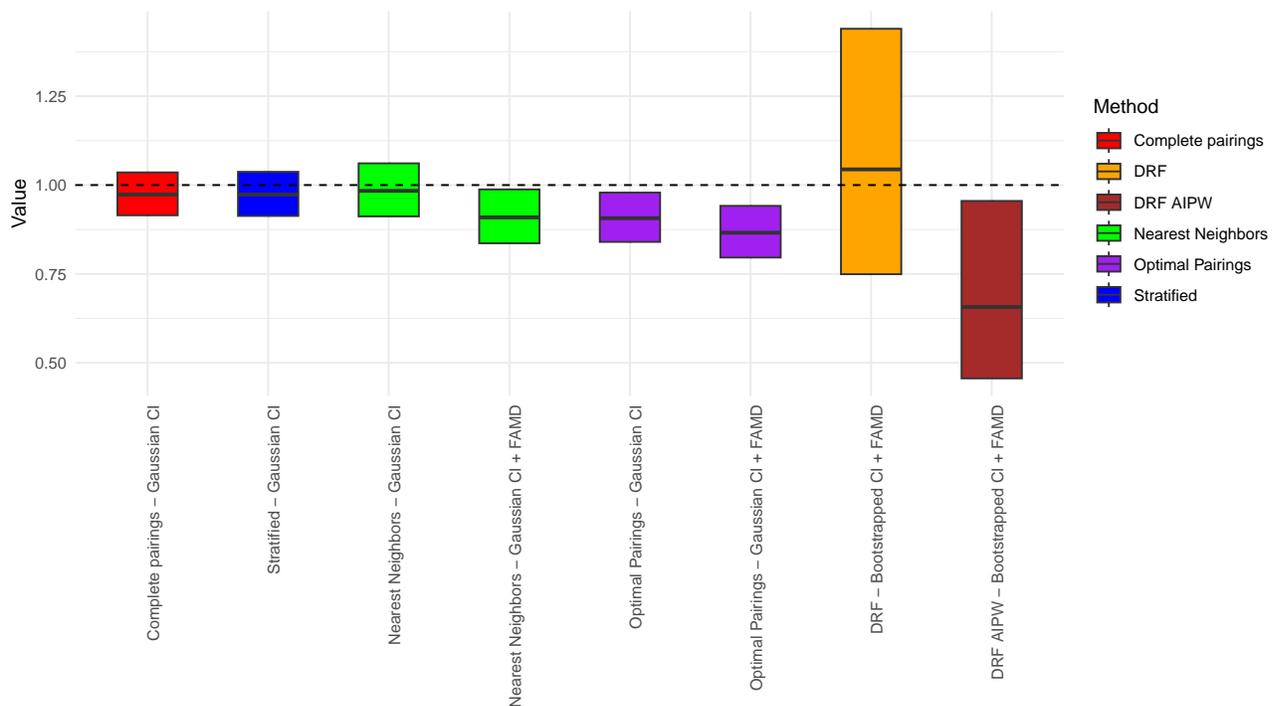


Figure 7: Win ratio computed on the CRASH-3 dataset using different methodologies.

### 6.3 Analysis of Figure 7.

The four pairing methods (complete pairings, stratified pairings according to the time since injury, nearest neighbors and optimal pairing) are compared in Figure 7.

**Win Ratio with complete and stratified pairings.** These two pairings do not show significant results: 1 is in the confidence intervals, meaning that the Win Ratio analysis with these methods cannot conclude for or against treatment.

**Nearest neighbors and optimal pairings on the covariates with the Mahalanobis distance.** Nearest neighbors have the same behavior if performed on the covariates. However, optimal pairings offer significant results against arm 1 in favor of placebo.

**Nearest neighbors and optimal pairings on the FAMD latent space.** However, when performed on the FAMD latent space, Win Ratio computed using nearest neighbors and optimal pairings are *both* significant (in favor of placebo). This suggests that computing matchings and distances in different ways may lead to slightly different results. We believe that in our case, due to mixed categorical and numerical features, the FAMD approach with distances in this latent space seems to be the right approach.

**Distributional regression approaches.** The boxplots of the two distributional regression approaches (DRF and AIPW DRF) are however much larger, reminiscent of the larger variability of forest based approaches. As such, results for the DRF approach are not significant. However, the augmented approach, AIPW DRF, offers significant results (in favor of placebo) despite its large variance, illustrating its double robustness properties.

**Conclusion from this study.** Our Win Ratio analysis of the CRASH-3 study suggests a preference over placebo rather than over treatment. This is not in contradiction with the conclusion of the CRASH-3 study, that were in favor of treatment only *on a subpopulation*. ATEs computed on the whole population are indeed not statistically significant, as showed above. This suggests is the strength of our methodologies, that can lead to significant discoveries, in scenarios where traditional or stratified Win Ratio fail to do so. Indeed, as shown in the simple synthetic example in Examples 1 and 2, in the presence of heterogeneity one must be very careful at which estimand is being targeted. Further investigations would be necessary to draw clinical conclusions from a Win Ratio analysis of the CRASH-3 dataset.

## 7 Conclusion and open directions

In this paper, we have introduced a causal inference framework for hierarchical outcome comparison methods like Win Ratio or Generalized Pairwise Comparisons. Our goal is to make such methods more grounded, by offering new perspectives and shedding light on the different causal effect measures that may be targeted when performing a Win Ratio or related analysis. In particular, we highlight the fact that if the population is heterogeneous, complete pairings (the historical and traditional way of forming pairs to compute the Win Ratio or the Net Benefit of a treatment) may target a population-level estimand that reverses treatment recommendations. The new causal effect measure  $\tau_\star$  we introduce in Definition 2 aims at answering this fallacy by taking into account covariate effects in the causal effect measure, thus being more robust to heterogeneous population. We stress the fact that this causal effect measure is related to stratified Win Ratio, since it can be estimated using an extreme form of stratification i.e., a Nearest Neighbors approach when forming pairs of treated-control patients. Sections 4 and 4.2 are then devoted to the estimation of our newly introduced estimand  $\tau_\star$ , in an effort to extend hierarchical outcome analyses and the Win Ratio methodology to observational settings and to handle missing covariates. We do so using a classical inverse propensity weighting approach to generalize our Nearest Neighbor pairing method, and using a less conventional distributional regression approach, that proves to be very efficient by leveraging recent Machine Learning tools such as distribution random forests.

Finally, our work paves the way to many open directions, among which we would like to highlight one: policy learning in the presence of hierarchical outcomes, an unexplored direction in the literature. A direct byproduct of our analysis and of Definition 2 is to define the value of a policy  $\pi : \mathcal{X} \rightarrow \{0, 1\}$  as:

$$V(\pi) \stackrel{\text{def}}{=} \mathbb{E} [\tau_\star^{(\pi(X_i))}(X_i)] ,$$

where

$$\tau_{\star}^{(t)}(x) \stackrel{\text{def}}{=} \mathbb{E} [w(Y^{(x)}(t)|Y_i(1-t))|X_i = x] .$$

An *optimal treatment rule* (OTR) is a policy that solves the following maximization problem:

$$\pi^{\star} \in \operatorname{argmax}_{\pi \in \Pi} V(\pi) ,$$

for  $\Pi$  a policy set. It then appears that for unconstrained policy estimation (i.e., when  $\Pi = \{0, 1\}^{\mathcal{X}}$ ), the OTR has a closed form expression that can easily be estimated using our developed tools:

$$\forall x \in \mathcal{X}, \quad \pi^{\star}(x) = \mathbb{1}_{\{\delta(x) > 0\}} ,$$

where:

$$\delta(x) = \mathbb{E} [w(Y^{(x)}(1)|Y_i(0))|X_i = x] - \mathbb{E} [w(Y^{(x)}(0)|Y_i(1))|X_i = x] .$$

## Acknowledgements

M. Even acknowledges funding from PEPR Santé Numérique SMATCH<sup>4</sup>. We would like to deeply thank Uri Shalit, Stefan Michiels and François Petit for very stimulating discussions in the writing of this paper, Dovid Parnas for suggesting Example 1 and for all the very the useful feedback, and Jeffrey Näf for the help in using the DRF package.

## References

drf: Distributional random forests, jun 2020. [36](#)

Ezimamaka Ajufo, Aditi Nayak, and Mandeep R. Mehra. Fallacies of using the win ratio in cardiovascular trials. *JACC: Basic to Translational Science*, 8(6):720–727, jun 2023. ISSN 2452-302X. [9](#), [14](#)

---

<sup>4</sup><https://pepr-santenum.fr/2023/11/07/smatch/>

Susan Athey and Stefan Wager. Policy learning with observational data. *Econometrica*, 89(1): 133–161, 2021. [28](#)

Mickael De Backer, Manju Sengar, Vikram Mathews, Samuel Salvaggio, Vaiva Deltuvaite-Thomas, Jean-Christophe Chiem, Everardo D Saad, and Marc Buyse. Design of a clinical trial using generalized pairwise comparisons to test a less intensive treatment regimen. *Clinical Trials*, 21(2):180–188, 2024. [8](#)

Clément Bénard, Jeffrey Näf, and Julie Josse. MMD-based variable importance for distributional random forest. In Sanjoy Dasgupta, Stephan Mandt, and Yingzhen Li, editors, *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pages 1324–1332. PMLR, 02–04 May 2024. [35](#), [39](#)

Gérard Biau and Luc Devroye. *Lectures on the nearest neighbor method*, volume 246. Springer, 2015. [26](#)

Matthias Boentert, Kenneth I Berger, Jordi Díaz-Manera, Mazen M Dimachkie, Alaa Hamed, Lionel Riou França, Nathan Thibault, Pragya Shukla, Jack Ishak, and J Jaime Caro. Applying the win ratio method in clinical trials of orphan drugs: an analysis of data from the comet trial of avalglucosidase alfa in patients with late-onset pompe disease. *Orphanet Journal of Rare Diseases*, 19(1):14, 2024. [9](#)

Marc Buyse. Generalized pairwise comparisons of prioritized outcomes in the two-sample problem. *Statistics in Medicine*, 29(30):3245–3257, dec 2010. ISSN 1097-0258. [1](#), [2](#), [3](#), [4](#), [8](#), [14](#), [15](#)

Domagoj Cevid, Loris Michel, Jeffrey Näf, Peter Bühlmann, and Nicolai Meinshausen. Distribu-

- tional random forests: Heterogeneity adjustment and multivariate distributional regression. *Journal of Machine Learning Research*, 23(333):1–79, 2022. [27](#), [33](#), [35](#), [36](#)
- Ruohui Chen, Tuo Lin, Lin Liu, Jinyuan Liu, Ruifeng Chen, Jingjing Zou, Chenyu Liu, Loki Natarajan, Wan Tang, Xinlian Zhang, and Xin Tu. A doubly robust estimator for the mann whitney wilcoxon rank sum test when applied for causal inference in observational studies. *Journal of Applied Statistics*, page 1–25, may 2024. ISSN 1360-0532. [6](#), [10](#), [18](#), [19](#), [29](#)
- Maria Vittoria Chiaruttini, Giulia Lorenzoni, Gaya Spolverato, and Dario Gregori. Win statistics in observational cancer research: Integrating clinical and quality-of-life outcomes. *Journal of Clinical Medicine*, 13(11):3272, may 2024. ISSN 2077-0383. [6](#), [10](#), [18](#), [19](#), [29](#)
- William G Cochran. Some methods for strengthening the common  $\chi^2$  tests. *Biometrics*, 10(4):417–451, 1954. [9](#)
- Bénédicte Colnet, Julie Josse, Gaël Varoquaux, and Erwan Scornet. Risk ratio, odds ratio, risk difference... which causal measure is easier to generalize?, 2024. [10](#), [19](#)
- The CRASH et al. Effects of tranexamic acid on death, disability, vascular occlusive events and other morbidities in patients with acute traumatic brain injury (crash-3): a randomised, placebo-controlled trial. *The Lancet*, 394(10210):1713–1723, 2019. [1](#), [8](#), [45](#), [46](#)
- Ying Cui and Bo Huang. Wins: The r wins package, dec 2021. [47](#)
- Jordi Diaz-Manera, Priya S Kishnani, Hani Kushlaf, Shafeeq Ladha, Tahseen Mozaffar, Volker Straub, Antonio Toscano, Ans T Van der Ploeg, Kenneth I Berger, Paula R Clemens, et al. Safety and efficacy of avalglucosidase alfa versus alglucosidase alfa in patients with late-onset pompe disease (comet): a phase 3, randomised, multicentre trial. *The Lancet Neurology*, 20(12):1012–1026, 2021. [9](#)

- Vanessa Didelez and Mats Julius Stensrud. On the logic of collapsibility for causal effect measures. *Biometrical Journal*, 64(2):235–242, feb 2021. ISSN 1521-4036. [11](#)
- Gaohong Dong, Junshan Qiu, Duolao Wang, and Marc Vandemeulebroecke. The stratified win ratio. *Journal of biopharmaceutical statistics*, 28(4):778–796, 2018. [4](#), [9](#), [47](#)
- Michael P Fay and Fan Li. Causal interpretation of the hazard ratio in randomized clinical trials. *Clinical Trials*, 21(5):623–635, apr 2024. ISSN 1740-7753. [10](#), [11](#), [18](#)
- Zijun Gao, Shu Ge, and Jian Qian. Bridging multiple worlds: multi-marginal optimal transport for causal partial-identification problem. *arXiv preprint arXiv:2406.07868*, 2024. [11](#)
- Rolf H.H. Groenwold, Karel G.M. Moons, Linda M. Peelen, Mirjam J. Knol, and Arno W. Hoes. Reporting of treatment effects from randomized trials: A plea for multivariable risk ratios. *Contemporary Clinical Trials*, 32(3):399–402, may 2011. ISSN 1551-7144. [11](#)
- Xiaohan Guo and Ai Ni. Contrast weighted learning for robust optimal treatment rule estimation. *Statistics in Medicine*, 41(27):5379–5394, sep 2022. ISSN 1097-0258. [6](#), [10](#), [18](#), [19](#), [29](#)
- Daniel G Horvitz and Donovan J Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American statistical Association*, 47(260):663–685, 1952. [28](#)
- Miloslav Jiřina. On regular conditional probabilities. *Czechoslovak Mathematical Journal*, 09(3):445–451, 1959. [21](#)
- Sébastien Lê, Julie Josse, and François Husson. Factominer: an r package for multivariate analysis. *Journal of statistical software*, 25:1–18, 2008. [48](#)

- Zhenhua Lin, Dehan Kong, and Linbo Wang. Causal inference on distribution functions. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 85(2):378–398, 2023. [11](#)
- H. B. Mann and D. R. Whitney. On a test of whether one of two random variables is stochastically larger than the other. *The Annals of Mathematical Statistics*, 18(1):50–60, mar 1947. ISSN 0003-4851. [8](#), [10](#)
- Nathan Mantel and William Haenszel. Statistical aspects of the analysis of data from retrospective studies of disease. *Journal of the national cancer institute*, 22(4):719–748, 1959. [9](#)
- Lu Mao. On causal estimation using  $u$ -statistics. *Biometrika*, 105(1):215–220, dec 2017. ISSN 1464-3510. [6](#), [10](#), [18](#), [19](#), [29](#)
- Lu Mao. Defining estimand for the win ratio: Separate the true effect from censoring. *Clinical Trials*, 21(5):584–594, jul 2024. ISSN 1740-7753. [9](#)
- Judea Pearl. *Causality*. Cambridge university press, 2009. [19](#)
- S. J. Pocock, C. A. Ariti, T. J. Collier, and D. Wang. The win ratio: a new approach to the analysis of composite endpoints in clinical trials based on clinical priorities. *European Heart Journal*, 33(2):176–182, sep 2011. ISSN 1522-9645. [1](#), [2](#), [3](#), [4](#), [8](#), [13](#), [14](#), [15](#), [38](#)
- Stuart J. Pocock, João Pedro Ferreira, Timothy J. Collier, Christiane E. Angermann, Jan Biegus, Sean P. Collins, Mikhail Kosiborod, Michael E. Nassif, Piotr Ponikowski, Mitchell A. Psotka, John R. Teerlink, Jasper Tromp, John Gregson, Jonathan P. Blatchford, Cordula Zeller, and Adriaan A. Voors. The win ratio method in heart failure trials: lessons learnt from <sc>empulse</sc>. *European Journal of Heart Failure*, 25(5):632–641, apr 2023. ISSN 1879-0844. [8](#)

- Björn Redfors, John Gregson, Aaron Crowley, Thomas McAndrew, Ori Ben-Yehuda, Gregg W Stone, and Stuart J Pocock. The win ratio approach for composite endpoints: practical guidance based on previous experience. *European Heart Journal*, 41(46):4391–4399, sep 2020. ISSN 1522-9645. [2](#), [9](#)
- James M Robins, Andrea Rotnitzky, and Lue Ping Zhao. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American statistical Association*, 89(427):846–866, 1994. [28](#)
- Edward H Simpson. The interpretation of interaction in contingency tables. *Journal of the Royal Statistical Society: Series B (Methodological)*, 13(2):238–241, 1951. [22](#)
- Jerzy Splawa-Neyman, Dorota M Dabrowska, and Terrence P Speed. On the application of probability theory to agricultural experiments. essay on principles. section 9. *Statistical Science*, pages 465–472, 1990. [5](#), [6](#), [12](#)
- Stef Van Buuren and Karin Groothuis-Oudshoorn. mice: Multivariate imputation by chained equations in r. *Journal of statistical software*, 45:1–67, 2011. [45](#)
- Clifford H Wagner. Simpson’s paradox in real life. *The American Statistician*, 36(1):46–48, 1982. [22](#)
- Frank Wilcoxon. Individual comparisons by ranking methods. *Biometrics Bulletin*, 1(6):80, dec 1945. ISSN 0099-4987. [8](#), [10](#)
- Anqi Yin, Ao Yuan, and Ming T. Tan. Highly robust causal semiparametric u-statistic with applications in biomedical studies. *The International Journal of Biostatistics*, 20(1):69–91, nov 2022. ISSN 1557-4679. [6](#), [10](#), [18](#), [19](#), [29](#)
- Di Zhang and Jong-Hyeon Jeong. Inference on win ratio for cluster-randomized semi-competing

risk data. *Japanese Journal of Statistics and Data Science*, 4:1263–1292, jun 2021. ISSN 2520-8764. [10](#)

Di Zhang, Stephen R. Wisniewski, and Jong-Hyeon Jeong. Causal inference on win ratio for observational data with dependent subjects, 2022. [6](#), [10](#), [18](#), [19](#), [29](#)

# A Proof of Theorem 1

## A.1 Proof of Theorem 1.2

*Proof of Theorem 1.2.* Let  $p = \hat{p}_W^{(n_0, n_1, \mathcal{C}_{\text{nn}})} = \frac{1}{n_0} \sum_{i \in \mathcal{N}_0} w(Y_{\sigma^*(i)} | Y_i)$  and  $\bar{p} = \mathbb{E} [w(Y^{(X_i)}(1) | Y_i(0))]$ .

We have, where  $X_{\mathcal{N}_1} = (X_k)_{k \in \mathcal{N}_1}$ :

$$\begin{aligned} \bar{p} - p &= \mathbb{E} \left[ \underbrace{\frac{1}{n_0} \sum_{i \in \mathcal{N}_0} \left\{ w(Y_i^{(X_i)}(1) | Y_i(0)) - w(Y_{\sigma^*(i)} | Y_i) \right\}}_{A_1} \middle| X_{\mathcal{N}_1} \right] \\ &\quad + \underbrace{\frac{1}{n_0} \sum_{i \in \mathcal{N}_0} \left\{ w(Y_{\sigma^*(i)} | Y_i) - \mathbb{E} [w(Y_{\sigma^*(i)} | Y_i) | i \in \mathcal{N}_0, X_{\mathcal{N}_1}] \right\}}_{A_2} \end{aligned}$$

**Control of  $A_2$ .** The term  $A_2$  is controlled by computing variance. Let

$$a_i = w(Y_{\sigma^*(i)} | Y_i) - \mathbb{E} [w(Y_{\sigma^*(i)} | Y_i) | i \in \mathcal{N}_0, X_{\mathcal{N}_1}] ,$$

and note that we have  $\mathbb{E} [a_i | i \in \mathcal{N}_0, X_{\mathcal{N}_1}] = 0$ . Let

$$p_k = \mathbb{P} (\sigma(i) = k | X_{\mathcal{N}_1})$$

be the (random) weights of the (random) Voronoi cells associated to elements of  $X_{\mathcal{N}_1}$ . We have,

where  $i \neq j \in \mathcal{N}_0$  are arbitrary (note that conditioned on  $X_{\mathcal{N}_1}$ ,  $\mathcal{N}_0$  is fixed):

$$\begin{aligned} \text{var} (A_2 | X_{\mathcal{N}_1}) &= \frac{\text{var} (a_i | X_{\mathcal{N}_1})}{n_0} \\ &\quad + \frac{n_0 - 1}{n_0} \left( \mathbb{E} [a_i a_j | \sigma(i) = \sigma(j), X_{\mathcal{N}_1}] \mathbb{P} (\sigma(i) = \sigma(j) | X_{\mathcal{N}_1}) \right. \\ &\quad \left. + \mathbb{E} [a_i a_j | \sigma(i) \neq \sigma(j), X_{\mathcal{N}_1}] \mathbb{P} (\sigma(i) \neq \sigma(j) | X_{\mathcal{N}_1}) \right) \\ &\leq \frac{1}{n_0} + \mathbb{P} (\sigma(i) = \sigma(j) | X_{\mathcal{N}_1}) + \frac{n_0 - 1}{n_0} \mathbb{E} [a_i a_j | \sigma(i) \neq \sigma(j), X_{\mathcal{N}_1}] \mathbb{P} (\sigma(i) \neq \sigma(j) | X_{\mathcal{N}_1}) . \end{aligned}$$

Conditionnally on  $X_{\mathcal{N}_1}$ ,  $\sigma(i)$  and  $\sigma(j)$  are independent random variables that assign  $k$  with probability  $p_k$ . Thus, we can prove that  $a_i, a_j$  are negatively correlated conditionally on  $\sigma(i) \neq$

$\sigma(j)$ :

$$\begin{aligned}
\mathbb{E} [a_i a_j \mathbf{1}_{\{\sigma(i) \neq \sigma(j)\}} | X_{\mathcal{N}_1}] &= \sum_{k \neq \ell \in \mathcal{N}_1} p_k p_\ell \mathbb{E} [a_i a_j | X_{\mathcal{N}_1}, \sigma(i) = k, \sigma(j) = \ell] \\
&= \sum_{k \neq \ell \in \mathcal{N}_1} p_k p_\ell \mathbb{E} [a_i | X_{\mathcal{N}_1}, \sigma(i) = k, \sigma(j) = \ell] \mathbb{E} [a_j | X_{\mathcal{N}_1}, \sigma(i) = k, \sigma(j) = \ell] \\
&\quad \text{since } a_i \perp\!\!\!\perp a_j | X_{\mathcal{N}_1}, \sigma(i) = k, \sigma(j) = \ell \\
&= \sum_{k \neq \ell \in \mathcal{N}_1} p_k p_\ell \mathbb{E} [a_i | X_{\mathcal{N}_1}, \sigma(i) = k] \mathbb{E} [a_j | X_{\mathcal{N}_1}, \sigma(j) = \ell] \\
&= \sum_{k, \ell \in \mathcal{N}_1} p_k p_\ell \mathbb{E} [a_i | X_{\mathcal{N}_1}, \sigma(i) = k] \mathbb{E} [a_j | X_{\mathcal{N}_1}, \sigma(j) = \ell] - \sum_{k \in \mathcal{N}_1} p_k^2 \mathbb{E} [a_i | X_{\mathcal{N}_1}, \sigma(i) = k]^2 \\
&\leq \left( \sum_{k \in \mathcal{N}_1} p_k \mathbb{E} [a_i | X_{\mathcal{N}_1}, \sigma(i) = k] \right)^2 \\
&= 0.
\end{aligned}$$

Using  $\mathbb{E} [a_i a_j \mathbf{1}_{\{\sigma(i) \neq \sigma(j)\}} | X_{\mathcal{N}_1}] = \mathbb{E} [a_i a_j | X_{\mathcal{N}_1}, \sigma(i) \neq \sigma(j)] \mathbb{P}(\sigma(i) \neq \sigma(j))$ ; we thus have that:

$$\mathbb{E} [a_i a_j | X_{\mathcal{N}_1}, \sigma(i) \neq \sigma(j)] \leq 0.$$

Thus, we have that  $\text{var}(A_2 | X_{\mathcal{N}_1}) \leq \frac{1}{n_0} + \mathbb{P}(\sigma(i) = \sigma(j))$ , and the last step of this first part of the proof is to show that  $\mathbb{P}(\sigma(i) = \sigma(j)) \rightarrow 0$ , the purpose of the following lemma.

**Lemma 1.** *We have, for  $i \neq j \in \mathcal{N}_0$ :*

$$\mathbb{P}(\sigma(i) = \sigma(j)) \rightarrow 0.$$

*Proof of Lemma 2.* We have:

$$\mathbb{P}(\sigma(i) = \sigma(j)) = \frac{1}{n_1} \sum_{k \in \mathcal{N}_1} \mathbb{P}(\sigma(i) = k | \sigma(j) = k).$$

Let  $k \in \mathcal{N}_1$  and  $x \in \text{Supp}(X)$ :  $\forall \varepsilon > 0, \mathbb{P}(X \in \mathcal{B}(x, \varepsilon)) > 0$ .

*First case:*  $\mathbb{P}(X = x) = p_x > 0$ . In that case, let  $N_x = |\{\ell \in \mathcal{N}_1, X_\ell = x\}|$ . We have that

$$\mathbb{P}(\sigma(i) = k | N_x, X_k = x, \sigma(j) = k) = \frac{1}{N_x},$$

and  $N_x$  is a binomial random variable of parameters  $(n_1, p_x)$ . This leads to:

$$\begin{aligned} \mathbb{P}(\sigma(i) = k | X_k = x, \sigma(j) = k) &= \sum_{N=1}^{n_1} 2^{-n_1} \frac{p_x^N (1-p_x)^{n_1-N}}{N} \binom{n_1}{N} \\ &= \sum_{N=1}^{n_1} \frac{p_x^N (1-p_x)^{n_1-N}}{n_1+1} \binom{n_1+1}{N+1} \\ &\leq \frac{1}{p_x(n_1+1)}. \end{aligned}$$

Thus,  $\mathbb{P}(\sigma(i) = k | X_k = x, \sigma(j) = k) \rightarrow 0$  as  $n_1 \rightarrow \infty$ .

*Second case:*  $\mathbb{P}(X = x) = 0$  (no Dirac mass). Let  $\delta \in (0, 1)$  and let  $R > \varepsilon > 0$  such that  $\mathbb{P}(X \in \mathcal{B}(x, \varepsilon)) < \delta$  and  $\mathbb{P}(X \in \mathcal{B}(x, R)) > 1 - \delta$ . We cover  $\mathcal{B}(0, R) \setminus \mathcal{B}(0, \varepsilon)$  with  $m$  balls of radius  $\varepsilon/2$ :  $\mathcal{B}(0, R) \setminus \mathcal{B}(0, \varepsilon) \subset \bigcup_{r=1}^m \mathcal{B}(z_r, \varepsilon/2)$ , where  $z_r \in \mathcal{B}(0, R) \setminus \mathcal{B}(0, \varepsilon)$ . We remove all  $z_r$  that satisfy  $\mathbb{P}(X \in \mathcal{B}(z_r, \varepsilon/2)) = 0$  from this union. Let  $\mathcal{E}$  be the event  $\{\forall r \in [m], \exists \ell \in \mathcal{N}_1 \setminus \{k\}, X_\ell \in \mathcal{B}(z_r, \varepsilon/2)\}$ .

We have that

$$\begin{aligned} \mathbb{P}(\sigma(i) = k | X_k = x, \sigma(j) = k, \mathcal{E}) &\leq \mathbb{P}(X_i \in \mathcal{B}(x, \varepsilon)) \\ &\leq \delta. \end{aligned}$$

Then,

$$\begin{aligned} \mathbb{P}(\mathcal{E}^C) &\leq \sum_{r=1}^m \mathbb{P}(\forall \ell \in \mathcal{N}_1 \setminus \{k\}, X_\ell \notin \mathcal{B}(z_r, \varepsilon/2)) \\ &\leq m(1 - p_{\min})^{n_1-1}, \end{aligned}$$

where  $p_{\min} = \min_{r \in [m]} \mathbb{P}(X \in \mathcal{B}(z_r, \varepsilon/2))$ . Thus,  $\mathbb{P}(\mathcal{E}) \rightarrow 1$ , and  $\mathbb{P}(\sigma(i) = k | X_k = x, \sigma(j) = k) \leq 1 - \mathbb{P}(\mathcal{E}) + \delta$ . We can thus conclude that  $\mathbb{P}(\sigma(i) = k | X_k = x, \sigma(j) = k) \rightarrow 0$  as  $n_1 \rightarrow \infty$ .

*Wrapping things up.* Using  $\mathbb{P}(\sigma(i) = k | \sigma(j) = k) = \int_{\mathcal{X}} \mathbb{P}(\sigma(i) = k | \sigma(j) = k, X_k = x) d\mathbb{P}(X_k = x)$ , we have  $\mathbb{P}(\sigma(i) = k | \sigma(j) = k) \rightarrow 0$  as  $n_1 \rightarrow \infty$ , using dominated convergence.  $\square$

Using this, we have  $\text{var}(A_2) \rightarrow 0$ , leading to  $A_2 \rightarrow 0$  in probability.

**Control of  $A_1$ .** We now control  $A_1$ , using the continuity assumption. Using unconfoundedness:

$$|A_1| \leq \mathbb{E} [\delta(X_i, X_{\sigma^*(i)}, Y_i(0)) | X_{\mathcal{N}_1}, T_i = 0], \quad \text{where}$$

$$\delta(x, x', y) \stackrel{\text{def}}{=} \left| \mathbb{E} [w(Y_i^{(X_i)}(1)|y) | X_i = x] - \mathbb{E} [w(Y_j(1)|y) | X_j = x'] \right|.$$

Let  $\varepsilon > 0$  and  $y$  fixed. Using our continuity and compactness assumptions,  $x, x' \mapsto \delta(x, x', y)$  is uniformly continuous on  $\mathcal{X} \times \mathcal{X}$ , so that there exists  $\eta > 0$  such that if  $\|x - x'\| \leq \eta$ , we have  $\delta(x, x', y) \leq \varepsilon$ . We are going to show that with high probability,  $\|X_i - X_{\sigma^*(i)}\| \leq \eta$ . Using compactness of  $\mathcal{X}$ , there exist  $u_1, \dots, u_p \in \mathcal{X}$  such that  $\mathcal{X} \subset \bigcup_{k=1}^p \mathcal{B}(u_k, \eta/2)$ . Let  $p_k = \mathbb{P}(X_i \in \mathcal{B}(u_k, \eta/2))$ : we assume that  $p_k > 0$  for all  $k$ , otherwise we remove this ball and the corresponding  $u_k$ . Let  $p_{\min} = \min_k p_k > 0$ . Let  $k_x \in [p]$  such that  $X_i \in \mathcal{B}(u_{k_x}, \eta/2)$ . We have, working conditionnally on  $\mathcal{N}_0, \mathcal{N}_1, i \in \mathcal{N}_0$ :

$$\begin{aligned} \mathbb{P}(\|X_i - X_{\sigma^*(i)}\| > \eta) &\leq \mathbb{E} [\mathbb{P}(X_{\sigma^*(i)} \notin \mathcal{B}(u_{k_x}, \eta/2) | k_x)] \\ &= \mathbb{E} [\mathbb{P}(\forall j \in \mathcal{N}_1, X_j \notin \mathcal{B}(u_{k_x}, \eta/2) | k_x)] \\ &= \mathbb{E} [(1 - p_{k_x})^{n_1}] \\ &\leq (1 - p_{\min})^{n_1} \\ &\xrightarrow[n_1 \rightarrow \infty]{} 0. \end{aligned}$$

This leads to:

$$\mathbb{P}(\delta(X_i, X_{\sigma(i)}, y) > \varepsilon) \leq (1 - p_{\min})^{n_1},$$

and thus  $\mathbb{P}(\delta(X_i, X_{\sigma(i)}, y) \rightarrow 0) = 1$  as  $n_1 \rightarrow \infty$ , leading to  $\mathbb{E}[\delta(X_i, X_{\sigma(i)}, Y_i(0))] \rightarrow 0$ . We thus have that  $\mathbb{E}[|A_1|] \rightarrow 0$ , and thus  $A_1 \rightarrow 0$  in probability, since  $|A_1| \leq 1$  almost surely. This concludes the proof.  $\square$

## A.2 Proof of Theorem 1.1

*Proof of Theorem 1.1.* Let  $\bar{p} = \mathbb{E}[w(Y_i(1)|Y_j(0))]$  for  $i \neq j$ . We now prove the second point, with complete pairings. Using our assumptions, we have that

$$\mathbb{E}[w(Y_i|Y_j)|T_i = 1, T_j = 0] = \mathbb{E}[w(Y_j(1)|Y_i(0))] \stackrel{\text{def}}{=} \bar{p},$$

so that

$$\mathbb{E}\left[\hat{p}_W^{(n_0, n_1, \mathcal{C}_{\text{Tot}})}\right] = \bar{p}.$$

Since  $0 \leq w \leq 1$ ,  $\text{var}(w(Y_j|Y_i)) \leq 1$ , leading to:

$$\begin{aligned} \text{var}\left(\hat{p}_W^{(n_0, n_1, \mathcal{C}_{\text{Tot}})}\right) &= \frac{1}{n_0^2 n_1^2} \sum_{i, i' \in \mathcal{N}_0, j, j' \in \mathcal{N}_1} \mathbb{E}[(w(Y_j|Y_i) - \bar{p})(w(Y_{j'}|Y_{i'}))] \\ &= \frac{1}{n_0^2 n_1^2} \sum_{i \in \mathcal{N}_0, j \in \mathcal{N}_1} \mathbb{E}[(w(Y_j|Y_i) - \bar{p})^2] \\ &\quad + \underbrace{\frac{1}{n_0^2 n_1^2} \sum_{i \neq i' \in \mathcal{N}_0, j, j' \neq \mathcal{N}_1} \mathbb{E}[(w(Y_j|Y_i) - \bar{p})(w(Y_{j'}|Y_{i'}))]}_{=0 \text{ (independence)}} \\ &\quad + \frac{1}{n_0^2 n_1^2} \sum_{i \in \mathcal{N}_0, j \neq j' \in \mathcal{N}_1} \mathbb{E}[(w(Y_j|Y_i) - \bar{p})(w(Y_{j'}|Y_{i'}))] \\ &\quad + \frac{1}{n_0^2 n_1^2} \sum_{i \neq i' \in \mathcal{N}_0, j \in \mathcal{N}_1} \mathbb{E}[(w(Y_j|Y_i) - \bar{p})(w(Y_{j'}|Y_{i'}))] \\ &\leq \frac{1}{n_0 n_1} + \frac{1}{n_0} + \frac{1}{n_1}. \end{aligned}$$

Thus,  $\hat{p}_W^{(n_0, n_1, \mathcal{C}_{\text{Tot}})} \rightarrow \bar{p}$  in probability as  $n_0, n_1 \rightarrow \infty$ . □

## B Proof of Theorem 2

*Proof.* Let

$$\hat{p} \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i \in \mathcal{N}_0} w(Y_{\sigma_1^*(i)}|Y_i)(1 - \hat{\pi}(X_i))^{-1}$$

be the IPW estimator,

$$\hat{p}^* \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i \in \mathcal{N}_0} w(Y_{\sigma_1^*(i)}|Y_i)(1 - \pi(X_i))^{-1}$$

be the IPW estimator with oracle propensities, and  $p = \tau_\star$  be the targeted estimand. We have:

$$\hat{p} - p = (\hat{p} - \hat{p}^\star) + (\hat{p}^\star - p).$$

For this first term,

$$\begin{aligned} |\hat{p} - \hat{p}^\star| &= \left| \frac{1}{n} \sum_{i=1}^n (1 - T_i) w(Y_{\sigma_1^\star(i)} | Y_i) \{ (1 - \hat{\pi}(X_i))^{-1} - (1 - \pi(X_i))^{-1} \} \right| \\ &\leq \frac{1}{n} \sum_{i=1}^n (1 - T_i) w(Y_{\sigma_1^\star(i)} | Y_i) |(1 - \hat{\pi}(X_i))^{-1} - (1 - \pi(X_i))^{-1}| \\ &\leq \frac{1}{n} \sum_{i=1}^n |(1 - \hat{\pi}(X_i))^{-1} - (1 - \pi(X_i))^{-1}| \\ &\leq \frac{1}{n} \sum_{i=1}^n \{ \mathbb{1}_{\{\hat{\pi}(X_i) < 1 - \eta'\}} |(1 - \hat{\pi}(X_i))^{-1} - (1 - \pi(X_i))^{-1}| \\ &\quad + \mathbb{1}_{\{\hat{\pi}(X_i) \geq 1 - \eta'\}} |(1 - \hat{\pi}(X_i))^{-1} - (1 - \pi(X_i))^{-1}| \}, \end{aligned}$$

for some given  $0 < \eta' < \eta$ . First, notice that  $u \mapsto (1 - u)^{-2}$  is  $\eta'^{-2}$ -Lipschitz on  $[0, 1 - \eta']$ , so that:

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{\{\hat{\pi}(X_i) < 1 - \eta'\}} |(1 - \hat{\pi}(X_i))^{-1} - (1 - \pi(X_i))^{-1}| &\leq \frac{1}{n\eta'^2} \sum_{i=1}^n \mathbb{1}_{\{\hat{\pi}(X_i) < 1 - \eta'\}} |\hat{\pi}(X_i) - \pi(X_i)| \\ &\leq \frac{1}{n\eta'^2} \sum_{i=1}^n |\hat{\pi}(X_i) - \pi(X_i)| \\ &= \eta'^{-2} \mathbb{E} [|\hat{\pi}(X_i) - \pi(X_i)|] \\ &\quad + \frac{1}{n\eta'^2} \sum_{i=1}^n \{ |\hat{\pi}(X_i) - \pi(X_i)| - \mathbb{E} [|\hat{\pi}(X_i) - \pi(X_i)|] \}. \end{aligned}$$

Here, we have that  $\mathbb{E} [|\hat{\pi}(X_i) - \pi(X_i)|] \rightarrow 0$  using mean consistency, while the second term converges to 0 in probability (sum of  $n$  centered bounded random variables). Then, using Cauchy-Schwarz inequality,

$$\begin{aligned} &\frac{1}{n} \sum_{i=1}^n \mathbb{1}_{\{\hat{\pi}(X_i) \geq 1 - \eta'\}} |(1 - \hat{\pi}(X_i))^{-1} - (1 - \pi(X_i))^{-1}| \\ &\leq \sqrt{\frac{1}{n} \sum_{i=1}^n \mathbb{1}_{\{\hat{\pi}(X_i) \geq 1 - \eta'\}} \times \frac{1}{n} \sum_{i=1}^n ((1 - \hat{\pi}(X_i))^{-1} - (1 - \pi(X_i))^{-1})^2}. \end{aligned}$$

The first factor in the square root satisfies

$$\frac{1}{n} \sum_{i=1}^n \mathbb{1}_{\{\hat{\pi}(X_i) \geq 1 - \eta'\}} = \mathbb{P}(\hat{\pi}(X_i) \geq 1 - \eta') + \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{\{\hat{\pi}(X_i) \geq 1 - \eta'\}} - \mathbb{P}(\hat{\pi}(X_i) \geq 1 - \eta').$$

The first term is deterministic and converges to zero almost surely thanks to pointwise convergence and dominated convergence, while the second term converges to zero as the averaged sum of  $n$  independent, centered and bounded random variables. All this leads to  $\hat{p} - \hat{p}^* \rightarrow 0$  in probability and almost surely.

For the second term  $\hat{p}^* - p$ , we adapt the proof of Theorem 1.2. We extend the definition of  $\sigma^*$  to  $\mathcal{N}_1$ : for  $i \in \mathcal{N}_1$  we have  $\sigma^*(i) = i$ . We have, using that  $p = \mathbb{E}[w(Y_i(1), Y_i(0))] = \mathbb{E}\left[\frac{1-T_i}{1-\pi(X_i)}w(Y_i(1), Y_i(0))\right]$  with unconfoundedness:

$$\begin{aligned} p - \hat{p}^* &= \mathbb{E} \left[ \underbrace{\frac{1}{n} \sum_{i \in \mathcal{N}_0} \frac{w(Y^{(X_i)}(1)|Y_i(0)) - w(Y_{\sigma^*(i)}|Y_i)}{1 - \pi(X_i)} \Big| X_{\mathcal{N}_1}}_{A_1} \right] \\ &\quad - \underbrace{\frac{1}{n} \sum_{i=1}^n \frac{w(Y_{\sigma^*(i)}|Y_i)(1 - T_i)}{1 - \pi(X_i)} - \mathbb{E} \left[ \frac{w(Y_{\sigma^*(i)}|Y_i)(1 - T_i)}{1 - \pi(X_i)} \Big| X_{\mathcal{N}_1} \right]}_{A_2}. \end{aligned}$$

**Control of  $A_2$ .** The term  $A_2$  is controlled computing its variance, as in the proof of Theorem 1.2.

Let  $a_i = \frac{w(Y_{\sigma^*(i)}|Y_i)(1 - T_i)}{1 - \pi(X_i)} - \mathbb{E}[w(Y_{\sigma^*(i)}|Y_i)|X_{\mathcal{N}_1}]$ . Since  $\mathbb{E}[a_i] = 0$ , we have  $\mathbb{E}[A_2|X_{\mathcal{N}_1}] = 0$ . Then, using overlap,  $|a_i| \leq 1/\eta$  almost surely, so that  $\mathbb{E}[a_i^2|\mathcal{N}_1] \leq 1/\eta^2$ .

Let  $p_k = \mathbb{P}(\sigma(i) = k|X_{\mathcal{N}_1})$  be the (random) weights of the (random) Voronoi cells associated to  $X_{\mathcal{N}_1}$ . We have:

$$\begin{aligned} \text{var}(A_2|X_{\mathcal{N}_1}) &= \frac{1}{n^2} \sum_{i=1}^n \text{var}(a_i|X_{\mathcal{N}_1}) \\ &\quad + \frac{1}{n^2} \sum_{i \neq j} \left( \mathbb{E}[a_i a_j | \sigma(i) = \sigma(j), X_{\mathcal{N}_1}] \mathbb{P}(\sigma(i) = \sigma(j)|X_{\mathcal{N}_1}) \right. \\ &\quad \left. + \mathbb{E}[a_i a_j | \sigma(i) \neq \sigma(j), X_{\mathcal{N}_1}] \mathbb{P}(\sigma(i) \neq \sigma(j)|X_{\mathcal{N}_1}) \right) \\ &\leq \frac{1}{\eta^2 n} + \frac{1}{n^2} \sum_{i \neq j} \mathbb{P}(\sigma(i) = \sigma(j)|X_{\mathcal{N}_1}) + \frac{n-1}{n} \mathbb{E}[a_i a_j | \sigma(i) \neq \sigma(j), X_{\mathcal{N}_1}] \mathbb{P}(\sigma(i) \neq \sigma(j)|X_{\mathcal{N}_1}) \end{aligned}$$

Conditionnally on  $X_{\mathcal{N}_1}$  and on  $i, j \in \mathcal{N}_0$ ,  $\sigma(i), \sigma(j)$  are independent random variables that assign  $k$  with probability  $p_k$ . Thus, we can prove that  $a_i, a_j$  are negatively correlated conditionally on  $\sigma(i) \neq \sigma(j)$ :

$$\begin{aligned}
\mathbb{E} [a_i a_j \mathbb{1}_{\{\sigma(i) \neq \sigma(j)\}} | X_{\mathcal{N}_1}] &= \sum_{k \neq \ell \in \mathcal{N}_1} p_k p_\ell \mathbb{E} [a_i a_j | X_{\mathcal{N}_1}, \sigma(i) = k, \sigma(j) = \ell] \\
&= \sum_{k \neq \ell \in \mathcal{N}_1} p_k p_\ell \mathbb{E} [a_i | X_{\mathcal{N}_1}, \sigma(i) = k, \sigma(j) = \ell] \mathbb{E} [a_j | X_{\mathcal{N}_1}, \sigma(i) = k, \sigma(j) = \ell] \\
&\quad \text{since } a_i \perp\!\!\!\perp a_j | X_{\mathcal{N}_1}, \sigma(i) = k, \sigma(j) = \ell \\
&= \sum_{k \neq \ell \in \mathcal{N}_1} p_k p_\ell \mathbb{E} [a_i | X_{\mathcal{N}_1}, \sigma(i) = k] \mathbb{E} [a_j | X_{\mathcal{N}_1}, \sigma(j) = \ell] \\
&= \sum_{k, \ell \in \mathcal{N}_1} p_k p_\ell \mathbb{E} [a_i | X_{\mathcal{N}_1}, \sigma(i) = k] \mathbb{E} [a_j | X_{\mathcal{N}_1}, \sigma(j) = \ell] - \sum_{k \in \mathcal{N}_1} p_k^2 \mathbb{E} [a_i | X_{\mathcal{N}_1}, \sigma(i) = k]^2 \\
&\leq \left( \sum_{k \in \mathcal{N}_1} p_k \mathbb{E} [a_i | X_{\mathcal{N}_1}, \sigma(i) = k] \right)^2 \\
&= 0.
\end{aligned}$$

Thus, we have that  $\text{var}(A_2) \leq \frac{1}{n_0} + \mathbb{P}(\sigma(i) = \sigma(j) | i, j \in \mathcal{N}_0)$  where  $i \neq j$ .

**Lemma 2.** *We have, for  $i \neq j$ :*

$$\mathbb{P}(\sigma(i) = \sigma(j) | i, j \in \mathcal{N}_0) \rightarrow 0.$$

*Proof of Lemma 2.* First, note that we have, for any measurable set  $\mathcal{S} \subset \mathcal{X}$ :

$$\frac{\mathbb{P}(X_i \in \mathcal{S} | T_i = 1)}{\mathbb{P}(X_i \in \mathcal{S} | T_i = 0)} \in \left[ \eta^2, \frac{1}{\eta^2} \right].$$

Indeed, we have for  $t = 0, 1$  and measurable set  $\mathcal{S} \subset \mathcal{X}$  such that  $\mathbb{P}(X_i \in \mathcal{S}) > 0$ :

$$\mathbb{P}(X_i \in \mathcal{S} | T_i = t) = \frac{\mathbb{P}(T_i = t | X_i \in \mathcal{S})}{\mathbb{P}(T_i = t)} \times \mathbb{P}(X_i \in \mathcal{S}).$$

Using overlap, this leads to  $\frac{\mathbb{P}(X_i \in \mathcal{S} | T_i = 1)}{\mathbb{P}(X_i \in \mathcal{S} | T_i = 0)} \in \left[ \eta^2, \frac{1}{\eta^2} \right]$ . A consequence is that for all measurable  $\mathcal{S} \subset \mathcal{X}$ , we have  $\mathbb{P}(X \in \mathcal{S}) = 0 \iff \mathbb{P}(X \in \mathcal{S} | T = 0) = 0 \iff \mathbb{P}(X \in \mathcal{S} | T = 1) = 0$ .

We have  $\mathbb{P}(\sigma(i) = \sigma(j)) = \frac{1}{n_1} \sum_{k \in \mathcal{N}_1} \mathbb{P}(\sigma(i) = k | \sigma(j) = k)$ . We work conditionally on  $\mathcal{N}_1, \mathcal{N}_0$ . Note that we have  $n_0, n_1 \rightarrow \infty$  almost surely as  $n \rightarrow \infty$ . Let  $k \in \mathcal{N}_1$  and  $x \in \text{Supp}(X)$ :  $\forall \varepsilon > 0, \mathbb{P}(X \in \mathcal{B}(x, \varepsilon)) > 0$ .

*First case:*  $\mathbb{P}(X = x | T = 1) = p_x > 0$ . In that case, let  $N_x = |\{\ell \in \mathcal{N}_1, X_\ell = x\}|$ . We have that

$$\mathbb{P}(\sigma(i) = k | N_x, X_k = x, \sigma(j) = k) = \frac{1}{N_x},$$

and  $N_x$  is a binomial random variable of parameters  $(n_1, p_x)$ . This leads to:

$$\begin{aligned} \mathbb{P}(\sigma(i) = k | X_k = x, \sigma(j) = k) &= \sum_{N=1}^{n_1} 2^{-n_1} \frac{p_x^N (1-p_x)^{n_1-N}}{N} \binom{n_1}{N} \\ &= \sum_{N=1}^{n_1} \frac{p_x^N (1-p_x)^{n_1-N}}{n_1+1} \binom{n_1+1}{N+1} \\ &\leq \frac{1}{p_x(n_1+1)}. \end{aligned}$$

Thus,  $\mathbb{P}(\sigma(i) = k | X_k = x, \sigma(j) = k) \rightarrow 0$  as  $n_1 \rightarrow \infty$ .

*Second case:*  $\mathbb{P}(X = x) = 0$ . Let  $\delta \in (0, 1)$ . Let  $R > \varepsilon > 0$  such that  $\mathbb{P}(X \in \mathcal{B}(x, \varepsilon)) < \delta$  and  $\mathbb{P}(X \in \mathcal{B}(x, R)) > 1 - \delta$ . We cover  $\mathcal{B}(0, R) \setminus \mathcal{B}(0, \varepsilon)$  with  $m$  balls of radius  $\varepsilon/2$ :  $\mathcal{B}(0, R) \setminus \mathcal{B}(0, \varepsilon) \subset \bigcup_{r=1}^m \mathcal{B}(z_r, \varepsilon/2)$ , where  $z_r \in \mathcal{B}(0, R) \setminus \mathcal{B}(0, \varepsilon)$ . We remove all  $z_r$  that satisfy  $\mathbb{P}(X \in \mathcal{B}(z_r, \varepsilon/2)) = 0$  from this union. Let  $\mathcal{E}$  be the event  $\{\forall r \in [m], \exists \ell \in \mathcal{N}_1 \setminus \{k\}, X_\ell \in \mathcal{B}(z_r, \varepsilon/2)\}$ . We have that

$$\begin{aligned} \mathbb{P}(\sigma(i) = k | X_k = x, \sigma(j) = k, \mathcal{E}) &\leq \mathbb{P}(X_i \in \mathcal{B}(x, \varepsilon)) \\ &\leq \delta / \eta^2. \end{aligned}$$

Then,

$$\begin{aligned} \mathbb{P}(\mathcal{E}^C) &\leq \sum_{r=1}^m \mathbb{P}(\forall \ell \in \mathcal{N}_1 \setminus \{k\}, X_\ell \notin \mathcal{B}(z_r, \varepsilon/2)) \\ &\leq m(1 - p_{\min})^{n_1-1}, \end{aligned}$$

where  $p_{\min} = \min_{r \in [m]} \mathbb{P}(X \in \mathcal{B}(z_r, \varepsilon/2) | T = 1) > 0$ . Thus,  $\mathbb{P}(\mathcal{E}) \rightarrow 1$ , and  $\mathbb{P}(\sigma(i) = k | X_k = x, \sigma(j) = k) \leq 1 - \mathbb{P}(\mathcal{E}) + \delta$ . We can thus conclude that  $\mathbb{P}(\sigma(i) = k | X_k = x, \sigma(j) = k) \rightarrow 0$  as  $n_1 \rightarrow \infty$ .

*Wrapping things up.* Using  $\mathbb{P}(\sigma(i) = k | \sigma(j) = k) = \int_{\mathcal{X}} \mathbb{P}(\sigma(i) = k | \sigma(j) = k, X_k = x) d\mathbb{P}(X_k = x)$ , we have  $\mathbb{P}(\sigma(i) = k | \sigma(j) = k) \rightarrow 0$ , using dominated convergence.  $\square$

Using this, we have  $\text{var}(A_2) \rightarrow 0$ , leading to  $A_2 \rightarrow 0$  in probability.

**Control of  $A_1$ .** Using unconfoundedness and overlap:

$$|A_1| \leq \delta^{-1} \mathbb{E} \left[ \delta(X_i, X_{\sigma^*(i)}, Y_i(0)) | X_{\mathcal{N}_1}, T_i = 0 \right], \quad \text{where}$$

$$\delta(x, x', y) \stackrel{\text{def}}{=} |\mathbb{E}[w(Y_i(1)|y) | X_i = x] - \mathbb{E}[w(Y_j(1)|y) | X_j = x']|.$$

Let  $\varepsilon > 0$  and  $y$  fixed. Using our continuity and compactness assumptions,  $x, x' \mapsto \delta(x, x', y)$  is uniformly continuous on  $\mathcal{X} \times \mathcal{X}$ , so that there exists  $\eta > 0$  such that if  $\|x - x'\| \leq \eta$ , we have  $\delta(x, x', y) \leq \varepsilon$ . We are going to show that with high probability,  $\|X_i - X_{\sigma^*(i)}\| \leq \eta$ . Using compactness of  $\mathcal{X}$ , there exist  $u_1, \dots, u_p \in \mathcal{X}$  such that  $\mathcal{X} \subset \bigcup_{k=1}^p \mathcal{B}(u_k, \eta/2)$ . Let  $p_k = \mathbb{P}(X_i \in \mathcal{B}(u_k, \eta/2))$ : we assume that  $p_k > 0$  for all  $k$ , otherwise we remove this ball and the corresponding  $u_k$ . Let  $p_{\min} = \min_k p_k > 0$ . Let  $k_x \in [p]$  such that  $X_i \in \mathcal{B}(u_{k_x}, \eta/2)$ . We have, working conditionnally on  $\mathcal{N}_0, \mathcal{N}_1, i \in \mathcal{N}_0$ :

$$\begin{aligned} \mathbb{P}(\|X_i - X_{\sigma^*(i)}\| > \eta) &\leq \mathbb{E} \left[ \mathbb{P}(X_{\sigma^*(i)} \notin \mathcal{B}(u_{k_x}, \eta/2) | k_x) \right] \\ &= \mathbb{E} \left[ \mathbb{P}(\forall j \in \mathcal{N}_1, X_j \notin \mathcal{B}(u_{k_x}, \eta/2) | k_x) \right] \\ &= \mathbb{E} \left[ (1 - \eta p_{k_x})^{n_1} \right] \\ &\leq (1 - \eta p_{\min})^{n_1} \\ &\xrightarrow[n_1 \rightarrow \infty]{} 0. \end{aligned}$$

This leads to:

$$\mathbb{P}(\delta(X_i, X_{\sigma(i)}, y) > \varepsilon) \leq (1 - \eta p_{\min})^{n_1},$$

and thus  $\mathbb{P}(\delta(X_i, X_{\sigma(i)}, y) \rightarrow 0) = 1$  as  $n_1 \rightarrow \infty$ , leading to  $\mathbb{E}[\delta(X_i, X_{\sigma(i)}, Y_i(0))] \rightarrow 0$  using dominated convergence. We thus have that  $\mathbb{E}[|A_1|] \rightarrow 0$ , and thus  $A_1 \rightarrow 0$  in probability, since  $|A_1| \leq 1$  almost surely.

□

## C Proof of Theorem 3

*Proof of Theorem 3.* We have, for the estimator  $\hat{\tau}$  defined in Equation (12):

$$\begin{aligned}\hat{\tau} - \tau_\star &= \frac{1}{n} \sum_{i=1}^n (1 - T_i) \hat{q}_1(X_i, Y_i) + T_i \hat{q}_0(X_i, Y_i) - \mathbb{E} [w(Y^{(X_i)}(1), Y_i(0))] \\ &= \frac{1}{n} \sum_{i=1}^n (1 - T_i) q_1(X_i, Y_i) + T_i q_0(X_i, Y_i) - \mathbb{E} [w(Y^{(X_i)}(1), Y_i(0))] \\ &\quad + \frac{1}{n} \sum_{i=1}^n (1 - T_i) (\hat{q}_1(X_i, Y_i) - q_1(X_i, Y_i)) + T_i (\hat{q}_0(X_i, Y_i) - q_0(X_i, Y_i)).\end{aligned}$$

For the first term, we have that  $(1 - T_i)q_1(X_i, Y_i) + T_i q_0(X_i, Y_i)$  are *i.i.d.* bounded random variables, of mean  $\mathbb{E} [w(Y^{(X_i)}(1), Y_i(0))]$ , so that the first sum converges almost surely to 0. We even have, using the central limit theorem, that:

$$\frac{1}{\sqrt{n}} \sum_{i=1}^n (1 - T_i) q_1(X_i, Y_i) + T_i q_0(X_i, Y_i) - \mathbb{E} [w(Y^{(X_i)}(1), Y_i(0))] \xrightarrow{\mathbb{P}} \mathcal{N}(0, \sigma_\infty^2),$$

where  $\sigma_\infty^2 = \text{var} ((1 - T_i)q_1(X_i, Y_i) + T_i q_0(X_i, Y_i) - \mathbb{E} [w(Y^{(X_i)}(1), Y_i(0))])$ . For the second term, we have:

$$\begin{aligned}& \left| \frac{1}{n} \sum_{i=1}^n (1 - T_i) (\hat{q}_1(X_i, Y_i) - q_1(X_i, Y_i)) + T_i (q_0(X_i, Y_i) - \hat{q}_0(X_i, Y_i)) \right| \\ & \leq \frac{1}{n} \sum_{i=1}^n |\hat{q}_1(X_i, Y_i) - q_1(X_i, Y_i)| + |\hat{q}_0(X_i, Y_i) - q_0(X_i, Y_i)| \\ & \leq \frac{1}{n} \sum_{i=1}^n (1 - T_i) (|\hat{q}_1(X_i, Y_i) - q_1(X_i, Y_i)| - \mathbb{E} [|\hat{q}_1(X_i, Y_i) - q_1(X_i, Y_i)| | T_i = 0]) \\ & \quad + T_i (|\hat{q}_0(X_i, Y_i) - q_0(X_i, Y_i)| - \mathbb{E} [|\hat{q}_1(X_i, Y_i) - q_1(X_i, Y_i)| | T_i = 1]) \\ & \quad + \mathbb{E} [|\hat{q}_1(X_i, Y_i) - q_1(X_i, Y_i)| | T_i = 0] + \mathbb{E} [|\hat{q}_0(X_i, Y_i) - q_0(X_i, Y_i)| | T_i = 1].\end{aligned}$$

These last two terms are deterministic and converge to zero due to our assumptions. The big sum converges almost surely to zero, as the average of  $n$  *i.i.d.* centered and bounded random

variables. This leads to the consistency of our estimator. For the asymptotic normality, it remains to prove that

$$A_i = \frac{1}{n} \sum_{i=1}^n (1 - T_i) (|\hat{q}_1(X_i, Y_i) - q_1(X_i, Y_i)| - \mathbb{E}[|\hat{q}_1(X_i, Y_i) - q_1(X_i, Y_i)| | T_i = 0]) \\ + T_i (|\hat{q}_0(X_i, Y_i) - q_0(X_i, Y_i)| - \mathbb{E}[|\hat{q}_1(X_i, Y_i) - q_1(X_i, Y_i)| | T_i = 1])$$

is  $o(1/\sqrt{n})$ . We have

$$\text{var} \left( (1 - T_i) (|\hat{q}_1(X_i, Y_i) - q_1(X_i, Y_i)| - \mathbb{E}[|\hat{q}_1(X_i, Y_i) - q_1(X_i, Y_i)| | T_i = 0]) \right) \\ \leq \text{var} (|\hat{q}_1(X_i, Y_i) - q_1(X_i, Y_i)| | T_i = 0) \\ \leq \mathbb{E} [|\hat{q}_1(X_i, Y_i) - q_1(X_i, Y_i)|^2 | T_i = 0] \\ \leq \mathbb{E} [|\hat{q}_1(X_i, Y_i) - q_1(X_i, Y_i)| | T_i = 0],$$

since  $|\hat{q}_1(X_i, Y_i) - q_1(X_i, Y_i)| \leq 1$  (these are probabilities). Thus, under our assumption, the variance of each term of  $A_i$  is  $o(\sqrt{1/n})$ . Then,  $\mathbb{P}(|1/n \sum_i A_i| > \varepsilon/\sqrt{n}) \leq \text{var}(A_i) \rightarrow 0$ . This leads to  $1/n \sum_i A_i = o_{\mathbb{P}}(1/\sqrt{n})$ , and concludes the proof.  $\square$

## D Proof of Theorem 4

*Proof of Theorem 4.* Assume first that (i) holds. Let

$$\tau_{\text{AIPW}}^* \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n \left\{ (1 - T_i) \hat{q}_1(X_i, Y_i) - \lambda (\hat{q}_1(X_i, Y_i) - w(Y_{\sigma_0(i)} | Y_i)) \frac{1 - T_i}{(1 - \pi(X_i))} \right\} \\ + \frac{1}{n} \sum_{i=1}^n \left\{ T_i \hat{q}_0(X_i, Y_i) - (1 - \lambda) (\hat{q}_0(X_i, Y_i) - w(Y_i | Y_{\sigma_1(i)})) \frac{T_i}{\pi(X_i)} \right\}.$$

We have:

$$\hat{\tau}_{\text{AIPW}} - \tau_{\text{AIPW}}^* = \frac{\lambda}{n} \sum_{i=1}^n (\hat{q}_1(X_i, Y_i) - w(Y_{\sigma_0(i)} | Y_i)) \left\{ \frac{1 - T_i}{(1 - \pi(X_i))} - \frac{1 - T_i}{(1 - \hat{\pi}(X_i))} \right\} \\ + \frac{1 - \lambda}{n} \sum_{i=1}^n (\hat{q}_0(X_i, Y_i) - w(Y_i | Y_{\sigma_1(i)})) \left\{ \frac{T_i}{\pi(X_i)} - \frac{T_i}{\hat{\pi}(X_i)} \right\},$$

so that:

$$\begin{aligned}
|\hat{\tau}_{\text{AIPW}}^* - \tau_{\text{AIPW}}| &\leq \frac{\lambda}{n} \sum_{i=1}^n \left| (\hat{q}_1(X_i, Y_i) - w(Y_{\sigma_0(i)}|Y_i)) \left\{ \frac{1 - T_i}{(1 - \pi(X_i))} - \frac{1 - T_i}{(1 - \hat{\pi}(X_i))} \right\} \right| \\
&\quad + \frac{1 - \lambda}{n} \sum_{i=1}^n \left| (\hat{q}_0(X_i, Y_i) - w(Y_i|Y_{\sigma_1(i)})) \left\{ \frac{T_i}{\pi(X_i)} - \frac{T_i}{\hat{\pi}(X_i)} \right\} \right| \\
&\leq \frac{\lambda}{n} \sum_{i=1}^n \left| \frac{1 - T_i}{(1 - \pi(X_i))} - \frac{1 - T_i}{(1 - \hat{\pi}(X_i))} \right| + \frac{1 - \lambda}{n} \sum_{i=1}^n \left| \frac{T_i}{\pi(X_i)} - \frac{T_i}{\hat{\pi}(X_i)} \right|.
\end{aligned}$$

We showed in the proof of Theorem 2 that this quantity converges to 0 under our assumptions.

We thus are left with proving that  $\tau_{\text{AIPW}}^*$  converges in probability towards  $\tau_*$ .<sup>4</sup> We have:

$$\begin{aligned}
\tau_{\text{AIPW}}^* - \tau_* &= \underbrace{\frac{\lambda}{n} \sum_{i=1}^n \left\{ w(Y_{\sigma_0(i)}|Y_i) \frac{1 - T_i}{(1 - \pi(X_i))} - \tau_* \right\}}_{(I)} + \underbrace{\frac{1 - \lambda}{n} \sum_{i=1}^n \left\{ w(Y_i|Y_{\sigma_1(i)}) \frac{T_i}{\pi(X_i)} - \tau_* \right\}}_{(I)} \\
&\quad + \underbrace{\frac{1}{n} \sum_{i=1}^n \left\{ (1 - T_i) \hat{q}_1(X_i, Y_i) - \lambda \hat{q}_1(X_i, Y_i) \frac{1 - T_i}{(1 - \pi(X_i))} \right\}}_{(II)} \\
&\quad + \underbrace{\frac{1}{n} \sum_{i=1}^n \left\{ T_i \hat{q}_0(X_i, Y_i) - (1 - \lambda) \hat{q}_0(X_i, Y_i) \frac{T_i}{\pi(X_i)} \right\}}_{(III)}.
\end{aligned}$$

Using Theorem 2, we directly have that (I) converges in probability towards 0. We now need to prove the same for (II) – (III). We start with (II). First,

$$\begin{aligned}
\mathbb{E}[(1 - T_i) \hat{q}_1(X_i, Y_i) | \hat{q}_1] &= \mathbb{P}(T_i = 0) \mathbb{E}[(1 - T_i) \hat{q}_1(X_i, Y_i) | T_i = 0, \hat{q}_1], \\
\mathbb{E} \left[ \hat{q}_1(X_i, Y_i) \frac{1 - T_i}{(1 - \pi(X_i))} | \hat{q}_1 \right] &= \mathbb{E}[(1 - T_i) \hat{q}_1(X_i, Y_i) | T_i = 0, \hat{q}_1],
\end{aligned}$$

and thus,

$$\mathbb{E} \left[ (1 - T_i) \hat{q}_1(X_i, Y_i) - \lambda \hat{q}_1(X_i, Y_i) \frac{1 - T_i}{(1 - \pi(X_i))} | \hat{q}_1 \right] = (\mathbb{P}(T_i = 0) - \lambda) \mathbb{E}[(1 - T_i) \hat{q}_1(X_i, Y_i) | T_i = 0, \hat{q}_1].$$

Now, since each term in the sum that defines  $(II)$  is bounded by  $1/\eta$ , we get that:

$$\begin{aligned}
\mathbb{E}[(II)^2|\hat{q}_1] &= (\mathbb{P}(T_i = 0) - \lambda)^2 \mathbb{E}[(1 - T_i)\hat{q}_1(X_i, Y_i)|T_i = 0, \hat{q}_1]^2 \\
&\quad + \frac{\text{var}\left(\left(1 - T_i\right)\hat{q}_1(X_i, Y_i) - \lambda\hat{q}_1(X_i, Y_i)\frac{1-T_i}{(1-\hat{\pi}(X_i))}\middle|\hat{q}_1\right)}{n} \\
&\leq (\mathbb{P}(T_i = 0) - \lambda)^2 \mathbb{E}[(1 - T_i)\hat{q}_1(X_i, Y_i)|T_i = 0, \hat{q}_1]^2 + \frac{1}{n} \\
&\leq (\mathbb{P}(T_i = 0) - \lambda)^2 + \frac{1}{n} \\
&\longrightarrow 0,
\end{aligned}$$

under our assumption on  $\lambda$ . The same then applies to  $(III)$ , concluding the proof for our first point.

Assume now that  $(ii)$  holds. Let

$$\begin{aligned}
\tau_{\text{AIPW}}^* &\stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n \left\{ (1 - T_i)q_1(X_i, Y_i) - \lambda(q_1(X_i, Y_i) - w(Y_{\sigma_0(i)}|Y_i))\frac{1 - T_i}{(1 - \hat{\pi}(X_i))} \right\} \\
&\quad + \frac{1}{n} \sum_{i=1}^n \left\{ T_i q_0(X_i, Y_i) - (1 - \lambda)(q_0(X_i, Y_i) - w(Y_i|Y_{\sigma_1(i)}))\frac{T_i}{\hat{\pi}(X_i)} \right\}.
\end{aligned}$$

We have:

$$\begin{aligned}
|\hat{\tau}_{\text{AIPW}} - \tau_{\text{AIPW}}^*| &\leq \frac{1}{n} \sum_{i=1}^n \left| 1 - T_i - \lambda \frac{1 - T_i}{(1 - \hat{\pi}(X_i))} \right| |\hat{q}_1(X_i, Y_i) - q_1(X_i, Y_i)| \\
&\quad + \frac{1}{n} \sum_{i=1}^n \left| T_i - (1 - \lambda) \frac{T_i}{\hat{\pi}(X_i)} \right| |\hat{q}_0(X_i, Y_i) - q_0(X_i, Y_i)| \\
&\leq \frac{1}{\eta'n} \sum_{i=1}^n |\hat{q}_1(X_i, Y_i) - q_1(X_i, Y_i)| \\
&\quad + \frac{1}{\eta'n} \sum_{i=1}^n |\hat{q}_0(X_i, Y_i) - q_0(X_i, Y_i)|,
\end{aligned}$$

and we have shown in Theorem 3 that under our assumptions, these two sums converge to 0 in

probability. Hence, we are left with proving that  $|\tau_{\text{AIPW}}^* - \tau_\star| \rightarrow 0$ .

$$\begin{aligned} |\tau_{\text{AIPW}}^* - \tau_\star| &\leq \left| \frac{1}{n} \sum_{i=1}^n (1 - T_i) q_1(X_i, Y_i) + \frac{1}{n} \sum_{i=1}^n T_i q_0(X_i, Y_i) - \tau_\star \right| \\ &\quad + \frac{\lambda}{n} \left| \sum_{i=1}^n \frac{1 - T_i}{1 - \hat{\pi}(X_i)} (q_1(X_i, Y_i) - w(Y_{\sigma_0(i)} | Y_i)) \right| \\ &\quad + \frac{1 - \lambda}{n} \left| \sum_{i=1}^n \frac{T_i}{\hat{\pi}(X_i)} (q_0(X_i, Y_i) - w(Y_i | Y_{\sigma_1(i)})) \right|. \end{aligned}$$

For the first term, we already know that it converges to 0 thanks to the proof of Theorem 3.

We thus need to control the second and third sums. We have, for the second sum:

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \frac{1 - T_i}{1 - \hat{\pi}(X_i)} (q_1(X_i, Y_i) - w(Y_{\sigma_0(i)} | Y_i)) &= \frac{1}{n} \sum_{i=1}^n \frac{1 - T_i}{1 - \hat{\pi}(X_i)} (\mathbb{E} [w(Y_i(1) | Y_i(0)) | X_i, Y_i(0)] - w(Y_{\sigma_0(i)} | Y_i(0))) \\ &= \frac{1}{n} \sum_{i=1}^n \frac{1 - T_i}{1 - \hat{\pi}(X_i)} (\mathbb{E} [w(Y_i^{(X_i)}(1), Y_i(0)) | X_i, Y_i(0)] - \mathbb{E} [w(Y_{\sigma_0(i)} | Y_i(0)) | X_i, Y_i(0), X_{\mathcal{N}_1}]) \\ &\quad + \frac{1}{n} \sum_{i=1}^n \frac{1 - T_i}{1 - \hat{\pi}(X_i)} (\mathbb{E} [w(Y_{\sigma_0(i)} | Y_i(0)) | X_i, Y_i(0), X_{\mathcal{N}_1}] - w(Y_{\sigma_0(i)} | Y_i(0))). \end{aligned}$$

Here,

$$\frac{1}{n} \sum_{i=1}^n \frac{1 - T_i}{1 - \hat{\pi}(X_i)} (\mathbb{E} [w(Y_{\sigma_0(i)} | Y_i(0)) | X_i, Y_i(0), X_{\mathcal{N}_1}] - w(Y_{\sigma_0(i)} | Y_i(0)))$$

is the mean of  $n$  bounded and centered random variables. These random variables are not independent, but conditionally on  $X_{\mathcal{N}_1}$  they are. Thus, this sum converges almost surely to 0 conditionally on  $X_{\mathcal{N}_1}$ , and thus converges almost surely to 0. We now are left with controlling:

$$A \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n \frac{1 - T_i}{1 - \hat{\pi}(X_i)} (\mathbb{E} [w(Y_i^{(X_i)}, Y_i(0)) | X_i, Y_i(0)] - \mathbb{E} [w(Y_{\sigma_0(i)} | Y_i(0)) | X_i, Y_i(0), X_{\mathcal{N}_1}]).$$

Note that using unconfoundedness and our uniform boundedness assumption on  $\hat{\pi}$ :

$$\begin{aligned} |A| &\leq \delta'^{-1} \mathbb{E} [\zeta(X_i, X_{\sigma_0(i)}, Y_i(0)) | X_{\mathcal{N}_1}, T_i = 0], \quad \text{where} \\ \zeta(x, x', y) &\stackrel{\text{def}}{=} \left| \mathbb{E} [w(Y_i^{(X_i)}(1) | y) | X_i = x] - \mathbb{E} [w(Y_j(1) | y) | X_j = x'] \right|. \end{aligned}$$

Let  $\varepsilon > 0$  and  $y$  fixed. Using our continuity and compactness assumptions,  $x, x' \mapsto \zeta(x, x', y)$  is uniformly continuous on  $\mathcal{X} \times \mathcal{X}$ , so that there exists  $\eta > 0$  such that if  $\|x - x'\| \leq \eta$ , we

have  $\zeta(x, x', y) \leq \varepsilon$ . We are going to show that with high probability,  $\|X_i - X_{\sigma_0(i)}\| \leq \eta$ . Using compactness of  $\mathcal{X}$ , there exist  $u_1, \dots, u_p \in \mathcal{X}$  such that  $\mathcal{X} \subset \bigcup_{k=1}^p \mathcal{B}(u_k, \eta/2)$ . Let  $p_k = \mathbb{P}(X_i \in \mathcal{B}(u_k, \eta/2))$ : we assume that  $p_k > 0$  for all  $k$ , otherwise we remove this ball and the corresponding  $u_k$ . Let  $p_{\min} = \min_k p_k > 0$ . Let  $k_x \in [p]$  such that  $X_i \in \mathcal{B}(u_{k_x}, \eta/2)$ . We have, working conditionnally on  $\mathcal{N}_0, \mathcal{N}_1, i \in \mathcal{N}_0$ :

$$\begin{aligned}
\mathbb{P}(\|X_i - X_{\sigma_0(i)}\| > \eta) &\leq \mathbb{E}[\mathbb{P}(X_{\sigma_0(i)} \notin \mathcal{B}(u_{k_x}, \eta/2) | k_x)] \\
&= \mathbb{E}[\mathbb{P}(\forall j \in \mathcal{N}_1, X_j \notin \mathcal{B}(u_{k_x}, \eta/2) | k_x)] \\
&= \mathbb{E}[(1 - \eta p_{k_x})^{n_1}] \\
&\leq (1 - \eta p_{\min})^{n_1} \\
&\xrightarrow{n_1 \rightarrow \infty} 0.
\end{aligned}$$

This leads to:

$$\mathbb{P}(\zeta(X_i, X_{\sigma_0(i)}, y) > \varepsilon) \leq (1 - \eta p_{\min})^{n_1},$$

and thus  $\mathbb{P}(\zeta(X_i, X_{\sigma_0(i)}, y) \rightarrow 0) = 1$  as  $n_1 \rightarrow \infty$ , leading to  $\mathbb{E}[\zeta(X_i, X_{\sigma_0(i)}, Y_i(0))] \rightarrow 0$  using dominated convergence. We thus have that  $\mathbb{E}[|A|] \rightarrow 0$ , and  $A \rightarrow 0$  in probability, since  $|A| \leq 1$  almost surely. We proceed in the same way for the remaining term:

$$\frac{1}{n} \sum_{i=1}^n \frac{1 - T_i}{1 - \hat{\pi}(X_i)} (\mathbb{E}[w(Y_{\sigma_0(i)} | Y_i(0)) | X_i, Y_i(0), X_{\mathcal{N}_1}] - w(Y_{\sigma_0(i)} | Y_i(0))),$$

concluding the proof. □