

Article

# Normalized Augmented Inverse Probability Weighting with Neural Network Predictions

Mehdi Rostami \*  and Olli Saarela 

Dalla Lana School of Public Health, University of Toronto, 155 College st., Toronto, ON M5T 3M7, Canada; olli.saarela@utoronto.ca

\* Correspondence: mehdi.rostamiforooshani@mail.utoronto.ca

**Abstract:** The estimation of average treatment effect (ATE) as a causal parameter is carried out in two steps, where in the first step, the treatment and outcome are modeled to incorporate the potential confounders, and in the second step, the predictions are inserted into the ATE estimators such as the augmented inverse probability weighting (AIPW) estimator. Due to the concerns regarding the non-linear or unknown relationships between confounders and the treatment and outcome, there has been interest in applying non-parametric methods such as machine learning (ML) algorithms instead. Some of the literature proposes to use two separate neural networks (NNs) where there is no regularization on the network's parameters except the stochastic gradient descent (SGD) in the NN's optimization. Our simulations indicate that the AIPW estimator suffers extensively if no regularization is utilized. We propose the normalization of AIPW (referred to as nAIPW) which can be helpful in some scenarios. nAIPW, provably, has the same properties as AIPW, that is, the double-robustness and orthogonality properties. Further, if the first-step algorithms converge fast enough, under regulatory conditions, nAIPW will be asymptotically normal. We also compare the performance of AIPW and nAIPW in terms of the bias and variance when small to moderate  $L_1$  regularization is imposed on the NNs.

**Keywords:** causal inference; instrumental variables; neural networks; doubly robust estimation; semi-parametric theory



**Citation:** Rostami, M.; Saarela, O. Normalized Augmented Inverse Probability Weighting with Neural Network Predictions. *Entropy* **2022**, *24*, 179. <https://doi.org/10.3390/e24020179>

Academic Editor: Kateřina Hlaváčková-Schindler

Received: 13 November 2021

Accepted: 27 December 2021

Published: 25 January 2022

**Publisher's Note:** MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Copyright:** © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

## 1. Introduction

Estimation of causal parameters such as the average treatment effect (ATE) in observational data requires confounder adjustment. The estimation and inference are carried out in two steps: In step 1, the treatment and outcome are predicted by a statistical models or machine learning (ML) algorithm, and in the second step the predictions are inserted into the causal effect estimator. If ML algorithms are employed in step 1, the non-linear relationships can potentially be taken into account. The relationship between the confounders and the treatment and outcome can be non-linear which make the application of machine learning (ML) algorithms, which are non-parametric models, appealing. Farrell et al. [1] proposed to use two separate neural networks (double NNs or dNNs) where there is no regularization on the network's parameters except the stochastic gradient descent (SGD) in the NN's optimization [2–5]. They derive the generalization bounds and prove that the NN's algorithms are fast enough so that the asymptotic distribution of causal estimators such as the augmented inverse probability weighting (AIPW) estimator [6–8] will be asymptotically linear, under regulatory conditions and the utilization of cross-fitting [9].

Farrell et al. [1] argue that the fact that SGD-type algorithms control the complexity of the NN algorithm to some extent [2,10] is sufficient for the first step. Our initial simulations and analyses, however, contradict this claim in scenarios where strong confounders and instrumental variables (IVs) exist in the data.

Conditioning on IVs is harmful to the performance of the causal effect estimators such as ATE (Myers et al. [11]) but there may be no prior knowledge about which covariates

are IVs, confounders or otherwise. The harm comes from the fact that the complex NNs can provide near-perfect prediction in the treatment model which violates the empirical positivity assumption [12].

The positivity assumption (Section 2) is fundamental to hold to have an identifiable causal parameter in a population. However, in a finite sample, although the parameter is identifiable by making the positivity assumption, the bias and variance of the estimator can be inflated if the estimated propensity scores are close to zero or one bounds (or become zero or one by rounding errors). This is referred to as the empirical positivity assumption which is closely related to the concept of sparsity studied in Chapter 10 of Van der Laan and Rose [8]. The violation of the empirical positivity assumption can cause the inflation of the bias and variance of inverse probability weighting (IPW)-type and AIPW-type estimators.

The inverse probability weighting method dates at least back to Horvitz and Thompson [13] in the literature of sampling with unequal selection probabilities in sub-populations. IPW-type and matching methods have been extensively studied Lunceford and Davidian [7], Rubin [14], Rosenbaum and Rubin [15,16], Busso et al. [17]. IPW is proven to be a consistent estimator of ATE if the propensity scores (that are the conditional probability of treatment assignments) are estimated by a consistent parameter or non-parametric model. The other set of ATE estimators include those involving the modeling of the outcome and inserting the predictions directly into the ATE estimator (Section 2). They are referred to as single robust (SR) estimators as they provide  $\sqrt{n}$ -consistent estimators for ATE if the outcome model is  $\sqrt{n}$ -consistent. In this sense, IPW is also single robust as it is consistent if the treatment (or the propensity score) model is  $\sqrt{n}$ -consistent. The focus of this work is to study the **augmented** IPW-type methods as they involve modeling both treatment and outcome and can be  $\sqrt{n}$ -consistent estimators of ATE if either of the models is consistent.

We propose and study a simple potential remedy to the empirical positivity violation issue by studying the normalization of the AIPW estimator (similar to the normalization of IPW [7]), here referred to as nAIPW. In fact, both AIPW and nAIPW can be viewed as a more general estimator which is derived via the efficient influence function of ATE [18,19].

A general framework of estimators that includes nAIPW as a special case was proposed by [20]. In their work, the authors did not consider machine learning algorithms for the first-step estimation, but rather assumed parametric statistical models estimated by likelihood-based approaches. They focused on how to consistently estimate ATE within different sub-populations imposed by the covariates. There is a lack of numerical experimentation on these estimators especially when IVs and strong confounders exist in the set of candidate covariates.

To the best of our knowledge, the performance of nAIPW has not been previously studied in the machine learning context, with the assumption that strong confounders and IVs exist in the data. We will prove that this estimator has the doubly robust [6] and the rate doubly robust [19] property, and illustrate that it is robust against extreme propensity score values. Further, nAIPW (similar to AIPW), has the orthogonality property [9] which means that it is robust against small variations in the predictions of the outcome and treatment assignment predictions. One theoretical difference is that AIPW is the most efficient estimator among all the double robust estimators of ATE given both treatment and outcome models are correctly specified [21]. In practice, however, often there is no a priori knowledge about the true outcome and propensity score relationships with the input covariates and thus this feature of AIPW is probably of less practical use.

We argue that for causal parameter estimation, dNN with no regularization may lead to high variance for the causal estimator used in the second step. We compare AIPW and nAIPW through a simulation study where we allow for moderate to strong confounding and instrumental variable effects, that is, we allow for possible violation of the empirical positivity assumption. Further, a comparison between AIPW and nAIPW is made on the Canadian Community Health Survey (CCHS) dataset where the intervention/treatment is the food security vs. food insecurity and the outcome is individuals' body mass index (BMI).

Our contributions include presenting the proof for the orthogonality, doubly robust and rate doubly robust property of nAIPW. Further, it is proven that, under certain assumptions, nAIPW is asymptotically normal and we provide its consistent variance estimator. We analyze the estimation of ATE in the presence of not only confounders, but also IVs, y-predictors and noise variables. We demonstrate that in the presence of strong confounders and IVs, if complex neural networks without  $L_1$  regularizations are used in the step 1 estimation, both AIPW and nAIPW estimators and their asymptotic variances perform poorly, but, relatively speaking, nAIPW performs better. In this paper, the NNs are mostly used as means of estimating the outcome and treatment predictions.

Organization of the article is as follows. In Section 2 we will formally introduce the nAIPW estimator to the readers and state its double robustness property, and in Section 3 we present the first-step prediction model, double neural networks. In Sections 4 and 5 we will present the theoretical aspects of the paper, including the asymptotic normality, doubly robustness and rate doubly robustness orthogonality of the proposed estimator (nAIPW) and the asymptotic normality. We will present the simulation scenarios and results of comparing the nAIPW estimator with other conventional estimators in Section 6. We apply the estimators on a real dataset in Section 7. The article will be concluded with a short discussion on the findings in Section 8. The proofs are straightforward but long and thus are included in Appendix A.

## 2. Normalized Doubly Robust Estimator

Let data  $\mathbf{O} = (O_1, O_2, \dots, O_n)$  be generated by a data generating process  $P$ , where  $O_i$  is a finite dimensional vector  $O_i = (Y_i, A_i, W_i)$ , with  $\mathbf{W}$  being the adjusting factors.  $P$  is the true observed data distribution,  $\hat{P}_n$  is the distribution of  $\mathbf{O}$  such that its marginal distribution with respect to  $W$  is its empirical distribution and the expectation of the conditional distribution  $Y|A = a, W$ , for  $a = 0, 1$ , can be estimated. We denote the prediction function of the observed outcome given explanatory variables in the treated group  $Q^1 := Q(1, W) = \mathbb{E}[Y|A = 1, W]$ , and that in the untreated group  $Q^0 := Q(0, W) = \mathbb{E}[Y|A = 0, W]$ , and the propensity score as  $g(W) = \mathbb{E}[A|W]$ . Throughout, the expectations  $\mathbb{E}$  are with respect to  $P$ . The symbol  $\hat{\cdot}$  on the population-level quantities indicates the corresponding finite sample estimator, and  $P$  is replaced by  $\hat{P}_n$ .

Let the causal parameter of interest be the average treatment effect (ATE)

$$\beta_{ATE} = \mathbb{E}[Y^1 - Y^0] = \mathbb{E}[\mathbb{E}[Y^1 - Y^0|W]] = \mathbb{E}[\mathbb{E}[Y|A = 1, W]] - \mathbb{E}[\mathbb{E}[Y|A = 0, W]], \quad (1)$$

where  $Y^1$  and  $Y^0$  are the potential outcomes of the treatment and controls [6].

For identifiability of the parameter, the following assumptions must hold true. The first assumption is the conditional independence, or unconfoundedness stating that, given the confounders, the potential outcomes are independent of the treatment assignments ( $Y^0, Y^1 \perp A|W$ ). The second assumption is positivity which entails that the assignment of treatment groups is not deterministic ( $0 < Pr(A = 1|W) < 1$ ). The third assumption is consistency which states that the observed outcomes equal their corresponding potential outcomes ( $Y^A = y$ ). There are other modeling assumptions made such as time order (i.e., the covariates  $W$  are measured before the treatment), IID subjects and a linear causal effect.

A list of first candidates to estimate ATE are

$$\begin{aligned}
 \text{naive ATE} \quad \hat{\beta}_{naiveATE} &= \frac{1}{n_1} \sum_{i \in A_1} \hat{Q}_i^1 - \frac{1}{n_0} \sum_{i \in A_0} \hat{Q}_i^0, \\
 \text{SR} \quad \hat{\beta}_{SR} &= \hat{\mathbb{E}}[\hat{\mathbb{E}}[Y^1 - Y^0|W]] = \frac{1}{n} \sum_{i=1}^n \hat{Q}_i^1 - \hat{Q}_i^0, \\
 \text{IPW} \quad \hat{\beta}_{IPW} &= \hat{\mathbb{E}}\left[\frac{Y^1}{\hat{\mathbb{E}}[A|W]} - \frac{Y^0}{1 - \hat{\mathbb{E}}[A|W]}\right] = \frac{1}{n} \sum_{i=1}^n \left(\frac{A_i y_i}{\hat{g}_i} - \frac{(1 - A_i)y_i}{1 - \hat{g}_i}\right), \\
 \text{nIPW} \quad \hat{\beta}_{nIPW} &= \sum_{i=1}^n \left(\frac{A_i w_i^{(1)} y_i}{\sum_{j=1}^n A_j w_j^{(1)}} - \frac{(1 - A_i) w_i^{(0)} y_i}{\sum_{j=1}^n (1 - A_j) w_j^{(0)}}\right).
 \end{aligned} \tag{2}$$

The naive average treatment effect (naive ATE) is a biased (due to the selection bias) estimator of ATE [22] and is the poorest estimator among all the candidates. The single robust (SR) is not an orthogonal estimator [9] and if ML algorithms which do not belong to the Donsker class ([23], Section 19.2) or have entropy that grows with the sample size are used, this estimator also becomes biased and is not asymptotically normal. The inverse probability weighting (IPW) [13] and its normalization versions adjust (or weight) the observations in the treatment and control groups. IPW and nIPW are also not orthogonal estimators and are similar to SR in this respect. In addition, both  $\hat{\beta}_{SR}$  and  $\hat{\beta}_{IPW}$  (and  $\hat{\beta}_{nIPW}$ ) are single robust, that is, they are consistent estimators of ATE if the models used are  $\sqrt{n}$ -consistent [7]. IPW is an unbiased estimator of ATE if  $g$  is correctly specified, but nIPW is not unbiased, but is less sensitive to extreme predictions. The augmented inverse probability weighting (AIPW) estimator [21] is an improvement over SR, IPW and nIPW, which involves the predictions for both treatment (the propensity score), and the causal parameter can be expressed as:

$$\beta = \mathbb{E}\left[\left(\frac{AY - Q(1, W)(A - \mathbb{E}[A|W])}{\mathbb{E}[A|W]} - \frac{(1 - A)Y + Q(0, W)(A - \mathbb{E}[A|W])}{1 - \mathbb{E}[A|W]}\right)\right], \tag{3}$$

and the sample version estimator of (3) is

$$\begin{aligned}
 \hat{\beta}_{AIPW} &= \frac{1}{n} \sum_{i=1}^n \left[\left(\frac{A_i Y_i - \hat{Q}(1, W_i)(A_i - \hat{\mathbb{E}}[A_i|W_i])}{\hat{\mathbb{E}}[A_i|W_i]} - \frac{(1 - A_i)Y_i + \hat{Q}(0, W_i)(A_i - \hat{g}_i)}{1 - \hat{\mathbb{E}}[A_i|W_i]}\right)\right] = \\
 &= \frac{1}{n} \sum_{i=1}^n \left(\frac{A_i(y_i - \hat{Q}_i^1)}{\hat{g}_i} - \frac{(1 - A_i)(y_i - \hat{Q}_i^0)}{1 - \hat{g}_i}\right) + \hat{\beta}_{SR}, \tag{4}
 \end{aligned}$$

where  $\hat{Q}_i^k = \hat{Q}(k, W_i) = \hat{\mathbb{E}}[Y_i|A_i = k, W_i]$  and  $\hat{g}_i = \hat{\mathbb{E}}[A_i|W_i]$ .

Among all the doubly robust estimators of ATE, AIPW is the most efficient estimator if both of the propensity score or outcome models are correctly specified, but is not necessarily efficient under incorrect model specification. In fact, this nice feature of AIPW may be less relevant in real-life problems as we might not have a priori knowledge about the predictors of the propensity score and outcome and we cannot correctly model them. Further, in practice, perfect or near-perfect prediction of the treatment assignment can inflate the variance of the AIPW estimator [8]. As a remedy, similar to the normalization of the IPW estimator, we can define a normalized version of the AIPW estimator which is less

sensitive to extreme values of the predicted propensity score, referred to as the normalized augmented inverse probability weighting (nAIPW) estimator:

$$\hat{\beta}_{nAIPW} = \sum_{i=1}^n \left( \frac{A_i(y_i - \hat{Q}_i^1)w_i^{(1)}}{\sum_{j=1}^n A_j w_j^{(1)}} - \frac{(1 - A_i)(y_i - \hat{Q}_i^0)w_i^{(0)}}{\sum_{j=1}^n (1 - A_j)w_j^{(0)}} \right) + \hat{\beta}_{SR}, \tag{5}$$

where  $w_k^{(1)} = \frac{1}{\hat{g}_k}$  and  $w_k^{(0)} = \frac{1}{1-\hat{g}_k}$ . Both AIPW and nAIPW estimators add adjustment factors to the SR estimator which involve both models of the treatment and the outcome.

Both AIPW and nAIPW are examples of a class of estimators where

$$\hat{\beta}_{GDR} = \frac{1}{n} \sum_{i=1}^n \left( \frac{A_i(y_i - \hat{Q}_i^1)}{\hat{h}_i^1} - \frac{(1 - A_i)(y_i - \hat{Q}_i^0)}{\hat{h}_i^0} \right) + \hat{\beta}_{SR}, \tag{6}$$

where we refer to this general class as the general doubly robust (GDR) estimator. Letting  $\hat{h}^1 = \hat{g}$  and  $\hat{h}^0 = 1 - \hat{g}$  gives the AIPW estimators and letting  $\hat{h}^1 = \hat{g}\hat{\mathbb{E}}\frac{A}{\hat{g}}$  and  $\hat{h}^0 = (1 - \hat{g})\hat{\mathbb{E}}\frac{1-A}{1-\hat{g}}$  gives the nAIPW estimator.

The GDR estimator can also be written as

$$\hat{\beta}_{GDR} = \hat{\mathbb{E}} \left( \left[ \frac{A}{\hat{h}^1} - \frac{1-A}{\hat{h}^0} \right] y - (A - \hat{h}^1)\hat{Q}^1 + (1 - A - \hat{h}^0)\hat{Q}^0 \right), \tag{7}$$

If  $h^1$  and  $h^0$  are chosen so that

$$\mathbb{E}[A - h^1] = 0, \mathbb{E}[1 - A - h^0] = 0, \tag{8}$$

by the total law of expectation  $\hat{\beta}_{GDR}$  is an unbiased estimator of  $\beta$ .

### 3. Outcome and Treatment Predictions

The causal estimation and inference when utilizing the AIPW and nAIPW is carried out in two steps. In step 1, the treatment and outcome are predicted by a statistical or machine learning (ML) algorithm, and in the second step the predictions are inserted into the estimator. The ML algorithms in step 1 can capture the linear and non-linear relationships between the confounders and the treatment and the outcome.

Neural networks (NNs) [2–4] are a class of non-linear and non-parametric complex algorithms that can be employed to model the relationship between any set of inputs and some outcome. There has been a tendency to use NNs as they have achieved great success in the most complex artificial intelligence (AI) tasks such as computer vision and natural language understanding [2].

Farrell et al. [1] used two independent NNs for modeling the propensity score model and the outcome with the rectified linear unit (RELU) activation function [2], here referred to as the double NN or dNN:

$$\begin{aligned} \mathbb{E}[Y|A, W] &= \beta_0 + \beta A + \mathbf{W}\alpha + \mathbf{H}\Gamma_Y \\ \mathbb{E}[A|W] &= \beta'_0 + \mathbf{W}'\alpha' + \mathbf{H}'\Gamma_A, \end{aligned} \tag{9}$$

where two separate neural nets model  $y$  and  $A$  (no parameter sharing). Farrell et al. [1] proved that dNN algorithms almost attain  $n^{\frac{1}{4}}$ -rates. By employing the cross-fitting method and theory developed by Chernozhukov et al. [9], an orthogonal causal estimator is asymptotically normal, under some regularity and smoothing conditions, if the dNN is used in the first step (see Theorem 1 in [1]).

These results assume no regularizations imposed on the NNs' weights, and only the stochastic gradient descent (SGD) is used. Farrell et al. claim that the fact that SGD controls the complexity of the NN algorithm to some extent [2,10] is sufficient for the first step. Our initial simulations, however, contradict this claim and we hypothesize that for causal

parameter estimation, a dNN with no regularization leads to high variance for the causal estimator used in the second step. Our initial experiments indicate that  $L_2$  regularization and dropout do not perform well in terms of the mean square error (MSE) of AIPW. The loss functions we use contain  $L_1$  regularization (in addition to SGD during the optimization):

$$\begin{aligned} L_Y(\mathcal{P}_Y, \beta, \alpha) &= \sum_{i=1}^n \left[ y_i - \alpha' - \beta A_i - \mathbf{W}_i \alpha - H_i^T \Gamma_Y \right]^2 + C_{L_1} \sum_{\omega \in \mathcal{P}} |\omega|, \\ L_A(\mathcal{P}_A, \alpha') &= \sum_{i=1}^n \left[ A_i \log \left( g(H_i^T \Gamma_A) \right) + (1 - A_i) \log \left( 1 - g(H_i^T \Gamma_A) \right) \right] + \\ &C'_{L_1} \sum_{\omega \in \mathcal{P}} |\omega|, \end{aligned} \quad (10)$$

where  $C_{L_1}, C'_{L_1}$  are hyperparameters that can be set before training or be determined by cross-validation, that can cause the training to pay more attention to one part of the output layer. The dNN can have an arbitrary number of hidden layers, or the width of the network ( $\mathcal{HL}$ ) can be another hyperparameter. For a three-layer network,  $\mathcal{HL} = [l_1, l_2, \dots, l_h]$ , where  $l_j$  is the number neurons in layer  $j$ ,  $j = 1, 2, \dots, h$ .  $\mathcal{P}_Y, \mathcal{P}_A$  are the connection parameters in the non-linear part of the networks, with  $\Omega$ s being shared for the two outcome and propensity models. Note that the gradient descent-type optimizations in the deep learning platforms (such as pytorch in our case) do not cause the NN parameters to shrink to zero.

#### 4. GDR Estimator Properties

In this section we will see that nAIPW (5) is doubly robust, that is, if either of the outcome or propensity score models are  $\sqrt{n}$ -consistent, nAIPW will be consistent. Further, nAIPW is orthogonal [9] and is asymptotically linear under certain assumptions and we calculate its asymptotic variance.

##### 4.1. Consistency and Asymptotic Distribution of nAIPW

In causal inference, estimating the causal parameter and drawing inference on the parameter are two major tasks. Employing a machine learning algorithm to estimate  $Q$  and  $g$  in (5) is a means to estimate and draw inference on the causal parameter; the ultimate goal is the relationship between the treatment and the outcome. This allows people to use blackbox ML models with no explanation how these models have learned from the explanatory features. The question is if the consistency and asymptotic normality of the second step causal estimator are preserved if complex ML algorithms are utilized twice for the treatment and outcome models, each with a convergence rate smaller than  $\sqrt{n}$ , and entropy that grows with  $n$ .

Chernozhukov et al. [24] provide numerical experiments illustrating that some estimators are not consistent or asymptotically normal if complex ML models are used that do not belong to the Donsker class and have entropy that grows with  $n$ . They further provide a solution by introducing “orthogonal” estimators that, under some regulatory conditions and cross-fitting, are asymptotically normal even if complex ML models can be used as long as their rates of convergence are as small as  $n^{\frac{1}{4}}$ .

The next two subsections provide an overview of the general theory and prove that nAIPW is asymptotically normal.

##### 4.2. The Efficient Influence Function

Hahn [18] derives the efficient influence function (EIF) of  $\beta = \beta_1 - \beta_0$  as

$$\phi(O, P) = \left( \frac{A}{g} (Y - Q^1) + Q^1 - \beta_1 \right) - \left( \frac{1-A}{1-g} (Y - Q^0) + Q^0 - \beta_0 \right) \quad (11)$$

To study the asymptotic behaviour of nAIPW, we write the scaled difference

$$\sqrt{n}(\hat{\beta} - \beta) = \frac{1}{\sqrt{n}} \sum_{i=1}^n \phi(O_i, P) - \frac{1}{\sqrt{n}} \sum_{i=1}^n \phi(O_i, \hat{P}_n) + \sqrt{n}(P_n - P)[\phi(O_i, \hat{P}_n) - \phi(O_i, P)] - \sqrt{n}R(P, \hat{P}_n), \quad (12)$$

where the first term is a normal distribution by the central limit theorem, and the third and fourth terms are controlled if the class of functions are Donsker and standard smoothing conditions are satisfied ([9,23], Theorem 19.26). If the nuisance parameters are not Donsker, data splitting and cross-fitting guarantees plus the regulatory conditions are needed to control these two terms [1,9]. It is unclear, however, how the second term behaves, i.e.,

$$-\frac{1}{\sqrt{n}}\phi(O, \hat{P}_n) = -\frac{1}{\sqrt{n}} \sum_{i=1}^n \left[ \frac{A_i}{g_i}(Y_i - \hat{Q}_i^1) - \frac{1 - A_i}{1 - g_i}(Y_i - \hat{Q}_i^0) + \hat{Q}_i^1 - \hat{Q}_i^0 \right] - \hat{\beta}, \quad (13)$$

where  $\hat{\beta} = \beta(\hat{P}_n)$ , as it contains data-adaptive nuisance parameter estimations. There are different tricks to get rid of this term. One method is the one-step method in which we move this term to the left to create a new estimator which is exactly the same as the AIPW estimator with known propensity scores:

$$\sqrt{n}\left(\hat{\beta} + \frac{1}{n}\phi(O, \hat{P}_n) - \beta\right) = \sqrt{n}\left(\frac{1}{n} \sum_{i=1}^n \left[ \frac{A_i}{g_i}(Y_i - \hat{Q}_i^1) - \frac{1 - A_i}{1 - g_i}(Y_i - \hat{Q}_i^0) + \hat{Q}_i^1 - \hat{Q}_i^0 \right] - \beta\right). \quad (14)$$

Another trick is to let this term vanish which results in estimating equations whose solution is exactly the same as the one-step estimator. The targeted learning strategy is to manipulate the data generating process which results in a different estimator [8,19] (which we do not study here).

The requirement in the above estimator is that the propensity score is known, which is unrealistic. In reality, this quantity should be estimated using the data. However, replacing  $g$  with a data-adaptive estimator changes the remainder term in (12) that needs certain assumptions to achieve asymptotic properties such as consistency. We replace  $g$  and  $1 - g$  in (14) by  $\hat{h}^1$  and  $\hat{h}^0$ , respectively, which provides a more general view of the above one-step estimator.

### 4.3. Doubly Robustness and Rate Doubly Robustness Properties of GDR

One of the appealing properties of AIPW is its doubly robust property which partially relaxes the restrictions of IPW and SR which require the consistency of the treatment and outcome models, respectively. This property is helpful when the first-step algorithms are  $\sqrt{n}$ -consistent. The following theorem states that the nAIPW estimator (5) actually possesses the doubly robustness property.

**Theorem 1** (nAIPW Double Robustness). *The DR estimator (5) is consistent if  $\hat{Q}^k \xrightarrow{P} Q^k$ ,  $k = 0, 1$  or  $\hat{g} \xrightarrow{P} g$ .*

The proof is left to the appendix. Theorem 1 is useful when we *a priori* knowledge about the propensity scores (such as in the experimental studies) or we estimate the propensity scores with  $\sqrt{n}$ -rate converging algorithms. In practice, however, the correct specification is infeasible in the observational data, but  $\sqrt{n}$ -rate algorithms such as parametric models, generalized additive models (GAMs) or the models that assume sparsity might be used [25].

This is restrictive and these model assumptions might not hold in practice which is why non-parametric ML algorithms such as NNs are used. As mentioned before, the NN we utilize here does not offer a  $\sqrt{n}$ -consistent prediction model in the first step of the estimation [1]. This reduces the usefulness of the double robustness property of the GDR estimator when using complex ML algorithms. A more useful property when using complex ML algorithms is the *rate double robustness (RDR)* property [26]. RDR does not require either of the prediction models to be  $\sqrt{n}$ -consistent; it suffices that they are consistent at any rate but together become  $\sqrt{n}$ -consistent; that is, if the propensity score and outcome model are consistent at  $n^{r_A}$  and  $n^{r_Y}$ , respectively ( $r_Y, r_A < 0$ ), we must have  $r_A + r_Y = \frac{1}{2}$ . To see that the DR has this property (as does DR [25]), note that the remainder (12) can be written as

$$-\sqrt{n}R(P, \hat{P}_n) = \sqrt{n}\mathbb{E}\left[\left(\frac{g}{\hat{h}^1} - 1\right)(Q^1 - \hat{Q}^1)\right] + \sqrt{n}\mathbb{E}\left[\left(\frac{1-g}{\hat{h}^0} - 1\right)(Q^0 - \hat{Q}^0)\right], \quad (15)$$

which, by the Hölder inequality, is upper bounded:

$$-\sqrt{n}R(P, \hat{P}_n) \leq \left[\mathbb{E}\left[\frac{g}{\hat{h}^1} - 1\right]^2\right]^{\frac{1}{2}} \left[\mathbb{E}\left[Q^1 - \hat{Q}^1\right]^2\right]^{\frac{1}{2}} + \left[\mathbb{E}\left[\frac{1-g}{\hat{h}^0} - 1\right]^2\right]^{\frac{1}{2}} \left[\mathbb{E}\left[Q^0 - \hat{Q}^0\right]^2\right]^{\frac{1}{2}} \quad (16)$$

Making the standard assumptions that

$$\begin{aligned} \left[\mathbb{E}\left[g - \hat{h}^k\right]^2\right]^{\frac{1}{2}} \left[\mathbb{E}\left[Q^k - \hat{Q}^k\right]^2\right]^{\frac{1}{2}} &= o(n^{-\frac{1}{2}}), \quad k = 0, 1, \\ \mathbb{E}\left[g - \hat{h}^k\right]^2 &= o(1), \quad \mathbb{E}\left[Q^k - \hat{Q}^k\right]^2 = o(1), \quad k = 0, 1, \\ \text{Empirical Positivity} \quad &c_1 < \hat{h}^k < 1 - c_2, \text{ for some } c_1, c_2 > 0, \end{aligned} \quad (17)$$

implies

$$-\sqrt{n}R(P, \hat{P}_n) = o(n^{-\frac{1}{2}}), \quad (18)$$

that is, the GDR has the rate double robustness property.

The assumptions in (17) are less restrictive than needing at least one of the prediction models to be  $\sqrt{n}$ -consistent for the double robust property [19,25]. This means that the outcome and propensity score models can be at least as fast as  $o(n^{-\frac{1}{4}})$  (which is an attainable generalization bound for many complex machine learning algorithms [9]), and the GDR estimator is still consistent. Farrell et al. [1] proves that two neural networks without regularization (except the one imposed by the stochastic gradient descent optimization) satisfy such bounds and can provide a convenient first-step prediction algorithm (when they utilize the AIPW estimator and the cross-fitting strategy proposed by Chernozhukov et al. [9]).

In order for a special case of GDR estimator to outperform the AIPW estimator, we must have  $Ah^1 \geq Ag$  and  $(1 - A)h^0 \geq (1 - A)(1 - g)$ , in addition to conditions in (17). Note that these two conditions are satisfied for nAIPW; replacing  $h^1$  and  $h^0$  with  $\hat{g}\hat{\mathbb{E}}_{\frac{A}{g}}$  and  $(1 - \hat{g})\hat{\mathbb{E}}_{\frac{1-A}{1-g}}$  can help stabilize the bias and variance magnitude and help shrink the remainder (15) to zero. The scenario analysis performed in Section 4.4 provides an insight about the reduction in the sensitivity to the violation of the empirical positivity assumption.



4.4. Robustness of nAIPW against Extreme Propensity Scores

There are two scenarios in which the empirical positivity is violated, where the probability of receiving the treatment for the people who are treated is 1, that is,  $A_k = 1$  and  $P(A_k = 1|W) = 1$  (or vice versa for the untreated group  $A_k = 0$  and  $P(A_k = 0|W) = 0$ ), and where there are a handful of treated subjects whose probability of receiving the treatment is 0, that is,  $A_k = 1$  and  $P(A_k = 1|W) = 0$  (and vice versa for the untreated group, that is,  $A_k = 0$  and  $P(A_k = 0|W) = 1$ ). Although the identifiability assumptions guarantee that such scenarios do not occur, in practice, extremely small or large probabilities similar to the second scenario above, that is, where there exists a treated individual who has a near-zero probability of receiving the treatment, can impact the performance of the estimators that involve propensity score weighting. For example, replacing  $h^1$  with  $\hat{g}$  and  $h^0$  with  $1 - \hat{g}$  in practice can increase both the bias and variance of AIPW [8]. This can be seen by viewing the bias and variance of these weighting terms. As noted before, the AIPW and nAIPW add adjustments to the single robust estimator  $\mathbb{E}Q^1 - Q^0$ . The adjustments involve weightings  $\frac{A}{g}$  or  $\frac{A}{g\mathbb{E}\frac{A}{g}}$  to the residuals of  $Y$  and  $Q^k, k = 0, 1$ . Under a correct specification of the propensity score  $g$ , these weights have the same expectations. The difference is in their variances:

$$\begin{aligned} \text{Var}\left(\frac{A}{g}\right) &= \frac{1}{g} - 1, \\ \text{Var}\left(\frac{A}{g\mathbb{E}\frac{A}{g}}\right) &= \frac{1}{\mathbb{E}^2\frac{A}{g}}\left(\frac{1}{g} - 1\right), \end{aligned} \tag{19}$$

under the correct specification of the propensity score  $g$ . By letting  $g$  tend to zero in violation of the empirical positivity assumption, it can be seen that the nAIPW is less volatile than the AIPW estimator. That is, the weights in AIPW might have a larger variance than those in nAIPW.

4.5. Scenario Analysis

A scenario analysis is performed to see how nAIPW stabilizes the estimator: Assume that the empirical positivity is violated, that is, there is at least an observation  $k$  where  $A_k = 1$  where  $\hat{g}_k$  is extremely close to zero, such as  $\hat{g}_k = 10^{-s}$  for  $s \gg 0$ . AIPW will blow up in this case:

$$\begin{aligned} \beta_{1,AIPW} &= \frac{1}{n}\left(10^s(Y_k^1 - Q_k^1) + \sum_{i \in I_{-k}^1} \frac{Y_i^1 - Q_i^1}{g_i}\right) + \frac{1}{n} \sum_{i=1}^n Q_i^1, \\ \beta_{0,AIPW} &= \frac{1}{n}\left(\sum_{i \in I^0} \frac{Y_i^0 - Q_i^0}{1 - g_i}\right) + \frac{1}{n} \sum_{i=1}^n Q_i^0, \end{aligned} \tag{20}$$

where  $I^a = \{j : A_j = a\}$ ,  $I_{-k}^a = \{j : A_j = a\}$ , and subscripts  $a = 1$  and  $a = 0$  refer to the estimators of the first and the second components in ATE (1). However, nAIPW is robust against this empirical positivity violation:

$$\beta_{1,nAIPW} = \left(\frac{Y_k^1 - Q_k^1}{10^{-s}(10^s + \sum_{j \neq k} \frac{A_j}{g_j})} + \sum_{i \in I_{-k}^1} \frac{Y_i^1 - Q_i^1}{g_i(10^s + \sum_{j \neq k} \frac{A_j}{g_j})}\right) + \frac{1}{n} \sum_{i=1}^n Q_i^1, \tag{21}$$

and

$$\beta_{0,nAIPW} = \left(\frac{0 \times (Y_k^1 - Q_k^0)}{\star} + \sum_{i \in I_{-k}^0} \frac{Y_i^0 - Q_i^0}{(1 - g_i)(\sum_{j=1}^n \frac{1 - A_j}{1 - g_j})}\right) + \frac{1}{n} \sum_{i=1}^n Q_i^0. \tag{22}$$

Thus

$$\beta_{1,nAIPW} \approx \left( \frac{Y_k^1 - Q_k^1}{1 + 10^{-s}(n-1)} + \sum_{i \in I_{-k}^1} \frac{Y_i^1 - Q_i^1}{g_i 10^s + g_i(n-1)} \right) + \frac{1}{n} \sum_{i=1}^n Q_i^1, \tag{23}$$

The factor  $10^s$  in (20) can blow up the AIPW if  $10^s \gg n$  (and the outcome estimation is not close enough to the observer outcome), but this factor does not appear in the numerator of the nAIPW estimator. For such large factors, (23) can be simplified to

$$\beta_{1,nAIPW} \approx Y_k^1 - Q_k^1 + \frac{1}{n} \sum_{i=1}^n Q_i^1. \tag{24}$$

Thus, the extreme probability does not make  $\beta_{1,nAIPW}$  blow up, but the adjustment to the  $\beta_{1,SR}$  that accounts for confounding effects. The second factor  $\beta_{0,nAIPW}$  is not impacted in this scenario.

Considering a scenario that there is another treated individual with extremely small probability, such as  $\hat{g}_l = 10^{-t}$ , such that, without loss of generality,  $t > s \gg 0$ , we will have:

$$\beta_{1,nAIPW} \approx \frac{Y_k^1 - Q_k^1}{1 + 10^{t-s} + 10^{-s}(n-2)} + \frac{Y_l^1 - Q_l^1}{1 + 10^{s-t} + 10^{-t}(n-2)} + \frac{1}{n} \sum_{i=1}^n Q_i^1. \tag{25}$$

Depending on the values  $s$  and  $t$ , one of the first two terms in (25) might vanish, but the estimator does not blow up. There is at most only a handful of treated individuals with extremely small probabilities and, based on the above observation, the nAIPW estimator does not blow up. That said, nAIPW might not sufficiently correct the  $\beta_{SR}$  for the confounding effects, although confounders have been taken into account in the calculation of  $\beta_{SR}$  to some extent.

The same observation can be made in the asymptotic variance of these estimators. This shows how extremely small probabilities for treated individuals (or extremely large probabilities for untreated individuals) can result in a biased and unstable estimator, while neither of the bias or variance of nAIPW suffer as much. Although not performed, the same observation can be made for the untreated individuals with extremely large probabilities.

The above scenario analysis indicates the bias and variance of nAIPW might go up in cases of the violation of empirical positivity, but it still is less biased and more stable than AIPW. The remainder term (15) is also more likely to be  $o(n^{-\frac{1}{2}})$  in nAIPW versus AIPW as it contains  $k$ 's where  $A_k = 1, g_k \mathbb{E}_n \frac{A_k}{g_k} \geq g_k$ .

### 5. Asymptotic Sampling Distribution of nAIPW

Replacing  $g$  in the denominator of the von Mises expansion (12) with the normalizing terms is enough to achieve the asymptotic distribution of the nAIPW and its asymptotic standard error. However, we can see that nAIPW is also the solution to (extended) estimating equations. The solution to the estimating equations is important as van der Vaart (Chapters 19 and 25) proves that under certain regularity conditions, if the prediction models belong to the Donsker class, the solutions to Z-estimators are consistent and asymptotically normal ([23], Theorem 19.26). Thus, nAIPW that is the solution to a Z-estimator (also referred to an M-estimator) will inherit the consistency and asymptotic normality,

assuming certain regulatory conditions and that the first-step prediction models belong to the Donsker class:

$$\begin{aligned} \mathbb{E} \left[ \frac{A(Y^1 - Q^1)}{\gamma g} - \frac{(1 - A)(Y^0 - Q^0)}{\lambda(1 - g)} + (Q^1 - Q^0 - \beta) \right] &= 0, \\ \mathbb{E} \left[ \frac{A}{g} - \gamma \right] &= 0, \\ \mathbb{E} \left[ \frac{1 - A}{1 - g} - \lambda \right] &= 0. \end{aligned} \tag{26}$$

The Donsker class assumption prevents too complex algorithms in the first step, algorithms such as tree-based models, NNs, cross-hybrid algorithms or their aggregations [19,27]. The Donsker class assumption can be relaxed if sample splitting (or cross-fitting) is utilized and the target parameter is orthogonal [9]. In the next section we see that nAIPW is orthogonal and, thus, theoretically, we can relax the Donsker class assumption under certain smoothing regulatory conditions. Before seeing the orthogonality property of nAIPW, let us review the smoothing regularity conditions necessary for asymptotic normality. Let  $\beta$  be the causal parameter,  $\eta \in T$  be the infinite dimensional nuisance parameters where  $T$  is a convex set with a norm. Additionally, let the score function  $\phi : \mathbb{O} \times \mathcal{B} \times T \rightarrow \mathbb{R}$  be a measurable function,  $\mathbb{O}$  be the measurable space of all random variables  $O$  with probability distribution  $P \in \mathcal{P}_n$  and  $\mathcal{B}$  be an open subset of  $\mathbb{R}$  containing the true causal parameter. Let the sample  $O = (O_1, O_2, \dots, O_n)$  be observed and the set of probability measures  $\mathcal{P}_n$  expand with sample size  $n$ . In addition, let  $\beta \in \mathcal{B}$  be the solution to the estimating equation  $\mathbb{E}\phi(\mathbb{O}, \beta, \eta) = 0$ . The assumptions that guarantee that the second-step orthogonal estimator  $\hat{\beta}$  is asymptotically normal are [9]: (1)  $\beta$  does not fall on the boundary of  $\mathcal{B}$ ; (2) the map  $(\beta, \eta) \rightarrow \mathbb{E}_P\phi(\mathbb{O}, \beta, \eta)$  is twice Gateaux differentiable (this holds by the positivity assumption).  $\beta$  is identifiable; (3)  $\mathbb{E}_P\phi(\mathbb{O}, \beta, \eta)$  is smooth enough; (4)  $\hat{\eta} \in T$  with high probability and  $\eta \in T$ .  $\hat{\eta}$  converges to  $\eta_0$  at least as fast as  $n^{-\frac{1}{4}}$  (similar but slightly stronger than first two assumptions in (17)); (5) score function(s)  $\phi(\cdot, \beta, \eta)$  has finite second moment for all  $\beta \in \mathcal{B}$  and all nuisance parameters  $\eta \in T$ ; (6) the score function(s)  $\phi(\cdot, \beta, \eta)$  is measurable; (7) the number of folds increases by sample size.

### 5.1. Orthogonality and the Regulatory Conditions

The orthogonality condition [9] is a property related to the estimating equations

$$\mathbb{E}\phi(\mathbb{O}, \beta, \eta) = 0. \tag{27}$$

We refer to an estimator drawn from the estimating Equation (27) as an orthogonal estimator.

Let  $\eta \in T$ , where  $T$  is a convex set with a norm. Additionally, let the score functions  $\phi : \mathbb{O} \times \mathcal{B} \times T \rightarrow \mathbb{R}$  be a measurable function,  $\mathbb{O}$  is measurable space of all random variables  $O$  with probability distribution  $P \in \mathcal{P}_n$  and  $\mathcal{B}$  is an open subset of  $\mathbb{R}$  containing the true causal parameter. Let the sample  $O = (O_1, O_2, \dots, O_n)$  be observed and the set of probability measures  $\mathcal{P}_n$  can expand with sample size  $n$ . The score function  $\phi$  follows the Neyman orthogonality condition with respect to  $\mathcal{T} \subseteq T$ , if the Gateaux derivative operator exists for all  $\epsilon \in [0, 1)$ :

$$\partial_{\tilde{\eta}} \mathbb{E}_P \phi(\mathbb{O}, \beta_0, \tilde{\eta}) \Big|_{\tilde{\eta}=\eta} [\tilde{\eta} - \eta] := \partial_{\epsilon} \mathbb{E}_P \phi(\mathbb{O}, \beta_0, \eta + \epsilon(\tilde{\eta} - \eta)) \Big|_{\epsilon=0} = 0. \tag{28}$$

Chernozhukov et al. [24] presents a few examples of orthogonal estimating equations including the AIPW estimator (4). Utilizing cross-fitting, under standard regulatory conditions, the asymptotic normality of estimators with orthogonal estimating equations is guaranteed even if the nuisance parameters are estimated by ML algorithms not belonging to the Donsker class and without finite entropy conditions [24]. The regulatory conditions to be satisfied are (1)  $\beta$  does not fall on the boundary of  $\mathcal{B}$ ; (2) the map  $(\beta, \eta) \rightarrow \mathbb{E}_P\phi(\mathbb{O}, \beta, \eta)$

is twice Gateaux differentiable.  $\beta$  is identifiable; (3)  $\mathbb{E}_P\phi(\mathbf{O}, \beta, \eta)$  is smooth enough; (4)  $\hat{\eta} \in \mathcal{T}$  with high probability and  $\eta \in \mathcal{T}$ .  $\hat{\eta}$  converges to  $\eta_0$  at least as fast as  $n^{-\frac{1}{4}}$ ; (5) score function(s)  $\phi(\cdot, \beta, \eta)$  has finite second moment for all  $\beta \in \mathcal{B}$  and all nuisance parameters  $\eta \in \mathcal{T}$ ; (6) the score function(s)  $\phi(\cdot, \beta, \eta)$  is measurable; (7) the number of folds increases by sample size.

By replacing  $\lambda$  and  $\gamma$  in the first line of (26) with their solutions in the second and third equations:

$$\mathbb{E}_P\phi(\mathbf{O}, \beta, Q^1, Q^0, g) = \mathbb{E} \left[ \frac{A(Y^1 - Q^1)}{g\mathbb{E}\frac{A}{g}} - \frac{(1 - A)(Y^0 - Q^0)}{(1 - g)\mathbb{E}\frac{1-A}{1-g}} + (Q^1 - Q^0 - \beta) \right] = 0, \quad (29)$$

Implementing the orthogonality condition (28), it can be verified that nAIPW (5) is also an example of an orthogonal estimator. To see this, we apply the definition of orthogonality [9]:

$$\begin{aligned} \partial_\eta \mathbb{E}_P\phi(\mathbf{O}, \beta, \eta) \Big|_{\eta=\eta_0} [\eta - \eta_0] &= \\ \partial_\eta \mathbb{E}_P \left( Q^1 + \frac{A(Y^1 - Q^1)}{g\mathbb{E}\frac{A}{g}} - Q^0 - \frac{(1 - A)(Y^0 - Q^0)}{(1 - g)\mathbb{E}\frac{1-A}{1-g}} - \beta \right) \Big|_{\eta=\eta_0} [\eta - \eta_0] &= \\ \propto \partial_\epsilon \mathbb{E}_P \left( Q_\epsilon^1 + \frac{A(Y^1 - Q_\epsilon^1)}{g_\epsilon\mathbb{E}\frac{A}{g_\epsilon}} - Q_\epsilon^0 - \frac{(1 - A)(Y^0 - Q_\epsilon^0)}{(1 - g_\epsilon)\mathbb{E}\frac{1-A}{1-g_\epsilon}} - \beta \right) \Big|_{\epsilon=0} &= \\ \mathbb{E} \left( (\tilde{Q}^1 - Q^1) + \frac{A}{g\mathbb{E}\frac{A}{g}} (- (\tilde{Q}^1 - Q^1)) + A(Y - Q^1)a(g, \tilde{g} - g) \right) - &= \\ \mathbb{E} \left( (\tilde{Q}^0 - Q^0) + \frac{1 - A}{(1 - g)\mathbb{E}\frac{1-A}{1-g}} (- (\tilde{Q}^0 - Q^0)) + \right. &= \\ \left. (1 - A)(Y - Q^0)b(g, \tilde{g} - g) \right) &= 0, \quad (30) \end{aligned}$$

where  $Q_\epsilon^k = \epsilon\tilde{Q}^k + (1 - \epsilon)Q^k$ ,  $k = 0, 1$ , and  $g_\epsilon = \epsilon\tilde{g} + (1 - \epsilon)g$ , and for some functions  $a$  and  $b$ . The last equality is because  $\mathbb{E}A(Y - Q^1) = 0$ ,  $\mathbb{E}(1 - A)(Y - Q^0) = 0$ ,  $\mathbb{E}\frac{A}{g\mathbb{E}\frac{A}{g}} = 1$  and  $\mathbb{E}\frac{1-A}{(1-g)\mathbb{E}\frac{1-A}{1-g}} = 1$ , under correct specification of the propensity score  $g$ .

Thus, nAIPW is orthogonal, and by utilizing cross-fitting for the estimation, nAIPW is consistent and asymptotically normal, under certain regulatory conditions.

### 5.2. Asymptotic Variance of nAIPW

To evaluate the asymptotic variance of nAIPW, we employ the M-estimation theory [23,28]. For causal inference for M-estimators, the bootstrap for the estimation of causal estimator variance is not generally valid even if the nuisance parameter estimators are  $\sqrt{n}$ -convergent. However, sub-sampling  $m$  out of  $n$  observations [29] can be shown to be universally valid, provided  $m \rightarrow \infty$  and  $\frac{m}{n} \rightarrow 0$ . In practice, however, we can face computational issues since nuisance parameters must be separately estimated (possibly with ML models) for each subsample/bootstrap sample.

The variance estimator of AIPW (4) is [7]

$$\begin{aligned} \hat{\sigma}_{AIPW}^2 &= \\ \frac{1}{n^2} \sum_{i=1}^n \left( \frac{A_i Y_i - \hat{Q}_i^1(A_i - \hat{g}_i)}{\hat{g}_i} - \frac{(1 - A_i) Y_i + \hat{Q}_i^0(A_i - \hat{g}_i)}{1 - \hat{g}_i} - \hat{\beta}_{AIPW} \right)^2 &= \\ \frac{1}{n^2} \sum_{i=1}^n \left( \frac{A_i (y_i - \hat{Q}_i^1)}{\hat{g}_i} - \frac{(1 - A_i)(y_i - \hat{Q}_i^0)}{1 - \hat{g}_i} + \hat{\beta}_{SR} - \hat{\beta}_{AIPW} \right)^2. & \quad (31) \end{aligned}$$

The theorem below states that the variance estimator of AIPW (31) can intuitively extend to calculate the variance estimator of nAIPW (5) by moving the denominator  $n^2$  to the square term in the summation and replacing it with  $\hat{g}\hat{\mathbb{E}}(\frac{A}{\hat{g}})$  or  $(1 - \hat{g})\hat{\mathbb{E}}(\frac{1-A}{1-\hat{g}})$  in the terms containing  $g$  and  $1 - g$  in the denominator, respectively.

**Theorem 2.** *The asymptotic variance of the nAIPW (5) is*

$$\hat{\sigma}_{nAIPW}^2 = \sum_{i=1}^n \left( \frac{A_i(y_i - \hat{Q}_i^1)w_i^{(1)}}{\sum_{j=1}^n A_j w_j^{(1)}} - \frac{(1 - A_i)(y_i - \hat{Q}_i^0)w_i^{(0)}}{\sum_{j=1}^n (1 - A_j)w_j^{(0)}} + \frac{1}{n}(\hat{\beta}_{SR} - \hat{\beta}_{nAIPW}) \right)^2, \quad (32)$$

where  $\hat{Q}_i^k = \hat{Q}(k, W_i)$  and  $\hat{g}_i = \hat{\mathbb{E}}[A_i|W_i]$ .

The proof utilizing the estimating equation technique is straightforward and is left to Appendix A. The same result can be seen when deriving the estimator in the one-step method (see (12) and (14)). Since nAIPW is orthogonal,  $\hat{\sigma}_{nAIPW}^2$  is consistent by applying the theories of [1,9], if the assumptions are met, cross-fitting is used, and the step 1 ML algorithms have the required convergence rates.

The above theorem states that the variance estimator of AIPW (31) can intuitively extend to calculate the variance estimator of nAIPW (5) by moving the denominator  $n^2$  to the square term in the summation and replacing it with  $\hat{g}\hat{\mathbb{E}}(\frac{A}{\hat{g}})$  or  $(1 - \hat{g})\hat{\mathbb{E}}(\frac{1-A}{1-\hat{g}})$  in the terms containing  $g$  and  $1 - g$  in the denominator, respectively. This is intuitive because, by the law of total probability,  $\mathbb{E}$  the first two terms is  $n$ .

### 6. Monte Carlo Experiments

A Monte Carlo simulation study (with 100 iterations) was performed to compare AIPW and nAIPW estimators, where the dNN is used for the first-step prediction. There are a total of 2 case scenarios according to the size of the data. We fixed the sample sizes to be  $n = 750$  and  $n = 7500$ , with the number of covariates being  $p = 32$  and  $p = 300$ , respectively. The predictors include four types of covariates: The confounders,  $X_c$ , instrumental variables,  $X_{iv}$ , the outcome predictors,  $X_y$ , and the noise or irrelevant covariates,  $X_{irr}$ . Their sizes for the scenarios are  $\#X_c = \#X_{iv} = \#X_y = \#X_{irr} = 8,75$  and they are independent from each other and drawn from the multivariate normal (MVN) distribution as  $X \sim \mathcal{N}(\mathbf{0}, \Sigma)$ , with  $\Sigma_{kj} = \rho^{j-k}$  and  $\rho = 0.5$ . The models to generate the treatment assignment and outcome were specified as

$$\begin{aligned} A &\sim Ber\left(\frac{1}{1 + e^{-\eta}}\right), \text{ with } \eta = f_a(X_c)\gamma_c + g_a(X_{iv})\gamma_{iv}, \\ y &= 3 + A + f_y(X_c)\gamma'_c + g_y(X_y)\gamma_y + \epsilon, \end{aligned} \quad (33)$$

and  $\beta = 1$ . The functions  $f_a, g_a, f_y, g_y$  select 20% of the columns and apply interactions and non-linear functions listed below (35). The strength of the instrumental variable and confounding effects were chosen as  $\gamma_c, \gamma'_c, \gamma_y \sim Unif(r_1, r_2)$  where  $(r_1 = r_2 = 0.25)$ , and  $\gamma_{iv} \sim Unif(r_3, r_4)$  where  $(r_3 = r_4 = 0.25)$ .

The non-linearities are randomly selected from among the following functions:

$$\begin{aligned} l(x_1, x_2) &= e^{-\frac{x_1 x_2}{2}} \\ l(x_1, x_2) &= \frac{x_1}{1 + e^{x_2}} \\ l(x_1, x_2) &= \left(\frac{x_1 x_2}{10} + 2\right)^3 \\ l(x_1, x_2) &= (x_1 + x_2 + 3)^2 \\ l(x_1, x_2) &= g(x_1) \times h(x_2) \end{aligned} \quad (34)$$


where  $g(x) = -2I(x \leq -1) - I(-1 \leq x \leq 0) + I(0 \leq x \leq 2) + 3I(x \geq 2)$ , and  $h(x) = -5I(x \leq 0) - 2I(0 \leq x \leq 1) + 3I(x \geq 1)$ , or  $g(x) = I(x \geq 0)$  and  $h(x) = I(x \geq 1)$ .

The networks' activation function is rectified linear unit (ReLU), with 3 hidden layers as large as the input size ( $p$ ), with  $L_1$  regularization and batch size equal to  $3 * p$  and 200 epochs. The adaptive moment estimation (Adam) optimizer [30] with learning rate 0.01 and momentum 0.95 was used to estimate the network's parameters, including the causal parameter (ATE).

*Simulation Results*

The oracle estimations are plotted in all the graphs to compare the real-life situations with the truth. In almost all the scenarios we cannot obtain perfect causal effect estimation and inference.

Figure 1 shows the distribution of AIPW and nAIPW for different hyperparameter settings of NNs. The nAIPW estimator outperforms AIPW in almost all the scenarios. As the AIPW gives huge values in some simulation iterations, the log of the estimation is taken in Figure 1.



**Figure 1.** The distribution of log of the estimated AIPW and nAIPW in the 100 simulated iterations. The performance of nAIPW is clearly superior to the performance of AIPW as it is less dispersed and is more stable in terms of different hyperparameter settings.  $p$  is either 32 or 300 for the small or large datasets and  $q \approx \frac{p}{10}$ , that is, 3 or 30.

We also compare the estimators in different scenarios with bias, variance and their tradeoff measures:


$$\begin{aligned}
 \text{Bias} \quad \hat{\delta} &= \beta - \frac{1}{m} \sum_{j=1}^m \hat{\beta}_j \\
 \text{MC std} \quad \hat{\sigma}_{MC} &= \sqrt{\frac{1}{m} \sum_{j=1}^m (\hat{\beta}_j - \hat{\mu})^2} \\
 \text{MC RMSE} \quad \text{RMSE} &= \sqrt{\hat{\sigma}_{MC}^2 + \hat{\delta}^2} \\
 \text{Asymptotic StdErr} \quad \hat{\sigma}_{SE} &= \frac{1}{m} \sum_{j=1}^m \hat{\sigma}_j,
 \end{aligned} \tag{35}$$

where  $\beta = 1$ , with  $\hat{\beta}_j$ s being the AIPW or nAIPW estimations in the  $j$ th simulation round,  $\hat{\mu} = \frac{1}{m} \sum_{j=1}^m \hat{\beta}_j$  and  $m = 100$  being the number of simulation rounds and  $\hat{\sigma}$  being the square root of (31) or (32).


Figure 2 demonstrates the bias, MC standard deviation (MC std) and the root mean square error (RMSE) of AIPW and nAIPW estimators for the scenarios where  $n = 750$  and  $n = 7500$ , and for four hyperparameter sets ( $L_1$  regularization and width of the dNN). In general, in each figure of the panel, the hyperparameter scenarios in the left imply a more complex model (with less regularization or a narrower network). In these graphs, the lower the values, the better the estimator. For the smaller data size  $n = 750$  in the left three panels, the worst results are attributed to AIPW when there is the least regularization and the hidden layers are as wide as the number of inputs. To have more clear plots for comparison, we skipped plotting the upper bounds as they were large numbers; the lower bounds are enough to show the significance of the results. In the scenarios where there are smaller numbers of hidden neurons with  $0.01 L_1$  regularization, the bias, variance and their tradeoff (here measured by RMSE) are more stable. By increasing the  $L_1$  regularization, these measures go down which indicates the usefulness of regularization and AIPW normalization for causal estimation and inference. Almost the same pattern is seen for the larger size ( $n = 7500$ ) scenario, except for the bump in all the three measures in the hyperparameter scenario where regularization remains the same ( $L_1 = 0.01$ ) and the numbers of neurons in the first and last hidden layers are small too. In all three measures of bias, standard deviation and RMSE, nAIPW is superior to AIPW, or at least there is no statistically significant difference between AIPW and nAIPW.

We have noted that the results of the step 1 NN architecture without  $L_1$  regularization are too unstable and cannot be visually presented in the graphs. To avoid that, we have allowed a span of values for the  $L_1$  regularization strengths:  $L_1 = 0.01$  and  $L_1 = 0.1$ . The former case is close to no regularization. So, if the results of the latter are better than the former's, this is evidence that enough  $L_1$  must be imposed.

Figure 3 illustrates how the theoretical standard error formulas perform in MC experiments, and how accurately they estimate the MC standard deviations. In these two graphs, smaller does not necessarily imply superiority. In fact, the best results will be achieved as long as the confidence intervals of asymptotic SEs and MC SDs intersect. In the left two scenarios where the NN's complexity is high, the MC std and SE are far from each other. Additionally, in the hyperparameter scenarios where both the width of the NNs is small and regularization is higher, the MC std and SE are well separated. The scenario with largest regularization and wide NN architecture seems to be the best scenario. That said, none of the scenarios confirm the consistency of SEs, which would likely also result in low coverage probability of the resulting confidence intervals.



**Figure 2.** The bias, MC standard error and the root mean square error of the AIPW and nAIPW estimators for different data sizes and NN hyperparameters ( $L_1$  regularization and width of the network).  $p$  is either 32 or 300 for the small or large datasets and  $q \approx \frac{p}{10}$ , that is, 3 or 30. The estimates are capped at  $-10$  and  $10$ .



**Figure 3.** The MC standard deviation and the standard error of the AIPW and nAIPW estimators for different data sizes and NN hyperparameters ( $L_1$  regularization and width of the network).

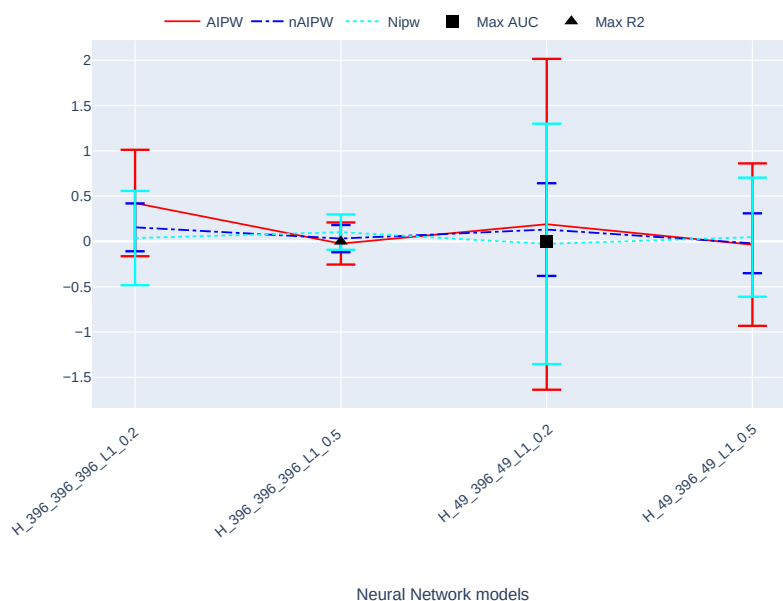


### 7. Application: Food Insecurity and BMI

The Canadian Community Health Survey (CCHS) is a cross-sectional survey that collects data related to health status, health care utilization and health determinants for the Canadian population in multiple cycles. The 2021 CCHS covers the population 12 years of age and over living in the ten provinces and the three territorial capitals. Excluded from the survey’s coverage are: Persons living on reserves and other Aboriginal settlements in the provinces and some other sub-populations that altogether represent less than 3% of the Canadian population aged 12 and over. Examples of modules asked in most cycles are: General health, chronic conditions, smoking and alcohol use. For the 2021 cycle, thematic content on food security, home care, sedentary behavior and depression, among many others, was included. In addition to the health component of the survey are questions about respondent characteristics such as labor market activities, income and socio-demographics.

In this article, we use the CCHS dataset to investigate the causal relationship of food insecurity and body mass index (BMI). Other gathered information in the CCHS is used which might contain potential confounders, y-predictors and instrumental variables. The data are from a survey and need special methods such as the resampling or bootstrap methods to estimate the standard errors. However, here, we use the data to illustrate the utilization of a dNN on the causal parameters in the case of empirical positivity violation. In order to reduce the amount of variability in the data, we have focused on the sub-population 18–65 years of age.

Figure 4 shows the ATE estimates and their 95% asymptotic confidence intervals with nIPW, DR and nDR methods, with four different neural networks which vary in terms of width and strength of  $L_1$  regularization. The scenario that results in the largest  $R^2$  (as a measure of outcome prediction performance) outperforms the other scenarios. The scenario that results in the largest AUC (as a measure of treatment model performance) results in the largest confidence intervals. This is because of more extreme propensity scores in this scenario. It is worth noting that the normalized IPW has smaller confidence intervals as compared to AIPW. However, as we do not know the truth about the ATE in this dataset, we can never know which estimator outperforms the other. To gain insight about this using the input matrix of this data, we simulated multiple treatments and outcomes with small to strong confounders and IVs and compared AIPW and nAIPW. In virtually all of them, the nAIPW is the best one. We do not present these results in this paper, but they can be provided to readers upon request.



**Figure 4.** The ATE estimates and their asymptotically calculated 95% confidence intervals with NIPW, AIPW and nAIPW methods.

## 8. Discussion

Utilizing machine learning algorithms such as NNs in the first-step estimation process is comforting as the concerns with regard to the non-linear relationships between the confounders and the treatment and outcome are addressed. However, there is no free lunch, and using NNs has its own caveats including theoretical as well as numerical challenges. Farrell et al. [1] addressed the theoretical concerns where they calculated the generalization bounds when two separate NNs are used to model the treatment and the outcome. However, they did not use or take into account regularization techniques such as  $L_1$  or  $L_2$  regularization. As NNs are complex algorithms, they provide perfect prediction for the treatment when the predictors are strong enough (or might overfit). Through Monte Carlo (MC) simulations, we illustrated that causal estimation and inference with double NNs can fail without the usage of regularization techniques such as  $L_1$  and/or extreme propensity scores are not taken care of. If  $L_1$  regularization is not used, the normalization of the AIPW estimator (i.e., nAIPW) is advised to be employed as it dilutes the extreme predictions of the propensity score model and provides better bias, variance and RMSE. Our scenario analysis also showed that in the case of violation of the empirical positivity assumption in AIPW, normalization helps avoid blowing up the estimator (and standard error), but might be ineffective in taking into account confounding effects for some observations.

We note that the nAIPW estimator cannot perform better when the empirical positivity is violated as compared to when it is not. However, when the empirical positivity is violated, nAIPW can perform better than AIPW. If the empirical positivity is not violated, our results indicated that AIPW outperforms nAIPW.

An alternative estimator might be trimming the propensity scores to avoid extreme values. However, the causal effect estimator will no longer be consistent and there is no determined method for where to trim. We hypothesize that  $\hat{h}^1 = \hat{g} \mathbb{E}_{\frac{A}{g}} \times I(\hat{g} \in (0, \epsilon)) + \hat{g} \times I(\hat{g} \in (\epsilon, 1))$  and  $\hat{h}^0 = (1 - \hat{g}) \mathbb{E}_{\frac{1-A}{1-g}} \times I(\hat{g} \in (1 - \epsilon, 1)) + (1 - \hat{g}) \times I(\hat{g} \in (0, 1 - \epsilon))$  where  $\epsilon = \frac{1}{n}$  will result in a consistent estimator, making the right assumptions, and will outperform both AIPW and nAIPW in the case of the empirical positivity violation. We will study this hypothesis in a future article.

Another reason why NNs without regularization fail in the causal estimation and inference is that the networks are not targeted, and are not directly designed for these tasks. NNs are complex algorithms with strong predictive powers. This does not accurately serve the purpose of causal parameter estimation, where the empirical positivity assumption can be violated if strong confounders and/or instrumental variables [22] exist in the data. Ideally, the network should target the confounders and should be able to automatically limit the strength of predictors so that the propensity scores are not extremely close to 1 or 0. This was not investigated in this article and a solution to this problem is postponed to another study.

In Section 7, we applied the asymptotic standard errors of both AIPW and nAIPW, where the latter achieves smaller standard errors. That said, we acknowledge the fact that the asymptotic standard errors when using complex ML are not reliable and, in fact, they underestimate the calculated MC standard deviations, as illustrated in the simulations Section 6. This is partly because of the usage of complex algorithms such as NNs for estimation of the nuisance parameters in the first step. Further, the asymptotic distributions of the estimators are not symmetric (and thus are not normal). However, nAIPW is more symmetric than AIPW, according to the simulations, while both estimators suffer from outliers. We will investigate the reasons and possible remedies for both the asymptotic distribution and standard errors of the estimators in a future paper. The consistency of the variance of nAIPW (and AIPW) relies on meeting the assumptions. More investigations are needed on how to achieve consistent and asymptotically normal estimators for ATE with a consistent variance estimator. Potential avenues can include proposing alternative estimators or improving the step 1 ML algorithms.

**Author Contributions:** Data curation, M.R.; Formal analysis, M.R.; Investigation, M.R.; Methodology, M.R. and O.S.; Project administration, M.R. and O.S.; Resources, M.R.; Software, M.R.; Supervision, O.S.; Validation, M.R. and O.S.; Visualization, M.R.; Writing—original draft, M.R.; Writing—review & editing, O.S. All authors have read and agreed to the published version of the manuscript.

**Funding:** The work of Olli Saarela was supported by a Discovery Grant from the Natural Sciences and Engineering Research Council of Canada.

**Data Availability Statement:** The simulated data can be regenerated using the codes, which can be provided to the interested user via an email request to the correspondence author. The CCHS data is not publicly available and only the authorized people can access and perform analyses on it.

**Conflicts of Interest:** The authors declare no conflict of interest.

### Appendix A

First, let us review the proof sketch of the AIPW double robustness: (3) can be consistently estimated by

$$\begin{aligned} \hat{\beta}_{AIPW} &= \frac{1}{n} \sum_{i=1}^n \left[ \left( \frac{A_i Y_i - \hat{Q}(1, W_i)(A_i - \hat{\mathbb{E}}[A_i|W_i])}{\hat{\mathbb{E}}[A_i|W_i]} - \right. \right. \\ &\quad \left. \left. \left( \frac{(1 - A_i) Y_i + \hat{Q}(0, W_i)(A_i - \hat{g}_i)}{1 - \hat{\mathbb{E}}[A_i|W_i]} \right) \right] = \\ &= \frac{1}{n} \sum_{i=1}^n \left( \left[ \frac{A_i}{\hat{g}_i} - \frac{1 - A_i}{1 - \hat{g}_i} \right] y_i - \frac{A_i - \hat{g}_i}{\hat{g}_i(1 - \hat{g}_i)} [(1 - \hat{g}_i)\hat{Q}_i^1 + \hat{g}_i\hat{Q}_i^0] \right) = \\ &= \frac{1}{n} \sum_{i=1}^n \left( \frac{A_i(y_i - \hat{Q}_i^1)}{\hat{g}_i} - \frac{(1 - A_i)(y_i - \hat{Q}_i^0)}{1 - \hat{g}_i} \right) + \frac{1}{n} \sum_{i=1}^n (\hat{Q}_i^1 - \hat{Q}_i^0) \quad (A1) \end{aligned}$$

The second formula guarantees the consistency of AIPW if  $\hat{g}$  is consistent, and the third expression shows that the consistency of  $\hat{Q}_i^0$  and  $\hat{Q}_i^1$  is consistent.

**Theorem A1** (nAIPW double robustness). *Let the nAIPW estimator of risk difference be*

$$\hat{\beta}_{nAIPW} = \hat{\mathbb{E}}(\hat{Q}^1 - \hat{Q}^0) + \hat{\mathbb{E}}\left(\frac{A(Y - \hat{Q}^1)}{\hat{g}\hat{\mathbb{E}}[\frac{A}{\hat{g}}]} - \frac{(1 - A)(Y - \hat{Q}^0)}{(1 - \hat{g})\hat{\mathbb{E}}[\frac{1-A}{1-\hat{g}}]}\right). \quad (A2)$$

Then,  $\hat{\beta}_{nAIPW}$  is a consistent estimator of  $\beta$  if  $\hat{g} \xrightarrow{p} g$  or  $\hat{Q}^k \xrightarrow{p} Q^k, k = 0, 1$ .

**Proof.** From (A2),  $\hat{\beta}_{nAIPW}$  is a consistent estimator of  $\beta$  if  $\hat{Q}_i^0$  and  $\hat{Q}_i^1$  are consistent. This is because the first term  $\hat{\mathbb{E}}(\hat{Q}^1 - \hat{Q}^0)$  converges to  $\beta$ , while the second term tends to zero.

By re-expressing (A2), the other argument is clear. Letting  $\hat{w}^1 = \hat{\mathbb{E}}[\frac{A}{\hat{g}}]$  and  $\hat{w}^0 = \hat{\mathbb{E}}[\frac{1-A}{1-\hat{g}}]$ , we have:

$$\begin{aligned} \hat{\beta}_{nAIPW} &= \frac{1}{n} \sum_{i=1}^n \left( \left[ \frac{A_i}{\hat{g}_i \hat{w}_i^1} - \frac{1 - A_i}{(1 - \hat{g}_i) \hat{w}_i^0} \right] y_i \right) + \\ &= \hat{\mathbb{E}}\left(\hat{Q}^1 - \hat{Q}^0 - \frac{A_i \hat{Q}_i^1}{\hat{g}_i \hat{w}_i^1} + \frac{(1 - A_i) \hat{Q}_i^0}{(1 - \hat{g}_i) \hat{w}_i^0}\right) = \frac{1}{n} \sum_{i=1}^n \left( \left[ \frac{A_i}{\hat{g}_i \hat{w}_i^1} - \frac{1 - A_i}{(1 - \hat{g}_i) \hat{w}_i^0} \right] y_i - \right. \\ &\quad \left. \hat{Q}_i^1 (A_i - \hat{g}_i \hat{w}_i^1) + \hat{Q}_i^0 (1 - A_i - (1 - \hat{g}_i) \hat{w}_i^0) \right) \quad (A3) \end{aligned}$$

The first expression in (A3) is the same as the nIPW estimator which is a consistent estimator of  $\beta$  [7]. Now, under the consistency of  $\hat{g}$ , the second term tends to zero, as  $\hat{w}_1 \xrightarrow{p} 1$  and  $\hat{w}_0 \xrightarrow{p} 1$ .

In the theorem below, it is shown that there is an M-estimation equivalent to  $\beta_{nAIPW}$  and  $w^1$  and  $w^0$ . This, plus the continuous mapping theorem, proves that  $\sum_{i=1}^n \frac{A_i}{\hat{g}_i}$  converges in probability to  $n$  if  $\hat{g} \xrightarrow{p} g$ .

□

**Theorem A2.** *The asymptotic variance of the nAIPW (5) is*

$$\hat{\sigma}_{nAIPW}^2 = \sum_{i=1}^n \left( \frac{A_i(y_i - \hat{Q}_i^1)w_i^{(1)}}{\sum_{j=1}^n A_j w_j^{(1)}} - \frac{(1 - A_i)(y_i - \hat{Q}_i^0)w_i^{(0)}}{\sum_{j=1}^n (1 - A_j)w_j^{(0)}} + \frac{1}{n} (\hat{\beta}_{SR} - \hat{\beta}_{nAIPW}) \right)^2, \tag{A4}$$

where  $\hat{Q}_i^k = \hat{Q}(k, W_i)$  and  $\hat{g}_i = \hat{\mathbb{E}}[A_i|W_i]$ .

**Proof.** Let us define a few notations first:

$$\begin{aligned} q &= Q^1 - Q^0, \\ g &= \mathbb{E}[A|W], \\ f &= y - Q^1, \\ h &= y - Q^0, \\ v &= \frac{A}{g}, \\ u &= \frac{1 - A}{1 - g}. \end{aligned} \tag{A5}$$

With this set of notations, the nAIPW estimator (5) can be written as

$$\hat{\beta}_{nAIPW} = \sum_{i=1}^n \left( \frac{v_i f_i}{\sum_{j=1}^n v_j} - \frac{u_i h_i}{\sum_{j=1}^n u_j} + \frac{q_i}{n} \right), \tag{A6}$$

Following the methods in [28], to find an estimating equation whose solution is  $\hat{\beta}_{nAIPW}$ , we introduce two more estimating equations. Employing the M-estimation theory, we will prove that nAIPW is asymptotically normal, and we will calculate its standard error.

It can be seen that (A6) is not a solution to an M-estimator directly. However, by defining two more parameters and concatenating their estimating equations, we obtain 3-dim multivariate estimating equations:

$$\begin{aligned} \sum_{i=1}^n \left( \frac{v_i f_i}{\gamma} - \frac{u_i h_i}{\lambda} + \frac{1}{n} (q_i - \beta) \right) &= 0, \\ \sum_{i=1}^n \left( v_i - \frac{\gamma}{n} \right) &= 0, \\ \sum_{i=1}^n \left( u_i - \frac{\lambda}{n} \right) &= 0. \end{aligned} \tag{A7}$$

To ease the calculations, we modify the first estimating equation with an equivalent one, but the results will not differ:

$$\begin{aligned} \sum_{i=1}^n \lambda v_i f_i - \gamma u_i h_i + \frac{\gamma \lambda}{n} (q_i - \beta) &= 0, \\ \sum_{i=1}^n v_i - \frac{\gamma}{n} &= 0, \\ \sum_{i=1}^n u_i - \frac{\lambda}{n} &= 0. \end{aligned} \tag{A8}$$

By defining the following notations,

$$\psi = \begin{pmatrix} \phi \\ \eta \\ \Omega \end{pmatrix} = \begin{pmatrix} \lambda v f - \gamma u h + \frac{\gamma \lambda}{n} (q - \beta) \\ v - \frac{\gamma}{n} \\ u - \frac{\lambda}{n} \end{pmatrix},$$

we have  $\sum_{i=1}^n \psi_i = 0$ , or

$$\begin{aligned} \sum_{i=1}^n \phi_i &= 0, \\ \sum_{i=1}^n \eta_i &= 0, \\ \sum_{i=1}^n \Omega_i &= 0. \end{aligned} \tag{A9}$$

The M-estimation theory implies that under regulatory conditions, the solutions to these estimating equations converge in distribution to a multivariate normal distribution:

$$\sqrt{n} \begin{pmatrix} \hat{\beta}_{nAIPW} \\ \hat{\gamma} \\ \hat{\lambda} \end{pmatrix} \sim MVN(\theta, \mathbf{I}^{-1}(\theta) \mathbf{B}(\theta) \mathbf{I}^{-1}(\theta)^T)$$

where

$$\begin{aligned} \theta &= \begin{pmatrix} \beta \\ \gamma \\ \lambda \end{pmatrix}, \\ \mathbf{I}(\theta) &= -\mathbb{E} \frac{\partial \psi}{\partial \theta^T} = \frac{1}{n} \begin{pmatrix} \frac{\lambda \gamma}{n} & \mathbb{E}(u h - \frac{\lambda}{n} (q - \beta)) & -\mathbb{E}(v f + \frac{\gamma}{n} (q - \beta)) \\ 0 & \frac{1}{n} & 0 \\ 0 & 0 & \frac{1}{n} \end{pmatrix}, \end{aligned} \tag{A10}$$

whose inverse is

$$\mathbf{I}^{-1}(\theta) = \frac{n}{\gamma \lambda} \begin{pmatrix} 1 & -n \mathbb{E}(u h - \lambda (q - \beta)) & n \mathbb{E}(v f + \frac{\gamma}{n} (q - \beta)) \\ 0 & \gamma \lambda & 0 \\ 0 & 0 & \gamma \lambda \end{pmatrix}, \tag{A11}$$

and,

$$\mathbf{B}(\theta) = \mathbb{E} \psi \psi^T = \begin{pmatrix} \mathbb{E} \phi^2 & \mathbb{E} \phi \eta & \mathbb{E} \phi \Omega \\ \mathbb{E} \phi \eta & \mathbb{E} \eta^2 & \mathbb{E} \eta \Omega \\ \mathbb{E} \phi \Omega & \mathbb{E} \eta \Omega & \mathbb{E} \Omega^2 \end{pmatrix}. \tag{A12}$$

In order to estimate the variance of  $\hat{\beta}_{nAIPW}$ , we do not need to calculate all entries of the variance–covariance matrix, only the first entry:

$$\frac{1}{n} \left( \frac{n^2}{(\gamma\lambda)^2} \right) \begin{pmatrix} \mathbb{E}\phi^2 + \epsilon & \star & \star \\ \star & \star & \star \\ \star & \star & \star \end{pmatrix}. \quad (\text{A13})$$

The  $\star$  entries are irrelevant to the calculation of variance of nAIPW and the term  $\epsilon$  is a very long expression which involves terms converging to zero faster than the actual estimating Equation (A9) [19] (also verified by simulations):

$$\begin{aligned} \epsilon = & -\mathbb{E}\phi\eta(n\mathbb{E}uh + \lambda(\beta - q)) + \mathbb{E}\phi\Omega(n\mathbb{E}vf - \gamma(\beta - q)) - \\ & (n\mathbb{E}uh + \lambda(\beta - q))(-\mathbb{E}\eta^2(n\mathbb{E}uh + \lambda(\beta - q)) + \mathbb{E}\eta\Omega(n\mathbb{E}vf - \gamma(\beta - q)) + \\ & \mathbb{E}\phi\eta) + (n\mathbb{E}vf - \gamma(\beta - q))(-\mathbb{E}\eta\Omega(n\mathbb{E}uh + \lambda(\beta - q)) + \\ & \mathbb{E}\Omega^2(n\mathbb{E}vf - \gamma(\beta - q)) + \mathbb{E}\phi\Omega). \end{aligned} \quad (\text{A14})$$

Further,

$$\sqrt{n} \begin{pmatrix} \hat{\beta}_{nAIPW} \\ \hat{\gamma} \\ \hat{\Omega} \end{pmatrix} \sim MVN(\theta, \hat{\mathbf{I}}^{-1}(\hat{\theta})\hat{\mathbf{B}}(\hat{\theta})\hat{\mathbf{I}}^{-1}(\hat{\theta})^T) \quad (\text{A15})$$

where we replace  $\mathbb{E}$  with sample averages in Expressions (A10)–(A12) and  $\theta$  with their corresponding solutions to Equation (A8). Following this recipe, we obtain

$$\hat{\sigma}_{nAIPW}^2 = \frac{1}{n} \left( \frac{n^2}{(\gamma\lambda)^2} \right) \hat{\mathbb{E}}\phi^2 + \hat{\epsilon} \approx \sum_{i=1}^n \left( \frac{v_i f_i}{\hat{\gamma}} - \frac{u_i h_i}{\hat{\lambda}} + \frac{1}{n} q_i - \hat{\beta}_{nAIPW} \right)^2, \quad (\text{A16})$$

which is the same as (A4).

□

## References

1. Farrell, M.H.; Liang, T.; Misra, S. Deep neural networks for estimation and inference: Application to causal effects and other semiparametric estimands. *arXiv* **2018**, arXiv:1809.09953.
2. Goodfellow, I.; Bengio, Y.; Courville, A.; Bengio, Y. *Deep Learning*; MIT Press: Cambridge, UK, 2016; Volume 1.
3. Cheng, X.; Khomtchouk, B.; Matloff, N.; Mohanty, P. Polynomial regression as an alternative to neural nets. *arXiv* **2018**, arXiv:1806.06850.
4. Xu, B.; Wang, N.; Chen, T.; Li, M. Empirical evaluation of rectified activations in convolutional network. *arXiv* **2015**, arXiv:1505.00853.
5. Yan, S.; Nguang, S.K.; Gu, Z. H\_infinity Weighted Integral Event-Triggered Synchronization of Neural Networks With Mixed Delays. *IEEE Trans. Ind. Inform.* **2020**, *17*, 2365–2375. [[CrossRef](#)]
6. Robins, J.M.; Rotnitzky, A.; Zhao, L.P. Estimation of regression coefficients when some regressors are not always observed. *J. Am. Stat. Assoc.* **1994**, *89*, 846–866. [[CrossRef](#)]
7. Lunceford, J.K.; Davidian, M. Stratification and weighting via the propensity score in estimation of causal treatment effects: A comparative study. *Stat. Med.* **2004**, *23*, 2937–2960. [[CrossRef](#)]
8. Van der Laan, M.J.; Rose, S. *Targeted Learning: Causal Inference for Observational and Experimental Data*; Springer Science & Business Media: Cham, Switzerland, 2011.
9. Chernozhukov, V.; Chetverikov, D.; Demirer, M.; Duflo, E.; Hansen, C.; Newey, W.; Robins, J. Double/debiased machine learning for treatment and structural parameters. *Econom. J.* **2018**, *21*, C1–C68. [[CrossRef](#)]
10. Zhang, C.; Bengio, S.; Hardt, M.; Recht, B.; Vinyals, O. Understanding deep learning (still) requires rethinking generalization. *Commun. ACM* **2021**, *64*, 107–115. [[CrossRef](#)]
11. Myers, J.A.; Rassen, J.A.; Gagne, J.J.; Huybrechts, K.F.; Schneeweiss, S.; Rothman, K.J.; Joffe, M.M.; Glynn, R.J. Effects of adjusting for instrumental variables on bias and precision of effect estimates. *Am. J. Epidemiol.* **2011**, *174*, 1213–1222. [[CrossRef](#)]
12. Díaz, I. Doubly robust estimators for the average treatment effect under positivity violations: Introducing the  $e$ -score. *arXiv* **2018**, arXiv:1807.09148.
13. Horvitz, D.G.; Thompson, D.J. A generalization of sampling without replacement from a finite universe. *J. Am. Stat. Assoc.* **1952**, *47*, 663–685. [[CrossRef](#)]

14. Rubin, D.B. Using multivariate matched sampling and regression adjustment to control bias in observational studies. *J. Am. Stat. Assoc.* **1979**, *74*, 318–328.
15. Rosenbaum, P.R.; Rubin, D.B. The central role of the propensity score in observational studies for causal effects. *Biometrika* **1983**, *70*, 41–55. [[CrossRef](#)]
16. Rosenbaum, P.R.; Rubin, D.B. Constructing a control group using multivariate matched sampling methods that incorporate the propensity score. *Am. Stat.* **1985**, *39*, 33–38.
17. Busso, M.; DiNardo, J.; McCrary, J. New evidence on the finite sample properties of propensity score reweighting and matching estimators. *Rev. Econ. Stat.* **2014**, *96*, 885–897. [[CrossRef](#)]
18. Hahn, J. On the role of the propensity score in efficient semiparametric estimation of average treatment effects. *Econometrica* **1998**, *66*, 315–331. [[CrossRef](#)]
19. Hines, O.; Dukes, O.; Diaz-Ordaz, K.; Vansteelandt, S. Demystifying statistical learning based on efficient influence functions. *arXiv* **2021**, arXiv:2107.00681.
20. Słoczyński, T.; Wooldridge, J.M. A general double robustness result for estimating average treatment effects. *Econom. Theory* **2018**, *34*, 112–133. [[CrossRef](#)]
21. Scharfstein, D.O.; Rotnitzky, A.; Robins, J.M. Adjusting for nonignorable drop-out using semiparametric nonresponse models. *J. Am. Stat. Assoc.* **1999**, *94*, 1096–1120. [[CrossRef](#)]
22. Angrist, J.D.; Pischke, J.S. *Mostly Harmless Econometrics: An Empiricist's Companion*; Princeton University Press: Princeton, NJ, USA, 2008.
23. Van der Vaart, A.W. *Asymptotic Statistics*; Cambridge University Press: Cambridge, UK, 2000; Volume 3.
24. Chernozhukov, V.; Chetverikov, D.; Demirer, M.; Duflo, E.; Hansen, C.; Newey, W.K. *Double Machine Learning for Treatment and Causal Parameters*; Technical Report, Cemmap Working Paper; 2016. Available online: <https://ifs.org.uk/uploads/cemmap/wps/cwp491616.pdf> (accessed on 12 November 2021).
25. Farrell, M.H. Robust inference on average treatment effects with possibly more covariates than observations. *J. Econom.* **2015**, *189*, 1–23. [[CrossRef](#)]
26. Smucler, E.; Rotnitzky, A.; Robins, J.M. A unifying approach for doubly-robust l1 regularized estimation of causal contrasts. *arXiv* **2019**, arXiv:1904.03737.
27. Friedman, J.; Hastie, T.; Tibshirani, R. *The Elements of Statistical Learning*; Springer Series in Statistics; Springer: New York, NY, USA, 2001; Volume 1.
28. Stefanski, L.A.; Boos, D.D. The calculus of M-estimation. *Am. Stat.* **2002**, *56*, 29–38. [[CrossRef](#)]
29. Politis, D.N.; Romano, J.P. The stationary bootstrap. *J. Am. Stat. Assoc.* **1994**, *89*, 1303–1313. [[CrossRef](#)]
30. Kingma, D.P.; Ba, J. Adam: A method for stochastic optimization. *arXiv* **2014**, arXiv:1412.6980.