

## Research Article

Aaron Hudson\*, Elvin H. Geng, Thomas A. Odeny, Elizabeth A. Bukusi, Maya L. Petersen, and Mark J. van der Laan

# An approach to nonparametric inference on the causal dose–response function

<https://doi.org/10.1515/jci-2024-0001>

received January 08, 2024; accepted July 01, 2024

**Abstract:** The causal dose–response curve is commonly selected as the statistical parameter of interest in studies where the goal is to understand the effect of a continuous exposure on an outcome. Most of the available methodology for statistical inference on the dose-response function in the continuous exposure setting requires strong parametric assumptions on the probability distribution. Such parametric assumptions are typically untenable in practice and lead to invalid inference. It is often preferable to instead use nonparametric methods for inference, which only make mild assumptions about the data-generating mechanism. We propose a nonparametric test of the null hypothesis that the dose-response function is equal to a constant function. We argue that when the null hypothesis holds, the dose-response function has zero variance. Thus, one can test the null hypothesis by assessing whether there is sufficient evidence to claim that the variance is positive. We construct a novel estimator for the variance of the dose-response function, for which we can fully characterize the null limiting distribution and thus perform well-calibrated tests of the null hypothesis. We also present an approach for constructing simultaneous confidence bands for the dose-response function by inverting our proposed hypothesis test. We assess the validity of our proposal in a simulation study. In a data example, we study, in a population of patients who have initiated treatment for HIV, how the distance required to travel to an HIV clinic affects retention in care.

**Keywords:** dose–response function, continuous exposure, nonparametric testing, targeted minimum loss-based estimation

**MSC 2020:** 62G05, 62G10

## 1 Introduction

In many scientific studies, one of the main objectives is to use observational data to make inferences about the causal relationship between a treatment or exposure variable and some outcome. Many commonly-used for causal inference make strong, and often untenable, parametric assumptions about the data's probability distribution. Consequently, using such methods in practice can result in invalid inference under model misspecification. Interest has grown in instead using more robust nonparametric and semiparametric approaches for causal inference, which only make mild assumptions about the data-generating mechanism.

Our objective is to make nonparametric inference about the probability distribution of the counterfactual outcomes, or the collection of potential outcomes that would have been observed if a subject had received any

---

\* **Corresponding author: Aaron Hudson**, Vaccine and Infectious Disease Division, Fred Hutchinson Cancer Center, Seattle, WA, USA, e-mail: [ahudson@fredhutch.org](mailto:ahudson@fredhutch.org)

**Elvin H. Geng:** Division of Infectious Diseases, Department of Medicine, Washington University in St. Louis, St. Louis, Missouri, USA

**Thomas A. Odeny, Elizabeth A. Bukusi:** Research Care Training Program, Centre for Microbiology Research, Kenya Medical Research Institute, Nairobi, Kenya

**Maya L. Petersen, Mark J. van der Laan:** Division of Biostatistics, School of Public Health, University of California, Berkeley, California, United States

possible level of the exposure [1]. In many conventional analyses, it is of primary interest to make comparisons about the mean of the counterfactual outcome for different levels of the exposure. The target estimand is the function that maps any given exposure level to the corresponding mean counterfactual outcome. We commonly refer to this estimand as the *causal dose–response function*.

The literature on nonparametric inference on mean counterfactual outcomes most commonly focuses on the setting in which the exposure is binary. In this case, the mean counterfactual outcomes can, under standard causal assumptions, be characterized as pathwise differentiable estimands, or smooth functionals of the unknown probability distribution. For estimands that have this characterization, there exist many strategies for constructing nonparametric estimators that converge to the true quantity at the parametric rate and achieve a tractable sampling distribution. Commonly used estimators for the mean counterfactual outcomes with binary exposures that satisfy these properties include the augmented inverse probability weighted estimator [2] and the targeted maximum likelihood estimator [3]. When such estimators are available, it is straightforward to perform statistical inference on the mean counterfactual outcomes by constructing confidence sets or performing hypothesis tests. With minor modifications, these approaches for inference can be readily extended to the more general setting in which the exposure is a discrete random variable.

Nonparametric inference on the causal dose–response function is more challenging when the exposure is a continuous random variable. In this setting, the mean counterfactual outcomes are nonsmooth functions of the probability distribution, and consequently, nonparametric estimation is not possible at the parametric rate. There nonetheless exist proposals for consistent and rate-optimal estimation of the dose–response function in a nonparametric model. For instance, Díaz *et al.* [4] introduced a cross-validated targeted minimum loss (TML)-based estimator for the dose–response function, Kennedy *et al.* [5] and Colangelo and Lee [6] proposed nonparametric kernel smoothing estimators, and Westling *et al.* [7] proposed an estimator based on isotonic regression. Rate-optimal nonparametric estimators, such as those described above, typically retain a non-negligible asymptotic bias. Due to this bias retention, these estimators attain nonstandard limiting distributions, and constructing hypothesis tests or confidence sets based on the estimators can be challenging. In order to obtain an estimator for the dose–response curve that converges at the parametric rate, it is necessary to make parametric assumptions about either the conditional distribution of the exposure given the covariates [8–10] or the conditional mean of the outcome given exposure and covariates [11,12]. These approaches may not be preferred because specifying a parametric model correctly can be challenging in practice.

Nonparametric inference in the continuous treatment setting is a growing research area, and there have been many recent developments. For instance, Doss *et al.* [13] presented an omnibus and doubly robust test of the null of no treatment effect based on local linear regression, and an approach to confidence band construction using debiased local linear regression is proposed by Takatsu and Westling [14]. There also exist alternative approaches to nonparametric causal inference with continuous exposures which are based on the study of a finite-dimensional parameter that summarizes the effect of the exposure on the outcome [15–17]. While such summaries can be informative and carry meaningful interpretations, they may not suffice in settings where an investigator is primarily interested in learning about the dose–response function in its entirety.

Our work draws upon some recent advances in nonparametric inference on the dose–response function and inference on nonpathwise differentiable estimands more broadly. Westling [18] demonstrated that one can construct an omnibus nonparametric test of the hypothesis that the dose–response function is flat by simultaneously estimating primitive functions, or integrals, of the dose–response function. Hudson *et al.* [19] uses a generalization of this idea to develop a framework for testing hypotheses about nonpathwise differentiable estimands by estimating a suitably large collection of pathwise differentiable estimands that can effectively summarize the target. Hudson *et al.* [19] also propose a method for constructing simultaneous confidence sets for the target estimand by inverting the proposed hypothesis test.

In this article, we propose a novel method for inference on the dose–response function using the framework of Hudson *et al.* [19]. We develop a test of the null hypothesis that a mean-centered version of the dose–response function is equal to any given null function. Of particular importance is the instance where the null function is zero, in which case we perform a test of the null hypothesis that the dose–response function is flat. We also discuss the construction of simultaneous confidence sets for the centered dose–response function, and we describe a method for summarizing the confidence sets.

The remainder of the article is organized as follows. In Section 2, we provide a high-level overview of our proposed methodology, and in Section 3, we discuss our inferential procedure in more detail and describe its theoretical properties. In Section 4, we discuss implementation of our method. In Section 5, we study the behavior of our proposal in a simulation study. In Section 6, we apply our method to data from the Adaptive Strategies for Preventing and Treating Lapses of Retention in HIV Care (ADAPT-R) trial [20], which studies retention in care for people who initiated treatment for HIV. We conclude with a brief discussion in Section 7.

## 2 Overview

### 2.1 Identification of the dose–response function

Let  $O_1, \dots, O_n$  be *i.i.d.* random vectors drawn from a probability distribution  $P_0$ , which resides in a model  $\mathcal{M}$ . We allow  $\mathcal{M}$  to be a rich nonparametric model that is essentially unrestricted and must only satisfy some mild regularity conditions. We express our data as  $O = (W, A, Y)$ , where  $Y$  is a bounded real-valued outcome with sample space  $\mathcal{Y} \subset \mathbb{R}$ ,  $A$  is a bounded exposure variable with sample space  $\mathcal{A} \subset \mathbb{R}$ , and  $W$  is a  $q$ -dimensional vector of covariates, with sample space  $\mathcal{W} \subset \mathbb{R}^q$ . Throughout, we write  $E_P$  to denote the expectation under any probability distribution  $P \in \mathcal{M}$ , and we use the shorthand notation  $E_0 = E_{P_0}$ .

Let  $Y(a)$  denote the counterfactual outcome under exposure level  $a$ , or the potential outcome that would have been observed if an individual had been observed with exposure level  $A = a$ . We define the counterfactual mean outcome as the mapping  $P \mapsto \theta_P(a) = E_P[Y(a)]$  for  $a \in \mathcal{A}$ . Our objective is to perform inference on  $\theta_0 = \theta_{P_0}$ , which is commonly called the *causal dose–response function*.

The stochastic process of counterfactual outcomes for subject  $i$ ,  $\{Y_i(a) : a \in \mathcal{A}\}$  is not observable, as we only observe the outcome under one single exposure level  $A_i$ . However, under some standard causal assumptions, the counterfactual mean outcomes can nonetheless be estimated from the data. Let  $Q_P(w, a) = E_P[Y|W = w, A = a]$  denote the conditional mean of  $Y$  given the exposure and covariates, and let  $g_P(a|w) = \frac{d}{da}P(A \leq a|W = w)$  be the conditional density of the exposure given the covariates. Similarly as above, we use the shorthand notation  $Q_0 = Q_{P_0}$  and  $g_0 = g_{P_0}$ . The dose–response function can be estimated from the data under the following assumptions:

**Assumption A1 (Consistency).**  $A = a$  implies  $Y = Y(a)$ .

**Assumption A2 (No unmeasured confoundedness).**  $Y(a)$  is conditionally dependent on  $A$ , given  $W$ .

**Assumption A3 (Positivity).** There exists  $g_{\min} > 0$  so that  $\inf_{a \in \mathcal{A}} g_0(a|w) > g_{\min}$  for all  $w \in \mathcal{W}$ .

The consistency assumption states that the observed outcome for each subject is the potential outcome under their observed exposure level. The no unmeasured confoundedness assumption says that all confounding variables are contained within  $W$ . The positivity assumption means that it is possible for subjects with any covariate measurement to receive any treatment  $a \in \mathcal{A}$ . When Assumptions A1–A3 hold, the counterfactual mean outcome at any exposure level  $a$  can be expressed as

$$\theta_0(a) = E_{0,W}[Q(W, a)],$$

where  $E_{0,W}$  denotes the expectation with respect to the marginal distribution of  $W$  under  $P_0$ .

Assumptions A1–A3 typically hold in randomized trials, but they are generally unverifiable in observational studies. However, even when the causal assumptions fail,  $\theta_0$  retains an interpretation as a mean regression function, which can be used to study the conditional association between the exposure and the outcome, given the covariates.

## 2.2 Summary of proposed methodology

We first derive a test of the null hypothesis that the dose–response function, centered about its mean, is equal to a given candidate null function. Let  $\bar{\theta}_P : a \mapsto \theta_P(a) - E_P[\theta_P(A)]$  be the mean-centered dose–response function. For an arbitrary function  $\theta^*$  from  $\mathcal{A}$  to  $\mathbb{R}$ , let  $\bar{\theta}_P^* = \theta^* - E_P[\theta^*(A)]$ , and let  $\bar{\theta}_0^* := \bar{\theta}_{P_0}^*$ . We are interested in testing the null hypothesis that  $\bar{\theta}_0$  is equal to the candidate null parameter  $\bar{\theta}_0^*$ ,

$$H_0 : \bar{\theta}_0(a) = \bar{\theta}_0^*(a) \quad \text{for all } a \in \mathcal{A}. \quad (1)$$

Of particular interest in many applications is the case where  $\bar{\theta}_0^* = 0$ , which corresponds to the null hypothesis that the dose–response function is flat, or that the counterfactual mean outcomes are the same at every level of the exposure.

We also propose an approach for constructing confidence sets for the centered dose–response function. That is, for any  $\alpha \in (0, 1)$ , we construct a set  $C_n^\alpha$  that contain  $\bar{\theta}_0$  with probably at least  $1 - \alpha$  as  $n$  tends to infinity, i.e.,

$$\liminf_{n \rightarrow \infty} P_0(\bar{\theta}_0 \in C_n^\alpha) \geq 1 - \alpha.$$

When it is possible to obtain an estimator for  $\theta_0$  that has negligible asymptotic bias and a characterizable limiting distribution, it is straightforward to derive hypothesis tests and construct confidence sets based on the estimator. When the dose–response function is pathwise differentiable, meaning that  $\theta_P(a)$  changes smoothly in  $P$  with respect to local fluctuations around  $P_0$ , there are well-established theory and methodology for constructing  $n^{1/2}$ -consistent and asymptotically efficient estimators [21]. However,  $\theta_P(a)$  is typically only pathwise differentiable in a nonparametric model when the event  $\{A = a\}$  occurs with probability greater than zero. While this condition can be satisfied when  $A$  has a discrete support under  $P_0$ , it is generally not met when the exposure variable is continuous.

While the dose–response function is not pathwise differentiable when the exposure is continuous, one can perform inference by instead estimating a collection of pathwise differentiable functionals of the dose–response function, as suggested by, e.g., Westling [18] and Hudson et al. [19]. In this work, we present an approach based on the estimation of linear transformations of the causal dose–response function. In what follows, we briefly summarize our approach to inference.

For any function  $h$  from  $\mathcal{A}$  to  $\mathbb{R}$ , we define

$$\psi_{P, \theta^*}(h) := E_P[\{\bar{\theta}_P(A) - \bar{\theta}_P^*(A)\}h(A)]$$

as the  $L_2(P)$  inner product of  $\bar{\theta}_P - \bar{\theta}_P^*$  and  $h$ , and let  $\psi_{0, \theta^*}(h) := \psi_{P_0, \theta^*}(h)$ . Observe that if  $\bar{\theta}_0$  is equal to  $\bar{\theta}_0^*$  almost everywhere,  $\psi_{0, \theta^*}(h) = 0$  for all functions  $h$ , and if  $\bar{\theta}_0$  is not almost everywhere equal to  $\bar{\theta}_0^*$ , there exists a function  $h_*$  such that  $\psi_{0, \theta^*}(h_*) \neq 0$ . Therefore, one could test the null hypothesis in (1) by assessing whether, for a large class  $\mathcal{H}$  of bounded functions from  $\mathcal{A}$  to  $\mathbb{R}$ ,

$$\Psi_{0, \theta^*}(\mathcal{H}) := \sup_{h \in \mathcal{H}} |\psi_{0, \theta^*}(h)| = 0, \quad (2)$$

or in other words, whether there exists a linear transformation of  $\bar{\theta}_0 - \bar{\theta}_0^*$  that is nonzero.

The manner in which the supremum in (2) is calculated will depend on the construction of  $\mathcal{H}$ . For instance, if  $\mathcal{H}$  is countable, one may approximate the supremum as the maximum of a finite approximation of  $\{\psi_{0, \theta^*}(h) : h \in \mathcal{H}\}$ . We also consider the setting in which  $\mathcal{H}$  is constructed using a basis expansion, and the supremum corresponds to an optimal weighting of the coefficients for the basis functions.

We develop a test of the null hypothesis based on the estimation of  $\Psi_{0, \theta^*}(\mathcal{H})$ . We later show that  $\psi_{0, \theta^*}(h)$  is pathwise differentiable for any fixed  $h$  and can therefore be estimated at the parametric rate of  $n^{1/2}$ . If, in addition,  $\mathcal{H}$  is not too complex, one can construct an estimator  $\psi_{n, \theta^*}$  such that the standardized process  $\{n^{1/2}[\psi_{n, \theta^*}(h) - \psi_{0, \theta^*}(h)] : h \in \mathcal{H}\}$  converges weakly to a Gaussian process  $\mathbb{G} = \{\mathbb{G}(h) : h \in \mathcal{H}\}$ . As a consequence, we have that  $\Psi_{n, \theta^*}(\mathcal{H}) := \sup_{h \in \mathcal{H}} |\psi_{n, \theta^*}(h)|$  is a consistent estimator for  $\Psi_{0, \theta^*}(\mathcal{H})$ , and under the null, when  $\psi_{0, \theta^*}(h) = 0$  for all  $h$ ,  $n^{1/2}\Psi_{n, \theta^*}(\mathcal{H})$  converges weakly to  $\sup_{h \in \mathcal{H}} |\mathbb{G}(h)|$ . Because  $\Psi_{n, \theta^*}(\mathcal{H})$  has a tractable limiting distribution, it can be used as a test statistic. A hypothesis test that rejects the null when  $n^{1/2}\Psi_{n, \theta^*}(\mathcal{H})$

is larger than the  $1 - \alpha$  quantile of the null limiting distribution would achieve the type-1 error level  $\alpha$  asymptotically.

We can invert our proposed hypothesis test to obtain a confidence set for  $\bar{\theta}_0$ . Let  $\Theta$  be a large nonparametric class of functions from  $\mathcal{A}$  to  $\mathbb{R}$  that serve as candidate values for the dose–response function. Consider the set

$$C_n^\alpha = \{\bar{\theta}_0^* : \theta^* \in \Theta, \text{ and we fail to reject } \bar{\theta}_0 = \bar{\theta}_0^* \text{ at the level } \alpha \text{ based on } O_1, \dots, O_n\}. \quad (3)$$

We can interpret  $C_n^\alpha$  as the set of functions in  $\Theta$  that are in accordance with the observed data. If  $\theta_0$  belongs to  $\Theta$ , then  $\bar{\theta}_0$  belongs to  $C_n^\alpha$  with probability tending  $1 - \alpha$ , and  $C_n^\alpha$  is an asymptotically valid confidence set for  $\bar{\theta}_0$ . Even when  $\theta_0$  does not belong to  $\Theta$ ,  $C_n^\alpha$  retains a nice interpretation as long as  $\Theta$  contains a good approximation for  $\theta_0$ .

The confidence set  $C_n^\alpha$  can be difficult to visualize as it is a complex set of infinite-dimensional objects. We propose to simply summarize  $C_n^\alpha$  by displaying the smallest band that contains all functions in  $C_n^\alpha$ , which can be defined point-wise as

$$C_n^\alpha(a) = [\inf\{\bar{\theta}_0^*(a) : \bar{\theta}_0^* \in C_n^\alpha\}, \sup\{\bar{\theta}_0^*(a) : \bar{\theta}_0^* \in C_n^\alpha\}]. \quad (4)$$

As we discuss later in Section 3.2, this confidence band is generally conservative for correctly specified  $\Theta$ .

The inferential procedure described above requires selection of the class of functions  $\mathcal{H}$ , the elements of which index the set of linear transformations of  $\bar{\theta}_0 - \bar{\theta}_0^*$  that we need to estimate. Asymptotic type-1 error control is preserved using any choice  $\mathcal{H}$  that satisfies some mild regularity conditions, to be discussed later. The choice of  $\mathcal{H}$  does, however, influence the test's statistical power. We use the following intuition for constructing an  $\mathcal{H}$  so that our test is well-powered. Because the inner product  $\psi_{0,\theta^*}(h)$  is a measure of orthogonality of  $\bar{\theta}_0 - \bar{\theta}_0^*$  to a given  $h$ , the amount of evidence  $\psi_{0,\theta^*}(h)$  provided against the null depends on the *shape* of  $h$ , rather than the *scale*. It is therefore sensible to construct  $\mathcal{H}$  as a class of equally scaled functions that contains the function that is least orthogonal  $\bar{\theta}_0$ . We scale each  $h$  to have unit variance; while other measures of scale could be chosen, we scale by the standard deviation for simplicity. It can be shown by an application of the Cauchy-Schwarz inequality that the maximizer of  $|\psi_{0,\theta^*}(h)|$ , among all  $h$  with unit standard deviation, can be expressed as

$$h_0 = \text{Var}_0(\theta_0(A) - \theta^*(A))^{-1/2}(\bar{\theta}_0 - \bar{\theta}_0^*), \quad (5)$$

where  $\text{Var}_0$  is the population variance under  $P_0$ . The maximizer  $h_0$  is a scaled difference between the true dose–response function and the candidate null parameter, and the maximal inner product  $\psi_{0,\theta^*}(h_0)$  is equal to the standard deviation of  $\theta_0(A) - \theta_0^*(A)$ . We construct  $\mathcal{H}$  as a class of smooth functions that contain a good approximation of  $h_0$  so that our target estimand  $\Psi_{0,\theta^*}(\mathcal{H})$  can be interpreted as an approximation of the standard deviation of the difference between the truth and the null.

The hypothesis testing procedure proposed by Westling [18] can be viewed as a special case of our method with  $\mathcal{H}$  taken as the class of binary indicators of whether an input exposure level  $a$  is greater than any given cutoff  $a_0 \in \mathcal{A}$ . While Westling [18] showed that this choice results in an omnibus test of the null hypothesis, we later show that in finite samples, performance can be improved by considering a class of functions that contain the function that is least orthogonal to the difference between the true dose–response function and the null.

Our proposal is closely related to the nonparametric score test presented in the study of Hudson et al. [19]. The authors use the representation of function-valued nonpathwise differentiable parameters as the minimizer of a population risk functional to derive a set of estimating equations, indexed by functions  $h$  in a large class  $\mathcal{H}$ , that the true population parameter must satisfy. They then construct hypothesis tests that assesses whether the candidate null parameter also satisfies these estimating equations. Their proposal rejects the null hypothesis when the data provide contrary evidence. Our proposal can be viewed as a special case of this approach, where we assess whether  $\theta^*$  satisfies the estimating equations  $\psi_{0,\theta^*}(h) = 0$  for all  $h \in \mathcal{H}$ .

### 3 Inferential procedure

Having provided a brief overview of our inferential procedure in the previous section, we now provide theoretical details. In this section, we first discuss the estimation of  $\psi_{0,\theta^*}(h)$ , and we subsequently describe an approach for determining whether there is sufficient evidence to conclude whether  $\Psi_{0,\theta^*}(\mathcal{H})$  is zero.

#### 3.1 Estimation of $\psi_{0,\theta^*}(h)$

Recall that our proposal requires us to have an estimator  $\psi_n(h)$  of  $\psi_0(h)$  and a class of functions  $\mathcal{H}$  such that the process  $\{n^{1/2}[\psi_n(h) - \psi_0(h)] : h \in \mathcal{H}\}$  converges weakly to a Gaussian process. In this subsection, we specify conditions on  $\psi_{n,\theta^*}$  and  $\mathcal{H}$  such that the weak convergence property is satisfied, and we discuss the construction of a weakly convergent estimator.

As noted in Section 2, for any estimand that is pathwise differentiable in the sense of Bickel et al. [21], one can construct an estimator that, when centered around the true target estimand, converges weakly to a Gaussian distribution at the parametric rate of  $n^{1/2}$ . Constructing such an estimator and establishing efficiency typically require knowledge of the *efficient influence function* of the estimand of interest. The following lemma states that  $\psi_{p,\theta^*}(h)$  is pathwise differentiable in a nonparametric model and provides the form of the efficient influence function.

**Lemma 1.** *The parameter  $\psi_{p,\theta^*}(h)$  is pathwise differentiable in a nonparametric model, and its nonparametric efficient influence function is*

$$\begin{aligned} \phi_{p,\theta^*}(w, a, y; h) := & \left[ \bar{\theta}_p(a) - \bar{\theta}_p^*(a) + \frac{E_p[g_p(a|W)]}{g_p(a|w)} \{y - Q_p(w, a)\} \right] \{h(a) - E_p[h(A)]\} \\ & + E_p[Q_p(w, A)\{h(A) - E_p[h(A)]\}] - \{2E_p[\bar{\theta}_p(A)h(A)] - E_p[\bar{\theta}_p^*(A)h(A)]\}. \end{aligned}$$

As before, we use the shorthand notation  $\phi_{0,\theta^*} = \phi_{P_0,\theta^*}$  to denote the efficient influence function at  $P_0$ . Lemma 1 generalizes a result presented in [18] that characterizes the efficient influence function for the special case in which  $h$  is an indicator of whether the observed treatment level is greater than a specified cutoff, and  $\bar{\theta}_0^* = 0$ .

Because  $\psi_{0,\theta^*}(h)$  is pathwise differentiable, it is possible to construct an estimator  $\psi_{n,\theta^*}(h)$  that is asymptotically linear in the sense that

$$\psi_{n,\theta^*}(h) - \psi_{0,\theta^*}(h) = \frac{1}{n} \sum_{i=1}^n \phi_{P_0,\theta^*}(O_i; h) + r_n(h), \quad (6)$$

where  $r_n(h) = o_p(n^{-1/2})$ . In view of the central limit theorem and the fact that  $\phi_{0,\theta^*}$  has zero mean and finite variance, an asymptotically linear estimator  $\psi_{n,\theta^*}(h)$  is asymptotically Gaussian for any fixed  $h$ . However, our proposal requires a stronger notion of uniform convergence of  $\psi_{n,\theta^*}(h)$  for  $h$  in a large collection  $\mathcal{H}$ , so some additional conditions are needed. The following lemma, which is a consequence of Slutsky's theorem, provides conditions under which the desired uniform convergence holds.

**Lemma 2.** *Let  $\psi_{n,\theta^*}(h)$  be an asymptotically linear estimator of  $\psi_0(h)$  that has the representation in (6) for any  $h \in \mathcal{H}$ , and let  $\ell^\infty(\mathcal{H})$  denote the vector space of bounded real-valued functionals on  $\mathcal{H}$ . Assume the following conditions hold:*

- $\{\phi_{0,\theta^*}(\cdot; h) : h \in \mathcal{H}\}$  is a  $P_0$ -Donsker class,
- $\sup_{h \in \mathcal{H}} |r_n(h)| = o_p(n^{-1/2})$ .

*Then,  $\{n^{1/2}[\psi_{n,\theta^*}(h) - \psi_{0,\theta^*}(h)] : h \in \mathcal{H}\}$  converges weakly to a tight Gaussian process  $\mathbb{G}$  as an element of  $\ell^\infty(\mathcal{H})$ , where  $\mathbb{G}$  has mean zero and covariance  $\Sigma : (h_1, h_2) \mapsto E_0[\phi_{0,\theta^*}(O; h_1)\phi_{0,\theta^*}(O; h_2)]$ .*



The first condition of Lemma 2 is a constraint on the complexity of  $\mathcal{H}$  and typically holds when  $\mathcal{H}$  is a  $P_0$ -Donsker class. The second condition directly involves the estimator  $\psi_n$  and requires that the remainder term is asymptotically negligible in a uniform sense.

In what follows, we present two strategies for constructing weakly convergent estimators for  $\{\psi_{0,\theta^*}(h) : h \in \mathcal{H}\}$ . We begin by considering a naïve plug-in estimator of  $\psi_0(h)$ . Suppose we have a consistent estimator  $\hat{P}_n$  for  $P_0$ . In practice, we do not need to estimate the entire distribution  $P_0$ , and we must only estimate nuisance parameters upon which  $\psi_{0,\theta^*}$  and  $\phi_{0,\theta^*}$  depend. In our setting, there are four nuisance parameters: (i)  $F_{0,W}$ , the marginal cumulative distribution for  $W$ ; (ii)  $F_{0,A}$ , the marginal cumulative distribution function for  $A$ ; (iii)  $Q_0$ , the conditional mean of  $Y$  given  $A$  and  $W$ ; and (iv)  $g_0$ , the conditional density of  $A$  given  $W$ . One can obtain estimators  $F_{n,W}$  and  $F_{n,A}$  for  $F_{0,W}$  and  $F_{0,A}$  nonparametrically using the empirical distribution function, and one typically requires machine learning to construct consistent nonparametric estimators  $Q_n$  and  $g_n$  for  $Q_0$  and  $g_0$ . Given the initial estimator  $\hat{P}_n$ , one can obtain the naïve plug-in estimator  $\psi_{\hat{P}_n}(h)$ , which can be expressed as

$$\psi_{\hat{P}_n,\theta^*}(h) = n^{-1} \sum_{i=1}^n \left[ \theta_n(A_i) - \theta^*(A_i) - n^{-1} \sum_{j=1}^n \{ \theta_n(A_j) - \theta^*(A_j) \} \right] h(A_i),$$

where  $\theta_n : a \mapsto n^{-1} \sum_{i=1}^n Q_n(a, W_i)$  is the plug-in estimator for the dose–response function.

The plug-in estimator  $\psi_{\hat{P}_n,\theta^*}(h)$  will typically retain non-negligible asymptotic bias for  $\psi_{0,\theta^*}(h)$  and consequently will not be asymptotically linear. This bias is attributable to the fact that nonparametric estimators  $Q_n$  of  $Q_0$  are usually obtained by balancing a bias-variance tradeoff that is sub-optimal for the objective of estimating  $\psi_{0,\theta^*}(h)$ . We discuss two widely used strategies for correcting the bias of the naïve plug-in: one-step estimation [22] and TML-based estimation [23,24].

The estimation strategies we discuss require the following assumptions:

**Assumption B1.** There exists a  $P_0$ -Donsker class  $\Phi$  that contains  $\phi_{0,\theta^*}(\cdot; h)$  and  $\phi_{\hat{P}_n,\theta^*}(\cdot; h)$  for each  $h \in \mathcal{H}$  with probability tending to one.

**Assumption B2.** The nuisance parameter estimators satisfy

$$\begin{aligned} \int \{Q_n(w, a) - Q_0(w, a)\}^2 dP_0(w, a) &= o_p(1), \\ \int \{g_n(a|w) - g_0(a|w)\}^2 dP_0(w, a) &= o_p(1), \\ \int \{Q_n(w, a) - Q_0(w, a)\} \{g_n(a|w) - g_0(a|w)\} dP_0(w, a) &= o_p(n^{-1/2}). \end{aligned}$$

Assumptions B1 and B2 impose conditions on the estimators for the nuisance parameters upon which the target estimand and the efficient influence function depend. Assumption B1 places a constraint on the complexity of the family of candidate estimators for the nuisance components. Many flexible nonparametric estimators, e.g., those constructed via the highly adaptive LASSO [25], satisfy this condition. Assumption B2 places a requirement on the rates of convergence that the conditional mean and conditional density estimators must achieve. This condition holds when both estimators are consistent, and the product of the convergence rates is greater than  $n^{1/2}$ . In view of Assumption B2, the estimators of  $\psi_{0,\theta^*}(h)$  we develop are doubly robust in the sense that consistency and asymptotic normality are achieved if one of the nuisance parameters is estimated at a slow rate as long as the other nuisance is estimated at a fast enough rate to compensate.

We first construct a one-step estimator for  $\psi_{0,\theta^*}(h)$ . Given an estimator for the nuisance parameters upon which  $\phi_{0,\theta^*}(h)$  depends, we can obtain an estimator  $\phi_{\hat{P}_n,\theta^*}(\cdot; h)$  for the efficient influence function  $\phi_{0,\theta^*}(\cdot; h)$ . The empirical average of the estimator of the influence function  $\phi_{\hat{P}_n,\theta^*}$  can be shown to serve as a first-order approximation to the bias of the naïve plug-in [22]. This allows one to perform a, so-called, one-step bias correction. We define the one-step estimator as

$$\begin{aligned} \psi_{n,\theta^*}^I(h) &= \psi_{\hat{P}_n,\theta^*}(h) + n^{-1} \sum_{i=1}^n \phi_{\hat{P}_n,\theta^*}(O_i; h) \\ &= \psi_{\hat{P}_n,\theta^*}(h) + n^{-1} \sum_{i=1}^n \left[ \frac{n^{-1} \sum_{j=1}^n g_n(A_i|W_j)}{g_n(A_i|W_i)} \{Y_i - Q_n(W_i, A_i)\} \right] \left[ h(A_i) - n^{-1} \sum_{j=1}^n h(A_j) \right]. \end{aligned} \quad (7)$$

The following theorem states that, under mild regularity conditions, the one-step estimator is asymptotically linear and hence asymptotically Gaussian.

**Theorem 1.** *Under assumptions B1 and B2, the one-step estimator  $\psi_{n,\theta^*}^1(h)$  has the asymptotically linear representation in (6), with  $\sup_{h \in \mathcal{H}} |r_n(h)| = o_P(n^{-1/2})$ .*

While one-step estimators are asymptotically efficient, they are not guaranteed to be compatible in the sense that there exists a probability distribution  $P$  in the model  $\mathcal{M}$  such that  $\psi_{p,\theta^*}(h) = \psi_{n,\theta^*}^1(h)$  for all  $h$  in  $\mathcal{H}$ . TML-based estimation is an appealing alternative strategy that can be used to construct an estimator for  $\{\psi_{0,\theta^*}(h) : h \in \mathcal{H}\}$  that is compatible in this sense. TML estimators correct the bias of the naïve plug-in by updating the initial estimator  $\hat{P}_n$  of  $P_0$  to obtain a new estimator  $\tilde{P}_n$  such that the updated plug-in  $\{\psi_{\tilde{P}_n,\theta^*} : h \in \mathcal{H}\}$  that takes as input  $\tilde{P}_n$  has reduced bias for the target estimand  $\{\psi_{0,\theta^*}(h) : h \in \mathcal{H}\}$ . Therefore, as long as  $\tilde{P}_n$  resides within  $\mathcal{M}$ , the TML estimator is compatible. In our presentation, we only briefly summarize some of the main principles of targeted learning, and we refer readers to previous studies [23,24] for a comprehensive discussion.

The main idea behind TML-based estimation is to construct an updated estimator  $\tilde{P}_n$  based on the initial estimator  $\hat{P}_n$  so that the following efficient influence function estimating equations are satisfied:

$$\sup_{h \in \mathcal{H}} \left| n^{-1} \sum_{i=1}^n \phi_{\tilde{P}_n,\theta^*}(O_i; h) \right| = o_P(n^{-1/2}),$$

and such that  $\tilde{P}_n$  remains sufficiently close to  $\hat{P}_n$ , so as to remain a good estimator for  $P_0$ . Because estimating a marginal distribution function using the empirical distribution function does not generate bias for the target estimand, we do not need to update the initial estimators  $F_{A,n}$  and  $F_{W,n}$  for  $F_{A,0}$  and  $F_{W,0}$ . We only need to obtain an updated estimator  $\tilde{Q}_n$  for the conditional mean  $Q_0$ , since the initial estimator makes a bias-variance trade-off that is suboptimal for the estimation of the target parameter.

It can be verified algebraically that for any choice  $\tilde{Q}_n$ , the empirical average of the efficient influence function evaluated at  $\tilde{P}_n$  can be expressed as

$$n^{-1} \sum_{i=1}^n \phi_{\tilde{P}_n,\theta^*}(O_i; h) = n^{-1} \sum_{i=1}^n \{Y_i - \tilde{Q}_n(W_i, A_i)\} Z_n(W_i, A_i; h),$$

where we define  $Z_n(w, a; h)$  as

$$Z_n(w, a; h) = \frac{n^{-1} \sum_{i=1}^n g_n(a|W_i)}{g_n(a|w)} \left\{ h(a) - n^{-1} \sum_{i=1}^n h(A_i) \right\}. \quad (8)$$

Thus,  $\tilde{Q}_n$  satisfies the efficient influence function estimating equations at an adequate level if

$$\sup_{h \in \mathcal{H}} \left| n^{-1} \sum_{i=1}^n \{Y_i - \tilde{Q}_n(W_i, A_i)\} Z_n(W_i, A_i; h) \right| = o_P(n^{-1/2}). \quad (9)$$

We now discuss how to construct a  $\tilde{Q}_n$  that satisfies (9). Let  $Q_{n,\beta}$  be a parametric working model indexed by a scalar parameter  $\beta \in \mathbb{R}$ , for which  $Q_{n,\beta} = Q_n$  when  $\beta = 0$ . We construct the working model so that the derivative of the squared error loss is equal in magnitude to the supremum over  $\mathcal{H}$  of the empirical average of  $\phi_{\tilde{P}_n}(\cdot; h)$  with  $\tilde{Q}_n = Q_{n,\beta}$ , i.e.,

$$\left| \frac{d}{d\beta} \left[ (2n)^{-1} \sum_{i=1}^n \{Y_i - Q_{n,\beta}(W_i, A_i)\}^2 \right] \right| = \sup_{h \in \mathcal{H}} \left| n^{-1} \sum_{i=1}^n \{Y_i - Q_{n,\beta}(W_i, A_i)\} Z_n(W_i, A_i; h) \right|. \quad (10)$$

We then take  $\tilde{Q}_n = Q_{n,\beta_n}$ , where  $\beta_n$  is a near minimizer of the squared error loss and satisfies

$$n^{-1} \sum_{i=1}^n \{Y_i - Q_{n,\beta_n}(W_i, A_i)\}^2 = \inf_{\beta \in \mathbb{R}} \left[ n^{-1} \sum_{i=1}^n \{Y_i - Q_{n,\beta}(W_i, A_i)\}^2 \right] + \delta_n,$$



for a small positive sequence  $\delta_n \downarrow 0$ . Because  $\tilde{Q}_n$  is a near minimizer of the loss, we can see that this choice of  $\tilde{Q}_n$  satisfies (9) for  $\delta_n$  sufficiently small. We note that a sub-model satisfying (10) is referred to as a *universal least favorable submodel* and provides the maximal reduction of the bias of  $\{\psi_{\tilde{p}_n, \theta^*}(h) : h \in \mathcal{H}\}$  as  $Q_{n, \beta}$  moves away from  $Q_n$  toward  $Q_{n, \beta_n}$ . Strategies based on a locally least favorable submodel, which would satisfy (10) only when  $\beta = 0$ , could alternatively have been considered, but such approaches tend to perform worse in small samples, in particular when the target estimand is multidimensional or infinite-dimensional [26]. We recursively define the universal least favorable sub-model  $Q_{n, \beta}$  point-wise as

$$Q_{n, \beta}(w, a) = \begin{cases} Q_n(w, a) + \int_0^\beta Z_n(a, w; h_{n, b}) db, & \beta \geq 0 \\ Q_n(w, a) - \int_\beta^0 Z_n(w, a; h_{n, b}) db, & \beta < 0, \end{cases} \quad (11)$$

where  $h_{n, \beta}$  is a solution to

$$\left| n^{-1} \sum_{i=1}^n \{Y_i - Q_{n, \beta}(W_i, A_i)\} Z_n(W_i, A_i; h_{n, \beta}) \right| = \sup_{h \in \mathcal{H}} \left| n^{-1} \sum_{i=1}^n \{Y_i - Q_{n, \beta}(W_i, A_i)\} Z_n(W_i, A_i; h) \right|.$$

It can be verified using the fundamental theorem of calculus that the above working model satisfies (10).

After obtaining the updated estimator  $\tilde{Q}_n$ , we can construct the TML estimator as

$$\psi_{n, \theta^*}^{\text{II}}(h) = \psi_{\tilde{p}_n, \theta^*}(h) = n^{-1} \sum_{i=1}^n \left[ \theta_{\tilde{p}_n}(A_i) - \theta^*(A_i) - n^{-1} \sum_{j=1}^n \{ \theta_{\tilde{p}_n}(A_j) - \theta^*(A_j) \} \right] h(A_i), \quad (12)$$

where  $\theta_{\tilde{p}_n} : a \mapsto n^{-1} \sum_{i=1}^n \tilde{Q}_n(W_i, a)$  is the updated TML estimator for the dose-response function. The following theorem states that the TML estimator satisfies the conditions of Lemma 2.

**Theorem 2.** Under Assumptions B1 and B2, the TML-based estimator  $\psi_{n, \theta^*}^{\text{II}}(h)$  has the asymptotically linear representation in (6) with  $\sup_{h \in \mathcal{H}} |r_n(h)| = o_p(n^{-1/2})$ .

### 3.2 Inference on $\Psi_{0, \theta^*}(\mathcal{H})$

We are at this point prepared to discuss inference on  $\Psi_{0, \theta^*}(\mathcal{H}) = \sup_{h \in \mathcal{H}} |\psi_{0, \theta^*}(h)|$ . Suppose that we have an estimator  $\{\psi_{n, \theta^*}(h) : h \in \mathcal{H}\}$  for  $\{\psi_{0, \theta^*}(h) : h \in \mathcal{H}\}$  and a function class  $\mathcal{H}$  that satisfy the conditions of Lemma 2. Such an estimator could be obtained using either of the strategies presented in Section 3.1. Consider the plug-in estimator  $\Psi_{n, \theta^*}(\mathcal{H}) = \sup_{h \in \mathcal{H}} |\psi_{n, \theta^*}(h)|$ . The continuous mapping theorem implies that  $n^{1/2} \sup_{h \in \mathcal{H}} |\psi_{n, \theta^*}(h) - \psi_{0, \theta^*}(h)|$  converges weakly to  $\sup_{h \in \mathcal{H}} |\mathbb{G}(h)|$  where  $\mathbb{G}$  is the Gaussian process in Lemma 2. This in combination with the reverse triangle inequality together imply that

$$|\Psi_{n, \theta^*}(\mathcal{H}) - \Psi_{0, \theta^*}(\mathcal{H})| = O_p(n^{-1/2}).$$

Furthermore, because when the null hypothesis holds,  $\psi_{0, \theta^*}(h) = 0$  for all  $h$ ,  $n^{1/2} \Psi_{n, \theta^*}(\mathcal{H})$  converges weakly to  $\sup_{h \in \mathcal{H}} |\mathbb{G}(h)|$ . Thus,  $\Psi_{n, \theta^*}(\mathcal{H})$  is a consistent estimator for  $\Psi_{0, \theta^*}(\mathcal{H})$  that has a fully characterizable null limiting distribution. This makes it possible to construct an asymptotically valid hypothesis test based on the estimator.

We make this explicit by providing the form of our proposed test below. Letting  $t_{1-\alpha}^*$  denote the  $1 - \alpha$  quantile of the null limiting distribution of  $n^{1/2} \Psi_{n, \theta^*}(\mathcal{H})$ , we consider the hypothesis test

$$T(O_1, \dots, O_n) = \begin{cases} \text{"Reject"} & n^{1/2} \Psi_{n, \theta^*}(\mathcal{H}) > t_{1-\alpha}^* \\ \text{"Fail to reject"} & n^{1/2} \Psi_{n, \theta^*}(\mathcal{H}) \leq t_{1-\alpha}^*. \end{cases}$$

This test achieves type-1 error control at the level  $\alpha$  when the conditions of Lemma 2 are satisfied. Additionally, when  $\mathcal{H}$  contains a function that is nonorthogonal to  $\bar{\theta}_0^*$  (so that  $\Psi_{0,\theta^*}(\mathcal{H}) > 0$ ), the above test is consistent in the sense that the probability of rejecting the null tends to one. To see this, we observe that  $n^{1/2}\Psi_{n,\theta^*}(\mathcal{H})$  tends to infinity because  $\Psi_{n,\theta^*}(\mathcal{H})$  converges to  $\Psi_{0,\theta^*}(\mathcal{H})$  in probability.

Moreover, the fact that our proposed test achieves nominal type-1 error control asymptotically implies that the confidence band in (4) is a uniform confidence band. This follows from the facts that

$$P_0(\bar{\theta}_0(a) \in C_n^a(a) \text{ for all } a \in \mathcal{A}) \geq P_0(\bar{\theta}_0 \in C_n^a),$$

and that  $P_0(\bar{\theta}_0 \in C_n^a)$  approaches  $1 - \alpha$  in the limit of large  $n$  when  $\Theta$  contains  $\theta_0$ . This confidence band will in general achieve a coverage rate not lower than  $1 - \alpha$ , though this lower bound is not tight. This type of confidence band construction has been shown to result in over-coverage by Hudson et al. [19].

While the null limiting distribution of  $\Psi_{n,\theta^*}(\mathcal{H})$  can indeed be characterized, a closed form expression may not be available. It may therefore be necessary to use an approximation. We use the multiplier bootstrap method presented in the study of Hudson et al. [19], which makes use of the asymptotic linearity of  $\psi_{n,\theta^*}(h)$ . We first note that due to the uniform asymptotic linearity of  $\psi_{n,\theta^*}$ , under the null hypothesis,  $\Psi_{n,\theta^*}(\mathcal{H})$  can be expressed as the sum of the supremum of an empirical process and an asymptotically negligible remainder. That is,

$$\Psi_{n,\theta^*}(\mathcal{H}) = \sup_{h \in \mathcal{H}} \left| n^{-1} \sum_{i=1}^n \phi_{0,\theta^*}(O_i; h) \right| + o_P(n^{-1/2}). \quad (13)$$

We can approximate the null distribution of  $\Psi_{n,\theta^*}(\mathcal{H})$  as the supremum of a bootstrapped empirical process that attains the same limiting distribution as the empirical process in (13), conditional on the observed data. For  $m = 1, \dots, M$  and  $M$  large, let  $\xi_1^m, \dots, \xi_n^m$  be i.i.d. random variables, independent of  $O_1, \dots, O_n$ , with mean zero, unit variance, and  $E_0[|\xi_i^m|^{2+u}] < \infty$  for some  $u > 0$ , and let  $\phi_{n,\theta^*}$  be an estimator for the efficient influence function. We define the  $m$ th bootstrap sample of  $\Psi_{n,\theta^*}(\mathcal{H})$  as

$$\Psi_{n,\theta^*}^m(\mathcal{H}) = \sup_{h \in \mathcal{H}} \left| n^{-1} \sum_{i=1}^n \xi_i^m \phi_{n,\theta^*}(O_i; h) \right|. \quad (14)$$

It is shown in the study of Hudson et al. [19] that if  $\phi_{n,\theta^*}$  is a consistent estimator for  $\phi_{0,\theta^*}$ , and if  $\mathcal{H}$  is not overly complex, the multiplier bootstrap statistic converges weakly to  $\sup_{h \in \mathcal{H}} |\mathbb{G}(h)|$ , conditional on  $O_1, \dots, O_n$ . Thus, the distribution of the multiplier bootstrap samples closely approximates the null limiting distribution of  $\Psi_{n,\theta^*}(\mathcal{H})$  in the limit of large  $n$ .

To estimate the efficient influence function, one approach is to use the plug-in estimator

$$\begin{aligned} \phi_{n,\theta^*}(w, a, y; h) = & \left[ \theta_n(a) - \theta^*(a) - n^{-1} \sum_{i=1}^n \{ \theta_n(A_i) - \theta^*(A_i) \} \right] \left[ h(a) - n^{-1} \sum_{i=1}^n h(A_i) \right] \\ & + \left[ \frac{n^{-1} \sum_{i=1}^n g_n(a|W_i)}{g_n(a|w)} \{ y - Q_n(w, a) \} \right] \left[ h(a) - n^{-1} \sum_{i=1}^n h(A_i) \right] \\ & + n^{-1} \sum_{i=1}^n Q_n(w, A_i) \left[ h(A_i) - n^{-1} \sum_{j=1}^n h(A_j) \right] - n^{-1} \sum_{i=1}^n \{ Y_i - Q_n(W_i, A_i) \} Z_n(W_i, A_i; h), \end{aligned} \quad (15)$$

where we recall that  $\theta_n : a \mapsto n^{-1} \sum_{i=1}^n Q_n(W_i, a)$  is the plug-in estimator for  $\theta_0$ , and  $Z_n$  is as defined in (8). Alternatively, we can observe that because  $\bar{\theta}_0 = \bar{\theta}^*$  under the null, in (15), we can replace  $\theta^*$  with an estimator of  $\theta_0$ , such as the plug-in  $\theta_n$ . We note that when we substitute  $\theta^*$  by  $\theta_n$  in (15), some cancellation occurs, and the first line in the above expression vanishes. This strategy of replacing  $\theta^*$  with an estimator for  $\theta_0$  is appealing because the bootstrap approximation of the limiting distribution no longer depends on  $\theta^*$ , so one can test any hypothesis of the form (1) using the same bootstrap sample. This is particularly useful when we are interested

in constructing a confidence set for  $\bar{\theta}_0$  by inverting our proposed test, as this requires us to test a large collection of hypotheses.

In our presentation so far, we have assumed that the class  $\mathcal{H}$  is fixed. We acknowledge that in practice, fixing a class *a priori* may be challenging, so data-adaptive approaches for selecting  $\mathcal{H}$  may be preferable. It has been shown by Hudson et al. [19] that data-adaptive selection of  $\mathcal{H}$  does not affect the type-1 error rate of our proposed test as long as the data-adaptive choice converges to a fixed class. In Section 4, we propose an approach for data-adaptive selection of  $\mathcal{H}$ , and we later show in simulations that our approach is asymptotically valid.

## 4 Implementation

### 4.1 Construction of $\mathcal{H}$

Recall that it is our objective to construct  $\mathcal{H}$  as a model for  $h_0$  in (5) so that  $\Psi_{0,\theta^*}(\mathcal{H})$  can be interpreted as the standard deviation of the difference between the true dose–response curve and the candidate null parameter. As discussed in Section 3.2, we can choose a flexible nonparametric model, so long as the model is not overly complex, and the conditions of Lemma 2 are satisfied. In what follows, we describe a practical approach for selecting such a class. Our approach is similar to that used by Hudson et al. [19] to implement their proposed nonparametric score test. In what follows, we describe an approach for constructing  $\mathcal{H}$  as a subspace of a reproducing kernel Hilbert space (RKHS). This construction of  $\mathcal{H}$  is justified because RKHSs are known to be Donsker classes [27].

For a positive-semidefinite kernel function  $K$  from  $\mathcal{A} \times \mathcal{A}$  to  $\mathbb{R}^+$ , let  $\mathcal{S}_K$  denote its unique RKHS, endowed with the inner product  $\langle \cdot, \cdot \rangle_{\mathcal{S}_K}$ . The kernel function  $K$  has the eigen-decomposition

$$(a_1, a_2) \mapsto K(a_1, a_2) = \sum_{d=1}^{\infty} \gamma_d \eta_d(a_1) \eta_d(a_2),$$

where the eigenfunctions  $\{\eta_1, \eta_2, \dots\}$  are orthogonal with respect to the RKHS inner product  $\langle \cdot, \cdot \rangle_{\mathcal{S}_K}$ , and  $0 \leq \gamma_1 \leq \gamma_2 \leq \dots$  are the eigenvalues. Any function  $s$  in the RKHS can be expressed as a linear combination of the eigenfunctions. That is, there exist coefficients  $c_1, c_2, \dots$  such that  $s(a) = \sum_{d=1}^{\infty} c_d \eta_d(a)$  for all  $a$ . The roughness of  $s$  can be measured by the RKHS norm as

$$J(s) = \langle s, s \rangle_{\mathcal{S}_K} = \sum_{d=1}^{\infty} \frac{c_d^2}{\gamma_d},$$

with higher values of  $J(s)$  corresponding to greater roughness. We construct  $\mathcal{H}$  as a subset of functions in  $\mathcal{S}_K$  with bounded roughness and unit variance. That is

$$\mathcal{H}_K = \left\{ h = \sum_{d=1}^{\infty} c_d \eta_d : c_1, c_2, \dots \in \mathbb{R}, J(h) \leq \kappa, \text{Var}_n(h(A)) = 1 \right\},$$

where  $\text{Var}_n(h(A))$  is the empirical variance of  $h(A)$ , and  $\kappa > 0$  is a tuning parameter. To facilitate computation, we truncate the eigenbasis at some large level  $D$ .

In our implementation, we select  $\mathcal{S}_K$  as the second-order Sobolev space on  $[0, 1]$ , which can be defined as an RKHS endowed with the inner product  $(s_1, s_2) \mapsto \langle s_1, s_2 \rangle_{\mathcal{S}_K} = \int_0^1 \tilde{s}_1(u) \tilde{s}_2(u) du$ , where  $\tilde{s}$  denotes the second derivative of any given function  $s$ . In this case, the eigenfunctions and eigenvalues are available in closed form and can be expressed as

$$\eta_{2d-1} : a \mapsto \sqrt{2} \cos(2\pi da), \quad \eta_{2d} : a \mapsto \sqrt{2} \sin(2\pi da), \quad \gamma_{2d-1} = \gamma_{2d} = (2\pi d)^{-4},$$

for  $d = 1, 2, \dots$  [28].

We conclude by discussing selection of the tuning parameter  $\kappa$ . Our goal is to select  $\kappa$  large enough so that  $\mathcal{H}_\kappa$  contains a good approximation of an  $h_0$ . Suppose  $h_0$  belongs into the RKHS  $\mathcal{S}_K$ . With prior knowledge on  $h_0$  of available, a natural choice would be to set  $\kappa = \kappa_0$ , where we define

$$\kappa_0 = J(h_0) = \frac{J(\theta_0 - \theta^*)}{\text{Var}_0(\theta_0(A) - \theta^*(A))}.$$

Because  $\kappa_0$  is typically unknown as it depends on  $\theta_0 - \theta^*$ , we may in practice rely upon an estimate.

We propose to use a simple plug-in estimator for  $\kappa_0$ . Consider the following transformation of the observed data:

$$f_n(O_i) = n^{-1} \sum_{j=1}^n Q_n(W_j, A_i) + \frac{n^{-1} \sum_{j=1}^n g_n(A_i | W_j)}{g_n(A_i | W_i)} \{Y_i - Q_n(W_i, A_i)\} - \theta^*(A_i).$$

It is shown in [5] that one can consistently estimate  $\theta_0 - \theta^*$  by regressing  $f_n(O_i)$  on  $A$ . We estimate  $\theta_0 - \theta^*$  as  $c_{0,n} + \sum_{d=1}^D c_{d,n} \eta_d$ , where the coefficients are the minimizers of a penalized least squares loss, namely

$$c_{0,n}, c_{1,n}, \dots, c_{d,n} = \arg \min_{c_0, c_1, \dots, c_d \in \mathbb{R}} n^{-1} \sum \left\{ f_n(O_i) - c_0 - \sum_{d=1}^D c_d \eta_d(A_i) \right\}^2 + \lambda \sum_{d=1}^D \frac{c_d^2}{\gamma_d}, \quad (16)$$

where  $\lambda > 0$  is a tuning parameter. The penalty term in (16) controls the RKHS norm of the resulting estimate, with smaller values of  $\lambda$  corresponding to a less smooth estimate. To select  $\lambda$ , we perform cross-validation for a large set of candidate values, and we choose the largest candidate for which the cross-validation error is within one standard error of the minimum cross-validation error. This strategy provides a parsimonious estimate of  $\theta_0 - \theta^*$  that fits the observed data well. In practice, the resulting estimate will often be less rough than  $\theta_0 - \theta^*$  [29]. Finally, we estimate  $h_0$  as

$$h_n = \frac{\sum_{d=1}^D c_{d,n} \left\{ \eta_j - n^{-1} \sum_{i=1}^n \eta_j(A_i) \right\}}{\text{Var}_n^{1/2} \left( \sum_{d=1}^D c_{d,n} \eta_d(A) \right)},$$

and we estimate  $\kappa_0$  as

$$\kappa_n = J(h_n) = \frac{\sum_{d=1}^D \frac{c_{d,n}^2}{\gamma_d}}{\text{Var}_n \left( \sum_{d=1}^D c_{d,n} \eta_d(A) \right)}.$$

## 4.2 Calculation of $\Psi_{n,\theta^*}(\mathcal{H})$

We now describe how to calculate  $\psi_{n,\theta^*}(h)$  and  $\Psi_{n,\theta^*}(\mathcal{H})$ . It is first necessary to estimate the nuisance parameters  $Q_0$  and  $g_0$ . The conditional mean  $Q_0$  can be estimated using any of a wide variety of flexible nonparametric estimators, such as artificial neural networks [30], the highly adaptive LASSO [25], or the Super Learner [31]. In this work, we use the highly adaptive LASSO, which is implemented in the publicly available R package `hal9001`.

To construct a nonparametric estimator for the conditional density function  $g_0$ , we first observe that  $g_0$  can be approximated by a conditional mean function. Let  $\pi$  be a non-negative and symmetric function from  $\mathbb{R}$  to  $\mathbb{R}^+$  for which  $\int_{\mathbb{R}} \pi(u) du = 1$ . We define  $g_{0,r}$  as

$$g_{0,r}(a|w) = E_0[r^{-1} \pi(r^{-1}\{A - a\}) | W = w],$$

where  $r > 0$  is a bandwidth. It can be shown that  $g_{0,r}(a|w)$  tends to  $g_0(a|w)$  as  $r$  tends to zero. It is therefore sensible to estimate  $g_0(a|w)$  using an estimator for the conditional mean  $g_{0,r}(a|w)$  for sufficiently small  $r$ . We treat the bandwidth  $r$  as a tuning parameter that modulates the smoothness of the conditional density estimate

in  $a$ , with smaller  $r$  corresponding to lesser smoothness. This estimator can be viewed as a generalization of the kernel density estimator for learning a marginal density function.

For a given bandwidth  $r$  and a fine grid of fixed points  $a_1 < a_2 < \dots$ , one can estimate each  $g_{0,r}(a_j|\cdot)$  using a flexible nonparametric estimator for the conditional mean. One can then obtain estimates at intermediate points via linear interpolation. To ensure that the conditional density estimate is non-negative, we fit a flexible nonparametric model for  $\log g_{0,r}$  using the highly adaptive LASSO and transform the resulting model fit, similarly as one would estimate a conditional mean in a generalized linear model with a log link. The bandwidth  $r$  can be selected by performing cross-validation using the log loss function. We note that performing cross-validation can be very slow as estimating  $g_{0,r}(a_j|\cdot)$  at a large number of grid points for several different choices of bandwidth can be computationally intensive. An alternative strategy that we suggest is to consider a kernel density estimator for the marginal density of  $A$  and to select  $r$  as the bandwidth for the kernel density estimator that minimizes the cross-validation error for the marginal density. We expect this approach to perform reasonably well as long as the conditional density of  $A$  at any given  $w$  is not much less smooth than the marginal density.

We now discuss calculation of the one-step and TML estimators for  $\psi_{0,\theta^*}(h)$ , given that estimators  $Q_n$  and  $g_n$  for  $Q_0$  and  $g_0$  are available. One can calculate the one-step estimator at  $h = \eta_d$  as

$$\begin{aligned} \psi_{n,\theta^*}^I(\eta_d) = & n^{-1} \sum_{i=1}^n \left[ \theta_n(A_i) - \theta^*(A_i) - n^{-1} \sum_{k=1}^n \{ \theta_n(A_k) - \theta^*(A_k) \} \right] \left[ \eta_j(A_i) - n^{-1} \sum_{k=1}^n \eta_d(A_k) \right] \\ & + n^{-1} \sum_{i=1}^n \left[ \frac{n^{-1} \sum_{k=1}^n g_n(A_i|W_k)}{g_n(A_i|W_i)} \right] \{ Y_i - Q_n(W_i, A_i) \} \left[ \eta_j(A_i) - n^{-1} \sum_{k=1}^n \eta_d(A_k) \right], \end{aligned}$$

where, as in Section 3.1,  $\theta_n : a \mapsto n^{-1} \sum_{i=1}^n Q_n(a, W_i)$  is the plug-in estimator for the dose-response function. Because the one-step estimator is linear in  $h$ , it easy to see that for any  $h = \sum_j c_d \eta_d$ ,  $\psi_n^I(h)$  can be expressed as

$$\psi_n^I(h) = \sum_{d=1}^D c_d \psi_n^I(\eta_d).$$

Computing the TML estimator is more involved than computing the one-step estimator as we need to calculate the TML update  $\tilde{Q}_n$  of the initial conditional mean estimator  $Q_n$ . Recall from Section 3.1 that we take  $\tilde{Q}_n$  as the minimizer of the squared error loss along the parametric working model  $Q_{n,\beta}$  in (11). For a small  $\varepsilon > 0$  and a positive integer  $B$ , we approximate  $Q_{n,\beta}$  at  $\beta = B\varepsilon$  as

$$Q_{n,B\varepsilon}(w, a) = Q_n(w, a) + \varepsilon \sum_{b=1}^B Z_n(a, w; h_{n,b\varepsilon}),$$

where we define  $h_{n,b\varepsilon}$  as

$$h_{n,b\varepsilon} = \arg \max_{h \in \mathcal{H}_\kappa} n^{-1} \sum_{i=1}^n \{ Y_i - Q_{n,\varepsilon(b-1)}(A_i, W_i) \} Z_n(A_i, W_i; h). \quad (17)$$

The optimization problem in (18) is a quadratically constrained quadratic program and can be solved by using publicly available software such as the CVXR package in R [32]. One can observe that because  $\mathcal{H}_\kappa$  is symmetric in the sense that  $h \in \mathcal{H}_\kappa$  implies  $-h \in \mathcal{H}_\kappa$ , the derivative of the squared error loss,

$$D_n(\beta) = -n^{-1} \sum_{i=1}^n \{ Y_i - Q_{n,\beta}(A_i, W_i) \} Z_{i,n}(h_{n,\beta}), \quad (18)$$

is necessarily nonpositive. One can therefore find a near-minimizer of the squared error loss by calculating  $Q_{n,B\varepsilon}$  for incrementally increasing  $B$  until  $D_n(B\varepsilon)$  is sufficiently small. We take  $\tilde{Q}_n = Q_{B_n\varepsilon}$ , where  $B_n$  satisfies

$$D_n(B_n\varepsilon) = \{ n \log(n) \}^{-1/2} \text{Var}_n^{-1/2}(\{ Y - Q_n(W, A) \} Z_n(W, A; h_n)),$$

where  $h_n$  is the estimator for  $h_0$  described in Section 4.1. This choice of  $B_n$  ensures that  $D_n(B_n\epsilon)$  approaches zero at a rate faster than  $n^{-1/2}$ , which is a key condition for establishing asymptotic linearity of the TML estimator. Additionally, choosing  $B_n$  so that  $D_n(B_n\epsilon)$  is not much smaller than necessary and tends to zero at only a slightly faster rate than  $n^{1/2}$  helps to prevent  $\tilde{Q}_n$  from being an overfitted estimator of  $Q_0$ .

Now, for any  $\eta_d$ , the TML estimator of  $\psi_{0,\theta^*}(\eta_d)$  can be expressed as

$$\psi_{n,\theta^*}^{\Pi}(\eta_d) = n^{-1} \sum_{i=1}^n \left[ \theta_{\tilde{h}_n}(A_i) - \theta^*(A_i) - n^{-1} \sum_{k=1}^n \{ \theta_{\tilde{h}_n}(A_k) - \theta^*(A_k) \} \right] \eta_d(A_i),$$

where we recall that  $\theta_{\tilde{h}_n} : a \mapsto n^{-1} \sum_{i=1}^n \tilde{Q}_n(a, W_i)$  is the updated TML estimator for the dose-response function. Because the TML estimator is linear in  $h$ , for any  $h = \sum_{d=1}^D c_d \eta_d$ , we have

$$\psi_{n,\theta^*}^{\Pi}(h) = \sum_{d=1}^D c_d \psi_{n,\theta^*}^{\Pi}(\eta_d).$$

Having described how to compute the one-step and TML estimators for  $\psi_{0,\theta^*}(h)$ , we now discuss how to calculate  $\Psi_{n,\theta^*}(\mathcal{H}_K)$ . Observe that for  $h = \sum_{d=1}^D c_d \eta_d$ , the one-step and TML estimators are linear in the coefficient vector  $\mathbf{c} = (c_1, \dots, c_D)^T$  and can be expressed as  $U(\theta^*)^T \mathbf{c}$ , where  $U(\theta^*)$  is a  $D$ -dimensional vector for which the  $d$ th element contains an estimator for  $\psi_{0,\theta^*}(\eta_d)$ . Let  $V$  be a  $D \times D$  matrix where element  $(d_1, d_2)$  is

$$V_{d_1, d_2} = n^{-1} \sum_{i=1}^n \left\{ \eta_{d_1}(A_i) - n^{-1} \sum_{j=1}^n \eta_{d_1}(A_j) \right\} \left\{ \eta_{d_2}(A_i) - n^{-1} \sum_{j=1}^n \eta_{d_2}(A_j) \right\}$$

so that the empirical variance of  $h(A)$  is  $\text{Var}_n(h(A)) = \mathbf{c}^T V \mathbf{c}$ , and let  $\Gamma = \text{diag}(\gamma_1^{-1}, \dots, \gamma_D^{-1})$ , where  $\gamma_1, \dots, \gamma_D$  are the eigenvalues for the kernel  $K$ . We can express  $\Psi_{n,\theta^*}(\mathcal{H}_K)$  as  $U(\theta^*)^T \mathbf{c}_n$ , where  $\mathbf{c}_n$  is defined as

$$\mathbf{c}_n = \arg \max \{ U(\theta^*)^T \mathbf{c} : \mathbf{c}^T V \mathbf{c} = 1, \mathbf{c}^T \Gamma \mathbf{c} \leq \gamma \}. \quad (19)$$

The Karush-Kuhn-Tucker conditions for the optimization problem in (19) imply that  $\mathbf{c}_n$  is the solution to

$$U(\theta^*) - \lambda_1 (V + \lambda_2 \Gamma) \mathbf{c} = 0,$$

where  $\lambda_1 > 0$  and  $\lambda_2 > 0$  are chosen so that the constraints are satisfied. With some algebra, one can show that

$$\mathbf{c}_n = \lambda_{1,n}^{-1} (V + \lambda_{2,n} \Gamma)^{-1} U(\theta^*),$$

where  $\lambda_{2,n}$  satisfies

$$\frac{U(\theta^*)^T (V + \lambda_{2,n} \Gamma)^{-1} \Gamma (V + \lambda_{2,n} \Gamma)^{-1} U(\theta^*)}{U(\theta^*)^T (V + \lambda_{2,n} \Gamma)^{-1} V (V + \lambda_{2,n} \Gamma)^{-1} U(\theta^*)} = \gamma \quad (20)$$

and

$$\lambda_{1,n} = \{ U(\theta^*)^T (V + \lambda_{2,n} \Gamma)^{-1} V (V + \lambda_{2,n} \Gamma)^{-1} U(\theta^*) \}^{1/2}. \quad (21)$$

Finally, we can express  $\Psi_n(\mathcal{H}_K)$  as

$$\Psi_{n,\theta^*}(\mathcal{H}_K) = \lambda_{1,n}^{-1} U(\theta^*)^T (V + \lambda_{2,n} \Gamma)^{-1} U(\theta^*).$$

### 4.3 Bootstrap approximation of the null limiting distribution

We now discuss how to draw multiplier bootstrap samples to estimate the null limiting distribution of  $\Psi_{n,\theta^*}(\mathcal{H}_K)$ . For  $m = 1, \dots, M$  and  $M$  large, we draw  $\xi_1^{(m)}, \dots, \xi_n^{(m)}$  as independent standard normal random variables. Let  $\bar{\xi}^{(m)} = n^{-1} \sum_{i=1}^n \xi_i^{(m)}$ , and let  $U^{(m)}$  be a  $D$ -dimensional vector with  $d$ th element

$$U_d^{(m)} = n^{-1} \sum_{i=1}^n (\xi_i^{(m)} - \bar{\xi}^{(m)}) \phi_{n,\theta^*}(O_i; \eta_d),$$



where we recall that  $\phi_{n,\theta^*}$  is the plug-in estimator for the efficient influence function in (15) (and as noted in Section 3.2, we may wish to replace  $\theta^*$  with an estimator for  $\theta_0$ ). The  $m$ th sample from the multiplier bootstrap estimate of the null distribution (see (14)) can be calculated as

$$\Psi_{n,\theta^*}^{(m)}(\mathcal{H}_\kappa) = \sup_{\mathbf{c} \in \mathbb{R}^d} \{ (U^{(m)})^\top \mathbf{c} : \mathbf{c}^\top \Gamma \mathbf{c} \leq \gamma, \mathbf{c}^\top V \mathbf{c} = 1 \}.$$

The above optimization problem can be solved using the same routine described in Section 4.2, simply replacing  $U$  with  $U^{(m)}$ . Finally, for a realization  $t$  of  $\Psi_{n,\theta^*}(\mathcal{H}_\kappa)$ , a bootstrap  $p$ -value can be calculated as

$$\rho_M(t) = M^{-1} \sum_{m=1}^M \mathbb{1}(\Psi_{n,\theta^*}^{(m)}(\mathcal{H}_\kappa) > t).$$

#### 4.4 Confidence band construction

In this section, we discuss how to visualize the confidence set for  $\bar{\theta}_0$  obtained by inverting our proposed hypothesis test. Recall from Section 2.2 that we propose to report the smallest band that contains all functions belonging to the confidence set. This confidence band is defined as  $C_n^\alpha$ , and its form is provided in (4).

We first need to construct a function class  $\Theta$  that contains a collection of candidate values for  $\theta_0$ . While  $\Theta$  can be a rich class, it cannot be entirely unrestricted. In fact, if  $\Theta$  is too large, the confidence band  $C_n^\alpha$  can possibly have infinite width. To see this, note that the confidence set  $C_n^\alpha$  in (3) contains a set of functions  $\theta$  for which  $\Psi_{n,\theta}(\mathcal{H})$  is close to zero. It is possible to construct  $\theta$  so that at any given  $a$ ,  $\theta(a)$  takes an arbitrarily large positive or negative value, but  $\Psi_{n,\theta}(\mathcal{H}) = 0$ . For instance, if  $\Psi_{n,\theta}(\mathcal{H})$  is constructed using TML-based estimation, this could be achieved by setting  $\theta(A_i) = \theta_{\hat{p}_i}(A_i)$  for  $i = 1, \dots, n$  and allowing  $\theta$  to take *any value* at points where no data are observed. We would encounter the same issue if we instead used the one-step estimator. By selecting  $\Theta$  as, e.g., a class of smooth functions, we are able to avoid this problem. On the other hand, we note that if  $\Theta$  is not large enough to contain  $\theta_0$ ,  $C_n^\alpha$  is not guaranteed to achieve the nominal coverage rate. Given these considerations, we suggest selecting  $\Theta$  as a class of functions that is no less smooth than a reasonable approximation of  $\theta_0$ . In our implementation, we construct  $\Theta$  as a subset of functions belonging to an RKHS for which the RKHS norm bounded above by a constant. That is, we take  $\Theta = \Theta_\nu$ , where

$$\Theta_\nu := \left\{ \theta = \sum_{d=1}^D c_d \eta_d : \sum_{d=1}^D \frac{c_d^2}{\lambda_d} < \nu \right\}$$

for  $\nu > 0$ . We propose to set  $\nu$  as the RKHS norm of a consistent estimate of  $\theta_0$ , which could be obtained using the method described in Section 4.1.

Let  $t_{1-\alpha}^*$  be the  $1 - \alpha$  quantile of the null limiting distribution of  $n^{1/2}\Psi_{n,\theta^*}(\mathcal{H}_\kappa)$  for a fixed  $\kappa$ . Given the above construction of  $\Theta$ , the confidence band  $C_n^\alpha$  takes the following form at any given point  $a_0$ :

$$C_n^\alpha(a_0) = \left[ \inf \left\{ \sum_d c_d \left( \eta_d(a_0) - n^{-1} \sum_{i=1}^n \eta_d(A_i) \right) : \sum_{d=1}^D \frac{c_d^2}{\lambda_d} \leq \nu, \Psi_{n, \sum_d c_d \eta_d}(\mathcal{H}_\kappa) \leq n^{-1/2} t_{1-\alpha}^* \right\}, \right. \\ \left. \sup \left\{ \sum_d c_d \left( \eta_d(a_0) - n^{-1} \sum_{i=1}^n \eta_d(A_i) \right) : \sum_{d=1}^D \frac{c_d^2}{\lambda_d} \leq \nu, \Psi_{n, \sum_d c_d \eta_d}(\mathcal{H}_\kappa) \leq n^{-1/2} t_{1-\alpha}^* \right\} \right]. \quad (22)$$

The optimization problems in (22) are challenging to solve because  $\Psi_{n, \sum_d c_d \eta_d}(\mathcal{H})$  does not have a closed form expression in the coefficients  $c_1, \dots, c_D$ . Recall from Section 4.2 that we can write

$$\Psi_{n, \sum_d c_d \eta_d}(\mathcal{H}_\kappa) = \lambda_{1,n}^{-1} U \left( \sum_{d=1}^D c_d \eta_d \right)^\top (V + \lambda_{2,n} \Gamma)^{-1} U \left( \sum_{d=1}^D c_d \eta_d \right),$$

where  $\lambda_{1,n}$  and  $\lambda_{2,n}$  are constants that depend on  $U(\sum_{d=1}^D c_d \eta_d)$ . By instead treating  $\lambda_{1,n}$  and  $\lambda_{2,n}$  as fixed, we are able to obtain a closed form approximation of  $\Psi_{n,\theta^*}(\mathcal{H}_\kappa)$ . When  $\Psi_{n,\theta^*}(\mathcal{H}_\kappa)$  is constructed using either TML

or one-step estimation,  $U(\sum_{d=1}^D c_d \eta_d)$  is linear in the coefficients. As a result, when the closed form approximation of the test statistic is used, the optimization problem in (22) becomes a quadratically constrained quadratic program. As noted previously, this type of problem can be solved using publicly available software such as the CVXR package in R.

We conclude by discussing how to select the tuning parameters  $\kappa$ ,  $\lambda_1$ , and  $\lambda_2$ . The choice of  $\kappa$  should have no bearing on the asymptotic coverage of the confidence set, though it may affect the confidence band's width. In order for the confidence band to have optimal width, we need to select  $\kappa$  as to maximize the power to reject any null hypothesis  $H_0 : \bar{\theta}_0 = \bar{\theta}_0^*$ . Although the optimal choice of  $\kappa$  generally depends on the specific null hypothesis being tested, the optimization problem in (22) would become complicated if  $\mathcal{H}_\kappa$  was not fixed. For computational ease, we fix  $\kappa$  as a single value that is large enough so that we have reasonable power to reject a large set of null hypotheses. We set  $\kappa$  as an estimate for  $J(\theta_0)/\text{Var}_0(\theta_0(A))$ , which can be obtained using the approach described in Section 4.1, so that we are well powered against nearly flat nulls when  $\theta_0$  is not very flat. Finally, we pick  $\lambda_{1,n}$  and  $\lambda_{2,n}$  to satisfy (20) and (21) with  $\theta^* = 0$ . For this choice of  $\lambda_{1,n}$  and  $\lambda_{2,n}$ , the closed form approximation of  $\Psi_{n,\theta^*}(\mathcal{H})$  will be fairly accurate when  $\theta^*$  is nearly flat. Although  $\lambda_{1,n}$  and  $\lambda_{2,n}$  are data-dependent, asymptotic coverage will be unaffected if  $\lambda_{1,n}$  and  $\lambda_{2,n}$  converge to fixed constants.

## 5 Simulation study

### 5.1 Simulation setting

We begin by describing our approach for generating synthetic data sets. We first generate  $W_1, \dots, W_n$  as independent bivariate normal random vectors with mean zero, unit variance, and correlation  $\frac{1}{2}$ . Given  $W$ , we then draw  $A_1, \dots, A_n$  from a conditional distribution with density function

$$g_0(a|w) = \frac{\text{expit}(\zeta(w)a)}{\int_{-1}^1 \text{expit}(\zeta(w)a) da} \mathbf{1}(-1 \leq a \leq 1),$$

where we define  $\zeta$  as

$$\zeta(w) = 3 \left\{ \text{expit}(w_1 + w_2) - \frac{1}{2} \right\}.$$

Random variables with the above conditional density can be generated via the inverse cumulative distribution function method.

We generate the outcome  $Y$  under the following settings.

*Setting 1:*

In the first setting, we construct the conditional distribution of the outcome given the exposure and covariates so that the centered dose–response curve is zero. We draw  $Y$  from the model


$$Y = -\zeta(W) \left[ 1 - \frac{A}{2} \right] + \varepsilon,$$

where  $\varepsilon$  is a uniform random variable on  $[-2, 2]$ . Because  $\zeta(W)$  has mean zero, it can be seen that  $E_0[Q_0(W, a)] = 0$  as desired.

*Setting 2:*

In the second setting, we consider the case where  $\bar{\theta}_0$  is nonzero. We construct a model for  $Y$  so that the dose–response function is

$$\theta_0 : a \mapsto \frac{2a + a^2 - 3a^3}{2}.$$



**Figure 1:** Plot of the dose–response function under the data-generating mechanism used in the simulation study.

A plot of the dose–response function is provided in Figure 1. We generate  $Y$  as

$$Y = \theta_0(A) - \zeta(W) \left\{ 1 - \frac{A}{2} \right\} + \varepsilon,$$

where  $\varepsilon$  is again a uniform random variable on  $[-2, 2]$ .

Under each of the above settings, we use our proposed methodology to perform a test of the null hypothesis that the dose function is flat, i.e.,  $H_0 : \bar{\theta}_0 = 0$ . In the first setting, the null holds, and we would expect our approach to achieve nominal type-1 error control in the limit of large  $n$ . Under the second setting, the alternative holds, allowing us to assess the power of our proposed test.

We also examine the behavior of the proposed confidence bands when  $\bar{\theta}_0 \neq 0$ . We compute the interval  $C_n^\alpha$  in (22) on an equally spaced grid of 50 points on the interval  $[-1, 1]$ . We assess whether the bands are appropriate in the sense that they roughly capture the shape of the unknown dose–response function, and we evaluate the confidence bands' coverage. We calculate both simultaneous coverage probability and the average of the pointwise coverage probability over the 50 evaluation points.

We study the behavior of the following four variations of our proposal:

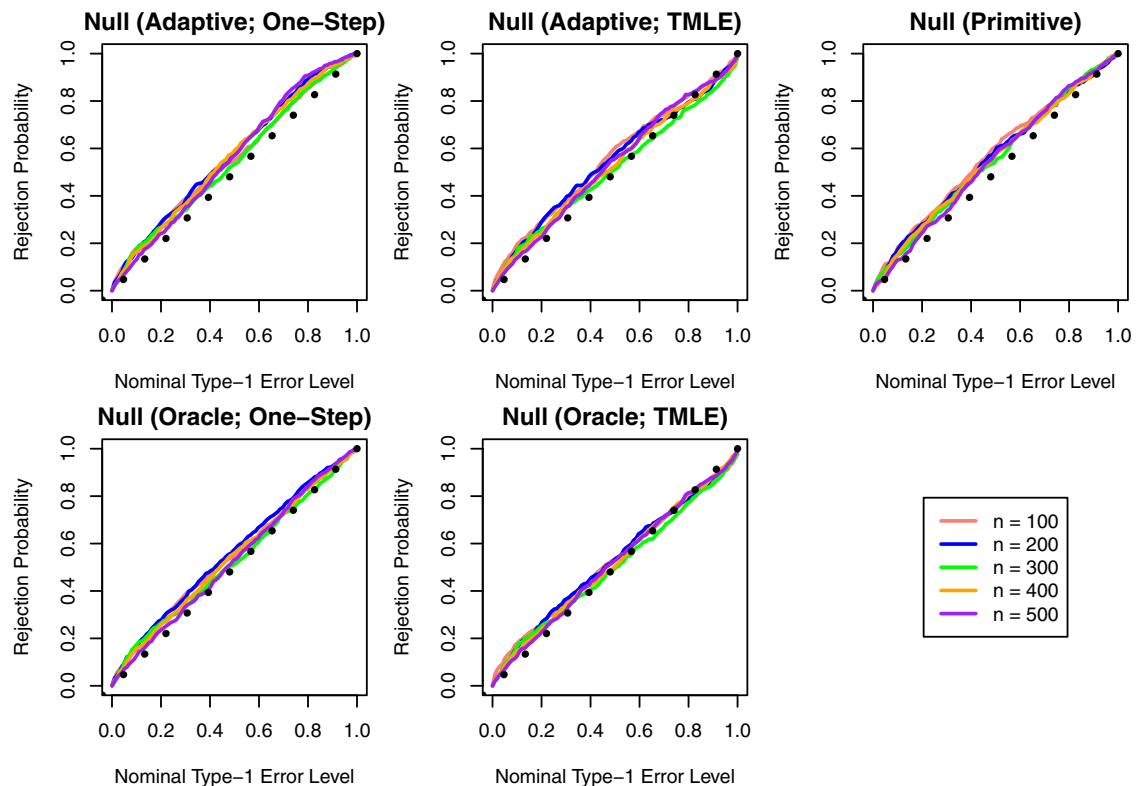
- A one-step estimator for  $\Psi_{0,\theta^*}$  is used, and we set  $\kappa$  as the oracle  $\kappa_0$ .
- A TML estimator for  $\Psi_{0,\theta^*}$  is used, and we set  $\kappa$  as the oracle  $\kappa_0$ .
- A one-step estimator for  $\Psi_{0,\theta^*}$  is used, and we set  $\kappa$  as the data-adaptive choice  $\kappa_n$ .
- A TML estimator for  $\Psi_{0,\theta^*}$  is used, and we set  $\kappa$  as the data-adaptive choice  $\kappa_n$ .

In each case, we use  $D = 20$  basis functions.

We compare our proposed hypothesis test with an approach similar to that described in the study of Westling [18], which is based on estimating primitive functions of the dose–response function. As noted above, their approach can be viewed as a variation of our proposal where we set  $\mathcal{H} = \{h(a) = \mathbb{1}(a \leq a_0) : a_0 \in \mathcal{A}\}$ . We use our own implementation of this procedure, which differs slightly in that we estimate  $\psi_{0,\theta^*}(h)$  using a one-step estimator, whereas Westling [18] uses a cross-fitted estimator. We apply each of the above methods to 500 synthetic data sets for  $n \in \{100, 200, \dots, 500\}$ .

## 5.2 Simulation results

Figure 2 shows the Monte Carlo estimate of the distribution function for the  $p$ -values produced from each method under Setting 1, where the flat null holds. When the type-1 error rate is well-controlled for any significance level  $\alpha$ , the distribution function should be linear. We find that our approach achieves type-1



**Figure 2:** Monte Carlo estimates of the empirical distribution of the  $p$ -values under the flat null.

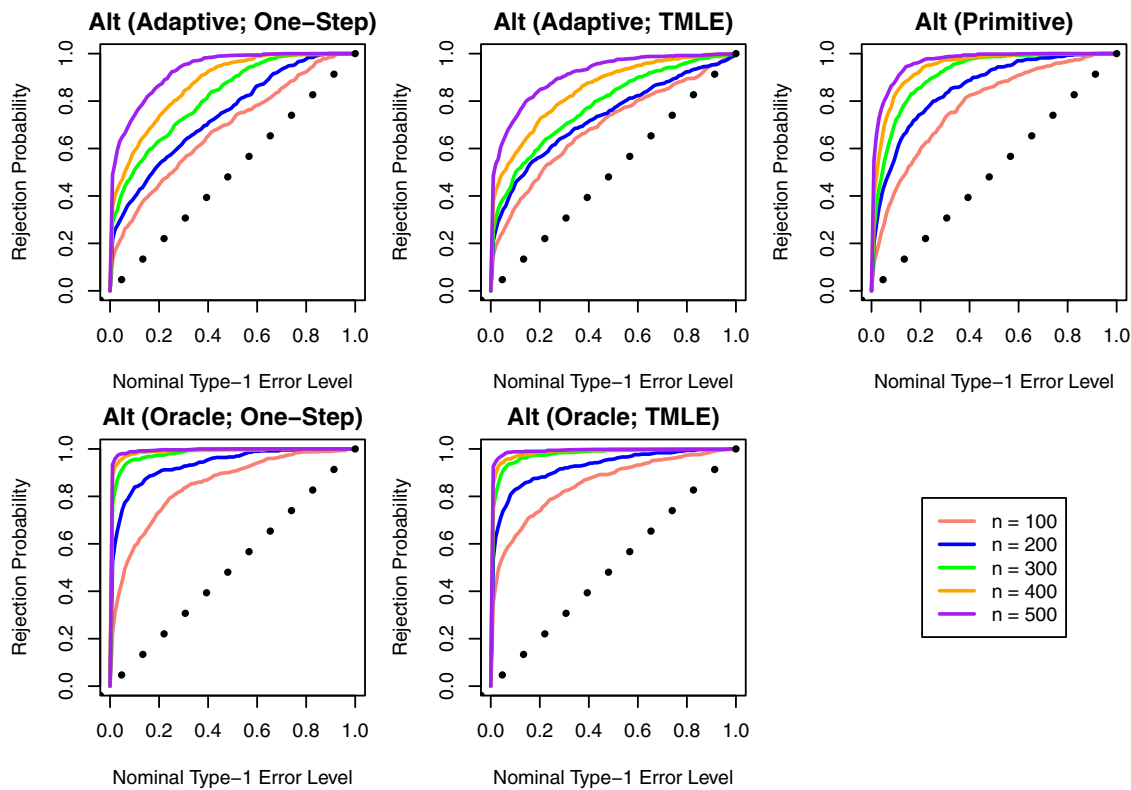
error control near the nominal error level when the oracle choice of  $\kappa$  is provided, and the data-adaptive choice results in some modest anti-conservatism. The approach based on estimation of primitive functions also achieves nearly nominal type-1 error control.

Figure 3 shows the Monte Carlo estimate of the distribution functions for the  $p$ -values under Setting 2, where the alternative holds. We find that our proposal has high power when the oracle choice of  $\kappa$  is supplied, and power declines when a data-adaptive choice is used. The approach proposed in the study of Westling [18] outperforms our approach when  $\kappa$  is chosen data-adaptively but performs worse than our approach when the oracle choice is used. This suggests that making use of known structure on  $\theta_0$  can help us improve power to reject some alternative hypotheses, though there is a notable decline in performance when we attempt to learn the structure from the data.

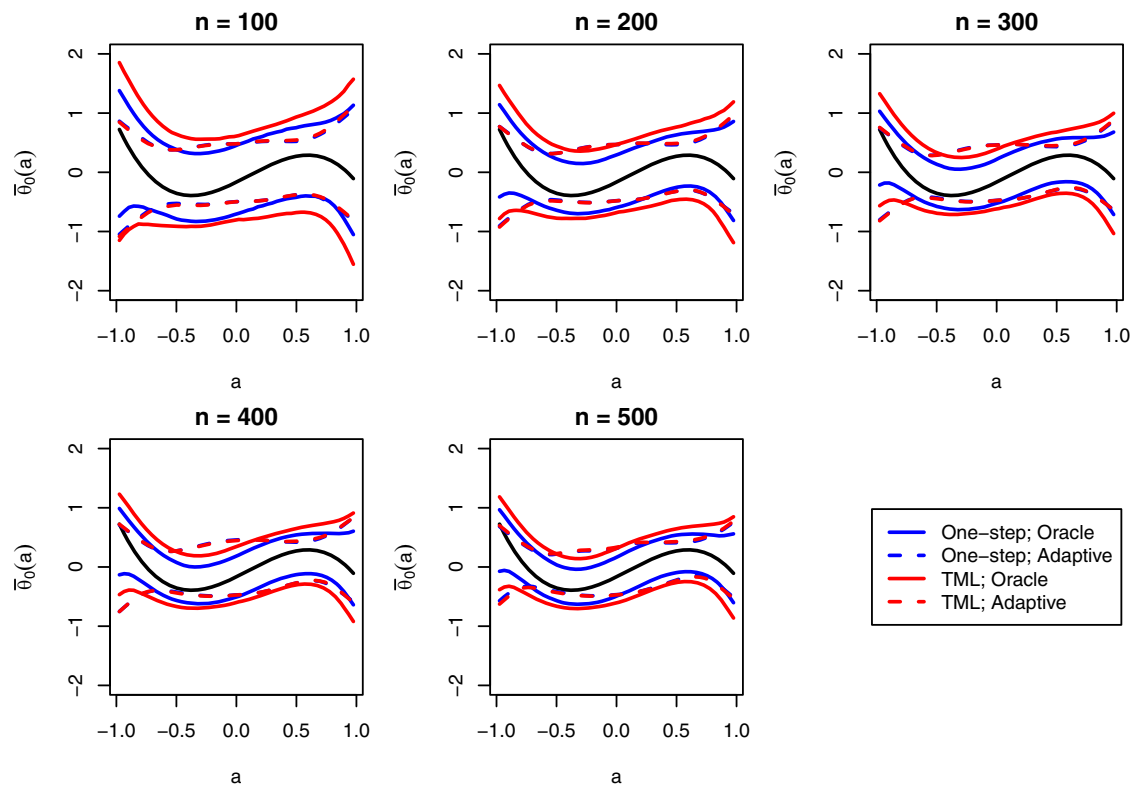
Figure 4 shows the median upper and lower limits of the confidence bands that were constructed using our proposal. We find that the confidence bands are able to capture the shape of the dose-response curve, and the width of the bands decreases as the sample size grows, as expected. Figure 5 shows the uniform and average coverage probabilities of the proposed confidence bands. We find that when  $\Theta$  is known, the confidence bands achieve simultaneous and average coverage at or above the nominal level. However, when  $\Theta$  is selected data-adaptively, the uniform coverage rate falls far below the nominal level, though the average coverage rate remains near the nominal level. That the confidence bands have poorer coverage with data-adaptive selection of  $\Theta$  is unsurprising as we only have theoretical coverage guarantees when  $\Theta$  is known to contain  $\theta_0$ , and our approach only assures that the class contains a close approximation for  $\theta_0$ . This is evidently insufficient, in particular for uniform coverage.

## 6 Data example

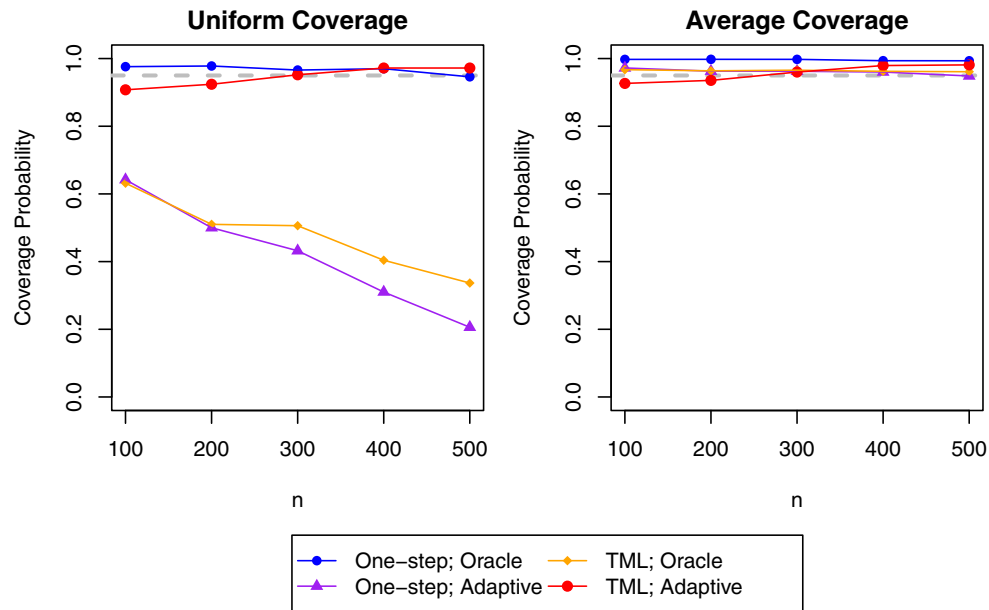
As an example, we use our method to analyze data from the Adaptive Strategies for Preventing and Treating Lapses of Retention in HIV Care (ADAPT-R) trial [20]. ADAPT-R was a sequential multiple assignment



**Figure 3:** Monte Carlo estimates of the empirical distribution of the  $p$ -values for a test against the flat null, under the alternative hypothesis.



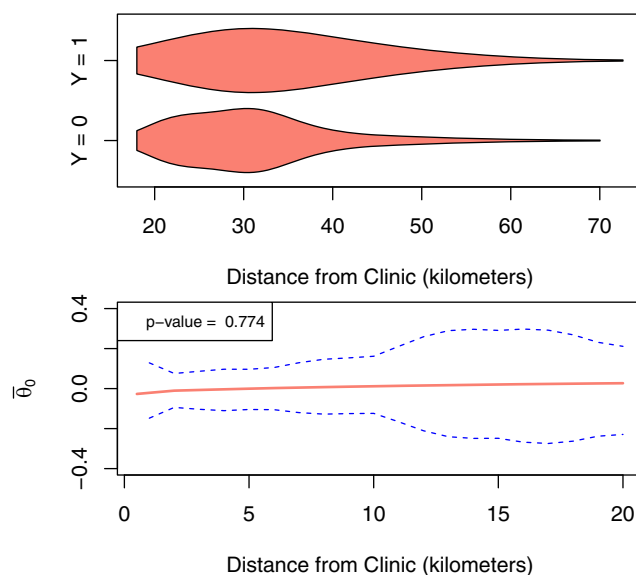
**Figure 4:** Monte Carlo estimates of the median upper and lower limits of the 95% bands obtained using our proposed methodology.



**Figure 5:** Monte Carlo estimates uniform coverage probability and average coverage probability of the 95% bands obtained using our proposed methodology. The dashed gray line indicates the nominal coverage rate 0.95.

randomized trial run in Kenya that studied the effectiveness of interventions for optimizing retention in HIV treatment in a population of people living with HIV who initiated care. In this study, a question of secondary interest is whether the distance a participant must travel to reach the nearest HIV clinic affects their retention in care. We perform an analysis to address this secondary aim, pooling across the trial's randomized arms.

We conduct our analysis using a sample of 1,815 participants from the ADAPT-R trial. We treat as the exposure of interest the distance from the nearest clinic. The distribution of the exposure variable is highly skewed. Approximately 95% of the study participants lived within 20 km of the nearest clinic, and among the remaining participants, distance ranges between 20 and 500 km. Because these extreme values are fairly rare,



**Figure 6:** (Top) Violin plot of distribution of the distance to furthest clinic, by status of retention in care. (Bottom) Plug-in estimate of centered dose-response function and 95% confidence bands.



there is concern about potential violation of the positivity assumption. To avoid this issue, we exclude from this analysis participants who lived more than 20 km from the nearest clinic, obtaining a final sample size of 1,600. Our outcome is a binary variable that is equal to one if a patient had neither experienced a lapse in care (defined as missing a scheduled clinic visit by at least 14 days) nor had unsuppressed HIV viral load 1 year after initiating care, and is zero otherwise. A total of 446 study participants experienced a lapse in care or had an unsuppressed HIV viral load within 1 year. As our exposure of interest is not randomized, we adjust for the following set of measured baseline variables that may either confound the exposure outcome relationship or predict the outcome and thus improve efficiency: age, sex assigned at birth, and a wealth index.

In Figure 6, we display the marginal distribution of distance to clinic by retention status. There does not appear to be a strong association between distance and retention, as the marginal distribution is nearly the same in both groups. To more formally assess the presence of an effect, we apply our method to perform a test of the flat null. We use the data-adaptive choice of  $\kappa$  described in Section 4.1, and we use a one-step estimator for  $\Psi_{0,\theta^*}(\mathcal{H})$ . Figure 6 shows a plug-in estimate of the centered dose–response function, in addition to 95% confidence bands and a  $p$ -value for a test of the flat null. The dose-response function appears to be nearly flat, and we are unable to reject the flat null hypothesis based as our  $p$ -value is quite large. These results suggest that there is no strong evidence to support that distance from clinic has a strong effect on retention in care in people living with HIV in this setting.

## 7 Conclusion

This work provides a novel approach to inference on the causal dose-response function. We show that, under mild regularity conditions, our nonparametric test achieves type-1 error control near the nominal level and is well-powered against the null. We also present a computationally tractable method for visualizing confidence sets constructed by inverting our proposed test. The recent proposal by Westling [18] also performs well under weak assumptions, though their work does not present a method for constructing or visualizing confidence sets. That we introduce a novel approach for constructing confidence sets is a key strength of our work.

The strategy for inference on the dose-response function we describe in this article can be adapted to address other problems of interest in the causal inference literature. For instance, one could use our approach to assess for treatment effect heterogeneity by testing the null hypothesis that a conditional average treatment effect curve is flat.

One of the main limitations of this work is that we require a specification for the function class  $\mathcal{H}$  for our test to be operational, though this class can be challenging to select in practice. When  $\mathcal{H}$  is specified *a priori* and contains  $h_0$ , our proposal performs very well, but when we attempt to select  $\mathcal{H}$  data-adaptively, we suffer a loss in performance. We note that in some settings, selecting  $\mathcal{H}$  *a priori* may be possible. For instance, if an independent data set is available (e.g., from a closely related study), one could use this data set to construct  $\mathcal{H}$  without looking at the data set they are primarily interested in analyzing. Alternatively, it is sensible in many settings to assume, without looking at the data, that the dose-response function is monotone. Therefore, one could consider implementing a version of our procedure where  $\mathcal{H}$  is a class of bounded monotone functions. In future work, we plan to develop improved strategies for tuning parameter selection.

Our proposal also requires that the nuisance parameter estimators are not overly complex. This condition is somewhat prohibitive and disallows us from using more flexible estimators, such as gradient-boosted trees [33]. To avoid this assumption, one could develop a slightly modified version of our procedure where  $\psi_{0,\theta^*}$  is estimated using cross-fitting [34,35].

**Acknowledgements:** We thank the Family AIDS Care and Education Services program operating HIV services in western Kenya, as well as the patients in these communities and the front-line health care workers in the region.

**Funding information:** This research was supported by grants (R01 MH104123, K24 AI134413, and R01 AI074345) from the National Institutes of Health.

**Author contributions:** AH developed this manuscript's theoretical contribution, conducted simulation experiments, led the data analysis, and prepared the manuscript. EHG, TAO, and EAB supported the data analysis. MLP and MJvdL jointly supervised this project.

**Conflict of interest:** Maya L. Petersen and Mark J. van der Laan serve on the Editorial Advisory Board for the Journal of Causal Inference, but neither was involved in the peer review process for this article.

**Data availability statement:** A de-identified data set from ADAPT-R can be shared upon request to the authors.

## References

- [1] Rubin DB. Estimating causal effects of treatments in randomized and nonrandomized studies. *J. Educ Psychol.* 1974;66(5):688.
- [2] Robins JM, Rotnitzky A. Semiparametric efficiency in multivariate regression models with missing data. *J Am Stat Assoc.* 1995;90(429):122–9.
- [3] van der Laan MJ, Rubin D. Targeted maximum likelihood learning. *Int J Biostat.* 2006;2(1).
- [4] Díaz I, van der Laan MJ. Targeted data adaptive estimation of the causal dose-response curve. *J Causal Inference.* 2013;1(2):171–92.
- [5] Kennedy EH, Ma Z, McHugh MD, Small DS. Non-parametric methods for doubly robust estimation of continuous treatment effects. *J R Stat Soc Ser B (Stat Meth).* 2017;79(4):1229–45.
- [6] Colangelo K, Lee Y-Y. Double debiased machine learning nonparametric inference with continuous treatments. 2020. arXiv: <http://arXiv.org/abs/arXiv:2004.03036>.
- [7] Westling T, Gilbert P, Carone M. Causal isotonic regression. *J R Stat Soc Ser B (Stat Meth).* 2020;82(3):719–47.
- [8] Hirano K, Imbens GW. The propensity score with continuous treatments. *Appl Bayesian Model Causal Inference Incomplete-data Perspectives.* 2004;226164:73–84.
- [9] Imai K, Van Dyk DA. Causal inference with general treatment regimes: Generalizing the propensity score. *J Am Stat Assoc.* 2004;99(467):854–66.
- [10] Galvao AF, Wang L. Uniformly semiparametric efficient estimation of treatment effects with a continuous treatment. *J Am Stat Assoc.* 2015;110(512):1528–42.
- [11] Robins JM, Hernan MA, Brumback B. Marginal structural models and causal inference in epidemiology. *Epidemiology* 2000;11(5):550–60.
- [12] Zhang Z, Zhou J, Cao W, Zhang J. Causal inference with a quantitative exposure. *Stat Methods Med Res.* 2016;25(1):315–35.
- [13] Doss CR, Weng G, Wang L, Moscovice I, Chantarat T. A nonparametric doubly robust test for a continuous treatment effect. *Ann. Statist.* 2024;52(4):1592–615.
- [14] Takatsu K, Westling T. Debiased inference for a covariate-adjusted regression function. *J R Stat Soc B: Stat. Methodol.* 2024;qkae041.
- [15] Neugebauer R, van der Laan M. Nonparametric causal effects based on marginal structural models. *J Stat Plan Inference.* 2007;137(2):419–34.
- [16] Muñoz ID, van der Laan M. Population intervention causal effects based on stochastic interventions. *Biometrics.* 2012;68(2):541–9.
- [17] Hines O, Diaz-Ordaz K, Vansteelandt S. Parameterising the effect of a continuous exposure using average derivative effects. 2021. arXiv: <http://arXiv.org/abs/arXiv:2109.13124>.
- [18] Westling T. Nonparametric tests of the causal null with nondiscrete exposures. *J Am Stat Assoc.* 2021;117:1–12.
- [19] Hudson A, Carone M, Shojaie A. Inference on function-valued parameters using a restricted score test. 2021. arXiv: <http://arXiv.org/abs/arXiv:2105.06646>.
- [20] Geng EH, Odeny TA, Montoya LM, Iguna S, Kulzer JL, Adhiambo HF, et al. Adaptive strategies for retention in care among persons living with HIV. *NEJM Evidence.* 2023;2(4):EVIDoa2200076.
- [21] Bickel PJ, Klaassen CAJ, Ritov Y, Wellner JA. Efficient and adaptive estimation for semiparametric models. New York, NY: Springer; 1998.
- [22] Pfanzagl J. Contributions to a general asymptotic statistical theory. New York, NY: Springer; 1982.
- [23] van der Laan MJ, Rose S. Targeted learning: causal inference for observational and experimental data. New York, NY: Springer Science & Business Media; 2011.
- [24] van der Laan MJ, Rose S. Targeted learning in data science. Cham: Springer International Publishing; 2018.
- [25] Benkeser D, van der Laan M. The highly adaptive LASSO estimator. In: 2016 IEEE international conference on data science and advanced analytics (DSAA). IEEE, 2016. pp. 689–96.
- [26] van der Laan M, Gruber S. One-step targeted minimum loss-based estimation based on universal least favorable one-dimensional submodels. *Int J Biostat.* 2016;12(1):351–78.

- [27] Hable R. Asymptotic normality of support vector machine variants and other regularized kernel methods. *J Multivariate Anal.* 2012;106:92–117.
- [28] Wahba G. *Spline models for observational data*. Philadelphia: SIAM; 1990.
- [29] Hastie T, Tibshirani R, Friedman JH, Friedman JH. *The elements of statistical learning: data mining, inference, and prediction*. vol. 2. New York, NY: Springer; 2009.
- [30] Barron AR. Statistical properties of artificial neural networks. In: *Proceedings of the 28th IEEE Conference on Decision and Control*. IEEE; 1989. pp. 280–5.
- [31] van der Laan MJ, Polley EC, Hubbard AE. Super learner. *Stat Appl Genetics Mol Biol.* 2007;6(1).
- [32] Fu A, Narasimhan B, Boyd S. CVXR: An r package for disciplined convex optimization. *J Stat Softw.* 2020;94(14):1–34.
- [33] Friedman JH. Stochastic gradient boosting. *Comput Stat Data Anal.* 2002;38(4):367–78.
- [34] Zheng W, van der Laan MJ. Cross-validated targeted minimum-loss-based estimation. In *Targeted learning*. New York, NY: Springer; 2011. pp. 459–74.
- [35] Chernozhukov V, Chetverikov D, Demirer M, Duflo E, Hansen C, Newey W, et al. Double/debiased machine learning for treatment and structural parameters. *Econometr J.* 2018;21(1):C1–68.
- [36] Kosorok MR. *Introduction to empirical processes and semiparametric inference*. New York, NY: Springer Science & Business Media; 2008.
- [37] van der Vaart A, Wellner J. *Weak convergence and empirical processes*. New York, NY: Springer; 1996.
- [38] van der Vaart A. *Asymptotic statistics*. vol. 3. Cambridge, England: Cambridge University Press; 2000.

## Appendix

### A Proofs of theoretical results

**Proof of Lemma 1.** Let  $P$  be a distribution function that satisfies conditions A1–A3, and let  $p$  denote the density of  $P$  with respect to a dominating measure  $\mu$ . Let  $\omega : \mathcal{W} \times \mathcal{A} \times \mathcal{Y} \rightarrow \mathbb{R}$  be an arbitrary function with zero mean and finite variance under  $P$ . We define  $P_\varepsilon$  as the parametric submodel with density function  $p_\varepsilon : (w, a, y) \mapsto p(w, a, y)\{1 + \varepsilon\omega(w, a, y)\}$  for  $\varepsilon$  small. We can observe that  $P_\varepsilon$  is equal to  $P$  at  $\varepsilon = 0$ , and the score function is  $\omega$ . Every regular parametric model that passes through  $P$  and has score function equal to  $\omega$  can be closely approximated using such a submodel. If there exists function  $\tilde{\phi}_p$  that satisfies

$$\frac{d}{d\varepsilon}\psi_{P_\varepsilon, \theta^*}(h)|_{\varepsilon=0} = E_P[\tilde{\phi}_p(O; h)\omega(O)], \quad (A1)$$

for any choice of  $\omega$ , then  $\psi_{P_\varepsilon, \theta^*}(h)$  is pathwise differentiable in a nonparametric model, and  $\tilde{\phi}_p(\cdot; h) - E_P[\tilde{\phi}_p(O; h)]$  is the nonparametric efficient influence function. We show that such a function exists, and when centered about its mean, it is equal to  $\phi_{P, \theta^*}(\cdot; h)$ .

We first evaluate the derivative of  $\psi_{P_\varepsilon}$  at  $\varepsilon = 0$ . We express  $\psi_{P_\varepsilon, \theta^*}(h)$  as the sum of three components:

$$\psi_{P_\varepsilon, \theta^*}(h) = G_{P_\varepsilon}^I(h) - G_{P_\varepsilon}^{II}(h) - G_{P_\varepsilon}^{III}(h),$$

where for any distribution  $P$ , we define

$$\begin{aligned} G_P^I(h) &:= E_P[\theta_P(A)h(A)] \\ G_P^{II}(h) &:= E_P[\theta_P(A)]E_P[h(A)] \\ G_P^{III}(h) &:= E_P[\theta^*(A)h(A)] - E_P[\theta^*(A)]E_P[h(A)]. \end{aligned}$$

To make calculation of the derivative of  $\psi_{P_\varepsilon, \theta^*}$  more manageable, we evaluate the derivative of the three additive components at  $\varepsilon = 0$ .

We first calculate  $\frac{d}{d\varepsilon}G_{P_\varepsilon}^I(h)|_{\varepsilon=0}$ . We can write  $G_{P_\varepsilon}^I(h)$  as

$$G_{P_\varepsilon}^I(h) = \int \theta_{P_\varepsilon}(a)h(a)p(w, a, y)\{1 + \varepsilon\omega(w, a, y)\}d\mu(w, a, y),$$

where  $\theta_{P_\varepsilon}$  is the dose-response function under  $P_\varepsilon$ . To proceed, we need an expression for the derivative of  $\theta_{P_\varepsilon}(a)$  with respect to  $\varepsilon$ . We can write

$$\theta_{P_\varepsilon}(a) = \left[ \int Q_{P_\varepsilon}(w, a) \left\{ \int p(w, a, y)\{1 + \varepsilon\omega(w, a, y)\}d\mu(dy, da) \right\} \mu(dw) \right],$$

where  $Q_{P_\varepsilon}$  is the conditional mean of  $Y$  given  $A$  and  $W$  under  $P_\varepsilon$ . Letting  $p^i(y|a, w)$  be the conditional density of  $Y$  given  $A$  and  $W$ , it can be shown that the derivative of  $Q_{P_\varepsilon}$  at  $\varepsilon = 0$  is

$$\frac{d}{d\varepsilon}Q_{P_\varepsilon}(w, a)|_{\varepsilon=0} = \int \{y - Q_P(w, a)\}\omega(w, a, y)p^i(y|a, w)d\mu(y).$$

Now, let the conditional density of  $(A, Y)$  given  $W$  be  $p^{ii}(a, y|w)$ , and let the marginal density of  $W$  be  $p^{iii}(w)$ . We now calculate the derivative of  $\theta_{P_\varepsilon}(a)$  at  $\varepsilon = 0$ :

$$\begin{aligned} \frac{d}{d\varepsilon}\theta_{P_\varepsilon}(a) \Big|_{\varepsilon=0} &= \left[ \frac{d}{d\varepsilon}Q_{P_\varepsilon}(w, a) \right]_{\varepsilon=0} \left\{ \int p(w, a, y)\mu(dy, da) \right\} \mu(dw) \\ &\quad + \int Q_P(w, a) \left\{ \int p(w, a, y)\omega(w, a, y)\mu(dy, da) \right\} \mu(dw) \\ &= \left[ \int \{y - Q_P(w, a)\}\omega(w, a, y)p^i(y|a, w)d\mu(y) \right] p^{iii}(w)\mu(dw) \\ &\quad + \int Q_P(w, a)p^{iii}(w) \left\{ \int p^{ii}(a, y|w)\omega(w, a, y)\mu(dy, da) \right\} \mu(dw). \end{aligned} \quad (A2)$$

We can now express the derivative of  $G_{P_\varepsilon}^I(h)$  as

$$\begin{aligned} \left. \frac{d}{d\varepsilon} G_{P_\varepsilon}^I(h) \right|_{\varepsilon=0} &= \int \frac{d}{d\varepsilon} [\theta_{P_\varepsilon}(a)]_{\varepsilon=0} h(a) p(w, a, y) d\mu(w, a, y) + \int \theta_P(a) p(w, a, y) \omega(w, a, y) d\mu(w, a, y) \\ &= \int \int \int \{y - Q_P(w, a)\} \omega(w, a, y) p^I(y|a, w) d\mu(y) \Big] p^{III}(w) \mu(dw) h(a) p(w, a, y) d\mu(w, a, y) \\ &\quad + \int \int Q_P(w, a) \Big[ \int p^{II}(a, y|w) \omega(w, a, y) \mu(dy, da) \Big] p^{III}(w) \mu(dw) h(a) p(w, a, y) d\mu(w, a, y) \\ &\quad + \int \theta_P(a) h(a) p(w, a, y) \omega(w, a, y) d\mu(w, a, y). \end{aligned}$$

Let the marginal density of  $A$  be  $p^{IV}(a)$ , and let the conditional density of  $W$  given  $A$  be  $p^V(w|a)$ . Observe that we can express the marginal density of  $W$  as

$$p^{III}(w) = p^V(w|a) \left[ \frac{p^{IV}(a)}{g_P(a|W)} \right].$$

Using this fact, we can write

$$\begin{aligned} \left. \frac{d}{d\varepsilon} G_{P_\varepsilon}^I(h) \right|_{\varepsilon=0} &= E_P \left[ E_P \left[ \left[ \frac{p^{IV}(A)}{g_P(A|W)} \right] E_0[\{Y - Q_P(W, A)\} \omega(W, A, Y) | W, A] | A \right] h(A) \right] \\ &\quad + E_P \left[ E_P \left[ \left[ \frac{p^{IV}(A)}{g_P(A|W)} \right] Q_P(W, A) E_P[\omega(W, A, Y) | W] | A \right] h(A) \right] + E_P[\theta_P(A) h(A) \omega(W, A, Y)]. \end{aligned} \quad (A3)$$

By applying the law of total expectation in the first two lines, we obtain

$$\begin{aligned} \left. \frac{d}{d\varepsilon} G_{P_\varepsilon}^I(h) \right|_{\varepsilon=0} &= E_P \left[ \left[ \frac{p^{IV}(A)}{g_P(A|W)} \right] \{Y - Q_P(W, A)\} h(A) \omega(W, A, Y) \right] \\ &\quad + E_P \left[ E_P \left[ \left[ \frac{p^{IV}(A)}{g_P(A|W)} \right] Q_P(W, A) h(A) | W \right] \omega(W, A, Y) \right] \\ &\quad + E_P[\theta_P(A) h(A) \omega(W, A, Y)]. \end{aligned}$$

Finally, by observing that in the second line above, the conditional densities of  $A$  given  $W$  cancel, we have

$$\begin{aligned} \left. \frac{d}{d\varepsilon} G_{P_\varepsilon}^I(h) \right|_{\varepsilon=0} &= E_P \left[ \left[ \frac{p^{IV}(A)}{g_P(A|W)} \right] \{Y - Q_P(W, A)\} h(A) \omega(W, A, Y) \right] \\ &\quad + E_P[E_{P,A}[Q_P(W, A) h(A)] \omega(W, A, Y)] \\ &\quad + E_P[\theta_P(A) h(A) \omega(W, A, Y)], \end{aligned}$$

where  $E_{P,A}$  is the expectation over the marginal distribution of  $A$  under  $P$ .

We now evaluate the derivative of  $G_{P_\varepsilon}^{II}(h)$ . We can express  $G_{P_\varepsilon}^{II}$  as

$$G_{P_\varepsilon}^{II}(h) = \left[ \int \theta_{P_\varepsilon}(a) p(w, a, y) \{1 + \varepsilon \omega(w, a, y)\} d\mu(w, a, y) \right] \left[ \int h(a) p(w, a, y) \{1 + \varepsilon \omega(w, a, y)\} d\mu(w, a, y) \right].$$

The evaluation of its derivative at zero is

$$\begin{aligned} \left. \frac{d}{d\varepsilon} G_{P_\varepsilon}^{II}(h) \right|_{\varepsilon=0} &= \left[ \int h(a) p(w, a, y) d\mu(w, a, y) \right] \left[ \int \frac{d}{d\varepsilon} [\theta_{P_\varepsilon}(a)]_{\varepsilon=0} p(w, a, y) d\mu(w, a, y) \right] \\ &\quad + \left[ \int h(a) p(w, a, y) d\mu(w, a, y) \right] \left[ \int \theta_P(a) p(w, a, y) \omega(w, a, y) d\mu(w, a, y) \right] \\ &\quad + \left[ \int h(a) \omega(w, a, y) p(w, a, y) d\mu(w, a, y) \right] \left[ \int \theta_P(a) p(w, a, y) d\mu(w, a, y) \right]. \end{aligned}$$

Performing similar steps as were used to calculate  $\frac{d}{d\varepsilon} G_{P_\varepsilon}^I(h)|_{\varepsilon=0}$ , it can be shown that

$$\begin{aligned} \left. \frac{d}{d\varepsilon} G_{P_\varepsilon}^{II}(h) \right|_{\varepsilon=0} &= E_P[h(A)] E_0 \left[ \frac{p^{iv}(A)}{g_P(A|W)} \{Y - Q_P(W, A)\} \omega(W, A, Y) \right] \\ &\quad + E_P[E_P[h(A)] E_{P,A}[Q_P(W, A)] \omega(W, A, Y)] \\ &\quad + E_P[E_P[h(A)] \theta_P(A) \omega(W, A, Y)] + E_P[E_P[\theta_P(A)] h(A) \omega(W, A, Y)]. \end{aligned} \quad (A4)$$

Now, we take the derivative of the remaining term  $G_{P_\varepsilon}^{III}(h)$ . Rather than perform this calculation directly, we recognize that  $G_P^{III}(h)$  is simply the covariance between  $h(A)$  and  $\theta^*(A)$  under  $P$ , and it is well known that the derivative can be expressed as

$$\left. \frac{d}{d\varepsilon} G_{P_\varepsilon}^{III}(h) \right|_{\varepsilon=0} = E_P[\{\theta^*(A)h(A) - E[h(A)]\theta^*(A) - E[\theta^*(A)]h(A)\} \omega(W, A, Y)]. \quad (A5)$$

Now, by (A3), (A4), and (A5), we can express  $\frac{d}{d\varepsilon} \psi_{P_\varepsilon, \theta^*}(h)|_{\varepsilon=0}$  as

$$\begin{aligned} \frac{d}{d\varepsilon} \psi_{P_\varepsilon, \theta^*}(h)|_{\varepsilon=0} &= \frac{d}{d\varepsilon} [G_{P_\varepsilon}^I(h) - G_{P_\varepsilon}^{II}(h) - G_{P_\varepsilon}^{III}(h)]|_{\varepsilon=0} \\ &= E_P[\phi_P^*(W, A, Y; h) \omega(W, A, Y)], \end{aligned}$$

where we define

$$\begin{aligned} \phi_P^*(w, a, y; h) &:= \left[ \bar{\theta}_P(a) - \bar{\theta}_P^*(a) + \frac{p^{iv}(a)}{g_P(a|w)} \{y - Q_P(w, a)\} \right] \{h(a) - E_P[h(A)]\} \\ &\quad + E_P[Q_P(w, A) \{h(A) - E_P[h(A)]\}]. \end{aligned}$$

The proof is completed by observing that  $p^{iv}(a) = E_P[g_P(a|W)]$ , and

$$\phi_{P, \theta^*}(\cdot; h) = \phi_P^*(\cdot; h) - E_P[\phi_P^*(W, A, Y; h)]. \quad \square$$

**Proof of Lemma 2.** This result is an immediate consequence of Slutsky's theorem [36] (see, e.g., Theorem 7.15 of [36]).  $\square$

**Proof of Theorem 1.** The one-step estimator can be expressed as

$$\psi_{n, \theta^*}^I(h) - \psi_{0, \theta^*}(h) = n^{-1} \sum_{i=1}^n \phi_{P_0}(O_i; h) + R_{1,n}(h) + R_{2,n}(h),$$

where we define

$$\begin{aligned} R_{1,n}(h) &:= - \left[ n^{-1} \sum_{i=1}^n \phi_{\hat{P}, \theta^*}(O_i; h) - \phi_{P_0, \theta^*}(O_i; h) - \left\{ \int \phi_{\hat{P}}(o; h) - \phi_{P_0}(o; h) dP_0(o) \right\} \right], \\ R_{2,n}(h) &:= \psi_{n, \theta^*}(h) - \psi_{0, \theta^*}(h) - \int \phi_{\hat{P}, \theta^*}(o; h) dP_0(o). \end{aligned}$$

By the triangle inequality, it is sufficient to show that  $\sup_{h \in \mathcal{H}} |R_{1,n}(h)|$  and  $\sup_{h \in \mathcal{H}} |R_{2,n}(h)|$  are both  $o_P(n^{-1/2})$ .

To see that  $\sup |R_{1,n}(h)| = o_P(n^{-1/2})$ , we note that by Assumption B1 and Theorem 2.10.6 of [37],  $\mathcal{H}$  is a  $P_0$ -Donsker class, and therefore,  $\sup_{h \in \mathcal{H}} |n^{-1} \sum_{i=1}^n h(A_i) - E_0[h(A)]| = O_P(n^{-1/2})$ . In view of this fact and Assumption B2, it can be concluded that the estimator  $\phi_{\hat{P}_n, \theta^*}$  of the efficient influence function is uniformly consistent over  $\mathcal{H}$ , that is

$$\sup_{h \in \mathcal{H}} \int \{\phi_{\hat{P}_n}(o; h) - \phi_{0, \theta^*}(o; h)\}^2 dP_0(o) = o_P(1).$$

It is shown in the proof of lemma 19.26 in [38] that when this condition is satisfied, and Assumption B2 holds,  $\sup_{h \in \mathcal{H}} |R_{1,n}(h)| = o_P(n^{-1/2})$ .



It remains to be shown that  $R_{2,n}(h)$  is asymptotically negligible. With some algebra, it can be shown that the second remainder term can be expressed as

$$R_{2,n}(h) = R_{2,n}^i(h) + R_{2,n}^{ii}(h) + R_{2,n}^{iii}(h) + R_{2,n}^{iv}(h) + R_{2,n}^v(h),$$

where we define

$$\begin{aligned} R_{2,n}^i(h) &:= \int \left\{ \frac{E_0[g_0(a|W)]}{g_0(a|w)} - \frac{n^{-1} \sum_{i=1}^n g_n(a|W_i)}{g_n(a|w)} \right\} \{Q_0(w, a) - Q_n(w, a)\} \{h(a) - E_0[h(A)]\} dP_0(w, a) \\ R_{2,n}^{ii}(h) &:= n^{-1} \sum_{i=1}^n \left\{ n^{-1} \sum_{j=1}^n Q_n(W_i, A_j) \{h(A_j) - E_0[h(A)]\} - \int Q_n(W_i, a) \{h(a) - E_0[h(A)]\} dP_0(a) \right\} \\ &\quad - \int \left\{ n^{-1} \sum_{j=1}^n Q_n(w, A_j) \{h(A_j) - E_0[h(A)]\} - \int Q_n(w, a) \{h(a) - E_0[h(A)]\} dP_0(a) \right\} dP_0(w) \\ R_{2,n}^{iii}(h) &:= \left\{ n^{-1} \sum_{i=1}^n h(A_i) - E_0[h(A)] \right\} \int \frac{n^{-1} \sum_{i=1}^n g_n(a|W_i)}{g_n(a|w)} \{Q_0(w, a) - Q_n(w, a)\} dP_0(a) \\ R_{2,n}^{iv}(h) &:= \left\{ E_0[h(A)] - n^{-1} \sum_{i=1}^n h(A_i) \right\} \left[ n^{-1} \sum_{i=1}^n \left\{ n^{-1} \sum_{j=1}^n Q_n(W_i, A_j) - \int Q_n(W_i, a) dP_0(a) \right\} \right] \\ &\quad - \left\{ E_0[h(A)] - n^{-1} \sum_{i=1}^n h(A_i) \right\} \left[ \int \left\{ n^{-1} \sum_{j=1}^n Q_n(w, A_j) - \int Q_n(w, a) dP_0(a) \right\} dP_0(w) \right] \\ R_{2,n}^v(h) &:= \left\{ n^{-1} \sum_{i=1}^n h(A_i) - E_0[h(A)] \right\} \left[ n^{-1} \sum_{i=1}^n \theta^*(A_i) - E_0[\theta^*(A)] \right]. \end{aligned}$$

By the triangle inequality, it suffices to argue that each of the above components converges to zero in probability at a rate of  $n^{1/2}$  uniformly in  $\mathcal{H}$ .

It follows directly from Assumption B1 that  $\sup_{h \in \mathcal{H}} |R_{2,n}^i(h)| = o_p(n^{-1/2})$ . We now argue that  $\sup_{h \in \mathcal{H}} |R_{2,n}^{ii}(h)| = o_p(n^{-1/2})$ . Observe that

$$\begin{aligned} &\sup_{h \in \mathcal{H}} \int \left\{ n^{-1} \sum_{i=1}^n Q_n(w, A_i) h(A_i) - \int Q_n(w, a) h(a) dP_0(a) \right\}^2 dP_0(w) \\ &\leq 2 \sup_{h \in \mathcal{H}} \int \left\{ n^{-1} \sum_{i=1}^n \{Q_n(w, A_i) - Q_0(w, A_i)\} h(A_i) - \int \{Q_n(w, a) - Q_0(w, a)\} h(a) dP_0(a) \right\}^2 dP_0(w) \\ &\quad + 2 \sup_{h \in \mathcal{H}} \int \left\{ n^{-1} \sum_{i=1}^n Q_0(w, A_i) h(A_i) - \int Q_0(w, a) h(a) dP_0(a) \right\}^2 dP_0(w) = o_p(1), \end{aligned}$$

where the convergence follows from the rate conditions and the Donsker assumption. It can thus be concluded using the argument in the proof of Lemma 19.26 of [38] that  $\sup_{h \in \mathcal{H}} |R_{2,n}^{ii}(h)| = o_p(n^{-1/2})$ . It can be seen that  $\sup_{h \in \mathcal{H}} |R_{2,n}^{iii}(h)| = o_p(n^{-1/2})$  and  $\sup_{h \in \mathcal{H}} |R_{2,n}^{iv}(h)| = o_p(n^{-1/2})$  by recalling that  $\sup_{h \in \mathcal{H}} |n^{-1} \sum_{i=1}^n h(A_i) - E_0[h(A)]| = O_p(n^{-1/2})$  and noting that  $Q_n$  is a consistent estimator for  $Q_0$ . Similarly, that  $\sup_{h \in \mathcal{H}} |R_{2,n}^v(h)| = o_p(n^{-1/2})$  follows from uniform consistency of  $n^{-1} \sum_{i=1}^n h(A_i)$ . This completes our argument to show that the one-step estimator is uniformly asymptotically linear.  $\square$

**Proof of Theorem 2.** The proof of Theorem 2 is nearly the same as the proof of Theorem 1, so we only provide a brief outline of the argument. The TML estimator can be expressed as

$$\psi_{n,\theta^*}^{\Pi}(h) - \psi_{0,\theta^*}(h) = n^{-1} \sum_{i=1}^n \phi_{P_0}(O_i; h) + R_{1,n}(h) + R_{2,n}(h) + R_{3,n}(h),$$

where we define

$$\begin{aligned} R_{1,n}(h) &= - \left[ n^{-1} \sum_{i=1}^n \phi_{\tilde{P}, \theta^*}(O_i; h) - \phi_{P_0, \theta^*}(O_i; h) - \left\{ \int \phi_{\tilde{P}}(o; h) - \phi_{P_0}(o; h) dP_0(o) \right\} \right], \\ R_{2,n}(h) &= \psi_{n, \theta^*}(h) - \psi_{0, \theta^*}(h) - \int \phi_{\tilde{P}, \theta^*}(o; h) dP_0(o) \\ R_{3,n}(h) &= -n^{-1} \sum_{i=1}^n \phi_{\tilde{P}_h}(O_i; h). \end{aligned}$$

Uniform asymptotic negligibility of  $R_{1,n}(h)$  and  $R_{2,n}(h)$  follows from the same arguments we presented for establishing uniform asymptotic linearity of the one-step estimator. Additionally, we have by construction that  $\sup_{h \in \mathcal{H}} |R_{3,n}(h)| = o_P(n^{-1/2})$ . From this, we can conclude that the TML estimator is uniformly asymptotically linear.  $\square$