

Research Article

Alec McClean*, Zach Branson, and Edward H. Kennedy

Nonparametric estimation of conditional incremental effects

<https://doi.org/10.1515/jci-2023-0024>

received April 25, 2023; accepted January 23, 2024

Abstract: Conditional effect estimation has great scientific and policy importance because interventions may impact subjects differently depending on their characteristics. Most research has focused on estimating the conditional average treatment effect (CATE). However, identification of the CATE requires that all subjects have a non-zero probability of receiving treatment, or positivity, which may be unrealistic in practice. Instead, we propose conditional effects based on incremental propensity score interventions, which are stochastic interventions where the odds of treatment are multiplied by some factor. These effects do not require positivity for identification and can be better suited for modeling scenarios in which people cannot be forced into treatment. We develop a projection approach and a flexible nonparametric estimator that can each estimate all the conditional effects we propose and derive model-agnostic error guarantees showing that both estimators satisfy a form of double robustness. Further, we propose a summary of treatment effect heterogeneity and a test for any effect heterogeneity based on the variance of a conditional derivative effect and derive a nonparametric estimator that also satisfies a form of double robustness. Finally, we demonstrate our estimators by analyzing the effect of intensive care unit admission on mortality using a dataset from the (SPOT)light study.

Keywords: conditional effects, nonparametric estimators, stochastic interventions

MSC 2020: 62D20

1 Introduction

Estimating causal effects has great scientific and policy importance, and often there is interest in understanding if the effectiveness of a treatment depends on the subjects' characteristics. Conditional or "heterogeneous," effects describe how a treatment effect varies with the subjects' characteristics, and can illustrate qualitatively important phenomena that would be disguised by average effects. Previous work has focused on estimating the conditional average treatment effect (CATE), which considers the difference between counterfactual mean outcomes when all subjects at some covariate level receive treatment and all subjects receive control (e.g., [1–7], among others). However, in many contexts, researchers cannot force subjects to receive treatment or prevent them from receiving treatment, thereby making the counterfactual interventions behind the CATE unrealistic in practice. As a concrete example, we will consider the effect of intensive care unit (ICU) admission on mortality for emergency room entrants [8]. Typically, the counterfactual interventions where everyone is admitted to the ICU and no one is admitted to the ICU are both practically infeasible because there

* **Corresponding author: Alec McClean**, Department of Statistics & Data Science, Carnegie Mellon University, Pittsburgh, United States of America, e-mail: alec@stat.cmu.edu

Zach Branson: Department of Statistics & Data Science, Carnegie Mellon University, Pittsburgh, United States of America, e-mail: zach@stat.cmu.edu

Edward H. Kennedy: Department of Statistics & Data Science, Carnegie Mellon University, Pittsburgh, United States of America, e-mail: edward@stat.cmu.edu

are a finite number of ICU beds and because hospitals have a duty of care towards sick patients. Instead, we may be interested in assessing the causal effect of an intervention that could more realistically be implemented in practice, such as an intervention that moderately increases or decreases the probability of admission to the ICU. For example, increasing or decreasing the number of ICU beds would likely increase or decrease the probability of admission for all patients. Generally, these interventions can best be described with stochastic interventions, which characterize counterfactual outcomes under a shift in the treatment distribution [9–15]. With a binary treatment, this shift can be characterized by an incremental propensity score intervention (“incremental intervention”), which multiplies the odds of treatment by a user-specified factor δ [11,16].

Recent research on stochastic interventions generally and incremental interventions specifically have focused on average effects [9,11,17]. In this study, we consider estimating conditional incremental effects (CIEs), where we assess to what extent an incremental effect depends on the subjects’ characteristics, which can uncover treatment effect heterogeneity that is obscured by average effects. Furthermore, as well as corresponding to more realistic interventions, there are two additional advantages in considering CIEs instead of the CATE. First, incremental effects are robust to positivity violations, in the following sense. When positivity is deterministically violated, such that subjects have zero or one probability of receiving treatment, incremental effects can still be identified. Moreover, even when positivity is not strictly violated but estimated propensity scores are nonetheless close to zero or one, confidence intervals for incremental effects will not be influenced by these extreme propensity scores. By contrast, if positivity is deterministically violated, the CATE may not be identifiable. And, if positivity is nearly violated, without strong parametric modeling assumptions, it can be difficult to estimate the average treatment effect (ATE) or the CATE, in the sense that variance estimates are large and confidence intervals are wide [18].

A second advantage of using CIEs instead of the CATE is the ability to describe a continuum of policies between treating all subjects and treating none, where the interventions behind the CATE are special cases at each end of the continuum. A researcher might presume that stochastic effects follow a roughly linear relationship from one end of the continuum to the other, with the slope of the line matching the sign of the CATE. As discussed in Remark 2 in Section 2, this assumption is reasonable when conditioning on all the covariates, as the CIE curve must be monotonic in the incremental parameter δ and so its slope will match the sign of the CATE. However, most analyses, including our ICU data analysis in Section 5, condition on only a few covariates of interest, allowing for the possibility of other incremental effect curves. For example, consider Figure 1 – a preview of the real data analysis in Section 5 – which shows CIE curves for several Intensive Care

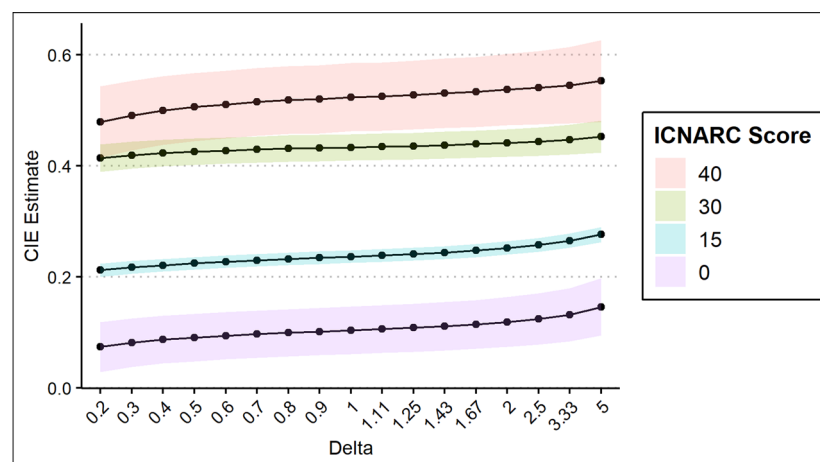


Figure 1: CIE curves for select ICNARC scores. The x-axis represents the incremental intervention parameter δ , where $\delta = 1$ corresponds to no intervention, and $\delta > 1$ and $\delta < 1$ correspond to increasing and decreasing the likelihood of admission to the ICU, respectively. The y-axis shows estimated mortality rate. The curves depict the estimated CIE for different ICNARC scores, which measure mortality risk. Our analysis shows that for each ICNARC score, the mortality rate decreases when fewer ($\delta < 1$) people are sent to the ICU and increases when more ($\delta > 1$) people are sent to the ICU, but mild changes in ICU admission rates lead to minimal changes in mortality rate.

National Audit & Research Centre (ICNARC) scores (a measure of mortality risk). The x -axis represents the incremental intervention parameter, where $\delta = 1$ corresponds to no intervention while $\delta > 1$ and $\delta < 1$ correspond to increasing and decreasing the likelihood, respectively, that patients are admitted to the ICU, y -axis shows estimated mortality rate. The curves illustrate that the estimated counterfactual mortality rate is higher at $\delta = 5$ than at $\delta = 0.2$, in agreement with prior research indicating that admitting everyone to the ICU is harmful compared to admitting no one [8]. However, the full curves also illustrate that mild interventions correspond to small changes in mortality rate. Thus, stochastic effects can be evaluated over a continuum of interventions, which may reveal additional information that would otherwise be hidden by only studying the always-treated vs never-treated counterfactual effects. This echoes previous findings that it can be useful to target additional effects beyond these two extreme counterfactuals [19].

1.1 Contribution and structure

Motivated by these observations, in this work, we describe how to estimate conditional causal effects for incremental interventions and illustrate how these effects facilitate a more nuanced understanding of treatment effect heterogeneity than the usual CATE. We focus on incremental interventions for two reasons. First, the incremental intervention has an intuitive parameterization for binary treatment since it corresponds to multiplying the odds of treatment by some factor. Second, the intervention demonstrates favorable properties because it is anchored at the observed treatment distribution and considers a smooth shift from the observed distribution. For example, identifying effects with this intervention does not require the positivity assumption that the probability of treatment is bounded between zero and one for all subjects, which is required for identifying the CATE. This allows estimation of CIEs to still be precise even in the face of positivity violations, unlike estimation of the CATE.

We consider three conditional effects in this work. First, we describe the CIE, which is the conditional analog to the average incremental effect. As shown in Figure 1, the CIE is described by a curve for each covariate value; this makes quantifying treatment effect heterogeneity challenging, because we have to consider how much these curves vary across covariate values. As a preliminary extension of the CIE, we describe the conditional incremental *contrast* effect (CICE), which considers a contrast between two incremental interventions, and is the incremental analog to the CATE. The CICE can enable better understanding of treatment effect heterogeneity than the CIE, but it requires specifying two incremental δ parameters, and it may not immediately be clear which parameter values would be of most interest in a particular application. Therefore, we propose the conditional incremental *derivative* effect (CIDE), which corresponds to the change in the CIE under an infinitesimal shift of the treatment distribution. We find that the CIDE is particularly useful for quantifying treatment effect heterogeneity for incremental interventions because it allows the researcher to examine the spectrum of interventions like in Figure 1 and also construct estimators and tests to quantify treatment effect heterogeneity, as discussed in Section 4.

For the three conditional effects, we propose two estimators. Our first estimator, the Projection-Learner, estimates the projection of the true conditional effect onto a finite dimensional model. This added structure allows us to re-frame the estimator as the solution to a moment condition and derive an efficient influence function. Utilizing the properties of efficient influence functions, we provide double robust style error guarantees for the Projection-Learner and show that its bias scales as a product of errors of the nuisance function estimators (in this work, the nuisance functions are the propensity score and the outcome regression, and are defined in Section 2). As a result, the Projection-Learner can achieve parametric efficiency even when the nuisance functions are estimated nonparametrically. Our second conditional effect estimator, the I-DR-Learner, is a two-stage meta-learner that extends the DR-Learner studied by Kennedy [3] to incremental effects. For the I-DR-Learner, the first stage estimates the efficient influence function values for the relevant average effect and the second stage regresses those values against the conditioning covariates. We establish when the I-DR-Learner exhibits double robust style guarantees; in particular, the conditional effect must lie in a certain infinite dimensional function class, and the second stage regression must satisfy a form of stability. In this case, we demonstrate that the I-DR-Learner can attain oracle efficiency when the nuisance functions

are estimated nonparametrically. Therefore, the I-DR-Learner cannot obtain parametric efficiency like the Projection-Learner, but it can estimate a larger class of true conditional effect curves with oracle efficiency.

Both the Projection-Learner and the I-DR-Learner can be used to estimate conditional effect curves across variables of interest. A natural question is whether there is any treatment effect heterogeneity across the curve. Thus, researchers may also be interested in a one-dimensional summary of effect heterogeneity and a corresponding test for any effect heterogeneity. Therefore, in addition to the CIE, CICE, and CIDE curves, we also propose a fourth effect, the variance of the conditional incremental derivative effect (V-CIDE), which can be used to estimate the degree of effect heterogeneity and test for any effect heterogeneity. For the V-CIDE, we derive a novel double robust style estimator based on its efficient influence function, illustrate that our estimator attains parametric efficiency under weak conditions on the nuisance function estimators, and derive a corresponding test for any effect heterogeneity.

The structure of the work is as follows. In Section 1.2, we define relevant notation. In Section 2, we define the data setup and different estimands of interest, state the causal assumptions required for identification, and establish identification results for our conditional effects. In Section 3, we outline the Projection-Learner and I-DR-Learner and demonstrate their convergence properties in Sections 3.1 and 3.2, respectively. In Section 4, we outline a nonparametric estimator for the V-CIDE, demonstrate its convergence properties, and describe methods for inference. In Section 5, we analyze data on ICU admission from the (SPOT)light prospective cohort study. We estimate that, while greatly increasing subjects' odds of attending the ICU would increase mortality rates, mild to moderate changes in ICU admissions rates lead to minimal changes in mortality rates. While this agrees with what would be concluded for CATE estimation, CATE estimation for this application would not be reliable, because there are positivity violations for this dataset. Using our test, we do not find evidence that there is treatment effect heterogeneity. Finally, in Section 6 we conclude and discuss future extensions of this research.

All of our methods can be implemented using the `npcausal` package in R [20,21], as demonstrated in the replication materials at <https://github.com/alecmcclean/NPCIE>.

1.2 Notation

We use E for expectation and V for variance. We use $P_n(f) = P_n\{f(Z)\} = \frac{1}{n} \sum_{i=1}^n f(Z_i)$ as a shorthand for sample averages and $V_n\{f(Z)\} = \frac{1}{n-1} \sum_{i=1}^n \left\{ f(Z_i) - \frac{1}{n} \sum_{j=1}^n f(Z_j) \right\}^2$ as shorthand for the sample variance. When $x \in \mathbb{R}^d$, we let $\|x\|^2 = \sum_{j=1}^d x_j^2$ denote the squared Euclidean norm, and for generic possibly random functions f , we let $\|f\|^2 = \int_{\mathcal{Z}} f(z)^2 dP(z)$ denote the squared $\ell_2(P)$ norm. We use the notation $a \lesssim b$ to mean $a \leq Cb$ for some constant C , and $a \asymp b$ to mean $cb \leq a \leq Cb$ for some constants c and C , so that $a \lesssim b$ and $b \lesssim a$. We use \rightsquigarrow to denote convergence in distribution and \xrightarrow{P} for convergence in probability. We use \hat{E}_n to denote the predicted regression function estimate from n samples (e.g., if we considered a regression of Y against X , then $\hat{E}_n(Y|X=x)$ is the estimated regression function of Y against X at $X=x$ using n data points $\{(X_i, Y_i)\}_{i=1}^n$). We use the set notation $A \setminus B$ to indicate “ A and not B .”

2 Estimands and identification results for CIEs

In this section, we describe estimands for incremental effects, and establish assumptions for identifying these effects. Assume we observe $\{Z_1, \dots, Z_n\}$ with $Z_i \stackrel{iid}{\sim} \mathcal{P}$, where $Z = (X, A, Y)$, $X \in \mathbb{R}^d$ are the covariates, $A \in \{0, 1\}$ is the treatment status, and $Y \in \mathbb{R}$ is an outcome. We define potential outcomes Y^a as the outcome that would be observed when treatment $A = a$.

Much of the causal inference literature has focused on estimating the ATE and CATE, defined as follows:

$$\text{ATE: } \psi_{\text{ate}} = \mathbb{E}(Y^1 - Y^0) \quad \text{and} \quad (1)$$

$$\text{CATE: } \tau_{\text{cate}}(x) = \mathbb{E}(Y^1 - Y^0 | X = x). \quad (2)$$

To identify the ATE and the CATE, the following three causal assumptions are commonly used and are sufficient:

Assumption 1 (Consistency), $Y = Y^a$ if $A = a$

Assumption 2 (Exchangeability), $A \perp\!\!\!\perp Y^a | X$

Assumption 3 (Positivity), There exists $\varepsilon > 0$ such that $\mathbb{P}\{\varepsilon \leq \mathbb{P}(A = a | X) \leq 1 - \varepsilon\} = 1$ for $a \in \{0, 1\}$ and all X .

Consistency says that if an individual takes treatment a , we observe the potential outcome under that treatment regime. By contrast, consistency would be violated if, for example, there was interference between subjects, such that one subject's treatment status affected another's outcome. Exchangeability says that treatment is effectively randomized within covariate strata, in the sense that treatment is independent of subjects' potential outcomes after conditioning on covariates. Positivity says that all subjects have a non-zero chance of receiving treatment or control, and positivity may be unrealistic in practice. Although positivity is required to identify the ATE and the CATE, as we show next, only Assumptions 1 and 2 are required to identify CIEs.

Remark 1. Ideally, estimability should be balanced against scientific relevance when choosing a causal estimand. While in this work, we focus on incremental effects and highlight their benefits, including robustness to positivity violations and ability to describe a spectrum of interventions, in many settings the scientifically relevant effect may be the ATE or the CATE. In that scenario, even though the incremental effect is estimable, it may not describe the causal effect of interest.

2.1 Incremental propensity score interventions

The incremental intervention corresponds to multiplying each individual's odds of treatment by a user-specified parameter δ . We define the propensity score, the probability that an individual receives treatment, as $\pi(X) = \mathbb{P}(A = 1 | X)$, and then the shifted propensity score under an incremental intervention is defined as

$$q\{\pi(X); \delta\} = \frac{\delta\pi(X)}{\delta\pi(X) + 1 - \pi(X)}. \quad (3)$$

Then, the average incremental effect is

$$\mathbb{E}(Y^{Q_\delta}), \quad (4)$$

where Q_δ is drawn from a Bernoulli distribution with parameter $q\{\pi(X); \delta\}$. Unlike deterministic interventions, the incremental intervention is stochastic because it does not deterministically assign subjects to treatment or control – rather, it shifts their propensity score. The incremental intervention also corresponds to multiplying the odds of treatment by δ , since $\delta = \frac{q\{\pi(X); \delta\} / [1 - q\{\pi(X); \delta\}]}{\pi(X) / [1 - \pi(X)]}$. In the ICU application previously discussed, the average incremental effect is the counterfactual 28-day mortality rate across emergency room patients if their odds of admission to the ICU were multiplied by δ . Incremental interventions were first proposed by Kennedy [11], with double robust style estimators for average (possibly time-varying) incremental effects. The analysis of average effects has been extended to censored data [22] and resource-constrained settings [23], and used for estimating the effect of aspirin on the incidence of pregnancy [24]; a review is provided by Bonvini et al. [16].

While this intervention is not prescriptive, since it is unlikely a hospital would specifically admit patients to the ICU based on draws from a Bernoulli distribution, it can be useful for describing interventions that might be implemented in practice. For example, if ICU capacity increased by some number of beds, it is plausible that every patient's odds of ICU admission might increase by a factor δ . Such an intervention cannot

be described by the CATE. Meanwhile, even though such an intervention would likely not correspond to Bernoulli draws, an incremental intervention with $\delta > 1$ could nonetheless appropriately describe this counterfactual question. Furthermore, a spectrum of δ could appropriately describe the range of admission criteria changes that a hospital may implement.

The incremental intervention is also dynamic in the sense that the intervention changes with X if $\pi(X)$ changes with X . This occurs because the intervention is constant on the odds ratio scale rather than the unit scale. For example, if $\delta = 2$ and the propensity scores for two covariate values are $\{\pi(X = x_1), \pi(X = x_2)\} = \{0.25, 0.5\}$, then the intervention propensity scores are $\{q\{\pi(X = x_1) = 0.25; \delta = 2\}, q\{\pi(X = x_2) = 0.5; \delta = 2\}\} = \{0.4, 0.66\}$. Therefore, the propensity score increases by 0.15 when $\pi(x_1) = 0.25$ and increases by 0.16 when $\pi(x_2) = 0.5$. However, although the interventions are dynamic in the sense just outlined, they are not dynamic in the sense that the user-specified parameter δ changes with X . We leave this as an avenue for future exploration. If δ were allowed to vary with X , a natural question then might be: what is the “optimal” choice of δ at a particular value $X = x$? As in the deterministic intervention literature, finding an optimal intervention could fruitfully build on the conditional effect estimators proposed in this work [25,26].

Other interventions have also been considered in the literature, such as modified treatment policies, which shift a continuous treatment by a specified amount [10,14,27]; stochastic interventions that shift a continuous treatment distribution [13]; dynamic interventions that depend on some time-varying information about subjects [14,28]; and stochastic interventions which shift a discrete but possibly multi-valued treatment distribution [29], including exponential tilts [9]. The incremental effect can also be interpreted as an exponential tilt. Wen et al. [17] recently proposed a similar intervention to the incremental intervention, but their intervention is parameterized as a shift of the risk ratio $\frac{q\{\pi(X); \delta\}}{\pi(X)}$, rather than the odds ratio.

2.2 CIEs

Now we will consider CIEs. We denote $V \subseteq X$ as either one or a set of covariates, and define the CIE as the counterfactual mean under an incremental intervention conditional on covariates V ,

$$\text{CIE: } \tau_{\text{cie}}(v; \delta) = \mathbb{E}(Y^{Q_\delta} | V = v). \quad (5)$$

In the ICU application previously discussed, $\tau_{\text{cie}}(v; \delta)$ is the counterfactual average 28-day mortality rate among emergency room patients with covariates v , if their odds of ICU admission were multiplied by δ . Although effects often refer to contrasts in the causal inference literature, for simplicity, we use the general term “effect” to refer to a causal estimand involving potential outcomes. Therefore, we refer to the CIE as an effect. When the effect of interest is a contrast, we explicitly describe it as a contrast, as in the CICE.

The following proposition establishes that the CIE is identifiable as a function of the observed data distribution:

Proposition 1. *Let Q_δ denote the incremental intervention defined in equation (4). Under Assumptions 1–2, the mean counterfactual outcome of the given covariates $V = v$ is identified by*

$$\mathbb{E}(Y^{Q_\delta} | V = v) = \mathbb{E} \left[\frac{\delta \pi(X) \mu(1, X) + \{1 - \pi(X)\} \mu(0, X)}{\delta \pi(X) + 1 - \pi(X)} | V = v \right], \quad (6)$$

where $\mu(a, x) = \mathbb{E}(Y | A = a, X = x)$.

All proofs are given in Appendix. Proposition 1 is a straightforward corollary of Corollary 1 in Kennedy [11], which shows that the CIE is identified by a linear combination of the regression functions $\mu(A, X)$ where the weights depend on the probabilities of receiving treatment and control under the incremental intervention.

The CIE does not consider a contrast between two interventions, and so it does not immediately describe treatment effect heterogeneity. In this sense, it is similar to the conditional counterfactual mean under

treatment, $\mathbb{E}(Y^1|V = v)$. As a first approach in understanding treatment effect heterogeneity, we define a second estimand, the CICE, which considers the difference between two incremental effects.

$$\text{CICE: } \tau_{\text{cice}}(v; \delta_u, \delta_l) \equiv \mathbb{E}(Y^{Q_{\delta_u}} - Y^{Q_{\delta_l}}|V = v). \quad (7)$$

The CICE is the difference (conditional at $V = v$) between the average outcomes if we multiply the odds of treatment by δ_u and if we multiply the odds of treatment by δ_l . In the ICU application previously discussed, $\tau_{\text{cice}}(v; \delta_u, \delta_l)$ is the difference in counterfactual average 28-day mortality among emergency room patients with covariates v if their odds of ICU admission were multiplied by δ_u vs if they were multiplied by δ_l . Like the CIE, the CICE is a descriptive rather than prescriptive estimand – it is unlikely that a hospital would consider interventions where patients are admitted to the ICU by draws from a Bernoulli distribution, where the odds of admission depend on the patient's natural probability of admission, multiplied by the incremental parameter. Nonetheless, the CICE may describe interventions that can be implemented in practice (e.g., making it more or less likely that emergency room entrants are admitted to the ICU), and thus we can understand treatment effect heterogeneity by looking at how the CICE changes with V . The CICE is readily comparable to the CATE since both consider contrasts between two interventions. In fact, if positivity is satisfied in Assumption 3, then the CICE approaches the CATE as $\delta_u \rightarrow \infty$ and $\delta_l \rightarrow 0$ since $\lim_{\delta_u \rightarrow \infty, \delta_l \rightarrow 0} \tau_{\text{cice}}(v; \delta_u, \delta_l) = \mathbb{E}(Y^1 - Y^0|V = v)$. Identification of the CICE follows by Proposition 1 and linearity of expectation since $\tau_{\text{cice}}(v; \delta_u, \delta_l) = \tau_{\text{cie}}(v; \delta_u) - \tau_{\text{cie}}(v; \delta_l)$.

2.3 Derivative effects

A limitation of the CICE is that it requires specifying two parameters, δ_u and δ_l , and it may not immediately be clear which parameter values would be of most interest in a particular application. Instead, we can consider a derivative effect, which describes the change in counterfactual outcomes with an infinitesimally small change in the treatment distribution. To ease exposition, we re-parameterize the average incremental effect with t instead of δ , and define the average derivative effect with respect to t , and evaluated at δ , as

$$\left. \frac{\partial}{\partial t} \mathbb{E}(Y^{Q_t}) \right|_{t=\delta}$$

and the associated CIDE as

$$\text{CIDE: } \tau_{\text{cide}}(v; \delta) = \left. \frac{\partial}{\partial t} \tau_{\text{cie}}(v; t) \right|_{t=\delta}. \quad (8)$$

In the ICU application previously discussed, $\tau_{\text{cide}}(v; \delta)$ is the change in counterfactual average 28-day mortality among emergency room patients with covariates v if their odds of ICU admission were infinitesimally increased from δ . The CIDE demonstrates treatment effect heterogeneity if it varies across v . Thus, it can illustrate effect heterogeneity across a continuum of policies if it is evaluated at several values for δ . For example, as we investigate in Section 5, the CIDE can illustrate whether the effect of ICU admission on mortality varies across ICNARC scores.

If $\tau_{\text{cie}}(x; \delta)$ is continuous in x and δ with absolutely integrable partial derivative with respect to δ , allowing differentiation under the integral, the CIDE is identified according to the following result.

Proposition 2. *Let Q_δ denote the incremental intervention defined in equation (4). Under Assumptions 1 and 2 and if $\tau_{\text{cie}}(x; \delta)$ is continuous in x and δ , and its partial derivative with respect to δ is absolutely integrable, then the CIDE is identified by*

$$\tau_{\text{cide}}(v; \delta) = \mathbb{E}[\omega(X; \delta)\{\mu(1, X) - \mu(0, X)\}|V = v], \quad (9)$$

where $\mu(a, x) = \mathbb{E}(Y|A = a, X = x)$ and

$$\omega(X; \delta) = \frac{\pi(X)\{1 - \pi(X)\}}{\{\delta\pi(X) + 1 - \pi(X)\}^2}. \quad (10)$$

Proposition 2 shows that the CIDE is a weighted average of the difference in mean outcomes under treatment and control, where the weights depend on the propensity scores and the incremental propensity scores.

Remark 2. When conditioning on all the covariates (i.e., $V = X$), the right-hand side of the identification result in equation (9) is monotonic in δ . This is because $\omega(x; \delta)$ is always non-negative, while $\mu(1, x) - \mu(0, x)$ does not change with δ . However, when $V \subset X$, the right-hand side of the identification result in equation (9) need not be monotonic across δ ; if it is not monotonic, this indicates that $\mu(1, x) - \mu(0, x)$ changes sign as x varies, holding v fixed.

We also propose a one-dimensional functional to assess treatment effect heterogeneity. We consider the variance of the V-CIDE, defined as

$$\text{V-CIDE: } \mathbb{V}\{\tau_{\text{cide}}(V; \delta)\}. \quad (11)$$

When this variance equals zero, it implies that the CIDE is constant over V , and thus there is no treatment effect heterogeneity. As before, the V-CIDE depends on δ , so it can be estimated over a grid of δ to evaluate treatment effect heterogeneity over a continuum of policies. By Proposition 2, the V-CIDE is identified by

$$\mathbb{V}\{\tau_{\text{cide}}(V; \delta)\} = \mathbb{V}(\mathbb{E}[\omega(X; \delta)\{\mu(1, X) - \mu(0, X)\}|V]) \quad (12)$$

and when $V = X$ this simplifies to

$$\mathbb{V}\{\tau_{\text{cide}}(X; \delta)\} = \mathbb{V}[\omega(X; \delta)\{\mu(1, X) - \mu(0, X)\}]. \quad (13)$$

In Section 4, we will derive an efficient estimator for the V-CIDE and propose a test for whether there is any effect heterogeneity at all. In Section 3, efficient estimators for the CIE, the CICE, and the CIDE are derived.

3 Estimating conditional incremental effects

The identification results in Propositions 1 and 2 suggest straightforward “plug-in” estimators for the conditional effects. Given estimates for $\hat{\pi}(X)$, $\hat{\mu}(a, X)$, and $\hat{\mathbb{P}}(X|V = v)$, an estimator can be constructed by plugging these estimates into the identification formulae in equations (6) and (9). If models for the nuisance functions are parametric and correctly specified, this approach can be optimal as the plug-in estimator will converge to a normal distribution at a $n^{-1/2}$ -rate. However, if the parametric models are misspecified, then the plug-in estimator will be biased [30]. Given this, it is tempting to use flexible nonparametric models to estimate the nuisance functions, in order to alleviate issues of model misspecification. However, in this case, typically the plug-in estimator will inherit the slow rate of convergence for the nonparametric models.

This motivates estimators based on semiparametric efficiency theory [31–35]. The first-order bias of the nonparametric plug-in can be characterized by the efficient influence function of the estimand, which can be thought of as the first derivative in a von Mises expansion of the estimand [36]. Thus, a natural approach is to estimate the efficient influence function and subtract this estimate from the nonparametric plug-in estimate in order to “de-bias” the plug-in. A benefit of estimators based on the efficient influence function is that their bias is a second-order product of errors of the nuisance function estimators, such that the estimator can achieve $n^{-1/2}$ efficiency even when the nuisance functions are estimated at slower nonparametric rates [33,37,38]. We consider two estimators that utilize efficient influence functions; as a result, they both exhibit double robust style error guarantees.

Our first estimator, the Projection-Learner, targets the projection of the true conditional effect onto a finite dimensional working model. Projection approaches have a long history in statistics [39–43] and causal

inference [6,44–46], and our approach is closely related to working marginal structural models [47] and assumption-lean-inference [48], in the sense that they also leverage finite-dimensional models but without invoking parametric assumptions. This added structure allows us to re-frame the estimator as the solution to a moment condition, and derive an efficient influence function. We show that the Projection-Learner exhibits a version of double robustness, and attains parametric efficiency under weak model-agnostic $n^{-1/4}$ conditions on the nuisance function estimators, which are achievable for nonparametric estimators under suitable smoothness or sparsity.

Our second estimator, the I-DR-Learner (inspired by the “DR-Learner” in Kennedy [3]), instead targets the true conditional effect. The I-DR-Learner is an estimation procedure that, like many recent CATE estimation approaches, tries to estimate the true conditional effect as flexibly as possible [1–5,7,49,50]. Without any further assumptions, no efficient influence function exists for the true conditional effect because it is not pathwise differentiable [51]. So, it is not possible to construct an estimator directly from an efficient influence function for the conditional effect. Instead, the I-DR-Learner is a two-stage meta-learner, which estimates the efficient influence function values for the relevant average effect (e.g., the average incremental effect for the CIE) in the first stage, and then regresses these values against the conditioning covariates in the second stage. We show that if the second stage regression satisfies a generalization of the classic stochastic equicontinuity-type condition, the I-DR-Learner exhibits a form of double robustness and achieves oracle efficiency under weak model-agnostic conditions ($n^{-1/4}$ or slower convergence rates) on the nuisance function estimators.

3.1 The Projection-Learner

In this subsection, we illustrate the Projection-Learner. We first define the finite dimensional working model

$$g(v; \delta, \beta) \equiv g(v; \beta), \beta \in \mathbb{R}^p,$$

for incremental intervention parameter δ and model parameter β . This model could be for the CIE, the CICE, or the CIDE, in which case we would use δ for the CIDE and CIE, and δ_u and δ_l for the CICE. For ease of exposition, we suppress the dependence of $g(v; \delta, \beta)$ on δ (or δ_u and δ_l). In what follows, we present results in terms of the CIDE, but the results also apply to the CIE and the CICE.

Before developing the theory and methods for the Projection-Learner, we first discuss criteria for choosing the working model. There are at least two: (1) scientific context and (2) model interpretability vs misspecification. Ideally, researchers incorporate subject-specific knowledge to inform the choice of model based on the application at hand. If that does not decide what working model to use, then researchers should balance the tradeoff between model simplicity and model relevance. For instance, a researcher might choose a simple model; e.g., $g(v) = \beta_0 + \beta_1 v$. Because this is a simple model, one can easily interpret the parameter estimates of the model. However, if the model is severely misspecified, these parameters may have little bearing on the true underlying conditional effect and so the parameter estimates may be less scientifically relevant. By contrast, a researcher might choose a complex model, and believe this better captures the underlying data generating process, but estimates from such a model may be hard to interpret.

We define the projection of the CIDE onto $g(v; \beta)$ as the $g(v; \beta)$ closest to $\tau_{\text{cide}}(v; \delta)$ over weighted ℓ_2 distance. Specifically, we define β^* as the coefficients corresponding to the least-squares projection, and $g(v; \beta^*)$ as the projection. Mathematically, β^* is

$$\beta^* = \arg \min_{\beta} \int_{\mathcal{V}} \{\tau_{\text{cide}}(v; \delta) - g(v; \beta)\}^2 d\mathbf{P}(v). \quad (14)$$

One could also incorporate a weight function and use a different distance metric [45]. We set the weights to 1 and focus on ℓ_2 distance for ease of exposition, but all our results follow with other weights and could be extended to other distance metrics.

As long as $g(v; \beta)$ is differentiable with respect to β , β^* is the solution of a moment condition. The moment condition corresponds to the first derivative of the right-hand side of (14) with respect to β ,

$$m(\beta) \equiv 2 \int_{\mathcal{V}} \frac{\partial g(v; \beta)}{\partial \beta} \{\tau_{\text{cide}}(v; \delta) - g(v; \beta)\} dP(v). \quad (15)$$

Then, the solution β^* in (14) satisfies $m(\beta^*) = 0$ in (15), and the projection of the CIDE onto the working model is $g(v; \beta^*)$.

Remark 3. Our projection approach is different from the proper semiparametric approach, since the definition of β^* in equation (14) does not explicitly assume anything about the true conditional effect curve. By contrast, a proper semiparametric approach assumes a finite dimensional model is correctly specified for the conditional effect curve [52–55].

It is possible to derive an efficient influence function and thus a semiparametrically efficient estimator for the moment condition $m(\beta)$, and by extension for β^* and $g(v; \beta^*)$. We use this efficient influence function to construct the Projection-Learner. The primary building block for the efficient influence function of the moment condition is the un-centered efficient influence function for the relevant average effect. The efficient influence functions for the average incremental effect and the average incremental contrast effect were derived in Corollary 2 in Kennedy [11], and are stated in equations (45) and (46) in Appendix. Meanwhile, Lemma 1 establishes the efficient influence function for the average incremental derivative effect.

Lemma 1. Under Assumptions 1 and 2, the un-centered efficient influence function for the average incremental derivative effect, $\mathbb{E}\{\tau_{\text{cide}}(V; \delta)\}$, is

$$\begin{aligned} \xi(Z; \delta) = & \omega(X; \delta) \left[\frac{A}{\pi(X)} \{Y - \mu(1, X)\} - \frac{1 - A}{1 - \pi(X)} \{Y - \mu(0, X)\} \right] \\ & + \left[\frac{1}{\{\delta\pi(X) + 1 - \pi(X)\}^2} - \frac{2\delta\pi(X)}{\{\delta\pi(X) + 1 - \pi(X)\}^3} \right] \{A - \pi(X)\} \{\mu(1, X) - \mu(0, X)\} \\ & + \omega(X; \delta) \{\mu(1, X) - \mu(0, X)\}, \end{aligned} \quad (16)$$

where $\omega(X; \delta)$ is defined in equation (10).

The un-centered efficient influence function, $\xi(Z; \delta)$, depends only on the nuisance functions $\mu(A, X)$ and $\pi(X)$, and consists of three terms. The first term is a product of the weight term, $\omega(X; \delta)$, and an inverse weighted residual for the outcome model. The second term is a product of the difference in mean outcomes, $\mu(1, X) - \mu(0, X)$, and an inverse weighted residual for the propensity score. And, the third term is the “plug-in” for the CIDE.

Remark 4. Throughout, we invoke Assumptions 1 and 2 so that the target of estimation is some counterfactual quantity (e.g., the CIDE). If these assumptions do not hold, the results still apply if the targets of estimation are the observed data functionals on the right hand side of the identification results in Propositions 1 and 2.

From Lemma 1, and Corollary 2 in Kennedy [11], we can derive the efficient influence function for the moment condition $m(\beta)$ for estimating the projection of the CIDE, the CIE, or the CICE.

Corollary 1. Let $\xi(Z; \delta)$ denote the true influence function values of the relevant average effect, where $\xi(Z; \delta)$ is defined in (16) if the estimand is a projection of $\tau_{\text{cide}}(v; \delta)$, and is defined analogously, as shown in equations (45) and (46) in Appendix, if the estimand is a projection of $\tau_{\text{cie}}(v; \delta)$ or $\tau_{\text{cice}}(v; \delta_u, \delta_l)$. Under Assumptions 1 and 2, the un-centered efficient influence function for $m(\beta)$ with unknown propensity scores and a uniform weight function constructed over ℓ_2 distance is

$$\phi(Z; \delta, \beta) = \frac{\partial g(V; \beta)}{\partial \beta} \{\xi(Z; \delta) - g(V; \beta)\},$$

where $g(v; \beta)$ is the working model.

Corollary 1 motivates estimators for β^* and $g(v; \beta^*)$. The first step estimates the un-centered efficient influence function values for the relevant average effect; for example, when estimating the projection of the CIDE, we have

$$\begin{aligned} \hat{\xi}(Z; \delta) = & \left[\frac{\hat{\pi}(X)\{1 - \hat{\pi}(X)\}}{\{\delta\hat{\pi}(X) + 1 - \hat{\pi}(X)\}^2} \right] \left[\frac{A}{\hat{\pi}(X)} \{Y - \hat{\mu}(1, X)\} - \frac{1 - A}{1 - \hat{\pi}(X)} \{Y - \hat{\mu}(0, X)\} \right] \\ & + \left[\frac{1}{\{\delta\hat{\pi}(X) + 1 - \hat{\pi}(X)\}^2} - \frac{2\delta\hat{\pi}(X)}{\{\delta\hat{\pi}(X) + 1 - \hat{\pi}(X)\}^3} \right] \{A - \hat{\pi}(X)\} \{\hat{\mu}(1, X) - \hat{\mu}(0, X)\} \\ & + \left[\frac{\hat{\pi}(X)\{1 - \hat{\pi}(X)\}}{\{\delta\hat{\pi}(X) + 1 - \hat{\pi}(X)\}^2} \right] \{\hat{\mu}(1, X) - \hat{\mu}(0, X)\}, \end{aligned} \quad (17)$$

where $\hat{\mu}(0, X)$, $\hat{\mu}(1, X)$, and $\hat{\pi}(X)$ are (possibly nonparametric) estimates of the nuisance functions. The second step estimates the population moment condition by solving the empirical moment condition using the estimated un-centered efficient influence function values for $m(\beta)$,

$$\mathbb{P}_n \left[\frac{\partial g(V; \hat{\beta})}{\partial \beta} \{\hat{\xi}(Z; \delta) - g(V; \hat{\beta})\} \right] = 0.$$

We state the Projection-Learner formally in the following algorithm.

Algorithm 1. (*Projection-Learner*) Assume as inputs (D_1, D_2) , which denote two independent samples of n observations of $Z_i = (X_i, A_i, Y_i)$. Then,

- (1) On the training data D_1 , estimate the nuisance functions $\hat{\mu}(0, X)$, $\hat{\mu}(1, X)$, and $\hat{\pi}(X)$.
- (2) On the estimation data D_2 , estimate the un-centered influence function values $\hat{\xi}(Z; \delta)$ using the models $\hat{\mu}(0, X)$, $\hat{\mu}(1, X)$, and $\hat{\pi}(X)$ from step 1, where $\hat{\xi}(Z; \delta)$ is defined in (17) if the conditional effect of interest is τ_{cide} , and analogously for τ_{cie} and τ_{cice} in equations (45) and (46) in Appendix.
- (3) On the estimation data D_2 , estimate $\hat{\beta}$ by solving the empirical moment condition

$$\mathbb{P}_n \left[\frac{\partial g(V; \hat{\beta})}{\partial \beta} \{\hat{\xi}(Z; \delta) - g(V; \hat{\beta})\} \right] = 0.$$

Algorithm 1 is relatively straightforward. For example, if the working model is $g(v; \beta) = \beta_1 + \beta_2 \cdot v + \beta_3 \cdot v^2$, then Algorithm 1 solves the empirical moment condition

$$\mathbb{P}_n \left[\begin{pmatrix} 1 \\ V \\ V^2 \end{pmatrix} \{\hat{\xi}(Z; \delta) - (\beta_1 + \beta_2 V + \beta_3 V^2)\} \right] = 0.$$

This can be achieved in the R language and environment for statistical computing and graphics [20] by running the regression (which implicitly includes an intercept by convention)

```
model <- lm(formula = xihat ~ V + I(V^2))
```

where xihat is calculated from estimated nuisance functions $\hat{\mu}(A, X)$ and $\hat{\pi}(X)$.

Remark 5. The structure of Algorithm 1 and the example code illustrate that the Projection-Learner uses estimated un-centered efficient influence functions values for $\mathbb{E}\{\tau_{\text{cide}}(V; \delta)\}$ as pseudo-outcomes in a parametric least squares second stage regression. In Section 3.2, we show that the I-DR-Learner follows the same form, but with a nonparametric second stage regression.

To guarantee the convergence rates demonstrated in Theorem 1, we could assume Donsker-type or low-entropy conditions for the nuisance functions $\mu(A, X)$ and $\pi(X)$, which might restrict what types of flexible estimators we can use; e.g., this could preclude us from using the LASSO and many types of random forests and neural networks [34,56,57]. Instead, we use sample splitting in step 1 of Algorithm 1 to estimate the nuisance functions; i.e., we split our sample in two and estimate the nuisance functions on the training data, D_1 , and calculate $\hat{\xi}(Z; \delta)$ and solve the empirical moment condition on the estimation data, D_2 . Sample splitting allows us to condition on the training sample and treat the estimated nuisance functions as fixed functions, which allows for a large class of estimators for estimating the nuisance functions (including, e.g., the LASSO, random forests, and neural networks). A concern one might then have with Algorithm 1 is that it only estimates $\hat{\beta}$ on half the sample. To utilize the whole sample for inference, we can improve on Algorithm 1 with cross-fitting by estimating the nuisance functions on both folds (D_1 and D_2), constructing $\hat{\xi}(Z; \delta)$ values on the opposite fold (i.e., by estimating $\hat{\xi}(Z; \delta)$ in D_1 using nuisance functions constructed on D_2 , and vice versa), and solving the empirical moment condition on the whole dataset (D_1 and D_2 together) [37,58,59]. This cross-fitting approach is also compatible with more folds (“k-fold cross-fitting”), which can be more stable than two-fold cross-fitting.

The following theorem shows that the estimator $\hat{\beta}$ for β^* outlined in Algorithm 1 converges to an asymptotically linear expansion around β^* where the bias is expressed as a product of errors from estimating the nuisance functions $\hat{\mu}(0, X)$, $\hat{\mu}(1, X)$, and $\hat{\pi}(X)$. For this result, and the rest of Section 3.1, we let $\mu = \{\mu(0, X), \mu(1, X)\}$ and $\pi = \pi(X)$ denote the generic nuisance functions, $\hat{\mu}$ and $\hat{\pi}$ denote the nuisance function estimators, and define μ^* and π^* as the true nuisance functions (consistent with the projection notation β^*).

Theorem 1. Let $\varphi(Z; \beta, \mu, \pi) \equiv \phi(Z; \delta, \beta) - m(\beta)$ denote the centered efficient influence function from Corollary 1. With Assumptions 1 and 2, also assume

(a) $\mathbb{P}(|\hat{\mu}(1, X) - \hat{\mu}(0, X)| \leq C) = 1$ and $\mathbb{P}(|\mu^*(1, X) - \mu^*(0, X)| \leq C) = 1$ for some $C < \infty$.

(b) $\mathbb{P}\left\{\left|\frac{\partial g(\beta; v)}{\partial \beta}\right| \leq C\right\} = 1$ for all v .

(c) The function class $\varphi(Z; \beta, \mu, \pi)$ is Donsker in β for any fixed μ, π .

(d) The estimators are consistent in the sense that $\hat{\beta} - \beta^* = o_{\mathbb{P}}(1)$ and $\|\varphi(Z; \hat{\beta}; \hat{\mu}, \hat{\pi}) - \varphi(Z; \beta^*; \mu^*, \pi^*)\| = o_{\mathbb{P}}(1)$.

(e) The map $\beta \mapsto \mathbb{P}\{\varphi(Z; \beta, \mu, \pi)\}$ is differentiable at β^* uniformly in (μ, π) , with nonsingular derivative matrix

$$\frac{\partial}{\partial \beta} \mathbb{P}\{\varphi(Z; \beta, \mu, \pi)\}_{|\beta=\beta^*} = M(\beta^*, \mu, \pi), \text{ where } M(\beta^*, \hat{\mu}, \hat{\pi}) \xrightarrow{P} M(\beta^*, \mu^*, \pi^*).$$

Then,

$$\hat{\beta} - \beta^* = -M(\beta^*, \mu^*, \pi^*)^{-1}(\mathbb{P}_n - \mathbb{P})\{\varphi(Z; \beta^*, \mu^*, \pi^*)\} + O_{\mathbb{P}}\left\{R_n + o_{\mathbb{P}}\left(\frac{1}{\sqrt{n}}\right)\right\},$$

where

$$R_n = (\|\hat{\mu} - \mu^*\| + \|\hat{\pi} - \pi^*\|)\|\hat{\pi} - \pi^*\|.$$

Theorem 1 provides a convergence statement for the coefficient estimate $\hat{\beta}$ to the true projection parameter β^* under relatively weak conditions. Assumption (a) says that the CATE and the estimate of the CATE are bounded. Assumption (b) says that the derivative of the model $g(v; \beta)$ with respect to β is bounded, which is quite weak and can be enforced through choice of an appropriate model. Assumption (c) ensures that the influence function φ is not too complex as a function of β , but allows for arbitrary complexity in the nuisance functions; again, this can be enforced with appropriate choice of $g(v; \beta)$, and most reasonable choices will suffice. Assumption (d) requires that $\{\beta^*, \varphi(Z; \beta^*, \mu^*, \pi^*)\}$ is consistently estimated by $\{\hat{\beta}, \varphi(Z; \hat{\beta}, \hat{\mu}, \hat{\pi})\}$ at any rate. Finally, Assumption (e) requires some smoothness of $\mathbb{P}\{\varphi(Z; \beta, \mu, \pi)\}$ in β , to allow for use of the delta method. Assumptions (c)–(e) are standard in the literature (see, for example, [34] Theorem 5.31).

The convergence statement shows that $\hat{\beta}$ obtains a faster rate of convergence to β^* than the nuisance function estimators $\hat{\mu}$ and $\hat{\pi}$ obtain to μ^* and π^* , respectively. The first term, $M(\beta^*, \mu^*, \pi^*)^{-1}(\mathbb{P}_n - \mathbb{P})\{\varphi(Z; \beta^*, \mu^*, \pi^*)\}$, is a

sample average scaled by a constant, and so by the central limit theorem, it is asymptotically Gaussian. Therefore, if $R_n = o_P(n^{-1/2})$, then the remainder terms $O_P(R_n + o_P(n^{-1/2}))$ will be asymptotically negligible and so $\hat{\beta} - \beta^*$ will converge in distribution to a mean-zero Gaussian distribution with variance equal to the variance of $M(\beta^*, \mu^*, \pi^*)^{-1}(\mathbb{P}_n - \mathbb{P})\{\varphi(Z; \beta^*, \mu^*, \pi^*)\}$, as shown in the following result.

Corollary 2. *Under the same assumptions as Theorem 1, if*

$$(\|\hat{\mu} - \mu^*\| + \|\hat{\pi} - \pi^*\|)\|\hat{\pi} - \pi^*\| = o_P\left(\frac{1}{\sqrt{n}}\right),$$

then

$$\sqrt{n}(\hat{\beta} - \beta^*) \rightsquigarrow N(0, M^{-1}\mathbb{E}(\varphi\varphi^T)(M^{-1})^T),$$

and for any fixed v we have

$$\sqrt{n}\{g(v; \hat{\beta}) - g(v; \beta^*)\} \rightsquigarrow N\left(0, \left(\frac{\partial g(v; \beta^*)}{\partial \beta}\right)^T M^{-1}\mathbb{E}(\varphi\varphi^T)(M^{-1})^T \left(\frac{\partial g(v; \beta^*)}{\partial \beta}\right)\right),$$

where

$$M = M(\beta^*, \mu^*, \pi^*), \quad \text{and} \quad \mathbb{E}(\varphi\varphi^T) = \mathbb{E}\{\varphi(Z; \beta^*, \mu^*, \pi^*) \times \varphi(Z; \beta^*, \mu^*, \pi^*)^T\}.$$

Corollary 2 provides a way to construct an asymptotically valid Wald-style $1-\alpha$ confidence interval around $g(\hat{\beta}, v)$ with

$$g(v; \hat{\beta}) \pm \Phi^{-1}(1 - \alpha/2) \left(\frac{\hat{\sigma}(v)}{\sqrt{n}} \right),$$

where $\Phi(\cdot)$ is the cumulative distribution function for the standard normal,

$$\hat{\sigma}^2(v) = \left(\frac{\partial g(v; \hat{\beta})}{\partial \beta} \right)^T \hat{M}^{-1} \mathbb{E}(\hat{\varphi}\hat{\varphi}^T) (\hat{M}^{-1})^T \left(\frac{\partial g(v; \hat{\beta})}{\partial \beta} \right),$$

and $\hat{M} = \mathbb{P}_n(\partial \hat{\varphi} / \partial \beta)$ is an estimate of the derivative matrix. Furthermore, this corollary demonstrates that $\hat{\beta}$ and $g(v; \hat{\beta})$ converge at $n^{-1/2}$ rates to Gaussian distributions, centered at β^* and $g(v; \beta^*)$, respectively, with less stringent model-agnostic convergence conditions on the nuisance function estimators $\hat{\mu}(0, X)$, $\hat{\mu}(1, X)$, and $\hat{\pi}(X)$. Thus, $\hat{\beta}$ and $g(v; \hat{\beta})$ still attain $n^{-1/2}$ convergence rates if both nuisance functions are estimated at $n^{-1/4}$ rates, which are attainable with nonparametric estimators under relatively realistic assumptions such as smoothness or sparsity [60–63].

Remark 6. These results are doubly-robust in spirit since the remainder bias is expressed as a product of nuisance function errors. However, there is no “double robustness” in the traditional sense, which would only require $\|\hat{\mu} - \mu^*\| \|\hat{\pi} - \pi^*\| = o_P(n^{-1/2})$. Instead, Corollary 2 requires that the propensity score is estimated well enough that $\|\hat{\pi} - \pi^*\|^2 = o_P(n^{-1/2})$. Intuitively, this occurs because incremental interventions shift the observed propensity scores and thus require a good estimate of the propensity score. By contrast, the intervention corresponding to the CATE does not depend on the propensity score, so the convergence rate for the propensity score estimator is less critical, depending on that of the outcome regression.

As demonstrated in Theorem 1 and Corollary 2, the Projection-Learner can attain $n^{-1/2}$ convergence rates to the projection of the true CIDE (or CIE, or CICE) onto the chosen working model $g(v; \beta)$. If, instead, we wish to target the true conditional effect curve, and that curve does not coincide with the projection, then we need to use a different estimator, as described in Section 3.2.

3.2 The I-DR-Learner

In this section, we outline the I-DR-Learner and illustrate its convergence properties. The I-DR-Learner targets the true conditional effects. Since the conditional effects are not pathwise differentiable, no efficient influence function exists for them. Instead, the I-DR-Learner makes use of the efficient influence function values for the relevant average effect by regressing them against the covariates of interest to estimate the conditional effect. In this way, the I-DR-Learner is a two-stage meta-learner, where the first stage estimates the efficient influence function values for the relevant average effect, and the second stage uses these values as pseudo-outcomes in a second stage regression against the conditioning covariates. The I-DR-Learner is stated formally in the following algorithm:

Algorithm 2. (*I-DR-Learner*). Assume as inputs (D_1, D_2) , which denote two independent samples of n observations of $Z_i = (X_i, A_i, Y_i)$.

- (1) On the training data D_1 , estimate the nuisance functions $\hat{\mu}(0, X)$, $\hat{\mu}(1, X)$, and $\hat{\pi}(X)$.
- (2) On the estimation data D_2 , estimate the un-centered influence function values $\hat{\xi}(Z; \delta)$ using the models $\hat{\mu}(0, X)$, $\hat{\mu}(1, X)$, and $\hat{\pi}(X)$ from step 1, where $\hat{\xi}(Z; \delta)$ is defined as in (17) if the conditional effect of interest is τ_{cde} , and analogously for τ_{cie} and τ_{cice} in equations (45) and (46) in Appendix.
- (3) In the estimation sample D_2 , regress $\hat{\xi}(Z; \delta)$ on the conditioning covariates V to obtain the estimate

$$\hat{\tau}_{i-\text{dr}}(v; \delta) = \hat{\mathbb{E}}_n\{\hat{\xi}(Z; \delta) | V = v\}.$$

Like the Projection-Learner, the I-DR-Learner also uses sample splitting and estimates the nuisance functions on a separate sample to avoid imposing Donsker-type conditions on the nuisance function estimators. The I-DR-Learner is also compatible with cross-fitting.

The I-DR-Learner can estimate all three conditional effects – the CIE, CICE, and CIDE. Furthermore, the error of the estimator is asymptotically equal to that of an oracle estimator under certain conditions. Specifically, the second stage regression must satisfy the stability condition in Definition 1 [3]. This is a generalization of the classic stochastic equicontinuity condition to nonparametric regression (Lemma 19.24 [34]), and says that the second stage regression is *stable* with respect to a distance metric d if the difference between second stage regressions with estimated outcomes and true outcomes shrinks appropriately fast. For example, in our context, a regression estimator $\hat{\mathbb{E}}_n$ is stable if $\hat{\xi} \rightarrow \xi$ implies that $\hat{\mathbb{E}}_n\{\hat{\xi}(Z; \delta) | V = v\}$ converges to $\mathbb{E}_n\{\xi(Z; \delta) | V = v\}$ at an appropriately fast rate. We discuss Definition 1 of Kennedy [3] in more detail in Appendix. The stability condition is satisfied by the class of linear smoothers (see [3] Theorem 1), which includes nearest neighbors estimators, regression splines, kernel smoothers, series regression, and certain random forests. It is possible that other classes of estimators also satisfy the stability condition, although examining that question is beyond the scope of this work.

Under this stability condition, the error of the I-DR-Learner can be tied to the error of an oracle estimator, which would have access to the un-centered efficient influence function values for the relevant average effect and would estimate the conditional effect merely by running a regression of $\xi(Z; \delta)$ against V . This approach was considered in Kennedy [3] for estimating the CATE, and their Theorem 2 showed that, under certain assumptions, the error of their DR-Learner will only exceed the error of an oracle estimator by an amount that depends on the product of errors in estimating the nuisance functions. The same logic holds for the I-DR-Learner, and we formally state the convergence result in the following theorem. We slightly amend notation from Section 3.1, and allow ξ , μ , and π to denote the true efficient influence function values and nuisance functions.

Theorem 2. Let $\tau_{i-\text{dr}}$ stand in for τ_{cde} , τ_{cie} , or τ_{cice} , and let $\xi(Z; \delta)$ denote the true influence function values of the relevant average effect. Furthermore, let $\tilde{\tau}_{i-\text{dr}}(v; \delta) = \mathbb{E}_n\{\xi(Z; \delta) | V = v\}$ denote an oracle estimator that regresses $\xi(Z; \delta)$ on V , and let $\hat{\tau}_{i-\text{dr}}(v; \delta)$ denote the I-DR-Learner from Algorithm 2. With Assumptions 1 and 2,

and Assumption (a) from Theorem 1, also assume that the second stage regression is stable according to Definition 1 given in the study by Kennedy [3]. Then,

$$\hat{\tau}_{i-dr}(v; \delta) - \tau_{i-dr}(v; \delta) = \tilde{\tau}_{i-dr}(v; \delta) - \tau_{i-dr}(v; \delta) + \hat{\mathbb{E}}_n\{\hat{b}(X)|V = v\} + o_P(R^*(v; \delta))$$

for

$$\hat{b}(x) \leq (|\hat{\mu}(0, x) - \mu(0, x)| + |\hat{\mu}(1, x) - \mu(1, x)| + |\hat{\pi}(x) - \pi(x)|) \cdot |\hat{\pi}(x) - \pi(x)|$$

and

$$R^*(v; \delta) = \sqrt{\mathbb{E}\{\{\tilde{\tau}_{i-dr}(v; \delta) - \tau_{i-dr}(v; \delta)\}^2\}}.$$

Theorem 2 shows that error for the I-DR-Learner differs from the error for the oracle estimator by at most $\hat{\mathbb{E}}_n\{\hat{b}(X)|X = x\}$ plus other terms, captured by $o_P(R^*(v; \delta))$, that are asymptotically negligible compared to the error of the oracle estimator. Thus, whether the I-DR-Learner achieves oracle efficiency is driven by the asymptotic behavior of the smoothed bias term $\hat{\mathbb{E}}_n\{\hat{b}(X)|X = x\}$. This bias term is asymptotically less than the product of errors for estimating $\mu(x)$ and $\pi(x)$, $|\hat{\mu}(a, x) - \mu(a, x)| \cdot |\hat{\pi}(x) - \pi(x)|$, and the squared error for estimating $\pi(x)$, $\{\hat{\pi}(x) - \pi(x)\}^2$. Therefore, the convergence rate of the I-DR-Learner is faster than the convergence rate of the nuisance function estimators. For example, if the nuisance functions are estimated at $n^{-1/4}$ rates, then the bias term $\hat{b}(X)$ will converge to zero at a $n^{-1/2}$ rate. Importantly, Theorem 2 does not require any assumptions about how the estimators $\hat{\mu}$ and $\hat{\pi}$ are constructed, beyond the boundedness conditions from Assumption (a) from Theorem 1.

However, the performance of the I-DR-Learner is also constrained by the oracle convergence rate for the second stage regression. For example, if $\mathbb{E}\{\xi(Z; \delta)|V = v\}$ is Hölder-smooth with smoothness s , then the minimax rate in root mean squared error is $n^{-1/(2+\frac{4}{s})}$, which is slower than $n^{-1/2}$ [60]. This is not surprising, since the conditional effect is a regression function, if we are only willing to assume it lies in a large nonparametric class, then the minimax rate of convergence will be slower than $n^{-1/2}$. One can also think of the slower oracle convergence as a positive aspect to the I-DR-Learner, since it reduces how well the nuisance functions must be estimated to achieve oracle efficiency. For example, if the oracle convergence rate is “only” $n^{-1/4}$, then the I-DR-Learner can estimate each nuisance function at $n^{-1/8}$ convergence rates and still attain oracle efficiency. When the nuisance functions are estimated well enough and the I-DR-Learner is oracle efficient, confidence bands can be constructed following well-known processes for nonparametric regression [64].

Both the Projection-Learner and the I-DR-Learner can be used to estimate conditional effect curves across δ and V , thereby quantifying how causal effects vary across V . A natural question is whether there is any treatment effect heterogeneity across V . In Section 4, we outline how to quantify and test for treatment effect heterogeneity.

4 Understanding effect heterogeneity with the V-CIDE

There is a large literature for understanding treatment effect heterogeneity by summarizing the CATE (e.g., [65–68]). In this section, we demonstrate how the V-CIDE, defined in equation (11), can be used to understand effect heterogeneity. To ease exposition, we focus on the case where $V = X$, and examine effect heterogeneity across all covariates. The case where V is a strict subset of X (i.e., $V \subset X$) is outlined in Appendix. By Proposition 2, the V-CIDE is identified by

$$\mathbb{V}\{\tau_{cide}(X; \delta)\} = \mathbb{V}[\omega(X; \delta)\{\mu(1, X) - \mu(0, X)\}].$$

When the V-CIDE is zero, the derivative is constant across V , and so shifting the treatment distribution has the same effect on all subjects. If the V-CIDE is greater than zero, then there is treatment effect heterogeneity in the incremental effect.

We construct an estimator in two pieces by first noting that the V-CIDE is the difference between two effects since $\mathbb{V}\{\tau_{\text{cide}}(X; \delta)\} = \mathbb{E}\{\tau_{\text{cide}}(X; \delta)^2\} - \mathbb{E}\{\tau_{\text{cide}}(X; \delta)\}^2$. The first effect, $\mathbb{E}\{\tau_{\text{cide}}(X; \delta)^2\}$, admits an efficient influence function by the following lemma:

Lemma 2. Under Assumptions 1 and 2, the un-centered efficient influence function for $\mathbb{E}\{\tau_{\text{cide}}(X; \delta)^2\}$ is

$$2\omega(X; \delta)\{\mu(1, X) - \mu(0, X)\}[\omega(X; \delta)\phi(Z) + \phi(Z; \delta)\{\mu(1, X) - \mu(0, X)\}] + [\omega(X; \delta)\{\mu(1, X) - \mu(0, X)\}]^2$$

where $\omega(X; \delta)$ is defined in equation (10) and

$$\phi(Z) = \frac{A}{\pi(X)}\{Y - \mu(1, X)\} - \frac{1 - A}{1 - \pi(X)}\{Y - \mu(0, X)\}, \quad (18)$$

$$\phi(Z; \delta) = \left[\frac{1}{\{\delta\pi(X) + 1 - \pi(X)\}^2} - \frac{2\delta\pi(X)}{\{\delta\pi(X) + 1 - \pi(X)\}^3} \right] \{A - \pi(X)\}. \quad (19)$$

Lemma 2 shows that the un-centered efficient influence function for $\mathbb{E}\{\tau_{\text{cide}}(X; \delta)^2\}$ can be written as a weighted residual plus a plug-in. The second effect, $\mathbb{E}\{\tau_{\text{cide}}(X; \delta)\}^2$, is also pathwise differentiable and admits an efficient influence function. However, since it is a smooth transformation of an already pathwise differentiable function, we estimate it by squaring the estimator based on the efficient influence function for $\mathbb{E}\{\tau_{\text{cide}}(X; \delta)\}$ provided in Lemma 1. Therefore, informed by Lemmas 1 and 2, we propose the estimator

$$\widehat{\mathbb{V}}\{\tau_{\text{cide}}(X; \delta)\} = \underbrace{\mathbb{P}_n[2\widehat{\omega}(\widehat{\mu}_1 - \widehat{\mu}_0)\{\widehat{\omega}\widehat{\phi} + \widehat{\phi}(\widehat{\mu}_1 - \widehat{\mu}_0)\} + \{\widehat{\omega}(\widehat{\mu}_1 - \widehat{\mu}_0)\}^2]}_{\text{Estimator for } \mathbb{E}\{\tau_{\text{cide}}(X; \delta)^2\}}, \quad (20)$$

$$- \underbrace{[\mathbb{P}_n\{\widehat{\omega}\widehat{\phi} + \widehat{\phi}(\widehat{\mu}_1 - \widehat{\mu}_0) + \widehat{\omega}(\widehat{\mu}_1 - \widehat{\mu}_0)\}]^2}_{\text{Estimator for } \mathbb{E}\{\tau_{\text{cide}}(X; \delta)\}^2}, \quad (21)$$

where we omit δ , X , and Z arguments and let $\mu_a = \mu(a, X)$ for brevity, and where $\widehat{\omega}$, $\widehat{\phi}$, $\widehat{\phi}$ indicate the relevant formulae from (10), (18), and (19), but with the estimated nuisance functions (e.g., $\widehat{\omega} = \frac{\widehat{\pi}(X)\{1 - \widehat{\pi}(X)\}}{\{\delta\widehat{\pi}(X) + 1 - \widehat{\pi}(X)\}^2}$). Equation (20) is the estimator for $\mathbb{E}\{\tau_{\text{cide}}(X; \delta)^2\}$ motivated by Lemma 2, it takes the sample average of the estimated un-centered efficient influence function values for $\mathbb{E}\{\tau_{\text{cide}}(X; \delta)^2\}$. Equation (21) is the estimator for $\mathbb{E}\{\tau_{\text{cide}}(X; \delta)\}^2$ motivated by Lemma 1, it squares the estimator for $\mathbb{E}\{\tau_{\text{cide}}(X; \delta)\}$, which itself is just the sample average of the estimated un-centered efficient influence function values for $\mathbb{E}\{\tau_{\text{cide}}(X; \delta)\}$ in equation (17). Formally, we outline the estimator in the following algorithm:

Algorithm 3. (*V-CIDE Estimator*) Assume as inputs (D_1, D_2) , which denote two independent samples of n observations of $Z_i = (X_i, A_i, Y_i)$, then:

- (1) On the training data D_1 , estimate the nuisance functions $\widehat{\mu}(0, X)$, $\widehat{\mu}(1, X)$, and $\widehat{\pi}(X)$.
- (2) On the estimation data D_2 , estimate $\mathbb{V}\{\tau_{\text{cide}}(X; \delta)\}$ per equations (20) and (21), plugging in the estimates for $\widehat{\mu}(0, X)$, $\widehat{\mu}(1, X)$, and $\widehat{\pi}(X)$ using the models from step 1.

As before, Algorithm 3 uses sample splitting to estimate the nuisance functions, which allows for estimating the nuisance functions with flexible machine learning models. Again, this estimator could use cross-fitting by repeating the algorithm but with D_1 and D_2 reversed and then averaging the two estimates. We establish the error guarantees of the estimator in the following result.

Theorem 3. Let $\widehat{\psi}_n$ denote the estimator from Algorithm 3. With Assumptions 1 and 2, and Assumption (a) from Theorem 1, also assume that

- (a) $\mathbb{P}\{\omega(\mu_1 - \mu_0) + \omega\phi + \phi\tau \leq C\}$ and $\mathbb{P}\{\widehat{\omega}(\widehat{\mu}_1 - \widehat{\mu}_0) + \widehat{\omega}\widehat{\phi} + \widehat{\phi}\widehat{\tau} \leq C\} = 1$ for some $C < \infty$.

If

$$(\|\widehat{\mu} - \mu\| + \|\widehat{\pi} - \pi\|)^2 = o_p\left(\frac{1}{\sqrt{n}}\right),$$

then

$$\sqrt{n}[\hat{\psi}_n - \mathbb{V}\{\tau_{\text{cide}}(X; \delta)\}] \rightsquigarrow N(0, \sigma^2),$$

where

$$\begin{aligned} \sigma^2 = & \mathbb{V}[2\omega(\mu_1 - \mu_0)\{\omega\phi + \phi(\mu_1 - \mu_0)\} + \{\omega(\mu_1 - \mu_0)\}^2 \\ & - \mathbb{E}\{\omega\phi + \phi(\mu_1 - \mu_0) + \omega(\mu_1 - \mu_0)\}\{\omega\phi + \phi(\mu_1 - \mu_0) + \omega(\mu_1 - \mu_0)\}], \end{aligned} \quad (22)$$

$\mu_a = \mu(a, X)$, and $\omega = \omega(X; \delta)$, $\phi = \phi(Z)$, and $\phi = \phi(Z; \delta)$ as defined in equations (10), (18), and (19).

Theorem 3 shows that the estimator for the V-CIDE satisfies a version of double robustness under relatively weak conditions. Assumption (a) says that the efficient influence function for the average derivative and the estimate for the efficient influence function are bounded, which is a mild assumption. Then, if both nuisance function estimators converge at $n^{-1/4}$ rates, the standardized difference between the estimator and the V-CIDE has a Gaussian limiting distribution. This is a slightly stronger requirement than that of Corollary 2, since both nuisance functions must be estimated at $n^{-1/4}$ rates, not just the propensity score. This occurs due to the nonlinearity of $\mathbb{E}\{\tau_{\text{cide}}(X; \delta)\}^2$ in terms of μ . Nonetheless, this result is still model-agnostic about the nuisance function estimators, and the convergence requirement can be satisfied by non-parametric estimators under suitable smoothness or sparsity. Theorem 3 suggests constructing Wald-style $1 - \alpha$ confidence intervals with

$$\hat{\psi}_n \pm \Phi^{-1}(1 - \alpha/2) \sqrt{\frac{\hat{\sigma}^2}{n}}, \quad (23)$$

where $\hat{\sigma}^2$ is the sample variance estimator for σ^2 defined in equation (22); i.e.,

$$\begin{aligned} \hat{\sigma}^2 = & \mathbb{V}_n[2\hat{\omega}(\hat{\mu}_1 - \hat{\mu}_0)\{\hat{\omega}\hat{\phi} + \hat{\phi}(\hat{\mu}_1 - \hat{\mu}_0)\} + \{\hat{\omega}(\hat{\mu}_1 - \hat{\mu}_0)\}^2 \\ & - \mathbb{P}_n\{\hat{\omega}\hat{\phi} + \hat{\phi}(\hat{\mu}_1 - \hat{\mu}_0) + \hat{\omega}(\hat{\mu}_1 - \hat{\mu}_0)\}\{\hat{\omega}\hat{\phi} + \hat{\phi}(\hat{\mu}_1 - \hat{\mu}_0) + \hat{\omega}(\hat{\mu}_1 - \hat{\mu}_0)\}], \end{aligned} \quad (24)$$

where \mathbb{V}_n denotes the sample variance.

Unfortunately, the estimator in Algorithm 3 converges to a degenerate distribution when $\mathbb{V}\{\tau_{\text{cide}}(X; \delta)\} = 0$ because the efficient influence function values are identically zero, and $\sigma^2 = 0$ in equation (22). So, the confidence interval in (23) would under-cover the true parameter. Instead, we can construct a conservative estimate of the variance by noting that the efficient influence function values of $\mathbb{E}\{\tau_{\text{cide}}(X; \delta)\}^2$ and $\mathbb{E}\{\tau_{\text{cide}}(X; \delta)\}^2$ have non-negative covariance when $\mathbb{V}\{\tau_{\text{cide}}(X; \delta)\} = 0$ (which is stated formally in Proposition 5 in the appendix). This suggests a simple way to conservatively estimate the variance of $\hat{\psi}_n$ and construct a valid confidence interval with

$$\hat{\psi}_n \pm \Phi^{-1}(1 - \alpha/2) \sqrt{\frac{\hat{\sigma}_1^2 + \hat{\sigma}_2^2}{n}}, \quad (25)$$

where

$$\hat{\sigma}_1^2 = \hat{\mathbb{V}}_n[2\hat{\omega}(\hat{\mu}_1 - \hat{\mu}_0)\{\hat{\omega}\hat{\phi} + \hat{\phi}(\hat{\mu}_1 - \hat{\mu}_0)\} + \{\hat{\omega}(\hat{\mu}_1 - \hat{\mu}_0)\}^2], \text{ and} \quad (26)$$

$$\hat{\sigma}_2^2 = \mathbb{V}_n[\mathbb{P}_n\{\hat{\omega}\hat{\phi} + \hat{\phi}(\hat{\mu}_1 - \hat{\mu}_0) + \hat{\omega}(\hat{\mu}_1 - \hat{\mu}_0)\}\{\hat{\omega}\hat{\phi} + \hat{\phi}(\hat{\mu}_1 - \hat{\mu}_0) + \hat{\omega}(\hat{\mu}_1 - \hat{\mu}_0)\}] \quad (27)$$

are, respectively, consistent estimators of the variance of the estimators in equations (20) and (21) for $\mathbb{E}\{\tau_{\text{cide}}(X; \delta)\}^2$ and $\mathbb{E}\{\tau_{\text{cide}}(X; \delta)\}^2$. This confidence interval suggests the following one-sided test for treatment effect heterogeneity:

$$\begin{cases} \text{Reject } H_0 : \mathbb{V}\{\tau_{\text{cide}}(X; \delta)\} = 0 & \text{if } \hat{\psi}_n - \Phi^{-1}(1 - \alpha) \sqrt{\frac{\hat{\sigma}_1^2 + \hat{\sigma}_2^2}{n}} > 0, \\ \text{Fail to reject } H_0 : \mathbb{V}\{\tau_{\text{cide}}(X; \delta)\} = 0 & \text{otherwise.} \end{cases} \quad (28)$$

This test controls Type I error at the appropriate level, as shown in the following result.

Proposition 3. *Under Assumptions 1 and 2, Assumption (a) from Theorem 1, and Assumption (a) from Theorem 3, if*

$$(\|\hat{\mu} - \mu\| + \|\hat{\pi} - \pi\|)^2 = o_p\left(\frac{1}{\sqrt{n}}\right),$$

then the asymptotic Type I error rate of the test in (28) is less than or equal to α .

Remark 7. In the causal inference literature, at least two other solutions have been proposed for constructing confidence intervals when an estimator converges to a degenerate distribution. Our approach is similar to that reported in the study by Williamson *et al.* [69], where they focus on testing variable importance. Luedtke *et al.* [68] proposed a different approach – they derived the higher order influence function for their parameter and constructed an associated estimator that achieves n^{-1} convergence under $n^{-1/4}$ conditions on the nuisance function estimators.

Remark 8. When we do not have knowledge of the true parameter value, and we want to construct a valid confidence interval (rather than conduct a test), we can combine the confidence intervals in (23) and (25) and construct a conservative confidence interval with

$$\hat{\psi}_n \pm \Phi^{-1}(1 - \alpha/2) \sqrt{\frac{\max(\hat{\sigma}^2, \hat{\sigma}_1^2 + \hat{\sigma}_2^2)}{n}}.$$

Remark 9. A benefit of the V-CIDE is that it does not require positivity, unlike standard tests for effect heterogeneity like those based on the variance of the CATE (e.g., tests that use the statistic $\mathbb{V}\{\tau_{\text{cate}}(X)\}$ [68]). However, a drawback of our approach is that the test based on the V-CIDE may be less powerful than those based directly on the CATE, depending on δ and the distribution of the propensity score, which determine the magnitude of the weight that multiplies the CATE in the identified formula for the V-CIDE on the right-hand side of equation (13).

In the appendix, we illustrate several simulations that demonstrate the properties of the Projection-Learner and I-DR-Learner. In short, the Projection-Learner achieves correct coverage for the projection parameter and the I-DR-Learner achieves oracle efficiency when the nuisance functions are estimated well enough. In Section 5, we apply these estimators to real ICU data and demonstrate how they can uncover interesting phenomena that would be obscured by looking at effects with deterministic interventions, like the ATE and the CATE.

5 Data analysis of the effect of ICU admission on mortality

In this section we illustrate the I-DR-Learner and the estimator for the V-CIDE by analyzing data from the (SPOT)light prospective cohort study in which investigators collected data on ICU transfers and mortality. These data are a cohort study collected between November 1st, 2010 and December 31st, 2011 of 13,011 patients with deteriorating health who were assessed for critical care unit admission across 49 National Health Service hospitals in the UK [8,70].

Previous literature has considered whether admission to the ICU reduces mortality [71,72], where the relevant exposure of interest is a binary indicator of whether someone was admitted to the ICU. Recent analyses have estimated the ATE or used ICU bed availability as an instrumental variable to estimate the local average treatment effect (LATE) [8]. Flexible estimation of the ATE finds that the ICU is harmful, whereas estimates for the LATE find a null effect, albeit with wide confidence intervals. However, there are two




Figure 2: Propensity scores by ICNARC score.

limitations to focusing on the ATE or LATE for this application. First, the relevant counterfactual interventions where everyone is sent to the ICU or no one is sent to the ICU may not be feasible (e.g., the ICU might not have capacity to admit everyone), but an intervention where it is made more or less likely that people are sent to the ICU could be feasible. Second, one might expect a priori that the positivity assumption is violated, in the sense that some patients – depending on their condition – may be almost certain to be admitted or never be admitted to the ICU. Indeed, this is validated by the data, as shown in Figure 2; thus, an intervention that does not require positivity is desirable for this application. Finally, understanding effect heterogeneity would be of great interest in this application, since it may be the case that the ICU is helpful for some patients while unhelpful or even harmful for others. While there are other interventions one might consider for describing effect heterogeneity under realistic interventions that are robust to positivity violations [9,10,13,17], incremental interventions are a natural candidate because they take an intuitive parameterization for binary treatment and are robust to positivity violations when the propensity scores equal zero or one.

5.1 Data

The data contain 28-day mortality as an outcome variable and a binary indicator for whether someone was admitted to the ICU. The data also contain detailed demographic, physiological, comorbidity, and mortality information for all patients. In terms of demographic information, the data include age, sex, septic diagnosis (0/1), and peri-arrest (0/1). In terms of physiology data, there are three risk scores: the ICNARC physiology score [73], the NHS National Early Warning score [74], and the Sepsis-related Organ Failure Assessment score [75]. Finally, the data also record the patient's existing level of care at assessment and recommended level of care after assessment, which were defined using the UK Critical Care Minimum Dataset levels of care. We used all these covariates in our analysis, and also included ICU bed availability, which is a binary measure of whether <4 ICU beds were available at the time of assessment. We use a publicly available version of the dataset, which contains the same number of observations sampled with replacement from the original dataset – we provide this dataset with our code at <https://github.com/alecmcclean/NPCIE>.

5.2 Method

We consider the counterfactual 28-day mortality rate if we increased or decreased the odds of ICU admission according to an incremental intervention. We use the I-DR-Learner to nonparametrically estimate the CIE and the CIDE over the ICNARC physiology score. We focus on the ICNARC score because it is a measure of the health risk of the patient (higher being riskier), and a natural question is whether the ICU affects healthier and sicker patients differently. Then, we estimate the V-CIDE to test for treatment effect heterogeneity across a continuum of policies. The nuisance functions $\hat{\pi}$ and $\hat{\mu}$ – and $\hat{\tau}_{\text{cide}}(\delta)$ when estimating the V-CIDE – were estimated with random forests via the *ranger* package and the efficient influence function pseudo-outcomes for the CIE were constructed using the *npcausal* package in R [21,77]. The I-DR-Learner second stage regression was estimated with a smoothing spline via the *mgcv*: :GAM function in R [78]. R code demonstrating how our analyses were implemented is provided at <https://github.com/alecmcclean/NPCIE>, and functions for estimating the average incremental derivative effect and V-CIDE are available in the *npcausal* package [21].

5.3 Results

Figure 2 shows estimated propensity scores by ICNARC score, which confirms prior intuition that positivity might be violated with these data, since for most ICNARC scores, there are estimated propensity scores very near 0 and 1. Figure 3(a) shows that the CIE varies across δ for all ICNARC scores. Estimated counterfactual 28-day mortality is lowest under the observed treatment process (when $\delta = 1$), and increases when the odds of ICU admission increase ($\delta > 1$) or decrease ($\delta < 1$). We also see strong evidence that the CIE varies across ICNARC score, and mortality increases as the ICNARC score increases. This agrees with what one might expect, since the ICNARC is a risk measure where a higher ICNARC score denotes a patient with a higher risk of death. However, this does not necessarily suggest treatment effect heterogeneity, since one would need to consider a contrast between two levels of the CIE to understand effect heterogeneity.

To ease presentation, Figure 3(b) shows the CIE across δ for only four ICNARC scores (0, 15, 30, and 40 were chosen because they are roughly evenly spaced across the range of ICNARC scores). Examining only four curves shows more clearly that the shape of the CIE at each ICNARC score is very similar, suggesting that perhaps there is little heterogeneity. Figure 3(b) also illustrates a further nuance. Previous work has estimated the ATE and found that mortality rates would be higher if everyone were admitted to the ICU vs if no one was admitted [8]. Taken at face value, this suggests that hospitals ought to send fewer people to the ICU; however, due to positivity issues in the data, ATE estimates are likely invalid. The difference between the endpoints of the curves in Figure 3(b) (i.e., $\tau_{\text{cice}}(v; \delta_u = 5, \delta_l = 0.2)$) suggests a similar conclusion to that implied by ATE




Figure 3: Predicted CIE by δ and ICNARC score. (a) Predicted CIE for all ICNARC scores and (b) predicted CIE for select ICNARC scores.

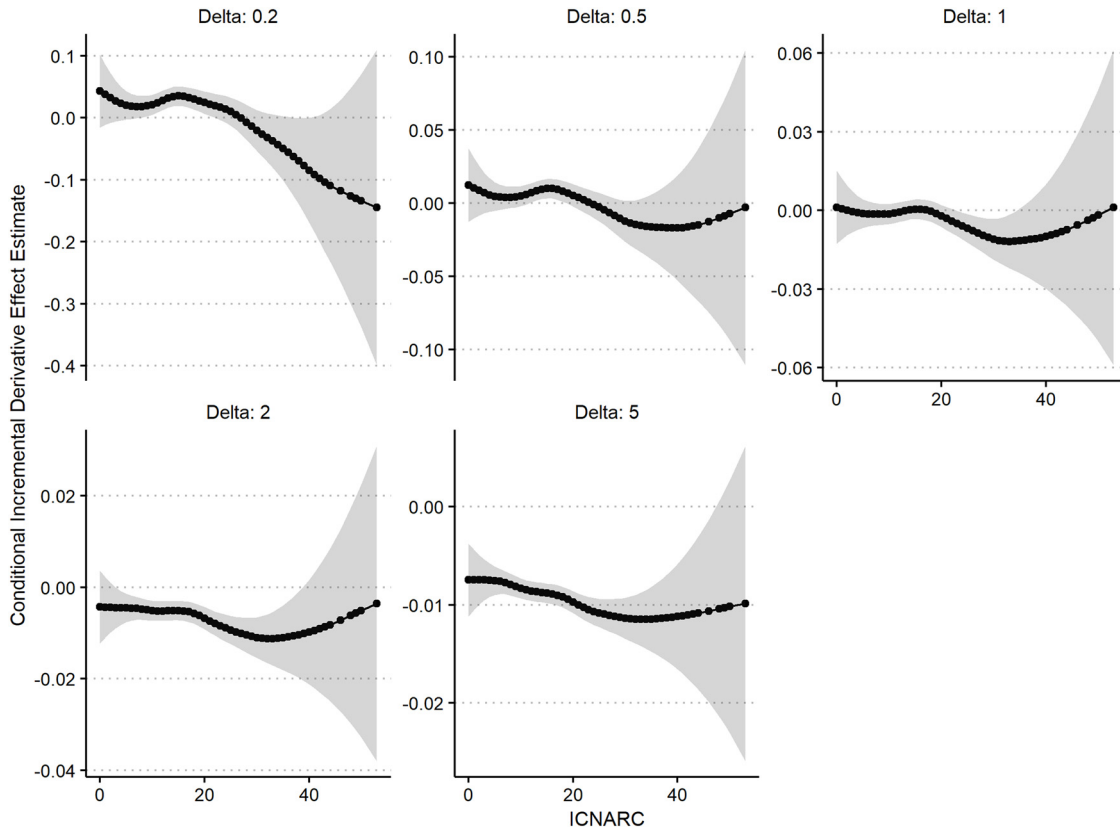


Figure 4: Predicted CIDE vs ICNARC score over δ .

estimates, since the mortality rate at $\delta = 5$ is higher than at $\delta = 0.2$. But, by examining the curve across the spectrum of interventions, one would also observe that mild interventions correspond to small changes in mortality rates. Therefore, our analysis validates previous research – in the sense that it estimates mortality to be lower when no one is admitted to the ICU, compared to everyone are admitted – but it also suggests a different practical implication, since one would conclude from our analysis that realistic interventions might only lead to small changes in mortality rates. This highlights how examining a spectrum of interventions can be more informative than examining a contrast like the ATE.

Meanwhile, Figure 4 shows the CIDE across ICNARC score for five δ values, and shows that the CIDE is generally very near to zero, and is only significantly different from zero at a few points across δ and ICNARC score. Figure 5 shows there is significant treatment effect heterogeneity across ICNARC score with 95% confidence intervals, but that the magnitude of the effect is very small, since the estimate for the V-CIDE is very close to zero for all δ values. Taken together, Figures 4 and 5 demonstrate that there is little effect heterogeneity across ICNARC scores, and increasing the odds of ICU admission affects subjects with different ICNARC scores similarly.

6 Discussion

In this work, we introduced three conditional effects based on incremental propensity score interventions – the conditional incremental effect (CIE), conditional incremental contrast effect (CICE), and conditional incremental derivative effect (CIDE). We proposed two estimators, the Projection-Learner and the I-DR-Learner, which can be used to estimate any of the three conditional effects. We showed that the Projection-Learner, a projection approach, achieves parametric efficiency under weak $n^{-1/4}$ conditions on the nuisance function estimators and that the I-DR-Learner, a nonparametric estimator, achieves oracle efficiency under similarly




Figure 5: V-CIDE vs δ .

weak conditions. We also proposed a fourth effect, the variance of the CIDE (V-CIDE), which is a one-dimensional summary of effect heterogeneity. For the V-CIDE, we proposed a new estimator also with double robust style properties, and outlined methods for inference and testing for treatment effect heterogeneity.

Finally, we illustrated our methods with a real data analysis of the effect of ICU admission on mortality conditional on a patient's risk score. This analysis demonstrated that estimating counterfactual mean outcomes across a spectrum of incremental interventions can be more informative than just estimating the ATE. We found evidence that the ATE is positive, suggesting that sending no one to the ICU is better than sending everyone to the ICU in terms of average mortality rates. However, by examining the spectrum of incremental interventions, we estimated that average mortality changes little with mild changes in ICU admission rates. Further, we found that there is indeed statistically significant treatment effect heterogeneity across patient risk scores, but the magnitude of heterogeneity is small. An additional limitation of this analysis is that it assumed there were no unmeasured confounders, which might be implausible. The sensitivity of the results to possible unmeasured confounders could be examined in future work. To our knowledge, sensitivity analyses for incremental propensity score interventions have not yet been developed.

There are other interesting avenues for future investigation. Here we proposed CIE estimators with the simplest data generating setup – one time point and binary treatment. Several natural extensions of this work to more complex frameworks are (i) time-varying data, (ii) incremental parameters that can depend on covariate data or past data, and (iii) multi-valued or continuous treatments with different stochastic interventions. Since positivity violations are almost guaranteed with time-varying data or multi-valued or continuous treatment, it would also be important to understand how nonparametric estimators behave and how projection approaches might be utilized to approximate ATE-style effects when positivity is violated.

Acknowledgements: The authors thank the Causal Inference Reading Group at Carnegie Mellon University and several anonymous reviewers for helpful discussion and comments, and Luke Keele for guidance on the (SPOT)light study data.

Funding information: The authors state no funding involved.

Conflict of interest: The authors state no conflict of interest.

References

- [1] Athey S, Imbens G. Recursive partitioning for heterogeneous causal effects. *Proc National Acad Sci*. 2016;113(27):7353–60.
- [2] Foster DJ, Syrgkanis V. Orthogonal statistical learning. *Ann Stat*. 2023;51(3):879–908.
- [3] Kennedy EH. Towards optimal doubly robust estimation of heterogeneous causal effects. 2020. arXiv: <http://arXiv.org/abs/arXiv:2004.14497>.
- [4] Künzel SR, Sekhon JS, Bickel PJ, Yu B. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proc Nat Acad Sci*. 2019;116(10):4156–65.
- [5] Nie X, Wager S. Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika*. 2021;108(2):299–319.
- [6] Semenova V, Chernozhukov V. Debiased machine learning of conditional average treatment effects and other causal functions. *Econom J*. 2021;24(2):264–89.
- [7] Shalit U, Johansson FD, Sontag D. Estimating individual treatment effect: generalization bounds and algorithms. In: *International Conference on Machine Learning*. PMLR; 2017. p. 3076–85.
- [8] Keele L, Harris S, Grieve R. Does transfer to intensive care units reduce mortality? A comparison of an instrumental variables design to risk adjustment. *Medical Care*. 2019;57(11):e73–9.
- [9] Díaz I, Hejazi NS. Causal mediation analysis for stochastic interventions. *J R Stat Soc Ser B Stat Methodol*. 2020;82(3):661–83.
- [10] Haneuse S, Rotnitzky A. Estimation of the effect of interventions that modify the received treatment. *Stat Med*. 2013;32(30):5260–77.
- [11] Kennedy EH. Nonparametric causal effects based on incremental propensity score interventions. *J Amer Stat Assoc*. 2019;114(526):645–56.
- [12] Moore KL, Neugebauer R, van der Laan MJ, Tager IB. Causal inference in epidemiological studies with strong confounding. *Stat Med*. 2012;31(13):1380–404.
- [13] Muñoz ID, van der Laan M. Population intervention causal effects based on stochastic interventions. *Biometrics*. 2012;68(2):541–9.
- [14] Young JG, Hernán MA, Robins JM. Identification, estimation and approximation of risk under interventions that depend on the natural value of treatment using observational data. *Epidemiol Methods*. 2014;3(1):1–9.
- [15] Zhou X, Opacic A. Marginal interventional effects. 2022. arXiv: <http://arXiv.org/abs/arXiv:2206.10717>.
- [16] Bonvini M, McClean A, Branson Z, Kennedy EH. Incremental causal effects: an introduction and review. In: *Handbook of Matching and Weighting Adjustments for Causal Inference*. New York, USA: Chapman and Hall/CRC; 2023. p. 349–72.
- [17] Wen L, Marcus JL, Young JG. Intervention treatment distributions that depend on the observed treatment process and model double robustness in causal survival analysis. *Stat Methods Med Res*. 2023;32(3):509–23.
- [18] Westreich D, Cole SR. Invited commentary: positivity in practice. *Amer J Epidemiol*. 2010;171(6):674–7.
- [19] Stensrud MJ, Laurendeau J, Sarvet AL. Optimal regimes for algorithm-assisted human decision-making. 2022. arXiv: <http://arXiv.org/abs/arXiv:2203.03020>.
- [20] R Core Team. A Language and Environment for Statistical Computing. 2023. R Foundation for Statistical Computing.
- [21] Kennedy EH. npcausal: Nonparametric Causal Inference Methods [Internet]. 2021. [cited 2023 Sept 20]. <https://github.com/ehkennedy/npcausal/>.
- [22] Kim K, Kennedy EH, Naimi AI. Incremental intervention effects in studies with dropout and many timepoints. *J Causal Infer*. 2021;9(1):302–44.
- [23] Sarvet AL, Wanis KN, Young JG, Hernandez-Alejandro R, Stensrud MJ. Longitudinal incremental propensity score interventions for limited resource settings. Wiley Online Library; 2023.
- [24] Rudolph JE, Kim K, Kennedy EH, Naimi AI. Estimation of the time-varying incremental effect of low-dose aspirin on incidence of pregnancy. *Epidemiology*. 2022;34(1):38–44.
- [25] Chakraborty B, Murphy SA. Dynamic treatment regimes. *Ann Rev Stat Appl*. 2014;1:447–64.
- [26] Murphy SA. Optimal dynamic treatment regimes. *J R Stat Soc Ser B Stat Methodol*. 2003;65(2):331–55.
- [27] Díaz I, Williams N, Hoffman KL, Schenck EJ. Nonparametric causal effects based on longitudinal modified treatment policies. *J Amer Stat Assoc*. 2023;118(542):846–57.

- [28] Taubman SL, Robins JM, Mittleman MA, Hernán MA. Intervening on risk factors for coronary heart disease: an application of the parametric g-formula. *Int J Epidemiol.* 2009;38(6):1599–611.
- [29] Robins JM, Hernán MA, Siebert U. Effects of multiple interventions. In: Ezzati M, Lopez AD, Rodgers AA, Murray CJ. *Comparative Quantification of Health Risks: Global and Regional Burden of Disease Attributable to Selected Major Risk Factors.* Geneva, Switzerland: World Health Organization; 2004. p. 2191–230.
- [30] Vansteelandt S, Bekaert M, Claeskens G. On model selection and model misspecification in causal inference. *Stat Meth Med Res.* 2012;21(1):7–30.
- [31] Bickel PJ, Klaassen CA, Ritov YA, Wellner JA. *Efficient and adaptive estimation for semiparametric models.* Baltimore: Johns Hopkins University Press; 1993.
- [32] Tsiatis AA. *Semiparametric theory and missing data.* New York: Springer; 2006.
- [33] van der Laan MJ, Robins JM. *Unified methods for censored longitudinal data and causality.* New York: Springer; 2003.
- [34] van der Vaart AW. *Asymptotic statistics.* Cambridge: Cambridge University Press; 2000.
- [35] van der Vaart AW. Semiparametric statistics. In: *Lectures on Probability Theory and Statistics.* Berlin: Springer; 2002. p. 331–457.
- [36] Mises RV. On the asymptotic distribution of differentiable statistical functions. *Ann Math Stat.* 1947;18(3):309–48.
- [37] Chernozhukov V, Chetverikov D, Demirer M, Dufo E, Hansen C, Newey W, et al. Double/debiased machine learning for treatment and structural parameters. *Econom J.* 2018;21(1):C1–C68.
- [38] Kennedy EH. Semiparametric doubly robust targeted double machine learning: a review. 2022. arXiv: <http://arXiv.org/abs/arXiv:2203.06469>.
- [39] Beran R. Minimum Hellinger distance estimates for parametric models. *Ann Stat.* 1977;5:445–63.
- [40] Berk R, Buja A, Brown L, George E, Kuchibhotla AK, Su W, et al. Assumption lean regression. *Amer Stat.* 2019;75(1):76–84.
- [41] Buja A, Brown L, Kuchibhotla AK, Berk R, George E, Zhao L. Models as approximations II. *Stat Sci.* 2019;34(4):545–65.
- [42] Huber PJ. The behavior of maximum likelihood estimates under nonstandard conditions. In: *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability.* 1967; (Vol. 1. Issue No. 1) pp. 221–33.
- [43] White H. Using least squares to approximate unknown regression functions. *Int Econom Rev.* 1980;21(1):149–70.
- [44] Cuellar M, Kennedy EH. A non-parametric projection-based estimator for the probability of causation, with application to water sanitation in Kenya. *J R Stat Soc Ser A Stat Soc.* 2020;183(4):1793–818.
- [45] Kennedy EH, Balakrishnan S, Wasserman LA. Semiparametric counterfactual density estimation. *Biometrika.* 2023;110:asad017.
- [46] Petersen M, Schwab J, Gruber S, Blaser N, Schomaker M, van der Laan MJ. Targeted maximum likelihood estimation for dynamic and static longitudinal marginal structural working models. *J Causal Infer.* 2014;2(2):147–85.
- [47] Neugebauer R, van der Laan M. Nonparametric causal effects based on marginal structural models. *J Stat Plan Inference.* 2007;137(2):419–34.
- [48] Vansteelandt S, Dukes O. Assumption-lean inference for generalised linear model parameters. *J R Stat Soc Ser B Stat Methodol.* 2022;84(3):657–85.
- [49] Hahn PR, Murray JS, Carvalho CM. Bayesian regression tree models for causal inference: Regularization, confounding, and heterogeneous effects (with discussion). *Bayesian Anal.* 2020;15(3):965–1056.
- [50] Zimmert M, Lechner M. Nonparametric estimation of causal heterogeneity under high-dimensional confounding. 2019. arXiv: <http://arXiv.org/abs/arXiv:1908.08779>.
- [51] Hines O, Dukes O, Diaz-Ordaz K, Vansteelandt S. Demystifying statistical learning based on efficient influence functions. *Amer Stat.* 2022;76(3):292–304.
- [52] Robins JM. Correcting for non-compliance in randomized trials using structural nested mean models. *Commun Stat-Theory Methods.* 1994;23(8):2379–412.
- [53] Robins JM, Mark SD, Newey WK. Estimating exposure effects by modelling the expectation of exposure conditional on confounders. *Biometrics.* 1992;48:479–95.
- [54] Robinson PM. Root-N-consistent semiparametric regression. *Econometrica: J Econometric Soc.* 1988;56:931–54.
- [55] Vansteelandt S, Joffe M. Structural nested models and g-estimation: the partially realized promise. *Stat Sci.* 2014;29(4):707–31.
- [56] Chen Q, Syrgkanis V, Austern M. Debiased machine learning without sample-splitting for stable estimators. *Adv Neural Inform Process Syst.* 2022;35:3096–109.
- [57] van der Vaart AW, Wellner JA. *Weak convergence and empirical processes.* New York, USA: Springer; 1996.
- [58] Robins J, Li L, Tchetgen Tchetgen E, van der Vaart A. Higher order influence functions and minimax estimation of nonlinear functionals. In: *Probability and statistics: essays in honor of David A. Freedman.* Beechwood, Ohio, USA: Institute of Mathematical Statistics. 2008. Vol. 2. p. 335–422.
- [59] Zheng W, van der Laan MJ. Asymptotic theory for cross-validated targeted maximum likelihood estimation. U.C. Berkeley Division of Biostatistics Working Paper Series. 2010. Paper 273.
- [60] Birgé L, Massart P. Estimation of integral functionals of a density. *Ann Stat.* 1995;23(1):11–29.
- [61] Farrell MH. Robust inference on average treatment effects with possibly more covariates than observations. *J Econometrics.* 2015;189(1):1–23.
- [62] Tibshirani R. Regression shrinkage and selection via the lasso. *J R Stat Soc Ser B Stat Methodol.* 1996;58(1):267–88.
- [63] Tsybakov AB. *Introduction to nonparametric estimation.* New York: Springer; 2009.
- [64] Wasserman L. *All of nonparametric statistics.* New York, USA: Springer Science & Business Media; 2006.

- [65] Crump RK, Hotz VJ, Imbens GW, Mitnik OA. Nonparametric tests for treatment effect heterogeneity. *Rev Econom Stat.* 2008;90(3):389–405.
- [66] Ding P, Feller A, Miratrix L. Randomization inference for treatment effect variation. *J R Stat Soc Ser B Stat Methodol.* 2016;78(3):655–71.
- [67] Ding P, Feller A, Miratrix L. Decomposing treatment effect variation. *J Amer Stat Assoc.* 2019;114(525):304–17.
- [68] Luedtke A, Carone M, van der Laan MJ. An omnibus non-parametric test of equality in distribution for unknown functions. *J R Stat Soc Ser B Stat Methodol.* 2019;81(1):75–99.
- [69] Williamson BD, Gilbert PB, Simon NR, Carone M. A general framework for inference on algorithm-agnostic variable importance. *J Amer Stat Assoc.* 2023;118(543):1645–58.
- [70] Harris S, Singer M, Sanderson C, Grieve R, Harrison D, Rowan K. Impact on mortality of prompt admission to critical care for deteriorating ward patients: an instrumental variable analysis using critical care bed strain. *Intensive Care Med.* 2018;44:606–15.
- [71] Gabler NB, Ratcliffe SJ, Wagner J, Asch DA, Rubenfeld GD, Angus DC, et al. Mortality among patients admitted to strained intensive care units. *Am J Respiratory Critical Care Med.* 2013;188(7):800–6.
- [72] Renaud B, Santin A, Coma E, Camus N, Van Pelt D, Hayon J, et al. Association between timing of intensive care unit admission and outcomes for emergency department patients with community-acquired pneumonia. *Critical Care Medicine.* 2009;37(11):2867–74.
- [73] Harrison DA, Parry GJ, Carpenter JR, Short A, Rowan K. A new risk prediction model for critical care: the Intensive Care National Audit & Research Centre (ICNARC) model. *Crit Care Med.* 2007;35(4):1091–8.
- [74] Williams B, Alberti G, Ball C, Ball D, Binks R, Durham L. National Early Warning Score (NEWS). Standardising the assessment of acute-illness severity in the NHS. London, UK: Royal College of Physicians; 2012.
- [75] Vincent JL, Moreno R, Takala J, Willatts S, De Mendonça A, Bruining H, et al. The SOFA (Sepsis-related Organ Failure Assessment) score to describe organ dysfunction/failure: On behalf of the Working Group on Sepsis-Related Problems of the European Society of Intensive Care Medicine. 1996.
- [76] Kang H, Jiang Y, Zhao Q, Small D. ivmodel: Statistical Inference and Sensitivity Analysis for Instrumental Variables Model [Internet]. 2023 [cited 2023 Oct 25]. <https://cran.r-project.org/web/packages/ivmodel/>.
- [77] Wright MN, Ziegler A. Ranger: a fast implementation of random forests for high dimensional data in C++ and R. *J Stat Software.* 2017;77(1):1–17.
- [78] Wood S. mgcv: Mixed GAM computation vehicle with GCV/AIC/REML smoothness estimation. 2012.
- [79] Kennedy EH, Balakrishnan S, G'Sell M. Sharp instruments for classifying compliers and generalizing causal effects. *Ann Stat.* 2020;48(4):2008–30.

Appendix

A Stability condition for Theorem 2

In this section, we state the stability condition invoked in Section 3.2 and Theorem 2. This stability condition is described in detail in Section 3 of Kennedy [3], and can be viewed as a form of stochastic equicontinuity for nonparametric regression.

Definition 1. (*Stability*) Suppose $D_1 = \{Z_i\}_{i=1}^n$ and $D_2 = \{Z_i\}_{i=n+1}^{2n}$ are independent training and estimation samples of n observations where $X \subset Z$ are covariates (e.g., $Z_i = (X_i, A_i, Y_i)$). Let

- (1) $\hat{f}(z) = \hat{f}(z; D_1)$ be an estimate of some function of the data, $f(z)$, using the training data D_1 ,
- (2) $\hat{b}(x) = \hat{b}(x; D_1) \equiv \mathbb{E}\{\hat{f}(X) - f(X) | D_1, X = x\}$, the conditional bias of the estimator \hat{f} ,
- (3) $\hat{\mathbb{E}}_n(Y|X = x)$ denote a generic regression estimator that regresses outcomes (Y_{n+1}, \dots, Y_{2n}) on predictors (X_{n+1}, \dots, X_{2n}) in the estimation sample D_2 .

Then, the regression estimator $\hat{\mathbb{E}}_n$ is defined as *stable* (with respect to a distance metric d) if

$$\frac{\hat{\mathbb{E}}_n\{\hat{f}(Z)|X = x\} - \hat{\mathbb{E}}_n\{f(Z)|X = x\} - \hat{\mathbb{E}}_n\{\hat{b}(X)|X = x\}}{\sqrt{\mathbb{E}([\hat{\mathbb{E}}_n\{f(Z)|X = x\} - \mathbb{E}\{f(Z)|X = x\}]^2)}} \xrightarrow{p} 0,$$

wherever $d(\hat{f}, f) \xrightarrow{p} 0$.

Definition 1 says that the difference between the regression estimate with estimated outcomes ($\hat{\mathbb{E}}_n\{\hat{f}(Z)|X = x\}$) and the oracle regression ($\hat{\mathbb{E}}_n\{f(Z)|X = x\}$) converges to zero appropriately fast. This definition can be viewed as a generalization of the classic stochastic equicontinuity condition

$$\frac{(\mathbb{P}_n - \mathbb{P})(\hat{f} - f)}{1/\sqrt{n}} \xrightarrow{p} 0,$$

where $\mathbb{P}_n(\hat{f} - f)$ is replaced by $\hat{\mathbb{E}}_n\{\hat{f}(Z)|X = x\} - \hat{\mathbb{E}}_n\{f(Z)|X = x\}$, $\mathbb{P}(\hat{f} - f)$ is replaced by the conditional bias term, $\hat{\mathbb{E}}_n\{\hat{b}(X)|X = x\}$, and the denominator $1/\sqrt{n}$ is replaced by the pointwise root mean squared error of the oracle estimator, $\sqrt{\mathbb{E}([\hat{\mathbb{E}}_n\{f(Z)|X = x\} - \mathbb{E}\{f(Z)|X = x\}]^2)}$. This stability condition is satisfied by linear smoothers, as is demonstrated in Kennedy [3] Theorem 1, and may be satisfied by more classes of estimators.

B Estimator for the variance of the conditional incremental derivative effect (V-CIDE) when the conditioning covariate is a strict subset of all covariates

In this section, we briefly outline an estimator for the V-CIDE when the conditioning covariates V are a strict subset of all covariates X (i.e., $V \subset X$) and discuss the convergence properties of the associated estimator. As a reminder, the variance of the CIDE is identified by

$$\mathbb{V}\{\tau_{\text{cide}}(V; \delta)\} = \mathbb{V}\left[\mathbb{E}\left[\frac{\pi(X)\{1 - \pi(X)\}}{\{\delta\pi(X) + 1 - \pi(X)\}^2}\{\mu(1, X) - \mu(0, X)\} \middle| V\right]\right].$$

By the definition of the variance and iterated expectation,

$$\mathbb{V}\{\tau_{\text{cide}}(V; \delta)\} = \mathbb{E}\{\tau_{\text{cide}}(V; \delta)^2\} - \mathbb{E}\{\tau_{\text{cide}}(V; \delta)\}^2 = \mathbb{E}\{\tau_{\text{cide}}(V; \delta)^2\} - \mathbb{E}\{\tau_{\text{cide}}(X; \delta)\}^2.$$

The squared expectation term, $\mathbb{E}\{\tau_{\text{cide}}(X; \delta)\}^2$, is the same as what appears when $V = X$, and therefore, can be estimated as in equation (21) in Section 4. The expected square term is new, and we derive the efficient influence function in the following result:

Lemma 3. *Under Assumptions 1 and 2, the un-centered efficient influence function for $\mathbb{E}\{\tau_{\text{cide}}(V; \delta)\}^2$ is*

$$\tau_{\text{cide}}(V; \delta)^2 + 2\tau_{\text{cide}}(V; \delta)\{\omega\varphi + \phi(\mu_1 - \mu_0) + \omega(\mu_1 - \mu_0) - \tau_{\text{cide}}(V; \delta)^2\},$$

where $\mu_a = \mu(a, X)$, and $\omega = \omega(X; \delta)$, $\varphi = \varphi(Z)$, and $\phi = \phi(Z; \delta)$ as defined in equations (10), (18), and (19).

This result suggests the following estimator:

$$\mathbb{P}_n[\hat{\tau}_{\text{cide}}(V; \delta)^2 + 2\hat{\tau}_{\text{cide}}(V; \delta)\{\hat{\omega}\hat{\varphi} + \hat{\phi}(\hat{\mu}_1 - \hat{\mu}_0) + \hat{\omega}(\hat{\mu}_1 - \hat{\mu}_0) - \hat{\tau}_{\text{cide}}(V; \delta)^2\}]. \quad (29)$$

And, combined with the results in Sections 3.2 and 4, suggests the following estimator for the V-CIDE:

Algorithm 4. (V-CIDE Estimator when $V \subset X$) Assume as inputs (D_1, D_2) , which denote two independent samples of n observations of $Z_i = (X_i, A_i, Y_i)$, then

- (1) On the training data D_1 , estimate the nuisance functions $\hat{\mu}(0, X)$, $\hat{\mu}(1, X)$, and $\hat{\pi}(X)$.
- (2) On the estimation data D_2 , estimate the un-centered influence function values $\hat{\xi}(Z; \delta)$ using the models $\hat{\mu}$ and $\hat{\pi}$ from step 1, where $\hat{\xi}(Z; \delta)$ is defined in (17) if the conditional effect of interest is τ_{cide} , and analogously for τ_{cie} and τ_{cice} in equations (45) and (46).
- (3) In the estimation sample D_2 , regress $\hat{\xi}(Z; \delta)$ on the conditioning covariates V to obtain the estimate

$$\hat{\tau}_{\text{cide}}(v; \delta) = \hat{\mathbb{E}}_n\{\hat{\xi}(Z; \delta) | V = v\}.$$

- (4) On the estimation data, estimate $\mathbb{V}\{\tau_{\text{cide}}(V; \delta)\}$ per equations (29) and (21), plugging in the estimates for $\hat{\mu}$, $\hat{\pi}$, $\hat{\tau}_{\text{cide}}$ from above.

This estimator satisfies a similar double robustness condition to the estimator outlined in Section 4, but includes a dependence on $\hat{\tau}_{\text{cide}}(V; \delta)$.

Theorem 4. *Let $\hat{\psi}_n$ denote the estimator from Algorithm 4. Under Assumptions 1 and 2, Assumption (a) from Theorem 1, and Assumption (a) from Theorem 3, if*

$$\|\hat{\pi} - \pi\|(\|\hat{\mu} - \mu\| + \|\hat{\pi} - \pi\|) + \|\hat{\tau}_{\text{cide}} - \tau_{\text{cide}}\|^2,$$

then

$$\sqrt{n}[\hat{\psi}_n - \mathbb{V}\{\tau_{\text{cide}}(V; \delta)\}] \rightsquigarrow N(0, \sigma^2),$$

where

$$\begin{aligned} \sigma^2 = & \mathbb{V}[\tau_{\text{cide}}^2 + 2\tau_{\text{cide}}\{\omega\varphi + \phi(\mu_1 - \mu_0) + \omega(\mu_1 - \mu_0) - \tau_{\text{cide}}^2\} \\ & - \mathbb{E}\{\omega\varphi + \phi(\mu_1 - \mu_0) + \omega(\mu_1 - \mu_0)\}\{\omega\varphi + \phi(\mu_1 - \mu_0) + \omega(\mu_1 - \mu_0)\}], \end{aligned} \quad (30)$$

where $\tau_{\text{cide}} \equiv \tau_{\text{cide}}(V; \delta)$, and $\mu_a = \mu(a, X)$, $\omega = \omega(X; \delta)$, $\varphi = \varphi(Z)$, and $\phi = \phi(Z; \delta)$ as defined in equations (10), (18), and (19).

The theorem shows that the estimator for the V-CIDE satisfies a version of double robustness under relatively weak conditions. The result shows that our estimator attains $n^{-1/2}$ convergence to the V-CIDE under model-agnostic $n^{-1/4}$ convergence rates for the nuisance function estimators and the I-DR-Learner. This is a different result from Theorem 2, since it is required that the CIDE is estimated at $n^{-1/4}$ rates, but it is no longer required that $\hat{\mu}$ is estimated at $n^{-1/4}$ rates. As discussed in the body of the article, $n^{-1/4}$ rates are achievable with nonparametric estimators under suitable smoothness or sparsity.

Like the estimator from Algorithm 3, the estimator in Algorithm 4 converges to a degenerate distribution when the V-CIDE equals zero. As discussed at the end of Section 4, we can construct a valid test for any treatment effect heterogeneity by overestimating the variance of the estimator $\hat{\psi}_n$. When $V \subset X$, we can construct the following asymptotically valid $1 - \alpha$ test

$$\begin{cases} \text{Reject } H_0 : \mathbb{V}\{\tau_{\text{cide}}(V; \delta)\} = 0 & \text{if } \hat{\psi}_n - \Phi^{-1}(1 - \alpha) \sqrt{\frac{\hat{\sigma}_1^2 + \hat{\sigma}_2^2}{n}} > 0, \\ \text{Fail to reject } H_0 : \mathbb{V}\{\tau_{\text{cide}}(V; \delta)\} = 0 & \text{otherwise,} \end{cases} \quad (31)$$

where

$$\begin{aligned} \hat{\sigma}_1^2 &= \hat{\mathbb{V}}_n[\hat{\tau}_{\text{cide}}^2 + 2\hat{\tau}_{\text{cide}}\{\hat{\omega}\hat{\phi} + \hat{\phi}(\hat{\mu}_1 - \hat{\mu}_0) + \hat{\omega}(\hat{\mu}_1 - \hat{\mu}_0) - \hat{\tau}_{\text{cide}}^2\}], \text{ and} \\ \hat{\sigma}_2^2 &= \mathbb{V}_n[\mathbb{P}_n\{\hat{\omega}\hat{\phi} + \hat{\phi}(\hat{\mu}_1 - \hat{\mu}_0) + \hat{\omega}(\hat{\mu}_1 - \hat{\mu}_0)\} \cdot \{\hat{\omega}\hat{\phi} + \hat{\phi}(\hat{\mu}_1 - \hat{\mu}_0) + \hat{\omega}(\hat{\mu}_1 - \hat{\mu}_0)\}]. \end{aligned}$$

C Simulations for the Projection-Learner and I-DR-learner

In this section, we study the performance of the Projection-Learner and I-DR-Learner for estimating the conditional incremental contrast effect (CICE) with $\delta_u = 5$ and $\delta_l = 0.2$; R code for these simulations is provided at <https://github.com/alecmcclean/NPCIE>. As a reminder, the CICE is defined as

$$\tau_{\text{cide}}(v; \delta_u = 5, \delta_l = 0.2) \equiv \mathbb{E}(Y^{Q_5} - Y^{Q_{0.2}} | V = v),$$

which corresponds to the difference between the counterfactual mean outcomes when the odds of treatment are multiplied by 5 minus the counterfactual mean outcome when the odds of treatment are divided by 5. For all the analyses, we simulate 1,000 times datasets of sizes $n \in \{1,000, 10,000\}$. In each dataset, we generate $\{(X_i, A_i, Y_i)\}$ from $i = 1$ to $n = 1,000$, where $X \in \mathbb{R}$, $A \in \{0, 1\}$, and $Y \in \mathbb{R}$. For each dataset, we specify a quadratic CICE,

$$\tau_{\text{cide}}(\delta_u = 5, \delta_l = 0.2; X) = 1 + 0.5X - 0.2X^2,$$

and the CATE is defined implicitly as

$$\tau_{\text{cate}}(X) = \frac{\tau_{\text{cide}}(X; \delta_u = 5, \delta_l = 0.2)}{q\{\pi(X); \delta_u = 5\} - q\{\pi(X); \delta_l = 0.2\}},$$

which follows from the identification result in Proposition 1. Each simulated dataset $\{(X_i, A_i, Y_i)\}_{i=1}^{n=1,000}$ is then constructed in the following manner:

$$X \sim \text{Unif}(-4, 4), \quad (32)$$

$$\pi(X) = \text{expit}\left(\frac{X}{2}\right), \quad (33)$$

$$\mu(0, X) = \mathbb{1}(X < -3) \cdot 2 + \mathbb{1}(X > -2) \cdot 2.55 + \mathbb{1}(X > 0) \cdot -2, \quad (34)$$

$$+ \mathbb{1}(X > 2) \cdot 4 + \mathbb{1}(X > 3) \cdot -1 \quad (35)$$

$$\mu(1, X) = \mu(0, X) + \tau_{\text{cate}}(X), \quad (36)$$

$$A \sim \text{Bernoulli}(\pi), \quad (37)$$

$$Y \sim A \cdot \mu(1, X) + (1 - A) \cdot \mu(0, X) + N(0, 1). \quad (38)$$

The data generating process is also illustrated in Figure A1. The covariate data X is one dimensional and uniform over $[-4, 4]$. The propensity score, shown in the top panel of Figure A1, follows a logistic model, which remains within reasonable bounds on the support of X , since $\pi(-4) \approx 0.12$ and $\pi(4) \approx 0.88$. The outcome of




Figure A1: Data generating process.

regressions are complicated discontinuous functions and are shown in the middle panel of Figure A1. The treatment A and outcome Y are defined implicitly from the propensity score and regression functions. The second panel also shows the CATE, which is a smooth function. This is because the CICE and the propensity scores are smooth functions. The CICE is shown in the bottom panel of Figure A1.

We simulated estimates for the propensity scores and regression functions by adding noise, parameterized by α , to the true nuisance functions:

$$\hat{\pi}(X) \sim \text{expit}[\text{logit}\{\pi(X) + N(n^{-\alpha_\pi}, n^{-2\alpha_\pi})\}], \quad (39)$$

$$\hat{\mu}(a, X) \sim \mu(a, X) + N\left[\left\{\max_x \mu(a, x) - \min_x \mu(a, x)\right\} \cdot n^{-\alpha_\mu}, \left\{\max_x \mu(a, x) - \min_x \mu(a, x)\right\}^2 \cdot n^{-2\alpha_\mu}\right]. \quad (40)$$

The α parameter allows us to control how well the nuisance functions are “estimated.” For example, when $\alpha = 0.1$, this corresponds to estimating a nuisance function with error converging at $n^{-1/10}$. We scale the error for the regression functions μ by the range of regression function values; this is purely a computing trick so that the error for neither nuisance function dominates the other, and this does not affect the convergence rates of the estimators.

First, we compare the I-DR-Learner to the oracle estimator (“Oracle I-DR-Learner”) and a baseline learner (“Baseline CICE”) in terms of integrated mean squared error (MSE). The oracle estimator constructs the true influence function value from $\mu(a, X)$ and $\pi(X)$ and regresses them against X . Both the oracle estimator and the I-DR-Learner use the `smooth.spline` function in R for the second stage regression. The baseline estimator is a plug-in estimator which calculates

$$\hat{\tau}_{\text{cice}}(X; \delta_u = 5, \delta_l = 0.2) = \{\hat{\mu}(1, X) - \hat{\mu}(0, X)\} \cdot [q\{\hat{\pi}(X); \delta_u = 5\} - q\{\hat{\pi}(X); \delta_l = 0.2\}].$$

This is motivated by causal identification, such as the result in Proposition 1, and does not make use of the efficient influence function for the relevant average effect. The baseline estimator for the CATE was previously examined in the literature, and has been referred to as the “T-Learner” [4].




Figure A2: Comparing CICE estimators.

The results of these simulations are summarized in Figure A2. Each different panel corresponds to a different convergence rate α_μ for estimating $\hat{\mu}$ and sample size, and α_μ increases from left to right while $n = 1,000$ is in the top row and $n = 10,000$ is in the bottom row. The x-axis shows the convergence rate α_π for estimating $\hat{\pi}$, and the y-axis shows the integrated MSE for each estimator. Finally, each estimator is denoted by a different color, and the points and whiskers show the sample mean value and 95% confidence interval for the MSE over 1,000 simulations.

Figure A2 illustrates the phenomenon anticipated by Theorem 2. The oracle estimator performs the best, which we would expect since it has access to the true nuisance functions. The I-DR-Learner performs the next best, and its error approaches that of the oracle estimator as $\min(2\alpha_\pi, \alpha_\pi + \alpha_\mu)$ increases, while the baseline learners, whose error is additive in terms of the nuisance function estimators' errors, fare the worst.

C.1 Coverage of the Projection-Learner

In this subsection, we outline results for the Projection-Learner. Specifically, we show that the projection-learner achieves approximately correct coverage for the true coefficients in the model. The true model is

$$\tau_{\text{cice}}(X; \delta_u = 5, \delta_l = 0.2) = 1 + 0.5X - 0.2X^2,$$

and the working model is

$$g(\beta; X) = \beta^* + \beta_1 X + \beta_2 X^2.$$

Since the working model is well-specified, the Projection-Learner estimates the true coefficients. Figure A3 shows the coverage of 95% confidence intervals constructed for each coefficient using the sandwich variance as in Corollary 2. When the nuisance function estimators have large error, such that $\alpha_\pi + \alpha_\mu < 0.5$ or $2\alpha_\pi < 0.5$ then the confidence intervals have poor coverage. As the errors of the nuisance function estimators decrease, the coverage of the confidence intervals for each coefficient approach 95%.




Figure A3: Coverage of Projection-Learner confidence intervals for coefficient estimates.

D Proofs for results in Section 2

Proposition 2. Let Q_δ denote the incremental intervention defined in equation (4). Under Assumptions 1 and 2 and if $\tau_{\text{cie}}(x; \delta)$ is continuous within x and δ and its partial derivative with respect to δ is absolutely integrable, then the CIDE is identified by

$$\tau_{\text{cie}}(v; \delta) = \mathbb{E}[\omega(X; \delta)\{\mu(1, X) - \mu(0, X)\}|V = v], \quad (9)$$

where $\mu(a, x) = \mathbb{E}(Y|A = a, X = x)$ and

$$\omega(X; \delta) = \frac{\pi(X)\{1 - \pi(X)\}}{\{\delta\pi(X) + 1 - \pi(X)\}^2}. \quad (10)$$

Proof. We have

$$\begin{aligned} \tau_{\text{cie}}(v; \delta) &= \frac{\partial}{\partial t} \mathbb{E}\{Y^{Q_t}|V = v\}_{t=\delta} \\ &= \frac{\partial}{\partial t} \mathbb{E}\left[\left[\frac{t\pi(X)}{t\pi(X) + 1 - \pi(X)}\right]\mu(1, X) + \left[\frac{1 - \pi(X)}{t\pi(X) + 1 - \pi(X)}\right]\mu(0, X)|V = v\right]_{t=\delta} \\ &= \mathbb{E}\left[\left[\frac{\pi(X)\{1 - \pi(X)\}}{\{\delta\pi(X) + 1 - \pi(X)\}^2}\right] \cdot \{\mu(1, X) - \mu(0, X)\}|V = v\right], \end{aligned}$$

where the first line follows by definition, the second by Assumptions 1 and 2, and the third and final line by exchanging expectation and derivative, taking the derivative with respect to t , rearranging, and setting $t = \delta$. \square

E Proofs for results in Section 3

Lemma 1. Under Assumptions 1 and 2, the un-centered efficient influence function for the average incremental derivative effect, $\mathbb{E}\{\tau_{\text{cide}}(V; \delta)\}$, is

$$\begin{aligned} \xi(Z; \delta) = \omega(X; \delta) & \left[\frac{A}{\pi(X)} \{Y - \mu(1, X)\} - \frac{1-A}{1-\pi(X)} \{Y - \mu(0, X)\} \right] \\ & + \left[\frac{1}{\{\delta\pi(X) + 1 - \pi(X)\}^2} - \frac{2\delta\pi(X)}{\{\delta\pi(X) + 1 - \pi(X)\}^3} \right] \{A - \pi(X)\} \{\mu(1, X) - \mu(0, X)\} \\ & + \omega(X; \delta) \{\mu(1, X) - \mu(0, X)\}, \end{aligned} \quad (16)$$

where $\omega(X; \delta)$ is defined in equation (10).

Proof. We prove the result by showing that $\mathbb{E}\{\tau_{\text{cide}}(v; \delta)\}$ satisfies a von Mises expansion where $\xi(Z; \delta)$ is the influence function and there is a second order remainder term. At the end of the proof, we will relate this to smooth parametric submodels and submodel scores to prove that $\xi(Z; \delta)$ is the efficient influence function.

Let $\xi(\mathbb{P}) \equiv \xi(Z; \delta)$ and $\psi(\mathbb{P}) = \mathbb{E}\{\tau_{\text{cide}}(v; \delta)\}$. Then, by a von Mises expansion,

$$\psi(\bar{\mathbb{P}}) = \psi(\mathbb{P}) + \int_Z \xi(\bar{\mathbb{P}}) d(\bar{\mathbb{P}} - \mathbb{P}) + R_2(\bar{\mathbb{P}}, \mathbb{P}), \quad (41)$$

where \mathbb{P} and $\bar{\mathbb{P}}$ are two different distributions at which the functional ψ is evaluated. Rearranging, we can see that

$$R_2(\bar{\mathbb{P}}, \mathbb{P}) = \psi(\bar{\mathbb{P}}) - \psi(\mathbb{P}) - \int_Z \xi(\bar{\mathbb{P}}) d(\bar{\mathbb{P}} - \mathbb{P}) = \int_Z \xi(\bar{\mathbb{P}}) - \xi(\mathbb{P}) d\mathbb{P} \equiv \mathbb{E}_{\mathbb{P}}\{\xi(\bar{\mathbb{P}}) - \xi(\mathbb{P})\}, \quad (42)$$

where $\mathbb{E}_{\mathbb{P}}$ denotes expectation under the distribution \mathbb{P} . By the definition of ξ ,

$$\begin{aligned} R_2(\bar{\mathbb{P}}, \mathbb{P}) = \mathbb{E}_{\mathbb{P}} & \left[\left[\frac{\bar{\pi}(X)\{1 - \bar{\pi}(X)\}}{\{\delta\bar{\pi}(X) + 1 - \bar{\pi}(X)\}^2} \right] \cdot \left[\frac{A}{\bar{\pi}(X)} \{Y - \bar{\mu}(1, X)\} - \frac{1-A}{1-\bar{\pi}(X)} \{Y - \bar{\mu}(0, X)\} \right] \right] \\ & + \mathbb{E}_{\mathbb{P}} \left[\left[\frac{1}{\{\delta\bar{\pi}(X) + 1 - \bar{\pi}(X)\}^2} - \frac{2\delta\bar{\pi}(X)}{\{\delta\bar{\pi}(X) + 1 - \bar{\pi}(X)\}^3} \right] \cdot \{A - \bar{\pi}(X)\} \{\bar{\mu}(1, X) - \bar{\mu}(0, X)\} \right] \\ & + \mathbb{E}_{\mathbb{P}} \left[\left[\frac{\bar{\pi}(X)\{1 - \bar{\pi}(X)\}}{\{\delta\bar{\pi}(X) + 1 - \bar{\pi}(X)\}^2} \right] \cdot \{\bar{\mu}(1, X) - \bar{\mu}(0, X)\} \right] \\ & - \mathbb{E}_{\mathbb{P}} \left[\left[\frac{\pi(X)\{1 - \pi(X)\}}{\{\delta\pi(X) + 1 - \pi(X)\}^2} \right] \cdot \left[\frac{A}{\pi(X)} \{Y - \mu(1, X)\} - \frac{1-A}{1-\pi(X)} \{Y - \mu(0, X)\} \right] \right] \\ & - \mathbb{E}_{\mathbb{P}} \left[\left[\frac{1}{\{\delta\pi(X) + 1 - \pi(X)\}^2} - \frac{2\delta\pi(X)}{\{\delta\pi(X) + 1 - \pi(X)\}^3} \right] \cdot \{A - \pi(X)\} \{\mu(1, X) - \mu(0, X)\} \right] \\ & - \mathbb{E}_{\mathbb{P}} \left[\left[\frac{\pi(X)\{1 - \pi(X)\}}{\{\delta\pi(X) + 1 - \pi(X)\}^2} \right] \cdot \{\mu(1, X) - \mu(0, X)\} \right]. \end{aligned}$$

By iterated expectations and rearranging,

$$\begin{aligned} R_2(\bar{\mathbb{P}}, \mathbb{P}) = \mathbb{E}_{\mathbb{P}} & \left[\bar{\omega}(X; \delta) \left[\left[\frac{\pi(X) - \bar{\pi}(X)}{\bar{\pi}(X)} \right] \{\mu(1, X) - \bar{\mu}(1, X)\} + \left[\frac{\pi(X) - \bar{\pi}(X)}{1 - \bar{\pi}(X)} \right] \{\mu(0, X) - \bar{\mu}(0, X)\} \right] \right] \\ & + \mathbb{E}_{\mathbb{P}} \{ \{\bar{\mu}(1, X) - \bar{\mu}(0, X)\} [\mathbb{E}\{\bar{\phi}(Z; \delta)|X\} + \bar{\omega}(X; \delta) - \omega(X; \delta)] \} \\ & + \mathbb{E}_{\mathbb{P}} \{ \{\bar{\omega}(X; \delta) - \omega(X; \delta)\} [\mu(1, X) - \bar{\mu}(1, X) - \{\mu(0, X) - \bar{\mu}(0, X)\}] \}, \end{aligned} \quad (43)$$

where

$$\omega(X; \delta) = \frac{\pi(X)\{1 - \pi(X)\}}{\{\delta\pi(X) + 1 - \pi(X)\}^2}, \text{ and}$$

$$\phi(Z; \delta) = \left[\frac{1}{\{\delta\pi(X) + 1 - \pi(X)\}^2} - \frac{2\delta\pi(X)}{\{\delta\pi(X) + 1 - \pi(X)\}^3} \right] \cdot \{A - \pi(X)\}.$$

It can be shown that $R_2(\bar{\mathbb{P}}, \mathbb{P})$ is second order, since

$$R_2(\bar{\mathbb{P}}, \mathbb{P}) = \mathbb{E}_{\mathbb{P}}[g_0(X)\{\pi(X) - \hat{\pi}(X)\}\{\mu(0, X) - \hat{\mu}(0, X)\} + g_1(X)\{\pi(X) - \hat{\pi}(X)\}\{\mu(1, X) - \hat{\mu}(1, X)\} \\ + h(X)\{\pi(X) - \hat{\pi}(X)\}^2].$$

We show this in the postscript to this proof. Here we provide some intuition. It is clear that the first line of (43) is already a second order. The second line of (43) is second order because ϕ is the un-centered efficient influence function of ω , and so the second multiplicand on the second line is the error term for estimating ω , which we would expect to be second order. The third line of (43) is more intuitively second order because it is the product of the errors of two plug-ins, and that product can be expressed as a product of errors.

Now, we relate $\xi(Z; \delta)$ back to scores of smooth parametric submodels. Recall from semiparametric efficiency theory that the nonparametric efficiency bound for a functional is given by the supremum of Cramer–Rao lower bounds for that functional across smooth parametric submodels [31,35]. The efficient influence function is the unique mean-zero function ξ that is a valid submodel score satisfying pathwise differentiability, i.e.,

$$\left. \frac{d}{d\varepsilon} \psi(\mathbb{P}_\varepsilon) \right|_{\varepsilon=0} = \int_Z \xi(\mathbb{P}) \left(\frac{d}{d\varepsilon} \log d\mathbb{P}_\varepsilon \right) \Big|_{\varepsilon=0} d\mathbb{P} \quad (44)$$

for \mathbb{P}_ε any smooth parametric submodel. To see that $\xi(Z; \delta)$ is the efficient influence function for $\mathbb{E}\{\tau_{\text{cide}}(v; \delta)\}$, observe that the von Mises expansion in (41) implies

$$\begin{aligned} \frac{d}{d\varepsilon} \psi(\mathbb{P}_\varepsilon) &= \frac{d}{d\varepsilon} \left(\psi(\mathbb{P}) - \int_Z \xi(\mathbb{P}) d(\mathbb{P} - \mathbb{P}_\varepsilon) - R_2(\mathbb{P}, \mathbb{P}_\varepsilon) \right) \\ &= \frac{d}{d\varepsilon} \int_Z \xi(\mathbb{P}) d(\mathbb{P}_\varepsilon - \mathbb{P}) - \frac{d}{d\varepsilon} R_2(\mathbb{P}, \mathbb{P}_\varepsilon) \\ &= \int_Z \xi(\mathbb{P}) \frac{d}{d\varepsilon} d\mathbb{P}_\varepsilon - \frac{d}{d\varepsilon} R_2(\mathbb{P}, \mathbb{P}_\varepsilon) \\ &= \int_Z \xi(\mathbb{P}) \left(\frac{d}{d\varepsilon} \log d\mathbb{P}_\varepsilon \right) d\mathbb{P}_\varepsilon - \frac{d}{d\varepsilon} R_2(\mathbb{P}, \mathbb{P}_\varepsilon) \end{aligned}$$

with R_2 defined in (43), and where the second line follows because $\psi(\mathbb{P})$ does not depend on ε , the third because $\int \xi(\mathbb{P}) d\mathbb{P} = 0$, and the fourth and final line since $\frac{d}{d\varepsilon} \log d\mathbb{P}_\varepsilon = \frac{1}{d\mathbb{P}_\varepsilon} \frac{d}{d\varepsilon} d\mathbb{P}_\varepsilon$. Evaluating this expression at $\varepsilon = 0$, we have

$$\int_Z \xi(\mathbb{P}) \left(\frac{d}{d\varepsilon} \log d\mathbb{P}_\varepsilon \right) d\mathbb{P}_\varepsilon - \frac{d}{d\varepsilon} R_2(\mathbb{P}, \mathbb{P}_\varepsilon) \Big|_{\varepsilon=0} = \int_Z \xi(\mathbb{P}) \left(\frac{d}{d\varepsilon} \log d\mathbb{P}_\varepsilon \right) \Big|_{\varepsilon=0} d\mathbb{P} = 0,$$

since

$$\left. \frac{d}{d\varepsilon} R_2(\mathbb{P}, \mathbb{P}_\varepsilon) \right|_{\varepsilon=0} = 0,$$

which shows that ξ satisfies the property in (44). The last equation involving R_2 follows because R_2 consists of only second-order products of errors between \mathbb{P}_ε and \mathbb{P} . Therefore, the derivative is composed of a sum of

terms, each of which is a product of either a derivative term that may not equal zero and an error term involving the differences of components of \mathbb{P} and \mathbb{P}_ε , which will be zero when $\varepsilon = 0$ since $\mathbb{P} = \mathbb{P}_\varepsilon$.

Since the model is nonparametric, the tangent space is the entire Hilbert space of mean-zero finite-variance functions, and so there is only one influence function satisfying (44) and it is the efficient one [32]. Therefore, $\xi(Z; \delta)$ is the efficient influence function for $\mathbb{E}\{\tau_{\text{cide}}(V; \delta)\}$. \square

Below, we show the algebra for why $R_2(\bar{\mathbb{P}}, \mathbb{P})$ is second order as stated above. Starting with the third line in (43), and omitting arguments, we have

$$\begin{aligned}\bar{\omega} - \omega &= \frac{\pi(1 - \bar{\pi})}{(\delta\bar{\pi} + 1 - \bar{\pi})^2} - \frac{\pi(1 - \pi)}{(\delta\pi + 1 - \pi)^2} \\ &= \frac{\pi(1 - \bar{\pi})(\delta\pi + 1 - \pi)^2 - \pi(1 - \pi)(\delta\bar{\pi} + 1 - \bar{\pi})^2}{(\delta\pi + 1 - \pi)^2(\delta\bar{\pi} + 1 - \bar{\pi})^2} \\ &= (\pi - \bar{\pi}) \left[\frac{(\delta + 1)(\delta - 1)\pi\bar{\pi} + \pi + \bar{\pi} - 1}{(\delta\pi + 1 - \pi)^2(\delta\bar{\pi} + 1 - \bar{\pi})^2} \right].\end{aligned}$$

Therefore, letting $\mu_a = \mu(a, X)$, we have

$$\begin{aligned}\{\bar{\omega}(X; \delta) - \omega(X; \delta)\}[\mu(1, X) - \bar{\mu}(1, X) - \{\mu(0, X) - \bar{\mu}(0, X)\}] \\ &= (\pi - \bar{\pi}) \left[\frac{(\delta + 1)(\delta - 1)\pi\bar{\pi} + \pi + \bar{\pi} - 1}{(\delta\pi + 1 - \pi)^2(\delta\bar{\pi} + 1 - \bar{\pi})^2} \right] \cdot \{\mu_1 - \bar{\mu}_1 - (\mu_0 - \bar{\mu}_0)\} \\ &= \left[\frac{(\delta + 1)(\delta - 1)\pi\bar{\pi} + \pi + \bar{\pi} - 1}{(\delta\pi + 1 - \pi)^2(\delta\bar{\pi} + 1 - \bar{\pi})^2} \right] \{(\pi - \bar{\pi})(\mu_1 - \bar{\mu}_1)\} + \left[\frac{(\delta + 1)(1 - \delta)\pi\bar{\pi} + \pi + \bar{\pi} - 1}{(\delta\pi + 1 - \pi)^2(\delta\bar{\pi} + 1 - \bar{\pi})^2} \right] \{(\pi - \bar{\pi})(\mu_0 - \bar{\mu}_0)\}.\end{aligned}$$

For the second line in (43), we have

$$\begin{aligned}\mathbb{E}(\bar{\phi}|X) + \bar{\omega} - \omega &= \left[\frac{1}{(\delta\bar{\pi} + 1 - \bar{\pi})^2} - \frac{2\delta\bar{\pi}}{(\delta\bar{\pi} + 1 - \bar{\pi})^3} \right] \cdot (\pi - \bar{\pi}) + (\pi - \bar{\pi}) \left[\frac{(\delta + 1)(\delta - 1)\pi\bar{\pi} + \pi + \bar{\pi} - 1}{(\delta\pi + 1 - \pi)^2(\delta\bar{\pi} + 1 - \bar{\pi})^2} \right] \\ &= (\pi - \bar{\pi}) \left[\frac{1}{(\delta\bar{\pi} + 1 - \bar{\pi})^2} - \frac{2\delta\bar{\pi}}{(\delta\bar{\pi} + 1 - \bar{\pi})^3} + \frac{(\delta + 1)(\delta - 1)\pi\bar{\pi} + \pi + \bar{\pi} - 1}{(\delta\pi + 1 - \pi)^2(\delta\bar{\pi} + 1 - \bar{\pi})^2} \right] \\ &= (\pi - \bar{\pi})^2 \left[\frac{(1 - \delta - \delta^2 - \delta^3)\pi\bar{\pi} + 1 - 2\delta + (1 - \delta)\{\pi + (1 - \delta)\pi\}}{(\delta\bar{\pi} + 1 - \bar{\pi})^3(\delta\pi + 1 - \pi)^2} \right].\end{aligned}$$

Finally, revisiting the first line of (43),

$$\begin{aligned}\bar{\omega}(X; \delta) &\left[\left[\frac{\pi(X) - \bar{\pi}(X)}{\pi(X)} \right] \{\mu(1, X) - \bar{\mu}(1, X)\} + \left[\frac{\pi(X) - \bar{\pi}(X)}{1 - \pi(X)} \right] \{\mu(0, X) - \bar{\mu}(0, X)\} \right] \\ &= (\pi - \bar{\pi})(\mu_1 - \bar{\mu}) \left[\frac{\bar{\omega}}{\bar{\pi}} \right] + (\pi - \bar{\pi})(\mu_0 - \bar{\mu}_0) \left[\frac{\bar{\omega}}{1 - \bar{\pi}} \right].\end{aligned}$$

E.1 Efficient influence functions for the average incremental effect and the average incremental contrast effect

By Corollary 2 in the study by Kennedy [11], the efficient influence function for the average incremental effect is

$$\begin{aligned}\xi_{ie}(Z; \delta) &= \frac{\delta\pi(X)\mu(1, X) + \{1 - \pi(X)\}\mu(0, X)}{\delta\pi(X) + 1 - \pi(X)} + \left[\frac{\delta A + 1 - A}{\delta\pi(X) + 1 - \pi(X)} \right] \{Y - \mu(A, X)\} \\ &\quad + \{\mu(1, X) - \mu(0, X)\} \frac{\delta\{A - \pi(X)\}}{\{\delta\pi(X) + 1 - \pi(X)\}^2},\end{aligned}\tag{45}$$

and the efficient influence function for the average incremental contrast effect is

$$\xi_{ice}(Z; \delta_u, \delta_l) = \xi_{ie}(Z; \delta_u) - \xi_{ie}(Z; \delta_l). \quad (46)$$

Corollary 1. Let $\xi(Z; \delta)$ denote the true influence function values of the relevant average effect, where $\xi(Z; \delta)$ is defined in (16) if the estimand is a projection of $\tau_{cde}(v; \delta)$, and is defined analogously, as shown in equations (45) and (46) in the appendix, if the estimand is a projection of $\tau_{cde}(v; \delta)$ or $\tau_{cde}(v; \delta_u, \delta_l)$. Under Assumptions 1 and 2, the un-centered efficient influence function for $m(\beta)$ with unknown propensity scores and a uniform weight function constructed over ℓ_2 distance is

$$\phi(Z; \delta, \beta) = \frac{\partial g(V; \beta)}{\partial \beta} \{ \xi(Z; \delta) - g(V; \beta) \},$$

where $g(v; \beta)$ is the working model.

Proof. Let $\phi(\mathbb{P}) \equiv \phi(Z; \delta; \beta)$. The functional $m(\beta)$ satisfies the following von Mises expansion

$$m(\beta, \bar{\mathbb{P}}) - m(\beta, \mathbb{P}) = \int_Z \phi(\bar{\mathbb{P}}) d(\bar{\mathbb{P}} - \mathbb{P}) + R_2(\bar{\mathbb{P}}, \mathbb{P}), \quad (47)$$

where

$$R_2(\bar{\mathbb{P}}, \mathbb{P}) = m(\beta, \bar{\mathbb{P}}) - m(\beta, \mathbb{P}) - \int_Z \phi(\bar{\mathbb{P}}) d(\bar{\mathbb{P}} - \mathbb{P}) = \int \phi(\bar{\mathbb{P}}) - \phi(\mathbb{P}) d\mathbb{P}.$$

Following essentially the same logic as in Lemma 1, by iterated expectations and rearranging,

$$\begin{aligned} R_2(\bar{\mathbb{P}}, \mathbb{P}) = \mathbb{E}_{\mathbb{P}} \left[\frac{\partial g(V; \beta)}{\partial \beta} \left(\bar{\omega}(X; \delta) \left[\frac{\pi(X) - \bar{\pi}(X)}{\pi(X)} \{ \mu(1, X) - \mu(0, X) \} + \frac{\pi(X) - \bar{\pi}(X)}{1 - \pi(X)} \{ \mu(0, X) - \bar{\mu}(0, X) \} \right] \right. \right. \\ \left. \left. + \{ \bar{\mu}(1, X) - \bar{\mu}(0, X) \} [\mathbb{E}(\bar{\phi}(Z; \delta) | X) + \bar{\omega}(X; \delta) - \omega(X; \delta)] \right. \right. \\ \left. \left. + \{ \bar{\omega}(X; \delta) - \omega(X; \delta) \} [\mu(1, X) - \bar{\mu}(1, X) - \{ \mu(0, X) - \bar{\mu}(0, X) \}] \right) \right]. \end{aligned}$$

This second order term can be expressed as a product of errors, as is shown in post script to the proof for Lemma 1. Therefore, since our model is nonparametric, $\phi(Z; \delta, \beta)$ is the efficient influence function for $m(\beta)$. \square

Theorem 1. Let $\varphi(Z; \beta, \mu, \pi) \equiv \phi(Z; \delta, \beta) - m(\beta)$ denote the centered efficient influence function from Corollary 1. With Assumptions 1 and 2, also assume

(a) $\mathbb{P}(|\hat{\mu}(1, X) - \hat{\mu}(0, X)| \leq C) = 1$ and $\mathbb{P}(|\mu^*(1, X) - \mu^*(0, X)| \leq C) = 1$ for some $C < \infty$.

(b) $\mathbb{P} \left\{ \left| \frac{\partial g(\beta; v)}{\partial \beta} \right| \leq C \right\} = 1$ for all v .

(c) The function class $\varphi(Z; \beta, \mu, \pi)$ is Donsker in β for any fixed μ, π .

(d) The estimators are consistent in the sense that $\hat{\beta} - \beta^* = o_{\mathbb{P}}(1)$ and $\|\varphi(Z; \hat{\beta}; \hat{\mu}, \hat{\pi}) - \varphi(Z; \beta^*; \mu^*, \pi^*)\| = o_{\mathbb{P}}(1)$.

(e) The map $\beta \mapsto \mathbb{P}\{\varphi(Z; \beta, \mu, \pi)\}$ is differentiable at β^* uniformly in (μ, π) , with nonsingular derivative matrix

$$\frac{\partial}{\partial \beta} \mathbb{P}\{\varphi(Z; \beta, \mu, \pi)\}|_{\beta=\beta^*} = M(\beta^*, \mu, \pi), \text{ where } M(\beta^*, \hat{\mu}, \hat{\pi}) \xrightarrow{p} M(\beta^*, \mu^*, \pi^*).$$

Then,

$$\hat{\beta} - \beta^* = -M(\beta^*, \mu^*, \pi^*)^{-1} (\mathbb{P}_n - \mathbb{P}) \{ \varphi(Z; \beta^*, \mu^*, \pi^*) \} + O_{\mathbb{P}} \left(R_n + o_{\mathbb{P}} \left(\frac{1}{\sqrt{n}} \right) \right),$$

where

$$R_n = (\|\hat{\mu} - \mu^*\| + \|\hat{\pi} - \pi^*\|) \|\hat{\pi} - \pi^*\|.$$

Proof. This proof follows closely both Lemma 3 from the study by Kennedy et al. [45] and Theorem 5.31 of the study by van der Vaart [34]. Since $(\hat{\beta}, \hat{\mu}, \hat{\pi})$ is an approximate solution to the empirical moment condition, $\mathbb{P}_n\{\varphi(Z; \hat{\beta}, \hat{\mu}, \hat{\pi})\} = o_p(1/\sqrt{n})$. Since (β^*, μ^*, π^*) is an exact solution to the population moment condition, $\mathbb{P}\{\varphi(Z; \beta^*, \mu^*, \pi^*)\} = 0$. Combining these two facts,

$$o_p(1/\sqrt{n}) = \mathbb{P}_n\{\varphi(Z; \hat{\beta}, \hat{\mu}, \hat{\pi})\} - \mathbb{P}\{\varphi(Z; \beta^*, \mu^*, \pi^*)\}.$$

By adding and subtracting on the right hand side of the equation above, omitting Z , and letting $\eta^* = (\mu^*, \pi^*)$ and $\hat{\eta} = (\hat{\mu}, \hat{\pi})$

$$\begin{aligned} o_p(1/\sqrt{n}) &= (\mathbb{P}_n - \mathbb{P})\{\varphi(\beta^*, \eta^*)\} + (\mathbb{P}_n - \mathbb{P})\{\varphi(\hat{\beta}, \hat{\eta}) - \varphi(\beta^*, \hat{\eta})\} \\ &\quad + (\mathbb{P}_n - \mathbb{P})\{\varphi(\beta^*, \hat{\eta}) - \varphi(\beta^*, \eta^*)\} + \mathbb{P}\{\varphi(\hat{\beta}, \hat{\eta}) - \varphi(\beta^*, \hat{\eta})\} + \mathbb{P}\{\varphi(\beta^*, \hat{\eta}) - \varphi(\beta^*, \eta^*)\}. \end{aligned}$$

The first term appears directly in the statement in the theorem, so we will not manipulate it. It is a sample average of a fixed function, and so by the central limit theorem it will be asymptotically Gaussian. The second and third terms are empirical process terms. The fourth term can be linearized in $(\hat{\beta} - \beta^*)$, and we will use this to rearrange and solve for the statement in the theorem. The fifth term captures the nuisance estimation error, and appears implicitly in the statement in the theorem if we define $R_n = \mathbb{P}\{\varphi(Z; \beta^*, \hat{\mu}, \hat{\pi}) - \varphi(Z; \beta^*, \mu^*, \pi^*)\}$.

First, we will tackle the second and third terms. Under the Donsker and consistency conditions in Assumptions (c) and (d), the second term is $o_p(1/\sqrt{n})$ by Lemma 19.24 of the study by van der Vaart [34]. Further, under the consistency of $\varphi(\hat{\beta}, \hat{\eta})$ in Assumption (d) and by sample splitting, the third term is $o_p(1/\sqrt{n})$ by Lemma 2 of the study by Kennedy et al. [79].

The fourth term, by the differentiability of the map $\beta \mapsto \mathbb{P}\{\varphi(\beta, \eta)\}$ in Assumption (e), can be expressed as

$$\begin{aligned} \mathbb{P}\{\varphi(\hat{\beta}, \hat{\eta}) - \varphi(\beta^*, \hat{\eta})\} &= M(\beta^*, \hat{\eta})(\hat{\beta} - \beta^*) + o_p(\|\hat{\beta} - \beta^*\|) \\ &= M(\beta^*, \eta^*)(\hat{\beta} - \beta^*) + o_p(\|\hat{\beta} - \beta^*\|), \end{aligned}$$

where the first line is a first-order Taylor expansion about β^* and the second line follows by the consistency of $M(\beta^*, \hat{\eta})$ in Assumption (e).

Bringing everything together,

$$\begin{aligned} o_p(1/\sqrt{n}) &= (\mathbb{P}_n - \mathbb{P})\{\varphi(\beta^*, \eta^*)\} + (\mathbb{P}_n - \mathbb{P})\{\varphi(\hat{\beta}, \hat{\eta}) - \varphi(\beta^*, \eta^*)\} \\ &\quad + (\mathbb{P}_n - \mathbb{P})\{\varphi(\beta^*, \hat{\eta}) - \varphi(\beta^*, \eta^*)\} \\ &\quad + \mathbb{P}\{\varphi(\hat{\beta}, \hat{\eta}) - \varphi(\beta^*, \hat{\eta})\} + \mathbb{P}\{\varphi(\beta^*, \hat{\eta}) - \varphi(\beta^*, \eta^*)\} \\ &= (\mathbb{P}_n - \mathbb{P})\{\varphi(\beta^*, \eta^*)\} + o_p(1/\sqrt{n}) + o_p(1/\sqrt{n}) \\ &\quad + M(\beta^*, \eta^*)(\hat{\beta} - \beta^*) + o_p(\|\hat{\beta} - \beta^*\|) + R_n. \end{aligned}$$

Re-arranging, we have that

$$o_p(1/\sqrt{n}) = (\mathbb{P}_n - \mathbb{P})\{\varphi(Z; \beta^*, \eta^*)\} + M(\beta^*, \eta^*)(\hat{\beta} - \beta^*) + o_p(\|\hat{\beta} - \beta^*\|) + R_n.$$

We can re-arrange and by the non-singularity of the derivative matrix M in Assumption (e) we can pre-multiply both sides by $M(\beta^*, \eta^*)^{-1}$ and see that

$$\hat{\beta} - \beta^* = -M(\beta^*, \eta^*)^{-1}(\mathbb{P}_n - \mathbb{P})\{\varphi(\beta^*, \eta^*)\} + O_p(R_n) + o_p(\|\hat{\beta} - \beta^*\|) + o_p(1/\sqrt{n}). \quad (48)$$

To address the $o_p(\|\hat{\beta} - \beta^*\|)$ term, notice the following: by Assumption (d), $\hat{\beta} - \beta^* = o_p(1)$, and by the Central Limit Theorem, $\varphi(Z; \beta^*, \eta^*) = O_p(1/\sqrt{n})$. Therefore, rearranging equation (48) above to put the $(\hat{\beta} - \beta^*)$ terms on the same side, we have

$$\hat{\beta} - \beta^* + o_p(\|\hat{\beta} - \beta^*\|) = O_p(1/\sqrt{n}) + O_p(R_n) + o_p(1/\sqrt{n}),$$

so that

$$\|\hat{\beta} - \beta^*\| \{1 + o_P(1)\} = O_P(1/\sqrt{n} + R_n)$$

and therefore

$$\|\hat{\beta} - \beta^*\| = O_P(1/\sqrt{n} + R_n)$$

and

$$o_P(\|\hat{\beta} - \beta^*\|) = o_P(O_P(1/\sqrt{n} + R_n)) = o_P(1/\sqrt{n}) + o_P(R_n).$$

Finally, plugging these results back into equation (48), we have

$$\begin{aligned} \hat{\beta} - \beta^* &= -M(\beta^*, \eta^*)^{-1}(\mathbb{P}_n - \mathbb{P})\{\varphi(\beta^*, \eta^*)\} + O_P(R_n) + o_P(\|\hat{\beta} - \beta^*\|) + o_P(1/\sqrt{n}) \\ &= -M(\beta^*, \eta^*)^{-1}(\mathbb{P}_n - \mathbb{P})\{\varphi(\beta^*, \eta^*)\} + O_P(R_n) + o_P(1/\sqrt{n}) + o_P(R_n) + o_P(1/\sqrt{n}) \\ &= -M(\beta^*, \eta^*)^{-1}(\mathbb{P}_n - \mathbb{P})\{\varphi(\beta^*, \eta^*)\} + O_P(R_n) + o_P(1/\sqrt{n}). \end{aligned}$$

And, adding back in all arguments, we can conclude

$$\hat{\beta} - \beta^* = -M(\beta^*, \mu^*, \pi^*)^{-1}(\mathbb{P}_n - \mathbb{P})\{\varphi(\beta^*, \mu^*, \pi^*; Z)\} + O_P(R_n + o_P(1/\sqrt{n})),$$

which provides the first statement of the theorem.

For the second statement, recall that $R_n = \mathbb{P}\{\varphi(\beta^*, \hat{\mu}, \hat{\pi}; Z) - \varphi(\beta^*, \mu^*, \pi^*; Z)\}$.

$$\begin{aligned} \mathbb{P}\{\varphi(Z; \beta^*, \hat{\mu}, \hat{\pi}) - \varphi(Z; \beta^*, \mu^*, \pi^*)\} &= \int_{\mathcal{Z}} 2 \frac{\partial g(v; \beta^*)}{\partial \beta} \{g(v; \beta^*) - \hat{\xi}(z; \delta)\} d\mathbb{P}(z) \\ &\quad - 2 \int_{\mathcal{Z}} \frac{\partial g(v; \beta^*)}{\partial \beta} \{g(v; \beta^*) - \xi_0(z; \delta)\} d\mathbb{P}(z) \\ &= \int_{\mathcal{Z}} 2 \frac{\partial g(v; \beta^*)}{\partial \beta} \{\xi_0(z; \delta) - \hat{\xi}(z; \delta)\} d\mathbb{P}(z). \end{aligned}$$

The final result follows by Cauchy–Schwartz and a boundedness condition, which depends on the target estimand. If the estimand is a projection of the CIDE, then the result follows by equivalent logic to the proof of Lemma 1 and the first part of Assumption (a), which says that the estimated CATE is bounded. If, instead, the estimand is a projection of the CIE or the CICE, the result follows by equivalent logic to the proofs of Lemmas 5 and 6 in the appendix of the study by Kennedy [11] and the second part of Assumption (a), which says that the true CATE is bounded. \square

Theorem 2. Let τ_{i-dr} stand in for τ_{cide} , τ_{cie} , or τ_{cice} , and let $\xi(Z; \delta)$ denote the true influence function values of the relevant average effect. Furthermore, let $\tilde{\tau}_{i-dr}(v; \delta) = \hat{\mathbb{E}}_n\{\xi(Z; \delta)|V = v\}$ denote an oracle estimator that regresses $\xi(Z; \delta)$ on V , and let $\hat{\tau}_{i-dr}(v; \delta)$ denote the I-DR-Learner from Algorithm 2. With Assumptions 1 and 2, and Assumption (a) from Theorem 1, also assume that the second stage regression is stable according to Definition 1 in the study by Kennedy [3]. Then,

$$\hat{\tau}_{i-dr}(v; \delta) - \tau_{i-dr}(v; \delta) = \tilde{\tau}_{i-dr}(v; \delta) - \tau_{i-dr}(v; \delta) + \hat{\mathbb{E}}_n\{\hat{b}(X)|V = v\} + o_P(R^*(v; \delta))$$

for

$$\hat{b}(x) \leq (|\hat{\mu}(0, x) - \mu(0, x)| + |\hat{\mu}(1, x) - \mu(1, x)| + |\hat{\pi}(x) - \pi(x)|) \cdot |\hat{\pi}(x) - \pi(x)|$$

and

$$R^*(v; \delta) = \sqrt{\mathbb{E}[\{\tilde{\tau}_{i-dr}(v; \delta) - \tau_{i-dr}(v; \delta)\}^2]}.$$

Proof. This follows from Proposition 1 of the study by Kennedy [3], the definition of $\hat{b}(x)$, and by iterated expectation. \square

The bounded functions in $\hat{b}(x)$ depend on the relevant estimand. In general,

$$\hat{b}(x) = g_0(x)\{\pi(x) - \hat{\pi}(x)\}\{\mu(0, x) - \hat{\mu}(0, x)\} + g_1(x)\{\pi(x) - \hat{\pi}(x)\}\{\mu(1, x) - \hat{\mu}(1, x)\} + h(x)\{\pi(x) - \hat{\pi}(x)\}^2,$$

where $g_0(x), g_1(x), h(x)$ are all bounded functions. Omitting arguments and letting $\mu_a = \mu(a, X)$, when $\tau_{i-dr} = \tau_{cie}$, then

$$\begin{aligned} g_0(x) &\equiv g_0(x; \delta) = \frac{1}{\delta\hat{\pi} + 1 - \hat{\pi}} \left(\frac{1 - \hat{\pi}}{1 - \pi} - \frac{\delta}{\delta\pi + 1 - \pi} \right), \\ g_1(x) &\equiv g_1(x; \delta) = \frac{1}{\delta\hat{\pi} + 1 - \hat{\pi}} \left(\frac{\delta}{\delta\pi + 1 - \pi} - \frac{\delta\hat{\pi}}{\pi} \right), \\ h(x) &\equiv h(x; \delta) = (\mu_1 - \mu_0) \left[\frac{\delta(1 - \delta)}{(\delta\pi + 1 - \pi)(\delta\hat{\pi} + 1 - \hat{\pi})^2} \right]. \end{aligned}$$

When $\tau_{i-dr} = \tau_{cice}$, then

$$\begin{aligned} g_0(x) &\equiv g_0(x; \delta_u, \delta_l) = g_0(x; \delta_u) - g_0(x; \delta_l), \\ g_1(x) &\equiv g_1(x; \delta_u, \delta_l) = g_1(x; \delta_u) - g_1(x; \delta_l), \\ h(x) &\equiv h(x; \delta_u, \delta_l) = h(x; \delta_u) - h(x; \delta_l), \end{aligned}$$

where $g_0(x; \delta), g_1(x; \delta), h(x; \delta)$ are defined above for the CIE. And, when $\tau_{i-dr} = \tau_{cide}$, then

$$\begin{aligned} g_0(x) &= \frac{\hat{\pi}}{(\delta\hat{\pi} + 1 - \hat{\pi})^2} + \left[\frac{(\delta + 1)(1 - \delta)\pi\bar{\pi} + \pi + \bar{\pi} - 1}{(\delta\pi + 1 - \pi)^2(\delta\bar{\pi} + 1 - \bar{\pi})^2} \right], \\ g_1(x) &= \frac{1 - \hat{\pi}}{(\delta\hat{\pi} + 1 - \hat{\pi})^2} + \left[\frac{(\delta + 1)(1 - \delta)\pi\bar{\pi} + \pi + \bar{\pi} - 1}{(\delta\pi + 1 - \pi)^2(\delta\bar{\pi} + 1 - \bar{\pi})^2} \right], \\ h(x) &= (\hat{\mu}_1 - \hat{\mu}_0) \cdot \left[\frac{(1 - \delta - \delta^2 - \delta^3)\pi\hat{\pi} + 1 - 2\delta + (1 - \delta)\{\hat{\pi} + (1 - \delta)\pi\}}{(\delta\hat{\pi} + 1 - \hat{\pi})^3(\delta\pi + 1 - \pi)^2} \right]. \end{aligned}$$

F Proofs for results in Section 4

Lemma 2. Under Assumptions 1 and 2, the un-centered efficient influence function for $E\{\tau_{cide}(X; \delta)^2\}$ is

$$2\omega(X; \delta)\{\mu(1, X) - \mu(0, X)\}\{\omega(X; \delta)\varphi(Z) + \phi(Z; \delta)\{\mu(1, X) - \mu(0, X)\}\} + [\omega(X; \delta)\{\mu(1, X) - \mu(0, X)\}]^2,$$

where $\omega(X; \delta)$ is defined in equation (10) and

$$\varphi(Z) = \frac{A}{\pi(X)}\{Y - \mu(1, X)\} - \frac{1 - A}{1 - \pi(X)}\{Y - \mu(0, X)\}, \quad (18)$$

$$\phi(Z; \delta) = \left[\frac{1}{\{\delta\pi(X) + 1 - \pi(X)\}^2} - \frac{2\delta\pi(X)}{\{\delta\pi(X) + 1 - \pi(X)\}^3} \right] \cdot \{A - \pi(X)\}. \quad (19)$$

Proof. As in Lemma 1, we prove this result by showing that the estimand admits a von Mises expansion where the second-order term is a product of errors.

Omitting arguments, let $\xi(\mathbb{P}) = 2\omega\tau(\omega\phi + \phi\tau) + (\omega\tau)^2$. Then,

$$\begin{aligned}
 R_2(\bar{\mathbb{P}}, \mathbb{P}) &= \int_Z \xi(\bar{\mathbb{P}}) - \xi(\mathbb{P}) d\mathbb{P} \\
 &= \mathbb{E}\{2\bar{\omega}\bar{\tau}(\bar{\omega}\bar{\phi} + \bar{\phi}\bar{\tau}) + (\bar{\omega}\bar{\tau})^2\} - \mathbb{E}\{2\omega\tau(\omega\phi + \phi\tau) + (\omega\tau)^2\} \\
 &= \mathbb{E}\{2\bar{\omega}\bar{\tau}(\bar{\omega}\bar{\phi} + \bar{\phi}\bar{\tau}) + (\bar{\omega}\bar{\tau})^2 - (\omega\tau)^2\} \\
 &= \mathbb{E}\{2\bar{\omega}\bar{\tau}(\bar{\omega}\bar{K} + \bar{\omega}(\tau - \bar{\tau}) + \bar{\phi}\bar{\tau}) + (\bar{\omega}\bar{\tau})^2 - (\omega\tau)^2\} \\
 &= \mathbb{E}[2\bar{\omega}\bar{\tau}\{\bar{\omega}\bar{K} + (\bar{\omega} - \omega)(\tau - \bar{\tau}) + (\bar{\phi} - \omega)\bar{\tau} + \omega\tau\} + (\bar{\omega}\bar{\tau})^2 - (\omega\tau)^2] \\
 &= \mathbb{E}[2\bar{\omega}\bar{\tau}\{\bar{\omega}\bar{K} + (\bar{\omega} - \omega)(\tau - \bar{\tau}) + (\bar{\phi} + \bar{\omega} - \omega)\bar{\tau} - \bar{\omega}\bar{\tau} + \omega\tau\} + (\bar{\omega}\bar{\tau})^2 - (\omega\tau)^2] \\
 &= \mathbb{E}(2\bar{\omega}\bar{\tau}[\bar{\omega}\bar{K} + (\bar{\omega} - \omega)(\tau - \bar{\tau}) + \bar{\tau}\{\mathbb{E}(\bar{\phi}|X) + \bar{\omega} - \omega\}] - 2(\bar{\omega}\bar{\tau})^2 + 2\bar{\omega}\bar{\tau}\omega\tau + (\bar{\omega}\bar{\tau})^2 - (\omega\tau)^2) \\
 &= \mathbb{E}(2\bar{\omega}\bar{\tau}[\bar{\omega}\bar{K} + (\bar{\omega} - \omega)(\tau - \bar{\tau}) + \bar{\tau}\{\mathbb{E}(\bar{\phi}|X) + \bar{\omega} - \omega\}] - (\bar{\omega}\bar{\tau} - \omega\tau)^2) \\
 &= \mathbb{E}(2\bar{\omega}\bar{\tau}[\bar{\omega}\bar{K} + (\bar{\omega} - \omega)(\tau - \bar{\tau}) + \bar{\tau}\{\mathbb{E}(\bar{\phi}|X) + \bar{\omega} - \omega\}] - \{\bar{\omega}(\bar{\tau} - \tau) + \tau(\bar{\omega} - \omega)\}^2) \\
 &\leq (\|\bar{\mu} - \mu\| + \|\bar{\pi} - \pi\|)^2,
 \end{aligned}$$

where

$$K = \sum_a (2a - 1) \left(\frac{\mathbb{P}_a - \bar{\mathbb{P}}_a}{\mathbb{P}_a} \right) (\mu_a - \bar{\mu}_a),$$

and $\mathbb{P}_a = \mathbb{P}(A = a|X)$. The last line follows by similar logic to the post script to Lemma 1. This shows that the second order term is a product of errors, and so $\xi(\mathbb{P})$ is the un-centered efficient influence function for $\mathbb{E}\{\tau_{\text{cide}}(X; \delta)^2\}$. \square

Theorem 3. Let $\hat{\psi}_n$ denote the estimator from Algorithm 3. With Assumptions 1 and 2, and Assumption (a) from Theorem 1, also assume that

(a) $\mathbb{P}\{\omega(\mu_1 - \mu_0) + \omega\phi + \phi\tau \leq C\}$ and $\mathbb{P}\{\hat{\omega}(\hat{\mu}_1 - \hat{\mu}_0) + \hat{\omega}\hat{\phi} + \hat{\phi}\hat{\tau} \leq C\} = 1$ for some $C < \infty$.

If

$$(\|\hat{\mu} - \mu\| + \|\hat{\pi} - \pi\|)^2 = o_p\left(\frac{1}{\sqrt{n}}\right),$$

then

$$\sqrt{n}[\hat{\psi}_n - \mathbb{V}\{\tau_{\text{cide}}(X; \delta)\}] \rightsquigarrow N(0, \sigma^2),$$

where

$$\begin{aligned}
 \sigma^2 &= \mathbb{V}[2\omega(\mu_1 - \mu_0)\{\omega\phi + \phi(\mu_1 - \mu_0)\} + \{\omega(\mu_1 - \mu_0)\}^2 - \mathbb{E}\{\omega\phi + \phi(\mu_1 - \mu_0) + \omega(\mu_1 - \mu_0)\} \\
 &\quad \cdot \{\omega\phi + \phi(\mu_1 - \mu_0) + \omega(\mu_1 - \mu_0)\}],
 \end{aligned} \tag{22}$$

$\mu_a = \mu(a, X)$, and $\omega = \omega(X; \delta)$, $\phi = \phi(Z)$, and $\phi = \phi(Z; \delta)$ as defined in equations (10), (18), and (19).

Proof. First, note that, by construction

$$\begin{aligned}
 \hat{\psi}_n &= \mathbb{P}_n\{2\hat{\omega}\hat{\tau}(\hat{\omega}\hat{\phi} + \hat{\phi}\hat{\tau}) + (\hat{\omega}\hat{\tau})^2 - (\hat{\omega}\hat{\tau} + \hat{\omega}\hat{\phi} + \hat{\phi}\hat{\tau})\mathbb{P}_n(\hat{\omega}\hat{\tau} + \hat{\omega}\hat{\phi} + \hat{\phi}\hat{\tau})\} \\
 &\equiv \mathbb{P}_n(\hat{\xi}_1) - \mathbb{P}_n(\hat{\xi}_2)^2,
 \end{aligned}$$

where the second line follows by defining

$$\begin{aligned}\hat{\xi}_1 &= 2\hat{\omega}\hat{\tau}(\hat{\omega}\hat{\phi} + \hat{\phi}\hat{\tau}) + (\hat{\omega}\hat{\tau})^2, \quad \text{and} \\ \hat{\xi}_2 &= (\hat{\omega}\hat{\tau} + \hat{\omega}\hat{\phi} + \hat{\phi}\hat{\tau}).\end{aligned}$$

Similarly, by construction, $\mathbb{V}\{\tau_{\text{cide}}(X; \delta)\} = \mathbb{E}(\xi_1) - \mathbb{E}(\xi_2)^2$ where

$$\begin{aligned}\xi_1 &= 2\omega\tau(\omega\phi + \phi\tau) + (\omega\tau)^2, \quad \text{and} \\ \xi_2 &= (\omega\tau + \omega\phi + \phi\tau).\end{aligned}$$

Therefore, if we define $\hat{\xi}_3 = \hat{\xi}_1 - \hat{\xi}_2\mathbb{P}_n(\hat{\xi}_2)$ and $\xi_3 = \xi_1 - \xi_2\mathbb{E}(\xi_2)$, we have

$$\hat{\psi}_n - \mathbb{V}\{\tau_{\text{cide}}(X; \delta)\} = \mathbb{P}_n(\hat{\xi}_1) - \mathbb{P}_n(\hat{\xi}_2)^2 - \mathbb{P}(\xi_1) + \mathbb{P}(\xi_2)^2 = \mathbb{P}_n(\hat{\xi}_3) - \mathbb{P}(\xi_3).$$

Then, by adding and subtracting terms, we have the usual expansion

$$\hat{\psi}_n - \mathbb{V}\{\tau_{\text{cide}}(X; \delta)\} = (\mathbb{P}_n - \mathbb{P})\xi_3 + (\mathbb{P}_n - \mathbb{P})(\hat{\xi}_3 - \xi_3) + \mathbb{P}(\hat{\xi}_3 - \xi_3), \quad (49)$$

where the first term on the right hand side of the final equation will follow a central limit theorem, the second term is an empirical process term, and the third term is a bias term.

Starting with the third term, we see

$$\begin{aligned}\mathbb{P}(\hat{\xi}_3 - \xi_3) &= \mathbb{P}(\hat{\xi}_1 - \xi_1) + \mathbb{P}\{\hat{\xi}_2\mathbb{P}_n(\hat{\xi}_2) - \xi_2\mathbb{P}(\xi_2)\} \\ &\leq (\|\hat{\pi} - \pi\| + \|\hat{\mu} - \mu\|)^2 + \mathbb{P}\{\hat{\xi}_2\mathbb{P}_n(\hat{\xi}_2) - \xi_2\mathbb{P}(\xi_2)\},\end{aligned}$$

where the second line follows by the proof of Lemma 2. For the second term on the final line above, we see that

$$\begin{aligned}\mathbb{P}(\hat{\xi}_2\mathbb{P}_n(\hat{\xi}_2) - \xi_2\mathbb{P}(\xi_2)) &= [\mathbb{E}\{(\hat{\omega}\hat{\tau} + \hat{\omega}\hat{\phi} + \hat{\phi}\hat{\tau})\mathbb{P}_n(\hat{\omega}\hat{\tau} + \hat{\omega}\hat{\phi} + \hat{\phi}\hat{\tau})\} - \mathbb{E}\{(\omega\tau + \omega\phi + \phi\tau)\mathbb{E}(\omega\tau + \omega\phi + \phi\tau)\}] \\ &= [\mathbb{E}\{(\hat{\omega}\hat{\tau} + \hat{\omega}\hat{\phi} + \hat{\phi}\hat{\tau})\mathbb{P}_n(\hat{\omega}\hat{\tau} + \hat{\omega}\hat{\phi} + \hat{\phi}\hat{\tau})\} - \mathbb{E}(\omega\tau)^2] \\ &= \frac{1}{n}[\mathbb{E}\{(\hat{\omega}\hat{\tau} + \hat{\omega}\hat{\phi} + \hat{\phi}\hat{\tau})^2\} - \mathbb{E}(\hat{\omega}\hat{\tau} + \hat{\omega}\hat{\phi} + \hat{\phi}\hat{\tau})^2] + 2\mathbb{E}(\hat{\omega}\hat{\tau} + \hat{\omega}\hat{\phi} + \hat{\phi}\hat{\tau})^2 - 2\mathbb{E}(\omega\tau)^2 \\ &= o_{\mathbb{P}}(1/n) + \mathbb{E}\{(\hat{\omega}\hat{\tau} + \hat{\omega}\hat{\phi} + \hat{\phi}\hat{\tau})^2\} - \mathbb{E}(\omega\tau)^2 \\ &= o_{\mathbb{P}}(1/n) + \mathbb{E}(\hat{\omega}\hat{\tau} + \hat{\omega}\hat{\phi} + \hat{\phi}\hat{\tau} + \omega\tau)\mathbb{E}(\hat{\omega}\hat{\tau} + \hat{\omega}\hat{\phi} + \hat{\phi}\hat{\tau} - \omega\tau) \\ &\leq o_{\mathbb{P}}(1/n) + (\|\hat{\pi} - \pi\| + \|\hat{\mu} - \mu\|)\|\hat{\pi} - \pi\|,\end{aligned}$$

where the fourth line follows by Assumption (a), the fifth line because $a^2 - b^2 = (a + b)(a - b)$, the sixth by Assumption (a) which implies $\omega\tau$ is bounded, and by the proof of Lemma 1. Therefore, if $(\|\hat{\pi} - \pi\| + \|\hat{\mu} - \mu\|)^2 = o_{\mathbb{P}}(n^{-1/2})$, then $\sqrt{n}\mathbb{P}(\hat{\xi}_3 - \xi_3) = o_{\mathbb{P}}(1)$.

The second term from equation (49), the empirical process term, is simpler to bound. By Lemma 2 of the study by Kennedy *et al.* [79] and by sample splitting, we have

$$(\mathbb{P}_n - \mathbb{P})(\hat{\xi}_3 - \xi_3) = O_{\mathbb{P}}\left(\frac{\|\hat{\xi}_3 - \xi_3\|}{n}\right).$$

By the triangle inequality

$$\begin{aligned}\|\hat{\xi}_3 - \xi_3\| &\leq \|\hat{\xi}_1 - \xi_1\| + \|\hat{\xi}_2\mathbb{P}_n(\hat{\xi}_2) - \xi_2\mathbb{P}(\xi_2)\| \\ &= \|\hat{\xi}_1 - \xi_1\| + \|\mathbb{P}_n(\hat{\xi}_2)(\hat{\xi}_2 - \xi_2) + \xi_2\{\mathbb{P}_n(\hat{\xi}_2) - \mathbb{P}(\xi_2)\}\| \\ &\leq \|\hat{\xi}_1 - \xi_1\| + \|\hat{\xi}_2 - \xi_2\| + \|\mathbb{P}_n(\hat{\xi}_2) - \mathbb{P}(\xi_2)\|,\end{aligned}$$

where the last line follows by Assumption (a), which says that ξ_2 and $\hat{\xi}_2$ are bounded. All three terms are $o_{\mathbb{P}}(1)$ since $(\|\hat{\pi} - \pi\| + \|\hat{\mu} - \mu\|)^2 = o_{\mathbb{P}}(n^{-1/2})$. Thus,

$$\psi_n - \mathbb{V}\{\tau_{\text{cide}}(X; \delta)\} = (\mathbb{P}_n - \mathbb{P})(\xi_3) + o_{\mathbb{P}}(n^{-1/2}).$$

Therefore, by the central limit theorem,

$$\sqrt{n}[\psi_n - \mathbb{V}\{\tau_{\text{cide}}(X; \delta)\}] \rightsquigarrow N(0, \mathbb{V}(\xi_3)) = N(0, \sigma^2)$$

with σ^2 as defined in the theorem. □

Proposition 4. *Let*

$$\begin{aligned}\hat{\psi}_{n,1} &= \mathbb{P}_n[2\hat{\omega}(\hat{\mu}_1 - \hat{\mu}_0)\{\hat{\omega}\hat{\phi} + \hat{\phi}(\hat{\mu}_1 - \hat{\mu}_0)\} + \{\hat{\omega}(\hat{\mu}_1 - \hat{\mu}_0)\}^2], \quad \text{and} \\ \hat{\psi}_{n,2} &= \mathbb{P}_n\{\hat{\omega}\hat{\phi} + \hat{\phi}(\hat{\mu}_1 - \hat{\mu}_0) + \hat{\omega}(\hat{\mu}_1 - \hat{\mu}_0)\}^2.\end{aligned}$$

Under Assumptions 1 and 2, Assumption (a) from Theorem 1, and Assumption (a) from Theorem 3. If

$$(\|\hat{\pi} - \pi\| + \|\hat{\mu} - \mu\|)^2 = o_{\mathbb{P}}\left(\frac{1}{\sqrt{n}}\right),$$

then

$$\sqrt{n}[\hat{\psi}_{n,1} - \mathbb{E}\{\tau_{\text{cide}}(X; \delta)^2\}] \rightsquigarrow N(0, \mathbb{V}(\zeta_1)), \text{ and} \quad (50)$$

$$\sqrt{n}[\hat{\psi}_{n,2} - \mathbb{E}\{\tau_{\text{cide}}(X; \delta)\}^2] \rightsquigarrow N(0, 4\mathbb{V}(\zeta_2)) \quad (51)$$

where

$$\zeta_1 = 2\omega(\mu_1 - \mu_0)\{\omega\phi + \phi(\mu_1 - \mu_0)\} + \{\omega(\mu_1 - \mu_0)\}^2, \quad \text{and} \quad (52)$$

$$\zeta_2 = \mathbb{E}\{\omega\phi + \phi(\mu_1 - \mu_0) + \omega(\mu_1 - \mu_0)\}\{\omega\phi + \phi(\mu_1 - \mu_0) + \omega(\mu_1 - \mu_0)\}. \quad (53)$$

Proof. This result follows by Lemmas 1 and 2, the conditions of the Proposition, and the Delta method (for the second convergence result). \square

Proposition 5. *Let ζ_1, ζ_2 be as defined in equations (52) and (53). Then, when $\mathbb{V}\{\tau_{\text{cide}}(X; \delta)\} = 0$,*

$$\text{cov}(\zeta_1, \zeta_2) \geq 0, \quad (54)$$

so that

$$\mathbb{V}(\zeta_1 - \zeta_2) = \mathbb{V}(\zeta_1) + \mathbb{V}(\zeta_2) - 2\text{cov}(\zeta_1, \zeta_2) \leq \mathbb{V}(\zeta_1) + \mathbb{V}(\zeta_2). \quad (55)$$

Proof. By the assumption that $\mathbb{V}\{\tau_{\text{cide}}(X; \delta)\} = 0$, it follows that $\tau_{\text{cide}}(X; \delta) \equiv \omega(\mu_1 - \mu_0) = C$ for some constant C . As a reminder,

$$\begin{aligned}\zeta_1 &= 2\omega(\mu_1 - \mu_0)\{\omega\phi + \phi(\mu_1 - \mu_0)\} + \{\omega(\mu_1 - \mu_0)\}^2, \quad \text{and} \\ \zeta_2 &= \mathbb{E}\{\omega\phi + \phi(\mu_1 - \mu_0) + \omega(\mu_1 - \mu_0)\}\{\omega\phi + \phi(\mu_1 - \mu_0) + \omega(\mu_1 - \mu_0)\}.\end{aligned}$$

Therefore, when $\omega(\mu_1 - \mu_0) = C$, then

$$\begin{aligned}\zeta_1 &= 2C\{\omega\phi + \phi(\mu_1 - \mu_0)\} + C^2, \text{ and} \\ \zeta_2 &= \mathbb{E}\{\omega\phi + \phi(\mu_1 - \mu_0) + C\}\{\omega\phi + \phi(\mu_1 - \mu_0) + C\} \\ &= \mathbb{E}\{C\}\{\omega\phi + \phi(\mu_1 - \mu_0) + C\} \\ &= C\{\omega\phi + \phi(\mu_1 - \mu_0)\} + C^2.\end{aligned}$$

And so,

$$2\text{cov}(\zeta_1, \zeta_2) = 2C^2\mathbb{V}\{\omega\phi + \phi(\mu_1 - \mu_0)\} \geq 0,$$

which implies the result. \square

Proposition 3. *Under Assumptions 1 and 2, Assumption (a) from Theorem 1, and Assumption (a) from Theorem 3, if*

$$(\|\hat{\mu} - \mu\| + \|\hat{\pi} - \pi\|)^2 = o_{\mathbb{P}}\left(\frac{1}{\sqrt{n}}\right),$$

then the asymptotic Type I error rate of the test in (28) is less than or equal to α .

Proof. By definition,

$$\hat{\psi}_n = \hat{\psi}_{n,1} - \hat{\psi}_{n,2},$$

where $\hat{\psi}_{n,1}$ and $\hat{\psi}_{n,2}$ are defined in Proposition 4. By Proposition 4, $\hat{\psi}_{n,1}$ and $\hat{\psi}_{n,2}$ converge to non-degenerate distributions,

$$\begin{aligned}\sqrt{n}[\hat{\psi}_{n,1} - \mathbb{E}\{\tau_{\text{cide}}(X; \delta)^2\}] &\rightsquigarrow N(0, \mathbb{V}(\zeta_1)), \text{ and} \\ \sqrt{n}[\hat{\psi}_{n,2} - \mathbb{E}\{\tau_{\text{cide}}(X; \delta)\}^2] &\rightsquigarrow N(0, 4\mathbb{V}(\zeta_2)),\end{aligned}$$

where ζ_1 and ζ_2 are defined in equations (F) and (53). And so,

$$\sqrt{n}[\hat{\psi}_n - \mathbb{V}\{\tau_{\text{cide}}(X; \delta)\}] \rightsquigarrow N(0, \mathbb{V}(\zeta_1) + 4\mathbb{V}(\zeta_2) - 4\text{cov}(\zeta_1, \zeta_2)).$$

By Proposition 5, when $\mathbb{V}\{\tau_{\text{cide}}(X; \delta)\} = 0$,

$$\mathbb{V}(\zeta_1) + 4\mathbb{V}(\zeta_2) - 4\text{cov}(\zeta_1, \zeta_2) \leq \mathbb{V}(\zeta_1) + 4\mathbb{V}(\zeta_2).$$

Therefore, since $\hat{\sigma}_1^2$ and $\hat{\sigma}_2^2$ in equations (26) and (27) are consistent estimators for $\mathbb{V}(\zeta_1)$ and $4\mathbb{V}(\zeta_2)$, respectively by the weak law of large numbers, $\hat{\sigma}_1^2 + \hat{\sigma}_2^2$ is a consistent estimator for $\mathbb{V}(\zeta_1) + 4\mathbb{V}(\zeta_2)$. And so, when $\mathbb{V}\{\tau_{\text{cide}}(X; \delta)\} = 0$,

$$\lim_{n \rightarrow \infty} \mathbb{P}(\sqrt{n}\hat{\psi}_n \leq \Phi^{-1}(1 - \alpha)\sqrt{\hat{\sigma}_1^2 + \hat{\sigma}_2^2}) \leq \alpha,$$

which implies that the asymptotic Type I error of the test in (28) is less than or equal to α . \square