
B-Learner: Quasi-Oracle Bounds on Heterogeneous Causal Effects Under Hidden Confounding

Miruna Oprescu¹ Jacob Dorn² Marah Ghoummaid³ Andrew Jesson⁴ Nathan Kallus¹ Uri Shalit³

Abstract

Estimating heterogeneous treatment effects from observational data is a crucial task across many fields, helping policy and decision-makers take better actions. There has been recent progress on robust and efficient methods for estimating the conditional average treatment effect (CATE) function, but these methods often do not take into account the risk of hidden confounding, which could arbitrarily and unknowingly bias any causal estimate based on observational data. We propose a meta-learner called the B-Learner, which can efficiently learn sharp *bounds* on the CATE function under limits on the level of hidden confounding. We derive the B-Learner by adapting recent results for sharp and valid bounds of the average treatment effect (Dorn et al., 2021) into the framework given by Kallus & Oprescu (2023) for robust and model-agnostic learning of conditional distributional treatment effects. The B-Learner can use any function estimator such as random forests and deep neural networks, and we prove its estimates are valid, sharp, efficient, and have a quasi-oracle property with respect to the constituent estimators under more general conditions than existing methods. Semi-synthetic experimental comparisons validate the theoretical findings, and we use real-world data demonstrate how the method might be used in practice.

1. Introduction

Using data to estimate the causal effect of actions is a fundamental task in medicine, economics, education research, and more. For instance, we might wish to use patient data

¹Cornell University and Cornell Tech ²Princeton University ³Technion, Israel Institute of Technology ⁴OATML, University of Oxford. Correspondence to: Miruna Oprescu <amo78@cornell.edu>.

to estimate which patients react well to a certain medication and which patients should avoid it. In many cases, due to economic and ethical considerations, the data available for these tasks is *observational data*, i.e. data that was not collected as part of a randomized experiment. Using such data carries the risk of *unobserved confounding*: correlations between the observed interventions and outcomes that are not accounted for in the available data. For example, patients with more social support might tend to receive certain interventions over others. If the level of a patient’s social support is not recorded in the data, the estimated effect of the intervention will be biased due to not observing the confounder of social support. Unobserved confounding cannot be detected from data, and its presence can lead to arbitrary and unknown bias in causal effect estimates. Such bias can lead to unreliable and potentially harmful policies.

In this work we are concerned with estimating causal effects on an individual level in the presence of a limited degree of unobserved confounding. Specifically, we give a method for effectively learning upper and lower bounds on the conditional average treatment effect (CATE) function that allows for flexible nuisance estimation and high-dimensional conditioning sets like patient medical records. The degree of allowed hidden confounding can be set by domain knowledge; alternatively, we can estimate what degree of hidden confounding is needed to significantly change our understanding of the CATE for any particular instance or sub-population.

We pursue the desirable treatment effect bound properties of validity, sharpness, efficiency, and robustness. A bound is called valid if it contains the true value of the causal estimand. A sharp bound is a valid bound that contains *only* those values of the causal estimand that could emerge from a plausible data generation process that could have produced the observed data (Ho & Rosen, 2017). Therefore, sharp bounds are the smallest possible bounds accounting for both observational data and domain knowledge (in the form of the degree of hidden confounding), a property which is important for precise decision making under hidden confounding. In contrast, a valid bound could in principle contain extraneous values, leading to overly cautious decision making. Efficient bounds converge to their target values using as lit

the data as possible. Typically, efficiency at best corresponds to quasi-oracle performance, where only slowly-consistent first-stage estimates are needed to achieve the same error bounds as we would obtain with access to oracle knowledge (Nie & Wager, 2021). Finally, a robust bound will be insensitive (within limits) to biases in the constituent estimators. We formalize these properties in Section 2.

In this paper, we present the B-Learner, for “bound-learner”, a scalable and flexible meta-learner for estimating *bounds* on the CATE function. The B-Learner is a meta-learner that uses a partially double-robust, Neyman-orthogonal estimating equation for the valid CATE bound characterization (Dorn et al., 2021). For unconfounded CATE estimation, there are several well-known meta-learners such as the X-Learner (Künzel et al., 2019), DR-Learner (Kennedy, 2020), and R-Learner (Nie & Wager, 2021). These methods allow the user to use any combination of learning methods (be it random forest, linear regression, or deep neural nets) and combine them efficiently to estimate the CATE function. In addition to flexibility, some of these methods have desirable rate and quasi-oracle properties. The B-Learner offers analogous flexibility, rate, and quasi-oracle guarantees for CATE *bounds* estimation, as well as novel bound validity and sharpness guarantees under appropriate conditions. We study the CATE bounds under Tan’s marginal sensitivity model (MSM) (Tan, 2006), which quantifies the degree of unobserved confounding through odds ratios. The properties of Tan’s MSM give the B-Learner validity under notably weak conditions.

We evaluate the B-Learner using synthetic and semi-synthetic experiments. In the synthetic experiments, the B-Learner displays quasi-oracle efficiency, requiring only a moderate amount of data for it to perform near-identically with estimated and oracle first-stage nuisances. The B-Learner also performs at least comparably to existing methods with analogous nuisances and can perform better with a well-tailored choice of second-stage regression function. In semi-synthetic experiments, we find the B-Learner is at least as effective as existing state-of-art models on a previously proposed benchmark. Finally, we illustrate the use of the B-Learner using real data to estimate the effect of 401(k) eligibility on financial wealth.

Related work. To the best of our knowledge, existing methods for CATE sensitivity analysis have yet to show all four of the properties of our proposed B-Learner. In particular, Kallus et al. (2019), Jesson et al. (2021) and Yin et al. (2022) each present methods that only achieve validity, and, to some degree, rate properties. These approaches start with estimators that have good properties under unconfoundedness and then optimize the estimated CATE or average treatment effect (ATE) bounds subject to a subset of the constraints implied by Tan’s MSM. Because these approaches

do not impose all implications of the MSM, they lack sharpness outside knife-edge cases. They also do not actively exploit the Neyman orthogonality of their solution, so they do not have the same rate guarantees that can be obtained under unconfoundedness. Lastly, these methods are tailored to specific learners and do not allow for the same flexibility as a meta-learner. Works such as Yadlowsky et al. (2022) and Chernozhukov et al. (2022) study bounds under other sensitivity assumptions. Yadlowsky et al. (2022) exploit Neyman orthogonality to obtain rate guarantees on CATE estimates and fast root- n guarantees on ATE estimates in Rosenbaum (2002)’s model, which they show are sharp when certain outcome symmetry properties hold. Chernozhukov et al. (2022) provide a method guaranteeing root- n consistency for *average* potential outcomes, treatment effects, and derivative bounds under limits on variance and covariance, and show that their estimates are sharp provided the bounds do not violate any implications of the observable data distribution. A rich literature on sensitivity analysis for ATEs exists, from as early as Cornfield et al. (1959) through Rosenbaum & Rubin (1983) to recent work like Colnet et al. (2022), but commentary on these methods is out of scope for this work.

2. Background and setup

We work in an observational data setting using the Neyman-Rubin potential outcomes framework. We assume data is drawn from an unobservable distribution P_{full} over $(X, A, Y(1), Y(0), U)$, where $A \in \{0, 1\}$ is a binary treatment, X is a set of baseline covariates in \mathbb{R}^d , $Y(1)$ and $Y(0)$ are the real-valued treated and untreated potential outcomes, respectively, and $U \in \mathbb{R}^k$ is an unobserved confounder. However, we face the fundamental problem of causal inference and only observe n draws from the coarsened distribution P over the observed variables $Z = (X, A, Y)$, where we assume that $Y = Y(A)$, i.e. (causal) consistency.

We are interested in learning about the conditional average treatment effect (CATE):

$$\tau(x) = \mathbb{E}_{P_{\text{full}}}[Y(1) - Y(0) \mid X = x].$$

The average treatment effect (ATE) is $\mathbb{E}[\tau(X)]$. When the (untestable) unconfoundedness assumption holds, formally $A \perp\!\!\!\perp Y(1), Y(0) \mid X$, then the CATE is equivalent to the difference in expected observed potential outcomes: $\tau(x) = \mathbb{E}_P[Y \mid X = x, A = 1] - \mathbb{E}_P[Y \mid X = x, A = 0]$. With the additional assumption of positivity, the CATE can be estimated with standard tools. However, the unconfoundedness assumption is untestable and often unrealistic, as we often have at least some degree of confounding unaccounted for by the observed covariates X . Therefore, we will assume unconfoundedness only holds with the addition of an unobserved $U \in \mathbb{R}^k$ for some k , such that $A \perp\!\!\!\perp Y(1), Y(0) \mid X, U$. In this case, it is possible to

bound $\tau(x)$ pointwise, for example, by assuming that unobserved confounding induces only a limited divergence between P and P_{full} .

We proceed under Tan’s Marginal Sensitivity Model (MSM) (Tan, 2006). Formally:

Assumption 1. Let $e(x, u) = P_{\text{full}}(A = 1 | X = x, U = u)$ and $e(x) = P(A = 1 | X = x)$ be the full and observed propensity scores, respectively. We assume $e(x), e(x, u) \in (0, 1)$ and that there exists $\Lambda \geq 1$ such that the following holds almost surely under P_{full} :

$$\Lambda^{-1} \leq \frac{e(x, u)}{1 - e(x, u)} \bigg/ \frac{e(x)}{1 - e(x)} \leq \Lambda.$$

The MSM imposes a bound on ratio between the full odds of treatment $e(x, u)/(1 - e(x, u))$ and the observed odds of treatment $e(x)/(1 - e(x))$. (The MSM is sometimes equivalently described using the log odds ratio bound $\log(\Lambda)$.) When $\Lambda = 1$, Assumption 1 is equivalent to the classic assumption of unconfoundedness with respect to the observed X . As Λ grows away from 1, greater unobserved confounding is allowed under the MSM and we can generally only estimate bounds on the CATE. In this paper our goal is to characterize these bounds, which describe a notion of “causal” uncertainty in the CATE estimate.

Remark 1. The sensitivity parameter Λ is a user-defined hyper-parameter as it specifies how much confounding to allow for. Choosing a suitable Λ is an ongoing area of study. Hsu & Small (2013) propose a procedure where we assess Λ values that correspond to dropping observed covariates and using domain knowledge to judge whether we omitted variables as important as these. Inversely, as our intervals increase with Λ , we can seek the Λ where a conclusion or decision would be overturned and judge whether the implied confounding is plausible. Ultimately, the choice of Λ is domain-specific and an analyst’s choice.

Notation. We now define the main notation, with a more detailed notation table available in Appendix A. To unify the analysis for upper (largest plausible CATE) and lower (smallest plausible CATE) bounds, we employ the convention that $+$, $-$ indicators symbolize upper and lower bounds, respectively. For nuisance functions (e.g. quantiles), these signs also encode the dependence on $\alpha = \Lambda/(\Lambda + 1)$ (and Λ) which we otherwise generally suppress in the remainder of the paper. We define the conditional outcome quantile and shorthand quantile notation:

$$q_c^*(x, a) = \inf\{\beta : F(\beta | x, a) \geq c\}$$

$$q_+^*(x, a) = q_\alpha^*(x, a), q_-^*(x, a) = q_{1-\alpha}^*(x, a).$$

The \pm and \mp symbols signal that an equation should be read twice, once with $\pm = +, \mp = -$ and once with $\pm = -, \mp = +$ (see example in Appendix A, Table 1). For

conciseness and clarity, we focus our main discussion on CATE upper bounds. In Appendix B, we provide a similar analysis of CATE lower bounds.

2.1. Properties of bound estimates

Our goal is to estimate the *identified set*: the set of CATEs that can be obtained in the unobserved distribution P_{full} generating the observed distribution P and satisfying the requirements of Assumption 1.

Definition 1. The *identified set* of estimands under Assumption 1 is the set of estimands that can be obtained for a distribution Q over $(X, A, Y(1), Y(0), U)$ such that the distribution of (X, A, Y) under Q matches the observed distribution P and $\Lambda^{-1} \leq \frac{Q(A=1|X=x,U=u)}{Q(A=0|X=x,U=u)} \bigg/ \frac{e^*(x)}{1-e^*(x)} \leq \Lambda$ almost surely. Let $\mathcal{M}(\Lambda)$ be the set of distributions Q that the observed data $Z = (X, Y, A)$ and Assumption 1 cannot rule out. Then, the *sharp* (upper) bounds on the identified set of conditional average potential outcomes and CATEs for a given point x are given by:

$$Y^+(x, a) \equiv \sup_{Q \in \mathcal{M}(\Lambda)} \mathbb{E}_Q[Y(a) | X = x]$$

$$\tau^+(x) \equiv \sup_{Q \in \mathcal{M}(\Lambda)} \mathbb{E}_Q[Y(1) - Y(0) | X = x].$$

Lower bounds follow symmetrically by replacing the suprema with infima. We note that the requirements of Assumption 1 decouple across x and are convex, so finding the identified set reduces to finding pointwise bounds. As we will see in Section 2.2, the CATE upper bounds $\tau^+(x)$ depend only on the observed distribution of data Z and the sensitivity parameter Λ . We can therefore ask what good properties we might want estimates $\hat{\tau}^+(x)$ to have. We suggest four desirable properties for bound estimation, of which the last two are closely linked:

Valid estimates. If $\hat{\tau}^+(x) < \tau^+(x) - o_p(1)$, then our estimated bounds would fail to cover the identified set and rule out plausible CATEs even asymptotically, which would be undesirable. Conversely, bound characterizations $\bar{\tau}$ satisfying $\bar{\tau}^+(x) \geq \tau^+(x)$ are called “valid” in the partial identification literature (Ho & Rosen, 2017), since Assumption 1 implies $\tau^+(x) \geq \mathbb{E}[Y(1) - Y(0) | X = x]$. Valid bounds (illustrated in Figure 1) give us some but not all information from our assumptions: every value they rule out is implausible, but some values they do not rule out may be implausible as well. We relax the notation and say that bound estimates $\hat{\tau}$ are *valid* if $\hat{\tau}^+(x) \geq \tau^+(x) - o_p(1)$.

Sharp estimates. If $\hat{\tau}^+(x) > \tau^+(x) + o_p(1)$, then our estimated bounds would fail to rule out impossible CATEs asymptotically under our assumptions. Exact characterizations of the identified sets are called “sharp” in the partial identification literature (Ho & Rosen, 2017). Sharpness is a




Figure 1. Example of a confounded CATE function with a true odds ratio Λ^* given by $\log(\Lambda^*) = 1.0$. The true $\tau(x)$ is the unobserved CATE in the full distribution, $\mathbb{E}_{P_{\text{full}}}[Y(1) - Y(0) | X = x]$. The confounded $\tau(x)$ is the biased estimand under assumed unconfoundedness, $\mathbb{E}_P[Y | X = x, A = 1] - \mathbb{E}_P[Y | X = x, A = 0]$. (1a): The sharp bounds from Result 1 for different levels of Λ . The true $\tau(x)$ is inside the sharp bounds for $\log(\Lambda) = 1.0$ (green). (1b): Example of *valid* and *sharp* CATE bounds as defined in Section 2.1.

stronger property than validity. We use lax notation to say that bound estimates $\hat{\tau}$ are *sharp* if $\hat{\tau}^+(x) = \tau^+(x) + o_p(1)$.

Efficient and robust estimates. We would like our bound estimates to converge to their limits at desirable rates and have multiple chances at sharp or valid limits. Ideally, we would be able to learn CATE bounds at the same rate as we could obtain under unconfoundedness. These properties relate to “double robust” estimators and may require constructing Neyman-orthogonal characterizations of valid, and ideally sharp, bounds.

2.2. Identification and estimation of sharp bounds

In this section, we use results from Dorn et al. (2021) to show how we can identify and estimate sharp CATE bounds from the observed data distribution, P . In order to express the sharp bounds, we introduce the following pseudo-outcomes from Dorn et al. (2021) that will correspond to the Conditional Value at Risk and the unobserved outcome bounds under Assumption 1:

$$\begin{aligned} H_{\pm}(z, \bar{q}) &= \bar{q}(x, a) + \frac{1}{1 - \alpha} \{y - \bar{q}(x, a)\}_{\pm} \\ R_{\pm}(z, \bar{q}) &= \Lambda^{-1}y + (1 - \Lambda^{-1})H_{\pm}(z, \bar{q}) \\ \rho_{\pm}^*(x, a, \bar{q}) &= \mathbb{E}[R_{\pm}(z, \bar{q}) | X = x, A = a]. \end{aligned}$$

We use the shorthand $\rho_{\pm}^*(x, a) = \rho_{\pm}^*(x, a, q_{\pm}^*)$ to write the ρ_{\pm}^* function evaluated at the true conditional quantiles q_{\pm}^* (which will end up corresponding to sharp bounds). The quantity $\text{CVaR}_{\pm}(x, a) := \mathbb{E}[H_{\pm}(z, q_{\pm}^*) | X = x, A = a]$ is known as the Conditional Value at Risk (Artzner et al., 1999; Kallus, 2022). In the distribution $Y | X = x, A = a$, $\text{CVaR}_{+}(x, a)$ is the expectation above the $(1 - \alpha)$ quantile, whereas $\text{CVaR}_{-}(x, a)$, is the expectation below the α quan-

tile. Hence, the pseudo-outcomes H and R correspond to the Conditional Value at Risk and conditional unobserved potential outcome, respectively.

Let $\mu^*(x, a) = \mathbb{E}[Y | X = x, A = a]$ be the conditional outcome regression in the observed data. Note that we can write the conditional potential outcome under Q as $\mathbb{E}_Q[Y(a) | X = x] = P[A = a | X = x]\mu^*(x, a) + P[A = 1 - a | X = x]\mathbb{E}_Q[Y(1 - a) | X = x, A = a]$ since Q must be consistent with the observed distribution P . Thus, it suffices to bound the conditional unobserved potential outcome $\mathbb{E}_Q[Y(1 - a) | X = x, A = a]$, which leads to the following result in terms of $\rho_{\pm}^*(x, a) = \rho_{\pm}^*(x, a, q_{\pm}^*)$:

Result 1 (Sharp bounds, (Dorn et al., 2021)). *The conditional average unobserved potential outcome $\mathbb{E}_Q[Y(1 - a) | X = x, A = a]$ has sharp upper and lower bounds under Assumption 1 given by $\rho_{+}^*(x, a)$ and $\rho_{-}^*(x, a)$, respectively. Thus, the sharp bounds on the conditional average potential outcomes can be written as:*

$$\begin{aligned} Y^+(x, 1) &= e^*(x)\mu^*(x, 1) + (1 - e^*(x))\rho_{+}^*(x, 1) \\ Y^-(x, 0) &= (1 - e^*(x))\mu^*(x, 0) + e^*(x)\rho_{-}^*(x, 0). \end{aligned}$$

The sharp CATE upper bound is further given by $\tau^+(x) = Y^+(x, 1) - Y^-(x, 0)$.

Thus, the bounds are a convex combination of the conditional outcome function $\mu^*(x, a)$ and the corresponding conditional CVaR terms, all of which can be estimated from P . As Λ grows, both the weight on the CVaR term in ρ^* grows and the CVaR term itself become more extreme. If the wrong putative quantile \bar{q} is used instead of the true q^* , the CVaR term moves the bound in a conservative yet valid direction. Finally, the difference between sharp conditional average potential outcome bounds $\tau^+(x)$ clearly

yields valid CATE bounds; those bounds are shown to be sharp by arguments outside the scope of this paper (Dorn & Guo, 2022). As we will see, this characterization of sharp and valid bounds alone will be insufficient for quasi-oracle estimation.

Pseudo-outcome regression for quasi-oracle estimation.

The expression of $\tau^+(x)$ suggests a plug-in strategy. We can estimate e , μ , and ρ through classification and regression and obtain bound estimates. However, such plug-in estimators are known to suffer from excessive bias due to the estimated nuisances (Kennedy, 2020; Kallus & Oprescu, 2023), especially when the nuisance functions are more complex than the CATE bounds. We follow the Kallus & Oprescu (2023) strategy and derive an efficient pseudo-outcome for the bounds based on the relevant influence function; we then regress that pseudo-outcome on X . We build on this literature to similarly provide an estimator for sharp CATE bounds with desirable properties beyond those of the plug-in estimators implied by Result 1.

Henceforth we will refer to e, q, ρ as nuisances as they will need to be estimated from data.

3. B-Learner: Pseudo-Outcome Regression for Doubly-Robust Sharp CATE Bounds

We propose a debiased learning procedure that consists of regressing a carefully constructed and nuisance-debiasing pseudo-outcome on covariates (Kennedy, 2020; Kallus & Oprescu, 2023).

Definition 2 (CATE Bounds Pseudo-Outcome). Let $\hat{\eta} = (\hat{e}, \hat{q}_-(\cdot, 0), \hat{q}_+(\cdot, 1), \hat{\rho}_-(\cdot, 0), \hat{\rho}_+(\cdot, 1)) \in \Xi$ be a set of estimated nuisances. We define the pseudo-outcome corresponding to the bounds for $Y^+(x, 1)$, $Y^-(x, 0)$ and $\tau^+(x)$ from Result 1 by:

$$\begin{aligned} \phi_1^+(Z, \hat{\eta}) &= AY + (1 - A)\hat{\rho}_+(X, 1) \\ &\quad + \frac{(1-\hat{e}(X))A}{\hat{e}(X)} \cdot (R_+(Z, \hat{q}_+(X, 1)) - \hat{\rho}_+(X, 1)), \\ \phi_0^-(Z, \hat{\eta}) &= (1 - A)Y + A\hat{\rho}_-(X, 0) \\ &\quad + \frac{\hat{e}(X)(1-A)}{(1-\hat{e}(X))} \cdot (R_-(Z, \hat{q}_-(X, 0)) - \hat{\rho}_-(X, 0)), \\ \phi_\tau^+(Z, \hat{\eta}) &= \phi_1^+(Z, \hat{\eta}) - \phi_0^-(Z, \hat{\eta}). \end{aligned}$$

The expressions in Definition 2 depend purely on the observed data distribution P , and so can be viewed as statistical estimands to be learned from the observed distribution.

When $\Lambda = 1$ and unconfoundedness holds, the expression for $\phi_\tau^+(Z, \hat{\eta})$ reduces to the familiar doubly-robust pseudo-outcome for CATE estimation, $\hat{\mu}(X, 1) - \hat{\mu}(X, 0) + \frac{A-\hat{e}(X)}{\hat{e}(X)(1-\hat{e}(X))}(Y - \hat{\mu}(X, A))$ (Kennedy, 2020; Knaus, 2022).

The pseudo-outcome is based on the efficient influence function of the estimand $\mathbb{E}[\tau^+(X)]$, so as we will see, small

errors in the nuisance estimation lead to ‘‘doubly small’’ (second-order) errors in the $\hat{\tau}^+(x)$ estimates. This special structure orthogonalizes the $\hat{\rho}_+$ estimation error in the plug-in bound estimand $AY + (1 - A)\hat{\rho}_+$ using the added term $\frac{(1-\hat{e}(X))A}{\hat{e}(X)}(R_+ - \hat{\rho}_+)$ that debiases $\hat{\rho}_+$ estimation error. The weighted CVaR terms $\rho_\pm^*(X, a, \bar{q})$ involve an objective which is sharpest when $\bar{q}_\pm = q_\pm^*$ and which turns out to have a second-order dependence on $\bar{q}_\pm - q_\pm^*$. Thus, quantile regression errors will move the pseudo-outcome in a conservative but still valid direction and consistent quantile regression errors will have favorable rate properties.

Algorithm 1 The B-Learner (detailed in Appendix E)

input Data $\{(X_i, A_i, Y_i) : i \in \{1, \dots, n\}\}$, folds $K \geq 2$, nuisance estimators, regression learner $\hat{\mathbb{E}}_n$

- 1: **for** $k \in \{1, \dots, K\}$ **do**
- 2: Use data $\{(X_i, A_i, Y_i) : i \neq k - 1 \pmod{K}\}$ to construct nuisance estimates $\hat{\eta}^{(k)} = (\hat{e}^{(k)}, \hat{q}^{(k)}, \hat{\rho}^{(k)})$
- 3: **for** $i = k - 1 \pmod{K}$ **do**
- 4: Set $\hat{\phi}_{\tau,i}^+ = \phi_\tau^+(Z_i, \hat{\eta}^{(k)})$
- 5: **end for**
- 6: **end for**

output $\hat{\tau}^+(x) = \hat{\mathbb{E}}_n[\hat{\phi}_\tau^+ | X = x]$

B-Learner. We call our full two-stage estimation procedure the *B-Learner*. Our procedure is summarized in Algorithm 1 (see Appendix E for a detailed version). In the first stage, we estimate the nuisances (outcome regression, propensity score, CVaR) with K -fold cross-fitting and construct Neyman-Orthogonal pseudo-outcome estimates based on Definition 2. In the second stage, we regress the estimated pseudo-outcomes on our covariates X , resulting in an estimated CATE bound function. As we will see now, the properties of this function depend on both the choice of nuisance estimators and the second-stage regressor.

Nuisance estimation. The propensity score $e^*(x)$ can be estimated using any standard probabilistic binary classifier. The quantiles q_\pm^* can be likewise estimated using any of several standard quantile regression methods (Yu & Jones, 1998; Meinshausen & Ridgeway, 2006; Athey et al., 2019). The modified outcome regression $\rho_\pm^*(x, a) = \Lambda^{-1}\mu^*(x, a) + (1 - \Lambda)^{-1}CVaR_\pm(x, a)$ is less standard, but it can be learned by either treating the CVaR pseudo-outcome R_\pm as an outcome, or separately learning the μ^* and $CVaR_\pm$ components of $\mathbb{E}[R_\pm | X = x, A = a]$. In the first approach, where we plug in the estimated quantiles into the expression for $R_\pm(Z, \bar{q})$ and then regress R_\pm onto X using any standard regressor, further sample splitting is theoretically required for estimating q^* and ρ^* . In the second approach, we can learn the μ^* and CVaR components on the same sample and then weight them accordingly to obtain estimates of ρ^* . The outcome regression $\mu^*(x, a)$ can be estimated via any regression learner and $CVaR_\pm$ can be

likewise estimated using several existing approaches (Athey et al., 2019; Kallus & Oprescu, 2023).

4. Theoretical properties of the B-Learner

We now describe the theoretical properties of our estimator. All proofs are in Appendix D. In Section 4.1, we use Kallus & Oprescu (2023)’s generic approach and Dorn et al. (2021)’s validity results to study the bias of the pseudo-outcome with first-stage nuisances. The pointwise bias from the **sharp** bounds is on the order of $|\widehat{e} - e^*| |\widehat{\rho} - \rho^*| + (\widehat{q}_- - q^*)^2$. When the quantiles are inconsistent, $(\widehat{q}_+ - q^*)^2$ and $(\widehat{q}_- - q^*)^2$ do not vanish. The pseudo-outcome bounds still remain **valid** in expectation, and any bias in the direction of failing to cover the identified CATE set disappears at a rate on the order of $|\widehat{e} - e^*| |\widehat{\rho} - \rho^*(\cdot, \widehat{q})|$. In Section 4.2, we characterize the second-stage regression and we show that we can learn CATE bounds at a **rate** dominated by the complexity of the target class. As a result, the estimator has robustness properties from the product-of-errors bias, with two chances at **sharp** bounds in L_2 norm and two chances at **valid** bounds on average. Our main text focuses on ERM-based second stage estimators with L_2 sharp bound guarantees. We show similar guarantees hold pointwise for linear smoother second-stage estimators in Appendix C.2.

4.1. Pseudo-outcome properties

We first analyze the bias in our proposed pseudo-outcomes.

Definition 3 (Conditional Pseudo-outcome Bias). Take $\widehat{\eta} \in \Xi$ be a set of estimated nuisances and let $\diamond \in \{0, 1, \tau\}$. We define the signed conditional pseudo-outcome bias: $\mathcal{E}_\diamond^+(x; \widehat{\eta}) = \mathbb{E}[\phi_\diamond^+(Z, \widehat{\eta}) - \phi_\diamond^+(Z, \eta^*) | X = x]$ and $\mathcal{E}_\diamond^-(x; \widehat{\eta}) = \mathbb{E}[\phi_\diamond^-(Z, \widehat{\eta}) - \phi_\diamond^-(Z, \eta^*) | X = x]$.

It immediately follows from Definition 3 that $\mathcal{E}_\tau^+(x; \widehat{\eta}) = \mathcal{E}_1^+(x; \widehat{\eta}) - \mathcal{E}_0^-(x; \widehat{\eta})$ and $|\mathcal{E}_\tau^+(x; \widehat{\eta})| \leq |\mathcal{E}_1^+(x; \widehat{\eta})| + |\mathcal{E}_0^-(x; \widehat{\eta})|$. The pseudo-outcome bias can be understood as the error incurred when performing pseudo-outcome regression with estimated nuisances rather than oracle nuisances. While any bias is undesirable, bias in one direction is worse. When $\mathcal{E}_\tau > 0$, the pseudo-outcomes are biased in a conservative but still valid direction. When $\mathcal{E}_\tau < 0$, the expected pseudo-outcomes are too aggressive and in expectation exclude plausible CATEs.

Our pseudo-outcomes fit into the framework of Kallus & Oprescu (2023) since the estimands and the nuisances are the solutions of conditional moment restrictions (see Proof of Theorem 1). Thus, under mild boundedness conditions, we can leverage their results to upper bound $|\mathcal{E}_\diamond^+|$.

Assumption 2 (Boundedness). Let $\widehat{\eta} \in \Xi$ be a set of estimated nuisances, and take $\overline{\eta} \in \text{conv}\{(\eta^*, \widehat{\eta})\}$.

- (i) $P(\epsilon \leq e^*(x), \widehat{e}(x) \leq 1 - \epsilon) = 1$ for some $\epsilon > 0$.
- (ii) $Y, \overline{q}_+(\cdot, 1), \overline{q}_-(\cdot, 0), \overline{\rho}_+(\cdot, 1), \overline{\rho}_-(\cdot, 0), f(\overline{q}_+(x, 1) | x, 1),$

$f(\overline{q}_-(x, 0) | x, 0)$ are all uniformly bounded.

The first condition in Assumption 2 is a standard requirement known as positivity, ensuring that both treatments and controls can be observed for any X with non-zero probability. The second condition is a common boundedness assumption often made in debiased machine learning for ATE and CATE in order to control the growth of $|\mathcal{E}_\tau^+|$. We now state the conditional Neyman orthogonality result we require, which we derive using the tools from Kallus & Oprescu (2023) and Dorn et al. (2021).

Theorem 1 (Pseudo-Outcome Conditional Neyman Orthogonality). *Suppose Assumption 2 holds. Then a Neyman-orthogonal characterization of the conditional outcome moment $\mathbb{E}[AY + (1 - A)\rho_+^*(X, 1) - Y^+(X, 1) | X] = 0$ has the form of ϕ_1^+ from Definition 2, and the symmetric result holds for ϕ_0^- . The absolute bias of the CATE upper bound has the product of rates bound:*

$$\begin{aligned} |\mathcal{E}_\tau^+(x; \widehat{\eta})| &\lesssim |\widehat{e}(x) - e^*(x)| |\widehat{\rho}_+(x, 1) - \rho_+^*(x, 1)| \\ &\quad + |\widehat{e}(x) - e^*(x)| |\widehat{\rho}_-(x, 0) - \rho_-^*(x, 0)| \\ &\quad + (\widehat{q}_+(x, 1) - q_+^*(x, 1))^2 \\ &\quad + (\widehat{q}_-(x, 0) - q_-^*(x, 0))^2. \end{aligned}$$

The undesirable direction of bias has the more favorable bound in terms of $\rho^(x, a, \widehat{q})$:*

$$\begin{aligned} \mathcal{E}_\tau^+(x; \widehat{\eta}) &\gtrsim -|\widehat{e}(x) - e^*(x)| |\widehat{\rho}_+(x, 1) - \rho_+^*(x, 1, \widehat{q}_+)| \\ &\quad - |\widehat{e}(x) - e^*(x)| |\widehat{\rho}_-(x, 0) - \rho_-^*(x, 0, \widehat{q}_-)|. \end{aligned}$$

Theorem 1 lets us characterize the pseudo-outcome biases.

Sharp pseudo-outcome bias. An immediate results is that the pseudo-outcome bias for the CATE (upper) bound is pointwise “doubly sharp” (Dorn et al., 2021): its bias tends to zero if \widehat{q}_\pm and one of \widehat{e} or $\widehat{\rho}_\pm$ are consistent, and the **rate** of bias goes to zero faster than the individual nuisances if all nuisances are consistent.

Valid pseudo-outcome bias. In some cases it may be difficult to estimate quantiles consistently or at a sufficient rate for the quantile error $(\widehat{q} - q^*)^2$ to vanish faster than $|\widehat{e} - e^*| |\widehat{\rho} - \rho^*|$. If so, the absolute value of pseudo-outcome bias relative to sharp bounds might be relevant to the second-stage estimates, but the level of bias in the direction of failing to cover the identified set still disappears at a product rate $|\widehat{e} - e^*| |\widehat{\rho} - \rho^*(\cdot, \widehat{q})|$. The pseudo-outcome estimator is therefore “doubly valid” (Dorn et al., 2021): its undesirable bias tends to zero if one of \widehat{e} or $\widehat{\rho}_\pm$ is consistent, and the **rate** of bias goes to zero faster than the individual nuisances if both are consistent.

Next, we leverage these results to illustrate the quasi-oracle properties of our B-Learner.

4.2. ERM-based estimators

We consider Algorithm 1 with an empirical risk minimization (ERM) algorithm as the second-stage estimator. In other words, given a class of functions $\mathcal{F} \subset [\mathcal{X} \rightarrow \mathbb{R}]$, the regression learner $\hat{\mathbb{E}}_n$ satisfies:

$$\hat{\mathbb{E}}_n [\hat{\phi}_\tau^+ | X = \cdot] \in \arg \min_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n (\hat{\phi}_{\tau,i}^+ - f(X_i))^2. \quad (1)$$

In this scenario, the error rates of our estimation procedure depend on the complexity of the class \mathcal{F} . These were studied in the context of learning with nuisance components in several works including Foster & Syrgkanis (2019); Kallus & Oprescu (2023). The implication of Theorem 1 is we can immediately apply Kallus & Oprescu (2023)’s Theorem 2 in our setting, employing bracketing entropy as a class complexity measure. We note that bracketing entropy is a *global* technique, with guarantees on the L_2 loss over the support of the estimand, in contrast with the *local* methods presented in Appendix C.2 which enable pointwise guarantees.

Corollary 1 (Rates for ERM Estimators, Theorem 2 from Kallus & Oprescu (2023)). *Suppose Assumption 2 holds for $\hat{\eta}^{(k)} \in \Xi, k \in \{1, \dots, K\}$. Let $\mathcal{E}_\tau^+(x) := \sum_{k=1}^K \mathcal{E}_\tau^+(x; \hat{\eta}^{(k)})$ and let $\hat{\mathbb{E}}_n[\cdot | X = x]$ be as in Equation (1). Further, suppose \mathcal{F} is convex and closed and has bracketing entropy $\log N_{[]}(\mathcal{F}, \epsilon) \lesssim \epsilon^{-r}$ with $0 < r < 2$ and that $|f(x)|$ is bounded $\forall f \in \mathcal{F}, x \in \mathcal{X}$. Then,*

$$\|\hat{\tau}^+(x) - \tau^+(x)\| \lesssim O_p(n^{-1/(2+r)}) + \|\mathcal{E}_\tau^+(x)\|.$$

Second-stage sharp consistency and robustness. When $\|\mathcal{E}_\tau^+(x)\| = o_p(1)$ and the conditions above hold, Corollary 1 shows that ERM estimates are L_2 consistent for the sharp CATE bounds. Learners satisfying the conditions of Corollary 1 include sparse linear models, neural networks, kernel classes (Foster & Syrgkanis, 2019), and Besov, Sobolev, Hölder-type function classes (Nickl & Pötscher, 2007). L_2 consistency of the pseudo-outcome bias follows if \hat{q} and one of \hat{e} or $\hat{\rho}$ are L_2 consistent.

Second-stage sharp rates. If $\|\mathcal{E}_\tau^+(x)\| = o_p(n^{-1/(2+r)})$ and the conditions of Corollary 1 hold, the pseudo-outcome bias has a negligible contribution to the CATE bounds estimation error. Thus, the estimation error is equivalent to the error as if the nuisances were known, a result known as the “quasi-oracle property” ((Nie & Wager, 2021)). Because the pseudo-outcome bias involves the product of rates, it will be sufficient to ask all pseudo-outcome nuisances to be consistent at an $o_p(n^{-1/4})$ rate. We give an example of sufficient conditions for our estimator to be oracle efficient (the property we synonymously call “quasi-oracle” in our main text) in Appendix C.1.

Second-stage validity. When the quantile estimates are inconsistent, we cannot apply Corollary 1 directly. Still, we

will have two chances to derive CATE bound estimates that are valid on average. In Appendix C.2, we show that linear smoothers can yield stronger pointwise validity guarantees.

Corollary 2 (ERM Validity on Average). *Assume the conditions of Corollary 1 are satisfied and for all $f \in \mathcal{F}$ and $c \in \mathbb{R}$ we have $f + c \in \mathcal{F}$. If $\|\hat{q}_+(\cdot, 1) - \bar{q}_+(\cdot, 1)\| = o_P(1)$ and $\|\hat{q}_-(\cdot, 0) - \bar{q}_-(\cdot, 0)\| = o_P(1)$ for a (potentially inconsistent) putative quantile function \bar{q} and either $\|\hat{e} - e^*\| = o_P(1)$ or both $\|\hat{\rho}_+(\cdot, 1) - \rho_+(\cdot, 1, \bar{q}_+)\| = o_p(1)$ and $\|\hat{\rho}_-(\cdot, 0) - \rho_-(\cdot, 0, \bar{q}_-)\| = o_p(1)$, then the estimated CATE bounds are valid on average in the sense that $\frac{1}{n} \sum_{i=1}^n \hat{\tau}^+(X_i) - \tau^+(X_i) \geq -o_P(1)$.*

5. Experiments

In this section, we demonstrate our method on synthetic and semi-synthetic datasets, as well as on a real-world case study. We first benchmark the B-Learner using a synthetic example similar to that in Kallus et al. (2019). We then illustrate how CATE bound estimators can be used for treatment deferral by using the hidden confounding variant of the IHDP dataset introduced by Jesson et al. (2021). For both sets of experiments, we compare with state-of-the-art methods proposed by Kallus et al. (2019) (*Sensitivity Kernel*) and Jesson et al. (2021) (*Quince*¹). We illustrate the usage of the B-Learner with real data through a case study of 401(k) eligibility effects on wealth. While we have focused our discussion on CATE upper bounds, our real data experiments also require estimating the CATE lower bounds we discuss in Appendix B. Details about the data generation processes, specific model implementation, hyperparameter selection and validation procedures used are given in Appendix F. We provide replication code at <https://github.com/CausalML/BLearner>.

While the *Sensitivity Kernel* approach uses Gaussian kernels and the *Quince* model uses Bayesian neural networks, the B-Learner (Algorithm 1) is flexible in the types of estimators allowed for both the first- and second-stage learners. We therefore compare three classes of nuisance and second-stage estimators: Random Forests (RF), Gaussian Kernels (GK), and Bayesian Neural Networks (NN). Whenever possible, we use the same hyperparameters and validation routine across models. For example, the B-Learner with NN nuisances uses the exact same neural networks as *Quince*.

We denote the upper bound given by the B-Learner output (Algorithm 1) by $\hat{\tau}^+(\{1^{\text{st}} \text{ stage}\}, \{2^{\text{nd}} \text{ stage}\})$ (e.g. $\hat{\tau}^+(RF, RF)$) to indicate the type of first- and second-stage learners used. For insight into the theoretical properties of our estimator, we also provide an oracle first-stage estimator

¹Jesson et al. (2021) train an ensemble of several models, which is a computationally intensive task. For the purposes of this section, we do not ensemble any of the compared methods.




Figure 2. Mean squared error (MSE) for different $\hat{\tau}^+$ learners. Shaded regions depict plus/minus one standard error over 50 simulations.

$\hat{\tau}^+(\text{Oracle}, \{2^{\text{nd}} \text{ stage}\})$ which uses the true nuisances in the pseudo-outcome calculation, as well as a “plug-in” estimator $\hat{\tau}^+(\{1^{\text{st}} \text{ stage}\}, \text{Plugin})$ which plugs in the estimated nuisances into the expressions from Result 1.

5.1. Simulated Data

Our synthetic dataset is sampled as follows:

$$X \sim \text{Unif}([-2, 2]^5), \quad A | X \sim \text{Bern}(\sigma(0.75X_0 + 0.5)), \\ Y \sim \mathcal{N}((2A - 1)(X_0 + 1) - 2 \sin((4A - 2)X_0), 1),$$

where σ is the sigmoid function. We wish to provide an estimate $\hat{\tau}^+(x)$ for the CATE upper bound under a level of confounding given by $\log \Lambda = 1$. With this simulation, it is straightforward to obtain the true nuisances e^*, μ^*, ρ^* . These, along with Result 1, allow us to determine the true value $\tau^+(x)$ of the upper bound. We run 50 simulations for sample sizes $n = 100, 200, 400, \dots, 12800$ and evaluate the different models on a fixed test set of 400 data points initially drawn at random. We compare the mean squared error (MSE) performance of each estimator with respect to the true bound and depict our findings in Figure 2.

In Figure 2a, we study the MSE convergence rates of the $\hat{\tau}^+(RF, RF)$ estimator, along with its oracle and plug-in variants. The convergence rate of our estimator matches the rate of the oracle estimator. That is, Algorithm 1 with more than a few hundred observations performs essentially as well as if the estimator had access to the true, oracle nuisances. This confirms our theoretical results from Corollary 1 in that small errors in the nuisance estimation lead to second-order errors in $\hat{\tau}(x)$. Moreover, we see that the simple plug-in estimator suffers from so-called plug-in bias for every value of n , as anticipated. The B-Learner MSE improvement slows for large n , which we expect reflects our use of rules-of-thumb to extrapolate hyperparameters to large samples.

In Figure 2b, we benchmark our estimator against *Sensitivity Kernel* and *Quince* for various first- and second- stage combinations. We see that using the same nuisances (GKs and NNs, respectively) leads to our method performing compa-




Figure 3. IHDP Hidden Confounding: Error recommendation rate for different values of the percentage of deferred points. The x-axis represents different levels of practitioner caution by varying the percentage of recommendations deferred.

rably with competitors. However, the B-Learner with NN or GK first stages and with RF second stage learners performs better than the state-of-the-art methods. This result underscores the importance of flexibility in choosing nuisance estimators, a key property of our method.

5.2. IHDP Hidden Confounding

We now show how the B-Learner can be used for other causal inference tasks, such as informing *deferral policies* for treatment recommendations. We replicate the experiment from Jesson et al. (2021) on IHDP Hidden Confounding. The dataset is multi-dimensional, has low overlap, and has hidden confounding due to a single covariate being hidden from the training models. The dataset contains synthetic potential outcomes generated according to the response surface B described by Hill (2011). We use the same deferral policy as in Jesson et al. (2021), namely, the policy simulates either recommending treatment or deferral to an expert. We make a treatment recommendation (either $A = 0$ or $A = 1$, according to the sign of CATE estimate) if and only if the predicted CATE interval excludes zero.




Figure 4. Example of lower and upper bound effect distributions for $\log \Lambda = 0.2$, as well as the percentage of negative effect lower bounds for as a function of $\log(\Lambda) = 0.1, \dots, 1.0$.

We measure model performance in terms of recommendation error rate across multiple deferral rates. The deferral rate is the fraction of observations for which we defer the action decision to the expert. The error rate is the percentage of observations for which we recommend the wrong treatment, among those in which we did not defer. Note that in this experiment we know the best treatment for each unit since we simulate both potential outcomes, although these effects do not correspond to a sharp bound under Assumption 1.

We compare two different variants of our B-Learner: $\hat{\tau}^{\pm}(RF, RF)$ and $\hat{\tau}^{\pm}(NN, NN)$ with *Sensitivity Kernel* and *Quince*. We see in Figure 3 that the RF B-Learner outperforms the GK-based *Sensitivity Kernel* method, and that the best performing methods are the NN B-Learner and *Quince* which perform very similarly.

5.3. Impact of 401(k) Eligibility on Wealth Distribution

We apply the B-Learner to illustrate the impact of hidden confounding in a study of 401(k) eligibility and its effects on financial wealth. We use the real-world dataset from Chernozhukov & Hansen (2004) that draws on the 1991 Survey of Income and Program Participation. The treatment of interest is 401(k) eligibility, while the target outcome is the net financial assets of an individual (taken as the aggregate of 401(k) balance, bank accounts and interest-earning assets minus non-mortgage debt).

This 401(k) eligibility dataset has been used in many analyses (Poterba et al., 1994; Chernozhukov & Hansen, 2004), often assuming unconfoundedness holds given observed covariates and finding a strong positive effect. However, unconfoundedness is an untestable assumption, so here we explore the uncertainty in the (conditional) treatment effects under varying degrees of hidden confounding. To that end, we apply the B-Learner algorithm repeatedly for different settings of Λ : $\log \Lambda = 0.1, 0.2, \dots, 1.0$. The nuisances are all estimated using Random Forest models with hyperparameters as in (Chernozhukov et al., 2018). We also estimate the CATE under unconfoundedness ($\log \Lambda = 0$ which corresponds to the DR-Learner (Kennedy, 2020)).

In Figure 4, we plot the distribution of predicted conditional effects on the 9,915 observations for $\log \Lambda = 0.2$ as well as the fraction of negative lower bound effects (frequency of $\mathbb{I}(\hat{\tau}^-(x) \leq 0)$) as we vary Λ . For lower values of Λ , the majority of lower bounds are still positive, which means that under those levels of confounding, most true conditional treatment effects are still positive. However, as we increase Λ , more and more of the true effects could be negative as the lower bound is comprised of mostly negative effects. For example, at $\log(\Lambda) = 0.6$, about half of the CATE lower bounds are negative which is to be interpreted as: if the data were truly confounded at this level, 50% of the effects measured as positive could in reality have been negative due to unobserved confounders. Regardless of what Λ level is most appropriate here, we see the B-Learner is a powerful tool for practitioners who wish to conduct what-if experiments for potential unobserved confounding.

6. Conclusion

We presented the B-Learner, a meta-learner for estimating bounds on the CATE function. The B-Learner can use any learning method as its base learners, including random forests and neural nets. We showed that the B-Learner provides bound estimates that are valid, sharp, robust, and have quasi-oracle rate properties, making it (to the best of our knowledge) the first CATE sensitivity analysis method with all these properties. Experiments validate our theoretical findings, show that the B-learner is comparable in performance to existing state-of-the-art methods, and demonstrate it can be used with real-world data to gain insight into the uncertainty of estimated causal effects.

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant No. 1846210. M.O. was supported by U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, under Award Number DE-SC0023112. M.G and U.S. were supported by Israeli Science Foundation, grant number 1950/19. Some of this material is based upon work by J.D. supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-2039656. The authors would like to thank the anonymous reviewers and Rebecca Dorn for useful discussion and feedback.

References

- Artzner, P., Delbaen, F., Eber, J.-M., and Heath, D. Coherent measures of risk. *Mathematical finance*, 9(3):203–228, 1999.
- Athey, S., Tibshirani, J., and Wager, S. Generalized random

- forests. *The Annals of Statistics*, 47(2):1148–1178, 2019.
- Chernozhukov, V. and Hansen, C. The effects of 401 (k) participation on the wealth distribution: an instrumental quantile regression analysis. *Review of Economics and Statistics*, 86(3):735–751, 2004.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., and Robins, J. Double/debiased machine learning for treatment and structural parameters, 2018.
- Chernozhukov, V., Cinelli, C., Newey, W., Sharma, A., and Syrgkanis, V. Long story short: Omitted variable bias in causal machine learning. Technical report, National Bureau of Economic Research, 2022.
- Colnet, B., Josse, J., Varoquaux, G., and Scornet, E. Causal effect on a target population: a sensitivity analysis to handle missing covariates. *Journal of Causal Inference*, 10(1):372–414, 2022.
- Cornfield, J., Haenszel, W., Hammond, E. C., Lilienfeld, A. M., Shimkin, M. B., and Wynder, E. L. Smoking and lung cancer: recent evidence and a discussion of some questions. *Journal of the National Cancer Institute*, 22(1):173–203, 1959.
- Dorn, J. and Guo, K. Sharp sensitivity analysis for inverse propensity weighting via quantile balancing. *Journal of the American Statistical Association*, 0(0):1–13, 2022. doi: 10.1080/01621459.2022.2069572.
- Dorn, J., Guo, K., and Kallus, N. Doubly-valid/doubly-sharp sensitivity analysis for causal inference with unmeasured confounding. *arXiv preprint arXiv:2112.11449*, 2021.
- Foster, D. J. and Syrgkanis, V. Orthogonal statistical learning. *arXiv preprint arXiv:1901.09036*, 2019.
- Hill, J. L. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240, 2011.
- Ho, K. and Rosen, A. Partial identification in applied research: Benefits and challenges. In Honore, B., Pakes, A., Piazzesi, M., and Samuelson, L. (eds.), *Advances in Economics and Econometrics: Eleventh World Congress (Econometrics Society Monographs)*, volume II, pp. 307–359. Cambridge University Press, Cambridge, 2017.
- Hsu, J. Y. and Small, D. S. Calibrating sensitivity analyses to observed covariates in observational studies. *Biometrics*, 69(4):803–811, 2013. ISSN 0006341X, 15410420.
- Jesson, A., Mindermann, S., Gal, Y., and Shalit, U. Quantifying ignorance in individual-level causal-effect estimates under hidden confounding. In *International Conference on Machine Learning*, pp. 4829–4838. PMLR, 2021.
- Kallus, N. Treatment effect risk: Bounds and inference. *arXiv preprint arXiv:2201.05893*, 2022.
- Kallus, N. and Oprescu, M. Robust and agnostic learning of conditional distributional treatment effects. In *International Conference on Artificial Intelligence and Statistics*, pp. 6037–6060. PMLR, 2023.
- Kallus, N., Mao, X., and Zhou, A. Interval estimation of individual-level causal effects under unobserved confounding. In *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS) 2019*, volume 89, 2019.
- Kennedy, E. H. Optimal doubly robust estimation of heterogeneous causal effects. *arXiv preprint arXiv:2004.14497*, 2020.
- Kennedy, E. H. Towards optimal doubly robust estimation of heterogeneous causal effects, 2022.
- Knaus, M. C. Double machine learning-based programme evaluation under unconfoundedness. *The Econometrics Journal*, 25(3):602–627, 2022.
- Künzel, S. R., Sekhon, J. S., Bickel, P. J., and Yu, B. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the national academy of sciences*, 116(10):4156–4165, 2019.
- Meinshausen, N. and Ridgeway, G. Quantile regression forests. *Journal of machine learning research*, 7(6), 2006.
- Nickl, R. and Pötscher, B. M. Bracketing metric entropy rates and empirical central limit theorems for function classes of Besov-and Sobolev-type. *Journal of Theoretical Probability*, 20(2):177–199, 2007.
- Nie, X. and Wager, S. Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika*, 108(2):299–319, 2021.
- Poterba, J. M., Venti, S. F., and Wise, D. A. 401 (k) plans and tax-deferred saving. *Studies in the Economics of Aging*, pp. 105–142, 1994.
- Rosenbaum, P. R. *Observational Studies*. Springer, 2002.
- Rosenbaum, P. R. and Rubin, D. B. Assessing sensitivity to an unobserved binary covariate in an observational study with binary outcome. *Journal of the Royal Statistical Society: Series B (Methodological)*, 45(2):212–218, 1983.
- Stone, C. J. Consistent nonparametric regression. *The Annals of Statistics*, pp. 595–620, 1977.

- Tan, Z. A distributional approach for causal inference using propensity scores. *Journal of the American Statistical Association*, 101(476):1619–1637, 2006.
- Wasserman, L. *All of nonparametric statistics*. Springer Science & Business Media, 2006.
- Yadlowsky, S., Namkoong, H., Basu, S., Duchi, J., and Tian, L. Bounds on the conditional and average treatment effect with unobserved confounding factors. *The Annals of Statistics*, 50(5):2587–2615, 2022.
- Yin, M., Shi, C., Wang, Y., and Blei, D. M. Conformal sensitivity analysis for individual treatment effects. *Journal of the American Statistical Association*, pp. 1–14, 2022.
- Yu, K. and Jones, M. Local linear quantile regression. *Journal of the American statistical Association*, 93(441):228–237, 1998.

Note: Throughout the appendix, we use \pm notation to encode either upper/ lower bounds results. This allows us to unify upper/ lower results and proofs at the cost of some readability.

A. Notation

We summarize the notation we use throughout this work in Table 1. In addition, note that we use upper case letters (e.g. X) to denote random variables and lower case letters (e.g. x) to refer to specific values of a random variable.

Table 1. Notation

X	The observed covariates in \mathbb{R}^d
A	A binary treatment ($A \in \{0, 1\}$)
Y	The outcome
Z	(X, A, Y) which is drawn from an observed distribution P
$Y(1), Y(0)$	Real-valued treated and untreated potential outcomes, respectively
U	The unobserved confounder in \mathbb{R}^k
P_{full}	An unobservable distribution over $(X, A, Y(1), Y(0), U)$
α	$\frac{\Lambda}{\Lambda+1} \in [0.5, 1)$ for $\Lambda \geq 1$
$\{b\}_+, \{b\}_-$	$\max\{b, 0\}, \min\{b, 0\}$ respectively, for a real number, b
$b \lesssim d$	$b \leq Cd$, for $b, d \in \mathbb{R}$, and for some universal constant C
g^*	The true value of a function g
\bar{g}	A putative value of a function g
\hat{g}	An estimated value of a function g from data
$\ g\ := \mathbb{E}_F[g(z)^2]^{1/2}$	The L_2 norm of g given a probability distribution $F(z)$ and a function $g(z)$
$+, -$	Indicators to symbolize upper and lower bounds, respectively
\pm, \mp	Symbols signal that an equation should be read twice, once with $\pm = +, \mp = -$ and once with $\pm = -, \mp = +$ E.g.: $a^\pm = b^\pm + c^\mp$ encodes two equalities: $a^+ = b^+ + c^-$ and $a^- = b^- + c^+$
$e(x)$	The observed propensity score $P(A = 1 X = x)$
$e(x, u)$	The full propensity score $P_{\text{full}}(A = 1 X = x, U = u)$
$F(y x, a)$	The conditional outcome distribution, $P(Y \leq y X = x, A = a)$
$f(y x, a)$	The conditional outcome density, $\frac{d}{dy}F(y x, a)$
$\mu^*(x, a)$	$\mathbb{E}[Y X = x, A = a]$, outcome regression
$q_c^*(x, a)$	$\inf\{\beta : F(\beta x, a) \geq c\}$, conditional outcome quantile
$q_+^*(x, a)$	$q_\alpha^*(x, a)$, shorthand α^{th} quantile notation
$q_-^*(x, a)$	$q_{1-\alpha}^*(x, a)$, shorthand $(1 - \alpha)^{\text{th}}$ quantile notation
$H_\pm(z, \bar{q})$	$\bar{q}(x, a) + \frac{1}{1-\alpha} \{y - \bar{q}(x, a)\}_\pm$, Conditional Value at Risk pseudo-outcome
$\text{CVaR}_\pm(x, a)$	$\mathbb{E}[H_\pm(z, q_\pm^*) X = x, A = a]$, the Conditional Value at Risk
$\text{CVaR}_+(x, a)$	The expectation above the $(1 - \alpha)$ quantile
$\text{CVaR}_-(x, a)$	The expectation below the α quantile
$R_\pm(z, \bar{q})$	$\Lambda^{-1}y + (1 - \Lambda^{-1})H_\pm(z, \bar{q})$, pseudo-outcome for the (conditional) unobserved potential outcome
$\rho_\pm^*(x, a, \bar{q})$	$\mathbb{E}[R_\pm(z, \bar{q}) X = x, A = a]$, the (conditional) expected unobserved potential outcome
$\rho_\pm^*(x, a)$	A shorthand for $\rho_\pm^*(x, a, q_\pm^*)$, the ρ_\pm^* function evaluated at the true conditional quantiles q_\pm^*
CATE Bounds Pseudo-Outcomes	
$\phi_1^+(Z, \hat{\eta})$	$AY + (1 - A)\hat{\rho}_+(X, 1) + \frac{(1-\hat{e}(X))A}{\hat{e}(X)} \cdot (R_+(Z, \hat{q}_+(X, 1)) - \hat{\rho}_+(X, 1))$
$\phi_0^-(Z, \hat{\eta})$	$(1 - A)Y + A\hat{\rho}_-(X, 0) + \frac{\hat{e}(X)(1-A)}{(1-\hat{e}(X))} \cdot (R_-(Z, \hat{q}_-(X, 0)) - \hat{\rho}_-(X, 0))$
$\phi_\tau^+(Z, \hat{\eta})$	$\phi_1^+(Z, \hat{\eta}) - \phi_0^-(Z, \hat{\eta})$

B. Results for CATE Lower Bounds

The results for the CATE lower bound $\tau^-(x)$ can be obtained by interchanging $+$ and $-$ symbols in the nuisances and/or replacing A with $1 - A$. We state them here for completeness.

CATE lower bounds identification (Result 1). The sharp CATE lower bound is given by $\tau^-(x) = Y^-(x, 1) - Y^+(x, 0)$, where the relevant bounds on the conditional average potential outcomes can be expressed as:

$$\begin{aligned} Y^-(x, 1) &= e^*(x)\mu^*(x, 1) + (1 - e^*(x))\rho_-^*(x, 1), \\ Y^+(x, 0) &= (1 - e^*(x))\mu^*(x, 0) + e^*(x)\rho_+^*(x, 0). \end{aligned}$$

Thus, the lower bounds can be expressed as a convex combination of quantities that can be estimated from the observed data alone, i.e. they are *identifiable* from data.

Pseudo-outcomes for CATE lower bounds (Definition 2). Let $\hat{\eta} = (\hat{e}, \hat{q}_+(\cdot, 0), \hat{q}_-(\cdot, 1), \hat{\rho}_+(\cdot, 0), \hat{\rho}_-(\cdot, 1)) \in \Xi$ be a set of nuisances. The pseudo-outcomes for the bounds $Y^-(x, 1)$, $Y^+(x, 0)$ and $\tau^-(x)$ are given by:

$$\begin{aligned} \phi_1^-(Z, \hat{\eta}) &= AY + (1 - A)\hat{\rho}_-(X, 1) + \frac{(1 - \hat{e}(X))A}{\hat{e}(X)} \cdot (R_-(Z, \hat{q}_-(X, 1)) - \hat{\rho}_-(X, 1)), \\ \phi_0^+(Z, \hat{\eta}) &= (1 - A)Y + A\hat{\rho}_+(X, 0) + \frac{\hat{e}(X)(1 - A)}{(1 - \hat{e}(X))} \cdot (R_+(Z, \hat{q}_+(X, 0)) - \hat{\rho}_+(X, 0)), \\ \phi_\tau^-(Z, \hat{\eta}) &= \phi_1^-(Z, \hat{\eta}) - \phi_0^+(Z, \hat{\eta}). \end{aligned}$$

Validity and sharpness for CATE lower bounds. We call lower bound estimates $\hat{\tau}^-(x)$ *valid* if $\hat{\tau}^-(x) - \tau^-(x) \leq -o_P(1)$. Similarly, the lower bound estimates $\hat{\tau}^-(x)$ are *sharp* if $\hat{\tau}^-(x) = \tau^-(x) + o_P(1)$.

Pseudo-outcome bias for CATE lower bounds. The absolute bias of the CATE lower bound pseudo-outcome has the form:

$$\begin{aligned} |\mathcal{E}_\tau^-(x; \hat{\eta})| &\lesssim |\hat{e}(x) - e^*(x)| |\hat{\rho}_-(x, 1) - \rho_-^*(x, 1)| \\ &\quad + |\hat{e}(x) - e^*(x)| |\hat{\rho}_+(x, 0) - \rho_+^*(x, 0)| \\ &\quad + (\hat{q}_-(x, 1) - q_-^*(x, 1))^2 \\ &\quad + (\hat{q}_+(x, 0) - q_+^*(x, 0))^2. \end{aligned}$$

whereas the signed bias bound is given by:

$$\begin{aligned} \mathcal{E}_\tau^-(x; \hat{\eta}) &\lesssim |\hat{e}(x) - e^*(x)| |\hat{\rho}_-(x, 1) - \rho_-^*(x, 1, \hat{q}_-)| \\ &\quad - |\hat{e}(x) - e^*(x)| |\hat{\rho}_+(x, 0) - \rho_+^*(x, 0, \hat{q}_+)|. \end{aligned}$$

□

The proofs of the theorems and corollaries in the paper (Appendix D) are unified across lower/upper bounds by using the \pm notation described above. For example, we will write the consolidated pseudo-outcome bias bounds as:

$$\begin{aligned} |\mathcal{E}_a^\pm(x; \hat{\eta})| &\lesssim |\hat{e}(x) - e^*(x)| |\hat{\rho}_\pm(x, a) - \rho_\pm^*(x, a)| \\ &\quad + (\hat{q}_\pm(x, a) - q_\pm^*(x, a))^2 \\ \mp \mathcal{E}_a^\pm(x; \hat{\eta}) &\lesssim |\hat{e}(x) - e^*(x)| |\hat{\rho}_\pm(x, a) - \rho_\pm^*(x, a, \hat{q}_\pm)|. \end{aligned}$$

which, together with $\mathcal{E}_\tau^\pm(x; \hat{\eta}) = \mathcal{E}_1^\pm(x; \hat{\eta}) - \mathcal{E}_0^\mp(x; \hat{\eta})$, yield the bias bounds for the lower and upper CATE bounds.

C. More Estimation Results

C.1. More ERM Results

Corollary 3 (Conditions for ERM Oracle Efficiency). *Let \mathcal{F} be a class of β -smooth functions in d dimensions (i.e. Hölder) and let e, ρ_\pm, q_\pm be γ_e, γ_ρ , and γ_q -smooth functions, respectively. Then, the L_2 error rate of Algorithm 1 is $O_p(n^{-1/(2+d/\beta)} + n^{-2/(2+d/\gamma_q)} + n^{-(1/(2+d/\gamma_e)+1/(2+d/\gamma_\rho))})$. Furthermore, if $\gamma_q \geq \frac{d/2}{1+d/\beta}$ and $\gamma_\rho \gamma_e \geq \frac{d^2}{4} - \frac{(\gamma_\rho+d/2)(\gamma_e+d/2)}{1+2\beta/d}$, our estimator is oracle efficient in the sense that the leading order error is that of the oracle estimator, $\widehat{\mathbb{E}}_n[\phi_\tau^\pm(Z, \eta^*) | X = x]$.*

C.2. Doubly Robust-Style Smoothing Estimators

We now study the behavior of Algorithm 1 with a DR Learner-style smoothing estimator as the second-stage learner. This technique was introduced in Kennedy (2020) and includes a wide range of estimators satisfying certain stability conditions, with linear smoothers as the archetype of this class. In this section, we analyze a generic linear smoother defined as follows:

$$\widehat{\mathbb{E}}_n \left[\widehat{\phi}_\tau^\pm \mid X = x \right] = \frac{1}{n} \sum_{i=1}^n w_i(x) \widehat{\phi}_{\tau,i}^\pm$$

where the $w_i(x)$'s are weights learned on a different sample than $\widehat{\phi}_{\tau,i}^\pm$ (which can be achieved by sample splitting). Under mild regularity assumptions, this estimator can yield stronger guarantees in the form of pointwise error bounds.

Theorem 2 (Rates for Linear Smoothing Estimators). *Assume the conditions of Assumption 2. Then:*

$$|\widehat{\tau}^\pm(x) - \tau^\pm(x)| \lesssim |\widetilde{\tau}^\pm(x) - \tau^\pm(x)| + b_n^\pm(x) + O_p \left(\left(\|\widehat{\phi}_\tau^\pm - \phi_\tau^\pm\|_{w^2} + o_p(1) \right) \left(\frac{1}{n^2} \sum_{i=1}^n w_i(x)^2 \right)^{1/2} \right)$$

where $\widetilde{\tau}^\pm(x)$ corresponds to the linear smoother procedure with oracle first-stage nuisances, $\|\cdot\|_{w^2}$ is the empirical $w_i(x)^2$ -weighted distance of Kennedy (2022), and the $b_n^\pm(x)$ bias function is of the form:

$$b_n^\pm(x) = \left| \frac{1}{n} \sum_{i=1}^n w_i(x) \mathcal{E}_\tau^\pm(X_i; \widehat{\eta}) \right|$$

Second-stage sharp consistency and robustness. $\widehat{\tau}$ consistency follows under weak conditions like $\frac{1}{n} \sum_{i=1}^n |w_i(X_i)| \leq C$ (Stone, 1977). Thus, we can state corollaries that prove consistent estimation of sharp bounds under either strong restrictions on weights or strong requirements on consistency. We show one such corollary for a wide class of linear smoothers that includes linear and ridge regression, local polynomial and RKHS regression, kernel estimators, and some tree methods (Wasserman, 2006). In this corollary, we ask for uniform nuisance consistency to make the bias term b_n^\pm tend to zero.

Corollary 4 (Pointwise Consistency of Sharp CATE Bounds). *Assume the conditions of Theorem 2 are satisfied, the $\frac{w_i(x)}{n}$ weighting functions satisfy the requirements of Stone (1977) Theorem 1, and $\frac{1}{n} \sum |w_i(x)| = O_p(1)$. If \widehat{q}_\pm and either \widehat{e} or $\widehat{\rho}$ are uniformly consistent, then $\widehat{\tau}^\pm(x)$ converges to the true pointwise sharp CATE bounds.*

Second-stage sharp rates. Take $\tau^\pm, e, \rho_\pm, q_\pm$ be γ_q to be Hölder with smoothness $\beta, \gamma_e, \gamma_\rho,$ and γ_q . Then, the pointwise error rate of Algorithm 1 is $O_p \left(n^{-1/(2+d/\beta)} + n^{-2/(2+d/\gamma_q)} + n^{-(1/(2+d/\gamma_e)+1/(2+d/\gamma_\rho))} \right)$ and the estimator will be oracle efficient, though the error bounds here are *pointwise* (local), whereas the ERM-based bounds are L_2 (global).

Second-stage validity. The linear smoothers also have pointwise validity. Unlike in the ERM case where the best model fit to conservative bounds might extrapolate to invalid bounds for some regions of the covariates, the linear smoothers will have pointwise validity guarantees.

Corollary 5 (Pointwise Validity of Lax CATE Bounds). *Assume the conditions of Theorem 2 are satisfied, $w_i(x)$ satisfies the requirements of Theorem 1 in Stone (1977), and $\frac{1}{n} \sum_{i=1}^n |w_i(x)| = O_p(1)$. If \widehat{q}_\pm is uniformly consistent to some limiting quantile \bar{q}_\pm and \widehat{e} is uniformly consistent for e^* or $\widehat{\rho}_\pm$ is uniformly consistent for $\rho_\pm^*(X, A, \bar{q}_\pm)$ and $f(\bar{q}_\pm(x, a) \mid x, a) > 0$. Then the estimated bounds are pointwise valid in the sense that $\pm (\widehat{\tau}^\pm(x) - \tau^\pm(x)) \geq -o_P(1)$.*

D. Proofs

Note: we assume throughout that $X, U, Y(0)$, and $Y(1)$ to have probability measures absolutely continuous w.r.t. the Lebesgue measure so that we can condition on the event $X = x$.

Proof of Theorem 1. We start with the bound for the unsigned bias. Consider the $Y^+(X, 1)$ bound for simplicity. We first show that our problem fits into the framework of Kallus & Oprescu (2023) since the estimand and the oracle nuisances are the solutions of following conditional moment restrictions:

$$\mathbb{E}[AY + (1 - A)\rho_+^*(X, 1) - Y^+(X, 1) \mid X] = 0 \quad (\text{Estimand moment})$$

$$\mathbb{E}[R_+(Z, q_+^*(X, 1)) - \rho_+^*(X, 1) \mid X, A = 1] = 0 \quad (\text{Modified outcome moment})$$

$$\mathbb{E}[\alpha - \mathbb{I}(Y \leq q_+^*(X, 1)) \mid X, A = 1] = 0 \quad (\text{Quantile moment})$$

Let ν_1 be the nuisance set corresponding to this set of moments (as defined in [Kallus & Oprescu \(2023\)](#)). Then $\nu_1^*(X) = (Y^+(X, 1), \rho_+^*(X, 1), q_+^*(X, 1))$. This is different from η^* since the propensity does not have an estimating conditional moment. The Jacobian of the moments with respect to ν_1^* is thus given by:

$$J_1^*(X) = \begin{pmatrix} -1 & 1 - e^*(X) & 0 \\ 0 & -1 & 0 \\ 0 & 0 & -f(q_+^*(X, 1) | X, 1) \end{pmatrix}$$

where $f(y | x, 1)$ be the conditional density at a point y for $a = 1$. The first row of the inverse is then given by $\alpha_1^*(X) := (J_1^*(X))_1^{-1} = (-1, e^*(X) - 1, 0)$. Thus, using the pseudo-outcome in Definition 3 of [Kallus & Oprescu \(2023\)](#), replacing ν_1^*, α_1^* with their estimated counterparts $\widehat{\nu}_1(X) = (\widehat{Y}^+(X, 1), \widehat{\rho}_+(X, 1), \widehat{q}_+(X, 1))$, $\widehat{\alpha}_1(X) = (-1, \widehat{e}(X) - 1, 0)$, and noting that the first moment is conditional only on X , we obtain the pseudo-outcome:

$$\phi_1^+(Z, \widehat{\eta}) = AY + (1 - A)\widehat{\rho}_+(X, 1) + \frac{(1 - \widehat{e}(X))A}{\widehat{e}(X)} \cdot (R_+(Z, \widehat{q}_+) - \widehat{\rho}_+(X, 1))$$

as desired. Therefore, our Assumption 2 is a direct application of the boundedness assumption (Assumption 1) in [Kallus & Oprescu \(2023\)](#) and the bound for the unsigned bias follows largely from their Theorem 1. We first note that the results of Theorem 1 in [Kallus & Oprescu \(2023\)](#) also hold pointwise (see the proof in their Appendix A). It now remains to calculate the H and G matrices in their Assumption 1:

$$G = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad H = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

since G is just a binary mask for $J_1^*(X)$ and H involves second order derivatives of the moments. Plugging these into the bound for the unsigned bias, we obtain:

$$\begin{aligned} |\mathcal{E}_1^+(x; \widehat{\eta})| &\lesssim \sum_{i=1}^3 \sum_{j=1}^3 G_{ij} |\widehat{\alpha}_{1,i}(x) - \alpha_{1,i}^*(x)| |\widehat{\nu}_{1,j}(x) - \nu_{1,j}^*(x)| + \sum_{i=1}^3 \sum_{j=1}^3 H_{ij} |\widehat{\nu}_{1,i}(x) - \nu_{1,i}^*(x)| |\widehat{\nu}_{1,j}(x) - \nu_{1,j}^*(x)| \\ &\lesssim |\widehat{e}(x) - e^*(x)| |\widehat{\rho}_+(x, 1) - \rho_+^*(x, 1)| + (\widehat{q}_+(x, 1) - q_+^*(x, 1))^2. \end{aligned}$$

The result for $\mathcal{E}_0^+(x; \widehat{\eta})$ follows from replacing $a = 1$ with $a = 0$ everywhere. The bound for $\mathcal{E}_a^-(x; \widehat{\eta})$ follows from writing the corresponding conditional moments for $Y^-(X, a)$.

We now study the bound for the signed bias. We first take the expectation of $\phi_1^+(Z, \widehat{\eta})$:

$$\begin{aligned} \mathbb{E} [\phi_1^+(Z, \widehat{\eta}) | X] &= \mathbb{E} \left[AY + \left(1 - \frac{A}{\widehat{e}(X)}\right) \widehat{\rho}_+(X, 1) + A \frac{1 - \widehat{e}(X)}{\widehat{e}(X)} R_+(Z, \widehat{q}_+) | X \right] \\ &= e^*(X) \mu^*(X, 1) + \frac{\widehat{e}(X) - e^*(X)}{\widehat{e}(X)} \widehat{\rho}_+(X, 1) + \left(\frac{e^*(X)}{\widehat{e}(X)} - e^*(X) \right) \rho_+^*(X, 1, \widehat{q}_+) \\ &= e^*(X) \mu^*(X, 1) + \left(1 - \frac{e^*(X)}{\widehat{e}(X)}\right) (\widehat{\rho}_+(X, 1) - \rho_+^*(X, 1, \widehat{q}_+)) + (1 - e^*(X)) \rho_+^*(X, 1, \widehat{q}_+) \end{aligned}$$

As a result, we can write write:

$$\mathbb{E} [\phi_1^+(Z, \widehat{\eta}) - \phi_1^+(Z, \eta^*) | X] = \frac{\widehat{e}(X) - e^*(X)}{\widehat{e}(X)} (\widehat{\rho}_+(X, 1) - \rho_+^*(X, 1, \widehat{q}_+)) + (\rho_+^*(X, 1, \widehat{q}_+) - \rho_+^*(X, 1)) (1 - e^*(X))$$

Recall the CVaR property that $\rho_+^*(X, 1) = \inf_{\bar{q}} \rho_+^*(X, 1, \bar{q})$, so that $\rho_+^*(X, 1, \widehat{q}_+) \geq \rho_+^*(X, 1)$.

Therefore we have:

$$\begin{aligned} -\mathcal{E}_1^+(x; \widehat{\eta}) &= \mathbb{E} [\phi_1^+(Z, \eta^*) - \phi_1^+(Z, \widehat{\eta}) | X = x] \\ &= -\frac{\widehat{e}(x) - e^*(x)}{\widehat{e}(x)} (\widehat{\rho}_+(x, 1) - \rho_+^*(x, 1, \widehat{q}_+)) - (\rho_+^*(x, 1, \widehat{q}_+) - \rho_+^*(x, 1)) (1 - e^*(x)) \\ &\leq -\frac{\widehat{e}(x) - e^*(x)}{\widehat{e}(x)} (\widehat{\rho}_+(x, 1) - \rho_+^*(x, 1, \widehat{q}_+)) \\ &\lesssim -|\widehat{e}(x) - e^*(x)| |\widehat{\rho}_+(x, 1) - \rho_+^*(x, 1; \widehat{q}_+)| \end{aligned}$$

The result for $Y^+(X, 0)$ follows by symmetry. The result for $Y^-(X, a)$ follows by negating Y , applying the argument, and negating the argument. Results for $\tau^\pm(X)$ follow by Result 1. \square

Proof of Corollary 1. Since we showed in the proof of Theorem 1 that our pseudo-outcomes fit into the framework of Kallus & Oprescu (2023) and our Assumption 2 maps to their Assumption 1, we can apply their Theorem 2 directly to our setting, yielding the statement of our theorem. \square

Proof of Corollary 2. Using lax notation, choose an estimated $\hat{f}^\pm \in \mathcal{F}$ to minimize Equation (1) and then an estimated $\hat{c}^\pm \in \mathbb{R}$ such that $\hat{f}^\pm + \hat{c}^\pm$ is a minimizer of Equation (1). By construction, we must have $\hat{f}^\pm + 0$ is an optimizer.

If we differentiate (1) with respect to \hat{c}^\pm , evaluate it at 0, and divide by 2, we obtain the requirement on any optimizer that $\frac{1}{n} \sum_{i=1}^n (\hat{\phi}_{\tau,i}^\pm - \hat{f}^\pm(X_i)) = 0$. As a result:

$$\pm \left(\frac{1}{n} \sum_{i=1}^n \hat{\tau}^\pm(X_i) - \tau^\pm(X_i) \right) = \pm \left(\frac{1}{n} \sum_{i=1}^n \hat{\phi}_{\tau,i}^\pm - \tau^\pm(X_i) \right)$$

By applying Chebyshev's inequality to the average of zero-meaned bounded random variables $\hat{\tau}^\pm(X_i) - \tau^\pm(X_i) - \mathcal{E}_\tau^\pm(X_i; \hat{\eta})$, we can further obtain:

$$\begin{aligned} \pm \left(\frac{1}{n} \sum_{i=1}^n \hat{\tau}^\pm(X_i) - \tau^\pm(X_i) \right) &= \pm \frac{1}{n} \sum_{i=1}^n \mathcal{E}_\tau^\pm(X_i; \hat{\eta}) - O_p(n^{-1/2}) \\ &\geq -\frac{1}{n} \sum O \left(|\hat{e}(X) - e^*(X)| \sum_a |\hat{\rho}_\pm(X, a) - \rho_\pm^*(X, a, \hat{q}_\pm)| \right) - O_p(n^{-1/2}) \\ &\geq -O \left(\|\hat{e} - e^*\| \sum_a \|\hat{\rho}_\pm(\cdot, a) - \rho_\pm^*(\cdot, a, \hat{q}_\pm)\| \right) - o_p(1) = -o_p(1), \end{aligned}$$

demonstrating the desired bound. \square

Proof of Corollary 3. The L_2 convergence rate for a Hölder β -smooth functions in d dimension is $O_P(n^{-1/(2+d/\beta)})$. Taking e, ρ_\pm, q_\pm to be $\gamma_e, \gamma_\rho,$ and γ_q -Hölder, we have that their convergence rates are $O_P(n^{-1/(2+d/\gamma_e)}), O_P(n^{-1/(2+d/\gamma_\rho)}),$ and $O_P(n^{-1/(2+d/\gamma_q)})$ respectively. Thus, the L_2 conditional bias in Theorem 1 is bounded above by a term that is $O_p(n^{-2/(2+d/\gamma_q)} + n^{-(1/(2+d/\gamma_e))})$. Applying Corollary 1 with a β -smooth function class \mathcal{F} , we obtain the desired rate $O_p(n^{-1/(2+d/\beta)} + n^{-2/(2+d/\gamma_q)} + n^{-(1/(2+d/\gamma_e)+1/(2+d/\gamma_\rho))})$. The rest follows by algebraic manipulation. \square

Proof of Theorem 2. We first derive:

$$\begin{aligned} |\hat{\tau}^\pm(x) - \tau^\pm(x)| &= |\hat{\tau}^\pm(x) - \tau^\pm(x) + \hat{\tau}^\pm(x) - \tilde{\tau}^\pm(x)| \\ &= \left| \hat{\tau}^\pm(x) - \tau^\pm(x) + \frac{1}{n} \sum_{i=1}^n w_i(x) (\phi_\tau^\pm(Z_i, \hat{\eta}) - \phi_\tau^\pm(Z_i, \eta^*)) \right| \\ &= \left| \hat{\tau}^\pm(x) - \tau^\pm(x) + \frac{1}{n} \sum_{i=1}^n w_i(x) (\mathcal{E}_\tau^\pm(X_i; \hat{\eta}) + \phi_\tau^\pm(Z_i, \hat{\eta}) - \phi_\tau^\pm(Z_i, \eta^*) - \mathcal{E}_\tau^\pm(X_i; \hat{\eta})) \right| \\ &\leq |\hat{\tau}^\pm(x) - \tau^\pm(x)| + \left| \frac{1}{n} \sum_{i=1}^n w_i(x) \mathcal{E}_\tau^\pm(X_i; \hat{\eta}) \right| + \left| \frac{1}{n} \sum_{i=1}^n w_i(x) (\phi_\tau^\pm(Z_i, \hat{\eta}) - \phi_\tau^\pm(Z_i, \eta^*) - \mathcal{E}_\tau^\pm(X_i; \hat{\eta})) \right| \end{aligned}$$

Since $\phi_\tau^\pm(Z_i, \hat{\eta}) - \phi_\tau^\pm(Z_i, \eta) - \mathcal{E}_\tau^\pm(X_i; \hat{\eta})$ is zero-meaned conditional on X and nuisances (including weights), we can apply Chebyshev's inequality to randomness in $(A, Y) | X$ to obtain:

$$|\hat{\tau}^\pm(x) - \tau^\pm(x)| \leq |\tilde{\tau}^\pm(x) - \tau^\pm(x)| + \left| \frac{1}{n} \sum_{i=1}^n w_i(x) \mathcal{E}_\tau^\pm(X_i; \hat{\eta}) \right|$$

$$+ O_p \left(\|\widehat{\phi}_\tau^\pm - \phi_\tau^\pm - \mathcal{E}_\tau^\pm\|_{w^2} \frac{1}{n^2} \sum_{i=1}^n w_i(x)^2 \right)$$

Since $\mathcal{E}_\tau^\pm(X; \widehat{\eta}) = \mathbb{E} [\widehat{\phi}_\tau^\pm - \phi_\tau^\pm | X]$, we can take advantage of the weighted L^2 norm and the weak law of large numbers to further bound $\|\widehat{\phi}_\tau^\pm - \phi_\tau^\pm - \mathcal{E}_\tau^\pm\|_{w^2} \leq \|\widehat{\phi}_\tau^\pm - \phi_\tau^\pm\|_{w^2} + o_p(1)$:

$$|\widehat{\tau}^\pm(x) - \tau^\pm(x)| \leq |\widetilde{\tau}^\pm(x) - \tau^\pm(x)| + b_n^\pm(x) + O_p \left((\|\widehat{\phi}_\tau^\pm - \phi_\tau^\pm\|_{w^2} + o_p(1)) \left(\frac{1}{n^2} \sum_{i=1}^n w_i(x)^2 \right)^{1/2} \right),$$

which is the desired inequality. □

Proof of Corollary 4. By Stone (1977) Theorem 1 and since $|\phi_\tau^\pm|$ is bounded, $\widetilde{\tau}(x) \xrightarrow{p} \mathbb{E}[\phi_\tau^\pm(Z, \eta) | X = x] = \tau^\pm(x)$.

For the second term, we use the Theorem 1 and the supremum assumptions to derive:

$$\begin{aligned} |b_n^\pm(x)| &\leq \frac{1}{n} \sum_{i=1}^n |w_i(x)| \sup_x |\mathcal{E}_\tau^\pm(x; \widehat{\eta})| \\ &\lesssim \left(\sup_x |\widehat{e}(x) - e^*(x)| \right) \left(\sup_{x,a} |\widehat{\rho}_\pm(x, a) - \rho_\pm^*(x, a)| \right) + \sup_{x,a} (\widehat{q}_\pm(x, a) - q_\pm^*(x, a))^2 \\ &= o_p(1) \end{aligned}$$

For the final term, the sup consistency implies $\|\widehat{\phi}_\tau^\pm - \phi_\tau^\pm\| = o_p(1)$. We also have:

$$\frac{1}{n^2} \sum_{i=1}^n w_i(x)^2 \leq \left(\frac{1}{n} \sum_{i=1}^n |w_i(x)| \right)^2 = O_p(1)$$

So that $O_p \left(\frac{1}{n^2} \sum_{i=1}^n w_i(x)^2 \|\widehat{\phi}_\tau^\pm - \phi_\tau^\pm\|_{w^2} \right) = o_p(1)$. □

Proof of Corollary 5. If we define $\widetilde{\tau}^\pm(x)$ for the linear smoother estimate that uses $\bar{\eta}$ as first-stage nuisances, we can similarly argue that:

$$\begin{aligned} |\widehat{\tau}^\pm(x) - \widetilde{\tau}^\pm(x)| &= |\widetilde{\tau}^\pm(x) - \tau^\pm(x) + \widehat{\tau}^\pm(x) - \widetilde{\tau}^\pm(x)| \\ &\leq \frac{1}{n} \sum_{i=1}^n |w_i(x)| |\mathcal{E}_\tau^\pm(X_i; \widehat{\eta}) - \mathcal{E}_\tau^\pm(X_i; \bar{\eta})| \\ &\quad + O_p \left((\|\phi_\tau^\pm(\cdot, \widehat{\eta}) - \phi_\tau^\pm(\cdot, \bar{\eta})\|_{w^2} + o_p(1)) \left(\frac{1}{n^2} \sum_{i=1}^n w_i(x)^2 \right)^{1/2} \right) \\ &= o_p(1) \end{aligned}$$

By Stone (1977) Theorem 1, $\widetilde{\tau}^\pm(x) \xrightarrow{p} \mathbb{E}[\phi_\tau^\pm(Z, \bar{\eta}) | X = x]$.

By Theorem 1, $\pm (\mathbb{E}[\phi_\tau^\pm(Z, \bar{\eta}) | X = x] - \tau^\pm(x)) \geq 0$.

Therefore $\pm (\widehat{\tau}^\pm(x) - \tau^\pm(x)) \geq -o_p(1)$. □

E. Detailed Algorithm

We present a more detailed version of the B-Learner pseudocode in Algorithm 1.

F. Additional Experimental Details

The replication code for all simulations is distributed under an MIT license.

Algorithm 1 The B-Learner: Detailed

input Data $\{(X_i, A_i, Y_i) : i \in \{1, \dots, n\}\}$, folds $K \geq 2$, sensitivity parameter $\Lambda \geq 1$, nuisance estimators, regression learner $\widehat{\mathbb{E}}_n$.

- 1: **for** $k \in \{1, \dots, K\}$ **do**
- 2: Set $\mathcal{S}_k = \{(X_i, A_i, Y_i) : i \neq k - 1 \pmod{K}\}$
 Using \mathcal{S}_k :
- 3: Learn outcome model: $\widehat{\mu}^{(k)}(x, a) = \widehat{\mathbb{E}}[Y | X = x, A = a]$
- 4: Learn propensity model: $\widehat{e}^{(k)}(x) = \widehat{p}(A = 1 | X = x)$
- 5: Learn conditional outcome quantile models:
 $\widehat{q}_+^{(k)}(x, a) = \widehat{\inf}\{\beta : F(\beta | X = x, A = a) \geq \Lambda/(\Lambda + 1)\}$
 $\widehat{q}_-^{(k)}(x, a) = \widehat{\inf}\{\beta : F(\beta | X = x, A = a) \geq 1/(\Lambda + 1)\}$
- 6: Learn conditional value at risk models:
 $\widehat{\text{CVaR}}_{\pm}^{(k)}(x, a) = \widehat{q}_{\pm}^{(k)}(x, a) + (\Lambda + 1)\widehat{\mathbb{E}}[\{Y - \widehat{q}_{\pm}^{(k)}(x, a)\}_{\pm} | X = x, A = a]$
- 7: Set $\widehat{\rho}_{\pm}^{(k)}(x, a) = \Lambda^{-1}\widehat{\mu}^{(k)}(x, a) + (1 - \Lambda^{-1})\widehat{\text{CVaR}}_{\pm}^{(k)}(x, a)$
- 8: **for** $i = k - 1 \pmod{K}$ **do**
- 9: Set $R_{\pm, i} = \Lambda^{-1}Y_i + (1 - \Lambda^{-1})\left(\widehat{q}_{\pm}^{(k)}(X_i, A_i) + \frac{1}{1-\alpha}\{Y_i - \widehat{q}_{\pm}^{(k)}(X_i, A_i)\}_{\pm}\right)$
- 10: Set pseudo-outcomes for $Y^{\pm}(X, 1)$:
 $\widehat{\phi}_{1, i}^{\pm} = A_i Y_i + (1 - A_i)\widehat{\rho}_{\pm}^{(k)}(X_i, 1) + \frac{(1 - \widehat{e}^{(k)}(X_i))A_i}{\widehat{e}^{(k)}(X_i)} \cdot (R_{\pm, i} - \widehat{\rho}_{\pm}^{(k)}(X_i, 1))$
- 11: Set pseudo-outcomes for $Y^{\pm}(X, 0)$:
 $\widehat{\phi}_{0, i}^{\pm} = (1 - A_i)Y_i + A_i\widehat{\rho}_{\pm}^{(k)}(X_i, 1) + \frac{\widehat{e}^{(k)}(X_i)(1 - A_i)}{(1 - \widehat{e}^{(k)}(X_i))} \cdot (R_{\pm, i} - \widehat{\rho}_{\pm}^{(k)}(X_i, 0))$
- 12: Set pseudo-outcomes for CATE:
 $\widehat{\phi}_{\tau, i}^{\pm} = \widehat{\phi}_{1, i}^{\pm} - \widehat{\phi}_{0, i}^{\pm}$
- 13: **end for**
- 14: **end for**
- 15: Create datasets $\mathcal{T}^{\pm} = \{(X_i, \widehat{\phi}_{\tau, i}^{\pm})\}$
- 16: Learn upper- and lower- bound functions $\widehat{\tau}^{\pm}(x) = \widehat{\mathbb{E}}_n[\widehat{\phi}_{\tau}^{\pm} | X = x]$ from the datasets \mathcal{T}^{\pm}

output $\widehat{\tau}^{\pm}$

F.1. Simulated Data

The results in Section 8 were obtained using an Amazon Web Services instance with 32 vCPUs and 64 GiB of RAM. For the Random Forest (RF) models, we use the `RandomForestRegressor` model from `scikit-learn`. For Gaussian Kernels (GK), we use RBF (radial basis function) method from `scikit-learn`. Finally, for the Bayesian Neural Networks (NN) we use several functions from the `PyTorch` package. The quantile estimators use weights from the nuisance regressors when RFs or GKs are used or are calculated from the sampled outcome distributions when NNs are used. We include the hyperparameters for the different models used with the synthetic data in Table 2.

Table 2. Hyperparameters for model choices in synthetic data experiments.

Model	Hyperparameter	Value
Random Forest (<code>scikit-learn</code>)	<code>max_depth</code>	6
	<code>min_samples_leaf</code>	0.05
RBF (<code>scikit-learn</code>)	<code>length_scale</code>	$0.9 \times n^{-\frac{1}{4+d}}$
Neural Network (<code>PyTorch</code>)	<code>hidden units</code>	100
	<code>network depth</code>	4
	<code>negative slope</code>	0.3
	<code>dropout rate</code>	0.2
	<code>batch size</code>	50
	<code>learning rate</code>	$5e-4$

Table 3. **IHDP Covariates** Binary covariates $x_9 - x_{18}$ are attributes of the mother. Mother’s education level “College” indicated by covariates $x_{10} - x_{12}$ all zero. Site 8 indicated by covariates $x_{19} - x_{25}$ all zero. We show the frequency of occurrence for each binary covariate $p(x = 1)$, as well as the adjusted mutual information $I(x; t)$ between the binary covariate and the treatment variable.

Continuous		Binary		$I(x; t)$	$p(x = 1)$
Covariate	Description	Covariate	Description		
x_1	birth weight	x_7	child’s gender (female=1)	0.00	0.51
x_2	head circumference	x_8	is child a twin	0.00	0.09
x_3	number of weeks pre-term	x_9	married when child born	0.02	0.52
x_4	birth order	x_{10}	left High School	0.00	0.36
x_5	“neo-natal health index”	x_{11}	completed High School	0.00	0.27
x_6	mom’s age	x_{12}	some College	0.00	0.22
		x_{13}	child is first born	0.00	0.36
		x_{14}	smoked cigarettes when pregnant	0.01	0.48
		x_{15}	consumed alcohol when pregnant	0.00	0.14
		x_{16}	used drugs when pregnant	0.00	0.96
		x_{17}	worked during pregnancy	0.01	0.59
		x_{18}	received any prenatal care	0.01	0.96
		x_{19}	site 1	0.00	0.14
		x_{20}	site 2	0.01	0.14
		x_{21}	site 3	0.00	0.16
		x_{22}	site 4	0.01	0.08
		x_{23}	site 5	0.02	0.07
		x_{24}	site 6	0.01	0.13
		x_{25}	site 7	0.02	0.16

F.2. IHDP Dataset

We use [Jesson et al. \(2021\)](#)’s hidden confounding version of the Infant Health and Development Program (IHDP) that was introduced by [Hill \(2011\)](#). The data comes from an experiment that targeted “low-birth-weight, premature infants, and provided the treatment group with both intensive high-quality child care and home visits from a trained provider” ([Hill, 2011](#)). For the purpose of simulating an observational study, [Hill \(2011\)](#) generates simulated outcomes using the following features: measurements on the child—birth weight, head circumference, weeks born preterm, birth order, firstborn, neonatal health index, sex, twin status—as well as behaviors engaged in during pregnancy—smoked cigarettes, drank alcohol, took drugs—and measurements on the mother at the time she gave birth—age, marital status, educational attainment (did not graduate from high school, graduated from high school, attended some college but did not graduate, graduated from college), whether she worked during pregnancy, whether she received prenatal care, and the site (8 total) in which the family resided at the start of the intervention. A non-random portion of the treatment group, the children of non-white mothers, are excluded from the study in order to mimic confounding in an otherwise randomized trial. Covariates consist of 6 continuous variables and 19 binary variables. We use the covariate descriptions from [Jesson et al. \(2021\)](#) which we replicate in Table 3 for completeness. The dataset consists of 747 samples, of which 139 are in the treatment group.

[Jesson et al. \(2021\)](#) create the Hidden Confounding of IHDP by hiding the covariate x_9 from models during training, however, the causal model depends on it for the data generation. Following is the data generation process of the Hidden Confounding version of response surface B ([Hill, 2011](#)), we restate the data generation process from [Jesson et al. \(2021\)](#):

$$u := N_u, \tag{2a}$$

$$\mathbf{x} := N_{\mathbf{x}}, \tag{2b}$$

$$t := N_t, \tag{2c}$$

$$y := (t - 1)(\exp(\beta_{\mathbf{x}}(\mathbf{x} + \mathbf{w}) + \beta_u(u + 0.5)) + N_{Y0}) + t(\beta_{\mathbf{x}}\mathbf{x} + \beta_u u - \omega^s + N_{Y1}), \tag{2d}$$

where $(N_u, N_{\mathbf{x}}, N_t) \sim p_{\mathcal{D}}(x_9, \{x_1, \dots, x_8, x_{10}, \dots, x_{25}\}, t)$, $N_{Y0} \sim \mathcal{N}(0, 1)$, and $N_{Y1} \sim \mathcal{N}(0, 1)$. The coefficient β_u is randomly sampled from $(0.1, 0.2, 0.3, 0.4, 0.5)$ with probabilities $(0.2, 0.2, 0.2, 0.2, 0.2)$, $\beta_{\mathbf{x}}$ is a vector of randomly sampled values $(0.0, 0.1, 0.2, 0.3, 0.4)$ with probabilities $(0.6, 0.1, 0.1, 0.1, 0.1)$, w is a vector with all the coordinates

equals 0.5, where ω^s was chosen as in Hill (2011): “for the s^{th} simulation, it was chosen in the overlap setting, where we estimate the effect of the treatment on the treated, such that CATT equals 4; similarly it was chosen in the incomplete setting, where we estimate the effect of the treatment on the controls so that CATC equals 4”.

Following Jesson et al. (2021)’s Hidden Confounding experiment, we generate 400 realizations of the IHDP dataset, such that the seed for each realization is the corresponding index of the realization, where the indices are 0, 1, ..., 400. Each realization is split into training ($n = 470$), validation ($n = 202$), and test ($n = 75$) subsets. For the B-Learner with NNs, we use the same models and hyperparameters used by Quince in Jesson et al. (2021). For the B-Learner with RF base estimators, we use the `RandomForestRegressor` from `scikit-learn` and `ForestRegressor` from `econml.grf` where we control for forest growth only through the `max_depth` (= 6) and `min_samples_leaf` (= 0.01) parameters. As for *Kernel Sensitivity* and *Quince*, to replicate the results from Jesson et al. (2021), we use the same models and hyperparameters they used for the Hidden Confounding IHDP experiment. *Note*: we exclude 5 of the 400 IHDP trials from the original analysis due to poor data quality (e.g. low overlap) that affects the NN training. These issues seem to be mitigated by ensembling which is why they do not pose a problem for the experiments in Jesson et al. (2021). We will perform a comparison of the ensembled methods in a future iteration of this work.

E.3. 401(k) Eligibility Study

The dataset includes 9,915 observations with 9 covariates such as age, income, education, family size, marital status, IRA participation, etc. We describe the features of the 401(k) dataset in Table 4. In order to replicate the CATEs obtained by (Chernozhukov et al., 2018), we use the same models (`RandomForestRegressor` and `RandomForestClassifier` from `scikit-learn`) and hyperparameters (`n_estimators` = 100, `max_depth` = 7, `max_features` = 3, `min_samples_leaf` = 10) for our nuisance estimators and second stage models.

Table 4. Features of 401(k) dataset.

Name	Description	Type
age	age	continuous
inc	income	continuous
educ	years of completed education	continuous
fsize	family size	continuous
marr	marital status	binary
two_earn	whether dual-earning household	binary
db	defined benefit pension status	binary
pira	IRA participation	binary
hown	home ownership	binary
e401	401 (k) eligibility	binary
net_tfa	net financial assets	continuous