

EFFICIENT AND SHARP OFF-POLICY LEARNING UNDER UNOBSERVED CONFOUNDING

Konstantin Hess^{1,2,*}, Dennis Frauen^{1,2}, Valentyn Melnychuk^{1,2}, Stefan Feuerriegel^{1,2}

¹LMU Munich ²Munich Center for Machine Learning

*Corresponding author: k.hess@lmu.de

ABSTRACT

We develop a novel method for personalized off-policy learning in scenarios with unobserved confounding. Thereby, we address a key limitation of standard policy learning: standard policy learning assumes unconfoundedness, meaning that no unobserved factors influence both treatment assignment and outcomes. However, this assumption is often violated, because of which standard policy learning produces biased estimates and thus leads to policies that can be harmful. To address this limitation, we employ causal sensitivity analysis and derive a *semi-parametrically efficient estimator for a sharp bound on the value function under unobserved confounding*. Our estimator has three advantages: (1) Unlike existing works, our estimator avoids unstable minimax optimization based on inverse propensity weighted outcomes. (2) Our estimator is semi-parametrically efficient. (3) We prove that our estimator leads to the optimal confounding-robust policy. Finally, we extend our theory to the related task of policy improvement under unobserved confounding, i.e., when a baseline policy such as the standard of care is available. We show in experiments with synthetic and real-world data that our method outperforms simple plug-in approaches and existing baselines. Our method is highly relevant for decision-making where unobserved confounding can be problematic, such as in healthcare and public policy.

1 INTRODUCTION

Policy learning is crucial in many areas such as healthcare (Feuerriegel et al., 2024; Kraus et al., 2023), education (Chan, 2023), and public policy (Ladi & Tsarouhas, 2020). However, collecting data through randomized experiments is often either infeasible or unethical. Instead, methods are needed that use observational data to inform decision-making. Here, we focus on *off-policy learning* to optimize decision policies from observational data (Athey & Wager, 2021).

The reliability of standard off-policy learning is compromised when *unobserved confounding* is present (Kallus et al., 2019). Unobserved confounding arises when factors affect both treatment choices and outcomes but are not recorded (Pearl, 2009). For example, the race of a patient – unfortunately – affect the access to treatments (Obermeyer et al., 2019), yet race is typically not recorded in patient records. Hence, standard off-policy learning that relies on the assumption of no unobserved confounding will lead to *biased* estimates and may even generate *harmful* policies.

As a remedy, *confounding-robust policy learning* aims to find the optimal policy under worst-case unobserved confounding. This is typically achieved using the marginal sensitivity model (MSM) (Tan, 2006), a framework from causal sensitivity analysis that bounds the effect of unobserved confounding. However, the existing method for confounding-robust policy learning under the MSM (see (Kallus & Zhou, 2018a) for the conference paper and (Kallus & Zhou, 2021) for the journal version) has notable shortcomings. First, it must numerically optimize the worst-case effect on the regret function due to unobserved confounding. Such minimax optimization is based on inverse propensity weighted outcomes and hence unstable. Second, this method is statistically suboptimal: it lacks the property of semi-parametric efficiency and thus suffers from suboptimal variance properties.

In this paper, we address the above shortcomings by developing a novel *semi-parametrically efficient and sharp estimator for personalized off-policy learning under unobserved confounding*. Here,

semi-parametric efficiency means the unbiased estimator with the lowest possible variance. Our key novelties are the following: (i) We derive a *closed-form expression* for a *sharp bound on the value function* of a candidate policy under unobserved confounding.¹ As a result, we can thereby directly minimize our closed-form bound and, unlike existing works, avoid an unstable minimax optimization based on inverse propensity weighted outcomes. (ii) We propose an estimator that is semi-parametrically efficient. Hence, our estimator is the first to achieve the *lowest variance* among all unbiased estimators for our task.

Methodologically, we proceed as follows. We first derive a *sharp bound* on the value function for scenarios with unobserved confounding and, hence, avoid the unstable minimax optimization as in other methods. We then propose a novel *one-step bias-corrected estimator* to achieve semi-parametric efficiency and thus guarantee that our estimator has the lowest variance among all unbiased estimators. For this, we derive the corresponding efficient influence function of the sharp bound on the value function. We finally provide theoretical guarantees that minimizing our estimated sharp bound on the value function ensures that our method yields the *optimal* confounding-robust policy. Such guarantees are particularly crucial in high-stakes applications such as medicine or public policy, where unreliable policies can lead to harmful consequences.

Our work makes the following **contributions**²: (i) We propose a novel *efficient estimator for our sharp bound on the value function*. (ii) We derive an estimator for our bounds that is *semi-parametrically efficient*. (iii) We generalize our theoretical findings to the related task of *confounding-robust policy improvement*. (iv) Through extensive experiments using synthetic and real-world datasets, we show that our method consistently *outperforms* simple plug-in estimators and existing baselines.

2 RELATED WORK

We provide an overview of three literature streams particularly relevant to our work, namely, standard off-policy learning (i) with and (ii) without unobserved confounding as well as (iii) causal sensitivity analysis. We provide an extended related work in Appendix A (where we distinguish our work from other streams such as, e.g., unobserved confounding in reinforcement learning).

Off-policy learning under unconfoundedness:

Off-policy learning aims to optimize the policy value, which needs to be estimated from data. For this, there are three major approaches: (i) the direct method (DM) (Qian & Murphy, 2011) leverages estimates of the response functions; (ii) inverse propensity weighting (IPW) (Swaminathan & Joachims, 2015) re-weights the data such that in order to resemble samples under the evaluation policy; and (iii) the doubly robust method (DR) (Athey & Wager, 2021; Dudik et al., 2011). The latter is based on the efficient influence function of the policy value (Robins et al., 1994) and is efficient (Chernozhukov et al., 2018; van der Vaart, 1998).

Several works aim at improving the finite sample performance of these methods, for instance, via re-weighting (Kallus, 2018; 2021) or TMLE (Bibaut et al., 2019). Further, several methods have been proposed for off-policy learning in specific settings involving, for example, distributional robustness (Kallus et al., 2022), fairness (Frauen et al., 2024b), interpretability (Tschernutter et al., 2022), and continuous treatments (Kallus & Zhou, 2018b; Schweisthal et al., 2023). However, all of the works assume unconfoundedness and, therefore, do **not** account for unobserved confounding.

Off-policy learning under unobserved confounding: In scenarios with unobserved confounding, standard approaches for off-policy learning are biased (Kallus & Zhou, 2018a; 2021), which can lead to harmful decisions. The reason is that, under unobserved confounding, the policy value *cannot* be

	Robust under unobserved conf.	Sharp bounds	Discrete treatments	Efficient for policy learning
Oprescu et al. (2023)	✓	✓	✗	✗
Swaminathan & Joachims (2015)	✗	✗	✓	✗
Athey & Wager (2021); Dudik et al. (2011)	✗	✗	✓	✓
Kallus & Zhou (2018a; 2021)	✓	✗	✓	✗
Efficient & sharp (ours)	✓	✓	✓	✓

Table 1: Overview of related methods. Oprescu et al. (2023) is designed for a different task, and standard policy learners ignore the issue of unobserved confounding. Only Kallus & Zhou (2018a; 2021) deals with our setting, but provides neither sharp bounds nor an efficient estimator. Only our work can deal with unobserved confounding, discrete treatments, provides sharp bounds, and an efficient estimator.

¹We use the term “sharp” as in earlier work from causal sensitivity analysis (Frauen et al., 2023): a valid upper (lower) bound of a causal quantity is *sharp* if there does not exist another valid bound that is strictly smaller (larger).

²Code is available at https://github.com/konstantinhess/Efficient_sharp_policy_learning.

identified from observational data. As a remedy, previous works leverage causal sensitivity analysis or related methods to obtain bounds on the unidentified policy value (Bellot & Chiappa, 2024; Guerdan et al., 2024; Huang & Wu, 2024; Joshi et al., 2024; Namkoong et al., 2020; Zhang & Bareinboim, 2024), which can then be used to learn an optimal worst-case policy. Optimizing such bounds is often termed “confounding-robust policy learning”. However, these works do *not* consider sharp bounds under unobserved confounding and do *not* provide semi-parametrically efficient estimators.

Closest to our work is (Kallus & Zhou, 2018a) with an extended version published in (Kallus & Zhou, 2021). Therein, the authors propose a method for confounding-robust policy improvement, yet with two notable shortcomings: (i) it is *not* based on closed-form solutions for the bounds, and (ii) it is *not* based on a semi-parametrically efficient estimator of these bounds. Therefore, Kallus & Zhou (2018a; 2021) require solving a minimax optimization problem that relies in inverse propensity weighted outcomes, which is *unstable*. Further, their estimator is suboptimal because it fails to achieve semi-parametric efficiency, meaning it does *not* achieve the lowest-possible variance among all unbiased estimators. We provide a thorough discussion on the work by Kallus & Zhou (2018a; 2021) and a detailed comparison to our method in Supplement B.

Causal sensitivity analysis: Causal sensitivity analysis (Cornfield et al., 1959) allows practitioners to account for unobserved confounding by using so-called sensitivity models (Jin et al., 2022; Rosenbaum, 1987), which incorporate domain knowledge on the strength of unobserved confounding. As a result, the sensitivity model allows to obtain bounds on a causal quantity of interest, which, if the sensitivity model is correctly specified, can then be used for consequential decision-making (Jesson et al., 2021).

A prominent sensitivity model is the MSM (Tan, 2006). The MSM gained popularity in recent years and, for instance, was used to obtain bounds on the conditional average treatment effect (CATE) through machine learning (Jesson et al., 2021; Kallus et al., 2019; Yin et al., 2022). Works like Aronow & Lee (2013); Miratrix & Wager (2018); Zhao et al. (2019) target mean effects under the MSM, rely on Hájek-style estimators which have finite-sample bias, and are not designed for policy learning. Yadlowsky et al. (2022) study ATE bounds, but their framework is limited to *binary* treatments. Finally, Bruns-Smith & Zhou consider *dynamic* policies in Markov decision processes and must estimate ratios of transition probabilities, which yields a fundamentally different estimator that is not sharp for action spaces with more than two actions. Only recently, sharp bounds on the CATE have been derived (Bonvini et al., 2022; Dorn & Guo, 2022; Frauen et al., 2024a; 2023; Jin et al., 2023). Other works have considered the estimation of such bounds (Dorn et al., 2024; Oprescu et al., 2023), or instruments for partial identification Schweisthal et al. (2025). However, these works only consider causal sensitivity analysis for CATE but **not** policy learning. Further, their works are limited to **binary** treatments. In contrast, our method can handle *discrete* treatments, and, therefore, requires entirely different influence functions.

Research gap: To the best of our knowledge, we are the first to derive a *semi-parametrically efficient estimator for a sharp bound on the value function using the MSM*. Thereby, we enable optimal confounding-robust off-policy learning.

3 PROBLEM SETUP

Data: Let $Y \in \mathcal{Y} \subset \mathbb{R}$ be our outcome of interest, such as the health condition of a patient. We follow the convention that w.l.o.g. *lower* values correspond to *better* outcomes (Kallus & Zhou, 2018a). Further, let $X \in \mathcal{X} \subset \mathbb{R}^{d_x}$ denote covariates that contain additional information, such as age, gender, or disease-related information. Finally, let $A \in \mathcal{A} = \{0, 1, \dots, d_a - 1\}$ be the assigned treatment (or action). Note that we do not restrict our setting to binary treatments but allow for arbitrary, discrete treatments. For the product space, we use $\mathcal{D} = \mathcal{Y} \times \mathcal{X} \times \mathcal{A}$. In the following, we assume that we have access to an observational dataset $\mathcal{D}_n = \{(Y_i, X_i, A_i)\}_{i=1}^n$ that consists of n i.i.d. copies of $(Y, X, A) \in \mathcal{D}$.

Policy value: Policy learning aims to find the best policy for assigning treatments, given covariates. Formally, a *policy* $\pi(a | x)$ is a conditional probability mass function $\pi : \mathcal{A} \times \mathcal{X} \rightarrow [0, 1]$ with $\sum_{a \in \mathcal{A}} \pi(a | x) = 1$, corresponding to the probability of receiving treatment $A = a$, given covariates




Figure 1: We can only block backdoor paths for observed confounders X . Hence, under unobserved confounding U , we cannot point-identify $V(\pi)$.

$X = x$. The *value* $V(\pi)$ of a policy is defined as $V(\pi) = \mathbb{E}\left[\sum_{a \in \mathcal{A}} \pi(a | X) Y[a]\right]$, where $Y[a]$ denotes the potential outcome for Y when intervening on treatment $A = a$ (Neyman, 1923; Rubin, 1978). Hence, the policy value $V(\pi)$ is the expected average potential outcome when adhering to the policy π .

Standard off-policy learning: Off-policy learning aims to find a policy π that has the best policy value among all $\pi \in \Pi$ for some policy class Π . Of note, it is standard in the literature (Frauen et al., 2024b; Hatt et al., 2022; Kallus & Zhou, 2018a) to restrict the analysis to policy classes Π with finite complexity such as neural networks.

The value function is identifiable under the following three assumptions (Rubin, 1978): (i) *Consistency*: $Y[A] = Y$; (ii) *Strong overlap*: There exists $c \in (0, 1)$ such that $c \leq p(A = a | X = x) \leq 1 - c \forall a \in \mathcal{A}, x \in \mathcal{X}$; (iii) *Unconfoundedness*: $Y[a] \perp\!\!\!\perp A | X \forall a \in \mathcal{A}$. Then, the policy value is identified from the observational data via $V(\pi) = \mathbb{E}\left[\sum_{a \in \mathcal{A}} \pi(a | X) Q(a, X)\right]$, where $Q(a, x) = \mathbb{E}[Y | X = x, A = a]$ is the conditional average potential outcome function.

The optimal policy can then be learned via

$$\pi_{\text{standard}}^* = \arg \min_{\pi \in \Pi} \hat{V}(\pi), \quad (1)$$

where $\hat{V}(\pi)$ is an estimator of the identified policy value. Recall that we follow the convention in (Kallus & Zhou, 2018a; 2021) that lower Y are better, so we aim to *minimize* the value function.

Allowing for unobserved confounding: The assumption of (iii) *unconfoundedness* is problematic and often unrealistic (Hemkens et al., 2018): Unconfoundedness requires that the observed covariates X capture *all* factors that affect both treatment choice and outcome. In this work, we do **not** rely on the *unconfoundedness* assumption, which is restrictive and oftentimes unrealistic. Instead, we allow for *unobserved confounding*, which we denote by a random variable $U \in \mathcal{U} \subset \mathbb{R}$ (see Figure 2).

Importantly, under unobserved confounding, we cannot point-identify the value function $V(\pi)$. Instead, we aim to *partially* identify the value function $V(\pi)$ by leveraging causal sensitivity analysis. Specifically, we adopt the MSM (Tan, 2006) to bound the ratio between the *nominal propensity score*

$$e(a, x) = p(A = a | X = x), \quad (2)$$

which can be estimated from data \mathcal{D}_n and the *true propensity score*

$$e(a, x, u) = p(A = a | X = x, U = u), \quad (3)$$

which is fundamentally unobserved. Formally, the MSM assumes

$$\Gamma^{-1} \leq \frac{e(a, x)}{1 - e(a, x)} \frac{1 - e(a, x, u)}{e(a, x, u)} \leq \Gamma \quad (4)$$

for some $\Gamma \geq 1$ that can be chosen by domain knowledge (Frauen et al., 2023; Kallus et al., 2019) or data-driven heuristics (Hatt et al., 2022) (see Supplement E).

Intuitively, Γ close to 1 implies that the impact of unobserved variables U on the treatment decision is small, whereas a large Γ means that observed variables X do not contain sufficient information to fully capture the treatment decision. In particular, $\Gamma = 1$ implies that the true propensity score coincides with the nominal propensity score. Hence, there is no unobserved confounding and our scenario simplifies to the naïve unconfoundedness setting. Conversely, if we let $\Gamma > 1$, the true and the nominal propensity scores differ, and, therefore, we account for additional unobserved confounding.

Formally, the marginal sensitivity model gives rise to a set of distributions $\mathcal{P}(\Gamma)$ over $\mathcal{D} \times \mathcal{U}$ that are compatible with the constraints in Equation 4. This set is defined as

$$\mathcal{P}(\Gamma) = \left\{ \tilde{p} \in \mathcal{P}(\mathcal{D} \times \mathcal{U}) : \int_{\mathcal{U}} \tilde{p}(d, u) du = p(d) \forall d \in \mathcal{D}, \Gamma^{-1} \leq \frac{\tilde{e}(A, X)}{1 - \tilde{e}(A, X)} \frac{1 - \tilde{e}(A, X, U)}{\tilde{e}(A, X, U)} \leq \Gamma \right\}, \quad (5)$$

where $\mathcal{P}(\mathcal{D} \times \mathcal{U})$ is the set of all possible joint distributions of the observables and the unobserved confounders, and where the nominal propensity score $\tilde{e}(A, X)$ and the true propensity score $\tilde{e}(A, X, U)$ result from \tilde{p} as in Equation 2 and Equation 3, respectively.

Proposition 4.1. Let $Q^{+,*}(a, x) = \sup_{\tilde{p} \in \mathcal{P}(\Gamma)} Q(a, x)$ and $Q^{-,*}(a, x) = \inf_{\tilde{p} \in \mathcal{P}(\Gamma)} Q(a, x)$ be the sharp upper and lower bound for the conditional average potential outcome, respectively, given our sensitivity constraints $\mathcal{P}(\Gamma)$. Then, the sharp upper bound $\sup_{\tilde{p} \in \mathcal{P}(\Gamma)} V(\pi) = V^{+,*}(\pi)$ and the sharp lower bound $\inf_{\tilde{p} \in \mathcal{P}(\Gamma)} V(\pi) = V^{-,*}(\pi)$ for the value function $V(\pi)$ are given by

$$V^{\pm,*}(\pi) = \int_{\mathcal{X}} \sum_a Q^{\pm,*}(a, x) \pi(a | x) dp(x). \quad (8)$$

Proof. See Supplement F.1. \square

Our above derivation of the closed-form solution has a crucial advantage over existing works (Kallus & Zhou, 2018a; 2021): we avoid an unstable minimax optimization that is based on inverse propensity weighted outcomes, and, instead, we can directly work with $V^{+,*}(\pi)$, which simplifies Equation 6 to $\pi^* = \arg \min_{\pi \in \Pi} V^{+,*}(\pi)$. As a result, we have reduced the original minimax problem to a much simpler *minimization task*.

4.2 SEMI-PARAMETRICALLY EFFICIENT ESTIMATOR FOR THE SHARP UPPER BOUND

In this section, we derive a semi-parametrically efficient estimator for our sharp upper bound $V^{+,*}(\pi)$ of the value function $V(\pi)$. semi-parametrically efficient estimators are desirable because they achieve the *lowest possible variance among all unbiased estimators* (Hines et al., 2022; Kennedy, 2022).

In order to derive such an estimator of $V^{+,*}(\pi)$, we first need to decompose the estimand $Q^{\pm,*}(a, x)$ in Proposition 4.1.

Definition 4.2 ((Dorn & Guo, 2022; Frauen et al., 2023)). Sharp bounds $Q^{\pm,*}(a, x)$ of the conditional average potential outcome $Q(a, x)$ function are given by

$$Q^{\pm,*}(a, x) = c^{\mp}(a, x) \underline{\mu}^{\pm}(a, x) + c^{\pm}(a, x) \bar{\mu}^{\pm}(a, x), \quad (9)$$

where we let

$$c^{\pm}(a, x) = b^{\pm} e(a, x) + \Gamma^{\pm 1}, \quad b^{\pm} = (1 - \Gamma^{\pm 1}) \quad (10)$$

and

$$\underline{\mu}^{\pm}(a, x) = \mathbb{E}[Y \underline{\Delta}^{\pm}(Y, A, X) | X = x, A = a], \quad \bar{\mu}^{\pm}(a, x) = \mathbb{E}[Y \bar{\Delta}^{\pm}(Y, A, X) | X = x, A = a] \quad (11)$$

with

$$\underline{\Delta}^{\pm}(y, a, x) = \mathbb{1}_{\{y \leq F_{x,a}^{-1}(\alpha^{\pm})\}}, \quad \bar{\Delta}^{\pm}(y, a, x) = \mathbb{1}_{\{y \geq F_{x,a}^{-1}(\alpha^{\pm})\}}, \quad (12)$$

where $\alpha^+ = \Gamma/(1 + \Gamma)$ and $\alpha^- = 1/(1 + \Gamma)$, and where $F_{x,a}^{-1}(q)$ is the conditional quantile function

$$F_{x,a}^{-1}(q) = \inf\{y \in \mathcal{Y} : p(Y \leq y | X = x, A = a) \geq q\}. \quad (13)$$

In order to achieve semi-parametric efficiency for the sharp upper bound $V^{+,*}(\pi)$, we need to carefully take into account the nuisance functions in Equation 8 and Equation 9, respectively. That is, the key difficulty lies in that $V^{+,*}(\pi)$ depends on several nuisance functions

$$\eta = \{e(a, x), F_{a,x}^{-1}(\alpha^{\pm}), \bar{\mu}^{\pm}(a, x), \underline{\mu}^{\pm}(a, x)\}. \quad (14)$$

If we followed a naïve plug-in approach (i.e., if we estimated $\hat{\eta}$ from data \mathcal{D}_n and plugged them into Equation 8 and thus Equation 9), our final estimator $\hat{V}^{+,*}(\pi)$ would suffer from **first-order bias** due to estimation errors in the nuisance functions. As a remedy, we present a *one-step bias-corrected* estimator. That is, we estimate the first-order bias and subtract it from our plug-in estimate (Hess et al., 2026; Frauen et al., 2025a; Kennedy, 2022; van der Vaart, 1998).

Theorem 4.3. Let $\mathbb{E}[|Y|] < \infty$ and $p(y | x, a)$ admit a continuous density bounded away from zero in a neighborhood of $F_{x,a}^{-1}(\alpha^+)$. This holds, for example, if the density $p(y | x, a)$ is strictly positive and continuous. Then, an estimator for the sharp upper bound of the value function is given by

$$\begin{aligned} \hat{V}^{+,*}(\pi) = & \mathbb{P}_n \left\{ \sum_a \pi_{a,X} \left[\hat{Q}_{a,X}^{+,*} - \hat{e}_{a,X} \left(b^- \hat{\underline{\mu}}_{a,X}^+ + b^+ \hat{\bar{\mu}}_{a,X}^+ \right) \right] + \pi_{A,X} \left(b^- \hat{\underline{\mu}}_{A,X}^+ + b^+ \hat{\bar{\mu}}_{A,X}^+ \right) \right. \\ & + \frac{\pi_{A,X}}{\hat{e}_{A,X}} \left[\left(\hat{e}_{A,X}^- - \hat{e}_{A,X}^+ \right) \left(\hat{F}_{X,A}^{-1}(\alpha^+) \left(\hat{\Delta}_{Y,A,X}^+ - \alpha^+ \right) \right) \right. \\ & \left. \left. + \hat{e}_{A,X}^- \left(Y \hat{\Delta}_{Y,A,X}^+ - \hat{\underline{\mu}}_{A,X}^+ \right) + \hat{e}_{A,X}^+ \left(Y \hat{\Delta}_{Y,A,X}^+ - \hat{\bar{\mu}}_{A,X}^+ \right) \right] \right\}. \quad (15) \end{aligned}$$

If, on top, $\mathbb{E}[|Y|^2] < \infty$, the above estimator is **semi-parametrically efficient**.

Proof. See Supplement F.2. □

We now have a *semi-parametrically efficient estimator* for the sharp upper bound of the value function under unobserved confounding. Algorithm 1 presents a flexible procedure to learn confounding-robust policies for parametric policy classes Π_θ (e.g., neural networks). For simplicity, Algorithm 1 only uses a single split, but we emphasize that cross-fitting can be used to increase sample efficiency.

4.3 LEARNING GUARANTEES

In this section, we provide asymptotic learning guarantees in the form of generalization bounds when learning the confounding-robust policy π via our Algorithm 1. Of note, it is *not* obvious that minimizing the *estimated* sharp upper bound $\hat{V}^{+,*}(\pi)$ provides a meaningful, confounding-robust policy π^* . Hence, we provide learning guarantees where we show that, with high probability, *minimizing our estimated sharp upper bound yields the optimal policy*.

For this, we show that minimizing the *estimated* sharp upper bound $\hat{V}^{+,*}(\pi)$ with respect to π indeed minimizes the true, unknown value function $V(\pi)$ on population level. Fortunately, our method only requires one additional assumption, namely, boundedness of the outcome $|Y| \leq C_y$. This is a very mild restriction and reasonable in practice.

We express the flexibility of our policy class Π in terms of the Rademacher complexity $\mathcal{R}_n(\pi)$, which is a common choice in the literature (Athey & Wager, 2021; Frauen et al., 2024b; Hatt et al., 2022; Kallus & Zhou, 2018a). For example, the class of neural networks with bounded size and smoothness, e.g., with fixed depth, bounded width, and weights controlled by simple norm constraints (so the networks are Lipschitz and outputs bounded), has Rademacher complexity $\mathcal{R}_n(\Pi) \in \mathcal{O}(n^{-1/2})$.

Theorem 4.4. *Assume $|Y| \leq C_y$. Let $C_v := 2C_y(1 + \Gamma^{-1} + \Gamma)$ and let $R_n(\Pi)$ denote the empirical Rademacher complexity of the policy class Π . Then, with probability at least $1 - \delta$,*

$$\sup_{\pi \in \Pi} \left\{ V(\pi) - \hat{V}^{+,*}(\pi) \right\} \leq 2C_v \left(R_n(\Pi) + \frac{5}{2} \sqrt{\frac{1}{2n} \log \frac{2}{\delta}} \right). \quad (16)$$

Equivalently, on the same high-probability event, simultaneously for all $\pi \in \Pi$,

$$V(\pi) \leq \hat{V}^{+,*}(\pi) + 2C_v \left(R_n(\Pi) + \frac{5}{2} \sqrt{\frac{1}{2n} \log \frac{2}{\delta}} \right). \quad (17)$$

Proof. See Supplement F.3. □

The above Theorem 4.4 has the following implication: given our sensitivity constraints $\mathcal{P}(\Gamma)$, our estimated sharp upper bound $\hat{V}^{+,*}(\pi)$ correctly bounds the true, unknown value function $V(\pi)$ on population level with high probability. Therefore, given sufficient data \mathcal{D}_n , minimizing $\hat{V}^{+,*}(\pi)$ with respect to π also minimizes $V(\pi)$ and, hence, yields the optimal π^* .

In sum, we have derived (i) a novel *sharp upper bound of the value function*, which circumvents unstable minimax optimization based on inverse propensity weighted outcomes. Further, we have proposed (ii) an estimator for this bound that is *semi-parametrically efficient*, i.e., an unbiased estimator with the lowest possible variance. Finally, we have derived (iii) *learning guarantees*, which show that minimizing our estimated bound via Algorithm 1 indeed optimizes the true, unknown population value.

Algorithm 1 Confounding-robust policy learning.

Input: Data $\mathcal{D}_n = \{(Y_i, A_i, X_i)\}_{i=1}^n$, sensitivity parameter $\Gamma \geq 1$, sample split $\rho \in (0, 1)$, learning rate λ , parametric policy class Π_θ , training iterations K

Output: $\hat{V}^{+,*}(\pi)$

- 1: Perform sample split $\mathcal{D}_{\lceil \rho n \rceil}^\eta, \mathcal{D}_{\lfloor (1-\rho)n \rfloor}^{V^{+,*}}$
(Employ cross-fitting to improve sample efficiency.)
 - 2: Estimate nuisance functions $\hat{\eta}$ on $\mathcal{D}_{\lceil \rho n \rceil}^\eta$
 - 3: Evaluate $\hat{\eta}$ on $\mathcal{D}_{\lfloor (1-\rho)n \rfloor}^{V^{+,*}}$
 - 4: Initialize policy $\pi_\theta^{(0)} \in \Pi_\theta$
 - 5: **for** $k = 0$ to $K - 1$ **do**
 - 6: Estimate $V^{+,*}(\pi_\theta^{(k)})$ as in Equation 38 (using evaluated $\hat{\eta}$)
 - 7: Update policy parameters:
 - 8: $\theta^{(k+1)} \leftarrow \theta^{(k)} - \lambda \nabla_{\theta} V^{+,*}(\pi_\theta^{(k)})$
 - 9: **end for**
 - 10: **Return:** Robust policy $\pi^* \leftarrow \pi_{\theta}^{(K)}$
-

4.4 EXTENSION TO POLICY IMPROVEMENT

Our main results from above focus on optimizing the value function $V(\pi)$, which is common in practice (e.g., Dudik et al., 2011; Hatt et al., 2022). However, in some scenarios, an established baseline policy π_0 may be available; then, one may aim to make a small relative improvement yet with certain guarantees. This setting is commonly termed as *policy improvement* (Kallus & Zhou, 2018a; 2021; Laroche et al., 2019; Thomas et al., 2015).

Hence, we no longer aim to minimize bounds on the value function $V(\pi)$ but, instead, bounds on the regret of a candidate policy against a baseline policy (Kallus & Zhou, 2018a). Specifically, we can define the *regret* of policy π over baseline π_0 as $R_{\pi_0}(\pi) = V(\pi) - V(\pi_0)$. Hence, a negative regret implies that policy π improves upon π_0 . Importantly, the optimal confounding-robust policy π^* in Equation 6 can also be defined as the policy π that achieves the best relative improvement over baseline π_0 in the worst-case scenario, that is,

$$\pi^* = \arg \min_{\pi \in \Pi} \sup_{\tilde{p} \in \mathcal{P}(\Gamma)} R_{\pi_0}(\pi). \quad (18)$$

This definition is *equivalent* to Equation 6. Nevertheless, the above objective may be preferred in practice when aiming at policy improvement.

We now show in the following three corollaries that our results directly generalize to policy improvement. First, we provide a closed-form solution for an upper bound of the regret function $R_{\pi_0}(\pi)$, given our sensitivity constraints $\mathcal{P}(\Gamma)$.

Corollary 4.5. *An upper bound for the regret function $R_{\pi_0}(\pi)$ is given by $R_{\pi}^+(\pi) = \int_{\mathcal{X}} \sum_a \left(Q^{+,*}(a, x)\pi(a | x) - Q^{-,*}(a, x)\pi_0(a | x) \right) dp(x)$.*

Proof. See Supplement F.4. □

Next, we derive a semi-parametrically efficient, one-step bias-corrected estimator, which is based on the efficient influence function.

Corollary 4.6. *A semi-parametrically efficient estimator for the upper regret bound is given by*

$$\begin{aligned} \hat{R}_{\pi_0}^+(\pi) = & \sum_{\pm \in \{-, +\}} \pm \mathbb{P}_n \left\{ \sum_a \pi_{a,X}^{\pm} \left[\hat{Q}_{a,X}^{\pm,*} - \hat{e}_{a,X} \left(b^{\mp} \hat{\underline{\mu}}_{a,X}^{\pm} + b^{\pm} \hat{\bar{\mu}}_{a,X}^{\pm} \right) \right] \right. \\ & + \pi_{A,X}^{\pm} \left(b^{\mp} \hat{\underline{\mu}}_{A,X}^{\pm} + b^{\pm} \hat{\bar{\mu}}_{A,X}^{\pm} \right) + \frac{\pi_{A,X}^{\pm}}{\hat{e}_{A,X}^{\pm}} \left[\left(\hat{c}_{A,X}^{\mp} - \hat{c}_{A,X}^{\pm} \right) \left(\hat{F}_{X,A}^{-1}(\alpha^{\pm}) \left(\hat{\Delta}_{Y,A,X}^{\pm} - \alpha^{\pm} \right) \right) \right. \\ & \left. \left. + \hat{c}_{A,X}^{\mp} \left(Y \hat{\Delta}_{Y,A,X}^{\pm} - \hat{\underline{\mu}}_{A,X}^{\pm} \right) + \hat{c}_{A,X}^{\pm} \left(Y \hat{\Delta}_{Y,A,X}^{\pm} - \hat{\bar{\mu}}_{A,X}^{\pm} \right) \right] \right\}, \end{aligned}$$

where we let $\pi^+ = \pi$ and $\pi^- = \pi_0$ for readability.

Proof. See Supplement F.5. □

Finally, we provide improvement guarantees: given a baseline policy π_0 (e.g., the standard of care), if the empirical estimator $\hat{R}_{\pi_0}(\pi)^+$ is *negative*, which we can check by evaluating it, we are *guaranteed* that π improves upon π_0 and introduces *no harm*.

Corollary 4.7. *Under the same assumption as in Theorem 4.4, for any policy $\pi \in \Pi$ and baseline policy $\pi_0 \in \Pi$, it holds, with probability $1 - \delta$, that $R_{\pi_0}(\pi) \leq \hat{R}_{\pi_0}^+(\pi) + 4C_v \left(\mathcal{R}_n(\Pi) + \frac{5}{2} \sqrt{\frac{1}{2n} \log \left(\frac{2}{\delta} \right)} \right)$.*

Proof. See Supplement F.6. □

	$\Gamma^* = 2$	$\Gamma^* = 4$	$\Gamma^* = 6$	$\Gamma^* = 8$	$\Gamma^* = 10$	$\Gamma^* = 12$	$\Gamma^* = 14$	$\Gamma^* = 16$
Standard IPW estimator	-1.31 ± 0.02	-0.60 ± 0.15	-0.09 ± 0.01	-0.07 ± 0.01	-0.06 ± 0.01	-0.06 ± 0.01	-0.05 ± 0.01	-0.03 ± 0.01
Standard DR estimator	-1.30 ± 0.04	-0.71 ± 0.02	-0.18 ± 0.13	-0.07 ± 0.01	-0.07 ± 0.01	-0.06 ± 0.01	-0.05 ± 0.01	-0.04 ± 0.01
Kallus & Zhou (2018a; 2021)	-1.21 ± 0.10	-0.70 ± 0.06	-0.40 ± 0.06	-0.22 ± 0.04	-0.16 ± 0.02	-0.14 ± 0.02	-0.10 ± 0.01	-0.08 ± 0.01
Efficient + sharp estimator (ours)	-1.12 ± 0.08	-1.00 ± 0.08	-0.89 ± 0.13	-0.66 ± 0.14	-0.64 ± 0.14	-0.58 ± 0.17	-0.50 ± 0.20	-0.30 ± 0.22
Absolute improvement	+0.19	-0.29	-0.49	-0.44	-0.48	-0.44	-0.40	-0.22

Table 2: **Varying confounding strength.** We vary the confounding parameter Γ^* in the DGP along with the sensitivity parameter Γ in both our efficient estimator and the baseline (Kallus & Zhou, 2018a; 2021). Then, we report the regret over a randomized policy (*lower values are better*). As confounding increases, our estimator is the only method that is robust and thus performs best.

5 EXPERIMENTS

In the following, we evaluate the performance of our method against: (1) the minimax optimization approach by Kallus & Zhou (2018a; 2021) and standard methods for policy learning, namely, (2) the IPW estimator (Swaminathan & Joachims, 2015) and (3) the DR estimator (Athey & Wager, 2021; Dudik et al., 2011). Importantly, the approach by Kallus & Zhou (2018a; 2021) is the **only** baseline that can deal with confounding-robust policy learning with the MSM and thus the **only** baseline for our task. To ensure a *fair comparison*, we use the same neural instantiations for all models in terms of (i) the nuisance functions $\hat{\eta}$ and (ii) the policy π_θ (see Supp. G). All results are averaged over 10 seeds.

• **Synthetic Data:** As is standard in causal inference literature (Hess & Feuerriegel, 2025; Hess et al., 2024; Kallus et al., 2019), we evaluate our method on synthetic data in order to have access to ground-truth counterfactuals. Here, we use an established data-generating process from the literature (Kallus et al., 2019): First, we simulate observed confounders $X \sim \text{Unif}[-2, 2]$ and unobserved confounders $U \sim \text{Ber}(1/2)$. The potential outcomes $Y[a]$ are then given by $Y[a] = (2a-1)X + (2a-1) - 2 \sin(2(2a-1)X) - 2(2U-1)(1+0.5X) + \varepsilon$, where $\varepsilon \sim \mathcal{N}(0, 1)$ is random noise. Further, we assume a binary treatment, i.e., $d_a = 2$. For this, we first fix a ground-truth Γ^* . Then, we let the *true propensity score* be given by $e(1, x, u) = \frac{u}{\rho(x; 1/\Gamma^*)} + \frac{1-u}{\rho(x; \Gamma^*)}$, where $\rho(x; \gamma) = 1 + (1/e(1, x) - 1)\gamma$, and $e(1, x) = \sigma(0.75x + 0.5)$ is the nominal propensity score.

Varying confounding strength: First, we demonstrate the performance of our method for increasing levels of unobserved confounding. For this, we increase the confounding parameter Γ^* in the data-generating process. We compare the regret of each method over a randomized baseline policy. In our method and (Kallus & Zhou, 2018a; 2021), we set the sensitivity parameter Γ equal to Γ^* .

Our results are shown in **Table 2**: (i) As expected, the standard methods for off-policy learning (i.e., a standard IPW estimator and a standard DR estimator) perform well for zero to very low levels of confounding. However, the standard methods are *biased* and thus become ineffective for $\Gamma^* > 1$. (ii) The method by Kallus & Zhou (2018a; 2021) performs well under low levels of confounding. Yet, the performance quickly deteriorates. (iii) Our proposed method performs clearly best for increasing Γ^* . Here, our method achieves a *relative performance gain by up to a factor of 4*.

Robustness analysis: Next, we show that our method is robust to mis-specification of the sensitivity parameter Γ . We thus fix the confounding strength to $\Gamma^* = 7$ in the DGP. We increase Γ from 1 (which corresponds to unconfoundedness) up to 100 (which mirrors almost assumption-free bounds). We again compute the regret of our learned policy over a randomized baseline policy to showcase the improvement. **Figure 3** shows that our approach yields *robust results even for mis-specified Γ* (i.e., when $\Gamma \neq 7$). In contrast, the method by Kallus & Zhou (2018a; 2021) improves for mis-specified values $\Gamma < \Gamma^*$, but performance *quickly deteriorates* for larger Γ , which renders it *unreliable*. This behavior is expected, as their method only minimizes a bound on the regret over a baseline policy, and quickly reverts to the baseline when the specified Γ increases (see Supplement B for a detailed




Figure 3: **Robustness analysis.** We set $\Gamma^* = 7$ in the data-generating process but use mis-specified sensitivity parameters Γ in our estimator (i.e., $\Gamma = 7$ is correctly specified, while $\Gamma \neq 7$ is mis-specified). We report the regret over a randomized policy (*lower values are better*). Our estimator significantly improves upon the standard DR estimator, even for a completely mis-specified Γ .

discussion). Further, even under the (almost) no assumptions constraint of $\Gamma = 100$, our method provides significant improvements over the biased DR estimator.

Semi-parametrically efficient estimation: Finally, we show the benefits of our efficient estimation strategy over simple plug-in estimators. A semi-parametrically efficient estimator is the unbiased estimator with the lowest variance (Kennedy, 2022; van der Vaart, 1998). Hence, policies based on efficiently estimated bounds are learned better in low sample settings than those based on plug-in approaches. Therefore, we report the performance when we vary the number of training samples \mathcal{D}_n . Here, we compare our method against a naïve plug-in estimator of our sharp bounds based on Proposition 4.1. We again report the regret over a randomized baseline policy. **Figure 4** shows that our method performs better in low sample settings, and achieves larger performance gains when increasing the sample size.

• **Real-world medical data:** We evaluate our method on a real-world medical case study. For this, we use data from the International Stroke Trial (Sandercock et al., 2011), which is a randomized control trial (RCT) that examines the outcomes for early administration of aspirin, heparin, a combination of both, or none on acute ischaemic stroke. Importantly, this means that there are *four different* treatments available. The advantage of an RCT over observational data is that we can estimate the ground truth value function without bias, as the true propensity score is known. Our aim is to find the optimal treatment strategy based on patient covariates in order to prolong the *time-to-death* (TD) outcome variable (in days).

For this, we artificially introduce unobserved confounding as follows: In the training dataset, we randomly drop 60% of the untreated patients whose diastolic blood pressure is larger than the average, as well as 60% of the patients who received aspirin and whose blood pressure is lower than average. Then, we remove the diastolic blood pressure variable. Thereby, we introduce *unobserved* confounding in the training dataset.

Results: We report the estimated improvement of the TD outcome of all methods over the randomized policy in Figure 5. We report regret relative to a fully randomized (RCT) policy because, under unobserved confounding, it is the only baseline whose value remains identified and therefore provides a neutral, assumption-free reference point for evaluating policy improvement. Here, we vary the sensitivity parameter for both the method by Kallus & Zhou (2018a; 2021) and our method. Our method performs best at $\Gamma = 24$. Further, our method has the overall best treatment strategy. Only for very small Γ , the performance of our method deteriorates – this is expected, as it does not guard against unobserved confounding and is therefore biased, and in contrast to standard methods, requires estimation of more complex nuisance functions. The baseline performance is generally poor. Specifically, we observe a similar pattern for the method by Kallus & Zhou (2018a; 2021) as in the robustness analysis; it only optimizes against a baseline policy, and quickly reverts to it as Γ increases (see Supplement B). In contrast, our method is highly stable for appropriate parameterizations of Γ , which confirms the effectiveness of our method, and shows its applicability to medical scenarios.

Conclusion: We develop a novel semi-parametrically efficient estimator for sharp bounds on the value function under unobserved confounding. Our results provide a principled way for reliable decision-making from observational data, and show that robust policy learning can improve decision-making in sensitive applications such as healthcare.

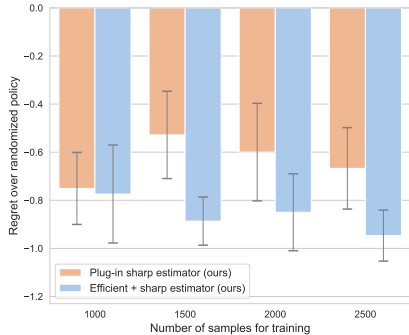


Figure 4: Property of semi-parametrically efficient estimation. We compare our efficient estimator with a simple plug-in estimator of our sharp upper bound from Proposition 4.1. We report the regret over a randomized policy (*lower values are better*). Our efficient estimator leads to a lower regret and benefits from increasing sample size due to its optimal estimation properties.




Figure 5: Real-world medical data. We compare our efficient estimator against the previous baselines based on data from the International Stroke Trial. Our method yields the best treatment policy and is robust over different Γ .

ACKNOWLEDGMENTS

This paper is supported by the DAAD program “Konrad Zuse Schools of Excellence in Artificial Intelligence”, sponsored by the Federal Ministry of Education and Research.

REFERENCES

- Peter M. Aronow and Donald K.K. Lee. Interval estimation of population means under unknown but bounded probabilities of sample selection. *Biometrika*, 100(1):235–240, 2013.
- Susan Athey and Stefan Wager. Policy learning with observational data. *Econometrica*, 89(1): 133–161, 2021.
- Alexis Bellot and Silvia Chiappa. Towards estimating bounds on the effect of policies under unobserved confounding. In *NeurIPS*, 2024.
- Andrew Bennett, Nathan Kallus, Lihong Li, and Ali Mousavi. Off-policy evaluation in infinite horizon reinforcement learning with latent confounders. In *AISTATS*, 2021.
- Andrew Bennett, Nathan Kallus, Miruna Oprescu, Wen Sun, and Kaiwen Wang. Efficient and sharp off-policy evaluation in robust Markov decision processes. In *NeurIPS*, 2024.
- Aurelien Bibaut, Ivana Malenica, Nikos Vlassis, and Mark van der Laan. More efficient off-policy evaluation through regularized targeted learning. In *ICML*, 2019.
- Matteo Bonvini, Edward Kennedy, Valerie Ventura, and Larry Wasserman. Sensitivity analysis for marginal structural models. *arXiv preprint*, arXiv:2210.04681, 2022.
- David Bruns-Smith and Angela Zhou. Robust fitted-q-evaluation and iteration under sequentially exogenous unobserved confounders. *arXiv preprint*, arXiv:2302.00662.
- Cecilia Ka Yuk Chan. A comprehensive AI policy framework for university teaching and learning. *International Journal of Educational Technology in Higher Education*, 20(38), 2023.
- Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James M. Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 2018.
- Carlos Cinelli and Chad Hazlett. Making sense of sensitivity: extending omitted variable bias. *Journal of the Royal Statistical Society*, 82(1):39–67, 2020.
- James Cornfield, William Haenszel, E. Cuyler Hammond, Abraham M. Lilienfeld, Michael B. Shimkin, and Ernst L. Wynder. Smoking and lung cancer: Recent evidence and a discussion of some questions. *Journal of the National Cancer Institute*, 22(1):173–203, 1959.
- Jacob Dorn and Kevin Guo. Sharp sensitivity analysis for inverse propensity weighting via quantile balancing. *Journal of the American Statistical Association*, 118(544):2645–2657, 2022.
- Jacob Dorn, Kevin Guo, and Nathan Kallus. Doubly-valid/ doubly-sharp sensitivity analysis for causal inference with unmeasured confounding. *Journal of the American Statistical Association*, 2024.
- Miroslav Dudik, John Langford, and Lihong Li. Doubly robust policy evaluation and learning. In *ICML*, 2011.
- Stefan Feuerriegel, Dennis Frauen, Valentyn Melnychuk, Jonas Schweisthal, Konstantin Hess, Alicia Curth, Stefan Bauer, Niki Kilbertus, Isaac S. Kohane, and Mihaela van der Schaar. Causal machine learning for predicting treatment outcomes. *Nature Medicine*, 2024.
- Dennis Frauen, Valentyn Melnychuk, and Stefan Feuerriegel. Sharp bounds for generalized causal sensitivity analysis. In *NeurIPS*, 2023.
- Dennis Frauen, Fergus Imrie, Alicia Curth, Valentyn Melnychuk, Stefan Feuerriegel, and Mihaela van der Schaar. A neural framework for generalized causal sensitivity analysis. In *ICLR*, 2024a.

- Dennis Frauen, Valentyn Melnychuk, and Stefan Feuerriegel. Fair off-policy learning from observational data. In *ICML*, 2024b.
- Dennis Frauen, Maresa Schröder, Konstantin Hess, and Stefan Feuerriegel. Orthogonal Survival Learners for Estimating Heterogeneous Treatment Effects from Time-to-Event Data. In *NeuIPS*, 2025a.
- Dennis Frauen, Jonas Schweisthal, Mihaela van der Schaar, and Stefan Feuerriegel. Treatment Effect Estimation for Optimal Decision Making. In *NeuIPS*, 2025b.
- Luke Guerdan, Amanda Coston, Kenneth Holstein, and Zhiwei Steven Wu. Predictive performance comparison of decision policies under confounding. In *ICML*, 2024.
- Tobias Hatt, Daniel Tschernutter, and Stefan Feuerriegel. Generalizing off-policy learning under sample selection bias. In *UAI*, 2022.
- Lars G. Hemkens, Hannah Ewald, Florian Naudet, Aviv Ladanie, Jonathan G. Shaw, Gautam Sajeew, and John P. A. Ioannidis. Interpretation of epidemiologic studies very often lacked adequate consideration of confounding. *Journal of Clinical Epidemiology*, 93, 2018.
- Konstantin Hess and Stefan Feuerriegel. Stabilized neural prediction of potential outcomes in continuous time. In *ICLR*, 2025.
- Konstantin Hess, Valentyn Melnychuk, Dennis Frauen, and Stefan Feuerriegel. Bayesian neural controlled differential equations for treatment effect estimation. In *ICLR*, 2024.
- Konstantin Hess, Dennis Frauen, Mihaela van der Schaar, and Stefan Feuerriegel. Overlap-weighted orthogonal meta-learner for treatment effect estimation over time. In *ICLR*, 2026.
- Jennifer L. Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):2017–2040, 2011.
- Oliver Hines, Oliver Dukes, Karla Diaz-Ordaz, and Stijn Vansteelandt. Demystifying statistical learning based on efficient influence functions. *The American Statistician*, 76(3):292–304, 2022.
- Wen Huang and Xintao Wu. Robustly improving bandit algorithms with confounded and selection biased offline data: A causal approach. In *AAAI*, 2024.
- Emil Javurek, Valentyn Melnychuk, Jonas Schweisthal, Konstantin Hess, Dennis Frauen, and Stefan Feuerriegel. An orthogonal learner for individualized outcomes in Markov decision processes. In *ICLR*, 2026.
- Andrew Jesson, Sören Mindermann, Yarin Gal, and Uri Shalit. Quantifying ignorance in individual-level causal-effect estimates under hidden confounding. In *ICML*, 2021.
- Ying Jin, Zhimei Ren, and Zhengyuan Zhou. Sensitivity analysis under the f -sensitivity models: A distributional robustness perspective. *arXiv preprint*, arXiv:2203.04373, 2022.
- Ying Jin, Zhimei Ren, and Emmanuel J. Candès. Sensitivity analysis of individual treatment effects: A robust conformal inference approach. *Proceedings of the National Academy of Sciences (PNAS)*, 120(6):e2214889120, 2023.
- Shalmali Joshi, Junzhe Zhang, and Elias Bareinboim. Towards safe policy learning under partial identifiability: A causal approach. In *AAAI*, 2024.
- Nathan Kallus. Balanced policy evaluation and learning. In *NeurIPS*, 2018.
- Nathan Kallus. More efficient policy learning via optimal retargeting. *Journal of the American Statistical Association*, 116(534):646–658, 2021.
- Nathan Kallus and Angela Zhou. Confounding-robust policy improvement. In *NeurIPS*, 2018a.
- Nathan Kallus and Angela Zhou. Policy evaluation and optimization with continuous treatments. In *AISTATS*, 2018b.

- Nathan Kallus and Angela Zhou. Confounding robust policy evaluation in infinite-horizon reinforcement learning. In *NeurIPS*, 2020.
- Nathan Kallus and Angela Zhou. Minimax-optimal policy learning under unobserved confounding. *Management Science*, 67(5):2870–2890, 2021.
- Nathan Kallus, Xiaojie Mao, and Angela Zhou. Interval estimation of individual-level causal effects under unobserved confounding. In *AISTATS*, 2019.
- Nathan Kallus, Xiaojie Mao, Kaiwen Wang, and Zhengyuan Zhou. Doubly robust distributionally robust off-policy evaluation and learning. In *ICML*, 2022.
- Chinmaya Kausik, Yangyi Lu, Kevin Tan, Maggie Maker, Yixin Wang, and Ambuj Tewari. Offline policy evaluation and optimization under confounding. In *AISTATS*, 2024.
- Edward H. Kennedy. Semiparametric doubly robust targeted double machine learning: A review. *arXiv preprint*, 2022.
- Mathias Kraus, Stefan Feuerriegel, and Maytal Saar-Tsechansky. Data-driven allocation of preventive care with application to diabetes mellitus type II. *Manufacturing & Service Operations Management*, 26(1):137–153, 2023.
- Stella Ladi and Dimitris Tsarouhas. EU economic governance and Covid-19: policy learning and windows of opportunity. *Journal of European Integration*, 42(8):1041–1056, 2020.
- Romain Laroche, Paul Trichelair, and Remi Tachet Des Combes. Safe policy improvement with baseline bootstrapping. In *ICML*, 2019.
- Michel Ledoux and Michel Talagrand. Comparison theorems, random geometry and some limit theorems for empirical processes. *Annals of Probability*, 17(2):596–631, 1989.
- Alex Luedtke. Simplifying debiased inference via automatic differentiation and probabilistic programming. *arXiv preprint*, 2405.08675, 2024.
- Colin McDiarmid. *On the method of bounded differences*, pp. 148–188. Surveys in Combinatorics, 1989: Invited Papers at the Twelfth British Combinatorial Conference. Cambridge University Press, 1989.
- Luke Miratrix and Stefan Wager. Shape-constrained partial identification of a population mean under unknown probabilities of sample selection. *Biometrika*, 105(1):103–114, 2018.
- Hongseok Namkoong, Ramtin Keramati, Steve Yadlowsky, and Emma Brunskill. Off-policy policy evaluation for sequential decisions under unobserved confounding. In *NeurIPS*, 2020.
- Jerzy Neyman. On the application of probability theory to agricultural experiments. *Annals of Agricultural Sciences*, 10:1–51, 1923.
- Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464):447–453, 2019.
- Miruna Oprescu, Jacob Dorn, Marah Ghoummaid, Andrew Jesson, Nathan Kallus, and Uri Shalit. B-learner: Quasi-oracle bounds on heterogeneous causal effects under hidden confounding. In *ICML*, 2023.
- Alizée Pace, Hugo Yèche, Bernhard Schölkopf, Gunnar Rätsch, and Guy Tennenholtz. Delphic offline reinforcement learning under nonidentifiable hidden confounding. In *ICLR*, 2024.
- Judea Pearl. *Causality*. Cambridge University Press, New York City, 2009. ISBN 9780521895606.
- Min Qian and Susan A. Murphy. Performance guarantees for individualized treatment rules. *Annals of Statistics*, 39(2):1180–1210, 2011.
- Ashesh Rambachan, Amanda Coston, and Edward Kennedy. Robust design and evaluation of predictive algorithms under unobserved confounding. *arXiv preprint*, arXiv:2212.09844, 2025.

- James M. Robins, Andrea Rotnitzky, and Lue Ping Zhao. Estimation of reversion coefficients when some regressors are not always observed. *Journal of the American Statistical Association*, 89(427): 846–688, 1994.
- Paul R. Rosenbaum. Sensitivity analysis for certain permutation inferences in matched observational studies. *Biometrika*, 74(1):13–26, 1987.
- Donald B. Rubin. Bayesian inference for causal effects: The role of randomization. *Annals of Statistics*, 6(1):34–58, 1978.
- Peter AG Sandercock, Maciej Niewada, and Anna Członkowska. The international stroke trial database. *Trials*, 12(101), 2011.
- Jonas Schweithal, Dennis Frauen, Valentyn Melnychuk, and Stefan Feuerriegel. Reliable off-policy learning for dosage combinations. In *NeurIPS*, 2023.
- Jonas Schweithal, Dennis Frauen, Maresa Schröder, Konstantin Hess, Niki Kilbertus, and Stefan Feuerriegel. Learning representations of instruments for partial identification of treatment effects. In *ICML*, 2025.
- Chengchun Shi, Masatoshi Uehara, Jiawei Huang, and Nan Jiang. A minimax learning approach to off-policy evaluation in confounded partially observable Markov decision processes. In *ICML*, 2022.
- Chengchun Shi, Jin Zhu, Ye Shen, Shikai Luo, Hongtu Zhu, and Rui Song. Off-policy confidence interval estimation with confounded Markov decision process. *Journal of the American Statistical Association*, 119(545):273–284, 2024.
- Adith Swaminathan and Thorsten Joachims. Counterfactual risk minimization: Learning from logged bandit feedback. In *ICML*, 2015.
- Zhiqiang Tan. A distributional approach for causal inference using propensity scores. *Journal of the American Statistical Association*, 101(476):1619–1637, 2006.
- Philip Thomas, Georgios Theodorou, and Mohammad Ghavamzadeh. High confidence policy improvement. In *ICML*, 2015.
- Daniel Tschernutter, Tobias Hatt, and Stefan Feuerriegel. Interpretable off-policy learning via hyperbox search. In *ICML*, 2022.
- Aart van der Vaart. *Asymptotic statistics*. Cambridge University Press, Cambridge, 1998. ISBN 0521496039.
- Lingxiao Wang, Zhuoran Yang, and Zhaoran Wang. Provably efficient causal reinforcement learning with confounded observational data. In *NeurIPS*, 2021.
- Steve Yadowsky, Hongseok Namkoong, Sanjay Basu, John Duchi, and Lu Tian. Bounds on the conditional and average treatment effect with unobserved confounding factors. *The Annals of Statistics*, 50(5):2587–2615, 2022.
- Mingzhang Yin, Claudia Shi, Yixin Wang, and David M. Blei. Conformal sensitivity analysis for individual treatment effects. *Journal of the American Statistical Association*, pp. 1–14, 2022.
- Junzhe Zhang and Elias Bareinboim. Eligibility traces for confounding robust off-policy evaluation. *OpenReview preprint*, 2024.
- Qingyuan Zhao, Dylan S. Small, and Bhaswar B. Bhattacharya. Sensitivity analysis for inverse probability weighting estimators via the percentile bootstrap. *Journal of the Royal Statistical Society: Series B*, 81(4):735–761, 2019.

A EXTENDED RELATED WORK

Offline reinforcement learning under unobserved confounding: Offline reinforcement learning deals with the problem of learning the optimal policy when the reward (value) function is defined over an infinite time horizon (Javurek et al., 2026). Therefore, these works rely upon techniques that are different from ours.

Some works focus on off-policy evaluation under unobserved confounding (Bennett et al., 2021; Kallus & Zhou, 2020) and even propose computationally efficient algorithms for this task (Kausik et al., 2024). However, these methods primarily focus on the identification of policy value bounds without semi-parametrically efficient estimation procedures. In contrast, Bennett et al. (2024) propose an efficient estimator for offline reinforcement learning. Different from our work, however, they require estimation of density ratios in order to evaluate the policy value. Further, Pace et al. (2024) propose a heuristic approach that learns representations of the unobserved confounders but does not provide theoretical guarantees for efficiency or unbiasedness. Shi et al. (2022) propose an approach that involves the approximation of bridge functions in partially observable Markov decision processes (POMDPs). Additionally, Shi et al. (2024) use mediators as auxiliary variables to construct confidence intervals for policy evaluation under unobserved confounding. Finally, Wang et al. (2021) improve sample efficiency in offline reinforcement learning under both observed and unobserved confounding.

Frauen et al. (2025b) develops a retargeting strategy for the conditional average treatment effect to balance estimation error and decision performance.

Further, Rambachan et al. (2025) study a slightly related but different problem setting: a *binary* selection decision $D \in \{0, 1\}$ with a *binary* selectively observed outcome Y^* , and they derive sharp bounds on *prediction and fairness* metrics (e.g. risk, ROC, precision–recall) under a *Mean Outcome Sensitivity Model (MOSM)*. The MOSM directly bounds the outcome gap

$$\delta(x) = P(Y^* = 1 \mid D = 0, X = x) - P(Y^* = 1 \mid D = 1, X = x), \quad (19)$$

and their sharpness guarantees hold relative to this MOSM.

Our work concerns *offline policy evaluation and learning* with (possibly) multi-valued actions $A \in \mathcal{A}$ and general outcomes Y , under a *Marginal Sensitivity Model (MSM)* on treatment assignment. We target the *policy value* $V(\pi)$ and derive *closed-form, sharp MSM bounds* together with a semiparametrically efficient EIF-based estimator.

Importantly, Rambachan et al. (2025) show that propensity-based models such as the MSM *imply a particular MOSM* (via outcome bounds derived from Bayes’ rule), but the resulting MOSM “may not exhaust all implications” of the original MSM. Thus MOSM bounds are generally *valid but not sharp* for the MSM. In contrast, our analysis works directly at the level of the MSM and characterizes the *sharp MSM-identified set* for $V(\pi)$.

B COMPARISON WITH KALLUS & ZHOU (2018A; 2021)

B.1 TECHNICAL SUMMARY OF KALLUS & ZHOU (2018A; 2021)

Kallus & Zhou (2018a; 2021) work under the same marginal sensitivity model as we do, specified via an odds-ratio bound Γ between the nominal propensity $e(A | X)$ and the true propensity $e(A | X, U)$. For a fixed policy π and baseline policy π_0 , they define an IPW/Hájek-type estimator for the *advantage* of π over π_0 , say $\lambda_i(\pi, \pi_0)$ for samples $i \in \{1, \dots, n\}$, and consider the set of feasible weights $w = (w_1, \dots, w_n)$ implied by the marginal sensitivity model:

$$U_\Gamma = \left\{ w : \underline{w}_i(\Gamma) \leq w_i \leq \bar{w}_i(\Gamma), \frac{1}{n} \sum_{i=1}^n w_i = 1 \right\}. \quad (20)$$

Given U_Γ , they define the worst-case regret of π relative to π_0 as

$$R_\Gamma(\pi; \pi_0) = \sup_{w \in U_\Gamma} \frac{\sum_{i=1}^n w_i \lambda_i(\pi, \pi_0)}{\sum_{i=1}^n w_i}, \quad (21)$$

and then select the policy that *minimizes* this worst-case regret. In the binary-action implementation, the inner optimization is solved by sorting the $\lambda_i(\pi, \pi_0)$ and exploiting the fact that the linear-fractional program over U_Γ attains its optimum at an extreme point: there exists a single index $k^* \in \{1, \dots, n\}$ such that units on one side of k^* are assigned their upper feasible weight and units on the other side their lower feasible weight. Hence the inner supremum reduces to a one-dimensional search over i to find k^* , which Kallus & Zhou (2018a; 2021) implement via ternary search.

B.2 OUR APPROACH AND KEY DIFFERENCES

Our work builds on the same marginal sensitivity model but differs in several technical aspects:

- *Objective: regret vs. value bounds.* Kallus & Zhou (2018a; 2021) optimize *worst-case regret* relative to a baseline policy, which is sharp for that regret quantity but does not provide closed-form sharp bounds on the *value* $V(\pi)$ itself. In contrast, we derive explicit *closed-form sharp upper and lower bounds* on $V(\pi)$ under the marginal sensitivity model and directly optimize these bounds over the policy class.
- *Propagation of IPW instability.* In Kallus & Zhou (2018a; 2021), the worst-case weights are determined by a single threshold k^* on the sorted IPW/Hájek-based scores λ_i . Because k^* is global, small estimation errors in the inverse propensities (and hence in the λ_i) can change which constraints bind in the inner problem and thus alter k^* . This then leads to a global change in the worst-case weights assigned to all units. Hence, local IPW noise can directly propagate into large changes in the estimated worst-case regret.
- *Closed-form sharp bounds and local error propagation.* Our estimator avoids this discrete minimax search and global thresholding. We derive sharp value bounds in closed form and estimate them with a semi-parametrically efficient, one-step bias-corrected estimator. That is, nuisance estimation errors then do **not** enter to the final bound estimate as first order bias, but only as lower order biases. Further, there is **no** single threshold index k^* in our method whose mis-estimation can flip the worst-case assignment for the entire sample. This yields substantially more stable behavior in practice.
- *Sharpness and efficiency.* While Kallus & Zhou (2018a; 2021) show sharpness for the regret interval relative to a baseline, our results provide the *globally sharp identified set* for $V(\pi)$ itself. Moreover, our one-step estimator based on the EIF is semi-parametrically efficient and Neyman-orthogonal with respect to the nuisances, whereas the Hájek/IPW-type estimators used in (Kallus & Zhou, 2018a; 2021) do not achieve this efficiency bound.

- *Sharp regret does not imply sharp value.* The Kallus–Zhou method yields the sharp identified set for the *regret* $R(P) = V(\pi; P) - V(\pi_0; P)$ over $P \in \mathcal{P}(\Gamma)$, but this does **not** imply sharp bounds on $V(\pi)$ itself. The distributions that maximize regret and those that maximize the baseline value generally do not coincide. Consequently, the implied value upper bound satisfies only the loose inequality

$$\sup_{P \in \mathcal{P}(\Gamma)} V(\pi; P) < \sup_{P \in \mathcal{P}(\Gamma)} (V(\pi; P) - V(\pi_0; P)) + \sup_{P \in \mathcal{P}(\Gamma)} V(\pi_0; P), \quad (22)$$

and this inequality is typically *strict*. Hence, even a sharp regret interval does not translate into the sharp identified set for $V(\pi)$. In contrast, our Theorem 4.1 directly characterizes the closed-form sharp bounds for $V(\pi)$ under the MSM.

Overall, we move from numerical minimax regret based on unstable IPW weights to closed-form sharp value bounds with an efficient, influence-function–based estimator that is both statistically and numerically more robust.

C ADDITIONAL RESULTS: COMPLEX PROPENSITY SCORES

In this section, we provide additional experimental results that illustrate the mechanism by which our efficient and sharp estimator improves over the method by Kallus & Zhou (2018a; 2021). In particular, we show how estimation errors in increasingly complex propensity score models propagate differently in the two approaches.

Experimental setup: We extend the synthetic data-generating process from Section 5 beyond binary actions and allow for

$$a \in \{0, 1, \dots, d_a - 1\}. \quad (23)$$

We fix the unobserved confounding at $\Gamma^* = 5$ and run both methods with the same correctly specified sensitivity parameter $\Gamma = 5$. Increasing d_a substantially complicates propensity estimation: estimating $e(a | X)$ becomes a multiclass classification problem, and the resulting inverse-propensity terms become noisier as d_a grows. This provides an ideal setting to study how each method handles nuisance complexity.

Results: Figure 6 summarizes the results. The findings align with our theoretical discussion in Supplement B:

- **Our method remains stable.** Our estimator maintains tight value bounds and stable performance as d_a increases. The one-step correction based on the efficient influence function ensures that first-order sensitivity to errors in the estimated nuisances is eliminated. Consequently, even as the propensity scores become harder to estimate, the resulting bounds remain well behaved and lead to an increasingly better, confounding-robust policy.
- **Kallus & Zhou (2018a; 2021) become increasingly conservative.** In contrast, the approach of Kallus & Zhou (2018a; 2021) again reverts to the baseline policy. Their method computes worst-case regret relative to π_0 via a minimax optimization over feasible weights U_Γ :

$$R_\Gamma(\pi; \pi_0) = \sup_{w \in U_\Gamma} \frac{\sum_{i=1}^n w_i \lambda_i(\pi, \pi_0)}{\sum_{i=1}^n w_i},$$

where the $\lambda_i(\pi, \pi_0)$ are IPW/Hájek-type advantage scores. The optimizer selects a single global threshold k^* (via ternary search) after sorting the λ_i to determine which observations receive lower versus upper feasible weights.

When d_a grows, the estimated λ_i become highly variable due to propensity estimation noise. This causes the inner maximization to lock onto overly pessimistic weights in which many non-baseline units receive their lower feasible weight. As a result, nearly all non-baseline policies appear unsafe, and the learned policy collapses back to π_0 , even though the true confounding level is correctly specified.

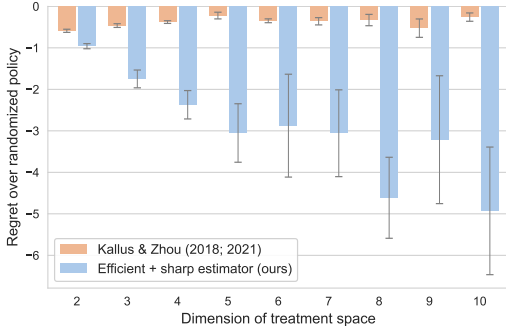


Figure 6: Performance comparison between our efficient and sharp method and Kallus & Zhou (2018a; 2021) under increasing dimension d_a of the action-space.

D IHDP DATASET

	$\Gamma=1.0$	1.1	1.2	1.3	1.4	1.5	1.6	1.7	1.8	1.9	2.0	2.1	2.2	2.3
Standard IPW estimator	-0.32 ± 0.15	-0.32 ± 0.15	-0.32 ± 0.15	-0.32 ± 0.15	-0.32 ± 0.15	-0.32 ± 0.15	-0.32 ± 0.15	-0.32 ± 0.15	-0.32 ± 0.15	-0.32 ± 0.15	-0.32 ± 0.15	-0.32 ± 0.15	-0.32 ± 0.15	-0.32 ± 0.15
Standard DR estimator	-2.19 ± 0.41	-2.19 ± 0.41	-2.19 ± 0.41	-2.19 ± 0.41	-2.19 ± 0.41	-2.19 ± 0.41	-2.19 ± 0.41	-2.19 ± 0.41	-2.19 ± 0.41	-2.19 ± 0.41	-2.19 ± 0.41	-2.19 ± 0.41	-2.19 ± 0.41	-2.19 ± 0.41
Kallus & Zhou (2018a; 2021)	-0.02 ± 0.15	-0.02 ± 0.14	-0.01 ± 0.14	-0.00 ± 0.15	-0.02 ± 0.14	0.02 ± 0.13	0.03 ± 0.13	0.02 ± 0.13	-0.01 ± 0.13	0.01 ± 0.14	0.02 ± 0.12	0.02 ± 0.12	0.01 ± 0.12	-0.01 ± 0.13
Efficient + sharp (ours)	-2.42 ± 0.31	-2.46 ± 0.32	-2.92 ± 0.22	-2.68 ± 0.27	-2.84 ± 0.24	-2.88 ± 0.26	-3.05 ± 0.25	-2.89 ± 0.27	-2.82 ± 0.30	-2.65 ± 0.28	-2.58 ± 0.25	-2.57 ± 0.25	-2.54 ± 0.30	-2.60 ± 0.31
Absolute improvement	-0.23	-0.27	-0.74	-0.49	-0.65	-0.70	-0.86	-0.70	-0.63	-0.46	-0.39	-0.38	-0.35	-0.41

Table 3: **Regret over randomized policy for varying sensitivity parameter Γ .** Lower values indicate better performance. Absolute improvement is computed against the best-performing baseline.

Dataset: The IHDP dataset is a widely used semi-synthetic benchmark based on the Infant Health and Development Program (IHDP) randomized trial. Here, the effect of home visits on cognitive test scores for infants are measured. Covariates are taken from the real RCT, and potential outcomes are simulated following Hill (2011).

Similar to our real-world data study in Section 5, because the original study is unconfounded, we can introduce a known amount of unobserved confounding in a controlled way: here, we select covariates *birth weight (indicator of neonatal health)*, *head circumference at birth (developmental marker)*, and *weeks of gestation (prematurity measure)*, and drop 60% of the observations whose value for any of the three covariates is below average and has received home visits, as well as 60% of the observations whose value is below the average and has received *no* home visits. Then, we drop the three covariates and, thereby, introduce unobserved confounding. This mimics a realistic setting where key health indicators are absent from the recorded data but nevertheless influence treatment decisions and the outcome variable.

Results: We evaluate the regret of all methods relative to a randomized policy because, under unobserved confounding, it is the only baseline whose value remains identified, whereas any non-randomized policy would implicitly rely on violated unconfoundedness assumptions and thus provide a misleading comparison.

The results in Table 3 closely mirror the patterns in our previous real-world experiment in Section 5. As expected, the standard IPW and DR baselines are biased and completely insensitive to the choice of Γ , since they assume unconfoundedness regardless of the true data-generating process.

The method of Kallus & Zhou (2018a; 2021) again behaves overly conservatively. In particular, it effectively reverts back toward its baseline policy and produces regret over the RCT baseline that cluster near zero. This, again, shows that the worst-case optimization collapses under even moderate confounding.

In contrast, our efficient and sharp estimator consistently achieves the lowest regret across all values of Γ . This demonstrates that our method computes meaningful bounds on the unidentified value function and, on top, finds optimizes toward the global best performing, confounding-robust policy.

E CHOOSING THE SENSITIVITY PARAMETER IN THE MSM

In this work, we adopt the MSM (Tan, 2006) in order to bound the ratio between the *nominal propensity score*

$$e(a, x) = p(A = a \mid X = x), \quad (24)$$

and the *true propensity score*

$$e(a, x, u) = p(A = a \mid X = x, U = u). \quad (25)$$

Here, the nominal propensity score can be estimated from data, whereas the true propensity score is fundamentally unobservable. In particular, the MSM is given by

$$\Gamma^{-1} \leq \frac{e(a, x)}{1 - e(a, x)} \frac{1 - e(a, x, u)}{e(a, x, u)} \leq \Gamma \quad (26)$$

for some sensitivity parameter $\Gamma \geq 1$.

Typically, the sensitivity constraints Γ are chosen via domain knowledge (Frauen et al., 2023; Kallus et al., 2019) or data-driven heuristics (Hatt et al., 2022). For example, in practical applications, one typically has a benchmark variable (e.g., hours with sunlight) that is a known cause of the outcome (e.g., vitamin D deficiency), and one then wants to study how strong a confounder (e.g., other ecological activities such as nutrition) must be to explain away the effect of the benchmark variables. (Cinelli & Hazlett, 2020) term this the robustness value, which quantifies the strength of unobserved confounding needed to change conclusions.

Hence, to achieve this, a commonly used strategy for selecting Γ is the following: We can search for the smallest Γ such that the partially identified interval for the causal quantity of interest includes 0. Then, we can interpret Γ as a measure of the minimal deviation from unconfoundedness required to invalidate the effect of an intervention (Jesson et al., 2021; Jin et al., 2023).

F PROOFS

F.1 SHARP BOUND OF THE VALUE FUNCTION

Proposition F.1. *Let $Q^{\pm,*}(a, x)$ be the sharp upper/lower bound for the conditional average potential outcome, given our sensitivity constraints $\mathcal{P}(\Gamma)$. Then, sharp bounds for the value function $V(\pi)$ are given by*

$$V^{\pm,*}(\pi) = \int_{\mathcal{X}} \sum_a Q^{\pm,*}(a, x) \pi(a | x) dp(x). \quad (27)$$

Proof. We provide the proof for the sharp upper bound $V^{+,*}(\pi)$. The lower bound follows completely analogously by swapping the signs and replacing the supremum with an infimum.

We start by noting that the upper bound on the value function depends on the set of admissible distributions $\mathcal{P}(\Gamma)$ induced by the sensitivity model, that is,

$$V^{+,*}(\pi) = V^{+,*}(\pi, \mathcal{P}(\Gamma)). \quad (28)$$

Hence, we can write

$$V^{+,*}(\pi) \quad (29)$$

$$= V^{+,*}(\pi, \mathcal{P}(\Gamma)) \quad (30)$$

$$= \sup_{\tilde{p} \in \mathcal{P}(\Gamma)} V(\pi, \tilde{p}) \quad (31)$$

$$= \sup_{\tilde{p} \in \mathcal{P}(\Gamma)} \int_{\mathcal{X}} \sum_a Q(a, x, \tilde{p}) \pi(a | x) d\tilde{p}(x) \quad (32)$$

$$= \sup_{\tilde{p} \in \mathcal{P}(\Gamma)} \int_{\mathcal{X}} \sum_a Q(a, x, \tilde{p}) \pi(a | x) dp(x), \quad (33)$$

where Equation 33 follows from the equality $p(\mathcal{D}) = \tilde{p}(\mathcal{D})$ for all $\tilde{p} \in \mathcal{P}(\Gamma)$.

Clearly, by definition of the optimal bounds $Q^{+,*}(a, x)$, we have that

$$Q(a, x, \tilde{p}) \leq \sup_{\tilde{p} \in \mathcal{P}(\Gamma)} Q(a, x, \tilde{p}) = Q^{+,*}(a, x, \mathcal{P}(\Gamma)) \quad (34)$$

for all $\tilde{p} \in \mathcal{P}(\Gamma)$, and since $Q^{+,*}(a, x) \in L^1(\pi, p)$, we know by dominated convergence (Frauen et al., 2023) that

$$\sup_{\tilde{p} \in \mathcal{P}(\Gamma)} \int_{\mathcal{X}} \sum_a Q(a, x, \tilde{p}) \pi(a | x) dp(x) = \int_{\mathcal{X}} \sum_a Q^{+,*}(a, x) \pi(a | x) dp(x). \quad (35)$$

□

F.2 EFFICIENT ESTIMATOR OF THE SHARP BOUND OF THE VALUE FUNCTION

On the construction of semi-parametrically efficient estimators. Let $\Psi(P) = V^{+,*}(\pi; P)$ denote our target functional in the nonparametric model, and let $\mathcal{T}(P)$ be the corresponding tangent space. For any regular parametric submodel $\{P_\varepsilon\}$ through P with score $s_\varepsilon(Z) = \partial_\varepsilon \log p_\varepsilon(Z)|_{\varepsilon=0}$, the pathwise derivative of Ψ at P in direction $s \in \mathcal{T}(P)$ is

$$\dot{\Psi}_P(s) := \left. \frac{d}{d\varepsilon} \Psi(P_\varepsilon) \right|_{\varepsilon=0}.$$

The *canonical gradient* $g^* \in L_0^2(P)$ is defined as the unique element such that

$$\dot{\Psi}_P(s) = \mathbb{E}_P[g^*(Z) s(Z)] \quad \text{for all } s \in \mathcal{T}(P), \quad (36)$$

i.e., g^* is the Riesz representer of the pathwise derivative functional (see, e.g., van der Vaart (1998), Chapter 25).

In the following, we derive the influence function of $V^{+,*}(\pi)$ via the chain rule for pathwise differentiability. Specifically, we decompose $V^{+,*}(\pi)$ into building blocks (conditional expectations, conditional distributions), plug in the known canonical gradients for each of these primitives, and then apply the chain rule as in Kennedy (2022) and Lemma S3 of Luedtke (2024). By construction, the resulting function $\varphi(Z)$ satisfies

$$\left. \frac{d}{d\varepsilon} \Psi(P_\varepsilon) \right|_{\varepsilon=0} = \mathbb{E}_P[\varphi(Z) s_\varepsilon(Z)] \quad \text{for every regular submodel } \{P_\varepsilon\}. \quad (37)$$

Hence φ is exactly the canonical gradient g^* , and therefore the efficient influence function in the nonparametric model. By standard semiparametric theory (e.g., van der Vaart (1998)), the one-step estimator based on this φ is semiparametrically efficient for $V^{+,*}(\pi)$.

Theorem F.2. *Let $\mathbb{E}[|Y|] < \infty$ and $p(y | x, a)$ admit a continuous density bounded away from zero in a neighborhood of $F_{x,a}^{-1}(\alpha^+)$. This holds, for example, if the density $p(y | x, a)$ is strictly positive and continuous. Then, an estimator for the sharp upper bound of the value function is given by*

$$\begin{aligned} \hat{V}^{+,*}(\pi) = \mathbb{P}_n \left\{ \sum_a \pi_{a,X} \left[\hat{Q}_{a,X}^{+,*} - \hat{e}_{a,X} \left(b^- \hat{\mu}_{a,X}^+ + b^+ \hat{\mu}_{a,X}^+ \right) \right] + \pi_{A,X} \left(b^- \hat{\mu}_{A,X}^+ + b^+ \hat{\mu}_{A,X}^+ \right) \right. \\ \left. + \frac{\pi_{A,X}}{\hat{e}_{A,X}} \left[\left(\hat{c}_{A,X}^- - \hat{c}_{A,X}^+ \right) \left(\hat{F}_{X,A}^{-1}(\alpha^+) \left(\hat{\Delta}_{Y,A,X}^+ - \alpha^+ \right) \right) \right. \right. \\ \left. \left. + \hat{c}_{A,X}^- \left(Y \hat{\Delta}_{Y,A,X}^+ - \hat{\mu}_{A,X}^+ \right) + \hat{c}_{A,X}^+ \left(Y \hat{\Delta}_{Y,A,X}^+ - \hat{\mu}_{A,X}^+ \right) \right] \right\}. \quad (38) \end{aligned}$$

If, on top, $\mathbb{E}[|Y|^2] < \infty$, the above estimator is *semi-parametrically efficient*.

Proof. The sharp upper bound of the value function is given by

$$V^{\pm,*}(\pi) = \int_{\mathcal{X}} \sum_a Q^{\pm,*}(a, x) \pi(a | x) dp(x). \quad (39)$$

In the following, we derive the efficient estimator for this quantity. Therein, we make use of the chain rule for deriving efficient influence function (Kennedy, 2022). A proof of the validity of the chain rule for deriving efficient influence functions is provided by Luedtke (2024) (**Lemma S3**).

In order to avoid notational overload and for the sake of clarity, we do not use additional variables such as $b^\pm, c^\pm, \Delta^\pm, \bar{\mu}^\pm$, etc. until the final steps, such that the derivation becomes easier to follow. Moreover, we make the dependency on nuisance functions $\eta \subseteq \{e(a, x), F_{a,x}^{-1}(\alpha^\pm), \bar{\mu}^\pm(a, x), \underline{\mu}^\pm(a, x)\}$ explicit by writing, for example, $V^{+,*}(\pi; \eta)$ for $V^{+,*}(\pi)$.

The influence function of $V^{+,*}(\pi; \eta)$ is given by

$$\mathbb{IF}\left(V^{+,*}(\pi; \eta)\right) \quad (40)$$

$$= \mathbb{IF}\left(\int_{\mathcal{X}} \sum_a Q^{+,*}(a, x; \eta) \pi(a | x) dp(x)\right) \quad (41)$$

$$= \sum_a \int_{\mathcal{X}} \pi(a | x) \mathbb{IF}\left(p(x) Q^{+,*}(a, x; \eta)\right) dx \quad (42)$$

$$= \sum_a \int_{\mathcal{X}} \pi(a | x) \mathbb{IF}\left(p(x)\right) Q^{+,*}(a, x; \eta) + \pi(a | x) p(x) \mathbb{IF}\left(Q^{+,*}(a, x; \eta)\right) dx \quad (43)$$

$$= \sum_a \int_{\mathcal{X}} \pi(a | x) \left(\mathbb{1}_{\{X=x\}} - p(x)\right) Q^{+,*}(a, x; \eta) dx + \sum_a \int_{\mathcal{X}} \pi(a | x) p(x) \mathbb{IF}\left(Q^{+,*}(a, x; \eta)\right) dx \quad (44)$$

$$= \sum_a \pi(a | X) Q^{+,*}(a, X; \eta) - V^{+,*}(\pi; \eta) + \sum_a \int_{\mathcal{X}} \pi(a | x) p(x) \mathbb{IF}\left(Q^{+,*}(a, x; \eta)\right) dx \quad (45)$$

Hence, in Equation 45, we are left to compute the efficient influence function of $Q^{+,*}(a, x)$, that is, the sharp upper bound of the CAPO. With $\alpha^+ = \Gamma/(1 + \Gamma)$, the sharp upper bound $Q^{+,*}(a, x)$ is given by

$$Q^{+,*}(a, x; \eta) \quad (46)$$

$$= \left((1 - \Gamma^{-1})e(a, x) + \Gamma^{-1}\right) \mathbb{E}\left[Y \mathbb{1}_{\{Y \leq F_{X,A}^{-1}(\alpha^+)\}} \mid X = x, A = a\right] \quad (47)$$

$$+ \left((1 - \Gamma)e(a, x) + \Gamma\right) \mathbb{E}\left[Y \mathbb{1}_{\{Y \geq F_{X,A}^{-1}(\alpha^+)\}} \mid X = x, A = a\right]. \quad (48)$$

Hence, the influence function is given by

$$\mathbb{IF}\left(Q^{+,*}(a, x; \eta)\right) \quad (49)$$

$$= \underbrace{\mathbb{IF}\left(\left((1 - \Gamma^{-1})e(a, x) + \Gamma^{-1}\right) \mathbb{E}\left[Y \mathbb{1}_{\{Y \leq F_{X,A}^{-1}(\alpha^+)\}} \mid X = x, A = a\right]\right)}_{(a)} \quad (50)$$

$$+ \underbrace{\left((1 - \Gamma^{-1})e(a, x) + \Gamma^{-1}\right) \mathbb{IF}\left(\mathbb{E}\left[Y \mathbb{1}_{\{Y \leq F_{X,A}^{-1}(\alpha^+)\}} \mid X = x, A = a\right]\right)}_{(c)} \quad (51)$$

$$+ \underbrace{\mathbb{IF}\left(\left((1 - \Gamma)e(a, x) + \Gamma\right) \mathbb{E}\left[Y \mathbb{1}_{\{Y \geq F_{X,A}^{-1}(\alpha^+)\}} \mid X = x, A = a\right]\right)}_{(b)} \quad (52)$$

$$+ \underbrace{\left((1 - \Gamma)e(a, x) + \Gamma\right) \mathbb{IF}\left(\mathbb{E}\left[Y \mathbb{1}_{\{Y \geq F_{X,A}^{-1}(\alpha^+)\}} \mid X = x, A = a\right]\right)}_{(d)}, \quad (53)$$

where we can apply the chain rule to the conditional expectation of the truncated first moments of Y , since Y is integrable by assumption. We start with (a) and (b). For (a), we obtain

$$\mathbb{IF}\left(\left((1 - \Gamma^{-1})e(a, x) + \Gamma^{-1}\right)\right) \quad (54)$$

$$= (1 - \Gamma^{-1}) \mathbb{IF}\left(e(a, x)\right) \quad (55)$$

$$= (1 - \Gamma^{-1}) \frac{\mathbb{1}_{\{X=x\}}}{p(x)} \left(\mathbb{1}_{\{A=a\}} - e(a, x)\right), \quad (56)$$

and, similarly for (b), we yield

$$\mathbb{IF}\left(\left((1 - \Gamma)e(a, x) + \Gamma\right)\right) \quad (57)$$

$$= (1 - \Gamma) \frac{\mathbb{1}_{\{X=x\}}}{p(x)} \left(\mathbb{1}_{\{A=a\}} - e(a, x)\right). \quad (58)$$

Next, we compute the influence function of (c) via

$$\mathbb{IF}\left(\mathbb{E}\left[Y\mathbb{1}_{\{Y\leq F_{x,A}^{-1}(\alpha^+)\}}\mid X=x, A=a\right]\right) \quad (59)$$

$$=\mathbb{IF}\left(\int_{\mathcal{Y}}\mathbb{1}_{\{y\leq F_{x,a}^{-1}(\alpha^+)\}}yp(y\mid x,a)dy\right) \quad (60)$$

$$=\int_{\mathcal{Y}}\mathbb{IF}\left(\mathbb{1}_{\{y\leq F_{x,a}^{-1}(\alpha^+)\}}\right)yp(y\mid x,a)+\mathbb{1}_{\{y\leq F_{x,a}^{-1}(\alpha^+)\}}y\mathbb{IF}\left(p(y\mid x,a)\right)dy \quad (61)$$

$$=\underbrace{\int_{\mathcal{Y}}\mathbb{IF}\left(\mathbb{1}_{\{y\leq F_{x,a}^{-1}(\alpha^+)\}}\right)yp(y\mid x,a)dy}_{(c_1)}+\underbrace{\int_{\mathcal{Y}}\mathbb{1}_{\{y\leq F_{x,a}^{-1}(\alpha^+)\}}y\mathbb{IF}\left(p(y\mid x,a)\right)dy}_{(c_2)}. \quad (62)$$

For (c₁), we first note that

$$\mathbb{IF}\left(F_{x,a}(\alpha^+)\right) \quad (63)$$

$$=\mathbb{IF}\left(\mathbb{P}(Y\leq y\mid X=x, A=a)\right) \quad (64)$$

$$=\mathbb{IF}\left(\mathbb{E}[\mathbb{1}_{\{Y\leq y\}}\mid X=x, A=a]\right) \quad (65)$$

$$=\frac{\mathbb{1}_{\{X=x, A=a\}}}{p(a,x)}\left(\mathbb{1}_{\{Y\leq y\}}-\mathbb{E}[\mathbb{1}_{\{Y\leq y\}}\mid X=x, A=a]\right) \quad (66)$$

$$=\underbrace{\frac{\mathbb{1}_{\{X=x, A=a\}}}{p(a,x)}\left(\mathbb{1}_{\{Y\leq y\}}-F_{x,a}(y)\right)}_{(*)}. \quad (67)$$

Then, we can simplify (c₁) via

$$\int_{\mathcal{Y}}\mathbb{IF}\left(\mathbb{1}_{\{y\leq F_{x,a}^{-1}(\alpha^+)\}}\right)yp(y\mid x,a)dy \quad (68)$$

$$=\int_{\mathcal{Y}}\delta\left(y-F_{x,a}^{-1}(\alpha^+)\right)\mathbb{IF}\left(F_{x,a}^{-1}(\alpha^+)\right)yp(y\mid x,a)dy \quad (69)$$

$$=\mathbb{IF}\left(F_{x,a}^{-1}(\alpha^+)\right)\int_{\mathcal{Y}}\delta\left(y-F_{x,a}^{-1}(\alpha^+)\right)yp(y\mid x,a)dy \quad (70)$$

$$=\mathbb{IF}\left(F_{x,a}^{-1}(F_{x,a}(y^*))\right)\int_{\mathcal{Y}}\delta\left(y-F_{x,a}^{-1}(\alpha^+)\right)yp(y\mid x,a)dy \quad (71)$$

$$=\frac{d}{dq}F_{x,a}^{-1}(q)\Big|_{q=F_{x,a}(y^*)}\mathbb{IF}\left(F_{x,a}(y^*)\right)\int_{\mathcal{Y}}\delta\left(y-F_{x,a}^{-1}(\alpha^+)\right)yp(y\mid x,a)dy \quad (72)$$

$$=\frac{1}{F'_{x,a}(F_{x,a}^{-1}(F_{x,a}(y^*)))}\mathbb{IF}\left(F_{x,a}(y^*)\right)\int_{\mathcal{Y}}\delta\left(y-F_{x,a}^{-1}(\alpha^+)\right)yp(y\mid x,a)dy \quad (73)$$

$$=\frac{1}{p(y^*\mid x,a)}\mathbb{IF}\left(F_{x,a}(y^*)\right)\int_{\mathcal{Y}}\delta\left(y-F_{x,a}^{-1}(\alpha^+)\right)yp(y\mid x,a)dy \quad (74)$$

$$\stackrel{(*)}{=} \frac{1}{p(y^*\mid x,a)}\frac{\mathbb{1}_{\{X=x, A=a\}}}{p(a,x)}\left(\mathbb{1}_{\{Y\leq y^*\}}-F_{x,a}(y^*)\right)\int_{\mathcal{Y}}\delta(y-y^*)yp(y\mid x,a)dy \quad (75)$$

$$=\frac{1}{p(y^*\mid x,a)}\frac{\mathbb{1}_{\{X=x, A=a\}}}{p(a,x)}\left(\mathbb{1}_{\{Y\leq y^*\}}-\alpha^+\right)\int_{\mathcal{Y}}\delta(y-y^*)yp(y\mid x,a)dy \quad (76)$$

$$=\frac{1}{p(y^*\mid x,a)}\frac{\mathbb{1}_{\{X=x, A=a\}}}{p(a,x)}\left(\mathbb{1}_{\{Y\leq y^*\}}-\alpha^+\right)y^*p(y^*\mid x,a) \quad (77)$$

$$=\frac{\mathbb{1}_{\{X=x, A=a\}}}{p(a,x)}F_{x,a}^{-1}(\alpha^+)\left(\mathbb{1}_{\{Y\leq F_{x,a}^{-1}(\alpha^+)\}}-\alpha^+\right), \quad (78)$$

for some $y^* \in \mathcal{Y}$ such that $F_{x,a}(y^*) = \alpha^+$, where $\delta(\cdot)$ is the Dirac-delta function, and where we use in Equation 73 that $p(y\mid x,a)$ is a continuous density strictly bounded away from zero in a

neighborhood of $F_{x,a}^{-1}(\alpha^+)$, which guarantees that the *conditional quantile is unique*, and where Equation 74 exists by positivity of the conditional density.

Next, we simplify (c₂) via

$$\int_{\mathcal{Y}} \mathbb{1}_{\{y \leq F_{x,a}^{-1}(\alpha^+)\}} y \mathbb{I}\mathbb{F}(p(y | x, a)) \, dy \quad (79)$$

$$= \int_{\mathcal{Y}} \mathbb{1}_{\{y \leq F_{x,a}^{-1}(\alpha^+)\}} y \mathbb{I}\mathbb{F}\left(\mathbb{E}[\mathbb{1}_{\{Y=y\}} | X=x, A=a]\right) \, dy \quad (80)$$

$$= \int_{\mathcal{Y}} \mathbb{1}_{\{y \leq F_{x,a}^{-1}(\alpha^+)\}} y \frac{\mathbb{1}_{\{X=x, A=a\}}}{p(a, x)} \left(\mathbb{1}_{\{Y=y\}} - p(y | x, a) \right) \, dy \quad (81)$$

$$= \frac{\mathbb{1}_{\{X=x, A=a\}}}{p(a, x)} \left(Y \mathbb{1}_{\{Y \leq F_{x,a}^{-1}(\alpha^+)\}} - \mathbb{E}[Y \mathbb{1}_{\{Y \leq F_{X,A}^{-1}(\alpha^+)\}} | X=x, A=a] \right). \quad (82)$$

Then, combining (c₁) and (c₂), we get

$$\mathbb{I}\mathbb{F}\left(\mathbb{E}\left[Y \mathbb{1}_{\{Y \leq F_{X,A}^{-1}(\alpha^+)\}} | X=x, A=a\right]\right) \quad (83)$$

$$= \frac{\mathbb{1}_{\{X=x, A=a\}}}{p(a, x)} \left(Y \mathbb{1}_{\{Y \leq F_{x,a}^{-1}(\alpha^+)\}} - \mathbb{E}[Y \mathbb{1}_{\{Y \leq F_{x,a}^{-1}(\alpha^+)\}} | X=x, A=a] + F_{x,a}^{-1}(\alpha^+) (\mathbb{1}_{\{Y \leq F_{x,a}^{-1}(\alpha^+)\}} - \alpha^+) \right). \quad (84)$$

Finally, we compute the influence function of (d) analogously to (c) via

$$\mathbb{I}\mathbb{F}\left(\mathbb{E}\left[Y \mathbb{1}_{\{Y \geq F_{X,A}^{-1}(\alpha^+)\}} | X=x, A=a\right]\right) \quad (85)$$

$$= \underbrace{\int_{\mathcal{Y}} \mathbb{I}\mathbb{F}\left(\mathbb{1}_{\{y \geq F_{x,a}^{-1}(\alpha^+)\}}\right) y p(y | x, a) \, dy}_{(d_1)} + \underbrace{\int_{\mathcal{Y}} \mathbb{1}_{\{y \geq F_{x,a}^{-1}(\alpha^+)\}} y \mathbb{I}\mathbb{F}(p(y | x, a)) \, dy}_{(d_2)}. \quad (86)$$

Again, we start with (d₁) via

$$\int_{\mathcal{Y}} \mathbb{I}\mathbb{F}\left(\mathbb{1}_{\{y \geq F_{x,a}^{-1}(\alpha^+)\}}\right) y p(y | x, a) \, dy \quad (87)$$

$$= \int_{\mathcal{Y}} \mathbb{I}\mathbb{F}\left((1 - \mathbb{1}_{\{y \leq F_{x,a}^{-1}(\alpha^+)\}})\right) y p(y | x, a) \, dy \quad (88)$$

$$= - \int_{\mathcal{Y}} \mathbb{I}\mathbb{F}\left(\mathbb{1}_{\{y \leq F_{x,a}^{-1}(\alpha^+)\}}\right) y p(y | x, a) \, dy \quad (89)$$

$$= \frac{\mathbb{1}_{\{X=x, A=a\}}}{p(a, x)} \left(- F_{x,a}^{-1}(\alpha^+) \right) (\mathbb{1}_{\{Y \leq F_{x,a}^{-1}(\alpha^+)\}} - \alpha^+), \quad (90)$$

using the result for (c₁) in Equation 78. Further, for (d₂), we get that

$$\int_{\mathcal{Y}} \mathbb{1}_{\{y \geq F_{x,a}^{-1}(\alpha^+)\}} y \mathbb{I}\mathbb{F}(p(y | x, a)) \, dy \quad (91)$$

$$= \frac{\mathbb{1}_{\{X=x, A=a\}}}{p(a, x)} \left(Y \mathbb{1}_{\{Y \geq F_{x,a}^{-1}(\alpha^+)\}} - \mathbb{E}[Y \mathbb{1}_{\{Y \geq F_{X,A}^{-1}(\alpha^+)\}} | X=x, A=a] \right). \quad (92)$$

Combining (d₁) and (d₂), we obtain that

$$\mathbb{I}\mathbb{F}\left(\mathbb{E}\left[Y \mathbb{1}_{\{Y \geq F_{X,A}^{-1}(\alpha^+)\}} | X=x, A=a\right]\right) \quad (93)$$

$$= \frac{\mathbb{1}_{\{X=x, A=a\}}}{p(a, x)} \left(Y \mathbb{1}_{\{Y \geq F_{x,a}^{-1}(\alpha^+)\}} - \mathbb{E}[Y \mathbb{1}_{\{Y \geq F_{X,A}^{-1}(\alpha^+)\}} | X=x, A=a] - F_{x,a}^{-1}(\alpha^+) (\mathbb{1}_{\{Y \leq F_{x,a}^{-1}(\alpha^+)\}} - \alpha^+) \right). \quad (94)$$

Finally, we can state the influence function of $Q^{+,*}(a, x)$ by combining (a)–(d) in Equation 56, Equation 58, Equation 84, and Equation 94. We get

$$\mathbb{IF}\left(Q^{+,*}(a, x; \eta)\right) \quad (95)$$

$$= \frac{\mathbb{1}_{\{X=x\}}}{p(x)} (1 - \Gamma^{-1}) \left(\mathbb{1}_{\{A=a\}} - e(a, x) \right) \mathbb{E}\left[Y \mathbb{1}_{\{Y \leq F_{X,A}^{-1}(\alpha^+)\}} \mid X = x, A = a \right] \quad (96)$$

$$+ \frac{\mathbb{1}_{\{X=x, A=a\}}}{p(a, x)} \left((1 - \Gamma^{-1})e(a, x) + \Gamma^{-1} \right) \quad (97)$$

$$\times \left(Y \mathbb{1}_{\{Y \leq F_{x,a}^{-1}(\alpha^+)\}} - \mathbb{E}[Y \mathbb{1}_{\{Y \leq F_{X,A}^{-1}(\alpha^+)\}} \mid X = x, A = a] + F_{x,a}^{-1}(\alpha^+) (\mathbb{1}_{\{Y \leq F_{x,a}^{-1}(\alpha^+)\}} - \alpha^+) \right) \quad (98)$$

$$+ \frac{\mathbb{1}_{\{X=x\}}}{p(x)} (1 - \Gamma) \left(\mathbb{1}_{\{A=a\}} - e(a, x) \right) \mathbb{E}\left[Y \mathbb{1}_{\{Y \geq F_{X,A}^{-1}(\alpha^+)\}} \mid X = x, A = a \right] \quad (99)$$

$$+ \frac{\mathbb{1}_{\{X=x, A=a\}}}{p(a, x)} \left((1 - \Gamma)e(a, x) + \Gamma \right) \quad (100)$$

$$\times \left(Y \mathbb{1}_{\{Y \geq F_{x,a}^{-1}(\alpha^+)\}} - \mathbb{E}[Y \mathbb{1}_{\{Y \geq F_{X,A}^{-1}(\alpha^+)\}} \mid X = x, A = a] - F_{x,a}^{-1}(\alpha^+) (\mathbb{1}_{\{Y \geq F_{x,a}^{-1}(\alpha^+)\}} - \alpha^+) \right). \quad (101)$$

In order to simplify the above lengthy equation as in our main paper, we introduce the following variables:

- $b^\pm = (1 - \Gamma^{\pm 1})$
- $c^\pm(a, x; \eta) = (b^\pm e(a, x) + \Gamma^{\pm 1})$
- $\underline{\Delta}^\pm(y, a, x; \eta) = \mathbb{1}_{\{y \leq F_{x,a}^{-1}(\alpha^\pm)\}}$
- $\bar{\Delta}^\pm(y, a, x; \eta) = \mathbb{1}_{\{y \geq F_{x,a}^{-1}(\alpha^\pm)\}}$
- $\underline{\mu}^\pm(a, x; \eta) = \mathbb{E}[Y \underline{\Delta}^\pm(Y, A, X; \eta) \mid X = x, A = a]$
- $\bar{\mu}^\pm(a, x; \eta) = \mathbb{E}[Y \bar{\Delta}^\pm(Y, A, X; \eta) \mid X = x, A = a]$

Then, Equation 101 simplifies to

$$\mathbb{IF}\left(Q^{+,*}(a, x; \eta)\right) \quad (102)$$

$$= \frac{\mathbb{1}_{\{X=x\}}}{p(x)} b^- \left(\mathbb{1}_{\{A=a\}} - e(a, x) \right) \underline{\mu}^+(a, x; \eta) \quad (103)$$

$$+ \frac{\mathbb{1}_{\{X=x, A=a\}}}{p(x)e(a, x)} c^-(a, x; \eta) \left(Y \underline{\Delta}^+(Y, a, x; \eta) - \underline{\mu}^+(a, x; \eta) + F_{x,a}^{-1}(\alpha^+) (\underline{\Delta}^+(Y, a, x; \eta) - \alpha^+) \right) \quad (104)$$

$$+ \frac{\mathbb{1}_{\{X=x\}}}{p(x)} b^+ \left(\mathbb{1}_{\{A=a\}} - e(a, x) \right) \bar{\mu}^+(a, x; \eta) \quad (105)$$

$$+ \frac{\mathbb{1}_{\{X=x, A=a\}}}{p(x)e(a, x)} c^+(a, x; \eta) \left(Y \bar{\Delta}^+(Y, a, x; \eta) - \bar{\mu}^+(a, x; \eta) - F_{x,a}^{-1}(\alpha^+) (\bar{\Delta}^+(Y, a, x; \eta) - \alpha^+) \right) \quad (106)$$

$$= \frac{\mathbb{1}_{\{X=x\}}}{p(x)} \left\{ \left(\mathbb{1}_{\{A=a\}} - e(a, x) \right) \left(b^- \underline{\mu}^+(a, x; \eta) + b^+ \bar{\mu}^+(a, x; \eta) \right) \right. \quad (107)$$

$$+ \frac{\mathbb{1}_{\{A=a\}}}{e(a, x)} \left[\left(c^-(a, x; \eta) - c^+(a, x; \eta) \right) \left(F_{x,a}^{-1}(\alpha^+) (\underline{\Delta}^+(Y, a, x; \eta) - \alpha^+) \right) \right. \quad (108)$$

$$\left. \left. + c^-(a, x; \eta) \left(Y \underline{\Delta}^+(Y, a, x; \eta) - \underline{\mu}^+(a, x; \eta) \right) + c^+(a, x; \eta) \left(Y \bar{\Delta}^+(Y, a, x; \eta) - \bar{\mu}^+(a, x; \eta) \right) \right] \right\} \quad (109)$$

Finally, we can combine Equation 45 and Equation 109. That is, the efficient influence function of $V^{+,*}(\pi)$ is given by

$$\mathbb{IF}\left(V^{+,*}(\pi; \eta)\right) \quad (110)$$

$$= \sum_a \pi(a | X) Q^{+,*}(a, X) - V^{+,*}(\pi) \quad (111)$$

$$+ \sum_a \pi(a | X) \left(\mathbb{1}_{\{A=a\}} - e(a, X) \right) \left(b^- \underline{\mu}^+(a, X; \eta) + b^+ \bar{\mu}^+(a, X; \eta) \right) \quad (112)$$

$$+ \sum_a \pi(a | X) \frac{\mathbb{1}_{\{A=a\}}}{e(a, X)} \left[\left(c^-(a, X; \eta) - c^+(a, X; \eta) \right) \left(F_{x,a}^{-1}(\alpha^+) (\underline{\Delta}^+(Y, a, X; \eta) - \alpha^+) \right) \right. \quad (113)$$

$$\left. + c^-(a, X; \eta) \left(Y \underline{\Delta}^+(Y, a, X; \eta) - \underline{\mu}^+(a, X; \eta) \right) + c^+(a, X; \eta) \left(Y \bar{\Delta}^+(Y, a, X; \eta) - \bar{\mu}^+(a, X; \eta) \right) \right] \quad (114)$$

$$= -V^{+,*}(\pi) + \sum_a \pi(a | X) \left[Q^{+,*}(a, X) - e(a, X) \left(b^- \underline{\mu}^+(a, X; \eta) + b^+ \bar{\mu}^+(a, X; \eta) \right) \right] \quad (115)$$

$$+ \pi(A | X) \left(b^- \underline{\mu}^+(A, X; \eta) + b^+ \bar{\mu}^+(A, X; \eta) \right) \quad (116)$$

$$+ \frac{\pi(A | X)}{e(A, X)} \left[\left(c^-(A, X; \eta) - c^+(A, X; \eta) \right) \left(F_{X,A}^{-1}(\alpha^+) (\underline{\Delta}_{\alpha^+}(Y, A, X; \eta) - \alpha^+) \right) \right. \quad (117)$$

$$\left. + c^-(A, X; \eta) \left(Y \underline{\Delta}^+(Y, A, X; \eta) - \underline{\mu}^+(A, X; \eta) \right) + c^+(A, X; \eta) \left(Y \bar{\Delta}^+(Y, A, X; \eta) - \bar{\mu}^+(A, X; \eta) \right) \right] \quad (118)$$

We can derive the efficient estimator for the bounds of the value function through one-step bias correction via

$$V^{+,*}(\pi; \hat{\eta}) + \mathbb{P}_n \left\{ V^{+,*}(\pi; \hat{\eta}) \right\} \quad (119)$$

$$= \mathbb{P}_n \left\{ \sum_a \pi(a | X) \left[Q^{+,*}(a, X; \hat{\eta}) - \hat{e}(a, X) \left(b^- \hat{\underline{\mu}}^+(a, X; \hat{\eta}) + b^+ \hat{\bar{\mu}}^+(a, X; \hat{\eta}) \right) \right] \right\} \quad (120)$$

$$+ \pi(A | X) \left(b^- \hat{\underline{\mu}}^+(A, X; \hat{\eta}) + b^+ \hat{\bar{\mu}}^+(A, X; \hat{\eta}) \right) \quad (121)$$

$$+ \frac{\pi(A | X)}{\hat{e}(A, X)} \left[\left(c^-(A, X; \hat{\eta}) - c^+(A, X; \hat{\eta}) \right) \left(\hat{F}_{X,A}^{-1}(\alpha^+) (\underline{\Delta}^+(Y, A, X; \hat{\eta}) - \alpha^+) \right) \right. \quad (122)$$

$$\left. + c^-(A, X; \hat{\eta}) \left(Y \underline{\Delta}^+(Y, A, X; \hat{\eta}) - \hat{\underline{\mu}}^+(A, X; \hat{\eta}) \right) + c^+(A, X; \hat{\eta}) \left(Y \bar{\Delta}^+(Y, A, X; \hat{\eta}) - \hat{\bar{\mu}}^+(A, X; \hat{\eta}) \right) \right] \} \quad (123)$$

$$= \mathbb{P}_n \left\{ \sum_a \pi_{a,X} \left[\hat{Q}_{a,X}^{+,*} - \hat{e}_{a,X} \left(b^- \hat{\underline{\mu}}_{a,X}^+ + b^+ \hat{\bar{\mu}}_{a,X}^+ \right) \right] + \pi_{A,X} \left(b^- \hat{\underline{\mu}}_{A,X}^+ + b^+ \hat{\bar{\mu}}_{A,X}^+ \right) \right\} \quad (124)$$

$$+ \frac{\pi_{A,X}}{\hat{e}_{A,X}} \left[\left(\hat{c}_{A,X}^- - \hat{c}_{A,X}^+ \right) \left(\hat{F}_{X,A}^{-1}(\alpha^+) (\hat{\Delta}_{Y,A,X}^+ - \alpha^+) \right) + \hat{c}_{A,X}^- \left(Y \hat{\Delta}_{Y,A,X}^+ - \hat{\underline{\mu}}_{A,X}^+ \right) + \hat{c}_{A,X}^+ \left(Y \hat{\Delta}_{Y,A,X}^+ - \hat{\bar{\mu}}_{A,X}^+ \right) \right] \}$$

using our short-hand notation from the main paper.

□

F.3 LEARNING GUARANTEE: VALUE FUNCTION

Theorem F.3. Assume $|Y| \leq C_y$. Let $C_v := 2C_y(1 + \Gamma^{-1} + \Gamma)$ and let $R_n(\Pi)$ denote the empirical Rademacher complexity of the policy class Π . Then, with probability at least $1 - \delta$,

$$\sup_{\pi \in \Pi} \left\{ V(\pi) - \hat{V}^{+,*}(\pi) \right\} \leq 2C_v \left(R_n(\Pi) + \frac{5}{2} \sqrt{\frac{1}{2n} \log \frac{2}{\delta}} \right). \quad (125)$$

Equivalently, on the same high-probability event, simultaneously for all $\pi \in \Pi$,

$$V(\pi) \leq \hat{V}^{+,*}(\pi) + 2C_v \left(R_n(\Pi) + \frac{5}{2} \sqrt{\frac{1}{2n} \log \frac{2}{\delta}} \right). \quad (126)$$

Proof. We start by bounding the sharp upper bound $V^{+,*}(\pi)$ of the value function. By assumption, we have that $|Y| \leq C_y$. Hence, for any $\pi \in \Pi$, we can bound $V^{+,*}(\pi)$ via

$$|V^{+,*}(\pi)| \quad (127)$$

$$= \left| \int_{\mathcal{X}} \sum_a Q^{+,*}(a, x) \pi(a | x) dp(x) \right| \quad (128)$$

$$\leq \left| \sup_{(x,a) \in \mathcal{X} \times \mathcal{A}} Q^{+,*}(a, x) \right| \quad (129)$$

$$= \sup_{(x,a) \in \mathcal{X} \times \mathcal{A}} \left| \left((1 - \Gamma^{-1})e(a, x) + \Gamma^{-1} \right) \mathbb{E} \left[Y \mathbb{1}_{\{Y \leq F_{x,a}^{-1}(\alpha^+)\}} \mid X = x, A = a \right] \right| \quad (130)$$

$$+ \left| \left((1 - \Gamma)e(a, x) + \Gamma \right) \mathbb{E} \left[Y \mathbb{1}_{\{Y \geq F_{x,a}^{-1}(\alpha^+)\}} \mid X = x, A = a \right] \right| \quad (131)$$

$$\leq \sup_{(x,a) \in \mathcal{X} \times \mathcal{A}} C_y \left(\left((1 - \Gamma^{-1})e(a, x) + \Gamma^{-1} \right) + \left((1 - \Gamma)e(a, x) + \Gamma \right) \right) \quad (132)$$

$$\leq C_y \left(2 + 2\Gamma^{-1} + 2\Gamma \right) \quad (133)$$

$$= C_v. \quad (134)$$

Now, for the main result, we seek to find an upper bound for

$$\sup_{\pi \in \Pi} \hat{V}^{+,*}(\pi) - V(\pi). \quad (135)$$

By adding a zero and sublinearity of the supremum operator, we have that

$$\sup_{\pi \in \Pi} \hat{V}^{+,*}(\pi) - V(\pi) \leq \sup_{\pi \in \Pi} \left(\hat{V}^{+,*}(\pi) - V^{+,*}(\pi) \right) + \sup_{\pi \in \Pi} \left(V^{+,*}(\pi) - V(\pi) \right). \quad (136)$$

Further, by validity of our bounds, we know that

$$\sup_{\pi \in \Pi} V^{+,*}(\pi) - V(\pi) \geq 0, \quad (137)$$

such that we only need to focus on

$$D = \sup_{\pi \in \Pi} \hat{V}^{+,*}(\pi) - V^{+,*}(\pi). \quad (138)$$

Since

$$\hat{V}^{+,*}(\pi) = \frac{1}{n} \sum_{i=1}^n V_i^{+,*}(\pi) = \frac{1}{n} \sum_{i=1}^n \sum_a Q^{+,*}(X_i, a) \pi(a | X_i) \quad (139)$$

is a sample average with $|V_i^{+,*}(\pi)| \leq C_v$, we know that D satisfies the bounded difference with C_v/n . Hence, we can apply McDiarmid's inequality (McDiarmid, 1989), which yields

$$\mathbb{P} \left(D - \mathbb{E}[D] \geq \epsilon \right) \leq \underbrace{\exp \left(- \frac{2\epsilon^2 n}{C_v^2} \right)}_{=p_1}. \quad (140)$$

Then, solving for ϵ gives us

$$\epsilon = \sqrt{\frac{C_v^2}{2n} \log\left(\frac{1}{p_1}\right)}. \quad (141)$$

Hence, we know, with probability at least $1 - p_1$, that

$$D \leq \mathbb{E}[D] + \sqrt{\frac{C_v^2}{2n} \log\left(\frac{1}{p_1}\right)}. \quad (142)$$

As in (Athey & Wager, 2021; Frauen et al., 2024b; Hatt et al., 2022; Kallus & Zhou, 2018a), we express the flexibility of our policy class Π in terms of the Rademacher complexity. For this, we first note that a standard symmetrization argument yields

$$\mathbb{E}[D] \leq \mathbb{E} \left[\frac{1}{2^n} \sum_{\sigma \in \{-1, +1\}} \sup_{\pi \in \Pi} \left| \frac{1}{n} \sum_{i=1}^n \sigma_i V_i^{+,*}(\pi) \right| \right], \quad (143)$$

where $\sigma_i \sim_{\text{iid}} \text{Unif}\{-1, +1\}$. Then, using the Rademacher-comparison lemma (Ledoux & Talagrand, 1989), we have that

$$\mathbb{E}[D] \leq 2C_v \mathbb{E} \left[\mathcal{R}_n(\Pi) \right], \quad (144)$$

where

$$\mathcal{R}_n(\Pi) = \mathbb{E}_\sigma \left[\sup_{\pi \in \Pi} \frac{1}{n} \sum_{i=1}^n \sigma_i V_i^{+,*} \right] \quad (145)$$

is the empirical Rademacher complexity of policy class Π . Again, $\mathcal{R}_n(\Pi)$ satisfies bounded difference with $2/n$, such that we can apply McDiarmid's inequality (McDiarmid, 1989). This gives

$$\mathbb{P} \left(\mathcal{R}_n(\Pi) - \mathbb{E}[\mathcal{R}_n(\Pi)] \geq \epsilon \right) \leq \underbrace{\exp\left(-\frac{\epsilon^2 n}{2}\right)}_{=p_2}. \quad (146)$$

Solving for ϵ yields

$$\epsilon = \sqrt{\frac{2}{n} \log\left(\frac{1}{p_2}\right)}, \quad (147)$$

such that, with probability at least $1 - p_2$, we have that

$$\mathbb{E} \left[\mathcal{R}_n(\Pi) \right] \leq \mathcal{R}_n(\Pi) + \sqrt{\frac{2}{n} \log\left(\frac{1}{p_2}\right)}. \quad (148)$$

Combining Equation 148 with our previous result in Equation 142, we have, with probability of at least $1 - p_1 - p_2$, that

$$\sup_{\pi \in \Pi} \left(\hat{V}^{+,*}(\pi) - V^{+,*}(\pi) \right) \quad (149)$$

$$\leq 2C_v \mathcal{R}_n(\Pi) + \sqrt{\frac{C_v^2}{2n} \log\left(\frac{1}{p_1}\right)} + 2C_v \sqrt{\frac{2}{n} \log\left(\frac{1}{p_2}\right)} \quad (150)$$

$$= 2C_v \left(\mathcal{R}_n(\Pi) + \sqrt{\frac{1}{8n} \log\left(\frac{1}{p_1}\right)} + \sqrt{\frac{2}{n} \log\left(\frac{1}{p_2}\right)} \right). \quad (151)$$

Now let $p_1 = p_2 = \delta/2$. Then, we know that with probability at least $1 - \delta$,

$$\sup_{\pi \in \Pi} \left(\hat{V}^{+,*}(\pi) - V^{+,*}(\pi) \right) \quad (152)$$

$$\leq 2C_v \left(\mathcal{R}_n(\Pi) + \sqrt{\frac{1}{8n} \log\left(\frac{1}{\delta}\right)} + \sqrt{\frac{2}{n} \log\left(\frac{1}{\delta}\right)} \right) \quad (153)$$

$$= 2C_v \left(\mathcal{R}_n(\Pi) + \frac{5}{2} \sqrt{\frac{1}{2n} \log\left(\frac{2}{\delta}\right)} \right), \quad (154)$$

or, equivalently,

$$V^{+,*}(\pi) \leq \hat{V}^{+,*}(\pi) + 2C_v \left(\mathcal{R}_n(\Pi) + \frac{5}{2} \sqrt{\frac{1}{2n} \log\left(\frac{2}{\delta}\right)} \right) \quad (155)$$

for all $\pi \in \Pi$, which concludes the proof. \square

F.4 BOUND OF THE REGRET FUNCTION

Corollary F.4. *Given $Q^{\pm,*}(a, x)$ and our sensitivity constraints $\mathcal{P}(\Gamma)$ as in Proposition 4.1, an upper bound for the regret function $R_{\pi_0}(\pi)$ is given by*

$$R_{\pi}^+(\pi) = \int_{\mathcal{X}} \sum_a \left(Q^{+,*}(a, x)\pi(a | x) - Q^{-,*}(a, x)\pi_0(a | x) \right) dp(x). \quad (156)$$

Proof. Our proof follows similar steps as in the proof of Proposition 4.1. For clarity, we repeat the same steps such that the proof is self-contained.

Again, we start by noting that the upper bound on the regret function depends on the choice of our sensitivity constraints and, hence, the set of distributions $\mathcal{P}(\Gamma)$. Therefore, we can write that

$$R_{\pi_0}^+(\pi) = R_{\pi_0}^+(\pi, \mathcal{P}(\Gamma)). \quad (157)$$

Following similar steps as in Proposition 4.1, we can write

$$R_{\pi_0}^+(\pi) \quad (158)$$

$$= R_{\pi_0}^+(\pi, \mathcal{P}(\Gamma)) \quad (159)$$

$$= \sup_{\tilde{p} \in \mathcal{P}(\Gamma)} R_{\pi_0}(\pi, \tilde{p}) \quad (160)$$

$$= \sup_{\tilde{p} \in \mathcal{P}(\Gamma)} \left(V(\pi, \tilde{p}) - V(\pi_0, \tilde{p}) \right) \quad (161)$$

$$= \sup_{\tilde{p} \in \mathcal{P}(\Gamma)} \int_{\mathcal{X}} \sum_a \left(Q(a, x, \tilde{p})\pi(a | x) - Q(a, x, \tilde{p})\pi_0(a | x) \right) d\tilde{p}(x) \quad (162)$$

$$= \sup_{\tilde{p} \in \mathcal{P}(\Gamma)} \int_{\mathcal{X}} \sum_a \left(Q(a, x, \tilde{p})\pi(a | x) - Q(a, x, \tilde{p})\pi_0(a | x) \right) dp(x), \quad (163)$$

where Equation 163 again follows from $p(\mathcal{D}) = \tilde{p}(\mathcal{D})$ for all $\tilde{p} \in \mathcal{P}(\Gamma)$.

Since the optimal bounds $Q^{\pm,*}(a, x)$ are those $Q(a, x, \tilde{p})$, $\tilde{p} \in \mathcal{P}(\Gamma)$, for which the supremum/infimum are attained, we have that

$$Q(a, x, \tilde{p}) \leq \sup_{\tilde{p} \in \mathcal{P}(\Gamma)} Q(a, x, \tilde{p}) = Q^{+,*}(a, x, \mathcal{P}(\Gamma)) \quad (164)$$

and

$$Q(a, x, \tilde{p}) \geq \inf_{\tilde{p} \in \mathcal{P}(\Gamma)} Q(a, x, \tilde{p}) = Q^{-,*}(a, x, \mathcal{P}(\Gamma)) \quad (165)$$

for all $\tilde{p} \in \mathcal{P}(\Gamma)$. Then, since $Q^{\pm,*}(a, x) \in L^1(\pi, p)$, it follows by dominated convergence that

$$\sup_{\tilde{p} \in \mathcal{P}(\Gamma)} \int_{\mathcal{X}} \sum_a \left(Q(a, x, \tilde{p})\pi(a | x) - Q(a, x, \tilde{p})\pi_0(a | x) \right) dp(x) \quad (166)$$

$$= \int_{\mathcal{X}} \sup_{\tilde{p} \in \mathcal{P}(\Gamma)} \sum_a \left(Q(a, x, \tilde{p})\pi(a | x) - Q(a, x, \tilde{p})\pi_0(a | x) \right) dp(x) \quad (167)$$

$$\leq \int_{\mathcal{X}} \sum_a \sup_{\tilde{p} \in \mathcal{P}(\Gamma)} Q(a, x, \tilde{p})\pi(a | x) dp(x) - \int_{\mathcal{X}} \sum_a \inf_{\tilde{p} \in \mathcal{P}(\Gamma)} Q(a, x, \tilde{p})\pi_0(a | x) dp(x) \quad (168)$$

$$= \int_{\mathcal{X}} \sum_a \left(Q^{+,*}(a, x)\pi(a | x) - Q^{-,*}(a, x)\pi_0(a | x) \right) dp(x), \quad (169)$$

where Equation 168 follows from the sublinearity of the supremum/infimum operator. \square

F.5 EFFICIENT ESTIMATOR OF THE REGRET BOUND

Corollary F.5. Let $\mathbb{E}[|Y|^2] < \infty$ and $p(y | x, a)$ admit a continuous density bounded away from zero in a neighborhood of $F_{x,a}^{-1}(\alpha^+)$. This holds, for example, if the density $p(y | x, a)$ is strictly positive and continuous. Then, an efficient estimator for the upper bound of the regret function is given by

$$\hat{R}_{\pi_0}^+(\pi) \quad (170)$$

$$= \sum_{\pm \in \{-, +\}} \pm \mathbb{P}_n \left\{ \sum_a \pi_{a,X}^{\pm} \left[\hat{Q}_{a,X}^{\pm,*} - \hat{e}_{a,X} \left(b^{\mp} \hat{\mu}_{a,X}^{\pm} + b^{\pm} \hat{\mu}_{a,X}^{\pm} \right) \right] + \pi_{A,X}^{\pm} \left(b^{\mp} \hat{\mu}_{A,X}^{\pm} + b^{\pm} \hat{\mu}_{A,X}^{\pm} \right) \right. \quad (171)$$

$$\left. + \frac{\pi_{A,X}^{\pm}}{\hat{e}_{A,X}} \left[\left(\hat{c}_{A,X}^{\mp} - \hat{c}_{A,X}^{\pm} \right) \left(\hat{F}_{X,A}^{-1}(\alpha^{\pm}) (\hat{\Delta}_{Y,A,X}^{\pm} - \alpha^{\pm}) \right) + \hat{c}_{A,X}^{\mp} \left(Y \hat{\Delta}_{Y,A,X}^{\pm} - \hat{\mu}_{A,X}^{\pm} \right) + \hat{c}_{A,X}^{\pm} \left(Y \hat{\Delta}_{Y,A,X}^{\pm} - \hat{\mu}_{A,X}^{\pm} \right) \right] \right\}, \quad (172)$$

where we use $\pi^+ = \pi$ and $\pi^- = \pi_0$ for readability.

Proof. The upper bound of the regret function is given by

$$R_{\pi_0}^+ = V^{+,*}(\pi) - V^{-,*}(\pi_0), \quad (173)$$

where

$$V^{\pm,*}(\pi) = \int_{\mathcal{X}} \sum_a Q^{\pm,*}(a, x) \pi(a | x) dp(x). \quad (174)$$

By additivity of the efficient influence function, we know that

$$\mathbb{IF}\left(R_{\pi_0}^{+,*}(\pi)\right) = \mathbb{IF}\left(V^{+,*}(\pi)\right) - \mathbb{IF}\left(V^{-,*}(\pi_0)\right), \quad (175)$$

such that we can focus on both terms separately in Supplements F.5.1 and F.5.2 and then plug them together in Supplement F.5.3 in order to obtain our efficient estimator.

F.5.1 EFFICIENT INFLUENCE FUNCTION OF $V^{+,*}(\pi)$

We already have the efficient influence function of $V^{+,*}(\pi)$ from the proof of Theorem 4.3 in Supplement F.2. It is given by

$$\mathbb{IF}\left(V^{+,*}(\pi; \eta)\right) \quad (176)$$

$$= -V^{+,*}(\pi) + \sum_a \pi(a | X) \left[Q^{+,*}(a, X) - e(a, X) \left(b^- \underline{\mu}^+(a, X; \eta) + b^+ \bar{\mu}^+(a, X; \eta) \right) \right] \quad (177)$$

$$+ \pi(A | X) \left(b^- \underline{\mu}^+(A, X; \eta) + b^+ \bar{\mu}^+(A, X; \eta) \right) \quad (178)$$

$$+ \frac{\pi(A | X)}{e(A, X)} \left[\left(c^-(A, X; \eta) - c^+(A, X; \eta) \right) \left(F_{X,A}^{-1}(\alpha^+) (\underline{\Delta}_{\alpha^+}(Y, A, X; \eta) - \alpha^+) \right) \right] \quad (179)$$

$$+ c^-(A, X; \eta) \left(Y \underline{\Delta}^+(Y, A, X; \eta) - \underline{\mu}^+(A, X; \eta) \right) + c^+(A, X; \eta) \left(Y \bar{\Delta}^+(Y, A, X; \eta) - \bar{\mu}^+(A, X; \eta) \right). \quad (180)$$

F.5.2 EFFICIENT INFLUENCE FUNCTION OF $V^{-,*}(\pi)$

We can derive the sharp lower bound for the value function analogously to $V^{+,*}$. Again, we make use of the chain rule for deriving efficient influence function (Kennedy, 2022), a proof of which can be found in (Luedtke, 2024) (**Lemma S3**).

For this, let $\alpha^- = 1/(1 + \Gamma)$. Then, the sharp lower bound of the CAPO is given by

$$Q^{-,*}(a, x; \eta) \quad (181)$$

$$= \left((1 - \Gamma)e(a, x) + \Gamma \right) \mathbb{E} \left[Y \mathbb{1}_{\{Y \leq F_{X,A}^{-1}(\alpha^-)\}} \mid X = x, A = a \right] \quad (182)$$

$$+ \left((1 - \Gamma^{-1})e(a, x) + \Gamma^{-1} \right) \mathbb{E} \left[Y \mathbb{1}_{\{Y \geq F_{X,A}^{-1}(\alpha^-)\}} \mid X = x, A = a \right]. \quad (183)$$

Hence, the influence function of $Q^{-,*}(a, x)$ is given by

$$\mathbb{IF} \left(Q^{-,*}(a, x; \eta) \right) \quad (184)$$

$$= \frac{\mathbb{1}_{\{X=x\}}}{p(x)} (1 - \Gamma) \left(\mathbb{1}_{\{A=a\}} - e(a, x) \right) \mathbb{E} \left[Y \mathbb{1}_{\{Y \leq F_{X,A}^{-1}(\alpha^-)\}} \mid X = x, A = a \right] \quad (185)$$

$$+ \frac{\mathbb{1}_{\{X=x, A=a\}}}{p(a, x)} \left((1 - \Gamma)e(a, x) + \Gamma \right) \quad (186)$$

$$\times \left(Y \mathbb{1}_{\{Y \leq F_{x,a}^{-1}(\alpha^-)\}} - \mathbb{E} \left[Y \mathbb{1}_{\{Y \leq F_{X,A}^{-1}(\alpha^-)\}} \mid X = x, A = a \right] + F_{x,a}^{-1}(\alpha^-) (\mathbb{1}_{\{Y \leq F_{x,a}^{-1}(\alpha^-)\}} - \alpha^-) \right) \quad (187)$$

$$+ \frac{\mathbb{1}_{\{X=x\}}}{p(x)} (1 - \Gamma^{-1}) \left(\mathbb{1}_{\{A=a\}} - e(a, x) \right) \mathbb{E} \left[Y \mathbb{1}_{\{Y \geq F_{X,A}^{-1}(\alpha^-)\}} \mid X = x, A = a \right] \quad (188)$$

$$+ \frac{\mathbb{1}_{\{X=x, A=a\}}}{p(a, x)} \left((1 - \Gamma^{-1})e(a, x) + \Gamma^{-1} \right) \quad (189)$$

$$\times \left(Y \mathbb{1}_{\{Y \geq F_{x,a}^{-1}(\alpha^-)\}} - \mathbb{E} \left[Y \mathbb{1}_{\{Y \geq F_{X,A}^{-1}(\alpha^-)\}} \mid X = x, A = a \right] - F_{x,a}^{-1}(\alpha^-) (\mathbb{1}_{\{Y \geq F_{x,a}^{-1}(\alpha^-)\}} - \alpha^-) \right) \quad (190)$$

$$= \frac{\mathbb{1}_{\{X=x\}}}{p(x)} \left\{ \left(\mathbb{1}_{\{A=a\}} - e(a, x) \right) \left(b^+ \underline{\mu}^-(a, x; \eta) + b^- \bar{\mu}^-(a, x; \eta) \right) \right. \quad (191)$$

$$+ \frac{\mathbb{1}_{\{A=a\}}}{e(a, x)} \left[\left(c^+(a, x; \eta) - c^-(a, x; \eta) \right) \left(F_{x,a}^{-1}(\alpha^-) (\underline{\Delta}^-(Y, a, x; \eta) - \alpha^-) \right) \right. \quad (192)$$

$$\left. \left. + c^+(a, x; \eta) \left(Y \underline{\Delta}^-(Y, a, x; \eta) - \underline{\mu}^-(a, x; \eta) \right) + c^-(a, x; \eta) \left(Y \bar{\Delta}^-(Y, a, x; \eta) - \bar{\mu}^-(a, x; \eta) \right) \right] \right\} \quad (193)$$

Following the same steps as for Equation 45, the influence function of $V^{-,*}(\pi)$ is given by

$$\mathbb{IF} \left(V^{-,*}(\pi; \eta) \right) \quad (194)$$

$$= -V^{-,*}(\pi; \eta) + \sum_a \pi(a \mid X) \left[Q^{-,*}(a, X; \eta) - e(a, X) \left(b^+ \underline{\mu}^-(a, X; \eta) + b^- \bar{\mu}^-(a, X; \eta) \right) \right] \quad (195)$$

$$+ \pi(A \mid X) \left(b^+ \underline{\mu}^-(A, X; \eta) + b^- \bar{\mu}^-(A, X; \eta) \right) \quad (196)$$

$$+ \frac{\pi(A \mid X)}{e(A, X)} \left[\left(c^+(A, X; \eta) - c^-(A, X; \eta) \right) \left(F_{X,A}^{-1}(\alpha^-) (\underline{\Delta}^-(Y, A, X; \eta) - \alpha^-) \right) \right. \quad (197)$$

$$\left. \left. + c^+(A, X; \eta) \left(Y \underline{\Delta}^-(Y, A, X; \eta) - \underline{\mu}^-(A, X; \eta) \right) + c^-(A, X; \eta) \left(Y \bar{\Delta}^-(Y, A, X; \eta) - \bar{\mu}^-(A, X; \eta) \right) \right] \right]. \quad (198)$$

F.5.3 EFFICIENT ESTIMATOR OF $R_{\pi_0}^+(\pi)$

We can derive the efficient estimator for the bounds of the value function through one-step bias correction using our results from Supplements F.5.1 and F.5.2 via

$$V^{\pm,*}(\pi; \hat{\eta}) + \mathbb{P}_n \left\{ V^{\pm,*}(\pi; \hat{\eta}) \right\} \quad (199)$$

$$= \mathbb{P}_n \left\{ \sum_a \pi(a | X) \left[Q^{\pm,*}(a, X; \hat{\eta}) - \hat{e}(a, X) \left(b^{\mp} \hat{\underline{\mu}}^{\pm}(a, X; \hat{\eta}) + b^{\pm} \hat{\bar{\mu}}^{\pm}(a, X; \hat{\eta}) \right) \right] \right\} \quad (200)$$

$$+ \pi(A | X) \left(b^{\mp} \hat{\underline{\mu}}^{\pm}(A, X; \hat{\eta}) + b^{\pm} \hat{\bar{\mu}}^{\pm}(A, X; \hat{\eta}) \right) \quad (201)$$

$$+ \frac{\pi(A | X)}{\hat{e}(A, X)} \left[\left(c^{\mp}(A, X; \hat{\eta}) - c^{\pm}(A, X; \hat{\eta}) \right) \left(\hat{F}_{X,A}^{-1}(\alpha^{\pm})(\underline{\Delta}^{\pm}(Y, A, X; \hat{\eta}) - \alpha^{\pm}) \right) \right] \quad (202)$$

$$+ c^{\mp}(A, X; \hat{\eta}) \left(Y \underline{\Delta}^{\pm}(Y, A, X; \hat{\eta}) - \hat{\underline{\mu}}^{\pm}(A, X; \hat{\eta}) \right) + c^{\pm}(A, X; \hat{\eta}) \left(Y \bar{\Delta}^{\pm}(Y, A, X; \hat{\eta}) - \hat{\bar{\mu}}^{\pm}(A, X; \hat{\eta}) \right) \quad (203)$$

$$= \mathbb{P}_n \left\{ \sum_a \pi_{a,X} \left[\hat{Q}_{a,X}^{\pm,*} - \hat{e}_{a,X} \left(b^{\mp} \hat{\underline{\mu}}_{a,X}^{\pm} + b^{\pm} \hat{\bar{\mu}}_{a,X}^{\pm} \right) \right] + \pi_{A,X} \left(b^{\mp} \hat{\underline{\mu}}_{A,X}^{\pm} + b^{\pm} \hat{\bar{\mu}}_{A,X}^{\pm} \right) \right\} \quad (204)$$

$$+ \frac{\pi_{A,X}}{\hat{e}_{A,X}} \left[\left(\hat{c}_{A,X}^{\mp} - \hat{c}_{A,X}^{\pm} \right) \left(\hat{F}_{X,A}^{-1}(\alpha^{\pm})(\hat{\Delta}_{Y,A,X}^{\pm} - \alpha^{\pm}) \right) + \hat{c}_{A,X}^{\mp} \left(Y \hat{\Delta}_{Y,A,X}^{\pm} - \hat{\underline{\mu}}_{A,X}^{\pm} \right) + \hat{c}_{A,X}^{\pm} \left(Y \hat{\bar{\Delta}}_{Y,A,X}^{\pm} - \hat{\bar{\mu}}_{A,X}^{\pm} \right) \right] \quad (205)$$

using our short-hand notation from the main paper. Hence, the efficient estimator of the upper bound of the regret function is given by

$$\hat{R}_{\pi_0}^+(\pi) \quad (206)$$

$$= R_{\pi_0}^+(\pi; \hat{\eta}) + \mathbb{P}_n \left\{ \mathbb{IF} \left(R_{\pi_0}^{\pm,*}(\pi; \hat{\eta}) \right) \right\} \quad (207)$$

$$= \left(V^{\pm,*}(\pi; \hat{\eta}) - V^{\mp,*}(\pi_0; \hat{\eta}) \right) + \mathbb{P}_n \left\{ \mathbb{IF} \left(V^{\pm,*}(\pi; \hat{\eta}) \right) - \mathbb{IF} \left(V^{\mp,*}(\pi_0; \hat{\eta}) \right) \right\} \quad (208)$$

$$= \sum_{\pm \in \{+, -\}} \pm \mathbb{P}_n \left\{ \sum_a \pi_{a,X}^{\pm} \left[\hat{Q}_{a,X}^{\pm,*} - \hat{e}_{a,X} \left(b^{\mp} \hat{\underline{\mu}}_{a,X}^{\pm} + b^{\pm} \hat{\bar{\mu}}_{a,X}^{\pm} \right) \right] + \pi_{A,X}^{\pm} \left(b^{\mp} \hat{\underline{\mu}}_{A,X}^{\pm} + b^{\pm} \hat{\bar{\mu}}_{A,X}^{\pm} \right) \right\} \quad (209)$$

$$+ \frac{\pi_{A,X}^{\pm}}{\hat{e}_{A,X}} \left[\left(\hat{c}_{A,X}^{\mp} - \hat{c}_{A,X}^{\pm} \right) \left(\hat{F}_{X,A}^{-1}(\alpha^{\pm})(\hat{\Delta}_{Y,A,X}^{\pm} - \alpha^{\pm}) \right) + \hat{c}_{A,X}^{\mp} \left(Y \hat{\Delta}_{Y,A,X}^{\pm} - \hat{\underline{\mu}}_{A,X}^{\pm} \right) + \hat{c}_{A,X}^{\pm} \left(Y \hat{\bar{\Delta}}_{Y,A,X}^{\pm} - \hat{\bar{\mu}}_{A,X}^{\pm} \right) \right] \quad (210)$$

where we let $\pi^+ = \pi$ and $\pi^- = \pi_0$ for readability. \square

F.6 IMPROVEMENT GUARANTEE: REGRET FUNCTION

Corollary F.6. *Under the same assumption as in Theorem 4.4, simultaneously for all $\pi \in \Pi$ and fixed baseline policy $\pi_0 \in \Pi$, it holds, with probability $1 - \delta$, that*

$$R_{\pi_0}(\pi) \leq \hat{R}_{\pi_0}^+(\pi) + 4C_v \left(\mathcal{R}_n(\Pi) + \frac{5}{2} \sqrt{\frac{1}{2n} \log \left(\frac{2}{\delta} \right)} \right), \quad (211)$$

where $C_v = 2C_y(1 + \Gamma^{-1} + \Gamma)$ and $\mathcal{R}_n(\pi)$ is the empirical Rademacher complexity of policy class Π .

Proof. In order to show the main result, we note that

$$R_{\pi_0}^+(\pi) = V^{+,*}(\pi) - V^{-,*}(\pi_0) \quad (212)$$

for arbitrary $\pi, \pi_0 \in \Pi$.

From Theorem 4.4, we know that

$$V^{+,*}(\pi) \leq \hat{V}^{+,*}(\pi) + 2C_v \left(\mathcal{R}_n(\Pi) + \frac{5}{2} \sqrt{\frac{1}{2n} \log \left(\frac{2}{\delta} \right)} \right). \quad (213)$$

Since $\pi, \pi_0 \in \Pi$ are arbitrary, we can repeat the same arguments for $V^{-,*}(\pi_0)$ and obtain

$$V^{-,*}(\pi_0) \geq \hat{V}^{-,*}(\pi_0) - 2C_v \left(\mathcal{R}_n(\Pi) + \frac{5}{2} \sqrt{\frac{1}{2n} \log \left(\frac{2}{\delta} \right)} \right). \quad (214)$$

Then, we conclude the proof by

$$R_{\pi_0}(\pi) \quad (215)$$

$$= V(\pi) - V(\pi_0) \quad (216)$$

$$\leq V^{+,*}(\pi) - V^{-,*}(\pi_0) \quad (217)$$

$$= \left[\hat{V}^{+,*}(\pi) + 2C_v \left(\mathcal{R}_n(\Pi) + \frac{5}{2} \sqrt{\frac{1}{2n} \log \left(\frac{2}{\delta} \right)} \right) \right] - \left[\hat{V}^{-,*}(\pi_0) - 2C_v \left(\mathcal{R}_n(\Pi) + \frac{5}{2} \sqrt{\frac{1}{2n} \log \left(\frac{2}{\delta} \right)} \right) \right] \quad (218)$$

$$= \hat{R}_{\pi_0}^+(\pi) + 4C_v \left(\mathcal{R}_n(\Pi) + \frac{5}{2} \sqrt{\frac{1}{2n} \log \left(\frac{2}{\delta} \right)} \right). \quad (219)$$

□

G IMPLEMENTATION DETAILS

We summarize the neural instantiations of all estimators in Section 5.

Runtime: Training the first and second stage models for our method took in total approximately 20 seconds using $n = 1000$ synthetic data samples and a standard computer with AMD Ryzen 7 Pro CPU and 32GB of RAM.634. All baselines have a comparable runtime.

Nuisance function	Hyperparameter	Configuration	Standard methods		Kallus & Zhou (2018a; 2021)	Plug-in sharp (ours)	Efficient sharp (ours)
			IPW estimator	DR estimator			
Propensity score	Hidden layers	3					
	Layer size	{64, 64, 32}					
	Hidden activation	ReLU	✓	✓	✓	✗	✓
	Learning rate	0.001					
	Number of epochs	300					
	Batch size	64					
Conditional quantile function	Hidden layers	3					
	Layer size	{64, 64, 32}					
	Hidden activation	ReLU	✗	✗	✗	✗	✓
	Learning rate	0.001					
	Number of epochs	300					
	Batch size	64					
(Truncated) outcome regression model	Hidden layers	3					
	Layer size	{64, 64, 32}					
	Hidden activation	ReLU	✗	✓	✗	✓	✓
	Learning rate	0.001					
	Number of epochs	300					
	Batch size	64					
Parametric policy	Hidden layers	3					
	Layer size	{64, 64, 32}					
	Hidden activation	ReLU	✓	✓	✓	✓	✓
	Learning rate	0.001					
	Number of epochs	300					
	Batch size	64					

Table 4: Neural instantiations of estimated nuisance functions $\hat{\eta}$ and parametric policy π_θ . To ensure a fair comparison, all methods share the same nuisance function where applicable. For all models, we set the split parameter for training the nuisance and the policy model to $\rho = 0.5$.