RISK-AWARE A/B TESTING THROUGH EFFICIENT INFERENCE ON TREATMENT EFFECT DISTRIBUTIONS WITHOUT A MARGIN ASSUMPTION

Anonymous authors
Paper under double-blind review

ABSTRACT

In many fields, including healthcare, marketing, and online platform design, A/B tests are used to evaluate new treatments and make launch decisions based on average treatment effect (ATE) estimates. But ATE-based decision-making overlooks distributional risks, such as a large fraction of individuals affected negatively by the treatment. These distributional risks can be detected by estimating the entire cumulative distribution function of the treatment effect, rather than only its mean. Prior work on estimating such treatment effect distributions has estimated partial identification bounds known as Makarov bounds under restrictive assumptions on the outcome distribution. In particular, these assumptions require that the treatment effect cannot be constant for any subset of experimental units, an assumption that is often either violated or nearly-violated in practice. In this paper, we propose a novel method for asymptotically optimal estimation of and statistically-valid inference on Makarov bounds under more realistic assumptions than existing methods. Our main technical contributions are to a. propose smoothed surrogates for the Makarov bounds, b. derive semiparametrically efficient estimators of these surrogates, and c. propose a procedure for optimal selection of the smoothing parameters. We show empirically on synthetic and semi-synthetic datasets that our estimators achieve a better bias-variance trade-off and lower mean-squared error than existing estimators. Finally, we deploy our method on real A/B test data from a large social media platform, and show how estimates of the treatment effect distribution can inform decision-making.

1 INTRODUCTION

Randomized controlled trials (RCTs) and A/B tests are crucial for evaluating the impact of treatments or product changes on a target outcome. Examples include medical treatments aimed at improving patient health (Feuerriegel et al., 2024), advertising placement to increase revenue (Varian, 2016), and interventions on digital platforms to boost user engagement (Swaminathan & Joachims, 2015). By randomly assigning units to treatment and control, A/B tests provide unbiased estimates of the average difference in outcome between receiving and not receiving treatment, also known as the average treatment effect (ATE). A common decision rule is to deploy the new treatment when the ATE is statistically significantly positive (Athey et al., 2020).

In many settings though, the ATE is not sufficient for decision-making purposes and practitioners in fact need estimates of distributional quantities more complex than averages (Kallus & Zhou, 2021). In sensitive applications, it is vital to understand the fraction of individuals that are affected negatively by the treatment, even when the average treatment effect is positive. For example, a change to a platform's ranking algorithm may lift overall engagement yet reduce engagement for new or low-activity users, increasing early churn. Decision-making in contexts like these requires estimating the *entire distribution of the treatment effect*, rather than merely its mean.

The problem of estimating the treatment effect distribution is complicated by the fact that, even in an RCT, the joint distribution of the potential outcomes is not identified (Fan & Park, 2010). Intuitively, because we never observe both the control and treatment outcome for the same unit, we can only estimate means and other functionals of the marginal control and marginal treatment distribution,

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

not of the joint distribution. As a result, the treatment effect distribution is only *partially identified*: its cumulative distribution function (c.d.f.) is guaranteed to lie between the sharp *Makarov bounds* (Makarov, 1982), and these bounds can be identified and estimated from RCT data. The Makarov bounds correspond to the best and worst case joint distributions of treatment and control that are consistent with the observed marginal distributions. As such, they are the solutions to optimization problems written in terms of suprema and infima over the outcome space, are non-smooth, and cannot be estimated using standard tools for efficient estimation of statistical functionals (Kennedy, 2024).

To address these issues, existing methods for estimating the Makarov bounds either address stylized binary-outcome settings (Kallus et al., 2022; Zhang & Richardson, 2025), or rely on a *margin assumption* (Semenova, 2025). The margin assumption requires the suprema and infima in the Makarov bounds to be attained at unique points, and is often violated in practice. For example, if the treatment effect is constant or nearly constant for a subset of users, the margin assumption will fail or nearly fail to hold. This would be the case for a medical treatment that works well on moderately ill patients but not at all on severely ill patients (since the treatment effect would be a constant zero for the severely ill patients) or for a platform change that increases engagement for very active users but has little to no effect on new or less active users (since the treatment effect would be a constant zero for the new and less active users). Alternatively, a notification treatment on an online platform may lead to one additional visit from all users, again violating this assumption. In settings like these, existing estimators are biased and high-variance, leading to inaccurate estimates and sub-optimal decisions.

In this paper, we propose a novel method for inference on Makarov bounds for the treatment-effect distribution that does not rely on a binary outcome assumption or a margin assumption, and is thus more reliable in practice. Our key idea is to replace the non-smooth suprema and infima that define the bounds with smooth surrogates, allowing us to apply standard semiparametric efficiency theory to obtain debiased estimators. We then derive an upper bound on the bias introduced by smoothing and incorporate it into both the selection of the smoothing parameters and the construction of our confidence intervals. Our method offers three key advantages: (i) it provides lower mean-squared error (MSE) estimates of Makarov bounds when the margin assumption is violated (as would be the case if treatment effects are constant for any units); (ii) it is model-agnostic and can be combined with arbitrary black-box learners to exploit covariates and tighten the bounds; and (iii) it admits data-driven procedures for selecting the smoothing level optimally.

Our contributions are[1]: (i) We propose debiased estimators of smoothed versions of the Makarov bounds and quantify the induced smoothing bias; (ii) We propose a new data-driven procedure to tune the smoothing parameters of our method; and (iii) We validate the effectiveness of our method using synthetic, semi-synthetic, and real-world data.

To summarize, methods for estimating treatment effect distributions are necessary for high-stakes decision making, but existing methods are sub-optimal in even the constant treatment effect model, arguably the most basic causal inference setting. This paper provides new methods that can be used for optimal estimation and high-stakes decision making in realistic settings excluded by previous work.

2 Related work

Treatment-effect distributions. Prior work either constructs predictive intervals for individual treatment effects using conformal methods (Lei & Candès, 2021; Alaa et al., 2023; Schröder et al., 2025), or performs inference on the cumulative distribution function of treatment effects. Because the joint distribution of potential outcomes is not identified under standard causal assumptions (Rubin, 1974; Robins, 1986; Fan & Park, 2010), the latter line typically either imposes additional identifying assumptions (Post & Van Den Heuvel, 2025) or estimates partial-identification bounds. For general outcomes, the sharp bounds are the Makarov bounds (Makarov, 1982); existing estimators are mostly plug-in procedures (Fan & Park, 2010; Ruiz & Padilla, 2022; Fava, 2024; Lee, 2024; Cui & Han, 2025; Liang & Wu, 2025) and therefore do not directly target semiparametric efficiency.

Non-smooth causal estimands. Makarov bounds involve suprema and infima, creating the same kind of non-differentiability that appears in policy-value, intersection-bound, and instrumental-variable

[1]Code: `https://anonymous.4open.science/r/SmoothedMakarovBounds-632B`.

problems (Luedtke & Van Der Laan, 2016; Semenova, 2025; Levis et al., 2025). One approach assumes a unique optimizer through a margin condition (Kitagawa & Tetenov, 2018; Melnychuk et al., 2024); another replaces the non-smooth operator with a smooth surrogate (Park, 2024; Levis et al., 2025). Our work applies the smoothing approach to Makarov bounds, avoiding the margin assumption that fails under constant or nearly constant treatment effects. Appendix Table 2 summarizes the closest methodological connections.

3 Problem setup

3.1 Setting

Data. We consider a causal inference setting with either randomized or observational data. In this setting there is a population distribution $Z = (X, A, Y) \sim \mathbb{P}$, where $X \in \mathcal{X} \subseteq \mathbb{R}^d$ are covariates, $A \in \{0, 1\}$ is a binary treatment, and $Y \in \mathbb{R}$ is a continuous outcome of interest that is observed after taking the action A. For example, X may be user demographics, A may be a binary decision of whether a policy is implemented, and Y may be a user engagement metric. We provide extensions of all our results to discrete outcomes in Appendix C. We also assume that we have access to a dataset $\mathcal{D} = \{(x_i, a_i, y_i)\}_{i=1}^n$ of size $n \in \mathbb{N}$ sampled i.i.d. from $\mathbb{P}$.

Notation. We define the *propensity score* as $\pi(x) = \mathbb{P}(A = 1 \mid X = x)$. The propensity score characterizes the treatment assignment mechanism and is often known in randomized experiments, e.g., $\pi(x) = 0.5$. Furthermore, we define the *response distributions* as the conditional outcome c.d.f.s $F_a(y|x) = \mathbb{P}(Y \leq y | X = x, A = a)$ for $a \in 0, 1$. We let $\eta = \{\pi, F_1, F_0\}$ be the collection of these *nuisance functions*. For any $v \in \mathbb{R}$, we write $v_+ = \max(v, 0)$ and $v_- = \min(v, 0)$.

Target estimand. We work in the potential outcomes framework of Rubin (1974) and denote $Y(a)$ as the potential outcome corresponding to the treatment $A = a$. The parameter of interest is the cumulative distribution function of the treatment effect evaluated at a single point,

$$\rho(\delta) = \mathbb{P}(Y(1) - Y(0) \leq \delta). \tag{1}$$

The c.d.f. $\rho(\delta)$ characterizes the entire treatment effect distribution. For example, $\rho(0)$ is the fraction of users negatively affected by the treatment (Kallus et al., 2022).

Partial identification bounds. We impose the following standard causal inference assumptions necessary for identification (Rubin, 1974).

Assumption 3.1 (Standard causal inference assumptions)**.** For all $a \in \{0, 1\}$ and $x \in \mathcal{X}$ we have: (i) *consistency*, $Y(a) = Y$ whenever $A = a$; (ii) *overlap*, $0 < \pi(x) < 1$ whenever $\mathbb{P}(X = x) > 0$; and (iii) *ignorability*, $A \perp Y(1), Y(0) \mid X = x$.

Conditions (ii) and (iii) are automatically satisfied in randomized experiments where the propensity score π is known and can be controlled (and is often the constant function $\pi(x) = 0.5$). Condition (i) prohibits interference between individuals and excludes the possibility of spillover effects. Together, these assumptions allow us to *partially identify* the treatment effect distribution $\rho(\delta)$ via the Makarov bounds (Makarov, 1982), so that

$$\rho^-(\delta) \leq \rho(\delta) \leq \rho^+(\delta), \tag{2}$$

where

$$\rho^-(\delta) = \mathbb{E}\left[\sup_{y \in \mathcal{Y}} \left(F_1(y|X) - F_0(y - \delta|X)\right)_+\right], \rho^+(\delta) = \mathbb{E}\left[\inf_{y \in \mathcal{Y}} \left(F_1(y|X) - F_0(y - \delta|X)\right)_-\right] + 1. \tag{3}$$

The Makarov bounds $\rho^-(\delta)$ and $\rho^+(\delta)$ only depend on the marginal response distributions F_1 and F_0 and are thus identified from the observed distribution $\mathbb{P}$. Intuitively, the Makarov bounds quantify the stochastically smallest and largest distributions that can be obtained by maximizing or minimizing over joint distributions of the potential outcomes that are compatible with the observed marginals F_1 and F_0. (They are also generalizations of the bounds from (Kallus, 2022) for the fraction of negatively affected users in the case of binary outcomes.)

3.2 Background on estimating Makarov bounds

A natural baseline is the *plug-in estimator*: estimate the nuisance functions $\hat{\eta} = (\hat{F}_1, \hat{F}_0, \hat{\pi})$ and substitute them into Eq. (3). Plug-in estimators are simple and flexible, but their nuisance-estimation

bias complicates valid confidence intervals and prevents semiparametric efficiency (Kennedy, 2024). Debiased estimators correct this by adding an empirical average of an efficient influence function, yielding asymptotically linear estimators when the target functional is sufficiently smooth.

For Makarov bounds, existing debiased *envelope estimators* choose an optimizer y^* of the supremum or infimum and plug the corresponding influence function into the one-step estimator (Melnychuk et al., 2024; Semenova, 2025). This construction relies on pathwise differentiability (van der Laan & Rubin, 2006), which in turn requires the optimizer to be unique. The resulting margin assumption is violated whenever the difference $F_1(y \mid X) - F_0(y - \delta \mid X)$ has a plateau, as happens for constant or nearly constant treatment effects, discrete outcomes, zero-inflated outcomes, or treatments that affect only a subset of the outcome range; see Fig. 1 for an illustration. In such cases the envelope estimator can be unstable, biased, and inefficient. Our goal is to retain the debiasing benefits while removing this uniqueness requirement.

4 Estimating Makarov bounds in the presence of margin violations

In this section, we present our methodology for estimating Makarov bounds when the margin assumption is violated. Our strategy proceeds in three steps. First, in Sec. 4.1, we introduce a smooth surrogate approximation of the Makarov bounds that is amenable to debiased estimation. Next, in Sec. 4.2, we derive debiased estimators for these smoothed bounds and construct confidence intervals that achieve valid coverage. Finally, in Sec. 4.3, we propose a data-driven procedure to select the smoothing parameters.

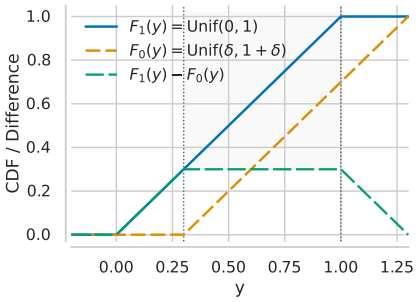

Figure 1: Example of margin assumption violation. The treatment effect is constant, and so the treatment c.d.f. is a shifted version of the control c.d.f.; the difference between the two has a plateau, so the supremum in Eq. (3) is not attained at a unique point.

4.1 A smooth approximation for the Makarov bounds

There are two features of the Makarov bounds in Eq. (3) that prevent debiased inference without the margin assumption: (i) the supremum and infumum operators, $\sup_{y\in\mathcal{Y}}$ and $\inf_{y\in\mathcal{Y}}$, and (ii) the positivity and negativity operators, $(\cdot)_+ = \max(\cdot, 0)$ and $(\cdot)_- = \min(\cdot, 0)$. Both components are non-smooth in the presence of margin violations. Intuitively, when the optimizer has multiple solutions, even small estimation errors in the nuisance functions can lead to large and unstable variation in the estimator, as well as bias in the wrong direction. This instability is what invalidates asymptotic efficiency and leads to inaccurate estimates.

Bound approximations. The key idea is to replace the problematic components identified above with smooth approximations that have unique values, even when the original functions do not. This smoothing restores pathwise differentiability, allowing us to construct a debiased estimator and valid confidence intervals without requiring the margin assumption. For the suprema and infima, we adapt the approximation from (Levis et al., 2025) for discrete maxima to continuous variables and define the continuous log-sum-exp operator (LSE) $g_{t_1} : \mathcal{F} \to \mathbb{R}$ as

$$g_{t_1}(f) = \frac{1}{t_1} \log\left(\int \exp(t_1 f(y))\, \mathrm{d}y\right). \tag{4}$$

For (ii), we replace v_+ and v_- with the softplus function $h_{t_2} : \mathbb{R} \to \mathbb{R}$ via $h_{t_2}(u) = \frac{1}{t_2}\log\left(1 + e^{t_2 u}\right)$. The following lemma shows that replacing the components in (i) and (ii) with g_{t_1} and h_{t_2} gives low-error approximations of the Makarov bounds.

Lemma 4.1 (Makarov bound approximations)**.** *Assume that $\mathcal{Y}$ is compact with finite positive Lebesgue measure $|\mathcal{Y}|$, and, for the fixed δ under consideration, that $y \mapsto F_1(y \mid x) - F_0(y - \delta \mid x)$ is continuous on $\mathcal{Y}$ for $\mathbb{P}$-almost every x. For any $t_1, t_2 > 0$, we define the smoothed Makarov bounds via*

$$\rho^-_{t_1,t_2}(\delta) = \mathbb{E}\left[\frac{1}{t_2}\log\left(1 + I_{t_1,\delta}(X)^{\frac{t_2}{t_1}}\right)\right] \text{ and } \rho^+_{t_1,t_2}(\delta) = 1 - \mathbb{E}\left[\frac{1}{t_2}\log\left(1 + I_{-t_1,\delta}(X)^{\frac{t_2}{t_1}}\right)\right], \tag{5}$$

where $I_{t,\delta}(x) = \int_{\mathcal{Y}} \exp\big(t[F_1(y|x) - F_0(y-\delta|x)]\big)\, dy$.

Then, it holds that

$$\rho^-_{t_1,t_2}(\delta) - b(t_1,t_2) \le \rho^-(\delta) \text{ and } \rho^+(\delta) \le \rho^+_{t_1,t_2}(\delta) + b(t_1,t_2), \tag{6}$$

where $b(t_1,t_2) = \frac{\log(2)}{t_2} + \frac{(\log|\mathcal{Y}|)_+}{t_1}$ *is a one-sided approximation correction. Furthermore,* $\rho^-_{t_1,t_2}(\delta) \to \rho^-(\delta)$ *and* $\rho^+_{t_1,t_2}(\delta) \to \rho^+(\delta)$ *as* $t_1, t_2 \to \infty$.

Proof. See Appendix A. □

Intuitively, this smoothing replaces the "hard max" operation of the supremum with a "soft max." Instead of arbitrarily selecting one of multiple equally optimal points, the approximation takes a smooth weighted average that changes gradually as the nuisance estimates shift. This stabilizes the estimator and leads to asymptotically normal debiased estimators.

The compact-support assumption can, for example, be enforced by winsorizing the observed outcome, as is common for heavy-tailed responses. In that case, the estimand should be interpreted as the treatment-effect distribution of the winsorized outcome.

4.2 Debiased estimators

We now derive the debiased estimators for the smoothed bounds in Lemma 4.1. The following result gives the efficient influence functions (EIF) of the lower and upper approximation.

Theorem 4.2. *Fix* $t_1, t_2 > 0$ *and* $\delta \in \mathbb{R}$. *The efficient influence functions of the smoothed Makarov bounds* $\rho^-_{t_1,t_2}(\delta)$ *and* $\rho^+_{t_1,t_2}(\delta)$ *are given by*

$$\Psi^-_{t_1,t_2,\delta}\left(Z,\eta,\rho^-_{t_1,t_2}(\delta)\right) = \sigma^-_{t_1,t_2,\delta}(X)\left[\frac{A}{\pi(X)}\,\Phi^1_{t_1,\delta}(X,Y) \;-\; \frac{1-A}{1-\pi(X)}\,\Phi^0_{t_1,\delta}(X,Y)\right] \tag{7}$$

$$+\frac{1}{t_2}\log\left(1 + I_{t_1,\delta}(X)^{\frac{t_2}{t_1}}\right) - \rho^-_{t_1,t_2}(\delta), \tag{8}$$

$$\Psi^+_{t_1,t_2,\delta}\left(Z,\eta,\rho^+_{t_1,t_2}(\delta)\right) = \sigma^+_{t_1,t_2,\delta}(X)\left[\frac{A}{\pi(X)}\,\Phi^1_{-t_1,\delta}(X,Y) \;-\; \frac{1-A}{1-\pi(X)}\,\Phi^0_{-t_1,\delta}(X,Y)\right] \tag{9}$$

$$+1 - \frac{1}{t_2}\log\left(1 + I_{-t_1,\delta}(X)^{\frac{t_2}{t_1}}\right) - \rho^+_{t_1,t_2}(\delta), \tag{10}$$

where we denote $\Phi^a_{t_1,\delta}(X,Y) = \int_{\mathcal{Y}} w_{t_1,\delta}(y \mid X)\left(\mathbf{1}\{Y \le y-(1-a)\delta\} - F_a(y-(1-a)\delta \mid X)\right) dy$ *and let*

$$w_{t_1,\delta}(y|x) = \frac{\exp\big(t_1(F_1(y|x) - F_0(y-\delta|x))\big)}{I_{t_1,\delta}(x)}, \sigma^\pm_{t_1,t_2,\delta}(x) = \frac{I_{\mp t_1,\delta}(x)^{\,t_2/t_1}}{1 + I_{\mp t_1,\delta}(x)^{\,t_2/t_1}}. \tag{11}$$

Proof. See Appendix A. □

The structure of the smoothed EIFs from Theorem 4.2 is similar to that of the envelope EIF. The crucial difference is that the terms involving the argmax/argmin y^* are replaced by smooth weighted integrals $\Phi^a_{t_1,\delta}(X,Y)$, and the indicator $\mathbf{1}(d_{\delta,\eta}(y^*|X) > 0)$ is replaced with one of the smooth approximations $\sigma^\pm_{t_1,t_2,\delta}(X)$.

Using Theorem 4.2, we can now obtain debiased estimators for the smoothed Makarov bounds. Following the usual one-step construction, these are given by

$$\hat\rho^-_{t_1,t_2}(\delta) = \frac{1}{n}\sum_{i=1}^n \frac{1}{t_2}\log\left(1 + \hat I_{t_1,\delta}(x_i)^{\frac{t_2}{t_1}}\right) + \hat\sigma^-_{t_1,t_2,\delta}(x_i)\left[\frac{a_i - \hat\pi(x_i)}{\hat\pi(x_i)(1-\hat\pi(x_i))}\,\hat\Phi^{a_i}_{t_1,\delta}(x_i,y_i)\right], \tag{12}$$

$$\hat\rho^+_{t_1,t_2}(\delta) = \frac{1}{n}\sum_{i=1}^n 1 - \frac{1}{t_2}\log\left(1 + \hat I_{-t_1,\delta}(x_i)^{\frac{t_2}{t_1}}\right) + \hat\sigma^+_{t_1,t_2,\delta}(x_i)\left[\frac{a_i - \hat\pi(x_i)}{\hat\pi(x_i)(1-\hat\pi(x_i))}\,\hat\Phi^{a_i}_{-t_1,\delta}(x_i,y_i)\right].$$

Correcting for approximation bias. The debiased estimators above are efficient and asymptotically normal, and thus admit standard normal confidence intervals. Those intervals, however are centered

around the smoothed Makarov bounds, and so it remains to translate these back to intervals for the non-smoothed Makarov bounds. We can do this by using the upper bound from Lemma 4.1 on the bias introduced by smoothing, $b(t_1, t_2)$. We then add this term to the one-sided confidence intervals for the upper and lower Makarov bounds, yielding valid confidence intervals as shown below.

Corollary 4.3 (Asymptotic confidence interval for the treatment effect distribution)**.** *Assume the nuisance estimators $\hat{\eta}$ are obtained on an independent sample from $\mathcal{D}$ (e.g., via sample splitting or cross-fitting). Under standard regularity and rate conditions for orthogonal/debiased estimation (see e.g., Chernozhukov et al. (2018)), it holds for each fixed δ that*

$$\lim_{n\to\infty} \mathbb{P}\Big(c^{-}_{t_1,t_2,\delta,\alpha}(\mathcal{D}_n) \le \rho(\delta) \le c^{+}_{t_1,t_2,\delta,\alpha}(\mathcal{D}_n)\Big) \;\ge\; 1-\alpha, \tag{13}$$

where $c^{\pm}_{t_1,t_2,\delta,\alpha}(\mathcal{D}_n)$ is given by

$$\hat{\rho}^{\pm}_{t_1,t_2}(\delta) \pm z_{1-\alpha/2}\sqrt{\frac{1}{n^2}\sum_{i=1}^{n}\Big(\Psi^{\pm}_{t_1,t_2,\delta}\big(z_i,\hat{\eta},\hat{\rho}^{\pm}_{t_1,t_2}(\delta)\big)\Big)^2} \pm b(t_1,t_2), \tag{14}$$

$b(t_1, t_2)$ is as in Lemma 4.1, $\Psi^{\pm}_{t_1,t_2,\delta}$ are given in Theorem 4.2, and $z_{1-\alpha/2}$ is the $1-\alpha/2$ standard normal quantile.

Proof. See Appendix A. □

4.3 Smoothing parameter selection

Bias-variance trade-off. The estimators in Eq. (12) depend on the smoothing parameters t_1 and t_2, which control both the amount of smoothing and the quality of the approximation. For fixed t_1, t_2, a useful leading-order criterion for understanding this trade-off is

$$\mathcal{R}^{-}_{n}(t_1,t_2;\delta) := \underbrace{\big(\rho^{-}_{t_1,t_2}(\delta)-\rho^{-}(\delta)\big)^2}_{\text{squared smoothing bias}} + \underbrace{\frac{1}{n}\,\mathbb{E}\Big[\Psi^{-}_{t_1,t_2,\delta}\big(Z,\eta,\rho^{-}_{t_1,t_2}(\delta)\big)^2\Big]}_{\text{asymptotic variance}/n}. \tag{15}$$

The first term decreases as the smoothed target approaches the original Makarov bound, while stronger smoothing can reduce the variance of the estimator. Because the true smoothing bias is unknown, we use data-driven proxies and stability criteria to choose t_1 and t_2. We consider two such procedures: (i) a bias-variance proxy based on the deterministic correction from Lemma 4.1, and (ii) a Lepski-style stability rule. Our full procedure is shown in Algorithm 1.

Bias-variance proxy. Using Lemma 4.1, we can replace the smoothing bias term in Eq. (15) with $b(t_1, t_2) = \frac{\log(2)}{t_2} + \frac{(\log|\mathcal{Y}|)_+}{t_1}$, leading to the data-driven tuning rule

$$\arg\min_{t_1,t_2} b(t_1,t_2)^2 + \frac{1}{n}\sum_{i=1}^{n}\Psi^{-}_{t_1,t_2,\delta}\big(z_i,\hat{\eta},\hat{\rho}^{-}_{t_1,t_2}(\delta)\big)^2, \tag{16}$$

which can be minimized via grid search.

Lepski's method. Lepski's method (Lepskii, 1992; 1993) selects the smallest amount of smoothing whose estimate is statistically indistinguishable from estimates obtained with *less* smoothing, thereby controlling approximation bias without paying unnecessary variance. We start by ranking all candidate pairs $(t_1^{(k)}, t_2^{(k)})$ within the grid by the size of their associated bias term $B_k = b(t_1^{(k)}, t_2^{(k)})$ in *descending* order (more smoothing first). Scanning in this order, for each k we compare the corresponding estimates $\hat{\rho}^{-}_{t_1^{(k)},t_2^{(k)}}(\delta)$ to a small set of less–smoothed candidates r with $B_r < B_k$ and accept the first k such that $\big|\hat{\rho}^{-}_{t_1^{(k)},t_2^{(k)}}(\delta) - \hat{\rho}^{-}_{t_1^{(r)},t_2^{(r)}}(\delta)\big| \;\le\; z_{1-\alpha/2}\,\widehat{\mathrm{SE}}(k,r)$ for all such r, where the tolerance uses EIF differences

$$\widehat{\mathrm{SE}}(k,r) = 1/\sqrt{n}\,\mathrm{sd}_n\Big(\widehat{\Psi}^{-}_{t_1^{(k)},t_2^{(k)},\delta}(Z_i,\hat{\eta},\hat{\rho}^{-}_{t_1^{(k)},t_2^{(k)}}(\delta)) - \widehat{\Psi}^{-}_{t_1^{(r)},t_2^{(r)},\delta}(Z_i,\hat{\eta},\hat{\rho}^{-}_{t_1^{(r)},t_2^{(r)}}(\delta))\Big), \tag{17}$$

and $\mathrm{sd}_n(\cdot)$ is the empirical standard deviation over $i = 1, \ldots, n$.

5 EXPERIMENTS

We now demonstrate the effectiveness of our proposed method empirically. As is standard in the literature (Shalit et al., 2017; Curth & van der Schaar, 2021), we evaluate our method on synthetic and semi-synthetic data where we have access to ground-truth values of causal quantities. We also provide experimental results using real-world A/B tests. Additional results are in Appendix I.

Algorithm 1: Smoothed estimation of Makarov bounds with sample splitting.

1: **Input:** Dataset $\mathcal{D} = \{(x_i, a_i, y_i)\}_{i=1}^n$
2: **Stage 0:** Split the data $\mathcal{D}$ randomly into two disjoint datasets $\mathcal{D}_1$ and $\mathcal{D}_2$.
3: **Stage 1a:** Estimate nuisance functions $\hat{\eta} = (\hat{F}_1, \hat{F}_0, \hat{\pi})$ on $\mathcal{D}_1$; obtain predictions on $\mathcal{D}_2$.
4: **Stage 1b:** Obtain tuned smoothing parameters $\hat{t}_1^{\pm}$ and $\hat{t}_2^{\pm}$ on $\mathcal{D}_1$.
5: **Stage 2:** Compute debiased estimators $\hat{\rho}^-_{\hat{t}_1^-,\hat{t}_2^-}(\delta)$ and $\hat{\rho}^+_{\hat{t}_1^+,\hat{t}_2^+}(\delta)$.
6: **Output:** $\hat{\rho}^-_{\hat{t}_1^-,\hat{t}_2^-}(\delta)$ and $\hat{\rho}^+_{\hat{t}_1^+,\hat{t}_2^+}(\delta)$.

Nuisance estimation. We estimate the treatment-specific conditional c.d.f.s $F_a(\cdot \mid x)$, $a \in \{0, 1\}$, via likelihood-based gradient boosting with sample-splitting. For continuous outcomes we fit covariate-dependent Gaussian mixtures; for discrete outcomes we use a multinomial classifier and obtain the c.d.f. by cumulatively summing predicted class probabilities; and for nonnegative counts with excess zeros we employ a zero-inflated Poisson with covariate-dependent rate and zero-inflation. Training uses early stopping on held-out log-likelihood. Each learner outputs $\hat{F}_a(y \mid x)$ on an arbitrary grid in $\mathcal{Y}$ of size $k = 300$ that is used for integration. We emphasize that this is a one-dimensional integration whose cost does not scale with the sample size or covariate dimension and whose error is negligible for moderate values of k. As such, this step is computationally cheaper and also more accurate than the nuisance estimation, and there is no need to tune the grid-size k. More details are in Appendix F.

Baselines. We compare our method with the *plug-in estimator* and the *envelope estimator*. To ensure a fair comparison, we use the same nuisance estimators for $F_1(y|x)$ and $F_0(y|x)$ from above. Additionally, we also report results for our method and the baselines targeting the *marginal* Makarov bounds (corresponding to not using any covariates X). For the marginal bounds, we use the standard empirical c.d.f. on a holdout dataset as a nuisance estimator. We use Lepski's method for selecting the smoothing parameters.

Evaluation. We evaluate our estimators as point-estimators of the upper and lower Makarov bounds, and compute the bias, variance, and mean-squared error of all estimators averaged across $K = 2500$ Monte-Carlo runs. We focus on these point-estimation metrics rather than coverage and interval width here because the desired one-sided coverage (i.e., is the lower confidence interval below the lower Makarov bound) is often trivially achieved by biased estimators (since a lower-bound estimator that is biased-downward will often trivially achieve one-sided coverage of the lower bound, and vice-versa for the upper bound), and all of the estimators we consider here are biased. As such, we focus on point estimation metrics here, while providing additional results on interval coverage and width Appendix I, as well as a comparison of interval width on real data in Section 5.2.

5.1 (SEMI)-SYNTHETIC DATA

Fully synthetic data. We simulate data from a synthetic data-generating process (DGP) that violates the margin assumption by having a constant treatment effect for some users. This leads to plateaus in the difference between the treatment and control c.d.f. , where the size of these plateaus is controlled by a parameter γ, with smaller values of γ corresponding to larger violations. For a detailed description of the DGP, we refer the interested reader to Appendix G.

The bias, variance, and MSE of the envelope and smoothing estimators with and without covariates for this setting are shown in the left panel of Figure 2. Our smoothing estimator has consistently lower variance than the envelope estimator and has smaller bias as well when the margin violation is large (i.e., when γ is small). As a result, our estimator consistently improves over baselines, and this improvement is larger when the margin violation is larger.

Semi-synthetic data. For our semi-synthetic experiments, we use covariates from the *Oregon health insurance experiment* (OHIE) (Finkelstein et al., 2012). The OHIE was a randomized experiment meant to assess the effect of health insurance on outcomes such as health or economic status (see

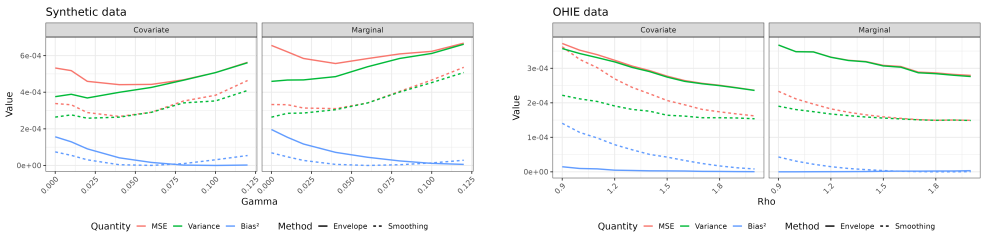

Figure 2: MSE, bias, and variance of our proposed smoothing estimator compared to the baseline envelope estimator for estimating the upper Makarov bound on synthetic and OHIE data, with and without covariates. Our estimator has a more efficient bias-variance trade-off than the baseline estimator, leading to consistently lower MSE across all settings.

Appendix H for details). We use the following covariates: age, gender, language, the number of emergency visits before the experiment, and the number of people the individual signed up with. We then simulate treatment and outcomes such that the margin assumption is violated. The extent of the margin violation is controlled by a parameter ρ, with larger values corresponding to larger violations; see Appendix H for details.

The bias, variance, and MSE of both estimators with and without covariates are shown in the right panel of Figure 2. As in the synthetic setting, our smoothing estimator has consistently lower variance than the envelope estimator; unlike in the synthetic setting, the smoothing estimator typically has higher bias than the envelope estimator. However, the smoothing estimator still trades off bias and variance more efficiently to obtain MSE improvements over the envelope estimator, and these improvements are larger when the margin violation is larger (i.e., when ρ is large).

Finally, Appendix Table 3 reports lower-bound results and plug-in comparisons, confirming that our estimator achieves the lowest mean-squared error across a range of settings.

5.2 Real-world data

A/B tests from a consumer technology company. Finally, we apply our method to three real A/B tests from a large consumer technology company. All three A/B tests have constant propensity score $\pi(x) = 0.5$ and have on the order of tens of millions of observations. The first experiment is an experiment that boosts certain content in a ranking context, and the outcome metric is a measure of engagement. The second experiment demotes certain content in a (different) ranking context, and the outcome metric is a (different) measure of engagement. The third experiment tests a treatment meant to increase visitation, and the outcome metric is a visitation metric. Basic statistics for these experiments are reported in Table 1. For each experiment, we estimate the Makarov bounds using our method to obtain a partial identification region for the c.d.f. of the treatment effect.

Table 1: Experiment statistics summary.

	Sample size, n	% ATE	Bound on $\rho(-1)^{(*)}$
Exp. 1	$\approx$ 30m	16.25%	0.37
Exp. 2	$\approx$ 5m	-1.76%	0.76
Exp. 3	$\approx$ 10m	0.45%	0.99

Results. The results are shown in Figure 3, with bias-corrected point estimates as solid lines and confidence intervals as shaded regions. We see that the confidence intervals are typically small relative to the width of the interval (and are in fact almost completely negligible in experiments 1 and 2), which is a function of the large sample sizes in our experiments.

For the first experiment, we see that the partial identification region is relatively narrow. We are confident that $\mathbb{P}(Y_i(1) - Y_i(0) \leq -1) = \mathbb{P}(Y_i(1) - Y_i(0) < 0)$ is no more than 0.25, meaning that there are very few users for whom the treatment is decreasing engagement. Since we are confident this treatment does not negatively affect a significant fraction of users, and the ATE estimate in Table 1 is positive, this treatment would be safe to launch. (On the other hand, the vertical jump in the identification region at 0 and the high lower bound on $\mathbb{P}(Y_i(1) - Y_i(0) \leq 0)$ suggest that—despite the positive ATE estimate—the treatment is actually having no effect on most users, likely due to the dynamics of the ranking, meaning that we may want to search for more effective treatments).

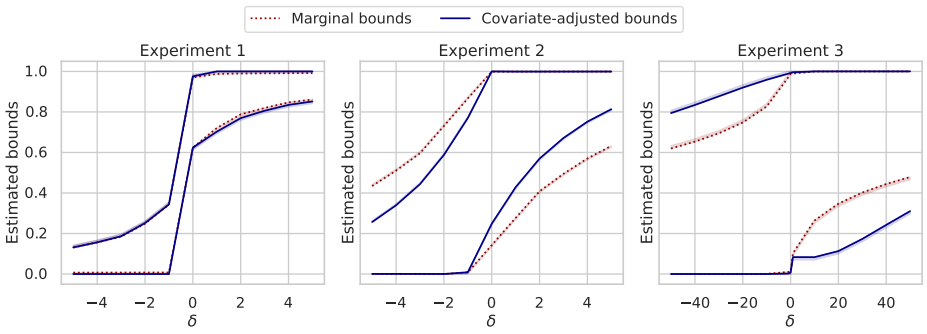

Figure 3: Bias-corrected point estimates of Makarov bounds and confidence intervals for three A/B tests from a consumer technology company. Bias-corrected point estimates are shown as solid lines, while confidence intervals are shown as shaded regions above and below the upper and lower bounds respectively. Based on these bounds, we are confident that Experiment 1 does not negatively impact many users, and thus is safe to launch, whereas Experiments 2 and 3 may negatively impact many users.

For the second experiment, the partial identification region is wider—our upper bound on $\mathbb{P}(Y_i(1) - Y_i(0) < 0)$ is now 0.74, suggesting that this treatment may affect a majority of users negatively, and we should conduct further analysis before launching it. Finally, for the third experiment, the partial identification region is extremely wide: our upper bound on $\mathbb{P}(Y_i(1) - Y_i(0) < 0)$ is now 0.99, suggesting that this treatment could have a negative effect on nearly all users, and would be inadvisable to launch. This experiment has a positive average treatment effect, as seen in Table 1, and so this recommendation to not launch contradicts the standard decision-making process. Taken together, these examples show that our method can be used in real world settings to characterize treatment effect distributions and inform launch decisions.

Figure 4 compares our smoothed estimator to the baseline envelope estimator on Experiment 1 and shows close agreement, especially for the upper bound, while retaining the point-estimation improvements observed in the simulations.

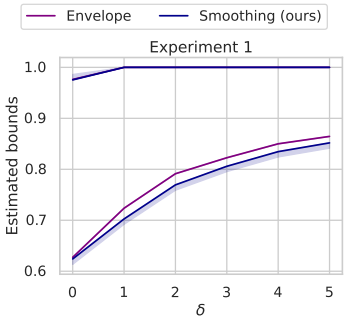

Figure 4: Comparison of our method to the baseline envelope estimator on a range of δ values for Experiment 1. Our bound is slightly more conservative than the baseline, but the two are overall similar, especially for the upper bound.

6 Discussion.

In this paper, we proposed a new method for inferring Makarov bounds on the treatment effect distribution without making a margin assumption. This method replaces non-differentiable terms in the Makarov bounds with smooth approximations, and then uses a debiased estimator based on the efficient influence function for the smooth approximation to estimate the bound, with smoothing levels chosen based on MSE upper bounds. We showed in synthetic and semi-synthetic simulations that this method achieves lower mean-squared-error than existing estimators, and showed in real-data that this estimator can be used to inform decision making from A/B test results.

There are several possible extensions, such as generalizing to multi-valued, continuous, or time-varying treatments, as well as to experiments with interference/spillovers. Perhaps the most important extension is understanding the downstream impact of risk-aware decision making. In practice, we recommend launching treatments only if the ATE is stat-sig positive and we are confident based on our methods that $\rho(-1)$ is smaller than a domain-specific threshold. Exploring how many decisions are changed when using this heuristic and how much downstream harm is prevented is an important step in further demonstrating the value of more sophisticated causal estimation methods like ours.

AI USE STATEMENT

In this work, we used generative AI tools to assist with code implementation, language polishing for readability and clarity, checking mathematical derivations and proofs for consistency, and suggesting the organization of the manuscript. All AI-assisted code was reviewed and tested by the authors, all mathematical arguments were independently verified by the authors, and all AI-assisted text was manually reviewed and edited. We take responsibility for the final content of this work, including text, claims, code, and other artifacts produced with the aid of generative AI.

REFERENCES

Alberto Abadie and Joshua Angrist. Instrumental variables estimation of quantile treatment effects, 1998.

Ahmed Alaa, Zaid Ahmad, and Mark van der Laan. Conformal meta-learners for predictive inference of individual treatment effects. In *NeurIPS*, 2023.

Susan Athey, Raj Chetty, and Guido Imbens. Combining experimental and observational data to estimate treatment effects on long term outcomes. *arXiv preprint*, arXiv:2006.09676, 2020. URL `http://arxiv.org/pdf/2006.09676v1`.

Vahid Balazadeh, Vasilis Syrgkanis, and Rahul G. Krishnan. Partial identification of treatment effects with implicit generative models. In *NeurIPS*, 2022.

Alexander Balke and Judea Pearl. Bounds on treatment effects from studies with imperfect compliance. *Journal of the American Statistical Association*, 92(439):1171–1176, 1997.

Kan Chen, Bingkai Wang, and Dylan S. Small. A differential effect approach to partial identification of treatment effects. *arXiv preprint*, arXiv:2303.06332, 2023.

Victor Chernozhukov and Christian Hansen. An iv model of quantile treatment effects. *Econometrica*, 73(1):245–261, 2005.

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. ISSN 1368-4221. doi: 10.1111/ectj.12097.

Yifan Cui and Sukjin Han. Policy learning with distributional welfare. *Journal of the American Statistical Association*, pp. 1–12, 2025.

Alicia Curth and Mihaela van der Schaar. Nonparametric estimation of heterogeneous treatment effects: From theory to learning algorithms. In *AISTATS*, 2021.

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, 2023.

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*, 120(549):331–342, 2025.

Guilherme Duarte, Noam Finkelstein, Dean Knox, Jonathan Mummolo, and Ilya Shpitser. An automated approach to causal inference in discrete settings. *Journal of the American Statistical Association*, 119(547):1778–1793, 2024.

Yanqin Fan and Sang Soo Park. Sharp bounds on the distribution of treatment effects and their statistical inference. *Econometric Theory*, 26(3):931–951, 2010.

Bruno Fava. Predicting the distribution of treatment effects: A covariate-adjustment approach. *arXiv preprint arXiv:2407.14635*, 2024.

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*, 30(4):958–968, 2024.

Amy Finkelstein, Sarah Taubman, Bill Wright, Mira Bernstein, Jonathan Gruber, Joseph P. Newhouse, Heidi Allen, and Katherine Baicker. The oregon health insurance experiment: Evidence from the first year. *The Quarterly Journal of Economics*, 127(3):1057–1106, 2012. doi: 10.1093/qje/qjs020.

Sergio Firpo. Efficient semiparametric estimation of quantile treatment effects. *Econometrica*, 75(1): 259–276, 2007.

Dylan J. Foster and Vasilis Syrgkanis. Orthogonal statistical learning. *The Annals of Statistics*, 53(3): 879–908, 2023. ISSN 0090-5364. URL http://arxiv.org/pdf/1901.09036v3.

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*, 2024.

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 Emmanuel J. Candès. Sensitivity analysis of individual treatment effects: A robust conformal inference approach. *Proceedings of the National Academy of Sciences (PNAS)*, 120(6), 2023.

Nathan Kallus. What's the harm? sharp bounds on the fraction negatively affected by treatment. In *NeurIPS*, 2022.

Nathan Kallus. Treatment effect risk: Bounds and inference. *Management Science*, 69(8):4579–4590, 2023.

Nathan Kallus and Angela Zhou. Fairness, welfare, and equity in personalized pricing. In *FAccT*, 2021.

Nathan Kallus, Xiaojie Mao, Kaiwen Wang, and Zhengyuan Zhou. Doubly robust distributionally roust off-policy evaluation and learning. In *ICML*, 2022. URL http://arxiv.org/pdf/2202.09667v1.

Edward H. Kennedy. Towards optimal doubly robust estimation of heterogeneous causal effects. *Electronic Journal of Statistics*, 17(2):3008–3049, 2023.

Edward H Kennedy. Semiparametric doubly robust targeted double machine learning: a review. *Handbook of statistical methods for precision medicine*, pp. 207–236, 2024.

Edward H. Kennedy, Sivaraman Balakrishnan, and Larry Wasserman. Semiparametric counterfactual density estimation. *Biometrika*, 110(4):875–896, 2023.

Niki Kilbertus, Matt J. Kusner, and Ricardo Silva. A class of algorithms for general instrumental variable models. In *NeurIPS*, 2020.

Toru Kitagawa and Aleksey Tetenov. Who should be treated? empirical welfare maximization methods for treatment choice. *Econometrica*, 86(2):591–616, 2018.

Sungwon Lee. Partial identification and inference for conditional distributions of treatment effects. *Journal of Applied Econometrics*, 39(1):107–127, 2024.

Lihua Lei and Emmanuel J Candès. Conformal inference of counterfactuals and individual treatment effects. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 83(5):911–938, 2021.

OV Lepskii. Asymptotically minimax adaptive estimation. i: Upper bounds. optimally adaptive estimates. *Theory of Probability & Its Applications*, 36(4):682–697, 1992.

OV Lepskii. Asymptotically minimax adaptive estimation. ii. schemes without optimal adaptation: Adaptive estimators. *Theory of Probability & Its Applications*, 37(3):433–448, 1993.

Alexander W Levis, Matteo Bonvini, Zhenghao Zeng, Luke Keele, and Edward H Kennedy. Covariate-assisted bounds on causal effects with instrumental variables. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 87(5):1508–1527, 2025.

Wei Liang and Changbao Wu. Estimation of treatment harm rate via partitioning. *arXiv preprint arXiv:2505.12209*, 2025.

Christos Louizos, Uri Shalit, Joris Mooij, David Sontag, Richard Zemel, and Max Welling. Causal effect inference with deep latent-variable models. In *NeurIPS*, 2017.

Alexander R Luedtke and Mark J Van Der Laan. Statistical inference for the mean outcome under a possibly non-unique optimal treatment strategy. *Annals of statistics*, 44(2):713, 2016.

GD Makarov. Estimates for the distribution function of a sum of two random variables when the marginal distributions are fixed. *Theory of Probability & its Applications*, 26(4):803–806, 1982.

Charles F. Manski. Nonparametric bounds on treatment effects. *The American Economic Review*, 80 (2):319–323, 1990.

Valentyn Melnychuk, Stefan Feuerriegel, and Mihaela van der Schaar. Quantifying aleatoric uncertainty of the treatment effect: A novel orthogonal learner. In *NeurIPS*, 2024.

Kirtan Padh, Jakob Zeitler, David Watson, Matt Kusner, Ricardo Silva, and Niki Kilbertus. Stochastic causal programming for bounding treatment effects. In *CLeaR*, 2023.

Gyungbae Park. Debiased machine learning when nuisance parameters appear in indicator functions. *arXiv preprint arXiv:2403.15934*, 2024.

Richard AJ Post and Edwin R Van Den Heuvel. Beyond conditional averages: Estimating the individual causal effect distribution. *Journal of Causal Inference*, 13(1):20240007, 2025.

James M. Robins. A new approach to causal inference in mortality studies with a sustained exposure period: Application to control of the healthy worker survivor effect. *Mathematical Modelling*, 7: 1393–1512, 1986.

James M. Robins, Andrea Rotnitzky, and Lue Ping Zhao. Estimation of reression coefficients when some regressors are not always observed. *Journal of the American Statistical Association*, 89(427): 846–688, 1994.

Donald B. Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5):688–701, 1974. ISSN 0022-0663. doi: 10.1037/h0037350.

Gabriel Ruiz and Oscar Hernan Madrid Padilla. High confidence inference on the probability an individual benefits from treatment using experimental or observational data with known propensity scores. *arXiv preprint arXiv:2205.09094*, 2022.

Maresa Schröder, Dennis Frauen, Jonas Schweisthal, Konstantin Heß, Valentyn Melnychuk, and Stefan Feuerriegel. Conformal prediction for causal effects of continuous treatments. In *NeurIPS*, 2025.

Vira Semenova. Aggregated intersection bounds and aggregated minimax values. *arXiv preprint arXiv:2303.00982*, 2025.

Uri Shalit, Fredrik D. Johansson, and David Sontag. Estimating individual treatment effect: Generalization bounds and algorithms. In *ICML*, 2017.

Adith Swaminathan and Thorsten Joachims. Counterfactual risk minimization: Learning from logged bandit feedback. In *ICML*, 2015. URL `http://arxiv.org/pdf/1502.02362v2`.

Zhiqiang Tan. A distributional approach for causal inference using propensity scores. *Journal of the American Statistical Association*, 101(476):1619–1637, 2006.

Mark J. van der Laan and Susan Gruber. Targeted minimum loss based estimation of causal effects of multiple time point interventions. *The International Journal of Biostatistics*, 8(1), 2012. doi: 10.1515/1557-4679.1370.

Mark J. van der Laan and Donald B. Rubin. Targeted maximum likelihood learning. *The International Journal of Biostatistics*, 2(1), 2006.

Aad W Van Der Vaart and Jon A Wellner. Weak convergence. In *Weak convergence and empirical processes: with applications to statistics*, pp. 16–28. Springer, 1996.

Aart van der Vaart. *Asymptotic statistics*. Cambridge University Press, Cambridge, 1998. ISBN 0521496039.

Hal R. Varian. Causal inference in economics and marketing. *Proceedings of the National Academy of Sciences (PNAS)*, 113(27):7310–7315, 2016. doi: 10.1073/pnas.1510479113.

Mingzhang Yin, Claudia Shi, Yixin Wang, and David M Blei. Conformal sensitivity analysis for individual treatment effects. *Journal of the American Statistical Association*, 119(545):122–135, 2024.

Zhehao Zhang and Thomas S Richardson. Individual treatment effect: Prediction intervals and sharp bounds. *arXiv preprint*, arXiv:2506.07469, 2025.

A PROOFS.

A.1 PROOF OF LEMMA 4.1

Proof. Recall that the LSE operator is defined as

$$g_{t_1}(f) = \frac{1}{t_1} \log \left(\int \exp(t_1 f(y)) \, \mathrm{d}y \right) \tag{18}$$

and that softplus function is defined as

$$h_{t_2}(u) = \frac{1}{t_2} \log\bigl(1 + e^{t_2 u}\bigr). \tag{19}$$

We can now write the smoothed Makarov bounds as

$$\rho^{-}_{t_1,t_2}(\delta) = \mathbb{E}\left[\frac{1}{t_2} \log\left(1 + I_{t_1,\delta}(X)^{\frac{t_2}{t_1}}\right)\right] = \mathbb{E}\left[h_{t_2}\bigl(g_{t_1}(F_1(\cdot|X) - F_0(\cdot - \delta|X))\bigr)\right] \tag{20}$$

and

$$\rho^{+}_{t_1,t_2}(\delta) = 1 - \mathbb{E}\left[\frac{1}{t_2} \log\left(1 + I_{-t_1,\delta}(X)^{\frac{t_2}{t_1}}\right)\right] = 1 + \mathbb{E}\left[h_{-t_2}\bigl(g_{-t_1}(F_1(\cdot|X) - F_0(\cdot - \delta|X))\bigr)\right]. \tag{21}$$

We now employ standard bounds on the softmax and LSE functions. For all $u \in \mathbb{R}$, $t_2 > 0$, it holds that

$$h_{t_2}(u) - \frac{\log 2}{t_2} \leq (u)_+ \text{ and } (u)_- \leq h_{-t_2}(u) + \frac{\log 2}{t_2}. \tag{22}$$

For bounded $f : \mathcal{Y} \to \mathbb{R}$ on compact $\mathcal{Y}$ and $t_1 > 0$, it holds that

$$g_{t_1}(f) - \frac{\log|\mathcal{Y}|}{t_1} \leq \sup_{y \in \mathcal{Y}} f(y) \text{ and } \inf_{y \in \mathcal{Y}} f(y) \leq g_{-t_1}(f) + \frac{\log|\mathcal{Y}|}{t_1}. \tag{23}$$

In the following, we prove the theorem for the lower Makarov bound but the argument for the upper bound follows analogously. We apply the inequalities in (22) and (23) to Eq. (20) and obtain

$$\rho^{-}_{t_1,t_2}(\delta) \overset{(*)}{\leq} \mathbb{E}\left[h_{t_2}\left(\sup_{y \in \mathcal{Y}}(F_1(y|X) - F_0(y - \delta|X)) + \frac{\log|\mathcal{Y}|}{t_1}\right)\right] \tag{24}$$

$$\leq \mathbb{E}\left[\left(\sup_{y \in \mathcal{Y}}(F_1(y|X) - F_0(y - \delta|X)) + \frac{\log|\mathcal{Y}|}{t_1}\right)_+\right] + \frac{\log 2}{t_2} \tag{25}$$

$$\overset{(**)}{\leq} \rho^{-}(\delta) + \frac{\log|\mathcal{Y}|_+}{t_1} + \frac{\log 2}{t_2}, \tag{26}$$

where $(*)$ follows from monotonicity of h_{t_2} and $(**)$ follows from $(a+b)_+ \leq a_+ + b_+$. □

A.2 PROOF OF THEOREM 4.2

Proof. We start by deriving the efficinet influence functions (EIFs) of the component ${}_{t_1,\delta}(x)$. By employing the chain rule for influence functions (Kennedy et al., 2023), we obtain

$$EIF\bigl\{I_{t_1,\delta}(x)\bigr\} = t_1 \int_{\mathcal{Y}} \exp\Bigl(t_1 \left[F_1(y \mid x) - F_0(y - \delta \mid x)\right]\Bigr) \, EIF\bigl\{F_1(y \mid x) - F_0(y - \delta \mid x)\bigr\} \, dy. \tag{27}$$

Hence, it holds that

$$EIF\{I_{t_1,\delta}(x)^{\frac{t_2}{t_1}}\} = \frac{t_2}{t_1} I_{t_1,\delta}(x)^{\frac{t_2}{t_1} - 1} \, EIF\{I_{t_1,\delta}(x)\} \tag{28}$$

$$= t_2 I_{t_1,\delta}(x)^{\frac{t_2}{t_1}} \int_{\mathcal{Y}} w_{t_1,\delta}(y|x) \, EIF\bigl\{F_1(y \mid x) - F_0(y - \delta \mid x)\bigr\} \, dy, \tag{29}$$

and for the upper bound analogue

$$EIF\{I_{-t_1,\delta}(x)^{\frac{t_2}{t_1}}\} = \frac{t_2}{t_1} I_{-t_1,\delta}(x)^{\frac{t_2}{t_1}-1}\, EIF\{I_{-t_1,\delta}(x)\} \tag{30}$$

$$= -t_2 I_{-t_1,\delta}(x)^{\frac{t_2}{t_1}} \int_{\mathcal{Y}} w_{-t_1,\delta}(y|x)\, EIF\big\{F_1(y \mid x) - F_0(y-\delta \mid x)\big\}\, dy. \tag{31}$$

The efficient influence functions for the treatment and control c.d.f. are of standard form (Melnychuk et al., 2024), i.e.,

$$EIF\{F_a(y-\delta|X)\} = \frac{\mathbf{1}(A=a)\delta(X=x)}{\mathbb{P}(X=x, A=a)} \left(\mathbf{1}(Y \le y-\delta) - F_a(y-\delta|x)\right). \tag{32}$$

Plugging everything together, the EIF for the lower smoothed Makarov bound is

$$\Psi^{-}_{t_1,t_2,\delta}(Z,\eta) \tag{33}$$

$$= \int \frac{\mathbb{P}(X=x)}{t_2} EIF \log\left(1 + I_{t_1,\delta}(x)^{\frac{t_2}{t_1}}\right) \mathrm{d}x \tag{34}$$

$$+ \frac{1}{t_2}\log\left(1 + I_{t_1,\delta}(X)^{\frac{t_2}{t_1}}\right) - \rho^{-}_{t_1,t_2}(\delta) \tag{35}$$

$$= \int \frac{\mathbb{P}(X=x)}{t_2} \frac{EIF I_{t_1,\delta}(x)^{\frac{t_2}{t_1}}}{1 + I_{t_1,\delta}(x)^{\frac{t_2}{t_1}}} \mathrm{d}x + \frac{1}{t_2}\log\left(1 + I_{t_1,\delta}(X)^{\frac{t_2}{t_1}}\right) - \rho^{-}_{t_1,t_2}(\delta) \tag{36}$$

$$= \int \Pr(X=x) \underbrace{\frac{I_{t_1,\delta}(x)^{\frac{t_2}{t_1}}}{1 + I_{t_1,\delta}(x)^{\frac{t_2}{t_1}}}}_{\sigma^{-}_{t_1,t_2,\delta}(x)} \Big[\int_{\mathcal{Y}} w_{t_1,\delta}(y|x)\, EIF\{F_1 - F_0\}\, dy\Big] dx \tag{37}$$

$$+ \frac{1}{t_2}\log\left(1 + I_{t_1,\delta}(X)^{\frac{t_2}{t_1}}\right) - \rho^{-}_{t_1,t_2}(\delta) \tag{38}$$

$$= \int \Pr(X=x)\, \sigma^{-}_{t_1,t_2,\delta}(x) \Big[\int_{\mathcal{Y}} w_{t_1,\delta}(y|x) \frac{A\delta(X=x)}{\mathbb{P}(X=x, A=1)} \left(\mathbf{1}(Y \le y) - F_1(y|x)\right)\, dy\Big] dx \tag{39}$$

$$- \int \Pr(X=x)\, \sigma^{-}_{t_1,t_2,\delta}(x) \Big[\int_{\mathcal{Y}} w_{t_1,\delta}(y|x) \frac{(1-A)\delta(X=x)}{\mathbb{P}(X=x, A=0)} \left(\mathbf{1}(Y \le y-\delta) - F_0(y-\delta|x)\right)\, dy\Big] dx \tag{40}$$

$$+ \frac{1}{t_2}\log\left(1 + I_{t_1,\delta}(X)^{\frac{t_2}{t_1}}\right) - \rho^{-}_{t_1,t_2}(\delta) \tag{41}$$

$$= \sigma^{-}_{t_1,t_2,\delta}(X) \Big[\int_{\mathcal{Y}} w_{t_1,\delta}(y|X) \frac{A}{\pi(X)} \left(\mathbf{1}(Y \le y) - F_1(y|X)\right)\, dy\Big] \tag{42}$$

$$- \sigma^{-}_{t_1,t_2,\delta}(X) \Big[\int_{\mathcal{Y}} w_{t_1}(y|X) \frac{(1-A)}{1-\pi(X)} \left(\mathbf{1}(Y \le y-\delta) - F_0(y-\delta|X)\right)\, dy\Big] \tag{43}$$

$$+ \frac{1}{t_2}\log\left(1 + I_{t_1,\delta}(X)^{\frac{t_2}{t_1}}\right) - \rho^{-}_{t_1,t_2}(\delta) \tag{44}$$

$$\tag{45}$$

Analogously, we obtain the EIF for the upper smoothed Makarov bound as

$$\Psi^{+}_{t_1,t_2,\delta}(Z,\eta) \tag{46}$$

$$= \int \mathbb{P}(X=x) EIF \left\{ 1 - \frac{1}{t_2} \log\left(1 + I_{-t_1,\delta}(x)^{\frac{t_2}{t_1}}\right) \right\} \mathrm{d}x \tag{47}$$

$$+ 1 - \frac{1}{t_2} \log\left(1 + I_{-t_1,\delta}(X)^{\frac{t_2}{t_1}}\right) - \rho^{+}_{t_1,t_2}(\delta) \tag{48}$$

$$= \int \frac{\mathbb{P}(X=x)}{-t_2} EIF \log\left(1 + I_{-t_1,\delta}(x)^{\frac{t_2}{t_1}}\right) \mathrm{d}x \tag{49}$$

$$+ 1 - \frac{1}{t_2} \log\left(1 + I_{-t_1,\delta}(X)^{\frac{t_2}{t_1}}\right) - \rho^{+}_{t_1,t_2}(\delta) \tag{50}$$

$$= \int \frac{\mathbb{P}(X=x)}{-t_2} \frac{EIF I_{-t_1,\delta}(x)^{\frac{t_2}{t_1}}}{1 + I_{-t_1,\delta}(x)^{\frac{t_2}{t_1}}} \mathrm{d}x + 1 - \frac{1}{t_2} \log\left(1 + I_{-t_1,\delta}(X)^{\frac{t_2}{t_1}}\right) - \rho^{+}_{t_1,t_2}(\delta) \tag{51}$$

$$= \int \Pr(X=x) \underbrace{\frac{I_{-t_1,\delta}(x)^{\frac{t_2}{t_1}}}{1 + I_{-t_1,\delta}(x)^{\frac{t_2}{t_1}}}}_{\sigma^{+}_{t_1,t_2,\delta}(x)} \Big[\int_{\mathcal{Y}} w_{-t_1,\delta}(y|x)\, EIF\{F_1 - F_0\}\, dy \Big] dx \tag{52}$$

$$+ 1 - \frac{1}{t_2} \log\left(1 + I_{-t_1,\delta}(X)^{\frac{t_2}{t_1}}\right) - \rho^{+}_{t_1,t_2}(\delta) \tag{53}$$

$$= \int \Pr(X=x)\, \sigma^{+}_{t_1,t_2,\delta}(x) \Big[\int_{\mathcal{Y}} w_{-t_1,\delta}(y|x) \frac{A\delta(X=x)}{\mathbb{P}(X=x, A=1)} \left(\mathbf{1}(Y \le y) - F_1(y|x)\right) dy \Big] dx \tag{54}$$

$$- \int \Pr(X=x)\, \sigma^{+}_{t_1,t_2,\delta}(x) \Big[\int_{\mathcal{Y}} w_{-t_1,\delta}(y|x) \frac{(1-A)\delta(X=x)}{\mathbb{P}(X=x, A=0)} \left(\mathbf{1}(Y \le y-\delta) - F_0(y-\delta|x)\right) dy \Big] dx \tag{55}$$

$$+ 1 - \frac{1}{t_2} \log\left(1 + I_{-t_1,\delta}(X)^{\frac{t_2}{t_1}}\right) - \rho^{+}_{t_1,t_2}(\delta) \tag{56}$$

$$= \sigma^{+}_{t_1,t_2,\delta}(X) \Big[\int_{\mathcal{Y}} w_{-t_1,\delta}(y|X) \frac{A}{\pi(X)} \left(\mathbf{1}(Y \le y) - F_1(y|X)\right) dy \Big] \tag{57}$$

$$- \sigma^{+}_{t_1,t_2,\delta}(X) \Big[\int_{\mathcal{Y}} w_{-t_1}(y|X) \frac{(1-A)}{1-\pi(X)} \left(\mathbf{1}(Y \le y-\delta) - F_0(y-\delta|X)\right) dy \Big] \tag{58}$$

$$+ 1 - \frac{1}{t_2} \log\left(1 + I_{-t_1,\delta}(X)^{\frac{t_2}{t_1}}\right) - \rho^{+}_{t_1,t_2}(\delta). \tag{59}$$

□

A.3 Proof of Corollary 4.3

Assumption A.1 (Regularity and rate conditions for Corollary 4.3)**.** Fix $\delta \in \mathbb{R}$ and smoothing parameters $t_1, t_2 > 0$. Let $Z = (X, A, Y) \sim \mathbb{P}$ denote a generic observation and let $\eta = (F_0, F_1, \pi)$ denote the collection of nuisance functions.

1. **Causal assumptions and strong overlap.** Assumption 3.1 holds and, for some constant $\epsilon > 0$,
$$\epsilon \le \pi(X) \le 1 - \epsilon \qquad \text{almost surely.}$$
2. **Outcome support and boundedness.** The conditions of Lemma 4.1 hold. The efficient influence functions $\Psi^{\pm}_{t_1,t_2,\delta}(Z, \eta, \theta^{\pm}(\delta))$ in Theorem 4.2 are square-integrable with strictly positive and finite variance
$$0 < V^{\pm}(\delta) := \mathrm{Var}\Big(\Psi^{\pm}_{t_1,t_2,\delta}(Z, \eta, \theta^{\pm}(\delta))\Big) < \infty.$$

3. **Cross-fitting.** The estimators $\hat{\eta} = (\hat{F}_0, \hat{F}_1, \hat{\pi})$ are obtained with sample splitting or K-fold cross-fitting: for each observation i, the corresponding nuisance estimates $\hat{\eta}^{(-k(i))}$ are trained without using the fold containing i.

4. **Rate conditions for nuisances.** Define
$$\|G\|_{2,\mathcal{Y}}^2 := \mathbb{E}\left[\int_{\mathcal{Y}} G(y,X)^2\, dy\right].$$
For $a \in \{0,1\}$,
$$\|\hat{F}_a - F_a\|_{2,\mathcal{Y}} = o_{\mathbb{P}}(n^{-1/4}), \qquad \|\hat{\pi} - \pi\|_2 = o_{\mathbb{P}}(n^{-1/4}).$$

5. **Estimated overlap.** With probability tending to one,
$$\epsilon/2 \le \hat{\pi}(X) \le 1 - \epsilon/2.$$

Proof of Corollary 4.3. Fix $\delta \in \mathbb{R}$ and abbreviate
$$\theta^{\pm}(\delta) \;=\; \rho^{\pm}_{t_1,t_2}(\delta), \qquad \hat{\theta}^{\pm}(\delta) \;=\; \hat{\rho}^{\pm}_{t_1,t_2}(\delta),$$
for the smoothed lower and upper Makarov bounds and their debiased estimators in (12). For readability we suppress the dependence on δ where no confusion arises.

Step 1: Oracle influence function representation. By Theorem 4.2 and Assumption A.1(v), the smoothed functionals $\theta^{\pm}$ are pathwise differentiable at $\mathbb{P}$with efficient influence functions $\Psi^{\pm}_{t_1,t_2,\delta}(Z,\eta,\theta^{\pm}(\delta))$. Consider the oracle estimator
$$\tilde{\theta}^{\pm}(\delta) \;:=\; \frac{1}{n}\sum_{i=1}^{n}\Big(\theta^{\pm}(\delta) + \Psi^{\pm}_{t_1,t_2,\delta}(Z_i,\eta,\theta^{\pm}(\delta))\Big),$$
which satisfies
$$\tilde{\theta}^{\pm}(\delta) - \theta^{\pm}(\delta) = \frac{1}{n}\sum_{i=1}^{n}\Psi^{\pm}_{t_1,t_2,\delta}(Z_i,\eta,\theta^{\pm}(\delta)).$$
By Assumption A.1(ii), the summands are i.i.d. with mean zero and variance $V^{\pm}(\delta) \in (0,\infty)$, so the Lindeberg–Feller central limit theorem yields
$$\sqrt{n}\big(\tilde{\theta}^{\pm}(\delta) - \theta^{\pm}(\delta)\big) \;\Rightarrow\; \mathcal{N}\big(0, V^{\pm}(\delta)\big). \tag{60}$$

Step 2: Effect of plug-in nuisance estimation with cross-fitting. Let the sample be partitioned into K folds and denote by $\hat{\eta}^{(-k)}$ the nuisance estimates trained on all folds except fold k. For each i, let $k(i)$ be the index of the fold containing Z_i. The debiased estimators with cross-fitting can be written as
$$\hat{\theta}^{\pm}(\delta) = \frac{1}{n}\sum_{i=1}^{n}\Big(\theta^{\pm}(\delta) + \Psi^{\pm}_{t_1,t_2,\delta}(Z_i,\hat{\eta}^{(-k(i))},\theta^{\pm}(\delta))\Big) \;+\; R_n^{\pm},$$
where $R_n^{\pm}$ collects the higher-order terms arising from replacing $\theta^{\pm}(\delta)$ by $\hat{\theta}^{\pm}(\delta)$ inside the influence function.

By Neyman orthogonality of the score in Theorem 4.2 and Assumption A.1(iii)–(iv), standard arguments for orthogonal / debiased estimation (see, e.g., Chernozhukov et al. (2018)) imply that
$$\left|\frac{1}{n}\sum_{i=1}^{n}\Big(\Psi^{\pm}_{t_1,t_2,\delta}(Z_i,\hat{\eta}^{(-k(i))},\theta^{\pm}(\delta)) - \Psi^{\pm}_{t_1,t_2,\delta}(Z_i,\eta,\theta^{\pm}(\delta))\Big)\right| \;=\; o_{\mathbb{P}}(n^{-1/2}),$$
and that $R_n^{\pm} = o_{\mathbb{P}}(n^{-1/2})$ by a Taylor expansion in θ. Hence we obtain the linear expansion
$$\hat{\theta}^{\pm}(\delta) - \theta^{\pm}(\delta) = \frac{1}{n}\sum_{i=1}^{n}\Psi^{\pm}_{t_1,t_2,\delta}(Z_i,\eta,\theta^{\pm}(\delta)) \;+\; r_n^{\pm}, \qquad r_n^{\pm} = o_{\mathbb{P}}(n^{-1/2}). \tag{61}$$
Comparing (61) with the oracle representation shows that the leading term is identical and the remainder is negligible at the $\sqrt{n}$ scale. Combining (60) and (61) and applying Slutsky's lemma gives, for each fixed δ,
$$\sqrt{n}\big(\hat{\theta}^{\pm}(\delta) - \theta^{\pm}(\delta)\big) \;\Rightarrow\; \mathcal{N}\big(0, V^{\pm}(\delta)\big). \tag{62}$$

Step 3: Consistency of the standard error estimator. Define

$$\widehat{V}^{\pm}(\delta) = \frac{1}{n}\sum_{i=1}^{n}\Big(\Psi^{\pm}_{t_1,t_2,\delta}(Z_i, \hat{\eta}^{(-k(i))}, \hat{\theta}^{\pm}(\delta)) - \bar{\Psi}^{\pm}(\delta)\Big)^2,$$

where

$$\bar{\Psi}^{\pm}(\delta) = \frac{1}{n}\sum_{i=1}^{n}\Psi^{\pm}_{t_1,t_2,\delta}(Z_i, \hat{\eta}^{(-k(i))}, \hat{\theta}^{\pm}(\delta)),$$

and set $\widehat{\mathrm{se}}^{\pm}(\delta) = \sqrt{\widehat{V}^{\pm}(\delta)/n}$.

By the same orthogonality and rate conditions as in Step 2, the difference between

$$\Psi^{\pm}_{t_1,t_2,\delta}(Z_i, \hat{\eta}^{(-k(i))}, \hat{\theta}^{\pm}(\delta)) \quad \text{and} \quad \Psi^{\pm}_{t_1,t_2,\delta}(Z_i, \eta, \theta^{\pm}(\delta))$$

is $o_{\mathbb{P}}(1)$ in L_2 uniformly in i. The law of large numbers under cross-fitting then yields

$$\widehat{V}^{\pm}(\delta) \xrightarrow{\mathbb{P}} V^{\pm}(\delta),$$

so $\widehat{\mathrm{se}}^{\pm}(\delta) \xrightarrow{\P} \sqrt{V^{\pm}(\delta)/n}$. Combining this with (62) and using Slutsky's lemma again, we obtain the studentized central limit theorem

$$\frac{\hat{\theta}^{\pm}(\delta) - \theta^{\pm}(\delta)}{\widehat{\mathrm{se}}^{\pm}(\delta)} \Rightarrow \mathcal{N}(0,1). \tag{63}$$

Step 4: One-sided intervals for the smoothed bounds. From (63), for each fixed δ,

$$\Pr\Big\{\theta^{-}(\delta) \geq \hat{\theta}^{-}(\delta) - z_{1-\alpha/2}\,\widehat{\mathrm{se}}^{-}(\delta)\Big\} \to 1-\alpha/2,$$

and

$$\Pr\Big\{\theta^{+}(\delta) \leq \hat{\theta}^{+}(\delta) + z_{1-\alpha/2}\,\widehat{\mathrm{se}}^{+}(\delta)\Big\} \to 1-\alpha/2,$$

where $z_{1-\alpha/2}$ is the $(1-\alpha/2)$ quantile of the standard normal distribution.

Step 5: Translating to the original Makarov bounds. By Lemma 4.1, for all fixed $t_1, t_2 > 0$,

$$\theta^{-}(\delta) - b(t_1,t_2) \;\leq\; \rho(\delta) \;\leq\; \theta^{+}(\delta) + b(t_1,t_2),$$

where $b(t_1,t_2)$ is the approximation bias bound. On the intersection of the two one-sided events from Step 4, we have

$$\hat{\theta}^{-}(\delta) - z_{1-\alpha/2}\,\widehat{\mathrm{se}}^{-}(\delta) - b(t_1,t_2) \;\leq\; \rho(\delta) \;\leq\; \hat{\theta}^{+}(\delta) + z_{1-\alpha/2}\,\widehat{\mathrm{se}}^{+}(\delta) + b(t_1,t_2).$$

By the definition of $c^{-}_{t_1,t_2,\delta,\alpha}(\mathcal{D}_n)$ and $c^{+}_{t_1,t_2,\delta,\alpha}(\mathcal{D}_n)$, this event is precisely

$$c^{-}_{t_1,t_2,\delta,\alpha}(\mathcal{D}_n) \;\leq\; \rho(\delta) \;\leq\; c^{+}_{t_1,t_2,\delta,\alpha}(\mathcal{D}_n).$$

A Bonferroni argument then yields

$$\lim_{n\to\infty}\Pr\Big\{c^{-}_{t_1,t_2,\delta,\alpha}(\mathcal{D}_n) \leq \rho(\delta) \leq c^{+}_{t_1,t_2,\delta,\alpha}(\mathcal{D}_n)\Big\} \;\geq\; 1-\alpha,$$

which proves the corollary. □

B EXTENDED RELATED WORK

Quantile treatment effects. One stream of literature focuses on estimating quantile treatment effects (QTEs) (Abadie & Angrist, 1998; Chernozhukov & Hansen, 2005; Firpo, 2007). These are contrasts of quantiles of the potential outcomes distributions, defined via

$$\text{QTE}(\tau) \;=\; F^{-1}_{Y(1)}(\tau) - F^{-1}_{Y(0)}(\tau),$$

which compares the τ-quantile of $Y(1)$ to the τ-quantile of $Y(0)$. Note that QTEs do not quantify quantiles of the treatment effect distribution: QTEs describe how treatment shifts the marginal distributions of potential outcomes across individuals (for example, how the τ-quantile under treatment differs from the τ-quantile under control), while the latter would characterize the distribution of individual-level causal effects $Y(1) - Y(0)$ in the population. In contrast to QTEs, the treatment effect distribution is generally not identifiable from observed data without additional assumptions, and is the focus of our paper.

Semiparametric efficient causal inference. Semiparametric efficiency theory (van der Vaart, 1998) and efficient-influence function-based estimators have a long tradition in causal inference (Kennedy, 2024). Examples include the AIPTW estimator (Robins et al., 1994), Targeted maximum-likelihood estimation (TMLE) (van der Laan & Rubin, 2006). These frameworks have been extended to various causal quantities and are the de-facto standard for modern causal effect estimation in many settings (van der Laan & Gruber, 2012; Chernozhukov et al., 2018; Foster & Syrgkanis, 2023; Kennedy, 2023).

Bounds for partially identified causal quantities In many situations, the causal parameter of interest is only *partially* identified. That is, we need to obtain bounds on the parameter of interest which we can then estimate with observational data Manski (1990). Several works have proposed methods for partial identification of causal quantities, including treatment effects in instrumental variable settings Balke & Pearl (1997); Kilbertus et al. (2020) and more general causal graphs Duarte et al. (2024); Balazadeh et al. (2022); Chen et al. (2023); Padh et al. (2023), and treatment effect risk Kallus (2023). A related stream of literature obtains bounds under so-called sensitivity models, which impose assumptions on the degree of non-identifiability Tan (2006); Jesson et al. (2021); Dorn & Guo (2023); Dorn et al. (2025); Yin et al. (2024); Frauen et al. (2023); Jin et al. (2023); Frauen et al. (2024).

C EXTENSION TO DISCRETE OUTCOMES.

Here we provide the modifications needed to extend our methodology to discrete outcomes. Let $\mathcal{Y}_d = \{y_1 < \cdots < y_M\}$ denote a finite ordered support. For $a \in \{0, 1\}$, write $F_a(\cdot \mid x)$ for the right-continuous conditional c.d.f. of $Y \mid (A = a, X = x)$. All other notation is as in the main text.

For discrete outcomes, the function

$$y \mapsto F_1(y \mid x) - F_0(y - \delta \mid x)$$

is piecewise constant, with possible jumps at points in

$$\mathcal{B}_\delta := \mathcal{Y}_d \cup \{y + \delta : y \in \mathcal{Y}_d\}.$$

To capture all values attained by this function, including values immediately to the left of a jump, define

$$F_a(u^+ \mid x) := F_a(u \mid x), \qquad F_a(u^- \mid x) := \mathbb{P}(Y < u \mid X = x, A = a).$$

Let

$$\mathcal{J}_\delta := \big(\mathcal{B}_\delta \times \{-, +\}\big) \cup \{\star\},$$

where the additional state $\star$ represents the limiting value zero in the tails. For $j = (b, s) \in \mathcal{B}_\delta \times \{-, +\}$ define

$$d_{\delta,j}(x) = F_1(b^s \mid x) - F_0((b - \delta)^s \mid x),$$

and set $d_{\delta,\star}(x) = 0$. Then

$$\sup_{y \in \mathbb{R}} \big\{F_1(y \mid x) - F_0(y - \delta \mid x)\big\} = \max_{j \in \mathcal{J}_\delta} d_{\delta,j}(x), \tag{64}$$

with the analogous identity for the infimum.

Discrete log-sum-exp and softplus. The softplus $h_{t_2}(u) = t_2^{-1} \log(1 + e^{t_2 u})$ is unchanged. We apply the discrete log-sum-exp to the finite collection of candidate states:

$$g^{(d)}_{t_1}(d_\delta(\cdot, x)) = \frac{1}{t_1} \log \left(\sum_{j \in \mathcal{J}_\delta} \exp\{t_1 d_{\delta,j}(x)\} \right). \tag{65}$$

Accordingly, define

$$I^{(d)}_{t,\delta}(x) = \sum_{j \in \mathcal{J}_\delta} \exp(t\, d_{\delta,j}(x)). \tag{66}$$

Smoothed Makarov bounds (discrete). Replacing the continuous log-sum-exp by (65) yields

$$\rho^{-,(d)}_{t_1,t_2}(\delta) = \mathbb{E}\left[\frac{1}{t_2} \log\Big(1 + \big(I^{(d)}_{t_1,\delta}(X)\big)^{t_2/t_1}\Big)\right], \qquad \rho^{+,(d)}_{t_1,t_2}(\delta) = 1 - \mathbb{E}\left[\frac{1}{t_2} \log\Big(1 + \big(I^{(d)}_{-t_1,\delta}(X)\big)^{t_2/t_1}\Big)\right]. \tag{67}$$

Let $K_\delta := |\mathcal{J}_\delta|$. Since $|\mathcal{B}_\delta| \le 2M$, we have $K_\delta \le 4M + 1$. The analogue of Lemma 4.1 therefore holds with

$$b_d(t_1, t_2; \delta) = \frac{\log 2}{t_2} + \frac{\log K_\delta}{t_1}. \tag{68}$$

In particular,

$$\rho^{-,(d)}_{t_1,t_2}(\delta) - b_d(t_1, t_2; \delta) \le \rho^-(\delta), \qquad \rho^+(\delta) \le \rho^{+,(d)}_{t_1,t_2}(\delta) + b_d(t_1, t_2; \delta),$$

and $\rho^{\pm,(d)}_{t_1,t_2}(\delta) \to \rho^\pm(\delta)$ as $t_1, t_2 \to \infty$.

Efficient influence functions (discrete). Define the candidate-state softmax weights

$$w^{(d)}_{t,\delta,j}(x) = \frac{\exp(t\, d_{\delta,j}(x))}{I^{(d)}_{t,\delta}(x)}, \qquad j \in \mathcal{J}_\delta. \tag{69}$$

For $j = (b, +)$ define

$$R^a_{\delta,j}(X, Y) = \mathbf{1}\{Y \le b - (1-a)\delta\} - F_a\big((b - (1-a)\delta)^+ \mid X\big),$$

while for $j = (b, -)$ define

$$R^a_{\delta,j}(X, Y) = \mathbf{1}\{Y < b - (1-a)\delta\} - F_a\big((b - (1-a)\delta)^- \mid X\big).$$

For the tail state, set $R^a_{\delta,\star}(X, Y) = 0$. The discrete analogue of $\Phi^a_{t,\delta}$ is then

$$\Phi^{a,(d)}_{t,\delta}(X, Y) = \sum_{j \in \mathcal{J}_\delta} w^{(d)}_{t,\delta,j}(X) R^a_{\delta,j}(X, Y). \tag{70}$$

With $\sigma^{\pm}_{t_1,t_2,\delta}(x)$ defined exactly as in the main text but using $I^{(d)}_{\pm t_1,\delta}(x)$, the EIFs have the same form as in Theorem 4.2, with I, Φ, and $\rho^{\pm}$ replaced by their discrete counterparts. The EIFs are

$$\begin{aligned}\Psi^{-,(d)}_{t_1,t_2,\delta}\big(Z, \eta, \rho^{-,(d)}_{t_1,t_2}(\delta)\big) = \sigma^{-}_{t_1,t_2,\delta}(X) \left[\frac{A}{\pi(X)}\, \Phi^{1,(d)}_{t_1,\delta}(X, Y) \;-\; \frac{1-A}{1-\pi(X)}\, \Phi^{0,(d)}_{t_1,\delta}(X, Y)\right] \\ + \frac{1}{t_2} \log\Big(1 + \big(I^{(d)}_{t_1,\delta}(X)\big)^{t_2/t_1}\Big) \;-\; \rho^{-,(d)}_{t_1,t_2}(\delta).\end{aligned} \tag{71}$$

$$\begin{aligned}\Psi^{+,(d)}_{t_1,t_2,\delta}\big(Z, \eta, \rho^{+,(d)}_{t_1,t_2}(\delta)\big) = \sigma^{+}_{t_1,t_2,\delta}(X) \left[\frac{A}{\pi(X)}\, \Phi^{1,(d)}_{-t_1,\delta}(X, Y) \;-\; \frac{1-A}{1-\pi(X)}\, \Phi^{0,(d)}_{-t_1,\delta}(X, Y)\right] \\ + 1 \;-\; \frac{1}{t_2} \log\Big(1 + \big(I^{(d)}_{-t_1,\delta}(X)\big)^{t_2/t_1}\Big) \;-\; \rho^{+,(d)}_{t_1,t_2}(\delta).\end{aligned} \tag{72}$$

Debiased estimators (discrete). The one-step estimators in Eq. (12) become

$$\hat\rho^{-,(d)}_{t_1,t_2}(\delta) = \frac{1}{n}\sum_{i=1}^{n}\left\{\frac{1}{t_2}\log\Big(1 + \big(\hat I^{(d)}_{t_1,\delta}(x_i)\big)^{t_2/t_1}\Big) + \hat\sigma^{-}_{t_1,t_2,\delta}(x_i)\, \frac{a_i - \hat\pi(x_i)}{\hat\pi(x_i)\{1-\hat\pi(x_i)\}}\, \hat\Phi^{a_i,(d)}_{t_1,\delta}(x_i, y_i)\right\}. \tag{73}$$

$$\hat\rho^{+,(d)}_{t_1,t_2}(\delta) = \frac{1}{n}\sum_{i=1}^{n}\left\{1 - \frac{1}{t_2}\log\Big(1 + \big(\hat I^{(d)}_{-t_1,\delta}(x_i)\big)^{t_2/t_1}\Big) + \hat\sigma^{+}_{t_1,t_2,\delta}(x_i)\, \frac{a_i - \hat\pi(x_i)}{\hat\pi(x_i)\{1-\hat\pi(x_i)\}}\, \hat\Phi^{a_i,(d)}_{-t_1,\delta}(x_i, y_i)\right\}. \tag{74}$$

Here $\hat I^{(d)}_{\pm t_1,\delta}$, $\hat w^{(d)}_{\pm t_1,\delta,j}$, and $\hat\Phi^{a_i,(d)}_{\pm t_1,\delta}$ replace (66), (69), and (70) with estimated nuisances. Cross-fitting and variance estimation via the sample variance of the estimated EIF proceed unchanged.

D EXTENSION TO UNIFORMLY VALID CONFIDENCE INTERVALS

In this section, we extend Corollary 4.3 to construct confidence intervals for the treatment effect distribution that are *simultaneously* valid for all values of δ in a compact set $\Delta \subset \mathbb{R}$. The construction proceeds in two steps. First, we establish a functional central limit theorem (CLT) for the smoothed debiased estimators $\hat{\rho}^{\pm}_{t_1,t_2}(\delta)$ viewed as stochastic processes in δ. Second, we use a multiplier bootstrap to approximate the distribution of the supremum of the corresponding Gaussian limit and combine this with the approximation bias bound $b(t_1, t_2)$ from Lemma 4.1 to obtain uniform confidence bands for the original Makarov bounds.

D.1 UNIFORM ASYMPTOTIC LINEARITY AND FUNCTIONAL CLT

We focus on a compact interval $\Delta \subset \mathbb{R}$ of values of δ that are of substantive interest. Throughout this section, we treat the smoothing parameters (t_1, t_2) as fixed and suppress their dependence where it is notationally convenient. We write

$$\mathbb{G}_n f \;=\; \frac{1}{\sqrt{n}} \sum_{i=1}^{n} \big(f(Z_i) - \mathbb{E}[f(Z)] \big)$$

for the empirical process indexed by a function f, and consider the class of efficient influence functions

$$\mathcal{F}^{\pm} \;=\; \big\{ \Psi^{\pm}_{t_1,t_2,\delta}(\,\cdot\,, \eta, \rho^{\pm}_{t_1,t_2}(\delta)) : \delta \in \Delta \big\}$$

as given in Theorem 4.2.

We impose the following regularity condition, which is a uniform version of the high-level conditions used in Corollary 4.3.

Assumption D.1 (Uniform regularity of the EIF process)**.** Let $\Delta \subset \mathbb{R}$ be compact and let $\Psi^{\pm}_{t_1,t_2,\delta}(Z, \eta, \rho^{\pm}_{t_1,t_2}(\delta))$ be as in Theorem 4.2. Assume:

1. **Uniform bounded second moments:**
$$\sup_{\delta \in \Delta} \mathbb{E}\big[\Psi^{\pm}_{t_1,t_2,\delta}(Z, \eta, \rho^{\pm}_{t_1,t_2}(\delta))^2 \big] < \infty.$$

2. **Stochastic equicontinuity in δ:** there exists a semi-metric d on Δ such that, for any $\varepsilon > 0$,
$$\lim_{\gamma \to 0} \limsup_{n \to \infty} \mathbb{P}\Big(\sup_{d(\delta,\delta') \le \gamma} \big| \mathbb{G}_n \big(\Psi^{\pm}_{t_1,t_2,\delta} - \Psi^{\pm}_{t_1,t_2,\delta'} \big) \big| > \varepsilon \Big) = 0.$$

3. **Donsker-type condition:** the class $\mathcal{F}^{\pm}$ is P-Donsker (or, more generally, satisfies an entropy condition that guarantees a functional CLT for $\mathbb{G}_n$ indexed by $\mathcal{F}^{\pm}$).

4. **Uniform orthogonality and nuisance rates:** let $\hat{\eta}$ be obtained via sample splitting as in Corollary 4.3. The asymptotic linear expansion
$$\hat{\rho}^{\pm}_{t_1,t_2}(\delta) = \rho^{\pm}_{t_1,t_2}(\delta) + \frac{1}{n} \sum_{i=1}^{n} \Psi^{\pm}_{t_1,t_2,\delta}(Z_i, \eta, \rho^{\pm}_{t_1,t_2}(\delta)) + r^{\pm}_n(\delta)$$
holds with a remainder satisfying
$$\sup_{\delta \in \Delta} \sqrt{n} \, \big| r^{\pm}_n(\delta) \big| \;\xrightarrow{P}\; 0.$$

Under Assumption D.1, we obtain a functional CLT for the smoothed lower and upper bounds following standard arguments in empirical process theory (Van Der Vaart & Wellner, 1996).

Theorem D.2 (Functional CLT for smoothed Makarov bounds)**.** *Let Assumption D.1 hold and let $\Delta \subset \mathbb{R}$ be compact. Then, in the space $\ell^\infty(\Delta)$ of bounded real-valued functions on Δ,*

$$\big\{ \sqrt{n} \big(\hat{\rho}^{-}_{t_1,t_2}(\delta) - \rho^{-}_{t_1,t_2}(\delta) \big) \big\}_{\delta \in \Delta} \rightsquigarrow \{ G^{-}(\delta) \}_{\delta \in \Delta}, \tag{75}$$

$$\big\{ \sqrt{n} \big(\hat{\rho}^{+}_{t_1,t_2}(\delta) - \rho^{+}_{t_1,t_2}(\delta) \big) \big\}_{\delta \in \Delta} \rightsquigarrow \{ G^{+}(\delta) \}_{\delta \in \Delta}, \tag{76}$$

where $G^{\pm}$ are mean-zero tight Gaussian processes with covariance functions

$$\mathrm{Cov}\big(G^{\pm}(\delta_1), G^{\pm}(\delta_2)\big) = \mathbb{E}\big[\Psi^{\pm}_{t_1,t_2,\delta_1}(Z,\eta,\rho^{\pm}_{t_1,t_2}(\delta_1))\,\Psi^{\pm}_{t_1,t_2,\delta_2}(Z,\eta,\rho^{\pm}_{t_1,t_2}(\delta_2))\big]$$

for all $\delta_1, \delta_2 \in \Delta$.

As a consequence of Theorem D.2, the random variables

$$\sup_{\delta\in\Delta}\big|\sqrt{n}\big(\hat{\rho}^{\pm}_{t_1,t_2}(\delta) - \rho^{\pm}_{t_1,t_2}(\delta)\big)\big|$$

converge in distribution to $\sup_{\delta\in\Delta}|G^{\pm}(\delta)|$. We now describe how to approximate the distribution of these suprema via a multiplier bootstrap.

D.2 Multiplier bootstrap for uniform bands of smoothed bounds

Define the estimated influence function contributions

$$\hat{\psi}^{\pm}_i(\delta) = \Psi^{\pm}_{t_1,t_2,\delta}\big(Z_i,\hat{\eta},\hat{\rho}^{\pm}_{t_1,t_2}(\delta)\big), \qquad i = 1,\dots,n, \;\; \delta\in\Delta,$$

where $\hat{\eta}$ and $\hat{\rho}^{\pm}_{t_1,t_2}(\delta)$ are obtained as in Eq. (12). Let

$$\bar{\hat{\psi}}^{\pm}_n(\delta) = \frac{1}{n}\sum_{i=1}^{n}\hat{\psi}^{\pm}_i(\delta)$$

denote the empirical mean of these contributions.

Let $\xi_1,\dots,\xi_n$ be i.i.d. multiplier weights with $\mathbb{E}[\xi_i] = 0$ and $\mathbb{E}[\xi_i^2] = 1$ (e.g., Rademacher or standard normal). The multiplier bootstrap process is defined by

$$\hat{G}^{\pm,*}_n(\delta) = \frac{1}{\sqrt{n}}\sum_{i=1}^{n}\xi_i\big(\hat{\psi}^{\pm}_i(\delta) - \bar{\hat{\psi}}^{\pm}_n(\delta)\big), \qquad \delta\in\Delta. \tag{77}$$

We then consider the sup-norm statistic

$$T^{\pm,*}_n = \sup_{\delta\in\Delta}\big|\hat{G}^{\pm,*}_n(\delta)\big|.$$

Repeating this construction B times with independent multipliers $\{\xi_i^{(b)}\}_{i=1}^n$, $b = 1,\dots,B$, yields bootstrap draws

$$T^{\pm,*}_{n,1},\dots,T^{\pm,*}_{n,B}.$$

Let $\hat{c}^{\pm}_{1-\alpha}$ denote the empirical $(1-\alpha)$ quantile of $\{T^{\pm,*}_{n,b}\}_{b=1}^B$. Under standard conditions for the validity of the multiplier bootstrap for suprema of empirical processes (which are implied by Assumption D.1 and mild additional technical assumptions), this quantity consistently estimates the $(1-\alpha)$ quantile of $\sup_{\delta\in\Delta}|G^{\pm}(\delta)|$.

The following corollary summarizes the resulting uniform confidence bands for the *smoothed* bounds $\rho^{\pm}_{t_1,t_2}(\delta)$, following standard multiplier bootstrap theory for empirical processes (Van Der Vaart & Wellner, 1996).

Corollary D.3 (Simultaneous confidence bands for smoothed Makarov bounds)**.** *Suppose the conditions of Theorem D.2 hold and the multiplier bootstrap described in Eq.* (77) *is valid for the processes in* (75)–(76)*. Let $\hat{c}^{\pm}_{1-\alpha}$ be the empirical $(1-\alpha)$ quantiles of $T^{\pm,*}_n$ as defined above. Then,*

$$\lim_{n\to\infty}\mathbb{P}\left(\forall\,\delta\in\Delta:\;\big|\hat{\rho}^{\pm}_{t_1,t_2}(\delta) - \rho^{\pm}_{t_1,t_2}(\delta)\big| \le \frac{\hat{c}^{\pm}_{1-\alpha}}{\sqrt{n}}\right) \;\ge\; 1-\alpha.$$

Equivalently, the bands

$$\delta\mapsto\Big[\hat{\rho}^{\pm}_{t_1,t_2}(\delta) - \tfrac{\hat{c}^{\pm}_{1-\alpha}}{\sqrt{n}},\; \hat{\rho}^{\pm}_{t_1,t_2}(\delta) + \tfrac{\hat{c}^{\pm}_{1-\alpha}}{\sqrt{n}}\Big]$$

are asymptotically valid $(1-\alpha)$ simultaneous confidence bands for the smoothed Makarov bounds $\rho^{\pm}_{t_1,t_2}(\delta)$, uniformly over $\delta\in\Delta$.

D.3 Uniformly valid bands for the original Makarov bounds

We now translate the simultaneous bands for the smoothed bounds $\rho^{\pm}_{t_1,t_2}(\delta)$ into simultaneous bands for the original Makarov bounds $\rho^{\pm}(\delta)$. Recall from Lemma 4.1 that

$$\rho^{-}_{t_1,t_2}(\delta) - b(t_1,t_2) \le \rho^{-}(\delta) \quad \text{and} \quad \rho^{+}(\delta) \le \rho^{+}_{t_1,t_2}(\delta) + b(t_1,t_2), \tag{78}$$

for every δ, where

$$b(t_1,t_2) \;=\; \frac{\log(2)}{t_2} + \frac{(\log|\mathcal{Y}|)_+}{t_1}$$

does not depend on δ. Hence, the inequalities in (78) hold uniformly for all $\delta \in \Delta$.

Combining Corollary D.3 with (78) yields the following result.

Corollary D.4 (Simultaneous confidence bands for the treatment effect distribution)**.** *Fix $\alpha \in (0,1)$ and a compact set $\Delta \subset \mathbb{R}$. Let $\hat{c}^{-,\mathrm{unif}}_{1-\alpha/2}$ and $\hat{c}^{+,\mathrm{unif}}_{1-\alpha/2}$ denote the bootstrap critical values obtained as in Corollary D.3 for the lower and upper smoothed bounds, respectively, with level $\alpha/2$. Define, for each $\delta \in \Delta$,*

$$\underline{c}_n(\delta) = \hat{\rho}^{-}_{t_1,t_2}(\delta) - \frac{\hat{c}^{-,\mathrm{unif}}_{1-\alpha/2}}{\sqrt{n}} - b(t_1,t_2), \tag{79}$$

$$\overline{c}_n(\delta) = \hat{\rho}^{+}_{t_1,t_2}(\delta) + \frac{\hat{c}^{+,\mathrm{unif}}_{1-\alpha/2}}{\sqrt{n}} + b(t_1,t_2). \tag{80}$$

Then

$$\liminf_{n\to\infty} \mathbb{P}\big(\forall\, \delta \in \Delta :\; \underline{c}_n(\delta) \le \rho(\delta) \le \overline{c}_n(\delta)\big) \;\ge\; 1-\alpha.$$

In particular, the pair of bands $\big\{\underline{c}_n(\delta)\big\}_{\delta\in\Delta}$ and $\big\{\overline{c}_n(\delta)\big\}_{\delta\in\Delta}$ defines an asymptotically valid $(1-\alpha)$ simultaneous confidence band for the treatment effect distribution $\rho(\delta)$ uniformly over $\delta \in \Delta$.

Practical implementation. In practice, the set Δ is approximated by a finite grid $\{\delta_1,\dots,\delta_K\} \subset \Delta$. The procedure then reduces to:

1. Compute $\hat{\rho}^{\pm}_{t_1,t_2}(\delta_k)$ and $\hat{\psi}^{\pm}_i(\delta_k)$ for all $k = 1,\dots,K$.
2. For each bootstrap replication $b = 1,\dots,B$, draw multipliers $\{\xi^{(b)}_i\}_{i=1}^n$ and form

$$\hat{G}^{\pm,*}_{n,b}(\delta_k) = \frac{1}{\sqrt{n}} \sum_{i=1}^{n} \xi^{(b)}_i \big(\hat{\psi}^{\pm}_i(\delta_k) - \bar{\hat{\psi}}^{\pm}_n(\delta_k)\big),$$

   and set $T^{\pm,*}_{n,b} = \max_{k=1,\dots,K} |\hat{G}^{\pm,*}_{n,b}(\delta_k)|$.
3. Let $\hat{c}^{\pm,\mathrm{unif}}_{1-\alpha/2}$ be the empirical $(1-\alpha/2)$ quantiles of $\{T^{\pm,*}_{n,b}\}_{b=1}^B$, and construct $\underline{c}_n(\delta_k)$ and $\overline{c}_n(\delta_k)$ as in (79)–(80).

For a sufficiently dense grid, the resulting discrete bands provide an accurate approximation to the uniform confidence bands over Δ.

E EXTENSION TO MULTIVALUED TREATMENTS

In this appendix, we briefly describe how our methodology extends to settings with multivalued treatments. Let the treatment take values in a finite set $\mathcal{A} \subset \mathbb{R}$ with $|\mathcal{A}| \geq 2$, and let

$$Z \;=\; (X, A, Y) \sim \mathbb{P}, \qquad A \in \mathcal{A},$$

where $X \in \mathcal{X} \subseteq \mathbb{R}^d$ and $Y \in \mathbb{R}$ as in the main text. For each $a \in \mathcal{A}$ we denote the potential outcome by $Y(a)$ and define the generalized propensity scores and response distributions by

$$\pi_a(x) \;=\; \mathbb{P}(A = a \mid X = x), \qquad F_a(y \mid x) \;=\; \mathbb{P}(Y \leq y \mid X = x, A = a).$$

E.1 TARGET ESTIMAND FOR A PAIR OF TREATMENT LEVELS

In many applications, the primary goal is to compare two specific treatment levels $a_1, a_0 \in \mathcal{A}$. For any such ordered pair (a_1, a_0) we define the treatment effect distribution

$$\rho_{a_1,a_0}(\delta) \;=\; \mathbb{P}\big(Y(a_1) - Y(a_0) \leq \delta\big), \qquad \delta \in \mathbb{R}. \tag{81}$$

This is directly analogous to the binary-treatment estimand in the main text, where $(a_1, a_0) = (1, 0)$.

Under the natural multivalued analogue of Assumption 3.1, namely:

1. *Consistency:* $Y(a) = Y$ whenever $A = a$ for all $a \in \mathcal{A}$,
2. *Overlap:* $0 < \pi_a(X) < 1$ almost surely for all $a \in \mathcal{A}$,
3. *Ignorability:* $A \perp \{Y(a) : a \in \mathcal{A}\} \mid X$,

the treatment effect distribution $\rho_{a_1,a_0}(\delta)$ is again partially identified by Makarov-type bounds based on the pair of marginals (F_{a_1}, F_{a_0}):

$$\rho^{-}_{a_1,a_0}(\delta) \;\leq\; \rho_{a_1,a_0}(\delta) \;\leq\; \rho^{+}_{a_1,a_0}(\delta), \tag{82}$$

where, for each fixed pair (a_1, a_0),

$$\rho^{-}_{a_1,a_0}(\delta) = \mathbb{E}\left[\sup_{y \in \mathcal{Y}} \big(F_{a_1}(y \mid X) - F_{a_0}(y - \delta \mid X)\big)_+\right], \tag{83}$$

and

$$\rho^{+}_{a_1,a_0}(\delta) = 1 + \mathbb{E}\left[\inf_{y \in \mathcal{Y}} \big(F_{a_1}(y \mid X) - F_{a_0}(y - \delta \mid X)\big)_-\right], \tag{84}$$

which are obtained by replacing F_1 and F_0 in Eq. (3) with F_{a_1} and F_{a_0}.

E.2 REDUCTION TO THE BINARY CASE

Our proposed estimators and theoretical results for the binary-treatment case extend directly to this multivalued setting by a simple recoding argument. For a fixed pair (a_1, a_0), define the binary indicators

$$\tilde{A}_{a_1,a_0} \;=\; \mathbf{1}\{A = a_1\}, \qquad 1 - \tilde{A}_{a_1,a_0} \;=\; \mathbf{1}\{A = a_0\},$$

and the corresponding generalized propensity scores

$$\pi_{a_1}(X) \;=\; \mathbb{P}(A = a_1 \mid X), \qquad \pi_{a_0}(X) \;=\; \mathbb{P}(A = a_0 \mid X).$$

Conditional on X, the distribution of $(\tilde{A}_{a_1,a_0}, Y)$ restricted to the subset $\{A \in \{a_1, a_0\}\}$ is algebraically identical to the binary-treatment setup in the main text with treatment $A = 1$ and control $A = 0$, after replacing:

$$A \;\rightsquigarrow\; \mathbf{1}\{A = a_1\}, \quad 1 - A \;\rightsquigarrow\; \mathbf{1}\{A = a_0\}, \quad \pi(X) \;\rightsquigarrow\; \pi_{a_1}(X), \quad 1 - \pi(X) \;\rightsquigarrow\; \pi_{a_0}(X),$$

and

$$F_1(\cdot \mid X) \;\rightsquigarrow\; F_{a_1}(\cdot \mid X), \qquad F_0(\cdot \mid X) \;\rightsquigarrow\; F_{a_0}(\cdot \mid X).$$

E.3 Influence functions and debiased estimators

Let $\rho^{\pm}_{a_1,a_0,t_1,t_2}(\delta)$ denote the smoothed Makarov bounds obtained by applying the smoothing construction from Lemma 4.1 to (F_{a_1}, F_{a_0}) instead of (F_1, F_0). The corresponding efficient influence functions follow by the same substitutions in Theorem 4.2.

Remark E.1 (EIF for multivalued treatments). Fix $a_1, a_0 \in \mathcal{A}$ and define

$$d_{a_1,a_0,\delta,\eta}(y \mid X) \;=\; F_{a_1}(y \mid X) - F_{a_0}(y - \delta \mid X),$$

and $I^{(a_1,a_0)}_{t,\delta}(x)$, $w^{(a_1,a_0)}_{t_1,\delta}(y \mid x)$, $\sigma^{\pm,(a_1,a_0)}_{t_1,t_2,\delta}(x)$ analogously to $I_{t,\delta}(x)$, $w_{t_1,\delta}(y \mid x)$ and $\sigma^{\pm}_{t_1,t_2,\delta}(x)$ in Theorem 4.2, but with F_1, F_0 replaced by F_{a_1}, F_{a_0}. Then the efficient influence functions for the smoothed multivalued bounds $\rho^{\pm}_{a_1,a_0,t_1,t_2}(\delta)$ are obtained from Theorem 4.2 by replacing

$$A \to \mathbf{1}\{A = a_1\},\;\; 1{-}A \to \mathbf{1}\{A = a_0\},\;\; \pi(X) \to \pi_{a_1}(X),\;\; 1{-}\pi(X) \to \pi_{a_0}(X),\;\; F_1 \to F_{a_1},\;\; F_0 \to F_{a_0}.$$

In particular, the debiased estimators in Eq. (12) extend verbatim to the comparison of any pair (a_1, a_0) after making these substitutions.

Because the derivation of the EIF in Theorem 4.2 relies only on the binary nature of the comparison (treated vs. control) and not on the cardinality of $\mathcal{A}$, the same arguments imply that all results on asymptotic normality, confidence intervals, and uniform confidence bands (Corollary 4.3 and Appendix D) carry over to the multivalued setting for any fixed pair (a_1, a_0).

In practice, a practitioner who wishes to compare the treatment levels a_1 and a_0 simply specifies these two values and applies our binary-treatment procedure to the recoded data with A replaced by $\mathbf{1}\{A = a_1\}$ and the nuisance functions (F_1, F_0, π) replaced by $(F_{a_1}, F_{a_0}, \pi_{a_1})$ and π_{a_0}, as described above.

F DETAILS ON NUISANCE ESTIMATION

We estimate the response c.d.f.s $F_a(y \mid x)$, $a \in \{0, 1\}$, by fitting separate conditional distribution models within each arm ($A = a$). All learners return the *entire* conditional c.d.f. evaluated on an arbitrary grid $y \in \mathcal{Y}$, which is required by the EIF in (4.2).

Training protocol (common to all learners). For each arm a we split the data into two folds as outlined in Algorithm 1. On each fold, we train on the complement and predict $\hat{F}_a(\cdot \mid x_i)$ for held-out i. Models are trained by gradient boosting (LightGBM) with early stopping on a validation set and a likelihood-based metric. Monotone transformations T of Y can be applied for stability; since T is monotone, $F_Y(y \mid x) = F_{T(Y)}(T(y) \mid x)$. In particular, we use standardization and log-transformation where appropriate. Given a grid $\{y_j\}_{j=1}^J$, each learner returns the matrix $\big(\hat{F}_a(y_j \mid x_i)\big)_{i,j}$. These values feed the smoothed operators in (5) and the weighted integrals $\hat{\Phi}^a_{t_1,\delta}(x_i, y_i)$ in (12).

Continuous outcomes: conditional Gaussian mixtures. For $Y \in \mathbb{R}$ we model $T(Y) \mid X = x$ as a K-component Gaussian mixture with covariate-dependent weights $\pi_k(x)$, means $\mu_k(x)$ and variances $\sigma_k^2(x)$. Boosting optimizes the (negative) log-likelihood with a custom objective

$$\ell_{\text{GM}}(x, y) \;=\; -\log\left(\sum_{k=1}^{K} \pi_k(x)\, \varphi\big(T(y); \mu_k(x), \sigma_k^2(x)\big)\right),$$

where $\varphi(y; \mu, \sigma^2)$ is the $\mathcal{N}(\mu, \sigma^2)$ density. The log-likelihood loss supplies per-parameter gradients and diagonal Hessians. Derivatives are rescaled to balance curvature across logits/means/variances. Initialization uses k-means on $T(Y)$ (cluster means/variances feed the intial score matrix. The resulting c.d.f. is

$$\hat{F}_a(y \mid x) = \sum_{k=1}^{K} \hat{\pi}_k(x)\, \Phi\left(\frac{T(y) - \hat{\mu}_k(x)}{\hat{\sigma}_k(x)}\right).$$

Discrete outcomes: multinomial classifier. For discrete Y taking finitely many values $\{c_1 < \cdots < c_M\}$ we fit a multiclass boosted classifier returning $\hat{p}_x(c_m) = \Pr(Y = c_m \mid X = x)$. The c.d.f. is the cumulative sum

$$\hat{F}_a(y \mid x) = \sum_{m:\, c_m \le y} \hat{p}_x(c_m),$$

implemented by summing predicted class probabilities over classes $\le y$.

Counts with excess zeros: zero–inflated Poisson (ZIP). For nonnegative counts we fit a covariate-dependent ZIP with rate $\lambda(x)$ and zero–inflation $\psi(x)$, learned via a custom objective for the ZIP log-likelihood and stabilized derivatives. Initialization uses the empirical mean for λ and the observed zero rate for ψ. The c.d.f. (right-continuous) is

$$\hat{F}_a(y \mid x) = \begin{cases} 0, & y < 0, \\ \psi(x) + (1 - \psi(x))\, \text{PoisCDF}(\lfloor y \rfloor; \lambda(x)), & y \ge 0. \end{cases}$$

G Details regarding synthetic data

Data generating process.

We sample uniform covariates $X \sim \mathrm{Unif}[0,1]$ and define $\rho(x) = \frac{1}{2}x^2 + \frac{1}{2}$. We then define the c.d.f. for the treatment arm as

$$Y(1) \mid X = x \sim \mathrm{Unif}\big([0,\rho(x)]\big), \qquad F_{1|x}(y) = \begin{cases} 0, & y < 0, \\ \dfrac{y}{\rho(x)}, & 0 \le y \le \rho(x), \\ 1, & y > \rho(x). \end{cases}$$

For the control arm, we define for a parameter γ

$$L_1(x) = -\tfrac{\gamma}{4}\,x, \qquad L_2(x) = \tfrac{1}{2}\rho(x) + \tfrac{\gamma}{4}\,x.$$

Then, we define the control arm c.d.f. as

$$F_{0|x}(y) = \begin{cases} 0, & y < L_1, \\ \dfrac{y - L_1}{\rho}, & L_1 \le y \le L_1 + \frac{1}{2}\rho, \\ \frac{1}{2}, & L_1 + \frac{1}{2}\rho < y < L_2, \\ \frac{1}{2} + \dfrac{y - L_2}{\rho}, & L_2 \le y \le L_2 + \frac{1}{2}\rho, \\ 1, & y > L_2 + \frac{1}{2}\rho, \end{cases}$$

where, for brevity, $L_j = L_j(x)$ and $\rho = \rho(x)$. The gap between the two control components is

$$G(x) \;=\; L_2 - \Big(L_1 + \tfrac{1}{2}\rho\Big) \;=\; \frac{\gamma}{2}\,x.$$

To model a joint distribution leveraging the marginals above, we draw (U_1, U_0) from a Gumbel copula C_θ with $\theta = 5$ (dimension 2), and set

$$Y(1) \mid X = F_{1|X}^{-1}(U_1), \qquad Y(0) \mid X = F_{0|X}^{-1}(U_0).$$

Finally, the observed outcome is $Y = A\,Y(1) + (1 - A)\,Y(0)$.

Quantifying margin violation. The difference $D_x(y) = F_{1|x}(y) - F_{0|x}(y)$ admits two plateaus that attain supremum and infimum and thus violating the margin assumption. The left, negative one is given for $0 \le y \le \frac{1}{2}\rho - \frac{\gamma}{4}x$ as

$$D_x(y) \;=\; \frac{y}{\rho} - \frac{y - L_1}{\rho} \;=\; \frac{L_1}{\rho} \;=\; -\frac{\gamma x}{4\,\rho}.$$

Its width is

$$W_{\mathrm{L}}(x) \;=\; \frac{\rho}{2} - \frac{\gamma}{4}x.$$

The right, positive plateau is given for $\frac{1}{2}\rho + \frac{\gamma}{4}x \le y \le \rho$ as

$$D_x(y) \;=\; \frac{y}{\rho} - \Big(\tfrac{1}{2} + \frac{y - L_2}{\rho}\Big) \;=\; \frac{L_2 - \frac{1}{2}\rho}{\rho} \;=\; +\frac{\gamma x}{4\,\rho}.$$

Its width is given by

$$W_{\mathrm{R}}(x) \;=\; \frac{\rho}{2} - \frac{\gamma}{4}x.$$

Averaging over X yields $\mathbb{E}[\rho(X)] = \frac{2}{3}$, $\mathbb{E}[X] = \frac{1}{2}$, and

$$\mathbb{E}\left[\frac{X}{\rho(X)}\right] = \int_0^1 \frac{2x}{x^2 + 1}\,dx = \ln 2.$$

Hence, obtain average plateau widths

$$\mathbb{E}\big[W_{\mathrm{L}}(X)\big] = \mathbb{E}\big[W_{\mathrm{R}}(X)\big] = \frac{1}{3} - \frac{\gamma}{8}, \qquad \mathbb{E}\big[G(X)\big] = \frac{\gamma}{4},$$

and normalized

$$\mathbb{E}\left[\frac{W_{\mathrm{L}}(X)}{\rho(X)}\right] = \mathbb{E}\left[\frac{W_{\mathrm{R}}(X)}{\rho(X)}\right] = \frac{1}{2} - \frac{\gamma}{4}\ln 2, \qquad \mathbb{E}\left[\frac{G(X)}{\rho(X)}\right] = \frac{\gamma}{2}\ln 2.$$

Thus, when descreasing γ, we maximize the average platenau width and correspondingly the degree of margin violation.

H DETAILS REGARDING SEMI-SYNTHETIC DATA

Real-world covariate data. The so-called *Oregon health insurance experiment*[2] (OHIE) (Finkelstein et al., 2012) is a randomized experiment that was intentionally conducted as a large-scale effort in public health to assess the effect of health insurance on several outcomes such as health or economic status. In 2008, a lottery draw offered low-income, uninsured adults in Oregon participation in a Medicaid program, providing health insurance. Individuals whose names were drawn could decide to sign up for the program.

In our analysis, we extract the following covariates X: age, gender, language, the number of emergency visits before the experiment, and the number of people the individual signed up with. The data collection in the OHIE was done as follows: after excluding individuals below the age of 19, above the age of 64, and individuals with residence outside of Oregon, 74,922 individuals were considered for the lottery. Among those, 29,834 were selected randomly and were offered participation in the program. Out of these, we randomly select $n = 3000$ data points.

Synthetic treatment and outcome generation. We define the empirical mean covariate

$$m(x) \;=\; \frac{1}{p}\sum_{j=1}^{p} x_j.$$

and introduce two covariate-dependent shape functions:

$$\rho(x) \;=\; \tfrac{1}{2}\, m(x)^2 + \rho_0, \qquad g(x) \;=\; m(x)^3,$$

where $\rho_0 > 0$ is a constant hyperparameter (in our experiments $\rho_0 = 0.5$). The treatment is randomized with constant propensity:

$$\pi(x) \equiv \mathbb{P}(A = 1 \mid X = x) = \tfrac{1}{2}, \qquad A_i \sim \text{Bernoulli}\big(\tfrac{1}{2}\big) \text{ independently of } X_i.$$

We define the treatment c.d.f. as

$$Y(1) \mid X = x \sim \text{Unif}\big([\,0,\ \rho(x)\,]\big), \qquad F_{1|x}(y) = \begin{cases} 0, & y < 0, \\ \dfrac{y}{\rho(x)}, & 0 \le y \le \rho(x), \\ 1, & y > \rho(x). \end{cases}$$

For the control arm c.d.f., we define

$$L_1(x) = -\tfrac{1}{4}\, g(x), \qquad L_2(x) = \tfrac{1}{2}\rho(x) + \tfrac{1}{4}\, g(x),$$

and

$$F_{0|x}(y) = \begin{cases} 0, & y < L_1, \\ \dfrac{y - L_1}{\rho}, & L_1 \le y \le L_1 + \tfrac{1}{2}\rho, \\ \tfrac{1}{2}, & L_1 + \tfrac{1}{2}\rho < y < L_2, \\ \tfrac{1}{2} + \dfrac{y - L_2}{\rho}, & L_2 \le y \le L_2 + \tfrac{1}{2}\rho, \\ 1, & y > L_2 + \tfrac{1}{2}\rho, \end{cases}$$

where, for brevity, $L_j = L_j(x)$ and $\rho = \rho(x)$.

To couple $Y(1)$ and $Y(0)$ while preserving the marginals above, we draw (U_1, U_0) from a Gumbel copula C_θ with $\theta = 5$ (dimension 2), and set

$$Y(1) = F^{-1}_{1|X}(U_1), \qquad Y(0) = F^{-1}_{0|X}(U_0).$$

Finally, the observed outcome is $Y = A\, Y(1) + (1 - A)\, Y(0)$.

[2]Data available here: https://www.nber.org/programs-projects/projects-and-centers/oregon-health-insurance-experiment

Quantifying margin violations. The c.d.f. difference $D_x(y) = F_{1|x}(y) - F_{0|x}(y)$ again admits two plateaus attaining the supremum and infemum, thus violating the margin assumption for both lower and upper bound.

For $0 \leq y \leq \frac{1}{2}\rho - \frac{1}{4}g$, the left plateau is

$$D_x(y) \;=\; \frac{y}{\rho} - \frac{y - L_1}{\rho} \;=\; \frac{L_1}{\rho} \;=\; -\frac{g}{4\,\rho},$$

with an associated width of $W_{\mathrm{L}}(x) = \dfrac{\rho}{2} - \dfrac{g}{4}$.

For $\frac{1}{2}\rho + \frac{1}{4}g \leq y \leq \rho$, the right plateau is

$$D_x(y) \;=\; \frac{y}{\rho} - \Big(\tfrac{1}{2} + \frac{y - L_2}{\rho}\Big) \;=\; \frac{L_2 - \frac{1}{2}\rho}{\rho} \;=\; +\frac{g}{4\,\rho}.$$

with a width of $W_{\mathrm{R}}(x) = \dfrac{\rho}{2} - \dfrac{g}{4}$. The average widths are given via

$$\mathbb{E}\big[W_{\mathrm{L}}(X)\big] = \mathbb{E}\big[W_{\mathrm{R}}(X)\big] = \mathbb{E}\Big[\big(\tfrac{1}{2}\rho_0 + \tfrac{1}{4}\big(m(X)^2 - m(X)^3\big)\big)_+\Big].$$

Hence, the widths of the plateaus and thus also the degrees of margin violations increasing in ρ_0.

Table 2: Approaches for debiased estimation of Makarov-bounds and other non-smooth causal estimands.

	Other causal estimands	**Makarov bounds**
Margin assumption	Policy value (Luedtke & Van Der Laan, 2016); Intersection bounds (Semenova, 2025);	(Melnychuk et al., 2024) (Semenova, 2025)
Smoothing	Optimal policy value (Park, 2024); IV bounds (Levis et al., 2025).	*Ours.*

I ADDITIONAL EXPERIMENTS

I.1 COVERAGE AND INTERVAL WIDTH FOR SYNTHETIC AND SEMI-SYNTHETIC SIMULATIONS

In this section, we present additional results in the synthetic and semi-synthetic settings of Section 5 studying the interval width and coverage properties of the envelope and smoothing estimators in Figure 5.

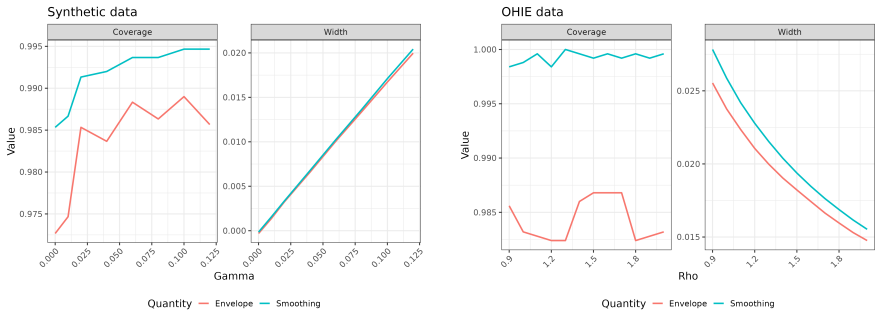

Figure 5: Coverage and interval width from the synthetic and semi-synthetic experiments in Section 5. Both estimators are conservative, with coverage above the target 95% level, due to the presence of bias, and the intervals for the smoothed estimator are the same or slightly wider than the intervals for the envelope estimator.

The results in Figure 5 show that both estimators are conservative, with coverage above the 95% level, and that the smoothed intervals are comparable to or slightly wider than the envelope intervals. The reason both estimators over-cover is that, as we saw in Figure 2, both estimators are biased, and a biased estimator can trivially achieve one-sided coverage. To see why, consider an estimator of the lower Makarov bound that is biased downward, so that it is typically less than the lower bound. Then, the corresponding interval will always have one-sided coverage of the Makarov lower bound, because the estimator itself is below the lower Makarov bound. Both estimators achieve high coverage for this reason, and so coverage does not provide a meaningful comparison between the two.

I.2 TWINS DATASET

Here, we use the binary outcome TWINS dataset from Louizos et al. (2017). The covariates X are measurements and demographic attributes of various twin pairs below 2000 grams. The treatment A is being born as the heavier twin, and the outcome Y denotes infant mortality within the first year of life. As both counterfactual outcomes are observed, we can estimate the ground-truth treatment effect c.d.f. , which we plot along the estimated Makarov bounds in Fig.6. We see that the estimated bounds are tight and cover the ground-truth treatment effect c.d.f. , thus confirming the validity of our estimator.

J ADDITIONAL TABLES AND FIGURES FROM THE MAIN TEXT

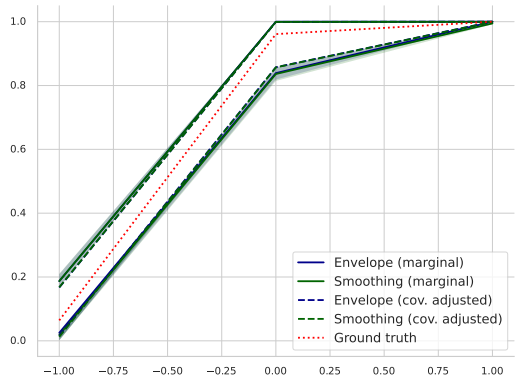

Figure 6: Estimated bounds and treatment effect c.d.f. for the TWINS dataset.

Data	Makarov type	Bound type	Small violations			Large violations		
			Plugin	Envelope	Ours	Plugin	Envelope	Ours
Synthetic	Marginal	lower	0.25 ± 0.01	0.06 ± 0.00	$\mathbf{0.04 \pm 0.00}$	0.41 ± 0.02	0.06 ± 0.01	$\mathbf{0.03 \pm 0.00}$
		upper	0.30 ± 0.01	0.07 ± 0.01	$\mathbf{0.05 \pm 0.00}$	0.41 ± 0.02	0.07 ± 0.01	$\mathbf{0.03 \pm 0.00}$
	Cov.-adjusted	lower	0.37 ± 0.02	0.05 ± 0.00	$\mathbf{0.04 \pm 0.00}$	0.54 ± 0.02	0.05 ± 0.00	$\mathbf{0.03 \pm 0.00}$
		upper	0.31 ± 0.01	0.06 ± 0.00	$\mathbf{0.05 \pm 0.00}$	0.51 ± 0.02	0.05 ± 0.00	$\mathbf{0.03 \pm 0.00}$
Semi-synthetic	Marginal	lower	0.06 ± 0.00	0.02 ± 0.00	$\mathbf{0.01 \pm 0.00}$	0.08 ± 0.00	0.02 ± 0.00	$\mathbf{0.01 \pm 0.00}$
		upper	0.09 ± 0.00	0.03 ± 0.00	$\mathbf{0.02 \pm 0.00}$	0.09 ± 0.00	0.03 ± 0.00	$\mathbf{0.01 \pm 0.00}$
	Cov.-adjusted	lower	0.29 ± 0.01	$\mathbf{0.03 \pm 0.00}$	$\mathbf{0.03 \pm 0.00}$	0.28 ± 0.01	0.02 ± 0.00	$\mathbf{0.01 \pm 0.00}$
		upper	0.33 ± 0.01	$\mathbf{0.03 \pm 0.00}$	$\mathbf{0.03 \pm 0.00}$	0.24 ± 0.01	$\mathbf{0.02 \pm 0.00}$	$\mathbf{0.02 \pm 0.00}$

Table 3: Additional comparison of mean-squared-error for different estimators, including results for the plugin estimator and the lower Makarov bound. All results are scaled by 100 for readability. Small and large assumption violations correspond to $\gamma = 0.12, \rho = 1.2$ and $\gamma = 0, \rho = 2$ in the synthetic and semi-synthetic settings respectively. Our estimator consistently achieves the best performance.