# ASSUMPTION-LEAN INFERENCE ON TREATMENT EFFECT DISTRIBUTIONS

**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 this workflow can overlook distributional risks, such as a large fraction of individuals affected negatively by the treatment. In this paper, we propose a novel method for inference on the treatment effect distribution. Prior work in this setting has estimated partial identification bounds—known as Makarov bounds—on the cumulative distribution function of the treatment effect by making restrictive assumptions on the outcome distribution. In contrast, our method guarantees accurate estimation and valid asymptotic inference of the Makarov bounds for any outcome distribution. Our main technical contributions are to develop smoothed surrogates for the Makarov bounds, derive semiparametrically efficient estimators of these surrogates, and propose a procedure for optimal selection of the smoothing parameters. We show empirically on synthetic and semi-synthetic datasets that, by not relying on the assumptions made by other methods, our estimators achieve a better bias-variance trade-off and lower mean-squared error. 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 enable unbiased comparisons of outcomes between groups. A common decision rule is to estimate the *average treatment effect* (ATE) and adopt the new treatment when the ATE is statistically significantly positive (Athey et al., 2020).

In many settings, the ATE is not sufficient for decision-making and practitioners in fact need estimates of distributional quantities beyond the average (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 minor tweak to a platform's ranking algorithm may lift overall engagement yet reduce engagement for new or low-activity users, increasing early churn. Such business-critical decisions motivate estimating the *entire distribution of the treatment effect*, rather than merely the mean.

Estimating the treatment effect distribution is challenging. Even in RCTs, where randomization rules out unobserved confounding, the joint distribution of potential outcomes is not identified (Fan & Park, 2010). Consequently, the treatment-effect distribution is only *partially* identified: its cumulative distribution function (c.d.f.) is bounded by the sharp *Makarov bounds* (Makarov, 1982). These bounds are obtained by optimizing over all possible joint distributions of treatment and control that are consistent with the observed marginal distributions of treatment and control. Hence the Makarov bounds are complex functionals of the data-generating distribution that involve suprema and infima over the outcome space. The non-smoothness of these suprema and infima complicates estimation of the bounds and especially complicates the application of standard tools from efficiency theory, such as debiased estimation (Kennedy, 2022).

Existing methods for efficient statistical inference on Makarov bounds either address stylized settings, such as a binary outcome (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. In this case, existing estimators are biased or high-variance, leading to inaccurate estimates and sub-optimal decisions.

In this paper, we propose a novel, assumption-lean, method for inference on Makarov bounds for the treatment-effect distribution. Our key idea is to replace the non-smooth suprema and infima that define the bounds with smooth surrogates, allowing us to apply 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 under violation of the margin assumption; (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<sup>1</sup>: (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.

# 2 RELATED WORK

**Distributions of treatment effects.** There are two main streams of work on statistical inference for treatment-effect distributions. The first stream aims at constructing *prediction intervals* for the individual treatment effect (ITE), typically based on conformal prediction (Lei & Candès, 2021; Alaa et al., 2023; Schröder et al., 2024). These methods provide valid predictive intervals for individual-level effects, but do not allow for inference on the full treatment-effect distribution (e.g., estimating its c.d.f. or density).

The second stream focuses on inference on the c.d.f. of the treatment-effect distribution. However, this distribution is not identifiable under standard assumptions (Rubin, 1974; Robins, 1986; Fan & Park, 2010), even with experimental data. Two approaches to overcome this obstacle have emerged. The first is to impose additional assumptions on the data-generating process to achieve point identification, as in Post & Van Den Heuvel (2025). Such assumptions, however, are unrealistic and untestable. The second is to take a partial identification approach and estimate upper and lower *bounds* on the cumulative distribution function. For example, sharp bounds and corresponding estimators have been proposed for binary outcomes (Kallus et al., 2022; Zhang & Richardson, 2025).

For general outcomes, sharp bounds are given by the Makarov bounds (Makarov, 1982). Estimation and inference methods for these bounds are mostly based on *plug-in* approaches (Fan & Park, 2010; Ruiz & Padilla, 2022; Fava, 2024; Lee, 2024; Cui & Han, 2023; Liang & Wu, 2025), which do not leverage semiparametric efficiency theory (Kennedy, 2022). The resulting estimators are not guaranteed to be asymptotically normal and confidence intervals are often constructed via invalid bootstrap procedures. Beyond Makarov bounds, Firpo & Ridder (2019) introduced uniformly sharp bounds for the treatment-effect distribution, but these are computationally challenging as no closed-form expressions are available. Similarly, Ji et al. (2023) proposed a dual-optimization framework for partial identification, but this framework can yield overly conservative bounds when strong duality fails (see their Sec. 6.2).

Efficient estimation of non-smooth causal estimands. A key technical difficulty in the estimation of the Makarov bounds is that the bounds involve suprema and infima over the outcome space. These suprema and

Table 1: Approaches for debiased estimation of Makarovbounds 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.

infima are non-differentiable, precluding the use of standard theory.

<sup>&</sup>lt;sup>1</sup>Code: https://anonymous.4open.science/r/SmoothedMakarovBounds-632B.

One literature stream attempts to remedy this by making a so-called *margin assumption* (Kitagawa & Tetenov, 2018), which essentially requires that the suprema and infima are attained at a single point only. This assumption has been used in the estimation of various other statistical estimands with similar non-differentiability issues: the optimal policy value (Luedtke & Van Der Laan, 2016), covariate-conditional Makarov bounds (Melnychuk et al., 2024), and more general intersection bounds (Semenova, 2025) which include Makarov bounds as a special case. However, the margin assumption is violated even in simple settings, like that of a constant treatment effect, making these methods unreliable in practice.

A second stream of work handles non-differentiability by approximating non-differentiable functions (e.g., the supremum) with smooth functions. This approach has been developed for inference on the optimal policy value Park (2024) and on bounds on average treatment effects in instrumental variable settings Levis et al. (2025). To the best of our knowledge no prior work has proposed a similar method for estimation and inference of Makarov bounds, which is the scope of our paper.

# 3 Problem setup

# 3.1 Setting

**Data:** We consider a causal inference setting using either randomized or observational data. That is, we consider a population  $Z=(X,A,Y)\sim \mathbb{P}$ , where  $X\in \mathcal{X}\subseteq \mathbb{R}^d$  are covariates,  $A\in \{0,1\}$  is a binary action, 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 use 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,

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

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 (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  $\pi$  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)

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

where

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

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 observation 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

**Plug-in estimation.** The simplest estimator of the Makarov bounds is the so-called *plug-in estimator* (Fan & Park, 2010): one first obtains nuisance estimators  $\hat{\eta} = (\hat{F}_1, \hat{F}_0, \hat{\pi})$  of the response distributions  $F_a$  and propensity score  $\pi$  (in the observational setting). Then, one substitutes the estimated nuisance functions into the expression for the bound from Eq. (3) to obtain

$$\hat{\rho}_{PI}^{-}(\delta) = \frac{1}{n} \sum_{i=1}^{n} \sup_{y \in \mathcal{Y}} \left( \hat{F}_{1}(y|x_{i}) - \hat{F}_{0}(y - \delta|x_{i}) \right)_{+}. \tag{4}$$

However, it turns out that the plugin estimator  $\hat{\rho}_{\rm PI}^{-}(\delta)$  has two drawbacks (Kennedy, 2022): (i) plug-in bias, which makes it asymptotically suboptimal, and (ii) difficulty establishing asymptotic normality under reasonable conditions on the nuisance estimators, which prevents reliable confidence interval construction.

**Semiparametric efficient estimators.** To address these limitations, a major line of work in causal inference leverages *semiparametric efficiency theory* to construct *debiased estimators*. Debiased estimators are of the form

$$\hat{\rho}_{\text{AIPTW}}^{-}(\delta) = \hat{\rho}_{\text{PI}}^{-}(\delta) + \frac{1}{n} \sum_{i=1}^{n} \Psi_{\delta, y^*}^{-} \left( z_i, \hat{\eta}, \hat{\rho}_{\text{PI}}^{-}(\delta) \right), \tag{5}$$

where  $\Psi^-_{\delta,y^*}(Z,\eta,\rho^-(\delta))$  is the so-called *efficient influence function* (EIF) of the target estimand. Adding the EIF term debiases the plug-in estimator, yielding an estimator that is asymptotically efficient and normally distributed, i.e., yielding the lowest possible asymptotic variance given by  $\mathbb{E}[\Psi^-_{\delta,y^*}(Z,\eta,\rho^-(\delta))^2]$ . Hence, debiasing using the EIF enables valid and efficient statistical inference.

**Envelope estimators for Makarov bounds.** For the lower Makarov bound, the efficient influence function is given by (Melnychuk et al., 2024; Semenova, 2025):

$$\Psi_{\delta,y^*}^{-}(Z,\eta,\rho^{-}(\delta)) = d_{\delta,\eta}(y^*|X) - \rho^{-}(\delta) + \mathbf{1}(d_{\delta,\eta}(y^*|X) > 0) \left(\frac{A}{\pi(X)} \left(\mathbf{1}(Y \le y^*) - F_1(y^*|X)\right)\right)$$
(6)  
$$-\frac{1-A}{1-\pi(X)} \left(\mathbf{1}(Y \le y^* - \delta) - F_0(y^* - \delta|X)\right),$$
(7)

where 
$$d_{\delta,\eta}(y|X) = F_1(y|X) - F_0(y - \delta|X)$$
 and  $y^* \in \arg\max_{y \in \mathcal{Y}} d_{\delta,\eta}(y|X)$ .

For each observation  $x_i$ ,  $y^*$  can be computed (e.g., via grid search) and plugged into the debiased estimator from Eq. (5). This approach is known as an *envelope estimator* (Semenova, 2025).

**Margin assumption.** Debiased estimators rely on several assumptions, including pathwise differentiability (van der Laan & Rubin, 2006). For Makarov bounds, this essentially translates to requiring the supremum and infimum  $y^*$  to yield unique solutions, which is known as the *margin assumption* (Kitagawa & Tetenov, 2018). However, this assumption is violated even for the simple case of constant treatment effects for a subset of users. Indeed, if the treatment effect is constant for users with some values of X, then the treatment c.d.f.  $F_1(y \mid X)$  for these users is a constant vertical shift of the control c.d.f.  $F_0(y \mid X)$ , and the supremum/ infimum  $y^*$  can be attained at many values of y (see Fig. 1 for an illustration). Thus the margin assumption fails and the envelope estimator is no longer efficient. This inefficiency can manifest a sub-optimal

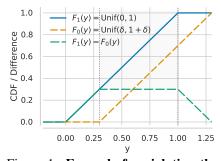


Figure 1: **Example for violating the margin assumption.** Depicted are two shifted (uniform) c.d.f.s that violate the margin assumption, leading to a plateau where the argmax in not unique.

mean-squared-error, because the margin assumption can inflate the variance of the estimator which can lead to suboptimal bias-variance trade-off as we will show later.

We emphasize that the failure of the margin assumption is not restricted to pathological edge cases. Constant treatment effects are a basic model of causal inference, and since effect sizes are often extremely small in digital experiments, we are often very nearly in the setting of a constant zero treatment effect. We cannot confidently use estimators that rely on a margin assumption when analyzing real-world data, necessitating the development of new methods.

# 4 ESTIMATING MAKAROV BOUNDS IN THE PRESENCE OF MARGIN VIOLATIONS

In this section, we present our methodology for valid inference on 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 in a principled way.

#### 4.1 A SMOOTH APPROXIMATION FOR THE MAKAROV BOUNDS

Two key components of the Makarov bounds in Eq. (3) prevent debiased inference in the presence of margin violations: (i) the supremum and infumum operators,  $\sup_{y \in \mathcal{Y}} \operatorname{and inf}_{y \in \mathcal{Y}}$ , 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 reliable confidence intervals.

**Bound approximations**. Our key idea is to replace these problematic components with smooth approximations that return unique values, even when the original functions are flat or multi-valued. This smoothing restores pathwise differentiability, enabling debiased estimation and valid inference 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)  $g_{t_1} \colon \mathcal{F} \to \mathbb{R}$  via

$$g_{t_1}(f) = \frac{1}{t_1} \log \left( \int \exp(t_1 f(y)) \, \mathrm{d}y \right) \tag{8}$$

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

**Lemma 4.1** (Makarov bound approximations). *Assume that*  $\mathcal{Y}$  *is compact with finite Lebesque measure*  $|\mathcal{Y}|$ . *For any*  $t_1, t_2 > 0$ , *we define the smoothed Makarov bounds via* 

$$\rho_{t_1,t_2}^-(\delta) = \mathbb{E}\left[\frac{1}{t_2}\log\left(1 + I_{t_1,\delta}(X)^{\frac{t_2}{t_1}}\right)\right] \text{ and } \rho_{t_1,t_2}^+(\delta) = 1 - \mathbb{E}\left[\frac{1}{t_2}\log\left(1 + I_{-t_1,\delta}(X)^{\frac{t_2}{t_1}}\right)\right],\tag{9}$$

where

$$I_{t,\delta}(x) = \int_{\mathcal{Y}} \exp(t[F_1(y|x) - F_0(y - \delta|x)]) dy. \tag{10}$$

Then, it holds that

$$\rho_{t_1,t_2}^-(\delta) - b(t_1,t_2) \le \rho^-(\delta) \text{ and } \rho^+(\delta) \le \rho_{t_1,t_2}^+(\delta) + b(t_1,t_2), \tag{11}$$

where  $b(t_1,t_2) = \frac{\log(2)}{t_2} + \frac{(\log |\mathcal{Y}|)_+}{t_1}$  quantifies the approximation bias. Furthermore,  $\rho_{t_1,t_2}^-(\delta) \to \rho^-(\delta)$  and  $\rho_{t_1,t_2}^+(\delta) \to \rho^+(\delta)$  as  $t_1,t_2 \to \infty$ .

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 makes the estimator more stable and allows for establishing asymptotic normality of debiased estimators.

#### 4.2 Debiased estimators

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

**Theorem 4.2.** 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

$$\Psi_{t_1,t_2,\delta}^-\left(Z,\eta,\rho_{t_1,t_2}^-(\delta)\right) = \sigma_{t_1,t_2,\delta}^-(X) \left[ \frac{A}{\pi(X)} \, \Phi_{t_1,\delta}^1(X,Y) \, - \, \frac{1-A}{1-\pi(X)} \, \Phi_{t_1,\delta}^0(X,Y) \right] \tag{12}$$

$$+\frac{1}{t_2}\log\left(1+I_{t_1,\delta}(X)^{\frac{t_2}{t_1}}\right)-\rho_{t_1,t_2}^{-}(\delta),\tag{13}$$

$$\Psi_{t_1,t_2,\delta}^+\left(Z,\eta,\rho_{t_1,t_2}^+(\delta)\right) = \sigma_{t_1,t_2,\delta}^+(X) \left[ \frac{A}{\pi(X)} \, \Phi_{-t_1,\delta}^1(X,Y) \, - \, \frac{1-A}{1-\pi(X)} \, \Phi_{-t_1,\delta}^0(X,Y) \right] \tag{14}$$

$$+1 - \frac{1}{t_2} \log \left(1 + I_{-t_1,\delta}(X)^{\frac{t_2}{t_1}}\right) - \rho_{t_1,t_2}^+(\delta), \tag{15}$$

where we denote

$$\Phi_{t_1,\delta}^a(X,Y) = \int_{\mathcal{Y}} w_{t_1,\delta}(y \mid X) \left( \mathbf{1}\{Y \le y - (1-a)\delta\} - F_a(y - (1-a)\delta \mid X) \right) dy, \quad (16)$$

$$w_{t_1,\delta}(y|x) = \frac{\exp(t_1(F_1(y|x) - F_0(y - \delta|x)))}{I_{t_1,\delta}(x)},$$
(17)

$$\sigma_{t_1,t_2,\delta}^-(x) = \frac{I_{t_1,\delta}(x)^{t_2/t_1}}{1 + I_{t_1,\delta}(x)^{t_2/t_1}}, \text{ and } \sigma_{t_1,t_2,\delta}^+(x) = \frac{I_{-t_1,\delta}(x)^{t_2/t_1}}{1 + I_{-t_1,\delta}(x)^{t_2/t_1}}.$$
 (18)

The smoothed EIFs from Theorem 4.2 are of similar structure as the envelope EIF from Eq. (6). Cruicially however, 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 Eq. (5), these are given by

$$\hat{\rho}_{t_1,t_2}^{-}(\delta) = \frac{1}{n} \sum_{i=1}^{n} \frac{1}{t_2} \log \left( 1 + \hat{I}_{t_1,\delta}(x_i)^{\frac{t_2}{t_1}} \right) + \hat{\sigma}_{t_1,t_2,\delta}^{-}(x_i) \left[ \frac{a_i - \hat{\pi}(x_i)}{\hat{\pi}(x_i)(1 - \hat{\pi}(x_i))} \, \hat{\Phi}_{t_1,\delta}^{a_i}(x_i, y_i) \right], \tag{19}$$

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

Correcting for approximation bias. Using above debiased estimators enables efficient and asymptotically normal inference, and thus construction of confidence intervals for the smoothed Makarov bounds. However, it remains to translate these intervals to the non-smoothed Makarov bounds, as the smoothing may introduce bias. Luckily, we can leverage Lemma 4.1 to upper bound this approximation bias via the term  $b(t_1,t_2)$ . We can then add this term to the one-sided confidence intervals for the upper and lower Makarov bounds, yielding valid confidence intervals as follows.

**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 (e.g., Chernozhukov et al., 2018), it holds for each fixed  $\delta$  that

$$\lim_{n \to \infty} \mathbb{P}\left(c_{t_1, t_2, \delta, \alpha}^-(\mathcal{D}_n) \le \rho(\delta) \le c_{t_1, t_2, \delta, \alpha}^+(\mathcal{D}_n)\right) \ge 1 - \alpha, \tag{20}$$

where

$$c_{t_1,t_2,\delta,\alpha}^{\pm}(\mathcal{D}_n) = \hat{\rho}_{t_1,t_2}^{\pm}(\delta) \pm z_{1-\alpha/2} \sqrt{\frac{1}{n^2} \sum_{i=1}^{n} \left( \Psi_{t_1,t_2,\delta}^{\pm}(z_i, \hat{\eta}, \hat{\rho}_{t_1,t_2}^{\pm}(\delta)) \right)^2} \pm b(t_1,t_2), \quad (21)$$

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

#### 4.3 SMOOTHING PARAMETER SELECTION

**Bias-variance trade-off.** Our smoothed debiased estimators from Eq. (19) depend on the smoothing parameters  $t_1$  and  $t_2$  that quantify the approximation quality of the non-smooth bound components. This implies a trade-off between the approximation-induced bias and the variance reduction under margin violation. Formally, a standard MSE decomposition yields

$$\mathbb{E}_{n}\left[\left(\hat{\rho}_{t_{1},t_{2}}^{-}(\delta) - \rho^{-}(\delta)\right)^{2}\right] = \underbrace{\left(\rho_{t_{1},t_{2}}^{-}(\delta) - \rho^{-}(\delta)\right)^{2}}_{\text{approximation (smoothing) bias}^{2}} + \underbrace{\frac{1}{n}\mathbb{E}\left[\Psi_{t_{1},t_{2},\delta}^{-}\left(Z,\eta,\rho_{t_{1},t_{2}}^{-}(\delta)\right)^{2}\right]}_{\text{asymptotic variance}/n} + \underbrace{0\left(\frac{1}{n}\right)}_{\text{remainder}}.$$
 (22)

As a consequence, choosing the smoothing parameters  $t_1$  and  $t_2$  correctly is crucial to obtain a low MSE that optimizes the bias-variance trade-off. We propose two different methods for data-driven smoothing-parameter selection: (i) minimizing an upper bound on the MSE, and (ii) using Lepski's method. Our full proposed procedure is shown in Algorithm 1.

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

- 1: **Input:** Dataset  $\mathcal{D} = \{(x_i, a_i, y_i)\}_{i=1}^n$
- Stage 0: Split the data D randomly into two disjoint datasets D<sub>1</sub> and D<sub>2</sub>.
- 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  $\hat{\rho}_{\tilde{t}_{1}^{-},\tilde{t}_{2}^{-}}^{-}(\delta)$  and  $\hat{\rho}_{\tilde{t}_{1}^{+},\tilde{t}_{2}^{+}}^{+}(\delta)$ .
- 6: **Output:**  $g_{\theta}$  and  $\alpha_{\phi}$ .

Minimizing an upper bound on the MSE. Using Lemma 4.1, we can replace the smoothing bias term in Eq. (22) with  $b(t_1,t_2)=\frac{\log(2)}{t_2}+\frac{(\log|\mathcal{Y}|)_+}{t_1}$ , which yields the datadriven tuning rule  $\arg\min_{t_1,t_2}b(t_1,t_2)^2+\frac{1}{n}\sum_{i=1}^n\Psi^-_{t_1,t_2,\delta}\big(z_i,\hat{\eta},\hat{\rho}^-_{t_1,t_2}(\delta)\big)^2$ , 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 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  $\left|\hat{\rho}_{t_1^{(k)},t_2^{(k)}}(\delta) - \hat{\rho}_{t_1^{(r)},t_2^{(r)}}(\delta)\right| \leq z_{1-\alpha/2} \widehat{\mathrm{SE}}(k,r)$  for all such r, where the tolerance uses EIF differences  $\widehat{\mathrm{SE}}(k,r) = 1/\sqrt{n} \operatorname{sd}_n\left(\widehat{\Psi}_{t_1^{(k)},t_2^{(k)},\delta}(Z_i,\hat{\eta},\hat{\rho}_{t_1^{(k)},t_2^{(k)}}(\delta)) - \widehat{\Psi}_{t_1^{(r)},t_1^{(r)},\delta}(Z_i,\hat{\eta},\hat{\rho}_{t_1^{(r)},t_2^{(r)}}(\delta))\right)$ , and  $\operatorname{sd}_n(\cdot)$  is the empirical standard deviation over  $i=1,\ldots,n$ .

			Smal	l assumption vi	olation	Large	e assumption vi	olation
Data	Bound	Side	Plugin	Envelope	Ours	Plugin	Envelope	Ours
Synthetic	Marginal Covadjusted	lower upper lower upper	$\begin{array}{c} 3.98 \pm 0.11 \\ 4.53 \pm 0.11 \\ 5.02 \pm 0.12 \\ 4.67 \pm 0.11 \end{array}$	$2.00 \pm 0.05$ $1.99 \pm 0.06$ $1.86 \pm 0.05$ $1.93 \pm 0.05$	$\begin{aligned} 1.81 &\pm 0.04 \\ 1.92 &\pm 0.05 \\ 1.81 &\pm 0.04 \\ 1.84 &\pm 0.04 \end{aligned}$	$\begin{array}{c} 5.59 \pm 0.11 \\ 5.55 \pm 0.11 \\ 6.42 \pm 0.13 \\ 6.29 \pm 0.12 \end{array}$	$\begin{array}{c} 1.35 \pm 0.08 \\ 1.40 \pm 0.08 \\ 1.28 \pm 0.07 \\ 1.25 \pm 0.07 \end{array}$	$\begin{array}{c} 0.82 \pm 0.06 \\ 0.83 \pm 0.06 \\ 0.90 \pm 0.06 \\ 0.86 \pm 0.06 \end{array}$
OHIE	Marginal Covadjusted	lower upper lower upper	$\begin{array}{c} 1.95 \pm 0.06 \\ 2.40 \pm 0.07 \\ 4.88 \pm 0.09 \\ 5.29 \pm 0.09 \end{array}$	$\begin{aligned} 1.21 &\pm 0.03 \\ 1.46 &\pm 0.04 \\ 1.46 &\pm 0.03 \\ 1.50 &\pm 0.04 \end{aligned}$	$egin{array}{l} 0.76 \pm 0.02 \ 1.09 \pm 0.03 \ 1.44 \pm 0.03 \ 1.41 \pm 0.03 \end{array}$	$\begin{array}{c} 2.22 \pm 0.06 \\ 2.52 \pm 0.07 \\ 4.71 \pm 0.09 \\ 4.31 \pm 0.09 \end{array}$	$1.14 \pm 0.03$ $1.31 \pm 0.04$ $1.09 \pm 0.03$ $1.28 \pm 0.03$	$\begin{array}{c} 0.77 \pm 0.02 \\ 0.97 \pm 0.03 \\ 1.00 \pm 0.02 \\ 1.08 \pm 0.03 \end{array}$

Table 2: Average mean-squared error of estimators of  $\rho(0)$  across settings. Error bars are 95% CIs across 2500 replications. Our method consistently attains the lowest mean-squared error.

# 5 EXPERIMENTS

We now confirm the effectiveness of our proposed method empirically. As is standard in causal inference (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 experimental results are in Appendix G.

**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 numerical integration. Details are in Appendix D.

**Baselines.** We compare our method with the *plug-in estimator* from Eq.(4) and the *envelope estimator* from Eq. (6). 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 empricial c.d.f. on a holdout dataset as a nuisance estimator. We use Lepski's method for selecting the smoothing parameters.

**Evaluation.** For synthetic data with available ground-truth we report the estimation error  $|\rho^{\pm}(\delta)| - \hat{\rho}^{\pm}(\delta)|$  averaged over K=2500 Monte Carlo runs for both the lower and the upper Makarov bound. For real-world data without ground-truth we report the estimated bounds and confidence intervals.

# 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 like those in Fig.1 in the difference between the treatment and control c.d.f. (see Appendix E).

We consider two variations of this DGP: (i) a variation with a smaller plateau width corresponding to only small assumption violation, and (ii) a variation with a larger plateau width corresponding to larger assumption violation. A comparison of our method to existing methods for both of these DGPs, for both the upper and lower bounds, with and without covariates, is given in Table. 2 (top) and shows that our **method achieves the best estimation performance** across all bound types and settings. Furthermore, the gap between our estimator and the baselines widens under stronger margin violation, highlighting that our method efficiently handles these violations when the baselines do not.

**Semi-synthetic data.** For our semi-synthetic experiments, we use covariate data 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 Appendix F 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. Again, we consider two different settings with different the severities of assumption violation. The results are

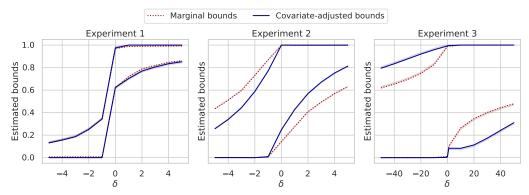


Figure 2: Estimated Makarov bounds for three A/B tests from a consumer technology company. Based on these bounds, we are confident that Experiment 1 does not negatively impact many users, and thus is safe to launch, whereas Experiments 2 and 3 may negatively impact many users.

shown in Table. 2 (bottom). Again, **our method achieves the best estimation performance** across all bound types and settings and the gap between the smoothing and envelope estimators widens when the degree of assumption violation increases.

#### 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 3. For each experiment, we estimate the Makarov bounds to obtain a partial identification region for the c.d.f. of the treatment effect.

Table 3: Experiment statistics summary.

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

*Results*. The results are shown in Figure 2. For the first experiment, we see that the partial identification region is relatively narrow. We are confident that  $\mathbb{P}(Y_i(1) - Y_i(0) \le -1) = \mathbb{P}(Y_i(1) - Y_i(0) < 0)$  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 3 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.)

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 in fact 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 potentially have a negative effect on nearly all users, and would be inadvisable to launch. Note that this experiment has a positive average treatment effect, as seen in Table 3, and so this recommendation to not launch contradicts the standard decision-making process. Taken together, these three examples highlight how our methods can be used in real world settings to characterize treatment effect distributions more completely than the average treatment effect, and thus inform launch decisions. We also provide additional comparisons with baselines in Appendix G.

**Discussion.** In this paper, we proposed a new method for inferring Makarov bounds on the treatment effect distribution without making a margin assumption. **Limitations and future work.** Our current approach is limited to static binary treatments. Future work may consider extending our approach to continuous or time-varying treatments and outcomes. **Broader impact.** Our method enables practitioners to make inferences about treatment effect distribution without relying on untestable assumptions, thus improving the reliability of established methods.

**Ethics statement.** We adhere to the ICLR Code of Ethics and acknowledge this during submission. Our work analyzes synthetic, semi-synthetic, and anonymized A/B-test data; no personal data are released or re-identified. The methods aim to improve safety by quantifying the fraction potentially harmed, but misuse is possible; we therefore emphasize risk-aware reporting and recommend subgroup audits when sensitive attributes are used.

**Reproducibility Statement.** An anonymized repository (linked in the submission) provides code to train/evaluate our estimators, select smoothing parameters, and compute one-sided confidence intervals. The paper specifies assumptions and estimands (Problem Setup), the smoothing/bias bounds and EIF-based estimators (Method), and evaluation protocols (Experiments), with full proofs and additional details in the appendix. We release synthetic and semi-synthetic generators, document preprocessing, model classes, hyperparameters, seeds, and scripts to recreate all synthetic and semi-synthetic figures/tables. Proprietary A/B data cannot be shared, but we plan to release anonymized data and code upon acceptence.

# REFERENCES

- 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.
  - 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, 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.
    - Yifan Cui and Sukjin Han. Policy learning with distributional welfare. *arXiv preprint* arXiv:2311.15878, 2023.
    - 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*, 2022.
    - Jacob Dorn, Kevin Guo, and Nathan Kallus. Doubly-valid/ doubly-sharp sensitivity analysis for causal inference with unmeasured confounding. *Journal of the American Statistical Association*, 2024.
    - 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*, 2023.
    - 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*, 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.
    - Sergio Firpo and Geert Ridder. Partial identification of the treatment effect distribution and its functionals. *Journal of Econometrics*, 213(1):210–234, 2019.
  - Dylan J. Foster and Vasilis Syrgkanis. Orthogonal statistical learning. *The Annals of Statistics*, 53(3): 879–908, 2023. ISSN 0090-5364.
    - 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.
  - Wenlong Ji, Lihua Lei, and Asher Spector. Model-agnostic covariate-assisted inference on partially identified causal effects. *arXiv* preprint arXiv:2310.08115, 2023.
    - 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.
  - Edward H. Kennedy. Semiparametric doubly robust targeted double machine learning: A review. *arXiv preprint*, 2022.
  - Edward H. Kennedy. Towards optimal doubly robust estimation of heterogeneous causal effects. *Electronic Journal of Statistics*, 17(2):3008–3049, 2023.
  - Edward H. Kennedy, Sivaraman Balakrishnan, and Larry Wasserman. Semiparametric counterfactual density estimation. *Biometrika*, 2023. ISSN 0006-3444.
  - 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*, pp. qkaf028, 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.
  - 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. *arXiv* preprint, arXiv:2407.03094, 2024.
  - 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.
  - 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.
  - Mark J. van der Laan and Donald B. Rubin. Targeted maximum likelihood learning. *The International Journal of Biostatistics*, 2(1), 2006.
  - 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.
  - Mingzhang Yin, Claudia Shi, Yixin Wang, and David M. Blei. Conformal sensitivity analysis for individual treatment effects. *Journal of the American Statistical Association*, pp. 1–14, 2022.
  - 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

$$g_{t_1}(f) = \frac{1}{t_1} \log \left( \int \exp(t_1 f(y)) \, \mathrm{d}y \right) \tag{23}$$

and that softplus function is defined as

$$h_{t_2}(u) = \frac{1}{t_2} \log(1 + e^{t_2 u}).$$
 (24)

We can now write the smoothed Makarov bounds as

$$\rho_{t_1,t_2}^{-}(\delta) = \mathbb{E}\left[\frac{1}{t_2}\log\left(1 + I_{t_1,\delta}(X)^{\frac{t_2}{t_1}}\right)\right] = \mathbb{E}\left[h_{t_2}\left(g_{t_1}(F_1(\cdot|X) - F_0(\cdot - \delta|X))\right)\right]$$
(25)

and

$$\rho_{t_1,t_2}^+(\delta) = 1 - \mathbb{E}\left[\frac{1}{t_2}\log\left(1 + I_{-t_1,\delta}(X)^{\frac{t_2}{t_1}}\right)\right] = 1 + \mathbb{E}\left[h_{-t_2}\left(g_{-t_1}(F_1(\cdot|X) - F_0(\cdot - \delta|X))\right)\right]. \tag{26}$$

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

$$h_{t_2}(u) - \frac{\log 2}{t_2} \le (u)_+ \text{ and } (u)_- \le h_{-t_2}(u) + \frac{\log 2}{t_2}.$$
 (27)

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

$$g_{t_1}(f) - \frac{\log |\mathcal{Y}|}{t_1} \le \sup_{y \in \mathcal{Y}} f(y) \text{ and } \inf_{y \in \mathcal{Y}} f(y) \le g_{-t_1}(f) + \frac{\log |\mathcal{Y}|}{t_1}. \tag{28}$$

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 (27) and (28) to Eq. (25) and obtain

$$\rho_{t_1,t_2}^{-}(\delta) \stackrel{(*)}{\leq} \mathbb{E}\left[h_{t_2}\left(\sup_{y \in \mathcal{Y}} (F_1(y|X) - F_0(y - \delta|X)) + \frac{\log|\mathcal{Y}|}{t_1}\right)\right]$$
(29)

$$\leq \mathbb{E}\left[\left(\sup_{y\in\mathcal{Y}}(F_1(y|X) - F_0(y - \delta|X)) + \frac{\log|\mathcal{Y}|}{t_1}\right)_{+}\right] + \frac{\log 2}{t_2} \tag{30}$$

$$\stackrel{(**)}{\leq} \rho^{-}(\delta) + \frac{\log |\mathcal{Y}|_{+}}{t_{1}} + \frac{\log 2}{t_{2}},\tag{31}$$

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

#### A.2 PROOF OF THEOREM 4.2

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

$$EIF\{I_{t_1,\delta}(x)\} = t_1 \int_{\mathcal{Y}} \exp(t_1 \left[F_1(y \mid x) - F_0(y - \delta \mid x)\right]) EIF\{F_1(y \mid x) - F_0(y - \delta \mid x)\} dy.$$
(32)

Hence, it holds that

$$EIF\{I_{t_1,\delta}(x)^{\frac{t_2}{t_1}}\} = \frac{t_2}{t_1}I_{t_1,\delta}(x)^{\frac{t_2}{t_1}-1}EIF\{I_{t_1,\delta}(x)\}$$
(33)

$$= t_2 I_{t_1,\delta}(x)^{\frac{t_2}{t_1}} \int_{\mathcal{V}} w_{t_1,\delta}(y|x) EIF \{ F_1(y \mid x) - F_0(y - \delta \mid x) \} dy, \quad (34)$$

and for the upper bound analogue

$$EIF\{I_{-t_{1},\delta}(x)^{\frac{t_{2}}{t_{1}}}\} = \frac{t_{2}}{t_{1}}I_{-t_{1},\delta}(x)^{\frac{t_{2}}{t_{1}}-1}EIF\{I_{-t_{1},\delta}(x)\}$$

$$= -t_{2}I_{-t_{1},\delta}(x)^{\frac{t_{2}}{t_{1}}}\int_{\mathcal{Y}}w_{-t_{1},\delta}(y|x)EIF\{F_{1}(y|x) - F_{0}(y-\delta|x)\}dy.$$
(36)

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

$$EIF\{F_a(y - \delta | X)\} = \frac{\mathbf{1}(A = a)\delta(X = x)}{\mathbb{P}(X = x, A = a)} \left( \mathbf{1}(Y \le y - \delta) - F_a(y - \delta | x) \right). \tag{37}$$

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

$$\Psi_{t_1,t_2,\delta}^-(Z,\eta) \tag{38}$$

$$= \int \frac{\mathbb{P}(X=x)}{t_2} EIF \log \left(1 + I_{t_1,\delta}(x)^{\frac{t_2}{t_1}}\right) dx \tag{39}$$

$$+\frac{1}{t_2}\log\left(1+I_{t_1,\delta}(X)^{\frac{t_2}{t_1}}\right)-\rho_{t_1,t_2}^-(\delta) \tag{40}$$

$$= \int \frac{\mathbb{P}(X=x)}{t_2} \frac{EIFI_{t_1,\delta}(x)^{\frac{t_2}{t_1}}}{1 + I_{t_1,\delta}(x)^{\frac{t_2}{t_1}}} \, \mathrm{d}x + \frac{1}{t_2} \log\left(1 + I_{t_1,\delta}(X)^{\frac{t_2}{t_1}}\right) - \rho_{t_1,t_2}^-(\delta) \tag{41}$$

$$= \int \Pr(X=x) \underbrace{\frac{I_{t_1,\delta}(x)^{\frac{t_2}{t_1}}}{1 + I_{t_1,\delta}(x)^{\frac{t_2}{t_1}}}}_{\sigma_{t_1,t_2,\delta}^-(x)} \left[ \int_{\mathcal{Y}} w_{t_1,\delta}(y|x) EIF\{F_1 - F_0\} \, dy \right] dx \tag{42}$$

$$+\frac{1}{t_2}\log\left(1+I_{t_1,\delta}(X)^{\frac{t_2}{t_1}}\right)-\rho_{t_1,t_2}^-(\delta) \tag{43}$$

$$= \int \Pr(X=x) \ \sigma_{t_1,t_2,\delta}^-(x) \Big[ \int_{\mathcal{Y}} w_{t_1,\delta}(y|x) \frac{A\delta(X=x)}{\mathbb{P}(X=x,A=1)} \left( \mathbf{1}(Y \le y) - F_1(y|x) \right) \ dy \Big] dx$$
(44)

$$-\int \Pr(X=x) \ \sigma_{t_1,t_2,\delta}^-(x) \Big[ \int_{\mathcal{Y}} w_{t_1,\delta}(y|x) \frac{(1-A)\delta(X=x)}{\mathbb{P}(X=x,A=0)} \left( \mathbf{1}(Y \le y-\delta) - F_0(y-\delta|x) \right) \ dy \Big] dx$$
(45)

$$+\frac{1}{t_2}\log\left(1+I_{t_1,\delta}(X)^{\frac{t_2}{t_1}}\right)-\rho_{t_1,t_2}^{-}(\delta) \tag{46}$$

$$= \sigma_{t_1, t_2, \delta}^-(X) \left[ \int_{\mathcal{Y}} w_{t_1, \delta}(y|X) \frac{A}{\pi(X)} \left( \mathbf{1}(Y \le y) - F_1(y|X) \right) dy \right]$$
 (47)

$$-\sigma_{t_1,t_2,\delta}^-(X) \left[ \int_{\mathcal{Y}} w_{t_1}(y|X) \frac{(1-A)}{1-\pi(X)} \left( \mathbf{1}(Y \le y - \delta) - F_0(y - \delta|X) \right) dy \right]$$
 (48)

$$+\frac{1}{t_2}\log\left(1+I_{t_1,\delta}(X)^{\frac{t_2}{t_1}}\right)-\rho_{t_1,t_2}^{-}(\delta) \tag{49}$$

(50)

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

$$\Psi_{t_1,t_2,\delta}^+(Z,\eta) \tag{51}$$

$$= \int \mathbb{P}(X=x)EIF\left\{1 - \frac{1}{t_2}\log\left(1 + I_{-t_1,\delta}(x)^{\frac{t_2}{t_1}}\right)\right\} dx$$
 (52)

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

$$= \int \frac{\mathbb{P}(X=x)}{-t_2} EIF \log \left(1 + I_{-t_1,\delta}(x)^{\frac{t_2}{t_1}}\right) dx \tag{54}$$

$$+1 - \frac{1}{t_2} \log \left(1 + I_{-t_1,\delta}(X)^{\frac{t_2}{t_1}}\right) - \rho_{t_1,t_2}^+(\delta) \tag{55}$$

$$= \int \frac{\mathbb{P}(X=x)}{-t_2} \frac{EIFI_{-t_1,\delta}(x)^{\frac{t_2}{t_1}}}{1 + I_{-t_1,\delta}(x)^{\frac{t_2}{t_1}}} dx + 1 - \frac{1}{t_2} \log\left(1 + I_{-t_1,\delta}(X)^{\frac{t_2}{t_1}}\right) - \rho_{t_1,t_2}^+(\delta)$$
 (56)

$$= \int \Pr(X=x) \underbrace{\frac{I_{-t_1,\delta}(x)^{\frac{t_2}{t_1}}}{1 + I_{-t_1,\delta}(x)^{\frac{t_2}{t_1}}}}_{\sigma_{t_1,t_2,\delta}^+(x)} \Big[ \int_{\mathcal{Y}} w_{-t_1,\delta}(y|x) EIF\{F_1 - F_0\} \, dy \Big] dx$$
 (57)

$$+1 - \frac{1}{t_2} \log \left(1 + I_{-t_1,\delta}(X)^{\frac{t_2}{t_1}}\right) - \rho_{t_1,t_2}^+(\delta)$$
(58)

$$= \int \Pr(X=x) \ \sigma_{t_1,t_2,\delta}^+(x) \Big[ \int_{\mathcal{Y}} w_{-t_1,\delta}(y|x) \frac{A\delta(X=x)}{\mathbb{P}(X=x,A=1)} \left( \mathbf{1}(Y \le y) - F_1(y|x) \right) \ dy \Big] dx$$
(59)

$$-\int \Pr(X=x) \ \sigma_{t_1,t_2,\delta}^+(x) \Big[ \int_{\mathcal{Y}} w_{-t_1,\delta}(y|x) \frac{(1-A)\delta(X=x)}{\mathbb{P}(X=x,A=0)} \left( \mathbf{1}(Y \le y-\delta) - F_0(y-\delta|x) \right) \ dy \Big] dx$$
(60)

$$+1 - \frac{1}{t_2} \log \left(1 + I_{-t_1,\delta}(X)^{\frac{t_2}{t_1}}\right) - \rho_{t_1,t_2}^+(\delta) \tag{61}$$

$$= \sigma_{t_1, t_2, \delta}^+(X) \left[ \int_{\mathcal{Y}} w_{-t_1, \delta}(y|X) \frac{A}{\pi(X)} \left( \mathbf{1}(Y \le y) - F_1(y|X) \right) dy \right]$$
 (62)

$$-\sigma_{t_1,t_2,\delta}^+(X) \left[ \int_{\mathcal{Y}} w_{-t_1}(y|X) \frac{(1-A)}{1-\pi(X)} \left( \mathbf{1}(Y \le y - \delta) - F_0(y - \delta|X) \right) dy \right]$$
 (63)

$$+1 - \frac{1}{t_2} \log \left(1 + I_{-t_1,\delta}(X)^{\frac{t_2}{t_1}}\right) - \rho_{t_1,t_2}^+(\delta). \tag{64}$$

### A.3 PROOF OF COROLLARY 4.3

*Proof (sketch)*. Fix  $\delta$  and abbreviate  $\theta^{\pm}(\delta) = \rho_{t_1,t_2}^{\pm}(\delta)$  and  $\hat{\theta}^{\pm}(\delta) = \hat{\rho}_{t_1,t_2}^{\pm}(\delta)$ . By Theorem 4.2 and standard semiparametric arguments (Neyman orthogonality with sample splitting / cross-fitting and the stated rate conditions; cf. Chernozhukov et al., 2018), the debiased estimators admit the linear expansion

$$\hat{\theta}^{\pm}(\delta) - \theta^{\pm}(\delta) = \frac{1}{n} \sum_{i=1}^{n} \Psi_{t_1, t_2, \delta}^{\pm} \left( Z_i, \eta, \theta^{\pm}(\delta) \right) + r_n^{\pm}, \qquad r_n^{\pm} = o_{\P}(n^{-1/2}).$$

Let  $V^\pm(\delta)=\mathrm{Var}ig(\Psi^\pm_{t_1,t_2,\delta}(Z,\eta,\theta^\pm(\delta))ig)\in(0,\infty)$ . A Lindeberg–Feller CLT yields

$$\sqrt{n}(\hat{\theta}^{\pm}(\delta) - \theta^{\pm}(\delta)) \Rightarrow \mathcal{N}(0, V^{\pm}(\delta)).$$

Define 
$$\widehat{V}^{\pm}(\delta) = \frac{1}{n} \sum_{i=1}^{n} \left( \Psi_{t_1,t_2,\delta}^{\pm}(Z_i,\hat{\eta},\hat{\theta}^{\pm}(\delta)) - \bar{\Psi}^{\pm} \right)^2$$
 with  $\bar{\Psi}^{\pm} = \frac{1}{n} \sum_{i=1}^{n} \Psi_{t_1,t_2,\delta}^{\pm}(Z_i,\hat{\eta},\hat{\theta}^{\pm}(\delta))$  and set  $\widehat{\operatorname{se}}^{\pm}(\delta) = \sqrt{\widehat{V}^{\pm}(\delta)/n}$ . By a law of large numbers

under cross-fitting,  $\widehat{V}^{\pm}(\delta) \xrightarrow{p} V^{\pm}(\delta)$ , hence by Slutsky

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

Consequently,

$$\mathbb{P}\!\!\left\{\theta^-(\delta) \geq \hat{\theta}^-(\delta) - z_{1-\alpha/2}\widehat{\operatorname{se}}^-(\delta)\right\} \to 1-\alpha/2, \qquad \mathbb{P}\!\!\left\{\theta^+(\delta) \leq \hat{\theta}^+(\delta) + z_{1-\alpha/2}\widehat{\operatorname{se}}^+(\delta)\right\} \to 1-\alpha/2.$$

By Lemma 4.1, for all fixed  $t_1, t_2 > 0$ ,

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

Therefore, on the intersection of the two one-sided events above, it follows that

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

i.e.,  $c_{t_1,t_2,\delta,\alpha}^-(\mathcal{D}_n) \leq \rho(\delta) \leq c_{t_1,t_2,\delta,\alpha}^+(\mathcal{D}_n)$ . A Bonferroni bound gives

$$\lim_{n \to \infty} \Pr \left\{ c_{t_1, t_2, \delta, \alpha}^-(\mathcal{D}_n) \le \rho(\delta) \le c_{t_1, t_2, \delta, \alpha}^+(\mathcal{D}_n) \right\} \ge 1 - \alpha,$$

which proves the result.

# B EXTENDED RELATED WORK

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, 2022). 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. (2023); 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 (2022); Dorn et al. (2024); Yin et al. (2022); Frauen et al. (2023); Jin et al. (2023); Frauen et al. (2024).

# C EXTENSION TO DISCRETE OUTCOMES.

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

**Discrete log-sum-exp and softplus.** The softplus  $h_{t_2}(u) = t_2^{-1} \log(1 + e^{t_2 u})$  is unchanged. We replace the continuous log-sum-exp with its discrete analogue

$$g_{t_1}^{(d)}(f) = \frac{1}{t_1} \log \left( \sum_{y \in \mathcal{Y}_d} \exp\{t_1 f(y)\} \right).$$
 (65)

Accordingly, define the discrete normalizer

$$I_{t,\delta}^{(d)}(x) = \sum_{y \in \mathcal{V}_d} \exp \Big( t \big[ F_1(y \mid x) - F_0(y - \delta \mid x) \big] \Big). \tag{66}$$

**Smoothed Makarov bounds (discrete).** Replacing the operators in Sec. 4.1 by (65)–(66) yields the discrete smoothed bounds

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

The analogue of Lemma 4.1 holds with the discrete approximation bias

$$b_d(t_1, t_2) = \frac{\log 2}{t_2} + \frac{\log M}{t_1},\tag{68}$$

i.e. 
$$\rho_{t_1,t_2}^{-,(d)}(\delta) - b_d(t_1,t_2) \leq \rho^-(\delta)$$
 and  $\rho^+(\delta) \leq \rho_{t_1,t_2}^{+,(d)}(\delta) + b_d(t_1,t_2)$ , with  $\rho_{t_1,t_2}^{\pm,(d)}(\delta) \to \rho^\pm(\delta)$  as  $t_1,t_2 \to \infty$ .

**Efficient influence functions (discrete).** The EIFs in Theorem 4.2 carry over after replacing all integrals over y by sums over  $y_d$ . Define the discrete softmax weights

$$w_{t_1,\delta}^{(d)}(y \mid x) = \frac{\exp(t_1[F_1(y \mid x) - F_0(y - \delta \mid x)])}{I_{t_1,\delta}^{(d)}(x)}, \quad y \in \mathcal{Y}_d,$$
 (69)

and the discrete analogue of  $\Phi^a_{t_1,\delta}$ ,

$$\Phi_{t_1,\delta}^{a,(d)}(X,Y) = \sum_{y \in \mathcal{Y}_d} w_{t_1,\delta}^{(d)}(y \mid X) \left( \mathbf{1} \{ Y \le y - (1-a)\delta \} - F_a \left( y - (1-a)\delta \mid X \right) \right). \tag{70}$$

With  $\sigma_{t_1,t_2,\delta}^{\pm}(x)$  defined exactly as in the main text but using  $I_{\pm t_1,\delta}^{(d)}(x)$ , the EIFs are

$$\Psi_{t_{1},t_{2},\delta}^{-,(d)}(Z,\eta,\rho_{t_{1},t_{2}}^{-,(d)}(\delta)) = \sigma_{t_{1},t_{2},\delta}^{-}(X) \left[ \frac{A}{\pi(X)} \Phi_{t_{1},\delta}^{1,(d)}(X,Y) - \frac{1-A}{1-\pi(X)} \Phi_{t_{1},\delta}^{0,(d)}(X,Y) \right] \\
+ \frac{1}{t_{2}} \log \left( 1 + \left( I_{t_{1},\delta}^{(d)}(X) \right)^{t_{2}/t_{1}} \right) - \rho_{t_{1},t_{2}}^{-,(d)}(\delta). \tag{71}$$

$$\Psi_{t_{1},t_{2},\delta}^{+,(d)}(Z,\eta,\rho_{t_{1},t_{2}}^{+,(d)}(\delta)) = \sigma_{t_{1},t_{2},\delta}^{+}(X) \left[ \frac{A}{\pi(X)} \Phi_{-t_{1},\delta}^{1,(d)}(X,Y) - \frac{1-A}{1-\pi(X)} \Phi_{-t_{1},\delta}^{0,(d)}(X,Y) \right] + 1 - \frac{1}{t_{2}} \log \left( 1 + \left( I_{-t_{1},\delta}^{(d)}(X) \right)^{t_{2}/t_{1}} \right) - \rho_{t_{1},t_{2}}^{+,(d)}(\delta).$$

$$(72)$$

**Debiased estimators (discrete).** The one-step estimators in Eq. (19) become

$$\hat{\rho}_{t_{1},t_{2}}^{-,(d)}(\delta) = \frac{1}{n} \sum_{i=1}^{n} \left\{ \frac{1}{t_{2}} \log \left( 1 + \left( \hat{I}_{t_{1},\delta}^{(d)}(x_{i}) \right)^{t_{2}/t_{1}} \right) + \hat{\sigma}_{t_{1},t_{2},\delta}^{-}(x_{i}) \frac{a_{i} - \hat{\pi}(x_{i})}{\hat{\pi}(x_{i})\{1 - \hat{\pi}(x_{i})\}} \hat{\Phi}_{t_{1},\delta}^{a_{i},(d)}(x_{i},y_{i}) \right\}.$$

$$\hat{\rho}_{t_{1},t_{2}}^{+,(d)}(\delta) = \frac{1}{n} \sum_{i=1}^{n} \left\{ 1 - \frac{1}{t_{2}} \log \left( 1 + \left( \hat{I}_{-t_{1},\delta}^{(d)}(x_{i}) \right)^{t_{2}/t_{1}} \right) + \hat{\sigma}_{t_{1},t_{2},\delta}^{+}(x_{i}) \frac{a_{i} - \hat{\pi}(x_{i})}{\hat{\pi}(x_{i})\{1 - \hat{\pi}(x_{i})\}} \hat{\Phi}_{-t_{1},\delta}^{a_{i},(d)}(x_{i},y_{i}) \right\}.$$

$$(74)$$

Here  $\hat{I}^{(d)}_{\pm t_1,\delta}$ ,  $\hat{w}^{(d)}_{t_1,\delta}$ , 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 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 numerical 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  $(\hat{F}_a(y_j \mid x_i))_{i,j}$ . These values feed the smoothed operators in (9) and the weighted integrals  $\hat{\Phi}_{t_1,\delta}^a(x_i,y_i)$  in (19).

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

$$\ell_{\text{GM}}(x,y) = -\log\Biggl(\sum_{k=1}^K \pi_k(x) \, \varphi\bigl(T(y); \mu_k(x), \sigma_k^2(x)\bigr)\Biggr),$$

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  $\leq y$ .

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  $\lambda$  and the observed zero rate for  $\psi$ . The c.d.f. (right-continuous) is

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

# E DETAILS REGARDING SYNTHETIC DATA

# Data generating process.

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

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

For the control arm, we define for a parameter  $\gamma$ 

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

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

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

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

$$G(x) = L_2 - \left(L_1 + \frac{1}{2}\rho\right) = \frac{\gamma}{2}x.$$

To model a joint distribution leveraging the marginals above, we draw  $(U_1, U_0)$  from a Gumbel copula  $C_{\theta}$  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 = AY(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

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

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

$$D_x(y) = \frac{y}{\rho} - \left(\frac{1}{2} + \frac{y - L_2}{\rho}\right) = \frac{L_2 - \frac{1}{2}\rho}{\rho} = +\frac{\gamma x}{4\rho}.$$

Its width is given by

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

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

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

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

# F DETAILS REGARDING SEMI-SYNTHETIC DATA

Real-world covariate data. The so-called *Oregon health insurance experiment*<sup>2</sup> (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:

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

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

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

We define the treatment c.d.f. as

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

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

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

and

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

where, for brevity,  $L_i = L_i(x)$  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  $C_{\theta}$  with  $\theta = 5$  (dimension 2), and set

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

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

 $<sup>^2</sup> Data\ available\ here:\ https://www.nber.org/programs-projects/projects-and-centers/oregon-health-insurance-experiment$ 

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  $0 \le y \le \frac{1}{2}\rho - \frac{1}{4}g$ , 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  $W_{\rm L}(x)=rac{
ho}{2}-rac{g}{4}.$ 

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

$$D_x(y) = \frac{y}{\rho} - \left(\frac{1}{2} + \frac{y - L_2}{\rho}\right) = \frac{L_2 - \frac{1}{2}\rho}{\rho} = +\frac{g}{4\rho}.$$

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

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

Hence, the widths of the plateaus and thus also the degrees of margin violations increasing in  $\rho_0$ .

# G ADDITIONAL EXPERIMENTS

#### G.1 (SEMI)-SYNTHETIC DATA

Here, we report full results for the synthetic and semi-synthetic datasets from Appendix E and Appendix F for a range of parameters  $\gamma$  (synthetic) and  $\rho$  (semi-synthetic) to complement the results from Table 2. Recall that increasing  $\gamma$  decreases margin violation, while higher  $\rho$  increase margin violation. The results for all methods are bound types are reported in Fig.3. For both datasets, larger margin violation leads to a larger gap between our method and the baselines. Importantly, our method consistently outperforms the baselines.

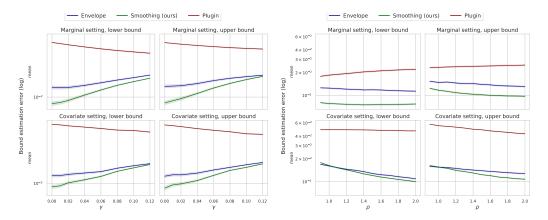


Figure 3: Results for synthetic and semi-synthetic data over a range of parameters quantifying margin violation.

# G.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.4. 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.

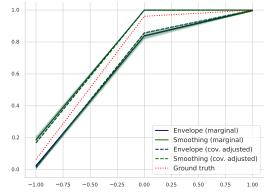


Figure 4: Estimated bounds and treatment effect c.d.f. for the TWINS dataset.

#### G.3 ADDITIONAL RESULTS FOR CONSUMER COMPANY A/B TESTS

In Figure 5, we compare both the results of our smoothed estimator with the results of the envelope baseline. The upper Makarov bound estimates mostly coincide, however, our estimator is more conservative and estimates a smaller lower bound than the baseline. As our estimator is obtained by

minimizing MSE over various smoothing parameters (including the special case of large smoothing parameters coinciding with the envelope), this indicates that the ground-truth Makarov bound is smaller than the envelope estimate. In particular, this may indicate envelope estimator is too confident and may undercover the ground-truth treatment effect c.d.f., leading to potentially wrong conclusions.

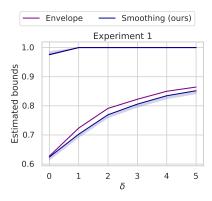


Figure 5: Estimated bounds and confidence intervals for Experiment 1.