WILL OUR TREATMENT POLICY REALLY WORK? MULTISCALE SEGMENTATION OF TREATMENT-BENEFIT REGIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

Heterogeneous treatment effect estimation informs a decision maker who benefits most, but not whether the resulting rule is deployable: how many distinct beneficiary groups must be handled, and what is lost when the segment budget is smaller. Existing policy learning methods impose constraints on capacity or the policy class, yet policies with identical value and coverage can have very different geometry and implementation cost. We study deployability through the topology of treatment-benefit superlevel sets. Connected components count distinct beneficiary segments, while persistence across nearby thresholds identifies segments stable to perturbations in the effect estimate and threshold. We propose a post-estimation deployability layer, compatible with any CATE estimator, that reports persistent zeroth Betti numbers before and after overlap trimming and, when more segments survive than the budget allows, returns a conservative policy targeting those with the largest estimated gains. We prove that persistent segmentation lower-bounds conservative policy complexity, insufficient segment capacity forces regret, and uniform accuracy of the effect estimate transfers to stable segmentation summaries. Empirical studies show that the proposed method uncovers deployment failure modes not captured by existing policy learning approaches.

1 INTRODUCTION

Heterogeneous treatment effect estimation characterizes how treatment benefits vary across subjects with different covariate profiles and provides a foundation for individualized treatment decisions. A common target is the conditional average treatment effect (CATE), whose estimate is widely used to guide treatment assignment.

One line of work considers an unrestricted policy class, where the welfare-maximizing treatment rule treats individuals with CATE above a threshold t, i.e., $\mathbf{1}\{\tau(x) > t\}$ (Hirano & Porter, 2009). The threshold may reflect a capacity constraint or treatment cost, with capacity-constrained treatment selection studied extensively in prior work (Bhattacharya & Dupas, 2012; Luedtke & van der Laan, 2016; VanderWeele et al., 2019; Imai & Li, 2023). In practice, this rule is implemented using an estimated CATE and, when needed, an estimated threshold. However, the estimated CATE surface can be highly irregular, and the resulting threshold policy may induce an arbitrarily complex treatment-benefit region consisting of many small, fragmented components in covariate space. When a policymaker requires a simple treatment rule that can be communicated, audited, and implemented based on a small number of distinct beneficiary segments, such a policy may be impractical to deploy even if it closely approximates the optimal policy.

Alternatively, another line of work obtains interpretable treatment rules by optimizing welfare over a prespecified structured class, including decision lists, policy trees, and unions of hyperboxes (Zhang et al., 2015; Lakkaraju & Rudin, 2017; Kitagawa & Tetenov, 2018; Athey & Wager, 2021; Zhou et al., 2023; Tschernutter et al., 2022). Most closely related, Tschernutter et al. (2022) represent the treated region as a union of at most a prespecified number of hyperboxes. Such approaches impose representational simplicity ex ante. In contrast, we start from the threshold policy induced by an arbitrary CATE estimator and ask whether its treatment-benefit region consists of a small number of

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.

stable, overlap-supported connected components, and what is lost when only a limited number can be deployed.

We propose evaluating the deployability of an estimated policy based on the (i) *structural simplicity*, (ii) *stability*, and (iii) *causal credibility* of its induced treatment-benefit region. Highly fragmented regions are difficult to operationalize; unstable beneficiary segments do not provide a reliable basis for treatment decisions; and high-benefit segments in low-overlap regions may be driven by extrapolation rather than by well-supported causal comparisons. Related subgroup-discovery methods learn interpretable or high-effect subgroups using rule sets or constrained partitioning (Wang & Rudin, 2022; Cai et al., 2022; Yang et al., 2026), but do not characterize the fragmentation, threshold persistence, or overlap-supported topology of an unrestricted threshold policy. Recent work on stable policy learning studies sample-to-sample stability of the learned policy's welfare (Barnhard et al., 2026), which is distinct from the persistence of beneficiary-region geometry studied here.

We tackle these questions by applying tools from topological data analysis (TDA) to the geometry of treatment-benefit regions (Chazal et al., 2016). Recent causal work has used persistent homology to define effects for topological outcomes (Kim & Lee, 2026); here, outcomes remain scalar and topology is applied instead to CATE superlevel sets as a policy-deployment object. Zeroth Betti numbers quantify fragmentation, persistence identifies threshold-stable segments, and overlap trimming restricts attention to causally supported regions. We combine these summaries in a post-estimation deployability layer that, when the surviving segments exceed the segment budget, returns a conservative policy targeting those with the largest estimated gains.

Our contribution. Our contributions are threefold. First, we introduce post-estimation diagnostics of treatment-benefit geometry—fragmentation, threshold persistence, and overlap-supported persistence—that apply to any CATE estimator without imposing a prespecified interpretable policy class. Second, we use these diagnostics to construct a segment-budgeted conservative policy and show that persistent fragmentation lower-bounds the segment complexity needed to retain stable, overlap-supported beneficiaries, with unavoidable regret when the segment budget is insufficient. Third, we establish stability-transfer results showing how estimation errors in the CATE, capacity threshold, and propensity score propagate to persistent segmentation summaries.

2 SETUP

Let $D_n = \{Z_i = (X_i, A_i, Y_i)\}_{i=1}^n$ be a sample of n i.i.d. observations, where $X \in \mathbb{R}^d$ is a covariate vector, $A \in \{0,1\}$ is a binary treatment indicator, and $Y \in \mathbb{R}$ is the observed outcome. The potential outcome that would have been observed under $A = a$ is denoted by Y^a. We further define the propensity score and outcome regression functions as

$$e(x) \coloneqq \mathbb{P}(A = 1 \mid X = x), \quad \mu_a(x) \coloneqq \mathbb{E}(Y \mid A = a, X = x), \quad a \in \{0,1\},$$

respectively. Under the standard causal assumptions of (i) *consistency*: $Y = Y^a$ if $A = a$, (ii) *no unmeasured confounding*: $Y^a \perp\!\!\!\perp A \mid X$, and (iii) *positivity*: $0 < e(x) < 1$ almost surely, the CATE $\mathbb{E}(Y^1 - Y^0 \mid X = x)$ is identified as

$$\tau(x) \coloneqq \mu_1(x) - \mu_0(x).$$

Henceforth, we refer to $\tau(x)$ as the CATE, noting that our subsequent results remain valid irrespective of whether the causal assumptions hold.

Notations. We write $\mathcal{X} = \mathrm{supp}(P_X)$ for the covariate support, where P_X is the distribution of X. We use $\mathbb{P}_n(f(Z)) = \frac{1}{n}\sum_{i=1}^n f(Z_i)$ to denote sample averages, $a \lesssim b$ to mean $a \le Cb$ for some finite constant $C > 0$, and $a \asymp b$ to mean $cb \le a \le Cb$ for some finite constants $0 < c \le C$. Finally, $a \vee b$ is used as a shorthand for $\max\{a, b\}$.

3 FRAMEWORK

We consider an unrestricted policy class, under which, for a fixed threshold $t \ge 0$, the *optimal treatment policy* and its corresponding *treatment-benefit region* are given by

$$\pi_t^*(x) \coloneqq \mathbf{1}\{\tau(x) > t\}, \qquad \mathcal{S}_t \coloneqq \mathcal{S}_t(\tau) = \{x \in \mathcal{X} : \tau(x) > t\},$$

respectively. Hence, for any threshold policy, the induced treatment-benefit region is simply the corresponding superlevel set in the covariate support. To further account for capacity constraints that limit treatment to at most a fraction p of the population, we introduce the following quantities: for a fixed capacity $p \in (0, 1)$, we define the *coverage curve* and the *p-coverage threshold* as

$$\mathcal{C}(t) \coloneqq P_X\{\tau(X) > t\}, \qquad t_p \coloneqq \inf\{t : \mathcal{C}(t) \leq p\},$$

respectively. The coverage curve $\mathcal{C}(t)$ reports the proportion of the population that is treated by the threshold policy at each threshold t, while the p-coverage threshold t_p corresponds to the upper p-quantile of $\tau(X)$. When $\tau(X)$ has no point mass at t_p, the optimal policy under the capacity constraint is given by $\pi^*_{t_p}(x) = \mathbf{1}\{\tau(x) > t_p\}$. We implicitly assume that $t_p > 0$, since otherwise the capacity constraint is nonbinding.

While recovering the p-coverage threshold t_p is important for enforcing the capacity constraint, we emphasize that our objective extends beyond estimating the optimal threshold itself. Our primary interest lies in establishing deployability conditions for an estimated capacity-compatible threshold policy based on the *structural simplicity, stability*, and *causal credibility* of its induced treatment-benefit region, and in developing a deployability layer that returns a conservative policy when these criteria are not met. Accordingly, Section 3.1 introduces the main tools for evaluating policy deployability, Section 3.2 develops the deployability layer, and Section 3.3 provides an interpretation for the target policy class underlying our framework.

Remark 1. *The coverage curve and p-coverage threshold are familiar quantities in the resource-constrained treatment selection literature (Bhattacharya & Dupas, 2012; Luedtke & van der Laan, 2016; Kitagawa & Tetenov, 2018; VanderWeele et al., 2019; Imai & Li, 2023), and are used primarily as auxiliary summaries in deriving our main segmentation results rather than as novel estimands of independent interest. In the present framework, coverage serves as the capacity coordinate along which the topology and overlap support of the treatment-benefit region are studied.*

3.1 Causal Segmentation of the Treatment-Benefit Region

Here, we define the causal segmentation summaries that characterize the geometry and overlap support of the treatment-benefit region and discuss their implications for deployable policy design. We note that all geometric objects are defined relative to $\mathcal{X}$, since regions outside the covariate support have no direct policy interpretation.

Segmentation summaries. The *zeroth Betti number* at a fixed threshold t is defined as

$$\beta_0(t) \coloneqq \operatorname{rank} \mathcal{H}_0(\mathcal{S}_t), \tag{1}$$

where $\mathcal{H}_0(\mathcal{S}_t)$ is the zeroth homology group of the superlevel set $\mathcal{S}_t = \{x \in \mathcal{X} : \tau(x) > t\}$, computed with coefficients in $\mathbb{Z}/2\mathbb{Z}$. This quantity counts the number of distinct beneficiary segments in $\mathcal{S}_t$ that a policy must handle separately, thus providing a measure of policy complexity. However, raw zeroth Betti numbers change abruptly when t crosses a topological transition level, and even small perturbations in the estimated treatment effect or threshold can cause the estimated zeroth Betti number to differ substantially from its true value. Therefore, for policy interpretation, we focus on components that remain distinguishable across a band of nearby thresholds by using the persistent zeroth Betti number (Edelsbrunner et al., 2002). Given a persistence level $\lambda \geq 0$, we define the *λ-persistent zeroth Betti number* as

$$\beta_{0,\lambda}(t) \coloneqq \operatorname{rank} \operatorname{im}(\mathcal{H}_0(\mathcal{S}_{t+\lambda}) \to \mathcal{H}_0(\mathcal{S}_t)), \tag{2}$$

where $\operatorname{im}(\mathcal{H}_0(\mathcal{S}_{t+\lambda}) \to \mathcal{H}_0(\mathcal{S}_t))$ is the image of the homomorphism induced by the inclusion $\mathcal{S}_{t+\lambda} \subset \mathcal{S}_t$. This is a multiscale analogue of the zeroth Betti number that counts the distinct connected components of $\mathcal{S}_t$ that remain represented at the stricter threshold $t + \lambda$, i.e., high-benefit segments whose maximum treatment effect exceeds the decision threshold by more than λ. In policy terms, any beneficiary segment counted by the λ-persistent zeroth Betti number persists over a threshold band of width λ, identifying stable segments that are robust to sufficiently small perturbations in the estimated treatment effect and threshold, as formalized in Section 4.2. Accordingly, $\beta_{0,\lambda}(t)$ provides a measure of complexity for the policy that treats only stable beneficiary segments that are not at risk of disappearing under perturbations of size λ.

Overlap-trimmed segmentation. For a fixed $\varepsilon \in (0, 1/2)$, we define the overlap-supported region as $\{x \in \mathcal{X} : \varepsilon < e(x) < 1 - \varepsilon\}$. Treatment effects estimated outside this region may rely

on extrapolation that is poorly supported by the observed data and are therefore less credible. To address this, we apply overlap trimming to exclude areas with insufficient overlap and define the resulting *overlap-trimmed treatment-benefit region* as the following two-parameter superlevel set:

$$\mathcal{S}_{t,\varepsilon} \coloneqq \mathcal{S}_{t,\varepsilon}(\tau, e) = \{x \in \mathcal{X} : \tau(x) > t,\ \min\{e(x), 1 - e(x)\} > \varepsilon\}.$$

After estimating τ and e, comparing the resulting treatment-benefit region with its overlap-trimmed counterpart distinguishes causally credible beneficiary segments supported by adequate overlap from those potentially driven by extrapolation. A rule of thumb for the trimming threshold, as suggested by Crump et al. (2009), is $\varepsilon = 0.1$. Then, given persistence levels $\lambda, \mu \geq 0$, the two-parameter (λ, μ)-persistent zeroth Betti number of the overlap-trimmed treatment-benefit region is defined as

$$\beta_{0,(\lambda,\mu)}(t, \varepsilon) \coloneqq \operatorname{rank} \operatorname{im}(\mathcal{H}_0(\mathcal{S}_{t+\lambda,\varepsilon+\mu}) \to \mathcal{H}_0(\mathcal{S}_{t,\varepsilon})). \tag{3}$$

We introduce this general form for completeness, but our analysis does not consider persistence with respect to the overlap parameter ε. Accordingly, we set $\mu = 0$ and focus on $\beta_{0,(\lambda,0)}(t, \varepsilon)$. This quantity counts the distinct connected components of $\mathcal{S}_{t,\varepsilon}$ that remain represented in $\mathcal{S}_{t+\lambda,\varepsilon}$. Thus, it measures the complexity of the policy that treats only stable, overlap-supported beneficiary segments.

3.2 Deployability Layer

We now present the deployability layer that evaluates an estimated threshold policy using the causal segmentation summaries from the previous section and, when the required criteria are not met, returns a conservative policy targeting beneficiary segments with the largest estimated gains subject to the segment budget. Our framework is applied after CATE estimation and is compatible with any CATE estimator, including causal forests, meta-learners, doubly robust learners, and inverse-variance weighting methods (Wager & Athey, 2018; Künzel et al., 2019; Nie & Wager, 2021; Kennedy, 2023; Fisher, 2024). We first need to set a target capacity $p \in (0, 1)$, a persistence level $\lambda > 0$, a trimming threshold $\varepsilon \in (0, 1/2)$, and a segment budget $m \in \mathbb{N}$. Given a CATE estimator $\hat{\tau}$, we obtain the empirical coverage curve and the estimated p-coverage threshold via a simple plug-in:

$$\widehat{\mathcal{C}}(t) \coloneqq \mathbb{P}_n\big[\mathbf{1}\{\hat{\tau}(X) > t\}\big], \qquad \hat{t}_p \coloneqq \inf\{t : \widehat{\mathcal{C}}(t) \leq p\}.$$

Then, the estimated optimal policy under the capacity constraint and its treatment-benefit region are

$$\hat{\pi}^*_{\hat{t}_p}(x) \coloneqq \mathbf{1}\{\hat{\tau}(x) > \hat{t}_p\}, \qquad \widehat{\mathcal{S}}_{\hat{t}_p} \coloneqq \mathcal{S}_{\hat{t}_p}(\hat{\tau}) = \{x \in \mathcal{X} : \hat{\tau}(x) > \hat{t}_p\}.$$

Given an estimated propensity score $\hat{e}$, the overlap-trimmed analogue is

$$\widehat{\mathcal{S}}_{\hat{t}_p,\varepsilon} \coloneqq \mathcal{S}_{\hat{t}_p,\varepsilon}(\hat{\tau}, \hat{e}) = \{x \in \mathcal{X} : \hat{\tau}(x) > \hat{t}_p,\ \min\{\hat{e}(x), 1 - \hat{e}(x)\} > \varepsilon\}.$$

Diagnostics. The final diagnostic for the estimated optimal policy $\hat{\pi}^*_{\hat{t}_p}(x) = \mathbf{1}\{\hat{\tau}(x) > \hat{t}_p\}$ reports

$$\widehat{\beta}_0(\hat{t}_p), \qquad \widehat{\beta}_{0,\lambda}(\hat{t}_p), \qquad \widehat{\beta}_{0,(\lambda,0)}(\hat{t}_p, \varepsilon),$$

which are the estimated segmentation summaries obtained by replacing the true treatment-benefit region in equation 1, equation 2, and equation 3 with their estimated counterparts defined above. Comparing the three summaries reveals how fragmentation changes after imposing the stability and overlap-support criteria; in particular, overlap trimming may either remove or split connected components. If $\widehat{\mathcal{S}}_{\hat{t}_p,\varepsilon} = \widehat{\mathcal{S}}_{\hat{t}_p}$ and $\widehat{\beta}_0(\hat{t}_p) = \widehat{\beta}_{0,(\lambda,0)}(\hat{t}_p, \varepsilon)$, then every beneficiary segment of the estimated optimal policy is overlap-supported and persists across the threshold band. If, in addition, $\widehat{\beta}_0(\hat{t}_p) \leq m$, where m is the segment budget, then the estimated optimal policy $\hat{\pi}^{(x)}_{\hat{t}_p}$ is deployable. When any of these conditions fail, the deployability layer returns a conservative alternative.

From diagnostics to a deployable policy. Let $\widehat{G}_1, \ldots, \widehat{G}_{\hat{q}}$ denote the distinct connected components of $\widehat{\mathcal{S}}_{\hat{t}_p,\varepsilon}$ that intersect $\widehat{\mathcal{S}}_{\hat{t}_p+\lambda,\varepsilon}$, where $\hat{q} = \widehat{\beta}_{0,(\lambda,0)}(\hat{t}_p, \varepsilon)$, each with estimated policy gain $\widehat{W}_j = \mathbb{P}_n\big[\{\hat{\tau}(X) - \hat{t}_p\}\mathbf{1}\{X \in \widehat{G}_j\}\big]$. Let $\widehat{\mathcal{J}}$ be the index set of the $\min\{\hat{q}, m\}$ components among $\widehat{G}_1, \ldots, \widehat{G}_{\hat{q}}$ with the largest estimated policy gains. Then, the deployability layer returns the modified conservative treatment rule

$$\hat{\pi}_{m,p,\lambda,\varepsilon}(x) = \mathbf{1}\{x \in \cup_{j \in \widehat{\mathcal{J}}} \widehat{G}_j\}.$$

Thus, the returned policy treats only stable, causally credible beneficiaries of the estimated treatment-benefit region $\widehat{\mathcal{S}}_{\hat{t}_p}$, subject to the segment budget. Note that all components counted by $\widehat{\beta}_{0,(\lambda,0)}(\hat{t}_p, \varepsilon)$ are treated when $\hat{q} \leq m$; otherwise, $\hat{q} - m$ components with the smallest estimated policy gains are excluded.

3.3 Policy Interpretation

Policies that treat only stable, overlap-supported segments of the optimal treatment-benefit region can be viewed as *conservative policies*. Beyond deployability, such policies are also natural when overtreatment is undesirable. Thus, our framework may be interpreted as policy learning over a restricted target class $\Pi_{m,p,\lambda,\varepsilon}$ of conservative policies formed by unions of at most m stable, overlap-supported segments within $\mathcal{S}_{t_p,\varepsilon}$. When the number q of such segments exceeds the segment budget m, no policy in $\Pi_{m,p,\lambda,\varepsilon}$ can coincide with the optimal conservative policy that treats all q segments, as at least $q-m$ such segments must be omitted. This, in turn, yields a welfare consequence: the (threshold-adjusted) policy value $W(\pi) = \mathbb{E}\big[\pi(X)\{\tau(X) - t_p\}\big]$ of any policy in $\Pi_{m,p,\lambda,\varepsilon}$ is necessarily smaller than that of the optimal conservative policy. The following proposition formalizes this point by deriving a lower bound on the regret of any $\pi \in \Pi_{m,p,\lambda,\varepsilon}$ relative to the optimal conservative policy when $q > m$.

Proposition 1. *Let $G_1, \ldots, G_q$ denote the distinct connected components of $\mathcal{S}_{t_p,\varepsilon}$ that intersect $\mathcal{S}_{t_p+\lambda,\varepsilon}$, where $q = \beta_{0,(\lambda,0)}(t_p, \varepsilon)$, each with policy gain $W_j = \mathbb{E}\big[\{\tau(X) - t_p\}\,\mathbf{1}\{X \in G_j\}\big]$. Let $\mathcal{J}$ be the index set of the $\min\{q, m\}$ components among $G_1, \ldots, G_q$ with the largest policy gains, and let $\pi^*_{p,\lambda,\varepsilon}$ be the optimal conservative policy that treats all q components. If $q > m$, then every conservative policy $\pi \in \Pi_{m,p,\lambda,\varepsilon}$ satisfies*

$$W(\pi^*_{p,\lambda,\varepsilon}) - W(\pi) \geq \sum_{j \notin \mathcal{J}} W_j.$$

The lower bound is attained by the policy $\pi^*_{m,p,\lambda,\varepsilon} = \mathbf{1}\{x \in \cup_{j\in\mathcal{J}} G_j\}$, which is optimal within the class $\Pi_{m,p,\lambda,\varepsilon}$.

4 Theoretical Development

This section presents the main theoretical results. Central to our theoretical development is the following assumption of uniform consistency for the CATE estimator.

Assumption 1 (Uniform consistency of CATE estimation)**.** *There exists a positive sequence $r_n \to 0$ such that*

$$\|\hat{\tau} - \tau\|_\infty \;=\; O_p(r_n).$$

Then, Section 4.1 establishes how this error propagates, under suitable assumptions, to the estimated coverage curve and p-coverage threshold. Our central contribution, presented in Section 4.2, is the stability transfer from uniform CATE estimation to persistent segmentation summaries. The proofs are provided in Appendix A.

4.1 Coverage Curve and p-coverage threshold estimation

We now show that, under suitable assumptions, the uniform CATE estimation error propagates to the estimated coverage curve and, in turn, to the estimated p-coverage threshold. We first consider the coverage curve $\mathcal{C}(t) = \mathbb{E}[\mathbf{1}\{\tau(X) > t\}]$, whose indicator form makes it sensitive to probability mass near each threshold t. To address this, we impose the following assumption to control the local mass around all relevant thresholds.

Assumption 2 (Margin condition)**.** *Let $\mathcal{T} \subset [0, \infty)$ be a compact interval of thresholds under consideration. Then, there exists a rate $\alpha > 0$ and a bandwidth $\epsilon_0 > 0$ such that, for all $0 \leq \epsilon \leq \epsilon_0$,*

$$\sup_{t\in\mathcal{T}} P_X\big\{|\tau(X) - t| \leq \epsilon\big\} \;\lesssim\; \epsilon^\alpha.$$

Assumption 2 is a standard margin condition, commonly adopted in the treatment selection literature (Qian & Murphy, 2011; Kitagawa & Tetenov, 2018), that prevents the local probability mass of $\tau(X)$

from being concentrated around any treatment decision threshold $t \in \mathcal{T}$. Specifically, it controls the proportion of the population whose treatment effect lies close enough to any relevant decision boundary that an estimation error of size ϵ in $\hat{\tau}$ may alter treatment eligibility, thereby ensuring that small estimation errors in $\hat{\tau}$ do not translate into arbitrarily large errors in the estimated coverage curve. A larger value of α indicates less probability mass near the relevant thresholds, whereas a smaller value of α indicates the opposite. In particular, the strong margin condition with $\alpha = 1$ holds when $\tau(X)$ has a density bounded above on $\mathcal{T}_{\epsilon_0} := \{t' \in \mathbb{R} : \min_{t\in\mathcal{T}} |t' - t| \le \epsilon_0\}$, which is the closed ϵ_0-neighborhood of $\mathcal{T}$. A representative example excluded by Assumption 2 occurs when the distribution of $\tau(X)$ has a point mass at some $t \in \mathcal{T}$; that is, the treatment effect surface contains a flat region $\tau(x) = t$ over a subset of $\mathcal{X}$. Under this assumption, the following theorem establishes uniform consistency of the estimated coverage curve over all $t \in \mathcal{T}$, with a convergence rate inherited from the uniform CATE estimation error.

Theorem 2. *Suppose that Assumptions 1 and 2 hold. Then,*

$$\sup_{t\in\mathcal{T}} \left|\widehat{\mathcal{C}}(t) - \mathcal{C}(t)\right| \;=\; O_p(r_n^{\alpha} \vee n^{-1/2}).$$

Theorem 2 shows that the estimated coverage curve $\widehat{\mathcal{C}}(t) = \mathbb{P}_n\big[\mathbf{1}\{\hat{\tau}(X) > t\}\big]$ is uniformly consistent over all thresholds under consideration. The term $O_p(r_n^{\alpha})$ arises from the uniform estimation error of $\hat{\tau}$, modulated by the anti-concentration exponent α, whereas $O_p(n^{-1/2})$ represents the empirical fluctuation arising from replacing the expectation in $\mathcal{C}$ with the empirical average $\mathbb{P}_n$ in $\widehat{\mathcal{C}}$. In the common case $\alpha = 1$, the rate r_n inherited from nonparametric CATE estimation is generally slower than $n^{-1/2}$, and hence the overall rate is governed by $O_p(r_n)$. Having established the uniform consistency of $\widehat{\mathcal{C}}$, we next derive how this rate transfers to the estimated p-coverage threshold $\hat{t}_p = \inf\{t : \widehat{\mathcal{C}}(t) \le p\}$. We first introduce an additional assumption needed to establish the desired result.

Assumption 3 (Regularity near p-coverage threshold)**.** *For the desired capacity level $p \in (0, 1)$, let $t_p = \inf\{t : \mathcal{C}(t) \le p\}$ be the true p-coverage threshold value such that $\mathcal{C}(t_p) = p$. Then, there exist constants $\eta > 0$ and $0 < c \le C < \infty$ such that $(t_p - \eta, t_p + \eta) \subset \mathcal{T}$ and $\tau(X)$ has a density f_τ on the η-neighborhood of t_p satisfying*

$$c \;\le\; f_\tau(t) \;\le\; C, \quad \textit{for all } t \in (t_p - \eta, t_p + \eta).$$

Assumption 3 essentially states that the coverage curve is strictly decreasing in the η-neighborhood of t_p with slope between $-C$ and $-c$, since $\mathcal{C}'(t) = -f_\tau(t)$ when $\tau(X)$ has a density f_τ. This ensures that the coverage curve is locally invertible around t_p and enables us to leverage the local geometry of $\mathcal{C}$ to characterize the horizontal error between $\hat{t}_p$ and t_p in terms of the vertical error between $\widehat{\mathcal{C}}(t)$ and $\mathcal{C}(t)$. Assumption 3 rules out, for example, settings where the treatment effect surface $\tau(x)$ exhibits only sharp spikes within the range $(t_p-\eta, t_p+\eta)$ over a subset of $\mathcal{X}$; that is, the coverage curve $\mathcal{C}(t)$ is flat around t_p, causing small vertical errors between $\widehat{\mathcal{C}}$ and $\mathcal{C}$ to translate into arbitrarily large horizontal errors between $\hat{t}_p$ and t_p. Under this additional assumption, Theorem 3 shows that the estimated coverage threshold $\hat{t}_p$ inherits the uniform estimation error of $\widehat{\mathcal{C}}$.

Theorem 3. *Let $\hat{t}_p = \inf\{t : \widehat{\mathcal{C}}(t) \le p\}$, and suppose that Assumptions 1, 2, and 3 hold. Then,*

$$|\hat{t}_p - t_p| \;=\; O_p(r_n^{\alpha} \vee n^{-1/2}).$$

Thus, the estimated coverage threshold $\hat{t}_p$ is consistent for t_p, inheriting the uniform estimation error of $\widehat{\mathcal{C}}$. In particular, when $\tau(X)$ has a density bounded above and away from zero over the interval $\mathcal{T}_{\epsilon_0}$, the anti-concentration exponent in Theorems 2 and 3 is given as $\alpha = 1$.

4.2 Stability transfer to persistent segmentation

This section establishes the main theoretical results on the stability transfer from uniform CATE estimation to persistent zeroth Betti number estimators. Before proceeding, we first define a regularity condition for functions $f : \mathcal{X} \to \mathbb{R}$ that induce filtrations of the underlying space $\mathcal{X}$.

Definition 1 (Tameness). *A function $f : \mathcal{X} \to \mathbb{R}$ is tame if it has finitely many homological critical values and the homology groups $\mathcal{H}_k\big(\mathcal{S}_t(f)\big)$ of the superlevel set $\mathcal{S}_t(f) = \{x \in \mathcal{X} : f(x) \geq t\}$ have finite rank for all $k \in \mathbb{Z}$ and $t \in \mathbb{R}$.*

This is a mild condition commonly used in the TDA literature (Cohen-Steiner et al., 2005; Edelsbrunner & Harer, 2010), which, in our setting, ensures that the (persistent) zeroth Betti numbers are well defined and finite at each threshold. Then, the following theorem presents the first stability transfer result for λ-persistent zeroth Betti number estimators in settings where the threshold is fixed and need not be estimated, such as when it is determined by a known treatment cost rather than a capacity constraint.

Theorem 4. *Let $\tau, \hat{\tau} : \mathcal{X} \to \mathbb{R}$ be tame functions and suppose that Assumption 1 holds. Then, for any fixed $\delta \in (0,1)$ and $\lambda > 0$, there exists a corresponding sequence $a_n(\delta) \asymp r_n$ such that, for all sufficiently large n, $\|\hat{\tau} - \tau\|_\infty \leq a_n(\delta)$ with probability at least $1 - \delta$ and $a_n(\delta) < \lambda/2$. Writing $a_n \coloneqq a_n(\delta)$, for any threshold t,*

$$\beta_{0,\lambda+2a_n}(t - a_n) \;\leq\; \widehat{\beta}_{0,\lambda}(t) \;\leq\; \beta_{0,\lambda-2a_n}(t + a_n)$$

for all sufficiently large n, with probability at least $1 - \delta$. Consequently, if the true persistent zeroth Betti number $\beta_{0,\ell}(s)$ is locally constant as a function of (s, ℓ) in a neighborhood of (t, λ), then

$$\widehat{\beta}_{0,\lambda}(t) \;=\; \beta_{0,\lambda}(t)$$

for all sufficiently large n, with probability at least $1 - \delta$.

Theorem 4 shows that the estimated λ-persistent zeroth Betti number $\widehat{\beta}_{0,\lambda}(t)$ is sandwiched between $\beta_{0,\lambda+2a_n}(t - a_n)$ and $\beta_{0,\lambda-2a_n}(t + a_n)$. Since a_n is a positive sequence converging to zero at the same asymptotic rate as r_n, the sandwich bound reduces to an equality $\widehat{\beta}_{0,\lambda}(t) = \beta_{0,\lambda}(t)$ as $n \to \infty$ if the map $(s, \ell) \mapsto \beta_{0,\ell}(s)$ is locally constant in a neighborhood of (t, λ). Note that Assumptions 2 and 3 are not required here; the persistent segmentation summary estimator inherits the uniform CATE estimation rate specified in Assumption 1, requiring only the mild tameness condition on τ and $\hat{\tau}$. The next corollary extends this result to the case of an estimated threshold, which is particularly relevant when the threshold is chosen to satisfy a capacity constraint.

Corollary 5. *Under the setting of Theorem 4, suppose further that Assumptions 2 and 3 hold with $\alpha = 1$. Then, for any fixed $\delta \in (0,1)$ and $\lambda > 0$, there exists a corresponding sequence $b_n(\delta) \asymp (r_n \vee n^{-1/2})$ such that, for all sufficiently large n, $|\hat{t}_p - t_p| \leq b_n(\delta)$ with probability at least $1 - \delta$ and $a_n + b_n(\delta) < \lambda/2$. Writing $b_n \coloneqq b_n(\delta)$, for the p-coverage threshold t_p and its estimate $\hat{t}_p$,*

$$\beta_{0,\lambda+2(a_n+b_n)}(t_p - a_n - b_n) \;\leq\; \widehat{\beta}_{0,\lambda}(\hat{t}_p) \;\leq\; \beta_{0,\lambda-2(a_n+b_n)}(t_p + a_n + b_n)$$

for all sufficiently large n, with probability at least $1 - \delta$. Consequently, if the true persistent zeroth Betti number $\beta_{0,\ell}(s)$ is locally constant as a function of (s, ℓ) in a neighborhood of (t_p, λ), then

$$\widehat{\beta}_{0,\lambda}(\hat{t}_p) = \beta_{0,\lambda}(t_p)$$

for all sufficiently large n, with probability at least $1 - \delta$.

Compared with Theorem 4, threshold estimation introduces an additional perturbation b_n, enlarging the sandwich width from a_n to $a_n + b_n$, so that the stability transfer is governed jointly by the CATE and threshold estimation errors. The additional Assumptions 2 and 3 with $\alpha = 1$ are required here to ensure consistency of the estimated p-coverage threshold, with rate $b_n \asymp r_n \vee n^{-1/2}$ as established in Theorem 3. In the usual nonparametric setting where $r_n \gtrsim n^{-1/2}$, we have $b_n \asymp r_n$ and, since $a_n \asymp r_n$, it follows that $a_n + b_n \asymp a_n$. Hence, asymptotically, the additional threshold estimation does not alter the rate at which the sandwich stability bound contracts, and we recover the equality $\widehat{\beta}_{0,\lambda}(\hat{t}_p) = \beta_{0,\lambda}(t_p)$ if the map $(s, \ell) \mapsto \beta_{0,\ell}(s)$ is locally constant in a neighborhood of (t_p, λ).

Before extending this result to the overlap-trimmed persistent zeroth Betti number, we introduce one additional assumption.

Assumption 4 (Local overlap topology stability). *There exists a $\mu > 0$ such that, for all relevant t and $r \in [0, \mu]$, the homomorphisms induced by the inclusions $\mathcal{S}_{t,\varepsilon+r} \subset \mathcal{S}_{t,\varepsilon}$ and $\widehat{\mathcal{S}}_{t,\varepsilon+r} \subset \widehat{\mathcal{S}}_{t,\varepsilon}$ are isomorphisms on $\mathcal{H}_0$.*

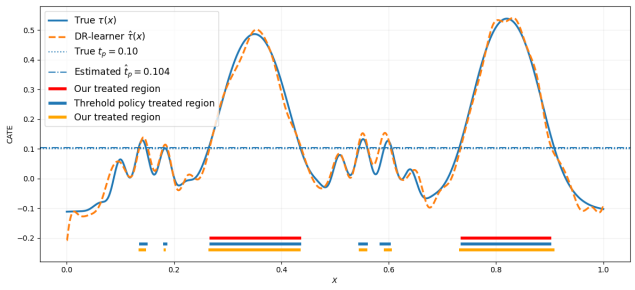


Figure 1: Estimated CATE and treatment regions.

Table 1: Policy quality summaries.

	Threshold	Ours
Policy value	0.090	0.089
Capacity	0.40	0.342
Segments	5.8	2
Jaccard	0.970	0.985

This condition ensures that small increases in the overlap trimming threshold ε do not alter the correspondence of connected components of the overlap-trimmed treatment-benefit region, thereby ruling out additional topological variation along the overlap direction. Under this assumption, the following theorem establishes the stability transfer result for overlap-trimmed persistent zeroth Betti numbers.

Theorem 6. *Assume the setting of Corollary 5. Suppose further that, for any fixed $\delta \in (0,1)$ and $\mu > 0$, there exists a corresponding positive sequence $c_n(\delta) \to 0$ such that, for all sufficiently large n, $\|\hat{e} - e\|_\infty \le c_n(\delta)$ with probability at least $1-\delta$ and $c_n(\delta) < \mu/2$. Writing $c_n := c_n(\delta)$,*

$$\begin{aligned}\beta_{0,(\lambda+2(a_n+b_n),\,\mu+2c_n)}(t_p - a_n - b_n, \varepsilon - c_n) \;&\le\; \widehat{\beta}_{0,(\lambda,\mu)}(\hat{t}_p, \varepsilon) \\ &\le\; \beta_{0,(\lambda-2(a_n+b_n),\,\mu-2c_n)}(t_p + a_n + b_n, \varepsilon + c_n)\end{aligned}$$

for all sufficiently large n, with probability at least $1-\delta$. Consequently, if the true persistent zeroth Betti number $\beta_{0,(\ell,\ell')}(s,\varepsilon)$ is locally constant as a function of (s,ℓ,ℓ') in a neighborhood of (t_p,λ,μ) and, in addition, Assumption 4 holds, then

$$\widehat{\beta}_{0,(\lambda,0)}(\hat{t}_p, \varepsilon) \;=\; \beta_{0,(\lambda,0)}(t_p, \varepsilon)$$

for all sufficiently large n, with probability at least $1-\delta$.

We note that, even under the local constancy of the true persistent zeroth Betti number, one can generally establish only $\widehat{\beta}_{0,(\lambda,\mu)}(\hat{t}_p) = \beta_{0,(\lambda,\mu)}(t_p)$ from the sandwich stability bound. Recovering the stronger identity $\widehat{\beta}_{0,(\lambda,0)}(\hat{t}_p) = \beta_{0,(\lambda,0)}(t_p)$ additionally relies on Assumption 4.

5 Empirical Evaluation

We now conduct empirical experiments to illustrate that our proposed method captures deployment failure modes that are invisible to existing treatment decision and policy learning methods. First, we specify deterministic models for $e(x) = \mathbb{P}(A|X=x)$ and $\mu_a(x) = \mathbb{E}(Y|A=a, X=x)$, and subsequently generate a synthetic dataset $\{X_i, A_i, Y_i\}_{i=1}^n$ with $n = 2000$, where the covariate $X \in \mathbb{R}$ is sampled from $\mathrm{Unif}(0,1)$. The treatment $A_i \in \{0,1\}$ is sampled from a binomial distribution with probabiility $e(X_i)$. Given (X_i, A_i), the corresponding outcome $Y_i \in \mathbb{R}$ is generated by adding Gaussian noise with standard deviation 0.05 to the output of $\mu_{A_i}(X_i)$. The true CATE is set as $\tau(x) = \mu_1(x) - \mu_0(x)$, and we specify the capacity level $p = 0.4$. The true threshold satisfying the capacity constraint is $t_p = 0.1$. We use the DR-Learner of Kennedy (2023) to estimate the CATE, where we use cubic spline to estimate the nuisance functions, with 50 knots for the regression model and 10 knots for the propensity score model. The second stage regression also uses cubic spline with 50 knots. Details of the data generating process are provided in Appendix B.

We run the experiment 20 times and evaluate the policy based on the following metrics: average policy value, average coverage, number of treated segments, and stability of the treatment-benefit region. Stability of the treatment-benefit region is measured by the average of pairwise Jaccard similarity of the treated sets across multiple runs.

Figure 1 illustrates the estimated CATE function together with the treatment regions induced by the competing policies. The DR-learner estimate closely tracks the true CATE over most of the

covariate domain, while exhibiting small local fluctuations in regions where the treatment effect is near the decision threshold. As a consequence, the standard threshold policy produces several short and fragmented treated regions in addition to the two dominant high-benefit regions. In contrast, the proposed policy retains only the two large, stable treatment-benefit regions, thereby yielding a substantially simpler treatment rule while preserving the regions associated with the largest treatment effects. This highlights how the proposed procedure removes small, potentially unstable components created by estimation noise and focuses on treatment of persistent beneficiary segments, thereby providing a foundation for constructing more deployable policies.

The proposed evaluation metrics for the standard threshold policy and are proposed method are summarized in Table 1. Both approaches achieve nearly identical policy values, with values of 0.090 and 0.089 for the threshold policy and our method, respectively, indicating that the proposed restriction entails only a negligible loss in expected utility. At the same time, our policy treats a smaller proportion of the population, with a capacity of 0.342 compared with 0.40 for the threshold policy. While both policies satisfy the capacity constraint on average, this indicates that some treatment rules for the threshold policy have slightly exceeded the designated coverage $p = 0.4$ over the 20 runs. More notably, our proposed substantially reduces the number of treated segments, from an average of 5.8 under the threshold policy to only 2 under our method, yielding a considerably simpler treatment rule. It also exhibits slightly greater stability by eliminating the unstable beneficiary segments, as reflected by the higher Jaccard similarity (0.985 versus 0.970).

6 Discussion

This proposal reframes heterogeneous treatment effect analysis around the structure of treatment-benefit regions. Existing treatment choice and policy learning methods already address important questions about whom to treat under capacity or policy class constraints. The present framework asks a different but complementary question: whether the induced beneficiary population is structurally simple, stable, and causally credible enough to support deployment.

Coverage determines where a policy lies along a familiar capacity axis, but coverage alone is not the contribution. The contribution is to analyze what the benefit region looks like at that capacity level. Fragmentation determines how many distinct beneficiary segments must be targeted. Persistence distinguishes stable beneficiary groups from threshold sensitive artifacts. Overlap trimming identifies which segments remain supported by adequate causal comparisons.

The resulting framework is intentionally focused on degree zero topology. While higher order features such as loops or cavities may be meaningful in special geometric settings, connected components have the clearest causal decision interpretation. They directly quantify the number of distinct beneficiary groups that a policy must be able to target.

The theoretical development clarifies the roles of effect estimation, anti concentration, and persistence. A standard doubly robust procedure may be used to estimate the conditional treatment effect surface. Threshold anti concentration controls coverage and capacity threshold estimation as auxiliary inputs. Persistence provides stability for the connected beneficiary segments themselves. Together, these ingredients yield a transfer principle from uniform conditional effect estimation to stable causal segmentation.

AI USE STATEMENT

(This section is **required** and does not count toward the page limit.)

In this work, we used generative AI tools for generating synthetic data for experiments and providing critical ingredients for proving mathematical claims. We have not used generative AI tools for developing theoretical models or conceptual frameworks and other tasks, such as proposing hypothesis or designing experiments, are not applicable to this work. Additionally, we used generative AI tools for editing our paper to improve readability. We have reviewed all AI-assisted work. We have manually completed all sanity checks to ensure that the synthetic data generated by AI is sound, and that mathematical claims are valid. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REFERENCES

Susan Athey and Stefan Wager. Policy learning with observational data. *Econometrica*, 89(1): 133–161, 2021. doi: 10.3982/ECTA15732.

Harvey Barnhard, Giacomo Opocher, and Rahul Singh. Stable policy learning, 2026.

Debopam Bhattacharya and Pascaline Dupas. Inferring welfare maximizing treatment assignment under budget constraints. *Journal of Econometrics*, 167(1):168–196, 2012. doi: 10.1016/j.jeconom.2011.11.007.

Hengrui Cai, Wenbin Lu, Rachel Marceau West, Devan V. Mehrotra, and Lingkang Huang. CAPITAL: Optimal subgroup identification via constrained policy tree search. *Statistics in Medicine*, 41(21):4227–4244, 2022. doi: 10.1002/sim.9507.

Frédéric Chazal, Vin De Silva, Marc Glisse, and Steve Oudot. *The structure and stability of persistence modules*, volume 10. Springer, 2016.

David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. In *Proceedings of the Twenty-First Annual Symposium on Computational Geometry*, SCG '05, pp. 263–271, New York, NY, USA, 2005. Association for Computing Machinery. doi: 10.1145/1064092.1064133.

Richard K. Crump, V. Joseph Hotz, Guido W. Imbens, and Oscar A. Mitnik. Dealing with limited overlap in estimation of average treatment effects. *Biometrika*, 96(1):187–199, 03 2009. doi: 10.1093/biomet/asn055.

Herbert Edelsbrunner and John Harer. *Computational Topology: An Introduction*. American Mathematical Society, 2010. doi: 10.1090/mbk/069.

Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. Topological persistence and simplification. *Discrete Computational Geometry*, 28(4):511–533, 2002. doi: 10.1007/s00454-002-2885-2.

Aaron Fisher. Inverse-variance weighting for estimation of heterogeneous treatment effects. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 13706–13733. PMLR, 21-27 Jul 2024.

Keisuke Hirano and Jack R. Porter. Asymptotics for statistical treatment rules. *Econometrica*, 77(5):1683–1701, 2009. doi: 10.3982/ECTA6630.

Kosuke Imai and Michael Lingzhi Li. Experimental evaluation of individualized treatment rules. *Journal of the American Statistical Association*, 118(541):242–256, 2023. doi: 10.1080/01621459.2021.1923511.

Edward H. Kennedy. Towards optimal doubly robust estimation of heterogeneous causal effects. *Electronic Journal of Statistics*, 17(2):3008–3049, 2023. doi: 10.1214/23-EJS2157.

Kwangho Kim and Hajin Lee. Topological causal effects. In *International Conference on Learning Representations*, 2026.

Toru Kitagawa and Aleksey Tetenov. Who should be treated? empirical welfare maximization methods for treatment choice. *Econometrica*, 86(2):591–616, 2018. doi: 10.3982/ECTA13288.

Sören R. Künzel, Jasjeet S. Sekhon, Peter J. Bickel, and Bin Yu. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences*, 116(10):4156–4165, 2019. doi: 10.1073/pnas.1804597116.

Himabindu Lakkaraju and Cynthia Rudin. Learning cost-effective and interpretable treatment regimes. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pp. 166–175. PMLR, 2017.

Alexander R. Luedtke and Mark J. van der Laan. Optimal individualized treatments in resource-limited settings. *The International Journal of Biostatistics*, 12(1):283–303, 2016. doi: 10.1515/ijb-2015-0007.

Xinkun Nie and Stefan Wager. Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika*, 108(2):299–319, 06 2021. doi: 10.1093/biomet/asaa076.

Min Qian and Susan A. Murphy. Performance guarantees for individualized treatment rules. *The Annals of Statistics*, 39(2):1180–1210, 2011. doi: 10.1214/10-AOS864.

Daniel Tschernutter, Tobias Hatt, and Stefan Feuerriegel. Interpretable off-policy learning via hyperbox search. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 21795–21827. PMLR, 2022.

Tyler J. VanderWeele, Alex R. Luedtke, Mark J. van der Laan, and Ronald C. Kessler. Selecting optimal subgroups for treatment using many covariates. *Epidemiology*, 30(3):334–341, 2019. doi: 10.1097/EDE.0000000000000991.

Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242, 2018. doi: 10.1080/01621459.2017.1319839.

Tong Wang and Cynthia Rudin. Causal rule sets for identifying subgroups with enhanced treatment effects. *INFORMS Journal on Computing*, 34(3):1626–1643, 2022. doi: 10.1287/ijoc.2021.1143.

Lincen Yang, Zhong Li, Matthijs van Leeuwen, and Saber Salehkaleybar. Learning subgroups with maximum treatment effects without causal heuristics. *Proceedings of the AAAI Conference on Artificial Intelligence*, 40(32):27565–27573, 2026. doi: 10.1609/aaai.v40i32.39976.

Yichi Zhang, Eric B. Laber, Anastasios Tsiatis, and Marie Davidian. Using decision lists to construct interpretable and parsimonious treatment regimes. *Biometrics*, 71(4):895–904, 2015. doi: 10.1111/biom.12354.

Zhengyuan Zhou, Susan Athey, and Stefan Wager. Offline multi-action policy learning: Generalization and optimization. *Operations Research*, 71(1):148–183, 2023. doi: 10.1287/opre.2022.2271.

A PROOFS

Before presenting the main proofs, we establish two supporting lemmas used throughout this section.

Lemma 1. *Suppose that $\|\hat{\tau} - \tau\|_\infty \le \epsilon$ for some $\epsilon > 0$. Then, for every $x \in \mathcal{X}$ and t,*

$$\big|\mathbf{1}\{\hat{\tau}(x) > t\} - \mathbf{1}\{\tau(x) > t\}\big| \;\le\; \mathbf{1}\{|\tau(x) - t| \le \epsilon\}.$$

Proof of Lemma 1. First, note that both the LHS and the RHS can only take values zero or one. The LHS equals one only when the two indicators disagree, in which case t lies between $\hat{\tau}(x)$ and $\tau(x)$. It then follows that $|\tau(x) - t| \le |\hat{\tau}(x) - \tau(x)| \le \epsilon$, and hence the RHS is always one when the LHS takes value one, which yields the desired inequality. □

Lemma 2. *Let $F_\tau(t) = P_X\{\tau(X) \le t\}$ and $\widehat{F}_\tau(t) = \mathbb{P}_n\big[\mathbf{1}\{\tau(X) \le t\}\big]$ be the cumulative distribution function (cdf) and the empirical distribution function of $\tau(X)$, respectively. Furthermore, define $F_\tau(t^-) \coloneqq P_X\{\tau(X) < t\}$ and $\widehat{F}_\tau(t^-) = \mathbb{P}_n\big[\mathbf{1}\{\tau(X) < t\}\big]$. Then,*

$$\sup_{t\in\mathbb{R}} \big|\widehat{F}_\tau(t^-) - F_\tau(t^-)\big| \;=\; \sup_{t\in\mathbb{R}} \big|\widehat{F}_\tau(t) - F_\tau(t)\big|.$$

A.1 PROOF OF RESULT IN SECTION 3

Proof of Proposition 1. Since the optimal conservative policy $\pi^*_{p,\lambda,\varepsilon}$ treats the whole region $\mathcal{S}_{t_p,\varepsilon}$,

$$W(\pi^*_{p,\lambda,\varepsilon}) = \sum_{j=1}^{q} W_j.$$

whereas, for $\pi \in \Pi_{m,p,\lambda,\varepsilon}$,

$$W(\pi) = \sum_{j\in\mathcal{J}(\pi)} W_j,$$

where $\mathcal{J}(\pi)$ denotes the index set of the components among $G_1, \ldots, G_q$ that are treated by π. Hence

$$W(\pi^*_{p,\lambda,\varepsilon}) - W(\pi) = \sum_{j\notin\mathcal{J}(\pi)} W_j.$$

Since at least $q-m$ connected components of $_{t_p}$ must be omitted, this sum is minimized by omitting the $q - m$ components having the smallest gains. Therefore, for any $\pi \in \Pi_{m,p,\lambda,\varepsilon}$,

$$W(\pi^*_{p,\lambda,\varepsilon}) - W(\pi) \ge \sum_{j\notin\mathcal{J}} W_j.$$

□

A.2 PROOF OF RESULTS IN SECTION 4.1

Proof of Theorem 2. We now present the main proof. Suppose that we work on the event $\|\hat{\tau} - \tau\|_\infty \le \epsilon$ for some $\epsilon > 0$, and consider the following decomposition:

$$\begin{aligned}
\big|\widehat{C}(t) - C(t)\big| \;&=\; \Big|\mathbb{P}_n\big[\mathbf{1}\{\hat{\tau}(X) \ge t\}\big] - \mathbb{P}\{\tau(X) \ge t\}\Big| \\
&\le\; \Big|\mathbb{P}_n\big[\mathbf{1}\{\tau(X) \ge t\}\big] - \mathbb{P}\{\tau(X) \ge t\}\Big| + \Big|\mathbb{P}_n\big[\mathbf{1}\{\hat{\tau}(X) \ge t\}\big] - \mathbb{P}_n\big[\mathbf{1}\{\tau(X) \ge t\}\big]\Big| \\
&\le\; \Big|\mathbb{P}_n\big[\mathbf{1}\{\tau(X) \ge t\}\big] - \mathbb{P}\{\tau(X) \ge t\}\Big| + \mathbb{P}_n\big[\,\big|\mathbf{1}\{\hat{\tau}(X) \ge t\} - \mathbf{1}\{\tau(X) \ge t\}\big|\,\big] \\
&\le\; \Big|\mathbb{P}_n\big[\mathbf{1}\{\tau(X) \ge t\}\big] - \mathbb{P}\{\tau(X) \ge t\}\Big| + \mathbb{P}_n\big[\mathbf{1}\{|\tau(X) - t| \le \epsilon\}\big] \\
&\le\; \Big|\mathbb{P}_n\big[\mathbf{1}\{\tau(X) \ge t\}\big] - \mathbb{P}\{\tau(X) \ge t\}\Big| \\
&\quad + \Big|\mathbb{P}_n\big[\mathbf{1}\{|\tau(X) - t| \le \epsilon\}\big] - \mathbb{P}\{|\tau(X) - t| \le \epsilon\}\Big| \\
&\quad + \mathbb{P}\{|\tau(X) - t| \le \epsilon\}.
\end{aligned}$$

The first and second inequalities follow from the triangle inequality, and the third inequality follows from Lemma 1. By the subadditivity of the supremum operator, it further holds that

$$\begin{aligned}\sup_{t\in\mathcal{T}}\left|\widehat{C}(t)-C(t)\right| \;\le\; &\underbrace{\sup_{t\in\mathcal{T}}\left|\mathbb{P}_n\big[\mathbf{1}\{\tau(X)\ge t\}\big]-\mathbb{P}\{\tau(X)\ge t\}\right|}_{\textbf{(i)}}\\ &+\underbrace{\sup_{t\in\mathcal{T}}\left|\mathbb{P}_n\big[\mathbf{1}\{|\tau(X)-t|\le\epsilon\}\big]-\mathbb{P}\{|\tau(X)-t|\le\epsilon\}\right|}_{\textbf{(ii)}}\\ &+\underbrace{\sup_{t\in\mathcal{T}}\mathbb{P}\{|\tau(X)-t|\le\epsilon\}}_{\textbf{(iii)}}.\end{aligned}$$

Thus, it suffices to show that the terms (i), (ii), and (iii) correspond to $O_p(n^{-1/2})$, $O_p(n^{-1/2})$, and $O_p(r_n^{\alpha})$, respectively.

Part (i). By Lemma 2, we have

$$\begin{aligned}\textbf{(i)} &= \sup_{t\in\mathcal{T}}\left|\big(1-\mathbb{P}_n\big[\mathbf{1}\{\tau(X)<t\}\big]\big)-\big(1-\mathbb{P}\{\tau(X)<t\}\big)\right|\\ &= \sup_{t\in\mathcal{T}}\left|\widehat{F}_\tau(t^-)-F_\tau(t^-)\right|\\ &= \sup_{t\in\mathcal{T}}\left|\widehat{F}_\tau(t)-F_\tau(t)\right|.\end{aligned}$$

Then, by the Dvoretzky-Kiefer-Wolfowitz inequality, for every $\varepsilon>0$,

$$\mathbb{P}\Big(\sup_{t\in\mathcal{T}}\left|\widehat{F}_\tau(t)-F_\tau(t)\right|>\varepsilon\Big)\;\le\;2e^{-2n\varepsilon^2}.$$

Setting $\lambda=2e^{-2n\varepsilon^2}$ yields

$$\mathbb{P}\Big(\sqrt{n}\sup_{t\in\mathcal{T}}\left|\widehat{F}_\tau(t)-F_\tau(t)\right|>M_\lambda\Big)\;\le\;\lambda,$$

where $M_\lambda=\sqrt{\frac{\ln(2/\lambda)}{2}}$. This is precisely the definition of $O_p(n^{-1/2})$, and thus,

$$\textbf{(i)}\;=\;O_p(n^{-1/2}).$$

Part (ii). We can further decompose each component of (ii) as

$$\mathbb{P}_n\big[\mathbf{1}\{|\tau(X)-t|\le\epsilon\}\big]\;=\;\mathbb{P}_n\big[\mathbf{1}\{\tau(X)\le t+\epsilon\}\big]-\mathbb{P}_n\big[\mathbf{1}\{\tau(X)<t-\epsilon\}\big],$$

and

$$\mathbb{P}\{|\tau(X)-t|\le\epsilon\}\;=\;\mathbb{P}\{\tau(X)\le t+\epsilon\}-\mathbb{P}\{\tau(X)<t-\epsilon\}.$$

Then, by the triangle inequality and the subadditivity of the supremum operator,

$$\begin{aligned}\textbf{(ii)} &\le \sup_{t\in\mathcal{T}}\left|\mathbb{P}_n\big[\mathbf{1}\{\tau(X)\le t+\epsilon\}\big]-\mathbb{P}\{\tau(X)\le t+\epsilon\}\right|+\sup_{t\in\mathcal{T}}\left|\mathbb{P}_n\big[\mathbf{1}\{\tau(X)<t-\epsilon\}\big]-\mathbb{P}\{\tau(X)<t-\epsilon\}\right|\\ &= \sup_{t\in\mathcal{T}}\left|\widehat{F}_\tau(t+\epsilon)-F_\tau(t+\epsilon)\right|+\sup_{t\in\mathcal{T}}\left|\widehat{F}_\tau((t-\epsilon)^-)-F_\tau((t-\epsilon)^-)\right|\\ &= \sup_{t\in\mathcal{T}}\left|\widehat{F}_\tau(t+\epsilon)-F_\tau(t+\epsilon)\right|+\sup_{t\in\mathcal{T}}\left|\widehat{F}_\tau(t-\epsilon)-F_\tau(t-\epsilon)\right|,\end{aligned}$$

where the last equality follows from Lemma 2. Applying the same argument used for (i), we obtain

$$\begin{aligned}\textbf{(ii)} &\le \sup_{t\in\mathcal{T}}\left|\widehat{F}_\tau(t+\epsilon)-F_\tau(t+\epsilon)\right|+\sup_{t\in\mathcal{T}}\left|\widehat{F}_\tau(t-\epsilon)-F_\tau(t-\epsilon)\right|\\ &= O_p(n^{-1/2})+O_p(n^{-1/2})=O_p(n^{-1/2}),\end{aligned}$$

which implies that

$$\textbf{(ii)}\;=\;O_p(n^{-1/2}).$$

Part (iii). Recall that we are operating on the event $\|\hat{\tau} - \tau\|_\infty \leq \epsilon$. By Assumption 1, for any $\delta \in (0,1)$, there exists a finite $M_\delta > 0$ such that, for all n large enough,

$$\mathbb{P}\big(\|\hat{\tau} - \tau\|_\infty > M_\delta\, r_n\big) \,<\, \delta.$$

Setting $\epsilon = M_\delta\, r_n$, the event $\|\hat{\tau} - \tau\|_\infty \leq M_\delta\, r_n$ occurs with probability at least $1-\delta$ for all sufficiently large n. Moreover, since $r_n \to 0$, it holds that $M_\delta\, r_n \leq \epsilon_0$ for all sufficiently large n, where ϵ_0 is the bandwidth in Assumption 2. Thus, by Assumption 2, for some finite constant $C > 0$,

$$\textbf{(iii)} \,=\, \sup_{t\in\mathcal{T}} \mathbb{P}\{|\tau(X) - t| \leq M_\delta\, r_n\} \,\leq\, C\, M_\delta^\alpha\, r_n^\alpha,$$

with probability at least $1-\delta$ for all sufficiently large n. Equivalently, we have

$$\mathbb{P}\left(\frac{\textbf{(iii)}}{r_n^\alpha} > C\, M_\delta^\alpha\right) \,<\, \delta,$$

for all sufficiently large n, which is the precise definition of $O_p(r_n^\alpha)$. Thus,

$$\textbf{(iii)} \,=\, O_p(r_n^\alpha).$$

Combining the results above completes the proof:

$$\sup_{t\in\mathcal{T}} \big|\widehat{\mathcal{C}}(t) - \mathcal{C}(t)\big| \,=\, O_p(r_n^\alpha) + O_p(n^{-1/2}).$$

□

Proof of Theorem 3. For all $t \in \mathcal{T}$,

$$\mathcal{C}(t) \,=\, \mathbb{P}\{\tau(X) > t\} \,=\, 1 - \mathbb{P}\{\tau(X) \leq t\} \,=\, 1 - F_\tau(t).$$

Combined with Assumption 3, for all $t \in (t_p - \eta, t_p + \eta)$,

$$\mathcal{C}'(t) \,=\, -f_\tau(t) \quad\Rightarrow\quad -\infty \,<\, -C \,\leq\, \mathcal{C}'(t) \,\leq\, -c \,<\, 0.$$

Thus, by the Mean Value Theorem, for any $0 < h \leq \eta$, there exists some $a \in (t_p - h, t_p)$ such that

$$\mathcal{C}(t_p) - \mathcal{C}(t_p - h) \,=\, \mathcal{C}'(a)h \,\leq\, -ch \quad\Rightarrow\quad \mathcal{C}(t_p - h) \,\geq\, \mathcal{C}(t_p) + ch \,=\, p + ch,$$

and there exists some $b \in (t_p, t_p + h)$ such that

$$\mathcal{C}(t_p + h) - \mathcal{C}(t_p) \,=\, \mathcal{C}'(b)h \,\leq\, -ch \quad\Rightarrow\quad \mathcal{C}(t_p + h) \,\leq\, \mathcal{C}(t_p) - ch \,=\, p - ch.$$

Now, for some $\delta > 0$, on the event that

$$\sup_{t\in[t_p-\eta, t_p+\eta]} \big|\widehat{\mathcal{C}}(t) - \mathcal{C}(t)\big| \,\leq\, \delta,$$

we have, for any $0 < h \leq \eta$,

$$\widehat{\mathcal{C}}(t_p - h) \,\geq\, \mathcal{C}(t_p - h) - \delta \,\geq\, p + ch - \delta$$

and

$$\widehat{\mathcal{C}}(t_p + h) \,\leq\, \mathcal{C}(t_p + h) + \delta \,\leq\, p - ch + \delta.$$

If δ is sufficiently small such that $\delta \leq ch$, then

$$\widehat{\mathcal{C}}(t_p + h) \,\leq\, p \,\leq\, \widehat{\mathcal{C}}(t_p - h),$$

which implies that $\hat{t}_p = \inf\{t : \widehat{\mathcal{C}}(t) \leq p\}$ must exist within the range $[t_p - h, t_p + h]$, i.e.,

$$|\hat{t}_p - t_p| \,\leq\, h.$$

By setting $h = 2\delta/c$ to satisfy $\delta \leq ch$ and letting $\delta = \sup_{t\in\mathcal{T}} \big|\widehat{\mathcal{C}}(t) - \mathcal{C}(t)\big|$, we obtain

$$|\hat{t}_p - t_p| \,\leq\, \frac{2}{c} \sup_{t\in\mathcal{T}} \big|\widehat{\mathcal{C}}(t) - \mathcal{C}(t)\big|,$$

for all sufficiently large n, and also $0 < h \leq \eta$ is satisfied. By Theorem 2, in a neighborhood of t_p,

$$\sup_{t\in\mathcal{T}} \big|\widehat{\mathcal{C}}(t) - \mathcal{C}(t)\big| \,=\, O_p(r_n^\alpha) + O_p(n^{-1/2}),$$

and thus,

$$|\hat{t}_p - t_p| \,=\, O_p(r_n^\alpha) + O_p(n^{-1/2}).$$

□

Proof of Theorem 4. We first show that there exists a sequence $a_n \asymp r_n$ satisfying the conditions stated in the theorem. By Assumption 1, for any fixed $\delta \in (0,1)$, there exists a finite constant $M_\delta > 0$ and a natural number N_δ such that

$$\mathbb{P}\big(\|\hat{\tau} - \tau\|_\infty \le M_\delta\, r_n\big) \ge 1 - \delta, \qquad \forall n \ge N_\delta.$$

Thus, we may set $a_n = M_\delta\, r_n$, where $a_n \to 0^+$ since $r_n \to 0^+$. Then, for any fixed $\lambda > 0$, there also exists a natural number N_λ such that $a_n < \lambda/2$ for all $n \ge N_\lambda$. Hence, for all sufficiently large n satisfying $n \ge \max\{N_\delta, N_\lambda\}$, it holds that $\|\hat{\tau} - \tau\|_\infty \le a_n$ with probability at least $1 - \delta$ and $a_n < \lambda/2$.

We now proceed to the main proof. Since we will be working on the event $\|\hat{\tau} - \tau\|_\infty \le a_n$, the subsequent results hold for all sufficiently large n, with probability at least $1 - \delta$.

Lower bound. We first derive the lower bound

$$\beta_{0,\lambda+2a_n}(t - a_n) \le \widehat{\beta}_{0,\lambda}(t)$$

of the sandwich inequality. On the event $\|\hat{\tau} - \tau\|_\infty \le a_n$, we have the inclusion

$$\mathcal{S}_{t+\lambda+a_n} \subseteq \widehat{\mathcal{S}}_{t+\lambda} \subseteq \widehat{\mathcal{S}}_t \subseteq \mathcal{S}_{t-a_n},$$

which induces the following homomorphisms f, g, h:

$$\mathcal{H}_0(\mathcal{S}_{t+\lambda+a_n}) \xrightarrow{f} \mathcal{H}_0(\widehat{\mathcal{S}}_{t+\lambda}) \xrightarrow{g} \mathcal{H}_0(\widehat{\mathcal{S}}_t) \xrightarrow{h} \mathcal{H}_0(\mathcal{S}_{t-a_n}).$$

Note that f, g, h are linear maps between vector spaces since homology is applied with coefficients in the field $\mathbb{Z}/2\mathbb{Z}$, and thus,

$$\operatorname{rank}\,\operatorname{im}(h \circ g \circ f) \le \operatorname{rank}\,\operatorname{im}(g).$$

Then, by the definition of persistent zeroth Betti numbers in equation 2,

$$\beta_{0,\lambda+2a_n}(t - a_n) = \operatorname{rank}\,\operatorname{im}(h \circ g \circ f) \le \operatorname{rank}\,\operatorname{im}(g) = \widehat{\beta}_{0,\lambda}(t).$$

Upper bound. For $a_n < \lambda/2$, an analogous argument can be made for the upper bound

$$\widehat{\beta}_{0,\lambda}(t) \le \beta_{0,\lambda-2a_n}(t + a_n)$$

of the sandwich inequality, by considering, on the event $\|\hat{\tau} - \tau\|_\infty \le a_n$, the inclusion

$$\widehat{\mathcal{S}}_{t+\lambda} \subseteq \mathcal{S}_{t+\lambda-a_n} \subseteq \mathcal{S}_{t+a_n} \subseteq \widehat{\mathcal{S}}_t,$$

which induces the upper bound using the same procedure.

Thus, combining the lower and the upper bound yields the desired result:

$$\beta_{0,\lambda+2a_n}(t - a_n) \le \widehat{\beta}_{0,\lambda}(t) \le \beta_{0,\lambda-2a_n}(t + a_n).$$

Since $a_n \to 0^+$, the equality statement follows immediately if the true persistent zeroth Betti number $\beta_{0,\ell}(s)$ is locally constant as a function of (s, ℓ) in a neighborhood of (t, λ), i.e., it is constant over $s \in [t - a_n, t + a_n]$ and $\ell \in [\lambda - 2a_n, \lambda + 2a_n]$. □

Proof of Corollary 5. The existence of a sequence $b_n \asymp (r_n \vee n^{-1/2})$ satisfying the conditions stated in the corollary can be verified in a similar manner to the proof of Theorem 4, following the result of Theorem 3.

We now proceed to the main proof. Since we will be working on the event $\|\hat{\tau} - \tau\|_\infty \le a_n$ and $|\hat{t}_p - t_p| \le b_n$, the subsequent results hold for all sufficiently large n, with probability at least $1 - \delta$.

Lower bound. We first derive the lower bound

$$\beta_{0,\lambda+2(a_n+b_n)}(t_p - a_n - b_n) \le \widehat{\beta}_{0,\lambda}(\hat{t}_p)$$

of the sandwich inequality. On the event $\|\hat{\tau} - \tau\|_\infty \le a_n$ and $|\hat{t}_p - t_p| \le b_n$, we have the inclusion

$$\mathcal{S}_{t_p+\lambda+a_n+b_n} \subseteq \widehat{\mathcal{S}}_{t_p+\lambda+b_n} \subseteq \widehat{\mathcal{S}}_{\hat{t}_p+\lambda} \subseteq \widehat{\mathcal{S}}_{\hat{t}_p} \subseteq \widehat{\mathcal{S}}_{t_p-b_n} \subseteq \mathcal{S}_{t_p-a_n-b_n},$$

which induces the following homomorphisms f, g, h:

$$\mathcal{H}_0(\mathcal{S}_{t_p+\lambda+a_n+b_n}) \xrightarrow{f} \mathcal{H}_0(\widehat{\mathcal{S}}_{\hat{t}_p+\lambda}) \xrightarrow{g} \mathcal{H}_0(\widehat{\mathcal{S}}_{\hat{t}_p}) \xrightarrow{h} \mathcal{H}_0(\mathcal{S}_{t_p-a_n-b_n}).$$

Then, by the same argument as that of the proof of Theorem 4,

$$\beta_{0,\lambda+2(a_n+b_n)}(t_p - a_n - b_n) = \operatorname{rank}\operatorname{im}(h \circ g \circ f) \leq \operatorname{rank}\operatorname{im}(g) = \widehat{\beta}_{0,\lambda}(\hat{t}_p).$$

Upper bound. For $a_n + b_n < \lambda/2$, an analogous argument can be made for the upper bound

$$\widehat{\beta}_{0,\lambda}(\hat{t}_p) \leq \beta_{0,\lambda-2(a_n+b_n)}(t_p + a_n + b_n)$$

of the sandwich inequality, by considering, on the event $\|\hat{\tau} - \tau\|_\infty \leq a_n$ and $|\hat{t}_p - t_p| \leq b_n$, the inclusion

$$\widehat{\mathcal{S}}_{\hat{t}_p+\lambda} \subseteq \widehat{\mathcal{S}}_{t_p+\lambda-b_n} \subseteq \mathcal{S}_{t_p+\lambda-a_n-b_n} \subseteq \mathcal{S}_{t_p+a_n+b_n} \subseteq \widehat{\mathcal{S}}_{t_p+b_n} \subseteq \widehat{\mathcal{S}}_{\hat{t}_p},$$

which induces the upper bound using the same procedure.

Thus, combining the lower and the upper bound yields the desired result:

$$\beta_{0,\lambda+2(a_n+b_n)}(t_p - a_n - b_n) \leq \widehat{\beta}_{0,\lambda}(\hat{t}_p) \leq \beta_{0,\lambda-2(a_n+b_n)}(t_p + a_n + b_n).$$

Since $a_n, b_n \to 0^+$, the equality statement follows immediately if the true persistent zeroth Betti number $\beta_{0,\ell}(s)$ is locally constant as a function of (s, ℓ) in a neighborhood of (t, λ), i.e., it is constant over $s \in [t - a_n - b_n, t + a_n + b_n]$ and $\ell \in [\lambda - 2(a_n + b_n), \lambda + 2(a_n + b_n)]$. □

Proof of Theorem 6. When doubly robust or inverse-probability estimators are used for CATE estimation, the existence of a sequence c_n satisfying the conditions stated in the corollary is implied by Assumption 1.

We now proceed to the main proof. Since we will be working on the event $\|\hat{\tau} - \tau\|_\infty \leq a_n$, $|\hat{t}_p - t_p| \leq b_n$, and $\|\hat{e} - e\|_\infty \leq c_n$, the subsequent results hold for all sufficiently large n, with probability at least $1 - \delta$.

Lower bound. We first derive the lower bound

$$\beta_{0,(\lambda+2(a_n+b_n),\,\mu+2c_n)}(t_p - a_n - b_n, \varepsilon - c_n) \leq \widehat{\beta}_{0,(\lambda,\mu)}(\hat{t}_p, \varepsilon).$$

of the sandwich inequality. Not yet accounting for the perturbation in $\hat{e}$, we know from the proof of Corollary 5 that, on the event $\|\hat{\tau} - \tau\|_\infty \leq a_n$ and $|\hat{t}_p - t_p| \leq b_n$,

$$\mathcal{S}_{t_p+\lambda+a_n+b_n,\,\varepsilon+\mu}(\tau, \hat{e}) \subseteq \mathcal{S}_{\hat{t}_p+\lambda,\,\varepsilon+\mu}(\hat{\tau}, \hat{e}) \subseteq \mathcal{S}_{\hat{t}_p,\varepsilon}(\hat{\tau}, \hat{e}) \subseteq \mathcal{S}_{t_p-a_n-b_n,\,\varepsilon}(\tau, \hat{e}).$$

Further accounting for the perturbation $\|\hat{e} - e\|_\infty \leq c_n$, we have

$$\begin{aligned}\mathcal{S}_{t_p+\lambda+a_n+b_n,\,\varepsilon+\mu+c_n}(\tau, e) &\subseteq \mathcal{S}_{t_p+\lambda+a_n+b_n,\,\varepsilon+\mu}(\tau, \hat{e}),\\ \mathcal{S}_{t_p-a_n-b_n,\varepsilon}(\tau, \hat{e}) &\subseteq \mathcal{S}_{t_p-a_n-b_n,\,\varepsilon-c_n}(\tau, e).\end{aligned}$$

Combining the above results yields the inclusion

$$\mathcal{S}_{t_p+\lambda+a_n+b_n,\,\varepsilon+\mu+c_n}(\tau, e) \subseteq \mathcal{S}_{\hat{t}_p+\lambda,\,\varepsilon+\mu}(\hat{\tau}, \hat{e}) \subseteq \mathcal{S}_{\hat{t}_p,\varepsilon}(\hat{\tau}, \hat{e}) \subseteq \mathcal{S}_{t_p-a_n-b_n,\,\varepsilon-c_n}(\tau, e),$$

which induces the following homomorphisms f, g, h:

$$\mathcal{H}_0(\mathcal{S}_{t_p+\lambda+a_n+b_n,\,\varepsilon+\mu+c_n}) \xrightarrow{f} \mathcal{H}_0(\widehat{\mathcal{S}}_{\hat{t}_p+\lambda,\,\varepsilon+\mu}) \xrightarrow{g} \mathcal{H}_0(\widehat{\mathcal{S}}_{\hat{t}_p,\varepsilon}) \xrightarrow{h} \mathcal{H}_0(\mathcal{S}_{t_p-a_n-b_n,\,\varepsilon-c_n}).$$

From an analogous argument as that of the proof of Theorem 4,

$$\begin{aligned}\beta_{0,(\lambda+2(a_n+b_n),\,\mu+2c_n)}(t_p - a_n - b_n, \varepsilon - c_n) &= \operatorname{rank}\operatorname{im}(h \circ g \circ f)\\ &\leq \operatorname{rank}\operatorname{im}(g) = \widehat{\beta}_{0,(\lambda,\mu)}(\hat{t}_p, \varepsilon).\end{aligned}$$

Upper bound. For $a_n + b_n < \lambda/2$ and $c_n < \mu/2$, an analogous argument can be made for the upper bound

$$\widehat{\beta}_{0,(\lambda,\mu)}(\hat{t}_p, \varepsilon) \leq \beta_{0,(\lambda-2(a_n+b_n),\mu-2c_n)}(t_p + a_n + b_n, \varepsilon + c_n)$$

of the sandwich inequality. Again, not yet accounting for the perturbation in $\hat{e}$, we know from the proof of Corollary 5 that, on the event $\|\hat{\tau} - \tau\|_\infty \le a_n$ and $|\hat{t}_p - t_p| \le b_n$,

$$\mathcal{S}_{\hat{t}_p+\lambda,\,\varepsilon+\mu}(\hat{\tau},\hat{e}) \;\subseteq\; \mathcal{S}_{t_p+\lambda-a_n-b_n,\,\varepsilon+\mu}(\tau,\hat{e}) \;\subseteq\; \mathcal{S}_{t_p+a_n+b_n,\,\varepsilon}(\tau,\hat{e}) \;\subseteq\; \mathcal{S}_{\hat{t}_p,\varepsilon}(\hat{\tau},\hat{e}).$$

Further accounting for the perturbation $\|\hat{e} - e\|_\infty \le c_n$, we have

$$\mathcal{S}_{t_p+\lambda-a_n-b_n,\,\varepsilon+\mu}(\tau,\hat{e}) \subseteq \mathcal{S}_{t_p+\lambda-a_n-b_n,\,\varepsilon+\mu-c_n}(\tau,e) \subseteq \mathcal{S}_{t_p+a_n+b_n,\,\varepsilon+c_n}(\tau,e) \subseteq \mathcal{S}_{t_p+a_n+b_n,\,\varepsilon}(\tau,\hat{e}).$$

Combining the above results yields the inclusion

$$\mathcal{S}_{\hat{t}_p+\lambda,\,\varepsilon+\mu}(\hat{\tau},\hat{e}) \;\subseteq\; \mathcal{S}_{t_p+\lambda-a_n-b_n,\,\varepsilon+\mu-c_n}(\tau,e) \;\subseteq\; \mathcal{S}_{t_p+a_n+b_n,\,\varepsilon+c_n}(\tau,e) \;\subseteq\; \mathcal{S}_{\hat{t}_p,\varepsilon}(\hat{\tau},\hat{e}),$$

which induces the upper bound using the same procedure.

Thus, combining the lower and the upper bound yields the desired result:

$$\begin{aligned}\beta_{0,(\lambda+2(a_n+b_n),\,\mu+2c_n)}(t_p - a_n - b_n, \varepsilon - c_n) \;&\le\; \widehat{\beta}_{0,(\lambda,\mu)}(\hat{t}_p,\varepsilon) \\ &\le\; \beta_{0,(\lambda-2(a_n+b_n),\mu-2c_n)}(t_p + a_n + b_n, \varepsilon + c_n).\end{aligned}$$

Since $a_n, b_n, c_n \to 0^+$, the equality $\widehat{\beta}_{0,(\lambda,\mu)}(\hat{t}_p) = \beta_{0,(\lambda,\mu)}(t_p)$ follows immediately if the true persistent zeroth Betti number $\beta_{0,(\ell,\ell')}(s)$ is locally constant as a function of (s,ℓ,ℓ') in a neighborhood of (t_p,λ,μ), i.e., it is constant over $s \in [t-a_n-b_n, t+a_n+b_n]$, $\ell \in [\lambda-2(a_n+b_n), \lambda+2(a_n+b_n)]$, and $\ell' \in [\mu - 2c_n, \mu + 2c_n]$. Under Assumption 4, it further holds that

$$\widehat{\beta}_{0,(\lambda,0)}(\hat{t}_p) \;=\; \widehat{\beta}_{0,(\lambda,\mu)}(\hat{t}_p) \;=\; \beta_{0,(\lambda,\mu)}(t_p) \;=\; \beta_{0,(\lambda,0)}(t_p),$$

which gives our desired result. □

B Experiment Details

We sample covariates from $X \sim \mathrm{Uniform}(0,1)$, and set the capacity $p = 0.4$ and p-coverage threshold $t_p = 0.1$. Then, we define a smooth window

$$w(x;a,b,s) = \mathrm{expit}\left(\frac{x-a}{s}\right)\mathrm{expit}\left(\frac{b-x}{s}\right), \qquad \mathrm{expit}(z) = \frac{1}{1+e^{-z}},$$

and model the CATE as

$$\begin{aligned}\tau(x) = c &+ w(x;0.085,0.205,0.014)\left[0.185 + 0.060\sin\left(2\pi\cdot 2.7\frac{x-0.085}{0.12}\right)\right] \\ &+ w(x;0.28,0.42,0.026)\left[0.62\left\{1 + 0.10\cos\left(\pi\frac{x-0.35}{0.14}\right)\right\}\right] \\ &+ w(x;0.495,0.645,0.014)\left[0.185 + 0.060\sin\left(2\pi\cdot 3.2\frac{x-0.495}{0.15} + 0.20\right)\right] \\ &+ w(x;0.75,0.89,0.026)\left[0.68\left\{1 + 0.09\cos\left(\pi\frac{x-0.82}{0.14}\right)\right\}\right].\end{aligned}$$

Here c is a constant chosen so that the capacity threshold satisfies the capacity constraint$P\{\tau(X) > t_p\} = p$. Equivalently, $c = t_p - Q_{1-p}(\tau_{\mathrm{base}}(X))$, where τ_{base} denotes the expression above without the constant c. Then, we can define the regression functions symmetrically as

$$\mu_0(x) = g(x) - \frac{1}{2}\tau(x), \qquad \mu_1(x) = g(x) + \frac{1}{2}\tau(x),$$

where

$$g(x) = 0.8 + 0.4x + 0.15\sin(2\pi x) + 0.08\cos(4\pi x).$$

Thus,

$$\mu_1(x) - \mu_0(x) = \tau(x).$$

The outcome Y is generated by adding mean-zero Gaussian noise with standard deviation 0.05 to μ_1 and μ_0, and the treatment assignment is modeled as

$$\mathbb{P}(A \mid X) \;=\; \mathrm{expit}(0.70\sin(2\pi x) + 0.30\cos(4\pi x)).$$