

Partial Identification of Policy-Relevant Treatment Effects with Instrumental Variables via Optimal Transport

Jiyuan Tan Jose Blanchet Vasilis Syrgkanis

`{jiyuantan, jose.blanchet, vsyrgk}@stanford.edu`

Management Science and Engineering, Stanford University

October 15, 2026

Abstract

Policy-Relevant Treatment Effects (PRTEs) are generally not point-identified under standard instrumental variable (IV) assumptions when the instrument generates limited support in treatment propensity. Existing approaches optimize over marginal treatment response functions subject to moment restrictions and can discard identifying distributional information. We show that PRTE partial identification in the generalized Roy model can instead be formulated as a Constrained Conditional Optimal Transport (CCOT) problem, which reduces analytically to separable one-dimensional OT problems with product costs, yielding sharp closed-form bounds. We also develop estimation and inference procedures: for discrete instruments, a Double Machine Learning approach based on Neyman-orthogonal scores that accommodates high-dimensional covariates while achieving the parametric $\sqrt{n}$ rate and asymptotic normality; for continuous instruments, we characterize the nonparametric convergence rates. The framework accommodates covariates as well as discrete and continuous instruments. In simulations and a bed-net subsidy application, our bounds are substantially tighter than existing moment-relaxation methods.

1 Introduction

Instrumental variables (IVs) are a standard tool for identifying causal effects when treatment assignment is endogenous in observational studies. Under the classical IV assumptions, Angrist and Imbens [1995] show that the Local Average Treatment Effect (LATE), defined as the average effect for the subpopulation of compliers, is point-identified. However, many policy questions are not about compliers themselves; they ask how outcomes would change under a counterfactual policy that shifts treatment take-up in the broader population. Such questions are naturally captured by Policy-Relevant Treatment Effects (PRTEs), which are central targets in the generalized Roy and Marginal Treatment Effect (MTE) framework [Heckman and Vytlacil, 1999, 2005, Heckman et al., 2006]. This framework encompasses a wide range of causal estimands, including the Average Treatment Effect (ATE); empirical policy evaluations within it include Carneiro et al. [2010, 2011]. Unlike LATE, PRTEs are generally not point-identified when the instrument induces only limited support in the treatment propensity.

This limited-support problem is central in the generalized Roy model [Heckman and Vytlacil, 2005], where treatment selection follows a threshold-crossing mechanism formally equivalent to the classical monotonicity condition [Angrist and Imbens, 1995, Vytlacil, 2002]. In this framework, PRTEs can be written as weighted averages of the MTE, so identification depends on how much of the latent resistance distribution is reached by the propensity score induced by the instrument. When the propensity score lacks full support on the unit interval, PRTEs and related quantities such as the ATE remain only partially identified, and a long line of work has studied such identified sets under various assumptions [Manski, 1990, 1997, Manski and Pepper, 2000, Heckman and Vytlacil, 2001, Mogstad et al., 2018].

Despite its importance, deriving tight, closed-form bounds for general PRTEs remains a significant challenge. Existing strategies fall into three broad approaches: moment-relaxation over MTR functions [Mogstad et al., 2018, Mogstad and Torgovitsky, 2018], second-order stochastic dominance constraints on the MTE function solved by rearrangement [Marx, 2024], and infinite-dimensional linear programming over the joint law of outcome and latent resistance solved by sieve approximation [Han and Yang, 2024]. Each faces limitations—discarded distributional information, restricted scope to weights on the complier interval under a binary instrument, lack of closed-form solutions or incompatibility with continuous instruments—that we discuss in detail in Section 1.1.

In this paper, we develop a unified Optimal Transport (OT) framework for partial

identification in the generalized Roy model and use it to derive sharp, closed-form bounds for PRTEs. Specifically, we recast PRTE partial identification as a Constrained Conditional Optimal Transport (CCOT) problem, a structured instance within the broader OT framework that is equivalent to the infinite-dimensional LP of [Han and Yang \[2024\]](#). Rather than summarizing the observed distributions through moment conditions or imposing distributional constraints on conditional mean (MTE) functions, we place the distributional constraints directly on the joint conditional law of the potential outcome and latent resistance. Then, we exploit the OT structure to prove that this heavily constrained, multidimensional CCOT problem analytically reduces to separable one-dimensional OT problems. This reduction completely bypasses computationally intensive linear programming and sieve approximation, yielding explicit, closed-form solutions for the sharp bounds.

Our main contributions can be summarized as follows.

1. **A Unified OT Framework for the Generalized Roy Model:** We formulate partial identification in the generalized Roy model as a CCOT problem. The framework imposes distributional constraints directly at the joint-law level, and we prove that the constraints decompose analytically into separable marginal constraints.
2. **Analytic Decomposition and Closed-Form PRTE Bounds:** Leveraging the decomposition, we prove the CCOT problem decomposes into separable one-dimensional OT problems with product costs, yielding explicit, closed-form sharp bounds and bypassing the sieve approximation required by existing computational approaches.
3. **Semiparametric Estimation and Inference:** We develop estimation results for the binary-treatment bounds. For discrete instruments, we leverage Double Machine Learning (DML) to construct a debiased estimator, achieving $\sqrt{n}$ -consistency and asymptotic normality. For continuous instruments, we explicitly characterize the nonparametric convergence rates.
4. **Empirical Validation:** We validate our theoretical results through both synthetic simulations and a real-world empirical application, demonstrating how our closed-form bounds tighten the identified set relative to the moment-relaxation methods.

Replacing threshold crossing by a conditional quantile representation of the selection equation yields sharp closed-form bounds for continuous and multi-valued treatments, under a rank-overlap condition and an algebraic condition on the instrument-induced latent intervals,

respectively. A companion note develops those results in full; the present paper treats the binary case throughout.

Organization of this paper. Section 2 introduces the setup and optimal transport background; Section 3 presents the main identification theory and closed-form bounds; Section 4 develops estimation and inference; and Section 5 reports simulation and empirical results. The main identification and inference proofs, the algorithms, and the orthogonal score formulas are in the Appendix; the Online Appendix collects the notation table, worked examples, supporting derivations, and full simulation and empirical documentation.

1.1 Related Work

Local IV and MTE. The Marginal Treatment Effect (MTE) of Heckman and Vytlacil [1999, 2005] is the building block of our identification analysis: it is the treatment effect at a given quantile of unobserved resistance, and all standard IV estimands—LATE, ATE, ATT, and PRTE—are weighted averages of it [Heckman and Vytlacil, 2005, Heckman et al., 2006]. Vytlacil [2002] establishes the equivalence between the threshold-crossing selection model and the monotonicity assumption of Angrist and Imbens [1995]. Point identification proceeds via the Local Instrumental Variables (LIV) estimand [Heckman and Vytlacil, 1999], which requires full propensity score support over $[0, 1]$ —the condition that fails under limited instrument variation, the regime our paper addresses.

Partial Identification. Manski [1990, 1997], Manski and Pepper [2000] pioneered partial identification, providing closed-form bounds for the IV problem under a range of assumptions; as a special case, our bound recovers the classical Manski bound for the ATE. In the econometric literature, partial identification is often cast as estimating the identified set defined by moment inequalities [Chernozhukov et al., 2007]; see Tamer [2010] for a comprehensive survey.

Balke and Pearl [1997] develop the linear programming approach for partial identification, which has recently been revisited and generalized using modern techniques [Balazadeh Meresht et al., 2022, Guo et al., 2022, Duarte et al., 2024, Levis et al., 2025]; in particular, Levis et al. [2025] use covariates to tighten the classical bound of Balke and Pearl [1997] and develop estimation theory in the IV setting. Several works connect partial identification to robust optimization [Guo et al., 2022, Balazadeh Meresht et al., 2022, Fan et al., 2025, Tan et al., 2024], and, in settings without unmeasured confounding, the problem can be formulated

as an optimal transport problem with marginal constraints derived from the observed-data distribution [Ji et al., 2024, Lin et al., 2025a,b]; see Gunsilius [2025] for a recent primer. In related but distinct settings, Fan et al. [2017] study partial identification of functionals of the joint distribution of potential outcomes when the conditional marginals are identified, and Ober-Reynolds [2023] characterizes sharp bounds on functionals of the joint complier outcome distribution via optimal transport under the binary-instrument LATE framework of Angrist and Imbens [1995]. The first bounds functionals of the joint law from identified marginals, the second an identified subpopulation; neither addresses the PRTE weights we study here.

PRTE Partial Identification. Three strategies have emerged for PRTE partial identification under limited support. Mogstad et al. [2018] optimize over MTR functions subject to moment constraints [Mogstad and Torgovitsky, 2018]. This moment relaxation discards distributional information beyond the first moment. Han and Yang [2024] preserve full distributional content by formulating an infinite-dimensional LP over the joint law of latent states, solved by sieve approximation for discrete instruments. Marx [2024] generalizes the approach of Mogstad et al. [2018] by imposing full distributional content on the MTE function via a second-order stochastic dominance (SSD) constraint, obtaining sharp closed-form bounds by rearrangement; the paper’s explicit theorems cover weights on the complier interval under a binary instrument, with extensions beyond compliers requiring rank similarity.

Our work differs from all three in scope: we impose distributional constraints directly on the joint conditional law of the potential outcome and latent resistance rather than on a conditional mean, cast the resulting program as a CCOT problem equivalent to the LP of Han and Yang [2024], and show it reduces analytically to separable one-dimensional OT problems with product costs. This delivers (i) sharp closed-form bounds for the full PRTE on all of $[0, 1]$ under multi-valued instruments via an interval decomposition across $K + 1$ compliance types, including boundary intervals where only one potential outcome marginal is identified, (ii) extensions to continuous and mixed instruments with sub-problems indexed by the propensity score continuum, and (iii) discrete-instrument inference together with continuous-instrument rate results, which are limited in either Marx [2024] or Han and Yang [2024]. On the complier interval under a binary instrument, the CCOT solution and the SSD rearrangement of Marx [2024] coincide. Beyond PRTE, our sharpness result characterizes the identified set at the level of the full joint latent law and sheds light on the joint-identification question raised in Marx [2024] through its reduction to independent one-dimensional marginal

constraints.

Remark (notation). K is a population count in propensity decompositions and a chosen function when written as the local-linear smoothing kernel $K(t)$. In Theorem 3.4: $K = |\{p(z) : z \in \mathcal{Z}, p(z) > 0\}|$, excluding the structural boundary $p_0 = 0$. In Theorem 3.7: K is the number of ordered maximal closed intervals in $p(\mathcal{Z})$ after merging overlaps, including degenerate singleton intervals and any interval containing zero. In Proposition 4.3, Proposition 4.4, Lemma D.1, Theorem 4.7: K counts the distinct positive baseline propensity levels at each x and does not depend on x ; $k = 0, \dots, K - 1$ indexes the identified intervals and $k = K$ indexes the residual tail. In Corollary O5.1: $K(t)$ is the base one-dimensional smoothing kernel in $K_{h_n}(t) = K(t/h_n)/h_n$, rather than a propensity-level count.

2 Preliminaries

2.1 IV Model and Policy-Relevant Treatment Effects

For most of this paper, we work with the generalized Roy model [Heckman and Vytlacil, 2005]. In the model, the observable variables are treatment $W \in \mathcal{W} = \{0, 1\}$, instrument $Z \in \mathcal{Z}$, covariates $X \in \mathcal{X}$, and outcome $Y \in \mathcal{Y}$, where $\mathcal{X} \subset \mathbb{R}^{d_x}$, $\mathcal{Z}$, and $\mathcal{Y} \subset \mathbb{R}$ are the domains of the corresponding variables. We consider both discrete and continuous instrument settings in this paper. Throughout the paper, we assume that all domains are compact, carry their Borel σ -fields, and are standard Borel spaces. Write $y_{\min} = \min \mathcal{Y}$ and $y_{\max} = \max \mathcal{Y}$; compactness guarantees that these extrema exist, and we treat the outcome domain and hence its feasible endpoints as known. The unobservables are the potential outcomes $Y(0), Y(1)$, and the variable U in the selection equation (2.1). Thus, Z is the instrument, W is the treatment, X is observed covariate information, Y is the observed outcome, and U is an unobserved latent resistance variable.

Assumption 2.1 (Structural Assumptions). We impose the following structural assumptions on our IV model. Throughout, we assume $Y(w) \in \mathcal{Y}$ almost surely.

1. (Consistency) $Y = WY(1) + (1 - W)Y(0)$.
2. (Conditional Instrumental Exogeneity) $Z \perp\!\!\!\perp (U, Y(0), Y(1)) \mid X$, where $\perp\!\!\!\perp$ denotes statistical independence.

3. (Threshold Crossing) The treatment is selected by

$$W = \mathbb{1}(U \leq p(Z, X)), \quad (2.1)$$

where $p(Z, X) = \mathbb{P}(W = 1 \mid Z, X)$. Moreover, $U \mid X \sim \text{Unif}(0, 1)$.

4. (Common Instrument Support) The conditional support of Z given X does not depend on X : $\text{Supp}(Z \mid X = x) = \mathcal{Z}$ for all $x \in \mathcal{X}$.

Remark. $Y = WY(1) + (1 - W)Y(0)$ almost surely under the structural population law. $W = \mathbb{1}(U \leq p(Z, X))$ almost surely under the structural population law. $\mathcal{L}(U \mid X = x) = \text{Unif}(0, 1)$ for $\mathbb{P}_X$ -almost every x .

Remark (notation). $p(z, x) = \mathbb{P}(W = 1 \mid Z = z, X = x)$ denotes the baseline population conditional treatment propensity, with values in $[0, 1]$. In Theorem 3.4: Since X is almost surely constant, p is written as $p(z)$. In Theorem 3.7: Since X is almost surely constant, p is written as $p(z)$, using its continuous version on each connected component of $\mathcal{Z}$. In Theorem 3.8, Proposition 4.8, Theorem 4.10, Corollary O5.1: For almost every x , $z \mapsto p(z, x)$ is the version continuous on each connected component of $\mathcal{Z}$, so its range and extrema are taken over that version throughout the conditional support. $\perp\!\!\!\perp$ denotes independence, with $A \perp\!\!\!\perp B \mid C$ denoting conditional independence given C , namely factorization of the conditional joint law into the two conditional marginal laws almost surely. $\mathcal{Y}$ is the fixed, known, compact outcome domain contained in $\mathbb{R}$, equipped with its Borel σ -field, with known feasible endpoints $y_{\min} = \min \mathcal{Y}$ and $y_{\max} = \max \mathcal{Y}$. Y denotes the observed outcome taking values in the known compact Borel domain $\mathcal{Y} \subset \mathbb{R}$, with $y_{\min} \leq Y \leq y_{\max}$. In Proposition 4.4, Lemma D.1, Theorem 4.7: In part (ii), $Y \in \{0, 1\}$. $\mathcal{Z}$ is the compact standard Borel instrument domain equipped with its Borel σ -field and is the common conditional support $\text{Supp}(Z \mid X = x) = \mathcal{Z}$ for all $x \in \mathcal{X}$. In Theorem 3.4, Proposition 4.3, Proposition 4.4, Lemma D.1, Theorem 4.7: $\mathcal{Z}$ is finite. In Theorem 3.7, Theorem 3.8: $\mathcal{Z}$ has finitely many connected components. In Proposition 4.8, Theorem 4.10, Corollary O5.1: $\mathcal{Z}$ is connected. $\mathcal{X} \subset \mathbb{R}^{d_X}$ is the compact covariate domain, equipped with its Borel σ -field and regarded as a standard Borel space. X denotes observed covariate information taking values in the compact Borel domain $\mathcal{X} \subset \mathbb{R}^{d_X}$; x denotes a conditioning value, and expectations indexed by X average over its marginal law. In Theorem 3.4, Theorem 3.7: X is almost surely constant, so conditional covariate arguments are suppressed. U denotes the latent resistance in the threshold-crossing model, invariant under the policy change, with interval-uniform uses denoting its law restricted to the indicated resistance interval. In Theorem 3.4: In

defining $Q_{\omega,i}$, $U \sim \text{Unif}(p_i, p_{i+1})$ denotes the resistance draw conditional on $U \in (p_i, p_{i+1}]$, for $i = 0, \dots, K-1$ with $p_i < p_{i+1}$. In Theorem 3.7: In defining $Q_{\omega,k}$, $U \sim \text{Unif}(G_k)$ denotes the resistance draw restricted to the nonempty gap G_k , for $k = 0, \dots, K-1$; no gap-specific law is defined for an empty gap. In Theorem 3.8: For fixed x and a nonempty $G_k(x)$, the draw $U \sim \text{Unif}(G_k(x))$ used to define $Q_{\omega,k|x}$ is the resistance law conditional on $X = x$ and $U \in G_k(x)$, for $k = 0, \dots, K_x-1$; in the randomized rank construction this gap draw is independent of $S \sim \text{Unif}(0, 1)$.

Because $Y(w) \in \mathcal{Y}$ and $\mathcal{Y}$ is compact, $y_{\min} \leq Y(w) \leq y_{\max}$ almost surely; boundedness therefore follows from the maintained setup. It is without loss of generality to impose $U | X \sim \text{Unif}(0, 1)$ when the original latent resistance has a continuous conditional distribution, because we can normalize that distribution.¹ The common instrument support condition lets us use the same instrument domain at each covariate value. The observed law identifies conditional probabilities almost surely; the continuity condition introduced in ?? 3.6 selects their continuous propensity-score version throughout the conditional support.

Remark (notation). In transport notation, c is the specified cost function on the two transported state spaces; a conditional cost $c(y, u, x)$ is read at fixed x as the function $(y, u) \mapsto y\omega(x, u)$. In Theorem 2.3: $\mathcal{A} = \mathcal{B} = \mathbb{R}$ and $c(a, b) = a \cdot b$.

The fundamental building block for evaluating causal parameters in this framework is the MTE, originally introduced by Heckman and Vytlačil [1999, 2005]. The MTE is defined as the average treatment effect for individuals with covariates $X = x$ and latent variable $U = u$:

$$\text{MTE}(x, u) = \mathbb{E}[Y(1) - Y(0) | X = x, U = u].$$

In policy analysis, we are often interested in the effect of an intervention that changes treatment take-up from a baseline status to a new regime. Let the baseline policy be characterized by the current propensity score $p(Z, X)$ and the corresponding treatment status W . Consider an alternative policy that induces a counterfactual propensity score $q(Z, X)$ and a new treatment status $W^q = \mathbf{1}(U \leq q(Z, X))$, where we assume the latent variable U is invariant to the policy change, as is standard in the MTE framework.

Remark (notation). $q : \mathcal{Z} \times \mathcal{X} \rightarrow [0, 1]$ denotes the alternative-policy propensity function determining $W^q = \mathbf{1}(U \leq q(Z, X))$ with the same latent resistance U as under the baseline policy. In Proposition 4.3, Proposition 4.4, Lemma D.1, Theorem 4.7, Proposition 4.8,

¹Let V denote the original latent resistance, let $c(Z, X)$ be its cutoff, and suppose $F_{V|X}(\cdot | x)$ is continuous. Define $\tilde{U} = F_{V|X}(V | X)$ and the transformed cutoff $\tilde{p}(z, x) = F_{V|X}(c(z, x) | x)$. The conditional probability integral transform gives $\tilde{U} | X \sim \text{Unif}(0, 1)$, and $\mathbf{1}\{V \leq c(Z, X)\} = \mathbf{1}\{\tilde{U} \leq \tilde{p}(Z, X)\}$ almost surely.

Theorem 4.10, Corollary O5.1: $q(z, x) = \phi(z, x, p(z, x))$, where the known map $\phi : \mathcal{Z} \times \mathcal{X} \times [0, 1] \rightarrow [0, 1]$ is jointly measurable. $W^q = \mathbb{1}(U \leq q(Z, X))$ uses the same latent resistance U as baseline treatment W ; U is invariant to the policy change.

The PRTE evaluates the average per-person impact of shifting from the baseline policy to the alternative policy. Define the outcome realized under the alternative policy by $Y^q = W^q Y(1) + (1 - W^q) Y(0)$. The PRTE is formally defined as

$$\text{PRTE} = \frac{\mathbb{E}[Y^q - Y]}{\mathbb{E}[W^q - W]}.$$

Many common treatment effect parameters can be expressed as a weighted average of the MTE. Specifically, the PRTE can be rewritten as:

$$\text{PRTE} = \mathbb{E}_{X,U}[\text{MTE}(X, U) \omega_{\text{PRTE}}(X, U)],$$

where the policy-specific weight function is given by:

$$\omega_{\text{PRTE}}(x, u) = \frac{\mathbb{P}(q(Z, X) \geq u \mid X = x) - \mathbb{P}(p(Z, X) \geq u \mid X = x)}{\mathbb{E}[q(Z, X) - p(Z, X)]},$$

provided $\mathbb{E}[q(Z, X) - p(Z, X)] \neq 0$. Following Heckman and Vytlacil [2005], Mogstad et al. [2018], we consider the following more general target parameter.

$$\theta_\omega = \mathbb{E}_X \int_0^1 \text{MTE}(X, u) \omega(X, u) du, \quad (2.2)$$

where ω is the target-specific weight. We impose the following standing condition on this weight.

Assumption 2.2 (Target-Weight Regularity). The function $\omega : \mathcal{X} \times [0, 1] \rightarrow \mathbb{R}$ is jointly measurable and identifiable from the observed law, and

$$\mathbb{E}_X \left[\int_0^1 |\omega(X, u)| du \right] < \infty.$$

Remark. Any two structural laws satisfying Structural Assumptions and inducing the same $\mathbb{P}_{\text{obs}}$ give the same $\omega(x, u)$ for $\mathbb{P}_{\text{obs}}(dx) du$ -almost every (x, u) .

Together with the compact outcome domain, ?? 2.2 ensures that θ_ω is finite. We can obtain different PRTEs by taking

$$\omega(x, u) \propto \mathbb{E}[\mathbb{1}(u \leq q_1(Z, X)) - \mathbb{1}(u \leq q_0(Z, X)) \mid X = x], \quad (2.3)$$

Table 1: Common Causal Parameters as Weighted Averages of the MTE

Target Parameter	Notation	Weight Function $\omega(x, u)$
Average Treatment Effect	ATE	1
Average Treatment Effect on the Treated	ATT	$\frac{\mathbb{E}[\mathbb{1}(u \in (0, p(Z, X))) X=x]}{\mathbb{E}[p(Z, X)]}$
Average Treatment Effect on the Untreated	ATU	$\frac{\mathbb{E}[\mathbb{1}(u \in (p(Z, X), 1)) X=x]}{\mathbb{E}[1-p(Z, X)]}$
Local Average Treatment Effect ^a	LATE(z_1, z_0)	$\frac{\mathbb{1}(u \in [p(z_0, x), p(z_1, x)])}{\mathbb{E}[p(z_1, X) - p(z_0, X)]}$
Policy-Relevant Treatment Effect	PRTE	$\frac{\mathbb{P}(q(Z, X) \geq u X=x) - \mathbb{P}(p(Z, X) \geq u X=x)}{\mathbb{E}[q(Z, X) - p(Z, X)]}$

^a Evaluated for a shift from instrument state z_0 to z_1 , assuming $p(z_1, x) \geq p(z_0, x)$.

for given propensity functions q_0 and q_1 , where $q_0 = p$ and $q_1 = q$ correspond to the baseline and alternative propensity score functions, respectively, in the PRTE case. The coefficient hidden inside the proportionality symbol is identifiable. Table 1 shows a variety of common causal parameters under different choices of ω .

Remark (notation). $\omega : \mathcal{X} \times [0, 1] \rightarrow \mathbb{R}$ is the population target weight, jointly measurable and identifiable from the observed law, with $\mathbb{E}_X[\int_0^1 |\omega(X, u)| du] < \infty$. In Theorem 3.4, Theorem 3.7: With X almost surely constant, ω is written as a function $\omega(u)$ of latent resistance alone. In Proposition 4.3, Proposition 4.4, Lemma D.1, Theorem 4.7, Proposition 4.8, Theorem 4.10, Corollary O5.1: $\omega(x, u) = \mathbb{E}_{\text{obs}}[\mathbb{1}(u \leq q(Z, X)) - \mathbb{1}(u \leq p(Z, X)) | X = x]$, where $q(z, x) = \phi(z, x, p(z, x))$ and ϕ is known; the target is the unscaled numerator $\mathbb{E}[Y^q - Y]$. $\theta_\omega = \mathbb{E}_X \int_0^1 \text{MTE}(X, u) \omega(X, u) du$ denotes the target under the latent law being considered, equivalently $\mathbb{E}[Y(1)\omega(X, U)] - \mathbb{E}[Y(0)\omega(X, U)]$. In Proposition 3.1: When written with candidate-law arguments, $\theta_\omega(\pi_0, \pi_1) = \mathbb{E}_{\pi_1}[Y(1)\omega(X, U)] - \mathbb{E}_{\pi_0}[Y(0)\omega(X, U)]$, with $(\pi_0, \pi_1) \in \Gamma(\mathbb{P}_{\text{obs}})$; its value varies with the candidate marginal laws.

While our bounds recover the classical bounds for the first four quantities in Table 1, we focus in particular on PRTEs. As a leading example, consider a uniform policy shift of the form $q(Z, X) = \min(p(Z, X) + \alpha, 1)$ for some $\alpha > 0$. This corresponds to a policy that uniformly raises the probability of treatment uptake by α . In our empirical application

(Section 5), the instrument is the offered price of an insecticide-treated bed net [Dupas, 2014], and $q = p + \alpha$ models a uniform price subsidy that lowers the cost of purchase and thereby increases take-up probability by α .

Heckman and Vytlačil [2001] shows that if the support of $p(Z, x)$ spans the entire interval $(0, 1)$ for all x , then the MTE is identifiable. However, because the common support of $p(Z, X)$ is often a strict subset of $(0, 1)$, the MTE is generally not identifiable. For instance, when both X and Z are discrete, the supports of $p(Z, x)$ are discrete and this full-support condition is violated.

2.2 A Primer on Optimal Transport

To build the framework for our partial identification results, we briefly introduce the relevant concepts from OT theory. OT provides a mathematical framework for coupling two probability distributions at minimum cost, and has emerged as a powerful toolkit for bounding joint distributions when only their marginals are known [Galichon, 2016].

Let μ and ν be two probability measures on state spaces $\mathcal{A}$ and $\mathcal{B}$, respectively. In the context of partial identification, we often observe the marginal distributions of two variables but not their joint distribution. We define $\Pi(\mu, \nu)$ as the set of all couplings between μ and ν . Formally, a coupling $\pi \in \Pi(\mu, \nu)$ is a joint probability measure on the product space $\mathcal{A} \times \mathcal{B}$ such that its marginals are exactly μ and ν .

Remark (notation). μ denotes a probability measure on the state space appropriate to the coupling or quantile expression in which it occurs. In Theorem 2.3: $\mu \in \mathcal{P}(\mathbb{R})$, and either μ has bounded support and ν has a finite first moment, or vice versa. ν denotes a probability measure in transport notation, a locally bound candidate real-valued function of X inside the conditional quantile-regression minimization, or an unperturbed scalar quantile boundary in a local boundary-perturbation calculation. In Theorem 2.3: $\nu \in \mathcal{P}(\mathbb{R})$, and either μ has bounded support and ν has a finite first moment, or vice versa. In Lemma D.1, Theorem 4.7: In the continuous-outcome quantile-regression minimization, ν is a candidate function $\mathcal{X} \rightarrow \mathbb{R}$; the resulting population threshold is $\nu_{j,k}(x) = Q_{Y,k|x}(\kappa_{j,k}(x))$ for $j = 1, \dots, l_k$, while $\nu_{0,k}(x) = y_{\min}$ is fixed. In Lemma D.1: In the continuous-outcome proof's indicator-boundary calculation, unindexed ν denotes the scalar value of the threshold being perturbed at a given covariate value: $a_t = \nu + t \delta \nu$, and $\int_{\nu}^{a_t} (a_t - y) dF(y)$ uses the oriented-integral convention when $a_t < \nu$. $\mathcal{A}$ is reused contextually for the first marginal's state space in the OT primer and for the explicitly specified countable measure-determining class in the LIV argument. In Theorem 2.3: $\mathcal{A} = \mathbb{R}$ is the state space carrying the first marginal μ . In Theorem 3.8: In the LIV proof,

$\mathcal{A} = \{\mathcal{Y} \cap (-\infty, r] : r \in \mathbb{Q}\} \cup \{\mathcal{Y}\}$ is the countable measure-determining class of outcome events used to obtain a common exceptional null set. The meaning of π is determined by its arguments: transport notation denotes a joint probability measure, whereas $\pi(z, x)$ denotes the conditional instrument assignment probability $\mathbb{P}(Z = z \mid X = x)$. In Theorem 2.3: π ranges over $\Pi(\mu, \nu) \subseteq \mathcal{P}(\mathbb{R}^2)$. In Lemma D.1, Theorem 4.7: In instrument-specific probability expressions, $\pi(z, x) = \mathbb{P}_{\text{obs}}(Z = z \mid X = x)$, rather than a transport coupling. The notation $\mathcal{B}$ denotes an OT state space when standing alone and the Borel σ -field operator when applied to an outcome domain. In Theorem 2.3: The second marginal state space is $\mathcal{B} = \mathbb{R}$. In Theorem 3.8: In the kernel definition, $A \in \mathcal{B}(\mathcal{Y})$ means that A is a Borel subset of $\mathcal{Y}$.

Given a cost function $c : \mathcal{A} \times \mathcal{B} \rightarrow \mathbb{R} \cup \{+\infty\}$ that represents the cost of associating state a with state b , the standard Kantorovich optimal transport problem seeks the coupling that minimizes the expected cost:

$$\inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathcal{A} \times \mathcal{B}} c(a, b) d\pi(a, b). \quad (2.4)$$

Generally, solving (2.4) in multidimensional spaces requires computationally intensive linear programming. However, a celebrated result in the OT literature demonstrates that when the spaces are one-dimensional and the cost function satisfies specific structural conditions—such as supermodularity—the optimal coupling admits a sharp, closed-form solution.

We formalize this one-dimensional closed-form solution in the following theorem, which we use to derive the analytic PRTE bounds.

Theorem 2.3 (1D Optimal Transport with Product Cost; see, e.g., [Villani, 2021, Galichon, 2016]). *Let μ and ν be probability measures on $\mathbb{R}$ with quantile functions Q_μ and Q_ν . Suppose either that μ has bounded support and ν has a finite first moment, or vice versa, and let the cost be the simple product $c(a, b) = a \cdot b$. Then, the minimum and maximum of the expected cost over all valid couplings are completely characterized by the countermonotonic and comonotonic couplings, respectively:*

$$\begin{aligned} \min_{\pi \in \Pi(\mu, \nu)} \int_{\mathbb{R}^2} a \cdot b d\pi(a, b) &= \int_0^1 Q_\mu(t) Q_\nu(1 - t) dt, \\ \max_{\pi \in \Pi(\mu, \nu)} \int_{\mathbb{R}^2} a \cdot b d\pi(a, b) &= \int_0^1 Q_\mu(t) Q_\nu(t) dt. \end{aligned}$$

The minimum is attained by the countermonotonic coupling $(Q_\mu(\xi), Q_\nu(1 - \xi))$, and the maximum by the comonotonic coupling $(Q_\mu(\xi), Q_\nu(\xi))$, where $\xi \sim \text{Unif}(0, 1)$.

Remark. For $0 < t < 1$, $(Q_\mu(t), Q_\nu(t)) = (\inf\{a \in \mathbb{R} : \mu((-\infty, a]) \geq t\}, \inf\{b \in \mathbb{R} : \nu((-\infty, b]) \geq t\})$. For $\rho \in \{\mu, \nu\}$ and $t \in (0, 1)$, $Q_\rho(t) = \inf\{a \in \mathbb{R} : \rho((-\infty, a]) \geq t\}$.

Remark (notation). In transport notation, $\xi \sim \text{Unif}(0, 1)$ is the auxiliary variable shared by both coordinates of $(Q_\mu(\xi), Q_\nu(1 - \xi))$ or $(Q_\mu(\xi), Q_\nu(\xi))$.

Theorem 2.3 provides a closed-form characterization of the optimal value of 1D OT problems. We will leverage this characterization to derive tight bounds for PRTEs later.

2.3 Notation

We use $\mathbb{P}$ and $\mathbb{E}$ for probability and expectation, and $\perp\!\!\!\perp$ for (conditional) independence. $\mathbb{1}(\cdot)$ denotes the indicator function. We write d for the differential or Lebesgue measure. For a probability measure μ on $\mathbb{R}$, $Q_\mu(t)$ denotes its quantile function. $\Pi(\mu, \nu)$ denotes the set of all couplings of two probability measures μ and ν , and $W_1(\mu, \nu) = \int_0^1 |Q_\mu(t) - Q_\nu(t)| dt$ their 1-Wasserstein distance. Calligraphic letters $\mathcal{X}, \mathcal{Y}, \mathcal{Z}, \mathcal{W}$ denote domains of the corresponding variables. We write $O_P(\cdot)$ and $o_P(\cdot)$ for stochastic order notation: $X_n = O_P(a_n)$ means X_n/a_n is bounded in probability, and $X_n = o_P(a_n)$ means $X_n/a_n \xrightarrow{P} 0$. A table collecting the variables and symbols used throughout the paper is given in Section O1 of the Online Appendix.

3 Partial Identification via Optimal Transport

3.1 CCOT Formulation of Partial Identification

In this section, we demonstrate how to derive tight, closed-form bounds for the PRTE. We begin by formally mapping the partial identification problem into an optimal transport framework.

The core intuition is to search over all possible joint distributions of the latent and observable variables $(Y(0), Y(1), U, X, Z, W)$ that are compatible with both the observed data and the structural conditions in ?? 2.1. Because the potential outcomes are conditionally independent of the instrument given X , the identifying power of our model is entirely captured by the joint distributions of $(Y(w), U, X)$ for $w \in \{0, 1\}$.

Let π_w denote the joint probability measure of $(Y(w), U, X)$ and let $\pi_w(\cdot \mid x, u)$ denote the conditional distribution of $Y(w)$ given x, u . Recall that our target parameter is a linear

functional of these measures:

$$\theta_w(\pi_0, \pi_1) = \mathbb{E}_{(X,U,Y(1)) \sim \pi_1}[Y(1)\omega(X,U)] - \mathbb{E}_{(X,U,Y(0)) \sim \pi_0}[Y(0)\omega(X,U)].$$

To achieve sharp bounds, we must constrain π_w using all available observational information. Let $\mathbb{P}_{obs}$ denote the true observed joint distribution of (Y, W, Z, X) . Under the threshold-crossing model in ?? 2.1, while U is generally unobserved, its conditional distribution is explicitly tied to the treatment status and the propensity score. Specifically, conditional on X and Z , the event $W = 1$ perfectly corresponds to $U \leq p(Z, X)$, and $W = 0$ corresponds to $U > p(Z, X)$.

Remark (notation). For $w \in \{0, 1\}$, π_w denotes a joint probability law of $(Y(w), U, X)$, including a candidate such law when varied in the CCOT program, and $\pi_w(\cdot \mid u, x)$ denotes a measurable regular conditional outcome law determined only π_w -almost everywhere with respect to its (U, X) -marginal.

Together with the observed distribution of (X, Z) , the joint distributions π_0 and π_1 fully determine the observational distribution of (X, Z, W, Y) . To see this, we can link the unknown structural distributions directly to the observed conditional distribution of the data by exploiting the consistency equation $Y = WY(1) + (1 - W)Y(0)$ and the selection mechanism $W = \mathbb{1}(U \leq p(Z, X))$. For any measurable set $A \subseteq \mathcal{Y}$, the following identity holds for $\mathbb{P}_{Z,X}$ -almost every (z, x) :

$$\mathbb{P}_{obs}(Y \in A, W = 1 \mid Z = z, X = x) = \mathbb{P}(Y(1) \in A, U \leq p(z, x) \mid Z = z, X = x).$$

By the conditional independence assumption $Z \perp\!\!\!\perp (Y(1), U) \mid X$, we can omit Z in the second conditioning:

$$\mathbb{P}(Y(1) \in A, U \leq p(z, x) \mid Z = z, X = x) = \mathbb{P}(Y(1) \in A, U \leq p(z, x) \mid X = x).$$

Since instruments that generate the same propensity value $p(z, x)$ yield the same right-hand side conditional on $X = x$, we can replace conditioning on the full instrument value $Z = z$ with conditioning on the scalar propensity value $p(Z, X) = p(z, x)$ together with $X = x$. Next, we disintegrate the joint measures π_w into the conditional distribution of the potential outcome given the covariates and the latent variable. The threshold-crossing rule selects $U \leq p(z, x)$ into treatment and $U > p(z, x)$ out of it, so write $\mathcal{R}_1(p) := (0, p]$ and $\mathcal{R}_0(p) := (p, 1]$ for the two selected regions. Because $U \mid X \sim \text{Unif}(0, 1)$, integrating over $\mathcal{R}_w(p)$ yields, for $P \mid X = x$ -almost every p and each $w \in \{0, 1\}$,

$$\mathbb{P}_{obs}(Y \in A, W = w \mid p(Z, X) = p, X = x) = \int_{\mathcal{R}_w(p)} \pi_w(A \mid u, x) du.$$

Write $P = p(Z, X)$. Because a regular conditional law at an exact value of a continuous propensity score is defined only almost surely, first define, for $P \mid X = x$ -almost every p , the observed finite-measure-valued maps

$$H_{w,x}(p, dy) := \mathbb{P}_{\text{obs}}(Y \in dy, W = w \mid P = p, X = x), \quad w \in \{0, 1\}. \quad (3.1)$$

To select canonical versions at every point of the conditional propensity support, let $K_{w,x}(dy \mid u)$ be a regular conditional law of $Y(w)$ given $(U, X) = (u, x)$. Under ?? 2.1, the preceding derivation gives

$$H_{w,x}(p, dy) = \int_{\mathcal{R}_w(p)} K_{w,x}(dy \mid u) du, \quad w \in \{0, 1\},$$

for $P \mid X = x$ -almost every p . The right-hand sides are Lipschitz in total variation and hence extend the observed equivalence classes in (3.1) continuously to $\text{supp}(P \mid X = x)$. This extension is unique: two continuous versions that agree almost surely under the conditional law of P agree everywhere on its support. We henceforth use $H_{w,x}$ for these canonical observed versions. Thus $H_{w,x}$ is identified from $\mathbb{P}_{\text{obs}}$ on the conditional propensity support, including its endpoints even when the corresponding singleton events have probability zero; the latent kernels $K_{w,x}$ themselves need not be identified. With no covariates we write H_w . Furthermore, by construction, the marginal distribution of (U, X) under π_w must equal the known population distribution $\mathbb{P}_{U,X}$.

Remark (notation). $P = p(Z, X)$ is the observed baseline propensity index obtained by evaluating the population propensity function at (Z, X) . In Proposition 4.8, Theorem 4.10, Corollary O5.1: The conditional propensity support is $\text{supp}(P \mid X = x) = [\underline{p}(x), \bar{p}(x)]$ in the connected-instrument setting. In Theorem 3.4, Theorem 3.7: With X almost surely constant, the covariate argument is suppressed and $P = p(Z)$. $\mathcal{R}_w$ is the selected-region map for $w \in \{0, 1\}$ and $p \in [0, 1]$, with $\mathcal{R}_1(p) = (0, p]$ and $\mathcal{R}_0(p) = (p, 1]$, including empty intervals at the corresponding boundary propensities. For $w \in \{0, 1\}$ and almost every x , $H_{w,x}$ is the unique total-variation-continuous version on $\text{supp}(P \mid X = x)$ of $\mathbb{P}_{\text{obs}}(Y \in dy, W = w \mid P = p, X = x)$, including support endpoints; its total mass is p for $w = 1$ and $1 - p$ for $w = 0$, and without covariates it is written H_w . H_w , $w \in \{0, 1\}$, denotes the no-covariate canonical continuous version of $\mathbb{P}_{\text{obs}}(Y \in dy, W = w \mid P = p)$ on $\text{supp}(P)$, uniquely determined there by its almost-sure observed values. In Theorem 3.7: $H_1(0, \cdot) = 0$ is specified even when 0 lies outside propensity support, so the initial gap uses a structural boundary submeasure. $K_{w,x}(dy \mid u)$, for $w \in \{0, 1\}$, is a regular conditional probability kernel for $Y(w)$ given $(U, X) = (u, x)$ under a compatible latent law, determined only $\mathbb{P}_{U,X}$ -almost

everywhere and not necessarily identified, with $H_{w,x}(p, dy) = \int_{\mathcal{R}_w(p)} K_{w,x}(dy | u) du$ on the conditional propensity support for $\mathbb{P}_X$ -almost every x .

This allows us to formulate the partial identification of the PRTE as the following infinite-dimensional linear programming problem, in which every constraint is imposed for both $w \in \{0, 1\}$:

$$\begin{aligned} & \max_{\pi_0, \pi_1} / \min_{\pi_0, \pi_1} \quad \theta_\omega(\pi_0, \pi_1) \\ & \text{subject to} \quad \int_{\mathcal{R}_w(p(z,x))} \pi_w(dy | u, x) du = H_{w,x}(p(z, x), dy), \quad \mathbb{P}_{Z,X}\text{-a.e. } (z, x), \\ & \quad \pi_w(dx, du) = du \mathbb{P}_{obs}(dx), \\ & \quad \pi_w \text{ is supported on } \mathcal{Y} \times [0, 1] \times \mathcal{X}. \end{aligned} \tag{3.2}$$

We call the optimization in (3.2) the Constrained Conditional Optimal Transport (CCOT) problem. The feasible set is constrained by the observational compatibility conditions in (3.2); the problem is conditional in that, for fixed covariates x and latent resistance u , it seeks a conditional measure $\pi_w(\cdot | u, x)$ on $\mathcal{Y}$; and it is a transport problem because the objective is linear in $\pi_w(\cdot | u, x)$ with effective cost $c(y, u, x) = y \cdot \omega(x, u)$, so that mass is transported from the latent space U to the outcome space $Y(w)$ conditional on X .

To formally guarantee that the CCOT formulation in (3.2) yields sharp bounds for the PRTE, we must show that any pair of measures (π_0, π_1) satisfying the CCOT constraints corresponds to a valid data-generating process that satisfies our structural assumptions and perfectly reproduces the observed data. Let $\Gamma(\mathbb{P}_{obs})$ be the set of all measure pairs (π_0, π_1) satisfying the constraints in (3.2). The following proposition characterizes the joint distribution of $(Y(0), Y(1), U, X)$, from which observational equivalence of the marginal pair (π_0, π_1) follows immediately.

Proposition 3.1 (Sharp Identified Set of the Joint Latent Law). *Suppose the observed distribution $\mathbb{P}_{obs}$ is generated by a true structural model satisfying ?? 2.1. Let $\tilde{\pi}$ be any probability measure on $\mathcal{Y} \times \mathcal{Y} \times [0, 1] \times \mathcal{X}$ whose $(Y(w), U, X)$ -marginals $\tilde{\pi}_w$ satisfy $(\tilde{\pi}_0, \tilde{\pi}_1) \in \Gamma(\mathbb{P}_{obs})$. Then there exists a joint distribution of the latent and observable variables $(Y(0), Y(1), U, X, Z, W)$ such that:*

- (i) *the joint marginal distribution of $(Y(0), Y(1), U, X)$ is exactly $\tilde{\pi}$;*
- (ii) *?? 2.1 holds for $(Y(0), Y(1), U, X, Z, W)$;*
- (iii) *the induced distribution of (Z, X, W, Y) exactly matches $\mathbb{P}_{obs}$.*

Conversely, every joint marginal $\tilde{\pi}$ arising from a distribution of $(Y(0), Y(1), U, X, Z, W)$ satisfying (ii)–(iii) has $(Y(w), U, X)$ -marginals in $\Gamma(\mathbb{P}_{obs})$ for $w \in \{0, 1\}$.

Remark. Every constraint in (3.2) is imposed simultaneously for both $w \in \{0, 1\}$.

Remark (notation). $\tilde{\pi}$ is a probability measure on $\mathcal{Y}^2 \times [0, 1] \times \mathcal{X}$, interpreted as a candidate joint law of $(Y(0), Y(1), U, X)$. In Proposition 3.1: Its marginals satisfy $(\tilde{\pi}_0, \tilde{\pi}_1) \in \Gamma(\mathbb{P}_{obs})$, hence its (U, X) -marginal is $du \mathbb{P}_{obs}(dx)$; the coupling of $Y(0)$ and $Y(1)$ conditional on (U, X) is unrestricted apart from those marginal constraints. For $w \in \{0, 1\}$, $\tilde{\pi}_w$ is the $(Y(w), U, X)$ -marginal of $\tilde{\pi}$. In Proposition 3.1: The pair satisfies $(\tilde{\pi}_0, \tilde{\pi}_1) \in \Gamma(\mathbb{P}_{obs})$, so $\tilde{\pi}_w(du, dx) = du \mathbb{P}_{obs}(dx)$ and its conditional outcome kernels satisfy the observational constraints in (3.2); these kernels are defined up to $(du \otimes \mathbb{P}_{obs}(dx))$ -null sets.

The content of Proposition 3.1 is twofold. First, the CCOT constraints act only on the two marginals $\tilde{\pi}_0$ and $\tilde{\pi}_1$: any coupling of these marginals—including every joint law of $(Y(0), Y(1))$ given (U, X) —is observationally equivalent. Second, the converse makes the characterization sharp: no joint law outside this set can be generated by a structural model matching $\mathbb{P}_{obs}$. Consequently, the identified set for the PRTE target is $\{\theta_\omega(\pi_0, \pi_1) : (\pi_0, \pi_1) \in \Gamma(\mathbb{P}_{obs})\}$, and the maximum and minimum of $\theta_\omega(\pi_0, \pi_1)$ over $\Gamma(\mathbb{P}_{obs})$ are the sharp upper and lower bounds on the PRTE. The proof is given in Section A.

Remark 3.2. The fact that the moment relaxation approach of Mogstad et al. [2018] can yield arbitrarily loose bounds compared to our CCOT formulation is well-recognized; see also Han and Yang [2024]. A detailed comparison example is given in Example O2.2 of the Online Appendix.

Remark 3.3 (Beyond PRTE: CCOT for general joint functionals). The CCOT perspective is not specific to the PRTE. At a conceptual level, one could consider a bounded measurable cost $f : \mathcal{Y}^2 \rightarrow \mathbb{R}$ and the joint target $\theta_f := \mathbb{E}[f(Y(1), Y(0))]$ by replacing the PRTE objective in (3.2) with $\mathbb{E}_\pi[f(Y(1), Y(0))]$ and enlarging the decision variable from the pair (π_0, π_1) to a joint conditional kernel $\pi(dy_1, dy_0 \mid u, x)$ whose marginals satisfy the same observational constraints. A sharp characterization should again follow from the lifting logic behind Proposition 3.1. What is special about the PRTE is the separable product cost $y \cdot \omega(x, u)$: it allows the CCOT problem to collapse into one-dimensional OT problems for each potential outcome separately. For a general joint functional, that separability is lost, and the resulting interval-specific problems are instead OT problems over the coupling between the identified marginals of $Y(1)$ and $Y(0)$, with no comparable closed-form reduction in general. Since the present paper is about the PRTE case, we do not pursue this extension here.

Next, we show how to obtain a closed-form solution to the CCOT problem (3.2). Notice that in (3.2), the objective is the difference of two linear functionals of π_0 and π_1 , and the constraints on π_0 and π_1 are separate. Therefore, we only need to consider one side of the problem, as the other side follows identically. In what follows, we only consider the $W = 1$ side minimization problem, i.e.,

$$\begin{aligned}
& \min_{\pi_1} \quad \mathbb{E}_{\pi_1}[Y(1)\omega(X, U)] \\
& \text{subject to} \quad \int_0^{p(z,x)} \pi_1(dy \mid u, x) du = H_{1,x}(p(z, x), dy), \quad \mathbb{P}_{Z,X}\text{-a.e. } (z, x) \\
& \quad \pi_1(dx, du) = du \mathbb{P}_{obs}(dx), \\
& \quad \pi_1 \text{ is supported on } \mathcal{Y} \times [0, 1] \times \mathcal{X}.
\end{aligned} \tag{3.3}$$

The solution for the maximization problem can be derived similarly. Writing $\theta_\omega = \theta_{\omega,1} - \theta_{\omega,0}$, let $[\underline{\theta}_{\omega,w}, \bar{\theta}_{\omega,w}]$ denote the sharp identified interval for the w -th component. Because the feasible sets for π_0 and π_1 are Cartesian products, the sharp identified interval for the full target is

$$[\underline{\theta}_\omega, \bar{\theta}_\omega] = [\underline{\theta}_{\omega,1} - \bar{\theta}_{\omega,0}, \bar{\theta}_{\omega,1} - \underline{\theta}_{\omega,0}].$$

Accordingly, we state the main derivations for $\theta_{\omega,1}$; the formulas for $\theta_{\omega,0}$ are obtained by the same arguments with the untreated analogues.

3.2 Closed-Form Bounds without Covariates

To better illustrate the idea, we first derive the results without covariates and generalize them in the next subsection. Therefore, in this subsection, ω and p do not depend on x .

3.2.1 Discrete Instrument Setting

We first consider the setting where there are no covariates and the instrument Z takes values in a finite discrete set $\mathcal{Z}$. Let the unique positive values of the propensity score be ordered as

$$\{p(z) : z \in \mathcal{Z}, p(z) > 0\} = \{p_1, \dots, p_K\},$$

with the conventions $p_0 := 0$ and $p_{K+1} := 1$, such that $0 = p_0 < p_1 < p_2 < \dots < p_K \leq p_{K+1} = 1$. Any instrument value with $p(z) = 0$ is represented by the boundary level p_0 and contributes only the zero treated measure. We allow distinct instrument values to map to the same propensity value. This partitions the unit interval of the latent variable U into the

K nonempty identified subintervals $I_i = (p_i, p_{i+1}]$, $i = 0, \dots, K-1$, and the residual interval $I_K = (p_K, 1]$, which is empty when $p_K = 1$.

Remark (notation). p_i denotes an ordered unique positive baseline propensity level for $1 \leq i \leq K$, with structural integration boundaries $p_0 = 0$ and $p_{K+1} = 1$; a structural boundary need not be an attained positive propensity level. In Theorem 3.4: The levels satisfy $0 = p_0 < p_1 < \dots < p_K \leq p_{K+1} = 1$ and are the unique positive values of $p(z)$; $p_K = 1$ makes the residual interval empty. In Proposition 4.3, Proposition 4.4, Lemma D.1, Theorem 4.7: The levels are functions $p_i(x)$ of the covariates, with $p_0(x) = 0$ and $p_{K+1}(x) = 1$. The positive level count, equality classes S_i , and their strict ordering are invariant in x under the maintained separation condition, although the numerical levels may vary with x ; p_i without a hat denotes the population level. $I_i = (p_i, p_{i+1}]$ for $i = 0, \dots, K-1$, and $I_K = (p_K, 1]$; the first K intervals are nonempty and the residual interval is empty when $p_K = 1$.

Because the objective function (3.3) is an integral over U , it decomposes additively across these subintervals into one conditional-mean term per I_i plus a residual integral over $(p_K, 1]$; the decomposed objective is displayed in (3.6) below.

The fundamental advantage of the discrete setting is that the global observational constraints can be differenced, perfectly isolating the conditional distribution of $Y(1)$ within each interval I_i . Recall that from (3.3),

$$\int_0^{p_i} \pi_1(dy \mid u) du = \mathbb{P}_{obs}(dy, W = 1 \mid p(Z) = p_i), \quad i = 1, \dots, K. \quad (3.4)$$

The key step is a staircase reduction of the cumulative constraints in (3.4). Subtracting adjacent constraints converts the cumulative restrictions into interval-wise marginal restrictions, which in turn converts (3.3) into separate standard one-dimensional OT problems. For $i = 1, \dots, K-1$, subtracting the i -th constraint from the $(i+1)$ -th in (3.4) isolates the integral over $(p_i, p_{i+1}]$; for $i = 0$, the first constraint directly governs $(0, p_1]$. Since cumulative sums recover the original constraints, this invertible transformation shows that the system (3.4) is equivalent to constraining the average of $\pi_1(dy \mid u)$ over each I_i , $i = 0, \dots, K-1$, to equal the identified measure

$$\mu_{1,i}(dy) = \begin{cases} \frac{\mathbb{P}_{obs}(dy, W=1 \mid p(Z)=p_1)}{p_1}, & \text{if } i = 0, \\ \frac{\mathbb{P}_{obs}(dy, W=1 \mid p(Z)=p_{i+1}) - \mathbb{P}_{obs}(dy, W=1 \mid p(Z)=p_i)}{p_{i+1} - p_i}, & \text{if } 0 < i < K. \end{cases} \quad (3.5)$$

Remark (notation). $\mu_{1,i}$ is the identified conditional probability law of $Y(1)$ given $U \in I_i$, defined for $i = 0, \dots, K-1$.

When the final interval $I_K = (p_K, 1]$ is nonempty, the potential outcome $Y(1)$ is never observed there because no individual in the population has a propensity score high enough to induce treatment when $U > p_K$. By (3.4), since π_1 is nonnegative, $\mathbb{P}_{obs}(dy, W = 1 \mid p(Z) = p)$ is nondecreasing in p . Moreover, $\mu_{1,i}(\mathcal{Y}) = 1$ since $\mathbb{P}_{obs}(W = 1 \mid p(Z) = p) = p$. Therefore, the $\mu_{1,i}$ are valid distributions. Mourifié and Wan [2017] provide a simple method to test the monotonicity of $\mathbb{P}_{obs}(dy, W = 1 \mid p(Z) = p)$.

For $0 < i < K$ and any measurable set A , (3.5) gives

$$\begin{aligned}\mu_{1,i}(A) &= \frac{\mathbb{P}(Y(1) \in A, U \in (0, p_{i+1}] \mid p(Z) = p_{i+1}) - \mathbb{P}(Y(1) \in A, U \in (0, p_i] \mid p(Z) = p_i)}{p_{i+1} - p_i} \\ &= \mathbb{P}(Y(1) \in A \mid U \in (p_i, p_{i+1}]),\end{aligned}$$

where we use ?? 2.1(3) and consistency in the first equality and ?? 2.1(2) in the second. Therefore, $\mu_{1,i}$ can be interpreted as the distribution of the potential outcome under treatment for the i -th complier group. The i -th complier group consists of units with latent resistance $U \in (p_i, p_{i+1}]$: they refuse treatment at propensity values weakly below p_i and take treatment once the propensity exceeds their latent threshold. In particular, when the instrument is binary, this recovers the identification results of Imbens and Rubin [1997].

Combining the additive decomposition of the objective with the interval-wise constraints just derived, the original problem (3.3) can be equivalently written as

$$\min_{\pi_1} \sum_{i=0}^{K-1} (p_{i+1} - p_i) \mathbb{E}_{\pi_1}[Y(1)\omega(U) \mid U \in I_i] + \int_{p_K}^1 \int_{\mathcal{Y}} y\omega(u) \pi_1(dy \mid u) du \quad (3.6)$$

subject to the K independent marginal constraints

$$\begin{aligned}\frac{1}{p_{i+1} - p_i} \int_{p_i}^{p_{i+1}} \pi_1(dy \mid u) du &= \mu_{1,i}(dy), \quad i = 0, \dots, K-1, \\ \pi_1(du) &= du,\end{aligned} \quad (3.7)$$

with the support constraint $\text{supp}(\pi_1(\cdot \mid u)) \subseteq \mathcal{Y}$ for all $u \in [0, 1]$, and no further distributional constraint on $\pi_1(\cdot \mid u)$ for $u \in I_K = (p_K, 1]$. Crucially, each constraint in (3.7) involves π_1 only on the single interval I_i , and the objective (3.6) is additively separable across these intervals. Therefore, the global problem decomposes into K constrained OT subproblems and one unconstrained residual problem, the latter contributing zero when $p_K = 1$.

Specifically, for each interval $i = 0, \dots, K-1$, the data pin down the marginal distribution of $Y(1)$ as $\mu_{1,i}$ and the marginal distribution of U as uniform on $(p_i, p_{i+1}]$. Let ν_i denote the uniform probability measure $\text{Unif}(p_i, p_{i+1})$. The set of valid joint distributions for $(Y(1), U)$

conditional on $U \in I_i$ is exactly the set of all couplings $\Pi(\mu_{1,i}, \nu_i)$. Therefore, for each of these K intervals, we solve a standard 1D optimal transport problem:

$$\min_{\gamma_i \in \Pi(\mu_{1,i}, \nu_i)} \int_{\mathcal{Y} \times I_i} y \cdot \omega(u) d\gamma_i(y, u). \quad (3.8)$$

By contrast, when the final interval $I_K = (p_K, 1]$ is nonempty, the propensity score never reaches a value high enough to induce treatment there, so the data impose no constraint whatsoever on the distribution of $Y(1)$ for this subpopulation. Without distributional constraints, the optimal transport problem degenerates into a trivial pointwise minimization. Because $y_{\min}$ and $y_{\max}$ are the extrema of the compact outcome domain, for each $u \in I_K$ the optimal solution assigns all probability mass to $y_{\min}$ if $\omega(u) \geq 0$, or to $y_{\max}$ if $\omega(u) < 0$:

$$\min_{\pi_1} \int_{p_K}^1 \int_{\mathcal{Y}} y \omega(u) \pi_1(dy | u) du = \int_{p_K}^1 \left(y_{\min} \max\{0, \omega(u)\} + y_{\max} \min\{0, \omega(u)\} \right) du. \quad (3.9)$$

When $p_K = 1$, this is an integral over the empty interval and equals zero. Summing the optimal values of the K constrained subproblems in (3.8) scaled by their respective interval lengths $(p_{i+1} - p_i)$ and adding the trivial unconstrained bound from (3.9), we obtain the global minimum. Applying Theorem 2.3 to each subproblem in (3.8) yields the countermonotonic coupling, which leads directly to the following closed-form sharp bounds.

Theorem 3.4 (Closed-Form Sharp Bounds). *Suppose ?? 2.1 and ?? 2.2 hold, X is almost surely constant, and $\mathcal{Z}$ is finite. For $i = 0, \dots, K-1$, let $Q_{Y,i}(t)$ be the quantile function of the identified conditional distribution $\mu_{1,i}$ defined in (3.5). Let $Q_{\omega,i}(t)$ be the quantile function of the random variable $\omega(U)$ where $U \sim \text{Unif}(p_i, p_{i+1})$. The sharp lower bound for the $W = 1$ component of the PRTE is given by:*

$$\begin{aligned} \underline{\theta}_{\omega,1} = & \sum_{i=0}^{K-1} (p_{i+1} - p_i) \int_0^1 Q_{Y,i}(t) Q_{\omega,i}(1-t) dt \\ & + \int_{p_K}^1 \left(y_{\min} \max\{0, \omega(u)\} + y_{\max} \min\{0, \omega(u)\} \right) du. \end{aligned} \quad (3.10)$$

Similarly, the sharp upper bound is given by replacing $Q_{\omega,i}(1-t)$ with $Q_{\omega,i}(t)$ in the OT term and swapping $y_{\min}$ and $y_{\max}$ in the trivial tail term.

Remark. For $i = 0, \dots, K-1$ and $0 < t < 1$, $Q_{Y,i}(t) = \inf\{y \in \mathbb{R} : \mu_{1,i}(\mathcal{Y} \cap (-\infty, y]) \geq t\}$ and $Q_{\omega,i}(t) = \inf\{v \in \mathbb{R} : \mathbb{P}(\omega(U) \leq v) \geq t\}$, where $U \sim \text{Unif}(p_i, p_{i+1})$.

Remark (notation). $Q_{Y,i} : (0, 1) \rightarrow \mathcal{Y}$ is the quantile function of the identified conditional law $\mu_{1,i}$ of $Y(1)$ given $U \in (p_i, p_{i+1}]$, for $i = 0, \dots, K-1$. $Q_{\omega,i}$ is the quantile function of the law

of $\omega(U)$ under $U \sim \text{Unif}(p_i, p_{i+1})$, for $t \in (0, 1)$ and $i = 0, \dots, K-1$, using the nonempty identified baseline intervals.

We defer the formal proof of this theorem to that of Theorem 3.8 in Section A.2, since the present theorem is the constant- X special case with every LIV interval degenerate. Theorem 3.4 provides a closed-form, analytically tractable expression for the sharp bounds, bypassing any need for linear programming and requiring only estimation of the conditional quantile functions from the observed data.

As a special case, setting $\omega(u) = 1$ makes every $Q_{\omega,i} \equiv 1$, so the OT term telescopes into the observed treated mean at p_K and

$$\underline{\theta}_{\omega,1} = p_K \mathbb{E}[Y \mid W = 1, p(Z) = p_K] + y_{\min}(1 - p_K),$$

the classical bound of Manski [1990] for $\mathbb{E}[Y(1)]$, which is known to be tight [Heckman and Vytlacil, 2001]. Figure 1 illustrates the three identification regions that compose the bound in Theorem 3.4; Example O2.1 of the Online Appendix works through a concrete binary-instrument example with a constant propensity score, deriving the explicit three-part decomposition.


By Theorem 3.4, the width of the identified interval is

$$\begin{aligned} \bar{\theta}_{\omega,1} - \underline{\theta}_{\omega,1} &= \sum_{i=0}^{K-1} (p_{i+1} - p_i) \int_0^1 (Q_{Y,i}(t) - Q_{Y,i}(1-t)) Q_{\omega,i}(t) dt \\ &\quad + (y_{\max} - y_{\min}) \int_{p_K}^1 |\omega(u)| du, \end{aligned}$$

which depends on the range of the propensity function and on the distributions $\mu_{1,i}$. If p_K is close to 1 and the support of each $\mu_{1,i}$ is small, the width is also small. In particular, if each $Q_{Y,i}$ is constant and $p_K = 1$, then $\theta_{\omega,1}$ is point-identified from the data.

Remark 3.5 (Decomposition into 1D marginal constraints and Marx [2024]). Marx [2024] raises the open question of characterizing the distribution of treatment effects for the complier group. Although Proposition 3.1 states sharpness at the joint level of $(Y(0), Y(1), U, X)$, the CCOT constraints defining $\Gamma(\mathbb{P}_{obs})$ reduce to independent one-dimensional marginal constraints on the conditional kernels $\pi_w(dy \mid u, x)$, so once these marginals are pinned down the remaining freedom lies entirely in the coupling between $Y(1)$ and $Y(0)$ given (U, X) . In the binary-instrument, no-covariate setting with $p_1 < p_2$, the complier interval is $I_c = (p_1, p_2]$ and (3.5) identifies the treated-outcome distribution for compliers as $\mu_{1,1}$; differencing the untreated component, which is nonincreasing in p , identifies

$$\mu_{0,1}(dy) := \frac{\mathbb{P}_{obs}(dy, W = 0 \mid p(Z) = p_1) - \mathbb{P}_{obs}(dy, W = 0 \mid p(Z) = p_2)}{p_2 - p_1}.$$



The identified set for the complier treatment-effect distribution is therefore exactly $\mathcal{I}_{\Delta,c} = \{(y_1 - y_0)_\# \gamma : \gamma \in \Pi(\mu_{1,1}, \mu_{0,1})\}$, so Marx's question becomes: once the identified complier marginals are fixed, characterize all pushforwards of their couplings under $(y_1, y_0) \mapsto y_1 - y_0$. Characterizing $\mathcal{I}_{\Delta,c}$ therefore reduces to determining the treatment-effect distributions generated by these couplings.

3.2.2 General Instrument Setting

We now consider a general instrument setting under a single regularity condition on the instrument domain and the propensity score. This framework nests the finite discrete case (each component a singleton), the connected continuous case studied in the local instrumental variables (LIV) literature [Heckman and Vytlacil, 1999], and mixtures of the two.

Assumption 3.6 (Finite-Component Propensity Continuity). The instrument domain $\mathcal{Z}$ has finitely many connected components. For almost every $x \in \mathcal{X}$, the conditional propensity score function $z \mapsto p(z, x) = \mathbb{P}(W = 1 \mid Z = z, X = x)$ is continuous on each connected component.

Remark. The continuity condition holds for $\mathbb{P}_X$ -almost every $x \in \mathcal{X}$.

Under the standing compactness of $\mathcal{Z}$ and Assumption 3.6, the image of each connected component under $p(\cdot)$ is a closed interval (possibly a singleton if p is constant on the component or if the component is a single point). Merging overlaps and ordering from left to right, the range of the propensity score admits a maximal-interval decomposition

$$p(\mathcal{Z}) = \bigcup_{k=1}^K [\underline{p}_k, \bar{p}_k], \quad 0 \leq \underline{p}_1 \leq \bar{p}_1 < \underline{p}_2 \leq \bar{p}_2 < \cdots < \underline{p}_K \leq \bar{p}_K \leq 1, \quad (3.11)$$

where we explicitly allow degenerate intervals $\underline{p}_k = \bar{p}_k$ corresponding to isolated propensity values, as in the finite discrete case. Adopting the boundary conventions $\bar{p}_0 := 0$ and $\underline{p}_{K+1} := 1$, this decomposition partitions the unit interval of the latent variable U into three types of regions: the LIV intervals $L_k := [\underline{p}_k, \bar{p}_k]$ for $k = 1, \dots, K$, the OT gaps $G_k := (\bar{p}_k, \underline{p}_{k+1})$ for $k = 0, \dots, K - 1$, and the trivial tail $(\bar{p}_K, 1]$.

Remark (notation). $\underline{p}_k$ denotes the left endpoint of the k -th maximal closed interval in the propensity range ordered from left to right, allowing singleton intervals. In Theorem 3.7: The endpoints are indexed by $k = 1, \dots, K$ in $p(\mathcal{Z}) = \bigcup_{k=1}^K [\underline{p}_k, \bar{p}_k]$, with the additional boundary convention $\underline{p}_{K+1} = 1$. In Theorem 3.8: For almost every x , the endpoint is $\underline{p}_k(x)$, indexed by $k = 1, \dots, K_x$ in the conditional range $p(\mathcal{Z}, x)$, with $\underline{p}_{K_x+1}(x) = 1$; both endpoints and the

component count may vary with x . $\bar{p}_k$ denotes the right endpoint of the k -th maximal closed interval in the propensity range ordered from left to right, allowing singleton intervals and using zero as the auxiliary initial endpoint. In Theorem 3.7: The endpoints are indexed by $k = 1, \dots, K$ in $p(\mathcal{Z}) = \bigcup_{k=1}^K [\underline{p}_k, \bar{p}_k]$, with $\bar{p}_0 = 0$. In Theorem 3.8: For almost every x , the endpoint is $\bar{p}_k(x)$, indexed by $k = 1, \dots, K_x$ in $p(\mathcal{Z}, x)$, with $\bar{p}_0(x) = 0$; endpoints and the component count may vary with x . In Proposition 4.3, Proposition 4.4: Each conditional range component is a singleton: $\underline{p}_k(x) = \bar{p}_k(x) = p_k(x)$ for $k = 1, \dots, K$, with the structural initial endpoint $p_0(x) = \bar{p}_0(x) = 0$. L_k denotes the k -th closed maximal interval of the propensity range, ordered from left to right, including singleton intervals. In Theorem 3.7: $L_k = [\underline{p}_k, \bar{p}_k]$ for $k = 1, \dots, K$; singleton intervals contribute a Lebesgue-null integral. In Theorem 3.8: For almost every x , $L_k(x) = [\underline{p}_k(x), \bar{p}_k(x)]$ for $k = 1, \dots, K_x$; the intervals are those of $p(\mathcal{Z}, x)$ and may be singletons. G_k denotes the open interval between consecutive ordered propensity-range components, with the initial gap starting at zero and with the final unconstrained tail excluded from the gap indices. In Theorem 3.7: $G_k = (\bar{p}_k, \underline{p}_{k+1})$ for $k = 0, \dots, K - 1$, with $\bar{p}_0 = 0$; empty gaps are omitted from gap-specific definitions and contribute zero. In Theorem 3.8: $G_k(x) = (\bar{p}_k(x), \underline{p}_{k+1}(x))$ for $k = 0, \dots, K_x - 1$ and almost every x , with $\bar{p}_0(x) = 0$; empty conditional gaps contribute zero. In Proposition 4.3, Proposition 4.4: The propensity-range components are singletons, so $G_k(x) = (p_k(x), p_{k+1}(x))$ for $k = 0, \dots, K - 1$, with $p_0(x) = 0$; the interval beyond $p_K(x)$ is the separate trivial tail. $\underline{p}(x) = \inf_{z \in \mathcal{Z}} p(z, x)$ is the population lower propensity endpoint at covariate value x , with values in $[0, 1]$ and with zero allowed. In Proposition 4.8, Theorem 4.10, Corollary O5.1: $\underline{p}(x)$ is the lower endpoint of the connected conditional support $[\underline{p}(x), \bar{p}(x)]$ of $P \mid X = x$. $\bar{p}(x) = \sup_{z \in \mathcal{Z}} p(z, x)$ is the population upper propensity endpoint at covariate value x , with values in $[0, 1]$. In Proposition 4.8, Theorem 4.10, Corollary O5.1: $\bar{p}(x)$ is the upper endpoint of the connected conditional support $[\underline{p}(x), \bar{p}(x)]$ of $P \mid X = x$. $\bar{p}_2(z)$ is the sample-dependent average of the clipped first-stage fit $\hat{p}(z, \cdot)$ trained on I_1 over the covariates indexed by I_2 ; conditional on I_1 it remains an empirical average over evaluation covariates, rather than the population marginal mean.

The CCOT problem for the treated ($W = 1$) component becomes

$$\begin{aligned}
& \min_{\pi_1} \quad \mathbb{E}_{\pi_1}[Y(1)\omega(U)] \\
& \text{subject to} \quad \int_0^p \pi_1(dy \mid u) du = H_1(p, dy), \quad \forall p \in p(\mathcal{Z}), \\
& \quad \pi_1(du) = du.
\end{aligned} \tag{3.12}$$

Following the same strategy as in Section 3.2.1, we rewrite (3.12) into a manifestly separable

form across the three region types. The objective splits additively as

$$\begin{aligned}\mathbb{E}_{\pi_1}[Y(1)\omega(U)] &= \sum_{k=0}^{K-1} \int_{G_k} \int_{\mathcal{Y}} y \omega(u) \pi_1(dy | u) du + \sum_{k=1}^K \int_{L_k} \int_{\mathcal{Y}} y \omega(u) \pi_1(dy | u) du \\ &\quad + \int_{\bar{p}_K}^1 \int_{\mathcal{Y}} y \omega(u) \pi_1(dy | u) du.\end{aligned}$$

On each gap G_k , the same telescoping argument that produced (3.5)—differencing the observational constraint in (3.12) at $p = \bar{p}_k$ and $p = \underline{p}_{k+1}$ —fixes the marginal of $Y(1)$ on G_k to be

$$\mu_{1,k}(dy) := \frac{H_1(\underline{p}_{k+1}, dy) - H_1(\bar{p}_k, dy)}{\underline{p}_{k+1} - \bar{p}_k}, \quad k : \bar{p}_k < \underline{p}_{k+1}, \quad (3.13)$$

where $H_1(0, \cdot) = 0$ by definition. Thus the $k = 0$ formula applies when $\underline{p}_1 > 0$; when $\underline{p}_1 = 0$, G_0 is empty and has zero contribution. More generally, empty gaps are omitted from all gap-specific definitions and sums below. Let $g_1^c(p) := \int_{\mathcal{Y}} y H_1(p, dy)$, and fix the countable measure-determining class $\mathcal{A} := \{\mathcal{Y} \cap (-\infty, r] : r \in \mathbb{Q}\} \cup \{\mathcal{Y}\}$. On each non-degenerate LIV interval L_k , the constraint in (3.12) holds for every $p \in [\underline{p}_k, \bar{p}_k]$. For each $A \in \mathcal{A}$, the scalar map $p \mapsto H_1(p, A)$ is absolutely continuous on L_k and the ordinary fundamental theorem of calculus gives

$$\pi_1(A | u) = \left. \frac{\partial}{\partial p} \right|_{p=u} H_1(p, A), \quad A \in \mathcal{A}, \quad u \in L_k \text{ a.e.} \quad (3.14)$$

Countability of $\mathcal{A}$ gives one common null set, so (3.14) identifies $\pi_1(\cdot | u)$ almost everywhere on L_k . Applying the same scalar argument to the bounded function y gives $\partial_u g_1^c(u) = \int_{\mathcal{Y}} y \pi_1(dy | u)$ almost everywhere. The trivial tail $(\bar{p}_K, 1]$ carries no constraint. Because cumulative sums of the gap marginals plus integration of the identified LIV kernels recover the original constraint at every $p \in p(\mathcal{Z})$, this transformation is invertible and the two constraint systems are equivalent. Each rewritten constraint involves π_1 only on its own region and the objective is additively separable, so the global problem decomposes into three types of subproblems, one for each region type.

Remark (notation). $\mu_{1,k}$ is the identified marginal law $\mathcal{L}(Y(1) | U \in G_k)$ on each nonempty gap G_k , $k \in \{0, \dots, K-1\}$, obtained from the canonical treated submeasures, including the structural zero submeasure at the initial boundary. $g_1^c(p, x) = \int_{\mathcal{Y}} y H_{1,x}(p, dy)$ uses the canonical population observed submeasure on the conditional propensity support, including its endpoints; it is absolutely continuous on each non-degenerate LIV interval, and its derivatives mean intervalwise Lebesgue-a.e. derivatives. In Theorem 3.7: With no covariates, $g_1^c(p) = \int_{\mathcal{Y}} y H_1(p, dy)$, and derivatives are taken on the non-degenerate intervals L_k . In

Proposition 4.8, Theorem 4.10: $g_1 = g_1^c$ on the connected conditional support $[\underline{p}(x), \bar{p}(x)]$. $g_1 = g_1^c$ denotes the canonical population treated regression obtained by integrating y against the canonical observed treated submeasure, identified on the conditional propensity support including its endpoints, with intervalwise Lebesgue-a.e. derivatives on non-degenerate LIV intervals. In Theorem 3.7: With no covariates, $g_1(u) = \int_{\mathcal{Y}} y H_1(u, dy)$; derivatives are taken separately on the non-degenerate intervals L_k . In Proposition 4.8, Theorem 4.10, Corollary O5.1: The identified conditional support is $[\underline{p}(x), \bar{p}(x)]$, and target evaluations use $q^{\text{tr}}(Z, X)$ and P within this support.

1. **The OT gaps** G_k , $k = 0, \dots, K - 1$: the $Y(1)$ -marginal is fixed to $\mu_{1,k}$ of (3.13) and the U -marginal is $\text{Unif}(G_k)$, so the contribution of G_k is a standard 1D OT problem, solved by Theorem 2.3. Empty gaps contribute zero.
2. **The LIV intervals** L_k , $k = 1, \dots, K$: for non-degenerate L_k , the LIV identity (3.14) point-identifies $\mathbb{E}[Y(1) \mid U = u] = \partial_u g_1^c(u)$ for a.e. $u \in L_k$, so the contribution is uniquely determined. Degenerate L_k contribute a Lebesgue-null integral and drop out.
3. **The trivial tail** $(\bar{p}_K, 1]$: $Y(1)$ is never observed and the data impose no constraint, so the minimization over the compact outcome domain assigns all mass to $y_{\min}$ where $\omega(u) \geq 0$ and to $y_{\max}$ where $\omega(u) < 0$.

Combining the solutions across these regions yields the closed-form sharp bounds for the general instrument setting, which are illustrated in Figure 6 of the Appendix.

Theorem 3.7 (General Closed-Form Sharp Bounds). *Suppose ?? 2.1, ?? 2.2, and ?? 3.6 hold, and X is almost surely constant. Let $p(\mathcal{Z}) = \bigcup_{k=1}^K [\underline{p}_k, \bar{p}_k]$ be the maximal-interval decomposition in (3.11). For each nonempty gap G_k , let $Q_{Y,k}(t)$ be the quantile function of $\mu_{1,k}$ defined in (3.13), and $Q_{\omega,k}(t)$ the quantile function of $\omega(U)$ for $U \sim \text{Unif}(G_k)$. In the displayed sums, a term indexed by an empty gap is defined to be zero. The sharp lower bound for the $W = 1$ component of the target parameter is given by*

$$\begin{aligned} \theta_{\omega,1} = & \underbrace{\sum_{k=0}^{K-1} (\underline{p}_{k+1} - \bar{p}_k) \int_0^1 Q_{Y,k}(t) Q_{\omega,k}(1-t) dt}_{\text{OT Bound on the Gaps}} + \underbrace{\sum_{k=1}^K \int_{\underline{p}_k}^{\bar{p}_k} \frac{d}{du} g_1^c(u) \omega(u) du}_{\text{Point-Identified LIV on the Intervals}} \\ & + \underbrace{\int_{\bar{p}_K}^1 \left(y_{\min} \max\{0, \omega(u)\} + y_{\max} \min\{0, \omega(u)\} \right) du}_{\text{Trivial Lower Bound on the Tail}}. \end{aligned} \quad (3.15)$$

Similarly, the sharp upper bound is given by replacing $Q_{\omega,k}(1-t)$ with $Q_{\omega,k}(t)$ in the OT term and swapping $y_{\min}$ and $y_{\max}$ in the trivial tail term.

Remark. The partition assertion means $\int_{[0,1] \setminus (\cup_{k=1}^K L_k \cup \cup_{k=0}^{K-1} G_k \cup (\bar{p}_K, 1])} du = 0$. For each probability law μ used here, $Q_\mu(t) = \inf\{a \in \mathbb{R} : \mu((-\infty, a]) \geq t\}$ for $t \in (0, 1)$, with the displayed quantile integrals interpreted as Lebesgue integrals over $(0, 1)$.

Remark (notation). $Q_{Y,k}$ is the quantile function of $\mu_{1,k}$ for $t \in (0, 1)$ and $k \in \{0, \dots, K-1\}$ such that $\bar{p}_k < \underline{p}_{k+1}$; an empty gap has zero contribution and requires no quantile function. $Q_{\omega,k}$ is the quantile function of the law of $\omega(U)$ under $U \sim \text{Unif}(G_k)$, defined for $t \in (0, 1)$ and $k \in \{0, \dots, K-1\}$ such that $\bar{p}_k < \underline{p}_{k+1}$; an empty gap has zero contribution and requires no quantile function.

The proof combines the subtraction-of-constraints decomposition from Theorem 3.4 (applied region-by-region to identify $\mu_{1,k}$ on each gap) with the LIV-derivative argument on each non-degenerate interval L_k , followed by the 1D OT coupling of Theorem 2.3. We defer the formal proof to that of Theorem 3.8 in Section A.2, which restates the result in the unified form of the covariate subsection below. Theorem 3.7 subsumes both classical settings as limit cases. When every LIV interval is degenerate ($\underline{p}_k = \bar{p}_k$ for all k), the LIV sum vanishes identically, each $\mu_{1,k}$ coincides with the discrete complier distribution in (3.5), and (3.15) reduces exactly to Theorem 3.4. Conversely, when $K = 1$ with $\underline{p}_1 < \bar{p}_1$, the gap sum contains a single OT term on $(0, \underline{p}_1)$, the LIV sum contributes a single integral on $[\underline{p}_1, \bar{p}_1]$, and we recover the connected-continuous case; in particular, we match the identification results of Heckman and Vytlacil [1999] when additionally $\underline{p}_1 = 0$ and $\bar{p}_1 = 1$, so the propensity score has full support with no gaps. The width of the identified interval depends on the gap widths, the identified gap distributions $\mu_{1,k}$, and the tail beyond $\bar{p}_K$:

$$\begin{aligned} \bar{\theta}_{\omega,1} - \underline{\theta}_{\omega,1} &= \sum_{k=0}^{K-1} (\underline{p}_{k+1} - \bar{p}_k) \int_0^1 \left(Q_{Y,k}(t) - Q_{Y,k}(1-t) \right) Q_{\omega,k}(t) dt \\ &\quad + (y_{\max} - y_{\min}) \int_{\bar{p}_K}^1 |\omega(u)| du. \end{aligned}$$

3.3 Closed-Form Bounds with Covariates

Generalizing the previous results to covariates is mathematically straightforward. Because ?? 2.1 imposes $Z \perp\!\!\!\perp (Y(w), U) \mid X$, the global CCOT problem disintegrates into a family of conditional one-dimensional optimal transport problems indexed by $x \in \mathcal{X}$: we solve conditional on $X = x$ and aggregate the conditional bounds over the marginal distribution of

X by iterated expectations.

Under ?? 3.6, for $\mathbb{P}_{obs}$ -a.e. $x \in \mathcal{X}$, the conditional range of the propensity score admits a maximal-interval decomposition

$$p(\mathcal{Z}, x) = \bigcup_{k=1}^{K_x} [\underline{p}_k(x), \bar{p}_k(x)], \quad 0 \leq \underline{p}_1(x) \leq \bar{p}_1(x) < \cdots < \underline{p}_{K_x}(x) \leq \bar{p}_{K_x}(x) \leq 1, \quad (3.16)$$

with degenerate intervals $\underline{p}_k(x) = \bar{p}_k(x)$ corresponding to isolated conditional propensity values and boundary conventions $\bar{p}_0(x) := 0$ and $\underline{p}_{K_x+1}(x) := 1$. For each x , this partitions $[0, 1]$ into the conditional LIV intervals $L_k(x) := [\underline{p}_k(x), \bar{p}_k(x)]$, the OT gaps $G_k(x) := (\bar{p}_k(x), \underline{p}_{k+1}(x))$, and the trivial tail $(\bar{p}_{K_x}(x), 1]$. The conditional gap-identified measure of $Y(1)$ on $G_k(x)$ is

$$\mu_{1,k|x}(dy) := \frac{H_{1,x}(\underline{p}_{k+1}(x), dy) - H_{1,x}(\bar{p}_k(x), dy)}{\underline{p}_{k+1}(x) - \bar{p}_k(x)}, \quad (3.17)$$

for indices k such that $\bar{p}_k(x) < \underline{p}_{k+1}(x)$; $H_{1,x}(0, \cdot) = 0$ by definition. Empty conditional gaps are omitted and contribute zero. Finally, let $g_1^c(p, x) := \int_{\mathcal{Y}} y H_{1,x}(p, dy)$. This canonical version is absolutely continuous in p on each non-degenerate LIV interval $L_k(x)$, and every derivative below means its intervalwise Lebesgue-a.e. derivative.

Remark (notation). For almost every x , K_x is the finite number of closed maximal intervals, including singleton intervals, obtained by merging overlapping component images in $p(\mathcal{Z}, x)$; its value may depend on x . In Proposition 4.8: The instrument domain is connected, so $K_x = 1$ and the sole conditional range interval is $[\underline{p}(x), \bar{p}(x)]$, possibly a singleton. $\mu_{1,k|x}$ is the population probability law of $Y(1)$ conditional on $X = x$ and $U \in G_k(x)$, defined for nonempty gaps with $k = 0, \dots, K_x - 1$; empty-gap terms are omitted and contribute zero. In Proposition 4.3, Proposition 4.4, Lemma D.1: With finite $\mathcal{Z}$, the conditional propensity-range intervals are singletons, $G_k(x) = (p_k(x), p_{k+1}(x))$ for $k = 0, \dots, K - 1$, and $\mu_{1,k|x}$ is the corresponding conditional complier outcome law.

Theorem 3.8 (Conditional General Bound). *Suppose ?? 2.1, ?? 2.2, and ?? 3.6 hold. For each nonempty conditional gap $G_k(x)$, let $Q_{Y,k|x}(t)$ be the quantile function of $\mu_{1,k|x}$ defined in (3.17), and $Q_{\omega,k|x}(t)$ the quantile function of $\omega(x, U)$ for $U \sim \text{Unif}(G_k(x))$. In the displayed sums, a term indexed by an empty conditional gap is defined to be zero. The sharp lower*

bound for the $W = 1$ component of the target parameter is

$$\begin{aligned} \theta_{\omega,1} = \mathbb{E}_X \bigg[& \sum_{k=0}^{K_X-1} (\underline{p}_{k+1}(X) - \bar{p}_k(X)) \int_0^1 Q_{Y,k|X}(t) Q_{\omega,k|X}(1-t) dt \\ & + \sum_{k=1}^{K_X} \int_{\underline{p}_k(X)}^{\bar{p}_k(X)} \frac{d}{du} g_1^c(u, X) \omega(X, u) du \\ & + \int_{\bar{p}_{K_X}(X)}^1 \left(y_{\min} \max\{0, \omega(X, u)\} + y_{\max} \min\{0, \omega(X, u)\} \right) du \bigg]. \end{aligned} \quad (3.18)$$

Similarly, the sharp upper bound is given by replacing $Q_{\omega,k|X}(1-t)$ with $Q_{\omega,k|X}(t)$ in the OT term and swapping $y_{\min}$ and $y_{\max}$ in the trivial tail term.

Additional condition. For every nonempty conditional gap $G_k(x)$, extend $Q_{Y,k|x}$ to $[0, 1]$ by $Q_{Y,k|x}(0) := y_{\min}$ and $Q_{Y,k|x}(1) := y_{\max}$.

Remark. For every nonempty conditional gap $G_k(x)$ and $0 < t < 1$, $(Q_{Y,k|x}(t), Q_{\omega,k|x}(t)) = (\inf\{y \in \mathbb{R} : \mu_{1,k|x}(\mathcal{Y} \cap (-\infty, y]) \geq t\}, \inf\{v \in \mathbb{R} : F_{\omega,k|x}(v) \geq t\})$. For $\mathbb{P}_X$ -almost every x , the intervals $L_k(x)$, $1 \leq k \leq K_x$, the gaps $G_k(x)$, $0 \leq k < K_x$, and the tail $(\bar{p}_{K_x}(x), 1]$ partition $[0, 1]$ up to Lebesgue-null endpoints. For $0 < t < 1$, use the lower generalized inverses $Q_{Y,k|x}(t) = \inf\{y \in \mathbb{R} : \mu_{1,k|x}(\mathcal{Y} \cap (-\infty, y]) \geq t\}$ and $Q_{\omega,k|x}(t) = \inf\{v \in \mathbb{R} : F_{\omega,k|x}(v) \geq t\}$. Interpret the kernel formula with any fixed jointly measurable $\mathcal{Y}$ -valued extension of $Q_{Y,k|x}$ to ranks 0 and 1; the asserted gap marginal and attaining objective hold for every such extension.

Remark (notation). $Q_{Y,k|x}$ is the population quantile function of $\mu_{1,k|x}$ for each nonempty conditional gap $G_k(x)$; it is required only on such gaps, and contributions from empty gaps are defined as zero. In Proposition 4.3, Proposition 4.4, Lemma D.1: In the finite-instrument setting, the gap is the baseline interval between $p_k(x)$ and $p_{k+1}(x)$, with $k = 0, \dots, K-1$, and $\mu_{1,k|x}$ is its identified conditional complier outcome law. $Q_{\omega,k|x}$ is defined for each nonempty conditional gap $G_k(x)$, $0 \leq k < K_x$, as the quantile function of the pushforward of $\text{Unif}(G_k(x))$ by $u \mapsto \omega(x, u)$; empty-gap contributions are zero without requiring a quantile function there. In Proposition 4.3: The conditional propensity-range intervals are singletons, so the relevant gap law is $\text{Unif}((p_k(x), p_{k+1}(x)))$ for $k = 0, \dots, K-1$, with $p_0(x) = 0$.

The proof of Theorem 3.8 is given in Section A.2. The two limit cases stated after Theorem 3.7 carry over verbatim under x -conditioning: degenerate conditional LIV intervals reduce (3.18) to the discrete covariate bound, and $K_X = 1$ with $\underline{p}_1(X) < \bar{p}_1(X)$ recovers the standard LIV covariate bound.

4 Estimation and Inference

In this section, we develop estimation and inference results for the closed-form bounds derived in Section 3. For the discrete instrument setting, we leverage DML [Chernozhukov et al., 2018] to accommodate high-dimensional covariates, constructing Neyman-orthogonal scores that yield $\sqrt{n}$ -consistent and asymptotically normal estimators. In the continuous instrument setting, we characterize the corresponding nonparametric convergence rates. Throughout, we observe an i.i.d. sample $\{O_i\}_{i=1}^n$ of $O = (X, Z, W, Y) \sim \mathbb{P}_{\text{obs}}$, and η collects the nuisance functions entering the bound, with $\hat{\eta}$ its estimate.

Remark (notation). η collects the nuisance functions entering the applicable bound and estimation formula, with its components determined by the instrument and outcome regime and $\hat{\eta}$ denoting their fitted counterparts. In Lemma D.1, Theorem 4.7: In part (i), $\eta(x) = \{p(z, x), p_k(x), \pi_k(x), \tau_{j,k}(x), a_{j,k}(z, x), \gamma_{\text{full},k}(x), \gamma_{j,k}(x), \gamma_K(x), \nu_{j,k}(x), M_{j,k}^\pm(x), J_{j,k}^\pm(x), J_{\text{full},k}^\pm(x)\}_{j,k}$; π_k and minus-side objects are included only for $k \geq 1$, and estimated quantiles have $j \geq 1$, with structural boundary values fixed. In part (ii), $\eta(x) = \{p(z, x), p_k(x), \pi_k(x), \tau_{j,k}(x), a_{j,k}(z, x), \gamma_{\text{full},k}(x), \gamma_{j,k}(x)\}$, with $p_0 = P_{1,0} = 0$ and no fitted π_0 ; within-set share directions satisfy $\sum_{z \in T_{j,k}} \delta a_{j,k}(z, x) = 0$. In Theorem 4.7: The unaccented η in $\sigma^2 = \mathbb{E}[(\psi(O; \eta) - \underline{\theta}_{\omega,1})^2]$ denotes η_0 , the true nuisance vector. In Theorem 4.10, Corollary O5.1: The nuisance functions are those entering the connected-instrument plug-in representation: p , the policy propensity and support endpoints derived from it, g_1 , and the treated boundary law or its quantile function. $\hat{\eta}$ is the vector of fitted nuisance functions entering the applicable bound estimator, with its components determined by the instrument and outcome regime. In Theorem 4.7: The first stage is fitted on I_1 and the remaining nuisances are fitted on I_1 given the recovered level sets, subject to the stated clipping conventions. Feasible recovery uses $\{X_i : i \in I_2\}$; in the oracle procedure the population level sets are used and $\hat{\eta}$ is fixed conditional on I_1 . In Theorem 4.10, Corollary O5.1: The fitted nuisances enter (4.8), including estimates of p , the derived policy and support endpoints, g_1 , and the treated boundary law or its quantile function; they are trained on $I_0 \cup I_1$ and held fixed conditional on these training samples for evaluation on I_2 . In Corollary O5.1: The concrete procedure fits p on I_0 and, conditional on that first stage, fits g_1 and the treated boundary law on I_1 , with predictions clipped to their stated supports.

We model the alternative policy as $q(z, x) = \phi(z, x, p(z, x))$, where the known map $\phi : \mathcal{Z} \times \mathcal{X} \times [0, 1] \rightarrow [0, 1]$ is jointly measurable. Common examples include propensity shifts $\phi(z, x, p) = p + \alpha$ or $p(1 + \alpha)$ on designs for which the shifted propensities remain in $[0, 1]$, or setting $\phi(z, x, p) = r(z, x)$ to evaluate a specific alternative policy r .

Remark (notation). $\phi : \mathcal{Z} \times \mathcal{X} \times [0, 1] \rightarrow [0, 1]$ is the known jointly measurable policy map specifying $q(z, x) = \phi(z, x, p(z, x))$ and the substitution $\hat{q}(z, x) = \phi(z, x, \hat{p}(z, x))$. In Lemma D.1, Theorem 4.7: $\phi(z, x, p)$ is differentiable in p , and $\partial\phi/\partial p$ is bounded and uniformly Lipschitz in p over $z \in \mathcal{Z}$ and $x \in \mathcal{X}$. In Corollary O5.1: $|\phi(z, x, p') - \phi(z, x, p)| \leq L_\phi |p' - p|$ for all $(z, x, p, p') \in \mathcal{Z} \times \mathcal{X} \times [0, 1]^2$, with $L_\phi < \infty$.

We focus our exposition on the bounds for the PRTE because the derivations for other causal quantities proceed similarly. For simplicity, we first develop theory for the unscaled numerator $\mathbb{E}[Y^q - Y]$. In the empirical designs considered below, the PRTE denominator is the known policy-shift size $\alpha > 0$, so the reported bounds and confidence intervals for PRTE_α are obtained by dividing the corresponding numerator bounds by α . More generally, when $\mathbb{E}[q(Z, X) - p(Z, X)]$ is not fixed by design, it is point-identified and can be estimated separately. Thus, our target weight function simplifies to

$$\omega(x, u) \equiv \mathbb{E}_{\text{obs}}[\mathbf{1}(u \leq q(Z, X)) - \mathbf{1}(u \leq p(Z, X)) \mid X = x].$$

4.1 Discrete Instrument Setting

For a fixed x , the weight function $\omega(x, u)$ is a step function that is non-increasing on each latent interval $u \in (p_i(x), p_{i+1}(x)]$, with jumps only at the alternative-policy values $q(z, x)$ falling inside that interval. This piecewise structure allows us to decompose the closed-form lower bound integral into a finite weighted sum over the instrument level sets.

Define $q_{j,i}(x)$ as the j -th smallest value of the alternative policy propensity $q(z, x)$ that falls in the baseline bin $[p_i(x), p_{i+1}(x))$ for $i < K$ and in the closed final bin $[p_K(x), 1]$ for $i = K$, ordered as follows:

$$p_i(x) \leq q_{1,i}(x) < \dots < q_{l_i,i}(x) < p_{i+1}(x) \quad (i < K),$$

with the boundary conventions $q_{0,i}(x) = p_i(x)$ and $q_{l_i+1,i}(x) = p_{i+1}(x)$. Let the baseline and alternative instrument level sets be

$$S_k = \{z \in \mathcal{Z} : p(z, x) = p_k(x)\}, \quad T_{j,k} = \{z \in \mathcal{Z} : q(z, x) = q_{j,k}(x)\}, \quad j = 1, \dots, l_k.$$

For the final bin, the last displayed strict inequality is replaced by $q_{l_K,K}(x) \leq p_{K+1}(x) = 1$. Here $q_{0,i}$ and $q_{l_i+1,i}$ are integration boundaries; set $T_{0,i} = T_{l_i+1,i} = \emptyset$. Thus, if $q_{1,i} = p_i$, instruments attaining that value belong only to $T_{1,i}$, while an alternative propensity equal to one belongs to $T_{l_K,K}$ and is retained in the tail term. Values of u exactly on a boundary are Lebesgue-null and do not affect the target integral.

Remark (notation). $q_{j,k}(x)$ denotes a population policy level for $k = 0, \dots, K$ and $j = 1, \dots, l_k$, ordered distinctly within its baseline bin, with $T_{j,k} = \{z \in \mathcal{Z} : q(z, x) = q_{j,k}(x)\}$; $K, l_k, T_{j,k}$ are invariant in x , and $q_{0,k}, q_{l_k+1,k}$ are integration boundaries with empty associated instrument sets. In Lemma D.1, Theorem 4.7: For every nonempty $T_{j,k}$ in a nonfinal bin $k < K$, $q_{j,k}(x) \neq p_k(x)$; the structural boundary $q_{0,k} = p_k$ is therefore distinct from every regular policy boundary used by the displayed scores. $l_k \in \mathbb{N}_0$, for $k = 0, \dots, K$, counts distinct population policy levels rather than instruments, is invariant in x under the stipulated invariant ordering, and excludes structural integration boundaries while retaining an actual policy value equal to one in the final bin. $S_k = \{z \in \mathcal{Z} : p(z, x) = p_k(x)\}$ denotes the population baseline instrument level set for $k = 1, \dots, K$, invariant in x under the maintained equality-class convention; $S_0 = \{z \in \mathcal{Z} : p(z, x) = 0\}$, when nonempty, is the structural-zero class, and $\widehat{S}_k$ denotes a recovered set rather than S_k itself. In Lemma D.1, Theorem 4.7: At $k = 0$, minus-side conditional expectations and residuals are omitted and assigned zero only for notation; no treated conditional regression on S_0 or division by $\pi_0(X)$ is used. $T_{j,k}$ denotes the population alternative-policy equality class for $k = 0, \dots, K$ and $j = 1, \dots, l_k$, invariant in x under the stipulated invariant ordering; structural integration-boundary sets are empty, an actual policy value equal to $p_k(x)$ belongs to $T_{1,k}$, and an actual value equal to one is retained in $T_{l_K, K}$. In Lemma D.1, Theorem 4.7: For $k = 0, \dots, K - 1$, every nonempty $T_{j,k}$ has $q_{j,k}(x) \neq p_k(x)$ for all $x \in \mathcal{X}$. S is an auxiliary $\text{Unif}(0, 1)$ draw independent of $U \sim \nu_{k|x}$ for each fixed x and nonempty gap $G_k(x)$, used to randomize ties in $R_{k,x}(U, S)$ and integrated out in the attaining outcome kernel.

The DML scores developed after the reduction differentiate $q(z, x) = \phi(z, x, p(z, x))$ with respect to the propensity score, so we impose the following condition only for inference.

Assumption 4.1 (Policy-Map Smoothness). The map $\phi(z, x, p)$ is differentiable with respect to p , and $\frac{\partial \phi}{\partial p}$ is bounded and uniformly Lipschitz in p for all $z \in \mathcal{Z}$ and $x \in \mathcal{X}$.

Remark. For every $z \in \mathcal{Z}$ and $x \in \mathcal{X}$, $p \mapsto \phi(z, x, p)$ is differentiable on $(0, 1)$ and has finite right and left derivatives at $p = 0$ and $p = 1$, respectively; these are its endpoint values of $\partial \phi / \partial p$. There exist finite constants $C, L \geq 0$ such that, for all $z \in \mathcal{Z}$, $x \in \mathcal{X}$, and $p, p' \in [0, 1]$, $|\partial \phi / \partial p(z, x, p)| \leq C$ and $|\partial \phi / \partial p(z, x, p) - \partial \phi / \partial p(z, x, p')| \leq L|p - p'|$.

This condition controls the pathwise derivative and its second-order remainder.

We impose the following separation condition to consistently estimate the level sets.

Assumption 4.2 (Propensity-Level Separation). (i) There exists a constant $c_{\text{gap}} > 0$ such that, for all $x \in \mathcal{X}$ and all $v, v' \in R_x = \{p(z, x), q(z, x)\}_{z \in \mathcal{Z}}$, either $v' = v$ or $|v - v'| >$

c_{gap} . Moreover, the equality classes and their strict ordering are invariant in x : the sets S_k , $T_{j,k}$ and their indices K, l_k do not depend on x .

- (ii) For every $k = 0, \dots, K-1$ and every nonempty $T_{j,k}$, $q_{j,k}(x) \neq p_k(x)$ for all x . Under part (i), this implies $q_{j,k}(x) > p_k(x) + c_{\text{gap}}$ for $k \geq 1$.

Remark (notation). $R_x = \{p(z, x), q(z, x) : z \in \mathcal{Z}\}$ is the set of distinct baseline and alternative-policy propensity values at the covariate value x . In Proposition 4.3, Lemma D.1, Theorem 4.7: Distinct elements of R_x are separated by more than c_{gap} , uniformly in x ; the equality classes and their strict ordering are invariant in x , although the numerical propensity values may vary with x . c_{gap} is one common positive separation constant for distinct values in the combined baseline and alternative-policy range, uniformly over $x \in \mathcal{X}$. In Theorem 4.7: Level recovery joins averaged fitted values lying within $c_{\text{gap}}/2$ of each other, where c_{gap} denotes this same separation constant.

Part (i) implies the stable level sets needed by the finite-dimensional reduction and the clustering step below. We therefore suppress the dependence on x in S_k , $T_{j,k}$, K , and l_k from now on. Part (ii) keeps a nonempty pooled policy boundary distinct from the structural boundary $q_{0,k} = p_k$ in the displayed continuous-outcome quantile score and binary active-set score. We impose this condition only to simplify the clustering labels and score formulas; it is not required for identification. It can be removed by using a more elaborate clustering method that detects the coinciding boundaries.

Because $\mathcal{Z}$ is finite, ?? 3.6 holds trivially and every conditional LIV interval in Theorem 3.8 is degenerate, so that theorem specializes to the covariate version of Theorem 3.4; we refer to it in this form throughout the subsection. The resulting bound admits a representation as an expectation of a finite weighted sum over instrument level sets and their sub-intervals. Before stating this representation, we introduce the nuisance components that carry the observable content of the bound.

Nuisance components. The representation below is assembled from five objects, defined for each interval index $k \in \{0, \dots, K\}$ and sub-interval index $j \in \{0, \dots, l_k\}$. The first two are the aggregate weights that the step weight ω assigns to instruments,

$$\begin{aligned}\gamma_{\text{full},k}(x) &:= \mathbb{P}_{\text{obs}}\left(Z \in \bigcup_{i=k+1}^K \bigcup_{l=0}^{l_i} T_{l,i} \mid X = x\right) - \mathbb{P}_{\text{obs}}\left(Z \in \bigcup_{i=k+1}^K S_i \mid X = x\right), \\ \gamma_{j,k}(x) &:= \mathbb{P}_{\text{obs}}\left(Z \in \bigcup_{l=j}^{l_k} T_{l,k} \mid X = x\right),\end{aligned}$$

where $\gamma_{\text{full},k} \in [-1, 1]$ because it differences two conditional probabilities. The next two carry the outcome distribution: with the relative threshold $\kappa_{j,k}(x) := (q_{j,k}(x) - p_k(x))/(p_{k+1}(x) - p_k(x)) \in [0, 1]$,

$$J_{\text{full},k}(x) := (p_{k+1}(x) - p_k(x)) \mathbb{E}_{\mu_{1,k}|x}[Y],$$

$$J_{j,k}(x) := (p_{k+1}(x) - p_k(x)) \int_{\kappa_{j,k}(x)}^{\kappa_{j+1,k}(x)} Q_{Y,k|x}(u) du,$$

where $\mu_{1,k|x}$ is the identified conditional complier distribution of (3.17) and $Q_{Y,k|x}$ is its quantile function. The fifth is the trivial-bound boundary term

$$\Delta_K(x) := y_{\min} \mathbb{E}_{\text{obs}} \left[\mathbf{1} \left(Z \in \bigcup_{j=1}^{l_K} T_{j,K} \right) (q(Z, X) - p_K(X)) \mid X = x \right].$$

This is the tail term of Theorem 3.8 under the PRTE weight: since $p(z, x) \leq p_K(x)$ for every z , the weight satisfies $\omega(x, u) = \mathbb{E}_{\text{obs}}[\mathbf{1}(u \leq q(Z, X)) \mid X = x] \geq 0$ on $(p_K(x), 1]$, so the $y_{\max}$ branch of the tail integral is inactive, and Fubini converts $y_{\min} \int_{p_K(x)}^1 \omega(x, u) du$ into the displayed expectation.

Remark (notation). $\gamma_{\text{full},k}(x)$ is the population difference of the two conditional probabilities in its defining display, taking values in $[-1, 1]$ for $k = 0, \dots, K$, with $\gamma_{\text{full},K} \equiv 0$ by empty unions; the full-interval products use $k = 0, \dots, K - 1$. $\kappa_{j,k}(x)$ is the population relative threshold for a nonempty baseline interval, indexed by $k = 0, \dots, K - 1$ and $j = 0, \dots, l_k + 1$, with $\kappa_{0,k}(x) = 0$ and $\kappa_{l_k+1,k}(x) = 1$. $J_{\text{full},k}(x)$ is the population full-interval contribution $(p_{k+1}(x) - p_k(x)) \mathbb{E}_{\mu_{1,k}|x}[Y]$ for $k = 0, \dots, K - 1$. In Proposition 4.4, Lemma D.1, Theorem 4.7: In the binary-outcome case, $J_{\text{full},k}(x) = P_{1,k+1}(x) - P_{1,k}(x)$, where $P_{1,k}(x) = \mathbb{E}_{\text{obs}}[YW \mid Z \in S_k, X = x]$ for $k \geq 1$ and $P_{1,0}(x) = 0$. $\Delta_K(x) = y_{\min} \mathbb{E}_{\text{obs}}[\max\{0, q(Z, X) - p_K(X)\} \mid X = x]$ is the conditional lower-bound contribution on $(p_K(x), 1]$, with the final policy bin $[p_K(x), 1]$ retaining instruments whose alternative propensity equals one.

Estimating these five requires three further level-set quantities that do not appear in the representation itself—the tail weight $\gamma_K(x) := \mathbb{P}_{\text{obs}}(Z \in \bigcup_{j=1}^{l_K} T_{j,K} \mid X = x)$, the within-set assignment shares $a_{j,k}(z, x) := \mathbb{P}_{\text{obs}}(Z = z \mid Z \in T_{j,k}, X = x)$, and the pooled representation $q_{j,k}(x) = \sum_{z \in T_{j,k}} a_{j,k}(z, x) q(z, x)$ that extends the policy level regularly when estimated instrument-specific levels are not exactly tied. The latter two are detailed in Section B.1, and all three enter only through the scores.

Remark (notation). $\gamma_K(x)$ is the population conditional probability of $\bigcup_{j=1}^{l_K} T_{j,K}$ in the final closed policy bin $[p_K(x), 1]$, taking values in $[0, 1]$ and equaling zero when $l_K = 0$; actual

policy levels equal to one are included. $a_{j,k}(z, x) = \mathbb{P}_{\text{obs}}(Z = z \mid Z \in T_{j,k}, X = x)$ is the population assignment-share function for each nonempty policy set, indexed by $k = 0, \dots, K$ and $j = 1, \dots, l_k$, with $a_{j,k}(z, x) = 0$ outside $T_{j,k}$ and $\sum_{z \in T_{j,k}} a_{j,k}(z, x) = 1$ wherever the conditioning probability is positive.

Proposition 4.3 (Discrete-Instrument Reduction to Nuisance Functional). *Suppose ?? 2.1 and part (i) of ?? 4.2 hold and $\mathcal{Z}$ is finite. Then the sharp lower bound in Theorem 3.8 admits the representation*

$$\underline{\theta}_{\omega,1} = \mathbb{E} \left[\sum_{k=0}^{K-1} \left(\gamma_{\text{full},k}(X) J_{\text{full},k}(X) + \sum_{j=1}^{l_k} \gamma_{j,k}(X) J_{j-1,k}(X) \right) + \Delta_K(X) \right]. \quad (4.1)$$

The sharp upper bound admits an analogous representation, obtained by replacing $Q_{Y,k|x}(u)$ by $Q_{Y,k|x}(1-u)$ inside $J_{j,k}$ and swapping $y_{\min}$ with $y_{\max}$ in Δ_K .

Remark. The identified outcome components $J_{\text{full},k}$ and $J_{j,k}$ are defined for $0 \leq k < K$, with $0 \leq j \leq l_k$ for $J_{j,k}$; the final interval is represented by Δ_K , without an identified law $\mu_{1,K|x}$.

Remark (notation). $\gamma_{j,k}(x)$ is the population conditional probability of $\bigcup_{l=j}^{l_k} T_{l,k}$, taking values in $[0, 1]$ for $k = 0, \dots, K$ and $j = 0, \dots, l_k$, with $T_{0,k} = \emptyset$; the sub-interval products use $k = 0, \dots, K-1$ and $j = 1, \dots, l_k$. $J_{j,k}(x) = (p_{k+1}(x) - p_k(x)) \int_{\kappa_{j,k}(x)}^{\kappa_{j+1,k}(x)} Q_{Y,k|x}(u) du$ is the population lower-bound subinterval contribution on a nondegenerate identified bin, with structural thresholds $\kappa_{0,k} = 0$ and $\kappa_{l_k+1,k} = 1$. In Proposition 4.3, Proposition 4.4, Lemma D.1, Theorem 4.7: The bound uses $J_{j-1,k}$ for $0 \leq k < K$ and $1 \leq j \leq l_k$. In the continuous-outcome representation, for $0 \leq j < l_k$ and under the stated continuity conditions, $J_{j,k} = (p_{k+1} - p_k) \mathbb{E}_{\mu_{1,k|x}}[Y \mathbb{1}(\nu_{j,k} < Y \leq \nu_{j+1,k})]$; in the binary-outcome case, $J_{j-1,k} = h_{j,k}^-$. In Lemma D.1, Theorem 4.7: In the continuous-outcome score construction the plug-in component is $J_{j,k} = p_{k+1} J_{j,k}^+ - p_k J_{j,k}^-$; at $k = 0$, the minus-side conditional expectation is omitted and set to zero for notation.

Proposition 4.3 is the key analytical step that converts the optimal-transport integral in Theorem 3.8 into a Neyman-orthogonalizable functional of standard nuisance quantities—conditional probabilities, conditional means, and conditional quantile integrals. The proof, given in Section O3.1 of the Online Appendix, exploits the piecewise-constant, monotone structure of $\omega(X, \cdot)$ on each baseline interval and identifies the step-function value attached to each sub-interval $(q_{j,k}(X), q_{j+1,k}(X))$.

The sub-interval term $J_{j,k}$ still involves the conditional quantile function $Q_{Y,k|x}$, which is delicate to estimate directly. The next proposition replaces it with conditional expectations in the two outcome regimes of interest.

Proposition 4.4 (Specialization of the Sub-Interval Contribution). *Fix $k \in \{0, \dots, K-1\}$. Set the lower boundary $\nu_{0,k}(x) \equiv y_{\min}$, and let $\nu_{j,k}(x) := Q_{Y,k|x}(\kappa_{j,k}(x))$ for $j \in \{1, \dots, l_k\}$.*

(i) **Continuous outcome.** *If $\mu_{1,k|X}$ is continuous at $\nu_{j,k}(X)$ and $\nu_{j+1,k}(X)$ almost surely, then, for $j \in \{0, \dots, l_k - 1\}$,*

$$J_{j,k}(X) = (p_{k+1}(X) - p_k(X)) \mathbb{E}_{\mu_{1,k|X}}[Y \mathbb{1}(\nu_{j,k}(X) < Y \leq \nu_{j+1,k}(X))]. \quad (4.2)$$

(ii) **Binary outcome.** *If $Y \in \{0, 1\}$, then $J_{full,k}(X) = P_{1,k+1}(X) - P_{1,k}(X)$ with $P_{1,k}(x) := \mathbb{E}_{obs}[YW \mid Z \in S_k, X = x]$ for $k \geq 1$ and the boundary convention $P_{1,0}(x) \equiv 0$, and for $j \in \{1, \dots, l_k\}$,*

$$J_{j-1,k}(X) = h_{j,k}^-(X) := \max\left(0, q_{j,k}(X) - \max(q_{j-1,k}(X), p_{k+1}(X) - J_{full,k}(X))\right). \quad (4.3)$$

Consequently, (4.1) collapses to

$$\underline{\theta}_{\omega,1} = \mathbb{E} \left[\sum_{k=0}^{K-1} \left(\gamma_{full,k}(X) J_{full,k}(X) + \sum_{j=1}^{l_k} \gamma_{j,k}(X) h_{j,k}^-(X) \right) + \Delta_K(X) \right]. \quad (4.4)$$

Remark. The boundary quantile uses the bounded-domain lower generalized inverse $Q_{Y,k|x}(t) = \inf\{y \in \mathcal{Y} : \mu_{1,k|x}(\mathcal{Y} \cap (-\infty, y]) \geq t\}$ for $0 \leq t < 1$, so $Q_{Y,k|x}(0) = y_{\min}$. For each fixed $k \in \{0, \dots, K-1\}$ and $j \in \{0, \dots, l_k - 1\}$ to which part (i) is applied, its continuity hypothesis is $\mu_{1,k|X}(\{\nu_{j,k}(X)\}) = \mu_{1,k|X}(\{\nu_{j+1,k}(X)\}) = 0$ almost surely under the marginal law of X . For $0 < t < 1$, use the standard lower generalized quantile $Q_{Y,k|x}(t) = \inf\{y \in \mathbb{R} : \mu_{1,k|x}(\mathcal{Y} \cap (-\infty, y]) \geq t\}$. For $j = 1, \dots, l_k$, interpret $\nu_{j,k}(x) = Q_{Y,k|x}(\kappa_{j,k}(x))$ using $Q_{Y,k|x}(t) = \inf\{y \in \mathbb{R} : \mu_{1,k|x}(\mathcal{Y} \cap (-\infty, y]) \geq t\}$ for $0 < t < 1$, with $\nu_{0,k}(x) = y_{\min}$.

Remark (notation). $\nu_{j,k}$ denotes the population boundary function for $k = 0, \dots, K-1$ and $j = 0, \dots, l_k$, with $\nu_{0,k}(x) \equiv y_{\min}$ and $\nu_{j,k}(x) = Q_{Y,k|x}(\kappa_{j,k}(x))$ for $j \geq 1$; $\hat{\nu}_{j,k}$ denotes its estimate. In Proposition 4.4: In part (i), the gap law $\mu_{1,k|X}$ is continuous at the two boundaries $\nu_{j,k}(X)$ and $\nu_{j+1,k}(X)$ used in each band, almost surely. In Lemma D.1, Theorem 4.7: In the continuous-outcome part, the relevant treated conditional laws have uniformly bounded densities in fixed neighborhoods of $\nu_{j,k}(X)$ for $j = 1, \dots, l_k$, and no atom at the fixed boundary $\nu_{0,k} = y_{\min}$; only the boundaries with $j \geq 1$ are estimated. For $k = 1, \dots, K$, $P_{1,k}(x) = \mathbb{E}_{obs}[YW \mid Z \in S_k, X = x]$ is the population conditional treated-success regression in the binary-outcome setting, while $P_{1,0}(x) \equiv 0$ is a structural boundary convention rather than a conditional regression on S_0 . $h_{j,k}^-(x)$ is the binary-outcome population subinterval contribution for $k = 0, \dots, K-1$ and $j = 1, \dots, l_k$, with $q_{0,k}(x) = p_k(x)$, and equals $J_{j-1,k}(x)$.

Part (i) converts the sub-interval quantile integral into a plain conditional expectation of Y on the quantile-defined band $(\nu_{j,k}, \nu_{j+1,k}]$, so only the $l_k + 1$ boundary quantiles $\nu_{j,k}$ —rather than the full quantile process—enter the estimation. Part (ii) leverages the step-function form of $Q_{Y,k|X}$ when Y is binary: the complier mean $\mathbb{E}_{\mu_{1,k|X}}[Y]$ reduces to a difference of observed joint probabilities, and the integral of the step function over $[\kappa_{j-1,k}, \kappa_{j,k}]$ evaluates to the closed-form $h_{j,k}^-$. Both parts are proved in Section O3.2 of the Online Appendix.

Together, Propositions 4.3 and 4.4 expose (4.1) as a functional of conditional expectations, boundary quantiles, and instrument-level-set probabilities—the form required for DML-style $\sqrt{n}$ -inference. In both outcome regimes the estimator takes the same aggregated Neyman-orthogonal form,

$$\hat{\underline{\theta}}_{\omega,1} = \frac{1}{|I_2|} \sum_{i \in I_2} \left[\sum_{k=0}^{K-1} \left(\hat{\psi}_{\text{full},k}^{\text{prod}}(O_i; \hat{\eta}) + \sum_{j=1}^{l_k} \hat{\psi}_{j,k}^{\text{prod}}(O_i; \hat{\eta}) \right) + \hat{\psi}_{\Delta_K}(O_i; \hat{\eta}) \right], \quad (4.5)$$

where each $\hat{\psi}$ is a Neyman-orthogonal score that corrects for the estimation error of its associated nuisance. What differs across the two regimes is the nuisance list $\hat{\eta}$, and hence the score formulas: these are given in Section C.1 for a continuous outcome and in Section C.2 for a discrete one. We treat the two regimes in turn and then state a single asymptotic normality result covering both.

Remark (notation). $\hat{\underline{\theta}}_{\omega,1}$ denotes the sample-based estimator of the treated-component sharp lower bound, formed by averaging fitted score contributions over I_2 . In Theorem 4.7: $\hat{\underline{\theta}}_{\omega,1} = |I_2|^{-1} \sum_{i \in I_2} \left[\sum_{k=0}^{K-1} \left(\hat{\psi}_{\text{full},k}^{\text{prod}}(O_i; \hat{\eta}) + \sum_{j=1}^{l_k} \hat{\psi}_{j,k}^{\text{prod}}(O_i; \hat{\eta}) \right) + \hat{\psi}_{\Delta_K}(O_i; \hat{\eta}) \right]$, with the nuisance list and score formulas appropriate to the continuous or binary outcome regime. In Theorem 4.10, Corollary O5.1: $\hat{\underline{\theta}}_{\omega,1} = |I_2|^{-1} \sum_{i \in I_2} (\hat{\psi}_1(O_i) + \hat{\psi}_2(O_i) + \hat{\psi}_3(O_i))$, where the three contributions substitute estimated nuisances into the three terms of equation 4.7, with predictions clipped to their known supports. In Corollary O5.1: The nuisance estimates are those supplied by Algorithm 3, including the local-linear regression and treated boundary conditional quantile construction. n is the total observed sample size indexing the estimation sequences and their asymptotics. In Theorem 4.7: n ranges over sample sizes permitting the stipulated equal-half split, so n is even. $\hat{\psi}$ denotes the total fitted integrand averaged over I_2 to form $\hat{\underline{\theta}}_{\omega,1}$. In Theorem 4.7: $\hat{\psi}(O; \hat{\eta}) = \sum_{k=0}^{K-1} \left(\hat{\psi}_{\text{full},k}^{\text{prod}}(O; \hat{\eta}) + \sum_{j=1}^{l_k} \hat{\psi}_{j,k}^{\text{prod}}(O; \hat{\eta}) \right) + \hat{\psi}_{\Delta_K}(O; \hat{\eta})$, using the continuous-outcome or binary-outcome score formulas according to the outcome regime. In the oracle procedure used in the proof, this fitted function is fixed conditional on I_1 ; the feasible procedure uses recovered level sets. In Theorem 4.10: $\hat{\psi} = \hat{\psi}_1 + \hat{\psi}_2 + \hat{\psi}_3$, obtained by substituting fitted nuisances into the three terms of (4.7); this is the plug-in integrand, fixed

conditional on I_0, I_1 , rather than the finite-instrument orthogonal score.

4.1.1 Continuous Outcome

We first treat the case where the conditional complier distribution $\mu_{1,k|X}$ is continuous at the boundary quantiles $\nu_{j,k}(X)$, so Proposition 4.4(i) applies. Recall that $\nu_{j,k}(x) = Q_{Y,k|x}(\kappa_{j,k}(x))$ denotes the boundary quantile introduced in Proposition 4.4. The representation (4.2) converts each sub-interval term $J_{j,k}(X)$ into a conditional expectation over the band $(\nu_{j,k}(X), \nu_{j+1,k}(X)]$, thereby reducing the nuisance list from the full quantile process to the boundary quantiles $\nu_{j,k}$ alone. The Neyman-orthogonal score construction additionally requires the conditional indicator expectations

$$\begin{aligned} M_{j,k}^+(x) &:= \mathbb{E}[\mathbf{1}(Y \leq \nu_{j,k}(x)) \mid Z \in S_{k+1}, W = 1, X = x], \\ M_{j,k}^-(x) &:= \mathbb{E}[\mathbf{1}(Y \leq \nu_{j,k}(x)) \mid Z \in S_k, W = 1, X = x], \end{aligned}$$

which enter the Riesz representers $\alpha_{p_k}, \alpha_{p_{k+1}}$ for the pathwise derivatives with respect to the interval boundaries p_k, p_{k+1} , and whose own estimation error is corrected by the residuals $IF_{M_{j,k}^\pm}$ (see Section C.1).

Remark (notation). $M_{j,k}^+$ is the population conditional indicator-mean function at the true boundary $\nu_{j,k}$, indexed by $k = 0, \dots, K-1$ and $j = 0, \dots, l_k$, with conditioning on $Z \in S_{k+1}$, $W = 1$, and $X = x$. In Lemma D.1, Theorem 4.7: In the continuous-outcome part, the no-atom condition at $y_{\min}$ and $\nu_{0,k} = y_{\min}$ imply $M_{0,k}^+(X) = 0$ almost surely. $M_{j,k}^-(x)$ is the population conditional indicator mean among treated observations with $Z \in S_k$ at the continuous-outcome boundary $\nu_{j,k}(x)$ for $k \geq 1$; at $k = 0$ this minus-side regression is omitted and defined as zero solely for notational convenience. α_{p_k} is the population lower-boundary representer for the particular component $J_{j,k}$, with the dependence on j suppressed in its notation, for $k = 1, \dots, K-1$ and $j = 0, \dots, l_k - 1$; there is no p_0 perturbation or correction. $\alpha_{p_{k+1}}$ denotes the upper interval-boundary representer for the continuous-outcome component $J_{j,k}$, with suppressed dependence on j , indexed by $k = 0, \dots, K-1$ and $j = 0, \dots, l_k - 1$; it remains an upper-side representer at $k = 0$.

The functional form in (4.1) and (4.2) fixes the nuisance list $\hat{\eta}$ entering (4.5); the rest of this subsection develops the corresponding estimators, trained on an auxiliary sample I_1 .

Estimation Procedure. Following the standard DML template, we randomly partition the sample into an auxiliary set I_1 and a main estimation set I_2 . On I_1 we estimate the propensity score p , clip it to $[0, 1]$, and set $\hat{q} = \phi(\cdot, \cdot, \hat{p})$; using only the covariates in I_2 , we

average $\hat{p}$ and $\hat{q}$ over X and recover the level sets $\hat{S}_k$ and $\hat{T}_{j,k}$ as connected components of the graph that joins two averaged values lying within $c_{\text{gap}}/2$ of each other. Given those sets we fit the remaining nuisances on I_1 —the interval boundaries $\hat{p}_k$, the probability weights $\hat{\gamma}$, $\hat{\pi}_k$, $\hat{\tau}_{j,k}$ and within-set shares $\hat{a}_{j,k}$, the boundary quantiles $\hat{\nu}_{j,k}$, and the conditional expectations $\hat{M}_{j,k}^\pm$, $\hat{J}_{j,k}^\pm$, $\hat{J}_{\text{full},k}^\pm$ —clipping each to its known range, and finally average the plug-in orthogonal scores over I_2 to form (4.5). Algorithm 1 states the procedure step by step; the explicit graph construction, the moment equation defining $\nu_{j,k}$, the conditional-expectation components, and the Neyman-orthogonal scores ψ^{prod} are given in Section B.1 and Section C.1.

Remark (notation). ψ denotes the uncentered total population integrand associated with the estimator of $\theta_{\omega,1}$, with its nuisance arguments supplied explicitly when needed. In Lemma D.1, Theorem 4.7: $\psi(O; \eta) = \sum_{k=0}^{K-1} (\psi_{\text{full},k}^{\text{prod}}(O; \eta) + \sum_{j=1}^{l_k} \psi_{j,k}^{\text{prod}}(O; \eta)) + \psi_{\Delta_K}(O; \eta)$, with the component formulas in Section C.1 for continuous outcomes and Section C.2 for binary outcomes; η_0 supplies the true nuisance values. In Theorem 4.10, Corollary O5.1: $\psi(O) = -\int_{\min\{\underline{p}(X), q(Z, X)\}}^{\underline{p}(X)} Q_{1, \underline{p}|X}(v/\underline{p}(X)) dv + g_1(q^{\text{tr}}(Z, X), X) - g_1(P, X) + y_{\min}(q(Z, X) - \bar{p}(X))_+$, with the first integral zero when $\underline{p}(X) = 0$; this is the population plug-in integrand. $\hat{p} : \mathcal{Z} \times \mathcal{X} \rightarrow [0, 1]$ is a fitted estimator of the conditional treatment propensity $p(z, x)$, with fitted values clipped to $[0, 1]$. In Lemma D.1, Theorem 4.7: $\hat{p}$ is the pre-recovery first-stage fit trained on I_1 ; the level-recovery averages evaluate this fit at the covariates in I_2 . In Theorem 4.7: $\max_{z \in \mathcal{Z}} \|\hat{p}(z, \cdot) - p(z, \cdot)\|_{L_2} = o_P(n^{-1/4})$. In Theorem 4.10, Corollary O5.1: $\hat{p}$ is trained on the propensity sample I_0 and is held fixed conditional on I_0 when constructing the generated propensity index and localization rules on I_1 . In Corollary O5.1: $\|\hat{p} - p\|_\infty = O_P(r_{p,n})$. $\hat{q}(z, x) = \phi(z, x, \hat{p}(z, x)) \in [0, 1]$ is the fitted alternative propensity obtained by substituting the clipped baseline propensity estimator into the known policy map. In Lemma D.1, Theorem 4.7: $\hat{q}$ is formed from the pre-recovery propensity fit $\hat{p}$ trained on I_1 ; the level-recovery average evaluates this fitted function at the covariates in I_2 . In Theorem 4.10, Corollary O5.1: $\hat{q}$ is formed from $\hat{p}$ trained on I_0 ; the notation $\hat{q} = \phi(Z, X, \hat{p}(Z, X))$ denotes evaluation of this fitted function at (Z, X) . $\hat{\pi}_k(x)$ estimates $\mathbb{P}_{\text{obs}}(Z \in S_k \mid X = x)$ by fitting the indicator of the recovered set $\hat{S}_k$ on I_1 , for $k = 1, \dots, K$, and is clipped first to $[0, 1]$ and then to $[c_\pi/2, 1]$; no denominator estimator $\hat{\pi}_0$ is required at the structural boundary. $\hat{\tau}_{j,k}(x)$ estimates $\mathbb{P}_{\text{obs}}(Z \in T_{j,k} \mid X = x)$ by fitting the indicator of the recovered set $\hat{T}_{j,k}$ on I_1 , for nonempty policy level sets with $k = 0, \dots, K$ and $j = 1, \dots, l_k$, and is clipped first to $[0, 1]$ and then to $[c_\pi/2, 1]$. $\hat{S}_k$ is the sample-dependent baseline level-set estimate obtained from the pre-recovery fits on I_1 and their averages over the covariates in I_2 , with a threshold-defined structural-zero set and positively indexed baseline components ordered

by their averages. $\hat{T}_{j,k}$ is the sample-dependent policy level-set estimate obtained from the pre-recovery fits on I_1 and their averages over the covariates in I_2 ; it depends on both of these inputs and is indexed by the recovered component ordering and baseline bins.

Let $\pi_k(x) := \mathbb{P}_{\text{obs}}(Z \in S_k \mid X = x)$ and $\tau_{j,k}(x) := \mathbb{P}_{\text{obs}}(Z \in T_{j,k} \mid X = x)$ denote the conditional probabilities of the baseline and alternative-policy level sets. The following overlap assumption bounds the inverse probability weights in the debiased terms.

Remark (notation). $\pi_k(x) = \mathbb{P}_{\text{obs}}(Z \in S_k \mid X = x)$ is a population conditional probability taking values in $[0, 1]$, with fitted baseline nuisances indexed by $k = 1, \dots, K$; no π_0 regression or inverse-probability correction is fitted. In Lemma D.1, Theorem 4.7: $\pi_k(x) \geq c_\pi > 0$ for every $x \in \mathcal{X}$ and $k = 1, \dots, K$. $\tau_{j,k}(x) = \mathbb{P}_{\text{obs}}(Z \in T_{j,k} \mid X = x)$ is a population conditional probability in $[0, 1]$, with fitted policy-set nuisances indexed by nonempty $T_{j,k}$, $k = 0, \dots, K$ and $j = 1, \dots, l_k$; structural boundary sets are empty and have probability zero. In Lemma D.1, Theorem 4.7: $\tau_{j,k}(x) \geq c_\pi > 0$ for every $x \in \mathcal{X}$ and every nonempty $T_{j,k}$.

Assumption 4.5 (Overlap). There exists a constant $c_\pi > 0$ such that $\pi_k(x) \geq c_\pi$ for every $x \in \mathcal{X}$ and observed positive-propensity level $k \in \{1, \dots, K\}$, and $\tau_{j,k}(x) \geq c_\pi$ for every nonempty alternative-policy level set $T_{j,k}$.

Remark. For every $x \in \mathcal{X}$ and every nonempty $T_{j,k}$, $\tau_{j,k}(x) \geq c_\pi$, where $c_\pi > 0$ is the same constant used for π_k in ?? 4.5.

Remark (notation). c_π is a common positive overlap constant for the specified nonempty level sets, with no overlap requirement on the structural index $k = 0$. In Theorem 4.7: The same c_π determines the clipping interval $[c_\pi/2, 1]$ for $\hat{\pi}_k$ and $\hat{\tau}_{j,k}$.

The index $k = 0$ is the structural convention $p_0 \equiv 0$ and need not correspond to a nonempty instrument level set. Accordingly, every minus-side conditional expectation, inverse-probability correction, and propensity residual indexed by $k = 0$ is omitted from the scores; it is multiplied by the known boundary mass $p_0 = 0$.

4.1.2 Discrete Outcome

When the outcome Y takes finitely many values—in particular, when $Y \in \{0, 1\}$ is binary—the conditional quantile function $Q_{Y,k|X}$ becomes a step function with atoms, and the continuity hypothesis of Proposition 4.4(i) fails. Instead, Proposition 4.4(ii) applies: the representation (4.4) expresses $J_{\text{full},k}$ as a difference of observed joint probabilities $P_{1,k+1} - P_{1,k}$ and evaluates the sub-interval contribution in closed form through $h_{j,k}^-$, so the nuisance list

collapses to conditional expectations alone and quantile estimation is avoided entirely. We present the binary case; the extension to general discrete outcomes is analogous.

The estimator is again (4.5), now with the scores of Section C.2. The nuisance-estimation procedure mirrors Algorithm 1 with the conditional-quantile and indicator-expectation steps removed, and with one addition: because the $\hat{P}_{1,k}$ are fitted on different instrument subsets, they need not be monotone in k in finite samples, so they are clipped to $[0, 1]$ and rearranged before $\hat{J}_{\text{full},k} = \hat{P}_{1,k+1}^* - \hat{P}_{1,k}^*$ is formed. Algorithm 2 states the procedure and Section B.2 and Section C.2 give the full details and the orthogonal scores.

Remark (notation). $\hat{P} = \hat{p}(Z, X)$ denotes the fitted baseline propensity index, with observation-indexed notation $\hat{P}_i = \hat{p}(Z_i, X_i)$ and values in $[0, 1]$. In Theorem 4.10, Corollary O5.1: The fitted map $\hat{p}$ is held fixed conditional on the training samples I_0, I_1 when evaluating expressions involving $\hat{P} = \hat{p}(Z, X)$ under the target distribution of (Z, X) . In Corollary O5.1: The fitted map $\hat{p}$ is trained on I_0 , so the selection rule involving $\hat{P}$ is fixed conditional on I_0 .

To ensure differentiability of the target functional with respect to the nuisance parameters, we impose the following gap condition, which replaces the conditional-density smoothness required by part (i) of Theorem 4.7.

Assumption 4.6 (Lower-Boundary Discrete Outcome Gap). There exists $c_h > 0$ such that, for every $k = 0, \dots, K - 1$ with $l_k \geq 1$ and $j = 0, \dots, l_k$, almost surely,

$$|p_{k+1}(X) - J_{\text{full},k}(X) - q_{j,k}(X)| \geq c_h.$$

Remark (notation). c_h is one positive constant uniform over the indicated indices and almost every covariate value, including the fixed lower boundary $q_{0,k} = p_k$.

This separation makes branch-selection errors for the lower bound second order under the stated L_2 rates, including at the fixed lower boundary $q_{0,k} = p_k$. It guarantees that the lower target is pathwise-differentiable in the nuisances despite the max operators appearing in $h_{j,k}^-$. The symmetric upper estimator, when used, requires the analogous separation $|p_k + J_{\text{full},k} - q_{j,k}| \geq c_h$ for $j = 1, \dots, l_k + 1$.

4.1.3 Asymptotic Normality

The two regimes share a common set of conditions—the structural assumptions, the level-set separation and overlap conditions, and $n^{-1/4}$ nuisance rates with clipping—and differ only in what is required at the sub-interval boundaries: a conditional density near each boundary quantile when the outcome is continuous, and the gap condition ?? 4.6 when it is discrete. A single theorem therefore covers both; both parts are proved in Section D.

Theorem 4.7 (Asymptotic Normality). *Suppose ?? 2.1, ?? 4.1, ?? 4.2, and ?? 4.5 hold and $\mathcal{Z}$ is finite. Suppose the pre-recovery first stage fitted on I_1 satisfies $\max_{z \in \mathcal{Z}} \|\widehat{p}(z, \cdot) - p(z, \cdot)\|_{L_2} = o_P(n^{-1/4})$. Given the recovered level sets, suppose the remaining score nuisances fitted on I_1 are clipped to their known bounded ranges, every estimated conditional probability is clipped to $[0, 1]$ before any difference is formed, $\widehat{\pi}_k$ and $\widehat{\tau}_{j,k}$ are clipped further to $[c_\pi/2, 1]$, and the estimators converge componentwise at rate $o_P(n^{-1/4})$ in L_2 . Suppose in addition that one of the following holds.*

- (i) **Continuous outcome.** *The nuisances include the boundary quantiles $\nu_{j,k}$, and, uniformly in X , each relevant conditional law $Y \mid X, Z \in S_{k'}, W = 1$ has a density bounded in a fixed neighborhood of every true boundary $\nu_{j,k}(X)$, $j = 1, \dots, l_k$, for $k' = k + 1$ and also for $k' = k$ when $k \geq 1$; at the fixed boundary $\nu_{0,k} = y_{\min}$, these conditional laws have no atom at $y_{\min}$.*
- (ii) **Discrete outcome.** *The outcome is binary, ?? 4.6 holds, and the nuisances are p_k , $P_{1,k}$, π_k , $\tau_{j,k}$, $a_{j,k}$, $\gamma_{\text{full},k}$, $\gamma_{j,k}$, and γ_K .*

Then the sample-split estimator (4.5) is $\sqrt{n}$ -consistent and asymptotically normal:

$$\sqrt{n/2}(\widehat{\theta}_{\omega,1} - \underline{\theta}_{\omega,1}) \xrightarrow{d} \mathcal{N}(0, \sigma^2),$$

where the asymptotic variance is the variance of the true orthogonal score, $\sigma^2 = \mathbb{E}[(\psi(O; \eta) - \underline{\theta}_{\omega,1})^2]$. It can be consistently estimated by its sample analog $\widehat{\sigma}^2 = \frac{1}{|I_2|} \sum_{i \in I_2} (\widehat{\psi}(O_i; \widehat{\eta}) - \widehat{\theta}_{\omega,1})^2$.

Additional condition. $K \geq 1$.

Remark. Every nuisance-error norm in this theorem is the population norm $\|\widehat{f} - f\|_{L_2}^2 = \mathbb{E}_X[|\widehat{f}(X) - f(X)|^2]$, with the fitted function held fixed during integration. For a fresh X with the population marginal law and independent of I_1 , $\|\widehat{p}(z, \cdot) - p(z, \cdot)\|_{L_2}^2 = \mathbb{E}[(\widehat{p}(z, X) - p(z, X))^2 \mid I_1]$, and $n^{1/4} \max_{z \in \mathcal{Z}} \|\widehat{p}(z, \cdot) - p(z, \cdot)\|_{L_2} \xrightarrow{P} 0$ as $n \rightarrow \infty$. Pointwise in x , the fitted probability, propensity, share, and indicator-mean nuisances lie in $[0, 1]$, $\widehat{\gamma}_{\text{full},k}(x) \in [-1, 1]$, $\widehat{\nu}_{j,k}(x)$, $\widehat{J}_{\text{full},k}^\pm(x) \in [y_{\min}, y_{\max}]$, and $\widehat{J}_{j,k}^\pm(x) \in [\min\{0, y_{\min}\}, \max\{0, y_{\max}\}]$, for each applicable fitted component; denominator clipping is applied further as stated. Componentwise L_2 errors integrate over the marginal law of X under $\mathbb{P}_{\text{obs}}$, holding each fitted function and any instrument argument z fixed; $o_P(n^{-1/4})$ refers to convergence in probability over the observed sample and random partition as $n \rightarrow \infty$. $Y \in \{0, 1\}$ almost surely. In this definition, $\psi_{\text{full},k}^{\text{prod}}$ and $\psi_{j,k}^{\text{prod}}$ are the binary-outcome components of Section C.2, and ψ_{Δ_K} is the common tail component of Section C.1. The binary nuisance list supplements the separately fitted first-stage propensity, so $\eta(x) = \{p(z, x), p_k(x), P_{1,k}(x), \pi_k(x), \tau_{j,k}(x), a_{j,k}(z, x), \gamma_{\text{full},k}(x), \gamma_{j,k}(x), \gamma_K(x)\}_{j,k}$, with

$p_0 = P_{1,0} = 0$ fixed and no fitted π_0 . The estimated baseline bins used to index $\widehat{T}_{j,k}$ include their left endpoints and exclude their right endpoints for $k < K$; the final bin for $k = K$ includes both endpoints, including 1. Two remaining labeled vertices are connected exactly when their averaged values differ by less than $c_{\text{gap}}/2$. The structural-zero assignment in Algorithm 1 is $\widehat{S}_0 = \{z : \bar{p}_2(z) \leq \rho_n\}$, with deterministic $\rho_n \downarrow 0$ satisfying $n^{1/4}\rho_n \rightarrow \infty$. At each evaluation covariate X_i , $(\widehat{P}_{1,1}^*(X_i), \dots, \widehat{P}_{1,K}^*(X_i))$ is the nondecreasing sort of the clipped vector $(\widehat{P}_{1,1}(X_i), \dots, \widehat{P}_{1,K}(X_i))$, with $\widehat{P}_{1,0}^*(X_i) = 0$; this same vector is used in every score plug-in and residual, without a separate gap projection. In this definition, all components of $\psi(O; \eta)$ are those of Section C.1, in particular $\psi_{j,k}^{\text{prod}} = \gamma_{j,k}\psi_{j-1,k}^J + J_{j-1,k}(\mathbb{1}(Z \in \bigcup_{l=j}^k T_{l,k}) - \gamma_{j,k})$.

4.2 Continuous Instrument Setting

We now consider a connected instrument domain $\mathcal{Z}$. Write $P = p(Z, X)$ and define

$$\begin{aligned} \underline{p}(x) &:= \inf_{z \in \mathcal{Z}} p(z, x), & \bar{p}(x) &:= \sup_{z \in \mathcal{Z}} p(z, x), \\ q^{\text{tr}}(z, x) &:= \min\{\bar{p}(x), \max\{\underline{p}(x), q(z, x)\}\}. \end{aligned}$$

The standing compactness of $\mathcal{Z}$ and continuity of $z \mapsto p(z, x)$ imply that the conditional support of P is the interval $[\underline{p}(x), \bar{p}(x)]$. Define $g_1(u, x) := g_1^c(u, x) = \int_{\mathcal{Y}} y H_{1,x}(u, dy)$, the canonical absolutely continuous version of $\mathbb{E}_{\text{obs}}[YW \mid P = u, X = x]$. The truncated policy propensity q^{tr} is needed because this function is identified only on the conditional support. For $\underline{p}(x) > 0$, let

$$\mu_{1,\underline{p}|x}(dy) := \frac{H_{1,x}(\underline{p}(x), dy)}{\underline{p}(x)}, \quad (4.6)$$

and denote its quantile function by $Q_{1,\underline{p}|x}$. Equivalently, it is the canonical boundary limit of $\mathcal{L}(Y \mid W = 1, P = p, X = x)$ as $p \downarrow \underline{p}(x)$ through the conditional support: it averages the law of $Y(1)$ over $0 < U \leq \underline{p}(x)$. When $\underline{p}(x) = 0$, choose $\mu_{1,\underline{p}|x}$ arbitrarily; every expression below multiplies its contribution by $\underline{p}(x)$ and is therefore invariant to this choice.

The general bound (3.18) contains the derivative $\partial_u g_1(u, X)$. The next proposition integrates this derivative analytically, so the estimator never differentiates a nonparametric regression.

Proposition 4.8 (Continuous-Instrument Reduction via Fubini). *Suppose ?? 2.1 and ?? 3.6 hold, and $\mathcal{Z}$ is connected. Then the sharp lower bound in Theorem 3.8 admits the representa-*

tion

$$\begin{aligned} \underline{\theta}_{\omega,1} = \mathbb{E}_{Z,X} \left[- \int_{\min\{\underline{p}(X), q(Z,X)\}}^{\underline{p}(X)} Q_{1,\underline{p}|X}(v/\underline{p}(X)) \, dv \right. \\ \left. + g_1(q^{\text{tr}}(Z, X), X) - g_1(P, X) \right. \\ \left. + y_{\min} \max\{0, q(Z, X) - \bar{p}(X)\} \right], \end{aligned} \quad (4.7)$$

where the first integral is defined as zero when $\underline{p}(X) = 0$. The sharp upper bound is obtained by replacing $Q_{1,\underline{p}|X}(v/\underline{p}(X))$ with $Q_{1,\underline{p}|X}(1 - v/\underline{p}(X))$ and $y_{\min}$ with $y_{\max}$.

Remark. For $\underline{p}(x) > 0$ and $0 < t < 1$, $Q_{1,\underline{p}|x}(t) = \inf\{y \in \mathbb{R} : \mu_{1,\underline{p}|x}(\mathcal{Y} \cap (-\infty, y]) \geq t\}$; at $t = 0, 1$ use the endpoint extension of the initial-gap outcome quantile in the referenced bound. For $0 < t < 1$, use the standard lower generalized quantile $Q_{1,\underline{p}|x}(t) = \inf\{y \in \mathbb{R} : \mu_{1,\underline{p}|x}(\mathcal{Y} \cap (-\infty, y]) \geq t\}$.

The first term uses the treated boundary law in (4.6); the second evaluates g_1 only at P and at the support projection $q^{\text{tr}}(Z, X)$. This lower truncation matters for policies with $q(Z, X) < \underline{p}(X)$: the movement below $\underline{p}(X)$ belongs to the OT gap term, whereas the LIV contribution stops at $\underline{p}(X)$. The third term covers $q(Z, X) > \bar{p}(X)$. The proof is given in Section O3.3 of the Online Appendix.

Proposition 4.8 motivates a plug-in estimator, obtained by replacing each nuisance in (4.7) with a machine-learning estimate trained on auxiliary samples. We use three-way sample splitting with $|I_j|/n \rightarrow \rho_j \in (0, 1)$ for $j = 0, 1, 2$: the sample I_0 estimates p , q , $\underline{p}$, and $\bar{p}$; conditional on I_0 , the sample I_1 estimates g_1 and the treated boundary law; and I_2 evaluates the target. The estimator is

$$\hat{\underline{\theta}}_{\omega,1} = \frac{1}{|I_2|} \sum_{i \in I_2} \left(\hat{\psi}_1(O_i) + \hat{\psi}_2(O_i) + \hat{\psi}_3(O_i) \right), \quad (4.8)$$

where $\hat{\psi}_1, \hat{\psi}_2, \hat{\psi}_3$ substitute estimated nuisances into the three terms of (4.7). Throughout, $\hat{P} = \hat{p}(Z, X)$ and $\hat{q} = \phi(Z, X, \hat{p}(Z, X))$, while $\hat{\underline{p}}, \hat{\bar{p}}, \hat{q}^{\text{tr}}, \min\{\hat{\underline{p}}(X), \hat{q}(Z, X)\}$, and $(\hat{q}(Z, X) - \hat{\bar{p}}(X))_+$ are obtained by substituting these estimates into the corresponding definitions. We clip estimated propensity scores and nuisance predictions to their known supports.

Remark (notation). $\hat{\underline{p}}(x) = \inf_{z \in \mathcal{Z}} \hat{p}(z, x)$ is the fitted lower propensity endpoint, a data-dependent function of x taking values in $[0, 1]$, held fixed conditional on the training samples I_0, I_1 . In Corollary O5.1: $\hat{\underline{p}}$ is obtained from the propensity fit trained on I_0 and is therefore fixed conditional on I_0 in the selected-law definition. $\hat{\bar{p}}(x) = \sup_{z \in \mathcal{Z}} \hat{p}(z, x) \in [0, 1]$ is the fitted

upper propensity endpoint, with the fitted function held fixed in target evaluation conditional on I_0, I_1 . In Corollary [O5.1](#): The endpoint is computed from $\hat{p}$ fitted on I_0 . $\hat{q}^{\text{tr}}(z, x) = \min\{\hat{p}(x), \max\{\hat{p}(x), \hat{q}(z, x)\}\}$ is the fitted policy propensity projected onto $[\hat{p}(x), \hat{p}(x)]$, with its fitted map held fixed in target evaluation conditional on I_0, I_1 . In Corollary [O5.1](#): The fitted map is constructed on I_0 from the clipped $\hat{p}$ and $\hat{q}(z, x) = \phi(z, x, \hat{p}(z, x))$. $\hat{\psi}_1(O)$ denotes the estimated lower-bound boundary contribution $-\int_{\min\{\hat{p}(X), \hat{q}(Z, X)\}}^{\hat{p}(X)} \hat{Q}_{1,p|X}(v/\hat{p}(X)) dv$, with value zero when $\hat{p}(X) = 0$. In Corollary [O5.1](#): For $\hat{p}(X_i) > 0$, the concrete implementation evaluates $-\hat{p}(X_i)[\hat{C}(1 | X_i) - \hat{C}(\hat{\tau}_i | X_i)]$, where $\hat{\tau}_i = \min\{1, \hat{q}(Z_i, X_i)/\hat{p}(X_i)\}$; cumulative-integral grid values are computed by the trapezoid rule, and off-grid lookup uses linear interpolation between adjacent cumulative-integral grid values. For $\hat{p}(X_i) = 0$, it returns zero without evaluating $\hat{\tau}_i$. $\hat{\psi}_3(O) = y_{\min}(\hat{q}(Z, X) - \hat{p}(X))_+$ is the fitted lower-bound tail score, evaluated at O_i for $i \in I_2$ with the fitted score map held fixed conditional on I_0, I_1 . For $j \in \{0, 1, 2\}$, $\rho_j \in (0, 1)$ is the limiting split proportion $|I_j|/n \rightarrow \rho_j$, with $\rho_0 + \rho_1 + \rho_2 = 1$.

We state the convergence rate under high-level nuisance-rate assumptions: [\(4.8\)](#) is well defined for any nuisance estimators trained on $I_0 \cup I_1$. Algorithm [3](#) then provides one concrete construction, and Corollary [O5.1](#) establishes its rate under conditions on those particular learners, collected in Section [O5](#) of the Online Appendix.

Assumption 4.9 (Targeted Nuisance Rates). Conditional on the training samples $I_0 \cup I_1$, suppose:

1. Each of the three maps entering the boundary integral and the trivial tail,

$$\Lambda \in \left\{ \underline{p}(X), \min\{\underline{p}(X), q(Z, X)\}, (q(Z, X) - \bar{p}(X))_+ \right\},$$

satisfies

$$\mathbb{E}_{Z, X} \left[|\hat{\Lambda} - \Lambda| \mid I_0, I_1 \right] = O_P(r_{e,n}),$$

with $r_{e,n} \rightarrow 0$.

2. The clipped estimator $\hat{g}_1$ satisfies the target-evaluation rate

$$\begin{aligned} \mathbb{E}_{Z, X} \left[|\hat{g}_1(\hat{q}^{\text{tr}}(Z, X), X) - g_1(q^{\text{tr}}(Z, X), X)| \right. \\ \left. + |\hat{g}_1(\hat{P}, X) - g_1(P, X)| \mid I_0, I_1 \right] = O_P(r_{g,n}). \end{aligned}$$

Here $r_{g,n} \rightarrow 0$.

3. Let $\hat{\mu}_{1,p|x}$ be the estimated treated boundary law with quantile function $\hat{Q}_{1,p|x}$. It satisfies the target-weighted rate

$$\mathbb{E}_X \left[\underline{p}(X) W_1 \left(\hat{\mu}_{1,p|X}, \mu_{1,p|X} \right) \mid I_0, I_1 \right] = O_P(r_{\mu,n}).$$

Here $r_{\mu,n} \rightarrow 0$.

Remark (notation). $\hat{g}_1(u, x)$ is an estimated regression on $[0, 1] \times \mathcal{X}$, clipped to the known range $[\min\{0, y_{\min}\}, \max\{0, y_{\max}\}]$ of YW , and evaluated at the estimated target indices $\hat{q}^{\text{tr}}(Z, X)$ and $\hat{P}$. In Corollary [O5.1](#): At $u_m = m/M_g$, $m = 0, \dots, M_g$, the fitted value is the intercept in $(\hat{g}_1(u_m, \cdot), \hat{\beta}_m) \in \arg \min_{g, \beta \in \mathcal{F}_n} \sum_{i \in I_1} K_{h_n}(\hat{P}_i - u_m) [Y_i W_i - g(X_i) - \beta(X_i)(\hat{P}_i - u_m)]^2$; between grid points it is the clipped linear interpolant. The learner loss compares $\hat{g}_1(\hat{U}_\ell, X)$ with $g_1(\Pi_X(\hat{U}_\ell), X)$ for $\ell \in \{1, 2\}$. $\hat{\mu}_{1,p|x}$ is an estimated probability law on $\mathcal{Y}$, indexed by $x \in \mathcal{X}$, whose quantile function is $\hat{Q}_{1,p|x}$ and whose target is the canonical treated boundary law; its target-error loss is weighted by $\underline{p}(x)$, so its value where $\underline{p}(x) = 0$ does not affect the target. In Corollary [O5.1](#): The law is generated by conditional check-function fits on $\mathcal{I}_{\delta,1} = \{i \in I_1 : W_i = 1, \hat{P}_i \leq \hat{p}(X_i) + \delta_n\}$ at $\tau_m = m/M_Q$, $m = 1, \dots, M_Q - 1$, clipped to $[y_{\min}, y_{\max}]$, sorted at each x , and linearly interpolated between adjacent quantile levels; endpoint values may be $y_{\min}, y_{\max}$ or the smallest and largest fitted values. $\hat{Q}_{1,p|x}$ denotes the quantile function of $\hat{\mu}_{1,p|x}$, with values in $[y_{\min}, y_{\max}]$ and endpoint values supplied by the chosen endpoint convention. In Corollary [O5.1](#): At $\tau_m = m/M_Q$, $m = 1, \dots, M_Q - 1$, fit $\tilde{Q}(\tau_m \mid \cdot) \in \arg \min_{f \in \mathcal{F}_Q} \sum_{i \in \mathcal{I}_{\delta,1}} \rho_{\tau_m}(Y_i - f(X_i))$, setting every grid fit to $y_{\min}$ if the selected sample is empty; clip to $[y_{\min}, y_{\max}]$, sort the fitted values for each x , and interpolate linearly in τ , with endpoints either $y_{\min}, y_{\max}$ or the smallest and largest fitted values. $r_{e,n}$ is a nonnegative rate sequence tending to zero controlling each of the three training-conditional absolute-error losses for the boundary and tail maps. In Corollary [O5.1](#): $r_{e,n} = O(r_{p,n})$ for the concrete nuisance construction. $r_{g,n}$ is a nonnegative rate sequence tending to zero for the sum of the two training-conditional target-evaluation losses, comparing fitted regressions at estimated indices with population regressions at the corresponding true indices. In Corollary [O5.1](#): $r_{g,n} = O(r_{g,\text{loc},n} + r_{p,n})$ for the concrete nuisance construction. $r_{\mu,n}$ is a nonnegative deterministic rate sequence tending to zero for the training-conditional, $\underline{p}(X)$ -weighted Wasserstein loss against the canonical treated boundary law. In Corollary [O5.1](#): For the concrete nuisance construction, one may take $r_{\mu,n} = O(r_{Q,n} + \delta_n + r_{p,n} + \varepsilon_n + M_Q^{-1})$. Λ is population notation whose meaning is determined by its local introduction: an evaluated endpoint or tail expression, or the logistic function in the simulation index. In Theorem [4.10](#): Λ ranges separately over $\underline{p}(X)$, $\min\{\underline{p}(X), q(Z, X)\}$, and $(q(Z, X) - \bar{p}(X))_+$, with the error

condition required for each choice; these are evaluated population expressions under the law of (Z, X) . $\widehat{\Lambda}$ denotes the corresponding estimated endpoint or tail expression evaluated at (Z, X) ; in $\mathbb{E}_{Z,X}[\widehat{\Lambda} - \Lambda \mid I_0, I_1]$, the fitted functions are held fixed and the expectation integrates the evaluation law of (Z, X) .

Part (1) controls the two limits of the boundary integral and the positive-part tail, averaged under the target distribution. Part (2) addresses the more demanding generated-regressor problem. The estimated propensity score plays two roles: $\widehat{P}$ is used in place of the latent index P when learning g_1 , and it determines the random arguments at which $\widehat{g}_1$ is evaluated. Part (3) controls the boundary OT contribution in its natural loss. Indeed, in one dimension W_1 equals the integrated absolute quantile error, so this assumption essentially requires that the estimated boundary quantile function $\widehat{Q}_{1,\underline{p}|x}$ converge in L_1 to the true boundary quantile function $Q_{1,\underline{p}|x}$, weighted by the target $\underline{p}(X)$.

Theorem 4.10 (Nonparametric Convergence Rate). *Suppose ?? 2.1, ?? 3.6, and ?? 4.9 hold. Then*

$$|\widehat{\underline{\theta}}_{\omega,1} - \underline{\theta}_{\omega,1}| = O_P(n^{-1/2} + r_{e,n} + r_{g,n} + r_{\mu,n}).$$

Remark. For every probability measure μ on $\mathbb{R}$ and $0 < t < 1$, $Q_\mu(t) = \inf\{y \in \mathbb{R} : \mu((-\infty, y]) \geq t\}$; the integral defining W_1 is a Lebesgue integral over $(0, 1)$ and requires no endpoint values.

Theorem 4.10 provides a high-level theoretical guarantee, and Section B.3 gives a concrete construction: on I_0 we regress W on (Z, X) ; on I_1 we estimate g_1 by boundary-adaptive local-linear smoothing in the scalar index P combined with a flexible regression in X , and we estimate the treated boundary law by conditional quantile regression on the treated observations whose estimated propensity lies within δ_n of the estimated lower boundary. Under conditions (L1)–(L3), Corollary O5.1 relates its learner errors to the target error:

$$|\widehat{\underline{\theta}}_{\omega,1} - \underline{\theta}_{\omega,1}| = O_P(n^{-1/2} + r_{p,n} + r_{g,\text{loc},n} + r_{Q,n} + \delta_n + \varepsilon_n + M_Q^{-1}),$$

where $r_{p,n}$ is a sup-norm first-stage rate, $r_{g,\text{loc},n}$ is an L_1 prediction rate for $\widehat{g}_1$ at the two arguments at which the target evaluates it, and $r_{Q,n}$ is a $\underline{p}(X)$ -weighted L_1 rate for the conditional quantile fit relative to the law selected by the estimated boundary window. The threshold-crossing model gives the linear localization error $\delta_n + r_{p,n}$ directly; ε_n and M_Q^{-1} are the trimmed quantile-tail and grid-interpolation errors.

Remark (notation). The notation δ denotes a nuisance perturbation direction, whereas the indexed notation δ_n denotes a fixed boundary-localization window sequence. In Lemma D.1,

Theorem 4.7: δ lies in the same function space as η and satisfies that $\eta_0 + t\delta$ is admissible for every t near zero; fixed structural components have zero perturbation, and within-set share components satisfy $\sum_{z \in T_{j,k}} \delta a_z(X) = 0$. In Corollary O5.1: $\delta_n \rightarrow 0$ is the localization window in $\mathcal{I}_{\delta,1} = \{i \in I_1 : W_i = 1, \hat{P}_i \leq \hat{p}(X_i) + \delta_n\}$, indexed by sample size n . $r_{p,n}$ is a nonnegative deterministic sample-size rate for uniform propensity estimation, rather than an integrated prediction loss. In Corollary O5.1: $r_{p,n} \rightarrow 0$. $r_{g,\text{loc},n}$ is a nonnegative deterministic rate sequence indexed by sample size for prediction under both target evaluation-index laws, using the projected estimated index in the population comparator. In Corollary O5.1: $r_{g,\text{loc},n} \rightarrow 0$. $r_{Q,n}$ is a nonnegative deterministic rate sequence indexed by sample size, measuring the interior-grid quantile-fit loss relative to the selected-law target, rather than the canonical boundary law. In Corollary O5.1: $r_{Q,n} \rightarrow 0$. $\delta_n > 0$ is the fixed, sample-size-indexed localization window used in the fitted selection rule $\hat{P} \leq \hat{p}(X) + \delta_n$. In Corollary O5.1: $\delta_n \rightarrow 0$ as $n \rightarrow \infty$. ε_n is the quantile-analysis trimming sequence satisfying $M_Q^{-1} \leq \varepsilon_n < 1/2$ and $\varepsilon_n \rightarrow 0$; it restricts the learner analysis to intermediate quantiles without changing the estimator. M_Q is the integer resolution of the quantile grid $\tau_m = m/M_Q$, $m = 1, \dots, M_Q - 1$, and may vary with n . In Corollary O5.1: $M_Q^{-1} \rightarrow 0$, so the quantile grid resolution diverges with sample size.

Applying the same construction to the reversed system $W^0 = 1 - W$, $P^0 = 1 - P$, $q^0 = 1 - q$, $U^0 = 1 - U$ gives the untreated component, and the two bounds for the full PRTE converge at the maximum of the two componentwise rates; see Section B.3.

5 Simulation

We validate our theoretical results through synthetic experiments and a real-data empirical application. All experiments compare our method (hereafter IVOT) bounds against the moment-relaxation bounds of Mogstad et al. [2018] (hereafter IVMTE). Detailed numerical tables, the full data-generating processes, and additional results are deferred to Section O6 of the Online Appendix.

5.1 Synthetic Experiments

We consider three synthetic settings: one with a continuous instrument, one with a discrete instrument, and one with a discrete instrument and covariates, each using a distinct data-generating process (DGP). Throughout the target is $\theta_\alpha := \mathbb{E}[Y^{q_\alpha} - Y]$ under the policy

$q_\alpha = \text{clip}(p + \alpha, 0, 1)$; the first two designs report $\alpha \in [-0.12, 0.12]$, and the covariate design a window of larger shifts. We give IVMTE a flexible, near-nonparametric u -spline sieve. Section [O6](#) gives the DGPs and the comparison choices; the replication package lists the numerical settings.

Continuous instrument. The continuous-instrument DGP uses a logistic propensity $p(Z) = \text{logistic}(-1 + 2Z)$ with limited support $p(Z) \in [0.27, 0.73]$ and a constant marginal treatment effect $\text{MTE}(u) = 0.5$; since $p(Z) + \alpha$ never leaves $[0, 1]$ on the reported grid, the ground truth is exactly $\theta_\alpha = 0.5\alpha$. Figure [2](#) shows the identified sets at $n = 5,000$, together with the population sharp set, which is closed form here because both potential-outcome means are affine in u ; Section [O6](#) of the Online Appendix derives it and confirms the formula numerically using a discretized linear program built from the primitive problem [\(3.2\)](#). IVOT is tighter at every grid point, by a width ratio ranging from $3.4\times$ at $|\alpha| = 0.12$ to $23.3\times$ at $|\alpha| = 0.01$, and the true θ_α falls inside the IVOT interval at 24 of the 25 grid points. The IVMTE bounds are markedly asymmetric—at $\alpha = 0.05$ the upper bound reaches twice the truth while the lower bound crosses zero. The single miss is at $\alpha = 0.01$, where the sharp set is itself only 4×10^{-4} wide: the unscaled identified-set width is of order α^2 near zero, so small estimation errors can displace the interval away from the truth. Section [O6](#) reports a 20-replication study of the inclusion rate.

Discrete instrument. The discrete-instrument DGP uses $Z \sim \text{Bernoulli}(0.5)$ with two-point propensity score $\{0.25, 0.75\}$, to which the closed-form bounds of Theorem [3.4](#) apply, and a heterogeneous $\text{MTE}(u) = 0.48 + 0.18u$. The outcome support is $[0, 1.08]$ and the exact policy-effect truth is $\theta_\alpha = 0.57\alpha + 0.09\alpha^2$; both are derived in Section [O6](#). Figure [3](#) shows the identified sets at $n = 5,000$ alongside the population sharp set. The truth falls inside the IVOT interval at all 25 grid points, and the IVOT interval is $2.4\text{--}3.6\times$ narrower than IVMTE throughout: at $\alpha = 0.05$ the IVOT interval is $[-0.001, 0.042]$ (width 0.043) versus the IVMTE interval $[-0.051, 0.054]$ (width 0.105). The gap between the two methods is enough to change the qualitative conclusion at large shifts: the estimated IVMTE identified set contains zero at every α , whereas the estimated IVOT identified set excludes zero once $\alpha \geq 0.08$.

Discrete instrument with covariates. The third design adds five continuous covariates to a three-valued instrument, with $\text{MTE}(u, x)$ varying in both arguments, and is built so




Figure 2: Continuous-instrument IVOT and IVMTE identified sets for $\theta_\alpha = \mathbb{E}[Y^{q_\alpha} - Y]$ across policy shifts $\alpha \in [-0.12, 0.12]$ at $n = 5,000$. Grey dotted curves: the population sharp identified set, computed in closed form. Black line: ground truth. Blue and red curves: estimated IVOT and IVMTE bounds (solid = lower, dashed = upper). Shaded ribbon: pointwise 95% backward confidence interval for the IVMTE bounds.




Figure 3: Discrete-instrument IVOT and IVMTE identified sets for $\theta_\alpha = \mathbb{E}[Y^{q_\alpha} - Y]$ across policy shifts $\alpha \in [-0.12, 0.12]$ at $n = 5,000$. Grey dotted curves: the population sharp identified set, computed in closed form. Black line: ground truth. Blue and red curves: estimated IVOT and IVMTE bounds (solid = lower, dashed = upper). Shaded ribbons: pointwise 95% confidence intervals for the bounds (blue = IVOT delta method, red = IVMTE backward CI).

that ?? 4.2 holds exactly with $c_{\text{gap}} = 0.04$ and every component of the score is active on the reported window $\alpha \in [0.195, 0.255]$, at $n = 4,000$. As in the two designs above the comparison is against IVMTE in its nonparametric form, so the u -spline sieve is interacted with the covariates and the MTR is left unrestricted in shape as well as in level. Section O6 of the Online Appendix gives the DGP, the policy window, the comparison choices, and the numerical tables.

Figure 4 shows the estimated IVOT interval tracking the population sharp set—closed form here up to the expectation over X —and containing the truth at all 11 grid points, at 0.79–1.06 times the sharp width. IVOT is narrower than IVMTE by 1.10–1.46 $\times$ at degree 3 and 1.34–1.78 $\times$ at degree 7, and the difference is qualitative: at $\alpha = 0.219$ the sharp set is $[0.019, 0.134]$ and IVOT $[0.015, 0.137]$, while IVMTE at degree 7 returns $[-0.012, 0.156]$. Both that interval and its pointwise 95% backward confidence band contain zero throughout the policy-shift grid, so the moment relaxation loses the sign of the policy effect.

The estimation theory of Section 4 is tested separately by Monte Carlo. Over 400 replications at each of five sample sizes from 1,000 to 16,000 the sampling standard deviation contracts at the parametric rate, the estimated standard errors closely match the Monte Carlo standard deviations at the larger sample sizes—their ratios rise to 1.00 and 1.02 for the two bounds—and the level-set recovery step returns the correct configuration in every replication once $n \geq 8,000$. Section O6 of the Online Appendix gives the full Monte Carlo results.




Figure 4: Covariate design: IVOT and IVMTE identified sets for $\theta_\alpha = \mathbb{E}[Y^{q_\alpha} - Y]$ at $n = 4,000$, with five continuous covariates and a three-valued instrument. Grey dotted curves: the population sharp identified set, computed in closed form. Black line: ground truth. Blue: IVOT (solid = lower, dashed = upper), with its pointwise 95% interval shaded. Orange and red: IVMTE with the u -spline sieve interacted with the covariates, at degrees 3 and 7. The red shading is the pointwise 95% backward bootstrap confidence band for the degree-7 IVMTE bounds. As in Figure 5, the richer sieve gives the wider bound.

5.2 Effect of Price Subsidies for Bed Nets

Background and setup. We apply our method to the bed net subsidy experiment of Dupas [2014]. The policy question is how the price of a bed net affects its usage, which is what a PRTE measures. For $n = 1,078$ Kenyan households the data contain the offered price as instrument Z (17 values spanning 0–250 KSh, randomized across distribution points), purchase as treatment W , and one-year follow-up usage as outcome Y . We compare two fixed-price policies: the baseline assigns every household the reference price $z_0 = 150$ KSh, and for each $\alpha \in [0.05, 0.62]$ the alternative raises the purchase propensity from $\hat{p}(z_0)$ to $q_\alpha := \min(\hat{p}(z_0) + \alpha, 1)$. The target is $\text{PRTE}_\alpha = \mathbb{E}[Y^{q_\alpha} - Y^{p(z_0)}]/\alpha$; because both policies are constant in Z , the mean policy shift equals α exactly, so the denominator is the shift size rather than an estimated quantity. The baseline is the fixed-price policy $Y^{p(z_0)}$, not the observed Y under randomized prices, and the grid is itself constructed from the first stage, its endpoint $\alpha_{\max} \approx 0.621$ being the estimated propensity at zero price minus the estimated baseline propensity. Because the instrument is discrete, the closed-form bounds of Theorem 3.4 apply directly to the numerator, after which we divide by α . For IVMTE, we approximate the marginal treatment response (MTR) functions with degree-10 and degree-20 u -spline sieves on $[0, 1]$. Full data, propensity-estimation, and policy-construction details are in Section O6.

Results. Figure 5 plots the IVOT and IVMTE identified sets for PRTE_α across subsidy levels. IVOT consistently yields a tighter identified set than IVMTE across the reported range of α and at both spline degrees. At the smallest subsidy ($\alpha = 0.05$), the IVOT interval is 3.3 times narrower than IVMTE at degree 10 and 5.3 times narrower at degree 20.

Both methods indicate a positive policy effect across most of the subsidy range—the IVOT lower bound is positive throughout, and the degree-10 IVMTE lower bound is positive for all α , with the more conservative degree-20 lower bound turning positive once $\alpha \geq 0.14$. IVOT achieves near-point identification (width < 0.05) for the majority of α values. Those narrow identified sets do not translate one-for-one into narrower confidence intervals. Section O6 gives more details and explanations for the wide confidence intervals.

6 Conclusion

We connected the partial identification of PRTEs in IV models to OT: formulating it as a CCOT problem over the joint laws compatible with the observed data, we showed that the




Figure 5: IVOTE versus IVMTE identified sets for $\text{PRTE}_\alpha = \mathbb{E}[Y^{q_\alpha} - Y^{p(z_0)}]/\alpha$ on the bed net data from Dupas [2014]. IVMTE is shown for degree-10 (orange) and degree-20 (red) u -spline sieves, with solid curves the lower bounds and dashed the upper; degree 20 widens the bound, reflecting the more nonparametric MTR class. Shaded ribbons: pointwise 95% confidence intervals for the estimated bounds (red = IVMTE deg. 20, blue = IVOTE). IVOTE substantially tightens the identified set and achieves near-point identification at several subsidy levels.

multidimensional optimization reduces analytically to one-dimensional OT problems with product costs, yielding closed-form sharp bounds and bypassing the numerical optimization that existing computational approaches require. In a companion note we replace threshold crossing by a conditional quantile representation of the selection equation and obtain sharp closed-form bounds for continuous and multi-valued treatments, under a rank-overlap condition in the strictly monotonic case and an algebraic condition on the instrument-induced latent intervals in the ordered-choice case. The estimation theory developed here is for binary treatments, delivering $\sqrt{n}$ -consistency and asymptotic normality with discrete instruments under the stated regularity conditions, and nonparametric convergence rates with continuous ones.

The central insight, demonstrated theoretically and empirically, is that preserving full distributional information is tractable and valuable for tight partial identification: moment-relaxation approaches reduce the observed data to conditional expectations and can produce substantially wider bounds, whereas CCOT enforces compatibility with the entire observed distribution while keeping the optimization transparent. We expect this to extend beyond the IV setting studied here.

Acknowledgments. Jose Blanchet gratefully acknowledges support from the U.S. Department of Defense through the Air Force Office of Scientific Research award FA9550-20-1-0397 and ONR award 1398311, as well as support from NSF grants 2229012, 2312204, and 2403007. Vasilis Syrgkanis gratefully acknowledges support from NSF Award IIS-2337916, an Amazon Research Award, and a Google Research Award.

References

- Joshua D. Angrist and Guido W. Imbens. Identification and estimation of local average treatment effects. Technical Working Paper 118, National Bureau of Economic Research, Cambridge, MA, 1995.
- Susan Athey, Julie Tibshirani, and Stefan Wager. Generalized random forests. *The Annals of Statistics*, 47(2):1148–1178, 2019.
- Vahid Balazadeh Meresht, Vasilis Syrgkanis, and Rahul G. Krishnan. Partial identification of treatment effects with implicit generative models. In *Advances in Neural Informa-*

- tion Processing Systems 35 (NeurIPS)*, 2022. URL <https://openreview.net/pdf?id=8cUGfg-zUnh>.
- 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.
- Alexandre Belloni, Victor Chernozhukov, Denis Chetverikov, and Kengo Kato. Some new asymptotic theory for least squares series: Pointwise and uniform results. *Journal of Econometrics*, 186(2):345–366, 2015. doi: 10.1016/j.jeconom.2015.02.014.
- Alexandre Belloni, Victor Chernozhukov, Denis Chetverikov, and Iván Fernández-Val. Conditional quantile processes based on series or many regressors. *Journal of Econometrics*, 213(1):4–29, 2019.
- Pedro Carneiro, James J Heckman, and Edward J Vytlacil. Evaluating marginal policy changes and the average effect of treatment for individuals at the margin. *Econometrica*, 78(1):377–394, 2010.
- Pedro Carneiro, James J Heckman, and Edward J Vytlacil. Estimating marginal returns to education. *American Economic Review*, 101(6):2754–2781, 2011.
- Victor Chernozhukov, Han Hong, and Elie Tamer. Estimation and confidence regions for parameter sets in econometric models. *Econometrica*, 75(5):1243–1284, 2007.
- Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 2018. doi: 10.1111/ectj.12097.
- Guilherme Duarte, Noam Finkelstein, Dean Knox, Jonathan Mummolo, and Ilya Shpitser. An automated approach to causal inference in discrete settings. *Journal of the American Statistical Association*, 119(547):1778–1793, 2024.
- Pascaline Dupas. Short-run subsidies and long-run adoption of new health products: Evidence from a field experiment. *Econometrica*, 82(1):197–228, 2014.
- Yanqin Fan, Emmanuel Guerre, and Dongming Zhu. Partial identification of functionals of the joint distribution of “potential outcomes”. *Journal of Econometrics*, 197(1):42–59, 2017.

- Yanqin Fan, Brendan Pass, and Xuetao Shi. Partial identification in moment models with incomplete data via optimal transport, 2025.
- Alfred Galichon. *Optimal transport methods in economics*. Princeton University Press, 2016.
- Florian Gunsilius. A primer on optimal transport for causal inference with observational data, 2025.
- Wenshuo Guo, Mingzhang Yin, Yixin Wang, and Michael I. Jordan. Partial identification with noisy covariates: A robust optimization approach. In *Proceedings of CLear 2022*, 2022. URL <https://proceedings.mlr.press/v177/guo22a/guo22a.pdf>. Also available as arXiv:2202.10665.
- Sukjin Han and Shenshen Yang. A computational approach to identification of treatment effects for policy evaluation. *Journal of Econometrics*, 240(1):105680, 2024.
- James J Heckman and Edward Vytlacil. Structural equations, treatment effects, and econometric policy evaluation. *Econometrica*, 73(3):669–738, 2005. doi: 10.1111/j.1468-0262.2005.00594.x.
- James J Heckman and Edward J Vytlacil. Local instrumental variables and latent variable models for identifying and bounding treatment effects. *Proceedings of the National Academy of Sciences*, 96(8):4730–4734, 1999. doi: 10.1073/pnas.96.8.4730.
- James J Heckman and Edward J Vytlacil. Instrumental variables, selection models, and tight bounds on the average treatment effect. In *Econometric Evaluation of Labour Market Policies*, pages 1–15. Springer, 2001.
- James J Heckman, Sergio Urzua, and Edward Vytlacil. Understanding instrumental variables in models with essential heterogeneity. *The Review of Economics and Statistics*, 88(3): 389–432, 2006.
- Guido W Imbens and Donald B Rubin. Estimating outcome distributions for compliers in instrumental variables models. *The Review of Economic Studies*, 64(4):555–574, 1997.
- Wenlong Ji, Lihua Lei, and Asher Spector. Model-agnostic covariate-assisted inference on partially identified causal effects, 2024.
- Edward H. Kennedy. Towards optimal doubly robust estimation of heterogeneous causal effects. *Electronic Journal of Statistics*, 17(2):3008–3049, 2023.

- Ying-Ying Lee. Partial mean processes with generated regressors: Continuous treatment effects and nonseparable models. *arXiv preprint arXiv:1811.00157*, 2018.
- Alexander W Levis, Matteo Bonvini, Zhenghao Zeng, Luke Keele, and Edward H Kennedy. Covariate-assisted bounds on causal effects with instrumental variables. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 87(5):1508–1527, 2025.
- Sirui Lin, Zijun Gao, Jose Blanchet, and Peter Glynn. Estimation of optimal causal bounds via covariate-assisted optimal transport. *arXiv preprint arXiv:2506.00257*, 2025a.
- Sirui Lin, Zijun Gao, José Blanchet, and Peter Glynn. Tightening causal bounds via covariate-aware optimal transport. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 37815–37845. PMLR, 2025b. URL <https://proceedings.mlr.press/v267/lin25f.html>.
- Enno Mammen, Christoph Rothe, and Melanie Schienle. Nonparametric regression with nonparametrically generated covariates. *The Annals of Statistics*, 40(2):1132–1170, 2012. doi: 10.1214/12-AOS995.
- Charles F. Manski. Nonparametric bounds on treatment effects. *American Economic Review*, 80(2):319–323, 1990.
- Charles F. Manski. Monotone treatment response. *Econometrica*, 65(6):1311–1334, 1997.
- Charles F. Manski and John V. Pepper. Monotone instrumental variables: With an application to the returns to schooling. *Econometrica*, 68(4):997–1010, 2000.
- Philip Marx. Sharp bounds in the latent index selection model. *Journal of Econometrics*, 238(2):105561, 2024. doi: 10.1016/j.jeconom.2023.105561.
- Magne Mogstad and Alexander Torgovitsky. Identification and extrapolation of causal effects with instrumental variables. *Annual Review of Economics*, 10:577–613, 2018.
- Magne Mogstad, Andres Santos, and Alexander Torgovitsky. Using instrumental variables for inference about policy relevant treatment parameters. *Econometrica*, 86(5):1589–1619, 2018. doi: 10.3982/ECTA15463.
- Ismael Mourifié and Yuanyuan Wan. Testing local average treatment effect assumptions. *Review of Economics and Statistics*, 99(2):305–313, 2017. doi: 10.1162/REST_a_00622.

- Daniel Ober-Reynolds. Estimating functionals of the joint distribution of potential outcomes with optimal transport. *arXiv preprint arXiv:2311.09435*, 2023.
- Elie Tamer. Partial identification in econometrics. *Annual Review of Economics*, 2(1):167–195, 2010.
- Jiyuan Tan, Jose Blanchet, and Vasilis Syrgkanis. Consistency of neural causal partial identification. *Advances in Neural Information Processing Systems*, 37:68956–68999, 2024.
- Cédric Villani. *Topics in optimal transportation*, volume 58 of *Graduate Studies in Mathematics*. American Mathematical Society, reprint edition, 2021. doi: 10.1090/gsm/058.
- Edward Vytlačil. Independence, monotonicity, and latent index models: An equivalence result. *Econometrica*, 70(1):331–341, 2002.

A Proofs in Section 3

A.1 Proof of Sharpness

Proof of Proposition 3.1. Direct direction. Let $\tilde{\pi}$ be a probability measure on $\mathcal{Y}^2 \times [0, 1] \times \mathcal{X}$ whose $(Y(w), U, X)$ -marginals satisfy $(\tilde{\pi}_0, \tilde{\pi}_1) \in \Gamma(\mathbb{P}_{\text{obs}})$. By the constraints in (3.2), each $\tilde{\pi}_w$ is a probability measure on $\mathcal{Y} \times [0, 1] \times \mathcal{X}$ satisfying:

$$(C1) \quad \int_0^{p(z,x)} \tilde{\pi}_1(dy \mid u, x) du = H_{1,x}(p(z, x), dy), \text{ for } \mathbb{P}_{\text{obs}}\text{-a.e. } (z, x);$$

$$(C2) \quad \int_{p(z,x)}^1 \tilde{\pi}_0(dy \mid u, x) du = H_{0,x}(p(z, x), dy), \text{ for } \mathbb{P}_{\text{obs}}\text{-a.e. } (z, x);$$

$$(C3) \quad \tilde{\pi}_w(du, dx) = du \mathbb{P}_{\text{obs}}(dx) \text{ for } w \in \{0, 1\}.$$

Since $\mathcal{Y}^2 \subseteq \mathbb{R}^2$ is a Polish space and the (U, X) -marginal of $\tilde{\pi}$ is $du \mathbb{P}_{\text{obs}}(dx)$ by (C3) applied to either marginal, $\tilde{\pi}$ admits a regular conditional kernel $\tilde{\pi}(dy_0, dy_1 \mid u, x)$ for $(du \otimes \mathbb{P}_{\text{obs}}(dx))$ -a.e. (u, x) , whose $Y(0)$ - and $Y(1)$ -marginals recover $\tilde{\pi}_0(dy_0 \mid u, x)$ and $\tilde{\pi}_1(dy_1 \mid u, x)$ respectively.

Construction. We construct the joint distribution $\mathbb{P}$ of $(Y(0), Y(1), U, X, Z, W)$ on the product space $\mathcal{Y}^2 \times [0, 1] \times \mathcal{X} \times \mathcal{Z} \times \{0, 1\}$ via the following disintegration:

$$\mathbb{P}(dy_0, dy_1, du, dx, dz, \{w\}) = \tilde{\pi}(dy_0, dy_1 \mid u, x) du \mathbb{P}_{\text{obs}}(dz \mid x) \mathbb{P}_{\text{obs}}(dx) \delta_{\mathbb{1}(u \leq p(z,x))}(\{w\}), \quad (\text{A.1})$$

where $\delta_{\mathbb{1}(u \leq p(z,x))}(\{w\})$ is the Dirac mass placing all weight on $w = \mathbb{1}(u \leq p(z, x))$, enforcing the threshold-crossing selection rule deterministically. Crucially, $(Y(0), Y(1))$ are drawn jointly from $\tilde{\pi}(\cdot \mid u, x)$, so the coupling prescribed by $\tilde{\pi}$ is preserved conditional on (U, X) ; the independent coupling $\tilde{\pi}_0(dy_0 \mid u, x) \tilde{\pi}_1(dy_1 \mid u, x)$ is one feasible choice among many. The measure in (A.1) is well-defined because the kernels are measurable and all component measures are probability measures.

Proof of (i): joint marginal distribution. We show that the $(Y(0), Y(1), U, X)$ -marginal of $\mathbb{P}$ equals $\tilde{\pi}$. Marginalizing (A.1) over (Z, W) :

$$\begin{aligned} \mathbb{P}(dy_0, dy_1, du, dx) &= \tilde{\pi}(dy_0, dy_1 \mid u, x) du \mathbb{P}_{\text{obs}}(dx) \underbrace{\sum_{w \in \{0,1\}} \int_{\mathcal{Z}} \delta_{\mathbb{1}(u \leq p(z,x))}(\{w\}) \mathbb{P}_{\text{obs}}(dz \mid x)}_{=1} \\ &= \tilde{\pi}(dy_0, dy_1 \mid u, x) du \mathbb{P}_{\text{obs}}(dx). \end{aligned}$$

By (C3) the (U, X) -marginal of $\tilde{\pi}$ is $du \mathbb{P}_{\text{obs}}(dx)$, so the right-hand side equals

$$\tilde{\pi}(dy_0, dy_1, du, dx),$$

establishing (i).

Proof of (ii): structural assumptions. We establish each part of ?? 2.1.

- **Consistency.** Consistency holds trivially by setting the observed outcome to be $Y = WY(1) + (1 - W)Y(0)$.
- **Conditional instrumental exogeneity.** By (A.1), conditional on $X = x$, the variables $(Y(0), Y(1), U)$ are drawn from $\tilde{\pi}(dy_0, dy_1 \mid u, x) du$ and Z is drawn independently from $\mathbb{P}_{\text{obs}}(dz \mid x)$. These appear as independent factors in (A.1), so $Z \perp\!\!\!\perp (Y(0), Y(1), U) \mid X$ under $\mathbb{P}$.
- **Threshold crossing.** The selection equation $W = \mathbb{1}(U \leq p(Z, X))$ holds $\mathbb{P}$ -a.s. by the Dirac mass in (A.1). Moreover, (C3) gives $U \mid X \sim \text{Unif}(0, 1)$. Conditional instrumental exogeneity then yields $\mathbb{P}(W = 1 \mid Z = z, X = x) = p(z, x)$ for $\mathbb{P}_{Z,X}$ -almost every (z, x) .
- **Common instrument support.** The construction preserves the observed conditional law $\mathbb{P}_{\text{obs}}(dz \mid x)$. Because the proposition assumes that $\mathbb{P}_{\text{obs}}$ is generated by a model satisfying ?? 2.1, its common-support property is inherited by the constructed model.

Proof of (iii): observational equivalence. We must show that the distribution of the observables (Y, Z, X, W) induced by $\mathbb{P}$ coincides with $\mathbb{P}_{\text{obs}}$. Since the marginal of (Z, X) under $\mathbb{P}$ equals $\mathbb{P}_{\text{obs}}(Z, X)$ by construction, it suffices to show that the conditional laws of (Y, W) given (Z, X) agree.

Fix (z, x) and consider the event $W = 1$. Under $\mathbb{P}$, marginalizing (A.1) over $(Y(0), Y(1))$ reduces the joint kernel to its $Y(1)$ -marginal $\tilde{\pi}_1(\cdot \mid u, x)$; using $W = 1 \Leftrightarrow U \leq p(z, x)$ and $Y = Y(1)$ when $W = 1$,

$$\mathbb{P}(Y \in A, W = 1 \mid Z = z, X = x) = \int_0^{p(z, x)} \tilde{\pi}_1(A \mid u, x) du.$$

By constraint (C1), this equals $H_{1,x}(p(z, x), A)$, which is the observed conditional subdistribution $\mathbb{P}_{\text{obs}}(Y \in A, W = 1 \mid Z = z, X = x)$ for $\mathbb{P}_{Z,X}$ -almost every (z, x) . Similarly, for $W = 0$ we have $Y = Y(0)$ and $U > p(z, x)$, so marginalizing over $(Y(0), Y(1))$ leaves the $Y(0)$ -marginal $\tilde{\pi}_0$ and

$$\mathbb{P}(Y \in A, W = 0 \mid Z = z, X = x) = \int_{p(z, x)}^1 \tilde{\pi}_0(A \mid u, x) du,$$

which equals $H_{0,x}(p(z, x), A)$ by constraint (C2), and hence equals the corresponding observed conditional subdistribution almost surely. Thus the conditional laws of $(Y, W) \mid (Z, X)$ under $\mathbb{P}$ and $\mathbb{P}_{\text{obs}}$ agree for both values of W , which proves (iii).

Converse direction. Suppose $(Y(0), Y(1), U, X, Z, W)$ is a joint distribution satisfying (ii)–(iii), and let $\tilde{\pi}$ denote its $(Y(0), Y(1), U, X)$ -marginal. Fix $w \in \{0, 1\}$. The (U, X) -marginal of $\tilde{\pi}_w$ equals the (U, X) -marginal of $\tilde{\pi}$; by ?? 2.1 (threshold crossing), $U \mid X \sim \text{Unif}(0, 1)$ and X has marginal $\mathbb{P}_{\text{obs}}(dx)$ under (iii), so $\tilde{\pi}_w(du, dx) = du \mathbb{P}_{\text{obs}}(dx)$, establishing (C3). For (C1), using consistency and the selection rule $W = \mathbf{1}(U \leq p(Z, X))$ together with conditional instrumental exogeneity,

$$\begin{aligned} \mathbb{P}_{\text{obs}}(Y \in A, W = 1 \mid Z = z, X = x) &= \mathbb{P}(Y(1) \in A, U \leq p(z, x) \mid Z = z, X = x) \\ &= \int_0^{p(z, x)} \tilde{\pi}_1(A \mid u, x) du, \end{aligned}$$

where the second equality uses $Z \perp\!\!\!\perp (Y(1), U) \mid X$ and $U \mid X \sim \text{Unif}(0, 1)$. Property (iii) identifies the left-hand side with the observational conditional distribution. By the propensity-score sufficiency argument preceding (3.1), its canonical propensity-indexed version is $H_{1,x}(p(z, x), A)$, yielding (C1). The argument for (C2) is symmetric. Therefore $(\tilde{\pi}_0, \tilde{\pi}_1) \in \Gamma(\mathbb{P}_{\text{obs}})$, completing the converse. $\square$

A.2 Proof of Theorem 3.8

Proof of Theorem 3.8. We prove the sharp lower bound; the upper bound follows by replacing the countermonotonic coupling with the comonotonic one and swapping $y_{\min}$ and $y_{\max}$ in the trivial bound term.

Disintegration over X . Recall that the $W = 1$ side minimization problem (3.3) reads $\min_{\pi_1 \in \Gamma_1(\mathbb{P}_{\text{obs}})} \mathbb{E}_{\pi_1}[Y(1) \omega(X, U)]$, where $\Gamma_1(\mathbb{P}_{\text{obs}})$ is the set of probability measures on $\mathcal{Y} \times [0, 1] \times \mathcal{X}$ satisfying

$$\int_0^{p(z, x)} \pi_1(dy \mid u, x) du = H_{1,x}(p(z, x), dy), \quad \mathbb{P}_{\text{obs}}\text{-a.e. } (z, x), \quad (\text{A.2})$$

$$\pi_1(du, dx) = du \mathbb{P}_{\text{obs}}(dx). \quad (\text{A.3})$$

The observational constraint (A.2) couples $\pi_1(dy \mid u, x)$ across different u only through the integral $\int_0^{p(z, x)} \pi_1(dy \mid u, x) du$, which for fixed x involves only the kernel values at that same




Figure 6: Decomposition of the unit interval $[0, 1]$ for the sharp lower bound in the general instrument setting (Theorem 3.7), illustrated with $K = 2$ in the order OT / LIV / OT / trivial: the non-degenerate LIV interval $L_1 = [\underline{p}_1, \bar{p}_1]$ (teal strip) lies between two OT gaps, and the degenerate LIV singleton $L_2 = \{p_2\}$ (teal line) sits at the boundary of the trivial tail. **OT gaps** $G_0 = (0, \underline{p}_1)$ and $G_1 = (\bar{p}_1, p_2)$: only the aggregate distributions $\mu_{1,0}, \mu_{1,1}$ are identified, so the tight bound requires countermonotone quantile coupling. **LIV interval** L_1 : the propensity score sweeps a full subinterval, so the standard LIV derivative point-identifies $\mathbb{E}[Y(1) | U = u]$ exactly; the singleton L_2 contributes a Lebesgue-null integral and drops out. **Trivial tail** $(p_2, 1]$: $Y(1)$ is never observed; the lower bound assigns $y_{\min}$ where $\omega(u) \geq 0$ and $y_{\max}$ where $\omega(u) < 0$. Collapsing every LIV interval to a singleton recovers the discrete case (Figure 1); taking $K = 1$ with $\underline{p}_1 < \bar{p}_1$ and no trailing singleton recovers the connected-continuous case.

x . Hence the constraints decouple across x . For every feasible π_1 , disintegration with respect to the (U, X) -marginal (A.3) gives

$$\mathbb{E}_{\pi_1}[Y(1)\omega(X, U)] = \mathbb{E}_X \left[\int_0^1 \left(\int_{\mathcal{Y}} y \pi_{1,X}(\mathrm{d}y \mid u) \right) \omega(X, u) \mathrm{d}u \right]. \quad (\text{A.4})$$

The conditional kernel $\pi_{1,x}$ belongs to the conditional feasible set $\Gamma_{1,x}$ for $\mathbb{P}_X$ -a.e. x . Writing $B_Y = \max\{|y_{\min}|, |y_{\max}|\}$, the absolute value of its conditional objective is bounded by $B_Y \int_0^1 |\omega(x, u)| \mathrm{d}u$, whose expectation is finite by ?? 2.2. We derive a lower bound for every such conditional kernel and then attain the integrated bound with an explicit jointly measurable kernel.

Remark (notation). $\Gamma_{1,x}$ is the conditional treated feasible class defined for $\mathbb{P}_X$ -almost every x , with its members understood through their Lebesgue-almost-everywhere kernel values in u . In Theorem 3.8: For $\mathbb{P}_X$ -almost every x , the observational constraint equivalently holds as $\int_0^p \pi_{1,x}(\mathrm{d}y \mid u) \mathrm{d}u = H_{1,x}(p, \mathrm{d}y)$ for every $p \in p(\mathcal{Z}, x)$, using the canonical continuous version of $H_{1,x}$. B_Y is a fixed, finite, nonnegative constant determined by $y_{\min}$ and $y_{\max}$. In Theorem 3.8: $B_Y = \max\{|y_{\min}|, |y_{\max}|\}$ bounds the absolute magnitude of feasible outcomes. In Theorem 4.7, Corollary O5.1: $B_Y = y_{\max} - y_{\min}$ is the outcome range, used to bound differences between values in $[y_{\min}, y_{\max}]$.

Maximal-interval decomposition for fixed x . Fix $x \in \mathcal{X}$. Under the standing compactness of $\mathcal{Z}$ and ?? 3.6, the map $z \mapsto p(z, x)$ is continuous on each connected component, so the image of each component is a closed interval (possibly degenerate). Merging overlaps and ordering from left to right, the conditional range admits the maximal-interval decomposition (3.16),

$$p(\mathcal{Z}, x) = \bigcup_{k=1}^{K_x} [\underline{p}_k(x), \bar{p}_k(x)], \quad 0 \leq \underline{p}_1(x) \leq \bar{p}_1(x) < \cdots < \underline{p}_{K_x}(x) \leq \bar{p}_{K_x}(x) \leq 1,$$

with boundary conventions $\bar{p}_0(x) := 0$ and $\underline{p}_{K_x+1}(x) := 1$. Adopting the notation of Section 3.2.2, these sets partition $[0, 1]$ up to their Lebesgue-null endpoints into the conditional LIV intervals $L_k(x) = [\underline{p}_k(x), \bar{p}_k(x)]$ ($k = 1, \dots, K_x$), the OT gaps $G_k(x) = (\bar{p}_k(x), \underline{p}_{k+1}(x))$ ($k = 0, \dots, K_x - 1$), and the trivial tail $(\bar{p}_{K_x}(x), 1]$. For fixed x , both $p \mapsto \int_0^p \pi_{1,x}(\mathrm{d}y \mid u) \mathrm{d}u$ and the canonical map $p \mapsto H_{1,x}(p, \mathrm{d}y)$ are continuous in total variation. Because they agree for the conditional law of P -almost every p by (A.2), they agree at every point of the conditional support. Hence the observational constraint at the x -slice reads, for every $p \in p(\mathcal{Z}, x)$,

$$\int_0^p \pi_{1,x}(\mathrm{d}y \mid u) \mathrm{d}u = H_{1,x}(p, \mathrm{d}y), \quad (\text{A.5})$$

where $H_{1,x}(0, \cdot) = 0$.

The OT gaps $G_k(x)$. For each $k = 0, \dots, K_x - 1$ such that $\bar{p}_k(x) < \underline{p}_{k+1}(x)$, differencing (A.5) at $p = \underline{p}_{k+1}(x)$ and at $p = \bar{p}_k(x)$ isolates the integral over the nonempty gap $G_k(x) = (\bar{p}_k(x), \underline{p}_{k+1}(x))$:

$$\int_{G_k(x)} \pi_{1,x}(\mathrm{d}y \mid u) \mathrm{d}u = (\underline{p}_{k+1}(x) - \bar{p}_k(x)) \mu_{1,k|x}(\mathrm{d}y), \quad (\text{A.6})$$

where $\mu_{1,k|x}$ is the conditional gap-identified measure in (3.17). Dividing by $|G_k(x)| = \underline{p}_{k+1}(x) - \bar{p}_k(x) > 0$, the marginal distribution of $Y(1)$ conditional on $U \in G_k(x)$ is fixed to $\mu_{1,k|x}$, while the U -marginal is $\nu_{k|x} := \text{Unif}(G_k(x))$. The joint distribution of $(Y(1), U)$ over $G_k(x)$ is not otherwise pinned down, so the contribution of this gap to the inner minimization is a standard 1D optimal transport problem. Empty gaps are omitted and contribute zero:

$$|G_k(x)| \min_{\gamma_k \in \Pi(\mu_{1,k|x}, \nu_{k|x})} \int_{\mathcal{Y} \times G_k(x)} y \cdot \omega(x, u) \mathrm{d}\gamma_k(y, u). \quad (\text{A.7})$$

Remark (notation). $\nu_{k|x}$ is the normalized Lebesgue probability law on $G_k(x)$, equivalently the law of U conditional on $X = x$ and $U \in G_k(x)$, for almost every x and $k = 0, \dots, K_x - 1$ with $|G_k(x)| > 0$.

The LIV intervals $L_k(x)$. Fix a non-degenerate $L_k(x)$ (with $\underline{p}_k(x) < \bar{p}_k(x)$) and the countable measure-determining class $\mathcal{A} := \{\mathcal{Y} \cap (-\infty, r] : r \in \mathbb{Q}\} \cup \{\mathcal{Y}\}$. For every $A \in \mathcal{A}$, the constraint (A.5) gives the scalar identity

$$H_{1,x}(p, A) = H_{1,x}(\underline{p}_k(x), A) + \int_{\underline{p}_k(x)}^p \pi_{1,x}(A \mid u) \mathrm{d}u, \quad p \in L_k(x).$$

Thus $p \mapsto H_{1,x}(p, A)$ is absolutely continuous on $L_k(x)$, and the scalar fundamental theorem of calculus yields

$$\pi_{1,x}(A \mid u) = \left. \frac{\partial}{\partial p} H_{1,x}(p, A) \right|_{p=u}, \quad A \in \mathcal{A}, \quad u \in L_k(x) \text{ a.e.} \quad (\text{A.8})$$

The countable class $\mathcal{A}$ lets us pass from these scalar identities to a statement about the conditional law. For each fixed A , (A.8) holds outside an A -dependent Lebesgue-null set; countability permits us to remove their union and obtain one exceptional set valid for every $A \in \mathcal{A}$. Using all Borel subsets of $\mathcal{Y}$ directly would not justify this step, because an uncountable union of null sets need not be null. Since $\mathcal{A}$ nevertheless determines probability measures, (A.8) identifies the conditional distribution of $Y(1)$ given $(U, X) = (u, x)$ for almost

every $u \in L_k(x)$. For the objective, set $m_{1,x}(u) := \int_{\mathcal{Y}} y \pi_{1,x}(\mathrm{d}y \mid u)$. Integrating (A.5) against the bounded scalar function y gives

$$g_1^c(p, x) = g_1^c(\underline{p}_k(x), x) + \int_{\underline{p}_k(x)}^p m_{1,x}(u) \mathrm{d}u, \quad p \in L_k(x).$$

Consequently, $g_1^c(\cdot, x)$ is absolutely continuous on this LIV interval and $\partial_u g_1^c(u, x) = m_{1,x}(u)$ almost everywhere. Every feasible kernel therefore has the same contribution:

$$\int_{\underline{p}_k(x)}^{\bar{p}_k(x)} \left(\int_{\mathcal{Y}} y \pi_{1,x}(\mathrm{d}y \mid u) \right) \omega(x, u) \mathrm{d}u = \int_{\underline{p}_k(x)}^{\bar{p}_k(x)} \frac{\partial}{\partial u} g_1^c(u, x) \omega(x, u) \mathrm{d}u. \quad (\text{A.9})$$

Moreover, $|y \omega(x, u)| \leq B_Y |\omega(x, u)|$. By ?? 2.2, $\int_0^1 |\omega(x, u)| \mathrm{d}u < \infty$ for $\mathbb{P}_X$ -a.e. x , so the integral in (A.9) is well-defined. Degenerate $L_k(x)$ ($\underline{p}_k(x) = \bar{p}_k(x)$) contribute a Lebesgue-null integral and drop out.

Remark (notation). The notation $m(X; \eta)$ denotes the binary lower-bound integrand, whereas bare m in u_m or τ_m denotes a grid index. In Lemma D.1, Theorem 4.7: In the binary-outcome case, $m(X; \eta)$ is the displayed integrand with nuisance argument η ; evaluating at η_0 gives the population target integrand. In Corollary O5.1: m is an integer grid index: $m = 0, \dots, M_g$ for $u_m = m/M_g$ and $m = 1, \dots, M_Q - 1$ for the fitted quantiles at $\tau_m = m/M_Q$.

Compatibility of the region-wise constraints. We show that the region-wise constraints (A.6) and (A.8) are jointly equivalent to the original constraint family (A.5). Fix $p \in p(\mathcal{Z}, x)$ and let k^* be the unique index with $p \in L_{k^*}(x)$. Then

$$\begin{aligned} \int_0^p \pi_{1,x}(\mathrm{d}y \mid u) \mathrm{d}u &= \sum_{k=0}^{k^*-1} \underbrace{\int_{G_k(x)} \pi_{1,x}(\mathrm{d}y \mid u) \mathrm{d}u}_{\text{fixed by (A.6)}} + \sum_{k=1}^{k^*-1} \underbrace{\int_{L_k(x)} \pi_{1,x}(\mathrm{d}y \mid u) \mathrm{d}u}_{\text{determined by (A.8)}} \\ &\quad + \underbrace{\int_{\underline{p}_{k^*}(x)}^p \pi_{1,x}(\mathrm{d}y \mid u) \mathrm{d}u}_{\text{determined by (A.8)}}. \end{aligned}$$

For every $A \in \mathcal{A}$, the telescoping definition of $\mu_{1,k|x}$ in (3.17) and the scalar fundamental theorem of calculus applied to (A.8) show that this sum evaluated at A equals $H_{1,x}(p, A)$. Since $\mathcal{A}$ determines finite measures, the sum equals $H_{1,x}(p, \mathrm{d}y)$, so (A.5) is satisfied. Conversely, the gap constraints (A.6) and the scalar LIV identities (A.8) follow from (A.5). Hence the constraints on the gaps, LIV intervals, and trivial tail are decoupled, and the conditional objective decomposes across regions.

The trivial tail $(\bar{p}_{K_x}(x), 1]$. For $u > \bar{p}_{K_x}(x)$, no value of the instrument produces a propensity score exceeding $\bar{p}_{K_x}(x)$, so $Y(1)$ is never observed. The constraint family (A.5) imposes no restriction on $\pi_{1,x}(\mathrm{d}y \mid u)$ there. Because $y_{\min}$ and $y_{\max}$ are the extrema of the compact outcome domain, the pointwise minimum of $\int_{\mathcal{Y}} y \pi_{1,x}(\mathrm{d}y \mid u) \cdot \omega(x, u)$ is attained by $\pi_{1,x}(\cdot \mid u) = \delta_{y_{\min}}$ when $\omega(x, u) \geq 0$ and by $\delta_{y_{\max}}$ when $\omega(x, u) < 0$:

$$\begin{aligned} \min \int_{\bar{p}_{K_x}(x)}^1 \left(\int_{\mathcal{Y}} y \pi_{1,x}(\mathrm{d}y \mid u) \right) \omega(x, u) \mathrm{d}u \\ = \int_{\bar{p}_{K_x}(x)}^1 \left(y_{\min} \max\{0, \omega(x, u)\} + y_{\max} \min\{0, \omega(x, u)\} \right) \mathrm{d}u. \end{aligned} \quad (\text{A.10})$$

Solving each OT sub-problem. Fix a nonempty gap $G_k(x)$ and let $T_x(u) = \omega(x, u)$ and $\rho_{k|x} = (T_x)_{\#} \nu_{k|x}$. Pushing any $\gamma_k \in \Pi(\mu_{1,k|x}, \nu_{k|x})$ forward under $(y, u) \mapsto (y, T_x(u))$ produces a coupling in $\Pi(\mu_{1,k|x}, \rho_{k|x})$ with the same objective value. Conversely, disintegrate $\nu_{k|x}$ with respect to T_x as a regular conditional law $\nu_{k|x}(\mathrm{d}u \mid v)$. Any coupling of (Y, V) with marginals $(\mu_{1,k|x}, \rho_{k|x})$ lifts through this conditional law to a coupling of (Y, U) with marginals $(\mu_{1,k|x}, \nu_{k|x})$ and $V = T_x(U)$ almost surely. Thus the two infima coincide. Applying Theorem 2.3 to the product cost yv gives

$$|G_k(x)| \min_{\gamma_k \in \Pi(\mu_{1,k|x}, \nu_{k|x})} \int y \cdot \omega(x, u) \mathrm{d}\gamma_k = (\underline{p}_{k+1}(x) - \bar{p}_k(x)) \int_0^1 Q_{Y,k|x}(t) Q_{\omega,k|x}(1-t) \mathrm{d}t, \quad (\text{A.11})$$

where $Q_{Y,k|x}$ is the quantile function of $\mu_{1,k|x}$ and $Q_{\omega,k|x}$ is the quantile function of $\omega(x, U)$ for $U \sim \text{Unif}(G_k(x))$. The prefactor $|G_k(x)| = \underline{p}_{k+1}(x) - \bar{p}_k(x)$ is simply $\mathbb{P}(U \in G_k(x))$: the coupling γ_k is a probability measure, so the region's contribution to the objective carries its Lebesgue mass as a weight. Empty gaps are omitted and their terms are defined as zero.

Remark (notation). $T_x : [0, 1] \rightarrow \mathbb{R}$, $T_x(u) = \omega(x, u)$, is the target-weight map with the covariate argument fixed at x . $\rho_{k|x} = (T_x)_{\#} \text{Unif}(G_k(x))$ is the population probability law of $\omega(x, U)$ for an auxiliary $U \sim \text{Unif}(G_k(x))$, defined for fixed x and $k = 0, \dots, K_x - 1$ with $|G_k(x)| > 0$.

Aggregation over X . Summing (A.11) over $k = 0, \dots, K_x - 1$, adding the LIV contribution (A.9) over $k = 1, \dots, K_x$, and adding the trivial tail (A.10) shows through (A.4) that

every feasible π_1 satisfies

$$\mathbb{E}_{\pi_1}[Y(1)\omega(X, U)] \geq \mathbb{E}_X \left[\sum_{k=0}^{K_X-1} (\underline{p}_{k+1}(X) - \bar{p}_k(X)) \int_0^1 Q_{Y,k|X}(t) Q_{\omega,k|X}(1-t) dt \right. \\ \left. + \sum_{k=1}^{K_X} \int_{\underline{p}_k(X)}^{\bar{p}_k(X)} \frac{\partial}{\partial u} g_1^c(u, X) \omega(X, u) du \right. \\ \left. + \int_{\bar{p}_{K_X}(X)}^1 \left(y_{\min} \max\{0, \omega(X, u)\} + y_{\max} \min\{0, \omega(X, u)\} \right) du \right].$$

The outer expectation is well-defined because the absolute value of the bracketed expression is bounded by $B_Y \int_0^1 |\omega(x, u)| du$, and this bound is integrable by ?? 2.2.

Sharpness. We now construct an attaining kernel. On a nonempty gap $G_k(x)$, define

$$F_{\omega,k|x}(v) := \nu_{k|x}\{u : \omega(x, u) \leq v\} \\ R_{k,x}(u, s) := F_{\omega,k|x}(\omega(x, u)-) \\ + s \left[F_{\omega,k|x}(\omega(x, u)) - F_{\omega,k|x}(\omega(x, u)-) \right].$$

For independent $U \sim \nu_{k|x}$ and $S \sim \text{Unif}(0, 1)$, the randomized probability integral transform gives $R_{k,x}(U, S) \sim \text{Unif}(0, 1)$ and $\omega(x, U) = Q_{\omega,k|x}(R_{k,x}(U, S))$ almost surely. Set, for $u \in G_k(x)$,

$$\pi_{1,x}^*(A | u) := \int_0^1 \mathbb{1}\{Q_{Y,k|x}(1 - R_{k,x}(u, s)) \in A\} ds, \quad A \in \mathcal{B}(\mathcal{Y}). \quad (\text{A.12})$$

This kernel has interval-averaged $Y(1)$ -marginal $\mu_{1,k|x}$ and attains (A.11), including when the measurable weight has atoms or is nonmonotone in u . On every non-degenerate LIV interval, set $\pi_{1,x}^*(dy | u) = K_{1,x}(dy | u)$ for the fixed compatible kernel used in (3.1). On the trivial tail, use $\delta_{y_{\min}}$ where $\omega(x, u) \geq 0$ and $\delta_{y_{\max}}$ where $\omega(x, u) < 0$; values on the finitely many region boundaries may be assigned arbitrarily.

Remark (notation). $F_{\omega,k|x}$ is the population distribution function of the pushforward of $\nu_{k|x}$ under $u \mapsto \omega(x, u)$, using a jointly measurable version in (x, v) for almost every x and each nonempty gap indexed by $k = 0, \dots, K_x - 1$. $R_{k,x}$ is a measurable map of $(u, s) \in G_k(x) \times [0, 1]$ for almost every x and each $k \in \{0, \dots, K_x - 1\}$ with $|G_k(x)| > 0$, using the left limit of $F_{\omega,k|x}$ to distribute each weight atom across its rank interval. $\pi_{1,x}^*$ is a jointly measurable population probability kernel for $Y(1)$ given $(X, U) = (x, u)$, defined for almost every x and

pasted across the nonempty gaps, non-degenerate LIV intervals and trivial tail, with arbitrary probability-law values at region boundaries.

The resulting kernel is jointly measurable. On each fixed instrument component C , continuity and common support identify the propensity extrema with the conditional essential extrema of P given $(X, Z \in C)$, which are measurable in x . Finite ordering and merging give measurable interval endpoints. The conditional gap laws, their distribution functions, and their generalized inverses admit jointly measurable versions on the standard Borel outcome space. The map $F_{\omega, k|x}(v)$ is jointly measurable by integration of a jointly measurable indicator over the measurable gap, and its left limit is measurable as a supremum over rational arguments. Equation (A.12) and finite pasting then give a probability kernel in (x, u) . Its gap marginals satisfy (A.6), its LIV pieces satisfy (A.8), and the compatibility argument above yields $\pi_1^* \in \Gamma_1(\mathbb{P}_{\text{obs}})$. It attains the displayed lower bound, so Proposition 3.1 makes the bound sharp. Replacing $Q_{Y, k|x}(1 - R_{k, x})$ by $Q_{Y, k|x}(R_{k, x})$ on each gap and swapping the two tail endpoints constructs the measurable upper-bound optimizer. $\square$

B Estimation Procedure Details

This appendix collects the estimation procedures in full: the three algorithms described in prose in Section 4 (Algorithms 1 to 3) and their nuisance-estimation steps. The closed-form rewritings (4.1) and (4.7) that these procedures start from are stated as Propositions 4.3, 4.4 and 4.8 in Section 4; their derivations are purely algebraic and are given in Section O3 of the Online Appendix.

B.1 Discrete Instrument, Continuous Outcome: Nuisance Estimation

Algorithm 1 states the estimation procedure, and the remainder of this subsection gives the nuisance-estimation steps in full.

Randomly partition the observed dataset of size n into two disjoint sets: an auxiliary sample I_1 and a main estimation sample I_2 , each of size $n/2$. We first estimate p and q on I_1 , use the covariates in I_2 only to recover their finite level sets, and then fit all remaining nuisance functions on I_1 given those sets.

1. **Propensity Scores:** Estimate the baseline propensity score $p(z, x)$ by regressing the treatment W on covariates X conditional on each $z \in \mathcal{Z}$. Clip $\hat{p}(z, x)$ to $[0, 1]$, then

Algorithm 1 DML estimator: discrete instrument, continuous outcome

- 1: Split sample into I_1 (auxiliary) and I_2 (main).
 - 2: **On I_1 :** estimate propensity $\hat{p}(z, x)$, clip it to $[0, 1]$, and set $\hat{q}(z, x) = \phi(z, x, \hat{p}(z, x))$.
 - 3: **Using the covariates in I_2 :** form $\bar{p}_2(z) = |I_2|^{-1} \sum_{i \in I_2} \hat{p}(z, X_i)$ and $\bar{q}_2(z) = |I_2|^{-1} \sum_{i \in I_2} \hat{q}(z, X_i)$. Assign structural-zero categories to $\hat{S}_0$ as in Step 1 below and remove their baseline vertices. Connect two remaining labeled values when their distance is less than $c_{\text{gap}}/2$ and take graph connected components to obtain and order the positive-baseline sets $\hat{S}_k$ and the policy sets $\hat{T}_{j,k}$.
 - 4: **On I_1 , given the recovered sets:** set $\hat{p}_0 \equiv 0$ and, for $k \geq 1$, estimate $\hat{p}_k(x) = \widehat{\mathbb{E}}[W \mid Z \in \hat{S}_k, X = x]$. Estimate the two conditional probabilities defining $\gamma_{\text{full},k}(x)$ from their set indicators and subtract them; estimate $\hat{\pi}_k(x)$, $\hat{\tau}_{j,k}(x)$, $\hat{\gamma}_{j,k}(x)$, and $\hat{\gamma}_K(x)$ from the corresponding set indicators. Clip all estimated conditional probabilities to $[0, 1]$ and $\hat{\pi}_k, \hat{\tau}_{j,k}$ further to $[c_\pi/2, 1]$. Within each $\hat{T}_{j,k}$, estimate the shares $\{\hat{a}_{j,k}(z, x) : z \in \hat{T}_{j,k}\}$ jointly with a simplex-valued multinomial learner, so they are nonnegative and sum to one, and set $\hat{q}_{j,k}(x) = \sum_{z \in \hat{T}_{j,k}} \hat{a}_{j,k}(z, x) \hat{q}(z, x)$.
 - 5: Estimate boundary quantiles $\hat{v}_{j,k}(x)$ via the moment-equation regression in Step 4 below and clip them to $[y_{\min}, y_{\max}]$.
 - 6: Estimate conditional expectations $\hat{M}_{j,k}^\pm(x)$, $\hat{J}_{j,k}^\pm(x)$, $\hat{J}_{\text{full},k}^\pm(x)$ by flexible ML, omitting every minus-side object at the structural boundary $k = 0$.
 - 7: **On I_2 :** compute orthogonal scores $\hat{\psi}_{\text{full},k}^{\text{prod}}, \hat{\psi}_{j,k}^{\text{prod}}, \hat{\psi}_{\Delta_K}$ from Section C.1.
 - 8: Return $\hat{\theta}_{\omega,1}$ as in (4.5) and $\hat{\sigma}^2 = |I_2|^{-1} \sum_{i \in I_2} (\hat{\psi}(O_i; \hat{\eta}) - \hat{\theta}_{\omega,1})^2$.
-

compute $\hat{q}(z, x) = \phi(z, x, \hat{p}(z, x))$.

2. **Level Sets:** Using all covariates in I_2 , compute

$$\bar{p}_2(z) = \frac{1}{|I_2|} \sum_{i \in I_2} \hat{p}(z, X_i), \quad \bar{q}_2(z) = \frac{1}{|I_2|} \sum_{i \in I_2} \hat{q}(z, X_i).$$

Choose a deterministic threshold $\rho_n \downarrow 0$ such that $n^{1/4} \rho_n \rightarrow \infty$. Define $\hat{S}_0 = \{z : \bar{p}_2(z) \leq \rho_n\}$ and remove the corresponding baseline vertices from the clustering graph. Form an undirected graph from the remaining baseline vertices and all policy vertices. Connect two vertices when the corresponding scalar values among $\{\bar{p}_2(z), \bar{q}_2(z) : z \in \mathcal{Z}\}$ differ by less than $c_{\text{gap}}/2$, and define the estimated positive score levels as the graph's connected components. The baseline sets $\hat{S}_k$, $k \geq 1$, collect instruments whose remaining (p, z) vertices share a component, ordered by their component averages. The policy sets

$\widehat{T}_{j,k}$ collect instruments whose (q, z) vertices share a component and are indexed by the baseline bin containing that component average, using the half-open bins for $k < K$ and the closed final bin for $k = K$.

3. **Baseline Levels, Instrument Assignment, and Pooled Policy Levels:** Given the recovered sets, return to I_1 . Set $\widehat{p}_0(x) \equiv 0$. For $k = 1, \dots, K$, estimate

$$p_k(x) = \mathbb{E}[W \mid Z \in S_k, X = x], \quad \pi_k(x) = \mathbb{P}_{\text{obs}}(Z \in S_k \mid X = x)$$

by regressing W and $\mathbb{1}(Z \in \widehat{S}_k)$, respectively, on X ; no p_0 or π_0 regression is fitted. For every nonempty $\widehat{T}_{j,k}$, estimate $\tau_{j,k}(x) = \mathbb{P}(Z \in T_{j,k} \mid X = x)$ and estimate the within-set assignment shares $a_{j,k}(z, x) = \mathbb{P}(Z = z \mid Z \in T_{j,k}, X = x)$ jointly with a simplex-valued multinomial learner. Thus $\widehat{a}_{j,k}(z, x) \geq 0$ and $\sum_{z \in \widehat{T}_{j,k}} \widehat{a}_{j,k}(z, x) = 1$. Clip $\widehat{\pi}_k$ and $\widehat{\tau}_{j,k}$ to $[c_\pi/2, 1]$. Form the pooled level by the plug-in identity

$$q_{j,k}(x) = \mathbb{E}[q(Z, x) \mid Z \in T_{j,k}, X = x]$$

and its sample analog $\widehat{q}_{j,k}(x) = \sum_{z \in \widehat{T}_{j,k}} \widehat{a}_{j,k}(z, x) \widehat{q}(z, x)$. At the population law, every $q(z, x)$ in $T_{j,k}$ equals $q_{j,k}(x)$, so pooling preserves the boundary while giving it a regular extension when estimated instrument-specific levels are not exactly tied. Estimate the two conditional probabilities defining $\gamma_{\text{full},k}$ from their respective set indicators, clip each estimate to $[0, 1]$, and subtract. Estimate $\gamma_{j,k}$ and the tail weight

$$\gamma_K(x) := \mathbb{P}_{\text{obs}} \left(Z \in \bigcup_{j=1}^{l_K} T_{j,K} \mid X = x \right)$$

from their set indicators and also clip them to $[0, 1]$.

4. **Conditional Quantiles:** Set the structural lower-boundary value $\nu_{0,k}(x) \equiv y_{\min}$. For $j = 1, \dots, l_k$, $\nu_{j,k}(x) = Q_{Y,k|x}(\kappa_{j,k}(x))$ is the $\kappa_{j,k}(x)$ -quantile of the conditional distribution $\mu_{1,k|x}$, where $\mu_{1,k|x}$ is identified in (3.17). Applying the quantile characterization $\mathbb{P}_{\mu_{1,k|x}}(Y \leq \nu_{j,k}(x)) = \kappa_{j,k}(x) = \frac{q_{j,k}(x) - p_k(x)}{p_{k+1}(x) - p_k(x)}$ and expanding $\mathbb{P}_{\mu_{1,k|x}}$ in terms of the observed conditional distributions yields the following moment equation:

$$\begin{aligned} p_{k+1}(X) \mathbb{E}[\mathbb{1}(Y \leq \nu_{j,k}(X)) \mid X, Z \in S_{k+1}, W = 1] \\ - p_k(X) \mathbb{E}[\mathbb{1}(Y \leq \nu_{j,k}(X)) \mid X, Z \in S_k, W = 1] - (q_{j,k}(X) - p_k(X)) \nu(X) = 0. \end{aligned}$$

For $j = 1, \dots, l_k$, we estimate $\nu_{j,k}$ by

$$\nu_{j,k}(x) = \arg \min_{\nu} \mathbb{E} [\Omega(Z, X, W) \max\{0, \nu(X) - Y\} - (q_{j,k}(X) - p_k(X)) \nu(X)],$$

where

$$\Omega(Z, X, W) = W \cdot \left(\frac{\mathbb{1}(Z \in S_{k+1})}{\pi_{k+1}(X)} - \frac{\mathbb{1}(Z \in S_k)}{\pi_k(X)} \right),$$

with the second term omitted when $k = 0$. Clip each fitted $\hat{\nu}_{j,k}$ to $[y_{\min}, y_{\max}]$. At $j = 0$, $q_{0,k} \equiv p_k$ and $\kappa_{0,k} = 0$, so the displayed moment equation does not point-identify a quantile: every value below the essential infimum solves it. The fixed convention $\nu_{0,k} = y_{\min}$ removes this nonidentified nuisance. Under the no-atom condition at $y_{\min}$ used for the continuous-outcome score, $M_{0,k}^{\pm} = 0$ and the associated indicator residuals vanish, so the score is invariant to this boundary convention. By part (ii) of ?? 4.2, every nonempty policy set in a nonfinal bin has $q_{j,k} > p_k$ and therefore corresponds to a regular estimated threshold with $j \geq 1$; the structural threshold $\nu_{0,k}$ is never estimated. Then, we estimate the conditional expectations of the indicators within the moment constraints:

$$\begin{aligned} M_{j,k}^+(X) &:= \mathbb{E}[\mathbb{1}(Y \leq \nu_{j,k}(X)) \mid Z \in S_{k+1}, W = 1, X], \\ M_{j,k}^-(X) &:= \mathbb{E}[\mathbb{1}(Y \leq \nu_{j,k}(X)) \mid Z \in S_k, W = 1, X]. \end{aligned}$$

The minus-side regression is estimated only for $k \geq 1$; at $k = 0$ it is omitted under the structural-boundary convention.

5. **Estimate $J_{j,k}(X)$ and $J_{\text{full},k}(X)$:** We estimate $J_{j,k}(x)$ by

$$\begin{aligned} J_{j,k}^+(X) &= \mathbb{E}[Y \mathbb{1}(\nu_{j,k}(X) < Y \leq \nu_{j+1,k}(X)) \mid X, Z \in S_{k+1}, W = 1], \\ J_{j,k}^-(X) &= \mathbb{E}[Y \mathbb{1}(\nu_{j,k}(X) < Y \leq \nu_{j+1,k}(X)) \mid X, Z \in S_k, W = 1], \\ J_{j,k}(X) &= p_{k+1}(X) J_{j,k}^+(X) - p_k(X) J_{j,k}^-(X), \end{aligned}$$

and estimate $J_{\text{full},k}(X)$ by

$$\begin{aligned} J_{\text{full},k}^+(X) &= \mathbb{E}[Y \mid X, Z \in S_{k+1}, W = 1], \\ J_{\text{full},k}^-(X) &= \mathbb{E}[Y \mid X, Z \in S_k, W = 1], \\ J_{\text{full},k}(X) &= p_{k+1}(X) J_{\text{full},k}^+(X) - p_k(X) J_{\text{full},k}^-(X). \end{aligned}$$

For $k = 0$, the two minus-side regressions in these displays are omitted and their products with p_0 are set to zero.

Remark (notation). $J_{j,k}^+$ is the population treated conditional band-mean function at the true boundaries $\nu_{j,k}$ and $\nu_{j+1,k}$, indexed by $k = 0, \dots, K - 1$ and $j = 0, \dots, l_k - 1$, with

conditioning on $Z \in S_{k+1}$, $W = 1$, and $X = x$. $J_{j,k}^-(x)$ is the population truncated first moment $\mathbb{E}[Y \mathbb{1}(\nu_{j,k}(x) < Y \leq \nu_{j+1,k}(x)) \mid X = x, Z \in S_k, W = 1]$ for $k \in \{1, \dots, K-1\}$ and $j \in \{0, \dots, l_k - 1\}$, using population quantile boundaries, with its $k = 0$ value defined as zero solely for notational convenience and every corresponding minus-side score term omitted. $J_{\text{full},k}^+(x)$ is the population conditional mean $\mathbb{E}[Y \mid X = x, Z \in S_{k+1}, W = 1]$ for $k \in \{0, \dots, K-1\}$. $J_{\text{full},k}^-(x)$ is the population conditional mean $\mathbb{E}[Y \mid X = x, Z \in S_k, W = 1]$ for $k \in \{1, \dots, K-1\}$, with its $k = 0$ value defined as zero solely for notational convenience and every corresponding minus-side score term omitted.

In summary, the freely estimated nuisance functions are

$$\eta(x) = \{p(z, x), p_k(x), \pi_k(x), \tau_{j,k}(x), a_{j,k}(z, x), \gamma_{\text{full},k}(x), \gamma_{j,k}(x), \gamma_K(x), \nu_{j,k}(x), M_{j,k}^\pm(x), J_{j,k}^\pm(x), J_{\text{full},k}^\pm(x)\}_{j,k}.$$

Here π_k and every minus-side object are included only for $k \geq 1$, while $\nu_{j,k}$ is estimated only for $j \geq 1$; the boundary values given above are fixed.

The Neyman-orthogonal score for the target (4.1) is constructed by combining empirical influence functions for the conditional expectations $J_{j,k}^\pm(X)$ and $J_{\text{full},k}^\pm(X)$ with pathwise derivatives arising from the quantile thresholds $\nu_{j,k}(X)$ and propensity score boundaries $p_k(X)$. In addition, the score must correct for the estimation of (i) the conditional quantile indicators $M_{j,k}^\pm(X)$ via their Riesz representers, (ii) the propensity scores $p_k(X)$ through IPW residuals, and (iii) the alternative policy values $q_{j,k}(X)$ through the derivative $\frac{\partial \phi}{\partial p}$. The weighting regressions $\gamma(X)$ are handled by forming product scores using the chain rule for influence functions. The complete formulas for all score components— $\psi_{j,k}^J$, $\psi_{\text{full},k}^J$, $\psi_{j,k}^{\text{prod}}$, $\psi_{\text{full},k}^{\text{prod}}$, and ψ_{Δ_K} —are given in Section C.1.

Remark (notation). $\psi_{j,k}^J(O; \eta)$ is the continuous-outcome uncentered lower-bound score for $\mathbb{E}[J_{j,k}(X)]$, used for $0 \leq k < K$ and $0 \leq j < l_k$, with $R_{q,0,k} = R_{q,l_k+1,k} = 0$, all minus-side terms omitted at $k = 0$, and no standalone $+\nu_{0,k}$ term in α_{p_k} at $j = 0$. $\psi_{\Delta_K}(O; \eta)$ is the uncentered lower-bound score for $\mathbb{E}[\Delta_K(X)]$ on the final bin $[p_K(X), 1]$, with the same formula for continuous and binary outcomes and with the union taken over the actual policy levels $j = 1, \dots, l_K$, including a policy value equal to one.

B.2 Discrete Instrument, Binary Outcome: Nuisance Estimation

The closed-form reduction (4.4) is proved as part (ii) of Proposition 4.4 in Section O3.2 of the Online Appendix. Algorithm 2 states the estimation procedure, and the remainder of this subsection gives the full nuisance-estimation steps.

Algorithm 2 DML estimator: discrete instrument, binary outcome

- 1: Split sample into I_1, I_2 . Estimate $\hat{p}$ and $\hat{q}$ on I_1 , and recover $\hat{S}_k, \hat{T}_{j,k}$ from $\{\bar{p}_2(z), \bar{q}_2(z)\}_{z \in \mathcal{Z}}$ using the I_2 covariates and the graph rule in Algorithm 1.
 - 2: **On I_1 , given the recovered sets:** set $\hat{p}_0(x) = \hat{P}_{1,0}(x) = 0$. For $k = 1, \dots, K$, estimate $\hat{p}_k(x) = \mathbb{E}[W \mid Z \in \hat{S}_k, X = x]$, $\hat{P}_{1,k}(x) = \mathbb{E}[YW \mid Z \in \hat{S}_k, X = x]$, and $\hat{\pi}_k(x)$. For every nonempty $\hat{T}_{j,k}$, estimate $\hat{\tau}_{j,k}(x)$ and the within-set shares $\{\hat{a}_{j,k}(z, x) : z \in \hat{T}_{j,k}\}$ jointly with a simplex-valued multinomial learner, and form $\hat{q}_{j,k}(x) = \sum_{z \in \hat{T}_{j,k}} \hat{a}_{j,k}(z, x) \hat{q}(z, x)$. Estimate the two conditional probabilities defining $\gamma_{\text{full},k}$ from their set indicators and subtract them; estimate $\hat{\gamma}_{j,k}, \hat{\gamma}_K$ from their set indicators. Clip all estimated conditional probabilities to $[0, 1]$ and $\hat{\pi}_k, \hat{\tau}_{j,k}$ further to $[c_\pi/2, 1]$.
 - 3: Clip each $\hat{P}_{1,k}$ to $[0, 1]$ and, for each i , sort $(\hat{P}_{1,1}(X_i), \dots, \hat{P}_{1,K}(X_i))$ to obtain $(\hat{P}_{1,1}^*(X_i), \dots, \hat{P}_{1,K}^*(X_i))$, retaining $\hat{P}_{1,0}^*(X_i) = 0$. Set $\hat{J}_{\text{full},k} = \hat{P}_{1,k+1}^* - \hat{P}_{1,k}^*$ and use the same rearranged $\hat{P}^*$ vector in every score plug-in and residual.
 - 4: **On I_2 :** compute orthogonal scores $\hat{\psi}^{\text{prod}}, \hat{\psi}_{\Delta_K}$ from Section C.2 and average to obtain $\hat{\theta}_{\omega,1}$ as in (4.5).
-

The procedure follows the same two-step sample-splitting strategy as the continuous outcome case, but with a simpler nuisance list: the conditional quantile estimation (Step 4 of the continuous-outcome procedure) and the indicator expectations $M_{j,k}^\pm$ are no longer needed. The full nuisance list, for all interval indices $k \in \{0, \dots, K\}$ and sub-interval indices $j \in \{0, \dots, l_k\}$, is

$$\eta(x) = \left\{ p(z, x), p_k(x), \pi_k(x), \tau_{j,k}(x), a_{j,k}(z, x), \gamma_{\text{full},k}(x), \gamma_{j,k}(x), \gamma_K(x), P_{1,k}(x) \right\}_{j,k}.$$

Step 1: Nuisance Estimation.

1. **Propensity Scores, Level Sets, and Assignment Probabilities:** Follow Steps 1–3 of Section B.1: estimate $\hat{p}(z, x)$ and $\hat{q}(z, x)$ on I_1 ; construct $\hat{S}_k, \hat{T}_{j,k}$ from $\bar{p}_2, \bar{q}_2$ using the covariates in I_2 ; then, given those sets, fit $\hat{p}_k, \hat{\pi}_k, \hat{\tau}_{j,k}, \hat{a}_{j,k}$, the pooled $\hat{q}_{j,k}$, the two conditional probabilities entering $\hat{\gamma}_{\text{full},k}$, and the probability nuisances $\hat{\gamma}_{j,k}, \hat{\gamma}_K$ on I_1 . Form $\hat{\gamma}_{\text{full},k}$ by subtracting its two constituent probability estimates. Set $\hat{p}_0 \equiv 0$, fit no π_0 , and clip all conditional probability estimates to $[0, 1]$, with $\hat{\pi}_k, \hat{\tau}_{j,k}$ clipped further to $[c_\pi/2, 1]$.
2. **Conditional Joint Probability:** Set $\hat{P}_{1,0}(x) = 0$. For $k = 1, \dots, K$, estimate $P_{1,k}(x) = \mathbb{E}[YW \mid Z \in S_k, X = x]$ on I_1 by regressing YW on X conditional on $Z \in \hat{S}_k$,

using any flexible machine learning method, and clip the fitted conditional probabilities to $[0, 1]$.

3. **Rearrangement:** Set the structural boundary value $\hat{P}_{1,0}(x) = \hat{P}_{1,0}^*(x) = 0$. In the population, $P_{1,k}(X)$ is monotone non-decreasing in k because $J_{\text{full},k}(X) = P_{1,k+1}(X) - P_{1,k}(X) \geq 0$. Because the $\hat{P}_{1,k}$ are fitted on different instrument subsets, finite-sample estimates need not preserve monotonicity. Clip them to $[0, 1]$, sort only $(\hat{P}_{1,1}(X_i), \dots, \hat{P}_{1,K}(X_i))$ to obtain $(\hat{P}_{1,1}^*(X_i), \dots, \hat{P}_{1,K}^*(X_i))$, and retain the fixed zero boundary. Define $\hat{J}_{\text{full},k} = \hat{P}_{1,k+1}^* - \hat{P}_{1,k}^*$ and use this same rearranged $\hat{P}^*$ vector in every score plug-in and residual; no separate gap projection is applied.

The identification formulas allow a nonempty policy set to lie at $q_{j,k} = p_k$. For the displayed inference formulas, part (ii) of ?? 4.2 excludes this collision in nonfinal bins, so the regular residual $R_{q,j,k}$ never represents the same boundary as the structural convention $q_{0,k} = p_k$.

Remark (notation). $R_{q,j,k}(O)$ pools the propensity residuals of the instruments sharing one nonempty regular policy level, normalizing by the total level-set probability $\tau_{j,k}(X)$, while its structural integration-boundary indices satisfy $R_{q,0,k} = R_{q,l_k+1,k} \equiv 0$. In Lemma D.1, Theorem 4.7: For $k < K$, a nonempty regular policy level is distinct from $q_{0,k} = p_k$ under part (ii) of Assumption 4.2; multiple instruments may nevertheless share that regular policy level.

Step 2: Neyman-Orthogonal Score Construction. The target functional for the lower bound is $\underline{\theta}_{\omega,1} = \mathbb{E}[m(X; \eta)]$, where

$$m(X; \eta) = \sum_{k=0}^{K-1} \left(\gamma_{\text{full},k}(X) J_{\text{full},k}(X) + \sum_{j=1}^{l_k} \gamma_{j,k}(X) h_{j,k}^-(X) \right) + \Delta_K(X),$$

and the Neyman-orthogonal score takes the aggregated form

$$\psi(O; \eta) = \sum_{k=0}^{K-1} \left(\psi_{\text{full},k}^{\text{prod}}(O; \eta) + \sum_{j=1}^{l_k} \psi_{j,k}^{\text{prod}}(O; \eta) \right) + \psi_{\Delta_K}(O; \eta). \quad (\text{B.1})$$

The construction of each component score proceeds by computing the pathwise derivatives of $h_{j,k}^-(X)$ with respect to $P_{1,k}(X)$, $p_k(X)$, and $q_{j,k}(X)$, and combining them with standard IPW residuals. Because $h_{j,k}^-$ involves a max operator, the pathwise derivatives depend on active-set, boundary, and dominance indicators that track which constraint binds in the max. The product scores account for the estimation of $\gamma(X)$ via the chain rule for influence

functions, and the trivial-bound score ψ_{Δ_K} is identical to the continuous-outcome case. The complete formulas are given in Section C.2. Averaging the plug-in scores over I_2 yields the estimator previewed in (4.5); the upper-bound score is derived similarly.

Remark (notation). $\psi_{j,k}^{\text{prod}}(O; \eta)$ is the uncentered lower-bound fractional-interval product score for $0 \leq k < K$ and $1 \leq j \leq l_k$, using the conditional union-indicator residual $\mathbb{1}(Z \in \bigcup_{l=j}^{l_k} T_{l,k}) - \gamma_{j,k}(X)$. In Lemma D.1, Theorem 4.7: In the continuous-outcome case, $\psi_{j,k}^{\text{prod}} = \gamma_{j,k} \psi_{j-1,k}^J + J_{j-1,k}(\mathbb{1}(Z \in \bigcup_{l=j}^{l_k} T_{l,k}) - \gamma_{j,k})$; in the binary-outcome case, $\psi_{j,k}^{\text{prod}} = \gamma_{j,k} \psi_{j,k}^h + h_{j,k}^-(\mathbb{1}(Z \in \bigcup_{l=j}^{l_k} T_{l,k}) - \gamma_{j,k})$, with covariate and observation arguments as in the defining displays. $\psi_{\text{full},k}^{\text{prod}}(O; \eta)$ is the uncentered lower-bound product score for $\mathbb{E}[\gamma_{\text{full},k}(X)J_{\text{full},k}(X)]$, for $0 \leq k < K$, using the full-interval component score appropriate to the outcome setting and the difference of policy and baseline higher-bin indicators.

B.3 Continuous Instrument: Localized g_1 and Boundary Quantile Estimation

We give the details behind Algorithm 3. Randomly partition the data into a propensity sample I_0 , a nuisance sample I_1 , and an evaluation sample I_2 , with $|I_j|/n \rightarrow \rho_j \in (0, 1)$ for $j = 0, 1, 2$. Estimating p on I_0 makes every localization rule used on I_1 fixed conditional on I_0 . The selected observations are therefore conditionally i.i.d. Cross-fitting repeats the construction over several partitions and averages the resulting scores.

Propensity scores, support, and policy. On I_0 , regress W on (Z, X) and clip the fitted values to $[0, 1]$. Set $\hat{q}(z, x) = \phi(z, x, \hat{p}(z, x))$ and compute

$$\hat{\underline{p}}(x) = \inf_{z \in \mathcal{Z}} \hat{p}(z, x), \quad \hat{\bar{p}}(x) = \sup_{z \in \mathcal{Z}} \hat{p}(z, x), \quad \hat{q}^{\text{tr}}(z, x) = \min\{\hat{\bar{p}}(x), \max\{\hat{\underline{p}}(x), \hat{q}(z, x)\}\}.$$

Conditional expectation g_1 via localized local-linear learning. The target evaluates $g_1(u, x) = \mathbb{E}[YW \mid P = u, X = x]$ at the boundaries of the conditional support. A local-constant regression generally has first-order boundary bias. We instead use a local-linear fit in the scalar index and allow flexible learners in x , following the same localization principle as local-polynomial DR-learning [Kennedy, 2023] and generalized random forests [Athey et al., 2019].

Let $K_{h_n}(t) = K(t/h_n)/h_n$, set $\hat{P}_i = \hat{p}(Z_i, X_i)$, and take a grid $u_m = m/M_g$, $m = 0, \dots, M_g$.

At each grid point, Step 3 of Algorithm 3 fits functions $(\hat{g}_1(u_m, \cdot), \hat{\beta}_m)$ on I_1 by

$$(\hat{g}_1(u_m, \cdot), \hat{\beta}_m) \in \arg \min_{g, \beta \in \mathcal{F}_n} \sum_{i \in I_1} K_{h_n}(\hat{P}_i - u_m) \left[Y_i W_i - g(X_i) - \beta(X_i)(\hat{P}_i - u_m) \right]^2.$$

Here $\mathcal{F}_n$ is a user-chosen class of functions of x . With a linear class, the display is weighted least squares on the augmented design $(1, x, \hat{P}_i - u_m, (\hat{P}_i - u_m)x)$. The fitted intercept $\hat{g}_1(u_m, x)$ estimates $g_1(u_m, x)$. Define $\hat{g}_1(u, x)$ between grid points by linear interpolation, and then clip it to the known range of YW . For $u \in [u_m, u_{m+1}]$,

$$\hat{g}_1(u, x) = \hat{g}_1(u_m, x) + \frac{u - u_m}{u_{m+1} - u_m} (\hat{g}_1(u_{m+1}, x) - \hat{g}_1(u_m, x)).$$

Condition (L2) of ?? O5.2 controls its prediction error only at the two estimated arguments used by the target.

Remark (notation). M_g is a fixed positive integer tuning choice for the regression grid at each sample size, giving spacing M_g^{-1} and $M_g + 1$ grid points u_m , $m = 0, \dots, M_g$. $u_m = m/M_g$ for $m = 0, \dots, M_g$ is the fixed regression index grid spanning $[0, 1]$, including both endpoints. $\mathcal{F}_n$ is a user-chosen, potentially n -dependent class of real-valued functions of x , used for both the intercept function g and the slope function β in each local-linear fit.

Treated boundary law. Equation (4.6) shows that the required distribution is the normalized canonical treated submeasure at $P = \underline{p}(X)$, equivalently the one-sided boundary limit of the treated outcome law. Choose $\delta_n \rightarrow 0$ and form

$$\mathcal{I}_{\delta,1} = \{i \in I_1 : W_i = 1, \hat{P}_i \leq \underline{p}(X_i) + \delta_n\}.$$

Conditional on I_0 , this is an i.i.d. selected sample. On the quantile grid $\tau_m = m/M_Q$, $m = 1, \dots, M_Q - 1$, estimate the conditional quantiles of Y given X on this sample by check-function regression over a function class $\mathcal{F}_Q$,

$$\tilde{Q}(\tau_m \mid \cdot) \in \arg \min_{f \in \mathcal{F}_Q} \sum_{i \in \mathcal{I}_{\delta,1}} \rho_{\tau_m}(Y_i - f(X_i)), \quad \rho_{\tau}(u) = u(\tau - \mathbb{1}(u < 0)). \quad (\text{B.2})$$

When there are no covariates, $\mathcal{F}_Q$ is the set of constants and (B.2) returns the empirical τ_m -quantile of $\{Y_i\}_{i \in \mathcal{I}_{\delta,1}}$. When X has finite support and $\mathcal{F}_Q$ is saturated in X , the same optimization decomposes by covariate stratum and returns the empirical quantile within each stratum. With general covariates, fit the conditional quantiles by (B.2) over $\mathcal{F}_Q$. Clip the fitted quantiles to $[y_{\min}, y_{\max}]$, monotonically rearrange them—that is, sort $\tilde{Q}(\tau_1 \mid x), \dots, \tilde{Q}(\tau_{M_Q-1} \mid$

x) in increasing order for each x —to prevent quantile crossing, and interpolate linearly in τ between adjacent grid levels. The endpoint values at $\tau = 0, 1$ may be set either to $y_{\min}, y_{\max}$ or to the smallest and largest fitted values; the two conventions differ by $O(M_Q^{-1})$ in the discretized integral below, and neither affects the population Lebesgue integrals. Denote the resulting law and quantile function by $\hat{\mu}_{1,p|x}$ and $\hat{Q}_{1,p|x}$.

Remark (notation). $\tau_m = m/M_Q$ for $m = 1, \dots, M_Q - 1$ denotes the interior quantile grid; the interpolation endpoint levels 0 and 1 are separate from these indexed grid points. $\mathcal{F}_Q$ is the chosen class of real-valued functions of x over which the boundary conditional quantile check-function regression is minimized; with no covariates it is the class of constants. $\tilde{Q}$ is the raw interior-grid check-regression fit on the treated localized sample from I_1 , whose selection rule is determined by the first stage on I_0 , before clipping, rearrangement, and interpolation. $\rho_\tau(u) = u(\tau - \mathbb{1}(u < 0))$ is the check loss for $u \in \mathbb{R}$ at the quantile levels $\tau = \tau_m = m/M_Q$, $m = 1, \dots, M_Q - 1$, used in the boundary quantile regression.

Target Evaluation. Using the main evaluation sample I_2 , we evaluate the empirical analog of the simplified target functional (4.7) by plugging in the nuisance estimators from I_0 and I_1 . For each observation $i \in I_2$, we compute the three components:

The first term (quantile tail integral) is evaluated over the estimated limits:

$$\hat{\psi}_1(O_i) = - \int_{\min\{\hat{p}(X_i), \hat{q}(Z_i, X_i)\}}^{\hat{p}(X_i)} \hat{Q}_{1,p|X_i} \left(\frac{v}{\hat{p}(X_i)} \right) dv,$$

with value zero when $\hat{p}(X_i) = 0$. The substitution $v = \hat{p}(X_i) s$ removes the observation-specific integration limits and turns this into a lookup in the cumulative quantile integral. Writing

$$\hat{C}(\tau | x) := \int_0^\tau \hat{Q}_{1,p|x}(s) ds, \quad \hat{\tau}_i := \min\{1, \hat{q}(Z_i, X_i)/\hat{p}(X_i)\} \quad \text{when } \hat{p}(X_i) > 0,$$

we obtain, for the lower and upper bounds respectively,

$$\hat{\psi}_1(O_i) = -\hat{p}(X_i) [\hat{C}(1 | X_i) - \hat{C}(\hat{\tau}_i | X_i)], \quad \hat{\psi}_1^+(O_i) = -\hat{p}(X_i) \hat{C}(1 - \hat{\tau}_i | X_i).$$

When $\hat{p}(X_i) = 0$, both first-term contributions are set to zero without evaluating $\hat{\tau}_i$. Because $\hat{Q}_{1,p|x}$ is piecewise linear on the τ grid, the trapezoid rule computes $\hat{C}(\tau_m | x)$ exactly; off-grid values of $\hat{\tau}_i$ are obtained by interpolating between adjacent $\hat{C}(\tau_m | x)$, which contributes to the M_Q^{-1} grid term of Corollary O5.1. When $\hat{p}(X_i) > 0$ and $\hat{q}(Z_i, X_i) \geq \hat{p}(X_i)$, the OT region is empty and $\hat{\tau}_i = 1$.

Remark (notation). $\widehat{C}(\tau \mid x) = \int_0^\tau \widehat{Q}_{1,p|x}(s) ds$ for $\tau \in [0, 1]$ and $x \in \mathcal{X}$, with grid values computed exactly by the trapezoid rule and off-grid evaluation implemented by linear interpolation of adjacent grid values. For $i \in I_2$ with $\widehat{p}(X_i) > 0$, $\widehat{\tau}_i = \min\{1, \widehat{q}(Z_i, X_i)/\widehat{p}(X_i)\} \in [0, 1]$ is the normalized boundary-integration cutoff; when $\widehat{p}(X_i) = 0$ it is not evaluated, and when $\widehat{q}(Z_i, X_i) \geq \widehat{p}(X_i) > 0$ it equals one.

The second term evaluates the nonparametric regression $\widehat{g}_1$ at the truncated limits:

$$\widehat{\psi}_2(O_i) = \widehat{g}_1(\widehat{q}^{\text{tr}}(Z_i, X_i), X_i) - \widehat{g}_1(\widehat{P}_i, X_i).$$

The third term (trivial bound) is the straightforward threshold excess:

$$\widehat{\psi}_3(O_i) = y_{\min} \max\{0, \widehat{q}(Z_i, X_i) - \widehat{p}(X_i)\}.$$

Averaging these three components over I_2 yields the estimator previewed in (4.8).

B.3.1 The Concrete Estimator

A Concrete Estimation Procedure. We now provide one nuisance-estimation method. On I_0 we regress W on (Z, X) ; on I_1 we estimate g_1 by boundary-adaptive local-linear smoothing in the scalar index P combined with a flexible regression in X , and we estimate the treated boundary law by conditional quantile regression on the treated observations whose estimated propensity lies within δ_n of the estimated lower boundary. Algorithm 3 states the procedure; the kernel $K_{h_n}(t) = K(t/h_n)/h_n$, the function classes $\mathcal{F}_n$ and $\mathcal{F}_Q$, and the interpolation conventions are described in Section B.3.

Remark. Off-grid evaluations of $\widehat{C}(\tau \mid x)$ use linear interpolation between adjacent cumulative-integral values on $\{0, \tau_1, \dots, \tau_{M_Q-1}, 1\}$, whose values are computed by the trapezoid rule.

Theorem 4.10 applies to this construction when its learners deliver the three losses in ?? 4.9. Section O5 of the Online Appendix states the learner conditions (?? O5.2), relates those learner errors to the target rate in Corollary O5.1, and remarks on series learners that could meet the individual requirements.

Full PRTE bound. The displayed estimator concerns the treated component because the main identification results use that component as their running case. For the untreated component, define the reversed system

$$W^0 = 1 - W, \quad P^0 = 1 - P, \quad q^0 = 1 - q, \quad U^0 = 1 - U.$$

Algorithm 3 Estimator: continuous instrument

- 1: Split the sample into I_0, I_1, I_2 . Fix a bandwidth h_n , a localization window δ_n , and the grids $u_m = m/M_g$ ($m = 0, \dots, M_g$) and $\tau_m = m/M_Q$ ($m = 1, \dots, M_Q - 1$).
- 2: **On** I_0 : regress W on (Z, X) and clip the fitted values to $[0, 1]$ to obtain $\hat{p}$; set $\hat{q}(z, x) = \phi(z, x, \hat{p}(z, x))$, $\hat{p}(x) = \inf_z \hat{p}(z, x)$, $\hat{\bar{p}}(x) = \sup_z \hat{p}(z, x)$, and $\hat{q}^{\text{tr}}(z, x) = \min\{\hat{p}(x), \max\{\hat{p}(x), \hat{q}(z, x)\}\}$. Write $\hat{P}_i = \hat{p}(Z_i, X_i)$.
- 3: **On** I_1 : at each u_m , fit the local-linear-in-index regression

$$(\hat{g}_1(u_m, \cdot), \hat{\beta}_m) \in \arg \min_{g, \beta \in \mathcal{F}_n} \sum_{i \in I_1} K_{h_n}(\hat{P}_i - u_m) \left[Y_i W_i - g(X_i) - \beta(X_i)(\hat{P}_i - u_m) \right]^2,$$

interpolate the fitted values $\hat{g}_1(u_m, x)$ linearly in u , and clip to the range of YW .

- 4: Form the treated localized sample $\mathcal{I}_{\delta,1} = \{i \in I_1 : W_i = 1, \hat{P}_i \leq \hat{p}(X_i) + \delta_n\}$. If it is empty, set $\tilde{Q}(\tau_m | x) = y_{\min}$ for every (m, x) . Otherwise, at each τ_m , fit the check-function regression

$$\tilde{Q}(\tau_m | \cdot) \in \arg \min_{f \in \mathcal{F}_Q} \sum_{i \in \mathcal{I}_{\delta,1}} \rho_{\tau_m}(Y_i - f(X_i)), \quad \rho_{\tau}(u) = u(\tau - \mathbf{1}(u < 0)),$$

Clip the fitted values to $[y_{\min}, y_{\max}]$, sort them across m to remove quantile crossing, and interpolate linearly in τ , using the endpoint convention of Section B.3 at $\tau \in \{0, 1\}$, to obtain $\hat{Q}_{1,p|x}$.

- 5: **On** I_2 : with $\hat{C}(\tau | x) = \int_0^\tau \hat{Q}_{1,p|x}(s) ds$ evaluated by the trapezoid rule on the τ grid, set $\hat{\tau}_i = \min\{1, \hat{q}(Z_i, X_i)/\hat{p}(X_i)\}$ whenever $\hat{p}(X_i) > 0$, and then set

$$\begin{aligned} \hat{\psi}_1(O_i) &= -\hat{p}(X_i) [\hat{C}(1 | X_i) - \hat{C}(\hat{\tau}_i | X_i)], \\ \hat{\psi}_2(O_i) &= \hat{g}_1(\hat{q}^{\text{tr}}(Z_i, X_i), X_i) - \hat{g}_1(\hat{P}_i, X_i), \\ \hat{\psi}_3(O_i) &= y_{\min}(\hat{q}(Z_i, X_i) - \hat{p}(X_i))_+, \end{aligned}$$

When $\hat{p}(X_i) = 0$, set $\hat{\psi}_1(O_i) = 0$ without evaluating $\hat{\tau}_i$. Average as in (4.8). For the upper bound, use the same zero branch; otherwise replace $\hat{\psi}_1(O_i)$ by $-\hat{p}(X_i)\hat{C}(1 - \hat{\tau}_i | X_i)$ and $y_{\min}$ by $y_{\max}$.

Its policy weight satisfies $\omega^0(x, v) = -\omega(x, 1 - v)$, so its treated-component target is

$$\vartheta_0 := \mathbb{E}[Y(0)\omega^0(X, U^0)] = -\theta_{\omega,0}.$$

Apply the same construction to obtain $[\widehat{\vartheta}_0, \widehat{\vartheta}_0]$. The full target is $\theta_\omega = \theta_{\omega,1} + \vartheta_0$, and its estimated sharp interval is therefore

$$[\widehat{\theta}_{\omega,1} + \widehat{\vartheta}_0, \widehat{\theta}_{\omega,1} + \widehat{\vartheta}_0].$$

This is equivalent to $[\widehat{\theta}_{\omega,1} - \widehat{\theta}_{\omega,0}, \widehat{\theta}_{\omega,1} - \widehat{\theta}_{\omega,0}]$. Its two bounds converge at the maximum of the corresponding componentwise rates; Corollary [O5.1](#) applies to the untreated component under the analogous learner conditions for the reversed system.

C Orthogonal Score Constructions

This section presents the detailed Neyman-orthogonal score formulas for the estimation procedures described in Section [4](#).

C.1 Continuous Outcome Score Components

We provide the complete lower-bound score components for the continuous-outcome setting of Section [4.1.1](#). Recall the target functional:

$$\theta_{\omega,1} = \mathbb{E} \left[\sum_{k=0}^{K-1} \left(\gamma_{\text{full},k}(X) J_{\text{full},k}(X) + \sum_{j=1}^{l_k} \gamma_{j,k}(X) J_{j-1,k}(X) \right) + \Delta_K(X) \right].$$

Influence-function residuals. The scaled influence-function residuals used to debias the conditional-expectation terms, evaluated at an observation $O = (X, Z, W, Y)$, are:

$$\begin{aligned} IF_{J^+}(O) &= \frac{1}{\pi_{k+1}(X)} \mathbb{1}(Z \in S_{k+1}) W \left(Y \mathbb{1}(\nu_{j,k}(X) < Y \leq \nu_{j+1,k}(X)) - J_{j,k}^+(X) \right) \\ IF_{M_{j,k}^+}(O) &= \frac{1}{\pi_{k+1}(X)} \mathbb{1}(Z \in S_{k+1}) W \left(\mathbb{1}(Y \leq \nu_{j,k}(X)) - M_{j,k}^+(X) \right), \end{aligned}$$

with IF_{J^-} and $IF_{M_{j,k}^-}$ defined analogously over S_k when $k \geq 1$. When $k = 0$, $p_0 \equiv 0$ is a structural boundary: all minus-side conditional expectations and residuals are omitted and are defined as zero solely for notational convenience. In particular, the score never divides by $\pi_0(X)$.

Remark (notation). IF_{J^+} denotes the plus-side observation-level residual component indexed implicitly by $k \in \{0, \dots, K-1\}$ and $j \in \{0, \dots, l_k-1\}$, evaluated with the nuisance functions supplied to the continuous-outcome score. IF_{J^-} denotes the minus-side observation-level residual component indexed implicitly by $k \in \{0, \dots, K-1\}$ and $j \in \{0, \dots, l_k-1\}$,

evaluated with the nuisance functions supplied to the continuous-outcome score, with its structural $k = 0$ component identically zero. $IF_{M_{j,k}^+}$ is the observation-level residual indexed by $k = 0, \dots, K - 1$ and boundary index $j = 0, \dots, l_k$. In Lemma D.1, Theorem 4.7: In the continuous-outcome part, the no-atom condition at $y_{\min}$ makes the true residual $IF_{M_{0,k}^+}$ vanish almost surely. $IF_{M_{j,k}^-}$ is the observation-level residual indexed by $k = 0, \dots, K - 1$ and boundary index $j = 0, \dots, l_k$, with the structural zero convention at $k = 0$. In Lemma D.1, Theorem 4.7: In the continuous-outcome part, the no-atom condition at $y_{\min}$ makes the true residual $IF_{M_{0,k}^-}$ vanish almost surely.

Riesz representers. Because $p_k(X)$ and $p_{k+1}(X)$ act as structural interval boundaries, the augmented functional yields non-zero derivatives with respect to these parameters. The exact Riesz representers for the marginal probabilities are:

$$\begin{aligned}\alpha_{p_k}(X) &= -J_{j,k}^-(X) - \nu_{j,k}(X)M_{j,k}^-(X) + \nu_{j+1,k}(X)M_{j+1,k}^-(X) \\ &\quad + \mathbb{1}(j \geq 1)\nu_{j,k}(X) - \nu_{j+1,k}(X), \\ \alpha_{p_{k+1}}(X) &= J_{j,k}^+(X) + \nu_{j,k}(X)M_{j,k}^+(X) - \nu_{j+1,k}(X)M_{j+1,k}^+(X).\end{aligned}$$

The first expression applies to $k \geq 1$; the p_0 correction is omitted because p_0 is known. The indicator is essential at the lower index: $q_{0,k} \equiv p_k$, so $q_{0,k} - p_k \equiv 0$ has zero derivative with respect to p_k . Thus the representer for the score component actually used at $j = 0$ contains no standalone $+\nu_{0,k}$ term.

Pooled policy-level residuals. For every nonempty alternative-policy level set, define

$$R_{q,j,k}(O) := \frac{1}{\tau_{j,k}(X)} \sum_{z \in T_{j,k}} \frac{\partial \phi}{\partial p}(z, X, p(z, X)) \mathbb{1}(Z = z)(W - p(z, X)).$$

Set $R_{q,0,k} = R_{q,l_k+1,k} \equiv 0$ for the structural integration boundaries. The normalization treats a tied policy value as one boundary; summing category-specific inverse-probability residuals without this normalization would count the same boundary once for every instrument in $T_{j,k}$. Multiple instruments may still share a regular policy level. Part (ii) of ?? 4.2 only prevents such a nonempty policy level from coinciding with the structural left boundary $q_{0,k} = p_k$ when $k < K$.

Score for $\mathbb{E}[J_{j,k}(X)]$. Aggregating the expectation residuals and the structural parameter influences, the complete uncentered orthogonal score for the integral component $J_{j,k}(X)$ is:

$$\begin{aligned}
\psi_{j,k}^J(O; \eta) = & J_{j,k}(X) + IF_{J^+}(O) - IF_{J^-}(O) \\
& + \nu_{j,k}(X) \left(p_{k+1}(X) M_{j,k}^+(X) - p_k(X) M_{j,k}^-(X) - (q_{j,k}(X) - p_k(X)) \right. \\
& \quad \left. + IF_{M_{j,k}^+}(O) - IF_{M_{j,k}^-}(O) \right) \\
& - \nu_{j+1,k}(X) \left(p_{k+1}(X) M_{j+1,k}^+(X) - p_k(X) M_{j+1,k}^-(X) - (q_{j+1,k}(X) - p_k(X)) \right. \\
& \quad \left. + IF_{M_{j+1,k}^+}(O) - IF_{M_{j+1,k}^-}(O) \right) \\
& + \alpha_{p_k}(X) \frac{1}{\pi_k(X)} \mathbb{1}(Z \in S_k) (W - p_k(X)) \\
& + \alpha_{p_{k+1}}(X) \frac{1}{\pi_{k+1}(X)} \mathbb{1}(Z \in S_{k+1}) (W - p_{k+1}(X)) \\
& - \nu_{j,k}(X) R_{q,j,k}(O) \\
& + \nu_{j+1,k}(X) R_{q,j+1,k}(O).
\end{aligned}$$

Score for $\mathbb{E}[J_{\text{full},k}(X)]$. The analogous score for the full-interval component omits the quantile-constraint terms:

$$\begin{aligned}
\psi_{\text{full},k}^J(O; \eta) = & J_{\text{full},k}(X) \\
& + \frac{1}{\pi_{k+1}(X)} \mathbb{1}(Z \in S_{k+1}) W \left(Y - J_{\text{full},k}^+(X) \right) \\
& - \frac{1}{\pi_k(X)} \mathbb{1}(Z \in S_k) W \left(Y - J_{\text{full},k}^-(X) \right) \\
& + J_{\text{full},k}^+(X) \frac{1}{\pi_{k+1}(X)} \mathbb{1}(Z \in S_{k+1}) (W - p_{k+1}(X)) \\
& - J_{\text{full},k}^-(X) \frac{1}{\pi_k(X)} \mathbb{1}(Z \in S_k) (W - p_k(X)).
\end{aligned}$$

Remark (notation). $\psi_{\text{full},k}^J(O; \eta)$ is the uncentered full-interval component score for $\mathbb{E}[J_{\text{full},k}(X)]$, for $0 \leq k < K$, with no quantile-constraint correction and with structural-zero rather than conditional-regression terms at $k = 0$. In Lemma D.1, Theorem 4.7: In the continuous-outcome case, use the displayed treated-outcome and propensity residual corrections, omitting all minus-side terms at $k = 0$; in the binary-outcome case, use $\psi_{\text{full},k}^J = J_{\text{full},k} + R_{P,k+1} - R_{P,k}$ with $R_{P,0} = 0$.

Every minus-side term in the two preceding score displays, including the p_k boundary correction, is omitted when $k = 0$, according to the structural-boundary convention above.

Product scores. To account for the estimation of the weighting regressions $\gamma(X)$, we construct the orthogonal product scores using the chain rule for influence functions. The uncentered product score for the fractional interval component is:

$$\psi_{j,k}^{\text{prod}}(O; \eta) = \gamma_{j,k}(X) \psi_{j-1,k}^J(O; \eta) + J_{j-1,k}(X) \left(\mathbb{1}\left(Z \in \bigcup_{l=j}^{l_k} T_{l,k}\right) - \gamma_{j,k}(X) \right),$$

and similarly for the full interval component:

$$\begin{aligned} \psi_{\text{full},k}^{\text{prod}}(O; \eta) &= \gamma_{\text{full},k}(X) \psi_{\text{full},k}^J(O; \eta) \\ &+ J_{\text{full},k}(X) \left(\mathbb{1}\left(Z \in \bigcup_{m=k+1}^K \bigcup_{l=0}^{l_m} T_{l,m}\right) - \mathbb{1}\left(Z \in \bigcup_{m=k+1}^K S_m\right) - \gamma_{\text{full},k}(X) \right). \end{aligned}$$

Trivial bound score. The uncentered orthogonal score for the trivial bound component on the final interval, $\Delta_K(X)$, is:

$$\begin{aligned} \psi_{\Delta_K}(O; \eta) &= y_{\min} \mathbb{1}\left(Z \in \bigcup_{j=1}^{l_K} T_{j,K}\right) (q(Z, X) - p_K(X)) \\ &+ y_{\min} \mathbb{1}\left(Z \in \bigcup_{j=1}^{l_K} T_{j,K}\right) \frac{\partial \phi}{\partial p}(Z, X, p(Z, X)) (W - p(Z, X)) \\ &- y_{\min} \gamma_K(X) \frac{1}{\pi_K(X)} \mathbb{1}(Z \in S_K) (W - p_K(X)). \end{aligned}$$

C.2 Binary Outcome Score Components

We provide the complete score components for the binary outcome setting of Section 4.1.2. For $k \geq 1$, the standard IPW residuals for $P_{1,k}(X)$ and $p_k(X)$ are:

$$\begin{aligned} R_{P,k}(O) &= \frac{\mathbb{1}(Z \in S_k)}{\pi_k(X)} (YW - P_{1,k}(X)), \\ R_{p,k}(O) &= \frac{\mathbb{1}(Z \in S_k)}{\pi_k(X)} (W - p_k(X)). \end{aligned}$$

At the known structural boundary, set $R_{P,0} = R_{p,0} \equiv 0$; neither residual involves π_0 or a conditional regression on S_0 . The key pathwise derivative of m with respect to $J_{\text{full},k}(X)$ is:

$$\frac{\partial m}{\partial J_{\text{full},k}}(X) = \gamma_{\text{full},k}(X) + \sum_{j=1}^{l_k} \gamma_{j,k}(X) \mathbb{1}(q_{j-1,k}(X) < p_{k+1}(X) - J_{\text{full},k}(X) < q_{j,k}(X)),$$

which equals the weight function evaluated at the lower integration boundary. This has a natural interpretation: an infinitesimal increase in $J_{\text{full},k}$ extends the integration region $[p_{k+1} - J_{\text{full},k}, p_{k+1}]$ to the left, and the marginal contribution is precisely ω at the boundary. *Remark (notation).* $R_{P,k}(O)$ is the binary-outcome treated-success IPW residual for observed positive-propensity levels $1 \leq k \leq K$, with $R_{P,0} \equiv 0$ at the fixed structural boundary, without evaluating π_0 or a conditional regression on S_0 .

Auxiliary indicators. To state the orthogonal score for $\mathbb{E}[h_{j,k}^-(X)]$, we introduce the following notation. Let

$$I_{j,k}^-(X) := \mathbb{1}(q_{j,k}(X) > \max(q_{j-1,k}(X), p_{k+1}(X) - J_{\text{full},k}(X)))$$

be the active set indicator for $h_{j,k}^-$, i.e., $I_{j,k}^- = \mathbb{1}(h_{j,k}^- > 0)$. Define the boundary indicator

$$D_{j,k}^-(X) := \mathbb{1}(q_{j-1,k}(X) < p_{k+1}(X) - J_{\text{full},k}(X) < q_{j,k}(X)),$$

which equals one precisely when the lower integration boundary $p_{k+1} - J_{\text{full},k}$ falls strictly inside the sub-interval $(q_{j-1,k}, q_{j,k})$, and the dominance indicator

$$E_{j,k}^-(X) := \mathbb{1}(q_{j-1,k}(X) > p_{k+1}(X) - J_{\text{full},k}(X)),$$

which equals one when the inner max in $h_{j,k}^-$ selects $q_{j-1,k}$. The two terms on the right partition the active set: $I_{j,k}^- = D_{j,k}^- + E_{j,k}^- \cdot \mathbb{1}(q_{j,k} > q_{j-1,k})$ (the two indicators are mutually exclusive).

Remark (notation). $I_{j,k}^-(x) = \mathbb{1}(q_{j,k}(x) > \max\{q_{j-1,k}(x), p_{k+1}(x) - J_{\text{full},k}(x)\}) = \mathbb{1}(h_{j,k}^-(x) > 0)$ is the binary-score active-set indicator for $k = 0, \dots, K-1$ and $j = 1, \dots, l_k$, with $q_{0,k}(x) = p_k(x)$. $D_{j,k}^-(x) = \mathbb{1}(q_{j-1,k}(x) < p_{k+1}(x) - J_{\text{full},k}(x) < q_{j,k}(x))$ is the binary-score boundary indicator for $k = 0, \dots, K-1$ and $j = 1, \dots, l_k$, with $q_{0,k}(x) = p_k(x)$; it is zero when the integration boundary equals either sub-interval endpoint. $E_{j,k}^-(x) = \mathbb{1}(q_{j-1,k}(x) > p_{k+1}(x) - J_{\text{full},k}(x))$ is the binary-score dominance indicator for $k = 0, \dots, K-1$ and $j = 1, \dots, l_k$, with $q_{0,k}(x) = p_k(x)$; its strict inequality assigns zero at equality.

Pathwise derivatives. The pathwise derivatives of $h_{j,k}^-$ with respect to the nuisance parameters are as follows. Since $h_{j,k}^-$ depends on $P_{1,m}$ only through $J_{\text{full},k} = P_{1,k+1} - P_{1,k}$:

$$\alpha_{P,k+1}^{(j,k)}(X) = D_{j,k}^-(X), \quad \alpha_{P,k}^{(j,k)}(X) = -D_{j,k}^-(X), \quad \alpha_{P,m}^{(j,k)} = 0 \text{ for } m \notin \{k, k+1\}.$$

For the propensity scores, p_{k+1} enters $h_{j,k}^-$ directly through the boundary $p_{k+1} - J_{\text{full},k}$, and p_k enters only when $j = 1$ (since $q_{0,k} = p_k$):

$$\alpha_{p,k+1}^{(j,k)}(X) = -D_{j,k}^-(X), \quad \alpha_{p,k}^{(j,k)}(X) = \begin{cases} -E_{1,k}^-(X) I_{1,k}^-(X) & \text{if } j = 1, \\ 0 & \text{if } j \geq 2. \end{cases}$$

Remark (notation). $\alpha_{P,k+1}^{(j,k)}(x)$ is the coefficient of a perturbation of $P_{1,k+1}(x)$ in the binary component $h_{j,k}^-(x)$, indexed by $k = 0, \dots, K-1$ and $j = 1, \dots, l_k$, holding the propensity and alternative-policy levels and $P_{1,k}$ fixed. $\alpha_{p,k+1}^{(j,k)}(x)$ is the direct upper-propensity coefficient for the binary score component $h_{j,k}^-$, indexed by $k = 0, \dots, K-1$ and $j = 1, \dots, l_k$, holding $P_{1,k}$, $P_{1,k+1}$, and the alternative-policy levels fixed. $\alpha_{p,k}^{(j,k)}(x)$ denotes the direct lower-propensity coefficient in the binary score for $h_{j,k}^-$, indexed by $k = 0, \dots, K-1$ and $j = 1, \dots, l_k$, holding the success regressions and alternative-policy levels fixed; at $k = 0$ its residual correction is omitted because $p_0 \equiv 0$ is known.

Score for $h_{j,k}^-$. The complete orthogonal score for $\mathbb{E}[h_{j,k}^-]$ is:

$$\begin{aligned} \psi_{j,k}^h(O; \eta) &= h_{j,k}^-(X) + D_{j,k}^-(X) (R_{P,k+1}(O) - R_{P,k}(O)) \\ &\quad + \alpha_{p,k+1}^{(j,k)}(X) R_{p,k+1}(O) + \alpha_{p,k}^{(j,k)}(X) R_{p,k}(O) \\ &\quad + I_{j,k}^-(X) R_{q,j,k}(O) \\ &\quad - E_{j,k}^-(X) I_{j,k}^-(X) R_{q,j-1,k}(O). \end{aligned}$$

The third line corrects for the estimation of the upper boundary $q_{j,k}$ through instruments in $T_{j,k}$, and the fourth line corrects for the lower boundary $q_{j-1,k}$ through instruments in $T_{j-1,k}$; the latter matters only when $q_{j-1,k}$ dominates $p_{k+1} - J_{\text{full},k}$ in the inner max (i.e., $E_{j,k}^- = 1$). When $j = 1$, the lower boundary $q_{0,k} = p_k$ does not depend on individual instruments, so the fourth line vanishes and its correction is absorbed into $\alpha_{p,k}^{(1,k)}$.

Remark (notation). $R_{p,k}(O)$ is the baseline propensity IPW residual for observed positive-propensity levels $1 \leq k \leq K$, with $R_{p,0} \equiv 0$ at the fixed structural boundary $p_0 \equiv 0$, without evaluating π_0 or a conditional regression on S_0 . $\psi_{j,k}^h(O; \eta)$ is the uncentered binary-outcome score for $\mathbb{E}[h_{j,k}^-]$, indexed by $0 \leq k < K$ and $1 \leq j \leq l_k$, with $R_{P,0} = R_{p,0} = 0$ and $R_{q,0,k} = 0$, so the lower-policy-boundary correction vanishes at $j = 1$ and its propensity correction is included in $\alpha_{p,k}^{(1,k)}$.

Product scores. The product score for the fractional interval component is:

$$\psi_{j,k}^{\text{prod}}(O; \eta) = \gamma_{j,k}(X) \psi_{j,k}^h(O; \eta) + h_{j,k}^-(X) \left(\mathbb{1}\left(Z \in \bigcup_{l=j}^{l_k} T_{l,k}\right) - \gamma_{j,k}(X) \right).$$

For the full-interval component, the orthogonal score for $J_{\text{full},k}(X) = P_{1,k+1}(X) - P_{1,k}(X)$ is simply $\psi_{\text{full},k}^J(O; \eta) = J_{\text{full},k}(X) + R_{P,k+1}(O) - R_{P,k}(O)$, and the full-interval product score is:

$$\begin{aligned} \psi_{\text{full},k}^{\text{prod}}(O; \eta) &= \gamma_{\text{full},k}(X) \psi_{\text{full},k}^J(O; \eta) \\ &\quad + J_{\text{full},k}(X) \left(\mathbb{1}\left(Z \in \bigcup_{m=k+1}^K \bigcup_{l=0}^{l_m} T_{l,m}\right) - \mathbb{1}\left(Z \in \bigcup_{m=k+1}^K S_m\right) - \gamma_{\text{full},k}(X) \right). \end{aligned}$$

The trivial bound score $\psi_{\Delta_K}(O; \eta)$ is identical to the continuous outcome case (see Section C.1), since $\Delta_K(X)$ does not involve the quantile structure.

D Proofs in Section 4

D.1 Discrete Instrument

Both parts of Theorem 4.7 rest on the following two properties of the corresponding score. Both follow from term-by-term differentiation of the score formulas in Section C, which we carry out in Section O4 of the Online Appendix; the argument is self-contained there and is used nowhere else.

Lemma D.1 (Moment Condition and Neyman Orthogonality). *Let η_0 denote the vector of true nuisance parameter values, let $\theta_0 = \underline{\theta}_{\omega,1}$, and let $\mathcal{T}$ be the set of nuisance directions δ such that $\eta_0 + t\delta$ is an admissible nuisance value for all t in a neighborhood of 0.*

- (i) *Under the hypotheses of Theorem 4.7(i), the continuous-outcome score ψ assembled in Section C.1 satisfies*

$$\mathbb{E}[\psi(O; \eta_0)] = \theta_0 \quad \text{and} \quad \frac{d}{dt} \mathbb{E}[\psi(O; \eta_0 + t\delta)] \Big|_{t=0} = 0 \quad \text{for every } \delta \in \mathcal{T},$$

where θ_0 is the target (4.1).

- (ii) *Under the hypotheses of Theorem 4.7(ii), the same two properties hold for the binary-outcome score (B.1), with θ_0 the target (4.4).*

Additional condition. $K \geq 1$.

Remark (notation). η_0 is the vector of population nuisance values for the applicable score, including its known structural-boundary conventions, distinct from the fitted vector $\hat{\eta}$. In Lemma D.1, Theorem 4.7: In part (i), η_0 supplies the true nuisances for the continuous-outcome score, including boundary quantiles and conditional indicator and treated-outcome regressions; in part (ii), it supplies the binary-outcome nuisances p_k , $P_{1,k}$, π_k , $\tau_{j,k}$, $a_{j,k}$, $\gamma_{\text{full},k}$, $\gamma_{j,k}$, and γ_K . $\mathcal{T}$ consists of directions in the nuisance function space for which $\eta_0 + t\delta$ remains admissible for every t in a neighborhood of zero; these preserve fixed structural components and satisfy $\sum_{z \in T_{j,k}} \delta a_z(X) = 0$ for each within-set share vector.

Proof of Theorem 4.7(i). Let η_0 denote the vector of true nuisance parameter values and $\hat{\eta}$ the estimates trained on the auxiliary sample I_1 . Write $\theta_0 = \underline{\theta}_{\omega,1}$ for the true parameter value. Recall the target functional:

$$\theta_0 = \mathbb{E} \left[\sum_{k=0}^{K-1} \left(\gamma_{\text{full},k}(X) J_{\text{full},k}(X) + \sum_{j=1}^{l_k} \gamma_{j,k}(X) J_{j-1,k}(X) \right) + \Delta_K(X) \right],$$

and the total Neyman-orthogonal score:

$$\psi(O; \eta) = \sum_{k=0}^{K-1} \left(\psi_{\text{full},k}^{\text{prod}}(O; \eta) + \sum_{j=1}^{l_k} \psi_{j,k}^{\text{prod}}(O; \eta) \right) + \psi_{\Delta_K}(O; \eta),$$

where each product score debiases the corresponding term in the target and all components are given in Section C.1. By Lemma D.1(i), ψ has mean θ_0 at η_0 and is Neyman orthogonal there; what remains is the standard DML convergence argument with sample splitting.

Exact recovery of the level sets. For each z , let $m_p(z) = \mathbb{E}[p(z, X)]$, $m_q(z) = \mathbb{E}[q(z, X)]$, and define the training-sample conditional means $\hat{m}_p(z) = \mathbb{E}[\hat{p}(z, X) \mid I_1]$ and $\hat{m}_q(z) = \mathbb{E}[\hat{q}(z, X) \mid I_1]$. Conditional on I_1 , the covariates $\{X_i : i \in I_2\}$ are i.i.d., while $\hat{p}, \hat{q}$ are fixed and bounded. Hence, because $\mathcal{Z}$ is finite,

$$\max_{z \in \mathcal{Z}} |\bar{p}_2(z) - \hat{m}_p(z)| = O_P(n^{-1/2}), \quad \max_{z \in \mathcal{Z}} |\bar{q}_2(z) - \hat{m}_q(z)| = O_P(n^{-1/2})$$

by conditional Chebyshev. Moreover,

$$|\hat{m}_p(z) - m_p(z)| \leq \|\hat{p}(z, \cdot) - p(z, \cdot)\|_{L_2}, \quad |\hat{m}_q(z) - m_q(z)| \leq C_\phi \|\hat{p}(z, \cdot) - p(z, \cdot)\|_{L_2},$$

where $C_\phi < \infty$ follows from the uniformly bounded derivative in ?? 4.1. Because $p(z, x) \in [0, 1]$, clipping $\hat{p}$ to $[0, 1]$ cannot increase its L_2 error; the stated range of ϕ then keeps $\hat{q} = \phi(z, x, \hat{p})$ in $[0, 1]$. The pre-recovery first-stage rate stated in the theorem and finite $\mathcal{Z}$ therefore give uniform convergence of $(\bar{p}_2, \bar{q}_2)$ to (m_p, m_q) at rate $o_P(n^{-1/4})$.

For $z \in S_0$, $\bar{p}_2(z) = o_P(n^{-1/4}) = o_P(\rho_n)$; for each remaining baseline category, $m_p(z) > 0$ and $\rho_n \rightarrow 0$. Thus $\widehat{S}_0 = S_0$ with probability approaching one. After removing those baseline vertices, the equality classes and their ordering are invariant in x under part (i) of ?? 4.2; equal score functions have equal marginal means, whereas two distinct ordered score functions have marginal means separated by more than c_{gap} . With probability approaching one, vertices from the same population level are consequently connected by an edge of the $c_{\text{gap}}/2$ graph, and vertices from distinct levels have no edge between them. Its connected components and their order recover exactly $\widehat{S}_k = S_k$ and $\widehat{T}_{j,k} = T_{j,k}$.

Let $\mathcal{E}_n$ denote this exact-recovery event and let the oracle estimator use the true level sets while otherwise applying the same two-fold procedure. On $\mathcal{E}_n$, every remaining nuisance fit, evaluation score, point estimate, and variance estimate agrees exactly with its oracle counterpart. In particular, the root- n difference between the feasible and oracle point estimators is exactly zero on $\mathcal{E}_n$. Although the feasible estimator used all I_2 covariates in the finite classification step, this exact equality and $\mathbb{P}(\mathcal{E}_n^c) \rightarrow 0$ make all feasible statistics asymptotically equivalent to their oracle versions. We may therefore prove the remaining claims for the oracle estimator, whose non-set nuisances are trained on I_1 and evaluated on I_2 .

Remark (notation). $\mathcal{E}_n$ is the event of simultaneous exact recovery of all baseline and policy level sets with their population ordering and bin indices, including $\widehat{S}_0 = S_0$, by the procedure using first-stage fits on I_1 and their averages over $\{X_i : i \in I_2\}$.

Second-order remainder bound. Let $r_{\text{reg},n}$ be the largest L_2 error among the finitely many fitted outcome regressions, propensity regressions, pooled policy levels, and boundary quantiles, and let $r_{\text{wt},n}$ be the largest L_2 error among $\pi_k, \tau_{j,k}, a_{j,k}, \gamma_{j,k}, \gamma_{\text{full},k}$, and γ_K . Conditioning on I_1 and expanding the full score expectation gives four remainder types: regression error times an IPW-denominator error; a product-score error $(\widehat{\gamma} - \gamma)(\widehat{J} - J)$; the simplex pooling product $\sum_z (\widehat{a}_z - a_z)(\widehat{q}_z - q_{j,k})$; and Taylor remainders from ϕ and the conditional CDFs at the estimated quantile boundaries. Consequently,

$$|\mathbb{E}[\psi(O; \widehat{\eta}) \mid I_1] - \theta_0| \leq C(r_{\text{reg},n}^2 + r_{\text{reg},n}r_{\text{wt},n}). \quad (\text{D.1})$$

The clipped IPW denominators make reciprocal errors Lipschitz. To control a quantile-boundary remainder, let $e_{j,k}(X) = \widehat{\nu}_{j,k}(X) - \nu_{j,k}(X)$ and let $\delta_\nu > 0$ be the fixed radius of the neighborhood in the local bounded-density condition. On $\{|e_{j,k}| \leq \delta_\nu\}$, the integral bound above is at most $Ce_{j,k}^2$. Because both quantiles are in $[y_{\min}, y_{\max}]$, the remainder is bounded by $B_Y = y_{\max} - y_{\min}$ on the complementary event, whose probability is at most

$\delta_\nu^{-2} \|e_{j,k}\|_{L_2}^2$. Hence its expectation is $O(\|e_{j,k}\|_{L_2}^2)$. The uniformly Lipschitz derivative of ϕ gives the corresponding quadratic policy remainder. Since $r_{\text{reg},n}, r_{\text{wt},n} = o_P(n^{-1/4})$, the right-hand side of (D.1) is $o_P(n^{-1/2})$.

Remark (notation). $r_{\text{reg},n}$ is the maximum population L_2 error of the fitted outcome regressions, propensity regressions, pooled policy levels, and boundary quantiles, with integration over the marginal law of X while holding the fitted functions fixed. In Theorem 4.7: In the continuous-outcome proof, the errors are those of the oracle nuisance fits trained on I_1 using the population level sets; hence their magnitudes are fixed conditional on I_1 and satisfy $r_{\text{reg},n} = o_P(n^{-1/4})$. $r_{\text{wt},n}$ is the maximum of the componentwise $L_2(\mathbb{P}_{\text{obs}}(dx))$ errors of the fitted weights $\hat{\pi}_k, \hat{\tau}_{j,k}, \hat{a}_{j,k}(z, \cdot), \hat{\gamma}_{j,k}, \hat{\gamma}_{\text{full},k}, \hat{\gamma}_K$ relative to their population counterparts, over the finitely many fitted indices and instrument categories.

Asymptotic normality. For the oracle estimator, the sample-splitting construction makes $\hat{\eta}$ (trained on I_1) independent of $\{O_i\}_{i \in I_2}$. Decompose:

$$\sqrt{n/2} (\hat{\theta}_{\omega,1} - \theta_0) = \underbrace{\frac{1}{\sqrt{n/2}} \sum_{i \in I_2} (\psi(O_i; \eta_0) - \theta_0)}_{\text{(I)}} + \underbrace{\frac{1}{\sqrt{n/2}} \sum_{i \in I_2} (\psi(O_i; \hat{\eta}) - \psi(O_i; \eta_0))}_{\text{(II)}}. \quad (\text{D.2})$$

For term (I): the summands are i.i.d. with mean zero and finite variance $\sigma^2 = \text{Var}(\psi(O; \eta_0))$. Finiteness follows from the compact outcome domain, ?? 4.5 (bounded IPW weights), and boundedness of ω and $\frac{\partial \phi}{\partial p}$. By the Lindeberg–Lévy central limit theorem, (I) $\xrightarrow{d} \mathcal{N}(0, \sigma^2)$.

For term (II): condition on I_1 (which fixes $\hat{\eta}$) and write

$$\text{(II)} = \sqrt{n/2} \cdot B_n + \frac{1}{\sqrt{n/2}} \sum_{i \in I_2} \epsilon_i,$$

where $B_n = \mathbb{E}[\psi(O; \hat{\eta}) - \psi(O; \eta_0) \mid I_1] = \mathbb{E}[\psi(O; \hat{\eta})] - \theta_0$ and $\epsilon_i = (\psi(O_i; \hat{\eta}) - \psi(O_i; \eta_0)) - B_n$ are conditionally i.i.d. zero-mean random variables. By (D.1), $\sqrt{n/2} |B_n| = O_P(\sqrt{n} \|\hat{\eta} - \eta_0\|_{L_2}^2) = o_P(1)$. For the centered empirical term, the conditional variance satisfies

$$\text{Var}(\epsilon_1 \mid I_1) \leq \mathbb{E}[(\psi(O; \hat{\eta}) - \psi(O; \eta_0))^2 \mid I_1].$$

The score $\psi(O; \eta)$ is Lipschitz in the smooth nuisance components (conditional expectations, propensity scores, instrument probabilities) under ?? 4.5. For the non-smooth indicator terms, $|\mathbb{1}(\hat{\nu}_{j,k} \leq Y) - \mathbb{1}(\nu_{j,k}^0 \leq Y)|$ is nonzero only when Y lies between the two boundaries. On $\{|e_{j,k}| \leq \delta_\nu\}$, the local density bound makes this probability at most $C|e_{j,k}|$; on the complement it is at most one, and Markov's inequality bounds the probability of that

complement by $\delta_\nu^{-2} \|e_{j,k}\|_{L_2}^2$. Thus each indicator difference converges to zero in L_2 whenever $\|e_{j,k}\|_{L_2} = o_P(1)$. Together with the smooth-component rates, this gives $\text{Var}(\epsilon_i | I_1) = o_P(1)$, and $\frac{1}{\sqrt{n/2}} \sum_{i \in I_2} \epsilon_i = o_P(1)$ by Chebyshev's inequality. Combining, $(\text{II}) = o_P(1)$.

By Slutsky's theorem applied to (D.2), we conclude

$$\sqrt{n/2} (\hat{\theta}_{\omega,1} - \underline{\theta}_{\omega,1}) \xrightarrow{d} \mathcal{N}(0, \sigma^2).$$

Consistency of the variance estimator. For the oracle estimator, conditional on I_1 , the evaluation observations are i.i.d. and $\hat{\eta}$ is fixed. The preceding L_2 score-convergence argument gives $\|\psi(\cdot; \hat{\eta}) - \psi(\cdot; \eta_0)\|_{L_2} = o_P(1)$. A conditional law of large numbers for the squared, centered scores, followed by this L_2 convergence and $\hat{\theta}_{\omega,1} \xrightarrow{P} \theta_0$, therefore gives $\hat{\sigma}^2 = \frac{1}{|I_2|} \sum_{i \in I_2} (\hat{\psi}(O_i; \hat{\eta}) - \hat{\theta}_{\omega,1})^2 \xrightarrow{P} \sigma^2$. Exact equality on $\mathcal{E}_n$ transfers both conclusions to the feasible estimator and variance estimate. $\square$

Proof of Theorem 4.7(ii). Let η_0 denote the true nuisance parameter values and $\hat{\eta}$ the estimates trained on I_1 . Write $\theta_0 = \underline{\theta}_{\omega,1}$. Recall the target functional (4.4) and the orthogonal score (B.1), which decomposes as $\psi = \sum_k (\psi_{\text{full},k}^{\text{prod}} + \sum_j \psi_{j,k}^{\text{prod}}) + \psi_{\Delta_K}$. By Lemma D.1(ii), this score has mean θ_0 at η_0 and is Neyman orthogonal there; what remains is the standard DML convergence argument.

Exact recovery of the level sets. The conditional concentration, graph-separation, and oracle-equivalence argument in the continuous-outcome proof applies without change. Thus $\mathbb{P}(\mathcal{E}_n) \rightarrow 1$, and on $\mathcal{E}_n$ the feasible binary estimator equals the oracle estimator that uses $(S_k, T_{j,k})$ and fits all non-set nuisances on I_1 . We prove the remaining claims for that oracle estimator.

Rearrangement. The true values satisfy $0 = P_{1,0}(x) \leq \dots \leq P_{1,K}(x)$ for all x because $J_{\text{full},k}(x) \geq 0$. After clipping $\hat{P}_{1,k}$ to $[0, 1]$, sort only $(\hat{P}_{1,1}, \dots, \hat{P}_{1,K})$ and retain the fixed boundary $\hat{P}_{1,0}^* = 0$. By the classical rearrangement inequality, for each x the sorted nonboundary vector is at least as close in aggregate squared error to the true monotone nonboundary vector as the original vector:

$$\sum_{k=1}^K |\hat{P}_{1,k}^*(x) - P_{1,k}(x)|^2 \leq \sum_{k=1}^K |\hat{P}_{1,k}(x) - P_{1,k}(x)|^2 \quad \text{for all } x.$$

After integrating over x , this aggregate inequality and fixed K imply $\|\hat{P}_{1,k}^* - P_{1,k}\|_{L_2} = o_P(n^{-1/4})$ for every k ; no componentwise contraction claim is needed. Defining $\hat{J}_{\text{full},k} =$

$\widehat{P}_{1,k+1}^* - \widehat{P}_{1,k}^*$ then gives

$$\|\widehat{J}_{\text{full},k} - J_{\text{full},k}\|_{L_2} \leq \|\widehat{P}_{1,k+1}^* - P_{1,k+1}\|_{L_2} + \|\widehat{P}_{1,k}^* - P_{1,k}\|_{L_2} = o_P(n^{-1/4}).$$

Therefore, rearrangement and clipping preserve all stated convergence rates while keeping each gap coherent with the single $\widehat{P}^*$ vector used in the score.

Second-order remainder bound. For the oracle estimator, define the conditional score bias

$$B_n = \mathbb{E}[\psi(O; \widehat{\eta}) - \psi(O; \eta_0) \mid I_1].$$

This is the relevant DML remainder; expanding only the plug-in target $m(X; \eta)$ would omit the estimated residual corrections. Let $e_n(X)$ be the maximum absolute error among the finitely many affine arguments of the maxima in $h_{j,k}^-$, and set $c_* = \min\{c_h, c_{\text{gap}}\}$. The nuisance-rate assumption gives $\|e_n\|_{L_2} = o_P(n^{-1/4})$. The separations in ?? 4.2 and ?? 4.6 control the active branches. At $k = 0, j = 1$, the latter assumption gives $p_1 - J_{\text{full},0} \geq c_h$ and $|p_1 - J_{\text{full},0} - q_{1,0}| \geq c_h$, covering the structural zero. Thus a wrong active branch can occur only on $A_n = \{e_n \geq c_*/2\}$. Hence

$$\mathbb{P}(A_n) \leq 4c_*^{-2}\|e_n\|_{L_2}^2 = o_P(n^{-1/2}), \quad \mathbb{E}[e_n \mathbf{1}(A_n)] \leq 2c_*^{-1}\|e_n\|_{L_2}^2 = o_P(n^{-1/2}).$$

Thus branch-selection contributes $o_P(n^{-1/2})$ to B_n . Let $r_{\text{core},n}$ be the largest L_2 error among the finitely many $P_{1,k}$, p_k , instrument-specific $p(z, \cdot)$, and pooled $q_{j,k}$ nuisances, and let $r_{\text{wt},n}$ be the largest error among π_k , $\tau_{j,k}$, $a_{j,k}$, $\gamma_{j,k}$, $\gamma_{\text{full},k}$, and γ_K . On A_n^c , the active indicators are fixed at their population branches. Expanding the conditional expectation of the full score and using the cancellations above leaves only: a Riesz-coefficient error times a residual-mean error; an IPW-denominator error times a regression error; a product-score error $(\widehat{\gamma} - \gamma)(\widehat{h} - h)$ or $(\widehat{\gamma} - \gamma)(\widehat{J}_{\text{full}} - J_{\text{full}})$; the simplex pooling product $\sum_z (\widehat{a}_z - a_z)(\widehat{q}_z - q_{j,k})$; and the quadratic Taylor remainder of ϕ . Therefore, for a fixed C ,

$$|B_n| \leq C(r_{\text{core},n}^2 + r_{\text{core},n}r_{\text{wt},n}) + C\mathbb{E}[e_n \mathbf{1}(A_n)] = o_P(n^{-1/2}).$$

The clipping of $\widehat{\pi}_k, \widehat{\tau}_{j,k}$ makes the reciprocal maps uniformly Lipschitz, and the simplex restriction removes a standalone first-order share error. The same grouped expansion without taking its conditional mean gives $\|\psi(\cdot; \widehat{\eta}) - \psi(\cdot; \eta_0)\|_{L_2} = o_P(1)$; on A_n this follows from bounded clipped scores and $\mathbb{P}(A_n) = o_P(n^{-1/2})$. Conditional Chebyshev therefore makes the centered evaluation-sample fluctuation of the score difference $o_P(1)$ after root- n scaling.

Remark (notation). $r_{\text{core},n}$ is the maximum of the $L_2(\mathbb{P}_X)$ norms of fitted-minus-population errors for the finitely many $P_{1,k}$, p_k , instrument-specific $p(z, \cdot)$, and pooled $q_{j,k}$ nuisances used

by the binary oracle score, with the fitted $P_{1,k}$ taken after clipping and rearrangement and the structural zero boundaries retained.

Asymptotic normality. By the standard DML argument [Chernozhukov et al., 2018], the moment condition, Neyman orthogonality, and the $o_P(n^{-1/2})$ remainder bound together imply:

$$\sqrt{n/2}(\hat{\theta}_{\omega,1} - \theta_0) = \frac{1}{\sqrt{n/2}} \sum_{i \in I_2} (\psi(O_i; \eta_0) - \theta_0) + o_P(1) \xrightarrow{d} \mathcal{N}(0, \sigma^2),$$

where $\sigma^2 = \text{Var}(\psi(O; \eta_0))$.

Consistency of the variance estimator. For the oracle estimator, conditional on I_1 , the evaluation observations are i.i.d. and the fitted score is fixed. The gap condition makes the active branch locally constant with probability approaching one, and the preceding rate argument gives $\|\psi(\cdot; \hat{\eta}) - \psi(\cdot; \eta_0)\|_{L_2} = o_P(1)$. A conditional law of large numbers for the squared, centered scores, followed by this L_2 convergence and consistency of the point estimator, gives $\hat{\sigma}^2 \xrightarrow{P} \sigma^2$. Exact equality on $\mathcal{E}_n$ transfers the point-estimator and variance conclusions to the feasible binary procedure. $\square$

D.2 Continuous Instrument

Proof of Theorem 4.10. Let $\psi = \psi_1 + \psi_2 + \psi_3$, where

$$\begin{aligned} \psi_1(O) &= - \int_{\min\{\underline{p}(X), q(Z, X)\}}^{\underline{p}(X)} Q_{1, \underline{p}|X} \left(\frac{v}{\underline{p}(X)} \right) dv, \\ \psi_2(O) &= g_1(q^{\text{tr}}(Z, X), X) - g_1(P, X), \\ \psi_3(O) &= y_{\min}(q(Z, X) - \bar{p}(X))_+. \end{aligned}$$

The first term is zero when $\underline{p}(X) = 0$. Let the estimated components be

$$\begin{aligned} \hat{\psi}_1(O) &= - \int_{\min\{\hat{\underline{p}}(X), \hat{q}(Z, X)\}}^{\hat{\underline{p}}(X)} \hat{Q}_{1, \hat{\underline{p}}|X} \left(\frac{v}{\hat{\underline{p}}(X)} \right) dv, \\ \hat{\psi}_2(O) &= \hat{g}_1(\hat{q}^{\text{tr}}(Z, X), X) - \hat{g}_1(\hat{P}, X), \\ \hat{\psi}_3(O) &= y_{\min}(\hat{q}(Z, X) - \hat{\bar{p}}(X))_+, \end{aligned}$$

again setting the first term to zero when $\hat{\underline{p}}(X) = 0$.

Conditional on $I_0 \cup I_1$, the function $\hat{\psi}$ is fixed and bounded because every nuisance estimate is clipped to its known support. Hence

$$\hat{\theta}_{\omega,1} - \theta_{\omega,1} = (\mathbb{P}_{I_2} - \mathbb{P})\hat{\psi} + \mathbb{P}(\hat{\psi} - \psi),$$

where the first term is $O_P(|I_2|^{-1/2}) = O_P(n^{-1/2})$ by conditional Chebyshev and the split-proportion condition. It remains to bound $\mathbb{E}[\widehat{\psi}_j - \psi_j \mid I_0, I_1]$.

Part (1) of ?? 4.9 directly controls the endpoint $\widehat{p}$, the lower integration limit $\min\{\widehat{p}, \widehat{q}\}$, and the positive-part tail appearing in $\widehat{\psi}_3$.

Because $\widehat{\psi}_3 - \psi_3 = y_{\min}[(\widehat{q}(Z, X) - \widehat{p}(X))_+ - (q(Z, X) - \bar{p}(X))_+]$, Part (1) gives

$$\mathbb{E}[\widehat{\psi}_3 - \psi_3 \mid I_0, I_1] = O_P(r_{e,n}).$$

The two summands in $\widehat{\psi}_2$ are evaluated at $\widehat{q}^{\text{tr}}(Z, X)$ and $\widehat{P}$. Therefore ?? 4.9(2) directly gives

$$\mathbb{E}[\widehat{\psi}_2 - \psi_2 \mid I_0, I_1] = O_P(r_{g,n}).$$

For a bounded function Q define

$$T(A, D, Q) = \begin{cases} A \int_{D/A}^1 Q(t) dt, & A > 0, \\ 0, & A = 0, \end{cases} \quad 0 \leq D \leq A \leq 1.$$

If $\|Q\|_\infty \leq C_Y$, a change of variables shows

$$|T(A', D', Q) - T(A, D, Q)| \leq C_Y (2|A' - A| + |D' - D|). \quad (\text{D.3})$$

This bound remains valid when either A or A' is zero and avoids any lower bound on $\underline{p}(X)$. Apply (D.3) with $Q = \widehat{Q}_{1,\underline{p}|X}$, then use the one-dimensional identity

$$W_1(\widehat{\mu}_{1,\underline{p}|X}, \mu_{1,\underline{p}|X}) = \int_0^1 |\widehat{Q}_{1,\underline{p}|X}(t) - Q_{1,\underline{p}|X}(t)| dt.$$

It follows that

$$\begin{aligned} \mathbb{E}[|\widehat{\psi}_1 - \psi_1| \mid I_0, I_1] &\leq O_P(r_{e,n}) + \mathbb{E}_X \left[\underline{p}(X) W_1(\widehat{\mu}_{1,\underline{p}|X}, \mu_{1,\underline{p}|X}) \mid I_0, I_1 \right] \\ &= O_P(r_{e,n} + r_{\mu,n}). \end{aligned}$$

Combining the three component bounds with the evaluation-sample term proves

$$|\widehat{\underline{\theta}}_{\omega,1} - \underline{\theta}_{\omega,1}| = O_P(n^{-1/2} + r_{e,n} + r_{g,n} + r_{\mu,n}).$$

□

Online Appendix

O1 Notation

Table 2: Important Variables and Mathematical Symbols

Symbol	Description
W	Treatment variable (binary, discrete, or continuous), taking values in $\mathcal{W}$
Z	Instrumental variable, taking values in $\mathcal{Z}$
X	Covariates, taking values in $\mathcal{X}$
Y	Observed outcome variable, taking values in $\mathcal{Y}$
$Y(w)$	Potential outcome under counterfactual treatment $W = w$
U	Unobserved latent heterogeneity (resistance type), normalized to $U \mid X \sim \text{Unif}(0, 1)$
$p(Z, X)$	Baseline conditional propensity score or threshold function, $\mathbb{P}(W = 1 \mid Z, X)$
$q(Z, X)$	Counterfactual policy propensity score
$\omega(x, u)$	Policy-specific weight function for the target parameter
θ_ω	Target causal estimand (e.g., Policy-Relevant Treatment Effect)
$\underline{\theta}_{\omega,1}, \bar{\theta}_{\omega,1}$	Sharp lower and upper bounds for the $W = 1$ component of the target parameter
$y_{\min}, y_{\max}$	Lower and upper logical/support bounds of the potential outcome
π_w	Structural joint distribution (probability measure) of $(Y(w), U, X)$
$\mathcal{R}_w(p)$	Latent region selected into arm w at propensity p : $\mathcal{R}_1(p) = (0, p]$, $\mathcal{R}_0(p) = (p, 1]$
$\mathbb{P}_{\text{obs}}$	The true observed joint distribution of the data (Y, W, Z, X)
$H_{w,x}(p, \cdot)$	Canonical observed sub-measure $\mathbb{P}_{\text{obs}}(Y \in \cdot, W = w \mid P = p, X = x)$; written H_w without covariates

Continued on next page

Table 2 – continued from previous page

Symbol	Description
$g_1^c(p, x)$	Canonical treated regression $\int_{\mathcal{Y}} y H_{1,x}(p, dy) = \mathbb{E}_{\text{obs}}[YW \mid P = p, X = x]$; written g_1 in Section 4
p_k	k -th smallest unique positive value in the range of $p(z)$ or $p(z, x)$ for $k \geq 1$; $p_0 := 0$ is the structural boundary
$I_i, I_i(x)$	A disjoint sub-interval $(p_i, p_{i+1}]$ of the latent variable space for U
$\mu_{1,i}$	Identified conditional probability measure of $Y(1)$ on the interval I_i
$Q_{Y,i}(t)$	Quantile function of the identified conditional distribution $\mu_{1,i}$
$Q_{\omega,i}(t)$	Quantile function of the random variable $\omega(U)$ when $U \sim \text{Unif}(p_i, p_{i+1})$
$\underline{p}(x), \bar{p}(x)$	Conditional lower and upper endpoints of the connected propensity-score support, $\inf_z p(z, x)$ and $\sup_z p(z, x)$
$q^{\text{tr}}(z, x)$	Policy propensity truncated to the identified support, $\min\{\bar{p}(x), \max\{\underline{p}(x), q(z, x)\}\}$
$\mu_{1,\underline{p} x}, Q_{1,\underline{p} x}$	Treated outcome law at the lower propensity boundary $\underline{p}(x)$ and its quantile function
$\underline{p}_k(x), \bar{p}_k(x)$	Endpoints of the k -th connected component of the conditional propensity-score range
$L_k(x), G_k(x)$	k -th identified LIV interval and adjacent OT gap in the general continuous-IV decomposition
$\mu_{1,k x}$	Identified conditional probability measure of $Y(1)$ on the gap $G_k(x)$
$q_{j,i}(x)$	j -th smallest alternative-policy value in $[p_i(x), p_{i+1}(x))$, with the final bin closed at 1
S_k	Baseline instrument level set, $\{z \in \mathcal{Z} : p(z, x) = p_k(x)\}$
$T_{j,k}$	Alternative instrument level set, $\{z \in \mathcal{Z} : q(z, x) = q_{j,k}(x)\}$
$\tau_{j,k}(x)$	Conditional probability $\mathbb{P}(Z \in T_{j,k} \mid X = x)$ of an alternative-policy level set

Continued on next page

Table 2 – continued from previous page

Symbol	Description
$a_{j,k}(z, x)$	Conditional assignment share $\mathbb{P}(Z = z \mid Z \in T_{j,k}, X = x)$ within an alternative-policy level set
c_{gap}	Minimal gap separating true propensity score levels
$\kappa_{j,i}(x)$	Relative quantile position within the i -th baseline interval, $\frac{q_{j,i}(x) - p_i(x)}{p_{i+1}(x) - p_i(x)}$
$\nu_{j,i}(x)$	Conditional quantile function of the outcome evaluated at $\kappa_{j,i}(x)$
$\gamma_{\text{full},i}(X), \gamma_{j,i}(X), \gamma_K(X)$	Weights attached to instruments spanning or terminating within specific intervals; $\gamma_{j,i}, \gamma_K \in [0, 1]$ are conditional probabilities, whereas $\gamma_{\text{full},i} \in [-1, 1]$ differences two of them
$J_{\text{full},i}(X), J_{j,i}(X)$	Conditional expectations/integrals of the outcome within specific quantile bounds
ψ	Neyman-orthogonal score function used for Double Machine Learning estimation
$R_{q,j,k}$	Pooled influence residual for the common policy level $q_{j,k}$
I_1, I_2 (samples)	Auxiliary (nuisance) and main estimation samples for the discrete-instrument DML split; the continuous-instrument estimator adds a propensity sample I_0 . Not to be confused with the latent sub-intervals I_i above
$K(\cdot), h_n$	One-dimensional kernel function and bandwidth sequence used for localized machine learning
$\mathcal{F}_n$	Machine learning hypothesis class
$r_{p,n}$	Sup-norm convergence rate for the estimated propensity score
$r_{g,n}, r_{\mu,n}$	Targeted rates for the continuous-instrument regression evaluations and treated boundary law
$r_{g,\text{loc},n}, r_{Q,n}$	Learner-level L_1 rates for the localized evaluation of $\hat{g}_1$ and for the boundary quantile grid fit
δ_n, M_g, M_Q	Boundary-localization bandwidth, regression grid size, and quantile grid size
$F_{W Z,X}^{-1}$	Conditional quantile function of the treatment W (structural equation for continuous W)

Continued on next page

Table 2 – continued from previous page

Symbol	Description
$\mathcal{U}_{\text{id}}(x, w)$	Identifiable region (union of supports of U) for the continuous/multi-valued treatment
$I_{x,w}(z)$	Interval support of the latent variable for a multi-valued treatment
$J_{x,w}(z)$	Isolated, disjoint sub-region of $I_{x,w}(z)$ derived from the inclusion DAG
$P_{1,k}(X)$	Conditional expectation $\mathbb{E}_{\text{obs}}[YW \mid Z \in S_k, X]$ in the discrete outcome case; in the binary case, $J_{\text{full},k}(X) = P_{1,k+1}(X) - P_{1,k}(X)$
$h_{j,k}^-(X), h_{j,k}^+(X)$	Binary-specialized integrals for the lower and upper bounds in the discrete outcome case

O2 Worked Examples

This section works through two examples that illustrate the identification results of Section 3: a binary-instrument calculation that makes the three-part decomposition behind Theorem 3.4 explicit, and a comparison quantifying how much wider the moment-relaxation bounds can be.

O2.1 Illustrative Example for Theorem 3.4

Example O2.1. Consider a binary instrument with $\mathbb{P}(Z = 1) = 1/2$ and a constant propensity score $p(0) = p(1) = 1/2$. Let the target weight function be $\omega(u) = \mathbb{E}[\mathbf{1}(u \in (0, q_1(Z)))]$, where $q_1(0) = 1/4$ and $q_1(1) = 3/4$.

Simple calculation gives the explicit piecewise form of the weight function:

$$\omega(u) = \begin{cases} 1, & \text{if } u \in (0, 1/4] \\ 1/2, & \text{if } u \in (1/4, 3/4] \\ 0, & \text{if } u \in (3/4, 1] \end{cases}$$

Because the propensity score is constant $1/2$, the data identifies a single constrained interval $I_0 = (0, 1/2]$ where $Y(1)$ is observed (corresponding to $W = 1$), and an unconstrained interval $I_1 = (1/2, 1]$ where $Y(1)$ is entirely unobserved.

Plugging this into our closed-form optimal transport solution, the sharp lower bound elegantly decomposes into three distinct economic components. On I_0 , the optimal coupling pairs the highest weights in $\omega(u)$ with the lowest quantiles of the observed outcome distribution $Y \mid W = 1$. For the unconstrained interval I_1 , because $\omega(u) \geq 0$ here, the optimal solution trivially assigns the logical minimum $y_{\min}$.

Evaluating the integral exactly yields:

$$\begin{aligned} \underline{\theta}_{\omega,1} &= \int_0^{1/2} Q_{Y|W=1}(2u)\omega(u) \, du + \int_{1/2}^1 y_{\min}\omega(u) \, du \\ &= \underbrace{\frac{1}{4}\mathbb{E}[Y \mid W = 1]}_{\text{Identifiable Mean}} + \underbrace{\frac{1}{8}\text{CVaR}_{0.5}(Y \mid W = 1)}_{\text{Optimal Transport (CVaR)}} + \underbrace{\frac{1}{8}y_{\min}}_{\text{Trivial Lower Bound}} \end{aligned} \quad (\text{O2.1})$$

where $\text{CVaR}_\alpha(Y) = \frac{1}{\alpha} \int_0^\alpha Q_Y(s) \, ds$ is the lower-tail Conditional Value at Risk. The integral allocates exactly probability α to the tail, including the appropriate fraction of any atom at its cutoff.

This decomposition also makes precise what is lost under moment relaxation. Moment relaxation enforces only the first moment over the entire observed interval: $\mathbb{E}[m_1(U) \mid U \in (0, 1/2)] = \mathbb{E}[Y \mid W = 1]$.

To minimize the objective, their linear program evaluates:

$$\min_{m_1} \left(\frac{1}{4}\mathbb{E}[Y \mid W = 1] + \frac{1}{2} \int_0^{1/4} m_1(u) \, du \right) + \frac{1}{8}y_{\min}$$

subject to the moment constraint and the box restriction $y_{\min} \leq m_1(u) \leq y_{\max}$. Writing $E_1 = \mathbb{E}[Y \mid W = 1]$, these restrictions imply

$$\min_{m_1} \int_0^{1/4} m_1(u) \, du = \max \left\{ \frac{1}{4}y_{\min}, \frac{1}{2}E_1 - \frac{1}{4}y_{\max} \right\}.$$

Hence the moment-relaxation lower bound is

$$\frac{1}{4}E_1 + \frac{1}{2} \max \left\{ \frac{1}{4}y_{\min}, \frac{1}{2}E_1 - \frac{1}{4}y_{\max} \right\} + \frac{1}{8}y_{\min}.$$

This is weakly below the CCOT lower bound because the lower-half average obeys $\text{CVaR}_{0.5}(Y \mid W = 1) \geq \max\{y_{\min}, 2E_1 - y_{\max}\}$. The inequality can be strict, but need not be: if $y_{\min} = 0$, $y_{\max} = 1$, and $Y \mid W = 1 \equiv 1$, both methods give $3/8$. CCOT tightens the relaxation precisely when the observed distribution contains tail information not already implied by the mean and support restrictions.

We further remark that a decomposition similar to (O2.1) exists in general (see Section 4 for details), and Figure 1 illustrates the three identification regions that compose the bound.

O2.2 Comparison with Moment Relaxation Approach

Example O2.2 (Comparison with Moment Relaxation Approach). Suppose there are no covariates and the weight function is $\omega(u) = \mathbf{1}(u \in (1/4, 1/2))$. Our target parameter simplifies to:

$$\mathbb{E}[(Y(1) - Y(0))\mathbf{1}(U \in (1/4, 1/2))].$$

Consider a data-generating process with a binary instrument $Z \in \{0, 1\}$, propensity scores $p(0) = 1/4$ and $p(1) = 3/4$, and the observed outcome is constantly zero, $Y(w) = 0$ almost surely for $w \in \{0, 1\}$.

Under our CCOT formulation, this parameter is point-identified and its value is exactly 0. Because the conditional distribution of the observed outcome is a Dirac measure at zero, $\mathbb{P}(Y \in \cdot \mid W = 1, Z = z) = \delta_0(\cdot)$ for each z , the observational constraint forces $\pi_1(\cdot \mid u) = \delta_0$ for almost all $u \in (0, 3/4)$. Similarly, $\mathbb{P}(Y \in \cdot \mid W = 0, Z = z) = \delta_0(\cdot)$ forces $\pi_0(\cdot \mid u) = \delta_0$ for almost all $u \in (1/4, 1)$. Since ω is supported on $(1/4, 1/2) \subset (0, 3/4) \cap (1/4, 1)$, both $Y(1)$ and $Y(0)$ are point-identified on the support of ω , and the target is sharply 0.

In contrast, [Mogstad et al. \[2018\]](#) relax the distributional constraints into moment constraints on the MTR, defined as $m_w(u) = \mathbb{E}[Y(w) \mid U = u]$. Under their framework, the data restrict the MTR only through its conditional expectation:

$$\begin{aligned} \frac{1}{p(z)} \int_0^{p(z)} m_1(u) du &= \mathbb{E}[Y \mid W = 1, Z = z] = 0, \\ \frac{1}{1 - p(z)} \int_{p(z)}^1 m_0(u) du &= \mathbb{E}[Y \mid W = 0, Z = z] = 0. \end{aligned}$$

Evaluating at $z = 0$ and $z = 1$ gives the four constraints

$$\int_0^{1/4} m_1 = \int_0^{3/4} m_1 = 0, \quad \int_{1/4}^1 m_0 = \int_{3/4}^1 m_0 = 0,$$

or equivalently, after differencing,

$$\int_0^{1/4} m_1 = 0, \quad \int_{1/4}^{3/4} m_1 = 0, \quad \int_{3/4}^1 m_0 = 0, \quad \int_{1/4}^{3/4} m_0 = 0.$$

Their method bounds the target parameter by solving:

$$\begin{aligned} \max_{m_0, m_1 \in \mathcal{M}} / \min_{m_0, m_1 \in \mathcal{M}} \quad & \int_{1/4}^{1/2} (m_1(u) - m_0(u)) du \\ \text{s.t.} \quad & \int_0^{1/4} m_1 = \int_{1/4}^{3/4} m_1 = 0, \quad \int_{3/4}^1 m_0 = \int_{1/4}^{3/4} m_0 = 0. \end{aligned}$$




Figure 7: Bounds of the moment relaxation approach and the CCOT approach by different sample sizes. Each data point is an average of 20 runs. In the experiment $M = 1$, i.e., $y_{\max} = 1, y_{\min} = -1$.

If the function class $\mathcal{M}$ is sufficiently rich, this relaxation discards the point-mass nature of the outcome and yields unnecessarily wide bounds. The moment constraints only pin the integral of each MTR over $(1/4, 3/4)$ to zero, but leave the shape within this interval unrestricted. An adversarial optimizer exploits this by setting $m_1(u) = M$ on $(1/4, 1/2)$ and $m_1(u) = -M$ on $(1/2, 3/4)$, leaving $m_1 \equiv 0$ on $(0, 1/4)$; this satisfies every constraint above yet contributes $M/4$ to the objective. Applying the symmetric construction to m_0 yields a total upper bound of $M/2$ instead of 0. Across the family of known outcome bounds $[-M, M]$, the resulting moment bounds $[-M/2, M/2]$ grow with M , while the distributional sharp value remains 0. By strictly enforcing the full distributional constraint, the CCOT approach eliminates these mathematically feasible but structurally impossible counterfactuals. Figure 7 illustrates this gap numerically.

O3 Proofs of the Reduction Propositions

This section proves the three closed-form rewritings on which the estimation procedures of Section 4 are built. Each is an algebraic reduction of an expression already displayed in the main text, and none of the arguments below is used anywhere else.

O3.1 Proof of Proposition 4.3

The nuisance components $\gamma_{\text{full},k}$, $\gamma_{j,k}$, γ_K , $J_{\text{full},k}$, $J_{j,k}$, Δ_K and the relative threshold $\kappa_{j,k}$ are defined in Section 4 immediately preceding Proposition 4.3; we refer to those formulas throughout this proof.

Proof of Proposition 4.3. By Theorem 3.8 specialized to a finite instrument (every conditional LIV interval degenerate), it suffices to rewrite the conditional integral $\int_0^1 Q_{Y,k|X}(u) Q_{\omega,k|X}(1-u) du$ as an observable finite weighted sum. Fix X . On each baseline interval $(p_k(X), p_{k+1}(X))$ the weight function $\omega(X, u)$ is a non-increasing step function. On the sub-interval $(q_{j-1,k}(X), q_{j,k}(X))$, $j = 1, \dots, l_k$, exactly the current-bin instruments in $\bigcup_{l=j}^{l_k} T_{l,k}$ and all instruments whose alternative propensity lies in a strictly higher bin contribute to the first probability in ω , whereas baseline propensities in strictly higher bins contribute to the second. Thus, up to the irrelevant boundary values,

$$\begin{aligned} \omega(X, u) = \mathbb{P}_{\text{obs}} \left(Z \in \bigcup_{l=j}^{l_k} T_{l,k} \mid X \right) + \mathbb{P}_{\text{obs}} \left(Z \in \bigcup_{i=k+1}^K \bigcup_{l=0}^{l_i} T_{l,i} \mid X \right) \\ - \mathbb{P}_{\text{obs}} \left(Z \in \bigcup_{i=k+1}^K S_i \mid X \right), \quad u \in (q_{j-1,k}, q_{j,k}). \end{aligned}$$

Regrouping, the step-function value on the relevant slice is $\gamma_{\text{full},k}(X) + \gamma_{j,k}(X)$: the first term acts on the entire baseline interval and the second is the additional current-bin mass. Multiplying by $Q_{Y,k|X}$ and summing the nested current-bin contributions over the normalized sub-intervals yields

$$\begin{aligned} (p_{k+1}(X) - p_k(X)) \int_0^1 Q_{Y,k|X}(u) Q_{\omega,k|X}(1-u) du \\ = \gamma_{\text{full},k}(X) J_{\text{full},k}(X) + \sum_{j=1}^{l_k} \gamma_{j,k}(X) J_{j-1,k}(X), \end{aligned}$$

where we used

$$J_{\text{full},k}(X) = (p_{k+1}(X) - p_k(X)) \mathbb{E}_{\mu_{1,k|X}}[Y] = (p_{k+1}(X) - p_k(X)) \int_0^1 Q_{Y,k|X}(u) du$$

and the definition of $J_{j-1,k}$.

The trivial-bound term on the tail $(p_K(X), 1]$ simplifies for the PRTE weight: since $\omega(X, u) = \mathbb{P}_{\text{obs}}(u < q(Z, X) \mid X) \geq 0$ for $u > p_K(X)$, we have $\int_{p_K(X)}^1 (y_{\min} \max\{0, \omega\} +$

$y_{\max} \min\{0, \omega\} du = y_{\min} \int_{p_K(X)}^1 \omega(X, u) du$, which equals $\Delta_K(X)$ by Fubini. Taking the outer expectation over X and summing over k yields (4.1).

The corresponding representation for the upper bound is obtained by the symmetric argument with $Q_{\omega, k|X}(1-t)$ replaced by $Q_{\omega, k|X}(t)$ in the OT term and $y_{\min}, y_{\max}$ swapped in the tail. $\square$

O3.2 Proof of Proposition 4.4

Proof of part (i). Fix k and $j \in \{0, \dots, l_k - 1\}$. By the change of variables $u = F_{\mu_{1,k|X}}(y)$, continuity of $\mu_{1,k|X}$ at $\nu_{j,k}(X)$ and $\nu_{j+1,k}(X)$ gives $F_{\mu_{1,k|X}}(\nu_{j,k}(X)) = \kappa_{j,k}(X)$ and $F_{\mu_{1,k|X}}(\nu_{j+1,k}(X)) = \kappa_{j+1,k}(X)$ almost surely. Therefore,

$$\begin{aligned} \int_{\kappa_{j,k}(X)}^{\kappa_{j+1,k}(X)} Q_{Y,k|X}(u) du &= \int_{\nu_{j,k}(X)}^{\nu_{j+1,k}(X)} y dF_{\mu_{1,k|X}}(y) \\ &= \mathbb{E}_{\mu_{1,k|X}}[Y \mathbb{1}(\nu_{j,k}(X) < Y \leq \nu_{j+1,k}(X))], \end{aligned}$$

where the second equality uses the continuity of $\mu_{1,k|X}$ at the two boundary points to include/exclude boundary atoms without ambiguity. Multiplying by $p_{k+1}(X) - p_k(X)$ yields (4.2). $\square$

Proof of part (ii). Suppose $Y \in \{0, 1\}$. Let $\theta_k(X) := \mathbb{E}_{\mu_{1,k|X}}[Y]$ denote the complier success probability on the k -th interval. By the definition of the identified complier measure [cf. Angrist and Imbens, 1995], $(p_{k+1}(X) - p_k(X))\theta_k(X) = \mathbb{E}_{\text{obs}}[YW \mid Z \in S_{k+1}, X] - \mathbb{E}_{\text{obs}}[YW \mid Z \in S_k, X] = P_{1,k+1}(X) - P_{1,k}(X)$, so $J_{\text{full},k}(X) = P_{1,k+1}(X) - P_{1,k}(X)$.

Since $Y \in \{0, 1\}$, the conditional quantile function is the step function $Q_{Y,k|X}(u) = \mathbb{1}(u \geq 1 - \theta_k(X))$. Equivalently, $Q_{Y,k|X}(u) = 1$ iff $u(p_{k+1}(X) - p_k(X)) + p_k(X) \geq p_{k+1}(X) - J_{\text{full},k}(X)$, i.e., the corresponding propensity value $v = p_k(X) + u(p_{k+1}(X) - p_k(X))$ lies in the interval $[p_{k+1}(X) - J_{\text{full},k}(X), p_{k+1}(X)]$. Writing $J_{j-1,k}(X)$ as a length in the original propensity scale, it equals the length of the intersection of $[q_{j-1,k}(X), q_{j,k}(X)]$ with $[p_{k+1}(X) - J_{\text{full},k}(X), p_{k+1}(X)]$:

$$J_{j-1,k}(X) = \max\left(0, q_{j,k}(X) - \max(q_{j-1,k}(X), p_{k+1}(X) - J_{\text{full},k}(X))\right) = h_{j,k}^-(X).$$

Substituting into (4.1) yields (4.4). $\square$

O3.3 Proof of Proposition 4.8

Proof of Proposition 4.8. Because $\mathcal{Z}$ is compact and connected and $z \mapsto p(z, X)$ is continuous by ?? 3.6, its image is a closed interval for almost every X , so Theorem 3.8 specializes with

$K_X = 1$ and conditional LIV interval $[\underline{p}(X), \bar{p}(X)]$. We derive the three terms of (4.7) in turn. Substituting the explicit form of the weight function,

$$\omega(x, u) = \mathbb{E}_{Z|X}[\mathbb{1}(u \leq q(Z, X)) - \mathbb{1}(u \leq p(Z, X)) \mid X = x],$$

and applying Fubini's theorem to swap the order of integration, we handle each term separately.

First term (quantile integral on the gap $(0, \underline{p}(X))$). For $\underline{p}(X) > 0$, substitute the weight function, swap expectations, and use the change of variables $v = \underline{p}(X)u$. Because $P \geq \underline{p}(X)$ conditional on X , the baseline indicator equals one throughout this gap. Thus

$$\begin{aligned} & \mathbb{E}_X \left[\underline{p}(X) \int_0^1 Q_{1, \underline{p}|X}(u) \omega(X, \underline{p}(X)u) du \right] \\ &= \mathbb{E}_{Z, X} \left[\underline{p}(X) \int_0^1 Q_{1, \underline{p}|X}(u) (\mathbb{1}(\underline{p}(X)u \leq q(Z, X)) - 1) du \right] \\ &= -\mathbb{E}_{Z, X} \left[\int_{\min\{\underline{p}(X), q(Z, X)\}}^{\underline{p}(X)} Q_{1, \underline{p}|X}(v/\underline{p}(X)) dv \right]. \end{aligned}$$

Both sides are zero when $\underline{p}(X) = 0$, which justifies the convention in the proposition.

Second term (g_1 difference on the LIV interval). The structural integral representation gives $g_1(u, x) = \int_0^u \mathbb{E}[Y(1) \mid U = s, X = x] ds$, so $g_1(\cdot, x)$ is absolutely continuous. Hence, for any $r \in [0, 1]$,

$$\int_{\underline{p}(X)}^{\bar{p}(X)} \partial_u g_1(u, X) \mathbb{1}(u \leq r) du = g_1(\min\{\bar{p}(X), \max\{\underline{p}(X), r\}\}, X) - g_1(\underline{p}(X), X).$$

Applying this identity first with $r = q(Z, X)$ and then with $r = P$ is essential when $q(Z, X) < \underline{p}(X)$. Since $P \in [\underline{p}(X), \bar{p}(X)]$, the two lower-endpoint terms cancel and give

$$\mathbb{E}_X \left[\int_{\underline{p}(X)}^{\bar{p}(X)} \frac{\partial g_1(u, X)}{\partial u} \omega(X, u) du \right] = \mathbb{E}_{Z, X} [g_1(q^{\text{tr}}(Z, X), X) - g_1(P, X)].$$

Third term (trivial bound on $(\bar{p}(X), 1]$). Since $P \leq \bar{p}(X)$, the baseline indicator vanishes on this tail and $\omega(X, u) = \mathbb{P}(u \leq q(Z, X) \mid X) \geq 0$. Therefore

$$y_{\min} \mathbb{E}_X \left[\int_{\bar{p}(X)}^1 \omega(X, u) du \right] = y_{\min} \mathbb{E}_{Z, X} [(q(Z, X) - \bar{p}(X))_+].$$

Aggregating the three components yields (4.7). The upper-bound representation replaces the countermonotonic boundary coupling by the comonotonic coupling and swaps $y_{\min}$ for $y_{\max}$ in the trivial tail. $\square$

O4 Moment Condition and Neyman Orthogonality of the Discrete-Instrument Scores

This section proves Lemma D.1. Both parts are term-by-term calculations that repeatedly apply the tower property to inverse-probability-weighted residuals with zero conditional mean.

Throughout, η_0 denotes the vector of true nuisance parameter values and $\theta_0 = \underline{\theta}_{\omega,1}$ the true parameter value. For a functional $\Phi(\eta)$ and a perturbation direction δ in the same function space as η , we write $\partial_\delta \Phi(\eta) := \frac{d}{dt} \Phi(\eta + t\delta)|_{t=0}$ for the Gateaux (directional) derivative; when the perturbation is along a single nuisance component η_α , we write ∂_{η_α} for brevity.

O4.1 Continuous Outcome

Recall the target functional:

$$\theta_0 = \mathbb{E} \left[\sum_{k=0}^{K-1} \left(\gamma_{\text{full},k}(X) J_{\text{full},k}(X) + \sum_{j=1}^{l_k} \gamma_{j,k}(X) J_{j-1,k}(X) \right) + \Delta_K(X) \right], \quad (\text{O4.1})$$

and the total Neyman-orthogonal score assembled from the components of Section C.1:

$$\psi(O; \eta) = \sum_{k=0}^{K-1} \left(\psi_{\text{full},k}^{\text{prod}}(O; \eta) + \sum_{j=1}^{l_k} \psi_{j,k}^{\text{prod}}(O; \eta) \right) + \psi_{\Delta_K}(O; \eta), \quad (\text{O4.2})$$

where each product score debiases the corresponding term in the target. Specifically, $\psi_{j,k}^J(O; \eta)$ is the uncentered orthogonal score for the conditional integral component $\mathbb{E}[J_{j,k}(X)]$, and the product score for $\mathbb{E}[\gamma_{j,k}(X) J_{j-1,k}(X)]$ is:

$$\psi_{j,k}^{\text{prod}}(O; \eta) = \gamma_{j,k}(X) \psi_{j-1,k}^J(O; \eta) + J_{j-1,k}(X) \left(\mathbb{1}(Z \in \bigcup_{l=j}^{l_k} T_{l,k}) - \gamma_{j,k}(X) \right), \quad (\text{O4.3})$$

and symmetrically for $\psi_{\text{full},k}^{\text{prod}}(O; \eta)$.

Proof of Lemma D.1(i). Moment condition. We show $\mathbb{E}[\psi(O; \eta_0)] = \theta_0$. All influence function correction terms in the score are inverse-probability-weighted residuals of the general form

$$\frac{\mathbb{1}(Z \in S_k)}{\pi_k(X)} W(g(Y, X) - \mathbb{E}[g(Y, X) \mid X, Z \in S_k, W=1]).$$

By the tower property, conditioning first on (X, Z) : when $Z \in S_k$, we have $p(Z, X) = p_k(X)$, so $\mathbb{E}[W \mid X, Z \in S_k] = p_k(X)$. It follows that

$$\mathbb{E} \left[\frac{\mathbb{1}(Z \in S_k)}{\pi_k(X)} W(g - \mathbb{E}[g \mid X, S_k, W=1]) \right] = \mathbb{E} \left[\frac{\mathbb{1}(Z \in S_k)}{\pi_k(X)} p_k(X) \cdot 0 \right] = 0.$$

Applying this identity to each correction term, $\mathbb{E}[IF_{J^\pm}(O)] = 0$ and $\mathbb{E}[IF_{M_{j,k}^\pm}(O)] = 0$ at η_0 . The quantile correction term in $\psi_{j,k}^J$ contains the factor $p_{k+1}(X)M_{j,k}^+(X) - p_k(X)M_{j,k}^-(X) - (q_{j,k}(X) - p_k(X))$, which vanishes at η_0 by the conditional moment equation

$$\begin{aligned} p_{k+1}(X) \mathbb{E}[\mathbb{1}(Y \leq \nu_{j,k}(X)) \mid X, S_{k+1}, W=1] \\ - p_k(X) \mathbb{E}[\mathbb{1}(Y \leq \nu_{j,k}(X)) \mid X, S_k, W=1] - (q_{j,k}(X) - p_k(X)) = 0 \end{aligned} \quad (\text{O4.4})$$

which defines $\nu_{j,k}$ for $j \geq 1$; the boundary value $\nu_{0,k} = y_{\min}$ is fixed. The propensity score correction terms $\frac{\mathbb{1}(Z \in S_k)}{\pi_k(X)}(W - p_k(X))$ and the pooled policy-level residuals $R_{q,j,k}$ likewise have zero mean by the same tower-property argument. Consequently,

$$\mathbb{E}[\psi_{j,k}^J(O; \eta_0)] = \mathbb{E}[J_{j,k}(X)], \quad \mathbb{E}[\psi_{\text{full},k}^J(O; \eta_0)] = \mathbb{E}[J_{\text{full},k}(X)].$$

For the product scores, substituting (O4.3) and using $\mathbb{E}[\mathbb{1}(Z \in \bigcup_{l=j}^k T_{l,k}) - \gamma_{j,k}(X)] = 0$:

$$\begin{aligned} \mathbb{E}[\psi_{j,k}^{\text{prod}}(O; \eta_0)] &= \mathbb{E}[\gamma_{j,k}(X) \underbrace{\mathbb{E}[\psi_{j-1,k}^J(O; \eta_0) \mid X]}_{=J_{j-1,k}(X)}] + \mathbb{E}[J_{j-1,k}(X) \underbrace{\mathbb{E}[\mathbb{1}(\cdots) - \gamma_{j,k}(X) \mid X]}_{=0}] \\ &= \mathbb{E}[\gamma_{j,k}(X) J_{j-1,k}(X)], \end{aligned}$$

and symmetrically $\mathbb{E}[\psi_{\text{full},k}^{\text{prod}}(O; \eta_0)] = \mathbb{E}[\gamma_{\text{full},k}(X) J_{\text{full},k}(X)]$. The same reasoning yields $\mathbb{E}[\psi_{\Delta_K}(O; \eta_0)] = \mathbb{E}[\Delta_K(X)]$. Summing over all components and comparing with (O4.1) gives $\mathbb{E}[\psi(O; \eta_0)] = \theta_0$.

Neyman orthogonality of $\psi_{j,k}^J$. We show that the Gateaux derivative $\partial_t \mathbb{E}[\psi_{j,k}^J(O; \eta_0 + t\delta)]|_{t=0}$ vanishes for all perturbation directions δ . We group the nuisance parameters by type and check orthogonality with respect to each group.

Conditional expectations $(J_{j,k}^\pm, M_{j,k}^\pm)$. Consider a perturbation $J_{j,k}^+(X) \rightarrow J_{j,k}^+(X) + t\delta J^+(X)$. The plug-in term $J_{j,k} = p_{k+1} J_{j,k}^+ - p_k J_{j,k}^-$ contributes $+\mathbb{E}[p_{k+1}(X)\delta J^+(X)]$, while the IPW correction IF_{J^+} contributes $-\mathbb{E}\left[\frac{\mathbb{1}(Z \in S_{k+1})}{\pi_{k+1}(X)} p_{k+1}(X)\delta J^+(X)\right]$. These cancel because

$$\mathbb{E}\left[\frac{\mathbb{1}(Z \in S_{k+1})}{\pi_{k+1}(X)} \mid X\right] = \frac{\mathbb{P}(Z \in S_{k+1} \mid X)}{\pi_{k+1}(X)} = 1.$$

For $J_{j,k}^-$ the plug-in derivative is $-p_k\delta J^-$, while differentiating $-IF_{J^-}$ gives $+p_k\delta J^-$. For $M_{j,k}^+$, the plug-in part of the quantile moment gives $\nu_{j,k}p_{k+1}\delta M^+$ and differentiating $\nu_{j,k}IF_{M_{j,k}^+}$ gives its negative; the minus-side and $j+1$ terms cancel in the same way, with their displayed signs. Thus every conditional-expectation direction cancels explicitly.

Instrument assignment probabilities (π_k, π_{k+1}) . Perturbing $\pi_{k+1}(X) \rightarrow \pi_{k+1}(X) + t \delta\pi(X)$ introduces a factor $-\delta\pi(X)/\pi_{k+1}(X)^2$ in the IPW terms. For any such IPW term $\mathbb{1}(Z \in S_{k+1}) \pi_{k+1}(X)^{-1} W(g(O) - \bar{g}(X))$, the first-order contribution is

$$-\mathbb{E} \left[\frac{\mathbb{1}(Z \in S_{k+1})}{\pi_{k+1}(X)^2} \delta\pi(X) W(g(O) - \bar{g}(X)) \right].$$

Conditioning on (X, Z) with $Z \in S_{k+1}$, we have $\mathbb{E}[W(g(O) - \bar{g}(X)) \mid X, Z \in S_{k+1}] = p_{k+1}(X) \cdot 0 = 0$ at η_0 , so the Gateaux derivative vanishes. The argument for $\delta\pi_k$ is identical when $k \geq 1$; there is no π_0 direction.

Propensity scores (p_k, p_{k+1}) . For $k \geq 1$, a perturbation $p_k(X) \rightarrow p_k(X) + t \delta p_k(X)$ generates first-order contributions from three sources (all multiplied by $\delta p_k(X)$ inside $\mathbb{E}[\cdot]$). There is no such direction at $k = 0$ because $p_0 \equiv 0$ is fixed:

- (i) The plug-in term $J_{j,k} = p_{k+1} J_{j,k}^+ - p_k J_{j,k}^-$, contributing $-J_{j,k}^-(X)$;
- (ii) The quantile correction terms $\nu_{j,k}(\cdots) - \nu_{j+1,k}(\cdots)$, contributing $\nu_{j,k}(X)(\mathbb{1}(j \geq 1) - M_{j,k}^-(X)) - \nu_{j+1,k}(X)(1 - M_{j+1,k}^-(X))$. At $j = 0$, the first constant is absent because $q_{0,k} - p_k \equiv 0$;
- (iii) The Riesz representer correction $\alpha_{p_k}(X) \frac{\mathbb{1}(Z \in S_k)}{\pi_k(X)} (W - p_k(X))$, contributing $-\alpha_{p_k}(X)$.

Substituting the definition

$$\alpha_{p_k} = -J_{j,k}^- - \nu_{j,k} M_{j,k}^- + \nu_{j+1,k} M_{j+1,k}^- + \mathbb{1}(j \geq 1) \nu_{j,k} - \nu_{j+1,k}$$

and collecting terms:

$$\begin{aligned} & (-J_{j,k}^-) + (\mathbb{1}(j \geq 1) \nu_{j,k} - \nu_{j,k} M_{j,k}^- - \nu_{j+1,k} + \nu_{j+1,k} M_{j+1,k}^-) \\ & - (-J_{j,k}^- - \nu_{j,k} M_{j,k}^- + \nu_{j+1,k} M_{j+1,k}^- + \mathbb{1}(j \geq 1) \nu_{j,k} - \nu_{j+1,k}) = 0. \end{aligned}$$

The calculation for δp_{k+1} , using $\alpha_{p_{k+1}} = J_{j,k}^+ + \nu_{j,k} M_{j,k}^+ - \nu_{j+1,k} M_{j+1,k}^+$, is identical.

Quantile thresholds $(\nu_{j,k}, \nu_{j+1,k})$. A perturbation is considered only for estimated thresholds. In particular, $\nu_{0,k} = y_{\min}$ is fixed and has no nuisance direction. A perturbation $\nu_{j,k}(X) \rightarrow \nu_{j,k}(X) + t \delta\nu(X)$ affects the score through two channels.

Channel (a): indicator boundaries. The indicator $\mathbb{1}(\nu_{j,k}(X) \leq Y)$ appears inside $IF_{J\pm}$ and $IF_{M_{j,k}^\pm}$. For a generic relevant conditional law F and $a_t = \nu + t \delta\nu$, the combined boundary change from the truncated first-moment correction and its quantile-moment correction is, up to its orientation,

$$\int_{\nu}^{a_t} (a_t - y) dF(y).$$

If the density of F is bounded by C in the stated neighborhood, then

$$\left| \int_{\nu}^{a_t} (a_t - y) dF(y) \right| \leq C |a_t - \nu|^2.$$

The oriented-integral convention covers $a_t < \nu$. Applying this bound to the plus and minus conditional laws shows that the combined boundary contribution is $O(t^2 \delta \nu^2)$ and therefore has derivative zero at $t = 0$.

Channel (b): explicit multiplicative factor. The term $\nu_{j,k}(X)(p_{k+1}M_{j,k}^+ - p_k M_{j,k}^- - (q_{j,k} - p_k))$ contributes $\delta \nu(X)$ times the conditional moment equation (O4.4), which vanishes at η_0 . The dependence of α_{p_k} and $\alpha_{p_{k+1}}$ on $\nu_{j,k}$ is absorbed by the zero-mean residuals: for any L_2 function $h(X)$,

$$\mathbb{E} \left[h(X) \frac{\mathbb{1}(Z \in S_k)}{\pi_k(X)} (W - p_k(X)) \right] = \mathbb{E} \left[h(X) \frac{\pi_k(X)}{\pi_k(X)} \cdot 0 \right] = 0.$$

The local density bound gives the quadratic control above. The treatment of $\delta \nu_{j+1,k}$ is symmetric.

Alternative policy propensity ($q_{j,k}$). For $j \geq 1$, use the pooled representation $q_{j,k}(X) = \mathbb{E}[q(Z, X) \mid Z \in T_{j,k}, X]$. At the population law, all the instrument-specific policy propensities in $T_{j,k}$ equal this common value. Consequently, perturbations of the conditional instrument probabilities have no first-order effect on the pooled value, while perturbing the propensity functions induces

$$\delta q_{j,k}(X) = \frac{1}{\tau_{j,k}(X)} \sum_{z \in T_{j,k}} \pi(z, X) \frac{\partial \phi}{\partial p}(z, X, p(z, X)) \delta p(z, X).$$

At $j = 0$, $T_{0,k} = \emptyset$ and the boundary $q_{0,k} = p_k$ is handled by the corrected p_k representer above. For $j \geq 1$, the first-order contributions are:

$$\begin{aligned} \text{from quantile correction :} & \quad -\mathbb{E}[\nu_{j,k}(X) \delta q_{j,k}(X)], \\ \text{from differentiating the policy residual :} & \quad +\mathbb{E}[\nu_{j,k}(X) \delta q_{j,k}(X)]. \end{aligned}$$

The pooled correction itself has mean zero at the truth. Differentiating its residuals gives

$$\partial_t \mathbb{E}[-\nu_{j,k}(X) R_{q,j,k}(O)]|_{t=0} = \mathbb{E}[\nu_{j,k}(X) \delta q_{j,k}(X)],$$

because $\mathbb{E}[\mathbb{1}(Z = z) \mid X] = \pi(z, X)$. Derivatives of the coefficients multiply zero-mean residuals and therefore vanish. The two displayed contributions cancel, and the treatment of $q_{j+1,k}$ is symmetric.

The simplex restriction is essential for the preceding pooled representation. Along a within-set share perturbation $a_{j,k}(z, X) \mapsto a_{j,k}(z, X) + t \delta a_z(X)$, admissible directions satisfy $\sum_{z \in T_{j,k}} \delta a_z(X) = 0$. Since $q(z, X) = q_{j,k}(X)$ for every $z \in T_{j,k}$ at the population law,

$$\partial_t \sum_{z \in T_{j,k}} (a_{j,k}(z, X) + t \delta a_z(X)) q(z, X) \Big|_{t=0} = q_{j,k}(X) \sum_{z \in T_{j,k}} \delta a_z(X) = 0.$$

With estimated instrument-specific policy values, the remaining share term is $\sum_z (\hat{a}_z - a_z)(\hat{q}(z, X) - q_{j,k}(X))$, a product of two nuisance errors. A perturbation of $\tau_{j,k}$ affects only the factor $1/\tau_{j,k}$ in $R_{q,j,k}$; its derivative multiplies the conditionally mean-zero numerator $\sum_{z \in T_{j,k}} \phi_p(z, X, p(z, X)) \mathbb{1}(Z = z)(W - p(z, X))$ and is therefore zero at the truth.

This establishes orthogonality for $\psi_{j,k}^J$.

Neyman orthogonality of the remaining components. For $\psi_{\text{full},k}^J$, a perturbation of $J_{\text{full},k}^+$ contributes $p_{k+1} \delta J^+$ through the plug-in and $-p_{k+1} \delta J^+$ through its outcome residual. A perturbation of p_{k+1} contributes $J_{\text{full},k}^+ \delta p_{k+1}$ through the plug-in and its negative through the propensity residual. The minus-side terms cancel identically when $k \geq 1$, and denominator perturbations in π_k, π_{k+1} multiply conditionally mean-zero residuals. Hence $\psi_{\text{full},k}^J$ is orthogonal in every nuisance direction.

For the product scores, we invoke the chain rule for influence functions. The preceding cancellations hold conditionally on X : at η_0 , $\mathbb{E}[\psi_{j-1,k}^J(O; \eta_0) \mid X] = J_{j-1,k}(X)$, and the Gateaux derivative of this conditional moment is zero in every nuisance direction already considered. Moreover, $\mathbb{1}(Z \in \bigcup_{l=j}^{l_k} T_{l,k}) - \gamma_{j,k}(X)$ is the conditional influence function for $\gamma_{j,k}(X) = \mathbb{P}(Z \in \bigcup_{l=j}^{l_k} T_{l,k} \mid X)$. Therefore, the product score (O4.3)

$$\psi_{j,k}^{\text{prod}} = \gamma_{j,k}(X) \psi_{j-1,k}^J(O; \eta) + J_{j-1,k}(X) (\mathbb{1}(Z \in \bigcup_{l=j}^{l_k} T_{l,k}) - \gamma_{j,k}(X))$$

is orthogonal for the product functional $\mathbb{E}[\gamma_{j,k}(X) J_{j-1,k}(X)]$. We treat the two groups of nuisance directions separately. First, let η_α be any component other than $\gamma_{j,k}$. Then $\gamma_{j,k}$ and the indicator residual do not vary, so

$$\begin{aligned} \partial_{\eta_\alpha} \mathbb{E}[\psi_{j,k}^{\text{prod}}] &= \mathbb{E} \left[\gamma_{j,k}(X) \underbrace{\partial_{\eta_\alpha} \mathbb{E}[\psi_{j-1,k}^J \mid X]}_{=0} \right] \\ &\quad + \mathbb{E} \left[\partial_{\eta_\alpha} J_{j-1,k}(X) \underbrace{\mathbb{E}[\mathbb{1}(\cdot \cdot \cdot) - \gamma_{j,k}(X) \mid X]}_{=0} \right] = 0, \end{aligned}$$

where the first expectation vanishes by the conditional orthogonality just established. Second, for the direction $\gamma_{j,k} \mapsto \gamma_{j,k} + t \delta \gamma$, the derivative contributions of the explicit prefactor and

the residual cancel:

$$\partial_t \mathbb{E}[(\gamma_{j,k} + t \delta \gamma) \psi_{j-1,k}^J + J_{j-1,k}(\mathbb{1}(\cdots) - \gamma_{j,k} - t \delta \gamma)] \Big|_{t=0} = \mathbb{E}[\delta \gamma J_{j-1,k}] - \mathbb{E}[\delta \gamma J_{j-1,k}] = 0,$$

using $\mathbb{E}[\psi_{j-1,k}^J | X] = J_{j-1,k}(X)$ at η_0 . For the full-interval product, the residual is

$$\mathbb{1}\left(Z \in \bigcup_{m=k+1}^K \bigcup_{l=0}^{l_m} T_{l,m}\right) - \mathbb{1}\left(Z \in \bigcup_{m=k+1}^K S_m\right) - \gamma_{\text{full},k}(X),$$

whose conditional mean is zero by the definition of $\gamma_{\text{full},k}$. Therefore its coefficient and residual derivatives cancel exactly as in (O4.3), even though $\gamma_{\text{full},k}$ is signed.

For ψ_{Δ_K} , write

$$A_K = \left\{ Z \in \bigcup_{j=1}^{l_K} T_{j,K} \right\}.$$

A perturbation of p_K contributes $-y_{\min} \mathbb{E}[\gamma_K(X) \delta p_K(X)]$ through the plug-in term and the opposite quantity through differentiation of the p_K residual. For an instrument-specific direction $p(z, X) \mapsto p(z, X) + t \delta p(z, X)$, the plug-in derivative is

$$y_{\min} \mathbb{E}[\mathbb{1}(A_K) \phi_p(Z, X, p(Z, X)) \delta p(Z, X)],$$

and differentiation of the policy residual gives its negative. Perturbations of π_K and of the coefficient γ_K multiply the conditionally mean-zero p_K residual, while derivatives of ϕ_p multiply $W - p(Z, X)$ and also have conditional mean zero. Thus every first-order direction in ψ_{Δ_K} cancels. □

O4.2 Binary Outcome

Proof of Lemma D.1(ii). Recall the target functional (4.4) and the orthogonal score (B.1), which decomposes as $\psi = \sum_k (\psi_{\text{full},k}^{\text{prod}} + \sum_j \psi_{j,k}^{\text{prod}}) + \psi_{\Delta_K}$. **Moment condition.** We show $\mathbb{E}[\psi(O; \eta_0)] = \theta_0$. The residuals $R_{P,k}$ and $R_{p,k}$ have conditional mean zero because $P_{1,k}(X) = \mathbb{E}[YW | Z \in S_k, X]$ and $p_k(X) = \mathbb{E}[W | Z \in S_k, X]$. The pooled policy residual also satisfies

$$\begin{aligned} & \mathbb{E}[R_{q,j,k}(O) | X] \\ &= \frac{1}{\tau_{j,k}(X)} \sum_{z \in T_{j,k}} \phi_p(z, X, p(z, X)) \mathbb{P}(Z = z | X) \mathbb{E}[W - p(z, X) | Z = z, X] = 0. \end{aligned}$$

Thus every correction in $\psi_{j,k}^h$ and $\psi_{\text{full},k}^J$ has conditional mean zero, so $\mathbb{E}[\psi_{j,k}^h(O; \eta_0) \mid X] = h_{j,k}^-(X)$ and $\mathbb{E}[\psi_{\text{full},k}^J(O; \eta_0) \mid X] = J_{\text{full},k}(X)$. Writing $A_{j,k} = \{Z \in \bigcup_{l=j}^{l_k} T_{l,k}\}$, another application of the tower property gives

$$\begin{aligned}\mathbb{E}[\psi_{j,k}^{\text{prod}}(O; \eta_0)] &= \mathbb{E}[\gamma_{j,k}(X) \mathbb{E}[\psi_{j,k}^h(O; \eta_0) \mid X] + h_{j,k}^-(X) \mathbb{E}[\mathbb{1}(A_{j,k}) - \gamma_{j,k}(X) \mid X]] \\ &= \mathbb{E}[\gamma_{j,k}(X) h_{j,k}^-(X)].\end{aligned}$$

The same conditional calculation applies to the full-interval product score. Summing all components therefore gives $\mathbb{E}[\psi(O; \eta_0)] = \mathbb{E}[m(X; \eta_0)] = \theta_0$.

Neyman orthogonality. We must show $\partial_{\eta_\alpha} \mathbb{E}[\psi(O; \eta)]|_{\eta=\eta_0} = 0$ for each nuisance component η_α . By the product score structure, it suffices to establish orthogonality of the component scores $\psi_{j,k}^h$ and $\psi_{\text{full},k}^J$ with respect to $(P_{1,m}, p_m, p(z, X), \pi_m, \tau_{j,k}, a_{j,k})$, and then show that the product scores inherit orthogonality with respect to the γ terms.

Orthogonality of $\psi_{j,k}^h$ with respect to $P_{1,m}$. Consider a perturbation $P_{1,m} \rightarrow P_{1,m} + t \delta P_{1,m}(X)$. Since $h_{j,k}^-$ depends on $P_{1,m}$ only through $J_{\text{full},k} = P_{1,k+1} - P_{1,k}$, the pathwise derivative is

$$\partial_t \mathbb{E}[h_{j,k}^-] \big|_{t=0} = \mathbb{E}[\alpha_{P,m}^{(j,k)}(X) \delta P_{1,m}(X)],$$

where $\alpha_{P,k+1}^{(j,k)} = D_{j,k}^-$, $\alpha_{P,k}^{(j,k)} = -D_{j,k}^-$, and $\alpha_{P,m}^{(j,k)} = 0$ for $m \notin \{k, k+1\}$. Under ?? 4.6, $D_{j,k}^- = \mathbb{1}(q_{j-1,k} < p_{k+1} - J_{\text{full},k} < q_{j,k})$ is well-defined. The IPW correction $D_{j,k}^-(R_{P,k+1} - R_{P,k})$ in $\psi_{j,k}^h$ contributes

$$\begin{aligned}\partial_t \mathbb{E} \left[D_{j,k}^- \left(\frac{\mathbb{1}(Z \in S_{k+1})}{\pi_{k+1}} (YW - P_{1,k+1}) - \frac{\mathbb{1}(Z \in S_k)}{\pi_k} (YW - P_{1,k}) \right) \right] \bigg|_{t=0} \\ = -\mathbb{E} \left[\alpha_{P,m}^{(j,k)} \delta P_{1,m} \right],\end{aligned}$$

using $\mathbb{E}[\frac{\mathbb{1}(Z \in S_m)}{\pi_m} \mid X] = 1$, so the two contributions cancel. The same argument gives $\partial_{P_{1,m}} \mathbb{E}[\psi_{\text{full},k}^J] = 0$ (with $\alpha_{P,k+1}^{(\text{full},k)} = 1$ and $\alpha_{P,k}^{(\text{full},k)} = -1$).

Remark (notation). For $k \in \{0, \dots, K-1\}$ and $j \in \{1, \dots, l_k\}$, $\alpha_{P,k}^{(j,k)}(x) = -D_{j,k}^-(x)$ is the coefficient of the $P_{1,k}$ perturbation in the derivative of $h_{j,k}^-$, with the other nuisance arguments held fixed.

Orthogonality of $\psi_{j,k}^h$ with respect to p_m and $p(z, X)$. The propensity scores enter $h_{j,k}^-$ through three channels:

- (a) the integration boundary $p_{k+1} - J_{\text{full},k}$, corrected by $\alpha_{p,k+1}^{(j,k)} R_{p,k+1}$;
- (b) the lower level-set boundary $q_{0,k} = p_k$ when $j = 1$, corrected by $\alpha_{p,k}^{(1,k)} R_{p,k}$;

- (c) the alternative policy values $q_{j,k}$ and $q_{j-1,k}$ through the mapping $\phi(z, X, p(z, X))$, corrected by the policy terms involving $T_{j,k}$ and $T_{j-1,k}$.

For channel (a), a perturbation $p_{k+1} \rightarrow p_{k+1} + t \delta p_{k+1}(X)$ shifts the boundary by δp_{k+1} , giving $\partial_t h_{j,k}^- = -D_{j,k}^- \delta p_{k+1}$. The correction $\alpha_{p,k+1}^{(j,k)} R_{p,k+1} = -D_{j,k}^- \cdot \frac{\mathbb{1}(Z \in S_{k+1})}{\pi_{k+1}} (W - p_{k+1})$ contributes $\partial_t \mathbb{E}[\alpha_{p,k+1}^{(j,k)} R_{p,k+1}] = D_{j,k}^- \delta p_{k+1}$, canceling the plug-in derivative.

For channel (b), when $j = 1$ and $E_{1,k}^- = 1$ (i.e., $p_k > p_{k+1} - J_{\text{full},k}$), a perturbation of p_k shifts $q_{0,k} = p_k$, giving $\partial_t h_{1,k}^- = -E_{1,k}^- I_{1,k}^- \delta p_k$. The correction $\alpha_{p,k}^{(1,k)} R_{p,k}$ contributes the matching cancellation.

For channel (c), the pooled derivative displayed in the continuous-outcome proof gives $\partial_t h_{j,k}^- = I_{j,k}^- \delta q_{j,k}$ for the upper boundary. The derivative of $I_{j,k}^- R_{q,j,k}$ is $-I_{j,k}^- \delta q_{j,k}$ and therefore cancels it. Symmetrically, when $E_{j,k}^- = 1$, the correction through $R_{q,j-1,k}$ cancels the derivative through the lower boundary $q_{j-1,k}$.

Denominators and within-set shares. A perturbation of π_m differentiates only denominators in $R_{P,m}$ and $R_{p,m}$; each resulting term is a function of X times a conditionally mean-zero residual, so its expectation is zero. The same argument applies to a perturbation of $\tau_{j,k}$ in $R_{q,j,k}$. Finally, admissible share directions satisfy $\sum_{z \in T_{j,k}} \delta a_z(X) = 0$. At the truth, all $q(z, X)$ in $T_{j,k}$ equal $q_{j,k}(X)$, and hence

$$\partial_t \sum_{z \in T_{j,k}} (a_{j,k}(z, X) + t \delta a_z(X)) q(z, X) \Big|_{t=0} = q_{j,k}(X) \sum_{z \in T_{j,k}} \delta a_z(X) = 0.$$

Thus the simplex-share direction has no first-order effect; with estimated $q(z, X)$ it contributes only the product $\sum_z (\hat{a}_z - a_z)(\hat{q}_z - q_{j,k})$.

Inheritance by product scores. Recall that

$$\psi_{j,k}^{\text{prod}} = \gamma_{j,k} \psi_{j,k}^h + h_{j,k}^- (\mathbb{1}(\cdots) - \gamma_{j,k}), \quad \mathbb{E}[\mathbb{1}(\cdots) - \gamma_{j,k} \mid X] = 0 \text{ at } \eta_0,$$

while the preceding component-score cancellations hold conditionally on X . Hence any perturbation of a component-score nuisance $(P_{1,m}, p_m, p(z, X), \pi_m, \tau_{j,k}, a_{j,k})$ yields

$$\partial_\eta \mathbb{E}[\psi_{j,k}^{\text{prod}}] = \mathbb{E} \left[\gamma_{j,k}(X) \underbrace{\partial_\eta \mathbb{E}[\psi_{j,k}^h \mid X]}_{=0} + \partial_\eta h_{j,k}^-(X) \underbrace{\mathbb{E}[\mathbb{1}(\cdots) - \gamma_{j,k}(X) \mid X]}_{=0} \right] = 0.$$

For $\psi_{\text{full},k}^{\text{prod}}$, the same calculation uses the signed residual

$$\mathbb{1} \left(Z \in \bigcup_{m=k+1}^K \bigcup_{l=0}^{l_m} T_{l,m} \right) - \mathbb{1} \left(Z \in \bigcup_{m=k+1}^K S_m \right) - \gamma_{\text{full},k}(X),$$

whose conditional mean is zero. For ψ_{Δ_K} , the p_K plug-in derivative cancels the derivative of its $R_{p,K}$ correction, while each instrument-specific derivative through $q(z, X) = \phi(z, X, p(z, X))$ cancels the corresponding policy residual. Perturbations of π_K , γ_K , and the policy-derivative coefficient multiply conditionally mean-zero residuals. Hence ψ_{Δ_K} is also orthogonal in every nuisance direction.

Perturbation of $\gamma_{j,k}$ and $\gamma_{\text{full},k}$. The derivative contributions of the two product-score terms cancel. Along $\gamma_{j,k} \mapsto \gamma_{j,k} + t \delta\gamma$,

$$\partial_t \mathbb{E}[(\gamma_{j,k} + t \delta\gamma) \psi_{j,k}^h + h_{j,k}^-(\mathbb{1}(\cdots) - \gamma_{j,k} - t \delta\gamma)] \Big|_{t=0} = \mathbb{E}[\delta\gamma \mathbb{E}[\psi_{j,k}^h | X]] - \mathbb{E}[\delta\gamma h_{j,k}^-] = 0,$$

because $\mathbb{E}[\psi_{j,k}^h | X] = h_{j,k}^-(X)$ at η_0 . Replacing the set-indicator residual by the signed residual displayed above gives the identical cancellation for $\gamma_{\text{full},k}$. □

O5 Rate and Learner Conditions for the Concrete Continuous-Instrument Estimator

Theorem 4.10 accepts any nuisance construction satisfying its three target losses. We now express those losses in terms of the learner prediction errors for Algorithm 3, stated in Section B.3.1 of the Appendix. ?? O5.2 uses the following quantities:

- $r_{p,n}$, a sup-norm first-stage rate, $\|\hat{p} - p\|_\infty = O_P(r_{p,n})$;
- $r_{g,\text{loc},n}$, an L_1 prediction rate for $\hat{g}_1$, required only at the two arguments $\hat{q}^{\text{tr}}(Z, X)$ and $\hat{P}$ at which the target evaluates it;
- $r_{Q,n}$, a $p(X)$ -weighted L_1 rate for the conditional quantile fit relative to the law generated by the estimated boundary window, and ε_n , the width of the two analytically trimmed quantile tails;
- $\delta_n + r_{p,n}$, the structural distance between that selected law and the canonical treated boundary law.

The first stage remains uniform because $\hat{p}$ determines conditional support extrema and membership in the shrinking boundary sample in addition to serving as a generated regressor. The same uniform first-stage requirement arises in nonparametric problems with generated covariates; see, for example, [Mammen et al. \[2012, Assumption 2\]](#) and [Lee \[2018, Assumption 7\]](#).

Corollary O5.1 (Rate for the Concrete Estimator). *Suppose ?? 2.1 and ?? 3.6 hold. Suppose additionally that ϕ is uniformly Lipschitz in its last argument: for some $L_\phi < \infty$,*

$$|\phi(z, x, p') - \phi(z, x, p)| \leq L_\phi |p' - p|$$

for all $(z, x, p, p') \in \mathcal{Z} \times \mathcal{X} \times [0, 1]^2$. Suppose that ?? O5.2 holds with rates $r_{p,n}$, $r_{g,\text{loc},n}$, $r_{Q,n}$ and trimming sequence ε_n , that $|I_j|/n \rightarrow \rho_j \in (0, 1)$ for $j = 0, 1, 2$, and that

$$r_{p,n}, r_{g,\text{loc},n}, r_{Q,n}, \delta_n, \varepsilon_n, M_Q^{-1} \rightarrow 0.$$

Then the nuisance estimators produced by Algorithm 3 satisfy ?? 4.9 with

$$\begin{aligned} r_{e,n} &= O(r_{p,n}), \\ r_{g,n} &= O(r_{g,\text{loc},n} + r_{p,n}), \\ r_{\mu,n} &= O(r_{Q,n} + \delta_n + r_{p,n} + \varepsilon_n + M_Q^{-1}). \end{aligned}$$

and consequently

$$|\widehat{\theta}_{\omega,1} - \underline{\theta}_{\omega,1}| = O_P(n^{-1/2} + r_{p,n} + r_{g,\text{loc},n} + r_{Q,n} + \delta_n + \varepsilon_n + M_Q^{-1}).$$

Remark. The endpoint-invariance assertion for population Lebesgue integrals concerns changes only at $\tau \in \{0, 1\}$ of a fixed population quantile; the two permitted fitted interpolants may differ on $[0, M_Q^{-1}] \cup [1 - M_Q^{-1}, 1]$.

Remark (notation). $L_\phi < \infty$ is a uniform Lipschitz bound satisfying $|\phi(z, x, p') - \phi(z, x, p)| \leq L_\phi |p' - p|$ for every $(z, x, p, p') \in \mathcal{Z} \times \mathcal{X} \times [0, 1]^2$.

The rate separates into three sources. The learner errors $r_{g,\text{loc},n}$ and $r_{Q,n}$ are the statistical cost of the outcome and interior boundary-quantile regressions. The first-stage error $r_{p,n}$ enters through the generated regressor, the support estimates, the two evaluation arguments, and the width of the selected boundary region on the true propensity scale. The threshold-crossing structure bounds the last contribution by $O(\delta_n + r_{p,n})$. The terms ε_n and M_Q^{-1} bound the two extreme quantile tails and quantile-grid interpolation.

The rest of this section states ?? O5.2 and proves Corollary O5.1. Remarks O5.3 and O5.4 then comment informally on series constructions that plausibly deliver (L2) and (L3), drawing on least-squares series prediction bounds from Belloni et al. [2015], generated-covariate expansions from Mammen et al. [2012], Lee [2018], and the interior-quantile process theory of Belloni et al. [2019].

Set $U_1 = q^{\text{tr}}(Z, X)$, $U_2 = P$, $\widehat{U}_1 = \widehat{q}^{\text{tr}}(Z, X)$, and $\widehat{U}_2 = \widehat{P}$, and define

$$\mathcal{S}_x = [\underline{p}(x), \bar{p}(x)], \quad \Pi_x(u) = \min\{\bar{p}(x), \max\{\underline{p}(x), u\}\}.$$

At a grid point u_m , $\widehat{g}_1(u_m, \cdot)$ denotes the fitted local-linear intercept in Algorithm 3; between grid points, $\widehat{g}_1$ denotes its clipped linear interpolant. All stochastic learner rates below are conditional on the training samples $I_0 \cup I_1$.

Remark (notation). $\mathcal{S}_x = [\underline{p}(x), \bar{p}(x)]$ is the population propensity interval at covariate value x , used as the domain of the population comparison after projection by Π_x . $\Pi_x : \mathbb{R} \rightarrow \mathcal{S}_x$ is projection onto the population support interval $\mathcal{S}_x = [\underline{p}(x), \bar{p}(x)]$, given by $\Pi_x(u) = \min\{\bar{p}(x), \max\{\underline{p}(x), u\}\}$.

Assumption O5.2 (Continuous-Instrument Learner Conditions). (L1) **First stage.** $\|\widehat{p} - p\|_\infty = O_P(r_{p,n})$.

(L2) **Outcome learner.** The fitted learner satisfies the target prediction rate

$$\max_{\ell \in \{1,2\}} \mathbb{E}_{Z,X} \left[\left| \widehat{g}_1(\widehat{U}_\ell, X) - g_1(\Pi_X(\widehat{U}_\ell), X) \right| \mid I_0, I_1 \right] = O_P(r_{g,\text{loc},n}).$$

(L3) **Boundary learner.** Conditional on I_0 , define the selected-law target

$$\mu_{1,n|x}^{\text{sel}} = \mathcal{L}\left(Y \mid W = 1, \widehat{P} \leq \widehat{p}(x) + \delta_n, X = x, I_0\right).$$

At any x for which $\mathbb{P}(W = 1, \widehat{P} \leq \widehat{p}(x) + \delta_n \mid X = x, I_0) = 0$, set $\mu_{1,n|x}^{\text{sel}} = \mu_{1,\underline{p}|x}$. Let $Q_{1,n|x}^{\text{sel}}$ be its quantile function. Let

$$\widetilde{Q}_m^c(x) = \min\{y_{\max}, \max\{y_{\min}, \widetilde{Q}(\tau_m \mid x)\}\}$$

be the clipped grid fit before rearrangement. For a sequence $M_Q^{-1} \leq \varepsilon_n < 1/2$ with $\varepsilon_n \rightarrow 0$, define $\mathcal{J}_{Q,n} = \{m : \varepsilon_n \leq \tau_m \leq 1 - \varepsilon_n\}$ and suppose

$$\frac{1}{M_Q} \sum_{m \in \mathcal{J}_{Q,n}} \mathbb{E}_X \left[\left| \underline{p}(X) | \widetilde{Q}_m^c(X) - Q_{1,n|X}^{\text{sel}}(\tau_m) \right| \mid I_0, I_1 \right] = O_P(r_{Q,n}).$$

Remark. For $0 < t < 1$, $Q_{1,n|x}^{\text{sel}}(t) = \inf\{y \in [y_{\min}, y_{\max}] : \mu_{1,n|x}^{\text{sel}}((-\infty, y]) \geq t\}$. $\|\widehat{p} - p\|_\infty = \sup_{(z,x) \in \mathcal{Z} \times \mathcal{X}} |\widehat{p}(z, x) - p(z, x)|$. In (L2), $O_P(r_{g,\text{loc},n})$ describes stochastic order over the training observations and random partition, while $\mathbb{E}_{Z,X}[\cdot \mid I_0, I_1]$ averages an independent population draw of (Z, X) conditional on the partition and the observations indexed by $I_0 \cup I_1$. In (L3), $O_P(r_{Q,n})$ describes stochastic order over the training observations and random partition, while $\mathbb{E}_X[\cdot \mid I_0, I_1]$ averages an independent population draw of X conditional on the partition and the observations indexed by $I_0 \cup I_1$. For $0 < t < 1$, $Q_{1,n|x}^{\text{sel}}(t) = \inf\{y \in \mathbb{R} : \mu_{1,n|x}^{\text{sel}}(\mathcal{Y} \cap (-\infty, y]) \geq t\}$. The superscript c in $\widetilde{Q}_m^c$ labels clipping to $[y_{\min}, y_{\max}]$ and does not denote the transport cost c .

Remark (notation). For $\ell \in \{1, 2\}$, $\widehat{U}_1 = \widehat{q}^{\text{tr}}(Z, X)$ and $\widehat{U}_2 = \widehat{P}$ are the two target-evaluation indices, with evaluation over (Z, X) conditional on I_0, I_1 , and the population comparison for $\widehat{g}_1(\widehat{U}_\ell, X)$ is $g_1(\Pi_X(\widehat{U}_\ell), X)$. $\mu_{1,n|x}^{\text{sel}}$ is the conditional population outcome law under the fitted selection event $\{W = 1, \widehat{p}(Z, X) \leq \widehat{p}(x) + \delta_n\}$ given $X = x$ and I_0 , with the canonical boundary law used when that event has zero conditional probability. $Q_{1,n|x}^{\text{sel}} : (0, 1) \rightarrow \mathcal{Y}$ is the quantile function of the I_0 -conditional selected law $\mu_{1,n|x}^{\text{sel}}$, indexed by sample size n and covariate value x , including its stipulated zero-selection version. $\widetilde{Q}_m^c(x)$ is the clipped pre-rearrangement fit at $\tau_m = m/M_Q$, for $m = 1, \dots, M_Q - 1$, obtained from the boundary quantile regression on $\mathcal{I}_{\delta,1} \subseteq I_1$ using the selection rule determined on I_0 . $\mathcal{J}_{Q,n} = \{m \in \{1, \dots, M_Q - 1\} : \varepsilon_n \leq \tau_m \leq 1 - \varepsilon_n\}$ is the interior grid subset used for analytical quantile trimming; it does not trim the estimator.

Condition (L2) uses the two random arguments that appear in the estimand. The learner is trained under the generated design $(\widehat{P}, X)$, whereas (L2) measures prediction under the evaluation laws induced by $(\widehat{U}_\ell, X)$. For $\ell = 2$, $\widehat{U}_2 = \widehat{P}$, so the training and evaluation index distributions agree up to sample splitting. For $\ell = 1$, $\widehat{U}_1 = \widehat{q}^{\text{tr}}(Z, X)$ may have a different distribution and may place mass at a support endpoint after truncation, so a learner meeting (L2) must predict well under both evaluation laws, including one-sidedly at such an endpoint; see Remark [O5.3](#). Projection keeps the population comparison inside its identified support. The structural representation makes $g_1(\cdot, x)$ automatically Lipschitz, which controls the difference between $\Pi_X(\widehat{U}_\ell)$ and U_ℓ .

Remark (notation). For $\ell \in \{1, 2\}$, $U_1 = q^{\text{tr}}(Z, X)$ and $U_2 = P = p(Z, X)$ are the population evaluation indices, with $U_\ell \in \mathcal{S}_X$; their estimated counterparts are $\widehat{U}_1 = \widehat{q}^{\text{tr}}(Z, X)$ and $\widehat{U}_2 = \widehat{P}$.

In (L3), the weight $\underline{p}(X)$ comes from the boundary term in [\(4.7\)](#), which weights the conditional law at $X = x$ by $\underline{p}(x)$. Condition (L3) compares the grid fit to $Q_{1,n|x}^{\text{sel}}$, the law generated by the selection rule actually applied to I_1 . The structural localization inequality derived below transfers this selected law to the canonical boundary law at cost $O(\delta_n + r_{p,n})$. The interior set $\mathcal{J}_{Q,n}$ confines the quantile regression analysis to intermediate quantiles. Because both quantile functions lie in $[y_{\min}, y_{\max}]$, the excluded tails have total Lebesgue measure at most $2\varepsilon_n + O(M_Q^{-1})$. This analytical trimming leaves the estimator unchanged and costs $O(\varepsilon_n + M_Q^{-1})$ in the target.

Proof of Corollary [O5.1](#). Condition (L1) implies Part (1) of [?? 4.9](#). Indeed, the policy-map Lipschitz condition in Corollary [O5.1](#) gives $\|\widehat{q} - q\|_\infty \leq L_\phi \|\widehat{p} - p\|_\infty$, while infimum, supremum, minimum, projection, and positive part are Lipschitz in their arguments. Hence $r_{e,n} = O(r_{p,n})$

and, with $\widehat{U}_1 = \widehat{q}^{\text{tr}}(Z, X)$ and $\widehat{U}_2 = \widehat{P}$,

$$\max_{\ell \in \{1, 2\}} \|\widehat{U}_\ell - U_\ell\|_\infty = O_P(r_{p,n}).$$

Let $C_Y = \max\{|y_{\min}|, |y_{\max}|\}$. The structural representation

$$g_1(v, x) - g_1(u, x) = \int_u^v \mathbb{E}[Y(1) \mid U = s, X = x] \, ds$$

gives $|g_1(v, x) - g_1(u, x)| \leq C_Y |v - u|$. Thus g_1 is uniformly Lipschitz because the outcome domain is compact. Fix $\ell \in \{1, 2\}$. Because projection onto $\mathcal{S}_X$ is nonexpansive and $U_\ell \in \mathcal{S}_X$,

$$|\Pi_X(\widehat{U}_\ell) - U_\ell| \leq |\widehat{U}_\ell - U_\ell|.$$

The structural Lipschitz bound therefore gives

$$|\widehat{g}_1(\widehat{U}_\ell, X) - g_1(U_\ell, X)| \leq |\widehat{g}_1(\widehat{U}_\ell, X) - g_1(\Pi_X(\widehat{U}_\ell), X)| + C_Y |\widehat{U}_\ell - U_\ell|.$$

Condition (L2) and the first-stage bound thus imply Part (2) of ?? 4.9 with rate

$$O_P(r_{g, \text{loc}, n} + r_{p, n}).$$

For the boundary term, fix x and write $a = \underline{p}(x)$. For $a > 0$ and $p > a$, the threshold-crossing representation gives

$$\nu_{p, x} = \frac{a}{p} \nu_{a, x} + \frac{p-a}{p} \bar{\nu}_{a, p, x}, \quad \bar{\nu}_{a, p, x} = \frac{1}{p-a} \int_a^p K_{1, x}(\cdot \mid u) \, du.$$

Both component laws are supported on $[y_{\min}, y_{\max}]$. Coupling the common a/p share identically and coupling the remaining share arbitrarily yields

$$W_1(\nu_{p, x}, \nu_{a, x}) \leq \frac{p-a}{p} B_Y, \quad a W_1(\nu_{p, x}, \nu_{a, x}) \leq B_Y(p-a),$$

which establishes the structural localization inequality; its weighted statement is immediate at $a = 0$.

Set $e_n = \|\widehat{p} - p\|_\infty$. At every x where the selection event has positive probability, the selected outcome law in (L3), conditional on I_0 , is a mixture of the laws $\nu_{P, x}$. Indeed, the selection event is a fixed function of (W, Z, X) , and conditional instrumental exogeneity makes the treated outcome law at any selected $Z = z$ equal to $\nu_{p(z, x), x}$. Every selected instrument value satisfies

$$a \leq P \leq \widehat{P} + e_n \leq \underline{\widehat{p}}(x) + \delta_n + e_n \leq a + \delta_n + 2e_n.$$

Convexity of W_1 and the structural localization inequality now give, pointwise in x ,

$$aW_1\left(\mu_{1,n|x}^{\text{sel}}, \mu_{1,p|x}\right) \leq B_Y(\delta_n + 2e_n).$$

The same bound holds at the remaining x values by the version convention in (L3). After averaging over X , (L1) makes this term $O_P(\delta_n + r_{p,n})$.

For each x , sorting the clipped fitted vector is an L_1 contraction toward the selected-law population quantile vector, which is monotone. Hence condition (L3), clipping, and the bound B_Y on each excluded-grid error give

$$\frac{1}{M_Q} \sum_{m=1}^{M_Q-1} \mathbb{E}_X \left[\underline{p}(X) |\hat{Q}_m(X) - Q_{1,n|x}^{\text{sel}}(\tau_m)| \mid I_0, I_1 \right] = O_P(r_{Q,n} + \varepsilon_n + M_Q^{-1}).$$

Both $Q_{1,n|x}^{\text{sel}}$ and the rearranged fitted quantile curve are nondecreasing and take values in $[y_{\min}, y_{\max}]$. Let $h_Q = M_Q^{-1}$ and write $e_m(x) = |\hat{Q}_m(x) - Q_{1,n|x}^{\text{sel}}(\tau_m)|$. On an interior cell $[\tau_m, \tau_{m+1}]$, linear interpolation and monotonicity give

$$\begin{aligned} \int_{\tau_m}^{\tau_{m+1}} |\hat{Q}(\tau, x) - Q_{1,n|x}^{\text{sel}}(\tau)| d\tau \\ \leq \frac{h_Q}{2} \{e_m(x) + e_{m+1}(x)\} + h_Q \{Q_{1,n|x}^{\text{sel}}(\tau_{m+1}) - Q_{1,n|x}^{\text{sel}}(\tau_m)\}. \end{aligned}$$

Summing the first term averages the grid-point errors, while the second telescopes and is at most $B_Y h_Q$. Together with the excluded tails, this gives

$$\mathbb{E}_X \left[\underline{p}(X) W_1(\hat{\mu}_{1,p|x}, \mu_{1,n|x}^{\text{sel}}) \mid I_0, I_1 \right] = O_P(r_{Q,n} + \varepsilon_n + M_Q^{-1}).$$

The triangle inequality thus implies Part (3) with

$$r_{\mu,n} = O(r_{Q,n} + \delta_n + r_{p,n} + \varepsilon_n + M_Q^{-1}).$$

Substitution into Theorem 4.10 proves the result. $\square$

Remark O5.3 (Series construction for the outcome learner). Condition (L2) restricts prediction accuracy rather than the choice of learner, and series methods are a natural candidate for meeting it. Take $\mathcal{F}_n$ to be a $K_{g,n}$ -dimensional series class in x , so that the local-linear step at grid point u_m regresses YW on the augmented regressor

$$\left(b_{K_{g,n}}(X), \frac{\hat{P} - u_m}{h_n} b_{K_{g,n}}(X) \right)$$

under kernel weights of bandwidth h_n . Suppose that g_1 has the smoothness in u needed for the usual local-linear bias bound, the series approximates $g_1(u, \cdot)$ and $\partial_u g_1(u, \cdot)$ with error $a_{g,n}$ under the relevant localized and evaluation laws, and the corresponding designs are uniformly well conditioned—including one-sidedly at a support endpoint, where $\widehat{U}_1$ may place mass after truncation. If also $h_n \rightarrow 0$ and $(r_{p,n} + M_g^{-1})/h_n \rightarrow 0$, the series prediction bounds of Belloni et al. [2015] and the local-linear generated-covariate expansions of Mammen et al. [2012], Lee [2018] suggest that one may take

$$r_{g,\text{loc},n} \lesssim a_{g,n} + h_n^2 + \sqrt{\frac{K_{g,n}}{nh_n}} + \frac{r_{p,n}}{h_n} + M_g^{-1}.$$

Here nh_n is the effective local sample size, h_n^2 the local-linear bias, $r_{p,n}/h_n$ the generated-regressor contribution, and M_g^{-1} the grid interpolation error. The scale separation $(r_{p,n} + M_g^{-1})/h_n \rightarrow 0$ keeps both the generated-regressor displacement and the bracketing grid points inside the $O(h_n)$ neighborhood on which the localized design is controlled, and the L_1 requirement in (L2) follows from an L_2 bound by Cauchy–Schwarz.

Remark O5.4 (Series construction for the boundary learner). For (L3), the analogous candidate is series conditional quantile regression, with $\mathcal{F}_Q$ a $K_{Q,n}$ -dimensional series class, run on the treated observations that the estimated boundary window selects. To see the effective sample size, first ignore first-stage error and write $a(x) = \underline{p}(x)$. If $f_{P|X}(\cdot | x)$ is uniformly bounded above and away from zero in a right neighborhood of $a(x)$, the oracle-window selection probability satisfies

$$\begin{aligned} s_{\delta_n}(x) &= \mathbb{P}(W = 1, P \leq a(x) + \delta_n \mid X = x) \\ &= \int_{a(x)}^{a(x) + \delta_n} p f_{P|X}(p \mid x) \, dp \asymp \delta_n \{a(x) + \delta_n\}. \end{aligned}$$

Thus the unconditional selection probability is of order $\delta_n \mathbb{E}[\underline{p}(X)] + \delta_n^2$. When $\mathbb{E}[\underline{p}(X)] > 0$, the effective sample size is therefore $n\delta_n$, and the normalized selected- X density, which is proportional to $s_{\delta_n}(x)$, dominates the target weight $\underline{p}(x)$ up to a constant. An unweighted prediction bound on the localized sample then transfers to the $\underline{p}(X)$ -weighted loss in (L3). If $\mathbb{E}[\underline{p}(X)] = 0$, the boundary term has zero target weight and no boundary learner is needed. Under $r_{p,n} = o(\delta_n)$, the estimated window inherits these sample-size and design comparisons with probability approaching one.

Conditional on I_0 , the quantile regression directly targets $Q_{1,n|x}^{\text{sel}}$, the selected-law quantile in (L3). Under the usual QR-series approximation and design conditions, the interior quantile-process bounds of Belloni et al. [2019], applied on $\tau \in [\varepsilon_n, 1 - \varepsilon_n]$, therefore suggest that one

may take

$$r_{Q,n} \lesssim a_{Q,n}(\varepsilon_n) + \sqrt{\frac{K_{Q,n} \log M_Q}{n\delta_n}},$$

where $a_{Q,n}(\varepsilon_n)$ is the series approximation error for the selected law over the trimmed quantile range. Alternatively, one may verify these conditions by comparison with the oracle window. The estimated window then changes probability content by $O_P(r_{p,n})$ and its normalized law by $O_P(r_{p,n}/\delta_n)$. If, in addition, the conditional outcome density is uniformly bounded away from zero at the relevant interior quantiles, this comparison gives the conservative bound

$$r_{Q,n} \lesssim a_{Q,n}(\varepsilon_n) + \sqrt{\frac{K_{Q,n} \log M_Q}{n\delta_n}} + \frac{r_{p,n}}{\delta_n}.$$

In either formulation, a narrower window brings the selected law closer to the boundary law at the cost of a slower quantile fit and a more demanding first stage. The sampling term vanishes when $K_{Q,n} \log M_Q = o(n\delta_n)$.

O6 Additional Simulation Details

This section presents the experimental designs, estimands, and supporting numerical results. The replication package gives the random seeds, learner tuning, numerical grids, solver options, and run instructions used to produce every exhibit.

O6.1 Synthetic Experiment: Data-Generating Process

Continuous instrument. The DGP has the following components:

- **Instrument and latent resistance.** $Z \sim \text{Unif}(0, 1)$, independent of everything else, and $U \sim \text{Unif}(0, 1)$.
- **Propensity score and treatment.** The propensity score is $p(Z) = \text{logistic}(\beta_0 + \beta_1 Z)$ with $\beta_0 = -1$ and $\beta_1 = 2$, so $p(Z)$ ranges from $\text{logistic}(-1) \approx 0.27$ to $\text{logistic}(1) \approx 0.73$. Treatment follows the threshold-crossing rule $W = \mathbb{1}(U \leq p(Z))$.
- **Potential outcomes.** With $\varepsilon \sim \text{Unif}(-0.5, 0.5)$, define $Y(1) = aU + \theta + \varepsilon$ and $Y(0) = aU + \varepsilon$, where $a = 0.5$ and $\theta = 0.5$. The common shock is therefore part of both potential outcomes, so consistency gives $Y = WY(1) + (1 - W)Y(0)$. Thus $Y(0)$ ranges over $[-0.5, 1]$ and $Y(1)$ over $[0, 1.5]$, so the tightest common bound is $[y_{\min}, y_{\max}] = [-0.5, 1.5]$, and both endpoints are attained.

The MTE is constant: $\text{MTE}(u) = \mathbb{E}[Y(1) - Y(0) \mid U = u] = \theta = 0.5$ for all u . The policy of interest is $q_\alpha(Z) = \text{clip}(p(Z) + \alpha, 0, 1)$, and the target parameter is

$$\theta_\alpha = \mathbb{E}[Y^{q_\alpha} - Y] = \theta \cdot \mathbb{E}[\mathbb{1}(U \leq q_\alpha(Z)) - \mathbb{1}(U \leq p(Z))].$$

Over the reported range $|\alpha| \leq 0.12$ we have $p(Z) + \alpha \in [0.15, 0.85] \subset [0, 1]$, so clipping is inactive and $\mathbb{E}[\mathbb{1}(U \leq q_\alpha(Z)) - \mathbb{1}(U \leq p(Z))] = \alpha$. The ground truth is therefore exact:

$$\theta_\alpha = \theta \cdot \alpha = 0.5\alpha.$$

Discrete instrument. The DGP has the following components, with a propensity score that takes only two values:

- **Instrument and latent resistance.** $Z \sim \text{Bernoulli}(0.5)$ (binary) and $U \sim \text{Unif}(0, 1)$, independent of each other.
- **Propensity score and treatment.** The propensity score is piecewise constant with two levels:

$$p(Z) = \begin{cases} 0.25 & \text{if } Z = 0, \\ 0.75 & \text{if } Z = 1. \end{cases}$$

Treatment follows the threshold-crossing rule $W = \mathbb{1}(U \leq p(Z))$.

- **Potential outcomes.** Let $\varepsilon \sim \text{Unif}(-0.1, 0.1)$ be independent of (U, Z) and define $Y(0) = 0.1 + 0.22U + \varepsilon$ and $Y(1) = 0.58 + 0.40U + \varepsilon$. The common shock is therefore part of both potential outcomes, so consistency gives $Y = WY(1) + (1 - W)Y(0)$. The treatment effect remains $\tau(U) = Y(1) - Y(0) = 0.48 + 0.18U$, and hence $\text{MTE}(u) = 0.48 + 0.18u$, ranging from 0.48 to 0.66. The resulting outcome support is contained in $[y_{\min}, y_{\max}] = [0, 1.08]$.

The target parameter is $\theta_\alpha = \mathbb{E}[Y^{q_\alpha} - Y]$ under the same uniform policy shift $q_\alpha(Z) = \text{clip}(p(Z) + \alpha, 0, 1)$. For the reported range $|\alpha| \leq 0.12$, clipping is inactive at both propensity values. Thus the following ground truth is exact (with oriented integrals when $\alpha < 0$):

$$\theta_\alpha = \frac{1}{2} \int_{0.25}^{0.25+\alpha} (0.48 + 0.18u) du + \frac{1}{2} \int_{0.75}^{0.75+\alpha} (0.48 + 0.18u) du = 0.57\alpha + 0.09\alpha^2.$$

Because the propensity score support consists of only two points $\{0.25, 0.75\}$, the data do not identify $\text{MTE}(u)$ at any single rank u . What is identified is the aggregate over the

complier block: differencing the two instrument arms gives $\mathbb{E}[Y \mid Z = 1] - \mathbb{E}[Y \mid Z = 0] = \int_{0.25}^{0.75} \text{MTE}(u) du$, that is, the complier average LATE $= \frac{1}{0.5} \int_{0.25}^{0.75} \text{MTE}(u) du$, together with the marginal laws of $Y(1)$ and $Y(0)$ for the block $U \in (0.25, 0.75]$ as a whole. The policy shift q_α moves mass over the sub-block $(0.25, 0.25 + \alpha]$, which is strictly finer than the identified block, and over $(0.75, 0.75 + \alpha]$, on which nothing beyond the outcome support is known.

Discrete instrument with covariates. The DGP has the following components, with five covariates and a three-valued instrument:

- **Covariates and indices.** $X \sim \text{Unif}(-1, 1)^5$, and two scalar indices drive the design: $\xi(X) = \Lambda(X'\beta_1)$ with $\beta_1 = (1.5, -1.2, 0.9, -0.6, 0.3)$ and Λ the logistic function, and $\zeta(X) = X'\beta_2$ with β_2 the β_1 -orthogonalized, unit-normalized version of $(-0.6, 0.9, 1.2, 0.3, -1.5)$. Neither index is known to the estimator, which sees only the five raw covariates, so every nuisance function is a genuinely nonparametric object in five dimensions.
- **Instrument.** $\mathcal{Z} = \{0, 1, 2\}$ with $\mathbb{P}_{\text{obs}}(Z = z \mid X = x) \propto \exp(d_z \zeta(x))$ and $d = (-0.35, 0, 0.35)$, so that every $\pi_k(x)$ lies in $[0.16, 0.54]$ and ?? 4.5 holds with room to spare, while π_k remains a nondegenerate function of X .
- **Propensity score and treatment.** $p(z, x) = a_z + 0.10 \xi(x)$ with $a = (0.20, 0.35, 0.50)$, and $W = \mathbb{1}(U \leq p(Z, X))$ with $U \mid X \sim \text{Unif}(0, 1)$. The covariate shift is common across instrument values, so every pairwise difference among the values in $R_x = \{p(z, x), q(z, x)\}_{z \in \mathcal{Z}}$ does not depend on x . Part (i) of ?? 4.2 therefore holds exactly: the level sets S_k , $T_{j,k}$ and their ordering are literally invariant in x , and the separation constant can be read off a single calculation. We take $c_{\text{gap}} = 0.04$.
- **Potential outcomes.** With $\varepsilon \sim \text{Unif}(-0.05, 0.05)$ independent of (U, Z, X) , $Y(0) = 0.30 + 0.20 \xi(X) + 0.15U + \varepsilon$ and $Y(1) = 0.50 + 0.40 \xi(X) + 0.40U + \varepsilon$, so that $\text{MTE}(u, x) = 0.30 + 0.25u + 0.20(\xi(x) - \frac{1}{2})$ varies in both arguments. We use the known outcome bound $[y_{\min}, y_{\max}] = [0.25, 1.35]$, which contains the support of both potential outcomes. The measurement error gives every relevant conditional law a bounded density, as part (i) of Theorem 4.7 requires at the boundary quantiles.

The policy is the uniform shift $q(z, x) = p(z, x) + \alpha$. We report $\alpha \in [0.195, 0.255]$, a window that satisfies the separation condition and activates the full-interval, sub-interval, product, and tail components of the orthogonal score. The Monte Carlo study uses $\alpha = 0.22$.

The affine outcome means again yield population bounds in closed form up to the outer expectation over X ; the replication package compares them with solutions to the primitive CCOT problem (3.2).

Sample sizes and comparison methods. The continuous and discrete instrument experiments both use $n = 5,000$ observations, and the covariate design uses $n = 4,000$ for the reported identified sets and n ranging from 1,000 to 16,000 for the Monte Carlo study. IVMTE bounds are computed using the `ivmte` R package with moment constraints derived from the observed data. The PRTE weight function $\omega(u)$ corresponding to the q_α -policy is used for both methods.

We approximate the unrestricted MTR with flexible u -spline sieves. In the covariate design, the sieve is interacted with all five covariates, and we report degrees 3 and 7. Degree 7 is the richest sieve that solves reliably across the policy grid without numerical error. The replication package provides the spline bases, the fitted propensity model, target-weight construction, solver settings, and bootstrap configuration.

O6.2 Synthetic Experiment: Numerical Results

Table 3 reports the continuous-design bounds for different values of $\alpha > 0$, together with the IVMTE 95% backward confidence interval; the full grid is plotted in Figure 2. The IVOT bounds are tighter at every point of the 25-point grid, by a factor that grows from $3.4\times$ at $|\alpha| = 0.12$ to $23.3\times$ at $|\alpha| = 0.01$: the IVOT width shrinks at the rate at which the policy perturbation itself shrinks, whereas the IVMTE width is driven by the unidentified part of the MTE and stays comparatively wide. Table 4 reports the corresponding discrete-design bounds under the outcome support $[0, 1.08]$, and adds the IVOT 95% delta-method confidence interval; the full grid is plotted in Figure 3. There the IVOT advantage is more modest and roughly flat in α : 2.4 – $2.7\times$ over the positive shifts displayed in the table, and 2.4 – $3.6\times$ over the full grid, the upper end being attained at $\alpha = -0.05$. The two-point propensity support confines both methods to bounding the same two objects, the sub-block $(0.25, 0.25 + \alpha]$ of the identified complier block and the never-taker block $(0.75, 0.75 + \alpha]$.

On the single draw of Table 3 the truth falls inside the estimated IVOT interval at 24 of the 25 grid points, with the exception at $\alpha = 0.01$. The population IVOT interval has width of order α^2 near zero: its width is 4×10^{-4} at $|\alpha| = 0.01$ and 0.045 at $|\alpha| = 0.12$. Across the grid the estimated width is 0.92 – 0.97 times the sharp width. Table 5 repeats the experiment over 20 independent draws. The inclusion rate rises with $|\alpha|$ on both sides of zero, from

Table 3: Continuous-instrument design ($n = 5,000$): selected IVOT and IVMTE bounds for θ_α . The truth is 0.5α ; $\dagger$ marks the single grid point where the estimated IVOT interval misses it. Width ratios use unrounded bounds, and the IVMTE column reports a 95% backward confidence interval.

α	Truth	IVOT $[\ell, u]$	IVOT width	IVMTE $[\ell, u]$	IVMTE 95% CI	Width ratio
0.01 †	0.0050	[0.0051, 0.0054]	0.0004	[−0.0037, 0.0043]	[−0.0038, 0.0043]	22.4
0.03	0.0150	[0.0145, 0.0176]	0.0032	[−0.0129, 0.0333]	[−0.0170, 0.0357]	14.6
0.05	0.0250	[0.0226, 0.0310]	0.0085	[−0.0130, 0.0506]	[−0.0178, 0.0542]	7.5
0.07	0.0350	[0.0289, 0.0451]	0.0162	[−0.0181, 0.0707]	[−0.0275, 0.0784]	5.5
0.10	0.0500	[0.0353, 0.0669]	0.0316	[−0.0273, 0.1002]	[−0.0409, 0.1123]	4.0
0.12	0.0600	[0.0376, 0.0820]	0.0445	[−0.0319, 0.1196]	[−0.0446, 0.1349]	3.4

Table 4: Discrete-instrument design ($n = 5,000$): selected IVOT and IVMTE bounds for θ_α under outcome support $[0, 1.08]$. The analytic truth is $0.57\alpha + 0.09\alpha^2$ and lies in every estimated IVOT interval. Width ratios use unrounded bounds; the reported intervals use the delta method for IVOT and backward inference for IVMTE.

α	Truth	IVOT $[\ell, u]$	IVOT width	IVOT 95% CI	IVMTE $[\ell, u]$	IVMTE 95% CI	Width ratio
0.01	0.0057	[−0.0007, 0.0089]	0.0096	[−0.0062, 0.0128]	[−0.0114, 0.0118]	[−0.0118, 0.0118]	2.4
0.03	0.0172	[−0.0010, 0.0259]	0.0269	[−0.0069, 0.0307]	[−0.0340, 0.0362]	[−0.0358, 0.0362]	2.6
0.05	0.0287	[−0.0007, 0.0422]	0.0430	[−0.0069, 0.0476]	[−0.0511, 0.0543]	[−0.0525, 0.0544]	2.5
0.07	0.0403	[−0.0001, 0.0581]	0.0581	[−0.0063, 0.0640]	[−0.0667, 0.0759]	[−0.0688, 0.0759]	2.5
0.10	0.0579	[0.0017, 0.0810]	0.0794	[−0.0047, 0.0875]	[−0.0812, 0.1158]	[−0.0880, 0.1160]	2.5
0.12	0.0697	[0.0032, 0.0959]	0.0926	[−0.0032, 0.1026]	[−0.0932, 0.1575]	[−0.0969, 0.1578]	2.7

0.15–0.25 at $|\alpha| = 0.01$, where the mean interval width is 4×10^{-4} , to 1.00 at $|\alpha| = 0.12$, where the mean width is 0.045; inclusion is lowest where the population interval is narrowest.

Table 5: Continuous-instrument design ($n = 5,000$): inclusion rate of the true θ_α in the estimated IVOT interval over 20 draws. Mean width averages the two signs of α .

$ \alpha $	0.01	0.02	0.04	0.06	0.08	0.10	0.12
Inclusion rate ($\alpha < 0$)	0.15	0.35	0.70	0.80	0.90	0.95	1.00
Inclusion rate ($\alpha > 0$)	0.25	0.45	0.70	0.80	0.85	0.90	1.00
Mean IVOT width	0.0004	0.0015	0.0058	0.0125	0.0215	0.0325	0.0454

Table 6 reports the covariate-design identified sets for different values of α ; the full grid is plotted in Figure 4. The estimated IVOT interval is 0.79–1.06 times as wide as the population sharp set and contains the truth at all 11 grid points. Its width is within 6% of the sharp width at nine points. At the two leftmost shifts, the estimated policy level merges with a neighboring baseline level because the design is close to the separation boundary, producing the narrower intervals marked in the table. The baseline level sets are recovered throughout.

Against IVMTE, the estimated IVOT interval is 1.10–1.46 times narrower at degree 3 and 1.34–1.78 times narrower at degree 7. The difference is qualitative: the sharp set excludes zero at every α on this window, and so does the estimated IVOT interval, whereas the degree-7 IVMTE interval and its 95% confidence band contain zero throughout, so the moment relaxation does not recover the sign of the policy effect.

Monte Carlo study. We draw 400 independent samples at each of $n = 1,000, 2,000, 4,000, 8,000$ and 16,000, and apply the procedure of Algorithm 1 at $\alpha = 0.22$, with the nuisance functions fitted by gradient-boosted trees. We report the $W = 1$ component studied in Theorem 4.7; Table 7 collects the results.

The scaled Monte Carlo standard deviation $\sqrt{n}$ SD lies in the range 0.30–0.38 from $n = 2,000$ on. The estimated standard errors approach the Monte Carlo standard deviations over the displayed sample sizes: the ratio of the mean estimated standard error to the Monte Carlo standard deviation rises from 0.78 to 1.00 for the lower bound and from 0.82 to 1.02 for the upper. The level-set recovery step succeeds in every replication once $n \geq 8,000$, having failed in 29% of them at $n = 1,000$; this is the finite-sample cost of ?? 4.2.

Table 6: Covariate design ($n = 4,000$): population sharp set and estimated IVOT and IVMTE bounds for θ_α . The population quantities are analytic up to the expectation over X .

A $\ddagger$ marks policy-level merging near the separation boundary. Width ratios compare degree-3 IVMTE with IVOT using unrounded bounds.

α	Truth	Sharp set	IVOT $[\ell, u]$	IVOT width	IVMTE deg. 3	IVMTE deg. 7	Width ratio
0.195 $\ddagger$	0.0830	[0.0207, 0.1172]	[0.0277, 0.1076]	0.0799	[0.0175, 0.1279]	[-0.0033, 0.1320]	1.38
0.207	0.0883	[0.0199, 0.1259]	[0.0176, 0.1298]	0.1122	[0.0153, 0.1386]	[-0.0067, 0.1438]	1.10
0.219	0.0938	[0.0192, 0.1345]	[0.0150, 0.1370]	0.1220	[0.0128, 0.1497]	[-0.0118, 0.1563]	1.12
0.231	0.0993	[0.0186, 0.1429]	[0.0181, 0.1443]	0.1262	[0.0102, 0.1607]	[-0.0160, 0.1685]	1.19
0.243	0.1048	[0.0182, 0.1513]	[0.0174, 0.1516]	0.1343	[0.0073, 0.1727]	[-0.0215, 0.1810]	1.23
0.255	0.1103	[0.0178, 0.1596]	[0.0161, 0.1634]	0.1473	[0.0043, 0.1836]	[-0.0261, 0.1932]	1.22

Table 7: Covariate design: Monte Carlo properties of the DML estimator for the $W = 1$ bound, with 400 replications per sample size and $\alpha = 0.22$. The population bounds are (0.1301, 0.2380). “SE/SD” compares the mean estimated standard error with the Monte Carlo standard deviation; “Sets” is the correct level-set recovery rate.

n	Lower bound			Upper bound			Sets
	$\sqrt{n}$ SD	SE/SD	Coverage	$\sqrt{n}$ SD	SE/SD	Coverage	
1,000	0.444	0.78	0.905	0.403	0.82	0.858	0.71
2,000	0.379	0.93	0.920	0.371	0.90	0.915	0.89
4,000	0.384	0.91	0.915	0.354	0.92	0.915	0.98
8,000	0.334	0.99	0.933	0.312	1.01	0.923	1.00
16,000	0.323	1.00	0.930	0.304	1.02	0.928	1.00

O6.3 Bed Net Application: Data and Setup

Dataset. We use data from Dupas [2014], which studies insecticide-treated bed net (ITN) take-up among Kenyan households. The instrument is the offered price Z taking 17 distinct values spanning 0–250 Kenyan shillings. The treatment $W \in \{0, 1\}$ indicates ITN purchase and the outcome $Y \in \{0, 1\}$ indicates ITN usage at a one-year follow-up. The estimation sample consists of $n = 1078$ observations after merging the purchase and follow-up datasets. Follow-up usage was not collected for households that did not purchase a net; we set $Y = 0$ for these households, since a household that acquired no net through the program cannot use a study net. Purchasers with a missing follow-up response are dropped. Coding non-purchasers as $Y = 0$ makes the outcome use of a study net, so that $Y(0) = 0$ holds by construction: a household that buys no net through the program has no study net to use, whatever it does about malaria protection otherwise. Dropping purchasers without a follow-up response requires follow-up to be missing at random among purchasers, conditional on the offered price:

$$\mathbb{P}(R = 1 \mid Y(1), Z, W = 1) = \mathbb{P}(R = 1 \mid Z, W = 1),$$

where R indicates an observed follow-up. The reported intervals condition on this assumption, which we do not test: attrition that varies with usage would differ across price arms, since purchase depends on the instrument.

Propensity score estimation. We estimate the propensity score by logistic regression of W on the offered price Z . The fitted propensity decreases from 0.854 at zero price to 0.232 at the reference price of 150 KSh. Relative to cell-by-cell estimation across the 17 price levels, the logistic fit smooths the first stage and imposes monotonicity through the fitted price coefficient. The replication package provides the fitted model and diagnostics.

Policy and target. Both policies compared here are fixed-price policies. The baseline policy assigns every household the reference price $z_0 = 150$ KSh, giving the constant propensity $\hat{p}(z_0) \approx 0.232$; the alternative policy q_α raises that constant to $q_\alpha = \min(\hat{p}(z_0) + \alpha, 1)$ for $\alpha \in [0.05, \alpha_{\max}]$. The target is accordingly

$$\text{PRTE}_\alpha = \frac{\mathbb{E}[Y^{q_\alpha} - Y^{p(z_0)}]}{\alpha},$$

the average per-unit effect of raising the purchase probability by α above its level at the reference price. Because both policies are constant in Z , the mean policy shift $\mathbb{E}[q_\alpha - p(z_0)]$

equals α exactly. The baseline here is the fixed-price policy $p(z_0)$; the other experiments instead use the observed assignment $p(Z)$.

O6.4 Bed Net Application: Numerical Results

Table 8 reports the IVOT and IVMTE bounds for PRTE_α at selected subsidy levels, together with 95% confidence intervals for both methods, the IVMTE interval being paired with the degree-20 point bounds it was computed from. IVOT yields tighter bounds throughout, and its estimated lower bound is positive at every reported subsidy level. The IVOT 95% delta-method CI remains inconclusive for the three smallest reported shifts and rules out non-positive per-unit effects from $\alpha = 0.211$ onward. The IVMTE 95% backward CI is likewise inconclusive at small α . At large α (near the maximum feasible shift), the IVOT bounds nearly point-identify PRTE_α , whereas the IVMTE identified set and its 95% CI remain substantially wider. The near-point identification arrives in a periodic pattern: the bounds tighten to near-zero width and then widen slightly as additional complier groups enter the policy margin.

The narrow estimated IVOT identified sets do not translate one-for-one into narrower confidence intervals. IVOT relies on treated-outcome quantities estimated separately across the 17 price cells, several of which contain few purchasers, whereas IVMTE pools information through its moment restrictions and sieve; the resulting variability in the IVOT bounds remains substantial, and division by α magnifies it at small policy shifts.

Table 8: Bed net application: IVOT and IVMTE bounds for PRTE_α at selected policy shifts. The IVMTE 95% backward interval corresponds to the degree-20 bounds. The IVOT 95% delta-method interval is truncated to the logical range $[-1, 1]$.

α	IVOT lower	IVOT upper	IVOT 95% CI	IVMTE deg. 10 $[\ell, u]$	IVMTE deg. 20 $[\ell, u]$	IVMTE deg. 20 95% CI
0.050	0.740	1.000	$[-1.000, 1.000]$	$[0.129, 0.990]$	$[-0.372, 1.002]$	$[-0.636, 1.002]$
0.094	0.774	0.774	$[-0.033, 1.000]$	$[0.167, 0.949]$	$[-0.219, 1.000]$	$[-0.380, 1.000]$
0.153	0.703	0.800	$[-0.084, 1.000]$	$[0.262, 0.830]$	$[0.116, 0.975]$	$[-0.164, 1.000]$
0.211	0.580	0.601	$[0.158, 1.000]$	$[0.406, 0.703]$	$[0.357, 0.824]$	$[-0.050, 1.000]$
0.299	0.588	0.637	$[0.290, 0.909]$	$[0.575, 0.674]$	$[0.524, 0.761]$	$[0.183, 1.000]$
0.402	0.622	0.640	$[0.306, 0.960]$	$[0.615, 0.753]$	$[0.571, 0.820]$	$[0.393, 0.987]$
0.504	0.603	0.624	$[0.410, 0.808]$	$[0.656, 0.784]$	$[0.625, 0.847]$	$[0.518, 0.933]$
0.621	0.597	0.597	$[0.457, 0.737]$	$[0.651, 0.782]$	$[0.627, 0.796]$	$[0.501, 0.861]$