SIMULTANEOUS INFERENCE FOR PROXIMAL CAUSAL LEARNING WITH WEAK PROXIES

Anonymous authors
Paper under double-blind review

ABSTRACT

Proximal causal learning uses the treatment-inducing proxy and the outcome-inducing proxy through the outcome bridge and the treatment bridge to provide a way to address unobserved confounding in observational studies. This paper studies simultaneous inference for the marginal counterfactual distribution and quantile treatment effect under weak proxies. First, we propose a measure of weak-proxy strength and tightens the population-bias bound for the counterfactual CDF estimate. The baseline Cauchy–Schwarz bound introduces the inverse of weak-proxy strength twice, whereas the proposed bound introduces the inverse of weak-proxy strength only once and is sharp for the bridge estimators, thereby relaxing the restriction on the decay rate of weak-proxy strength. Second, we prove that, even when outcome thresholds vary continuously, the number of distinct estimated outcome bridges is at most one greater than the number of observations in the training sample, thereby establishing a uniform expansion across treatment arms and outcome thresholds. Finally, we prove that the critical value remains valid after bridge estimation, thereby obtaining simultaneous CDF bands, and constructs quantile treatment effect bands through confidence-band inversion. The proposed method provides deterministic band-width bounds and does not require counterfactual density estimation. In the numerical study, the baseline Cauchy–Schwarz bound overestimates the true population bias by a factor of approximately 50, whereas the proposed bound is very close to the true population bias. The quantile treatment effect bands constructed by the proposed method achieve coverage of 99.8%–100.0% while being only approximately 20% wider than the oracle 95% bands, compared with baseline coverage of 68.4%–89.2%.

1 INTRODUCTION

Proximal causal learning uses the indirect information provided by the treatment-inducing proxy and the outcome-inducing proxy to estimate causal effects under appropriate identification conditions by estimating the outcome bridge and the treatment bridge, providing a way to address confounding in observational studies where the unobserved latent variable cannot be directly observed (Cui et al., 2024; Bozkurt et al., 2026). This paper focuses on the marginal counterfactual distribution and quantile treatment effect to characterize the effects of treatment at different positions of the outcome distribution, rather than being limited to mean effects (Chernozhukov et al., 2020).

Moving from bridge estimation to simultaneous inference for the entire marginal counterfactual distribution requires addressing three challenges:

First, weak proxies amplify bridge estimation errors (Bennett et al., 2023); directly multiplying the two bridge estimation error bounds to obtain a Cauchy–Schwarz bound overestimates the population bias of the counterfactual CDF estimate, thereby imposing overly stringent requirements on weak-proxy strength. *Second*, the treatment bridge is the same across all outcome thresholds, whereas the outcome bridge varies with outcome thresholds, requiring control of errors other than population bias across treatment arms and outcome thresholds (Tchetgen Tchetgen et al., 2024). *Third*, commonly used methods for constructing confidence bands for the quantile treatment effect require estimating the counterfactual density (Chernozhukov et al., 2013), and may be unreliable when the counterfactual density is close to zero or does not satisfy smoothness conditions (Firpo, 2007).

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

First, this paper complements existing proximal doubly robust analyses (Cui et al., 2024; Bennett et al., 2023) by retaining the stabilizing effect of ridge and ridge bias in the bridge estimation error bounds (Proposition 3.1) and tightening the population-bias bound (Corollary 3.1). The population-bias bound in this paper is sharp and cannot generally be further reduced. Consequently, weak-proxy strength is only required to decay more slowly than the inverse square root of the sample size, a condition of the same order as that required by the bridge estimation error bounds themselves.

Second, this paper complements the existing proximal inference literature (Ying et al., 2022; Wu et al., 2024) by proving that, even when outcome thresholds vary continuously, the number of distinct estimated outcome bridges is at most one greater than the number of observations in the training sample (Lemma 3.1). Consequently, the conditional entropy of the estimated outcome bridge class is at most of the same order as the logarithm of the number of observations in the training sample, so the complexity cost introduced by estimated outcome bridges is controlled. This paper establishes a uniform expansion (Theorem 3.1), showing that, under the stated conditions, bridge estimation does not alter the leading sampling behavior of the marginal counterfactual distribution.

Finally, this paper establishes simultaneous coverage of simultaneous CDF bands across treatment arms and outcome thresholds (Theorem 3.2). The contribution of this paper lies in ensuring the validity of the underlying simultaneous CDF bands under weak proxies and estimated bridges, thereby extending existing proximal inference that primarily focuses on mean effects and counterfactual survival curves (Cui et al., 2024; Ying et al., 2022) to simultaneous inference for the quantile treatment effect. The resulting quantile treatment effect bands admit deterministic band-width bounds and require neither counterfactual density estimation nor conditions on the counterfactual density.

The numerical results support the three contributions. In Figure 1, the true population bias is 11.00, the Bound underlying Corollary 3.1 is 20.97, and the baseline Cauchy–Schwarz bound is 571.75; the baseline overestimates it by a factor of approximately 50. In Figure 2, as the sample size increases, the median of the errors other than population bias decreases progressively in every setting of proxies, falling by approximately half in the strongest proxies setting considered. Across the 12 settings in Table 1, the simultaneous coverage of Proposed is 99.8%–100.0%, whereas that of Estimated-density delta is 68.4%–89.2%, below the nominal 95% in all settings; the mean band width of Proposed is only 0.06–0.27 wider than that of Oracle 95% and decreases as the sample size increases.

2 Background and Notation

Consider an observational study with n independent units. Let P_0 denote the true distribution of the observed data. For a generic unit, write $O = (Y, A, Z, W, X) \sim P_0$, and let $O_1, \ldots, O_n \overset{\text{iid}}{\sim} P_0$. Here, $Y \in \mathbb{R}$ is the outcome, $A \in \{0, 1\}$ is a binary observed treatment indicator, and X collects observed covariates. We allow treatment assignment and the outcome to be confounded by an unobserved latent variable U. Both Z and W serve as proxies that indirectly capture information about U. Specifically, Z is a treatment-inducing proxy and W is an outcome-inducing proxy.

For each treatment arm $a \in \{0, 1\}$, let $Y(a)$ denote the potential outcome under treatment a. For the unobserved latent variable U, we assume $Y(a) \perp A \mid U, X$, and $Y \perp Z \mid U, A, X$, and $W \perp (A, Z) \mid U, X$. We also assume latent positivity $0 < P(A = a \mid U, X) < 1$.

Our object is the marginal counterfactual distribution $F_a(y) = P\{Y(a) \le y\}$, from which the counterfactual quantile function and the quantile treatment effect are $Q_a(\tau) = \inf\{y \in \mathbb{R} : F_a(y) \ge \tau\}$ and $\Delta_Q(\tau) = Q_1(\tau) - Q_0(\tau)$.

2.1 Proximal Identification

The outcome bridge $h_{a,y}(W, X)$. Because U is unobserved, to identify $F_a(y)$ from the observed data, we introduce an outcome bridge $h_{a,y}(W, X)$. For each $a \in \{0, 1\}$ and $y \in \mathbb{R}$, suppose there exists a measurable function $h_{a,y}(W, X)$ satisfying

$$E\big[\mathbf{1}\{Y \le y\} - h_{a,y}(W, X) \mid A = a, Z, X\big] = 0. \tag{1}$$

Under these assumptions, the marginal counterfactual distribution $F_a(y)$ is identified as

$$F_a(y) = P\{Y(a) \le y\} = E\{h_{a,y}(W, X)\}. \tag{2}$$

Thus, the outcome bridge $h_{a,y}(W, X)$ uses the outcome-inducing proxy W to recover the outcome information that would otherwise be unavailable because U is unobserved (Xu et al., 2021).

The treatment bridge $q_a(Z, X)$. For each $a \in \{0, 1\}$, a treatment bridge is a measurable function $q_a(Z, X)$ satisfying $E\{\mathbf{1}\{A = a\}q_a^2(Z, X)\} < \infty$ and

$$E\big[\mathbf{1}\{A = a\}q_a(Z, X) - 1 \mid W, X\big] = 0. \tag{3}$$

When a treatment bridge $q_a(Z, X)$ exists, the marginal counterfactual distribution $F_a(y)$ also satisfies

$$F_a(y) = E\big[\mathbf{1}\{A = a\}q_a(Z, X)\mathbf{1}\{Y \le y\}\big].$$

The same treatment bridge $q_a(Z, X)$ can be used for every threshold y. Combining the equations for the outcome bridge $h_{a,y}(W, X)$ and the treatment bridge $q_a(Z, X)$ gives

$$F_a(y) = E\big[h_{a,y}(W, X) + \mathbf{1}\{A = a\}q_a(Z, X) \times \big\{\mathbf{1}\{Y \le y\} - h_{a,y}(W, X)\big\}\big]. \tag{4}$$

Thus, proximal identification replaces direct adjustment for the unobserved latent variable U with the outcome bridge $h_{a,y}(W, X)$ and the treatment bridge $q_a(Z, X)$ using observed Z and W.

2.2 Related Work and Challenges

Existing studies have shown that, under appropriate conditions, observed proxies can be used to identify causal effects in the presence of an unobserved latent variable U, and that outcome bridges and treatment bridges can be used for estimation (Miao et al., 2018; Cui et al., 2024). Subsequent work pointed out that identification of causal effects does not necessarily imply that these causal effects can be estimated stably, and that the existence of the treatment bridge is closely related to such estimation stability (Zhang et al., 2023). Related studies have also shown when bridge equations are difficult to solve, regularization and bridge estimation error can directly affect inference for the final target (Bennett et al., 2023). Meanwhile, existing proximal causal inference methods have studied estimation and inference for entire counterfactual survival curves (Ying et al., 2022).

Building on the above literature, we study three challenges in simultaneous inference when the outcome bridge $h_{a,y}(W, X)$ and treatment bridge $q_a(Z, X)$ are estimated from data:

Challenge 1: Bridge estimation errors under weak proxies. Weak proxies Z and W amplify sampling errors in $\widehat{h}_{a,y}(W, X)$ and $\widehat{q}_a(Z, X)$ (Kompa et al., 2022). Multiplying separate bridge error bounds can overstate the population bias of $\widehat{F}_a(y)$ because both bridges use the same sample matrix. The challenge is to *quantify* the bridge estimation errors and determine when using $\widehat{h}_{a,y}(W, X)$ and $\widehat{q}_a(Z, X)$ introduces only negligible population bias into the counterfactual CDF estimate $\widehat{F}_a(y)$.

Challenge 2: Keeping bridge estimation errors small across treatment arms and outcome thresholds. The counterfactual CDF estimate $\widehat{F}_a(y)$ combines the estimated outcome bridge $\widehat{h}_{a,y}(W, X)$ and the estimated treatment bridge $\widehat{q}_a(Z, X)$ (see equation (4)). $\widehat{q}_a(Z, X)$ is the same for all outcome thresholds y, whereas $\widehat{h}_{a,y}(W, X)$ changes with y. Bridge estimation can create additional sampling variation that must be controlled over the entire outcome threshold y (cf. Byambadalai et al., 2024). The challenge is therefore to *understand* how the estimated bridges vary with y and $a \in \{0, 1\}$, while controlling both the bridge estimation errors and how large the estimated bridges can become.

Challenge 3: Reliable confidence bands with estimated bridges. Errors in the estimated bridges can affect both the CDF estimate $\widehat{F}_a(y)$ and the width of its confidence band. The challenge is to show that the confidence bands still have the desired coverage simultaneously for both treatment arms $a \in \{0, 1\}$ and all outcome thresholds y. The usual method for the counterfactual quantile $Q_a(\tau)$ requires additional smoothness and positivity conditions on the counterfactual density near $Q_a(\tau)$ (Chernozhukov et al., 2013). The challenge is to *construct* simultaneous confidence bands for the quantile treatment effect $\Delta_Q(\tau)$ without introducing density estimation or density conditions.

3 Proximal Estimation and Simultaneous Inference

To address Challenge 1, we quantify weak-proxy strength in the chosen basis spaces by $\kappa_{a,n}$. Proposition 3.1 bounds the estimation errors of the outcome bridge $h_{a,y}(W, X)$ and treatment bridge $q_a(Z, X)$. Using the bridge moment equations rather than multiplying separate bridge error bounds, Corollary 3.1 establishes the sharper population-bias bound uniformly over $y \in \mathcal{Y}$ and gives explicit sufficient conditions for negligible population bias at the $n^{-1/2}$ scale. Corollary 3.1 further shows that the sharpness of $\kappa_{a,n}^{-1}$. We complement existing proximal doubly robust (Cui et al., 2024; Bennett et al., 2023) by quantifying the population bias of $\widehat{F}_a(y)$ across outcome thresholds y and establishing the sharp $\kappa_{a,n}^{-1}$ dependence for the specified bridge estimators.

To address Challenge 2, we exploit the shared estimating structure of the outcome bridges across outcome thresholds. For a training sample of size N_k, we prove that the estimated outcome bridges across outcome thresholds y contains at most $N_k + 1$ distinct functions in Lemma 3.1. Theorem 3.1 then shows that estimating the bridges does not change the leading sampling behavior of $\widehat{F}_a(y)$. Compared with existing proximal inference literature (Ying et al., 2022), our result gives an explicit control of the estimated bridge across outcome thresholds y for $\widehat{h}_{a,y}(W, X)$ and $\widehat{q}_a(Z, X)$.

To address Challenge 3, Theorem 3.2 establishes CDF coverage converging to $1-\alpha_{\text{sig}}$ simultaneously over both treatment arms $a \in \{0, 1\}$ and all outcome thresholds y. Existing proximal inference has focused primarily on mean effects and counterfactual survival curves (Cui et al., 2024; Ying et al., 2022). We also extend this line of work to simultaneous inference for the quantile treatment effect $\Delta_Q(\tau)$ in Theorem 3.2 and derive deterministic band-width bounds for it. Unlike standard method for $\Delta_Q(\tau)$ (Chernozhukov et al., 2013), require neither estimation of the counterfactual density nor additional density continuity conditions.

Estimating $F_a(y)$ over an entire outcome threshold range requires estimating the outcome bridge $h_{a,y}(W, X)$ for each treatment arm $a \in \{0, 1\}$, every outcome threshold y and estimating the treatment bridge $q_a(Z, X)$ for each treatment arm $a \in \{0, 1\}$. Under weak proxies W and Z, directly solving sample versions of (1) and (3) can amplify sampling errors (Kallus et al., 2021).

Let $\mathcal{Y} = [y_L, y_U]$ be a fixed compact threshold interval. Independently of $O_1, \ldots, O_n$, partition $\{1, \ldots, n\}$ into a fixed number $K \geq 2$ of disjoint folds $I_1, \ldots, I_K$, with $|I_k| \asymp n$. Define $N_k = n - |I_k|$ and $P_{n,-k} f = N_k^{-1} \sum_{i \notin I_k} f(O_i)$. Repeating this procedure over $k = 1, \ldots, K$ yields cross-fitted bridge evaluations for all observations $O_1, \ldots, O_n$, with each $O_i \in I_k$ evaluated using $\widehat{h}_{a,y}^{(-k)}$ and $\widehat{q}_a^{(-k)}$ estimated without the observations in I_k.

3.1 Regularized Bridge Estimation and Weak-Proxy Strength

Finite-dimensional bridge approximation. The bridge equations (1) and (3) characterize the unknown functions $h_{a,y}(W, X)$ and $q_a(Z, X)$ but do not directly provide estimators for them. We therefore approximate the two bridges $h_{a,y}(W, X)$ and $q_a(Z, X)$ by $b_W(W, X)^\top \theta$ and $b_Z(Z, X)^\top \alpha$, where the prespecified basis vectors $b_W(W, X), b_Z(Z, X) \in \mathbb{R}^{d_n}$ may have dimension d_n increasing with n. $b_W(W, X)$ and $b_Z(Z, X)$ can be constructed by evaluating prespecified polynomial or including interaction terms, at (W, X) and (Z, X), respectively (Tchetgen Tchetgen et al., 2024). Estimating the bridges $h_{a,y}(W, X)$ and $q_a(Z, X)$ then reduces to estimating the coefficient vectors θ and α.

Substituting $h_{a,y}(W, X) \approx b_W(W, X)^\top \theta$ and $q_a(Z, X) \approx b_Z(Z, X)^\top \alpha$ into (1) and (3) gives the sample moment conditions for the outcome and treatment bridges,

$$P_{n,-k}\big[\mathbf{1}\{A = a\} b_Z(Z, X)\{\mathbf{1}\{Y \leq y\} - b_W(W, X)^\top \theta\}\big] = 0, \tag{5}$$

$$P_{n,-k}\big[b_W(W, X)\{\mathbf{1}\{A = a\} b_Z(Z, X)^\top \alpha - 1\}\big] = 0. \tag{6}$$

Define

$$\widehat{\Sigma}_a^{(-k)} = P_{n,-k}\{\mathbf{1}\{A = a\} b_Z(Z, X) b_W(W, X)^\top\}, \qquad \widehat{\gamma}_a^{(-k)}(y) = P_{n,-k}\{\mathbf{1}\{A = a\} b_Z(Z, X)\mathbf{1}\{Y \leq y\}\},$$

$$\widehat{\mu}_W^{(-k)} = P_{n,-k}\{b_W(W, X)\}, \qquad (\widehat{\Sigma}_a^{(-k)})^\top = P_{n,-k}\{\mathbf{1}\{A = a\} b_W(W, X) b_Z(Z, X)^\top\}.$$

Expanding (5) and (6) then gives the *outcome and treatment bridge moment conditions*

$$\widehat{\Sigma}_a^{(-k)} \theta = \widehat{\gamma}_a^{(-k)}(y), \qquad (\widehat{\Sigma}_a^{(-k)})^\top \alpha = \widehat{\mu}_W^{(-k)}. \tag{7}$$

For treatment arm $a \in \{0, 1\}$, all outcome thresholds y share $\widehat{\Sigma}_a^{(-k)}$, only $\widehat{\gamma}_a^{(-k)}(y)$ changes with y.

Regularization under weak proxies. When $\widehat{\Sigma}_a^{(-k)}$ is nearly singular, directly solving $\widehat{\Sigma}_a^{(-k)} \theta = \widehat{\gamma}_a^{(-k)}(y)$ and $(\widehat{\Sigma}_a^{(-k)})^\top \alpha = \widehat{\mu}_W^{(-k)}$ can make the coefficient estimates highly sensitive to sampling errors. We therefore choose θ and α so that $\widehat{\Sigma}_a^{(-k)} \theta$ is close to $\widehat{\gamma}_a^{(-k)}(y)$ and $(\widehat{\Sigma}_a^{(-k)})^\top \alpha$ is close to $\widehat{\mu}_W^{(-k)}$, while preventing θ and α from becoming excessively large (Carrasco et al., 2007).

For $\lambda_h, \lambda_q > 0$, which may depend on n but are fixed across outcome thresholds y, define

$$\begin{aligned} \widehat{\theta}_{a,\lambda_h}^{(-k)}(y) &= \arg\min_{\theta \in \mathbb{R}^{d_n}} \left\{ \left\| \widehat{\Sigma}_a^{(-k)} \theta - \widehat{\gamma}_a^{(-k)}(y) \right\|_2^2 + \lambda_h \|\theta\|_2^2 \right\}, \\ \widehat{\alpha}_{a,\lambda_q}^{(-k)} &= \arg\min_{\alpha \in \mathbb{R}^{d_n}} \left\{ \left\| (\widehat{\Sigma}_a^{(-k)})^\top \alpha - \widehat{\mu}_W^{(-k)} \right\|_2^2 + \lambda_q \|\alpha\|_2^2 \right\}. \end{aligned} \tag{8}$$

The optimization problems are convex quadratic problems with unique solutions.

The estimated bridges are therefore

$$\widehat{h}_{a,y}^{(-k)}(W,X) = b_W(W,X)^\top \widehat{\theta}_{a,\lambda_h}^{(-k)}(y), \quad \widehat{q}_a^{(-k)}(Z,X) = b_Z(Z,X)^\top \widehat{\alpha}_{a,\lambda_q}^{(-k)}. \tag{9}$$

Quantifying weak-proxy strength. We now quantify how sampling errors can affect estimates of θ and α. Replacing the training-sample average $P_{n,-k}$ in $\widehat{\Sigma}_a^{(-k)}$ by expectation E under P_0, define

$$\Sigma_a = E\{\mathbf{1}\{A=a\} b_Z(Z,X) b_W(W,X)^\top\}. \tag{10}$$

We measure weak-proxy strength in the chosen basis spaces by

$$\kappa_{a,n} = \sigma_{\min}(\Sigma_a) > 0, \tag{11}$$

where $\sigma_{\min}(\Sigma_a)$ denotes the smallest singular value of Σ_a. A small $\kappa_{a,n}$ means that substantially different values of θ can produce nearly identical values of $\Sigma_a\theta$, making θ difficult to distinguish using noisy sample moments. Because $\Sigma_a^\top$ has the same smallest singular value as Σ_a, a small $\kappa_{a,n}$ can also make substantially different values of α difficult to distinguish through $\Sigma_a^\top \alpha$. Proposition 3.1 quantifies how $\kappa_{a,n}$ affects the estimation error bounds for $h_{a,y}(W,X)$ and $q_a(Z,X)$.

Proposition 3.1. *Under Assumption A.1. For every $\lambda_h, \lambda_q > 0$,*

$$\begin{aligned} r_{h,a,n} &= \max_k \sup_{y\in\mathcal{Y}} \left\| \widehat{h}_{a,y}^{(-k)} - h_{a,y} \right\|_W = O_p\left(\frac{1+M_{h,a,n}}{\sqrt{\kappa_{a,n}^2+\lambda_h}} \sqrt{\frac{d_n}{n}} + b_{h,a,n} + \frac{\lambda_h M_{h,a,n}}{\kappa_{a,n}^2+\lambda_h} \right), \\ r_{q,a,n} &= \max_k \left\| \widehat{q}_a^{(-k)} - q_a \right\|_{a,Z} = O_p\left(\frac{1+M_{q,a,n}}{\sqrt{\kappa_{a,n}^2+\lambda_q}} \sqrt{\frac{d_n}{n}} + b_{q,a,n} + \frac{\lambda_q M_{q,a,n}}{\kappa_{a,n}^2+\lambda_q} \right). \end{aligned} \tag{12}$$

The norms $M_{h,a,n}$, $M_{q,a,n}$, approximation errors, and coefficient norm bounds are defined in Appendix A.3. Appendix A.4 gives the exact factors involving the sample singular values.

3.2 Controlling $\widehat{F}_a(y) - F_a(y)$ simultaneously Across treatment arms and Outcome Thresholds

Motivated by the $F_a(y)$ in (4), we replace $h_{a,y}(W,X)$ and $q_a(Z,X)$ by their estimators. Define

$$\widehat{F}_a(y) = \frac{1}{n}\sum_{k=1}^K \sum_{i\in I_k} \left[\widehat{h}_{a,y}^{(-k)}(W_i,X_i) + \mathbf{1}\{A_i=a\}\widehat{q}_a^{(-k)}(Z_i,X_i)\{\mathbf{1}\{Y_i\le y\} - \widehat{h}_{a,y}^{(-k)}(W_i,X_i)\}\right]. \tag{13}$$

How bridge estimation errors enter $\widehat{F}_a(y)$. For each fold I_k, expectations containing estimated bridges integrate over $O\sim P_0$ with the training sample outside I_k held fixed. The outcome bridge equation (1), the treatment bridge equation (3), and (13) give

$$\begin{aligned} &E\left[\widehat{h}_{a,y}^{(-k)}(W,X) + \mathbf{1}\{A=a\}\widehat{q}_a^{(-k)}(Z,X)\times\{\mathbf{1}\{Y\le y\} - \widehat{h}_{a,y}^{(-k)}(W,X)\}\right] - F_a(y) \\ &\quad = -E\left[\mathbf{1}\{A=a\}\{\widehat{q}_a^{(-k)}(Z,X) - q_a(Z,X)\}\{\widehat{h}_{a,y}^{(-k)}(W,X) - h_{a,y}(W,X)\}\right]. \end{aligned} \tag{14}$$

Thus the bias of (14) depends on the product of the two bridge estimation errors rather than on either bridge estimation error alone (Bozkurt et al., 2025). Recall Proposition 3.1, by Cauchy–Schwarz, the absolute population bias in (14) is at most $r_{h,a,n}r_{q,a,n}$. The Cauchy–Schwarz bound

$$\max_k \sup_{y\in\mathcal{Y}} \left| E\left[\mathbf{1}\{A=a\}\{\widehat{q}_a^{(-k)} - q_a\}\{\widehat{h}_{a,y}^{(-k)} - h_{a,y}\}\right]\right| \le r_{h,a,n}r_{q,a,n}. \tag{15}$$

incurs the $\kappa_{a,n}^{-1}$ sampling-error factor twice.

Corollary 3.1. *Under the conditions of Proposition 3.1, suppose $M_{h,a,n}, M_{q,a,n} = O(1)$ and*

$$\lambda_h M_{h,a,n} + \lambda_q M_{q,a,n} = o\left(\kappa_{a,n}\sqrt{\frac{d_n}{n}}\right)$$

for each treatment arm $a\in\{0,1\}$. Then, for each treatment arm $a\in\{0,1\}$,

$$\max_k \sup_{y\in\mathcal{Y}} \left| E\left[\mathbf{1}\{A=a\}\{\widehat{q}_a^{(-k)}(Z,X) - q_a(Z,X)\}\{\widehat{h}_{a,y}^{(-k)}(W,X) - h_{a,y}(W,X)\}\right]\right| = O_p\left(\frac{d_n}{n\kappa_{a,n}} + b_{h,a,n}b_{q,a,n}\right). \tag{16}$$

Consequently,

$$\max_{a\in\{0,1\}} \left\{ \frac{d_n}{\kappa_{a,n}\sqrt{n}} + \sqrt{n}\, b_{h,a,n}b_{q,a,n} \right\} \to 0$$

is sufficient for the population bias in (14) to be negligible at the $n^{-1/2}$ scale.

Even in the ideal case where the true outcome bridge and treatment bridge can be exactly represented by the finite-dimensional bases $b_W(W,X)$ and $b_Z(Z,X)$, the population bias can still satisfy

$$\max_k \sup_{y\in\mathcal{Y}} \left| E\Big[\mathbf{1}\{A=1\}\{\widehat{q}_1^{(-k)}(Z,X)-q_1(Z,X)\}\{\widehat{h}_{1,y}^{(-k)}(W,X)-h_{1,y}(W,X)\}\Big]\right| \neq o_p\left(\frac{1}{n\kappa_{1,n}}\right). \tag{17}$$

Appendix B.1 gives the proof. Compared with the Cauchy–Schwarz bound, whose leading sampling term scales as $O_p\{d_n/(n\kappa_{a,n}^2)\}$, Corollary 3.1 improves this dependence to $O_p\{d_n/(n\kappa_{a,n}) + b_{h,a,n}b_{q,a,n}\}$, and (17) shows that the remaining $\kappa_{a,n}^{-1}$ is sharp, and cannot generally be improved in order for the specified bridge estimators.

Controlling bridge estimation errors across outcome thresholds y. Corollary 3.1 controls population bias of $\widehat{F}_a(y)$, but simultaneous inference also requires control of the sampling fluctuation caused by estimated bridges across $y\in\mathcal{Y}$. For each treatment arm $a\in\{0,1\}$, the treatment bridge estimate $\widehat{q}_a^{(-k)}(Z,X)$ is fixed across outcome thresholds y. The outcome bridge estimate $\widehat{h}_{a,y}^{(-k)}(W,X)$ varies with y, so we examine how many distinct functions occur in $\{\widehat{h}_{a,y}^{(-k)}(W,X): y\in\mathcal{Y}\}$.

Lemma 3.1. *Fix a treatment arm $a\in\{0,1\}$ and a fold I_k. Conditionally on the training sample $\{O_i : i\notin I_k\}$, the class*

$$\{\widehat{h}_{a,y}^{(-k)}(W,X): y\in\mathcal{Y}\}$$

contains at most N_k+1 distinct functions.

Appendix A.5 gives the proof. Thus, although y ranges over a continuum, the estimated outcome bridge class has conditional entropy of order $\log N_k$, providing an explicit simple complexity control for simultaneous inference with estimated bridges below.

Control of $\widehat{F}_a(y)-F_a(y)$ across both treatment arms and all outcome thresholds. To describe the estimation error that would remain if true $h_{a,y}(W,X)$ and $q_a(Z,X)$ were known, define

$$\varphi_{a,y}(O)=h_{a,y}(W,X)-F_a(y)+\mathbf{1}\{A=a\}q_a(Z,X)\{\mathbf{1}\{Y\le y\}-h_{a,y}(W,X)\}. \tag{18}$$

Equation (4) gives $E\{\varphi_{a,y}(O)\}=0$.

Theorem 3.1. *Assume Assumptions B.1 and B.2 hold. Corollary 3.1 gives sufficient population-bias conditions, and Lemma 3.1 verifies the estimated outcome bridge covering bound, these results give*

$$\max_{a\in\{0,1\}}\sup_{y\in\mathcal{Y}}\left|\sqrt{n}\{\widehat{F}_a(y)-F_a(y)\}-\frac{1}{\sqrt{n}}\sum_{i=1}^n\varphi_{a,y}(O_i)\right|=o_p(1). \tag{19}$$

Appendix B.4 gives the proof. Theorem 3.1 connects the bridge estimation error bounds in Proposition 3.1 to simultaneous inference over treatment arms $a\in\{0,1\}$ and outcome thresholds $y\in\mathcal{Y}$.

3.3 Simultaneous Confidence Bands for $F_a(y)$ and $\Delta_Q(\tau)$

Theorem 3.1 describes $\widehat{F}_a(y)-F_a(y)$ through $\varphi_{a,y}(O)$. Because $\varphi_{a,y}(O)$ contains unknown bridges, the critical value must remain valid after bridge estimation. The usual method for constructing confidence bands for the quantile treatment effect $\Delta_Q(\tau)$ requires estimating the counterfactual density at $Q_a(\tau)$ for both treatment arms (Chernozhukov et al., 2013). The counterfactual density is the density of the potential outcome $Y(a)$ and describes how quickly the marginal counterfactual distribution $F_a(y)$ increases with y. Estimating this density adds an additional estimation step, and can be unreliable when the counterfactual density near $Q_a(\tau)$ is close to zero or does not satisfy the required smoothness conditions (Firpo, 2007).

We instead construct simultaneous CDF bands using Theorem 3.1, validating the multiplier critical value after estimating the outcome bridge $h_{a,y}(W,X)$ and treatment bridge $q_a(Z,X)$ and replacing $F_a(y)$ by $\widehat{F}_a(y)$. Confidence-band inversion then gives simultaneous confidence bands for $\Delta_Q(\tau)$ and bounds the QTE band width by changes in the counterfactual quantile function $Q_a(\tau)$ without counterfactual density estimation.

Calculating the critical value from estimated bridges. To calculate the confidence-band width, we first estimate $\varphi_{a,y}(O_i)$ for each $i\in I_k$. In (18), replace $h_{a,y}(W_i,X_i)$, $q_a(Z_i,X_i)$, and $F_a(y)$ by $\widehat{h}_{a,y}^{(-k)}(W_i,X_i)$, $\widehat{q}_a^{(-k)}(Z_i,X_i)$, and $\widehat{F}_a(y)$, respectively, to define

$$\widehat{\varphi}_{a,y}^{(-k)}(O_i)=\widehat{h}_{a,y}^{(-k)}(W_i,X_i)-\widehat{F}_a(y)+\mathbf{1}\{A_i=a\}\widehat{q}_a^{(-k)}(Z_i,X_i)\{\mathbf{1}\{Y_i\le y\}-\widehat{h}_{a,y}^{(-k)}(W_i,X_i)\}. \tag{20}$$

We approximate the simultaneous distribution of $\sqrt{n}\{\widehat{F}_a(y) - F_a(y)\}$ over both treatment arms $a \in \{0,1\}$ and outcome thresholds $y \in \mathcal{Y}$ using a Rademacher multiplier bootstrap (van der Vaart and Wellner, 1996). Let $\xi_1, \ldots, \xi_n$ be independent Rademacher multipliers, with $P(\xi_i = 1) = P(\xi_i = -1) = 1/2$, and define

$$\widehat{\mathbb{G}}_a^{\xi}(y) = \frac{1}{\sqrt{n}} \sum_{k=1}^{K} \sum_{i \in I_k} \xi_i \widehat{\varphi}_{a,y}^{(-k)}(O_i). \tag{21}$$

We estimate $\widehat{c}_{1-\alpha_{\text{sig}}}$ by the conditional $(1-\alpha_{\text{sig}})$ quantile of

$$\max_{a \in \{0,1\}} \sup_{y \in \mathcal{Y}} |\widehat{\mathbb{G}}_a^{\xi}(y)| \tag{22}$$

over repeated multiplier draws.

Using $\widehat{c}_{1-\alpha_{\text{sig}}}/\sqrt{n}$ as the half-width around $\widehat{F}_a(y)$, define

$$L_a(y) = \max\{0, \widehat{F}_a(y) - \widehat{c}_{1-\alpha_{\text{sig}}}/\sqrt{n}\}, \quad U_a(y) = \min\{1, \widehat{F}_a(y) + \widehat{c}_{1-\alpha_{\text{sig}}}/\sqrt{n}\}. \tag{23}$$

Confidence bands for $\Delta_Q(\tau)$ without density estimation. Let $\mathcal{T} \subset (0,1)$ be a fixed compact set of quantile levels. We apply the confidence-band inversion method of Chernozhukov et al. (2020) to $[L_a(y), U_a(y)]$. The upper CDF bound $U_a(y)$ determines a lower quantile bound, and the lower CDF bound $L_a(y)$ determines an upper quantile bound:

$$Q_a^L(\tau) = \inf\{y \in \mathcal{Y} : U_a(y) \geq \tau\}, \quad Q_a^U(\tau) = \inf\{y \in \mathcal{Y} : L_a(y) \geq \tau\}, \tag{24}$$

If $L_a(y) \leq F_a(y) \leq U_a(y)$ for both treatment arms $a \in \{0,1\}$ and every outcome thresholds $y \in \mathcal{Y}$, then $Q_a^L(\tau) \leq Q_a(\tau) \leq Q_a^U(\tau)$ for every target quantile $Q_a(\tau)$ in the interior of $\mathcal{Y}$.

Since $\Delta_Q(\tau) = Q_1(\tau) - Q_0(\tau)$, define

$$\Delta_Q^L(\tau) = Q_1^L(\tau) - Q_0^U(\tau), \quad \Delta_Q^U(\tau) = Q_1^U(\tau) - Q_0^L(\tau). \tag{25}$$

Coverage of $F_0(y)$ and $F_1(y)$ over $y \in \mathcal{Y}$ therefore implies coverage of $\Delta_Q(\tau)$ simultaneously over $\tau \in \mathcal{T}$. Thus, the construction in (24) and (25) uses only $L_a(y)$ and $U_a(y)$, without estimating a counterfactual density or differentiating $Q_a(\tau)$.

Theorem 3.2. *Under the conditions of Theorem 3.1, suppose*

$$\max_{a \in \{0,1\}} \sup_{y \in \mathcal{Y}} E\{\varphi_{a,y}^2(O)\} > 0 \tag{26}$$

and $Q_a(\tau) \in (y_L, y_U)$ for every $a \in \{0,1\}$ and $\tau \in \mathcal{T}$. Then:

(i) The CDF bands satisfy

$$P_0\{L_a(y) \leq F_a(y) \leq U_a(y) \text{ for all } a \in \{0,1\},\ y \in \mathcal{Y}\} \longrightarrow 1 - \alpha_{\text{sig}}. \tag{27}$$

(ii) The QTE bands satisfy

$$\liminf_{n \to \infty} P_0\{\Delta_Q^L(\tau) \leq \Delta_Q(\tau) \leq \Delta_Q^U(\tau) \text{ for all } \tau \in \mathcal{T}\} \geq 1 - \alpha_{\text{sig}}. \tag{28}$$

(iii) On the simultaneous CDF-coverage event in (27), for both treatment arms,

$$\Delta_Q^U(\tau) - \Delta_Q^L(\tau) \leq \sum_{a=0}^{1} \left[Q_a\left(\tau + \frac{2\widehat{c}_{1-\alpha_{\text{sig}}}}{\sqrt{n}}\right) - Q_a\left(\tau - \frac{2\widehat{c}_{1-\alpha_{\text{sig}}}}{\sqrt{n}}\right) \right] \tag{29}$$

for every $\tau \in \mathcal{T}$ with $\tau \pm 2\widehat{c}_{1-\alpha_{\text{sig}}}/\sqrt{n} \in (0,1)$ and $Q_a(\tau \pm 2\widehat{c}_{1-\alpha_{\text{sig}}}/\sqrt{n}) \in (y_L, y_U)$.

Appendix C.6 gives the proof. Equation (28) establishes simultaneous coverage of $\Delta_Q(\tau)$, (29) bounds the QTE band width, without counterfactual density estimation.

4 Numerical Studies

We design three simulations to examine bridge estimation errors and population bias under weak proxies (*Challenge 1*), estimated-bridge sampling fluctuation and uniform expansion (*Challenge 2*), and simultaneous confidence bands for $\Delta_Q(\tau)$ without density estimation (*Challenge 3*), respectively.

In all three simulations, the observed covariates $X = (X_1, X_2)$ comprise two independent Bernoulli$(1/2)$ variables. The unobserved latent variable U, treatment-inducing proxy Z, outcome-inducing proxy W, and treatment indicator A are binary and generated from conditional Bernoulli distributions. The coefficient ρ controls Z and W dependence on U; the resulting $\kappa_{a,n}$ is calculated from the population matrix in (10) for the chosen basis. For each treatment arm $a \in \{0,1\}$, we generate the potential outcome $Y(a)$ by adding standardized noise. We use the prespecified basis

vectors $b_W(W, X)$ and $b_Z(Z, X)$ with $d_n = 8$ and $K = 3$ folds. We use 500 replications for each setting. See Appendix D for details.

Simulation 1: Bridge Estimation Errors and Population Bias under Weak Proxies.

We examine *Challenge 1* by connecting Proposition 3.1 with Corollary 3.1. Further details are provided in Appendix D.1. Figure 1 fixes $n = 3200$ and uses $\min_{a\in\{0,1\}} \kappa_{a,n}$ as the horizontal axis in all four panels. The first two panels plot $\max_{a\in\{0,1\}} r_{h,a,n}$ and $\max_{a\in\{0,1\}} r_{q,a,n}$, and the last two compare the $\sqrt{n}$-scaled population bias and its bounds without and with ridge. Population bias denotes the true absolute population bias in (14), maximized over both treatment arms, training folds, and outcome thresholds, and multiplied by $\sqrt{n}$. Cauchy–Schwarz bound denotes $\sqrt{n}\max_{a\in\{0,1\}} r_{h,a,n} r_{q,a,n}$ from (15). Bound underlying Corollary 3.1 denotes the right-hand side of (79), multiplied by $\sqrt{n}$. The additional bound in (85) is not plotted.

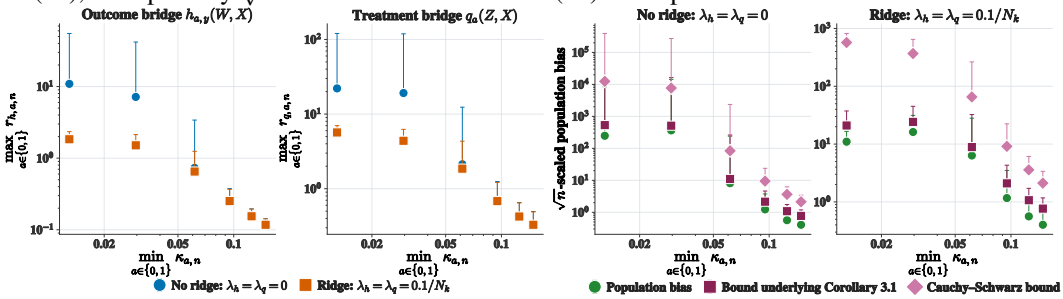

Figure 1: Bridge estimation errors and population bias under weak proxies at $n = 3200$. Markers show medians and upper bars show 90th percentiles among successful fits.

The first two panels of Figure 1 show that smaller $\min_{a\in\{0,1\}} \kappa_{a,n}$ is associated with larger bridge estimation errors, qualitatively consistent with Proposition 3.1. At the weakest-proxy setting $\rho = 0.2$, ridge in (8) reduces the median outcome- and treatment-bridge errors from approximately 10.90 to 1.84 (-83%) and from 22.09 to 5.68 (-74%), and their 90th percentiles from 55.12 to 2.35 and from 120.29 to 7.00. In the last two panels, the median population bias with ridge (11.00) stays close to the bound underlying Corollary 3.1 (20.97) but far below the Cauchy–Schwarz bound (571.75), supporting its sharper $\kappa_{a,n}^{-1}$ population-bias control; without ridge, the median population bias rises to 247.61, so ridge reduces it by about 96%.

Simulation 2: Estimated-Bridge Sampling Fluctuation and Uniform Expansion.

We examine *Challenge 2* by assessing Theorem 3.1, and Lemma 3.1. The estimated-bridge sampling fluctuation $V_n = \sqrt{n}\max_{a\in\{0,1\}} \sup_{y\in\mathcal{Y}} \left|\sum_{k=1}^{K} (n_k/n)(P_{n,k} - E) d_{a,y}^{(-k)}\right|$ measures the sampling-fluctuation component controlled in the proof of Theorem 3.1 (see (120)). The uniform expansion remainder $R_n = \max_{a\in\{0,1\}} \sup_{y\in\mathcal{Y}} \left|\sqrt{n}\{\widehat{F}_a(y) - F_a(y)\} - n^{-1/2}\sum_{i=1}^{n} \varphi_{a,y}(O_i)\right|$ measures the full remainder in Theorem 3.1, including both population bias and sampling fluctuation, and equals the left-hand side of (19). We use $n \in \{1600, 3200, 4800, 5400\}$ and $\rho \in \{0.75, 0.60, 0.45\}$. At every value of ρ, we compare Walsh, ReLU, and GB-leaf constructions of $b_W(W, X)$ and $b_Z(Z, X)$. Details of the three basis constructions are provided in Appendix D.2.

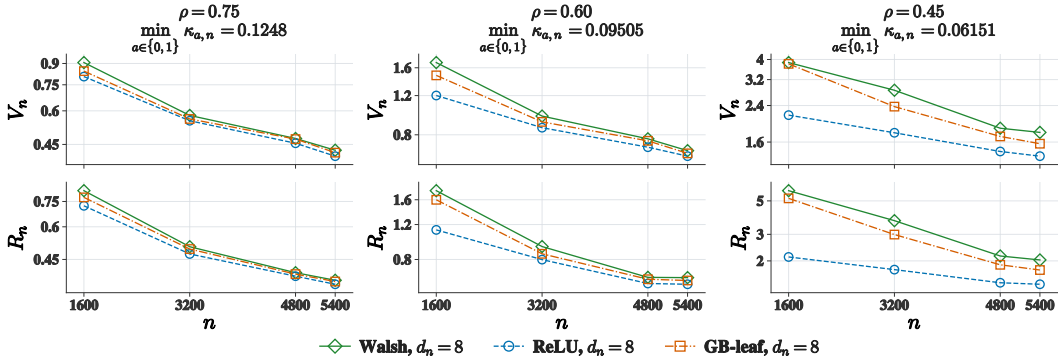

Figure 2: Estimated-bridge sampling fluctuation (top) and uniform expansion remainder (bottom).

At $\rho = 0.75$, Figure 2 shows that increasing n from 1600 to 5400 reduces the median V_n from approximately 0.80–0.91 to 0.41–0.43 and the median R_n from 0.72–0.82 to 0.36–0.37 across the three bases, i.e., reductions of 49%–53% and 50%–55%; these nearly identical decreases across Walsh, ReLU, and GB-leaf indicate robustness to the choice of basis construction. Within the fixed Walsh basis, $\rho = 0.75$, 0.60, and 0.45 correspond to $\min_{a\in\{0,1\}} \kappa_{a,n}$ of approximately 0.1248, 0.0951, and 0.0615. Over the same range of n, the median V_n and R_n decrease from 1.20–1.69 to 0.64–0.68 and from 1.13–1.78 to 0.60–0.65 at $\rho = 0.60$, and from 2.16–3.87 to 1.37–1.78 and from 2.12–5.84 to 1.40–2.03 at $\rho = 0.45$. Both medians decrease at each successive sample size for every basis and proxy setting, consistent with Theorem 3.1 and Lemma 3.1 with bounds on bridge estimation errors simultaneously over treatment arms and outcome thresholds.

Simulation 3: Simultaneous Confidence Bands for $\Delta_Q(\tau)$ without Density Estimation.

We examine *Challenge 3* by figuring whether confidence-band inversion maintains simultaneous coverage of the quantile treatment effect $\Delta_Q(\tau)$ when the usual normal approximation for the counterfactual quantile $Q_a(\tau)$ can be unreliable. The noise distribution changed to Double Weibull$(1/2)$, Beta$(1/4, 1/4)$, Beta$(1/4, 3/4)$, or χ^2_1. Further details are provided in Appendix D.3. We set $\lambda_h = \lambda_q = 0.01/\sqrt{N_k}$ on each training sample, $\alpha_{\text{sig}} = 0.05$ and $\mathcal{T} = \{0.10, 0.12, \ldots, 0.90\}$. We compare the simultaneous confidence bands $[\Delta_Q^L(\tau), \Delta_Q^U(\tau)]$, obtained by confidence-band inversion of the CDF bands $[L_a(y), U_a(y)]$, with Estimated-density delta, which uses the same counterfactual CDF estimate $\widehat{F}_a(y)$ as (13) and additionally estimates the counterfactual density.

Table 1 supports the Proposed method in addressing *Challenge 3*. Across all 12 combinations of noise distribution and sample size, Proposed achieves simultaneous coverage of 99.8%–100.0%, whereas Estimated-density delta achieves 68.4%–89.2%, below the nominal 95% level in every combination. The coverage advantage is 10.8–31.4 percentage points. Across 12 combinations, Proposed achieves mean band widths of 0.38–1.04, only exceeding the corresponding Oracle 95% widths by 0.06–0.27 (an average increase of approximately 20%), while maintaining simultaneous coverage of 99.8%–100.0%. Its mean band width decreases in every setting as n increases from 1600 to 4800, with the range across noise distributions decreasing from 0.76–1.04 to 0.38–0.57. These results indicate that the proposed method achieves conservative simultaneous coverage without an excessive increase in mean bandwidth, thereby striking a balance between efficiency and robustness.

Table 1: Simultaneous quantile treatment effects coverage and mean band width.

		Coverage (%)		Mean band width		
Noise	n	**Proposed**	Est.-density δ	**Proposed**	Est.-density δ	Oracle 95%
Double Weibull$(1/2)$	1600	**100.0**	81.6	**0.76**	0.41	0.58
	3200	**100.0**	81.4	**0.47**	0.26	0.39
	4800	**100.0**	74.6	**0.38**	0.21	0.32
Beta$(1/4, 1/4)$	1600	**100.0**	80.4	**1.04**	0.58	0.77
	3200	**100.0**	80.2	**0.70**	0.39	0.54
	4800	**100.0**	75.2	**0.57**	0.31	0.46
Beta$(1/4, 3/4)$	1600	**99.8**	80.2	**1.00**	0.55	0.81
	3200	**100.0**	68.8	**0.66**	0.36	0.54
	4800	**99.8**	68.4	**0.53**	0.29	0.45
χ^2_1	1600	**100.0**	89.2	**0.91**	0.48	0.73
	3200	**100.0**	86.6	**0.58**	0.31	0.49
	4800	**100.0**	85.4	**0.47**	0.25	0.41

Notes: Est.-density δ denotes Estimated-density delta.

5 Discussion

We develop simultaneous inference for the marginal counterfactual distribution $F_a(y)$ and the quantile treatment effect $\Delta_Q(\tau)$ under weak proxies: using the bridge moment equations, Corollary 3.1 establishes a population-bias bound for $\widehat{F}_a(y)$ with a sharp $\kappa_{a,n}^{-1}$ dependence rather than incurring the $\kappa_{a,n}^{-1}$ sampling-error factor twice as in the Cauchy–Schwarz bound, and Lemma 3.1 and Theorem 3.1 control the estimated bridges across treatment arms $a \in \{0, 1\}$ and outcome thresholds $y \in \mathcal{Y}$. Confidence-band inversion of the simultaneous CDF bands then yields simultaneous confidence bands for $\Delta_Q(\tau)$ without counterfactual density estimation. Future work includes inference when $\kappa_{a,n}$ is of the same order as the sample matrix error, data-driven choice of the basis vectors $b_W(W, X)$, $b_Z(Z, X)$ and the ridge parameters λ_h, λ_q, and less conservative confidence bands for $\Delta_Q(\tau)$.

AI USE STATEMENT

In this work, a generative AI assistant was used to check compliance with the ICLR 2027 template, edit and condense the manuscript language, and refine the LaTeX formatting, including the presentation of cited baseline methods and reported results.

REPRODUCIBILITY STATEMENT

The proposed implementation is available in our anonymous code repository.

ETHICS STATEMENT

The work does not recruit participants or use individual-level human-subject data.

REFERENCES

Andrew Bennett, Nathan Kallus, Xiaojie Mao, Whitney Newey, Vasilis Syrgkanis, and Masatoshi Uehara. Source condition double robust inference on functionals of inverse problems, 2023. arXiv:2307.13793.

Bariscan Bozkurt, Houssam Zenati, Dimitri Meunier, Liyuan Xu, and Arthur Gretton. Density ratio-free doubly robust proxy causal learning. In *Advances in Neural Information Processing Systems*, volume 38, 2025. doi: 10.52202/085713-1992.

Bariscan Bozkurt, Alexandre Galashov, Dimitri Meunier, Zikai Shen, Arthur Gretton, and Houssam Zenati. Doubly robust proxy causal learning with neural mean embeddings, 2026. Preprint.

Undral Byambadalai, Tatsushi Oka, and Shota Yasui. Estimating distributional treatment effects in randomized experiments: Machine learning for variance reduction. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 5082–5113. PMLR, 2024. URL https://proceedings.mlr.press/v235/byambadalai24a.html.

Marine Carrasco, Jean-Pierre Florens, and Eric Renault. Linear inverse problems in structural econometrics: Estimation based on spectral decomposition and regularization. In James J. Heckman and Edward E. Leamer, editors, *Handbook of Econometrics*, volume 6B, pages 5633–5751. Elsevier, 2007. doi: 10.1016/S1573-4412(07)06077-1.

Victor Chernozhukov, Iván Fernández-Val, and Blaise Melly. Inference on counterfactual distributions. *Econometrica*, 81(6):2205–2268, 2013. doi: 10.3982/ECTA10582.

Victor Chernozhukov, Iván Fernández-Val, Blaise Melly, and Kaspar Wüthrich. Generic inference on quantile and quantile effect functions for discrete outcomes. *Journal of the American Statistical Association*, 115(529):123–137, 2020. doi: 10.1080/01621459.2019.1611581.

Yifan Cui, Hongming Pu, Xu Shi, Wang Miao, and Eric J. Tchetgen Tchetgen. Semiparametric proximal causal inference. *Journal of the American Statistical Association*, 119:1348–1359, 2024. doi: 10.1080/01621459.2023.2191817.

Sergio Firpo. Efficient semiparametric estimation of quantile treatment effects. *Econometrica*, 75(1): 259–276, 2007. doi: 10.1111/j.1468-0262.2007.00738.x.

Peter Gaenssler, Péter Molnár, and Daniel Rost. On continuity and strict increase of the CDF for the sup-functional of a Gaussian process with applications to statistics. *Results in Mathematics*, 51: 51–60, 2007. doi: 10.1007/s00025-007-0257-1.

Amiremad Ghassami, Andrew Ying, Ilya Shpitser, and Eric Tchetgen Tchetgen. Minimax kernel machine learning for a class of doubly robust functionals with application to proximal causal inference. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pages 7210–7239. PMLR, 2022. URL https://proceedings.mlr.press/v151/ghassami22a.html.

Nathan Kallus, Xiaojie Mao, and Masatoshi Uehara. Causal inference under unmeasured confounding with negative controls: A minimax learning approach. *arXiv preprint arXiv:2103.14029*, 2021.

Benjamin Kompa, David R. Bellamy, Thomas Kolokotrones, James M. Robins, and Andrew L. Beam. Deep learning methods for proximal inference via maximum moment restriction. In *Advances in Neural Information Processing Systems*, volume 35, 2022. doi: 10.52202/068431-0813.

Wang Miao, Zhi Geng, and Eric J. Tchetgen Tchetgen. Identifying causal effects with proxy variables of an unmeasured confounder. *Biometrika*, 105(4):987–993, 2018. doi: 10.1093/biomet/asy038.

Andrea Rotnitzky, Ezequiel Smucler, and James M. Robins. Characterization of parameters with a mixed bias property. *Biometrika*, 108(1):231–238, 2021. doi: 10.1093/biomet/asaa054.

Eric J. Tchetgen Tchetgen, Andrew Ying, Yifan Cui, Xu Shi, and Wang Miao. An introduction to proximal causal inference. *Statistical Science*, 39(3):375–390, 2024. doi: 10.1214/23-STS911.

N. N. Vakhania. The topological support of Gaussian measure in Banach space. *Nagoya Mathematical Journal*, 57:59–63, 1975. doi: 10.1017/S002776300001655X.

Aad van der Vaart and Jon A. Wellner. A local maximal inequality under uniform entropy. *Electronic Journal of Statistics*, 5:192–203, 2011.

Aad W. van der Vaart and Jon A. Wellner. *Weak Convergence and Empirical Processes*. Springer, 1996.

Yong Wu, Yanwei Fu, Shouyan Wang, and Xinwei Sun. Doubly robust proximal causal learning for continuous treatments. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=TjGJFkU3xL.

Liyuan Xu, Heishiro Kanagawa, and Arthur Gretton. Deep proxy causal learning and its application to confounded bandit policy evaluation. In *Advances in Neural Information Processing Systems*, volume 34, pages 26264–26275, 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/hash/dcf3219715a7c9cd9286f19db46f2384-Abstract.html.

Andrew Ying, Yifan Cui, and Eric J. Tchetgen Tchetgen. Proximal causal inference for marginal counterfactual survival curves, 2022. arXiv:2204.13144.

Jeffrey Zhang, Wei Li, Wang Miao, and Eric Tchetgen Tchetgen. Proximal causal inference without uniqueness assumptions. *Statistics & Probability Letters*, 198:109836, 2023. doi: 10.1016/j.spl.2023.109836.

A Derivation and Proof for Regularized Finite-Dimensional Bridge Estimation

We derive the moment equations and regularized estimates in Section 3.1, define the finite-dimensional approximation errors and coefficient norm bounds, and prove Proposition 3.1. All expectations E are under the true distribution of the observed data P_0.

A.1 Derivation of the Finite-Dimensional Moment Equations

The folds $I_1, \ldots, I_K$ partition the observation indices $\{1, \ldots, n\}$ independently of $O_1, \ldots, O_n$. The number of folds K is fixed, and $|I_k| \asymp n$ for every k. The evaluation sample for fold I_k is $\{O_i : i \in I_k\}$, and the training sample is $\{O_i : i \notin I_k\}$. Define the training sample size and training sample average by

$$N_k = n - |I_k|, \qquad P_{n,-k} f = \frac{1}{N_k} \sum_{i \notin I_k} f(O_i). \tag{30}$$

Because K is fixed and every fold contains a number of observations proportional to n, $N_k \asymp n$. The superscript $(-k)$ indicates that estimation uses only the training sample outside I_k.

Fix a treatment arm $a \in \{0, 1\}$. The population counterparts of the sample moments in (5)–(6) are

$$\begin{aligned} \Sigma_a &= E\{\mathbf{1}\{A = a\} b_Z(Z, X) b_W(W, X)^\top\}, \\ \gamma_a(y) &= E\{\mathbf{1}\{A = a\} b_Z(Z, X) \mathbf{1}\{Y \le y\}\}, \\ \mu_W &= E\{b_W(W, X)\}. \end{aligned} \tag{31}$$

The matrix Σ_a has dimension $d_n \times d_n$, and $\gamma_a(y)$ and μ_W have dimension d_n.

A.1.1 From Proximal Identification to a Primal–Dual Inverse Problem

Although the representation in Section 2.1 identifies $F_a(y)$ using the outcome bridge $h_{a,y}(W, X)$ and the treatment bridge $q_a(Z, X)$, it does not by itself guarantee stable estimation. The outcome bridge and treatment bridge are defined indirectly through conditional expectation equations. In particular, recovering $h_{a,y}(W, X)$ can be unstable when the proxies Z and W provide limited information about the unobserved latent variable U.

We rewrite the two bridge equations in operator form. For each treatment arm $a \in \{0, 1\}$, define the conditional expectation operator T_a through

$$(T_a h_{a,y})(Z, X) = E\{h_{a,y}(W, X) \mid A = a, Z, X\}.$$

Also define

$$g_{a,y}(Z, X) = E\{\mathbf{1}\{Y \le y\} \mid A = a, Z, X\}.$$

The equation for the outcome bridge $h_{a,y}(W, X)$ can then be written as

$$(T_a h_{a,y})(Z, X) = g_{a,y}(Z, X),$$

or, more compactly,

$$T_a h_{a,y} = g_{a,y}.$$

For the outcome bridge $h_{a,y}(W, X)$ and the treatment bridge $q_a(Z, X)$,
$E\big[\mathbf{1}\{A = a\}(T_a h_{a,y})(Z, X) q_a(Z, X)\big] = E\big[\mathbf{1}\{A = a\} h_{a,y}(W, X) q_a(Z, X)\big] = E\big[h_{a,y}(W, X)(T_a^* q_a)(W, X)\big]$,
where the adjoint operator T_a^* acts on the treatment bridge $q_a(Z, X)$ as

$$(T_a^* q_a)(W, X) = E\{\mathbf{1}\{A = a\} q_a(Z, X) \mid W, X\}.$$

By the definition of the treatment bridge $q_a(Z, X)$,

$$E\{\mathbf{1}\{A = a\} q_a(Z, X) - 1 \mid W, X\} = 0.$$

Therefore,

$$(T_a^* q_a)(W, X) = E\{\mathbf{1}\{A = a\} q_a(Z, X) \mid W, X\} = 1,$$

which we write compactly as

$$T_a^* q_a = 1.$$

Together with the outcome bridge equation, the two bridge equations can now be written as the primal–dual pair

$$T_a h_{a,y} = g_{a,y}, \qquad T_a^* q_a = 1.$$

The two equations show how functions of (W, X) and functions of (Z, X) are linked in opposite directions through T_a and its adjoint operator T_a^*: T_a maps the outcome bridge $h_{a,y}(W, X)$ to the observed quantity $g_{a,y}(Z, X)$, while T_a^* maps the treatment bridge $q_a(Z, X)$ back to the constant function $(T_a^* q_a)(W, X) = 1$.

Recalling (2), the primal–dual pair is useful because the marginal counterfactual distribution $F_a(y)$ depends on the outcome bridge only through $F_a(y) = E\{h_{a,y}(W, X)\}$. Using the primal–dual pair

$T_a^* q_a = 1$ and $T_a h_{a,y} = g_{a,y}$,

$$\begin{aligned} F_a(y) &= E\{h_{a,y}(W,X)\} = E\big[h_{a,y}(W,X)\,(T_a^* q_a)(W,X)\big] \\ &= E\big[\mathbf{1}\{A=a\}(T_a h_{a,y})(Z,X)\,q_a(Z,X)\big] \\ &= E\big[\mathbf{1}\{A=a\} g_{a,y}(Z,X)\,q_a(Z,X)\big]. \end{aligned} \tag{32}$$

Thus, even when recovering the entire outcome bridge $h_{a,y}(W,X)$ is unstable, the marginal counterfactual distribution $F_a(y)$ may still be estimated stably through the primal–dual pair.

A.1.2 Finite-Dimensional Bridge Moment Equations

Outcome bridge moment equation. The outcome bridge equation $T_a h_{a,y} = g_{a,y}$ gives

$$E\{\mathbf{1}\{Y \le y\} - h_{a,y}(W,X) \mid A=a, Z, X\} = 0.$$

Multiplying the conditional equation by $b_Z(Z,X)$ and integrating over observations with $A=a$ yields

$$\begin{aligned} &E[\mathbf{1}\{A=a\} b_Z(Z,X)\{\mathbf{1}\{Y \le y\} - h_{a,y}(W,X)\}] \\ &\quad = E[\mathbf{1}\{A=a\} b_Z(Z,X) E\{\mathbf{1}\{Y \le y\} - h_{a,y}(W,X) \mid A, Z, X\}] = 0. \end{aligned} \tag{33}$$

For the finite-dimensional representation $b_W(W,X)^\top \theta$, the tested moment becomes

$$\begin{aligned} &E\big[\mathbf{1}\{A=a\} b_Z(Z,X)\{\mathbf{1}\{Y \le y\} - b_W(W,X)^\top \theta\}\big] \\ &\quad = E\{\mathbf{1}\{A=a\} b_Z(Z,X) \mathbf{1}\{Y \le y\}\} - E\big\{\mathbf{1}\{A=a\} b_Z(Z,X) b_W(W,X)^\top\big\}\,\theta \\ &\quad = \gamma_a(y) - \Sigma_a \theta. \end{aligned} \tag{34}$$

Setting the tested moment in (34) equal to zero gives the finite-dimensional equation $\Sigma_a \theta = \gamma_a(y)$. The finite-dimensional equation imposes the moments tested by $b_Z(Z,X)$, not every conditional moment in $T_a h_{a,y} = g_{a,y}$.

Treatment bridge moment equation. The treatment bridge equation $T_a^* q_a = 1$ gives

$$E\{\mathbf{1}\{A=a\} q_a(Z,X) - 1 \mid W, X\} = 0.$$

Multiplying the conditional equation by $b_W(W,X)$ and taking expectation yields

$$\begin{aligned} &E[b_W(W,X)\{\mathbf{1}\{A=a\} q_a(Z,X) - 1\}] \\ &\quad = E[b_W(W,X) E\{\mathbf{1}\{A=a\} q_a(Z,X) - 1 \mid W, X\}] = 0. \end{aligned} \tag{35}$$

For the finite-dimensional representation $b_Z(Z,X)^\top \alpha$, the tested moment becomes

$$\begin{aligned} &E\big[b_W(W,X)\{\mathbf{1}\{A=a\} b_Z(Z,X)^\top \alpha - 1\}\big] \\ &\quad = E\big\{\mathbf{1}\{A=a\} b_W(W,X) b_Z(Z,X)^\top\big\}\,\alpha - E\{b_W(W,X)\} \\ &\quad = \Sigma_a^\top \alpha - \mu_W. \end{aligned} \tag{36}$$

Setting the tested moment in (36) equal to zero gives the finite-dimensional equation $\Sigma_a^\top \alpha = \mu_W$.

Population finite-dimensional representations. When Σ_a is nonsingular, define

$$\begin{aligned} \theta_a^{(d_n)}(y) &= \Sigma_a^{-1} \gamma_a(y), \\ \alpha_a^{(d_n)} &= (\Sigma_a^\top)^{-1} \mu_W, \\ h_{a,y}^{(d_n)}(W,X) &= b_W(W,X)^\top \theta_a^{(d_n)}(y), \\ q_a^{(d_n)}(Z,X) &= b_Z(Z,X)^\top \alpha_a^{(d_n)}. \end{aligned} \tag{37}$$

The coefficients $\theta_a^{(d_n)}(y)$ and $\alpha_a^{(d_n)}$ solve

$$\Sigma_a \theta_a^{(d_n)}(y) = \gamma_a(y), \qquad \Sigma_a^\top \alpha_a^{(d_n)} = \mu_W. \tag{38}$$

The population finite-dimensional representations $h_{a,y}^{(d_n)}(W,X)$ and $q_a^{(d_n)}(Z,X)$ need not equal the true outcome bridge $h_{a,y}(W,X)$ and treatment bridge $q_a(Z,X)$.

Sample moment equations. Replacing E with $P_{n,-k}$ in (31) gives

$$\begin{aligned} \widehat{\Sigma}_a^{(-k)} &= \frac{1}{N_k} \sum_{i \notin I_k} \mathbf{1}\{A_i = a\} b_Z(Z_i,X_i) b_W(W_i,X_i)^\top, \\ \widehat{\gamma}_a^{(-k)}(y) &= \frac{1}{N_k} \sum_{i \notin I_k} \mathbf{1}\{A_i = a\} b_Z(Z_i,X_i) \mathbf{1}\{Y_i \le y\}, \\ \widehat{\mu}_W^{(-k)} &= \frac{1}{N_k} \sum_{i \notin I_k} b_W(W_i,X_i). \end{aligned} \tag{39}$$

The sample versions of the tested moments are therefore

$$\begin{aligned} &P_{n,-k}\big[\mathbf{1}\{A=a\}b_Z(Z,X)\{\mathbf{1}\{Y\le y\}-b_W(W,X)^\top\theta\}\big] \\ &\quad=\widehat{\gamma}_a^{(-k)}(y)-\widehat{\Sigma}_a^{(-k)}\theta, \\ &P_{n,-k}\big[b_W(W,X)\{\mathbf{1}\{A=a\}b_Z(Z,X)^\top\alpha-1\}\big] \\ &\quad=(\widehat{\Sigma}_a^{(-k)})^\top\alpha-\widehat{\mu}_W^{(-k)}. \end{aligned} \tag{40}$$

Exact sample moment estimation would set both expressions in (40) equal to zero. The regularized estimates in (8) instead minimize squared moment errors plus coefficient penalties. The regularized estimates need not satisfy the unregularized sample moment equations exactly.

A.2 Derivation of the Regularized Estimates

For the outcome bridge, differentiating the objective in (8) with respect to θ gives

$$\begin{aligned} &\nabla_\theta\Big\{\|\widehat{\Sigma}_a^{(-k)}\theta-\widehat{\gamma}_a^{(-k)}(y)\|_2^2+\lambda_h\|\theta\|_2^2\Big\} \\ &\quad=2(\widehat{\Sigma}_a^{(-k)})^\top\{\widehat{\Sigma}_a^{(-k)}\theta-\widehat{\gamma}_a^{(-k)}(y)\}+2\lambda_h\theta. \end{aligned} \tag{41}$$

For the treatment bridge, differentiating the objective with respect to α gives

$$\begin{aligned} &\nabla_\alpha\Big\{\|(\widehat{\Sigma}_a^{(-k)})^\top\alpha-\widehat{\mu}_W^{(-k)}\|_2^2+\lambda_q\|\alpha\|_2^2\Big\} \\ &\quad=2\widehat{\Sigma}_a^{(-k)}\{(\widehat{\Sigma}_a^{(-k)})^\top\alpha-\widehat{\mu}_W^{(-k)}\}+2\lambda_q\alpha. \end{aligned} \tag{42}$$

Let I_{d_n} denote the $d_n\times d_n$ identity matrix. The Hessians of the outcome bridge and treatment bridge objectives are twice the respective matrices

$$(\widehat{\Sigma}_a^{(-k)})^\top\widehat{\Sigma}_a^{(-k)}+\lambda_h I_{d_n},\qquad \widehat{\Sigma}_a^{(-k)}(\widehat{\Sigma}_a^{(-k)})^\top+\lambda_q I_{d_n}.$$

For every nonzero vector $v\in\mathbb{R}^{d_n}$,

$$\begin{aligned} v^\top\{(\widehat{\Sigma}_a^{(-k)})^\top\widehat{\Sigma}_a^{(-k)}+\lambda_h I_{d_n}\}v&=\|\widehat{\Sigma}_a^{(-k)}v\|_2^2+\lambda_h\|v\|_2^2>0, \\ v^\top\{\widehat{\Sigma}_a^{(-k)}(\widehat{\Sigma}_a^{(-k)})^\top+\lambda_q I_{d_n}\}v&=\|(\widehat{\Sigma}_a^{(-k)})^\top v\|_2^2+\lambda_q\|v\|_2^2>0. \end{aligned} \tag{43}$$

Positive λ_h and λ_q therefore give strictly convex quadratic objectives and unique minimizers, even when $\widehat{\Sigma}_a^{(-k)}$ is singular. Setting the gradients in (41) and (42) equal to zero gives

$$\begin{aligned} \{(\widehat{\Sigma}_a^{(-k)})^\top\widehat{\Sigma}_a^{(-k)}+\lambda_h I_{d_n}\}\widehat{\theta}_{a,\lambda_h}^{(-k)}(y)&=(\widehat{\Sigma}_a^{(-k)})^\top\widehat{\gamma}_a^{(-k)}(y), \\ \{\widehat{\Sigma}_a^{(-k)}(\widehat{\Sigma}_a^{(-k)})^\top+\lambda_q I_{d_n}\}\widehat{\alpha}_{a,\lambda_q}^{(-k)}&=\widehat{\Sigma}_a^{(-k)}\widehat{\mu}_W^{(-k)}. \end{aligned} \tag{44}$$

Inverting the positive definite matrices in (44) yields

$$\begin{aligned} \widehat{\theta}_{a,\lambda_h}^{(-k)}(y)&=\{(\widehat{\Sigma}_a^{(-k)})^\top\widehat{\Sigma}_a^{(-k)}+\lambda_h I_{d_n}\}^{-1}(\widehat{\Sigma}_a^{(-k)})^\top\widehat{\gamma}_a^{(-k)}(y), \\ \widehat{\alpha}_{a,\lambda_q}^{(-k)}&=\{\widehat{\Sigma}_a^{(-k)}(\widehat{\Sigma}_a^{(-k)})^\top+\lambda_q I_{d_n}\}^{-1}\widehat{\Sigma}_a^{(-k)}\widehat{\mu}_W^{(-k)}. \end{aligned} \tag{45}$$

Substituting the coefficients in (45) into the basis representations gives exactly the estimated bridges in (9). Within each training sample, $b_W(W,X)$, $b_Z(Z,X)$, $\widehat{\Sigma}_a^{(-k)}$, λ_h, and λ_q are fixed across thresholds. Consequently, the estimated outcome bridge depends on y only through $\widehat{\gamma}_a^{(-k)}(y)$, and the estimated treatment bridge does not depend on y.

A.3 Finite-Dimensional Approximation Errors and Assumptions

For functions $u(W,X)$ and $v(Z,X)$, the norms used in the main text are $\|u\|_W^2=E\{u(W,X)^2\}$ and $\|v\|_{a,Z}^2=E\{\mathbf{1}\{A=a\}v(Z,X)^2\}$. Using $\|\cdot\|_W$ and $\|\cdot\|_{a,Z}$, define

$$\begin{aligned} b_{h,a,n}&=\sup_{y\in\mathcal{Y}}\|h_{a,y}^{(d_n)}(W,X)-h_{a,y}(W,X)\|_W, \\ b_{q,a,n}&=\|q_a^{(d_n)}(Z,X)-q_a(Z,X)\|_{a,Z}, \\ M_{h,a,n}&=\sup_{y\in\mathcal{Y}}\|\theta_a^{(d_n)}(y)\|_2, \\ M_{q,a,n}&=\|\alpha_a^{(d_n)}\|_2. \end{aligned} \tag{46}$$

The finite-dimensional approximation errors $b_{h,a,n}$ and $b_{q,a,n}$ concern the population finite-dimensional representations in (37). The coefficient norm bounds $M_{h,a,n}$ and $M_{q,a,n}$ concern population coefficients. The quantities in (46) may depend on n through the basis dimension and the population moment equations. No zero finite-dimensional approximation error or uniformly bounded coefficient norm is assumed.

Assumption A.1 (Finite-dimensional concentration and weak-proxy strength)**.** The following conditions hold for both treatment arms $a \in \{0, 1\}$.

(i) The basis vectors $b_W(W, X)$ and $b_Z(Z, X)$ are prespecified, have the same dimension d_n, and are bounded or sub-Gaussian. There are constants $0 < c_b \le C_b < \infty$, independent of n and a, such that for every $v \in \mathbb{R}^{d_n}$,

$$\begin{aligned} c_b\|v\|_2^2 \le v^\top E\{b_W(W,X)b_W(W,X)^\top\}v \le C_b\|v\|_2^2, \\ c_b\|v\|_2^2 \le v^\top E\big\{\mathbf{1}\{A=a\}b_Z(Z,X)b_Z(Z,X)^\top\big\}\,v \le C_b\|v\|_2^2. \end{aligned} \tag{47}$$

(ii) The cross-moment matrix Σ_a is nonsingular, with $\kappa_{a,n} = \sigma_{\min}(\Sigma_a) > 0$. The finite-dimensional approximation errors and coefficient norm bounds in (46) are finite.

(iii) With the fixed folds and training sample sizes in (30), the cross-moment matrix and basis mean satisfy

$$\begin{aligned} \max_k \|\widehat{\Sigma}_a^{(-k)} - \Sigma_a\|_{\mathrm{op}} &= O_p\left(\sqrt{\frac{d_n}{n}}\right), \\ \max_k \|\widehat{\mu}_W^{(-k)} - \mu_W\|_2 &= O_p\left(\sqrt{\frac{d_n}{n}}\right). \end{aligned} \tag{48}$$

Here $\|\cdot\|_{\mathrm{op}}$ denotes the matrix operator norm induced by the Euclidean norm $\|\cdot\|_2$.

The moment-concentration conditions in (48) are assumptions on observed matrix and vector averages. The moment-concentration conditions are not deduced from boundedness or sub-Gaussianity alone.

Uniform concentration of the threshold moments. The upper Gram bound in Assumption A.1(i) implies

$$E\big\{\mathbf{1}\{A=a\}\|b_Z(Z,X)\|_2^2\big\} = \operatorname{tr} E\big\{\mathbf{1}\{A=a\}b_Z(Z,X)b_Z(Z,X)^\top\big\} \le C_b d_n.$$

Fix a treatment arm a and a fold I_k, and condition on the fold partition. The training observations remain iid from P_0, and N_k is fixed in the following calculation. Let $\{\varepsilon_i : i \notin I_k\}$ be independent Rademacher signs, independent of the observations and the fold partition, with $P(\varepsilon_i = 1) = P(\varepsilon_i = -1) = 1/2$. Write E_ε for expectation over the signs conditional on the training observations and the fold partition. All suprema below are measurable, since the threshold sums and $\gamma_a(y)$ are right-continuous and a fixed countable dense subset of $\mathcal{Y}$ together with y_U determines the suprema. Symmetrization gives

$$\begin{aligned} &E\left[\sup_{y\in\mathcal{Y}} \|\widehat{\gamma}_a^{(-k)}(y) - \gamma_a(y)\|_2\right] \\ &\quad\le \frac{2}{N_k} E E_\varepsilon \sup_{y\in\mathcal{Y}} \left\| \sum_{i\notin I_k} \varepsilon_i \mathbf{1}\{A_i = a\} b_Z(Z_i, X_i)\mathbf{1}\{Y_i \le y\} \right\|_2. \end{aligned}$$

Conditionally on the training observations, order the training indices as $i_1, \ldots, i_{N_k}$ so that $Y_{i_1} \le \cdots \le Y_{i_{N_k}}$, breaking ties by the original observation indices. Define

$$S_0 = 0, \qquad S_m = \sum_{j=1}^{m} \varepsilon_{i_j} \mathbf{1}\{A_{i_j} = a\} b_Z(Z_{i_j}, X_{i_j}), \qquad m = 1, \ldots, N_k.$$

Every threshold sum is one of the partial sums S_m, including when training outcomes are tied. The sequence S_m is a square-integrable vector martingale with respect to the successive Rademacher signs, and $\|S_m\|_2$ is a nonnegative submartingale. Doob's L_2 maximal inequality therefore gives

$$\begin{aligned} &E_\varepsilon \sup_{y\in\mathcal{Y}} \left\| \sum_{i\notin I_k} \varepsilon_i \mathbf{1}\{A_i = a\} b_Z(Z_i, X_i)\mathbf{1}\{Y_i \le y\} \right\|_2^2 \\ &\quad\le E_\varepsilon \max_{0\le m\le N_k} \|S_m\|_2^2 \le 4E_\varepsilon\|S_{N_k}\|_2^2 \\ &\quad= 4\sum_{i\notin I_k} \mathbf{1}\{A_i = a\}\|b_Z(Z_i, X_i)\|_2^2. \end{aligned}$$

Combining the symmetrization and maximal bounds, then applying Jensen's inequality, yields

$$
\begin{aligned}
&E\left[\sup_{y\in\mathcal{Y}}\|\widehat{\gamma}_a^{(-k)}(y)-\gamma_a(y)\|_2\right] \\
&\le \frac{4}{N_k}E\left[\left\{\sum_{i\notin I_k}\mathbf{1}\{A_i=a\}\|b_Z(Z_i,X_i)\|_2^2\right\}^{1/2}\right] \\
&\le 4\sqrt{\frac{E\{\mathbf{1}\{A=a\}\|b_Z(Z,X)\|_2^2\}}{N_k}} \le 4\sqrt{\frac{C_b d_n}{N_k}}.
\end{aligned}
$$

Markov's inequality, $N_k \asymp n$, and the fixed number of folds imply

$$
\max_k \sup_{y\in\mathcal{Y}}\|\widehat{\gamma}_a^{(-k)}(y)-\gamma_a(y)\|_2 = O_p\left(\sqrt{\frac{d_n}{n}}\right). \tag{49}
$$

The rate in (49) follows from the upper Gram bound and the threshold ordering, without an additional logarithmic factor. The argument applies to both treatment arms and does not require overlapping training samples to be independent. The matrix concentration condition in (48) remains a separate assumption.

A.4 Proof of the Uniform Bridge Estimation Error Bounds

Proof of Proposition 3.1. Step 1: Control the sample matrices. By (48) and $\sqrt{d_n/n} = o(\kappa_{a,n})$,

$$
\frac{\max_k \|\widehat{\Sigma}_a^{(-k)}-\Sigma_a\|_{\mathrm{op}}}{\kappa_{a,n}} = O_p\left(\frac{1}{\kappa_{a,n}}\sqrt{\frac{d_n}{n}}\right) = o_p(1).
$$

Because there are two treatment arms, the event

$$
\mathcal{E}_n = \bigcap_{a\in\{0,1\}}\left\{\max_k \|\widehat{\Sigma}_a^{(-k)}-\Sigma_a\|_{\mathrm{op}} \le \frac{\kappa_{a,n}}{2}\right\} \tag{50}
$$

has probability tending to one. On $\mathcal{E}_n$, for every a, k, and $v\in\mathbb{R}^{d_n}$,

$$
\begin{aligned}
\|\widehat{\Sigma}_a^{(-k)}v\|_2 &\ge \|\Sigma_a v\|_2 - \|(\widehat{\Sigma}_a^{(-k)}-\Sigma_a)v\|_2 \\
&\ge \left\{\kappa_{a,n} - \|\widehat{\Sigma}_a^{(-k)}-\Sigma_a\|_{\mathrm{op}}\right\}\|v\|_2 \\
&\ge \frac{\kappa_{a,n}}{2}\|v\|_2.
\end{aligned} \tag{51}
$$

Taking the infimum over unit vectors gives

$$
\sigma_{\min}(\widehat{\Sigma}_a^{(-k)}) \ge \frac{\kappa_{a,n}}{2}, \qquad \|\{\widehat{\Sigma}_a^{(-k)}\}^{-1}\|_{\mathrm{op}} \le \frac{2}{\kappa_{a,n}}.
$$

The eigenvalues of both $(\widehat{\Sigma}_a^{(-k)})^\top\widehat{\Sigma}_a^{(-k)}$ and $\widehat{\Sigma}_a^{(-k)}(\widehat{\Sigma}_a^{(-k)})^\top$ are squared singular values of $\widehat{\Sigma}_a^{(-k)}$. Every eigenvalue is therefore at least $\kappa_{a,n}^2/4$ on $\mathcal{E}_n$. Adding the positive regularization parameters and inverting gives

$$
\begin{aligned}
\left\|\{(\widehat{\Sigma}_a^{(-k)})^\top\widehat{\Sigma}_a^{(-k)}+\lambda_h I_{d_n}\}^{-1}\right\|_{\mathrm{op}} &\le \frac{1}{\kappa_{a,n}^2/4+\lambda_h}, \\
\left\|\{\widehat{\Sigma}_a^{(-k)}(\widehat{\Sigma}_a^{(-k)})^\top+\lambda_q I_{d_n}\}^{-1}\right\|_{\mathrm{op}} &\le \frac{1}{\kappa_{a,n}^2/4+\lambda_q}.
\end{aligned} \tag{52}
$$

Let $\sigma_j(\widehat{\Sigma}_a^{(-k)})$, $j=1,\dots,d_n$, denote the singular values of the sample matrix, including zero singular values. Singular value decomposition gives, for every $\lambda_h,\lambda_q>0$,

$$
\begin{aligned}
\left\|\{(\widehat{\Sigma}_a^{(-k)})^\top\widehat{\Sigma}_a^{(-k)}+\lambda_h I_{d_n}\}^{-1}(\widehat{\Sigma}_a^{(-k)})^\top\right\|_{\mathrm{op}} &= \max_j \frac{\sigma_j(\widehat{\Sigma}_a^{(-k)})}{\sigma_j(\widehat{\Sigma}_a^{(-k)})^2+\lambda_h}, \\
\left\|\{\widehat{\Sigma}_a^{(-k)}(\widehat{\Sigma}_a^{(-k)})^\top+\lambda_q I_{d_n}\}^{-1}\widehat{\Sigma}_a^{(-k)}\right\|_{\mathrm{op}} &= \max_j \frac{\sigma_j(\widehat{\Sigma}_a^{(-k)})}{\sigma_j(\widehat{\Sigma}_a^{(-k)})^2+\lambda_q}.
\end{aligned} \tag{53}
$$

The identities in (53) do not require an invertible sample matrix. On $\mathcal{E}_n$, the singular values are at least $\kappa_{a,n}/2$. For any $\lambda > 0$, elementary maximization gives

$$\sup_{t \geq \kappa_{a,n}/2} \frac{t}{t^2 + \lambda} = \begin{cases} \dfrac{\kappa_{a,n}/2}{\kappa_{a,n}^2/4 + \lambda}, & 0 < \lambda \leq \kappa_{a,n}^2/4, \\ \dfrac{1}{2\sqrt{\lambda}}, & \lambda > \kappa_{a,n}^2/4. \end{cases} \tag{54}$$

In particular,

$$\begin{aligned} \left\| \{(\widehat{\Sigma}_a^{(-k)})^\top \widehat{\Sigma}_a^{(-k)} + \lambda_h I_{d_n}\}^{-1} (\widehat{\Sigma}_a^{(-k)})^\top \right\|_{\mathrm{op}} &\leq \sup_{t \geq \kappa_{a,n}/2} \frac{t}{t^2 + \lambda_h} \leq \frac{2}{\sqrt{\kappa_{a,n}^2 + \lambda_h}}, \\ \left\| \{\widehat{\Sigma}_a^{(-k)} (\widehat{\Sigma}_a^{(-k)})^\top + \lambda_q I_{d_n}\}^{-1} \widehat{\Sigma}_a^{(-k)} \right\|_{\mathrm{op}} &\leq \sup_{t \geq \kappa_{a,n}/2} \frac{t}{t^2 + \lambda_q} \leq \frac{2}{\sqrt{\kappa_{a,n}^2 + \lambda_q}}. \end{aligned} \tag{55}$$

The last inequalities follow from $t/(t^2+\lambda) \leq (t^2+\lambda)^{-1/2}$ and $t \geq \kappa_{a,n}/2$. The proposition uses this simpler expression up to numerical constants, while (53)–(54) retain the exact factors. *Step 2: Bound the outcome bridge estimation error.* Subtracting $\{(\widehat{\Sigma}_a^{(-k)})^\top \widehat{\Sigma}_a^{(-k)} + \lambda_h I_{d_n}\}\theta_a^{(d_n)}(y)$ from both sides of the outcome bridge normal equation in (44) gives

$$\begin{aligned} &\{(\widehat{\Sigma}_a^{(-k)})^\top \widehat{\Sigma}_a^{(-k)} + \lambda_h I_{d_n}\}\{\widehat{\theta}_{a,\lambda_h}^{(-k)}(y) - \theta_a^{(d_n)}(y)\} \\ &\quad = (\widehat{\Sigma}_a^{(-k)})^\top \{\widehat{\gamma}_a^{(-k)}(y) - \widehat{\Sigma}_a^{(-k)} \theta_a^{(d_n)}(y)\} - \lambda_h \theta_a^{(d_n)}(y). \end{aligned} \tag{56}$$

The population equation $\Sigma_a \theta_a^{(d_n)}(y) = \gamma_a(y)$ gives the exact identity

$$\begin{aligned} \widehat{\gamma}_a^{(-k)}(y) - \widehat{\Sigma}_a^{(-k)} \theta_a^{(d_n)}(y) = \widehat{\gamma}_a^{(-k)}(y) - \gamma_a(y) \\ + (\Sigma_a - \widehat{\Sigma}_a^{(-k)}) \theta_a^{(d_n)}(y). \end{aligned}$$

Substitution into (56), followed by multiplication by the regularized inverse, gives

$$\begin{aligned} &\widehat{\theta}_{a,\lambda_h}^{(-k)}(y) - \theta_a^{(d_n)}(y) \\ &\quad = \{(\widehat{\Sigma}_a^{(-k)})^\top \widehat{\Sigma}_a^{(-k)} + \lambda_h I_{d_n}\}^{-1} (\widehat{\Sigma}_a^{(-k)})^\top \\ &\qquad \times \left\{ \widehat{\gamma}_a^{(-k)}(y) - \gamma_a(y) + (\Sigma_a - \widehat{\Sigma}_a^{(-k)}) \theta_a^{(d_n)}(y) \right\} \\ &\qquad - \lambda_h \{(\widehat{\Sigma}_a^{(-k)})^\top \widehat{\Sigma}_a^{(-k)} + \lambda_h I_{d_n}\}^{-1} \theta_a^{(d_n)}(y). \end{aligned} \tag{57}$$

Applying (52) and (55), then taking the supremum over thresholds and the maximum over folds, yields

$$\begin{aligned} &\max_k \sup_{y \in \mathcal{Y}} \|\widehat{\theta}_{a,\lambda_h}^{(-k)}(y) - \theta_a^{(d_n)}(y)\|_2 \\ &\leq \left(\sup_{t \geq \kappa_{a,n}/2} \frac{t}{t^2 + \lambda_h} \right) \left\{ \max_k \sup_{y \in \mathcal{Y}} \|\widehat{\gamma}_a^{(-k)}(y) - \gamma_a(y)\|_2 + M_{h,a,n} \max_k \|\widehat{\Sigma}_a^{(-k)} - \Sigma_a\|_{\mathrm{op}} \right\} \\ &\quad + \frac{\lambda_h M_{h,a,n}}{\kappa_{a,n}^2/4 + \lambda_h}. \end{aligned} \tag{58}$$

To convert coefficient error into outcome bridge estimation error, integrate over (W, X) with the estimated coefficients held fixed. By (9), (37), and (47),

$$\begin{aligned} &\|\widehat{h}_{a,y}^{(-k)}(W,X) - h_{a,y}^{(d_n)}(W,X)\|_W^2 \\ &\quad = \{\widehat{\theta}_{a,\lambda_h}^{(-k)}(y) - \theta_a^{(d_n)}(y)\}^\top E\{b_W(W,X) b_W(W,X)^\top\} \\ &\qquad \times \{\widehat{\theta}_{a,\lambda_h}^{(-k)}(y) - \theta_a^{(d_n)}(y)\} \\ &\quad \leq C_b \|\widehat{\theta}_{a,\lambda_h}^{(-k)}(y) - \theta_a^{(d_n)}(y)\|_2^2. \end{aligned} \tag{59}$$

The triangle inequality and the definition of $b_{h,a,n}$ imply

$$\begin{aligned} &\max_k \sup_{y \in \mathcal{Y}} \|\widehat{h}_{a,y}^{(-k)}(W,X) - h_{a,y}(W,X)\|_W \\ &\quad \leq \sqrt{C_b} \max_k \sup_{y \in \mathcal{Y}} \|\widehat{\theta}_{a,\lambda_h}^{(-k)}(y) - \theta_a^{(d_n)}(y)\|_2 + b_{h,a,n}. \end{aligned} \tag{60}$$

Step 3: Bound the treatment bridge estimation error. Subtracting $\{\widehat{\Sigma}_a^{(-k)}(\widehat{\Sigma}_a^{(-k)})^\top + \lambda_q I_{d_n}\}\alpha_a^{(d_n)}$ from both sides of the treatment bridge normal equation in (44) gives

$$\begin{aligned}&\{\widehat{\Sigma}_a^{(-k)}(\widehat{\Sigma}_a^{(-k)})^\top + \lambda_q I_{d_n}\}\{\widehat{\alpha}_{a,\lambda_q}^{(-k)} - \alpha_a^{(d_n)}\}\\ &\quad = \widehat{\Sigma}_a^{(-k)}\{\widehat{\mu}_W^{(-k)} - (\widehat{\Sigma}_a^{(-k)})^\top \alpha_a^{(d_n)}\} - \lambda_q \alpha_a^{(d_n)}.\end{aligned} \tag{61}$$

Because $\Sigma_a^\top \alpha_a^{(d_n)} = \mu_W$,

$$\begin{aligned}\widehat{\mu}_W^{(-k)} - (\widehat{\Sigma}_a^{(-k)})^\top \alpha_a^{(d_n)} &= \widehat{\mu}_W^{(-k)} - \mu_W\\ &\quad + \{\Sigma_a^\top - (\widehat{\Sigma}_a^{(-k)})^\top\}\alpha_a^{(d_n)}.\end{aligned}$$

Substituting the population equation into (61) and multiplying by the regularized inverse gives

$$\begin{aligned}&\widehat{\alpha}_{a,\lambda_q}^{(-k)} - \alpha_a^{(d_n)}\\ &\quad = \{\widehat{\Sigma}_a^{(-k)}(\widehat{\Sigma}_a^{(-k)})^\top + \lambda_q I_{d_n}\}^{-1}\widehat{\Sigma}_a^{(-k)}\\ &\qquad \times \left\{\widehat{\mu}_W^{(-k)} - \mu_W + \{\Sigma_a^\top - (\widehat{\Sigma}_a^{(-k)})^\top\}\alpha_a^{(d_n)}\right\}\\ &\qquad - \lambda_q\{\widehat{\Sigma}_a^{(-k)}(\widehat{\Sigma}_a^{(-k)})^\top + \lambda_q I_{d_n}\}^{-1}\alpha_a^{(d_n)}.\end{aligned} \tag{62}$$

Transposition preserves the matrix operator norm. Applying (52) and (55) therefore gives

$$\begin{aligned}&\max_k \|\widehat{\alpha}_{a,\lambda_q}^{(-k)} - \alpha_a^{(d_n)}\|_2\\ &\le \left(\sup_{t\ge\kappa_{a,n}/2}\frac{t}{t^2+\lambda_q}\right)\left\{\max_k\|\widehat{\mu}_W^{(-k)} - \mu_W\|_2 + M_{q,a,n}\max_k\|\widehat{\Sigma}_a^{(-k)} - \Sigma_a\|_{\mathrm{op}}\right\}\\ &\quad + \frac{\lambda_q M_{q,a,n}}{\kappa_{a,n}^2/4 + \lambda_q}.\end{aligned} \tag{63}$$

To convert coefficient error into treatment bridge estimation error, use the weighted norm $\|\cdot\|_{a,Z}$ with the estimated coefficients held fixed. By (47),

$$\begin{aligned}&\|\widehat{q}_a^{(-k)}(Z,X) - q_a^{(d_n)}(Z,X)\|_{a,Z}^2\\ &\quad = \{\widehat{\alpha}_{a,\lambda_q}^{(-k)} - \alpha_a^{(d_n)}\}^\top E\{\mathbf{1}\{A=a\}b_Z(Z,X)b_Z(Z,X)^\top\}\\ &\qquad \times \{\widehat{\alpha}_{a,\lambda_q}^{(-k)} - \alpha_a^{(d_n)}\}\\ &\quad \le C_b\|\widehat{\alpha}_{a,\lambda_q}^{(-k)} - \alpha_a^{(d_n)}\|_2^2.\end{aligned} \tag{64}$$

The triangle inequality and the definition of $b_{q,a,n}$ give

$$\begin{aligned}&\max_k\|\widehat{q}_a^{(-k)}(Z,X) - q_a(Z,X)\|_{a,Z}\\ &\quad \le \sqrt{C_b}\max_k\|\widehat{\alpha}_{a,\lambda_q}^{(-k)} - \alpha_a^{(d_n)}\|_2 + b_{q,a,n}.\end{aligned} \tag{65}$$

For later use, the same normal equations and singular value decomposition give the following bounds on every sample, without conditioning on $\mathcal{E}_n$:

$$\begin{aligned}&\|\widehat{\theta}_{a,\lambda_h}^{(-k)}(y) - \theta_a^{(d_n)}(y)\|_2\\ &\quad \le \max_j\frac{\sigma_j(\widehat{\Sigma}_a^{(-k)})}{\sigma_j(\widehat{\Sigma}_a^{(-k)})^2 + \lambda_h}\|\widehat{\gamma}_a^{(-k)}(y) - \widehat{\Sigma}_a^{(-k)}\theta_a^{(d_n)}(y)\|_2 + \frac{\lambda_h\|\theta_a^{(d_n)}(y)\|_2}{\sigma_{\min}(\widehat{\Sigma}_a^{(-k)})^2 + \lambda_h},\\ &\|\widehat{\alpha}_{a,\lambda_q}^{(-k)} - \alpha_a^{(d_n)}\|_2\\ &\quad \le \max_j\frac{\sigma_j(\widehat{\Sigma}_a^{(-k)})}{\sigma_j(\widehat{\Sigma}_a^{(-k)})^2 + \lambda_q}\|\widehat{\mu}_W^{(-k)} - (\widehat{\Sigma}_a^{(-k)})^\top\alpha_a^{(d_n)}\|_2 + \frac{\lambda_q\|\alpha_a^{(d_n)}\|_2}{\sigma_{\min}(\widehat{\Sigma}_a^{(-k)})^2 + \lambda_q}.\end{aligned} \tag{66}$$

Both denominators remain positive for singular sample matrices. The ridge-bias factor is at most one.

Step 4: Substitute the moment-concentration rates. Substituting (48) and (49) into (58) and then (60) gives

$$\begin{aligned}&\max_k\sup_{y\in\mathcal{Y}}\|\widehat{h}_{a,y}^{(-k)}(W,X) - h_{a,y}(W,X)\|_W\\ &\quad = O_p\left(\frac{1+M_{h,a,n}}{\sqrt{\kappa_{a,n}^2+\lambda_h}}\sqrt{\frac{d_n}{n}} + b_{h,a,n} + \frac{\lambda_h M_{h,a,n}}{\kappa_{a,n}^2+\lambda_h}\right).\end{aligned} \tag{67}$$

Substituting (48) into (63) and then (65) gives

$$\max_k \|\widehat{q}_a^{(-k)}(Z,X) - q_a(Z,X)\|_{a,Z} = O_p\left(\frac{1+M_{q,a,n}}{\sqrt{\kappa_{a,n}^2+\lambda_q}}\sqrt{\frac{d_n}{n}} + b_{q,a,n} + \frac{\lambda_q M_{q,a,n}}{\kappa_{a,n}^2+\lambda_q}\right). \tag{68}$$

The factor $1 + M_{q,a,n}$ combines the error in $\widehat{\mu}_W^{(-k)}$ with the error in $\widehat{\Sigma}_a^{(-k)}$ multiplied by $\alpha_a^{(d_n)}$. The bounds hold on $\mathcal{E}_n$, whose probability tends to one, so the bounds in (67) and (68) hold in probability without conditioning on $\mathcal{E}_n$. The inequalities $\lambda_h/(\kappa_{a,n}^2/4+\lambda_h) \le 4\lambda_h/(\kappa_{a,n}^2+\lambda_h)$ and its treatment-bridge counterpart give the displayed ridge terms. No small-ridge condition has been used. These upper bounds do not imply consistency when a right-hand side fails to vanish. The finite-dimensional approximation errors and coefficient norm bounds remain explicit in both rates. The argument uses concentration uniformly over the fixed number of folds and does not require overlapping training samples to be independent. □

A.5 Number of Estimated Outcome Bridges across Outcome Thresholds

The error bounds in Proposition 3.1 control estimation accuracy. The following lemma records a separate consequence of using the same sample matrix and regularization parameter across thresholds.

Lemma A.1 (Number of estimated outcome bridges)**.** *Fix a treatment arm a and a fold I_k. Suppose the basis vectors and $\lambda_h > 0$ are fixed across $y \in \mathcal{Y}$ within the training sample, and use the estimated outcome bridge in (9). Conditionally on the training sample, the class $\{\widehat{h}_{a,y}^{(-k)}(W,X) : y \in \mathcal{Y}\}$ contains at most $N_k + 1$ distinct functions. The L_2 covering number of the estimated outcome bridge class is therefore at most $N_k + 1$ for every positive radius and every finitely supported probability measure. On the envelope event in (97), for every finitely supported probability measure Q and $0 < \epsilon \le 1$,*

$$N\left\{\epsilon H_n, \{\widehat{h}_{a,y}^{(-k)}(W,X) : y \in \mathcal{Y}\}, L_2(\mathsf{Q})\right\} \le \min\left\{N_k + 1, \left(1+\frac{4}{\epsilon}\right)^{d_n}\right\}.$$

Proof. Condition on the training sample $\{O_i : i \notin I_k\}$. As y increases, each coordinate of the indicator vector $\left(\mathbf{1}\{Y_i \le y\}\right)_{i \notin I_k}$ changes from zero to one at most once. The indicator vector changes only at a distinct training outcome, and there are at most N_k distinct training outcomes. The indicator vector therefore takes at most $N_k + 1$ values.

By (39), $\widehat{\gamma}_a^{(-k)}(y)$ is a fixed weighted sum of the coordinates of the indicator vector. The vector $\widehat{\gamma}_a^{(-k)}(y)$ consequently takes at most $N_k + 1$ values. The outcome bridge coefficient in (45) is obtained by multiplying $\widehat{\gamma}_a^{(-k)}(y)$ by a matrix that does not depend on y. The coefficient $\widehat{\theta}_{a,\lambda_h}^{(-k)}(y)$ therefore takes at most $N_k + 1$ values. Since $b_W(W,X)$ is fixed across thresholds, the class $\{\widehat{h}_{a,y}^{(-k)}(W,X) : y \in \mathcal{Y}\}$ contains at most $N_k + 1$ distinct functions.

Use every distinct estimated outcome bridge as a covering center. Every estimated outcome bridge has zero distance from a covering center, so at most $N_k + 1$ centers cover the class at any positive L_2 radius under any finitely supported probability measure.

By (9), the same class lies in the linear span of the d_n components of $b_W(W,X)$. After identifying functions at zero $L_2(\mathsf{Q})$ distance, this span has dimension at most d_n. On the envelope event, every estimated outcome bridge has $L_2(\mathsf{Q})$ norm at most H_n. For $H_n > 0$, a maximal ϵH_n-separated subset of the class is an ϵH_n-net. The disjoint balls of radius $\epsilon H_n/2$ around its elements lie in a ball of radius $(1+\epsilon/2)H_n$ in that finite-dimensional space. Comparing volumes bounds the number of elements by $(1+2/\epsilon)^{d_n} \le (1+4/\epsilon)^{d_n}$. If $H_n = 0$, the class consists of the zero function on the corresponding support and the bound is immediate. Taking the smaller of this bound and $N_k + 1$ proves the stated covering bound. □

The bound on the number of estimated outcome bridges does not bound the magnitude of the estimated bridges. Envelope conditions are imposed separately when establishing the uniform expansion and simultaneous inference.

B Derivation and Proof for Uniform Expansion with Estimated Bridges

This section derives the uniform expansion in Theorem 3.1. The derivation separates the population bias caused by bridge estimation errors from the sampling fluctuation caused by using estimated bridges across outcome thresholds. The assumptions below concern bridge estimation errors, boundedness, and covering entropy.

For each fold I_k, let $n_k = |I_k|$ and define the evaluation-sample average

$$P_{n,k} f = \frac{1}{n_k} \sum_{i \in I_k} f(O_i). \tag{69}$$

Since the folds partition $\{1, \dots, n\}$,

$$\frac{1}{n} \sum_{i=1}^{n} f(O_i) = \sum_{k=1}^{K} \frac{n_k}{n} P_{n,k} f. \tag{70}$$

For each fixed k, condition on the training sample $\{O_i : i \notin I_k\}$. Then $\widehat{h}_{a,y}^{(-k)}(W, X)$ and $\widehat{q}_a^{(-k)}(Z, X)$ are fixed functions, while the observations $\{O_i : i \in I_k\}$ remain independent observations from P_0.

B.1 Derivation of the Population Bias

Fix $a \in \{0, 1\}$, $y \in \mathcal{Y}$, and a fold I_k. Starting from the expression used in (13),

$$\begin{aligned}
&E\Big[\widehat{h}_{a,y}^{(-k)}(W, X) + \mathbf{1}\{A = a\}\widehat{q}_a^{(-k)}(Z, X)\{\mathbf{1}\{Y \le y\} - \widehat{h}_{a,y}^{(-k)}(W, X)\}\Big] - F_a(y) \\
&\quad = E\{\widehat{h}_{a,y}^{(-k)}(W, X) - h_{a,y}(W, X)\} \\
&\qquad + E\Big[\mathbf{1}\{A = a\}\widehat{q}_a^{(-k)}(Z, X)\{\mathbf{1}\{Y \le y\} - h_{a,y}(W, X)\}\Big] \\
&\qquad - E\Big[\mathbf{1}\{A = a\}\widehat{q}_a^{(-k)}(Z, X)\{\widehat{h}_{a,y}^{(-k)}(W, X) - h_{a,y}(W, X)\}\Big].
\end{aligned} \tag{71}$$

The middle term in (71) is zero. Because $\widehat{q}_a^{(-k)}(Z, X)$ is fixed after conditioning on the training sample and is a function of (Z, X),

$$\begin{aligned}
&E\Big[\mathbf{1}\{A = a\}\widehat{q}_a^{(-k)}(Z, X)\{\mathbf{1}\{Y \le y\} - h_{a,y}(W, X)\}\Big] \\
&\quad = E\Big[\mathbf{1}\{A = a\}\widehat{q}_a^{(-k)}(Z, X) E\{\mathbf{1}\{Y \le y\} - h_{a,y}(W, X) \mid A, Z, X\}\Big] = 0,
\end{aligned} \tag{72}$$

where the last equality follows from the outcome bridge equation (1).

The treatment bridge equation (3) gives

$$E\{\mathbf{1}\{A = a\} q_a(Z, X) \mid W, X\} = 1.$$

Therefore,

$$\begin{aligned}
&E\{\widehat{h}_{a,y}^{(-k)}(W, X) - h_{a,y}(W, X)\} \\
&\quad = E\Big[\{\widehat{h}_{a,y}^{(-k)}(W, X) - h_{a,y}(W, X)\} E\{\mathbf{1}\{A = a\} q_a(Z, X) \mid W, X\}\Big] \\
&\quad = E\Big[\mathbf{1}\{A = a\} q_a(Z, X)\{\widehat{h}_{a,y}^{(-k)}(W, X) - h_{a,y}(W, X)\}\Big].
\end{aligned} \tag{73}$$

Substituting (72) and (73) into (71) gives

$$\begin{aligned}
&E\Big[\widehat{h}_{a,y}^{(-k)}(W, X) + \mathbf{1}\{A = a\}\widehat{q}_a^{(-k)}(Z, X)\{\mathbf{1}\{Y \le y\} - \widehat{h}_{a,y}^{(-k)}(W, X)\}\Big] - F_a(y) \\
&\quad = -E\Big[\mathbf{1}\{A = a\}\{\widehat{q}_a^{(-k)}(Z, X) - q_a(Z, X)\}\{\widehat{h}_{a,y}^{(-k)}(W, X) - h_{a,y}(W, X)\}\Big].
\end{aligned} \tag{74}$$

Equation (74) is the population bias identity stated in (14). Applying Cauchy–Schwarz gives

$$\begin{aligned}
&\Big|E\Big[\mathbf{1}\{A = a\}\{\widehat{q}_a^{(-k)}(Z, X) - q_a(Z, X)\}\{\widehat{h}_{a,y}^{(-k)}(W, X) - h_{a,y}(W, X)\}\Big]\Big| \\
&\quad \le \|\widehat{q}_a^{(-k)}(Z, X) - q_a(Z, X)\|_{a,Z} \Big[E\Big\{\mathbf{1}\{A = a\}\{\widehat{h}_{a,y}^{(-k)}(W, X) - h_{a,y}(W, X)\}^2\Big\}\Big]^{1/2} \\
&\quad \le \|\widehat{q}_a^{(-k)}(Z, X) - q_a(Z, X)\|_{a,Z} \|\widehat{h}_{a,y}^{(-k)}(W, X) - h_{a,y}(W, X)\|_W.
\end{aligned} \tag{75}$$

Taking the maximum over $k \in \{1, \dots, K\}$ and the supremum over $y \in \mathcal{Y}$ yields

$$\max_k \sup_{y \in \mathcal{Y}} \Big|E\Big[\widehat{h}_{a,y}^{(-k)}(W, X) + \mathbf{1}\{A = a\}\widehat{q}_a^{(-k)}(Z, X)\{\mathbf{1}\{Y \le y\} - \widehat{h}_{a,y}^{(-k)}(W, X)\}\Big] - F_a(y)\Big| \le r_{h,a,n} r_{q,a,n}. \tag{76}$$

The condition $r_{h,a,n} r_{q,a,n} = o_p(n^{-1/2})$ remains sufficient, but is not required by the sharper bound below.

Retaining the finite-dimensional moment equations. The refinement uses the mixed-bias property and a tested moment error rather than two strong-norm errors, as in mixed-bias and projected-error analyses (Rotnitzky et al., 2021; Kallus et al., 2021; Ghassami et al., 2022; Bennett et al., 2023). Related weak-norm guarantees for neural proxy estimators are developed by Bozkurt et al. (2026). For the population finite-dimensional bridges in (37), the bridge equations and (38) imply

$$\begin{aligned} E\Big[\mathbf{1}\{A=a\} b_Z(Z,X)\{h_{a,y}^{(d_n)}(W,X) - h_{a,y}(W,X)\}\Big] &= 0, \\ E\Big[\mathbf{1}\{A=a\} b_W(W,X)\{q_a^{(d_n)}(Z,X) - q_a(Z,X)\}\Big] &= 0. \end{aligned} \tag{77}$$

Expanding each estimated bridge around the corresponding population finite-dimensional bridge, the cross terms between a coefficient estimation error and an approximation error vanish by (77). Consequently, conditionally on the training sample,

$$\begin{aligned} &E\Big[\mathbf{1}\{A=a\}\{\widehat{q}_a^{(-k)}(Z,X) - q_a(Z,X)\}\{\widehat{h}_{a,y}^{(-k)}(W,X) - h_{a,y}(W,X)\}\Big] \\ &\quad = \{\widehat{\alpha}_{a,\lambda_q}^{(-k)} - \alpha_a^{(d_n)}\}^\top \Sigma_a \{\widehat{\theta}_{a,\lambda_h}^{(-k)}(y) - \theta_a^{(d_n)}(y)\} \\ &\qquad + E\Big[\mathbf{1}\{A=a\}\{q_a^{(d_n)}(Z,X) - q_a(Z,X)\}\{h_{a,y}^{(d_n)}(W,X) - h_{a,y}(W,X)\}\Big]. \end{aligned} \tag{78}$$

The population bias in (74) is the negative of (78). Cauchy–Schwarz applied after retaining Σ_a gives

$$\begin{aligned} &\max_k \sup_{y\in\mathcal{Y}} \Big| E\Big[\mathbf{1}\{A=a\}\{\widehat{q}_a^{(-k)}(Z,X) - q_a(Z,X)\}\{\widehat{h}_{a,y}^{(-k)}(W,X) - h_{a,y}(W,X)\}\Big]\Big| \\ &\le \max_k \Big\|\Sigma_a^\top\{\widehat{\alpha}_{a,\lambda_q}^{(-k)} - \alpha_a^{(d_n)}\}\Big\|_2 \max_k \sup_{y\in\mathcal{Y}} \Big\|\widehat{\theta}_{a,\lambda_h}^{(-k)}(y) - \theta_a^{(d_n)}(y)\Big\|_2 + b_{h,a,n} b_{q,a,n}. \end{aligned} \tag{79}$$

No independence between the outcome and treatment coefficient estimates is used in (79).

Bounding the tested treatment coefficient error. The population equation $\Sigma_a^\top \alpha_a^{(d_n)} = \mu_W$ gives the exact identity

$$\begin{aligned} \Sigma_a^\top\{\widehat{\alpha}_{a,\lambda_q}^{(-k)} - \alpha_a^{(d_n)}\} &= (\Sigma_a - \widehat{\Sigma}_a^{(-k)})^\top \widehat{\alpha}_{a,\lambda_q}^{(-k)} \\ &\quad + (\widehat{\Sigma}_a^{(-k)})^\top \widehat{\alpha}_{a,\lambda_q}^{(-k)} - \widehat{\mu}_W^{(-k)} + \widehat{\mu}_W^{(-k)} - \mu_W. \end{aligned} \tag{80}$$

For a bound that retains the factor $M_{q,a,n}$ in the ridge term, multiply (62) by $(\widehat{\Sigma}_a^{(-k)})^\top$. On the event $\mathcal{E}_n$ in (50),

$$\Big\|(\widehat{\Sigma}_a^{(-k)})^\top\{\widehat{\Sigma}_a^{(-k)}(\widehat{\Sigma}_a^{(-k)})^\top + \lambda_q I_{d_n}\}^{-1}\widehat{\Sigma}_a^{(-k)}\Big\|_{\mathrm{op}} \le 1,$$

and the transpose of the treatment matrix bound in (55) has norm at most $2/\kappa_{a,n}$. Thus

$$\begin{aligned} &\Big\|\Sigma_a^\top\{\widehat{\alpha}_{a,\lambda_q}^{(-k)} - \alpha_a^{(d_n)}\}\Big\|_2 \\ &\quad \le \|\widehat{\mu}_W^{(-k)} - \mu_W\|_2 + M_{q,a,n}\|\widehat{\Sigma}_a^{(-k)} - \Sigma_a\|_{\mathrm{op}} + \frac{2\lambda_q M_{q,a,n}}{\kappa_{a,n}} \\ &\qquad + \|\widehat{\Sigma}_a^{(-k)} - \Sigma_a\|_{\mathrm{op}}\|\widehat{\alpha}_{a,\lambda_q}^{(-k)} - \alpha_a^{(d_n)}\|_2. \end{aligned} \tag{81}$$

Using (48), (63), and $\sqrt{d_n/n} = o(\kappa_{a,n})$, the last term in (81) is of smaller order than the preceding sampling and ridge bounds. Since $P_0(\mathcal{E}_n) \to 1$,

$$\max_k \Big\|\Sigma_a^\top\{\widehat{\alpha}_{a,\lambda_q}^{(-k)} - \alpha_a^{(d_n)}\}\Big\|_2 = O_p\left((1 + M_{q,a,n})\sqrt{\frac{d_n}{n}} + \frac{\lambda_q M_{q,a,n}}{\kappa_{a,n}}\right). \tag{82}$$

The sampling term in (82) has no factor $\kappa_{a,n}^{-1}$.

Proof of Corollary 3.1. Combining (82) with the outcome coefficient bound in (58) and the moment concentration rates gives

$$\begin{aligned} &\max_k \Big\|\Sigma_a^\top\{\widehat{\alpha}_{a,\lambda_q}^{(-k)} - \alpha_a^{(d_n)}\}\Big\|_2 \max_k \sup_{y\in\mathcal{Y}} \Big\|\widehat{\theta}_{a,\lambda_h}^{(-k)}(y) - \theta_a^{(d_n)}(y)\Big\|_2 \\ &= O_p\left[\left\{(1 + M_{q,a,n})\sqrt{\frac{d_n}{n}} + \frac{\lambda_q M_{q,a,n}}{\kappa_{a,n}}\right\}\left\{\frac{1 + M_{h,a,n}}{\kappa_{a,n}}\sqrt{\frac{d_n}{n}} + \frac{\lambda_h M_{h,a,n}}{\kappa_{a,n}^2}\right\}\right]. \end{aligned} \tag{83}$$

Under the bounded coefficient norm condition and the $o\{\kappa_{a,n}\sqrt{d_n/n}\}$ ridge condition in Corollary 3.1, the first factor in (83) is $O_p(\sqrt{d_n/n})$ and the second is $O_p\{\sqrt{d_n/n}/\kappa_{a,n}\}$. Thus their

product is $O_p\{d_n/(n\kappa_{a,n})\}$. Equation (79) then proves (16). Multiplication by $\sqrt{n}$ and the fixed number of treatment arms prove the stated sufficient condition for negligible population bias. Approximation errors enter (79) only through $b_{h,a,n}b_{q,a,n}$, but individual approximation errors still enter the strong bridge errors used in the empirical-process bounds. The argument retains $\sqrt{d_n/n} = o(\kappa_{a,n})$ and does not establish inference when $\kappa_{a,n}$ is of the same order as the sample matrix error.

Retaining both ridge parameters in the population-bias bound. The coefficient product in (78) can also be bounded before separating the two regularized matrix factors. This calculation uses exactly the estimates in (8). In particular,

$$\begin{aligned}
&\{\widehat{\alpha}_{a,\lambda_q}^{(-k)} - \alpha_a^{(d_n)}\}^\top \Sigma_a \{\widehat{\theta}_{a,\lambda_h}^{(-k)}(y) - \theta_a^{(d_n)}(y)\} \\
&= \left[\widehat{\Sigma}_a^{(-k)}\{\widehat{\mu}_W^{(-k)} - (\widehat{\Sigma}_a^{(-k)})^\top \alpha_a^{(d_n)}\} - \lambda_q \alpha_a^{(d_n)}\right]^\top \\
&\quad \times \{\widehat{\Sigma}_a^{(-k)}(\widehat{\Sigma}_a^{(-k)})^\top + \lambda_q I_{d_n}\}^{-1} \widehat{\Sigma}_a^{(-k)} \{(\widehat{\Sigma}_a^{(-k)})^\top \widehat{\Sigma}_a^{(-k)} + \lambda_h I_{d_n}\}^{-1} \\
&\quad \times \left[(\widehat{\Sigma}_a^{(-k)})^\top \{\widehat{\gamma}_a^{(-k)}(y) - \widehat{\Sigma}_a^{(-k)} \theta_a^{(d_n)}(y)\} - \lambda_h \theta_a^{(d_n)}(y)\right] \\
&\quad - \{\widehat{\alpha}_{a,\lambda_q}^{(-k)} - \alpha_a^{(d_n)}\}^\top (\widehat{\Sigma}_a^{(-k)} - \Sigma_a)\{\widehat{\theta}_{a,\lambda_h}^{(-k)}(y) - \theta_a^{(d_n)}(y)\}.
\end{aligned} \tag{84}$$

Equation (84) follows by substituting (57) and (62) into the product containing $\widehat{\Sigma}_a^{(-k)}$, and then subtracting the product containing $\widehat{\Sigma}_a^{(-k)} - \Sigma_a$. Expanding the two square brackets gives four matrix products. Singular value decomposition shows that their operator norms are the maxima of the fractions with numerators σ_j^3, σ_j^2, σ_j^2, and σ_j, respectively, and denominator $(\sigma_j^2 + \lambda_h)(\sigma_j^2 + \lambda_q)$, with every singular value evaluated at $\widehat{\Sigma}_a^{(-k)}$. Thus

$$\begin{aligned}
&\max_k \sup_{y \in \mathcal{Y}} \left| E\left[\mathbf{1}\{A = a\}\{\widehat{q}_a^{(-k)}(Z,X) - q_a(Z,X)\}\{\widehat{h}_{a,y}^{(-k)}(W,X) - h_{a,y}(W,X)\}\right]\right| \\
&\le b_{h,a,n} b_{q,a,n} + \max_k \sup_{y \in \mathcal{Y}} \Bigg[\\
&\quad \max_j \frac{\sigma_j(\widehat{\Sigma}_a^{(-k)})^3}{\{\sigma_j(\widehat{\Sigma}_a^{(-k)})^2 + \lambda_h\}\{\sigma_j(\widehat{\Sigma}_a^{(-k)})^2 + \lambda_q\}} \\
&\qquad \times \|\widehat{\mu}_W^{(-k)} - (\widehat{\Sigma}_a^{(-k)})^\top \alpha_a^{(d_n)}\|_2 \|\widehat{\gamma}_a^{(-k)}(y) - \widehat{\Sigma}_a^{(-k)} \theta_a^{(d_n)}(y)\|_2 \\
&\quad + \max_j \frac{\sigma_j(\widehat{\Sigma}_a^{(-k)})^2}{\{\sigma_j(\widehat{\Sigma}_a^{(-k)})^2 + \lambda_h\}\{\sigma_j(\widehat{\Sigma}_a^{(-k)})^2 + \lambda_q\}} \\
&\qquad \times \Big\{\lambda_h \|\theta_a^{(d_n)}(y)\|_2 \|\widehat{\mu}_W^{(-k)} - (\widehat{\Sigma}_a^{(-k)})^\top \alpha_a^{(d_n)}\|_2 \\
&\qquad\qquad + \lambda_q \|\alpha_a^{(d_n)}\|_2 \|\widehat{\gamma}_a^{(-k)}(y) - \widehat{\Sigma}_a^{(-k)} \theta_a^{(d_n)}(y)\|_2\Big\} \\
&\quad + \lambda_h \lambda_q \max_j \frac{\sigma_j(\widehat{\Sigma}_a^{(-k)})}{\{\sigma_j(\widehat{\Sigma}_a^{(-k)})^2 + \lambda_h\}\{\sigma_j(\widehat{\Sigma}_a^{(-k)})^2 + \lambda_q\}} \|\theta_a^{(d_n)}(y)\|_2 \|\alpha_a^{(d_n)}\|_2 \\
&\quad + \|\widehat{\Sigma}_a^{(-k)} - \Sigma_a\|_{\mathrm{op}} \|\widehat{\alpha}_{a,\lambda_q}^{(-k)} - \alpha_a^{(d_n)}\|_2 \|\widehat{\theta}_{a,\lambda_h}^{(-k)}(y) - \theta_a^{(d_n)}(y)\|_2 \Bigg].
\end{aligned} \tag{85}$$

The coefficient errors in the last line can be replaced directly by (66). No favorable alignment of the two errors or independence of the two estimated bridges is required.

Both (84) and (85) hold for all positive ridge parameters, including singular sample matrices. The bound (79) also remains valid. Taking the smaller of the right-hand sides of (79) and (85) therefore never enlarges the population-bias bound in (79). The bound (85) alone need not be smaller on every sample.

For example, when $\lambda_h = \lambda_q$, elementary differentiation gives

$$\sup_{t \ge 0} \frac{t^3}{(t^2 + \lambda_h)^2} = \frac{3\sqrt{3}}{16\sqrt{\lambda_h}},$$

with the maximum attained at $t^2 = 3\lambda_h$. On $\mathcal{E}_n$,

$$\max_j \frac{\sigma_j(\widehat{\Sigma}_a^{(-k)})^3}{\{\sigma_j(\widehat{\Sigma}_a^{(-k)})^2 + \lambda_h\}^2} \le \min\left\{\frac{2}{\kappa_{a,n}}, \frac{3\sqrt{3}}{16\sqrt{\lambda_h}}\right\}. \tag{86}$$

The sampling-error product can consequently have a smaller amplification factor when ridge is not negligible. The ridge-bias and matrix-error terms in (85) cannot be omitted. These deterministic bounds do not establish root-n inference when $\kappa_{a,n}$ is of the same order as the sample matrix error.

A second-order expansion of the population bias. Consider the subcase of Corollary 3.1 satisfying

$$\lambda_h, \lambda_q = o(\kappa_{a,n}^2), \qquad \lambda_h M_{h,a,n} + \lambda_q M_{q,a,n} = o\left(\kappa_{a,n}\sqrt{\frac{d_n}{n}}\right). \tag{87}$$

The ridge conditions in (87) allow ridge terms to be omitted from the leading signed term.

The population moment equations give the exact centered-moment identities

$$\begin{aligned}\widehat{\gamma}_a^{(-k)}(y) - \widehat{\Sigma}_a^{(-k)}\theta_a^{(d_n)}(y) &= (P_{n,-k} - E)\left[\mathbf{1}\{A=a\}b_Z(Z,X)\{\mathbf{1}\{Y\le y\} - h_{a,y}^{(d_n)}(W,X)\}\right],\\ \widehat{\mu}_W^{(-k)} - (\widehat{\Sigma}_a^{(-k)})^\top \alpha_a^{(d_n)} &= (P_{n,-k} - E)\left[b_W(W,X)\{1 - \mathbf{1}\{A=a\}q_a^{(d_n)}(Z,X)\}\right].\end{aligned} \tag{88}$$

By (48), (49), and bounded $M_{h,a,n}, M_{q,a,n}$, both centered vectors are $O_p(\sqrt{d_n/n})$, uniformly over folds and, for the outcome vector, thresholds.

On $\mathcal{E}_n$, subtracting the sample inverse from the population inverse gives

$$\|\{\widehat{\Sigma}_a^{(-k)}\}^{-1} - \Sigma_a^{-1}\|_{\mathrm{op}} \le \frac{2}{\kappa_{a,n}^2}\|\widehat{\Sigma}_a^{(-k)} - \Sigma_a\|_{\mathrm{op}}.$$

The difference between the outcome ridge matrix factor and the sample inverse has operator norm at most $8\lambda_h/\kappa_{a,n}^3$. Applying these two bounds to (57) gives

$$\begin{aligned}&\max_k \sup_{y\in\mathcal{Y}} \left\|\widehat{\theta}_{a,\lambda_h}^{(-k)}(y) - \theta_a^{(d_n)}(y) - \Sigma_a^{-1}\{\widehat{\gamma}_a^{(-k)}(y) - \widehat{\Sigma}_a^{(-k)}\theta_a^{(d_n)}(y)\}\right\|_2\\ &\quad = O_p\left(\frac{d_n}{n\kappa_{a,n}^2} + \frac{\lambda_h}{\kappa_{a,n}^3}\sqrt{\frac{d_n}{n}} + \frac{\lambda_h M_{h,a,n}}{\kappa_{a,n}^2}\right) = o_p\left(\frac{1}{\kappa_{a,n}}\sqrt{\frac{d_n}{n}}\right).\end{aligned} \tag{89}$$

For the treatment coefficient error, multiply (62) by $(\widehat{\Sigma}_a^{(-k)})^\top$ and use

$$(\widehat{\Sigma}_a^{(-k)})^\top\{\widehat{\Sigma}_a^{(-k)}(\widehat{\Sigma}_a^{(-k)})^\top + \lambda_q I_{d_n}\}^{-1}\widehat{\Sigma}_a^{(-k)} = I_{d_n} - \lambda_q\{(\widehat{\Sigma}_a^{(-k)})^\top\widehat{\Sigma}_a^{(-k)} + \lambda_q I_{d_n}\}^{-1}.$$

The remaining difference between $\Sigma_a^\top$ and $(\widehat{\Sigma}_a^{(-k)})^\top$ multiplies the coefficient error. Consequently,

$$\begin{aligned}&\max_k \left\|\Sigma_a^\top\{\widehat{\alpha}_{a,\lambda_q}^{(-k)} - \alpha_a^{(d_n)}\} - \{\widehat{\mu}_W^{(-k)} - (\widehat{\Sigma}_a^{(-k)})^\top\alpha_a^{(d_n)}\}\right\|_2\\ &\quad = O_p\left(\frac{\lambda_q}{\kappa_{a,n}^2}\sqrt{\frac{d_n}{n}} + \frac{\lambda_q M_{q,a,n}}{\kappa_{a,n}} + \frac{d_n}{n\kappa_{a,n}}\right) = o_p\left(\sqrt{\frac{d_n}{n}}\right).\end{aligned} \tag{90}$$

The last equalities in (89) and (90) follow from (87) and $\sqrt{d_n/n} = o(\kappa_{a,n})$. Substituting both expansions into (78) yields

$$\begin{aligned}&E\left[\mathbf{1}\{A=a\}\{\widehat{q}_a^{(-k)}(Z,X) - q_a(Z,X)\}\{\widehat{h}_{a,y}^{(-k)}(W,X) - h_{a,y}(W,X)\}\right]\\ &= \{\widehat{\mu}_W^{(-k)} - (\widehat{\Sigma}_a^{(-k)})^\top\alpha_a^{(d_n)}\}^\top\Sigma_a^{-1}\{\widehat{\gamma}_a^{(-k)}(y) - \widehat{\Sigma}_a^{(-k)}\theta_a^{(d_n)}(y)\}\\ &\quad + E\left[\mathbf{1}\{A=a\}\{q_a^{(d_n)}(Z,X) - q_a(Z,X)\}\{h_{a,y}^{(d_n)}(W,X) - h_{a,y}(W,X)\}\right] + o_p\left(\frac{d_n}{n\kappa_{a,n}}\right).\end{aligned} \tag{91}$$

The remainder is uniform over thresholds and the fixed number of folds. Both treatment arms satisfy the same statement. The signed population bias in (14) is the negative of the right-hand side of (91). The leading term pairs the two centered training-sample moment errors through exactly one Σ_a^{-1}. The two centered errors use the same training observations and are not assumed independent.

Sharpness of the $\kappa_{a,n}^{-1}$ factor for the specified estimates. We construct a sequence of data distributions satisfying the conditions of Corollary 3.1. Let X be constant, let $U \in \{-1, 1\}$ have equal probabilities, and, for $n > 16$, generate $W, Z \in \{-1, 1\}$ independently conditional on U with

$$P(W = w \mid U = u) = \frac{1 + uw/2}{2}, \qquad P(Z = z \mid U = u) = \frac{1 + 2n^{-1/4}uz}{2}.$$

Then

$$P(W = w, Z = z) = \frac{1 + n^{-1/4}wz}{4}.$$

Generate $A \sim \text{Bernoulli}(1/2)$ independently of (U, W, Z). Independently generate $Y(0)$ and $Y(1)$ from the uniform distribution on $[0, 1]$, and set $Y = AY(1) + (1 - A)Y(0)$. Latent positivity and all three conditional independence statements in Section 2 follow from this construction. Also, the two conditional distributions of U given Z are distinct, so the corresponding binary conditional expectation matrix is invertible at every $n > 16$. The marginal counterfactual distribution is $F_a(y) = y$ on $[0, 1]$.

Choose $\mathcal{Y} = [1/4, 3/4]$, $d_n = 2$, $b_W(W, X) = (1, W)^\top$, and $b_Z(Z, X) = (1, Z)^\top$. For both treatment arms,

$$\Sigma_a = \frac{1}{2}\begin{pmatrix} 1 & 0 \\ 0 & n^{-1/4} \end{pmatrix}, \qquad \kappa_{a,n} = \frac{n^{-1/4}}{2}, \qquad \theta_a^{(2)}(y) = \begin{pmatrix} y \\ 0 \end{pmatrix}, \qquad \alpha_a^{(2)} = \begin{pmatrix} 2 \\ 0 \end{pmatrix}. \tag{92}$$

Thus $h_{a,y}(W, X) = y$, $q_a(Z, X) = 2$, $b_{h,a,n} = b_{q,a,n} = 0$, and the population coefficient norms remain bounded. The two Gram matrices are I_2 and $I_2/2$. Bounded fixed-dimensional observations give (48), and $\sqrt{d_n/n} = o(\kappa_{a,n})$. With $\lambda_h = \lambda_q = 0.1/N_k$ in each training sample, (87) holds because $N_k \asymp n$.

Fix $a = 1$, $y = 1/2$, and one fold I_k. The second coordinates of the centered vectors in (88) are

$$P_{n,-k}\{W(1 - 2A)\}, \qquad P_{n,-k}[AZ\{\mathbf{1}\{Y \le 1/2\} - 1/2\}].$$

Both summands have mean zero. Their variances are 1 and $1/8$, and their covariance is zero because Y is independent of (A, W, Z) and $E\{\mathbf{1}\{Y \le 1/2\} - 1/2\} = 0$. The bounded triangular-array central limit theorem gives

$$\sqrt{N_k}\begin{pmatrix} P_{n,-k}\{W(1 - 2A)\} \\ P_{n,-k}[AZ\{\mathbf{1}\{Y \le 1/2\} - 1/2\}] \end{pmatrix} \rightsquigarrow \begin{pmatrix} G_q \\ G_h \end{pmatrix} \sim \mathcal{N}\left(0, \begin{pmatrix} 1 & 0 \\ 0 & 1/8 \end{pmatrix}\right). \tag{93}$$

Here G_q and G_h denote the two coordinates of this limiting vector. The first coordinates of the two moment vectors contribute $O_p(N_k^{-1})$ to (91), which is negligible after multiplication by $N_k\kappa_{1,n}$. The second diagonal entry of Σ_1^{-1} equals $\kappa_{1,n}^{-1}$. Hence

$$N_k\kappa_{1,n}E\Big[\mathbf{1}\{A = 1\}\{\widehat{q}_1^{(-k)}(Z, X) - q_1(Z, X)\}\{\widehat{h}_{1,1/2}^{(-k)}(W, X) - h_{1,1/2}(W, X)\}\Big] \rightsquigarrow G_qG_h. \tag{94}$$

The limit is not concentrated at zero. Thus the absolute population bias is not $o_p\{(N_k\kappa_{1,n})^{-1}\}$, even for this fixed dimension and zero approximation error. Since $N_k \asymp n$ and the uniform maximum dominates the fixed fold and threshold, the factor $\kappa_{a,n}^{-1}$ in Corollary 3.1 cannot be replaced uniformly by a smaller order for these estimates. This example does not establish optimality of the factor d_n when d_n grows, or a lower bound over all estimators.

The true bridges in the construction also satisfy Assumption B.1. Proposition 3.1 gives bridge errors $O_p(n^{-1/4})$ and, because the basis vectors are bounded, fixed estimated-bridge envelopes are valid with probability tending to one. The population-bias bound is $O_p(n^{-3/4})$. Taking $\delta_n = n^{-1/8}$ and fixed sufficiently large H_n, Q_n verifies the remaining rate conditions in Assumption B.2. On $\mathcal{Y} = [1/4, 3/4]$, $E\{\varphi_{a,y}^2(O)\} = 2y(1 - y)$ is bounded away from zero.

B.2 Conditions for Uniform Control across Outcome Thresholds

For a class of functions $\mathcal{F}$, let

$$N\{\epsilon, \mathcal{F}, L_2(\mathsf{Q})\}$$

denote the $L_2(\mathsf{Q})$ covering number at radius ϵ, where Q is a finitely supported probability measure.

Assumption B.1 (Bridge function-class conditions)**.** For each treatment arm $a \in \{0, 1\}$, the following conditions hold.

(i) There exist fixed constants $H, Q < \infty$ such that

$$\sup_{y \in \mathcal{Y}} |h_{a,y}(W, X)| \le H$$

almost surely and

$$|q_a(Z, X)| \le Q$$

almost surely on $\{A = a\}$.

(ii) The class $\{h_{a,y}(W, X) : y \in \mathcal{Y}\}$ is pointwise measurable and there exist fixed constants $A_{\text{ent}}, v < \infty$ such that, for every finitely supported probability measure Q,

$$\log N\{\epsilon H, \{h_{a,y}(W, X) : y \in \mathcal{Y}\}, L_2(\mathsf{Q})\} \le v \log(A_{\text{ent}}/\epsilon), \qquad 0 < \epsilon \le 1. \tag{95}$$

Assumption B.2 (Cross-fitted bridge estimation conditions)**.** The number of folds K is fixed, the partition is independent of $O_1, \ldots, O_n$, and n_k/n is bounded away from zero for every k. The estimated bridges are measurable with respect to the corresponding training sample. The following conditions hold for both treatment arms and uniformly over folds.

(i) For the finite-dimensional bridge estimates in (9), the bridge estimation errors in (12) satisfy $r_{h,a,n} = o_p(1)$ and $r_{q,a,n} = o_p(1)$, and

$$\max_k \left\| \Sigma_a^\top \{\widehat{\alpha}_{a,\lambda_q}^{(-k)} - \alpha_a^{(d_n)}\} \right\|_2 \max_k \sup_{y \in \mathcal{Y}} \left\| \widehat{\theta}_{a,\lambda_h}^{(-k)}(y) - \theta_a^{(d_n)}(y) \right\|_2 + b_{h,a,n} b_{q,a,n} = o_p(n^{-1/2}). \tag{96}$$

(ii) There exist deterministic sequences $H_n \geq H$ and $Q_n \geq Q$ such that, on training-sample events whose probabilities tend to one,

$$\sup_{y \in \mathcal{Y}} |\widehat{h}_{a,y}^{(-k)}(W, X)| \leq H_n, \qquad |\widehat{q}_a^{(-k)}(Z, X)| \leq Q_n \tag{97}$$

almost surely on the corresponding supports. For these estimates, the covering bound is

$$\sup_{\mathsf{Q}} \log N\Big\{ \epsilon H_n, \{\widehat{h}_{a,y}^{(-k)}(W, X) : y \in \mathcal{Y}\}, L_2(\mathsf{Q}) \Big\} \leq \min\{\log(N_k + 1),\, d_n \log(1 + 4/\epsilon)\}, \qquad 0 < \epsilon \leq 1. \tag{98}$$

(iii) Define

$$B_n = (1 + Q_n)(H_n + H). \tag{99}$$

There exists a deterministic sequence $0 \leq \delta_n \leq \min\{1, B_n\}$ with $\delta_n \to 0$ such that

$$\max_{a \in \{0,1\}} (1 + Q_n) r_{h,a,n} \leq \delta_n \tag{100}$$

with probability tending to one. For a sufficiently large fixed constant $C \geq e$ and $\delta_n > 0$, define

$$\Lambda_n = 1 + \min\left\{ \log(n + 1),\, d_n \log\left(\frac{CB_n}{\delta_n}\right) \right\} + v \log\left(\frac{CB_n}{\delta_n}\right). \tag{101}$$

When $\delta_n = 0$, set $\Lambda_n = 0$. The quantities B_n, δ_n, and Λ_n control the first term in (104); the treatment bridge estimation error is controlled separately by $r_{q,a,n} = o_p(1)$ in (i). The sequences satisfy

$$\delta_n \sqrt{\Lambda_n} + \frac{B_n \Lambda_n}{\sqrt{n}} \longrightarrow 0. \tag{102}$$

Proposition 3.1 provides explicit sufficient rates for $r_{h,a,n}$ and $r_{q,a,n}$. The bound (83) and Corollary 3.1, with its stated ridge condition, provide sufficient conditions for (96), without requiring $r_{h,a,n} r_{q,a,n} = o_p(n^{-1/2})$. The deterministic bound (85) is an additional finite-sample refinement. In addition, Lemma A.1 gives (98) by combining the at most $N_k + 1$ distinct functions with their at most d_n-dimensional linear span. Thus this covering bound follows from the specified estimates and (97).

B.3 Derivation of the Estimated-Bridge Sampling Fluctuation

The function $\varphi_{a,y}(O)$ in (18) has mean zero. Indeed,

$$E\{\varphi_{a,y}(O)\} = E\{h_{a,y}(W, X) - F_a(y)\} + E\Big[\mathbf{1}\{A = a\} q_a(Z, X)\{\mathbf{1}\{Y \leq y\} - h_{a,y}(W, X)\}\Big] = 0, \tag{103}$$

where the first term is zero by (2) and the second term is zero by (1).

For a fixed fold I_k, define $d_{a,y}^{(-k)}(O)$ by subtracting (18) from the corresponding expression with estimated bridges:

$$\begin{aligned}
&\Big[\widehat{h}_{a,y}^{(-k)}(W, X) - F_a(y) + \mathbf{1}\{A = a\} \widehat{q}_a^{(-k)}(Z, X)\{\mathbf{1}\{Y \leq y\} - \widehat{h}_{a,y}^{(-k)}(W, X)\}\Big] - \varphi_{a,y}(O) \\
&= \{\widehat{h}_{a,y}^{(-k)}(W, X) - h_{a,y}(W, X)\} \\
&\quad + \mathbf{1}\{A = a\}\Big[\widehat{q}_a^{(-k)}(Z, X)\{\mathbf{1}\{Y \leq y\} - \widehat{h}_{a,y}^{(-k)}(W, X)\} - q_a(Z, X)\{\mathbf{1}\{Y \leq y\} - h_{a,y}(W, X)\}\Big] \\
&= \{1 - \mathbf{1}\{A = a\} \widehat{q}_a^{(-k)}(Z, X)\}\{\widehat{h}_{a,y}^{(-k)}(W, X) - h_{a,y}(W, X)\} \\
&\quad + \mathbf{1}\{A = a\}\{\widehat{q}_a^{(-k)}(Z, X) - q_a(Z, X)\}\{\mathbf{1}\{Y \leq y\} - h_{a,y}(W, X)\}.
\end{aligned} \tag{104}$$

Equation (104) shows separately how the two bridge estimation errors enter the sampling fluctuation.

Under (97) and Assumption B.1(i), the absolute value of the first term on the right side of (104) is at most

$$(1+Q_n)(H_n+H),$$

which is B_n in (99). The second term has the smaller, training-sample-dependent envelope

$$(1+H)\mathbf{1}\{A=a\}|\widehat{q}_a^{(-k)}(Z,X)-q_a(Z,X)| \le (Q_n+Q)(1+H).$$

Thus the whole difference is bounded by $B_n+(Q_n+Q)(1+H)$, but only its first term will use the bounded-envelope local maximal inequality.

The $L_2(P_0)$ norm of the first term on the right side of (104) satisfies

$$\Big[E\Big[\{1-\mathbf{1}\{A=a\}\widehat{q}_a^{(-k)}(Z,X)\}^2\{\widehat{h}_{a,y}^{(-k)}(W,X)-h_{a,y}(W,X)\}^2\Big]\Big]^{1/2}$$
$$\le (1+Q_n)\|\widehat{h}_{a,y}^{(-k)}(W,X)-h_{a,y}(W,X)\|_W. \tag{105}$$

The $L_2(P_0)$ norm of the second term satisfies

$$\Big[E\Big[\mathbf{1}\{A=a\}\{\widehat{q}_a^{(-k)}(Z,X)-q_a(Z,X)\}^2\{\mathbf{1}\{Y\le y\}-h_{a,y}(W,X)\}^2\Big]\Big]^{1/2}$$
$$\le (1+H)\|\widehat{q}_a^{(-k)}(Z,X)-q_a(Z,X)\|_{a,Z}. \tag{106}$$

Combining (105) and (106), then taking the maximum over folds and the supremum over thresholds, gives

$$\max_k \sup_{y\in\mathcal{Y}}\Big[E\{d_{a,y}^{(-k)}(O)^2\}\Big]^{1/2} \le (1+Q_n)r_{h,a,n}+(1+H)r_{q,a,n} \le \delta_n+(1+H)r_{q,a,n} \tag{107}$$

with probability tending to one.

It remains to control the two terms in (104) separately as y varies. For the second term, the class $\{\mathbf{1}\{Y\le y\}-h_{a,y}(W,X): y\in\mathcal{Y}\}$ contains only the threshold indicator and the true outcome bridge. The VC covering bound for the threshold indicators and (95) give

$$\sup_{\mathsf{Q}} \log N\{\epsilon(1+H),\{\mathbf{1}\{Y\le y\}-h_{a,y}(W,X): y\in\mathcal{Y}\},L_2(\mathsf{Q})\}$$
$$\le C(v+1)\log(C/\epsilon), \qquad 0<\epsilon\le 1. \tag{108}$$

Conditionally on the training sample, multiplication by $\mathbf{1}\{A=a\}\{\widehat{q}_a^{(-k)}(Z,X)-q_a(Z,X)\}$ preserves this polynomial entropy when the covering radius is normalized by the $L_2(\mathsf{Q})$ norm of the second term's envelope. Indeed, reweight any finitely supported Q by $\mathbf{1}\{A=a\}\{\widehat{q}_a^{(-k)}(Z,X)-q_a(Z,X)\}^2$ and normalize to a probability measure; weighted $L_2(\mathsf{Q})$ distances become unweighted distances under that measure, multiplied by the norm of the weight. If the normalizing constant is zero, every weighted function is zero Q-almost surely. The resulting uniform entropy integral over $(0,1)$ is bounded by a fixed constant. The uniform entropy maximal inequality (van der Vaart and Wellner, 1996, Section 2.14) therefore gives

$$E\Bigg[\sup_{y\in\mathcal{Y}}\Big|\sqrt{n_k}(P_{n,k}-E)\Big[\mathbf{1}\{A=a\}\{\widehat{q}_a^{(-k)}(Z,X)-q_a(Z,X)\}$$
$$\times\{\mathbf{1}\{Y\le y\}-h_{a,y}(W,X)\}\Big]\Big| \;\Big|\; \text{training sample}\Bigg]$$
$$\le C\|\widehat{q}_a^{(-k)}(Z,X)-q_a(Z,X)\|_{a,Z}.$$

In these maximal bounds, C depends only on the fixed true-bridge envelope and entropy constants, not on the number or dimension of the estimated outcome bridges. In particular, no estimated outcome bridge complexity multiplies the treatment bridge estimation error.

For the first term, the fixed multiplier $1-\mathbf{1}\{A=a\}\widehat{q}_a^{(-k)}(Z,X)$ is bounded in absolute value by $1+Q_n$. Covering the estimated and true outcome bridges separately, using (98) and (95), yields

$$\sup_{\mathsf{Q}} \log N\Big\{\epsilon B_n,\ \{(1-\mathbf{1}\{A=a\}\widehat{q}_a^{(-k)}(Z,X))(\widehat{h}_{a,y}^{(-k)}(W,X)-h_{a,y}(W,X)): y\in\mathcal{Y}\},\ L_2(\mathsf{Q})\Big\}$$
$$\le \min\{\log(N_k+1),\ d_n\log(C/\epsilon)\}+v\log(C/\epsilon), \qquad 0<\epsilon\le 1. \tag{109}$$

Condition on one training sample and use its event on which the estimated bridge bounds and $(1+Q_n)\sup_{y\in\mathcal{Y}}\|\widehat{h}_{a,y}^{(-k)}-h_{a,y}\|_W \le \delta_n$ hold. The evaluation observations remain iid from P_0.

For $\delta_n > 0$, apply the bounded-envelope local maximal inequality of van der Vaart and Wellner (2011, Theorem 2.1) only to the first term in (104), using B_n, δ_n, and (109). Combining it with the preceding treatment bridge bound gives

$$E\left[\sup_{y\in\mathcal{Y}}\left|\sqrt{n_k}(P_{n,k}-E)d_{a,y}^{(-k)}\right|\,\middle|\,\text{training sample}\right]$$
$$\le C\left\{\|\widehat{q}_a^{(-k)}(Z,X)-q_a(Z,X)\|_{a,Z}+\delta_n\sqrt{\Lambda_n}+\frac{B_n\Lambda_n}{\sqrt{n_k}}\right\}. \tag{110}$$

The difference $d_{a,y}^{(-k)}(O)$ is defined in (104). To see the dependence on the covering entropy explicitly, the normalized entropy integral appearing in the local maximal inequality is bounded by

$$\int_0^{\delta_n/B_n}\sqrt{1+\min\{\log(n+1),\,d_n\log(C/\epsilon)\}+v\log(C/\epsilon)}\,d\epsilon\le C\frac{\delta_n}{B_n}\sqrt{\Lambda_n}. \tag{111}$$

To obtain (111), bound the integral once using $\log(n+1)$ and once using $d_n\log(C/\epsilon)$, then take the smaller of the two bounds. Substitution into the local maximal inequality, followed by addition of the treatment bridge bound, yields (110). If $\delta_n = 0$, the first term is zero in $L_2(P_0)$ for every y on the specified training-sample event; pointwise measurability makes it zero simultaneously on the evaluation sample almost surely. Its contribution is then omitted without evaluating a logarithm or dividing by δ_n. Because $r_{q,a,n} = o_p(1)$, $n_k \asymp n$, and (102) holds, the conditional bound tends to zero in probability on events whose probabilities tend to one. Conditional Markov's inequality, with its probability bound truncated at one before averaging over the training sample, therefore gives

$$\max_{a\in\{0,1\}}\max_{1\le k\le K}\sup_{y\in\mathcal{Y}}\left|\sqrt{n_k}(P_{n,k}-E)d_{a,y}^{(-k)}\right|=o_p(1). \tag{112}$$

The maximum over folds in (112) is valid because K is fixed. The argument conditions on one training sample at a time and does not require the overlapping training samples to be independent.

B.4 Proof of Theorem 3.1

Proof. Assumption B.2(i) gives (96) and $r_{q,a,n} = o_p(1)$, while (100) controls the outcome bridge estimation error. Lemma A.1 verifies (98) on the estimated bridge envelope events. The separate bounds above establish (112) under (102). *Step 1: Rewrite $\widehat{F}_a(y)$ fold by fold.* By (13) and (69),

$$\widehat{F}_a(y)=\sum_{k=1}^K\frac{n_k}{n}P_{n,k}\left[\widehat{h}_{a,y}^{(-k)}(W,X)+\mathbf{1}\{A=a\}\widehat{q}_a^{(-k)}(Z,X)\{\mathbf{1}\{Y\le y\}-\widehat{h}_{a,y}^{(-k)}(W,X)\}\right]. \tag{113}$$

By (70),

$$\frac{1}{n}\sum_{i=1}^n\varphi_{a,y}(O_i)=\sum_{k=1}^K\frac{n_k}{n}P_{n,k}\varphi_{a,y}(O). \tag{114}$$

Subtracting $F_a(y)$ and (114) from (113) gives

$$\widehat{F}_a(y)-F_a(y)-\frac{1}{n}\sum_{i=1}^n\varphi_{a,y}(O_i)$$
$$=\sum_{k=1}^K\frac{n_k}{n}P_{n,k}d_{a,y}^{(-k)}. \tag{115}$$

Step 2: Separate the population bias and sampling fluctuation. Add and subtract the population expectation inside each fold in (115). Then

$$\widehat{F}_a(y)-F_a(y)-\frac{1}{n}\sum_{i=1}^n\varphi_{a,y}(O_i)$$
$$=\sum_{k=1}^K\frac{n_k}{n}E\{d_{a,y}^{(-k)}(O)\}$$
$$+\sum_{k=1}^K\frac{n_k}{n}(P_{n,k}-E)d_{a,y}^{(-k)}. \tag{116}$$

Step 3: Control the population bias. Because $E\{\varphi_{a,y}(O)\} = 0$, the first sum on the right side of (116) is exactly the fold-weighted population bias in (74). Hence (79) gives

$$\max_{a\in\{0,1\}} \sup_{y\in\mathcal{Y}} \sqrt{n} \left| \sum_{k=1}^{K} \frac{n_k}{n} E\{d_{a,y}^{(-k)}(O)\} \right|$$

$$\leq \sqrt{n} \max_{a\in\{0,1\}} \left\{ \max_k \left\| \Sigma_a^\top \{\widehat{\alpha}_{a,\lambda_q}^{(-k)} - \alpha_a^{(d_n)}\} \right\|_2 \right. \tag{117}$$

$$\left. \times \max_k \sup_{y\in\mathcal{Y}} \left\| \widehat{\theta}_{a,\lambda_h}^{(-k)}(y) - \theta_a^{(d_n)}(y) \right\|_2 + b_{h,a,n} b_{q,a,n} \right\} = o_p(1) \tag{118}$$

by (96). *Step 4: Control the sampling fluctuation caused by estimated bridges.* For the second sum on the right side of (116),

$$\sqrt{n} \left| \sum_{k=1}^{K} \frac{n_k}{n} (P_{n,k} - E) d_{a,y}^{(-k)} \right| \leq \sum_{k=1}^{K} \sqrt{\frac{n_k}{n}} \left| \sqrt{n_k} (P_{n,k} - E) d_{a,y}^{(-k)} \right|. \tag{119}$$

Since K is fixed, n_k/n is bounded, and (112) holds,

$$\max_{a\in\{0,1\}} \sup_{y\in\mathcal{Y}} \sqrt{n} \left| \sum_{k=1}^{K} \frac{n_k}{n} (P_{n,k} - E) d_{a,y}^{(-k)} \right| = o_p(1). \tag{120}$$

Step 5: Combine the two remainder terms. Combining (116), (118), and (120) gives

$$\max_{a\in\{0,1\}} \sup_{y\in\mathcal{Y}} \left| \sqrt{n}\{\widehat{F}_a(y) - F_a(y)\} - \frac{1}{\sqrt{n}} \sum_{i=1}^{n} \varphi_{a,y}(O_i) \right| = o_p(1), \tag{121}$$

which proves (19).

Step 6: Establish simultaneous weak convergence. Under Assumption B.1, the class

$$\{\mathbf{1}\{Y \leq y\} : y \in \mathcal{Y}\}$$

is a bounded VC class, the class

$$\{h_{a,y}(W, X) : y \in \mathcal{Y}\}$$

has the bounded polynomial covering entropy in (95), and $q_a(Z, X)$ is fixed and bounded for each treatment arm. Products and sums of these bounded classes preserve polynomial covering entropy. The centering constants $F_a(y) \in [0, 1]$ also form a bounded polynomial-entropy class. Therefore

$$\{\varphi_{a,y}(O) : a \in \{0, 1\}, y \in \mathcal{Y}\}$$

is a bounded P_0-Donsker class. Its tight Gaussian limit is a random element of $\ell^\infty(\{0, 1\} \times \mathcal{Y})$; its paths need not be continuous in the usual distance on $\mathcal{Y}$. The empirical-process central limit theorem then gives simultaneous weak convergence of

$$\left\{ \frac{1}{\sqrt{n}} \sum_{i=1}^{n} \varphi_{a,y}(O_i) : a \in \{0, 1\}, y \in \mathcal{Y} \right\}.$$

The uniform remainder in (121) is $o_p(1)$ in $\ell^\infty(\{0, 1\} \times \mathcal{Y})$. Slutsky's theorem therefore gives simultaneous weak convergence of $\{\sqrt{n}\{\widehat{F}_a(y) - F_a(y)\} : a \in \{0, 1\}, y \in \mathcal{Y}\}$. □

C Derivation and Proof for Simultaneous Inference for $F_a(y)$ and $\Delta_Q(\tau)$

We derive the multiplier process and confidence-band inversion in Section 3.3, then prove Theorem 3.2. The estimated bridges, $\widehat{F}_a(y)$, $\varphi_{a,y}(O)$, $\widehat{\varphi}_{a,y}^{(-k)}(O_i)$, and all simultaneous confidence bands retain the definitions in the main text. The quantities n_k, $P_{n,k}$, B_n, δ_n, and Λ_n retain the definitions in Appendix B. The multipliers $\xi_1, \ldots, \xi_n$ are mutually independent, independent of the observations and the fold partition, and satisfy $P(\xi_i = 1) = P(\xi_i = -1) = 1/2$. Write P_ξ and E_ξ for probability and expectation over the multipliers conditional on all observations and the fold partition. Population expectations containing estimated bridges integrate over $O \sim P_0$ with the corresponding training sample fixed. Probabilities under P_0 for sample-dependent events include the independent fold partition.

C.1 Derivation of the Multiplier Process and Simultaneous Confidence Bands

Centering $\widehat{\varphi}_{a,y}^{(-k)}(O_i)$. Substituting (20) and then using (13) gives

$$\begin{aligned}
&\frac{1}{n}\sum_{k=1}^{K}\sum_{i\in I_k}\widehat{\varphi}_{a,y}^{(-k)}(O_i)\\
&\quad=\frac{1}{n}\sum_{k=1}^{K}\sum_{i\in I_k}\Big[\widehat{h}_{a,y}^{(-k)}(W_i,X_i)+\mathbf{1}\{A_i=a\}\widehat{q}_a^{(-k)}(Z_i,X_i)\\
&\qquad\qquad\times\{\mathbf{1}\{Y_i\le y\}-\widehat{h}_{a,y}^{(-k)}(W_i,X_i)\}\Big]-\widehat{F}_a(y)\\
&\quad=\widehat{F}_a(y)-\widehat{F}_a(y)=0.
\end{aligned}\tag{122}$$

Equation (122) establishes centering over the full sample. The average of $\widehat{\varphi}_{a,y}^{(-k)}(O_i)$ within one fold need not be zero.

Conditional mean and covariance. The multiplier distribution gives

$$E_\xi(\xi_i)=0,\qquad E_\xi(\xi_i\xi_j)=\begin{cases}1, & i=j,\\ 0, & i\neq j.\end{cases}$$

Conditionally on the observations and the fold partition, every $\widehat{\varphi}_{a,y}^{(-k)}(O_i)$ is fixed. Consequently,

$$E_\xi\{\widehat{\mathbb{G}}_a^\xi(y)\}=0.\tag{123}$$

For $a,b\in\{0,1\}$ and $y,y'\in\mathcal{Y}$, multiplying the two sums in (21) gives

$$\begin{aligned}
&E_\xi\{\widehat{\mathbb{G}}_a^\xi(y)\widehat{\mathbb{G}}_b^\xi(y')\}\\
&\quad=\frac{1}{n}\sum_{k=1}^{K}\sum_{\ell=1}^{K}\sum_{i\in I_k}\sum_{j\in I_\ell}E_\xi(\xi_i\xi_j)\widehat{\varphi}_{a,y}^{(-k)}(O_i)\widehat{\varphi}_{b,y'}^{(-\ell)}(O_j)\\
&\quad=\frac{1}{n}\sum_{k=1}^{K}\sum_{i\in I_k}\widehat{\varphi}_{a,y}^{(-k)}(O_i)\widehat{\varphi}_{b,y'}^{(-k)}(O_i).
\end{aligned}\tag{124}$$

The last equality uses the disjoint folds and the identity $E_\xi(\xi_i\xi_j)=0$ for $i\neq j$. Using the same ξ_i for all a and y retains the products of $\widehat{\varphi}_{a,y}^{(-k)}(O_i)$ and $\widehat{\varphi}_{b,y'}^{(-k)}(O_i)$ evaluated at the same observation in (124). The conditional distribution of $\widehat{\mathbb{G}}_a^\xi(y)$ need not be Gaussian at a finite sample size. The Gaussian approximation is established below as a limit.

The conditional critical value. The definition of $\widehat{c}_{1-\alpha_{\mathrm{sig}}}$ in the main text is equivalently

$$\widehat{c}_{1-\alpha_{\mathrm{sig}}}=\inf\left\{t\ge 0: P_\xi\left(\max_{a\in\{0,1\}}\sup_{y\in\mathcal{Y}}|\widehat{\mathbb{G}}_a^\xi(y)|\le t\right)\ge 1-\alpha_{\mathrm{sig}}\right\}.\tag{125}$$

The conditional critical value is nonnegative because the maximum of absolute values is nonnegative. For every a and y, $0\le F_a(y)\le 1$ implies

$$\begin{aligned}
L_a(y)\le F_a(y)\le U_a(y)\quad&\Longleftrightarrow\quad \widehat{F}_a(y)-\frac{\widehat{c}_{1-\alpha_{\mathrm{sig}}}}{\sqrt{n}}\le F_a(y)\le\widehat{F}_a(y)+\frac{\widehat{c}_{1-\alpha_{\mathrm{sig}}}}{\sqrt{n}}\\
&\Longleftrightarrow\quad |\sqrt{n}\{\widehat{F}_a(y)-F_a(y)\}|\le\widehat{c}_{1-\alpha_{\mathrm{sig}}}.
\end{aligned}\tag{126}$$

The first equivalence follows directly from the maximum with zero defining $L_a(y)$ and the minimum with one defining $U_a(y)$. Thus proving simultaneous coverage of $F_a(y)$ reduces to proving that (125) consistently estimates the corresponding Gaussian critical value.

Confidence-band inversion. On the simultaneous coverage event for $F_a(y)$, the pointwise inequalities $L_a(y)\le F_a(y)\le U_a(y)$ imply

$$\{y\in\mathcal{Y}: L_a(y)\ge\tau\}\subseteq\{y\in\mathcal{Y}: F_a(y)\ge\tau\}\subseteq\{y\in\mathcal{Y}: U_a(y)\ge\tau\}.\tag{127}$$

For nonempty crossing sets, taking infima reverses the inclusions in (127). The infimum using $U_a(y)$ therefore gives the lower bound $Q_a^L(\tau)$, while the infimum using $L_a(y)$ gives the upper bound $Q_a^U(\tau)$. Subtracting the upper bound for $Q_0(\tau)$ from the lower bound for $Q_1(\tau)$ gives $\Delta_Q^L(\tau)$. Subtracting the lower bound for $Q_0(\tau)$ from the upper bound for $Q_1(\tau)$ gives $\Delta_Q^U(\tau)$. The proof of Theorem 3.2 below verifies the coverage implication for every $\tau\in\mathcal{T}$, using the convention $\inf\varnothing=y_U$ for the infima over $\mathcal{Y}$ in (24).

C.2 Derivation of the Bridge Replacement Error

Recall the training-sample-dependent difference

$$d_{a,y}^{(-k)}(O) = \widehat{h}_{a,y}^{(-k)}(W,X) - F_a(y) + \mathbf{1}\{A=a\}\widehat{q}_a^{(-k)}(Z,X)\{\mathbf{1}\{Y\le y\} - \widehat{h}_{a,y}^{(-k)}(W,X)\} - \varphi_{a,y}(O). \tag{128}$$

The function $d_{a,y}^{(-k)}(O)$ is the difference already expanded in (104). The definition uses $F_a(y)$ rather than $\widehat{F}_a(y)$ so that $d_{a,y}^{(-k)}(O)$ depends only on the training sample when viewed as a function of O. In contrast, $\widehat{\varphi}_{a,y}^{(-k)}(O_i)$ also uses $\widehat{F}_a(y)$, which depends on the full sample. Subtracting (18) from (20) therefore gives the exact identity

$$\widehat{\varphi}_{a,y}^{(-k)}(O_i) - \varphi_{a,y}(O_i) = d_{a,y}^{(-k)}(O_i) - \{\widehat{F}_a(y) - F_a(y)\}. \tag{129}$$

Multiply (129) by $\xi_i/\sqrt{n}$ and sum over all folds:

$$\begin{aligned}\widehat{\mathbb{G}}_a^{\xi}(y) - \frac{1}{\sqrt{n}}\sum_{i=1}^{n}\xi_i\varphi_{a,y}(O_i) &= \sum_{k=1}^{K}\sqrt{\frac{n_k}{n}}\frac{1}{\sqrt{n_k}}\sum_{i\in I_k}\xi_i d_{a,y}^{(-k)}(O_i) \\ &\quad - \{\widehat{F}_a(y) - F_a(y)\}\frac{1}{\sqrt{n}}\sum_{i=1}^{n}\xi_i.\end{aligned} \tag{130}$$

The first term on the right of (130) is caused by replacing the true bridges by estimated bridges. The second term is caused by replacing $F_a(y)$ by $\widehat{F}_a(y)$ in the centering. The two terms require separate bounds.

Expanding (128) using (18) yields

$$\begin{aligned}d_{a,y}^{(-k)}(O) &= \{1 - \mathbf{1}\{A=a\}\widehat{q}_a^{(-k)}(Z,X)\}\{\widehat{h}_{a,y}^{(-k)}(W,X) - h_{a,y}(W,X)\} \\ &\quad + \mathbf{1}\{A=a\}\{\widehat{q}_a^{(-k)}(Z,X) - q_a(Z,X)\}\{\mathbf{1}\{Y\le y\} - h_{a,y}(W,X)\}.\end{aligned} \tag{131}$$

The bounds in Appendix B consequently give, on the specified training-sample events,

$$\begin{aligned}&\sup_{y\in\mathcal{Y}}|d_{a,y}^{(-k)}(O)| \le B_n + (Q_n+Q)(1+H) \quad \text{almost surely,} \\ &\sup_{y\in\mathcal{Y}}\left[E\{d_{a,y}^{(-k)}(O)^2\}\right]^{1/2} \le (1+Q_n)r_{h,a,n} + (1+H)r_{q,a,n} \le \delta_n + (1+H)r_{q,a,n}.\end{aligned} \tag{132}$$

Write $\mathcal{D}_{a,n}^{(-k)} = \{d_{a,y}^{(-k)}(O) : y\in\mathcal{Y}\}$. Combining the separate covering arguments in (108) and (109) also gives the whole-class bound

$$\begin{aligned}&\sup_{\mathsf{Q}}\log N\Big(\epsilon\{B_n + (Q_n+Q)(1+H)\}, \mathcal{D}_{a,n}^{(-k)}, L_2(\mathsf{Q})\Big) \\ &\quad\le \min\{\log(N_k+1),\ d_n\log(C/\epsilon)\} + C(v+1)\log(C/\epsilon), \qquad 0<\epsilon\le 1.\end{aligned} \tag{133}$$

The proof below uses the two terms separately, rather than applying the whole-class entropy bound to the treatment bridge estimation error.

C.3 Proof of Multiplier Control for the Bridge Replacement Error

Lemma C.1 (Multiplier control for $d_{a,y}^{(-k)}(O)$)**.** *Under the conditions of Theorem 3.1 and pointwise measurability of the function classes, for every $\epsilon>0$,*

$$P_\xi\left\{\max_{a\in\{0,1\}}\max_{1\le k\le K}\sup_{y\in\mathcal{Y}}\left|\frac{1}{\sqrt{n_k}}\sum_{i\in I_k}\xi_i d_{a,y}^{(-k)}(O_i)\right| > \epsilon\right\} \longrightarrow 0 \tag{134}$$

in P_0-probability.

Proof. Step 1: Condition on one training sample. Fix a and k and condition on the fold partition and $\{O_i : i\notin I_k\}$. Let $\mathcal{E}_{a,n}^{(-k)}$ be the training-sample event on which the estimated bridge bounds, the covering entropy bound, and the following inequality hold:

$$(1+Q_n)\sup_{y\in\mathcal{Y}}\|\widehat{h}_{a,y}^{(-k)}(W,X) - h_{a,y}(W,X)\|_W \le \delta_n.$$

Assumption B.2 implies $P_0\{(\mathcal{E}_{a,n}^{(-k)})^c\} \to 0$. Only fold-specific errors appear in $\mathcal{E}_{a,n}^{(-k)}$, so the conditioning does not fix the errors from other training samples. Conditionally on the training sample, the pairs (O_i, ξ_i) for $i \in I_k$ are iid, and the functions $d_{a,y}^{(-k)}(O)$ are fixed.

Step 2: Check the multiplier function class. Multiply each of the two terms in (131) by ξ separately. Since $|\xi| = 1$, their respective envelopes remain B_n and $(1+H)\mathbf{1}\{A = a\}|\widehat{q}_a^{(-k)}(Z, X) - q_a(Z, X)|$. Independence and $E_\xi \xi = 0$ make every multiplier function mean zero. Their L_2 norms are unchanged: the first is at most δ_n on $\mathcal{E}_{a,n}^{(-k)}$, and the second is at most $(1 + H)\|\widehat{q}_a^{(-k)}(Z, X) - q_a(Z, X)\|_{a,Z}$. For a finitely supported probability measure Q on (O, ξ), let Q_O denote its marginal measure on O. For any $d, d' \in \mathcal{D}_{a,n}^{(-k)}$,

$$\|\xi d - \xi d'\|_{L_2(\mathsf{Q})}^2 = \int \xi^2 \{d(O) - d'(O)\}^2 \, d\mathsf{Q}(O, \xi) = \|d - d'\|_{L_2(\mathsf{Q}_O)}^2.$$

The same identity holds for either term separately. Thus the first multiplier class inherits (109), and the second inherits the fixed polynomial entropy from the reweighting argument following (108). Pointwise measurability is preserved by multiplication by ξ and by squaring.

Step 3: Apply the two maximal inequalities separately. For $0 < u \leq 1$, define the deterministic entropy-integral bound

$$J_n(u) = \int_0^u \sqrt{1 + \min\{\log(n + 1),\, d_n \log(C/t)\} + v \log(C/t)}\, dt.$$

Substituting $t = us$, using $\int_0^1 \sqrt{\log(1/s)}\, ds < \infty$, and bounding the minimum by each of its two entries gives

$$J_n(u) \leq Cu \min\left\{ \sqrt{1 + \log(n + 1) + v \log(C/u)},\ \sqrt{1 + (d_n + v) \log(C/u)} \right\}.$$

For $\delta_n > 0$, it follows that $J_n(\delta_n/B_n) \leq C(\delta_n/B_n)\sqrt{\Lambda_n}$. Normalize only the first multiplier class by B_n and apply the local maximal inequality of van der Vaart and Wellner (2011, Theorem 2.1), with normalized radius δ_n/B_n. For the second multiplier class, the uniform entropy maximal inequality (van der Vaart and Wellner, 1996, Section 2.14) uses its fixed entropy integral and the $L_2(P_0)$ norm of its envelope, giving the separate bound $C\|\widehat{q}_a^{(-k)}(Z, X) - q_a(Z, X)\|_{a,Z}$. Every multiplier function has mean zero, so no additional centering term is needed. Adding the bounds gives, on $\mathcal{E}_{a,n}^{(-k)}$,

$$\begin{aligned}
&E\left[E_\xi \left\{ \sup_{y \in \mathcal{Y}} \left| \frac{1}{\sqrt{n_k}} \sum_{i \in I_k} \xi_i d_{a,y}^{(-k)}(O_i) \right| \right\} \middle| \{I_j\}_{j=1}^K, \{O_i : i \notin I_k\} \right] \\
&\leq C\|\widehat{q}_a^{(-k)}(Z, X) - q_a(Z, X)\|_{a,Z} + CB_n \left\{ J_n(\delta_n/B_n) + \frac{J_n(\delta_n/B_n)^2}{(\delta_n/B_n)^2 \sqrt{n_k}} \right\} \qquad (135) \\
&\leq C \left\{ \|\widehat{q}_a^{(-k)}(Z, X) - q_a(Z, X)\|_{a,Z} + \delta_n \sqrt{\Lambda_n} + \frac{B_n \Lambda_n}{\sqrt{n_k}} \right\}.
\end{aligned}$$

If $\delta_n = 0$, the first multiplier class vanishes simultaneously on the evaluation sample almost surely, as in Appendix B.3. Its contribution is omitted, and the last bound in (135) remains valid under the convention $\Lambda_n = 0$.

Step 4: Obtain convergence conditional on all observations. Markov's inequality, (135), and $n_k \asymp n$ imply

$$\begin{aligned}
&E\left[P_\xi \left\{ \sup_{y \in \mathcal{Y}} \left| \frac{1}{\sqrt{n_k}} \sum_{i \in I_k} \xi_i d_{a,y}^{(-k)}(O_i) \right| > \epsilon \right\} \right] \\
&\leq P_0\{(\mathcal{E}_{a,n}^{(-k)})^c\} + E\left[\min\left\{ 1, \frac{C}{\epsilon} \|\widehat{q}_a^{(-k)}(Z, X) - q_a(Z, X)\|_{a,Z} \right\} \right] \qquad (136) \\
&\quad + \frac{C}{\epsilon} \left\{ \delta_n \sqrt{\Lambda_n} + \frac{B_n \Lambda_n}{\sqrt{n}} \right\} \longrightarrow 0.
\end{aligned}$$

The bounded expectation involving the treatment bridge error tends to zero because $r_{q,a,n} = o_p(1)$; no L_1 convergence of this error is required. The rate condition in (102) makes the last term vanish. A union bound over the two treatment arms and fixed K gives convergence of the expectation of the conditional probability in (134) to zero. That conditional probability lies in $[0, 1]$. A second application of Markov's inequality shows that the conditional probability converges to zero in P_0-

probability, proving (134). The argument conditions on one training sample at a time and never assumes independence of overlapping training samples. □

C.4 Proof of Conditional Multiplier Convergence

Let $\{\mathbb{G}_a : a \in \{0,1\}\}$ denote the centered tight Gaussian limit of the leading process $\{n^{-1/2}\sum_{i=1}^n \varphi_{a,y}(O_i) : a \in \{0,1\}, y \in \mathcal{Y}\}$ established in Appendix B. The covariance is

$$E\{\mathbb{G}_a(y)\mathbb{G}_b(y')\} = E\{\varphi_{a,y}(O)\varphi_{b,y'}(O)\}, \qquad a,b \in \{0,1\}, \quad y,y' \in \mathcal{Y}. \tag{137}$$

The notation $\rightsquigarrow_\xi$ denotes conditional weak convergence given all observations and the fold partition in the bounded-Lipschitz sense, in P_0-probability. The bounded-Lipschitz sense means uniform convergence of conditional expectations over real functionals bounded in absolute value by one with Lipschitz constant at most one for the supremum norm.

Lemma C.2 (Conditional multiplier convergence)**.** *Under the conditions of Lemma C.1,*

$$\{\widehat{\mathbb{G}}_a^\xi : a \in \{0,1\}\} \rightsquigarrow_\xi \{\mathbb{G}_a : a \in \{0,1\}\} \quad \text{in } \ell^\infty(\{0,1\} \times \mathcal{Y}), \tag{138}$$

in P_0-probability.

Proof. Step 1: Apply the multiplier central limit theorem to the mean-zero functions $\varphi_{a,y}(O)$. Appendix B establishes that $\{\varphi_{a,y}(O) : a \in \{0,1\}, y \in \mathcal{Y}\}$ is a bounded, pointwise measurable P_0-Donsker class and that $E\{\varphi_{a,y}(O)\} = 0$. The multipliers have mean zero and variance one. The multiplier integrability condition holds because $\int_0^\infty \sqrt{P(|\xi_i| > t)}\, dt = 1$. The conditional multiplier central limit theorem in van der Vaart and Wellner (1996, Section 2.9) therefore gives

$$\left\{ \frac{1}{\sqrt{n}} \sum_{i=1}^n \xi_i \left[\varphi_{a,y}(O_i) - \frac{1}{n}\sum_{j=1}^n \varphi_{a,y}(O_j) \right] : a \in \{0,1\},\ y \in \mathcal{Y} \right\} \tag{139}$$
$$\rightsquigarrow_\xi \{\mathbb{G}_a(y) : a \in \{0,1\},\ y \in \mathcal{Y}\}.$$

The empirical centering in (139) changes the process by

$$-\left\{ \frac{1}{n}\sum_{j=1}^n \varphi_{a,y}(O_j) \right\} \frac{1}{\sqrt{n}} \sum_{i=1}^n \xi_i.$$

The simultaneous empirical-process limit in Appendix B gives

$$\max_{a\in\{0,1\}} \sup_{y\in\mathcal{Y}} \left| \frac{1}{n}\sum_{j=1}^n \varphi_{a,y}(O_j) \right| = O_p(n^{-1/2}).$$

Also, independence and unit variance of the multipliers give the exact identity

$$E_\xi \left\{ \left(\frac{1}{\sqrt{n}} \sum_{i=1}^n \xi_i \right)^2 \right\} = 1. \tag{140}$$

Conditional Chebyshev's inequality therefore makes the empirical-centering difference negligible uniformly over a and y. Thus (139) remains valid with $n^{-1/2}\sum_i \xi_i \varphi_{a,y}(O_i)$ in place of the empirically centered process.

Step 2: Bound the bridge replacement term. The triangle inequality gives

$$\max_{a\in\{0,1\}} \sup_{y\in\mathcal{Y}} \left| \sum_{k=1}^K \sqrt{\frac{n_k}{n}} \frac{1}{\sqrt{n_k}} \sum_{i\in I_k} \xi_i d_{a,y}^{(-k)}(O_i) \right|$$
$$\le \left(\sum_{k=1}^K \sqrt{\frac{n_k}{n}} \right) \max_{a\in\{0,1\}} \max_{1\le k\le K} \sup_{y\in\mathcal{Y}} \left| \frac{1}{\sqrt{n_k}} \sum_{i\in I_k} \xi_i d_{a,y}^{(-k)}(O_i) \right|.$$

Cauchy–Schwarz and $\sum_k n_k/n = 1$ imply $\sum_k \sqrt{n_k/n} \le \sqrt{K}$. Lemma C.1 therefore makes the bridge replacement term in (130) negligible conditionally on all observations.

Step 3: Bound the replacement of $F_a(y)$ in the centering. Theorem 3.1 and tightness of the leading process imply

$$\max_{a\in\{0,1\}} \sup_{y\in\mathcal{Y}} |\widehat{F}_a(y) - F_a(y)| = O_p(n^{-1/2}). \tag{141}$$

By (140), for every $\epsilon > 0$,

$$P_\xi\left\{\max_{a\in\{0,1\}}\sup_{y\in\mathcal{Y}}\left|\{\widehat{F}_a(y)-F_a(y)\}\frac{1}{\sqrt{n}}\sum_{i=1}^{n}\xi_i\right|>\epsilon\right\}$$
$$\le \epsilon^{-2}\left[\max_{a\in\{0,1\}}\sup_{y\in\mathcal{Y}}|\widehat{F}_a(y)-F_a(y)|\right]^2\longrightarrow 0$$

in P_0-probability. The centering error is controlled after conditioning on all observations, not on a single training sample.

Step 4: Combine the conditional limits. Using (130), Steps 2 and 3 yield

$$P_\xi\left\{\max_{a\in\{0,1\}}\sup_{y\in\mathcal{Y}}\left|\widehat{\mathbb{G}}^{\xi}_a(y)-\frac{1}{\sqrt{n}}\sum_{i=1}^{n}\xi_i\varphi_{a,y}(O_i)\right|>\epsilon\right\}\longrightarrow 0 \tag{142}$$

in P_0-probability for every $\epsilon > 0$. For a functional bounded by one with Lipschitz constant one, the difference of conditional expectations under the two processes in (142) is bounded by ϵ plus twice the conditional probability in (142). The bound is uniform over the bounded-Lipschitz functionals. Combining Step 1 with (142) and then letting $\epsilon \downarrow 0$ proves (138). □

C.5 Proof of Continuity and Separation of the Gaussian Critical Value

Define the Gaussian critical value, corresponding to (125), by

$$c_{1-\alpha_{\mathrm{sig}}}=\inf\left\{t\ge 0: P\left(\max_{a\in\{0,1\}}\sup_{y\in\mathcal{Y}}|\mathbb{G}_a(y)|\le t\right)\ge 1-\alpha_{\mathrm{sig}}\right\}. \tag{143}$$

Probabilities and expectations of $\mathbb{G}_a$ are taken under the Gaussian limit law determined by P_0.

Lemma C.3 (Continuity and separation of $c_{1-\alpha_{\mathrm{sig}}}$)**.** *Under the conditions of Lemma C.2 and (26), the distribution of* $\max_{a\in\{0,1\}}\sup_{y\in\mathcal{Y}}|\mathbb{G}_a(y)|$ *is continuous and strictly increasing on* $[0,\infty)$*. The Gaussian critical value* $c_{1-\alpha_{\mathrm{sig}}}$ *is finite and positive, and for every* $0<\epsilon<c_{1-\alpha_{\mathrm{sig}}}$*,*

$$P\left\{\max_{a\in\{0,1\}}\sup_{y\in\mathcal{Y}}|\mathbb{G}_a(y)|\le c_{1-\alpha_{\mathrm{sig}}}-\epsilon\right\}<1-\alpha_{\mathrm{sig}}$$
$$<P\left\{\max_{a\in\{0,1\}}\sup_{y\in\mathcal{Y}}|\mathbb{G}_a(y)|\le c_{1-\alpha_{\mathrm{sig}}}+\epsilon\right\}. \tag{144}$$

Proof. Step 1: Establish boundedness and nondegeneracy of the Gaussian maximum. By Assumption B.1 and $0\le F_a(y)\le 1$,

$$|\varphi_{a,y}(O)|\le(1+Q)(1+H)\quad\text{almost surely.}$$

Equation (137) and the condition (26) imply

$$0<\max_{a\in\{0,1\}}\sup_{y\in\mathcal{Y}}E\{\mathbb{G}_a(y)^2\}\le(1+Q)^2(1+H)^2<\infty.$$

The bounded polynomial covering entropy established in Appendix B has a finite entropy integral. The Gaussian process therefore has a separable tight version with finite expected absolute supremum. At least one Gaussian coordinate has positive variance, so

$$P\left\{\max_{a\in\{0,1\}}\sup_{y\in\mathcal{Y}}|\mathbb{G}_a(y)|=0\right\}=0.$$

No positive lower bound is required for the variances of all coordinates.

Step 2: Prove continuity. Apply the Gaussian supremum continuity result of Gaenssler et al. (2007) to the centered Gaussian process containing both $\mathbb{G}_a(y)$ and $-\mathbb{G}_a(y)$ for every a and y. Its supremum equals $\max_a\sup_y|\mathbb{G}_a(y)|$, is almost surely finite, and is not degenerate at zero by Step 1. Its distribution is therefore continuous. Equivalently, uniform continuity of the resulting real distribution function gives

$$\lim_{\epsilon\downarrow 0}\sup_{x\in\mathbb{R}}P\left\{\left|\max_{a\in\{0,1\}}\sup_{y\in\mathcal{Y}}|\mathbb{G}_a(y)|-x\right|\le\epsilon\right\}=0. \tag{145}$$

This qualitative continuity is sufficient for critical-value consistency.

Step 3: Prove strict increase and separation. View the separable tight version as a centered Gaussian random element in a separable closed subspace of $\ell^\infty(\{0,1\}\times\mathcal{Y})$. The support of a centered Gaussian measure is a closed linear subspace (Vakhania, 1975, Theorem 1). Step 1 ensures that the support contains a nonzero function. For any $0\le r<s<\infty$, scaling a nonzero support function

gives a support function with supremum norm $(r+s)/2$. An open ball of radius less than $(s-r)/2$ around that function has positive probability and contains only functions whose norms lie between r and s. Thus

$$P\left\{r < \max_{a\in\{0,1\}} \sup_{y\in\mathcal{Y}} |\mathbb{G}_a(y)| < s\right\} > 0.$$

Every open ball about zero also has positive probability, and a nondegenerate Gaussian coordinate exceeds every finite absolute bound with positive probability. Together with continuity and almost sure finiteness, these facts imply that the distribution function is continuous and strictly increasing on $[0,\infty)$. For $0 < \alpha_{\text{sig}} < 1$, the critical value is finite and positive and satisfies

$$P\left\{\max_{a\in\{0,1\}} \sup_{y\in\mathcal{Y}} |\mathbb{G}_a(y)| \le c_{1-\alpha_{\text{sig}}}\right\} = 1-\alpha_{\text{sig}}.$$

Strict increase proves (144). □

C.6 Proof of Theorem 3.2

Proof. Step 1: Prove consistency of $\widehat{c}_{1-\alpha_{\text{sig}}}$. The maximum of absolute values is a Lipschitz functional for the supremum norm. Lemma C.2, the conditional continuous mapping theorem, and Lemma C.3 therefore imply that, for each fixed $t \in \mathbb{R}$,

$$\begin{aligned} &P_\xi\left\{\max_{a\in\{0,1\}} \sup_{y\in\mathcal{Y}} |\widehat{\mathbb{G}}_a^\xi(y)| \le t\right\} \\ &\qquad \longrightarrow P\left\{\max_{a\in\{0,1\}} \sup_{y\in\mathcal{Y}} |\mathbb{G}_a(y)| \le t\right\} \end{aligned} \tag{146}$$

in P_0-probability. Fix $0 < \epsilon < c_{1-\alpha_{\text{sig}}}$ and apply (146) at $c_{1-\alpha_{\text{sig}}} - \epsilon$ and $c_{1-\alpha_{\text{sig}}} + \epsilon$. By (144), with probability tending to one,

$$P_\xi\left\{\max_{a\in\{0,1\}} \sup_{y\in\mathcal{Y}} |\widehat{\mathbb{G}}_a^\xi(y)| \le c_{1-\alpha_{\text{sig}}} - \epsilon\right\} < 1-\alpha_{\text{sig}},$$

$$P_\xi\left\{\max_{a\in\{0,1\}} \sup_{y\in\mathcal{Y}} |\widehat{\mathbb{G}}_a^\xi(y)| \le c_{1-\alpha_{\text{sig}}} + \epsilon\right\} > 1-\alpha_{\text{sig}}.$$

The definition (125) then places $\widehat{c}_{1-\alpha_{\text{sig}}}$ between $c_{1-\alpha_{\text{sig}}} - \epsilon$ and $c_{1-\alpha_{\text{sig}}} + \epsilon$. Hence

$$\widehat{c}_{1-\alpha_{\text{sig}}} \longrightarrow c_{1-\alpha_{\text{sig}}} \quad \text{in } P_0\text{-probability.} \tag{147}$$

No continuity of the finite-sample conditional multiplier distribution is required.

Step 2: Prove simultaneous coverage of $F_a(y)$. Theorem 3.1 and the continuous mapping theorem give

$$\max_{a\in\{0,1\}} \sup_{y\in\mathcal{Y}} |\sqrt{n}\{\widehat{F}_a(y) - F_a(y)\}| \rightsquigarrow \max_{a\in\{0,1\}} \sup_{y\in\mathcal{Y}} |\mathbb{G}_a(y)|. \tag{148}$$

For every $\epsilon > 0$, the probability of the event with the conditional critical value is bounded by

$$\begin{aligned} &P_0\left\{\max_{a\in\{0,1\}} \sup_{y\in\mathcal{Y}} |\sqrt{n}\{\widehat{F}_a(y) - F_a(y)\}| \le c_{1-\alpha_{\text{sig}}} - \epsilon\right\} - P_0\{|\widehat{c}_{1-\alpha_{\text{sig}}} - c_{1-\alpha_{\text{sig}}}| > \epsilon\} \\ &\quad \le P_0\left\{\max_{a\in\{0,1\}} \sup_{y\in\mathcal{Y}} |\sqrt{n}\{\widehat{F}_a(y) - F_a(y)\}| \le \widehat{c}_{1-\alpha_{\text{sig}}}\right\} \\ &\quad \le P_0\left\{\max_{a\in\{0,1\}} \sup_{y\in\mathcal{Y}} |\sqrt{n}\{\widehat{F}_a(y) - F_a(y)\}| \le c_{1-\alpha_{\text{sig}}} + \epsilon\right\} + P_0\{|\widehat{c}_{1-\alpha_{\text{sig}}} - c_{1-\alpha_{\text{sig}}}| > \epsilon\}. \end{aligned} \tag{149}$$

Take $n \to \infty$ using (147) and (148), then let $\epsilon \downarrow 0$. Lemma C.3 gives

$$P_0\left\{\max_{a\in\{0,1\}} \sup_{y\in\mathcal{Y}} |\sqrt{n}\{\widehat{F}_a(y) - F_a(y)\}| \le \widehat{c}_{1-\alpha_{\text{sig}}}\right\} \longrightarrow 1-\alpha_{\text{sig}}.$$

The pointwise equivalence in (126), applied to every a and y, proves (27). The argument does not require independence between $\widehat{c}_{1-\alpha_{\text{sig}}}$ and $\widehat{F}_a(y)$.

Step 3: Derive simultaneous coverage of $Q_a(\tau)$. Fix a realization for which $L_a(y) \le F_a(y) \le U_a(y)$ for every $a \in \{0,1\}$ and $y \in \mathcal{Y}$. Let $\tau \in \mathcal{T}$. The definition of $Q_a(\tau)$ and right-continuity of the marginal counterfactual distribution $F_a(y)$ give

$$F_a\{Q_a(\tau)\} \ge \tau, \qquad F_a(y) < \tau \quad \text{for every } y < Q_a(\tau).$$

The assumed interior location of $Q_a(\tau)$ ensures $Q_a(\tau) \in \mathcal{Y}$. Since $U_a\{Q_a(\tau)\} \ge F_a\{Q_a(\tau)\} \ge \tau$, the set defining $Q_a^L(\tau)$ is nonempty and contains $Q_a(\tau)$. Taking the infimum gives $Q_a^L(\tau) \le Q_a(\tau)$.

For every $y \in \mathcal{Y}$ with $y < Q_a(\tau)$, the inequality $L_a(y) \le F_a(y) < \tau$ excludes y from the set defining $Q_a^U(\tau)$. When $\{y \in \mathcal{Y} : L_a(y) \ge \tau\}$ is nonempty, the infimum defining $Q_a^U(\tau)$ is at least $Q_a(\tau)$. When $\{y \in \mathcal{Y} : L_a(y) \ge \tau\}$ is empty, the convention $\inf \varnothing = y_U$ gives $Q_a^U(\tau) = y_U > Q_a(\tau)$. Therefore,

$$Q_a^L(\tau) \le Q_a(\tau) \le Q_a^U(\tau) \quad \text{for every } a \in \{0,1\},\ \tau \in \mathcal{T} \tag{150}$$

on the simultaneous coverage event for the marginal counterfactual distribution $F_a(y)$. The argument is deterministic on the simultaneous coverage event and applies to all quantile levels without a union bound over $\mathcal{T}$. Monotonicity of the lower band $L_a(y)$ and upper band $U_a(y)$ is not required for (150). The argument is the confidence-band inversion used in Chernozhukov et al. (2020).

Step 4: Derive simultaneous coverage of $\Delta_Q(\tau)$. On the same simultaneous coverage event, (150) yields

$$\begin{aligned}\Delta_Q^L(\tau) &= Q_1^L(\tau) - Q_0^U(\tau) \le Q_1(\tau) - Q_0(\tau) = \Delta_Q(\tau),\\ \Delta_Q^U(\tau) &= Q_1^U(\tau) - Q_0^L(\tau) \ge Q_1(\tau) - Q_0(\tau) = \Delta_Q(\tau).\end{aligned}$$

Consequently,

$$\begin{aligned}&P_0\{\Delta_Q^L(\tau) \le \Delta_Q(\tau) \le \Delta_Q^U(\tau) \text{ for all } \tau \in \mathcal{T}\}\\ &\quad \ge P_0\{L_a(y) \le F_a(y) \le U_a(y) \text{ for all } a \in \{0,1\},\ y \in \mathcal{Y}\}.\end{aligned} \tag{151}$$

The right side converges to $1 - \alpha_{\text{sig}}$ by Step 2. Taking the lower limit of the left side proves (28). The inequality can be strict because simultaneous coverage of $\Delta_Q(\tau)$ can occur without simultaneous coverage of the marginal counterfactual distributions $F_0(y)$ and $F_1(y)$. No continuity of F_a on $\mathcal{Y}$, uniqueness of a solution to $F_a(y) = \tau$, counterfactual density, or derivative of the counterfactual quantile function $Q_a(\tau)$ is used in Steps 3 and 4.

Step 5: Bound the widths without density conditions. Fix any realization of the data and any number

$$\epsilon \ge \max_{a \in \{0,1\}} \sup_{y \in \mathcal{Y}} |\widehat{F}_a(y) - F_a(y)| + \frac{\widehat{c}_{1-\alpha_{\text{sig}}}}{\sqrt{n}}.$$

The definitions in (23) imply $U_a(y) \le F_a(y) + \epsilon$ and $L_a(y) \ge F_a(y) - \epsilon$. Consequently,

$$\begin{aligned}\{y \in \mathcal{Y} : U_a(y) \ge \tau\} &\subseteq \{y \in \mathcal{Y} : F_a(y) \ge \tau - \epsilon\},\\ \{y \in \mathcal{Y} : F_a(y) \ge \tau + \epsilon\} &\subseteq \{y \in \mathcal{Y} : L_a(y) \ge \tau\}.\end{aligned}$$

For all crossing sets in this step, retain the convention that the infimum of an empty set is y_U. Taking infima gives

$$Q_a^L(\tau) \ge \inf\{y \in \mathcal{Y} : F_a(y) \ge \tau - \epsilon\}, \qquad Q_a^U(\tau) \le \inf\{y \in \mathcal{Y} : F_a(y) \ge \tau + \epsilon\}.$$

These inequalities also hold for empty crossing sets, since all other infima lie in $[y_L, y_U]$. Subtracting yields the deterministic bound

$$\begin{aligned}Q_a^U(\tau) - Q_a^L(\tau) \le\ &\inf\{y \in \mathcal{Y} : F_a(y) \ge \tau + \epsilon\}\\ &- \inf\{y \in \mathcal{Y} : F_a(y) \ge \tau - \epsilon\}.\end{aligned} \tag{152}$$

Because $\Delta_Q^U(\tau) - \Delta_Q^L(\tau) = \sum_{a=0}^{1}\{Q_a^U(\tau) - Q_a^L(\tau)\}$, the same argument gives

$$\begin{aligned}\Delta_Q^U(\tau) - \Delta_Q^L(\tau) \le \sum_{a=0}^{1} \Big[&\inf\{y \in \mathcal{Y} : F_a(y) \ge \tau + \epsilon\}\\ &- \inf\{y \in \mathcal{Y} : F_a(y) \ge \tau - \epsilon\}\Big].\end{aligned} \tag{153}$$

The right-hand sides are nonnegative, so the bounds also apply to the positive parts of the displayed endpoint differences. No CDF-coverage event has been assumed in deriving (152) and (153).

On the CDF-coverage event already established in Step 2, (126) gives $\max_a \sup_y |\widehat{F}_a(y) - F_a(y)| \le \widehat{c}_{1-\alpha_{\text{sig}}}/\sqrt{n}$. Thus $\epsilon = 2\widehat{c}_{1-\alpha_{\text{sig}}}/\sqrt{n}$ is valid on that event. More generally, choosing ϵ equal to the error-plus-half-width expression above gives $\epsilon = O_p(n^{-1/2})$ by (141) and (147). When $\tau - \epsilon, \tau + \epsilon \in (0,1)$ and the corresponding true quantiles lie in the interior of $\mathcal{Y}$, the right-hand side of (152) equals $Q_a(\tau+\epsilon) - Q_a(\tau-\epsilon)$. The crossing-set version remains valid without these interior restrictions. The bounds describe how changes in the probability level affect band width, rather than imposing smoothness of the counterfactual density. No universal root-n rate for the quantile-band widths follows without additional restrictions on the counterfactual distribution. □

The coverage statements use the conditional critical value defined in (125). A finite number of multiplier draws approximates $\widehat{c}_{1-\alpha_{\text{sig}}}$ numerically. The coverage statements do not assert that a fixed number of multiplier draws has negligible simulation error.

D MORE DETAILS FOR THE NUMERICAL STUDIES

This section specifies the observed data $O = (Y, A, Z, W, X)$ and settings shared by the three simulations, followed by the settings specific to each simulation. The observed covariates X, unobserved latent variable U, treatment-inducing proxy Z, outcome-inducing proxy W, treatment indicator A, outcome Y, and potential outcome $Y(a)$ retain their definitions in the main text. We write $X = (X_1, X_2)$ to identify the two components of the observed covariates X.

Observed data $O = (Y, A, Z, W, X)$. We draw X_1 and X_2 independently from Bernoulli$(1/2)$. Conditional on X, the unobserved latent variable U is binary, with

$$P(U = 1 \mid X) = \frac{1}{1 + \exp\{-(-0.25 + 0.55X_1 - 0.35X_2)\}}. \tag{154}$$

Conditional on (U, X), the treatment-inducing proxy Z and outcome-inducing proxy W are independent binary variables, with

$$\begin{aligned} P(Z = 1 \mid U, X) &= \frac{1}{1 + \exp\{-(-0.15 + 2.2\rho(2U - 1) + 0.25X_1 - 0.15X_2)\}}, \\ P(W = 1 \mid U, X) &= \frac{1}{1 + \exp\{-(0.10 + 2.2\rho(2U - 1) - 0.20X_1 + 0.20X_2)\}}. \end{aligned} \tag{155}$$

The coefficient 2.2ρ controls the dependence of each proxy on U. The resulting $\kappa_{a,n}$ is the smallest singular value of the population matrix in (10) and depends on the chosen basis. The values of ρ are specified separately below. We then draw the treatment indicator A from the binary distribution

$$P(A = 1 \mid U, Z, W, X) = \frac{1}{1 + \exp\{-(-0.20 + 0.80(2U - 1) + 0.45Z + 0.25X_1 - 0.20X_2)\}}. \tag{156}$$

The conditional probability in (156) does not depend on W. Together with the conditional independence of Z and W, this gives $W \perp (A, Z) \mid U, X$. The conditional probabilities are strictly between zero and one, so latent positivity $0 < P(A = a \mid U, X) < 1$ holds for both treatment arms $a \in \{0, 1\}$.

Potential outcome $Y(a)$ *and outcome* Y. For each treatment arm $a \in \{0, 1\}$, the conditional mean of the potential outcome $Y(a)$ is

$$E\{Y(a) \mid U, W, X\} = 0.35a + 0.35X_1 - 0.25X_2 + 0.75U + 0.25W + 0.20aU + 0.15aX_1. \tag{157}$$

Draw from the noise distribution specified for each simulation and standardize the draw to mean zero and variance one. Generate the potential outcome by adding this noise, multiplied by $0.70 + 0.10X_1 + 0.05X_2$, to the conditional mean:

$$Y(a) = E\{Y(a) \mid U, W, X\} + (0.70 + 0.10X_1 + 0.05X_2) \times \text{standardized noise}. \tag{158}$$

The draws used for $Y(0)$ and $Y(1)$ are mutually independent and independent of (U, Z, W, A, X) before multiplication by this scale. Thus, the additive noise has conditional mean zero and conditional variance $(0.70 + 0.10X_1 + 0.05X_2)^2$ in every simulation. The observed outcome is

$$Y = AY(1) + (1 - A)Y(0). \tag{159}$$

Thus, $Y(1)$ is observed when $A = 1$, and $Y(0)$ is observed when $A = 0$. The generation above satisfies $Y(a) \perp A \mid U, X$, $Y \perp Z \mid U, A, X$, and $W \perp (A, Z) \mid U, X$. Each replication consists of n independent observations $O_1, \ldots, O_n$. The unobserved latent variable U is used to generate the data and is excluded from bridge estimation.

Prespecified basis vectors $b_W(W, X)$ *and* $b_Z(Z, X)$. All three simulations use $d_n = 8$. For the Walsh construction, the prespecified basis vector $b_W(W, X)$ contains the constant 1, the three variables $2X_1 - 1$, $2X_2 - 1$, and $2W - 1$, their three pairwise products, and their product $(2X_1 - 1)(2X_2 - 1)(2W - 1)$. The prespecified basis vector $b_Z(Z, X)$ has the same form, with Z replacing W. These eight functions span all functions on the eight possible values of (W, X) or (Z, X). Simulations 1 and 3 use this Walsh construction. Simulation 2 also considers ReLU and GB-leaf constructions, fixed before the simulation replications as described in Appendix D.2. We use the bridge estimators in (8)–(9), with the ridge parameters specified separately below. The basis vectors and ridge parameters are fixed across outcome thresholds y within each training sample. We randomly divide each replication into $K = 3$ folds of sizes differing by at most one, and use 500 replications for each setting. All training-sample moments use the denominator N_k, including those containing $\mathbf{1}\{A = a\}$.

Marginal counterfactual distribution $F_a(y)$ *and population expectations.* We calculate the marginal counterfactual distribution $F_a(y)$ from

$$F_a(y) = E\big[P\{Y(a) \le y \mid U, W, X\}\big], \tag{160}$$

using the specified noise distribution for the conditional probability and summing over the 16 possible values of (U, W, X). For the bridge estimation errors and population bias in Simulations 1 and 2, we likewise evaluate population moments by summing over the binary variables and integrating the outcome indicator using its conditional probability. The true bridges $h_{a,y}(W, X)$ and $q_a(Z, X)$ are obtained from the population moment equations in (38). These population quantities are used only to evaluate the simulation results and are excluded from bridge estimation. Since all retained basis constructions span the complete eight-dimensional spaces, $b_{h,a,n} = b_{q,a,n} = 0$. All numerical suprema over outcome thresholds are evaluated on the grid specified for each simulation.

D.1 More Details for Simulation 1

We use $n \in \{1600, 3200, 4800, 5400\}$ and $\rho \in \{0.9, 0.75, 0.6, 0.45, 0.3, 0.2\}$ in (155). The standardized noise in (158) follows $N(0, 1)$. Only ρ changes the conditional distributions of Z and W; the other conditional distributions, the conditional mean of $Y(a)$, and the noise scale remain fixed. We evaluate the bridges at 401 equally spaced outcome thresholds in $\mathcal{Y} = [-1.5, 3.5]$. We compare $\lambda_h = \lambda_q = 0$ with $\lambda_h = \lambda_q = 0.1/N_k$ on each training sample, using the same observations and folds for both choices. The unregularized estimates solve the sample moment equations in (5)–(6) whenever the sample matrix is numerically nonsingular.

The first two panels of Figure 1 report $\max_{a\in\{0,1\}} r_{h,a,n}$ and $\max_{a\in\{0,1\}} r_{q,a,n}$. The population norms are calculated as defined in Appendix A.3, and $\kappa_{a,n}$ is calculated from (11). Figure 1 holds $n = 3200$ fixed and uses $\min_{a\in\{0,1\}} \kappa_{a,n}$ as the horizontal coordinate in all four panels.

The last two panels report the $\sqrt{n}$-scaled absolute population bias in (14), maximized over treatment arms, training folds, and outcome thresholds. The Bound underlying Corollary 3.1 is calculated from the right-hand side of (79), maximized over treatment arms and multiplied by $\sqrt{n}$. The additional bound (85) is not calculated in these panels. The Cauchy–Schwarz bound is $\sqrt{n}\max_{a\in\{0,1\}} r_{h,a,n} r_{q,a,n}$. Each point represents the median for one proxy strength and ridge choice at $n = 3200$; its upper bar is the 90th percentile, conditional on successful fits. Numerically singular unregularized fits are recorded as failures and excluded from these summaries.

D.2 More Details for Simulation 2

We use $\rho \in \{0.75, 0.60, 0.45\}$, corresponding to $2.2\rho \in \{1.65, 1.32, 0.99\}$ in (155), and $n \in \{1600, 3200, 4800, 5400\}$. We compare all three basis constructions at every value of ρ. The standardized noise follows $N(0, 1)$. We set $\lambda_h = \lambda_q = 0.1/N_k$ on each training sample and use 401 equally spaced outcome thresholds in $\mathcal{Y} = [-1.5, 3.5]$. Within each replication at each value of ρ, Walsh, ReLU, and GB-leaf use the same observations and folds.

Prespecified basis vectors $b_W(W, X)$ and $b_Z(Z, X)$. Walsh uses the construction above. For ReLU and GB-leaf, we first generate one independent pilot sample of 5,400 observations from the distribution with $\rho = 0.75$. For each construction, we separately regress the observed outcome Y on (X_1, X_2, W) and on (X_1, X_2, Z), pooling the two treatment arms. The three inputs are centered and scaled using their pilot means and standard deviations. Neither the unobserved latent variable U nor the simulation replications are used to train these regressions.

ReLU uses one hidden layer with 32 ReLU units, fitted by Adam with an initial learning rate of 0.001, an L_2 coefficient regularization parameter of 0.001, batch size 200, and at most 300 iterations. Early stopping uses 10% of the pilot sample for validation, tolerance 10^{-4}, and 20 iterations without improvement. The candidate basis functions are the fitted hidden-unit activations. GB-leaf uses 36 regression trees with squared-error loss, maximum depth three, learning rate 0.05, and at least 20 observations per leaf. Its candidate basis functions are the indicators of the fitted leaves across the trees.

For each input side, we evaluate the candidate functions on all eight possible binary input values, and center and scale each function using equal weights on those values. After removing numerically constant functions, column-pivoted QR selects seven linearly independent centered and scaled functions; adding the constant 1 gives $d_n = 8$. We retain the selected functions themselves, without replacing them by the orthogonal functions from QR, and check that the resulting evaluation matrix has rank eight. The two basis vectors are then held fixed across all sample sizes, proxy strengths, replications, training folds, treatment arms, and outcome thresholds. Thus, the three constructions span the same complete function spaces, while the ridge penalty is applied to the coefficients in each

construction. The pilot sample is excluded from n, and the comparison is conditional on the fixed pilot sample. Only the bridge coefficients are estimated on each training sample. The sample matrix and ridge parameters remain fixed across y, retaining the shared-threshold structure in Lemma 3.1.

Estimated-bridge sampling fluctuation V_n and uniform expansion remainder R_n. We use $d_{a,y}^{(-k)}(O)$ as defined in (128). For each training sample, its population expectation holds the estimated bridges fixed and is evaluated using the population moments above. We subtract this expectation before combining the evaluation-fold averages:

$$V_n = \sqrt{n} \max_{a\in\{0,1\}} \sup_{y\in\mathcal{Y}} \left| \sum_{k=1}^{K} \frac{n_k}{n} \{P_{n,k} d_{a,y}^{(-k)} - E d_{a,y}^{(-k)}\} \right|.$$

This removes the population-bias contribution from V_n. The uniform expansion remainder is

$$R_n = \max_{a\in\{0,1\}} \sup_{y\in\mathcal{Y}} \left| \sqrt{n}\{\widehat{F}_a(y) - F_a(y)\} - \frac{1}{\sqrt{n}} \sum_{i=1}^{n} \varphi_{a,y}(O_i) \right|,$$

using the same observations to evaluate the estimated bridges and the true bridges. Both quantities use the counterfactual CDF estimate in (13), without clipping or monotonicity adjustment; R_n is not divided by a critical value. The signed sampling fluctuation and population bias sum to the signed remainder before taking the maximum absolute value. Figure 2 reports the median V_n and R_n for each sample size and basis construction in columns corresponding to $\rho = 0.75$, 0.60, and 0.45. The top and bottom rows show V_n and R_n, respectively, and each panel uses its own logarithmic vertical range. The annotated $\min_{a\in\{0,1\}} \kappa_{a,n}$ values refer to the Walsh basis.

D.3 More Details for Simulation 3

We fix $\rho = 0.75$ and use $n \in \{1600, 3200, 4800\}$, with $\lambda_h = \lambda_q = 0.01/\sqrt{N_k}$ on each training sample. Only the noise distribution changes across the four settings; the conditional mean in (157), the conditional distributions of (U, Z, W, A, X), and the noise scale in (158) remain fixed. The Beta settings have zero influence-function variance at thresholds outside the outcome support, which is allowed by (26). Draw from the selected noise distribution in Table 2 and standardize the draw by subtracting the listed mean and dividing by the listed standard deviation.

Table 2: Noise distributions before standardization in Simulation 3.

Noise distribution	Mean	Standard deviation
Double Weibull$(1/2)$	0	$\sqrt{24}$
Beta$(1/4, 1/4)$	$1/2$	$1/\sqrt{6}$
Beta$(1/4, 3/4)$	$1/4$	$\sqrt{3/32}$
χ_1^2	1	$\sqrt{2}$

Before standardization, Double Weibull$(1/2)$ assigns an independent, equally likely positive or negative sign to a nonnegative draw whose cumulative probability at $y \geq 0$ is $1 - \exp(-\sqrt{y})$. We calculate the counterfactual quantile function $Q_a(\tau)$ by numerical inversion of $F_a(y)$ in (160) and obtain the quantile treatment effect $\Delta_Q(\tau) = Q_1(\tau) - Q_0(\tau)$.

The counterfactual CDF estimate $\widehat{F}_a(y)$ is evaluated at 501 equally spaced outcome thresholds in $\mathcal{Y} = [-3, 5]$. We use $\alpha_{\text{sig}} = 0.05$, 999 multiplier bootstrap draws, and $\mathcal{T} = \{0.10, 0.12, \ldots, 0.90\}$. Proposed obtains $[\Delta_Q^L(\tau), \Delta_Q^U(\tau)]$ by confidence-band inversion of $[L_a(y), U_a(y)]$. Estimated-density delta uses the same $\widehat{F}_a(y)$ and additionally estimates the counterfactual density. Coverage requires the band to contain $\Delta_Q(\tau)$ for every $\tau \in \mathcal{T}$; mean band width averages the widths over $\mathcal{T}$ and successful replications. Failed fits count as noncoverage. Oracle 95% is the width of a symmetric, constant-width band: twice the 4,751st smallest maximum absolute error in estimating $\Delta_Q(\tau)$ over $\mathcal{T}$ among 5,000 replications independent of the 500 replications above.