CAUSAL INFERENCE WITHOUT TRUSTING THE PREDICTOR: TARGET-PRESERVING CALIBRATION UNDER POPULATION SHIFT

Anonymous authors
Paper under double-blind review

ABSTRACT

Under population shift, averaging machine-learned counterfactual predictions over a large target sample can become nearly deterministic while retaining causal bias. We study target-population average treatment effects when outcomes appear only in a labeled source sample and the target population contributes covariates alone. We introduce transport-calibrated causal prediction-powered inference (TC-CPI). With correct fixed propensity and transport weights, its mixing weight reallocates two unbiased estimates of the same prediction mean, one averaged in the target sample and the other transported from the source, while labeled residual correction preserves the causal estimand for every fixed predictor and weight. The exactly quadratic conditional variance yields a closed-form oracle weight; independent calibration attains asymptotic variance no larger than either endpoint. Population shift also creates two limits: learned density ratios give generic interior weights a first-order imbalance term that cross-fitting does not remove, while the efficiency bound weights source-outcome noise by r^2, so unlimited target covariates cannot compensate for poor source coverage. These limits yield conditional-shift sensitivity intervals and optimal prospective label and treatment allocation. Known design weights and bounds give finite-sample concentration intervals; with learned nuisances, cross-fitted asymptotic inference is valid under an explicit drift condition. Simulations, a 500,000-row randomized advertising study, and Project STAR test the variance geometry and transport penalty.

1 INTRODUCTION

Machine-learning predictions are often available for many more units than causal outcomes. A hospital can score every patient at admission but observe endpoints months later; a platform can predict outcomes for every interaction but randomize only a small experimental cohort; a policy trial can observe outcomes in one population while deployment occurs in another. In each setting, predictions are cheap and causal labels are not. The inferential target is not predictive accuracy, but a target-population causal effect with calibrated uncertainty.

Consider a study conducted under a source law P and deployed under a target law Q, where target outcomes are absent. A learned model returns potential-outcome predictions $f_1(x)$ and $f_0(x)$. With millions of target covariates, the plug-in average of $\tau_f = f_1 - f_0$ can become nearly deterministic. If the model misses an effect modifier whose prevalence shifts, its causal bias remains: the target sample makes the wrong answer look certain. Transported AIPW remains valid through labeled residual correction, but does not directly use the large target sample to estimate the prediction mean. The central question is therefore:

How much predictor trust can be learned without changing the causal target?

Prediction-powered inference (PPI) corrects arbitrary predictions with gold-standard residuals (Angelopoulos et al., 2023a), and PPI++ calibrates prediction power for ordinary statistical estimands (Angelopoulos et al., 2023b). Their canonical formulations operate under a common population. Under covariate shift, the same prediction mean is observed through two population laws and the density ratio becomes an additional nuisance. Related work addresses semi-supervised treatment effects (Chakrabortty & Dai, 2022), trial generalization (Demirel et al., 2024), calibrated transfer (Cadei

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.

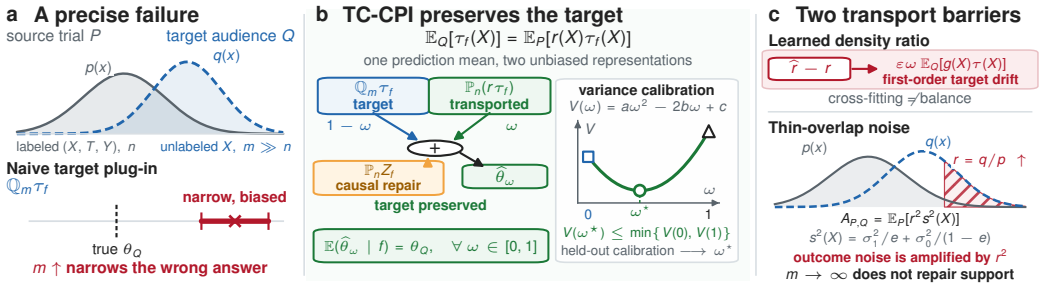


Figure 1: TC-CPI interpolates between target-sample prediction averaging and transported labeled estimation without changing the causal target. The target plug-in retains causal bias (a); the quadratic variance determines $\omega^\star$ (b); density-ratio error and r^2-weighted source noise limit the attainable gain (c).

et al., 2025), automated phenotypes under shift (Ying et al., 2025), and semi-supervised Riesz/DML estimation (Kato, 2026). TC-CPI studies how to combine the two representations, how an estimated transport ratio affects first-order validity, and how source support limits efficiency.

Transport-calibrated causal PPI (TC-CPI) uses this two-sample structure directly. Let $r = dQ_X/dP_X$. The identity $\mathbb{E}_Q\{\tau_f(X)\} = \mathbb{E}_P\{r(X)\tau_f(X)\}$ exposes two unbiased representations of one prediction mean. TC-CPI reallocates that mean between the target sample and transported source sample with $\omega \in [0,1]$, while retaining the full causal residual correction. The endpoints are external PPI ($\omega = 0$) and transported labeled AIPW ($\omega = 1$). With correct fixed design and transport weights, ω changes variance, not identification. The variance is exactly quadratic, giving a closed-form oracle and an independent held-out calibration rule.

Population shift creates two barriers absent from the common-population reduction. First, when r is learned, a generic interior weight has a first-order transport-imbalance term; cross-fitting controls sample reuse but does not enforce balance. The external endpoint is the ratio-orthogonal boundary under correct outcome regression. Second, even an oracle predictor faces labeled noise weighted by r^2, not r, in the efficiency bound. Thus predictive accuracy alone does not order causal precision. The covariance terms determine the variance-minimizing weight; transport balance and source coverage determine whether it can be estimated with valid uncertainty.

Contributions.

1. **Target-preserving calibration geometry.** TC-CPI connects external PPI and transported AIPW, derives the exact oracle allocation, and learns it independently without moving the estimand (Sections 3–4).
2. **Learned-ratio drift and support limits.** We characterize learned-weight drift and the unavoidable r^2 efficiency penalty, then derive conditional-shift sensitivity intervals and optimal prospective label allocation (Section 4).
3. **Empirical tests tied to the theory.** Simulations, randomized advertising data, and Project STAR test the variance geometry and the roles of gain, support, and balance (Section 5).

Section 2 defines the target and assumptions; Section 3 constructs TC-CPI; Section 4 establishes its gains and limits; Section 5 tests the theory; and Section 6 summarizes the scope and design implications.

2 Setup, Assumptions, and Prior Work

Two samples and one target. Let X be pretreatment covariates, $T \in \{0,1\}$ treatment, $Y(t)$ potential outcomes, and $Y = TY(1) + (1-T)Y(0)$. We observe independent samples

$$\mathcal{L}_n = \{W_i = (X_i, T_i, Y_i)\}_{i=1}^n \overset{\text{iid}}{\sim} P, \qquad \mathcal{U}_m = \{\widetilde{X}_j\}_{j=1}^m \overset{\text{iid}}{\sim} Q_X.$$

Write $e(x) = P(T = 1 \mid X = x)$, $r(x) = dQ_X/dP_X(x)$, $\mu_t(x) = \mathbb{E}\{Y(t) \mid X = x\}$, and $\tau(x) = \mu_1(x) - \mu_0(x)$; the target is $\theta_Q = \mathbb{E}_Q\{\tau(X)\}$. The predictor $f = (f_0, f_1)$ may be a causal forest, representation learner, foundation-model score, simulator, or simple regression; $\tau_f = f_1 - f_0$. It need not equal μ. For measurable $g(W)$ and $h(X)$, write $\mathbb{P}_n g = n^{-1}\sum_{i=1}^n g(W_i)$ and $\mathbb{Q}_m h = m^{-1}\sum_{j=1}^m h(\widetilde{X}_j)$, and let $\Pi_{[0,1]}(z) = \min\{1, \max(0, z)\}$.

Assumption 2.1 (Causal identification)**.** *Consistency holds and* $(Y(0), Y(1)) \perp\!\!\!\perp T \mid X$ *under* P.

Assumption 2.2 (Treatment overlap)**.** *For some* $\epsilon > 0$, $\epsilon \le e(X) \le 1 - \epsilon$ *almost surely under* P.

Assumption 2.3 (Transportability and support)**.** $Q_X \ll P_X$*, and the conditional potential-outcome means in the target equal* $\mu_t(x)$ *on the support of* Q_X*. The density ratio* r *is finite* P*-almost surely.*

Each assumption has an operational role. Exchangeability licenses causal residuals; randomized assignment provides it by design, while observational analyses require a substantive adjustment set and sensitivity analysis. Treatment overlap bounds inverse assignment weights; propensity histograms and effective sample sizes diagnose it, and an overlap-population target is the fallback. Transportability requires substantive justification. Density-ratio tails, effective sample sizes, and source–target classifiers can detect severe support violations; such violations require a restricted target, partial identification, or labels from unsupported strata. For exact finite-sample results below, e, r, and f are known or fixed independently of evaluation data. Section 4.3 states what survives when weights are learned.

Relation to prior work. PPI++ tunes a scalar path between labeled-only and prediction-powered mean estimation (Angelopoulos et al., 2023b); Calibrated PPI interprets this operation as empirical efficiency maximization and proves first-order equivalence between linear score calibration and PPI++ in that setting (van der Laan & van der Laan, 2026). Proposition 3.1 gives our exact common-population reduction. Under shift, TC-CPI instead mixes prediction means sampled under two laws without recalibrating the predictor; estimating their density ratio creates a new first-order drift. In a single randomized population, AI predictions may also enter interacted regression adjustment without worsening asymptotic efficiency and may improve it, although ordinary OLS retains an $O(n^{-1})$ finite-sample bias (Arbour et al., 2026). Our target is an external population, so the transport ratio, its imbalance, and the r^2 information term enter the problem.

Semi-supervised causal estimators already exploit unlabeled observations for ATEs and QTEs (Chakrabortty & Dai, 2022); related approaches study trial generalization (Demirel et al., 2024), conditional calibration across experiments (Cadei et al., 2025), automated phenotypes under shift (Ying et al., 2025), and semi-supervised Riesz/DML estimation (Kato, 2026). Recent predictive-shift analysis asks when observable covariate shift can quantify unobserved conditional shift (Jin et al., 2025); our extension converts a supplied conditional-shift radius into an estimand-specific sensitivity interval and a robust calibration rule. TC-CPI isolates the exact continuum between an external target average and a transported labeled average, retains arbitrary fixed predictors, and separates known-weight finite-sample validity from learned-weight asymptotics. Its drift connects to doubly robust missing-data scores (Robins et al., 1994; Bang & Robins, 2005): cross-fitting controls sample reuse (Chernozhukov et al., 2018), but not population transport imbalance. Causal representation learning (Johansson et al., 2016; Shalit et al., 2017; Johansson et al., 2022) supplies useful f's; here their quality affects precision rather than identification.

3 Transport-calibrated causal PPI

Define the transported residual

$$Z_f(W) = r(X)\left[\frac{T}{e(X)}\{Y - f_1(X)\} - \frac{1-T}{1-e(X)}\{Y - f_0(X)\}\right]. \tag{1}$$

For $\omega \in [0, 1]$, our estimator is

$$\widehat{\theta}_\omega = \omega\,\frac{1}{n}\sum_{i=1}^n r(X_i)\tau_f(X_i) + (1-\omega)\,\frac{1}{m}\sum_{j=1}^m \tau_f(\widetilde{X}_j) + \frac{1}{n}\sum_{i=1}^n Z_f(W_i). \tag{2}$$

At $\omega = 0$, predictions are averaged in the external target sample and labels only correct their bias. At $\omega = 1$, the target prediction average is replaced by its transported labeled analogue, giving transported

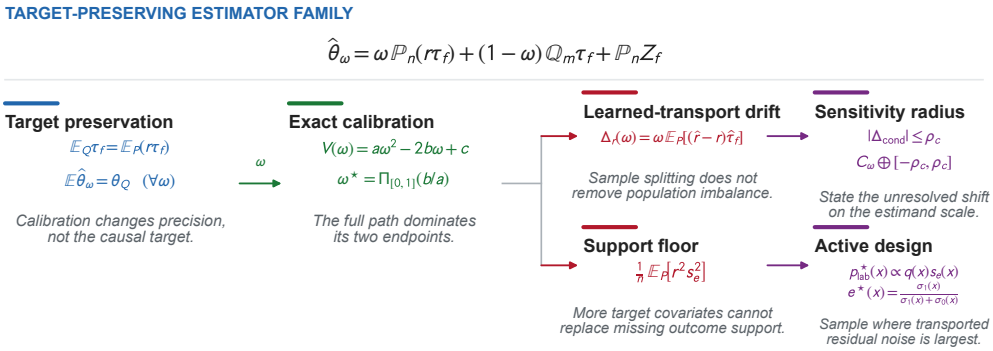


Figure 2: The main results form one mathematical chain. Target preservation separates identification from calibration; the quadratic variance selects $\omega^\star$. Learned transport adds a first-order drift, whereas limited source support leaves an r^2-weighted variance floor. These two limits lead, respectively, to sensitivity expansion and prospective label and treatment design.

AIPW. Interior values pool two unbiased estimates of the same prediction mean while retaining the full residual correction.

Running example. Let $X = 0$ denote light users and $X = 1$ heavy users, with $P_X = (0.8, 0.2)$, $Q_X = (0.2, 0.8)$, and $e = 1/2$. Suppose $\tau = (1, 3)$ but $\tau_f = (0.5, 3.5)$. Then $\theta_Q = 2.6$, whereas the increasingly precise target plug-in converges to $\mathbb{E}_Q\tau_f = 2.9$; the shifted heavy-user share exposes predictor bias rather than averaging it away.

Theorem 3.1 (Target preservation)**.** *Under Assumptions 2.1–2.3, suppose e, r, f are fixed independently of the evaluation samples and the terms in equation 2 are integrable. Then, for every $\omega \in \mathbb{R}$,*

$$\mathbb{E}(\widehat{\theta}_\omega \mid f) = \theta_Q.$$

No correctness condition on f is required.

The identity follows from $\mathbb{E}_P Z_f = \mathbb{E}_Q(\tau - \tau_f)$ and $\mathbb{E}_P(r\tau_f) = \mathbb{E}_Q\tau_f$. It is the reason calibration can optimize precision without changing the target. In the running example $r = (0.25, 4)$, so both prediction means equal 2.9 and the residual correction has mean -0.3: every ω is centered at $2.9 - 0.3 = 2.6$.

Proposition 3.1 (Common-population reduction)**.** *If $P_X = Q_X =: M$ and $r = 1$, define $H_f = \tau_f + Z_f$ and $\gamma = 1 - \omega$. Then*

$$\widehat{\theta}_\omega = \mathbb{P}_n H_f + \gamma\{\mathbb{Q}_m\tau_f - \mathbb{P}_n\tau_f\}. \tag{3}$$

This is exactly the PPI++ mean-estimation path applied to labeled score H_f and prediction τ_f. If $B = \mathrm{Var}_M(\tau_f) > 0$, its oracle power is $\gamma^\star = \Pi_{[0,1]}[\mathrm{Cov}_P\{H_f, \tau_f\}/\{B(1 + n/m)\}]$. The transport-specific phenomena begin when $P_X \neq Q_X$: estimation of r creates Proposition 4.1, while labeled outcome noise is multiplied by r^2 in Theorem 4.6.

4 Theory: Gain, Validity, and the Transport Barrier

Figure 2 organizes the argument in five steps. We first derive the exact variance, its oracle minimizer, and guarantees for learning that minimizer (Theorems 4.1–4.3). We then obtain finite-sample coverage under bounded summands and correct fixed propensity and transport weights (Theorem 4.4). Next, Proposition 4.1 isolates the first-order drift created by an estimated transport ratio, and Theorem 4.5 gives sufficient cross-fitted conditions for controlling it below the $n^{-1/2}$ scale. After relaxing effect transportability, Proposition 4.2 and Corollary 4.1 convert target displacement and ratio error into sensitivity intervals and robust calibration. Finally, Theorem 4.6 identifies the r^2-weighted efficiency floor, and Corollary 4.2 minimizes its source component through prospective label and treatment design.

4.1 EXACT VARIANCE AND OPTIMAL CALIBRATION

Theorem 3.1 identifies the same target for every fixed ω, so calibration can focus entirely on precision. The results proceed in three stages: Theorem 4.1 derives the exact quadratic variance and its optimal weight within the path; Theorem 4.2 shows that an independent split learns this weight without first-order loss; and Theorem 4.3 gives a high-probability finite-calibration bound on the resulting excess variance.

Condition on any external training sigma-field $\mathcal{T}$ that determines f and is independent of the evaluation samples. Let $U = Z_f(W)$, $R = r(X)\tau_f(X)$, and $S = \tau_f(\widetilde{X})$; set $A = \mathrm{Var}_P(U)$, $C_{UR} = \mathrm{Cov}_P(U, R)$, $B_P = \mathrm{Var}_P(R)$, and $B_{Q,f} = \mathrm{Var}_Q(S)$. These quantities are conditional on $\mathcal{T}$; we suppress that conditioning to keep the notation readable.

Theorem 4.1 (Oracle calibration and endpoint dominance)**.** *If the second moments are finite and the samples are independent, then*

$$V(\omega) := \mathrm{Var}(\widehat{\theta}_\omega \mid \mathcal{T}) = \frac{A + 2\omega C_{UR} + \omega^2 B_P}{n} + \frac{(1-\omega)^2 B_{Q,f}}{m}. \tag{4}$$

If $B_P/n + B_{Q,f}/m > 0$, the variance-minimizing member of $[0,1]$ is

$$\omega^\star = \Pi_{[0,1]}\left(\frac{B_{Q,f}/m - C_{UR}/n}{B_P/n + B_{Q,f}/m}\right), \tag{5}$$

and $V(\omega^\star) \le \min\{V(0), V(1)\}$. The two endpoints are external PPI and transported labeled AIPW, respectively.

The covariance C_{UR} determines whether transported predictions amplify or cancel residual noise, so neither endpoint uniformly minimizes V across data-generating laws.

For the running example, complete the variance model with $\mu_0 = f_0 = 0$, $\mu_1 = \tau$, $f_1 = \tau_f$, and unit-variance conditional Rademacher outcome noise. Then $(A, C_{UR}, B_P, B_{Q,f}) = (14.535, -4.7175, 30.8025, 1.44)$. At $(n, m) = (100, 1000)$, $V(0) = 0.1468$, $V(1) = 0.3590$, but $\omega^\star = 0.157$ gives $V(\omega^\star) = 0.1392$: the variance-minimizing member lies strictly inside the path, at neither endpoint.

Along a sequence with $m/n \to \lambda \in (0, \infty)$, write $D_\lambda = B_P + \lambda^{-1} B_{Q,f}$, $\omega^\star_\lambda = \Pi_{[0,1]}\{(\lambda^{-1} B_{Q,f} - C_{UR})/D_\lambda\}$, V_n for V at the current (n, m), and $\omega^\star_n \in \arg\min_{w \in [0,1]} V_n(w)$. Also let $v_\lambda(\omega) = A + 2\omega C_{UR} + \omega^2 B_P + \lambda^{-1}(1-\omega)^2 B_{Q,f}$.

Theorem 4.2 (Independent adaptive calibration)**.** *Let $\mathcal{C} = (\mathcal{C}_P, \mathcal{C}_Q)$ denote independent source and target calibration samples, independent of $\mathcal{T}$ and the inference samples. Replace $(C_{UR}, B_P, B_{Q,f})$ in equation 5 by consistent analogues from $\mathcal{C}$. Under uniformly bounded $2+\eta$ moments for some $\eta > 0$, $D_\lambda > 0$, and $v_\lambda(\omega^\star_\lambda) > 0$,*

$$\mathbb{E}(\widehat{\theta}_{\widehat{\omega}} \mid \mathcal{T}, \mathcal{C}) = \theta_Q, \quad \widehat{\omega} \to_p \omega^\star_\lambda, \quad n\{V_n(\widehat{\omega}) - V_n(\omega^\star_n)\} \to_p 0,$$

and the consistently studentized estimator converges to $N(0,1)$.

The preceding limit has a finite-calibration counterpart. Let $n_c, m_c \ge 2$ be the calibration sizes and let H_U, H_R, H_S denote almost-sure range lengths of U, R, S. Form $\widehat{\omega}$ from the unbiased calibration covariance and variances.

Theorem 4.3 (High-probability calibration regret)**.** *Let $N_c = \min\{\lfloor n_c/2 \rfloor, \lfloor m_c/2 \rfloor\}$ and $K_{\mathrm{cal}} = H_U H_R + H_R^2/2 + H_S^2/\underline{\lambda}$. If $m/n \ge \underline{\lambda} > 0$ and $D_{n,m} := B_P/n + B_{Q,f}/m$ satisfies $nD_{n,m} \ge \underline{d} > 0$, then, for every $0 < \delta < 1$, conditional on training, with probability at least $1-\delta$ over calibration,*

$$V(\widehat{\omega}) \le \min\{V(0), V(1)\} + \frac{2K_{\mathrm{cal}}^2 \log(6/\delta)}{\underline{d}\, n N_c}. \tag{6}$$

Appendix A.5 gives the flat-curve alternative; under curvature, calibration costs only $O\{\log(1/\delta)/(nN_c)\}$ variance.

4.2 FINITE-SAMPLE INFERENCE WITH KNOWN DESIGN WEIGHTS

The fixed-weight regime also admits exact finite-sample inference. Theorems 4.2–4.3 justify learning the variance-minimizing weight; Theorem 4.4 instead uses known summand ranges to replace asymptotic approximation by nonasymptotic coverage, without changing the target-preserving center in Theorem 3.1.

Theorem 4.4 (Finite-sample concentration intervals)**.** *Under the conditions of Theorem 3.1, suppose ω is fixed or independently calibrated. Let $S_{P,\omega}(W) = Z_f(W) + \omega r(X)\tau_f(X)$ and $S_{Q,\omega}(\widetilde{X}) = (1-\omega)\tau_f(\widetilde{X})$ lie in known intervals of range lengths $H_{P,\omega}$ and $H_{Q,\omega}$. For every $0 < \alpha < 1$, the Hoeffding and empirical-Bernstein intervals in Appendix A.6 have conditional and unconditional coverage at least $1-\alpha$, without predictor correctness or an asymptotic approximation.*

This guarantee requires the correct fixed propensity and density-ratio weights, plus known summand bounds. The running example has bounded covariates and Rademacher outcome noise, so its summand ranges are available without estimating them from the analysis sample.

4.3 THE TRANSPORT-ORTHOGONALITY BARRIER

The preceding results condition on correct nuisance functions. Proposition 4.1 first shows that estimating the transport ratio creates a first-order drift for every interior weight, even under correct outcome regression; cross-fitting alone does not cancel it. Theorem 4.5 then gives sufficient conditions for valid cross-fitted inference: foldwise score convergence in L_2 and an explicit $o_p(n^{-1/2})$ drift remainder.

For arbitrary candidate nuisances $(\widetilde{f}, \widetilde{e}, \widetilde{r})$ used in equation 2, let $e_1 = e$, $e_0 = 1-e$, and define $\widetilde{e}_1 = \widetilde{e}$, $\widetilde{e}_0 = 1-\widetilde{e}$. With $s_1 = 1, s_0 = -1$, the population drift is exactly

$$\sum_{t=0}^{1} s_t \mathbb{E}_P\left[\left\{\frac{\widetilde{r}e_t}{\widetilde{e}_t} - r\right\}(\mu_t - \widetilde{f}_t) + \omega(\widetilde{r}-r)\widetilde{f}_t\right] =: \mathfrak{D}(\omega; \widetilde{f}, \widetilde{e}, \widetilde{r}). \tag{7}$$

Proposition 4.1 (Transport-orthogonality barrier)**.** *Suppose $\widetilde{f}_t = \mu_t$ and let $\widetilde{r}_\varepsilon = r(1+\varepsilon g)$, where g is bounded, $\mathbb{E}_Q g = 0$, and $|\varepsilon|\,\|g\|_\infty < 1$. For any candidate propensity bounded away from zero and one,*

$$\mathfrak{D}(\omega; \mu, \widetilde{e}, \widetilde{r}_\varepsilon) = \varepsilon\omega\mathbb{E}_Q\{g(X)\tau(X)\}. \tag{8}$$

Thus $\omega = 0$ is ratio-orthogonal under correct outcome regression; if $\mathrm{Var}_Q(\tau) > 0$, every $\omega > 0$ has a bounded ratio direction with nonzero first derivative. Cross-fitting alone cannot remove it.

In the running example, grant oracle outcome regression and perturb r by $r(1+0.1g)$ with $g = (4, -1)$ and $\mathbb{E}_Q g = 0$. Then $\mathfrak{D}(\omega) = -0.16\omega$: the external endpoint cancels the error, while every interior weight inherits it. Thus cross-fitting controls sample reuse but does not enforce transport balance.

When r is known, the remaining drift is a product of propensity and outcome-regression errors. When r is learned and $\omega > 0$, root-n inference additionally requires $\mathbb{E}_P\{(\widehat{r}-r)\widehat{\tau}_f\} = o_p(n^{-1/2})$, or a balancing step that enforces this scalar condition.

Theorem 4.5 (Cross-fitted learned-weight inference)**.** *For a fixed K, partition both samples into deterministic balanced folds and let $(\widehat{f}_k, \widehat{e}_k, \widehat{r}_k)$, $1 \le k \le K$, denote nuisance fits trained off fold k. Use a common $\widehat{\omega}$ estimated on an independent calibration sample. Suppose e and all $\widehat{e}_k$ lie in $[\epsilon, 1-\epsilon]$ almost surely for some $\epsilon > 0$, $m/n \to \lambda \in (0, \infty)$, $\widehat{\omega} \to_p \omega^\star_\lambda$, and foldwise L_2 convergence to the fixed source and target limit scores generated by some $f^\star$, as made explicit in equation 30. Here $\omega^\star_\lambda$ is the oracle weight for those limit scores. Suppose further that the fitted and limit scores have uniformly bounded $2+\eta$ moments for some $\eta > 0$, the limiting variance is positive, and*

$$\max_{1 \le k \le K} \left|\mathfrak{D}(\widehat{\omega}; \widehat{f}_k, \widehat{e}_k, \widehat{r}_k)\right| = o_p(n^{-1/2}).$$

Then the estimator is asymptotically linear, its plug-in variance is consistent, and its Wald interval is valid.

Exact finite-sample validity remains tied to the correct fixed propensity and density-ratio weights; ω may still be independently calibrated. Learned nuisance weights use Theorem 4.5's explicit drift condition.

4.4 BEYOND EXACT TRANSPORTABILITY: PREDICTIVE-SHIFT SENSITIVITY

Proposition 4.1 concerns estimating the transport mechanism when conditional effects are shared. We now relax that shared-effect assumption itself: Proposition 4.2 isolates the resulting estimand displacement, and Corollary 4.1 turns it into a uniform interval expansion and a closed-form robust weight. Appendix A.8 establishes sharpness for the stated product-rectangular, common-noise model. The resulting interval adds the prescribed sensitivity radius to the transported estimand, while ratio uncertainty modifies the calibration criterion.

Let $\mu_t^P(x) = \mathbb{E}_P\{Y(t) \mid X = x\}$ and $\mu_t^Q(x) = \mathbb{E}_Q\{Y(t) \mid X = x\}$, and set $\tau_P = \mu_1^P - \mu_0^P$ and $\tau_Q = \mu_1^Q - \mu_0^Q$. Write $\theta_Q = \mathbb{E}_Q\tau_Q(X)$, $\theta_Q^{\mathrm{tr}} = \mathbb{E}_Q\tau_P(X)$, and $\Delta_{\mathrm{cond}} = \mathbb{E}_Q\{\tau_Q(X) - \tau_P(X)\}$.

Proposition 4.2 (Shift decomposition). *Under source identification and $Q_X \ll P_X$, allow $\mu_t^Q \neq \mu_t^P$. If e, r are correct and fixed, then for every fixed f, ω for which the displayed expectations are finite,*

$$\mathbb{E}(\widehat{\theta}_\omega \mid f) = \theta_Q^{\mathrm{tr}}, \qquad \mathbb{E}(\widehat{\theta}_\omega \mid f) - \theta_Q = -\Delta_{\mathrm{cond}}.$$

Conditional shift is therefore a common estimand displacement along the path.

Corollary 4.1 (Predictive-shift sensitivity and robust calibration). *Fix $0 < \alpha < 1$ and nonnegative sensitivity radii ρ_c, ρ_r. If C_ω is a $(1-\alpha)$ confidence interval for θ_Q^{tr}, then its Minkowski expansion $C_\omega \oplus [-\rho_c, \rho_c]$ covers θ_Q uniformly over $|\Delta_{\mathrm{cond}}| \le \rho_c$. For ratio errors, let $b_r(\omega)$ denote the ratio-induced centering error and suppose instead that C_ω covers $\theta_Q^{\mathrm{tr}} + b_r(\omega)$ uniformly with probability $1-\alpha$ and $|b_r(\omega)| \le \omega\rho_r$. Appendix A.8 gives the exact affine form $b_r(\omega) = b_0 + \omega b_1$; this corollary sets $b_0 = 0$. Then radius $\rho_c + \omega\rho_r$ gives uniform coverage. The bias–variance criterion $\mathcal{R}_{\mathrm{rob}}(\omega) = V(\omega) + (\rho_c + \omega\rho_r)^2$ has, when $D_{\mathrm{rob}} := B_P/n + B_{Q,f}/m + \rho_r^2 > 0$, the minimizer*

$$\omega_{\mathrm{rob}}^\star = \Pi_{[0,1]}\left[\frac{B_{Q,f}/m - C_{UR}/n - \rho_c\rho_r}{B_P/n + B_{Q,f}/m + \rho_r^2}\right]. \tag{9}$$

If $D_{\mathrm{rob}} = 0$, the criterion is constant and every $\omega \in [0,1]$ is optimal.

Under a nonadversarial random-shift model, a standardized observable covariate-shift measure can predict an estimand-specific standardized conditional shift (Jin et al., 2025). Appendix A.8 formalizes the model-based calibration $\rho_c = \Gamma d_X s_{\mathrm{cond}}$, propagates uncertain radii, and gives the exact minimax rule for a stated product scalar-bias/common-noise model. We use the relation to calibrate sensitivity; covariates alone do not identify conditional shift.

In the running example, a target-specific effect displacement $\delta_\tau = (d_0, d_1)$ moves the target by $\Delta_{\mathrm{cond}} = 0.2d_0 + 0.8d_1$ while leaving the entire calibrated path centered at $\theta_Q^{\mathrm{tr}} = 2.6$. A bound on this weighted contrast determines the interval expansion, while ratio uncertainty adds an ω-dependent penalty.

4.5 WHAT UNLABELED DATA CANNOT REMOVE

We next hold exact transportability fixed and study the information lost when target outcomes are absent. Theorem 4.6 locates the irreducible term through the r^2-weighted efficiency bound and proves attainability; Corollary 4.2 converts that bound into optimal label recruitment and treatment allocation.

Return to exact conditional-mean transportability in Assumption 2.3, so $\tau_P = \tau_Q = \tau$. Let $\sigma_t^2(x) = \mathrm{Var}_P\{Y(t) \mid X = x\}$, with $\sigma_t(x) = \sqrt{\sigma_t^2(x)}$, and assume $m/n \to \lambda \in (0, \infty)$. Consider the regular nonparametric two-sample model with known treatment propensity.

Theorem 4.6 (External-target efficiency bound). *Suppose the two terms below are finite. Every regular estimator of θ_Q has asymptotic variance under $\sqrt{n}$ scaling at least*

$$\underbrace{\mathbb{E}_P\left[r(X)^2\left\{\frac{\sigma_1^2(X)}{e(X)} + \frac{\sigma_0^2(X)}{1-e(X)}\right\}\right]}_{A_{P,Q}} + \lambda^{-1}\underbrace{\mathrm{Var}_Q\{\tau(X)\}}_{B_{Q,\tau}}. \tag{10}$$

The oracle score $f_t = \mu_t$, $\omega = 0$ attains the bound; a feasible cross-fitted one-step estimator does so under L_2 convergence to this efficient oracle score and $o_p(n^{-1/2})$ drift. Proposition A.3 gives the restricted common-marginal special case.

In the running example the bound is $13 + 0.64/\lambda$ (13.064 at $\lambda = 10$); its formal large-λ limit exposes the source floor 13.

Corollary 4.2 (Oracle label and treatment design)**.** *Let ν be a common σ-finite dominating measure, and let $q = dQ_X/d\nu$ and $p_{\mathrm{lab}} = dP_{\mathrm{lab}}/d\nu$ be the target and prospective label densities. Suppose recruitment changes only the marginal law of X and preserves the conditional potential-outcome laws. For $s_e = \{\sigma_1^2/e + \sigma_0^2/(1-e)\}^{1/2}$, assume $0 < \int q(x)s_e(x)\,d\nu(x) < \infty$ and $s_e > 0$ Q_X-almost surely. Among label designs with $p_{\mathrm{lab}} > 0$ Q_X-almost surely and finite bound, the optimum is*

$$p^\star_{\mathrm{lab}}(x) = \frac{q(x)s_e(x)}{\int q(u)s_e(u)\,d\nu(u)}. \tag{11}$$

If treatment allocation is also designable and $\sigma_0, \sigma_1 > 0$ Q_X-almost surely, then

$$e^\star(x) = \frac{\sigma_1(x)}{\sigma_1(x) + \sigma_0(x)}, \qquad p^\star_{\mathrm{lab}}(x) \propto q(x)\{\sigma_1(x) + \sigma_0(x)\}.$$

Given a target recruitment frame, an independent pilot makes the oracle rule implementable: Appendix A.10 proves a positivity-safe two-stage design, its regret and consistency, and constrained and cost-weighted variants. Minimizing equation 10 therefore yields the label-allocation rule. Here $s_e \equiv 2$, so $p^\star_{\mathrm{lab}} = Q$ lowers the source-noise term from 13 to 4 at fixed n: target covariates shrink $\lambda^{-1}B_{Q,\tau}$, but only better-placed labels reduce r^2-amplified noise. Appendix B extends the construction beyond means.

5 Experiments

Questions and protocol. We test four claims in order: the exact variance and oracle calibration; fixed-weight finite-sample validity; learned-nuisance drift; and the r^2 support, sensitivity, and design limits. The randomized-outcome studies test whether the same distinctions survive learned prediction and real outcomes. Main synthetic cells use 1,000 repetitions with disjoint calibration and inference blocks. Baselines are external PPI ($\omega = 0$), transported AIPW ($\omega = 1$), and an infeasible oracle; metrics are bias, variance ratio, 95% coverage, interval length, and $\mathbb{E}\widehat{\omega}$. The appendix records all settings; bundled scripts and metadata record fixed seeds, aggregate outputs, and selected replicate-level records.

Variance geometry and the calibration phase transition. We cross five ratios $m/n \in \{1, 2, 5, 10, 20\}$ with seven fixed predictor alignments $\kappa \in [-1, 1]$, from oracle ($\kappa = 1$) to reversed effect heterogeneity ($\kappa = -1$). This grid varies target-sample leverage and prediction quality while leaving the causal estimand fixed.

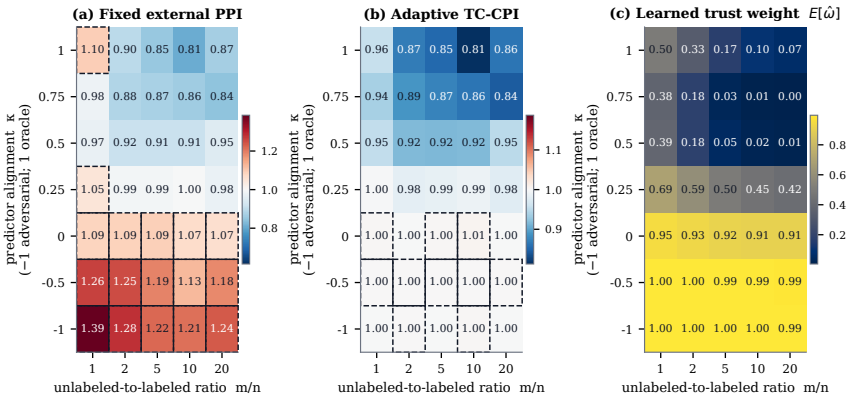


Figure 3: Calibration phase diagram (1,000 repetitions per cell). External PPI can lose under adversarial prediction (left); independent calibration adapts $\widehat{\omega}$ toward low variance (center/right). Dashed cells exceed labeled-AIPW variance.

Figure 3 shows calibration moving from external PPI toward transported AIPW as predictions deteriorate; Table 1 reports representative cells. At aligned $m/n = 10$, its variance ratio to labeled AIPW is 0.813. In a separate $\kappa = 1, m/n = 1$ variance-curve experiment, every point lies inside its 95% Monte Carlo band around equation 4; mean $\widehat{\omega}$ is 0.503 versus oracle 0.498. Across the 35 cells, adaptive Wald coverage is 0.920–0.960.

As target shift grows from 0 to 1, $\mathbb{E}_P(r^2)$ rises from 1.00 to 3.50 and residual noise from 4.67 to 22.44; the r surrogate reaches only 5.10. Empirical values stay within 1.02 Monte Carlo standard errors of the r^2 prediction (Appendix Figure 4), directly distinguishing the support limit in Theorem 4.6 from a first-moment heuristic.

Validity with fixed and learned nuisances. In an eight-stratum bounded law spanning $n \in \{80, 200, 500\}$, $m/n \in \{1, 5, 20\}$, and three predictors, 2,000 repetitions per cell give minimum known-support coverage 0.9995 and Wald coverage 0.9395–0.9580 (Appendix Figure 6). In a five-fold nonlinear Gaussian-mixture experiment (n = 1,000, m = 2,000, 500 repetitions), bias tracks equation 7 across 33 configurations (Pearson $r = 0.997$). With correct f and joint e, r misspecification, bias $(0, -0.051, -0.102)$ at $\omega = (0, 0.5, 1)$ matches drift $(0, -0.052, -0.103)$, while coverage falls from 0.960 to 0.720. A rate experiment (n = 4,000, m = 16,000, 3,000 repetitions) gives coverage 0.942, 0.903, and 0.076 below, at, and above $n^{-1/2}$ (Appendix Figures 8–9).

Ablations and sensitivity. Each comparison removes one ingredient. At a fixed calibrated weight, deleting residual correction changes bias from 0.006 (MCSE 0.004) to -0.632, the exact -0.633 omission bias, and coverage from 0.953 to zero. Replacing r^2 by r misses most source-noise inflation; omitting sensitivity expansion lowers worst-sign coverage from 0.971 to 0.424 at radius 0.4. Under ratio error, expansion raises coverage at $\omega = 1$ from 0.723 to 0.967. Historical, label-only, treatment-only, and joint designs have theoretical n-scaled variances 60.889, 21.051, 40.277, and 13.754, isolating both design modules; increasing the pilot from 8 to 512 observations per stratum–arm cell lowers median excess source variance from 0.697% to 0.017%. These tests map to target preservation, Theorem 4.6, and Corollaries 4.1–4.2; Appendix Figures 7–11 give the full checks.

Randomized outcomes. We use a fixed 500,000-row subsample of the 13,979,592-row CRITEO-UPLIFTv2 randomized-trial release (Diemert et al., 2021). Constant, logistic, boosted-tree, and MLP predictors use disjoint training, calibration, inference, and benchmark blocks. MLP minimizes unweighted arm-macro Brier score (0.0273), but logistic minimizes the repeated-split empirical variance ratio to labeled AIPW (0.997); calibration moves the boosted-tree ratio from 1.020 to 1.004 (Appendix Figure 12 and Table 8). The public `AER::STAR` extract of Project STAR (Krueger, 1999) yields a 2,627-pupil analysis cohort after our filtering of 2,785 complete two-arm records; 1,193 free-lunch pupils define the target. Across 500 honest splits, adaptive, labeled, and external SDs are 5.93, 5.94, and 5.96: the near-flat precision surface correctly yields no gain (Appendix Figure 13 and Table 9).

6 Discussion

PPI++ tunes prediction power under a common feature law (Angelopoulos et al., 2023b), whereas prediction-powered trial generalization transports effects across populations (Demirel et al., 2024). TC-CPI instead calibrates two representations of a common prediction mean across feature laws. Change of measure and full AIPW correction preserve the estimand; variance, drift, and variational arguments yield the oracle weight, the r^2 noise floor, and the optimal label law.

Three results overturn familiar intuitions: more target covariates can make a biased plug-in appear more precise; predictive accuracy need not rank causal precision; and cross-fitting alone does not correct ratio imbalance or weak source support. With correct nuisance weights and adequate overlap, an interior calibration can outperform both endpoints. Simulations reproduce these predictions; Criteo and STAR show that precision gains can be negligible.

These limits become design problems: estimated ratios require explicit balance control and sensitivity analysis; missing support requires target restriction or new labels. Prediction, inference, and label acquisition should be co-designed as deployment populations evolve.

AI USE STATEMENT

Generative AI tools were used for literature triage, manuscript restructuring and language editing, derivation checking, drafting experiment and figure code, and scaffolding Lean formalizations. Cited claims were checked against primary sources, numerical claims were regenerated from fixed-seed code, mathematical derivations were independently rechecked, and the reported Lean files were kernel-verified without placeholders. The authors take responsibility for the final content, including all AI-assisted text, claims, code, and artifacts.

ETHICS STATEMENT

This work introduces an inference method and uses simulated data plus de-identified advertising and education datasets. It collects no new human-subject data. Applied use requires domain evidence for transportability and overlap, together with explicit support checks; restricted-target or partial-identification analyses are appropriate when these conditions fail.

REPRODUCIBILITY STATEMENT

The submission source contains executable scripts, fixed seeds, environment information, aggregate outputs, selected replicate-level records, generated tables, vector figures, complete proofs, and a pinned Lean 4 project. Appendix C maps every experiment to a theorem and records the data-processing and splitting protocol. Licensed Criteo rows are not redistributed; the source records their official URL, license, checksum, and deterministic analysis command. Every reported numerical value is checked against a saved summary.

REFERENCES

Anastasios N. Angelopoulos, Stephen Bates, Clara Fannjiang, Michael I. Jordan, and Tijana Zrnic. Prediction-powered inference. *Science*, 382(6671):669–674, 2023a. doi: 10.1126/science.adi6000.

Anastasios N. Angelopoulos, John C. Duchi, and Tijana Zrnic. PPI++: Efficient prediction-powered inference. arXiv preprint arXiv:2311.01453, 2023b. URL https://arxiv.org/abs/2311.01453.

David Arbour, Eli Ben-Michael, Avi Feller, Apoorva Lal, and Lo-Hua Yuan. AI-assisted variance reduction in randomized experiments. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2*, pp. 21–32. ACM, 2026. doi: 10.1145/3770855.3817896.

Heejung Bang and James M. Robins. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61(4):962–973, 2005. doi: 10.1111/j.1541-0420.2005.00377.x.

Riccardo Cadei, Ilker Demirel, Piersilvio De Bartolomeis, Lukas Lindorfer, Sylvia Cremer, Cordelia Schmid, and Francesco Locatello. Prediction-powered causal inferences. In *Advances in Neural Information Processing Systems*, volume 38, pp. 82200–82229. Neural Information Processing Systems Foundation, Inc. (NeurIPS), 2025. doi: 10.52202/085713-2479.

Abhishek Chakrabortty and Guorong Dai. A general framework for treatment effect estimation in semi-supervised and high dimensional settings. arXiv preprint arXiv:2201.00468, 2022. URL https://arxiv.org/abs/2201.00468.

Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1):C1–C68, 2018. doi: 10.1111/ectj.12097.

Ilker Demirel, Ahmed Alaa, Anthony Philippakis, and David Sontag. Prediction-powered generalization of causal inferences. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 10385–10408. PMLR, 2024.

Eustache Diemert, Artem Betlei, Christophe Renaudin, Massih-Reza Amini, Théophane Gregoir, and Thibaud Rahier. A large scale benchmark for individual treatment effect prediction and uplift modeling. *arXiv preprint arXiv:2111.10106*, 2021. doi: 10.48550/arXiv.2111.10106. URL https://arxiv.org/abs/2111.10106.

Ying Jin, Naoki Egami, and Dominik Rothenhäusler. Beyond reweighting: On the predictive role of covariate shift in effect generalization. *Proceedings of the National Academy of Sciences*, 122(45): e2427181122, 2025. doi: 10.1073/pnas.2427181122. URL https://www.pnas.org/doi/10.1073/pnas.2427181122.

Fredrik Johansson, Uri Shalit, and David Sontag. Learning representations for counterfactual inference. In *Proceedings of the 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 3020–3029. PMLR, 2016.

Fredrik D. Johansson, Uri Shalit, Nathan Kallus, and David Sontag. Generalization bounds and representation learning for estimation of potential outcomes and causal effects. *Journal of Machine Learning Research*, 23(166):1–50, 2022.

Masahiro Kato. Prediction-powered causal inference by automatic debiased machine learning and semi-supervised Riesz regression. arXiv preprint arXiv:2606.12892, 2026. URL https://arxiv.org/abs/2606.12892.

Alan B. Krueger. Experimental estimates of education production functions. *The Quarterly Journal of Economics*, 114(2):497–532, 1999. doi: 10.1162/003355399556052.

Andreas Maurer and Massimiliano Pontil. Empirical Bernstein bounds and sample variance penalization. arXiv preprint arXiv:0907.3740, 2009. URL https://arxiv.org/abs/0907.3740.

James M. Robins, Andrea Rotnitzky, and Lue Ping Zhao. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association*, 89(427): 846–866, 1994. doi: 10.1080/01621459.1994.10476818.

Uri Shalit, Fredrik D. Johansson, and David Sontag. Estimating individual treatment effect: Generalization bounds and algorithms. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 3076–3085. PMLR, 2017.

Lars van der Laan and Mark van der Laan. Calibeating prediction-powered inference. arXiv preprint arXiv:2604.21260, 2026. URL https://arxiv.org/abs/2604.21260.

Chao Ying, Jun Jin, Yi Guo, Xiudi Li, Muxuan Liang, and Jiwei Zhao. Towards the efficient inference by incorporating automated computational phenotypes under covariate shift. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 72505–72534. PMLR, 2025.

Appendix Contents

A Proofs and technical details **12**

A.1 Proof conventions and change of measure 12

A.2 A conditional residual identity 13

A.3 Target preservation and exact variance 13

A.4 Independent calibration 15

A.5 Finite-calibration oracle inequality 16

A.6 Exact finite-sample intervals 18

A.7 Estimated nuisances: exact drift and asymptotic expansion 19

A.8 Approximate transportability and robust calibration 23

A.9 Efficiency lower bounds under transport 29

A.10 From the information bound to active data acquisition 32

A.11 Worked two-stratum calculation 40

B General causal moments and quantiles **40**

B.1 Exact decomposition for transported causal moments 40

B.2 Finite-sample transfer for vector moments 42

B.3 Finite-grid distribution bands 44

B.4 Quantile inversion with grid and support boundaries 45

C Experimental protocol and additional results **46**

D Machine-checked algebraic core **57**

A Proofs and technical details

This appendix separates exact finite-sample statements from asymptotic ones. Throughout, expectations are taken under the data-generating laws specified in Section 2. Estimated propensities and density ratios are treated explicitly in Section 4.3 and its proof below. We write $e = e_1$ and $e_0 = 1 - e$.

A.1 Proof conventions and change of measure

The labeled observations $W_i = (X_i, T_i, Y_i)$ are iid from P, the target covariates $\widetilde{X}_j$ are iid from Q_X, and the two arrays are independent. Whenever a predictor or nuisance estimate is trained on an external sample, all calculations are first conditional on the corresponding training sigma-field. This convention turns the fitted functions into fixed measurable maps; unconditional statements then follow by the tower property. This conditioning yields exact conditional statements; each asymptotic argument states its additional moment and convergence conditions explicitly.

We repeatedly use the following change-of-measure identity.

Lemma A.1 (Change of covariate law)**.** *Let $Q_X \ll P_X$ and let $r = dQ_X/dP_X$. If g is measurable and $\mathbb{E}_Q|g(X)| < \infty$, then $r(X)g(X)$ is P-integrable and*

$$\mathbb{E}_P\{r(X)g(X)\} = \mathbb{E}_Q\{g(X)\}.$$

If $g \in L_2(Q)$, then $\mathbb{E}_P\{r(X)g^2(X)\} = \mathbb{E}_Q\{g^2(X)\}$. By contrast, finiteness of $\mathbb{E}_P\{r^2(X)g^2(X)\}$ is a strictly stronger weighted-overlap condition and is not implied by absolute continuity alone.

Proof. By the defining property of the Radon–Nikodym derivative,

$$\int r(x)g^+(x)\,dP_X(x) = \int g^+(x)\,dQ_X(x), \qquad \int r(x)g^-(x)\,dP_X(x) = \int g^-(x)\,dQ_X(x).$$

Both right-hand sides are finite because $g \in L_1(Q)$. Subtracting the two equalities proves the claim. Applying the same argument to g^2 proves the $L_2(Q)$ identity. An r^2 moment requires an additional weighted-overlap condition. □

A.2 A CONDITIONAL RESIDUAL IDENTITY

Lemma A.2 (Transported residual identity)**.** *Under Assumptions 2.1–2.3, let f_0, f_1 be fixed measurable functions such that, for $t = 0, 1$, $\mathbb{E}_P[r(X)\mathbf{1}\{T = t\}|Y - f_t(X)|/e_t(X)] < \infty$ and $\mathbb{E}_Q\{|f_t(X)| + |\mu_t(X)|\} < \infty$. Then*

$$\mathbb{E}_P\left[r(X)\left\{\frac{T}{e(X)}(Y - f_1(X)) - \frac{1-T}{1-e(X)}(Y - f_0(X))\right\}\right] = \mathbb{E}_Q\{\tau(X) - \tau_f(X)\},$$

where $\tau = \mu_1 - \mu_0$ and $\tau_f = f_1 - f_0$.

Proof. Fix $t \in \{0, 1\}$. Positivity makes the following conditional ratio well-defined P_X-almost surely. By iterated expectation,

$$\begin{aligned}&\mathbb{E}_P\left[\frac{\mathbf{1}\{T = t\}}{e_t(X)}\{Y - f_t(X)\}\,\middle|\,X\right] \\ &\quad= \frac{1}{e_t(X)}\Pr_P(T = t \mid X)\mathbb{E}_P\{Y - f_t(X) \mid X, T = t\}.\end{aligned}$$

Consistency replaces Y by $Y(t)$ on $\{T = t\}$, and conditional exchangeability removes the conditioning on $T = t$. Since $\Pr_P(T = t \mid X) = e_t(X)$, the last display equals $\mathbb{E}_P\{Y(t) \mid X\} - f_t(X) = \mu_t(X) - f_t(X)$. The stated absolute-integrability condition justifies this application of the tower property after multiplication by $r(X)$. Consequently,

$$\begin{aligned}\mathbb{E}_P\left[r(X)\frac{\mathbf{1}\{T = t\}}{e_t(X)}\{Y - f_t(X)\}\right] &= \mathbb{E}_P\left[r(X)\{\mu_t(X) - f_t(X)\}\right] \\ &= \mathbb{E}_Q\{\mu_t(X) - f_t(X)\},\end{aligned}$$

where the second equality is Lemma A.1. Subtracting the $t = 0$ equality from the $t = 1$ equality gives $\mathbb{E}_Q\{\mu_1 - \mu_0 - (f_1 - f_0)\}$, which is the claimed right-hand side. □

A.3 TARGET PRESERVATION AND EXACT VARIANCE

Proof of Theorem 3.1. Recall $Z_f(W) = r(X)[T\{Y - f_1(X)\}/e(X) - (1-T)\{Y - f_0(X)\}/\{1 - e(X)\}]$. All three empirical averages in equation 2 are integrable under the theorem's hypotheses. Lemma A.2 and Lemma A.1 give, respectively,

$$\mathbb{E}_P Z_f = \theta_Q - \mathbb{E}_Q\tau_f, \qquad \mathbb{E}_P\{r(X)\tau_f(X)\} = \mathbb{E}_Q\tau_f(X). \tag{12}$$

The first equality says that the labeled residual term repairs the entire prediction error; it is not multiplied by ω. The second says that the two prediction-mean estimators have the same population target. Therefore, for every deterministic $\omega \in \mathbb{R}$,

$$\begin{aligned}\mathbb{E}(\widehat{\theta}_\omega) &= \omega\mathbb{E}_P(r\tau_f) + (1-\omega)\mathbb{E}_Q\tau_f + \mathbb{E}_P Z_f \\ &= \mathbb{E}_Q\tau_f + \theta_Q - \mathbb{E}_Q\tau_f = \theta_Q.\end{aligned}$$

No cancellation in the last line uses predictor correctness.

Now suppose f is random and let $\mathcal{T}$ be its training sigma-field. Independence of training and evaluation implies that, conditional on $\mathcal{T}$, the functions f_t are fixed and the evaluation samples still have laws $P^n \otimes Q^m$. Applying the deterministic calculation on every training realization for which the integrability conditions hold gives $\mathbb{E}(\widehat{\theta}_\omega \mid \mathcal{T}) = \theta_Q$ almost surely. The tower property then gives unconditional unbiasedness. Weights chosen from the evaluation outcomes fall outside this conditional argument; Theorem 4.2 restores it through an independent calibration sigma-field. □

Proof of Theorem 4.1. Condition on the external training sigma-field $\mathcal{T}$. Then f is fixed; write $U = Z_f(W)$, $R = r(X)\tau_f(X)$, and $S = \tau_f(\widetilde{X})$. Using equation 12, the centered estimator can be written exactly as

$$\widehat{\theta}_\omega - \theta_Q = \frac{1}{n}\sum_{i=1}^{n}\left[U_i + \omega R_i - \mathbb{E}_P\{U + \omega R\}\right] + \frac{1-\omega}{m}\sum_{j=1}^{m}\{S_j - \mathbb{E}_Q S\}. \tag{13}$$

The two sums are conditionally independent. Independence within each sample and the assumed second moments therefore imply

$$\mathrm{Var}(\widehat{\theta}_\omega \mid \mathcal{T}) = \frac{1}{n}\mathrm{Var}_P(U + \omega R) + \frac{(1-\omega)^2}{m}\mathrm{Var}_Q(S).$$

Expanding the first variance gives

$$V(\omega) = \frac{A + 2\omega C_{UR} + \omega^2 B_P}{n} + \frac{(1-\omega)^2 B_{Q,f}}{m},$$

with $A = \mathrm{Var}_P(U)$, $C_{UR} = \mathrm{Cov}_P(U, R)$, $B_P = \mathrm{Var}_P(R)$, and $B_{Q,f} = \mathrm{Var}_Q(S)$. Put $D = B_P/n + B_{Q,f}/m$ and $b = B_{Q,f}/m - C_{UR}/n$. The second derivative is $V''(\omega) = 2D \geq 0$. When $D > 0$, direct completion of the square yields

$$V(\omega) = V(b/D) + D(\omega - b/D)^2, \qquad V'(\omega) = 2D\omega - 2b.$$

Hence the unconstrained minimizer is b/D. A convex quadratic decreases up to that point and increases after it, so its minimizer over the closed interval $[0, 1]$ is the Euclidean projection $\omega^\star = \Pi_{[0,1]}(b/D)$. Equivalently, the three regimes are

$$\omega^\star = \begin{cases} 0, & b \leq 0, \\ b/D, & 0 < b < D, \\ 1, & b \geq D. \end{cases}$$

Because 0 and 1 are feasible, global optimality immediately gives $V(\omega^\star) \leq \min\{V(0), V(1)\}$.

If $D = 0$, nonnegativity of variances forces $B_P = B_{Q,f} = 0$. Cauchy–Schwarz gives $|C_{UR}|^2 \leq AB_P = 0$, hence $C_{UR} = 0$ and the displayed variance is the constant A/n. Thus every $\omega \in [0, 1]$ is optimal; the tie convention in the algorithm changes neither the estimand nor the variance.

Finally, $\widehat{\theta}_0 = \mathbb{Q}_m\tau_f + \mathbb{P}_n U$ is external-sample PPI, whereas $\widehat{\theta}_1 = \mathbb{P}_n(R + U)$ is transported labeled AIPW. If f is random but externally trained, the conditional mean is the constant θ_Q. The law of total variance then gives

$$\mathrm{Var}(\widehat{\theta}_\omega) = \mathbb{E}\{\mathrm{Var}(\widehat{\theta}_\omega \mid \mathcal{T})\} + \mathrm{Var}\{\mathbb{E}(\widehat{\theta}_\omega \mid \mathcal{T})\} = \mathbb{E}\{V(\omega)\}.$$

□

Proof of Proposition 3.1. When $P_X = Q_X = M$, $r = 1$ and $R = S = \tau_f$. Since $H_f = U + R$,

$$\widehat{\theta}_\omega = \mathbb{P}_n(U + R) + (1-\omega)\{\mathbb{Q}_m R - \mathbb{P}_n R\},$$

which is equation 3. Writing $B = \mathrm{Var}_M(R)$ and $\gamma = 1 - \omega$, the variance as a function of γ is

$$\frac{\mathrm{Var}_P(H_f)}{n} - \frac{2\gamma\mathrm{Cov}_P(H_f, R)}{n} + \gamma^2 B\left(\frac{1}{n} + \frac{1}{m}\right).$$

To obtain this display, note that the labeled contribution is $H_f - \gamma R$ and the target contribution is γR, so the covariance term appears only in the labeled variance. For $B > 0$, differentiating gives $-2\mathrm{Cov}_P(H_f, R)/n + 2\gamma B(1/n + 1/m) = 0$. The unconstrained solution is $\{m/(n+m)\}\mathrm{Cov}_P(H_f, R)/B$; projection onto $[0, 1]$ gives the stated $\gamma^\star$. If $B = 0$, then R is constant M-almost surely and Cauchy–Schwarz gives $\mathrm{Cov}_P(H_f, R) = 0$, so γ does not affect the variance. This is an algebraic reduction at a shared marginal law. It does not identify the unrestricted external-target experiment, whose tangent space allows the two marginal laws to vary separately, with the common-population PPI++ experiment; the distinct information bounds are derived below. □

A.4 INDEPENDENT CALIBRATION

Let $\mathcal{C} = (\mathcal{C}_P, \mathcal{C}_Q)$ denote independent calibration samples of sizes n_c, m_c. On $\mathcal{C}$, compute the usual sample covariance estimators $\widehat{A}, \widehat{C}_{UR}, \widehat{B}_P, \widehat{B}_{Q,f}$ and

$$\widehat{\omega} = \Pi_{[0,1]}\left[\frac{\widehat{B}_{Q,f}/m - \widehat{C}_{UR}/n}{\widehat{B}_P/n + \widehat{B}_{Q,f}/m}\right],$$

with a fixed endpoint tie rule if the denominator is zero. The inference samples used in $\widehat{\theta}_{\widehat{\omega}}$ are independent of $\mathcal{C}$.

Lemma A.3 (Stability of the projected calibration map)**.** *Suppose* $m/n \to \lambda \in (0, \infty)$ *and* $(\widehat{C}_{UR}, \widehat{B}_P, \widehat{B}_{Q,f}) \to_p (C_{UR}, B_P, B_{Q,f})$. *If* $D_\lambda = B_P + \lambda^{-1} B_{Q,f} > 0$, *then*

$$\widehat{\omega} \to_p \omega^\star_\lambda.$$

Moreover, if V_n *denotes the finite-*(n, m) *variance and* $\omega^\star_n \in \arg\min_{[0,1]} V_n$, *then*

$$n\{V_n(\widehat{\omega}) - V_n(\omega^\star_n)\} = o_p(1).$$

Proof. Multiplication of numerator and denominator in the unprojected rule by n shows that its random argument is $\widehat{a}_n = \{(n/m)\widehat{B}_{Q,f} - \widehat{C}_{UR}\}/\{\widehat{B}_P + (n/m)\widehat{B}_{Q,f}\}$. The denominator converges in probability to $D_\lambda > 0$, while the numerator converges to $\lambda^{-1} B_{Q,f} - C_{UR}$. Slutsky's theorem and the one-Lipschitz continuity of $\Pi_{[0,1]}$ prove the first claim.

For the second claim, define

$$v_n(\omega) = nV_n(\omega) = A + 2\omega C_{UR} + \omega^2 B_P + \frac{n}{m}(1-\omega)^2 B_{Q,f}.$$

Its coefficients converge to those of v_λ, hence $\sup_{\omega\in[0,1]} |v_n(\omega) - v_\lambda(\omega)| \to 0$. Because $D_\lambda > 0$, the limiting quadratic has the unique minimizer $\omega^\star_\lambda$. The explicit projection formula also gives $\omega^\star_n \to \omega^\star_\lambda$. Therefore both $\widehat{\omega}$ and $\omega^\star_n$ converge to the same point. Since v_n has uniformly bounded derivative on $[0, 1]$ for all sufficiently large n, the mean-value theorem yields $|v_n(\widehat{\omega}) - v_n(\omega^\star_n)| = o_p(1)$. Equivalently, the displayed scaled regret is $o_p(1)$. □

Proof of Theorem 4.2. Let $\mathcal{G} = \sigma(\mathcal{T}, \mathcal{C})$ collect external training and calibration information. Conditional on $\mathcal{G}$, $\widehat{\omega}$ is fixed and the inference samples retain product law $P^n \otimes Q^m$. Theorem 3.1 therefore gives the exact identity $\mathbb{E}(\widehat{\theta}_{\widehat{\omega}} \mid \mathcal{G}) = \theta_Q$. This is stronger than asymptotic unbiasedness and is the reason for using an independent calibration block.

The usual sample covariance and variance estimators are consistent by the law of large numbers under the stated moment condition. Lemma A.3 then proves both $\widehat{\omega} \to_p \omega^\star_\lambda$ and the scaled oracle-regret statement.

We next verify the random-weight central limit theorem. Let $V_n(w)$ denote the finite-(n, m) variance in equation 4, and, conditionally on $\mathcal{G}$, define the independent centered array

$$X_{n,i}(w) = \frac{U_i + wR_i - \mathbb{E}_P(U + wR)}{n\sqrt{V_n(w)}}, \qquad 1 \le i \le n,$$
$$Y_{n,j}(w) = \frac{(1-w)\{S_j - \mathbb{E}_Q S\}}{m\sqrt{V_n(w)}}, \qquad 1 \le j \le m.$$

Equation equation 13 gives the exact identity

$$\frac{\widehat{\theta}_{\widehat{\omega}} - \theta_Q}{\sqrt{V_n(\widehat{\omega})}} = \sum_{i=1}^{n} X_{n,i}(\widehat{\omega}) + \sum_{j=1}^{m} Y_{n,j}(\widehat{\omega}), \tag{14}$$

and the conditional variance of the right-hand side is one.

To check Lyapunov's condition, the convexity inequality $|a+b|^{2+\eta} \le 2^{1+\eta}(|a|^{2+\eta} + |b|^{2+\eta})$ and $0 \le \widehat{\omega} \le 1$ give a common conditional bound for the $(2+\eta)$ moments of $U + \widehat{\omega}R - \mathbb{E}_P(U + \widehat{\omega}R)$

and $(1-\widehat{\omega})(S-\mathbb{E}_Q S)$. Lemma A.3 and $v_\lambda(\omega_\lambda^\star)>0$ imply $nV_n(\widehat{\omega}) \to_p v_\lambda(\omega_\lambda^\star)>0$. Thus, on events whose probability tends to one, for a finite constant K,

$$\begin{aligned}\sum_{i=1}^{n}\mathbb{E}\{|X_{n,i}(\widehat{\omega})|^{2+\eta}\mid\mathcal{G}\}+\sum_{j=1}^{m}\mathbb{E}\{|Y_{n,j}(\widehat{\omega})|^{2+\eta}\mid\mathcal{G}\}\\ \le K\left\{\frac{n}{n^{2+\eta}V_n(\widehat{\omega})^{1+\eta/2}}+\frac{m}{m^{2+\eta}V_n(\widehat{\omega})^{1+\eta/2}}\right\}\\ =O_p(n^{-\eta/2})+O_p\{n^{1+\eta/2}m^{-1-\eta}\}=o_p(1),\end{aligned}$$

because $m/n\to\lambda\in(0,\infty)$. Conditional Lyapunov–Feller therefore gives, for every real t,

$$\mathbb{E}\left[\exp\left\{it\frac{\widehat{\theta}_{\widehat{\omega}}-\theta_Q}{\sqrt{V_n(\widehat{\omega})}}\right\}\middle|\mathcal{G}\right]\to_p e^{-t^2/2}.$$

The conditional characteristic functions have modulus at most one. Every subsequence has a further subsequence on which the last convergence holds almost surely; dominated convergence on that further subsequence gives the unconditional characteristic-function limit. The subsequence principle and Lévy's theorem yield the unconditional $N(0,1)$ limit.

For studentization, write

$$\widehat{V}_n=\frac{S_{P,n}^2(\widehat{\omega})}{n}+\frac{(1-\widehat{\omega})^2S_{Q,m}^2}{m},\tag{15}$$

where $S_{P,n}^2(w)$ is the Bessel-corrected sample variance of U_i+wR_i and $S_{Q,m}^2$ is that of S_j. Expanding the first sample variance as a quadratic polynomial in w reduces its coefficients to the sample variance of U, the sample covariance of (U,R), and the sample variance of R. Conditional laws of large numbers make these three coefficients converge to (A,C_{UR},B_P), while $S_{Q,m}^2\to_p B_{Q,f}$. Because $\widehat{\omega}\in[0,1]$, this coefficient convergence is uniform in w, and hence $n\{\widehat{V}_n-V_n(\widehat{\omega})\}=o_p(1)$. The positive limit of $nV_n(\widehat{\omega})$ implies $\widehat{V}_n/V_n(\widehat{\omega})\to_p 1$. Slutsky's theorem applied to equation 14 proves the studentized conclusion. If the calibration moments are $\sqrt{n_c\wedge m_c}$-consistent and the limiting oracle is interior, then the projection is locally inactive and $\widehat{\omega}-\omega_n^\star=O_p((n_c\wedge m_c)^{-1/2})$. The exact square-completion identity upgrades the scaled regret to $O_p((n_c\wedge m_c)^{-1})$. □

A.5 Finite-calibration oracle inequality

Lemma A.4 (Bounded order-two U-statistic)**.** *Let $X_1,\dots,X_N$ be iid, $N\ge 2$, and let h be a symmetric kernel whose range has length H. Then*

$$\Pr\left(\left|\binom{N}{2}^{-1}\sum_{i<j}h(X_i,X_j)-\mathbb{E}h(X_1,X_2)\right|\ge t\right)\le 2\exp\left\{-\frac{2\lfloor N/2\rfloor t^2}{H^2}\right\}.$$

Proof. Put $k=\lfloor N/2\rfloor$ and let Π be a uniform random permutation of $\{1,\dots,N\}$, independent of the observations. For a deterministic permutation π, define the paired average $T_\pi=k^{-1}\sum_{\ell=1}^{k}h(X_{\pi(2\ell-1)},X_{\pi(2\ell)})$. Every unordered pair $\{i,j\}$ appears in a uniformly permuted matching with the same probability. Counting the k positions in the matching therefore gives the exact symmetrization identity

$$U_N:=\binom{N}{2}^{-1}\sum_{i<j}h(X_i,X_j)=\mathbb{E}_\Pi(T_\Pi\mid X_1,\dots,X_N).$$

Let $\mu=\mathbb{E}h(X_1,X_2)$. Convexity of the exponential and the preceding identity imply, for $s>0$,

$$\begin{aligned}\mathbb{E}\exp\{s(U_N-\mu)\}&\le\mathbb{E}\mathbb{E}_\Pi\exp\{s(T_\Pi-\mu)\mid X_1,\dots,X_N\}\\&=\mathbb{E}\exp\{s(T_{\pi_0}-\mu)\},\end{aligned}$$

where the last equality uses exchangeability and any fixed permutation π_0. For fixed π_0, the k kernel evaluations use disjoint pairs of iid observations and are therefore independent. Each has range length at most H. Hoeffding's lemma applied to their average yields

$$\mathbb{E}\exp\{s(T_{\pi_0}-\mu)\}\le\exp\left(\frac{s^2H^2}{8k}\right).$$

Chernoff's method now gives

$$\Pr(U_N - \mu \geq t) \leq \inf_{s>0} \exp\left(-st + \frac{s^2 H^2}{8k}\right) = \exp\left(-\frac{2kt^2}{H^2}\right),$$

where the optimum is $s = 4kt/H^2$. Applying the same calculation to $-h$ and taking a union bound proves the two-sided inequality. □

The calibration estimators in Theorem 4.3 are

$$\widehat{C}_{UR} = \frac{1}{n_c(n_c-1)} \sum_{i<j} (U_i - U_j)(R_i - R_j),$$
$$\widehat{B}_P = \frac{1}{n_c(n_c-1)} \sum_{i<j} (R_i - R_j)^2,$$
$$\widehat{B}_{Q,f} = \frac{1}{m_c(m_c-1)} \sum_{i<j} (S_i - S_j)^2.$$

They are the usual unbiased sample covariance and variances. Define

$$\epsilon_{UR} = H_U H_R \sqrt{\frac{\log(6/\delta)}{2\lfloor n_c/2 \rfloor}}, \qquad \epsilon_P = H_R^2 \sqrt{\frac{\log(6/\delta)}{8\lfloor n_c/2 \rfloor}},$$
$$\epsilon_Q = H_S^2 \sqrt{\frac{\log(6/\delta)}{8\lfloor m_c/2 \rfloor}}. \tag{16}$$

Proof of Theorem 4.3. When written as U-statistics with normalization $\binom{N}{2}^{-1}$, the three kernels are $\frac{1}{2}(U_i - U_j)(R_i - R_j)$, $\frac{1}{2}(R_i - R_j)^2$, and $\frac{1}{2}(S_i - S_j)^2$. The first lies in an interval of length at most $H_U H_R$, and the last two in intervals of lengths at most $H_R^2/2$ and $H_S^2/2$. Their expectations are $C_{UR}, B_P, B_{Q,f}$ because, for independent copies, $\mathbb{E}\{(S_1 - S_2)(T_1 - T_2)/2\} = \mathrm{Cov}(S, T)$. Applying Lemma A.4 with failure probability $\delta/3$ to each kernel and taking a union bound gives, with probability at least $1 - \delta$,

$$|\widehat{C}_{UR} - C_{UR}| \leq \epsilon_{UR}, \qquad |\widehat{B}_P - B_P| \leq \epsilon_P, \qquad |\widehat{B}_{Q,f} - B_{Q,f}| \leq \epsilon_Q. \tag{17}$$

The A/n term in $V(\omega)$ is constant. Put

$$\mathcal{Q}(\omega) = \frac{2C_{UR}\omega + B_P\omega^2}{n} + \frac{B_{Q,f}(1-\omega)^2}{m},$$

and define $\widehat{\mathcal{Q}}$ analogously. On the event in equation 17, uniformly over $[0, 1]$,

$$|\widehat{\mathcal{Q}}(\omega) - \mathcal{Q}(\omega)| \leq \frac{2\epsilon_{UR}}{n} + \frac{\epsilon_P}{n} + \frac{\epsilon_Q}{m}.$$

Let the right-hand side of the last display be Δ_n. Empirical and population optimality give the deterministic comparison

$$\begin{aligned}\mathcal{Q}(\widehat{\omega}) - \mathcal{Q}(\omega^\star) &= \{\mathcal{Q} - \widehat{\mathcal{Q}}\}(\widehat{\omega}) + \{\widehat{\mathcal{Q}}(\widehat{\omega}) - \widehat{\mathcal{Q}}(\omega^\star)\} + \{\widehat{\mathcal{Q}} - \mathcal{Q}\}(\omega^\star) \\ &\leq 2\Delta_n,\end{aligned}$$

because the middle term is nonpositive. Since the omitted A/n term is constant in ω, this proves the always-valid bound

$$V(\widehat{\omega}) - V(\omega^\star) \leq \frac{4\epsilon_{UR}}{n} + \frac{2\epsilon_P}{n} + \frac{2\epsilon_Q}{m}. \tag{18}$$

For the curvature-dependent rate, write $D = B_P/n + B_{Q,f}/m$ and $b = B_{Q,f}/m - C_{UR}/n$, with hatted analogues. The sample variances ensure $\widehat{D} \geq 0$. The derivative of $\mathcal{Q}$ is $2(D\omega - b)$. First-order optimality on the convex set $[0, 1]$ says $\mathcal{Q}'(\omega^\star)(\omega - \omega^\star) \geq 0$ for every feasible ω. Taylor's identity for the quadratic therefore gives

$$\mathcal{Q}(\widehat{\omega}) - \mathcal{Q}(\omega^\star) \geq D(\widehat{\omega} - \omega^\star)^2.$$

Apply the analogous variational inequality to $\widehat{\mathcal{Q}}$ at $\widehat{\omega}$, add the two inequalities, and use $\{\mathcal{Q}'(w) - \widehat{\mathcal{Q}}'(w)\}/2 = (D - \widehat{D})w - (b - \widehat{b})$. Since $0 \le w \le 1$, this gives

$$\mathcal{Q}(\widehat{\omega}) - \mathcal{Q}(\omega^\star) \le 2\{|D - \widehat{D}| + |b - \widehat{b}|\}|\widehat{\omega} - \omega^\star|.$$

If $\widehat{\omega} = \omega^\star$ the fast bound is immediate. Otherwise, combining the last two displays and cancelling one distance gives

$$|\widehat{\omega} - \omega^\star| \le \frac{2}{D}\{|D - \widehat{D}| + |b - \widehat{b}|\}.$$

Substitution into the upper comparison yields, when $D > 0$,

$$V(\widehat{\omega}) - V(\omega^\star) \le \frac{4}{D}\left\{\frac{\epsilon_P + \epsilon_{UR}}{n} + \frac{2\epsilon_Q}{m}\right\}^2. \tag{19}$$

If $m/n \ge \underline{\lambda}$, $nD \ge \underline{d}$, and $N_c = \min\{\lfloor n_c/2 \rfloor, \lfloor m_c/2 \rfloor\}$, the braces are at most

$$\frac{1}{n}\sqrt{\frac{\log(6/\delta)}{2N_c}}\left(H_U H_R + \frac{H_R^2}{2} + \frac{H_S^2}{\underline{\lambda}}\right).$$

Substitution into equation 19, followed by $V(\omega^\star) \le \min\{V(0), V(1)\}$, proves equation 6. Equation equation 18 is the advertised flat-curve alternative. □

A.6 Exact finite-sample intervals

Retain the notation of Theorem 4.4 and condition on the external training data and, when applicable, the independent calibration sample. Then f and ω are fixed, while $S_{P,\omega}$ and $S_{Q,\omega}$ have known almost-sure range lengths $H_{P,\omega}$ and $H_{Q,\omega}$. The Hoeffding radius is

$$h_{\mathrm{H}} = \sqrt{\frac{\log(2/\alpha)}{2}\left(\frac{H_{P,\omega}^2}{n} + \frac{H_{Q,\omega}^2}{m}\right)}. \tag{20}$$

For the variance-adaptive alternative, assume $n, m \ge 2$, write $S_{P,\omega} \in [a_P, a_P + H_{P,\omega}]$ and $S_{Q,\omega} \in [a_Q, a_Q + H_{Q,\omega}]$ for the known conditional ranges, put $N = n + m$, and $R_0 = \max(H_{P,\omega}/n, H_{Q,\omega}/m)$, and normalize the N independent weighted summands to $Z_k \in [0, 1]$ by subtracting a_P/n or a_Q/m and dividing by R_0. If S_Z^2 is their Bessel-corrected sample variance, define

$$h_{\mathrm{EB}} = R_0 N\left\{\sqrt{\frac{2S_Z^2\log(4/\alpha)}{N}} + \frac{7\log(4/\alpha)}{3(N-1)}\right\}, \tag{21}$$

with $h_{\mathrm{EB}} = 0$ when $R_0 = 0$. The two confidence intervals asserted in Theorem 4.4 are

$$I_{\mathrm{H}} = [\widehat{\theta}_\omega - h_{\mathrm{H}}, \widehat{\theta}_\omega + h_{\mathrm{H}}], \qquad I_{\mathrm{EB}} = [\widehat{\theta}_\omega - h_{\mathrm{EB}}, \widehat{\theta}_\omega + h_{\mathrm{EB}}]. \tag{22}$$

The variance convention and the non-identically distributed form of the inequality are recorded explicitly below because the normalized array combines two samples with different weights.

Lemma A.5 (Two-sided empirical Bernstein inequality)**.** *Let $Z_1, \dots, Z_N$ be independent, not necessarily identically distributed, random variables in $[0, 1]$, where $N \ge 2$, and define*

$$\overline{Z}_N = \frac{1}{N}\sum_{k=1}^N Z_k, \qquad S_Z^2 = \frac{1}{N-1}\sum_{k=1}^N (Z_k - \overline{Z}_N)^2 = \frac{1}{N(N-1)}\sum_{k<\ell}(Z_k - Z_\ell)^2.$$

For every $0 < \alpha < 1$, with probability at least $1 - \alpha$,

$$\left|\overline{Z}_N - \frac{1}{N}\sum_{k=1}^N \mathbb{E}Z_k\right| \le \sqrt{\frac{2S_Z^2\log(4/\alpha)}{N}} + \frac{7\log(4/\alpha)}{3(N-1)}. \tag{23}$$

Proof. Theorem 11 of Maurer & Pontil (2009) applies to independent $[0, 1]$-valued variables without requiring a common distribution. With failure probability $\alpha/2$, it bounds $N^{-1}\sum_k \mathbb{E}Z_k - \overline{Z}_N$ by the right-hand side of equation 23. Apply the same theorem to $1 - Z_1, \dots, 1 - Z_N$. Their pairwise-difference variance is exactly S_Z^2, and the resulting inequality bounds the opposite tail. A union bound proves the two-sided statement. The equality between the two displayed definitions of S_Z^2 follows by expanding $\sum_{k<\ell}(Z_k - Z_\ell)^2 = N\sum_k (Z_k - \overline{Z}_N)^2$. □

Proof of Theorem 4.4. Condition as above and abbreviate $S_{P,i} = S_{P,\omega}(W_i)$ and $S_{Q,j} = S_{Q,\omega}(\widetilde{X}_j)$. By target preservation,

$$\widehat{\theta}_\omega - \theta_Q = \sum_{i=1}^{n} \frac{S_{P,i} - \mathbb{E}S_{P,i}}{n} + \sum_{j=1}^{m} \frac{S_{Q,j} - \mathbb{E}S_{Q,j}}{m}.$$

Hoeffding's inequality for independent, non-identically weighted bounded variables gives

$$\Pr\{|\widehat{\theta}_\omega - \theta_Q| \geq s \mid f, \omega\} \leq 2\exp\left\{-\frac{2s^2}{H_{P,\omega}^2/n + H_{Q,\omega}^2/m}\right\}.$$

Substituting equation 20 makes the right-hand side α.

For the second interval, suppose $R_0 > 0$ and define the normalized array

$$Z_i = \frac{S_{P,i}/n - a_P/n}{R_0}, \quad 1 \leq i \leq n, \qquad Z_{n+j} = \frac{S_{Q,j}/m - a_Q/m}{R_0}, \quad 1 \leq j \leq m.$$

The range assumptions imply $0 \leq S_{P,i}/n - a_P/n \leq H_{P,\omega}/n \leq R_0$ and $0 \leq S_{Q,j}/m - a_Q/m \leq H_{Q,\omega}/m \leq R_0$; hence these $N = n + m$ variables are independent and lie in $[0, 1]$. Moreover, the lower-endpoint constants cancel after centering, so

$$\begin{aligned} R_0 N \left\{ \frac{1}{N}\sum_{k=1}^{N} Z_k - \frac{1}{N}\sum_{k=1}^{N} \mathbb{E}Z_k \right\} &= \sum_{i=1}^{n} \frac{S_{P,i} - \mathbb{E}S_{P,i}}{n} + \sum_{j=1}^{m} \frac{S_{Q,j} - \mathbb{E}S_{Q,j}}{m} \\ &= \widehat{\theta}_\omega - \theta_Q. \end{aligned}$$

Lemma A.5, followed by multiplication by $R_0 N$, therefore gives the radius in equation 21. When $R_0 = 0$, every summand is constant and the zero radius is valid. Integrating either conditional bound proves the unconditional claim. Neither proof uses predictor correctness or an asymptotic approximation; bounded support and known correct weights are essential. □

The preceding argument has a useful pointwise set-valued form. Suppose a target moment satisfies $\Psi(\theta) = \Gamma_f(\theta) + \Delta_f(\theta)$ exactly. For a fixed root θ_0, if random sets $\widehat{\mathcal{G}}(\theta_0)$ and $\widehat{\mathcal{D}}(\theta_0)$ cover the two components with error probabilities α_G and α_D, then the set $\widehat{C} = \{\theta : 0 \in \widehat{\mathcal{G}}(\theta) + \widehat{\mathcal{D}}(\theta)\}$ covers that root θ_0 with probability at least $1 - \alpha_G - \alpha_D$. Indeed, on the intersection of the two component coverage events, $0 = \Gamma_f(\theta_0) + \Delta_f(\theta_0)$ belongs to the Minkowski sum. This is a pointwise transfer principle: simultaneous coverage of an entire root set requires component regions that are themselves uniform in θ.

A.7 Estimated nuisances: exact drift and asymptotic expansion

For arbitrary candidate functions $(\widetilde{f}, \widetilde{e}, \widetilde{r})$, write $\widetilde{e}_1 = \widetilde{e}$ and $\widetilde{e}_0 = 1 - \widetilde{e}$, and let the population version of the calibrated score replace (f, e, r) by those functions.

Lemma A.6 (Exact learned-nuisance drift)**.** *Assume $0 < \widetilde{e}(X) < 1$ P-almost surely and that every term below is integrable. For fixed ω, the difference between the population mean of the candidate score and θ_Q is*

$$\begin{aligned} &\Theta(\omega; \widetilde{f}, \widetilde{e}, \widetilde{r}) - \theta_Q \\ &= \sum_{t=0}^{1} s_t \mathbb{E}_P\left[\left\{\frac{\widetilde{r}e_t}{\widetilde{e}_t} - r\right\}(\mu_t - \widetilde{f}_t) + \omega(\widetilde{r} - r)\widetilde{f}_t\right], \qquad (s_1, s_0) = (1, -1). \end{aligned} \tag{24}$$

Proof. Write the candidate population functional arm by arm. Conditional on X, consistency and exchangeability give

$$\mathbb{E}_P\left[\frac{\mathbf{1}\{T = t\}}{\widetilde{e}_t(X)}\{Y - \widetilde{f}_t(X)\} \,\middle|\, X\right] = \frac{e_t(X)}{\widetilde{e}_t(X)}\{\mu_t(X) - \widetilde{f}_t(X)\}.$$

Hence the signed tth candidate contribution, after subtracting the signed target contribution $\mathbb{E}_P\{rs_t\mu_t\}$, is

$$\omega\mathbb{E}_P(\widetilde{r}s_t\widetilde{f}_t) + (1-\omega)\mathbb{E}_P(rs_t\widetilde{f}_t) + \mathbb{E}_P\left\{\widetilde{r}\frac{e_t}{\widetilde{e}_t}s_t(\mu_t - \widetilde{f}_t)\right\} - \mathbb{E}_P(rs_t\mu_t)$$
$$= s_t\mathbb{E}_P\left[\left\{\frac{\widetilde{r}e_t}{\widetilde{e}_t} - r\right\}(\mu_t - \widetilde{f}_t) + \omega(\widetilde{r}-r)\widetilde{f}_t\right].$$

The second line follows by adding and subtracting $\mathbb{E}_P(rs_t\widetilde{f}_t)$ and collecting the coefficient of $\mu_t - \widetilde{f}_t$. Summing over $t = 0, 1$ proves equation 24. All uses of iterated expectation are justified by the integrability assumption. □

Two boundaries follow immediately. If r is known, the drift is a product of propensity error and prediction error for every ω. At $\omega = 0$, the usual product-form robustness in the combined weight reappears. If r is estimated and $\omega > 0$, the extra term is the first-order scalar imbalance $\omega\mathbb{E}_P\{(\widetilde{r}-r)\widetilde{\tau}_f\}$; ordinary consistency of $\widetilde{r}$ need not make it $o_p(n^{-1/2})$.

Cauchy–Schwarz gives the sufficient bound

$$\left|\mathfrak{D}(\omega;\widetilde{f},\widetilde{e},\widetilde{r})\right| \le \sum_{t=0}^{1}\|\widetilde{r}e_t/\widetilde{e}_t - r\|_{P,2}\left\|\mu_t - \widetilde{f}_t\right\|_{P,2} + \left|\omega\mathbb{E}_P\{(\widetilde{r}-r)\widetilde{\tau}_f\}\right|.$$

Proof of Proposition 4.1. With $\widetilde{f}_t = \mu_t$, every product term in equation 24 vanishes, regardless of $\widetilde{e}$, and hence $\mathfrak{D}(\omega;\mu,\widetilde{e},\widetilde{r}) = \omega\mathbb{E}_P\{(\widetilde{r}-r)\tau\}$. The proposed path is normalized because $\mathbb{E}_P\widetilde{r}_\varepsilon = \mathbb{E}_Q(1+\varepsilon g) = 1$ and is nonnegative under the stated bound on ε. Substitution gives

$$\mathfrak{D}(\omega;\mu,\widetilde{e},\widetilde{r}_\varepsilon) = \varepsilon\omega\mathbb{E}_P(rg\tau) = \varepsilon\omega\mathbb{E}_Q(g\tau),$$

which is equation 8. At $\omega = 0$ the first display is zero for every valid candidate ratio, not only along local paths.

On $L_2^0(Q) = \{g \in L_2(Q) : \mathbb{E}_Q g = 0\}$, the derivative at $\varepsilon = 0$ is the continuous linear functional

$$\dot{R}_g = \omega\mathbb{E}_Q(g\tau) = \omega\langle g, \tau - \mathbb{E}_Q\tau\rangle_{L_2(Q)}.$$

The assumption $\mathrm{Var}_Q(\tau) > 0$ places $Z = \tau - \mathbb{E}_Q\tau$ in $L_2^0(Q)$ with nonzero norm. Therefore the kernel of $g \mapsto \langle g, Z\rangle$ is a proper closed codimension-one hyperplane.

It remains to show that an admissible bounded direction exists, rather than only an abstract L_2 direction. For $K > 0$ set $h_K = \mathrm{clip}(Z, -K, K)$ and $g_K = h_K - \mathbb{E}_Q h_K$. Then $g_K \in L_2^0(Q) \cap L_\infty(Q)$ and $\|g_K\|_\infty \le 2K$. Because $\mathbb{E}_Q Z = 0$,

$$\mathbb{E}_Q(g_K\tau) = \mathbb{E}_Q(g_K Z) = \mathbb{E}_Q\{Z\,\mathrm{clip}(Z, -K, K)\}.$$

The integrand is nonnegative and is strictly positive wherever $Z \neq 0$. Thus the expectation is strictly positive for every $K > 0$ such that $Q(0 < |Z| \le K) > 0$ or $Q(|Z| > K) > 0$; such a finite K exists because $Q(Z \neq 0) > 0$. For $|\varepsilon| < (2K)^{-1}$, $1 + \varepsilon g_K > 0$ and $\mathbb{E}_P\{r(1+\varepsilon g_K)\} = 1 + \varepsilon\mathbb{E}_Q g_K = 1$. Hence $r(1+\varepsilon g_K)$ is an admissible normalized density-ratio path on a nontrivial open interval and has nonzero derivative. If τ is constant Q-almost surely, by contrast, $\mathbb{E}_Q(g\tau) = \tau\mathbb{E}_Q g = 0$ for every normalized direction, which proves the stated exceptional case. □

Lemma A.7 (A held-out empirical-process bound)**.** *Let $Z_1, \ldots, Z_N$ be iid from a law M and independent of a sigma-field $\mathcal{G}_N$. Let h_N be a $\mathcal{G}_N$-measurable random function with $\|h_N\|_{M,2} = o_p(1)$. Then*

$$(\mathbb{M}_N - M)h_N = o_p(N^{-1/2}), \qquad \mathbb{M}_N = N^{-1}\sum_{i=1}^{N}\delta_{Z_i}. \tag{25}$$

The conclusion remains true for each member of a fixed finite collection of held-out samples and random functions.

Proof. Conditional on $\mathcal{G}_N$, the function h_N is fixed and the Z_i's remain iid from M. Consequently,

$$\mathbb{E}\left[\left\{\sqrt{N}(\mathbb{M}_N - M)h_N\right\}^2 \middle| \mathcal{G}_N\right] = \mathrm{Var}_M\{h_N(Z)\} \le \|h_N\|_{M,2}^2 .$$

For $\epsilon, \delta > 0$, conditional Chebyshev's inequality gives

$$\Pr\left(\sqrt{N}\,|(\mathbb{M}_N - M)h_N| > \epsilon\right) \le \Pr\{\|h_N\|_{M,2} > \delta\} + \frac{\delta^2}{\epsilon^2}.$$

First let $N \to \infty$ and then $\delta \downarrow 0$. This proves equation 25. A union bound proves the assertion for a fixed finite collection. □

Proof of Theorem 4.5. For a candidate nuisance tuple $\eta = (f, e, r)$ and a scalar ω, define the two uncentered score components

$$\psi_P(w; \eta, \omega) = r(x)\left[\frac{t}{e(x)}\{y - f_1(x)\} - \frac{1-t}{1-e(x)}\{y - f_0(x)\} + \omega\tau_f(x)\right], \tag{26}$$

$$\psi_Q(x; \eta, \omega) = (1-\omega)\tau_f(x), \tag{27}$$

where $\tau_f = f_1 - f_0$. Thus the sum of the two population means is the population version of equation 2.

Folds and fitted scores. Fix $K < \infty$. Partition the source indices into deterministic folds $I_{P,1}, \ldots, I_{P,K}$ and the target indices into deterministic folds $I_{Q,1}, \ldots, I_{Q,K}$. Write $n_k = |I_{P,k}|$ and $m_k = |I_{Q,k}|$ and use balanced folds, so $\max_k |n_k - n/K| \le 1$ and $\max_k |m_k - m/K| \le 1$. For fold k, train $\widehat{\eta}_k = (\widehat{f}_k, \widehat{e}_k, \widehat{r}_k)$ without using any observation in $I_{P,k} \cup I_{Q,k}$. The common weight $\widehat{\omega}$ is computed from an independent calibration sample. If $\mathcal{G}_k$ is generated by that calibration sample and all observations outside $I_{P,k} \cup I_{Q,k}$, then $\widehat{\eta}_k$ and $\widehat{\omega}$ are $\mathcal{G}_k$-measurable, whereas the two held-out folds are conditionally independent iid samples from P and Q.

Let $\mathbb{P}_{n,k}$ and $\mathbb{Q}_{m,k}$ denote the empirical measures in the two held-out folds and set

$$\widehat{\psi}_{P,k} = \psi_P(\,\cdot\,; \widehat{\eta}_k, \widehat{\omega}), \qquad \widehat{\psi}_{Q,k} = \psi_Q(\,\cdot\,; \widehat{\eta}_k, \widehat{\omega}). \tag{28}$$

The usual observation-weighted cross-fitted estimator is

$$\widehat{\theta}_{\mathrm{cf}} = \sum_{k=1}^{K} \frac{n_k}{n} \mathbb{P}_{n,k}\widehat{\psi}_{P,k} + \sum_{k=1}^{K} \frac{m_k}{m} \mathbb{Q}_{m,k}\widehat{\psi}_{Q,k}. \tag{29}$$

To spell out the theorem's score-convergence hypothesis, let $f^\star$ be the fixed L_2 limit of the foldwise predictors and put

$$\psi_P^\star(w) = r(x)\left[\frac{t}{e(x)}\{y - f_1^\star(x)\} - \frac{1-t}{1-e(x)}\{y - f_0^\star(x)\} + \omega_\lambda^\star \tau_{f^\star}(x)\right],$$

$$\psi_Q^\star(x) = (1-\omega_\lambda^\star)\tau_{f^\star}(x).$$

The assumption is precisely

$$\max_{k \le K} \left\|\widehat{\psi}_{P,k} - \psi_P^\star\right\|_{P,2} = o_p(1), \qquad \max_{k \le K} \left\|\widehat{\psi}_{Q,k} - \psi_Q^\star\right\|_{Q,2} = o_p(1). \tag{30}$$

This formulation includes the error from replacing $\omega_\lambda^\star$ by the independently calibrated $\widehat{\omega}$. Theorem 3.1 gives

$$P\psi_P^\star + Q\psi_Q^\star = \theta_Q. \tag{31}$$

Step 1: held-out empirical processes. Apply Lemma A.7, conditionally on $\mathcal{G}_k$, to $\widehat{\psi}_{P,k} - \psi_P^\star$ and $\widehat{\psi}_{Q,k} - \psi_Q^\star$. Because $n_k \asymp n$, $m_k \asymp m \asymp n$, and K is fixed, equation 30 yields

$$\sum_{k=1}^{K} \frac{n_k}{n}(\mathbb{P}_{n,k} - P)(\widehat{\psi}_{P,k} - \psi_P^\star) = o_p(n^{-1/2}), \tag{32}$$

$$\sum_{k=1}^{K} \frac{m_k}{m}(\mathbb{Q}_{m,k} - Q)(\widehat{\psi}_{Q,k} - \psi_Q^\star) = o_p(n^{-1/2}). \tag{33}$$

No Donsker or entropy condition is used: independence of each fitted score from its held-out fold reduces the argument to conditional variance.

Step 2: exact foldwise drift and its aggregation. Lemma A.6 gives the identity

$$P\widehat{\psi}_{P,k} + Q\widehat{\psi}_{Q,k} - \theta_Q = \mathfrak{D}(\widehat{\omega}; \widehat{f}_k, \widehat{e}_k, \widehat{r}_k) =: \mathfrak{D}_k \tag{34}$$

for every fold. Balanced folds imply $n_k/n = 1/K + O(n^{-1})$ and $m_k/m = 1/K + O(m^{-1})$ uniformly in k. The L_2 convergence in equation 30 implies $\max_k |P\widehat{\psi}_{P,k}| + \max_k |Q\widehat{\psi}_{Q,k}| = O_p(1)$. It follows that the population part of equation 29 is

$$\begin{aligned}\sum_{k=1}^{K} \frac{n_k}{n} P\widehat{\psi}_{P,k} + \sum_{k=1}^{K} \frac{m_k}{m} Q\widehat{\psi}_{Q,k} - \theta_Q \\ = \frac{1}{K}\sum_{k=1}^{K} \mathfrak{D}_k + O_p(n^{-1}) + O_p(m^{-1}) = o_p(n^{-1/2}),\end{aligned} \tag{35}$$

where the final equality is the theorem's maximum-drift assumption and $m/n \to \lambda \in (0, \infty)$.

Step 3: two-sample asymptotic linearity. Insert and subtract the two fixed limiting scores in equation 29. Equations equation 32, equation 33, and equation 35, together with the identities

$$\sum_{k=1}^{K} \frac{n_k}{n} \mathbb{P}_{n,k}\psi_P^\star = \mathbb{P}_n\psi_P^\star, \qquad \sum_{k=1}^{K} \frac{m_k}{m} \mathbb{Q}_{m,k}\psi_Q^\star = \mathbb{Q}_m\psi_Q^\star,$$

give the expansion

$$\widehat{\theta}_{\mathrm{cf}} - \theta_Q = (\mathbb{P}_n - P)\psi_P^\star + (\mathbb{Q}_m - Q)\psi_Q^\star + o_p(n^{-1/2}). \tag{36}$$

Equivalently, with $\phi_P^\star = \psi_P^\star - P\psi_P^\star$ and $\phi_Q^\star = \psi_Q^\star - Q\psi_Q^\star$, the leading term is $n^{-1}\sum_i \phi_P^\star(W_i) + m^{-1}\sum_j \phi_Q^\star(\widetilde{X}_j)$.

Step 4: limit distribution. The two leading sums in equation 36 are independent. The stated $2+\eta$ moment bound implies the Lindeberg condition, and hence

$$\sqrt{n}(\widehat{\theta}_{\mathrm{cf}} - \theta_Q) \rightsquigarrow N(0, \sigma_\star^2), \qquad \sigma_\star^2 = \mathrm{Var}_P(\psi_P^\star) + \lambda^{-1}\mathrm{Var}_Q(\psi_Q^\star) > 0. \tag{37}$$

For example, the target contribution follows from

$$\sqrt{n}(\mathbb{Q}_m - Q)\psi_Q^\star = \sqrt{n/m}\,\sqrt{m}(\mathbb{Q}_m - Q)\psi_Q^\star \rightsquigarrow N\{0, \lambda^{-1}\mathrm{Var}_Q(\psi_Q^\star)\}.$$

Step 5: consistency of the plug-in variance. For $i \in I_{P,k}$ and $j \in I_{Q,k}$, set $\widehat{\psi}_{P,i} = \widehat{\psi}_{P,k}(W_i)$ and $\widehat{\psi}_{Q,j} = \widehat{\psi}_{Q,k}(\widetilde{X}_j)$, and let $\overline{\psi}_P = n^{-1}\sum_i \widehat{\psi}_{P,i}$ and $\overline{\psi}_Q = m^{-1}\sum_j \widehat{\psi}_{Q,j}$. The natural $\sqrt{n}$-scale variance estimator is

$$\widehat{\sigma}^2 = \frac{1}{n}\sum_{i=1}^{n}(\widehat{\psi}_{P,i} - \overline{\psi}_P)^2 + \frac{n}{m}\frac{1}{m}\sum_{j=1}^{m}(\widehat{\psi}_{Q,j} - \overline{\psi}_Q)^2. \tag{38}$$

Consistency follows from conditional Markov's inequality and equation 30, which imply

$$\begin{aligned}\frac{1}{n}\sum_{i=1}^{n}\{\widehat{\psi}_{P,i} - \psi_P^\star(W_i)\}^2 = o_p(1), \\ \frac{1}{m}\sum_{j=1}^{m}\{\widehat{\psi}_{Q,j} - \psi_Q^\star(\widetilde{X}_j)\}^2 = o_p(1).\end{aligned}$$

Indeed, this holds fold by fold because the conditional expectation of each held-out average is the corresponding squared L_2 norm; summing a fixed number of folds preserves the order. Cauchy–Schwarz and the uniform $2+\eta$ moment bound then show that replacing each fitted score by its fixed limit changes neither its first nor its second empirical moment by more than $o_p(1)$. The ordinary law of large numbers applied to the fixed limiting scores therefore gives

$$\frac{1}{n}\sum_i(\widehat{\psi}_{P,i} - \overline{\psi}_P)^2 \to_p \mathrm{Var}_P(\psi_P^\star), \qquad \frac{1}{m}\sum_j(\widehat{\psi}_{Q,j} - \overline{\psi}_Q)^2 \to_p \mathrm{Var}_Q(\psi_Q^\star).$$

Since $n/m \to \lambda^{-1}$, $\widehat{\sigma}^2 \to_p \sigma_\star^2$. Combining this fact with equation 37 proves

$$\frac{\sqrt{n}(\widehat{\theta}_{\mathrm{cf}} - \theta_Q)}{\widehat{\sigma}} \rightsquigarrow N(0,1),$$

and establishes the validity of the stated Wald interval. This is an asymptotic guarantee for learned weights; Theorem 4.4 gives the separate fixed-weight finite-sample result. □

Corollary A.1 (Primitive sufficient rates for the drift)**.** *Suppose, in addition to the remaining conditions of Theorem 4.5, that for some constants $\epsilon, R > 0$, uniformly over folds,*

$$\epsilon \le e, \widehat{e}_k \le 1-\epsilon, \qquad 0 \le r, \widehat{r}_k \le R, \qquad 0 \le \widehat{\omega} \le 1.$$

Define

$$a_f = \max_{k\le K} \sum_{t=0}^{1} \left\| \widehat{f}_{t,k} - \mu_t \right\|_{P,2}, \qquad a_e = \max_{k\le K} \|\widehat{e}_k - e\|_{P,2},$$
$$a_r = \max_{k\le K} \|\widehat{r}_k - r\|_{P,2}, \qquad \beta_{\mathrm{bal}} = \max_{k\le K} |\widehat{\omega} P\{(\widehat{r}_k - r)\widehat{\tau}_{f,k}\}|.$$

If

$$(a_e + a_r)a_f = o_p(n^{-1/2}), \qquad \beta_{\mathrm{bal}} = o_p(n^{-1/2}), \tag{39}$$

then the maximum-drift condition of Theorem 4.5 holds. A stronger, purely norm-based sufficient condition for the second relation is

$$a_r \max_{k\le K} \|\widehat{\tau}_{f,k}\|_{P,2} = o_p(n^{-1/2}).$$

At the exact external endpoint $\widehat{\omega} = 0$, the quantity β_{bal} is identically zero.

Proof. For $t = 0, 1$, write $e_1 = e$, $e_0 = 1 - e$ and use the analogous convention for $\widehat{e}_{t,k}$. Algebra gives $\widehat{r}_k e_t/\widehat{e}_{t,k} - r = (\widehat{r}_k - r)e_t/\widehat{e}_{t,k} + r(e_t - \widehat{e}_{t,k})/\widehat{e}_{t,k}$. The uniform bounds therefore imply $\max_{k,t} \|\widehat{r}_k e_t/\widehat{e}_{t,k} - r\|_{P,2} \le \epsilon^{-1}(a_r + Ra_e)$. Apply Cauchy–Schwarz term by term in the exact drift formula equation 24 to obtain, for a constant C depending only on (ϵ, R), $\max_{k\le K} \left|\mathfrak{D}(\widehat{\omega}; \widehat{f}_k, \widehat{e}_k, \widehat{r}_k)\right| \le C(a_e + a_r)a_f + \beta_{\mathrm{bal}}$. Condition equation 39 proves the first claim. Finally, $\beta_{\mathrm{bal}} \le a_r \max_{k\le K} \|\widehat{\tau}_{f,k}\|_{P,2}$ because $|\widehat{\omega}| \le 1$, proving the norm-based sufficient condition. The identity $\beta_{\mathrm{bal}} = 0$ at $\widehat{\omega} = 0$ is immediate. Thus cross-fitting controls sample reuse, whereas interior-weight transport imbalance still requires $\beta_{\mathrm{bal}} = o_p(n^{-1/2})$. □

A.8 Approximate transportability and robust calibration

The sensitivity model underlying Proposition 4.2 and Corollary 4.1 decomposes error into conditional-shift displacement, a doubly robust remainder, and direct ω-dependent transport imbalance.

A.8.1 Domain, integrability, and the sensitivity class

Let $(\mathcal{X}, \mathcal{A})$ be the covariate space. The source and target covariate laws are probability measures P_X and Q_X on this space, with $Q_X \ll P_X$ and Radon–Nikodym derivative $r = dQ_X/dP_X$. To distinguish source and target outcome mechanisms, write $\mu_t^P(x) = \mathbb{E}_P\{Y(t) \mid X = x\}$ and $\mu_t^Q(x) = \mathbb{E}_Q\{Y(t) \mid X = x\}$ for $t \in \{0, 1\}$. Define $\tau_P = \mu_1^P - \mu_0^P$, $\tau_Q = \mu_1^Q - \mu_0^Q$, $\delta_t = \mu_t^Q - \mu_t^P$, and $\delta_\tau = \delta_1 - \delta_0$. All functions are identified up to the relevant almost-sure equivalence. The target ATE, transported source ATE, and conditional-shift remainder are

$$\theta_Q = \mathbb{E}_Q \tau_Q(X), \qquad \theta_Q^{\mathrm{tr}} = \mathbb{E}_Q \tau_P(X), \qquad \Delta_{\mathrm{cond}} = \mathbb{E}_Q \delta_\tau(X), \tag{40}$$

so that $\theta_Q = \theta_Q^{\mathrm{tr}} + \Delta_{\mathrm{cond}}$.

For fixed P, Q_X, let $\mathcal{F}_1(P, Q_X)$ contain the measurable predictor pairs $f = (f_0, f_1)$ for which, for $t = 0, 1$,

$$\mathbb{E}_Q\{|\mu_t^P(X)| + |\mu_t^Q(X)| + |f_t(X)|\} < \infty, \qquad \mathbb{E}_P\left[r(X)\frac{\mathbf{1}\{T = t\}}{e_t(X)}|Y - f_t(X)|\right] < \infty. \tag{41}$$

These conditions justify conditional expectation and the change of measure in the order used below. They can be weakened by uniform-integrability arguments, but mere existence of signed, non-absolutely-integrable expressions is insufficient for the identities.

For a supplied radius $\rho_c \geq 0$, define the estimand-specific departure class

$$\mathcal{D}_c(\rho_c; Q_X) = \left\{ \delta \in L_1(Q_X) : \left| \int \delta(x)\, dQ_X(x) \right| \leq \rho_c \right\}. \tag{42}$$

The class bounds the target-average conditional-effect departure and imposes no pointwise smoothness. Consequently, large positive and negative departures may cancel. A pointwise or norm-based scientific model may replace equation 42; only its induced bound on $|\Delta_{\mathrm{cond}}|$ is used by the results below.

A standardized predictive-shift model is one way to supply the radius. Fix a nonnegative observable covariate discrepancy $d_X(P_X, Q_X)$, a positive conditional-residual scale s_{cond}, and a sensitivity multiplier $\Gamma \geq 0$, and set

$$\mathcal{D}_{\mathrm{pred}}(\Gamma) = \left\{ \delta \in L_1(Q_X) : \frac{|\mathbb{E}_Q \delta(X)|}{s_{\mathrm{cond}}} \leq \Gamma d_X(P_X, Q_X) \right\}. \tag{43}$$

Then equation 42 holds with $\rho_c = \Gamma d_X(P_X, Q_X) s_{\mathrm{cond}}$. For example, after fixing a feature map $\varphi : \mathcal{X} \to \mathbb{R}^d$ satisfying $\mathbb{E}_P\|\varphi(X)\|_2 + \mathbb{E}_Q\|\varphi(X)\|_2 < \infty$ and a positive-definite standardization matrix S, one may use

$$d_X(P_X, Q_X) = \left\| S^{-1/2} \{ \mathbb{E}_Q \varphi(X) - \mathbb{E}_P \varphi(X) \} \right\|_2.$$

There is a direct specialization to the standardized measures of Jin et al. (2025). In their common-propensity randomized setting, let

$$\phi(T, Y) = \frac{T}{\pi} Y - \frac{1 - T}{1 - \pi} Y, \qquad \phi_P(X) = \mathbb{E}_P(\phi \mid X) = \tau_P(X),$$

and assume the same randomized propensity under the target law. Then $\mathbb{E}_Q\{\phi - \phi_P(X)\} = \Delta_{\mathrm{cond}}$. Take

$$s_{\mathrm{cond}} = \mathrm{sd}_P\{\phi - \phi_P(X)\}, \qquad d_X^{\mathrm{stab}}(P_X, Q_X) = \left\{ \frac{1}{L} \sum_{\ell=1}^{L} \frac{\{\mathbb{E}_Q X_\ell - \mathbb{E}_P X_\ell\}^2}{\mathrm{Var}_P(X_\ell)} \right\}^{1/2}. \tag{44}$$

Provided the displayed variances are positive and finite, the assumption

$$\frac{|\Delta_{\mathrm{cond}}|}{s_{\mathrm{cond}}} \leq \Gamma d_X^{\mathrm{stab}}(P_X, Q_X) \tag{45}$$

maps their dimensionless predictive relation to the ATE-scale radius $\rho_c = \Gamma d_X^{\mathrm{stab}} s_{\mathrm{cond}}$. The random-shift model of Jin et al. (2025) motivates this relation but does not imply it for arbitrary source–target pairs. We therefore treat equation 45 as a sensitivity assumption, with Γ supplied scientifically or calibrated on separate comparable shifts; Lemma A.9 propagates its sampling uncertainty. Absolute continuity alone places no such bound on conditional shift.

A.8.2 Calibration–transport separation

Proof of Proposition 4.2. Fix $f \in \mathcal{F}_1(P, Q_X)$ and condition on any external training sigma-field that determines f. The estimator contains outcomes only from the source population. By consistency, source exchangeability, and treatment overlap, for each $t \in \{0, 1\}$,

$$\mathbb{E}_P\left[r(X) \frac{\mathbf{1}\{T = t\}}{e_t(X)} \{Y - f_t(X)\} \,\middle|\, X \right] = r(X)\{\mu_t^P(X) - f_t(X)\}. \tag{46}$$

Absolute integrability in equation 41 permits integration of equation 46. Since $Q_X \ll P_X$,

$$\mathbb{E}_P Z_f = \mathbb{E}_Q(\tau_P - \tau_f), \qquad \mathbb{E}_P(r\tau_f) = \mathbb{E}_Q \tau_f. \tag{47}$$

Substitution in equation 2 gives, for every fixed $\omega \in \mathbb{R}$,

$$\begin{aligned}\mathbb{E}(\widehat{\theta}_\omega \mid f) &= \omega \mathbb{E}_Q \tau_f + (1-\omega)\mathbb{E}_Q \tau_f + \mathbb{E}_Q(\tau_P - \tau_f) \\ &= \mathbb{E}_Q \tau_P = \theta_Q^{\mathrm{tr}}.\end{aligned}$$

Combining this equality with equation 40 yields $\mathbb{E}(\widehat{\theta}_\omega \mid f) - \theta_Q = -\Delta_{\mathrm{cond}}$. Thus a change in ω cannot alter the structural displacement. Notice also that changing the unobserved target conditional-outcome law while holding P and Q_X fixed changes θ_Q but not the sampling law of $\widehat{\theta}_\omega$. This observation will make the conditional-shift minimax statement below exact rather than merely an upper bound. □

Proof of Corollary 4.1. Write $C_\omega = [L_\omega, U_\omega]$. On the event $\theta_Q^{\mathrm{tr}} \in C_\omega$, every $\delta_\tau \in \mathcal{D}_c(\rho_c; Q_X)$ satisfies

$$\theta_Q = \theta_Q^{\mathrm{tr}} + \Delta_{\mathrm{cond}} \in [L_\omega - \rho_c, U_\omega + \rho_c] = C_\omega \oplus [-\rho_c, \rho_c].$$

Because (P, Q_X) is fixed while only the unobserved target conditional-outcome law varies over $\mathcal{D}_c(\rho_c; Q_X)$, this same event has probability at least $1-\alpha$ for every member of that departure class. This proves uniformity over conditional shifts at the fixed observed-data law. Uniformity over a larger class of (P, Q_X) additionally requires the base interval to have uniform coverage over that larger class.

For the ratio-error statement, let $\psi_\omega = \theta_Q^{\mathrm{tr}} + b_r(\omega)$. On the assumed coverage event, $|\theta_Q - \psi_\omega| = |\Delta_{\mathrm{cond}} - b_r(\omega)| \le \rho_c + \omega\rho_r$. Therefore $\theta_Q \in C_\omega \oplus [-\rho_c - \omega\rho_r, \rho_c + \omega\rho_r]$. To minimize the bias–variance criterion, expand equation 4 and the squared radius. Up to a term independent of ω,

$$\begin{aligned}\mathcal{R}_{\mathrm{rob}}(\omega) = 2\omega &\left(\frac{C_{UR}}{n} - \frac{B_{Q,f}}{m} + \rho_c\rho_r\right) \\ &+ \omega^2\left(\frac{B_P}{n} + \frac{B_{Q,f}}{m} + \rho_r^2\right) + \mathrm{const.}\end{aligned} \tag{48}$$

If $D_{\mathrm{rob}} > 0$, completing the square and projecting the unconstrained minimizer onto $[0,1]$ gives equation 9. If $D_{\mathrm{rob}} = 0$, nonnegativity forces $B_P = B_{Q,f} = \rho_r = 0$. Cauchy–Schwarz gives $C_{UR} = 0$, and every coefficient depending on ω in equation 48 vanishes. Hence the criterion is constant on $[0,1]$. □

A.8.3 Exact joint drift with estimated nuisances

The joint drift decomposition identifies when the ratio contribution reduces to ωb_1.

Let $\widetilde{\eta} = (\widetilde{f}, \widetilde{e}, \widetilde{r})$ be fixed candidate nuisance functions, with $\widetilde{e}_1 = \widetilde{e}$, $\widetilde{e}_0 = 1 - \widetilde{e}$, and $\widetilde{\tau}_f = \widetilde{f}_1 - \widetilde{f}_0$. Assume the candidate propensities are strictly between zero and one, and assume absolute integrability of every term below. Let $\widehat{\theta}_\omega(\widetilde{\eta})$ denote equation 2 after replacing (f, e, r) by $(\widetilde{f}, \widetilde{e}, \widetilde{r})$. With $s_1 = 1, s_0 = -1$, define

$$b_0(\widetilde{\eta}) = \sum_{t=0}^{1} s_t \mathbb{E}_P\left[\left\{\frac{\widetilde{r}(X)e_t(X)}{\widetilde{e}_t(X)} - r(X)\right\}\{\mu_t^P(X) - \widetilde{f}_t(X)\}\right], \tag{49}$$

$$b_1(\widetilde{\eta}) = \mathbb{E}_P[\{\widetilde{r}(X) - r(X)\}\widetilde{\tau}_f(X)]\,. \tag{50}$$

Lemma A.8 (Exact structural–nuisance drift decomposition)**.** *Under source identification and the conditions above,*

$$\mathbb{E}\{\widehat{\theta}_\omega(\widetilde{\eta})\} - \theta_Q = b_0(\widetilde{\eta}) + \omega b_1(\widetilde{\eta}) - \Delta_{\mathrm{cond}} \tag{51}$$

for every $\omega \in \mathbb{R}$.

Proof. Conditional expectation under the source law gives

$$\mathbb{E}_P\left[\frac{\mathbf{1}\{T=t\}}{\widetilde{e}_t(X)}\{Y - \widetilde{f}_t(X)\}\,\middle|\, X\right] = \frac{e_t(X)}{\widetilde{e}_t(X)}\{\mu_t^P(X) - \widetilde{f}_t(X)\}.$$

Because $\mathbb{E}_Q\widetilde{\tau}_f = \mathbb{E}_P(r\widetilde{\tau}_f)$, the population expectation of the candidate estimator minus $\theta_Q^{\mathrm{tr}} = \mathbb{E}_P(r\tau_P)$ equals

$$(1-\omega)\mathbb{E}_P(r\widetilde{\tau}_f) + \omega\mathbb{E}_P(\widetilde{r}\widetilde{\tau}_f) + \sum_{t=0}^{1} s_t \mathbb{E}_P\left[\frac{\widetilde{r}e_t}{\widetilde{e}_t}(\mu_t^P - \widetilde{f}_t)\right] - \mathbb{E}_P(r\tau_P).$$

Add and subtract $\sum_t s_t\mathbb{E}_P\{r(\mu_t^P - \widetilde{f}_t)\}$ and collect terms. The result is $b_0(\widetilde{\eta}) + \omega b_1(\widetilde{\eta})$. Finally subtract $\Delta_{\mathrm{cond}} = \theta_Q - \theta_Q^{\mathrm{tr}}$. □

The decomposition yields three cases. First, if $\widetilde{f}_t = \mu_t^P$, then $b_0 = 0$ for every candidate propensity and

$$\mathbb{E}\{\widehat{\theta}_\omega(\widetilde{\eta})\} - \theta_Q = \omega\mathbb{E}_P\{(\widetilde{r} - r)\tau_P\} - \Delta_{\mathrm{cond}}.$$

This is the regime in which the ratio sensitivity radius enters only as $\omega\rho_r$. Second, if $\widetilde{r} = r$, then $b_1 = 0$ and

$$b_0 = \sum_{t=0}^{1} s_t \mathbb{E}_P\left[r\frac{e_t - \widetilde{e}_t}{\widetilde{e}_t}(\mu_t^P - \widetilde{f}_t)\right],$$

which is the usual product-form propensity–outcome remainder. Third, if $\widetilde{e} = e$ but both $\widetilde{r}$ and $\widetilde{f}$ may be inaccurate, then

$$b_0 = \mathbb{E}_P\{(\widetilde{r} - r)(\tau_P - \widetilde{\tau}_f)\}, \qquad b_1 = \mathbb{E}_P\{(\widetilde{r} - r)\widetilde{\tau}_f\}. \tag{52}$$

Cross-fitting separates empirical-process control from the population discrepancies b_0 and b_1, which remain in the centering error.

For $\rho_c, \rho_0, \rho_r \geq 0$, define the joint sensitivity class

$$\mathfrak{S}(\rho_c, \rho_0, \rho_r) = \left\{(P, Q, \widetilde{\eta}) : \begin{array}{l}\delta_r \in \mathcal{D}_c(\rho_c; Q_X), \\ |b_0(\widetilde{\eta})| \leq \rho_0, \end{array} \begin{array}{l} \\ |b_1(\widetilde{\eta})| \leq \rho_r \end{array}\right\}, \tag{53}$$

with source identification, support, and the integrability conditions understood. Lemma A.8 immediately gives the uniform bound

$$\sup_{\mathfrak{S}(\rho_c,\rho_0,\rho_r)} \left|\mathbb{E}\{\widehat{\theta}_\omega(\widetilde{\eta})\} - \theta_Q\right| \leq \rho_c + \rho_0 + |\omega|\rho_r. \tag{54}$$

For $\omega \in [0, 1]$, the right side is $\rho_c + \rho_0 + \omega\rho_r$. The main-text corollary is the $\rho_0 = 0$ case; equivalently, b_0 may already be included in the center covered by the base interval.

Proposition A.1 (Uniform sensitivity coverage)**.** *Suppose C_ω satisfies*

$$\inf_{(P,Q,\widetilde{\eta})\in\mathfrak{S}(\rho_c,\rho_0,\rho_r)} \Pr\{\theta_Q^{\mathrm{tr}} + b_0(\widetilde{\eta}) + \omega b_1(\widetilde{\eta}) \in C_\omega\} \geq 1 - \alpha.$$

Then

$$C_\omega^{\mathrm{sens}} = C_\omega \oplus [-\{\rho_c + \rho_0 + \omega\rho_r\}, \rho_c + \rho_0 + \omega\rho_r] \tag{55}$$

covers θ_Q with probability at least $1 - \alpha$ uniformly over $\mathfrak{S}(\rho_c, \rho_0, \rho_r)$. The same conclusion holds for an independently calibrated random $\widehat{\omega} \in [0, 1]$ if the displayed base coverage holds conditionally on the calibration sigma-field and the bounds in equation 53 hold simultaneously for all $\omega \in [0, 1]$.

Proof. On the base coverage event, let $\psi_\omega = \theta_Q^{\mathrm{tr}} + b_0 + \omega b_1 \in C_\omega$. Equation equation 40 and the triangle inequality yield

$$|\theta_Q - \psi_\omega| = |\Delta_{\mathrm{cond}} - b_0 - \omega b_1| \leq \rho_c + \rho_0 + \omega\rho_r.$$

This is exactly the Minkowski enlargement in equation 55. For random $\widehat{\omega}$, apply the same deterministic inclusion conditional on the calibration sigma-field, then integrate the conditional probability. □

A.8.4 EXACT MINIMAX CALIBRATION AND BOUNDARY CASES

The squared-radius criterion in the main text has an exact minimax interpretation when the sensitivity analysis varies the mean displacement while leaving the nominal noise law fixed. This is automatic for pure conditional shift because target outcomes are unobserved: with P and Q_X fixed, varying μ_t^Q leaves the distribution of the estimator unchanged. For a ratio perturbation, the same calculation is exact in a local-bias model or whenever $V(\omega)$ is the common conditional variance; if the perturbation also changes variance, replacing V by a uniform variance envelope gives a valid upper bound rather than an equality.

Proposition A.2 (Sharp rectangular minimax risk)**.** *Conditional on the external training sigma-field $\mathcal{T}$, suppose the sensitivity experiment is the product of a fixed centered noise law $\{\xi_\omega : \omega \in [0,1]\}$ and the scalar rectangle $[-\rho_c, \rho_c] \times [-\rho_0, \rho_0] \times [-\rho_r, \rho_r]$. For every point (Δ, b_0, b_1) in that rectangle, the conditional joint law of the entire process $\{\xi_\omega : \omega \in [0,1]\}$ given $\mathcal{T}$ is the same, and suppose*

$$\widehat{\theta}_\omega - \theta_Q = \xi_\omega + b_0 + \omega b_1 - \Delta, \qquad \mathbb{E}(\xi_\omega \mid \mathcal{T}) = 0, \quad \mathrm{Var}(\xi_\omega \mid \mathcal{T}) = V(\omega).$$

The three scalar departures range independently. Put $a = \rho_c + \rho_0$, $D_{\mathrm{rob}} = B_P/n + B_{Q,f}/m + \rho_r^2$, and $N_{\mathrm{rob}} = B_{Q,f}/m - C_{UR}/n - a\rho_r$. Then, for $\omega \in [0,1]$,

$$\sup_{|\Delta| \le \rho_c,\, |b_0| \le \rho_0,\, |b_1| \le \rho_r} \mathbb{E}\{(\widehat{\theta}_\omega - \theta_Q)^2 \mid \mathcal{T}\} = V(\omega) + (a + \omega\rho_r)^2. \tag{56}$$

If $D_{\mathrm{rob}} > 0$, every minimizer over $[0,1]$ equals

$$\omega^\star_{\mathrm{rob}} = \Pi_{[0,1]}\left(\frac{N_{\mathrm{rob}}}{D_{\mathrm{rob}}}\right) = \begin{cases} 0, & N_{\mathrm{rob}} \le 0, \\ N_{\mathrm{rob}}/D_{\mathrm{rob}}, & 0 < N_{\mathrm{rob}} < D_{\mathrm{rob}}, \\ 1, & N_{\mathrm{rob}} \ge D_{\mathrm{rob}}. \end{cases} \tag{57}$$

If $D_{\mathrm{rob}} = 0$, the risk is constant and every $\omega \in [0,1]$ is minimax. When $\rho_0 = 0$, equation 57 reduces exactly to equation 9.

Proof. All expectations in this proof are conditional on $\mathcal{T}$. For fixed (Δ, b_0, b_1), the mean-zero property of ξ_ω gives $\mathbb{E}(\widehat{\theta}_\omega - \theta_Q)^2 = V(\omega) + (b_0 + \omega b_1 - \Delta)^2$. The triangle inequality bounds the absolute bias by $a + \omega\rho_r$. Equality is attained by choosing, for example, $b_0 = \rho_0$, $b_1 = \rho_r$, and $\Delta = -\rho_c$. This proves equation 56; it also proves that the squared-radius term cannot be reduced over the rectangular class.

Expanding the right side of equation 56 yields $D_{\mathrm{rob}}\omega^2 - 2N_{\mathrm{rob}}\omega + \mathrm{const}$, up to an ω-independent constant. When $D_{\mathrm{rob}} > 0$, completion of the square gives the unconstrained minimizer $N_{\mathrm{rob}}/D_{\mathrm{rob}}$, and Euclidean projection gives the three branches in equation 57. If $D_{\mathrm{rob}} = 0$, then $B_P = B_{Q,f} = \rho_r = 0$ and Cauchy–Schwarz implies $C_{UR} = 0$, so both the variance and squared-radius terms are constant in ω. □

Under the product scalar model of Proposition A.2, the same aligned-sign construction proves set-theoretic sharpness of the interval enlargement. If the base center ψ_ω were observed without noise, the identified set under the rectangular class would be exactly

$$[\psi_\omega - (\rho_c + \rho_0 + \omega\rho_r), \psi_\omega + (\rho_c + \rho_0 + \omega\rho_r)].$$

Both endpoints are attained by aligned signs. Hence no strictly smaller symmetric enlargement can cover uniformly over this product class. This sharpness need not hold for a narrower nuisance model that links b_0 and b_1.

At the boundaries, the optimizer behaves as follows.

1. If $\rho_r = 0$, then the robust risk is $V(\omega) + a^2$ and the original oracle $\omega^\star$ remains minimax. Pure structural shift widens uncertainty but does not change calibration.
2. If $a = \rho_r = 0$, both the risk and optimizer reduce to Theorem 4.1.
3. If $a > 0$, increasing ρ_r moves the projected optimizer weakly toward zero. On the nondegenerate region $d + \rho^2 > 0$, let $z(\rho) = (b - a\rho)/(d + \rho^2)$, where $b = B_{Q,f}/m - C_{UR}/n$ and $d = B_P/n + B_{Q,f}/m$. If $b \le 0$, the projection is already zero. If $b > 0$, then on the region

$b - a\rho > 0$, $z'(\rho) = \{a\rho^2 - 2b\rho - ad\}/(d+\rho^2)^2 \le 0$, because $a\rho < b$ implies $a\rho^2 - 2b\rho - ad < b\rho - 2b\rho - ad \le 0$. Once $b - a\rho \le 0$, the projection is zero. For $a = 0$, the same conclusion follows from $z'(\rho) = -2b\rho/(d+\rho^2)^2$. At the sole degenerate point $d = \rho = 0$, Proposition A.2 already declares every ω minimax.

4. If the sensitivity radius is infinite, the target ATE is not bounded by this model and no finite symmetric enlargement is available. This is an identification failure, not a calibration failure.

A.8.5 Bounds for ratio error and random radii

The scalar ρ_r can be related to a norm discrepancy. Suppose $\widetilde{r} \ge 0$, $\mathbb{E}_P\widetilde{r} = \mathbb{E}_P r = 1$, and put $h = \widetilde{r} - r$. For conjugate exponents $p, q \in [1, \infty]$, Hölder's inequality gives, for every constant c,

$$|b_1| = |\mathbb{E}_P\{h(\widetilde{\tau}_f - c)\}| \le \|h\|_{L_p(P)}\|\widetilde{\tau}_f - c\|_{L_q(P)}. \tag{58}$$

Taking the infimum over c can be substantially tighter than using $\|\widetilde{\tau}_f\|_{L_q(P)}$. If $\widetilde{e} = e$, the same argument and equation 52 give

$$|b_0| \le \|h\|_{L_p(P)} \inf_{c\in\mathbb{R}} \|\tau_P - \widetilde{\tau}_f - c\|_{L_q(P)}.$$

For the normalized multiplicative perturbation $\widetilde{r} = r(1 + \varepsilon g)$ with $\mathbb{E}_Q g = 0$,

$$b_1 = \varepsilon\mathbb{E}_Q(g\widetilde{\tau}_f) = \varepsilon\mathbb{E}_Q\{g(\widetilde{\tau}_f - c)\}, \tag{59}$$

which connects the transport condition directly to balance of the predictor contrast.

When radii are estimated, nominal coverage requires accounting for their own failure probabilities. The following elementary statement records the exact bookkeeping.

Lemma A.9 (Coverage with random upper radii)**.** *Let $\widehat{\omega} \in [0, 1]$. Suppose a base interval $C_{\widehat{\omega}}$ covers $\theta_Q^{\mathrm{tr}} + b_0 + \widehat{\omega}b_1$ with probability at least $1 - \alpha_0$. Suppose random nonnegative radii satisfy*

$$\Pr(|\Delta_{\mathrm{cond}}| \le \widehat{\rho}_c) \ge 1 - \alpha_c, \quad \Pr(|b_0| \le \widehat{\rho}_0) \ge 1 - \alpha_0', \quad \Pr(|b_1| \le \widehat{\rho}_r) \ge 1 - \alpha_r.$$

All four probability statements are unconditional and include the randomness used to learn the nuisances and radii. Then

$$C_{\widehat{\omega}} \oplus [-\widehat{R}, \widehat{R}], \qquad \widehat{R} = \widehat{\rho}_c + \widehat{\rho}_0 + \widehat{\omega}\widehat{\rho}_r,$$

covers θ_Q with probability at least $1 - (\alpha_0 + \alpha_c + \alpha_0' + \alpha_r)$, without an independence assumption.

Proof. Intersect the four stated events. On their intersection, the deterministic triangle-inequality argument in Proposition A.1 applies with the realized radii. The union bound controls the complement of the intersection. □

Thus a plug-in value of d_X or s_{cond} becomes a coverage statement only after it is replaced by a simultaneous upper confidence bound or treated as a scientifically fixed sensitivity input. The deterministic-radius corollary is the special case in which the three radius failure probabilities are zero.

A.8.6 Complete two-stratum calculation under both departures

We now extend the running example in Appendix A.11. Let $P_X = (4/5, 1/5)$, $Q_X = (1/5, 4/5)$, $r = (1/4, 4)$, and $e = 1/2$, and let the source effect and imperfect predictor contrast be $\tau_P = (1, 3)$ and $\tau_f = (1/2, 7/2)$. Parameterize target conditional shift by $\delta_\tau = (d_0, d_1)$ and $\tau_Q = (1+d_0, 3+d_1)$. Then

$$\Delta_{\mathrm{cond}} = \frac{1}{5}d_0 + \frac{4}{5}d_1, qquad\theta_Q = \frac{13}{5} + \frac{1}{5}d_0 + \frac{4}{5}d_1. \tag{60}$$

With the correct ratio, $\mathbb{E}_Q\tau_f = (1/5)(1/2) + (4/5)(7/2) = 29/10$ and $\mathbb{E}_P Z_f = 13/5 - 29/10 = -3/10$. Consequently every ω has expectation $13/5$ and bias relative to the true target ATE equal to $-(d_0/5 + 4d_1/5) = -\Delta_{\mathrm{cond}}$. In particular, if $|d_0/5 + 4d_1/5| \le \rho_c$, the worst-case MSE is at most $V(\omega) + \rho_c^2$. It equals this bound, and hence leaves the variance oracle $\omega^\star$ unchanged, whenever the sensitivity class contains target conditional-outcome laws attaining both endpoint shifts $\Delta_{\mathrm{cond}} = \pm\rho_c$ while the observed-data law remains fixed.

Next perturb the ratio using $g = (4, -1)$, $\varepsilon = 1/10$, and $\widetilde{r} = r(1 + \varepsilon g) = (7/20, 18/5)$. Since $\mathbb{E}_Q g = (1/5)4 + (4/5)(-1) = 0$, this is normalized: $\mathbb{E}_P \widetilde{r} = 1$. With the correct propensity but the imperfect predictor, equation 52 gives

$$\begin{aligned} b_0 &= \varepsilon \mathbb{E}_Q\{g(\tau_P - \tau_f)\} \\ &= \frac{1}{10}\left\{\frac{1}{5}(4)\left(1 - \frac{1}{2}\right) + \frac{4}{5}(-1)\left(3 - \frac{7}{2}\right)\right\} = \frac{2}{25} = 0.08, \\ b_1 &= \varepsilon \mathbb{E}_Q(g\tau_f) \\ &= \frac{1}{10}\left\{\frac{1}{5}(4)\frac{1}{2} + \frac{4}{5}(-1)\frac{7}{2}\right\} = -\frac{6}{25} = -0.24. \end{aligned}$$

Therefore the exact expectation and target bias are

$$\mathbb{E}\{\widehat{\theta}_\omega(\widetilde{\eta})\} = \frac{13}{5} + 0.08 - 0.24\omega, \qquad \mathrm{Bias}_Q(\omega) = 0.08 - 0.24\omega - \Delta_{\mathrm{cond}}. \tag{61}$$

Thus b_0 contributes whenever both the ratio and predictor are inaccurate.

For the estimator actually computed with $\widetilde{r}$ and the outcome model specified in Appendix A.11, direct enumeration gives $\widetilde{A} = 60283/5000 = 12.0566$ and $\widetilde{C}_{UR} = -39263/10000 = -3.9263$. It also gives $\widetilde{B}_P = 247009/10000 = 24.7009$ and $B_{Q,f} = 36/25 = 1.44$. Consequently, at $(n, m) = (100, 1000)$, the actual conditional variance is

$$\widetilde{V}(\omega) = 0.122006 - 0.081406\omega + 0.248449\omega^2. \tag{62}$$

To evaluate the estimator beyond the one-dimensional relation between b_0 and b_1, consider the conservative rectangular scalar-bias class $|b_0| \le 0.08$, $|b_1| \le 0.24$. This class contains the displayed multiplicative perturbation and permits the two nuisance remainders to align. Exact sharpness holds for this enlarged scalar class, rather than for the original one-direction ratio path. Embedding the fixed noise law above in this product model and adding conditional-shift radius ρ_c, Proposition A.2 gives

$$\omega^\star_{\mathrm{rob}} = \Pi_{[0,1]}\left\{\frac{0.021503 - 0.24\rho_c}{0.306049}\right\}. \tag{63}$$

At $\rho_c = 0$, the robust weight is 0.070260. Once $\rho_c \ge 0.0895958$, the numerator in equation 63 is nonpositive and the robust rule selects the external endpoint $\omega = 0$, eliminating the direct b_1 term while leaving the bounded b_0 remainder.

Finally, if the arm-level predictor is oracle, $\widetilde{f}_t = \mu_t^P$, then $b_0 = 0$ and $b_1 = \varepsilon \mathbb{E}_Q(g\tau_P) = (1/10)\{(1/5)(4)(1) + (4/5)(-1)(3)\} = -4/25 = -0.16$. Thus the combined bias is exactly $-0.16\omega - \Delta_{\mathrm{cond}}$, recovering the main-text $\rho_c + \omega\rho_r$ geometry with $\rho_r = 0.16$. The two calculations are complementary: oracle outcome regression isolates the first-order ω-dependent transport barrier, whereas the imperfect-predictor version shows the additional product remainder that a joint sensitivity analysis must bound or absorb into the base interval.

A.9 Efficiency lower bounds under transport

Proof of Theorem 4.6. We calculate the canonical gradient in the dominated nonparametric two-sample experiment. The calculation also identifies exactly which part of the bound comes from labeled outcomes and which part comes from target covariates.

Step 1: dominated factorization and tangent spaces. Fix a common sigma-finite measure ν_X for covariates and a sigma-finite measure ν_Y for outcomes. At the distribution under consideration, write the source observed-data density as

$$p(w) = p_X(x)e(x)^t\{1 - e(x)\}^{1-t}p_t(y \mid x), \qquad w = (x, t, y), \tag{64}$$

and the target-covariate density as $q(x)$. The propensity e is fixed and known. Consider differentiable quadratic-mean submodels $(p_\varepsilon, q_\varepsilon)$ that preserve this factorization and the support condition locally. Their scores can be written

$$s_P(W) = s_X(X) + s_Y(Y, T, X), \qquad s_Q(X), \tag{65}$$

where

$$Ps_X = 0, \qquad \mathbb{E}_P\{s_Y(Y,T,X) \mid X,T\} = 0, \qquad Qs_Q = 0.$$

Conversely, bounded scores satisfying these restrictions generate regular submodels, and their L_2 closure is the tangent space used below. The source covariate and source outcome tangent spaces are orthogonal because $\mathbb{E}_P(s_Y \mid X) = 0$.

For $t \in \{0,1\}$, let $s_{Y,t}(y,x)$ denote $s_Y(y,t,x)$. Differentiating the conditional mean along the submodel gives

$$\dot{\mu}_t(x) = \mathbb{E}_P[\{Y - \mu_t(x)\}s_{Y,t}(Y,x) \mid X = x, T = t]. \tag{66}$$

The parameter $\theta(P,Q) = \int\{\mu_{1,P}(x) - \mu_{0,P}(x)\}q(x)d\nu_X(x)$ does not depend on the source marginal p_X when the conditional outcome laws and Q are held fixed. Its pathwise derivative is therefore

$$\dot{\theta}(s_P, s_Q) = \mathbb{E}_Q\{\dot{\mu}_1(X) - \dot{\mu}_0(X)\} + \mathbb{E}_Q[\{\tau(X) - \theta_Q\}s_Q(X)]. \tag{67}$$

Step 2: representation of the derivative. Define

$$\phi_P(W) = r(X)\left[\frac{T}{e(X)}\{Y - \mu_1(X)\} - \frac{1-T}{1-e(X)}\{Y - \mu_0(X)\}\right], \tag{68}$$

$$\phi_Q(X) = \tau(X) - \theta_Q. \tag{69}$$

Both functions have mean zero. Moreover, $\mathbb{E}_P(\phi_P \mid X,T) = 0$, so ϕ_P is orthogonal to the source-covariate tangent space and belongs to the $L_2(P)$ closure of the source-outcome tangent space. Conditioning on (X,T), using equation 66, and then changing measure from P_X to Q_X gives

$$\begin{aligned}\mathbb{E}_P\{\phi_P(W)s_P(W)\} &= \mathbb{E}_P\{\phi_P(W)s_Y(W)\} \\ &= \mathbb{E}_Q\{\dot{\mu}_1(X) - \dot{\mu}_0(X)\}.\end{aligned}$$

The second term in equation 67 is $\mathbb{E}_Q\{\phi_Q(X)s_Q(X)\}$. Hence

$$\dot{\theta}(s_P, s_Q) = \mathbb{E}_P(\phi_P s_P) + \mathbb{E}_Q(\phi_Q s_Q). \tag{70}$$

Step 3: the joint canonical gradient. Under $m/n \to \lambda \in (0,\infty)$ and $\sqrt{n}$ scaling, the LAN inner product for the product experiment $P^n \otimes Q^m$ is

$$\langle (g_P, g_Q), (h_P, h_Q)\rangle_\lambda = \mathbb{E}_P(g_P h_P) + \lambda\mathbb{E}_Q(g_Q h_Q). \tag{71}$$

Comparing equation 70 with equation 71 shows that the Riesz representer is

$$g_\lambda = (\phi_P, \lambda^{-1}\phi_Q). \tag{72}$$

It lies in the closure of the joint tangent space: the first coordinate lies in the source-outcome tangent space and the second is a centered element of $L_2(Q)$. It is therefore the canonical gradient. Indeed, if h is any other gradient representing the derivative on the joint tangent space, then $h - g_\lambda$ is orthogonal to that space. Since g_λ belongs to its closure, Pythagoras gives $\|h\|_\lambda^2 = \|g_\lambda\|_\lambda^2 + \|h - g_\lambda\|_\lambda^2 \geq \|g_\lambda\|_\lambda^2$. This proves the minimality of equation 72.

The squared norm of the canonical gradient is

$$\begin{aligned}\|g_\lambda\|_\lambda^2 &= \mathbb{E}_P(\phi_P^2) + \lambda^{-1}\mathbb{E}_Q(\phi_Q^2) \\ &= \mathbb{E}_P\left[r(X)^2\left\{\frac{\sigma_1^2(X)}{e(X)} + \frac{\sigma_0^2(X)}{1-e(X)}\right\}\right] + \lambda^{-1}\mathrm{Var}_Q\{\tau(X)\}.\end{aligned} \tag{73}$$

For the second equality, condition on X in the source term. The treated and control residual contributions are disjoint because $T(1-T) = 0$, while $\mathbb{E}_P[T\{Y - \mu_1(X)\}^2/e(X)^2 \mid X] = \sigma_1^2(X)/e(X)$, and analogously for control.

Step 4: lower bound and attainment. The convolution theorem for regular estimators in a differentiable experiment states that the weak limit of any regular $\sqrt{n}$-scaled estimator is the convolution of $N\{0, \|g_\lambda\|_\lambda^2\}$ with an independent residual law. Consequently every regular estimator with a finite limiting variance has limiting variance at least equation 73, which is exactly equation 10.

Now set $f_t = \mu_t$ and $\omega = 0$ in equation 2. The infeasible oracle then satisfies the exact identity

$$\widehat{\theta}_0^{\mathrm{or}} - \theta_Q = \mathbb{P}_n \phi_P + \mathbb{Q}_m \phi_Q. \tag{74}$$

The two sums are independent, so the variance of $\sqrt{n}(\widehat{\theta}_0^{\mathrm{or}} - \theta_Q)$ converges to equation 73; the oracle attains the bound. Finally, Theorem 4.5 shows that a feasible cross-fitted one-step estimator has the same expansion provided its estimated score converges to the efficient oracle score, specifically,

$$\max_{k \le K} \|\widehat{\psi}_{P,k} - \phi_P\|_{P,2} + \max_{k \le K} \|\widehat{\psi}_{Q,k} - \tau\|_{Q,2} = o_p(1),$$

and its exact population drift is $o_p(n^{-1/2})$. Equivalently, the limiting score in Theorem 4.5 must have $f^\star = \mu$ and $\omega_\lambda^\star = 0$; convergence to an arbitrary limiting score gives valid asymptotic inference but need not attain the information bound. Under the displayed efficient-score condition the feasible estimator also attains the bound. □

Proposition A.3 (Efficiency bound in the restricted common-marginal model)**.** *Suppose the statistical model itself imposes $P_X = Q_X =: M$, rather than merely evaluating the unrestricted external-target model at a distribution where the two marginals happen to coincide. Let*

$$A_0 = \mathbb{E}_M\left\{ \frac{\sigma_1^2(X)}{e(X)} + \frac{\sigma_0^2(X)}{1 - e(X)} \right\}, \qquad B_0 = \mathrm{Var}_M\{\tau(X)\}.$$

Assume $A_0 + B_0 < \infty$. If $m/n \to \lambda \in (0, \infty)$, the semiparametric efficiency bound under $\sqrt{n}$ scaling is

$$A_0 + (1 + \lambda)^{-1} B_0. \tag{75}$$

The oracle member of equation 2 with $f_t = \mu_t$ and $\omega_n = n/(n+m)$ attains this bound. Its exact variance under $\sqrt{n}$ scaling is $A_0 + \{n/(n+m)\}B_0$.

Proof. In this restricted experiment a single covariate density $m(x)$ appears in both samples. A regular submodel has source and target factorizations

$$p_\varepsilon(w) = m_\varepsilon(x) e(x)^t \{1 - e(x)\}^{1-t} p_{t,\varepsilon}(y \mid x), \qquad q_\varepsilon(x) = m_\varepsilon(x).$$

Write s_X for the shared covariate score and s_Y for the source outcome score, with $\mathbb{E}(s_Y \mid X, T) = 0$. The joint tangent vector is parameterized by (s_Y, s_X), and its $\sqrt{n}$-scale inner product is

$$\langle (g_Y, g_X), (h_Y, h_X) \rangle_{\mathrm{com}} = \mathbb{E}_P(g_Y h_Y) + (1 + \lambda)\mathbb{E}_M(g_X h_X). \tag{76}$$

The factor $1 + \lambda$ appears because the same covariate score is observed in both the n source units and the m target units.

Let $\theta_M = \mathbb{E}_M\{\tau(X)\}$ and define

$$\phi_Y(W) = \frac{T}{e(X)}\{Y - \mu_1(X)\} - \frac{1 - T}{1 - e(X)}\{Y - \mu_0(X)\}, \qquad \phi_X(X) = \tau(X) - \theta_M.$$

The pathwise derivative is

$$\dot{\theta}(s_Y, s_X) = \mathbb{E}_P(\phi_Y s_Y) + \mathbb{E}_M(\phi_X s_X).$$

Thus, relative to equation 76, the canonical gradient is $(\phi_Y, (1 + \lambda)^{-1}\phi_X)$. It belongs to the tangent-space closure, so the same orthogonal-projection argument used in the theorem proves minimality. Its squared norm is

$$\mathbb{E}_P(\phi_Y^2) + (1 + \lambda)^{-1}\mathbb{E}_M(\phi_X^2) = A_0 + (1 + \lambda)^{-1} B_0,$$

which proves equation 75 and the lower bound by the convolution theorem.

For clarity, the preceding pair uses the compressed coordinates (s_Y, s_X). In the original product tangent space, whose elements are $(s_Y + s_X, s_X)$ for the source and target samples, respectively, the same Riesz representer is

$$\left(\phi_Y + \frac{\phi_X}{1 + \lambda}, \frac{\phi_X}{1 + \lambda} \right).$$

Indeed, its weighted inner product with $(s_Y + s_X, s_X)$ equals $\mathbb{E}_P(\phi_Y s_Y) + \mathbb{E}_M(\phi_X s_X)$, and its squared weighted norm is again $A_0 + (1 + \lambda)^{-1} B_0$.

It remains to verify attainment. Set $\rho_n = n/(n+m)$. With $r = 1$, $f_t = \mu_t$, and $\omega_n = \rho_n$, the oracle estimator can be written

$$\widehat{\theta}_{\text{pool}} = \mathbb{P}_n \phi_Y + \frac{1}{n+m}\left\{\sum_{i=1}^{n} \tau(X_i) + \sum_{j=1}^{m} \tau(\widetilde{X}_j)\right\}. \tag{77}$$

The residual term is uncorrelated with every function of X because $\mathbb{E}(\phi_Y \mid X) = 0$. The pooled covariate average has variance $B_0/(n+m)$ and the residual average has variance A_0/n. Hence

$$n\text{Var}(\widehat{\theta}_{\text{pool}}) = A_0 + \frac{n}{n+m} B_0 \longrightarrow A_0 + (1+\lambda)^{-1} B_0.$$

This proves attainment and the exact prelimit variance formula. In contrast, substituting $r = 1$ into the unrestricted external-target bound gives $A_0 + \lambda^{-1} B_0$: in that larger model the source covariates carry no information about the independently varying target marginal. The two bounds correspond to different statistical experiments. □

A.10 From the information bound to active data acquisition

This section turns the efficiency bound in Theorem 4.6 into a prospective design problem. The target law and the conditional potential-outcome laws remain fixed; the investigator chooses where to acquire outcomes and, when randomization is available, how to allocate treatment. This formulation makes explicit the absolute-continuity, normalization, and integrability conditions that are implicit in the short statement of Corollary 4.2.

Let $(\mathcal{X}, \mathcal{A}, \nu)$ be a σ-finite measure space that dominates the target covariate law, and write $q = dQ_X/d\nu$. A prospective labeled covariate law P_p has density $p = dP_p/d\nu$. Recruitment is assumed to depend on X alone, so changing p leaves the conditional laws of $(Y(0), Y(1))$ given X unchanged. For a measurable treatment propensity $e : \mathcal{X} \to (0,1)$, define

$$v_e(x) := \frac{\sigma_1^2(x)}{e(x)} + \frac{\sigma_0^2(x)}{1-e(x)}, \qquad s_e(x) := \sqrt{v_e(x)}, \tag{78}$$

$$A(p,e) := \int_{\mathcal{X}} \frac{q(x)^2 s_e^2(x)}{p(x)}\, d\nu(x), \qquad K_e := \int_{\mathcal{X}} q(x) s_e(x)\, d\nu(x). \tag{79}$$

The integrand in $A(p,e)$ is interpreted as $+\infty$ wherever $q(x)s_e(x) > 0$ and $p(x) = 0$, and as zero when $q(x)s_e(x) = p(x) = 0$. For fixed e, the admissible label designs are

$$\mathfrak{P}_e := \left\{p \geq 0 : \int p\, d\nu = 1, \quad Q_X \ll P_p, \quad A(p,e) < \infty\right\}. \tag{80}$$

Thus a candidate design must be normalized, must place labels throughout the target support, and must give a finite source-noise bound.

Lemma A.10 (Design form of the information bound)**.** *Suppose the identification conditions of Theorem 4.6 hold after replacing the historical source law by the prospective law with covariate density p and propensity e. For n prospective labels and m target covariate observations, the efficient first-order variance is*

$$\frac{A(p,e)}{n} + \frac{B_{Q,\tau}}{m}, \qquad B_{Q,\tau} = \text{Var}_Q\{\tau(X)\}. \tag{81}$$

Equivalently, if $m/n \to \lambda \in (0,\infty)$, its value under $\sqrt{n}$ scaling is $A(p,e) + \lambda^{-1} B_{Q,\tau}$.

Proof. For the prospective design, the transport ratio is $r_p = dQ_X/dP_p = q/p$. The labeled component of the canonical gradient in the proof of Theorem 4.6 becomes

$$\phi_{p,e}(W) = \frac{q(X)}{p(X)}\left[\frac{T}{e(X)}\{Y - \mu_1(X)\} - \frac{1-T}{1-e(X)}\{Y - \mu_0(X)\}\right].$$

Its conditional mean given X is zero. Conditional randomization also makes the two arm-specific residual terms mutually exclusive, and therefore

$$\begin{aligned}\mathbb{E}_{P_p}\{\phi_{p,e}^2(W)\} &= \int p(x) \frac{q(x)^2}{p(x)^2}\left\{\frac{\sigma_1^2(x)}{e(x)} + \frac{\sigma_0^2(x)}{1-e(x)}\right\} d\nu(x) \\ &= A(p,e).\end{aligned}$$

The target-covariate component of the canonical gradient is $\tau(X) - \theta_Q$ and is unaffected by prospective source recruitment. Independence of the labeled and target samples gives equation 81. Multiplying by n and using $n/m \to \lambda^{-1}$ gives the second statement. □

Lemma A.10 isolates the part that data acquisition can change. With m fixed, collecting additional target covariates affects only $B_{Q,\tau}/m$; changing p or e acts directly on the source-noise term.

Lemma A.11 (Variational solution for label placement)**.** *Fix e and suppose*

$$0 < K_e < \infty, \qquad 0 < s_e(x) < \infty \quad \textit{for } Q_X\textit{-almost every } x. \tag{82}$$

Then

$$\inf_{p \in \mathfrak{P}_e} A(p, e) = K_e^2, \tag{83}$$

and the unique minimizing probability density, up to ν-null sets, is

$$p_e^\star(x) = \frac{q(x)s_e(x)}{K_e}. \tag{84}$$

At the optimum, the noise-adjusted transport leverage is constant:

$$\frac{q(x)s_e(x)}{p_e^\star(x)} = K_e \quad \textit{for } Q_X\textit{-almost every } x. \tag{85}$$

Proof. Let $S = \{x : q(x) > 0\}$. For any $p \in \mathfrak{P}_e$, the Cauchy–Schwarz inequality on S gives

$$\begin{aligned} K_e^2 &= \left\{ \int_S \frac{q(x)s_e(x)}{\sqrt{p(x)}} \sqrt{p(x)}\, d\nu(x) \right\}^2 \\ &\le \left\{ \int_S \frac{q(x)^2 s_e^2(x)}{p(x)}\, d\nu(x) \right\} \left\{ \int_S p(x)\, d\nu(x) \right\} \le A(p, e). \end{aligned} \tag{86}$$

The proposed density is nonnegative and, by the definition of K_e, integrates to one. Condition equation 82 implies $p_e^\star > 0$ Q_X-almost surely, hence $Q_X \ll P_{p_e^\star}$, and direct substitution gives

$$A(p_e^\star, e) = \int_S \frac{q(x)^2 s_e^2(x)}{q(x)s_e(x)/K_e}\, d\nu(x) = K_e \int_S q(x)s_e(x)\, d\nu(x) = K_e^2.$$

Thus the lower bound in equation 86 is attained.

Equality in the final inequality of equation 86 requires $\int_S p = 1$, so an optimizer places no mass outside the target support. Equality in Cauchy–Schwarz requires $q(x)s_e(x)/\sqrt{p(x)} = c\sqrt{p(x)}$ almost everywhere on S for a constant $c > 0$. Hence $p \propto qs_e$, and normalization forces $p = p_e^\star$. This proves uniqueness and equation 85.

An equivalent Euler-equation derivation is as follows. For a bounded relative perturbation h satisfying $\int h\, d\nu = 0$ and $p + th > 0$ for all sufficiently small $|t|$, differentiation under the integral gives

$$DA_p[h] = -\int_S \frac{q(x)^2 s_e^2(x)}{p(x)^2} h(x)\, d\nu(x), \tag{87}$$

$$D^2 A_p[h, h] = 2 \int_S \frac{q(x)^2 s_e^2(x)}{p(x)^3} h(x)^2\, d\nu(x) > 0 \tag{88}$$

for every admissible direction that is nonzero on S modulo ν. Stationarity under the normalization constraint therefore requires $q^2 s_e^2/p^2$ to be constant, which again gives $p \propto qs_e$; the positive second variation proves strict optimality. □

Condition equation 82 identifies the attainment boundary. If $Q_X\{s_e = 0\} > 0$ but $K_e > 0$, the same Cauchy–Schwarz argument yields the infimum K_e^2, but the formal optimizer in equation 84 assigns zero mass to part of the target support and is not in $\mathfrak{P}_e$. The support-preserving sequence $p_{e,\delta} = (1 - \delta)p_e^\star + \delta q$, $0 < \delta < 1$, approaches the infimum as $\delta \downarrow 0$. If $s_e = 0$ Q_X-almost surely, then $A(p, e) = 0$ for every support-preserving p, and label placement is not identified by the noise bound.

Proof of Corollary 4.2. Set $p = p_{\mathrm{lab}}$. Lemma A.10 identifies $A(p, e)$ as the design-dependent term. The corollary assumes Condition equation 82. Lemma A.11 therefore gives

$$p^\star_{\mathrm{lab}}(x) = \frac{q(x)s_e(x)}{\int q(u)s_e(u)\,d\nu(u)}, \qquad \min_p A(p, e) = \left\{\int q(x)s_e(x)\,d\nu(x)\right\}^2.$$

Now allow e to vary. For fixed x, put $a = \sigma_1(x)$ and $b = \sigma_0(x)$. For every $0 < e < 1$,

$$\frac{a^2}{e} + \frac{b^2}{1-e} - (a+b)^2 = \frac{\{a(1-e) - be\}^2}{e(1-e)} \geq 0. \tag{89}$$

When $a, b > 0$, equality holds uniquely at

$$e^\star(x) = \frac{a}{a+b} = \frac{\sigma_1(x)}{\sigma_1(x) + \sigma_0(x)}.$$

Consequently $s_{e^\star}(x) = \sigma_1(x) + \sigma_0(x)$. Applying Lemma A.11 at $e = e^\star$ yields

$$p^\star_{\mathrm{lab}}(x) = \frac{q(x)\{\sigma_1(x) + \sigma_0(x)\}}{\int q(u)\{\sigma_1(u) + \sigma_0(u)\}\,d\nu(u)},$$

which proves the joint design claim. Strictness in equation 89 and in the label-placement objective proves uniqueness up to null sets. □

Treatment-overlap constraints. When the trial enforces $\epsilon \leq e(x) \leq 1 - \epsilon$ for a fixed $\epsilon \in (0, 1/2)$, the unique pointwise solution is the clipped Neyman allocation

$$e^\star_\epsilon(x) = \mathrm{clip}\left(\frac{\sigma_1(x)}{\sigma_1(x) + \sigma_0(x)}, \epsilon, 1 - \epsilon\right). \tag{90}$$

Indeed, $e \mapsto \sigma_1^2/e + \sigma_0^2/(1-e)$ is strictly convex whenever $\sigma_0 + \sigma_1 > 0$, and its derivative changes sign at the unconstrained Neyman allocation. Projection onto the closed interval therefore gives the constrained minimizer. The corresponding optimal label density is $p^\star(x) \propto q(x)s_{e^\star_\epsilon}(x)$. Formula equation 90 also covers zero conditional variance in one arm, for which the unconstrained infimum occurs at a treatment-probability boundary.

A.10.1 A TWO-STAGE ADAPTIVE DESIGN WITH PILOT-ESTIMATED NOISE

The oracle rule depends on unknown conditional variances. We therefore formalize an implementable two-stage design: an independent pilot estimates the noise surface, and the resulting rule is frozen before the main labeled sample is recruited. The construction is a two-stage adaptive design with standard conditional independence; a fully sequential design requires a separate martingale sampling analysis.

Let $\mathcal{G}_k$ be the sigma-field generated by a pilot of size k, independent of the subsequent labeled and target samples. For fixed e, let $\widehat{s}_k$ be a nonnegative, $\mathcal{G}_k$-measurable estimate of s_e and put

$$\widehat{K}_k = \int q(x)\widehat{s}_k(x)\,d\nu(x), \qquad \widetilde{p}_k(x) = \frac{q(x)\widehat{s}_k(x)}{\widehat{K}_k}$$

when $0 < \widehat{K}_k < \infty$. On any pilot realization with an invalid normalizer, set $\widetilde{p}_k = q$. For deterministic $\eta_k \in (0, 1)$, define the positivity-safe rule

$$\widehat{p}_k(x) = (1 - \eta_k)\widetilde{p}_k(x) + \eta_k q(x). \tag{91}$$

It satisfies $\widehat{p}_k \geq \eta_k q$ and therefore $q/\widehat{p}_k \leq \eta_k^{-1}$ for every pilot realization.

Theorem A.1 (Two-stage adaptive label design)**.** *Suppose $0 < K_e < \infty$, $s_e > 0$ Q_X-almost surely, and $\mathbb{E}_Q(s_e^2) < \infty$. On the event*

$$\mathcal{E}_k(\delta) := \{(1-\delta)s_e \leq \widehat{s}_k \leq (1+\delta)s_e \quad Q_X\text{-almost surely}\}, \qquad 0 \leq \delta < 1, \tag{92}$$

the unregularized plug-in rule and its safety mixture satisfy

$$1 \le \frac{A(\widetilde{p}_k, e)}{A(p_e^\star, e)} \le \frac{1+\delta}{1-\delta}, \tag{93}$$

$$1 \le \frac{A(\widehat{p}_k, e)}{A(p_e^\star, e)} \le \frac{1+\delta}{(1-\eta_k)(1-\delta)}, \tag{94}$$

and

$$\|\widehat{p}_k - p_e^\star\|_{L_1(\nu)} \le \frac{2\delta}{1-\delta} + 2\eta_k. \tag{95}$$

If random $\delta_k = o_p(1)$ satisfy $\Pr\{\mathcal{E}_k(\delta_k)\} \to 1$ and $\eta_k \to 0$, then

$$\frac{A(\widehat{p}_k, e)}{A(p_e^\star, e)} \xrightarrow{p} 1, \qquad \|\widehat{p}_k - p_e^\star\|_{L_1(\nu)} \xrightarrow{p} 0. \tag{96}$$

Conditional on $\mathcal{G}_k$, the main sample is an ordinary fixed design with known realized ratio $q/\widehat{p}_k$. More precisely, let $k = k_n$ and suppose a main-sample estimator admits, conditionally on $\mathcal{G}_{k_n}$, the expansion

$$\widehat{\theta}_n - \theta_Q = \sum_{i=1}^{N_n} \xi_{n,i} + r_n, \qquad \mathbb{E}(\xi_{n,i} \mid \mathcal{G}_{k_n}) = 0, \tag{97}$$

where the $\xi_{n,i}$ are conditionally independent. Define $s_n^2 = \sum_{i=1}^{N_n} \mathbb{E}(\xi_{n,i}^2 \mid \mathcal{G}_{k_n})$. If $s_n > 0$ with probability tending to one, for some $\zeta > 0$

$$\frac{1}{s_n^{2+\zeta}} \sum_{i=1}^{N_n} \mathbb{E}(|\xi_{n,i}|^{2+\zeta} \mid \mathcal{G}_{k_n}) \to_p 0, \qquad \frac{r_n}{s_n} \to_p 0, \tag{98}$$

and a variance estimator satisfies $\widehat{s}_n^2 / s_n^2 \to_p 1$, then

$$\frac{\widehat{\theta}_n - \theta_Q}{\widehat{s}_n} \rightsquigarrow N(0,1). \tag{99}$$

Consequently the associated Wald interval has asymptotic unconditional coverage $1-\alpha$. The independent pilot contributes no additional first-order influence term.

Proof. On $\mathcal{E}_k(\delta)$,

$$(1-\delta)K_e \le \widehat{K}_k \le (1+\delta)K_e. \tag{100}$$

The plug-in density is therefore well defined and support preserving. Direct substitution gives

$$\begin{aligned} A(\widetilde{p}_k, e) &= \widehat{K}_k \int q(x) \frac{s_e^2(x)}{\widehat{s}_k(x)}\, d\nu(x) \\ &\le (1+\delta)K_e \frac{1}{1-\delta} \int q(x) s_e(x)\, d\nu(x) = \frac{1+\delta}{1-\delta} K_e^2. \end{aligned}$$

Lemma A.11 gives the matching lower bound $A(\widetilde{p}_k, e) \ge K_e^2 = A(p_e^\star, e)$, which proves equation 93. Since $\widehat{p}_k \ge (1-\eta_k)\widetilde{p}_k$, we have $A(\widehat{p}_k, e) \le (1-\eta_k)^{-1} A(\widetilde{p}_k, e)$, proving equation 94. Off the relative-error event, the safety floor still gives

$$A(\widehat{p}_k, e) \le \eta_k^{-1} \mathbb{E}_Q\{s_e^2(X)\}. \tag{101}$$

Furthermore,

$$\frac{\widetilde{p}_k(x)}{p_e^\star(x)} = \frac{\widehat{s}_k(x)}{s_e(x)} \frac{K_e}{\widehat{K}_k} \in \left[\frac{1-\delta}{1+\delta}, \frac{1+\delta}{1-\delta}\right]$$

on $\mathcal{E}_k(\delta)$. Integration against $p_e^\star$ yields $\|\widetilde{p}_k - p_e^\star\|_1 \le 2\delta/(1-\delta)$. The triangle inequality and the fact that two densities are at L_1 distance at most two now give

$$\begin{aligned} \|\widehat{p}_k - p_e^\star\|_1 &\le (1-\eta_k)\|\widetilde{p}_k - p_e^\star\|_1 + \eta_k \|q - p_e^\star\|_1 \\ &\le \frac{2\delta}{1-\delta} + 2\eta_k, \end{aligned}$$

proving equation 95. Substitution of $\delta_k = o_p(1)$ and $\eta_k \to 0$ into these deterministic bounds proves equation 96.

For the inference claim, condition on $\mathcal{G}_{k_n}$. The realized design is fixed and the main observations are independent. The conditional Lyapunov–Feller theorem applied to equation 98, followed by the conditional remainder and variance-ratio assumptions, implies, for every bounded continuous g,

$$\mathbb{E}\left[g\left(\frac{\widehat{\theta}_n - \theta_Q}{\widehat{s}_n}\right)\middle|\mathcal{G}_{k_n}\right] \xrightarrow{p} \mathbb{E}\{g(Z)\}, \qquad Z \sim N(0,1).$$

Boundedness and the subsequence principle imply convergence of these conditional expectations in L_1. Taking outer expectations gives $\mathbb{E}g\{(\widehat{\theta}_n - \theta_Q)/\widehat{s}_n\} \to \mathbb{E}g(Z)$, hence equation 99 holds. Continuity of the Gaussian distribution at its quantiles gives unconditional coverage tending to $1-\alpha$. Conditioning freezes the pilot-selected rule, so no pilot term enters the main-sample influence expansion. □

Corollary A.2 (Pilot-estimated joint label and treatment design)**.** *Fix $\epsilon \in (0, 1/2)$ and suppose $0 < \underline{\sigma} \le \sigma_t(x) \le \overline{\sigma} < \infty$ for $t = 0, 1$ and Q_X-almost every x. Suppose independent-pilot estimates are nonnegative Q_X-almost surely and satisfy*

$$\max_{t\in\{0,1\}} \left\| \frac{\widehat{\sigma}_{t,k}}{\sigma_t} - 1 \right\|_{L_\infty(Q)} = o_p(1).$$

Define

$$\widehat{e}_k(x) = \mathrm{clip}\left(\frac{\widehat{\sigma}_{1,k}(x)}{\widehat{\sigma}_{1,k}(x) + \widehat{\sigma}_{0,k}(x)}, \epsilon, 1-\epsilon\right),$$

$$\widehat{s}_k(x) = \left\{\frac{\widehat{\sigma}^2_{1,k}(x)}{\widehat{e}_k(x)} + \frac{\widehat{\sigma}^2_{0,k}(x)}{1-\widehat{e}_k(x)}\right\}^{1/2},$$

where the ratio defining $\widehat{e}_k$ is set to $1/2$ whenever its denominator is zero. If the construction is otherwise invalid on a pilot realization, set $\widehat{e}_k = 1/2$ and $\widehat{p}_k = q$ on that realization, and construct $\widehat{p}_k$ from equation 91 with $\eta_k \to 0$. Let $e^\star_\epsilon$ be equation 90 and $p^\star_\epsilon \propto q s_{e^\star_\epsilon}$. Then

$$\frac{A(\widehat{p}_k, \widehat{e}_k)}{A(p^\star_\epsilon, e^\star_\epsilon)} \xrightarrow{p} 1. \tag{102}$$

Conditional inference remains valid under the conditions in Theorem A.1.

Proof. The map $(a, b) \mapsto a/(a+b)$ is uniformly Lipschitz on $[\underline{\sigma}/2, 2\overline{\sigma}]^2$, and clipping is nonexpansive. Thus $\|\widehat{e}_k - e^\star_\epsilon\|_{L_\infty(Q)} = o_p(1)$. Because both propensities lie in $[\epsilon, 1-\epsilon]$, the same pilot condition gives $\|\widehat{s}_k/s_{\widehat{e}_k} - 1\|_{L_\infty(Q)} = o_p(1)$. Apply Theorem A.1 conditionally with $e = \widehat{e}_k$ to obtain $A(\widehat{p}_k, \widehat{e}_k) = \{1 + o_p(1)\}K^2_{\widehat{e}_k}$. The variance bounds and treatment overlap dominate s_e uniformly over $e \in [\epsilon, 1-\epsilon]$ by an integrable constant. Uniform convergence of $\widehat{e}_k$ and dominated convergence imply $K_{\widehat{e}_k} \to_p K_{e^\star_\epsilon}$. Finally, Lemma A.11 and pointwise convexity in e give $A(p^\star_\epsilon, e^\star_\epsilon) = K^2_{e^\star_\epsilon}$. Combining these displays proves equation 102. □

A.10.2 Recruitment caps, support floors, and truncated designs

Real recruitment systems often impose stratum quotas. The following result gives the exact constrained analogue of equation 84.

Proposition A.4 (Envelope-constrained label design)**.** *Let $S = \{x : q(x) > 0\}$ and $a(x) = q(x)s_e(x) > 0$ on S. Let measurable envelopes ℓ and u vanish on S^c, with $u > 0$ on S, and satisfy*

$$0 \le \ell(x) \le u(x) < \infty, \qquad \int_S \ell \, d\nu < 1 < \int_S u \, d\nu < \infty.$$

Define

$$\mathfrak{P}_e(\ell, u) = \{p \in \mathfrak{P}_e : p = 0 \text{ on } S^c,\ \ell \le p \le u \quad \nu\text{-almost everywhere on } S\}.$$

There exists $c > 0$ such that

$$\int \mathrm{clip}\{c\,a(x), \ell(x), u(x)\}\, d\nu(x) = 1, \tag{103}$$

and the unique minimizer is

$$p^\star_{\ell,u}(x) = \mathrm{clip}\{c\,q(x)s_e(x), \ell(x), u(x)\}. \tag{104}$$

Here $\mathrm{clip}(z, \ell, u) = \min\{u, \max(\ell, z)\}$ *pointwise.*

Proof. Set $F(c) = \int_S \mathrm{clip}\{ca, \ell, u\}\, d\nu$. The function F is continuous and nondecreasing by dominated convergence, with $F(c) \to \int_S \ell < 1$ as $c \downarrow 0$ and $F(c) \to \int_S u > 1$ as $c \uparrow \infty$. Hence a normalizing c in equation 103 exists.

Let $p^* = p^\star_{\ell,u}$ and put $\eta = c^{-2}$. The pointwise objective $z \mapsto a(x)^2/z$ is strictly convex for $z > 0$, so, for every feasible p,

$$\frac{a^2}{p} \geq \frac{a^2}{p^*} - \frac{a^2}{(p^*)^2}(p - p^*). \tag{105}$$

On the lower-bound set $\{p^* = \ell\}$, one has $ca \leq \ell$ and hence $-a^2/(p^*)^2 + \eta \geq 0$, while feasibility gives $p - p^* \geq 0$. On the upper-bound set $\{p^* = u\}$, these two inequalities reverse. On the interior set, $p^* = ca$ and $-a^2/(p^*)^2 + \eta = 0$. Therefore

$$\int \left\{ -\frac{a^2}{(p^*)^2} + \eta \right\} (p - p^*)\, d\nu \geq 0.$$

Because both densities integrate to one, the term involving η vanishes. Integrating equation 105 now gives $A(p, e) \geq A(p^*, e)$. Strict convexity on the target support makes the minimizer unique. □

A cap $r_p(x) = q(x)/p(x) \leq R$ is the lower envelope $\ell(x) = q(x)/R$. Thus a density-ratio cap is not an after-the-fact weight truncation: it is a prospective support floor. Upper envelopes encode limited recruitment capacity in particular strata. Proposition A.4 shows exactly which strata bind and renormalizes the remaining unconstrained allocation.

A.10.3 Cost-sensitive and budget-constrained acquisition

The fixed-n optimum treats every outcome label as equally costly. Let $c(x) \in (0, \infty)$ be the known expected cost of acquiring one labeled outcome at covariate value x. If labels are sampled independently from density p, then $n \int cp\, d\nu$ is their expected cost. Under a budget $\mathcal{B}$, the continuous relaxation uses $n = \mathcal{B}/\int cp$; the source component of the unscaled information bound is therefore

$$\mathcal{V}_{\mathcal{B}}(p, e) = \frac{1}{\mathcal{B}} \left\{ \int \frac{q^2 s_e^2}{p}\, d\nu \right\} \left\{ \int cp\, d\nu \right\}. \tag{106}$$

Proposition A.5 (Cost-weighted label design). *Fix e and suppose $s_e, c > 0$ Q_X-almost surely and*

$$L_c := \int \frac{q(x)s_e(x)}{\sqrt{c(x)}}\, d\nu(x) < \infty, \qquad K_c := \int q(x)s_e(x)\sqrt{c(x)}\, d\nu(x) < \infty. \tag{107}$$

Then the unique budget-optimal sampling density is

$$p^\star_c(x) = \frac{q(x)s_e(x)/\sqrt{c(x)}}{L_c}, \tag{108}$$

and

$$\min_p \mathcal{V}_{\mathcal{B}}(p, e) = \frac{K_c^2}{\mathcal{B}}. \tag{109}$$

At this design, $A(p^\star_c, e) = K_c L_c$, the expected cost per label is K_c/L_c, and the continuous optimal label count is $n^\star = \mathcal{B}L_c/K_c$.

Proof. Apply Cauchy–Schwarz to $qs_e/\sqrt{p}$ and $\sqrt{cp}$:

$$\left\{\int \frac{q^2 s_e^2}{p}\,d\nu\right\}\left\{\int cp\,d\nu\right\} \geq \left\{\int qs_e\sqrt{c}\,d\nu\right\}^2 = K_c^2. \tag{110}$$

Equality holds if and only if $q(x)s_e(x)/\sqrt{p(x)}$ is proportional to $\sqrt{c(x)p(x)}$, or equivalently $p(x) \propto q(x)s_e(x)/\sqrt{c(x)}$. The normalizer is L_c, which proves equation 108 and uniqueness. Direct substitution gives

$$A(p_c^\star, e) = L_c \int qs_e\sqrt{c}\,d\nu = K_c L_c, \qquad \int cp_c^\star\,d\nu = K_c/L_c.$$

Equations equation 109 and the expression for $n^\star$ follow. For an integer-valued sample size, rounding $n^\star$ changes the variance by $o(\mathcal{B}^{-1})$ as $\mathcal{B} \to \infty$. □

If the sample size n is fixed and only the average-cost constraint $\int cp\,d\nu \leq \bar{c}$ is imposed, the Karush–Kuhn–Tucker equation instead gives

$$p^\star(x) = \frac{q(x)s_e(x)}{\sqrt{\eta + \gamma c(x)}}, \tag{111}$$

where η normalizes the density, $\gamma \geq 0$, and $\gamma\{\int cp^\star - \bar{c}\} = 0$. This form reduces to equation 84 when the cost constraint is slack. It can be combined with the envelopes in Proposition A.4 by clipping the right-hand side of equation 111.

When treatment arms have different costs, it is cleaner to optimize arm-specific sampling intensities.

Proposition A.6 (Arm-specific cost allocation)**.** *Let $h_t(x)$ denote the expected number of recruited units per unit ν-mass assigned to arm t, and let $c_t(x) > 0$ be the corresponding per-outcome cost. Suppose*

$$0 < D_c := \sum_{t=0}^{1} \int q(x)\sigma_t(x)\sqrt{c_t(x)}\,d\nu(x) < \infty.$$

Among nonnegative intensities, interpreting $q^2\sigma_t^2/h_t$ as zero when $q\sigma_t = h_t = 0$ and as $+\infty$ when $q\sigma_t > 0 = h_t$, satisfying $\sum_t \int c_t h_t\,d\nu \leq \mathcal{B}$, the minimum of the source-noise bound $\sum_t \int q^2\sigma_t^2/h_t\,d\nu$ is $D_c^2/\mathcal{B}$ and is attained uniquely on $\{q\sigma_t > 0\}$ at

$$h_t^\star(x) = \frac{\mathcal{B}}{D_c}\frac{q(x)\sigma_t(x)}{\sqrt{c_t(x)}}. \tag{112}$$

The induced treatment propensity is

$$e_c^\star(x) = \frac{\sigma_1(x)/\sqrt{c_1(x)}}{\sigma_1(x)/\sqrt{c_1(x)} + \sigma_0(x)/\sqrt{c_0(x)}}. \tag{113}$$

On a set where both conditional standard deviations vanish, the propensity may be assigned any fixed value in $(0,1)$. Equal arm costs recover the Neyman allocation in Corollary 4.2.

Proof. On the product space $\{0,1\} \times \mathcal{X}$, Cauchy–Schwarz yields

$$\begin{aligned} D_c^2 &= \left\{\sum_{t=0}^{1} \int \frac{q(x)\sigma_t(x)}{\sqrt{h_t(x)}}\sqrt{c_t(x)h_t(x)}\,d\nu(x)\right\}^2 \\ &\leq \left\{\sum_{t=0}^{1} \int \frac{q(x)^2\sigma_t^2(x)}{h_t(x)}\,d\nu(x)\right\}\left\{\sum_{t=0}^{1} \int c_t(x)h_t(x)\,d\nu(x)\right\}. \end{aligned}$$

The budget constraint therefore makes the source-noise bound at least $D_c^2/\mathcal{B}$. Equality requires both a binding budget and $q\sigma_t/\sqrt{h_t}$ proportional to $\sqrt{c_t h_t}$, so $h_t \propto q\sigma_t/\sqrt{c_t}$. Enforcing the budget gives equation 112; direct substitution attains the lower bound. The equality condition also gives uniqueness wherever $q\sigma_t > 0$. Finally, a unit recruited at x is assigned to arm one with proportion $h_1^\star(x)/\{h_0^\star(x) + h_1^\star(x)\}$, which simplifies to equation 113. □

A.10.4 LOCAL SUPPORT AND THE VALUE OF THE NEXT LABEL

The same objective quantifies where weak support becomes an information barrier. For every measurable $H \subseteq \mathcal{X}$, Cauchy–Schwarz restricted to H gives

$$A(p,e) \geq \int_H \frac{q^2 s_e^2}{p}\, d\nu \geq \frac{\left\{\int_H q(x)s_e(x)\, d\nu(x)\right\}^2}{P_p(H)}. \tag{114}$$

If $P_p(H) = 0$ while $\int_H q s_e > 0$, the bound is infinite and $Q_X \not\ll P_p$. More generally, along any sequence of feasible designs for which $P_p(H) \downarrow 0$, the efficiency bound diverges at least at the rate $1/P_p(H)$. This is the quantitative form of the support barrier: target covariates quantify the mass of H, but only labels in H identify its conditional outcome means.

Equation equation 87 supplies a local acquisition rule. For a mass-preserving perturbation h, the marginal change is $DA_p[h] = -\int\{q(x)s_e(x)/p(x)\}^2 h(x)\, d\nu(x)$. Thus an additional infinitesimal amount of recruitment is most valuable where the current noise-adjusted transport leverage $r_p(x)s_e(x)$ is largest. The global optimum in equation 85 is exactly the design that equalizes this leverage over the target support. If the feasible design class forces $p = 0$ on a target-positive stratum, active design cannot repair the failure; the target must be restricted, the effect partially identified, or the recruitment constraint changed.

A.10.5 EXACT TWO-STRATUM DESIGN CALCULATION

Return to the running example with $P_X = (4/5, 1/5)$, $Q_X = (1/5, 4/5)$, $\sigma_0^2 = \sigma_1^2 = 1$, and $e = 1/2$. Then $s_e^2 = 4$ and $s_e = 2$ in both strata. Under the historical source design, $r = (1/4, 4)$ and

$$\begin{aligned} A(P_X, e) &= 4\left\{\frac{(1/5)^2}{4/5} + \frac{(4/5)^2}{1/5}\right\} \\ &= 4\left(\frac{1}{20} + \frac{16}{5}\right) = 13. \end{aligned} \tag{115}$$

The heavy-user stratum alone contributes $4(4/5)^2/(1/5) = 64/5 = 12.8$, which is also the local lower bound in equation 114 for $H = \{X = 1\}$.

Because s_e is constant, Lemma A.11 gives $p^\star = Q_X = (1/5, 4/5)$ and $r_{p^\star} \equiv 1$. The source-noise term becomes

$$A(p^\star, e) = 4\left\{\frac{(1/5)^2}{1/5} + \frac{(4/5)^2}{4/5}\right\} = 4. \tag{116}$$

Thus prospective target-matched labeling reduces the irreducible source term from 13 to 4 at the same labeled sample size. Since $B_{Q,\tau} = 16/25$, the complete efficient variance changes from $13/n + 16/(25m)$ to $4/n + 16/(25m)$. At $m/n = 10$, the corresponding $\sqrt{n}$-scale bounds are 13.064 and 4.064. Increasing m alone can remove only the common 0.064 term.

Suppose recruitment capacity limits the heavy-user share to at most $3/5$. The unconstrained solution violates this cap, so Proposition A.4 places the solution at $p^{\mathrm{cap}} = (2/5, 3/5)$. Its exact source term is

$$A(p^{\mathrm{cap}}, e) = 4\left\{\frac{(1/5)^2}{2/5} + \frac{(4/5)^2}{3/5}\right\} = 4\left(\frac{1}{10} + \frac{16}{15}\right) = \frac{14}{3}. \tag{117}$$

The cap prevents attainment of the ideal value 4, but the active design still removes $25/27 \approx 92.6\%$ of the excess historical penalty $(13 - 4)$.

Finally, suppose a heavy-user label costs four times a light-user label, $c = (1, 4)$. Proposition A.5 gives $p_c^\star \propto (2/5, 4/5)$, hence $p_c^\star = (1/3, 2/3)$. Here $A(p_c^\star, e) = 108/25$, $\mathbb{E}_{p_c^\star} c = 3$, and $\mathcal{V}_{\mathcal{B}}(p_c^\star, e) = 324/(25\mathcal{B})$. For comparison, the budget-normalized source variances under target-matched and historical sampling are, respectively,

$$\frac{4\{(1/5)1 + (4/5)4\}}{\mathcal{B}} = \frac{68}{5\mathcal{B}}, \qquad \frac{13\{(4/5)1 + (1/5)4\}}{\mathcal{B}} = \frac{104}{5\mathcal{B}}.$$

The cost-weighted solution therefore moves labels toward the high-leverage heavy-user stratum while accounting for its higher acquisition cost.

Why the factor r^2 is unavoidable. Under covariate shift, a rare labeled unit with large $r(X)$ represents many target units. Conditional outcome noise is therefore multiplied by $r(X)$ in the influence function and by $r(X)^2$ in its variance. Replacing r^2 by r would understate both the information bound and the uncertainty of any regular estimator.

A.11 Worked two-stratum calculation

Let $P(X = 0, 1) = (4/5, 1/5)$, $Q(X = 0, 1) = (1/5, 4/5)$, $e = 1/2$, $\tau = (1, 3)$, and $\tau_f = (1/2, 7/2)$. Then $r = (1/4, 4)$, $\theta_Q = 13/5$, $\mathbb{E}_Q \tau_f = \mathbb{E}_P(r\tau_f) = 29/10$, and $\mathbb{E}_P Z_f = -3/10$. This proves the displayed $2.9 - 0.3 = 2.6$ target-preservation arithmetic.

To determine variance, the arm-level predictor must be specified, not only its contrast. Set $\mu_0 = f_0 = 0$, $\mu_1 = \tau$, $f_1 = \tau_f$, and $Y(t) = \mu_t(X) + \varepsilon_t$, where the conditional errors are independent Rademacher variables and are independent of treatment given X. With $U = Z_f$ and $R = r\tau_f$,

$$R(0) = \frac{1}{8}, \quad R(1) = 14, \qquad \mathbb{E}(U \mid X = 0) = \frac{1}{8}, \quad \mathbb{E}(U \mid X = 1) = -2,$$

and $\mathbb{E}(U^2 \mid X = x) = \frac{9}{2} r(x)^2$. Direct calculation gives

$$A = \frac{2907}{200} = 14.535, \qquad C_{UR} = -\frac{1887}{400} = -4.7175,$$
$$B_P = \frac{12321}{400} = 30.8025, \quad B_{Q,f} = \frac{36}{25} = 1.44.$$

At $(n, m) = (100, 1000)$,

$$V(0) = 0.14679, \quad V(1) = 0.359025, \quad \omega^\star = \frac{3241}{20631} = 0.157094, \quad V(\omega^\star) = 0.139153.$$

For the ratio-only perturbation, instead grant the oracle outcome regressions and take $g = (4, -1)$ and $\varepsilon = 0.1$. Then $\mathbb{E}_Q g = 0$, $\widetilde{r} = r(1 + \varepsilon g) = (0.35, 3.6)$, and $\varepsilon\omega\mathbb{E}_Q(g\tau) = -0.16\omega$. Finally, unit arm variances give $s_e^2 = 4$, $A_{P,Q} = \mathbb{E}_P(r^2 s_e^2) = 13$, and $B_{Q,\tau} = \mathrm{Var}_Q(\tau) = 0.64$. Thus the efficiency bound is $13 + 0.64/\lambda$. The full oracle, constrained, pilot-adaptive, and cost-weighted acquisition calculations are given in Appendix A.10.5.

B General causal moments and quantiles

The target-preserving construction extends to finite-dimensional causal moments and, on a fixed finite grid, to distribution and quantile inference. All simultaneous probability statements below range over that fixed collection of indices.

B.1 Exact decomposition for transported causal moments

Let $\vartheta \in \Theta$ index a collection of moments. For $t \in \{0, 1\}$, let $h_{t,\vartheta} : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}^d$ and $a_\vartheta : \mathcal{X} \to \mathbb{R}^d$ be measurable. Define the target moment

$$\Psi(\vartheta) = \mathbb{E}_Q\left[a_\vartheta(X) + \sum_{t=0}^{1} h_{t,\vartheta}\{X, Y(t)\}\right]. \tag{118}$$

Signs are absorbed into $h_{t,\vartheta}$; for example, the ATE uses $h_1(x, y) = y$, $h_0(x, y) = -y$, and $a = 0$.

Assumption B.1 (Moment-specific transportability)**.** *For every (t, ϑ) under consideration and every coordinate $j \le d$, assume*

$$h_{t,\vartheta,j}\{X, Y(t)\} \in L_1(P) \cap L_1(Q). \tag{119}$$

The ordinary L_1 conditional means

$$\mu^P_{t,\vartheta}(x) = \mathbb{E}_P[h_{t,\vartheta}\{X, Y(t)\} \mid X = x],$$
$$\mu^Q_{t,\vartheta}(x) = \mathbb{E}_Q[h_{t,\vartheta}\{X, Y(t)\} \mid X = x]$$

then satisfy $\mu^P_{t,\vartheta}(x) = \mu^Q_{t,\vartheta}(x)$ for Q_X-almost every x. Their common version is denoted $\mu_{t,\vartheta}$. The function a_ϑ is also Q-integrable coordinatewise.

Assumption 2.3 implies Assumption B.1 for the linear ATE scores just given. For nonlinear scores it is a distinct condition. In particular, indicator scores for distribution functions require transportability of the associated conditional event probabilities, not only transportability of conditional outcome means.

Let $f_{t,\vartheta} : \mathcal{X} \to \mathbb{R}^d$ be arbitrary fixed measurable functions and put

$$g_{f,\vartheta}(x) = a_\vartheta(x) + \sum_{t=0}^{1} f_{t,\vartheta}(x). \tag{120}$$

Define the prediction and residual components

$$\Gamma_f(\vartheta) = \mathbb{E}_Q\{g_{f,\vartheta}(X)\}, \tag{121}$$

$$\Delta_f(\vartheta) = \mathbb{E}_P\left[r(X)\sum_{t=0}^{1}\frac{\mathbf{1}\{T=t\}}{e_t(X)}\{h_{t,\vartheta}(X,Y) - f_{t,\vartheta}(X)\}\right]. \tag{122}$$

Here and below vector equalities are coordinatewise.

Theorem B.1 (Exact causal-moment decomposition and calibrated path)**.** *Suppose Assumptions 2.1, 2.2, and B.1 hold, and suppose $Q_X \ll P_X$ with $r = dQ_X/dP_X$. Fix ϑ and suppose, for every coordinate $j \le d$,*

$$\mathbb{E}_Q\,|g_{f,\vartheta,j}(X)| < \infty, \tag{123}$$

$$\mathbb{E}_P\left[r(X)\sum_{t=0}^{1}\frac{\mathbf{1}\{T=t\}}{e_t(X)}\,|h_{t,\vartheta,j}(X,Y) - f_{t,\vartheta,j}(X)|\right] < \infty. \tag{124}$$

Then, for every fixed choice of the functions $f_{t,\vartheta}$,

$$\Psi(\vartheta) = \Gamma_f(\vartheta) + \Delta_f(\vartheta). \tag{125}$$

Moreover, for any fixed $\omega_\vartheta \in \mathbb{R}$, define

$$\begin{aligned}\widehat{\Psi}_{f,\omega}(\vartheta) = (1-\omega_\vartheta)\mathbb{Q}_m g_{f,\vartheta} + \omega_\vartheta \mathbb{P}_n\{r g_{f,\vartheta}\} \\ + \mathbb{P}_n\left[r(X)\sum_{t=0}^{1}\frac{\mathbf{1}\{T=t\}}{e_t(X)}\{h_{t,\vartheta}(X,Y) - f_{t,\vartheta}(X)\}\right].\end{aligned} \tag{126}$$

If f and ω_ϑ are fixed externally, then their evaluation-sample mean satisfies

$$\mathbb{E}_{\text{eval}}\left\{\widehat{\Psi}_{f,\omega}(\vartheta)\right\} = \Psi(\vartheta). \tag{127}$$

The same sampling-mean identity holds for almost every realization when f and ω_ϑ are measurable with respect to a training–calibration sigma-field $\mathcal{A}$ independent of the evaluation samples and equation 123– equation 124 hold almost surely. In that random case, if additionally, for every coordinate j,

$$\begin{aligned}\mathbb{E}[(1+|\omega_\vartheta|)\mathbb{E}_Q\,|g_{f,\vartheta,j}(X)|] < \infty, \\ \mathbb{E}\left[\mathbb{E}_P\left\{r(X)\sum_{t=0}^{1}\frac{\mathbf{1}\{T=t\}}{e_t(X)}\,|h_{t,\vartheta,j}(X,Y) - f_{t,\vartheta,j}(X)|\right\}\right] < \infty,\end{aligned} \tag{128}$$

then $\widehat{\Psi}_{f,\omega,j}(\vartheta)$ is jointly integrable, $\mathbb{E}\{\widehat{\Psi}_{f,\omega}(\vartheta) \mid \mathcal{A}\} = \Psi(\vartheta)$ almost surely, and $\mathbb{E}\{\widehat{\Psi}_{f,\omega}(\vartheta)\} = \Psi(\vartheta)$. No correctness condition on $f_{t,\vartheta}$ is required.

Proof. Fix t, ϑ, and a coordinate j. Treatment overlap makes the inverse propensity well-defined. Define the weighted arm residual

$$R_{t,\vartheta,j} = r(X)\frac{\mathbf{1}\{T=t\}}{e_t(X)}\{h_{t,\vartheta,j}(X,Y) - f_{t,\vartheta,j}(X)\}.$$

Condition equation 124 gives $R_{t,\vartheta,j} \in L_1(P)$. On $\{r(X) > 0\}$, conditional integrability of $R_{t,\vartheta,j}$ and treatment overlap justify the usual armwise calculation. Consistency and conditional exchangeability give

$$\begin{aligned}\mathbb{E}_P(R_{t,\vartheta,j} \mid X) &= r(X)\mathbb{E}_P\left[\frac{\mathbf{1}\{T=t\}}{e_t(X)}\{h_{t,\vartheta,j}(X,Y) - f_{t,\vartheta,j}(X)\}\,\middle|\, X\right] \\ &= r(X)\{\mu^P_{t,\vartheta,j}(X) - f_{t,\vartheta,j}(X)\}.\end{aligned}$$

On $\{r(X) = 0\}$ both sides are defined as zero, so the equality holds P-almost surely. Conditional Jensen now gives

$$\mathbb{E}_P\left[r(X)\left|\mu^P_{t,\vartheta,j}(X) - f_{t,\vartheta,j}(X)\right|\right] \leq \mathbb{E}_P\left[r(X)\frac{\mathbf{1}\{T=t\}}{e_t(X)}\left|h_{t,\vartheta,j}(X,Y) - f_{t,\vartheta,j}(X)\right|\right] < \infty.$$

Thus the change-of-law step below is an equality of finite quantities, rather than a formal subtraction of extended expectations. Assumption B.1 and Lemma A.1 therefore give

$$\begin{aligned}\mathbb{E}_P\left[r(X)\frac{\mathbf{1}\{T=t\}}{e_t(X)}\{h_{t,\vartheta,j}(X,Y) - f_{t,\vartheta,j}(X)\}\right] &= \mathbb{E}_P[r(X)\{\mu_{t,\vartheta,j}(X) - f_{t,\vartheta,j}(X)\}] \\ &= \mathbb{E}_Q[\mu_{t,\vartheta,j}(X) - f_{t,\vartheta,j}(X)]\,.\end{aligned}$$

Summing over t yields

$$\Delta_{f,j}(\vartheta) = \mathbb{E}_Q\left[\sum_{t=0}^{1}\{\mu_{t,\vartheta,j}(X) - f_{t,\vartheta,j}(X)\}\right].$$

Adding equation 121 cancels the arbitrary predictions and gives

$$\begin{aligned}\Gamma_{f,j}(\vartheta) + \Delta_{f,j}(\vartheta) &= \mathbb{E}_Q\left[a_{\vartheta,j}(X) + \sum_{t=0}^{1}\mu_{t,\vartheta,j}(X)\right] \\ &= \Psi_j(\vartheta),\end{aligned}$$

where the final equality is the tower property under the target law. This proves equation 125 in every coordinate.

For equation 127, first take fixed external functions and weight. Lemma A.1 gives $\mathbb{E}_P(rg_{f,\vartheta}) = \mathbb{E}_Q\{g_{f,\vartheta}\}$. Consequently, the expectation of the first two terms in equation 126 is $(1-\omega_\vartheta)\Gamma_f(\vartheta) + \omega_\vartheta\Gamma_f(\vartheta) = \Gamma_f(\vartheta)$, and the third term has expectation $\Delta_f(\vartheta)$ by the calculation above. Equation equation 125 proves the fixed evaluation-sample identity.

Now let the functions and weight be $\mathcal{A}$-measurable. Independence ensures that the conditional evaluation law remains $P^n \otimes Q_X^m$. The fixed calculation therefore applies for almost every external realization on which equation 123– equation 124 hold. Finally, equation 128, the triangle inequality, and

$$\mathbb{E}_P\{r\,|g_{f,\vartheta,j}|\} = \mathbb{E}_Q\,|g_{f,\vartheta,j}|$$

show that every term in equation 126 is jointly integrable. The sampling-mean identity is then the ordinary conditional-expectation equality given $\mathcal{A}$, and the tower property proves unconditional unbiasedness. □

B.2 Finite-sample transfer for vector moments

The next lemma records the set operation behind finite-sample root inversion. It is deterministic once the component coverage events occur and therefore does not require independence between the two component regions.

Lemma B.1 (Finite-family confidence-set transfer)**.** *Let $\Theta_0 \subseteq \Theta$ be finite. For every $\vartheta \in \Theta_0$, let $\widehat{\mathcal{G}}(\vartheta), \widehat{\mathcal{D}}(\vartheta) \subseteq \mathbb{R}^d$ be random sets satisfying*

$$\Pr\Big\{\Gamma_f(\vartheta) \in \widehat{\mathcal{G}}(\vartheta) \text{ for all } \vartheta \in \Theta_0\Big\} \geq 1 - \alpha_G,$$
$$\Pr\Big\{\Delta_f(\vartheta) \in \widehat{\mathcal{D}}(\vartheta) \text{ for all } \vartheta \in \Theta_0\Big\} \geq 1 - \alpha_D.$$

Define the Minkowski-sum region

$$\widehat{\mathcal{M}}(\vartheta) = \widehat{\mathcal{G}}(\vartheta) \oplus \widehat{\mathcal{D}}(\vartheta) = \{g + d : g \in \widehat{\mathcal{G}}(\vartheta), d \in \widehat{\mathcal{D}}(\vartheta)\}. \tag{129}$$

Then

$$\Pr\Big\{\Psi(\vartheta) \in \widehat{\mathcal{M}}(\vartheta) \text{ for all } \vartheta \in \Theta_0\Big\} \geq 1 - \alpha_G - \alpha_D. \tag{130}$$

In particular, with

$$\mathcal{R}_0 = \{\vartheta \in \Theta_0 : \Psi(\vartheta) = 0\},$$
$$\widehat{\mathcal{R}} = \{\vartheta \in \Theta_0 : 0 \in \widehat{\mathcal{M}}(\vartheta)\},$$

we have $\Pr(\mathcal{R}_0 \subseteq \widehat{\mathcal{R}}) \geq 1 - \alpha_G - \alpha_D$.

Proof. Let E_G and E_D denote the two simultaneous component-coverage events. The union bound gives $\Pr(E_G \cap E_D) \geq 1 - \alpha_G - \alpha_D$; no independence is used. On $E_G \cap E_D$, the exact identity $\Psi(\vartheta) = \Gamma_f(\vartheta) + \Delta_f(\vartheta)$ implies $\Psi(\vartheta) \in \widehat{\mathcal{M}}(\vartheta)$ simultaneously for every $\vartheta \in \Theta_0$. If additionally $\Psi(\vartheta) = 0$, then $0 \in \widehat{\mathcal{M}}(\vartheta)$, so $\vartheta \in \widehat{\mathcal{R}}$. This proves both claims. □

The following theorem supplies a direct rectangular region for the combined moment estimator. Let $\mathcal{J}$ be a finite collection of pairs $j = (\vartheta, \ell)$, where $\ell \in \{1, \ldots, d\}$ denotes a coordinate. For $j = (\vartheta, \ell)$, write

$$L_j(W) = r(X) \sum_{t=0}^{1} \frac{\mathbf{1}\{T = t\}}{e_t(X)} \{h_{t,\vartheta,\ell}(X, Y) - f_{t,\vartheta,\ell}(X)\} + \omega_\vartheta r(X) g_{f,\vartheta,\ell}(X), \tag{131}$$

$$G_j(\widetilde{X}) = (1 - \omega_\vartheta) g_{f,\vartheta,\ell}(\widetilde{X}). \tag{132}$$

Theorem B.2 (Finite-sample simultaneous coordinate region)**.** *Under the conditions of Theorem B.1, suppose the finite collection $\mathcal{J}$, all predictors, and all weights are fixed externally or by a training–calibration sample independent of the evaluation data. Suppose that, conditional on that external sample, each L_j lies almost surely in an interval of known length $H_{P,j}$ and each G_j lies almost surely in an interval of known length $H_{Q,j}$. Fix $\alpha \in (0, 1)$ and choose numbers $\alpha_j \in (0, 1)$ with $\sum_{j \in \mathcal{J}} \alpha_j \leq \alpha$ and define*

$$h_j = \sqrt{\frac{\log(2/\alpha_j)}{2} \left(\frac{H_{P,j}^2}{n} + \frac{H_{Q,j}^2}{m} \right)}. \tag{133}$$

Then

$$\Pr\Big[\Big|\widehat{\Psi}_{f,\omega,\ell}(\vartheta) - \Psi_\ell(\vartheta)\Big| \leq h_j \text{ for every } j = (\vartheta, \ell) \in \mathcal{J}\Big] \geq 1 - \alpha. \tag{134}$$

The result holds both conditionally on the external sample and unconditionally. For a fixed ϑ whose d coordinates belong to $\mathcal{J}$, the Cartesian product

$$\widehat{\mathcal{C}}(\vartheta) = \prod_{\ell=1}^{d} \Big[\widehat{\Psi}_{f,\omega,\ell}(\vartheta) - h_{(\vartheta,\ell)}, \widehat{\Psi}_{f,\omega,\ell}(\vartheta) + h_{(\vartheta,\ell)}\Big] \tag{135}$$

is therefore a simultaneous finite-sample confidence rectangle for $\Psi(\vartheta)$. Equal allocation $\alpha_j = \alpha/|\mathcal{J}|$ replaces the logarithm in equation 133 by $\log\{2|\mathcal{J}|/\alpha\}$.

Proof. Condition on the external training–calibration sigma-field. For $j = (\vartheta, \ell)$, Theorem B.1 gives the exact centered representation. Put $\overline{L}_j = \mathbb{E}_P\{L_j(W)\}$ and $\overline{G}_j = \mathbb{E}_Q\{G_j(\widetilde{X})\}$. Then

$$\widehat{\Psi}_{f,\omega,\ell}(\vartheta) - \Psi_\ell(\vartheta) = \frac{1}{n} \sum_{i=1}^{n} \{L_j(W_i) - \overline{L}_j\} + \frac{1}{m} \sum_{q=1}^{m} \{G_j(\widetilde{X}_q) - \overline{G}_j\}.$$

The $n+m$ summands are independent conditional on the external sample. Hoeffding's inequality for independent bounded variables, with their actual weights $1/n$ and $1/m$, gives

$$\Pr\left[\left|\widehat{\Psi}_{f,\omega,\ell}(\vartheta) - \Psi_\ell(\vartheta)\right| > s \;\middle|\; f, (\omega_\vartheta)\right] \leq 2\exp\left[-\frac{2s^2}{H_{P,j}^2/n + H_{Q,j}^2/m}\right].$$

Substituting $s = h_j$ makes this probability at most α_j. A union bound over $j \in \mathcal{J}$ proves the conditional version of equation 134. Taking expectations over the external sample proves the unconditional version. The rectangle statement is a restatement of the same coordinate event. □

Remark B.1 (Scope of root inversion). *Lemma B.1 and Theorem B.2 give exact inference over a fixed finite candidate set. A continuum-index, cross-fitted Z-estimation theorem requires additional structure: local root uniqueness, a continuous nonsingular Jacobian, uniform nuisance-drift control, and stochastic equicontinuity of the indexed empirical process. Those conditions are not implied by the pointwise decomposition equation 125.*

B.3 Finite-grid distribution bands

For $t \in \{0,1\}$, define the target potential-outcome distribution $F_t(u) = \Pr_Q\{Y(t) \leq u\}$. Fix a deterministic grid $\mathcal{U}_K = \{u_1 < \cdots < u_K\}$, or choose the grid using data independent of the evaluation samples. For every (t,k), assume the indicator version of Assumption B.1:

$$\Pr_Q\{Y(t) \leq u_k \mid X = x\} = \Pr_P\{Y(t) \leq u_k \mid X = x\} \quad \text{for } Q_X\text{-almost every } x. \tag{136}$$

Let $f_{t,k} : \mathcal{X} \to [0,1]$ be any fixed prediction of this conditional event probability and let $\omega_{t,k} \in [0,1]$ be fixed or independently calibrated. Set

$$\widehat{F}_{t,k} = (1-\omega_{t,k})\mathbb{Q}_m f_{t,k} + \omega_{t,k}\mathbb{P}_n(r f_{t,k}) + \mathbb{P}_n\left[r(X)\frac{\mathbf{1}\{T=t\}}{e_t(X)}\{\mathbf{1}(Y \leq u_k) - f_{t,k}(X)\}\right]. \tag{137}$$

For later use, denote its labeled and target summands by

$$L_{t,k}^F(W) = r(X)\left[\frac{\mathbf{1}\{T=t\}}{e_t(X)}\{\mathbf{1}(Y \leq u_k) - f_{t,k}(X)\} + \omega_{t,k} f_{t,k}(X)\right], \tag{138}$$

$$G_{t,k}^F(\widetilde{X}) = (1-\omega_{t,k}) f_{t,k}(\widetilde{X}). \tag{139}$$

Theorem B.3 (Finite-grid CDF bands). *Suppose Assumptions 2.1–2.2, $Q_X \ll P_X$, and equation 136 hold at every grid point. Suppose the quantities in equation 137 are integrable and, conditional on external training and calibration, the summands $L_{t,k}^F$ and $G_{t,k}^F$ lie almost surely in intervals of known lengths $H_{t,k}^P$ and $H_{t,k}^Q$. Fix $\alpha \in (0,1)$ and define*

$$h_{t,k}^F = \sqrt{\frac{\log(4K/\alpha)}{2}\left\{\frac{(H_{t,k}^P)^2}{n} + \frac{(H_{t,k}^Q)^2}{m}\right\}}, \tag{140}$$

and the clipped pointwise limits

$$L_{t,k}^{\mathrm{raw}} = \min\{1, \max\{0, \widehat{F}_{t,k} - h_{t,k}^F\}\}, \qquad U_{t,k}^{\mathrm{raw}} = \max\{0, \min\{1, \widehat{F}_{t,k} + h_{t,k}^F\}\}. \tag{141}$$

The monotone envelopes

$$\underline{F}_{t,k} = \max_{1 \leq \ell \leq k} L_{t,\ell}^{\mathrm{raw}}, \qquad overline F_{t,k} = \min_{k \leq \ell \leq K} U_{t,\ell}^{\mathrm{raw}} \tag{142}$$

satisfy

$$\Pr\left\{\underline{F}_{t,k} \leq F_t(u_k) \leq \overline{F}_{t,k} \text{ for every } t \in \{0,1\},\ k \leq K\right\} \geq 1-\alpha. \tag{143}$$

Coverage holds conditionally on external training and calibration and hence also unconditionally. Both envelope sequences are nondecreasing in k.

One sufficient boundedness condition is $0 \leq r(X) \leq \bar{r}$ and $e_t(X) \geq \epsilon > 0$ almost surely: then one may use

$$H_{t,k}^P = \bar{r}(2/\epsilon + 1), \qquad H_{t,k}^Q = 1 - \omega_{t,k}, \tag{144}$$

although sharper known ranges give shorter bands.

Proof. Apply Theorem B.1 with $d = 1$, $a = 0$, $h_{t,u_k}(x, y) = \mathbf{1}(y \le u_k)$ for the selected arm and the other arm score identically zero. Equation equation 136 is exactly the moment-specific transportability condition for this score. It follows that

$$\mathbb{E}_{\text{eval}}(\widehat{F}_{t,k}) = F_t(u_k) \quad \text{for every fixed external realization.} \tag{145}$$

There are $2K$ arm–threshold pairs. Apply Theorem B.2 with $\alpha_{t,k} = \alpha/(2K)$. Since $\log\{2/\alpha_{t,k}\} = \log(4K/\alpha)$, the radius is precisely equation 140. A union bound therefore gives, with conditional probability at least $1 - \alpha$,

$$\widehat{F}_{t,k} - h^F_{t,k} \le F_t(u_k) \le \widehat{F}_{t,k} + h^F_{t,k} \quad \text{for all } (t, k). \tag{146}$$

Because every CDF value lies in $[0, 1]$, clipping the two endpoints preserves these inequalities.

It remains to verify that monotone postprocessing preserves coverage. Fix t and k. For every $\ell \le k$, monotonicity of F_t and the lower half of equation 146 give $L^{\text{raw}}_{t,\ell} \le F_t(u_\ell) \le F_t(u_k)$. Taking the maximum over $\ell \le k$ yields $\underline{F}_{t,k} \le F_t(u_k)$. Similarly, for every $\ell \ge k$, $F_t(u_k) \le F_t(u_\ell) \le U^{\text{raw}}_{t,\ell}$; taking the minimum gives $F_t(u_k) \le \overline{F}_{t,k}$. These arguments hold simultaneously on equation 146, proving equation 143. The index sets in the maximum expand with k, while those in the minimum shrink, so both envelope sequences are nondecreasing.

Finally, under the sufficient condition in equation 144, $\mathbf{1}\{T = t\}\{\mathbf{1}(Y \le u_k) - f_{t,k}(X)\}/e_t(X)$ lies in $[-1/\epsilon, 1/\epsilon]$, while $\omega_{t,k} f_{t,k}(X) \in [0, 1]$. Multiplication by $r(X) \in [0, \bar{r}]$ places the labeled summand in $[-\bar{r}/\epsilon, \bar{r}(1/\epsilon + 1)]$, an interval of length $\bar{r}(2/\epsilon + 1)$. The target summand lies in $[0, 1 - \omega_{t,k}]$. This proves equation 144. □

B.4 Quantile inversion with grid and support boundaries

For $a \in (0, 1)$, define the left generalized inverse

$$q_t(a) = \inf\{u \in \mathbb{R} : F_t(u) \ge a\}. \tag{147}$$

This definition permits atoms, flat portions of the CDF, and nonunique solutions of $F_t(u) = a$. Let $s_{t,-} \in \mathbb{R} \cup \{-\infty\}$ and $s_{t,+} \in \mathbb{R} \cup \{+\infty\}$ be known support bounds satisfying

$$\Pr_Q\{s_{t,-} \le Y(t) \le s_{t,+}\} = 1. \tag{148}$$

When no finite support information is available, take $s_{t,-} = -\infty$ and $s_{t,+} = +\infty$.

Theorem B.4 (Finite-grid quantile and QTE brackets)**.** *On the simultaneous event in equation 143, define, for every $a \in (0, 1)$,*

$$\underline{q}_t(a) = \max\left\{s_{t,-}, \sup\{u_k : \overline{F}_{t,k} < a\}\right\}, \tag{149}$$

$$\overline{q}_t(a) = \min\left\{s_{t,+}, \inf\{u_k : \underline{F}_{t,k} \ge a\}\right\}, \tag{150}$$

where $\sup \varnothing = -\infty$ and $\inf \varnothing = +\infty$. Then, simultaneously for both arms and every $a \in (0, 1)$,

$$\underline{q}_t(a) \le q_t(a) \le \overline{q}_t(a). \tag{151}$$

Consequently, with probability at least $1 - \alpha$, equation 151 holds simultaneously for every requested quantile level without an additional union bound over levels.

Define the grid-bracketed quantile-treatment-effect set

$$\mathcal{I}_{\text{QTE}}(a) = \{x_1 - x_0 : x_t \in [\underline{q}_t(a), \overline{q}_t(a)] \cap \mathbb{R},\ t = 0, 1\}. \tag{152}$$

It contains $q_1(a) - q_0(a)$ on the same event. If all four bracket endpoints are finite, this set is the closed interval

$$\mathcal{I}_{\text{QTE}}(a) = \left[\underline{q}_1(a) - \overline{q}_0(a), \overline{q}_1(a) - \underline{q}_0(a)\right]. \tag{153}$$

Finite known support bounds guarantee finite brackets even when the grid does not cross the level a.

Proof. Fix an arm t, a level $a \in (0, 1)$, and an outcome of the data in the event equation 143. Because F_t is the CDF of a real-valued outcome and $a \in (0, 1)$, the set $\{u \in \mathbb{R} : F_t(u) \ge a\}$ is nonempty

and bounded below. Hence $q_t(a)$ in equation 147 is a finite real number. First consider any grid point u_k satisfying $\overline{F}_{t,k} < a$. Coverage implies $F_t(u_k) \le \overline{F}_{t,k} < a$. If $q_t(a) < u_k$, then the definition of the infimum in equation 147 supplies a point $v < u_k$ with $F_t(v) \ge a$; monotonicity would give $F_t(u_k) \ge F_t(v) \ge a$, a contradiction. Hence $u_k \le q_t(a)$. Taking the supremum over every such grid point gives

$$\sup\{u_k : \overline{F}_{t,k} < a\} \le q_t(a). \tag{154}$$

The support condition gives $s_{t,-} \le q_t(a)$: below $s_{t,-}$ the CDF is zero, whereas $a > 0$. Combining this fact with equation 154 proves $\underline{q}_t(a) \le q_t(a)$, including the case in which the supremum is over the empty set.

Next consider any grid point u_k satisfying $\underline{F}_{t,k} \ge a$. Coverage gives $F_t(u_k) \ge \underline{F}_{t,k} \ge a$, so u_k belongs to the set whose infimum defines $q_t(a)$. Therefore $q_t(a) \le u_k$, and

$$q_t(a) \le \inf\{u_k : \underline{F}_{t,k} \ge a\}. \tag{155}$$

The upper support condition gives $F_t(s_{t,+}) = 1$ when $s_{t,+} < \infty$, and hence $q_t(a) \le s_{t,+}$. Combining this with equation 155 proves $q_t(a) \le \overline{q}_t(a)$, including the empty-infimum case.

The proof used only the common CDF-band event, not a level-specific probability statement. It therefore holds simultaneously for every $a \in (0, 1)$ and both arms on that event. Since each true arm quantile lies in its bracket, their difference belongs to the set in equation 152. For finite endpoints, elementary interval arithmetic gives equation 153. □

The boundary conventions in Theorem B.4 keep the quantile bracket defined on a coarse grid. If no grid upper band falls below a, the lower bracket is the known lower support endpoint, or $-\infty$ when none is known. If no grid lower band reaches a, the upper bracket is the known upper support endpoint, or $+\infty$. Thus a coarse grid produces a wide but valid bracket instead of an undefined quantile. Atoms and CDF plateaus need no separate treatment because equation 147 always selects the leftmost crossing.

Remark B.2 (Finite grid versus a continuum-index band)**.** *The probability calculation in Theorem B.3 controls the $2K$ fixed arm–threshold estimators and no uncountable score class. The monotonicity of a CDF permits a deterministic coarse interpolation: between u_k and u_{k+1}, one may use $\underline{F}_{t,k}$ as a lower bound and $\overline{F}_{t,k+1}$ as an upper bound, with 0 and 1 outside the grid. This interpolation retains the unresolved between-grid variation. Continuum-index bands, mesh-free uniform rates, and outcome-dependent grids require a separate uniform concentration argument for the indexed, possibly learned, score class.*

C Experimental protocol and additional results

Monte Carlo design. The calibration experiment fixes a deterministic predictor family ex ante. For $x = (x_1, \dots, x_6)$ and alignment $\kappa \in [-1, 1]$, define

$$\begin{aligned} h(x) &= 0.9\{0.7x_2 - 0.5x_3 + 0.3\sin(x_5)\}, \\ \tau_f^{(\kappa)}(x) &= 1 + \kappa\{\tau(x) - 1\} + (1 - \kappa)h(x), \\ f_0^{(\kappa)}(x) &= \mu_0(x) + 0.45(1 - \kappa)\{0.6x_2 + 0.25x_3^2 - 0.25\}, \qquad f_1^{(\kappa)} = f_0^{(\kappa)} + \tau_f^{(\kappa)}. \end{aligned}$$

Thus $\kappa = 1$ is oracle, while negative κ reverses genuine effect heterogeneity and adds smooth irrelevant structure. The predictor is fixed rather than estimated from a simulated training sample. Every replicate instead draws independent calibration and inference samples, each with a labeled source block and a covariate-only target block. The code fixes the random seed, records every replicate, and computes uncertainty summaries from the raw CSV. The labeled blocks have binary treatment drawn from the stated propensity. Predictor quality, the target-to-labeled sample ratio, and Gaussian mean shift are varied factorially. The calibration sample estimates $(A, C_{UR}, B_P, B_{Q,f})$ and freezes $\widehat{\omega}$ before the inference sample is exposed. Oracle $\omega^\star$ is computed from a large Monte Carlo population and is used only as a reference.

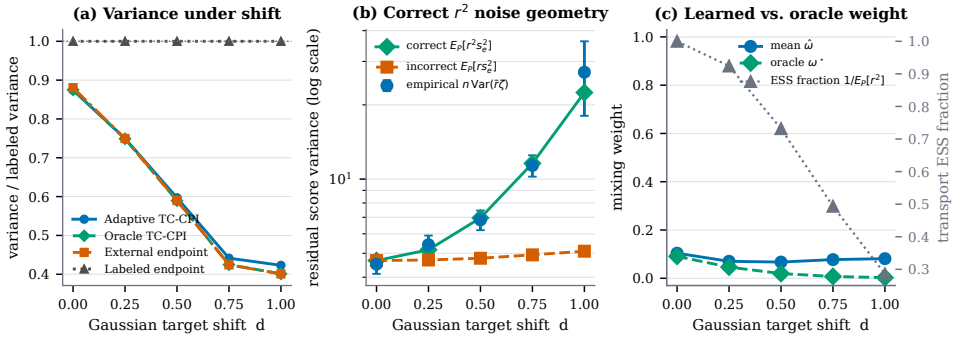


Figure 4: The transport barrier (1,000 repetitions per shift). Adaptive and oracle variances fall as target averaging becomes informative (left), but labeled residual noise grows with shift. Its empirical value tracks the correct r^2 expression rather than the r surrogate (center). Transport effective sample size collapses even though calibration remains near the external endpoint (right).

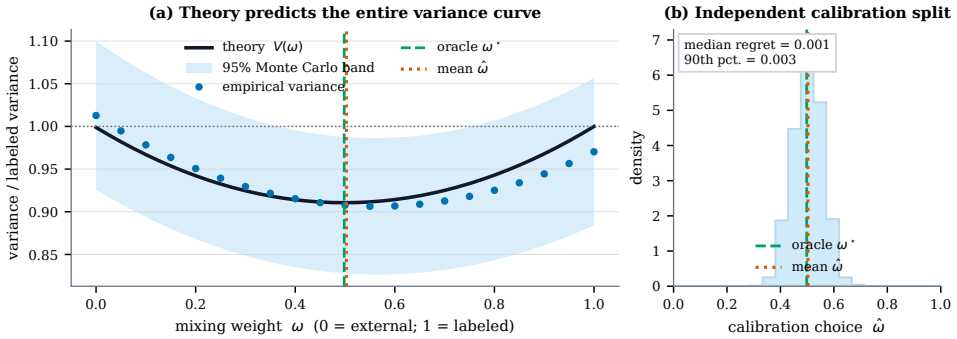


Figure 5: Numerical verification of Equation 4. The solid quadratic is computed from population moments; points and uncertainty bars are Monte Carlo variances over 1,000 independent inference samples. All evaluated weights fall inside their 95% Monte Carlo bands, and independent calibration concentrates near the analytic minimizer ($\mathbb{E}\widehat{\omega} = 0.503$ versus $\omega^\star = 0.498$).

Methods. We compare transported labeled AIPW ($\omega = 1$), external PPI ($\omega = 0$), independently calibrated $\widehat{\omega}$, and the infeasible oracle $\omega^\star$. All four use the same prediction functions and true design weights in the exact-theory experiment. A separate nuisance experiment replaces the propensity and density ratio by cross-fitted estimates and is reported as an asymptotic study. Exact finite-sample coverage is assessed only in the known-weight bounded experiment.

Residual-correction ablation. To isolate Theorem 3.1, we hold the independently calibrated $\widehat{\omega}$ fixed and delete only the labeled residual term. The shifted design uses $n = n_c = 480$, $m = m_c = 4{,}800$, alignment 0.25, and 2,000 independent repetitions. Full TC-CPI has bias 0.0057 (Monte Carlo standard error 0.0044) and coverage 0.953. Without residual correction, bias is -0.6321, matching the exact omission bias -0.6328, and coverage is zero. The two estimators have identical replicate-level calibration weights.

Target-shift mechanism check. Figure 5 checks the quadratic variance calculation against repeated sampling. Figure 4 separately compares $n\mathrm{Var}(\mathbb{P}_n\phi_P)$ to $\mathbb{E}_P[r^2\{\sigma_1^2/e + \sigma_0^2/(1-e)\}]$ and to the incorrect expression with only one power of r. The comparison shows the rapid uncertainty inflation produced by weak target support.

Bounded-summand finite-sample experiment. To isolate Theorem 4.4, we use a finite eight-stratum covariate distribution, binary outcomes, known e and r, and three fixed predictors. Each of 27 cells combines $n \in \{80, 200, 500\}$ and $m/n \in \{1, 5, 20\}$ with 2,000 independent calibra-

Table 1: Independent-split test of Theorems 4.1–4.3 at $m/n = 10$. Variance ratios use the labeled endpoint within the same predictor regime. Bold marks the best non-oracle ratio/length and coverage nearest 0.95, retaining displayed ties.

Alignment κ	Method	$E[\omega]$	Bias	Variance	Var. ratio	Coverage	Avg. length
-1.000	Adaptive TC-CPI	0.996	-0.012	0.053	1.002	0.935	**0.885**
-1.000	Oracle TC-CPI	1.000	-0.012	0.053	1.000	0.936	0.885
-1.000	External endpoint	0.000	-0.014	0.064	1.214	**0.947**	0.985
-1.000	Labeled endpoint	1.000	-0.012	0.053	**1.000**	0.936	**0.885**
0.250	Adaptive TC-CPI	0.451	0.002	0.017	**0.989**	**0.950**	**0.513**
0.250	Oracle TC-CPI	0.449	0.002	0.017	0.986	0.951	0.513
0.250	External endpoint	0.000	0.001	0.017	1.000	0.953	0.516
0.250	Labeled endpoint	1.000	0.003	0.017	1.000	0.953	0.517
1.000	Adaptive TC-CPI	0.101	-0.002	0.009	0.813	0.955	**0.390**
1.000	Oracle TC-CPI	0.091	-0.002	0.009	0.810	0.958	0.389
1.000	External endpoint	0.000	-0.002	0.009	**0.809**	0.956	**0.390**
1.000	Labeled endpoint	1.000	-0.001	0.011	1.000	**0.950**	0.426

Table 2: Test of Theorem 4.4 at $m/n = 5$. Hoeffding and empirical Bernstein (EB) use known supports after independent calibration; EB is a conservative finite-sample bound, and Wald is the asymptotic efficiency reference.

Predictor	n	$\mathbb{E}\hat{\omega}$	H. cov.	EB cov.	W. cov.	H. len.	EB len.	W. len.
Strong	80	0.490	1.000	1.000	0.947	0.884	1.843	0.428
	200	0.467	1.000	1.000	0.951	0.559	0.992	0.271
	500	0.434	1.000	1.000	0.955	0.353	0.557	0.171
Coarse	80	0.416	1.000	1.000	0.947	0.884	1.903	0.429
	200	0.386	1.000	1.000	0.954	0.559	1.030	0.272
	500	0.335	1.000	1.000	0.953	0.353	0.582	0.172
Adversarial	80	0.855	1.000	1.000	0.950	1.125	2.467	0.452
	200	0.933	1.000	1.000	0.948	0.711	1.330	0.286
	500	0.976	1.000	1.000	0.953	0.450	0.750	0.180

tion/inference repetitions. For every realized $\hat{\omega}$, the code enumerates the entire finite (X, T, Y) support to obtain $H_{P,\hat{\omega}}$ and $H_{Q,\hat{\omega}}$; sample extrema are never substituted for known ranges. Figure 6 reports all 27 cells, while Table 2 gives the $m/n = 5$ slice. The minimum empirical Hoeffding coverage is 0.9995 (minimum exact 95% Clopper–Pearson lower bound 0.9972), while the normalized-array Maurer–Pontil empirical-Bernstein construction covers in every cell (minimum Clopper–Pearson lower bound 0.9982). Its mean length is 1.40–2.27 times the Hoeffding length, so variance adaptation does not automatically improve worst-case constants at these sample sizes. Across the 27 adaptive cells, Wald coverage ranges from 0.9395 to 0.9580 and is shown only as an asymptotic efficiency reference.

Finite-grid distributions and quantiles. Figure 7 evaluates the joint CDF event in equation 143 and its quantile implication in equation 151. We use a four-stratum law with $Y(t) \mid X \sim \mathrm{Binomial}\{8, p_t(X)\}$, a nine-point outcome grid, nine QTE levels, $\omega = 0.35$, and $m = 4n$. Across 2,000 repetitions at each $n \in \{500, 2000, 8000\}$, the single simultaneous CDF event and the induced QTE family event both cover in every repetition (95% Wilson lower bound 0.9981); no monotonicity or band-crossing failure occurs. Median CDF-envelope width contracts from 0.360 to 0.102, and median QTE-bracket width from 8.000 to 3.556. Figure 7 and Table 3 report the full check.

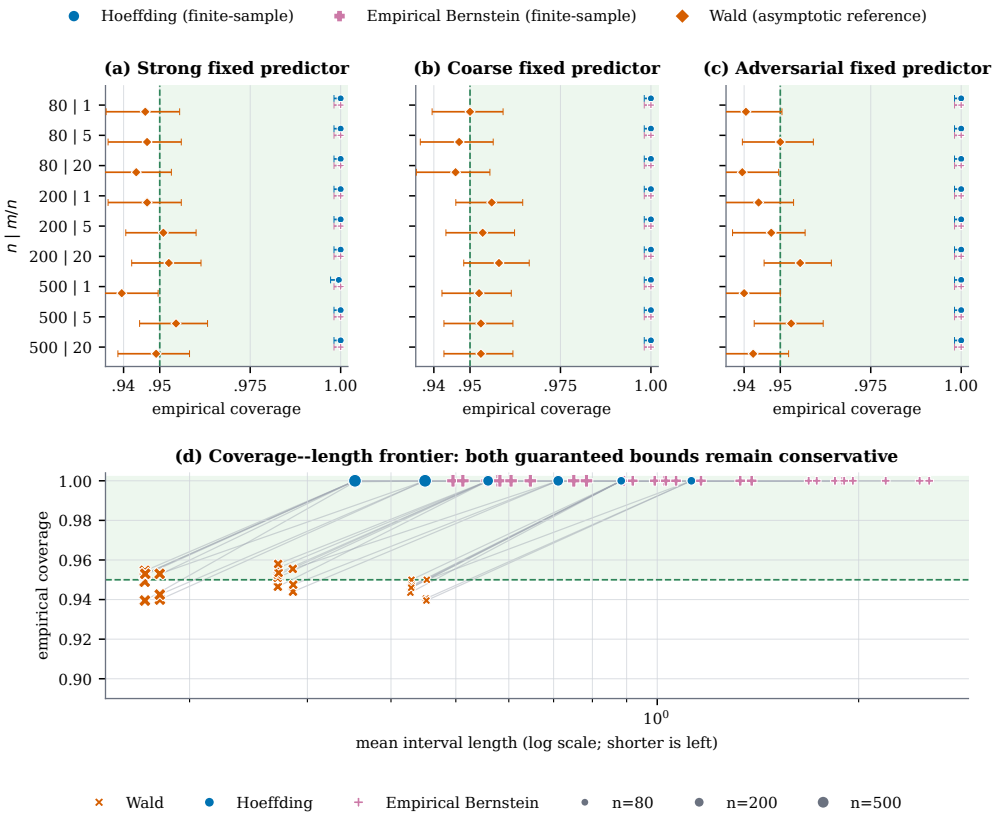


Figure 6: Finite-sample verification of Theorem 4.4. Panels (a)–(c) show empirical coverage with 95% exact binomial intervals for every sample-size and target-ratio cell. Panel (d) connects coverage to interval length. Both known-support concentration intervals are valid and conservative; here the empirical-Bernstein constants make it longer than Hoeffding, while Wald is an asymptotic efficiency reference.

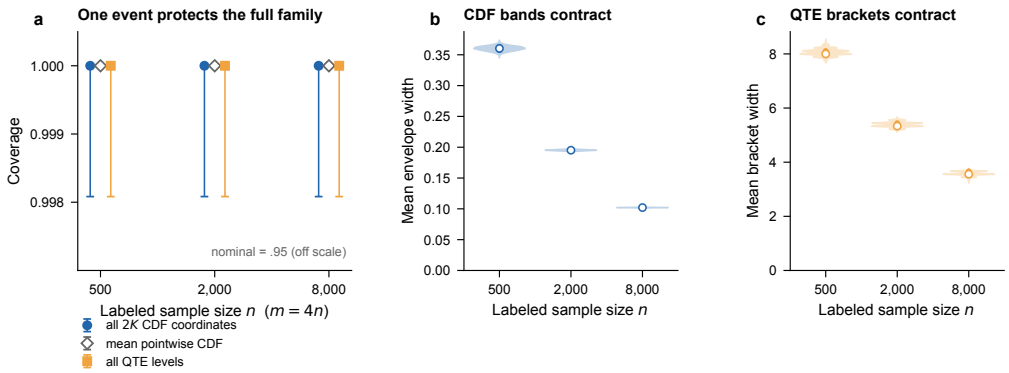


Figure 7: Finite-grid validation of Theorems B.3 and B.4. Left: one simultaneous known-range event covers all 18 arm–threshold CDF coordinates and all nine inverted QTE levels. Center and right: the monotone CDF envelopes and induced QTE brackets contract as the labeled sample grows. Coverage bars are 95% Wilson intervals; violins show replicate-level widths.

Table 3: Tests of Theorems B.3 and B.4 under a bounded four-stratum law (2,000 repetitions per row). Coverage is simultaneous across $2K = 18$ arm–threshold coordinates or across nine QTE levels; widths are replicate medians.

n	m	Simult. CDF	Min. coordinate	QTE family	Med. CDF width	Med. QTE width
500	2,000	1.000	1.000	1.000	0.360	8.000
2,000	8,000	1.000	1.000	1.000	0.195	5.333
8,000	32,000	1.000	1.000	1.000	0.102	3.556

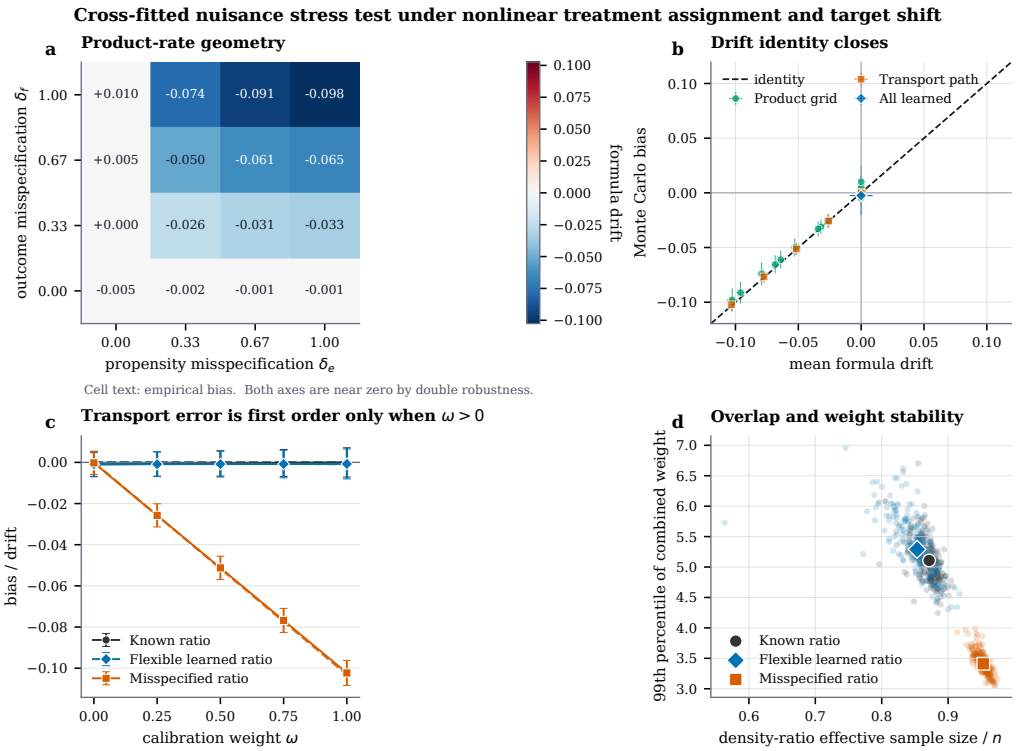


Figure 8: Cross-fitted nuisance stress test. The product surface is flat on either axis when r is known (a), and empirical bias tracks the exact drift over all 33 configurations (b). With a misspecified density ratio and correct outcome model, the $\omega = 0$ endpoint remains centered, whereas the first-order transport term grows linearly along the calibration path (c). Panel (d) relates that failure to density-ratio effective sample size and weight tails. Error bars are 95% Monte Carlo intervals; Wald coverage assesses asymptotic calibration.

Learned-nuisance stress test. The five-fold experiment behind Figure 8 uses a nonlinear propensity and a two-component Gaussian-mixture target. Controlled paths vary outcome and propensity misspecification while holding r known, then vary ω under known, flexible, and deliberately misspecified density ratios. Every fitted score is evaluated on held-out folds. Figure 8 reports all configurations and Table 4 gives representative rows; their cross-fitted Wald intervals assess asymptotic calibration.

Table 4: Tests of Proposition 4.1 and Theorem 4.5. Bias is compared with the exact drift in Equation 7; coverage reports cross-fitted 95% Wald intervals and evaluates the asymptotic guarantee.

Configuration	Bias	Drift	Gap	Coverage	Length	ESS_r/n
Correct e, r; bad f	0.010	-0.000	0.010	0.948	0.609	0.869
Correct f, r; bad e	-0.001	0.000	-0.001	0.956	0.287	0.869
Bad f, e; known r	-0.098	-0.103	0.005	0.860	0.469	0.869
Bad e, r; correct f; $\omega = 0$	-0.000	0.000	-0.000	0.960	0.268	0.952
Bad e, r; correct f; $\omega = .5$	-0.051	-0.052	0.000	0.894	0.266	0.952
All flexible cross-fitted; $\omega = 0$	-0.002	0.000	-0.003	0.958	0.580	0.849
All flexible cross-fitted; $\omega = .5$	-0.003	-0.000	-0.002	0.964	0.573	0.849

Results use 500 Monte Carlo repetitions with $n = 1000$ labeled-source and $m = 2000$ target-covariate observations. The gap is empirical bias minus mean formula drift.

Calibration and nuisance-rate checks. Figure 9 tests the rates that distinguish Theorems 4.3 and 4.5. With $(n, m) = (1000, 5000)$ fixed and 5,000 repetitions per calibration size, mean $n\{V(\widehat{\omega}) - V(\omega^\star)\}$ has fitted log–log slope -0.977 (bootstrap 95% interval $[-1.002, -0.952]$), matching the N_c^{-1} fast rate in equation 6. In the nuisance experiment, propensity and outcome-regression errors both scale as n^{-a}, so the exact product drift scales as n^{-2a}. At $n = 4000$, standardized bias is 0.110, 0.714, and 3.350 for $a = 0.35$, 0.25, and 0.15; the corresponding Wald coverage is 0.942, 0.903, and 0.076. Coverage remains near nominal when the product drift is $o(n^{-1/2})$, falls to 0.903 at the $n^{-1/2}$ boundary, and collapses to 0.076 when the drift is larger.

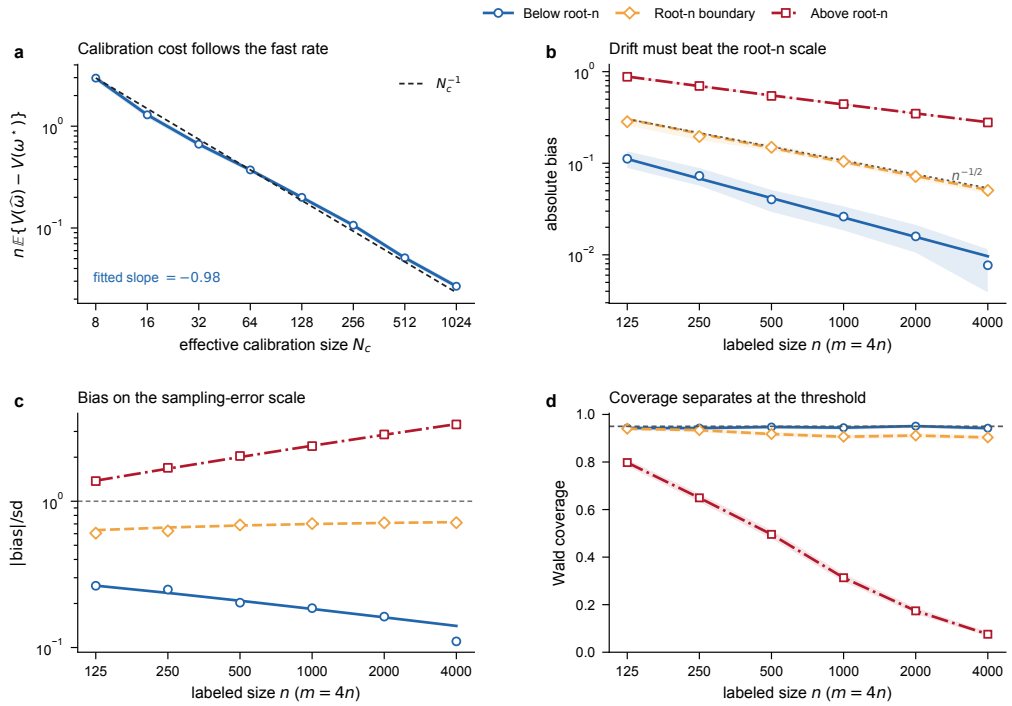


Figure 9: Rate validation for finite calibration and learned nuisances. (a) Excess variance follows the N_c^{-1} calibration rate. (b)–(c) Exact product drift either beats, matches, or exceeds the $n^{-1/2}$ sampling scale. (d) Wald coverage separates accordingly. The band in (a) and ribbons in (b) are simultaneous 95% Monte Carlo intervals over their displayed grids; ribbons in (d) are pointwise 95% Wilson intervals.

Table 5: Rate checks for Theorems 4.3 and 4.5. Calibration entries are mean n-scaled excess variance; Monte Carlo uncertainty is shown in Figure 9. Nuisance entries report standardized bias and Wald coverage at the endpoints of the sample-size grid.

Experiment	Setting	Quantity	Small grid	Large grid
Calibration	N_c	$n\,\mathbb{E}[\Delta V]$	2.9674	0.0268
Calibration	log–log	fitted slope	-0.977	
Nuisance	Below $n^{-1/2}$	$\lvert\text{bias}\rvert/\text{sd}$	0.264	0.110
		coverage	0.943	0.942
Nuisance	Root-n boundary	$\lvert\text{bias}\rvert/\text{sd}$	0.605	0.714
		coverage	0.940	0.903
Nuisance	Above $n^{-1/2}$	$\lvert\text{bias}\rvert/\text{sd}$	1.373	3.350
		coverage	0.798	0.076

Sensitivity and prospective-design experiments. We directly test the two extensions in Corollaries 4.1 and 4.2 in the two-stratum model. The conditional-shift experiment holds the observed-data law fixed, varies $|\Delta_{\text{cond}}| \in [0, 0.4]$, and sets $\rho_r = 0$ so that only structural shift remains. For each magnitude we retain the less favorable boundary sign over 4,000 repetitions. The ordinary 95% Wald interval loses coverage rapidly, from 0.945 at zero shift to 0.424 at $|\Delta_{\text{cond}}| = 0.4$, whereas expansion by $\rho_c = |\Delta_{\text{cond}}|$ gives worst-sign coverage 0.971 at $|\Delta_{\text{cond}}| = 0.4$.

The design experiment compares the historical source law with the oracle label and treatment allocations from Corollary 4.2. With equal conditional variances, changing only the label law from P to Q lowers the theoretical n-scaled variance from 13.064 to 4.064. Under heterogeneous noise, label redesign alone lowers it from 60.889 to 21.051, treatment redesign alone to 40.277, and their joint optimum to 13.754. All six theoretical variances lie inside their corresponding 95% Monte

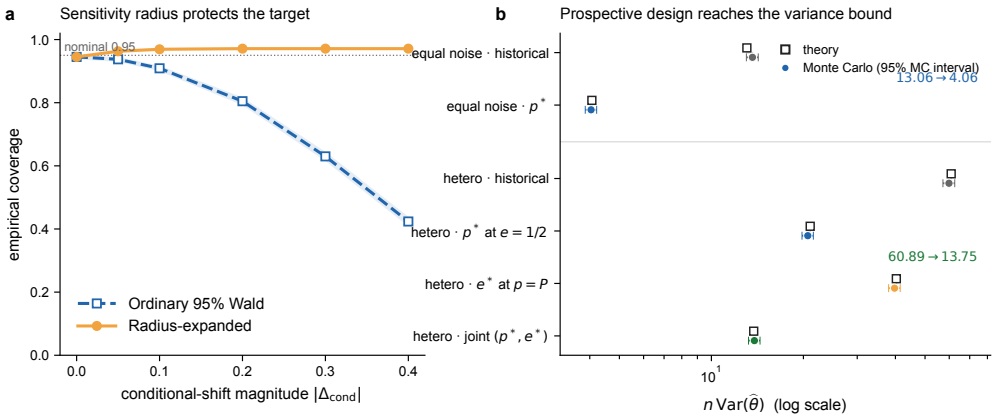


Figure 10: Direct empirical checks of sensitivity expansion and prospective design. Left: ordinary coverage collapses under conditional-effect shift, while the radius-expanded interval protects the target across both boundary signs. Right: Monte Carlo n-scaled variances, with 95% Monte Carlo intervals, track the efficiency bound under historical, label-only, treatment-only, and joint designs. Squares denote the theoretical values.

Table 6: Tests of Corollaries 4.1 and 4.2 in the two-stratum model. Coverage uses the less favorable boundary sign; parentheses in the design block are 95% Monte Carlo intervals.

(a) Conditional shift

$\lvert\Delta_{\mathrm{cond}}\rvert = \rho_c$	Ordinary	Expanded
0.00	0.945	0.945
0.20	0.805	0.971
0.40	0.424	0.971

(b) Prospective design

Noise	Design	Theory	Empirical
Equal	Historical	13.064	13.625 [13.015, 14.234]
Equal	Labels only	4.064	4.040 [3.866, 4.213]
Hetero	Historical	60.889	59.881 [57.222, 62.541]
Hetero	Labels only	21.051	20.671 [19.789, 21.554]
Hetero	Treatment only	40.277	39.736 [38.006, 41.467]
Hetero	Joint optimum	13.754	13.825 [13.225, 14.425]

Carlo intervals; the largest relative theory–simulation discrepancy is 4.29%. Figure 10 shows the two mechanisms and Table 6 gives the numerical values. This first experiment isolates structural sensitivity and oracle design; the next one tests ratio error and pilot-estimated design.

Ratio-robust calibration and independent-pilot design. The ratio experiment implements the normalized perturbation $\widetilde{r} = r(1 + 0.1g)$ from Proposition 4.1, with $b_0 = 0$, $b_1 = -0.16$, and $\rho_r = 0.16$. Across 4,000 repetitions per weight, ordinary coverage decreases from 0.945 at $\omega = 0$ to 0.723 at $\omega = 1$; expansion by $\omega\rho_r$ gives coverage between 0.945 and 0.972. The variance-only weight is 0.13694, whereas the analytic robust weight is 0.04295; dense-grid theoretical and empirical risk minimizers are 0.050 and 0.04505, respectively. The variance coefficients are computed under the actual perturbed-ratio score, and every theoretical robust risk lies inside its simultaneous 95% Monte Carlo interval.

For the two-stage design in Appendix A.10, an independent Gaussian pilot estimates the stratum-specific noise scale and applies the positivity mixture $\eta_k = k^{-1/2}$. As the pilot size per (x, t) cell grows from 8 to 512, median excess source variance falls from 0.697% to 0.017%, its 90th percentile from 3.920% to 0.091%, and median L_1 design error from 0.071 to 0.011. Over 28,000 pilots, the finite regret bound, L_1 bound, and support floor hold without a violation.

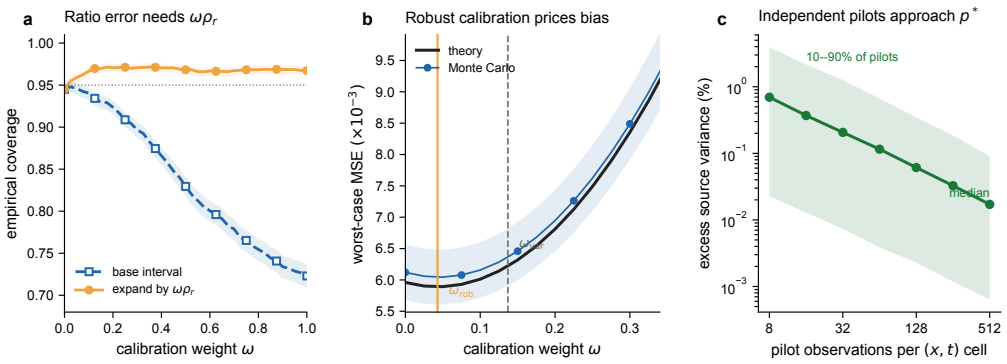


Figure 11: Ratio sensitivity and feasible prospective design. (a) Ratio error creates an ω-linear coverage loss that expansion by $\omega\rho_r$ removes. (b) Pricing this bias shifts calibration from the variance-only weight ω_{var} to the robust weight ω_{rob}. Theory and Monte Carlo risks use the perturbed-ratio score. (c) Independent pilots approach the oracle label density; the line and band show the median and 10–90% range of excess source variance on logarithmic axes.

Table 7: Tests of ratio-robust calibration in Proposition 4.1 and independent-pilot design in Corollary 4.2. The ratio block includes the grid points nearest the analytic robust and variance-only weights; pilot variance is on the $\sqrt{n}$ scale.

(a) Ratio error, $\rho_r = 0.16$

ω	Base cov.	Expanded	Theory $\times 10^3$	Emp. $\times 10^3$
0.000	0.945	0.945	5.960	6.120
0.050	0.946	0.959	5.893	6.045
0.125	0.934	0.970	6.142	6.284
0.500	0.830	0.968	13.682	13.830
1.000	0.723	0.967	40.052	40.346

(b) Independent pilot

k	Med. excess (%)	90th (%)	Med. L_1	Theory / emp. var.
8	0.697	3.920	0.071	21.362/21.174
32	0.206	1.268	0.039	21.147/21.451
128	0.061	0.347	0.021	21.078/20.878
512	0.017	0.091	0.011	21.058/21.086

Large-scale randomized-outcome protocol. The Criteo Uplift Prediction Dataset v2.1 contains 13,979,592 rows, 12 pretreatment features, a randomized advertising-targeting indicator, and the real binary visit outcome (Diemert et al., 2021). We verify the official archive against SHA-256 `2716e1bf0fd157a93b5bf86924d9088419dfbac2022c6cd90030220634f616dc` and draw a fixed 500,000-row sample. Post-treatment exposure is excluded. We create a controlled target domain using

$$s(x) = 0.15 + 0.70\ \mathrm{expit}\{1.10z(f_0) + 0.65z(f_6) - 0.40z(f_8)\}, \qquad r(x) = \frac{1-\pi_Q}{\pi_Q}\frac{s(x)}{1-s(x)},$$

where each z is median/IQR standardized and clipped to $[-3, 3]$. Hence the domain ratio is known and domain overlap is strict. The realized source and target samples contain 260,488 and 239,512 rows; the source 1st and 99th percentiles of r are 0.378 and 1.997.

An 80,000-row source block trains separate-arm constant, logistic, histogram-gradient-boosting, and one-hidden-layer MLP predictors; 20,000 additional source rows diagnose prediction. A 100,000-row target block gives the held-out randomized difference-in-means reference 0.009493 (standard error 0.001757). Each of 100 repetitions then draws mutually disjoint source calibration and inference blocks of 6,000 rows and target calibration and inference blocks of 18,000 rows. The treatment propensity 0.8500875 is estimated once on the disjoint source training block. This is a descriptive learned-propensity, data-adaptive-ω analysis, and its variance ratios summarize stability across 100 prespecified splits drawn from shared evaluation pools.

Figure 12 and Table 8 report the controlled-shift results.

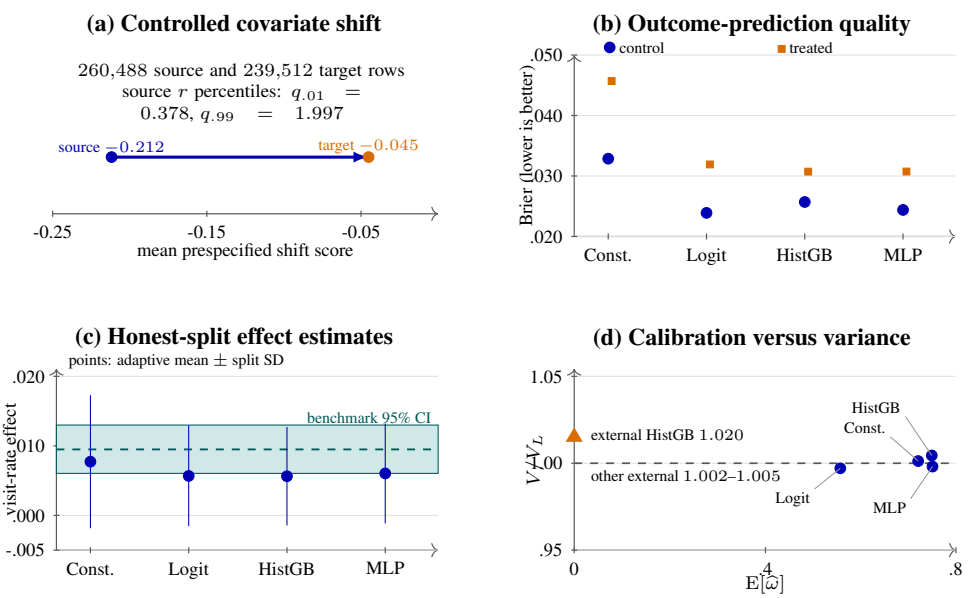


Figure 12: Independent calibration adapts predictor trust on the controlled Criteo shift. Panels report the prespecified domain displacement, held-out predictive quality, adaptive effect estimates against the held-out target benchmark, and variance ratios relative to the transported labeled endpoint. The ratios summarize stability across the prespecified repeated splits.

Table 8: Controlled-shift test of calibration and the prediction–precision distinction on Criteo Uplift v2.1. Ratios summarize repeated-split empirical variance relative to the transported labeled endpoint. Bold marks the lowest unweighted arm-macro Brier score and variance ratios below one.

Predictor	Macro Brier	$\mathbb{E}\hat{\omega}$	V_{adapt}/V_L	V_{ext}/V_L	Adapt. SD	Labeled SD
Constant	0.0390	0.720	1.001	1.002	0.0096	0.0096
Logistic	0.0276	0.557	**0.997**	1.005	0.0072	0.0072
HistGradientBoosting	0.0279	0.750	1.004	1.020	0.0071	0.0071
MLP	**0.0273**	0.751	**0.998**	1.002	0.0072	0.0072

Small-RCT protocol. The complementary study uses the public `AER::STAR` extract of the Tennessee STAR randomized class-size experiment (Krueger, 1999). Small classes are the treatment, regular classes are the control, and the outcome is the mean of the first-grade reading and mathematics scaled scores. The raw extract yields 2,785 complete two-arm records; school-level positivity filtering gives a 2,627-pupil analysis cohort in 69 schools, of whom 1,193 free-lunch pupils form the target population. Repeated four-way splitting reserves disjoint blocks for outcome-model training, calibration, labeled inference, and target covariates. The target block contributes X only: neither its treatment nor its outcome is used anywhere in fitting, calibration, propensity construction, or estimation. Because the public extract omits the original assignment probabilities, we estimate school-specific small-class rates on the training and calibration blocks with eight pseudo-observations of shrinkage toward their pooled rate. These rates range from 0.141 to 0.784 across the 500 splits. The public extract lacks pre–Grade 1 attrition information and classroom identifiers, so method intervals omit censoring weights and use student-level uncertainty. Figure 13 and Table 9 evaluate split stability and precision relative to the estimated full-cohort, school-stratified contrast, whose uncertainty uses a school-cluster sandwich. It estimates the observed free-lunch cohort contrast. The checked input has SHA-256 `0e8b179ea3d883730b25008ca1293c4c8f886aa8d5d299c03ab6ff05d0123ae6`.

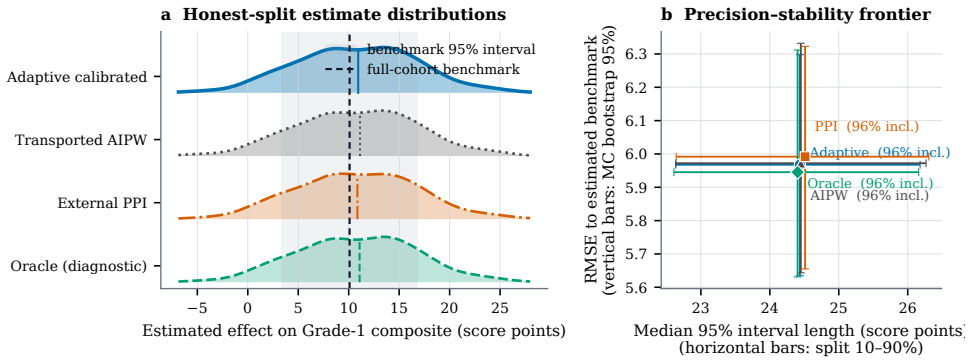


Figure 13: Project STAR uses disjoint splits in a small randomized trial. Split-to-split estimate distributions are shown with interval-length and calibration summaries; the full-data experimental contrast and its school-cluster uncertainty form the reference. TC-CPI uses only its calibration block to choose $\widehat{\omega}$.

Table 9: Project STAR repeated-split test of independent calibration. RMSE and inclusion are computed relative to the estimated full-cohort benchmark.

Method	Mean $\widehat{\theta}$	Split SD / SE	Median CI len.	RMSE to bench.	Bench. incl.	Mean ω
Adaptive TC-CPI	10.78	5.93	24.44	5.97	96.2%	0.50
Transported AIPW ($\omega = 1$)	10.78	5.94	24.45	5.97	96.4%	1.00
External PPI ($\omega = 0$)	10.76	5.96	24.51	5.99	96.4%	0.00
Oracle weight (diagnostic)	10.78	5.91	24.40	5.95	96.4%	0.52
Full-cohort benchmark	10.09	SE=3.44	13.49	–	–	–

Reproduction. The project contains the scripts, aggregate CSV files, selected compact replicate-level outputs, generated tables, vector figures, and machine-readable runtime records. The licensed Criteo archive is not redistributed; its official URL, license, checksum, and deterministic sampling command are recorded. No numerical entry in the paper is typed independently of the saved analysis artifacts.

D MACHINE-CHECKED ALGEBRAIC CORE

The directory `LeanProofs/` contains Lean 4/mathlib formalizations of the deterministic target-preservation identity and the completed-square argument behind the oracle calibration rule. The project is built with a pinned Lean toolchain, contains no `sorry`, `admit`, custom axioms, or unsafe definitions, and records `#print axioms` output. The Lean verification covers these two algebraic identities. Measure-theoretic identification, central limit theorems, and semiparametric tangent-space arguments are established by the conventional proofs above.