BEYOND A FIXED COMPOSITE ENDPOINT: GUARDED ORTHOGONAL LEARNING FOR PREFERENCE-ROBUST TARGET-GROUP AUDITS

Anonymous authors
Paper under double-blind review

ABSTRACT

Clinical trials use a prespecified primary endpoint, yet downstream benefit–risk decisions often combine several outcomes and value trade-offs. We ask whether the treatment conclusion for a prespecified target group remains stable over an admissible utility class. We introduce PROBE, which learns reusable cross-fitted orthogonal effects for interpretable basis outcomes. We compare complete effect vectors under a utility-aware patient-level loss and use a held-out guard to adapt among zero, constant, shared, coordinate-specific, and private structures. The selected vector supports post-fit prediction, whereas confirmatory inference uses cross-fitted basis scores to construct direct multiplier max-t bands over the prespecified group–utility family. These bands drive a four-way audit that returns robust support for either treatment, certified preference sensitivity, or evidence uncertainty. External preference information can then refine the admissible class without rerunning causal estimation. For this pipeline, we establish a uniform basis-approximation bound, an oracle inequality with an explicit exclusion price, false-admission control, asymptotic familywise validity, and an $\alpha+\beta$ error bound combining causal and preference uncertainty. In prespecified raw-data simulations, PROBE admits complex structure in 50/50 strong private and 44/50 strong mixed runs, while all 100 null runs retain an anchor. Two 1,000-run reference experiments attain familywise coverage of 0.945 and 0.949. We then apply the causal audit to the IST randomized trial and analyze dependence-preserving preference evidence using 394 same-respondent part-worth vectors from a tuberculosis discrete-choice study. Together, these results show how PROBE supports transparent target-group auditing across clinically meaningful value trade-offs.

1 INTRODUCTION

Clinical trials require a prespecified primary endpoint, yet downstream decisions often depend on multiple outcomes and on how benefits and harms are traded off. We ask whether, for a prespecified target group, the treatment conclusion remains stable when measured outcomes are evaluated over a clinically defensible utility class, while the original endpoint remains unchanged. This question arises when treatment improves disease control but worsens toxicity, burden, or function: a composite endpoint encodes one aggregation rule, whereas other defensible trade-offs may imply a different conclusion (Freemantle et al., 2003; Cordoba et al., 2010; McKenna & Heaney, 2020). Figure 1 summarizes how we keep the trial endpoint fixed and audit its downstream robustness across these trade-offs.

This audit first creates a representation problem: for nonlinear utilities, marginal component means need not determine expected utility: $\mathbb{E}\{u(\boldsymbol{Y}(a)) \mid X = x\} \neq u[\mathbb{E}\{\boldsymbol{Y}(a) \mid X = x\}]$. Rather than fitting a separate causal model for every utility, we learn cross-fitted orthogonal effects for a prespecified set of observable, clinically interpretable basis outcomes, including joint-event terms when needed. We then evaluate basis-representable utilities as linear post-fit queries.

We align the learning criterion with the audited utility surfaces by defining $G_\Omega = \mathbb{E}_\omega\{c(\omega)c(\omega)^\top\}$ and comparing complete basis-effect vectors under one patient-level quadratic loss. This criterion retains cross-outcome covariance and the utility-class geometry. An honest four-role design freezes

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.

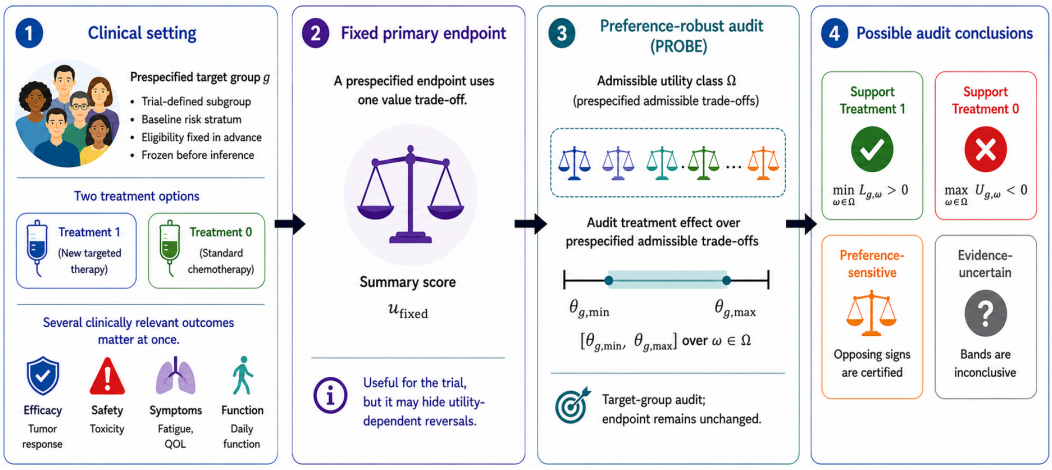


Figure 1: Conceptual overview of PROBE. The prespecified trial endpoint remains unchanged while a target group is evaluated over an admissible utility class. Its population envelope is $[\theta_{g,\min}, \theta_{g,\max}]$, and simultaneous utility-specific bands distinguish robust support for either treatment, certified preference sensitivity, and evidence uncertainty. The resulting output is a group-level assessment of robustness to clinically meaningful value trade-offs.

candidates before guarding and reserves untouched inference data for learned groups or reporting families. We retain zero, constant, and shared vectors as anchors and admit coordinate-wise or private departures only when they improve over every anchor on held-out data. With the group–utility family fixed, direct multiplier max-t bands use cross-fitted orthogonal scores rather than the selected learner. Their linear contrasts let external preference evidence restrict Ω and update the envelope without repeating causal estimation.

We evaluate these components in two complementary real-data settings. We use the IST aspirin comparison to audit a finite same-disease utility family with randomized treatment and patient-level outcomes. We then use same-respondent tuberculosis discrete-choice coefficients to preserve preference dependence; because they are not linked to trial outcomes, the contextualized analysis remains evidence-uncertain.

Our contributions are:

- **Preference-robust target-group auditing.** We formalize a downstream audit that keeps the prespecified trial endpoint fixed and evaluates whether a target-group conclusion is stable over an admissible utility class. The four-way output separates robust support, certified preference sensitivity, and evidence uncertainty.
- **Guarded reusable learning and direct inference.** PROBE learns cross-fitted basis-effect vectors under a utility-aware patient-level loss, adapts between shared and outcome-specific structures through anchor-relative validation, and constructs full-covariance max-t bands over a prespecified group–utility family. We establish approximation, oracle, false-admission, familywise, and two-source uncertainty guarantees.
- **Complementary empirical evaluation.** We use prespecified raw-data simulations to evaluate structural adaptation, fallback, coverage, and decision behavior. We then demonstrate the causal audit in the IST randomized trial and evaluate dependence-preserving preference inputs using the tuberculosis discrete-choice study.

Positioning relative to prior work. These contributions connect several adjacent literatures but target a different inferential object. Methods that estimate or specify composite utilities, conditional outcome distributions, preference-adaptive policies, and multi-task causal effects address related but distinct targets. Table 1 summarizes their primary objectives and outputs, and Appendix A provides a detailed comparison.

Table 1: Schematic comparison of the primary objectives and outputs of representative method classes.

Method class	Post-fit utilities	Joint outcomes	Guarded structure	Simultaneous audit	Preference layer
Learned latent-utility ITR (Luckett et al., 2021)	No	Yes	No	No	Learned/latent
Scalar distributional CATE (Park et al., 2021)	Yes	No	Model-dependent	Usually no	No
Preference-weighted endpoints/rules (Duc & Wolbers, 2017; Bakoyannis, 2023)	Varies	Yes	Varies	Weight/value bands	Specified
Multi-task causal learning (Aglietti et al., 2020; Chen et al., 2020; Jiang et al., 2023b)	No explicit audit	Method-dependent	Sometimes	No	No
PROBE	**Yes**	**Basis-defined**	**Yes**	**Frozen finite family**	**External/modular**

2 PROBLEM FORMULATION

2.1 OBSERVED DATA AND CAUSAL TARGETS

Let $O_i = (X_i, A_i, \boldsymbol{Y}_i)$, $i = 1, \ldots, n$, be i.i.d., with pretreatment covariates $X \in \mathcal{X}$, treatment $A \in \{0, 1\}$, and component outcomes $\boldsymbol{Y} = (Y_1, \ldots, Y_K)^\top$. Let $\boldsymbol{Y}(a)$ denote the potential outcome vector under treatment a, and orient outcomes and utilities so that larger values are preferable.

For a utility u_ω indexed by $\omega \in \Omega$, define

$$\tau_\omega(x) = \mathbb{E}[u_\omega\{\boldsymbol{Y}(1)\} - u_\omega\{\boldsymbol{Y}(0)\} \mid X = x].$$

We use the profile-level surface $\tau_\omega(x)$ for exploratory effect characterization and subgroup construction. Our confirmatory target is the average utility effect within each prespecified group $R_1, \ldots, R_G$:

$$\theta_{g,\omega} = \mathbb{E}\{\tau_\omega(X) \mid X \in R_g\}.$$

The set Ω collects the clinically defensible value trade-offs used in the audit. Under the identification conditions below, we conduct inference for these group-level targets and retain fitted profile surfaces for downstream exploration.

Assumption 1 (Causal identification)**.** *For $a \in \{0, 1\}$: (i) consistency, $\boldsymbol{Y} = \boldsymbol{Y}(A)$; (ii) no interference; (iii) conditional exchangeability, $\{\boldsymbol{Y}(0), \boldsymbol{Y}(1)\} \perp A \mid X$; and (iv) overlap, $\varepsilon \le e(X) \le 1-\varepsilon$ almost surely, where $e(x) = \mathbb{P}(A = 1 \mid X = x)$.*

2.2 APPROXIMATELY STRUCTURED UTILITIES

To evaluate many utilities with one fitted representation, we use a prespecified vector of observable, clinically interpretable basis outcomes, $\boldsymbol{b}(\boldsymbol{y}) = \{b_1(\boldsymbol{y}), \ldots, b_M(\boldsymbol{y})\}^\top$. Write $B_j = b_j(\boldsymbol{Y})$, $\tau_j(x) = \mathbb{E}\{B_j(1) - B_j(0) \mid X = x\}$, and $\boldsymbol{\tau}_B(x) = \{\tau_1(x), \ldots, \tau_M(x)\}^\top$.

Assumption 2 (Approximate utility basis)**.** *For every $\omega \in \Omega$, $u_\omega(\boldsymbol{y}) = c_0(\omega) + \boldsymbol{c}(\omega)^\top \boldsymbol{b}(\boldsymbol{y}) + r_\omega(\boldsymbol{y})$, and $\sup_{\omega \in \Omega,\, \boldsymbol{y}} |r_\omega(\boldsymbol{y})| \le \varepsilon_u$. Each $B_j(a)$ is square-integrable. There are finite constants C_1, C_2 such that $\sup_{\omega \in \Omega} \|\boldsymbol{c}(\omega)\|_1 \le C_1$ and $\sup_{\omega \in \Omega} \|\boldsymbol{c}(\omega)\|_2 \le C_2$.*

The corresponding basis-induced surface is

$$\tau_\omega^B(x) = \boldsymbol{c}(\omega)^\top \boldsymbol{\tau}_B(x).$$

For later bounds, if $\boldsymbol{f} = (f_1, \ldots, f_M)^\top$ is square-integrable, write $\|\boldsymbol{f}\|_{2,2} := \{\sum_j \|f_j\|^2_{L_2(P_X)}\}^{1/2}$.

Proposition 1 (Utility-basis misspecification)**.** *Under Assumption 2, for every x and $\omega \in \Omega$,*

$$|\tau_\omega(x) - \tau_\omega^B(x)| \le 2\varepsilon_u.$$

Thus the basis is exact when $\varepsilon_u = 0$, while omitted interactions incur an explicit approximation price.

For binary outcomes, this approximation becomes exact when the basis contains all joint-event monomials (Proposition 3, Appendix B.1); Proposition 1 covers a clinically chosen low-order basis.

2.3 TARGET-GROUP ENVELOPES AND DOWNSTREAM STABILITY

With the representation fixed, we now translate utility-specific effects into the target-group audit. For a fixed target group R_g, define the true utility envelopes

$$\theta_{g,\min} = \inf_{\omega \in \Omega} \theta_{g,\omega}, \qquad \theta_{g,\max} = \sup_{\omega \in \Omega} \theta_{g,\omega}.$$

At the population level, $\theta_{g,\min} > 0$ means that every admissible utility favors Treatment 1, whereas $\theta_{g,\max} < 0$ means that every admissible utility favors Treatment 0. A sign change in the population envelope makes the conclusion value dependent. Simultaneous intervals allow us to distinguish such a reversal from sampling uncertainty.

Let $[L_{g,\omega}, U_{g,\omega}]$ be simultaneous intervals for the prespecified target-group utility family. If intervals cover only the basis projection, Proposition 1 gives the expansion

$$[L_{g,\omega}, U_{g,\omega}] = [L^B_{g,\omega} - 2\varepsilon_u,\ U^B_{g,\omega} + 2\varepsilon_u]. \tag{1}$$

Suppose these intervals cover the frozen target-group–utility family simultaneously. On that event,

$$\inf_{\omega\in\Omega} L_{g,\omega} \le \theta_{g,\min} \le \inf_{\omega\in\Omega} U_{g,\omega}, \qquad \sup_{\omega\in\Omega} L_{g,\omega} \le \theta_{g,\max} \le \sup_{\omega\in\Omega} U_{g,\omega}. \tag{2}$$

Equation (2) connects the utility-specific bands to the lower and upper population envelopes. When a valid approximation radius ε_u is supplied, Equation (1) extends the bands to the corresponding true utilities; otherwise, we target the basis-projected utility effects.

Definition 1 (Target-group four-way audit)**.** For a frozen target group R_g and nonempty Ω, the audit returns:

(i) *robustly favor Treatment 1* if $\inf_{\omega\in\Omega} L_{g,\omega} > 0$;

(ii) *robustly favor Treatment 0* if $\sup_{\omega\in\Omega} U_{g,\omega} < 0$;

(iii) *certified preference sensitivity* if $\inf_{\omega\in\Omega} U_{g,\omega} < 0$ and $\sup_{\omega\in\Omega} L_{g,\omega} > 0$;

(iv) *evidence uncertainty* otherwise.

Accordingly, the third label identifies a statistically supported value-dependent reversal and motivates preference refinement, whereas the fourth records an unresolved sign pattern.

2.4 Modular refinement by external preference constraints

When the initial audit reveals value dependence, external preference evidence can refine the admissible class without changing the fitted causal model. Let $\Omega_0 = \Omega$. We encode information supplied at stage t as a transparent constraint set $\mathcal{C}_t$, giving

$$\Omega_{t+1} = \Omega_t \cap \mathcal{C}_t, \qquad \Omega_0 \supseteq \Omega_1 \supseteq \cdots. \tag{3}$$

Constraints may come from validated instruments, choice studies, clinician–patient conversations, published evidence, or advisory panels. We map each source to a prespecified constraint, update the envelope while reusing the fitted nuisance and basis models, and report an empty intersection explicitly. Under valid nested refinement, the envelope extrema tighten monotonically by Proposition 4.

3 Guarded Utility-Aware Orthogonal Learning

3.1 Orthogonal basis scores and a four-role sample split

Let $m_{a,j}(x) = \mathbb{E}(B_j \mid A = a, X = x)$ and $e(x) = \mathbb{P}(A = 1 \mid X = x)$. For basis coordinate j, we use the AIPW score

$$\psi_j(O;\eta) = m_{1,j}(X) - m_{0,j}(X) + \frac{A\{B_j - m_{1,j}(X)\}}{e(X)} - \frac{(1-A)\{B_j - m_{0,j}(X)\}}{1-e(X)}. \tag{4}$$

For full-pipeline studies, we use these scores within an outcome-independent four-role partition: construction $\mathcal{I}_c$, assembly $\mathcal{I}_s$, guarding $\mathcal{I}_v$, and inference $\mathcal{I}_{\inf}$. We cross-fit nuisances on $\mathcal{I}_c$, predict on $\mathcal{I}_s$ and $\mathcal{I}_v$, and use $\mathcal{I}_s$ to freeze the coordinate-adaptive candidate. We analyze the independent $\mathcal{I}_{\inf}$ only after selection, separating adaptation from confirmation. A prespecified trial-wide target may instead use full-trial cross-fitting because its family is fixed before outcome analysis.

Using these role-specific scores, we construct the vector library

$$\mathcal{H} = \{g^0, g^C, g^{\text{coord}}, g^{\text{shared}}, g^{\text{private}}\}. \tag{5}$$

Here $g^0 \equiv 0$, g^C is the construction-only constant effect vector, and the shared and private candidates use, respectively, a common multi-output partition and separate outcome learners. On $\mathcal{I}_s$, we assemble g^{coord} by choosing among zero, constant, shared, and private for each coordinate, then freeze it before evaluation on $\mathcal{I}_v$. The resulting library spans fully shared, fully private, mixed, constant, and null structure without patient–utility pseudo-replication.

3.2 The utility geometry is the selection target

To align the choice among these structures with the downstream audit, let Q be a prespecified distribution over admissible coefficient vectors. For a finite grid $\{c_\ell\}_{\ell=1}^L$ with weights q_ℓ, define

$$G_\Omega = \mathbb{E}_{\omega\sim Q}\{c(\omega)c(\omega)^\top\} = \sum_{\ell=1}^{L} q_\ell c_\ell c_\ell^\top, \qquad \sum_\ell q_\ell = 1.$$

For a vector function f, write $\|f\|_{G,2}^2 = \mathbb{E}\{f(X)^\top G_\Omega f(X)\}$. The selection target is exactly the integrated error of the *basis-induced* utility surfaces:

$$\mathcal{R}_\Omega^B(g) = \|g - \boldsymbol{\tau}_B\|_{G,2}^2 = \mathbb{E}_{X,\omega\sim Q}\left[\{c(\omega)^\top g(X) - \tau_\omega^B(X)\}^2\right].$$

This identity incorporates both coordinate covariance and utility-class geometry into the average utility-surface error under Q. For a finite grid with $q_{\min} = \min_\ell q_\ell > 0$,

$$\max_{\ell\le L} \|c_\ell^\top(g - \boldsymbol{\tau}_B)\|_{L_2}^2 \le q_{\min}^{-1}\|g - \boldsymbol{\tau}_B\|_{G,2}^2, \tag{6}$$

Equation (6) links this integrated criterion to the largest row-specific prediction error. As an alternative to a probability weighting Q, the finite-grid minimax loss $\max_\ell\{c_\ell^\top r_i\}^2$ provides a worst-row sensitivity analysis. All reported experiments use uniform weights fixed before fitting.

With this geometry fixed, we evaluate each complete candidate on $\mathcal{I}_v$. Patient i contributes one vector residual $r_i(g) = \widehat{\boldsymbol{\psi}}_i^v - g(X_i)$ and one loss

$$\ell_i(g) = r_i(g)^\top G_\Omega r_i(g), \qquad \widehat{Q}_\Omega(g) = n_v^{-1}\sum_{i\in\mathcal{I}_v}\ell_i(g). \tag{7}$$

Each patient contributes a single quadratic loss based on the complete residual vector, preserving cross-coordinate dependence during selection.

3.3 Shared-anchored vector selection

These losses support guarded selection between the anchor library $\mathcal{S} = \{g^0, g^C, g^{\text{shared}}\}$ and adaptive departures $\mathcal{D} = \{g^{\text{coord}}, g^{\text{private}}\}$. For $h \in \mathcal{D}$ and $s \in \mathcal{S}$, let $D_i^{s,h} = \ell_i(s) - \ell_i(h)$, with mean $\bar{D}^{s,h}$ and standard error $\widehat{\text{se}}^{s,h}$. Across the $N = |\mathcal{S}||\mathcal{D}| = 6$ comparisons, the primary one-sided guard is

$$\text{LCB}_{\text{T}}^{s,h} = \bar{D}^{s,h} - t_{n_v-1,\,1-\delta/N}\widehat{\text{se}}^{s,h}. \tag{8}$$

A departure is eligible only if its lower bound is positive against *every* anchor; among eligible departures and the empirically best anchor, we then minimize ordinary $\widehat{Q}_\Omega$. As a finite-sample bounded-loss option, we declare cap B before validation, replace ℓ_i by $\ell_i^B = \min(\ell_i, B)$ in the paired differences, let $\widehat{V}_B^{s,h}$ be the sample variance of $D_i^{B,s,h}$ with denominator $n_v - 1$, and use

$$\text{LCB}_{\text{EB}}^{s,h} = \bar{D}_B^{s,h} - \sqrt{2\widehat{V}_B^{s,h}\log(2N/\delta)/n_v} - \frac{14B\log(2N/\delta)}{3(n_v-1)}. \tag{9}$$

Under this option, the cap determines admission and the uncapped loss selects among retained candidates. Here, "guarded" means anchor-relative admission; the oracle deviation term in Section 4 represents finite-library selection among anchors.

Once the vector is frozen, its exploratory fitted surface for each utility is the linear post-fit query

$$\widehat{\tau}_\omega(x) = c(\omega)^\top\widehat{\boldsymbol{\tau}_B}(x),$$

Preference constraints can refine this surface without refitting the selected learner; confirmatory envelopes instead use the cross-fitted basis scores below.

4 GUARANTEES FOR SELECTION AND TARGET-GROUP AUDITS

We first connect the learning rule to its causal target and then construct direct simultaneous inference for the downstream audit.

Orthogonal pseudo-target. Under Assumption 1, $\mathbb{E}\{\psi(O;\eta_0) \mid X = x\} = \tau_B(x)$. Conditional on independently trained nuisance estimates with $\widehat{e} \in [\varepsilon, 1-\varepsilon]$, coordinate j has exact conditional score bias

$$R_{\eta,j}(x) = \{\widehat{e}(x) - e(x)\}\left\{\frac{\widehat{m}_{1,j}(x) - m_{1,j}(x)}{\widehat{e}(x)} + \frac{\widehat{m}_{0,j}(x) - m_{0,j}(x)}{1-\widehat{e}(x)}\right\}. \tag{10}$$

Its product form vanishes if the propensity is correct or both outcome regressions are correct. Appendix D gives the derivation in Theorem 3.

With this bias characterized, condition on $\mathcal{I}_c, \mathcal{I}_s$ so the finite vector library and G_Ω are fixed. Let $R_\eta = (R_{\eta,1}, \ldots, R_{\eta,M})^\top$ and $\mu = \tau_B + R_\eta$ be the conditional pseudo-target.

Theorem 1 (Anchor-relative utility-risk bound)**.** *Suppose the guard observations are independent and, with probability at least $1-\delta$, $|\widehat{Q}_\Omega(g) - Q_\Omega(g)| \le \epsilon_v$ simultaneously for $g \in \mathcal{H}$, where Q_Ω is fresh-draw pseudo-outcome risk. Because the empirical best anchor remains available, the selected vector satisfies*

$$\|\widehat{g} - \tau_B\|_{G,2} \le \min_{s\in\mathcal{S}} \|s - \tau_B\|_{G,2} + 2\|R_\eta\|_{G,2} + \sqrt{2\epsilon_v}.$$

The bound controls the selected estimator relative to the best zero, constant, or shared anchor, with explicit contributions from nuisance estimation and validation error.

To relate this result to the full library, define the retained set $\widehat{\mathcal{A}} = \mathcal{S} \cup \{h \in \mathcal{D} : \min_s \mathrm{LCB}^{s,h} > 0\}$ and its exclusion price $\gamma_\Omega = \min_{h\in\widehat{\mathcal{A}}} \|h - \tau_B\|^2_{G,2} - \min_{h\in\mathcal{H}} \|h - \tau_B\|^2_{G,2}$. Corollary 1 gives the corresponding full-library oracle inequality and states when $\gamma_\Omega = 0$; Theorem 4 controls false departure admission for the ordinary and capped losses. Their complete statements and proofs are in Appendix D.

The same utility-aware risk directly controls the average fitted surface:

$$\mathbb{E}_{\omega\sim Q}\|c(\omega)^\top(\widehat{g} - \tau_B)\|^2_{L_2(P_X)} = \|\widehat{g} - \tau_B\|^2_{G,2}.$$

For a uniform statement, Assumption 2 gives $\sup_{\omega\in\Omega} \|\widehat{\tau}_\omega - \tau_\omega\|_{L_2} \le C_2\|\widehat{g} - \tau_B\|_{2,2} + 2\varepsilon_u$.

4.1 DIRECT SIMULTANEOUS INFERENCE AT THE UTILITY LEVEL

The preceding guarantee concerns reusable surfaces, whereas certification targets group averages. The selected learner supports post-fit utility queries and exploratory subgroup construction; for certification, we instead compute cross-fitted AIPW basis scores within each prespecified group and project their centered influence vectors to the utility family. This preserves the group estimand independently of the conditional-effect model and lets a prespecified trial-wide audit use all cross-fitted observations.

Accordingly, for a prespecified target group R_g, define $\theta_g = \mathbb{E}\{\tau_B(X) \mid X \in R_g\}$, the projected utility effect $\theta^B_{g\ell} = c_\ell^\top \theta_g$, and the true effect $\theta_{g\ell} = \mathbb{E}\{\tau_{\omega_\ell}(X) \mid X \in R_g\}$. Proposition 1 gives $|\theta_{g\ell} - \theta^B_{g\ell}| \le 2\varepsilon_u$. Suppose an independent inference sample yields

$$\widehat{\theta}_g - \theta_g = n^{-1}\sum_{i=1}^n \phi_{ig} + o_p(n^{-1/2}). \tag{11}$$

We then project before studentizing: $\phi_{ig\ell} = c_\ell^\top \phi_{ig}$ and $\widehat{s}^2_{g\ell} = n^{-2}\sum_i \widehat{\phi}^2_{ig\ell}$. One shared multiplier draw preserves all basis, utility, and overlapping-group covariance:

$$T^* = \max_{g,\ell}\left|\frac{n^{-1}\sum_i \xi_i \widehat{\phi}_{ig\ell}}{\widehat{s}_{g\ell}}\right|. \tag{12}$$

The resulting interval uses $b_{g\ell} = 2\varepsilon_u$ for a declared true-utility misspecification bound and $b_{g\ell} = 0$ when only the basis projection is claimed: $I_{g\ell} = \left[c_\ell^\top \widehat{\theta}_g - \widehat{q}_{1-\alpha}\widehat{s}_{g\ell} - b_{g\ell},\ c_\ell^\top \widehat{\theta}_g + \widehat{q}_{1-\alpha}\widehat{s}_{g\ell} + b_{g\ell}\right]$. Thus, the default setting `bias_radius` $= 0$ reports intervals for the basis-projected utility effects. Supplying a valid value $2\varepsilon_u > 0$ expands these intervals to cover the corresponding true utility effects. Because $\widehat{s}_{g\ell}^2 = c_\ell^\top \widehat{\Sigma}_g c_\ell$, this direct construction retains contrast covariance; the older basis rectangle followed by ℓ_1 propagation is valid but generally wider.

Operational eligibility diagnostic. Before translating these intervals into a decisive label, we check minimum target-group size, fitted overlap, and concentration of the estimated influence contributions. In the reported experiments, every group must contain at least $\max\{40, \lceil 0.05 n_{\inf} \rceil\}$ observations, fitted propensities must lie in $[0.025, 0.975]$, and the maximum standardized influence concentration must not exceed 0.50. A failed check maps the result to evidence uncertainty.

To preserve the unconditional interpretation, we report coverage over all simulation repetitions, including those that fail this diagnostic.

Theorem 2 (Direct max-t validity and four-way error control)**.** *Let the basis, G target groups, and finite set of L utility rows be fixed independently of the inference outcomes, with GL fixed. Assume independent units; target-group mass bounded away from zero; overlap; uniform asymptotic linearity in* (11) *with nuisance-product remainder $o_p(n^{-1/2})$; bounded coefficient-row norms; contrast variances bounded away from zero and infinity; uniformly bounded $(2+\kappa)$ moments; consistent covariance; and vanishing maximum standardized leverage. If the number of multiplier draws diverges and $b_{g\ell} = 2\varepsilon_u$ uses a valid finite bound, then*

$$\mathbb{P}\{\theta_{g\ell} \in I_{g\ell} \text{ for all } g, \ell\} \geq 1 - \alpha - o(1). \tag{13}$$

On this event, every robust Treatment 1 or Treatment 0 audit has the declared sign for all rows in the frozen finite family, and every preference-sensitive audit contains two true opposing signs. Hence the familywise probability of a false decisive or false preference-sensitive output is at most $\alpha + o(1)$. When ε_u is unavailable, the coverage statement concerns the projected effects $\theta_{g\ell}^B$. For a polytopal utility class, calibration over all relevant vertices extends the conclusion to the full polytope; for a general utility class, the theorem applies to the reported finite grid.

4.2 Preference-set uncertainty as a second layer

The preceding theorem treats utility rows as frozen. To accommodate estimated preference sets, suppose an external sample $\mathcal{D}_p$ instead yields a random polytope $\widehat{\Omega}_\beta$ for a declared coefficient target $c_\star$, with

$$\mathbb{P}\{c_\star \in \widehat{\Omega}_\beta\} \geq 1 - \beta. \tag{14}$$

As preference information accumulates, it may produce nested regions $\widehat{\Omega}_0 \supseteq \cdots \supseteq \widehat{\Omega}_K$. We form one master family $\mathcal{V}_\star = \bigcup_k \operatorname{vert}(\widehat{\Omega}_k)$, including vertices created by later intersections, and calibrate direct max-t once over every prespecified target-group–vertex contrast. Conditional on the preference data, this causal multiplicity family remains fixed throughout downstream refinement.

Proposition 2 (Two-layer robust-sign validity)**.** *Suppose* (14) *holds and, conditional on $\mathcal{D}_p$, the causal bands cover all contrasts indexed by $\mathcal{V}_\star$ with probability at least $1-\alpha-o(1)$. This conditional coverage is justified when preference data are external to the causal outcomes or the master family is prespecified or constructed on an independent split. Then a reported robust sign for $c_\star$ is wrong with probability at most $\alpha + \beta + o(1)$. If both layers have finite-sample coverage, the remainder is absent. A certified conflict is a property of the admissible coefficient set: it establishes two coefficient vectors with opposing treatment-effect signs. The prevalence of those trade-offs among individual patients is a separate preference estimand.*

The proposition follows from linearity, which puts every polytope extremum at a vertex. On the event that the external set contains $c_\star$ and every master-vertex band covers, a unanimous strict sign holds for $c_\star$; a union bound proves the result. Independence is not required by that union bound, but it is sufficient for conditional causal calibration of a random family. Appendix G gives the construction and synthetic validation.

Finally, clustered data use cluster-summed influence vectors and require a diverging number of independent clusters and vanishing maximum cluster leverage. Adaptive target construction requires

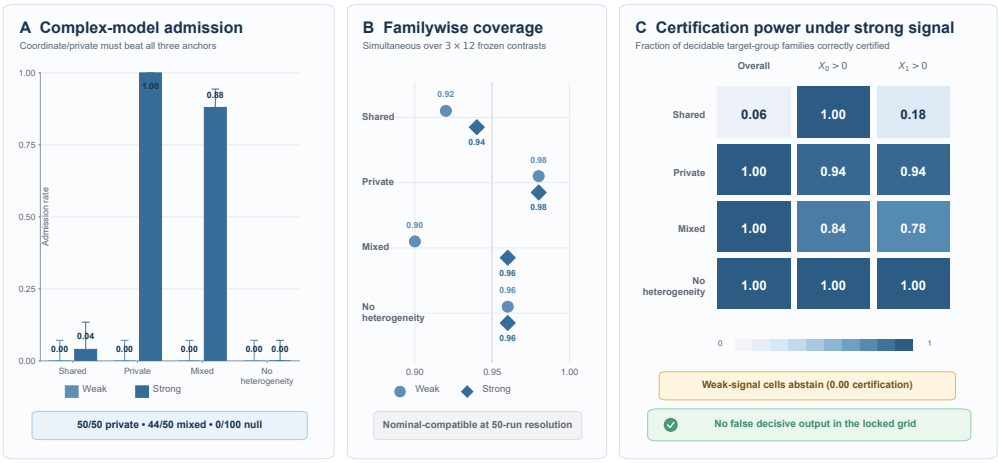


Figure 2: Prespecified full-pipeline results from raw (X, A, Y), with 50 repetitions per cell. **(A)** Complex-model admission: 50/50 strong-private, 44/50 strong-mixed, and anchor selection in all 100 no-heterogeneity runs. **(B)** Familywise coverage over the 3×12 target-group–utility contrasts. **(C)** Certification power under strong signals. Weak-signal cells return evidence uncertainty. Bars are Wilson 95% intervals; target groups may overlap.

a further independent split. In either extension, the guarantees concern target-group averages, while profile-level surfaces support exploration and subgroup construction.

5 Experiments

Prespecified evaluation protocol. We evaluate adaptation and decision power, representation, simultaneous inference, and real-data audits. The supplement adds a selector phase diagram, equal-risk admission calibration, raw- and score-level max-t stress tests, distributional scaling, and IST/ACTG analyses (Appendix F). Appendices H.1–H.4 detail the raw-data pipeline, high-K scaling, fail-closed eligibility, and tuberculosis preference audit; Appendix G validates joint causal–preference calibration and nested refinement. Synthetic comparisons share data, splits, utilities, and seeds; the artifact gives full configurations and run-level outputs.

5.1 Full-pipeline adaptation and power from raw data

Each repetition splits raw randomized-trial data (X, A, Y) into independent construction, assembly, guard, and inference roles. We cross four structural regimes with weak and strong signals at $n = 2{,}400$, $M = 8$, 12 utilities, and three target groups. Figure 2 shows this adaptation: complex departures are admitted in all 50 strong-private runs and 44 of 50 strong-mixed runs, while all 100 no-heterogeneity runs retain an anchor. Coverage ranges from 0.90 to 0.98. Strong signals yield high certification, whereas weak signals return evidence uncertainty, separating admission from certification power. Appendix Table 15 gives the complete grid and fixed-learner comparisons.

5.2 Representation trade-offs with distributional learners

Having evaluated the complete pipeline, we next isolate its representation trade-offs by comparing PROBE with fixed basis learners, a non-truncated observed-state forest, and an autoregressive chain. The distributional learners achieve lower predictive RMSE in the tested settings, and the joint-state model also fits faster for $K \leq 16$. At $K = 24$ and 32, the paired PROBE-minus-joint RMSE differences are $0.0035\ [-0.0020, 0.0092]$ and $-0.0003\ [-0.0052, 0.0055]$, while the autoregressive chain has the lowest mean RMSE.

The basis instead offers low-order attribution, smaller high-dimensional memory proxies, and rapid post-fit utility evaluation. We therefore report accuracy, computation, memory, and query cost separately; Appendix H.2 gives the high-K extension.

5.3 DIRECT MAX-t INFERENCE WITH ABSTENTION DIAGNOSTICS

Because prediction does not establish audit validity, we next evaluate simultaneous inference. Two 1,000-run references attain familywise coverage of 0.945 under known randomization and 0.949 under a correctly specified bounded observational design. Direct intervals are 23–25% narrower than basis rectangles, while pointwise intervals have familywise coverage of 0.622 and 0.694.

Stress tests then assess the diagnostic, which maps all weak-overlap and rare-group runs to evidence uncertainty. Dual nuisance misspecification passes with 0.910 coverage; coverage remains unconditional over all repetitions. A 400-run two-source simulation gives joint coverage of 0.9125, consistent with the nominal 0.10 union-error target.

5.4 RANDOMIZED-TRIAL AUDIT: IST SAME-DISEASE SENSITIVITY

We then apply the calibrated layer to the overall IST cohort (1997) ($n = 19{,}433$). Its basis combines survival, recurrent-stroke-free survival, and major bleeding, and we evaluate three analyst-specified same-disease utility rows. Their simultaneous intervals are all positive, indicating stable support for aspirin across this sensitivity family (Table 2). Appendix F.4 gives the utility construction and exploratory discovery–confirmation analysis.

Stroke anchor	Estimate	SE	Direct 95% interval
.05	.009497	.003577	[.002487, .016507]
.10	.009242	.003570	[.002245, .016238]
.20	.008718	.003561	[.001738, .015698]

Table 2: IST randomized-cohort audit over three analyst-specified same-disease utility rows. Direct max-t intervals are simultaneous across the family.

5.5 DEPENDENCE-PRESERVING PREFERENCE EVIDENCE FOR 3HP VERSUS 6H

Because IST fixes its utility rows, we next use empirical preference evidence to evaluate the external preference layer. The Uganda tuberculosis preventive-treatment DCE provides 394 complete same-respondent hierarchical-Bayes part-worth vectors over regimen burden and adverse-event attributes (Aschmann et al., 2026a;b). Our respondent-level bootstrap preserves cross-attribute dependence, while interpolation and nearest-level analyses address off-grid regimen features. With this construction, we reproduce the published choice pattern, and preserving the joint coefficient structure changes the estimated choice shares relative to a marginally permuted construction.

Preference construction	3HP (%)	6H (%)	No TPT (%)
Joint rows, linear	94.7 [93.5,95.9]	4.2 [3.3,5.3]	1.0 [0.4,1.8]
Joint rows, nearest	88.4	10.1	1.5
Marginal collapse	91.5 [90.8,92.3]	5.7 [5.1,6.3]	2.8 [2.2,3.4]

Table 3: Predicted shares from 394 complete same-respondent HB means. Joint-linear brackets are respondent-bootstrap intervals conditional on deposited means; marginal-collapse brackets are permutation ranges; nearest is mapping sensitivity.

To contextualize these preferences, randomized WHIP3TB arm-level estimates give 90.4% completion for combined 3HP and 50.5% for 6H, while the 12-month tuberculosis contrast is imprecise (Churchyard et al., 2021). Because these aggregate outcomes are not linked to DCE respondents, the configuration remains evidence-uncertain; Appendix H.4 gives the mappings and sensitivity analyses.

6 DISCUSSION AND LIMITATIONS

PROBE combines guarded learning with simultaneous inference for frozen target-group–utility families; Supplementary Discussion gives its scope, diagnostics, preference assumptions, and real-data limits.

REFERENCES

Virginia Aglietti, Theodoros Damoulas, Mauricio Álvarez, and Javier González. Multi-task causal learning with gaussian processes. In *Advances in Neural Information Processing Systems*, volume 33, pp. 6293–6304, 2020.

Ahmed M. Alaa and Mihaela van der Schaar. Bayesian inference of individualized treatment effects using multi-task gaussian processes. In *Advances in Neural Information Processing Systems*, volume 30, pp. 3424–3432, 2017.

Hélène E. Aschmann, Allan Musinguzi, Jillian L. Kadota, Catherine Namale, Juliet Kakeeto, Jane Nakimuli, Lydia Akello, Fred Welishe, Anne Nakitende, Christopher Berger, David W. Dowdy, Adithya Cattamanchi, Fred C. Semitala, and Andrew D. Kerkhoff. Predicted choice and acceptability of regimens for tuberculosis preventive treatment among people living with HIV in uganda: A discrete choice experiment. *PLOS Global Public Health*, 6(9):e0005630, 2026a. doi: 10.1371/journal.pgph.0005630.

Hélène E. Aschmann, Allan Musinguzi, Jillian L. Kadota, et al. Predicted choice and acceptability of regimens for tuberculosis preventive treatment among people living with HIV in uganda—a discrete choice experiment [dataset]. Dryad, 2026b. URL https://doi.org/10.5061/dryad.31zcrjf34. Dataset; CC0; version published 20 August 2026.

Giorgos Bakoyannis. Estimating optimal individualized treatment rules with multistate processes. *Biometrics*, 79(4):2830–2842, 2023. doi: 10.1111/biom.13864.

Henk Broekhuizen, Maarten J. IJzerman, A. Brett Hauber, and Catharina G. M. Groothuis-Oudshoorn. Weighing clinical evidence using patient preferences: An application of probabilistic multi-criteria decision analysis. *PharmacoEconomics*, 35(3):259–269, 2017. doi: 10.1007/s40273-016-0467-z.

Emily L. Butler, Eric B. Laber, Sonia M. Davis, and Michael R. Kosorok. Incorporating patient preferences into estimation of optimal individualized treatment rules. *Biometrics*, 74(1):18–26, 2018. doi: 10.1111/biom.12743.

Yuan Chen, Donglin Zeng, Tianchen Xu, and Yuanjia Wang. Representation learning for integrating multi-domain outcomes to optimize individualized treatment. In *Advances in Neural Information Processing Systems*, volume 33, pp. 17976–17986, 2020.

Victor Chernozhukov, Denis Chetverikov, and Kengo Kato. Gaussian approximations and multiplier bootstrap for maxima of sums of high-dimensional random vectors. *The Annals of Statistics*, 41 (6):2786–2819, 2013. doi: 10.1214/13-AOS1161.

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.

Gavin Churchyard, Vicky Cárdenas, Violet Chihota, Kathy Mngadi, Modulakgotla Sebe, William Brumskine, Neil Martinson, Getnet Yimer, Shu-Hua Wang, Alberto L. Garcia-Basteiro, et al. Annual tuberculosis preventive therapy for persons with HIV infection: A randomized trial. *Annals of Internal Medicine*, 174(10):1367–1376, 2021. doi: 10.7326/M20-7577.

Gloria Cordoba, Lisa Schwartz, Steven Woloshin, Harold Bae, and Peter C. Gøtzsche. Definition, reporting, and interpretation of composite outcomes in clinical trials: systematic review. *BMJ*, 341:c3920, 2010.

Anh Nguyen Duc and Marcel Wolbers. Weighted analysis of composite endpoints with simultaneous inference for flexible weight constraints. *Statistics in Medicine*, 36(3):442–454, 2017. doi: 10.1002/sim.7147.

Nerida Firth, Alison Pearce, Oyelola Adegboye, Kathryn S. Hayward, Julie Bernhardt, Robin Ray, and Ruth N. Barker. Stroke survivors' preferences for recovery-promoting drugs: Development and piloting of a discrete choice experiment. *Pilot and Feasibility Studies*, 12(1):38, 2026. doi: 10.1186/s40814-025-01760-z.

Nick Freemantle, Melanie Calvert, John Wood, Joanne Eastaugh, and Carl Griffin. Composite outcomes in randomized trials: greater precision but with greater uncertainty? *JAMA*, 289(19): 2554–2559, 2003.

Scott M. Hammer, David A. Katzenstein, Michael D. Hughes, Holly Gundacker, Robert T. Schooley, Richard H. Haubrich, W. Keith Henry, Michael M. Lederman, John P. Phair, Manette Niu, Martin S. Hirsch, and Thomas C. Merigan. A trial comparing nucleoside monotherapy with combination therapy in HIV-infected adults with CD4 cell counts from 200 to 500 per cubic millimeter. *The New England Journal of Medicine*, 335(15):1081–1090, 1996. doi: 10.1056/NEJM199610103351501.

Harold Hotelling. The generalization of student's ratio. *The Annals of Mathematical Statistics*, 2(3): 360–378, 1931. doi: 10.1214/aoms/1177732979.

International Stroke Trial Collaborative Group. The international stroke trial (IST): A randomised trial of aspirin, subcutaneous heparin, both, or neither among 19,435 patients with acute ischaemic stroke. *The Lancet*, 349(9065):1569–1581, 1997. doi: 10.1016/S0140-6736(97)04011-7.

Junguang Jiang, Baixu Chen, Junwei Pan, Ximei Wang, Dapeng Liu, Jie Jiang, and Mingsheng Long. Forkmerge: Mitigating negative transfer in auxiliary-task learning. In *Advances in Neural Information Processing Systems*, volume 36, pp. 30367–30389, 2023a. doi: 10.52202/075280-1322.

Ziyang Jiang, Zhuoran Hou, Yiling Liu, Yiman Ren, Keyu Li, and David Carlson. Estimating causal effects using a multi-task deep ensemble. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 15023–15040, 2023b.

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

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

Giwoong Lee, Eunho Yang, and Sung Ju Hwang. Asymmetric multi-task learning based on task relatedness and loss. In *Proceedings of the 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 230–238, 2016.

Hae Beom Lee, Eunho Yang, and Sung Ju Hwang. Deep asymmetric multi-task feature learning. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 2956–2964, 2018.

Daniel J. Luckett, Eric B. Laber, Siyeon Kim, and Michael R. Kosorok. Estimation and optimization of composite outcomes. *Journal of Machine Learning Research*, 22(167):1–40, 2021.

Andreas Maurer, Massimiliano Pontil, and Bernardino Romera-Paredes. The benefit of multitask representation learning. *Journal of Machine Learning Research*, 17(81):1–32, 2016.

Stephen P. McKenna and Alice Heaney. Composite outcome measurement in clinical research: the triumph of illusion over reality? *Journal of Medical Economics*, 23(10):1196–1204, 2020.

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

Junhyung Park, Uri Shalit, Bernhard Schölkopf, and Krikamol Muandet. Conditional distributional treatment effect with kernel conditional mean embeddings and u-statistic regression. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8401–8412, 2021.

Dovid Parnas, Mathieu Even, Julie Josse, and Uri Shalit. Preference-based conditional treatment effects and policy learning. In *Proceedings of the 29th International Conference on Artificial Intelligence and Statistics*, volume 300 of *Proceedings of Machine Learning Research*, pp. 226–234. PMLR, 2026. URL https://proceedings.mlr.press/v300/parnas26a.html.

Gabriela Rusin, Jadwiga Kubica, Krzysztof Piotr Malinowski, Agnieszka Słowik, and Anetta Undas. Factors affecting self-reported bleeding acceptance in acute ischemic stroke survivors on various types of antithrombotic therapy. *Journal of Stroke and Cerebrovascular Diseases*, 32(1):106894, 2023. doi: 10.1016/j.jstrokecerebrovasdis.2022.106894.

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, 2017.

Claudia Shi, David M. Blei, and Victor Veitch. Adapting neural networks for the estimation of treatment effects. In *Advances in Neural Information Processing Systems*, volume 32, pp. 2507–2517, 2019.

Karsten Bruins Slot and Eivind Berge. Thrombolytic treatment for stroke: Patient preferences for treatment, information, and involvement. *Journal of Stroke and Cerebrovascular Diseases*, 18(1): 17–22, 2009. doi: 10.1016/j.jstrokecerebrovasdis.2008.06.009.

Malik Tiomoko, Romain Couillet, and Frédéric Pascal. Pca-based multi-task learning: a random matrix approach. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 34280–34300, 2023.

Mark J. van der Laan, Eric C. Polley, and Alan E. Hubbard. Super learner. *Statistical Applications in Genetics and Molecular Biology*, 6(1):Article 25, 2007.

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

Liming Zhang, Ning Wang, Chunli Xian, Meng Fan, and Jianghua Xu. Preferences for advance care planning services among patients with recurrent stroke: A discrete choice experiment. *BMC Health Services Research*, 26:662, 2026. doi: 10.1186/s12913-026-14437-x.

Yu Zhang and Qiang Yang. A survey on multi-task learning. *IEEE Transactions on Knowledge and Data Engineering*, 34(12):5586–5609, 2022.

Joshua P. Zitovsky, Yating Zou, Leslie Wilson, and Michael R. Kosorok. Latent utility q-learning for preference-adaptive dynamic treatment regimes. arXiv:2307.12022, 2026. Version 3; first posted in 2023.

AI USE STATEMENT

Generative AI tools were used for language editing, LaTeX formatting, reference checking, and routine code assistance. The authors reviewed the mathematical arguments, executed and validated the analyses, and take responsibility for all claims, code, results, and accompanying artifacts.

ETHICS STATEMENT

This framework is intended for research auditing and group-level decision support. Applications should assess the causal assumptions, utility definitions, uncertainty calibration, data governance, and outcome coding appropriate to the study. Evidence-uncertain outputs call for additional review or information before a treatment conclusion is made.

REPRODUCIBILITY STATEMENT

We specify the targets, both guards, the estimator, and the proof assumptions. We link every reported value to archived raw artifacts and manifests and include the checksum-verified public DCE release together with participant-free derived summaries for the restricted datasets.

A Related Work

Composite outcomes and value trade-offs. Composite outcomes can increase event rates, but their interpretation may be dominated by frequent, less important components or obscured when component effects differ in direction (Freemantle et al., 2003; Cordoba et al., 2010; McKenna & Heaney, 2020). These limitations have motivated estimation and optimization under fixed or covariate-dependent latent utilities learned from clinician choices (Luckett et al., 2021). Related weighted composite-endpoint analyses provide simultaneous intervals over constrained weight families, and preference-weighted multistate rules provide simultaneous inference over finite preference sets (Duc & Wolbers, 2017; Bakoyannis, 2023). At the policy level, preference-informed ITRs and DTRs combine multivariate outcomes with elicited or latent subject-level trade-offs (Butler et al., 2018; Zitovsky et al., 2026), while probabilistic MCDA propagates uncertainty in clinical evidence and patient preferences (Broekhuizen et al., 2017). We instead freeze a target group and audit whether its downstream conclusion is stable over an admissible utility set, using guarded reusable basis learning and a four-way abstaining output rather than estimating an individual utility or policy.

Heterogeneous and distributional treatment effects. Within this broader problem, meta-learners, causal forests, balanced-representation methods, targeted-regularization networks, and Bayesian models estimate heterogeneous effects under different structural assumptions (Künzel et al., 2019; Wager & Athey, 2018; Shalit et al., 2017; Shi et al., 2019). Orthogonal and doubly robust constructions complement these learners by reducing first-order sensitivity to nuisance estimation (Chernozhukov et al., 2018; Nie & Wager, 2021; Kennedy, 2023). To evaluate many downstream functionals, scalar conditional distribution estimators offer one route (Park et al., 2021); we use a lower-complexity structured basis when the relevant utilities depend on prespecified component and joint-event transformations. Preference-based conditional effects can also use nonseparable comparisons (Parnas et al., 2026); here we study separable utility contrasts that retain transparent basis attribution.

Multi-task learning and guarded selection. Multi-task representation learning can improve sample efficiency when tasks share structure (Maurer et al., 2016; Zhang & Yang, 2022). Multi-task causal models share information across potential-outcome surfaces, intervention functions, or outcome domains using Gaussian processes, ensembles, and shared representations (Alaa & van der Schaar, 2017; Aglietti et al., 2020; Chen et al., 2020; Jiang et al., 2023b); among these, Chen et al. (2020) directly integrate multiple outcome domains. Negative transfer motivates task-selective structures and validation-driven adaptation (Lee et al., 2016; 2018; Jiang et al., 2023a; Tiomoko et al., 2023). Our guard treats coordinate-wise adaptation as one candidate in a finite vector library and evaluates the library through a patient-level quadratic loss induced by the utility geometry. The construction builds on cross-validated library aggregation and selection (van der Laan et al., 2007) while preserving the multi-outcome structure.

B Additional Representations and Utility Constructions

B.1 Binary-outcome basis completeness

For binary components, inclusion–exclusion already shows why joint-event coordinates matter: the indicator of any component event satisfies

$$\mathbf{1}\left(\max_k Y_k = 1\right) = \sum_{\emptyset \neq S \subseteq [K]} (-1)^{|S|+1} \prod_{k \in S} Y_k.$$

Proposition 3 (Multilinear representation for binary outcomes)**.** *For $\boldsymbol{y} \in \{0,1\}^K$, every function $u : \{0,1\}^K \to \mathbb{R}$ has the unique expansion*

$$u(\boldsymbol{y}) = \sum_{S \subseteq [K]} c_S \prod_{k \in S} y_k, \qquad c_S = \sum_{T \subseteq S} (-1)^{|S|-|T|} u(\mathbf{1}_T).$$

Proof of Proposition 3. Let $[K] = \{1, \ldots, K\}$. For each $S \subseteq [K]$, define the monomial

$$m_S(y) = \prod_{k \in S} y_k,$$

with the convention $m_\varnothing(y) = 1$. Every binary vector $y \in \{0,1\}^K$ can be written uniquely as $\mathbf{1}_T$ for some $T \subseteq [K]$, where

$$(\mathbf{1}_T)_k = \begin{cases} 1, & k \in T, \\ 0, & k \notin T. \end{cases}$$

At such a vector,

$$m_S(\mathbf{1}_T) = \prod_{k \in S} (\mathbf{1}_T)_k = \begin{cases} 1, & S \subseteq T, \\ 0, & S \not\subseteq T. \end{cases}$$

Consequently, a multilinear representation

$$u(y) = \sum_{S \subseteq [K]} c_S m_S(y)$$

must satisfy

$$u(\mathbf{1}_T) = \sum_{S \subseteq T} c_S \qquad \text{for every } T \subseteq [K].$$

Define

$$c_S = \sum_{R \subseteq S} (-1)^{|S|-|R|} u(\mathbf{1}_R).$$

We first verify that these coefficients reproduce u. For any $T \subseteq [K]$,

$$\begin{aligned} \sum_{S \subseteq T} c_S &= \sum_{S \subseteq T} \sum_{R \subseteq S} (-1)^{|S|-|R|} u(\mathbf{1}_R) \\ &= \sum_{R \subseteq T} u(\mathbf{1}_R) \sum_{S:\, R \subseteq S \subseteq T} (-1)^{|S|-|R|}. \end{aligned}$$

For fixed $R \subseteq T$, write $S = R \cup V$, where $V \subseteq T \setminus R$. Then

$$\begin{aligned} \sum_{S:\, R \subseteq S \subseteq T} (-1)^{|S|-|R|} &= \sum_{V \subseteq T \setminus R} (-1)^{|V|} \\ &= (1-1)^{|T \setminus R|}. \end{aligned}$$

This quantity is zero unless $R = T$, in which case it equals one. Therefore,

$$\sum_{S \subseteq T} c_S = u(\mathbf{1}_T).$$

Since every $y \in \{0,1\}^K$ equals $\mathbf{1}_T$ for a unique T, the proposed coefficients give the required representation on the entire binary hypercube.

To establish uniqueness, suppose that both $\{c_S\}$ and $\{\widetilde{c}_S\}$ represent u. Let $d_S = c_S - \widetilde{c}_S$. Then

$$\sum_{S \subseteq T} d_S = 0 \qquad \text{for every } T \subseteq [K].$$

Applying the same inversion formula to the identically zero function gives

$$d_S = \sum_{R \subseteq S} (-1)^{|S|-|R|} \left\{ \sum_{V \subseteq R} d_V \right\} = 0 \qquad \text{for every } S.$$

Thus $c_S = \widetilde{c}_S$ for every S, proving uniqueness. □

B.2 Optional shared–private neural candidate

As a representation-learning ablation, the shared vector candidate in (5) can be replaced coordinate by coordinate by

$$\widehat{g}_j^{SP}(x) = \rho_j \widehat{g}_j^{\text{shared}}(x) + (1-\rho_j)\widehat{g}_j^{\text{private}}(x), \qquad 0 \le \rho_j \le 1.$$

The implemented ablation uses shared and task-specific branches, coordinate-specific validation-selected mixing, Huber validation loss, pseudo-outcome standardization, and early stopping. The core experiments use the shared candidate in Section 3; this neural version is reported as a representation-learning ablation.

Weighted benefit–harm utility. The same basis framework also supports three clinically interpretable utility constructions. For an efficacy indicator Y_E and adverse-event indicators Y_T, Y_F, a higher-is-better utility may be written

$$u_\omega(\boldsymbol{Y}) = \omega_E Y_E - \omega_T Y_T - \omega_F Y_F, \qquad \omega_E, \omega_T, \omega_F \geq 0.$$

Ratio, ordering, or normalization constraints define a clinically interpretable region.

Interaction-sensitive utility. To represent nonadditive burden from simultaneous toxicity and functional decline, we include

$$u_\omega(\boldsymbol{Y}) = \omega_E Y_E - \omega_T Y_T - \omega_F Y_F - \omega_{TF} Y_T Y_F.$$

The basis coordinate $Y_T Y_F$ targets a joint-event effect that component marginals alone cannot identify.

External preference constraints. Preference information can then restrict these utility families. For example, a validated trade-off statement may imply a half-space such as

$$\omega_E \delta_E - \omega_T \delta_T \geq 0.$$

Intersecting such a constraint with the current admissible set implements (3). The external instrument or elicitation process supplies the clinical interpretation of this constraint.

B.3 Monotonicity of external preference refinement

Proposition 4 (Monotone post-fit refinement)**.** *If $\Omega_{t+1} \subseteq \Omega_t$, then*

$$\inf_{\omega \in \Omega_{t+1}} L_{g,\omega} \geq \inf_{\omega \in \Omega_t} L_{g,\omega}, \qquad \sup_{\omega \in \Omega_{t+1}} U_{g,\omega} \leq \sup_{\omega \in \Omega_t} U_{g,\omega}.$$

If a clinically relevant true utility index remains in every refined set, a robust Treatment 1 or Treatment 0 conclusion remains valid for that utility. Certified preference sensitivity is a set-level statement about two retained utilities with opposing signs. Individual-level interpretation requires individual preference data. Both statements condition on valid external constraints.

Proof of Proposition 4. Fix a target group R_g, and suppose that

$$\Omega_{t+1} \subseteq \Omega_t$$

with both sets nonempty. Because the function $\omega \mapsto L_{g,\omega}$ is unchanged when the admissible set is refined, restricting its domain cannot decrease its infimum. Formally, for every $\omega \in \Omega_{t+1}$,

$$L_{g,\omega} \geq \inf_{\widetilde{\omega} \in \Omega_t} L_{g,\widetilde{\omega}}.$$

Taking the infimum over $\omega \in \Omega_{t+1}$ on the left-hand side gives

$$\inf_{\omega \in \Omega_{t+1}} L_{g,\omega} \geq \inf_{\omega \in \Omega_t} L_{g,\omega}.$$

Similarly, for every $\omega \in \Omega_{t+1}$,

$$U_{g,\omega} \leq \sup_{\widetilde{\omega} \in \Omega_t} U_{g,\widetilde{\omega}},$$

and therefore

$$\sup_{\omega \in \Omega_{t+1}} U_{g,\omega} \leq \sup_{\omega \in \Omega_t} U_{g,\omega}.$$

Now let

$$\mathcal{E}_t = \{L_{g,\omega} \leq \theta_{g,\omega} \leq U_{g,\omega} \text{ for every } \omega \in \Omega_t\}$$

denote the simultaneous coverage event for the original set. Since $\Omega_{t+1} \subseteq \Omega_t$, the event $\mathcal{E}_t$ also implies simultaneous coverage for every utility retained in Ω_{t+1}.

For example, if the refined audit satisfies

$$\inf_{\omega \in \Omega_{t+1}} L_{g,\omega} > 0,$$

then, on $\mathcal{E}_t$,

$$\theta_{g,\omega} \geq L_{g,\omega} > 0 \qquad \text{for every } \omega \in \Omega_{t+1}.$$

The analogous conclusion holds for a uniformly negative upper envelope. The conclusion applies to clinically relevant utilities retained by the external constraint set; coverage of that containment is supplied by the preference procedure. □

C Complete Guarded Vector Algorithm

Algorithm 1 Guarded utility-aware orthogonal learning and target-group auditing

Require: Independent units; fixed basis b; coefficient grid $\{c_\ell, q_\ell\}$; candidate library; error levels δ, α; frozen target groups.

1: Draw the construction/outer-validation partition, then outcome-independently partition outer validation into assembly $\mathcal{I}_s$ and guard $\mathcal{I}_v$. Reserve an independent inference sample $\mathcal{I}_{\text{inf}}$ when targets or reported families are learned; prespecified targets may use cross-fitting on the full sample.
2: Cross-fit nuisance functions on $\mathcal{I}_c$; construction-only nuisance models separately score $\mathcal{I}_s$ and $\mathcal{I}_v$ by (4). Cross-fit scores on $\mathcal{I}_{\text{inf}}$ (or all observations for a prespecified target) for inference.
3: Fit zero, constant, shared, and private vector candidates using $\mathcal{I}_c$ only.
4: On $\mathcal{I}_s$, freeze the coordinate-adaptive vector before evaluation on $\mathcal{I}_v$.
5: Form $G_\Omega = \sum_\ell q_\ell c_\ell c_\ell^\top$ and one quadratic vector loss per guard patient by (7).
6: Treat zero, constant, and shared as anchors. Admit a coordinate or private departure only if (8) (primary) or (9) (bounded-loss option) is positive against all three anchors.
7: Select the lowest uncapped-risk candidate among the empirical best anchor and admitted departures; retain its construction fit.
8: For each frozen group, aggregate inference scores, center observation-level basis influence vectors, project them to every utility row, and calibrate (12) with shared multipliers.
9: Before a decisive label, apply the prespecified group-size, fitted-overlap, and influence-concentration diagnostic. Map failed eligibility checks to evidence uncertainty and report unconditional coverage.
10: Apply Definition 1; a group with any inactive zero-variance contrast returns evidence uncertainty. Refining Ω reuses the same basis estimates and influence vectors.

Our implementation preserves this separation between selection and inference. The vector selector accepts a map of frozen validation predictions and records all risks, paired lower bounds, eligibility decisions, sample sizes, utility geometry, and the selected vector. The direct inference routine then records the raw and pointwise-floored multiplier critical values, active and family contrast counts, covariance rank and condition number, maximum standardized leverage, multiplier type, seed, and draw count. We also record each eligibility check and its outcome without changing interval endpoints or coverage calculations. An explicit nonzero utility-basis misspecification allowance enters through *bias_radius*. We next establish the basis, selection, and inference results used by the algorithm.

D Proofs

D.1 Utility-basis results

Proof of Proposition 1. Fix x and $\omega \in \Omega$. By Assumption 2, for each treatment level $a \in \{0, 1\}$,

$$u_\omega\{\boldsymbol{Y}(a)\} = c_0(\omega) + c(\omega)^\top b\{\boldsymbol{Y}(a)\} + r_\omega\{\boldsymbol{Y}(a)\}.$$

Taking the difference between the two potential-treatment utilities gives

$$\begin{aligned} u_\omega\{\boldsymbol{Y}(1)\} - u_\omega\{\boldsymbol{Y}(0)\} = c(\omega)^\top \left[b\{\boldsymbol{Y}(1)\} - b\{\boldsymbol{Y}(0)\}\right] \\ + r_\omega\{\boldsymbol{Y}(1)\} - r_\omega\{\boldsymbol{Y}(0)\}, \end{aligned}$$

because the utility-specific intercept $c_0(\omega)$ cancels.

Taking conditional expectations given $X = x$ yields

$$\begin{aligned} \tau_\omega(x) = c(\omega)^\top \mathbb{E}\left[b\{\boldsymbol{Y}(1)\} - b\{\boldsymbol{Y}(0)\} \mid X = x\right] \\ + \mathbb{E}\left[r_\omega\{\boldsymbol{Y}(1)\} - r_\omega\{\boldsymbol{Y}(0)\} \mid X = x\right]. \end{aligned}$$

By the definition of $\boldsymbol{\tau}_B(x)$,

$$\mathbb{E}\left[b\{\boldsymbol{Y}(1)\} - b\{\boldsymbol{Y}(0)\} \mid X = x\right] = \boldsymbol{\tau}_B(x),$$

and therefore

$$\tau_\omega(x) - \tau_\omega^B(x) = \mathbb{E}\left[r_\omega\{\boldsymbol{Y}(1)\} - r_\omega\{\boldsymbol{Y}(0)\} \mid X = x\right].$$

Using Jensen's inequality for the absolute value, followed by the triangle inequality, gives

$$\begin{aligned}\left|\tau_\omega(x) - \tau_\omega^B(x)\right| &\le \mathbb{E}\left[\left|r_\omega\{\boldsymbol{Y}(1)\} - r_\omega\{\boldsymbol{Y}(0)\}\right| \mid X = x\right] \\ &\le \mathbb{E}\left[\left|r_\omega\{\boldsymbol{Y}(1)\}\right| + \left|r_\omega\{\boldsymbol{Y}(0)\}\right| \mid X = x\right].\end{aligned}$$

The uniform approximation condition

$$\sup_{\omega\in\Omega,\, y} |r_\omega(y)| \le \varepsilon_u$$

implies that each of the two terms inside the last conditional expectation is at most ε_u. Hence

$$\left|\tau_\omega(x) - \tau_\omega^B(x)\right| \le 2\varepsilon_u.$$

Because x and ω were arbitrary, the result holds uniformly over the stated target set. □

D.2 Orthogonality and guarded selection

To analyze the guard after assembly, let $\mathcal{F}_{cs}$ denote the σ-field generated by the construction and assembly roles. We condition on $\mathcal{F}_{cs}$ whenever analyzing the guard sample. Under this conditioning, the nuisance estimates, the finite candidate library, the assembled coordinate candidate, and G_Ω are fixed, while the guard observations remain independent fresh observations.

Theorem 3 (Exact orthogonal remainder)**.** *Under Assumption 1,*

$$\mathbb{E}\{\boldsymbol{\psi}(O;\eta_0) \mid X = x\} = \boldsymbol{\tau}_B(x).$$

Conditional on independently trained nuisance estimates with $\widehat{e} \in [\varepsilon, 1-\varepsilon]$, the conditional score bias equals (10) *coordinatewise. It vanishes if the propensity is correct or if both outcome regressions are correct.*

Corollary 1 (Full-library oracle with an explicit exclusion price)**.** *On the event of Theorem 1,*

$$\|\widehat{g} - \boldsymbol{\tau}_B\|_{G,2} \le \left\{\min_{h\in\mathcal{H}} \|h - \boldsymbol{\tau}_B\|_{G,2}^2 + \gamma_\Omega\right\}^{1/2} + 2\|R_\eta\|_{G,2} + \sqrt{2\epsilon_v}, \tag{15}$$

where γ_Ω is the exclusion price defined in the main text. It is zero if a target-risk oracle is an anchor or if at least one oracle departure is admitted. Under the studentized moment conditions below, an oracle departure is admitted with probability tending to one when its improvement over every anchor dominates the simultaneous paired standard-error radius.

Theorem 4 (False complex-admission control)**.** *For each of the fixed six paired comparisons, assume either positive variance with uniformly bounded $(2+\kappa)$ standardized moments, or an almost-surely constant difference evaluated by its exact mean. The studentized guard admits a departure whose ordinary pseudo-outcome risk fails to improve on at least one anchor with probability at most $\delta + o(1)$. Conditional on construction and assembly, the empirical-Bernstein guard has the corresponding finite-sample bound δ for the declared capped patient-vector loss. The corresponding interpretation for uncapped causal target risk additionally includes the nuisance remainder in Theorem 1.*

Proof of Theorem 3. We prove the result coordinatewise. Fix a basis coordinate j, and temporarily suppress the coordinate subscript. Write

$$m_a(x) = \mathbb{E}(B \mid A = a, X = x), \qquad e(x) = \Pr(A = 1 \mid X = x).$$

First consider the score evaluated at the true nuisance functions:

$$\psi(O;\eta_0) = m_1(X) - m_0(X) + \frac{A\{B - m_1(X)\}}{e(X)} - \frac{(1-A)\{B - m_0(X)\}}{1 - e(X)}.$$

Conditional on $X = x$,

$$\begin{aligned}\mathbb{E}\left[\frac{A\{B - m_1(x)\}}{e(x)}\middle| X = x\right] &= \frac{1}{e(x)}\Pr(A = 1 \mid X = x) \\ &\quad \times \mathbb{E}\{B - m_1(x) \mid A = 1, X = x\} \\ &= \mathbb{E}\{B - m_1(x) \mid A = 1, X = x\} \\ &= 0.\end{aligned}$$

Likewise,

$$\mathbb{E}\left[\frac{(1-A)\{B-m_0(x)\}}{1-e(x)}\middle| X=x\right] = \mathbb{E}\{B-m_0(x) \mid A=0, X=x\} = 0.$$

It follows that

$$\mathbb{E}\{\psi(O;\eta_0) \mid X=x\} = m_1(x) - m_0(x).$$

Under consistency and conditional exchangeability,

$$m_a(x) = \mathbb{E}\{B(a) \mid X=x\},$$

so

$$m_1(x) - m_0(x) = \mathbb{E}\{B(1) - B(0) \mid X=x\} = \tau_j(x).$$

Applying this argument to each basis coordinate proves

$$\mathbb{E}\{\boldsymbol{\psi}(O;\eta_0) \mid X=x\} = \boldsymbol{\tau}_B(x).$$

We next derive the exact conditional bias when estimated nuisances are used. Condition on the independently trained nuisance estimates and on $X=x$. The fitted score is

$$\widehat{\psi} = \widehat{m}_1 - \widehat{m}_0 + \frac{A(B-\widehat{m}_1)}{\widehat{e}} - \frac{(1-A)(B-\widehat{m}_0)}{1-\widehat{e}},$$

where all nuisance functions in this display are evaluated at x. Because the nuisance estimates are trained independently of the current observation,

$$\mathbb{E}\{A(B-\widehat{m}_1) \mid X=x\} = e(m_1-\widehat{m}_1),$$

and

$$\mathbb{E}\{(1-A)(B-\widehat{m}_0) \mid X=x\} = (1-e)(m_0-\widehat{m}_0).$$

Therefore,

$$\mathbb{E}(\widehat{\psi} \mid X=x, \widehat{\eta}) = \widehat{m}_1 - \widehat{m}_0 + \frac{e}{\widehat{e}}(m_1-\widehat{m}_1) - \frac{1-e}{1-\widehat{e}}(m_0-\widehat{m}_0). \tag{16}$$

Define the nuisance errors

$$\Delta_e = \widehat{e} - e, \qquad \Delta_1 = \widehat{m}_1 - m_1, \qquad \Delta_0 = \widehat{m}_0 - m_0.$$

Subtracting $m_1 - m_0$ from (16) gives

$$\begin{aligned}
&\mathbb{E}(\widehat{\psi} \mid X=x, \widehat{\eta}) - (m_1-m_0) \\
&= \Delta_1 - \Delta_0 - \frac{e}{\widehat{e}}\Delta_1 + \frac{1-e}{1-\widehat{e}}\Delta_0 \\
&= \Delta_1\left(1-\frac{e}{\widehat{e}}\right) + \Delta_0\left(-1+\frac{1-e}{1-\widehat{e}}\right) \\
&= \Delta_1\frac{\widehat{e}-e}{\widehat{e}} + \Delta_0\frac{\widehat{e}-e}{1-\widehat{e}} \\
&= (\widehat{e}-e)\left\{\frac{\widehat{m}_1-m_1}{\widehat{e}} + \frac{\widehat{m}_0-m_0}{1-\widehat{e}}\right\}.
\end{aligned}$$

This is exactly the coordinatewise remainder $R_{\eta,j}(x)$ in (10).

The expression also makes the double-robust property immediate. If $\widehat{e} = e$, then the factor $\widehat{e} - e$ is zero. Alternatively, if $\widehat{m}_1 = m_1$ and $\widehat{m}_0 = m_0$, then the expression inside braces is zero. Thus either a correct propensity model or two correct outcome-regression models removes the conditional score bias.

Finally, the same expression establishes Neyman orthogonality. Consider a regular path through the true nuisance functions,

$$e_t = e + th_e, \qquad m_{1,t} = m_1 + th_1, \qquad m_{0,t} = m_0 + th_0.$$

Along this path, the conditional remainder is

$$R_t(x) = t^2 h_e(x) \left\{ \frac{h_1(x)}{e(x) + t h_e(x)} + \frac{h_0(x)}{1 - e(x) - t h_e(x)} \right\}.$$

Overlap ensures that the denominators are bounded away from zero for all t in a sufficiently small neighborhood of zero. Hence

$$\frac{d}{dt} R_t(x) \Big|_{t=0} = 0.$$

Thus the first-order Gateaux derivative of the conditional score mean with respect to the nuisance functions vanishes at the truth. □

Proof of Theorem 1. With the pseudo-target characterized, condition throughout on $\mathcal{F}_{cs}$. Under this conditioning, every candidate $g \in \mathcal{H}$ and the matrix G_Ω are fixed.

Let $\boldsymbol{\psi}^v$ denote the vector pseudo-outcome for a fresh guard observation. By Theorem 3,

$$\mathbb{E}(\boldsymbol{\psi}^v \mid X, \mathcal{F}_{cs}) = \boldsymbol{\tau}_B(X) + R_\eta(X).$$

Define

$$\mu(X) = \boldsymbol{\tau}_B(X) + R_\eta(X).$$

For a fixed candidate g, its conditional fresh-draw pseudo-outcome risk is

$$Q_\Omega(g) = \mathbb{E}\left[\{\boldsymbol{\psi}^v - g(X)\}^\top G_\Omega \{\boldsymbol{\psi}^v - g(X)\} \middle| \mathcal{F}_{cs} \right].$$

Write

$$\boldsymbol{\psi}^v - g(X) = \{\boldsymbol{\psi}^v - \mu(X)\} + \{\mu(X) - g(X)\}.$$

Expanding the quadratic form gives

$$\begin{aligned} Q_\Omega(g) = \mathbb{E}\left[\{\boldsymbol{\psi}^v - \mu(X)\}^\top G_\Omega \{\boldsymbol{\psi}^v - \mu(X)\} \middle| \mathcal{F}_{cs} \right] \\ + 2\mathbb{E}\left[\{\boldsymbol{\psi}^v - \mu(X)\}^\top G_\Omega \{\mu(X) - g(X)\} \middle| \mathcal{F}_{cs} \right] \\ + \mathbb{E}\left[\{\mu(X) - g(X)\}^\top G_\Omega \{\mu(X) - g(X)\} \middle| \mathcal{F}_{cs} \right]. \end{aligned}$$

The middle term is zero. Indeed, by iterated expectation,

$$\begin{aligned} &\mathbb{E}\left[\{\boldsymbol{\psi}^v - \mu(X)\}^\top G_\Omega \{\mu(X) - g(X)\} \middle| \mathcal{F}_{cs} \right] \\ &= \mathbb{E}\left[\mathbb{E}\{\boldsymbol{\psi}^v - \mu(X) \mid X, \mathcal{F}_{cs}\}^\top G_\Omega \{\mu(X) - g(X)\} \middle| \mathcal{F}_{cs} \right] \\ &= 0. \end{aligned}$$

Consequently,

$$Q_\Omega(g) = C + \|g - \mu\|_{G,2}^2, \tag{17}$$

where

$$C = \mathbb{E}\left[\{\boldsymbol{\psi}^v - \mu(X)\}^\top G_\Omega \{\boldsymbol{\psi}^v - \mu(X)\} \middle| \mathcal{F}_{cs} \right]$$

does not depend on g.

Let

$$\mathcal{E}_v = \left\{ \sup_{g \in \mathcal{H}} \left| \widehat{Q}_\Omega(g) - Q_\Omega(g) \right| \le \epsilon_v \right\}.$$

By assumption,

$$\Pr(\mathcal{E}_v \mid \mathcal{F}_{cs}) \ge 1 - \delta.$$

We now work on $\mathcal{E}_v$.

Let

$$\widehat{s} \in \arg\min_{s \in \mathcal{S}} \widehat{Q}_\Omega(s)$$

be the empirically best anchor, and let $\widehat{\mathcal{D}} \subseteq \mathcal{D}$ be the set of eligible departures. By construction,

$$\widehat{g} \in \arg\min_{h \in \{\widehat{s}\} \cup \widehat{\mathcal{D}}} \widehat{Q}_\Omega(h).$$

Therefore, for every anchor $s \in \mathcal{S}$,

$$\widehat{Q}_\Omega(\widehat{g}) \le \widehat{Q}_\Omega(\widehat{s}) \le \widehat{Q}_\Omega(s).$$

Using the uniform-deviation event on both $\widehat{g}$ and s gives

$$\begin{aligned} Q_\Omega(\widehat{g}) &\le \widehat{Q}_\Omega(\widehat{g}) + \epsilon_v \\ &\le \widehat{Q}_\Omega(s) + \epsilon_v \\ &\le Q_\Omega(s) + 2\epsilon_v. \end{aligned}$$

Substituting the decomposition (17) and cancelling the common constant C yields

$$\|\widehat{g} - \mu\|_{G,2}^2 \le \|s - \mu\|_{G,2}^2 + 2\epsilon_v.$$

Since both sides are nonnegative,

$$\begin{aligned} \|\widehat{g} - \mu\|_{G,2} &\le \left\{\|s - \mu\|_{G,2}^2 + 2\epsilon_v\right\}^{1/2} \\ &\le \|s - \mu\|_{G,2} + \sqrt{2\epsilon_v}. \end{aligned}$$

The second inequality follows from $\sqrt{a^2 + b^2} \le a + b$ for $a, b \ge 0$.

Because G_Ω is positive semidefinite, it has a positive-semidefinite square root $G_\Omega^{1/2}$, and

$$\|f\|_{G,2} = \left[\mathbb{E}\left\{\|G_\Omega^{1/2} f(X)\|_2^2\right\}\right]^{1/2}.$$

Thus $\|\cdot\|_{G,2}$ is an L_2 seminorm and satisfies the triangle inequality. Since

$$\mu - \boldsymbol{\tau}_B = R_\eta,$$

we have

$$\begin{aligned} \|\widehat{g} - \boldsymbol{\tau}_B\|_{G,2} &\le \|\widehat{g} - \mu\|_{G,2} + \|R_\eta\|_{G,2} \\ &\le \|s - \mu\|_{G,2} + \|R_\eta\|_{G,2} + \sqrt{2\epsilon_v} \\ &\le \|s - \boldsymbol{\tau}_B\|_{G,2} + 2\|R_\eta\|_{G,2} + \sqrt{2\epsilon_v}. \end{aligned}$$

The final inequality uses

$$\|s - \mu\|_{G,2} = \|s - \boldsymbol{\tau}_B - R_\eta\|_{G,2} \le \|s - \boldsymbol{\tau}_B\|_{G,2} + \|R_\eta\|_{G,2}.$$

Because the displayed bound holds for every $s \in \mathcal{S}$, minimizing over the anchor library gives

$$\|\widehat{g} - \boldsymbol{\tau}_B\|_{G,2} \le \min_{s \in \mathcal{S}} \|s - \boldsymbol{\tau}_B\|_{G,2} + 2\|R_\eta\|_{G,2} + \sqrt{2\epsilon_v}.$$

This proves the claimed anchor-relative bound on $\mathcal{E}_v$ and hence with conditional probability at least $1 - \delta$.

If G_Ω is singular, functions that differ only in its null space have zero seminorm distance. The preceding argument remains valid because it uses only the quadratic form induced by G_Ω and the triangle inequality on its range. □

Proof of Corollary 1. To extend the bound to the full library, continue to condition on $\mathcal{F}_{cs}$ and work on the uniform deviation event $\mathcal{E}_v$ from the preceding proof. Define the retained library

$$\widehat{\mathcal{A}} = \mathcal{S} \cup \widehat{\mathcal{D}}.$$

Although the implementation carries only the empirically best anchor to the final comparison, minimizing over $\{\widehat{s}\} \cup \widehat{\mathcal{D}}$ is equivalent to minimizing empirical risk over $\widehat{\mathcal{A}}$, because $\widehat{s}$ has no larger empirical risk than any other anchor. Hence

$$\widehat{Q}_\Omega(\widehat{g}) = \min_{h \in \widehat{\mathcal{A}}} \widehat{Q}_\Omega(h).$$

Let

$$h_{\widehat{\mathcal{A}}} \in \arg\min_{h \in \widehat{\mathcal{A}}} \|h - \boldsymbol{\tau}_B\|_{G,2}^2$$

be a target-risk minimizer within the retained library. Since $h_{\widehat{\mathcal{A}}} \in \widehat{\mathcal{A}}$,

$$\widehat{Q}_\Omega(\widehat{g}) \le \widehat{Q}_\Omega(h_{\widehat{\mathcal{A}}}).$$

Using the uniform-deviation event on both candidates gives

$$Q_\Omega(\widehat{g}) \le Q_\Omega(h_{\widehat{\mathcal{A}}}) + 2\epsilon_v.$$

Applying (17) and cancelling C yields

$$\|\widehat{g} - \mu\|_{G,2}^2 \le \|h_{\widehat{\mathcal{A}}} - \mu\|_{G,2}^2 + 2\epsilon_v.$$

The same square-root and triangle-inequality argument as in the proof of Theorem 1 therefore gives

$$\|\widehat{g} - \boldsymbol{\tau}_B\|_{G,2} \le \|h_{\widehat{\mathcal{A}}} - \boldsymbol{\tau}_B\|_{G,2} + 2\|R_\eta\|_{G,2} + \sqrt{2\epsilon_v}. \tag{18}$$

By definition of the exclusion price,

$$\gamma_\Omega = \min_{h \in \widehat{\mathcal{A}}} \|h - \boldsymbol{\tau}_B\|_{G,2}^2 - \min_{h \in \mathcal{H}} \|h - \boldsymbol{\tau}_B\|_{G,2}^2.$$

Thus

$$\|h_{\widehat{\mathcal{A}}} - \boldsymbol{\tau}_B\|_{G,2}^2 = \min_{h \in \mathcal{H}} \|h - \boldsymbol{\tau}_B\|_{G,2}^2 + \gamma_\Omega.$$

Substituting this identity into (18) proves (15).

We next characterize when $\gamma_\Omega = 0$. Let

$$\mathcal{H}^\star = \arg\min_{h \in \mathcal{H}} \|h - \boldsymbol{\tau}_B\|_{G,2}^2$$

denote the set of full-library target-risk minimizers. If $\mathcal{H}^\star$ contains an anchor, then that minimizer belongs to $\widehat{\mathcal{A}}$ because every anchor is always retained. If $\mathcal{H}^\star$ contains a departure that is admitted, then that departure also belongs to $\widehat{\mathcal{A}}$. In either case,

$$\min_{h \in \widehat{\mathcal{A}}} \|h - \boldsymbol{\tau}_B\|_{G,2}^2 = \min_{h \in \mathcal{H}} \|h - \boldsymbol{\tau}_B\|_{G,2}^2,$$

and hence $\gamma_\Omega = 0$.

Finally, consider a target-risk oracle departure $h^\star \in \mathcal{D}$. For each anchor $s \in \mathcal{S}$, define its conditional pseudo-risk improvement

$$\Delta_{s,h^\star} = \mathbb{E}(D_i^{s,h^\star} \mid \mathcal{F}_{cs}) = Q_\Omega(s) - Q_\Omega(h^\star).$$

Suppose first that each comparison has positive conditional variance

$$\sigma_{s,h^\star}^2 = \mathrm{Var}(D_i^{s,h^\star} \mid \mathcal{F}_{cs}) > 0$$

and that

$$\min_{s \in \mathcal{S}} \left\{ \frac{\sqrt{n_v}\Delta_{s,h^\star}}{\sigma_{s,h^\star}} - t_{n_v-1,\,1-\delta/N} \right\} \longrightarrow +\infty. \tag{19}$$

The assumed studentized central limit theorem and standard-error consistency imply, simultaneously over the fixed anchor set,

$$\overline{D}_{s,h^\star} = \Delta_{s,h^\star} + O_{\mathrm{Pr}}\left(\frac{\sigma_{s,h^\star}}{\sqrt{n_v}}\right), \qquad \widehat{\mathrm{se}}_{s,h^\star} = \frac{\sigma_{s,h^\star}}{\sqrt{n_v}}\{1 + o_{\mathrm{Pr}}(1)\}.$$

Therefore,

$$\frac{\mathrm{LCB}_{s,h^\star}^T}{\sigma_{s,h^\star}/\sqrt{n_v}} = \frac{\sqrt{n_v}\Delta_{s,h^\star}}{\sigma_{s,h^\star}} - t_{n_v-1,\,1-\delta/N} + O_{\mathrm{Pr}}(1) + o_{\mathrm{Pr}}(1).$$

Under (19), the right-hand side converges to $+\infty$ uniformly over the finite anchor set. Hence

$$\Pr\left\{\mathrm{LCB}_{s,h^\star}^T > 0 \text{ for every } s \in \mathcal{S} \middle| \mathcal{F}_{cs}\right\} \longrightarrow 1.$$

Thus $h^\star$ is admitted with probability tending to one.

If a comparison has zero variance, then $D_i^{s,h^\star} = \Delta_{s,h^\star}$ almost surely. Under the implementation's exact-mean rule, a strictly positive $\Delta_{s,h^\star}$ passes deterministically. Consequently, the same admission conclusion holds when positive-variance comparisons satisfy (19) and zero-variance comparisons have strictly positive exact improvements. □

Proof of Theorem 4. It remains to control false admission. Condition on $\mathcal{F}_{cs}$. For each anchor–departure pair $(s, h) \in \mathcal{S} \times \mathcal{D}$, define

$$D_i^{s,h} = \ell_i(s) - \ell_i(h), \qquad \Delta_{s,h} = \mathbb{E}(D_i^{s,h} \mid \mathcal{F}_{cs}) = Q_\Omega(s) - Q_\Omega(h).$$

Thus h improves on anchor s in ordinary pseudo-outcome risk exactly when $\Delta_{s,h} > 0$. There are

$$N = |\mathcal{S}||\mathcal{D}| = 6$$

fixed comparisons.

Consider first a comparison with positive conditional variance

$$\sigma_{s,h}^2 = \mathrm{Var}(D_i^{s,h} \mid \mathcal{F}_{cs}) > 0.$$

Let

$$\widehat{\mathrm{se}}_{s,h} = \frac{\widehat{\sigma}_{s,h}}{\sqrt{n_v}}.$$

The assumed standardized moment condition gives the studentized central limit theorem

$$Z_{s,h} := \frac{\overline{D}_{s,h} - \Delta_{s,h}}{\widehat{\mathrm{se}}_{s,h}} \rightsquigarrow N(0,1),$$

and the Student critical value satisfies

$$t_{n_v-1,\,1-\delta/N} \longrightarrow z_{1-\delta/N}.$$

Suppose that the departure does not improve on this anchor, so that $\Delta_{s,h} \le 0$. Its lower confidence bound is

$$\mathrm{LCB}_{s,h}^T = \overline{D}_{s,h} - t_{n_v-1,\,1-\delta/N}\widehat{\mathrm{se}}_{s,h}.$$

The event that this non-improving comparison nevertheless passes can be written as

$$\left\{\mathrm{LCB}_{s,h}^T > 0\right\} = \left\{\frac{\overline{D}_{s,h} - \Delta_{s,h}}{\widehat{\mathrm{se}}_{s,h}} > t_{n_v-1,\,1-\delta/N} - \frac{\Delta_{s,h}}{\widehat{\mathrm{se}}_{s,h}}\right\}.$$

Because $\Delta_{s,h} \le 0$,

$$t_{n_v-1,\,1-\delta/N} - \frac{\Delta_{s,h}}{\widehat{\mathrm{se}}_{s,h}} \ge t_{n_v-1,\,1-\delta/N}.$$

It follows that

$$\begin{aligned}\Pr\left\{\mathrm{LCB}_{s,h}^T > 0 \middle| \mathcal{F}_{cs}\right\} &\le \Pr\left\{Z_{s,h} > t_{n_v-1,\,1-\delta/N} \middle| \mathcal{F}_{cs}\right\} \\ &\le \frac{\delta}{N} + o(1). \end{aligned} \tag{20}$$

Since the number of comparisons is fixed, the $o(1)$ terms may be taken uniformly over the six comparisons.

If instead $\mathrm{Var}(D_i^{s,h} \mid \mathcal{F}_{cs}) = 0$, then

$$D_i^{s,h} = \Delta_{s,h} \qquad \text{almost surely.}$$

The implementation evaluates this comparison by its exact mean. Hence a nonpositive constant $\Delta_{s,h} \le 0$ cannot produce a positive lower bound and cannot pass.

Now define the false-complex-admission event

$$\mathcal{F}_{\mathrm{adm}} = \left\{ \begin{array}{c} \text{there exists } h \in \mathcal{D} \text{ that is admitted, but} \\ \Delta_{s,h} \le 0 \text{ for at least one } s \in \mathcal{S} \end{array} \right\}.$$

For any departure appearing in this event, choose an anchor for which $\Delta_{s,h} \le 0$. Admission requires the lower confidence bound to be positive against every anchor, and hence in particular against this non-improved anchor. Therefore,

$$\mathcal{F}_{\mathrm{adm}} \subseteq \bigcup_{\substack{(s,h)\in\mathcal{S}\times\mathcal{D}:\\ \Delta_{s,h}\le 0}} \left\{\mathrm{LCB}_{s,h}^T > 0\right\}.$$

Using (20) and a union bound over at most N comparisons gives

$$\Pr(\mathcal{F}_{\mathrm{adm}} \mid \mathcal{F}_{cs}) \le N\left\{\frac{\delta}{N} + o(1)\right\} = \delta + o(1).$$

No independence among the six comparisons is required for this argument.

We now consider the capped empirical-Bernstein guard. Since

$$0 \le \ell_i^B(g) \le B,$$

the paired capped difference

$$D_i^{B,s,h} = \ell_i^B(s) - \ell_i^B(h)$$

lies in $[-B, B]$. Let

$$\Delta_{s,h}^B = \mathbb{E}(D_i^{B,s,h} \mid \mathcal{F}_{cs}).$$

The empirical-Bernstein inequality used in (9) implies that, for each fixed comparison,

$$\Pr\left[\Delta_{s,h}^B < \overline{D}_{s,h}^B - \sqrt{\frac{2\widehat{V}_{s,h}^B \log(2N/\delta)}{n_v}} - \frac{14B\log(2N/\delta)}{3(n_v-1)}\middle| \mathcal{F}_{cs}\right] \le \frac{\delta}{N}.$$

Equivalently,

$$\Pr\left\{\Delta_{s,h}^B < \mathrm{LCB}_{s,h}^{EB} \middle| \mathcal{F}_{cs}\right\} \le \frac{\delta}{N}.$$

A union bound therefore gives

$$\Pr\left\{\mathrm{LCB}_{s,h}^{EB} \le \Delta_{s,h}^B \text{ for every } (s,h) \middle| \mathcal{F}_{cs}\right\} \ge 1 - \delta.$$

On this simultaneous event, if a departure fails to improve on an anchor in capped risk, so that $\Delta_{s,h}^B \le 0$, then

$$\mathrm{LCB}_{s,h}^{EB} \le \Delta_{s,h}^B \le 0,$$

and that comparison cannot pass. Consequently, the probability of admitting a departure whose capped risk fails to improve on at least one anchor is at most δ.

This finite-sample conclusion concerns only the declared capped patient-vector pseudo-loss. The final candidate may be selected using uncapped empirical risk after admission, and neither the capped nor the ordinary admission result removes the nuisance remainder separating the pseudo-target $\mu = \boldsymbol{\tau}_B + R_\eta$ from the causal target $\boldsymbol{\tau}_B$. □

D.3 Direct utility-level simultaneous inference

Proof of Theorem 2. Turning from selection to inference, let

$$\mathcal{K} = \{(g,\ell) : g = 1,\dots,G,\ \ell = 1,\dots,L\}, \qquad d = |\mathcal{K}| = GL.$$

By assumption, d is fixed and the target groups and coefficient rows are frozen independently of the inference outcomes.

For each $(g,\ell) \in \mathcal{K}$, define the projected target and estimator

$$\theta_{g\ell}^B = c_\ell^\top \theta_g, \qquad \widehat{\theta}_{g\ell}^B = c_\ell^\top \widehat{\theta}_g,$$

and define the projected influence coordinate

$$\phi_{ig\ell} = c_\ell^\top \phi_{ig}.$$

Projecting the asymptotic linear representation (11) gives

$$\begin{aligned}\widehat{\theta}_{g\ell}^B - \theta_{g\ell}^B &= c_\ell^\top(\widehat{\theta}_g - \theta_g) \\ &= \frac{1}{n}\sum_{i=1}^n c_\ell^\top \phi_{ig} + r_{ng\ell} \\ &= \frac{1}{n}\sum_{i=1}^n \phi_{ig\ell} + r_{ng\ell}, \end{aligned} \tag{21}$$

where

$$r_{ng\ell} = c_\ell^\top r_{ng}.$$

Because the coefficient-row norms are uniformly bounded and the number of rows is fixed, the assumed uniform nuisance-product remainder implies

$$\max_{(g,\ell)\in\mathcal{K}} \sqrt{n}\,|r_{ng\ell}| = o_{\Pr}(1).$$

Let

$$\sigma_{g\ell}^2 = \mathbb{E}(\phi_{ig\ell}^2).$$

The theorem assumes that these variances are uniformly bounded away from zero and infinity. Define

$$\widehat{\sigma}_{g\ell}^2 = n\widehat{s}_{g\ell}^2 = \frac{1}{n}\sum_{i=1}^n \widehat{\phi}_{ig\ell}^2.$$

The assumed covariance consistency implies

$$\max_{(g,\ell)\in\mathcal{K}} \left|\widehat{\sigma}_{g\ell}^2 - \sigma_{g\ell}^2\right| = o_{\Pr}(1).$$

In particular,

$$\widehat{\sigma}_{g\ell}/\sigma_{g\ell} = 1 + o_{\Pr}(1)$$

uniformly over the fixed family.

Combining this standard-error consistency with (21) gives

$$\begin{aligned} Z_{n,g\ell} &:= \frac{\widehat{\theta}_{g\ell}^B - \theta_{g\ell}^B}{\widehat{s}_{g\ell}} \\ &= \frac{n^{-1/2}\sum_{i=1}^n \phi_{ig\ell}}{\sigma_{g\ell}} + o_{\Pr}(1), \end{aligned}$$

uniformly over $(g,\ell)\in\mathcal{K}$.

Stack the coordinates in

$$\boldsymbol{Z}_n = \{Z_{n,g\ell} : (g,\ell)\in\mathcal{K}\}.$$

The bounded $(2+\kappa)$ moment assumption and the fixed dimension d give the multivariate Lindeberg–Feller central limit theorem:

$$\boldsymbol{Z}_n \rightsquigarrow \boldsymbol{Z} \sim N_d(0,\mathcal{R}),$$

where the limiting correlation matrix has entries

$$\mathcal{R}_{(g,\ell),(g',\ell')} = \frac{\mathbb{E}(\phi_{ig\ell}\phi_{ig'\ell'})}{\sigma_{g\ell}\sigma_{g'\ell'}}.$$

This covariance matrix retains all dependence induced by common basis coordinates, common observations, multiple utilities, and overlapping target groups. No independence across target groups or utilities is assumed.

By the continuous mapping theorem,

$$T_n := \max_{g,\ell} |Z_{n,g\ell}| \rightsquigarrow T_\infty := \max_{g,\ell} |Z_{g\ell}|. \tag{22}$$

We next establish validity of the multiplier approximation. Let $\xi_1,\dots,\xi_n$ be the shared multipliers used in the procedure, independent of the data, with mean zero and variance one. The multiplier statistic can be written as

$$\begin{aligned} T^* &= \max_{g,\ell} \left| \frac{n^{-1}\sum_{i=1}^n \xi_i \widehat{\phi}_{ig\ell}}{\widehat{s}_{g\ell}} \right| \\ &= \max_{g,\ell} \left| \frac{\sum_{i=1}^n \xi_i \widehat{\phi}_{ig\ell}}{\left(\sum_{i=1}^n \widehat{\phi}_{ig\ell}^2\right)^{1/2}} \right|. \end{aligned} \tag{23}$$

Using the same multiplier ξ_i for every (g, ℓ) preserves the empirical cross-covariance among all contrasts.

For Gaussian multipliers, conditional on the data, the vector inside (23) is exactly multivariate Gaussian with empirical correlation matrix

$$\widehat{\mathcal{R}}_{(g,\ell),(g',\ell')} = \frac{\sum_{i=1}^n \widehat{\phi}_{ig\ell}\widehat{\phi}_{ig'\ell'}}{\left(\sum_{i=1}^n \widehat{\phi}_{ig\ell}^2\right)^{1/2}\left(\sum_{i=1}^n \widehat{\phi}_{ig'\ell'}^2\right)^{1/2}}.$$

For bounded mean-zero multipliers such as Rademacher multipliers, the assumed leverage condition

$$\max_{i,g,\ell} \frac{\widehat{\phi}_{ig\ell}^2}{\sum_{j=1}^n \widehat{\phi}_{jg\ell}^2} = o_{\Pr}(1)$$

is the conditional Lindeberg condition that prevents a single observation from dominating a studentized contrast.

Covariance consistency implies

$$\widehat{\mathcal{R}} \longrightarrow_{\Pr} \mathcal{R}.$$

The Gaussian multiplier argument in the first case, or the conditional multiplier central limit theorem in the second case, therefore yields

$$\sup_{t\in\mathbb{R}} |\Pr(T^* \le t \mid O_1, \ldots, O_n) - \Pr(T_\infty \le t)| \longrightarrow_{\Pr} 0. \tag{24}$$

Together with (22), this also implies

$$\sup_{t\in\mathbb{R}} |\Pr(T_n \le t) - \Pr(T^* \le t \mid O_1, \ldots, O_n)| \longrightarrow_{\Pr} 0.$$

This is the standard multiplier approximation for the maximum of a finite studentized vector; see Chernozhukov et al. (2013) and Chernozhukov et al. (2018).

Let $q^*_{1-\alpha}$ denote the conditional $(1-\alpha)$ quantile of T^*, and let $\widehat{q}_{1-\alpha}$ be its empirical approximation from the multiplier draws. Because the number of multiplier draws diverges, the conditional empirical distribution of the scalar draws converges uniformly to the conditional distribution of T^*. At every continuity point of the limiting Gaussian maximum distribution,

$$\widehat{q}_{1-\alpha} - q^*_{1-\alpha} = o_{\Pr}(1).$$

The marginal variance lower bound ensures that the relevant finite Gaussian maximum has no atom at its $(1-\alpha)$ quantile, after any exactly redundant coordinates are removed. It follows from (24) that

$$\Pr\left\{\max_{g,\ell} \frac{|\widehat{\theta}^B_{g\ell} - \theta^B_{g\ell}|}{\widehat{s}_{g\ell}} \le \widehat{q}_{1-\alpha}\right\} \ge 1 - \alpha - o(1).$$

Equivalently, with probability at least $1 - \alpha - o(1)$,

$$\widehat{\theta}^B_{g\ell} - \widehat{q}_{1-\alpha}\widehat{s}_{g\ell} \le \theta^B_{g\ell} \le \widehat{\theta}^B_{g\ell} + \widehat{q}_{1-\alpha}\widehat{s}_{g\ell} \tag{25}$$

simultaneously for every $(g, \ell) \in \mathcal{K}$.

It remains to transfer these projected-basis bands to the true utility effects. Proposition 1 gives, for every x,

$$|\tau_{\omega_\ell}(x) - \tau^B_{\omega_\ell}(x)| \le 2\varepsilon_u.$$

Because the target group has positive probability,

$$\begin{aligned} |\theta_{g\ell} - \theta^B_{g\ell}| &= \left|\mathbb{E}\left[\tau_{\omega_\ell}(X) - \tau^B_{\omega_\ell}(X) \mid X \in R_g\right]\right| \\ &\le \mathbb{E}\left[\left|\tau_{\omega_\ell}(X) - \tau^B_{\omega_\ell}(X)\right| \mid X \in R_g\right] \\ &\le 2\varepsilon_u. \end{aligned}$$

Thus, when a valid true-utility misspecification radius is declared, set

$$b_{g\ell} = 2\varepsilon_u.$$

On the simultaneous event (25),

$$\begin{aligned}\theta_{g\ell} &\geq \theta^B_{g\ell} - b_{g\ell} \\ &\geq \widehat{\theta}^B_{g\ell} - \widehat{q}_{1-\alpha}\widehat{s}_{g\ell} - b_{g\ell},\end{aligned}$$

and similarly

$$\begin{aligned}\theta_{g\ell} &\leq \theta^B_{g\ell} + b_{g\ell} \\ &\leq \widehat{\theta}^B_{g\ell} + \widehat{q}_{1-\alpha}\widehat{s}_{g\ell} + b_{g\ell}.\end{aligned}$$

Hence

$$\theta_{g\ell} \in I_{g\ell}$$

simultaneously for every g and ℓ, where

$$I_{g\ell} = \left[c_\ell^\top \widehat{\theta}_g - \widehat{q}_{1-\alpha}\widehat{s}_{g\ell} - b_{g\ell},\ c_\ell^\top \widehat{\theta}_g + \widehat{q}_{1-\alpha}\widehat{s}_{g\ell} + b_{g\ell}\right].$$

This proves

$$\Pr\left\{\theta_{g\ell} \in I_{g\ell} \text{ for all } g, \ell\right\} \geq 1 - \alpha - o(1),$$

which is (13).

We finally verify the four-way error-control statement. Let $\mathcal{E}_{\mathrm{sim}}$ denote the event that all intervals cover their corresponding true effects. On $\mathcal{E}_{\mathrm{sim}}$, suppose the audit reports robust support for Treatment 1 for group g. Then

$$\inf_{\ell \leq L} L_{g\ell} > 0,$$

so for every utility row ℓ,

$$\theta_{g\ell} \geq L_{g\ell} > 0.$$

Therefore every utility in the frozen family truly favors Treatment 1.

If the audit reports robust support for Treatment 0, then

$$\sup_{\ell \leq L} U_{g\ell} < 0,$$

and simultaneous coverage gives

$$\theta_{g\ell} \leq U_{g\ell} < 0$$

for every utility row.

If the audit reports certified preference sensitivity, then there exist ℓ_+ and ℓ_- such that

$$L_{g\ell_+} > 0, \qquad U_{g\ell_-} < 0.$$

On $\mathcal{E}_{\mathrm{sim}}$,

$$\theta_{g\ell_+} > 0, \qquad \theta_{g\ell_-} < 0,$$

so the frozen utility family contains two true effects with opposing signs.

It follows that any false robust conclusion or false preference-sensitive conclusion, for any one of the G target groups, must occur outside $\mathcal{E}_{\mathrm{sim}}$. Therefore,

$$\Pr\left\{\begin{array}{c}\text{at least one false robust or false} \\ \text{preference-sensitive output is reported}\end{array}\right\} \leq \Pr(\mathcal{E}^c_{\mathrm{sim}}) \leq \alpha + o(1).$$

The evidence-uncertain output makes no sign claim and therefore does not contribute to this false-decision event. Likewise, a fail-closed diagnostic that replaces a potentially decisive output by evidence uncertainty can only remove sign claims and cannot increase the familywise false-decision probability.

If ε_u is unknown and the analyst sets $b_{g\ell} = 0$, the same proof establishes simultaneous coverage and four-way error control only for the projected effects $\theta^B_{g\ell}$, not for the unrestricted true utility effects $\theta_{g\ell}$. □

Finite-dimensional scope. The preceding argument establishes Theorem 2 for a fixed family with $d = GL$.

Scenario	Final guard	Shared	Coordinate	Private	Final selection
Mixed, $M = 8$	0.0801	0.0801	0.1029	0.1221	shared 20/20
Nonlinear, $M = 16$	0.1090	0.1090	0.1382	0.1449	shared 20/20

Table 4: Utility RMSE for the complete shared-anchored estimator (20 paired repetitions; lower is better). The final guard selects the shared anchor in every run, so equality with the fixed shared row is exact rather than an independent refit.

E Additional Implementation Details

Finally, to stabilize finite multiplier runs, we report

$$\widetilde{q}_{1-\alpha} = \max\left\{\widehat{q}^{\mathrm{raw}}_{1-\alpha}, z_{1-\alpha/2}\right\}$$

and record the raw multiplier critical value alongside the floored value.

We serialize all pseudo-random seeds, split sizes, multiplier counts, coefficient rows, target definitions, and loss caps beside the raw results. We generate core tables from archived summaries and independently check the remaining literal tables against archived CSV artifacts. We use a modest common hyperparameter specification for the default forest across comparable methods. We estimate outcome standardization on construction data only and invert it before computing reported effects. We evaluate zero-variance comparisons exactly, making their admission decisions deterministic.

F Supplementary Experiments

F.1 Shared-anchored selector

Complete-estimator design. We randomize treatment at $1/2$ for Gaussian covariates and form continuous basis outcomes from a linear baseline, nonlinear perturbations, and treatment effects. Effect coordinates cycle through shared nonlinear, task-private nonlinear, constant, and exact-null tasks. We freeze a signed utility grid with $L = 41$ rows and orient every fourth coordinate as harm. We form AIPW pseudo-outcomes with three-fold cross-fitting and known propensity. The outer split reserves 70% of training observations for candidate construction and partitions the remainder between coordinate assembly and vector guarding; test observations never fit or select a model. We allocate all 480 outer-validation observations to coordinate selection for the coordinate baseline, whereas the final selector divides them equally between assembly and guarding. The coordinate baseline therefore receives a larger adaptive-selection sample.

Under this common design, we select the shared candidate in 20/20 mixed and 20/20 nonlinear runs. Its mixed guarded-minus-shared utility-RMSE contrast is exactly zero; guarded-minus-coordinate is $-0.02273\,[-0.03025, -0.01521]$, guarded-minus-private is $-0.04193\,[-0.04339, -0.04047]$, and guarded-minus-constant is $-0.04467\,[-0.04726, -0.04208]$. At $M = 16$, the corresponding differences are 0, $-0.02921\,[-0.03373, -0.02468]$, $-0.03592\,[-0.03681, -0.03503]$, and $-0.05187\,[-0.05312, -0.05062]$. These results match the shared oracle in both cells and improve over the coordinate, private, and constant candidates by the reported margins.

Selector-isolation phase diagram. We next isolate selection from fitting with pregenerated candidates and independent selection noise. Each cell uses $M = 12$, $L = 31$, $n_v = 900$, an independent test sample of 2,000, and 50 paired repetitions. We freeze the coordinate candidate on an assembly half before the final guard sees its validation half. The coordinate reference uses all 900 selection observations, whereas the final selector uses independent assembly and guard halves.

ρ_s	Final guard	Coordinate	Private	Shared	Selected
0.00	0.02981	0.02981	0.02981	0.19717	coordinate 50/50
0.25	0.03070	0.03045	0.04283	0.17304	coordinate 50/50
0.50	0.03080	0.03035	0.05336	0.13991	coordinate 50/50
0.75	0.03225	0.03059	0.06122	0.10221	coordinate 50/50
1.00	0.03035	0.03135	0.06885	0.03035	shared 50/50

Table 5: Utility RMSE in the selector-isolation phase diagram, with 50 paired repetitions per cell. Candidate prediction vectors are frozen before independent selection noise is generated. This design isolates the selection mechanism from learner fitting.

This isolated comparison characterizes structural adaptation across the shared-fraction continuum. The final guard and the legacy gate coincide through shared fraction .75. At the fully shared endpoint, the legacy gate chooses coordinate 22 times and shared 28 times, whereas the revised selector chooses shared in 50/50; the resulting revised-minus-legacy utility-RMSE interval is $[-0.00052, 0.00023]$ and includes zero. Simultaneous intervals use a two-sided Bonferroni paired-t critical value across all 90 reported phase–method–metric contrasts.

Admission calibration. We then calibrate the guard under an equal-risk null in which zero, constant, and shared have identical population risk and coordinate is a departure tested against all three anchors. We cross $n_v \in \{300, 800\}$, $M \in \{8, 16\}$, Gaussian/t_3 noise, and 300 independent repetitions per cell. Studentized false admission is 0–.0167 across the eight null cells. With signal increment .05 it rises from .473–.820 at $n_v = 300$ to .927–1.000 at $n_v = 800$; at increment .10 it is at least .987. The empirical-Bernstein guard admits no departure under the null for any tested cap $B \in \{.5, 1, 2, 5, 25\}$, with the expected loss of local power. This study isolates admission; the end-to-end studies then assess the causal target.

F.2 End-to-end utility-level inference

We use an independent design sample to freeze the subgroup rule before inference. Unless a model is explicitly labeled external-fit, we cross-fit nuisance models in three folds on the inference observations. We evaluate truth analytically in constant-effect cells and otherwise use independent Monte Carlo integration over 200,000 covariate draws. Each formal cell has 200 repetitions and 999 Gaussian multiplier draws. For direct max-t, we project the complete basis influence vector before studentization. The rectangle method first covers basis coordinates and propagates an ℓ_1 radius, while pointwise Wald intervals provide an unadjusted reference.

Regime	Direct	MC s.e.	Rectangle	Pointwise	Width	Abstain	Max leverage
randomized known	0.965	0.013	0.955	0.670	0.237	0.855	0.096
randomized external nuisance	0.960	0.014	0.940	0.615	0.226	0.858	0.085
calibration external constant	0.940	0.017	0.940	0.605	0.222	0.948	0.070
observational bounded correct	0.975	0.011	0.980	0.680	0.226	0.927	0.257
larger n correct	0.935	0.017	0.955	0.640	0.127	0.432	0.075
propensity wrong outcome correct	0.925	0.019	0.915	0.595	0.235	0.845	0.077
outcome wrong propensity correct	0.980	0.010	0.975	0.665	0.253	0.912	0.698
both nuisances wrong	0.910	0.020	0.905	0.565	0.235	0.770	0.077
weak overlap	0.965	0.013	0.960	0.685	0.385	0.993	0.957
rare group	0.915	0.020	0.920	0.435	0.768	0.957	0.939
heavy tail t3	0.960	0.014	0.940	0.625	0.246	0.917	0.957
larger family	0.905	0.021	0.925	0.085	0.317	0.944	0.921

Table 6: Complete 200-repetition end-to-end grid. Coverage is simultaneous over the full frozen family; max leverage is the largest squared standardized influence share recorded across repetitions.

We sharpen the reference estimates with 1,000 repetitions, 1,999 multipliers, analytic truth for the randomized constant-effect cell, and 400,000 truth draws for the heterogeneous observational cell. Both cells use cross-fitted inference scores. Direct coverage is $.945 \pm .0072$ under randomized treatment with known propensity and $.949 \pm .0070$ in the bounded-overlap observational design; rectangle coverage is .948 and .954. Maximum squared standardized-influence shares are .104 and .466. Pointwise familywise coverage is .622 and .694.

We then probe finite-sample stress by varying nuisance specification, group size, family size, tails, and overlap. Both-misspecified nuisances, a rare group, and the larger family show .905–.915 direct coverage. Heavy-tailed outcomes, weak overlap, and the rare group have maximum leverage .957, .957, and .939, respectively. These cells identify where more observations, fewer frozen contrasts, stronger overlap, or higher-order calibration may be needed.

Score-level layer isolation. Finally, a generated-score experiment separates calibration from nuisance learning. Every method receives the same score matrix, four frozen subgroups, utility rows, and multiplier seed; the preceding raw-data studies retain nuisance-learning variation.

Regime	(n, M, L)	Direct FW	Rectangle FW	Pointwise FW	Direct width	Width reduction
Near null	$(400, 4, 21)$	.948	.968	.513	.489	14.9%
Local	$(400, 4, 21)$	.950	.963	.480	.481	16.3%
Separated	$(400, 4, 21)$	.923	.948	.510	.477	16.9%
Larger n	$(800, 4, 21)$	.953	.963	.523	.344	16.2%
Higher M	$(800, 12, 21)$	.948	.995	.365	.275	39.6%
Many utilities	$(800, 4, 101)$	.943	.973	.485	.346	15.6%
Student-t_3	$(800, 4, 21)$	.950	.980	.518	.329	16.5%

Table 7: Score-level calibration with 400 independent repetitions per regime and 999 Rademacher multipliers.

F.3 Final selector versus distributional learners

We next connect these components by benchmarking the final representation against distributional learners. Within a repetition, we give every method the same realized data, outer train/test split, truth, profiles, and $L = 31$ utility rows; internal training-row use follows two declared panel contracts. The joint learner encodes every observed binary state and uses sparse multiplication over all observed classes. In the natural-runtime panel, we use method-appropriate budgets. In the matched-row panel, we equalize construction rows and nominal per-model stage caps while allowing total model count and wall-clock compute to remain method-specific. Concretely, the final selector fits candidates on 840 construction rows and uses another 360 held-out rows for assembly and guarding, whereas each distributional baseline fits on the 840 construction rows. Thus, the matched-row panel matches fit rows while retaining each method's native total data use.

K	Final−joint	Paired 95% CI	Final win rate
2	0.0279	[0.0164,0.0395]	0.10
4	0.0312	[0.0273,0.0350]	0.00
8	0.0338	[0.0312,0.0363]	0.00
12	0.0245	[0.0203,0.0286]	0.00
16	0.0135	[0.0079,0.0191]	0.10

Table 8: Paired natural-runtime utility-RMSE contrasts; positive values favor the joint learner. Intervals use the Monte Carlo repetition as the unit.

K	Final	Shared	Private	Coordinate	Joint	AR chain
2	0.1037	0.1401	0.1192	0.0861	0.0758	0.1299
4	0.0902	0.1026	0.0948	0.0863	0.0591	0.1004
8	0.0853	0.0826	0.0772	0.0866	0.0515	0.0793
12	0.0796	0.0740	0.0717	0.0866	0.0552	0.0712
16	0.0764	0.0683	0.0648	0.0872	0.0629	0.0634

Table 9: Utility-RMSE means for all six methods.

K	Final	Shared	Private	Coordinate	Joint	AR chain
2	0.1037	0.1401	0.1192	0.0861	0.0876	0.1259
4	0.0902	0.1026	0.0948	0.0863	0.0692	0.1004
8	0.0853	0.0826	0.0772	0.0866	0.0589	0.0800
12	0.0796	0.0740	0.0717	0.0866	0.0628	0.0729
16	0.0764	0.0683	0.0648	0.0872	0.0661	0.0653

Table 10: Construction rows and nominal per-model stage caps are matched in this panel.

K	Final fit	Final query	Joint fit	Joint query	Final MiB	Joint MiB
2	1.72	0.028	0.22	0.232	0.49	0.76
4	4.40	0.036	0.22	0.266	1.08	1.55
8	9.24	0.046	0.24	0.393	2.04	13.56
12	14.22	0.058	0.25	0.482	2.98	32.87
16	17.52	0.073	0.28	0.589	3.94	38.68

Table 11: Natural-runtime seconds and explicit peak-memory proxies. Memory is reported through explicit buffer and model accounting.

Under these two contracts, all selected models for $K \leq 16$ are zero, constant, or shared. Distributional models lead in predictive accuracy and fit time, while the basis representation uses less memory and supports faster repeated utility queries by $K = 16$. For $K = 2, 4, 8, 12, 16$, the zero/constant/shared counts are respectively $11/2/7$, $12/4/4$, $11/3/6$, $7/2/11$, and $5/3/12$; aggregated over the 100 selections this is 46% zero, 14% constant, and 40% shared.

Appendix H.2 reports the formal 20-repetition $K = 24, 32$ extension under the same paired contract.

F.4 IST AUDIT DETAILS

Frozen target and outcome orientation. We prespecify the overall randomized cohort as the primary target before inspecting audit outputs. Of 19,435 randomized records, we exclude two with unknown day-14 status and analyze the remaining 19,433 records. Treatment 1 is aspirin. We write A for alive by day 14, S for free of recurrent ischemic or unknown stroke, and M for any cerebral or major non-cerebral bleed. We use the death-anchored higher-is-better basis $b = (A, AS, -AM)$, so death maps to zero even if an event-form field is coded event-free. We form AIPW scores from the known randomization propensity and cross-fitted outcome regressions. We do not redistribute the raw trial data.

Full-simplex vertex	Estimate	SE	95% simultaneous interval
Alive	.00377	.00386	$[-.00507, .01261]$
Alive and recurrent-stroke-free	.01278	.00430	$[.00291, .02265]$
Negative alive-major-bleed indicator	$-.00540$	.00157	$[-.00899, -.00181]$

Table 12: Broad mathematical coefficient simplex on the death-anchored basis; direct max-t critical value 2.2923. The broad simplex serves as a mathematical sensitivity region.

On this common outcome orientation, the same-disease sensitivity family uses

$$c(u) = \frac{(u/(1-u),\ 1,\ 1/5)}{u/(1-u) + 1 + 1/5}, \qquad u \in \{0.05, 0.10, 0.20\},$$

equivalently, up to positive row-specific scale,

$$U_u = A - (1-u)A(1-S) - \frac{1-u}{5}AM.$$

This construction enforces $U(\text{death}) = 0$, a severe-stroke survivor ratio u, and one stroke loss equal to five major-bleed losses. We treat the resulting analyst-specified rows as fixed sensitivity coefficients constructed from the two elicitation quantities.

Stroke anchor	Estimate	SE	Direct 95% interval	Rectangle interval
.05	.009497	.003577	[.002487, .016507]	[.000632, .018362]
.10	.009242	.003570	[.002245, .016238]	[.000376, .018108]
.20	.008718	.003561	[.001738, .015698]	[−.000150, .017586]

Table 13: Same-disease finite sensitivity family. Direct max-t retains the full utility covariance; the coefficient rows are fixed rather than sampled.

The minimum pairwise utility-effect correlation is .9976, so direct max-t exploits strong covariance. For a separate exploratory question, we use a hash split with 9,837 participants for discovery and 9,596 for confirmation. The high-score confirmation group supports aspirin across all three rows, while the complement is evidence-uncertain; the prespecified overall target remains unchanged.

F.5 ACTG 175 REPLICATION AND LEVERAGE DIAGNOSTIC

We assess portability beyond IST in the frozen AZT-versus-AZT+ddI contrast, which contains 1,054 of 2,139 released ACTG 175 records (Hammer et al., 1996); the remaining 1,085 records belong to other randomized arms. Kaplan–Meier event-free survival at frozen 20- and 96-week horizons gives differences .0150 $[.0018, .0282]$ and .1281 $[.0740, .1823]$. Joint Greenwood covariance and Gaussian max-t provide simultaneous inference under independent censoring. We use this analysis as a second-trial implementation check for the survival-outcome extension. In a separate AIPW subgroup stress, the maximum single-observation squared-influence share is 0.206, and the subgroup audit returns evidence uncertainty.

F.6 ARTIFACT AND STATISTICAL CONTRACT

We serialize configurations, repetition-level CSVs, derived summaries, figures, and checksum manifests in each synthetic result directory, and we fix seeds before execution. We treat Monte Carlo repetitions as independent analysis units and utility rows and basis coordinates as within-repetition quantities. The IST and ACTG directories contain participant-free derived tables and metadata, while the separately licensed CC0 TPT bundle contains public respondent-level source and derived rows. We access IST source data through the checksum-pinned retrieval helper. We regenerate the seven active priority-4 fragments with *scripts/build_core_tables.py* and the five active revision fragments with *scripts/build_submission_tables.py*. Both builders reject smoke or under-replicated primary inputs. We include independent summary/hash checks and a cold-build submission test in the release.

G JOINT PREFERENCE UNCERTAINTY AND NESTED REFINEMENT

G.1 EXTERNAL CONFIDENCE LAYER

We represent external preference uncertainty with an i.i.d. joint sample $C_1, \dots, C_{n_p} \in \mathbb{R}^M$ having mean $c_\star$, sample mean $\bar{C}$, and covariance S. Our implementation allows S to be rank deficient, as when normalized coefficients lie on an affine simplex. If $r = \operatorname{rank}(S)$ and sampling is Gaussian in a fixed r-dimensional affine subspace, Hotelling's construction gives the exact ellipsoid (Hotelling, 1931)

$$\widehat{\mathcal{E}}_\beta = \left\{ c : (c - \bar{C})^\top S^+ (c - \bar{C}) \le h_{1-\beta}/n_p,\ c - \bar{C} \in \operatorname{range}(S) \right\},$$
$$h_{1-\beta} = \frac{r(n_p - 1)}{n_p - r} F_{r, n_p - r; 1-\beta}.$$

For direction d, its exact support limits are $d^\top \bar{C} \pm \sqrt{h_{1-\beta} d^\top S d / n_p}$. Intersecting finitely many corresponding halfspaces gives an outer polyhedral approximation without reducing coverage. We then intersect known scientific constraints guaranteed to contain $c_\star$; statistically estimated refinements spend their own β budgets. The resulting frequentist region retains the joint covariance structure.

G.2 One causal calibration for the refinement path

We next calibrate causal contrasts over the external region's vertices. For each $v \in \mathcal{V}_\star$ and frozen group g, we project the complete basis influence vector before studentization,

$$T^* = \max_{g,v} \frac{|n^{-1/2} \sum_i \xi_i v^\top \widehat{\phi}_{ig}|}{\{n^{-1} \sum_i (v^\top \widehat{\phi}_{ig})^2\}^{1/2}}.$$

We add $2\varepsilon_u$ to the resulting intervals only when a valid utility-basis misspecification allowance is supplied; otherwise the intervals cover the basis projection. Because intersections may create additional vertices, we include every vertex along the reported refinement path in the calibrated master family. Our software deduplicates the union of vertices at every reported level, verifies adjacent containment, and checks deterministic monotonicity of the point envelopes.

Proof of Proposition 2. Let $\mathcal{E}_p = \{c_\star \in \widehat{\Omega}_\beta\}$ and let $\mathcal{E}_c$ be simultaneous coverage of all master-family contrasts. Their complement probabilities are at most β and $\alpha + o(1)$. On $\mathcal{E}_p \cap \mathcal{E}_c$, a positive lower endpoint at every vertex implies a positive contrast for every convex combination, including $c_\star$; the negative case is symmetric. Hence a wrong robust sign lies in $\mathcal{E}_p^c \cup \mathcal{E}_c^c$, and the union bound gives $\alpha + \beta + o(1)$. □

G.3 Synthetic joint-preference validation

We validate both uncertainty layers with joint rank-one Gaussian samples on the affine hyperplane $c_1 + c_2 = 1$, namely $C_i = (W_i, 1 - W_i)$ with $W_i \sim N(0.62, 0.08^2)$. Individual Gaussian draws need not be nonnegative; the broad coefficient simplex is instead a known constraint containing the mean target $c_\star = (0.62, 0.38)$. We construct the external region as a 95% Hotelling outer polytope using 12 support directions and impose the simulation-known constraint $c_1 \geq 0.61$ in a final region. Four causal regimes encode robust Treatment 1, robust Treatment 0, preference sensitivity, and a near boundary. We run 100 repetitions per regime, with $n_p = 80$, causal $n = 400$, $\alpha = \beta = .05$, and 499 Rademacher draws. The master family has 4.82 unique vertices on average.

Event	Probability	MC s.e.
External region contains $c_\star$	.9625	.0095
Causal master family covers	.9475	.0112
Both layers cover	.9125	.0141
At least one layer fails	.0875	.0141
Adjacent containment diagnostic	1.0000	.0000
Point-envelope monotonicity	1.0000	.0000

Table 14: Two-layer validation over 400 independent repetitions. The nominal union-failure bound is $\alpha + \beta = .10$.

Across 400 runs, we observe no false non-abstaining certificates or wrong robust signs. Nested containment and point-envelope monotonicity hold in every repetition. In the preference-sensitive regime, replacing the broad simplex by the external Hotelling region yields Treatment 1 in 83% of runs; intersecting that region with $c_1 \geq 0.61$ raises the rate to 91%. Near the boundary, 95–98% of runs abstain.

H End-to-End and Preference-Evidence Audits

H.1 Full-pipeline adaptation and decision power

To evaluate the complete pipeline, we begin every repetition with raw observed data and construct cross-fitted scores within designated sample roles. Before fitting, we partition independent units into construction, assembly, guard, and inference roles with fractions $.50, .15, .15, .20$. We cross-fit AIPW nuisances and fit shared and private forests on the construction role, freeze the coordinate vector on assembly data, and compare the complete zero, constant, shared, coordinate, and private

Structure	Signal	Complex admitted	Utility RMSE	FW coverage	Power: all	Power: $X_0 > 0$	Power: $X_1 > 0$
Shared	Weak	0.00 [0.00,0.07]	0.090	0.92	0.00 [0.00,0.07]	0.00 [0.00,0.07]	0.00 [0.00,0.07]
Shared	Strong	0.04 [0.01,0.13]	0.203	0.94	0.06 [0.02,0.16]	1.00 [0.93,1.00]	0.18 [0.10,0.31]
Private	Weak	0.00 [0.00,0.07]	0.101	0.98	0.00 [0.00,0.07]	0.00 [0.00,0.07]	0.00 [0.00,0.07]
Private	Strong	1.00 [0.93,1.00]	0.166	0.98	1.00 [0.93,1.00]	0.94 [0.84,0.98]	0.94 [0.84,0.98]
Mixed	Weak	0.00 [0.00,0.07]	0.086	0.90	0.00 [0.00,0.07]	0.00 [0.00,0.07]	0.00 [0.00,0.07]
Mixed	Strong	0.88 [0.76,0.94]	0.157	0.96	1.00 [0.93,1.00]	0.84 [0.71,0.92]	0.78 [0.65,0.87]
No heterogeneity	Weak	0.00 [0.00,0.07]	0.021	0.96	–	–	–
No heterogeneity	Strong	0.00 [0.00,0.07]	0.041	0.96	1.00 [0.93,1.00]	1.00 [0.93,1.00]	1.00 [0.93,1.00]

Table 15: Prespecified full-pipeline results from raw $(X, A, \boldsymbol{Y})$, 50 independent repetitions per cell. Brackets are 95% Wilson intervals across repetitions. Power is computed separately for each fixed subgroup; overlapping groups are never treated as independent replicates. Familywise coverage is over all 3×12 frozen contrasts.

vectors on guard data under the declared utility geometry. On the inference role, we independently refit nuisances, compute target-group basis influences, and apply direct Gaussian multiplier max-t over the full frozen group–utility family. With this pipeline fixed, every repetition passes the frozen eligibility thresholds: the minimum group count is 206, treatment is randomized at .5, and the maximum standardized influence concentration is .110. Applying the operational gate therefore leaves every label unchanged.

We cross shared, private, mixed, and null task structures with weak and strong signals. Each of the eight cells has 50 independent repetitions, $n = 2400$, $M = 8$, 12 fixed utility rows, three frozen overlapping groups, and 999 multiplier draws. We give the shared and private candidates the same rows and per-forest hyperparameters; the private candidate fits one forest per coordinate and consequently uses more total compute. We define "complex-admission power" as the probability that the coordinate or private candidate passes the vector guard in private/mixed cells. This quantity differs from both selection probability and four-way decision power. We compute the latter as the fraction of population group–utility families that receive the correct certificate among those decidable at the prespecified practical tolerance, and report familywise coverage and FWER unconditionally over the complete inference family.

H.2 Formal $K = 24, 32$ extension of the distributional boundary

We next extend the paired representation benchmark to $K = 24, 32$ under the same contract as the $K \leq 16$ study: within a repetition, every method receives the same realized data, construction rows, test profiles, sparse-basis truth, and frozen utility rows. We distinguish the number of clinical outcomes K, basis coordinates M, and post-fit utilities L. The enumerated-state learner uses all observed bit-coded classes through sparse multiplication. We measure memory with model-plus-query buffers and report fit time on the same machine.

K	Final	Joint	AR	Δ(F–J) [95%]	Δ(F–AR) [95%]	Best
24	0.0709	0.0673	0.0563	0.0035 [-0.0020,0.0092]	0.0145 [0.0094,0.0201]	AR chain (0.0563)
32	0.0683	0.0686	0.0547	-0.0003 [-0.0052,0.0055]	0.0136 [0.0089,0.0192]	AR chain (0.0547)

Table 16: Natural-runtime utility RMSE in the formal high-K extension, with 20 paired repetitions per dimension. We report repetition-bootstrap intervals for the joint-state and AR-chain comparators.

K	Final fit	Joint fit	Final query	Joint query	Final MiB	Joint MiB
24	21.50	0.27	0.076	0.574	5.85	40.33
32	31.92	0.27	0.111	0.648	7.76	40.82

Table 17: Recorded fit time, query time, and model-plus-query-buffer memory proxies in the high-dimensional extension.

H.3 Inference eligibility and unconditional operating characteristics

We next assess inferential reliability with the frozen protocol `diagnostic_fail_closed_v1` on the archived 200- and 1,000-repetition inference CSVs. A repetition passes only when every group has at least $\max\{40, \lceil .05 n_{\text{inf}} \rceil\}$ observations, fitted propensities lie in $[.025, .975]$, and maximum standardized influence concentration is no larger than $.50$. We assign evidence uncertainty to a repetition that fails an eligibility check while leaving its interval endpoints unchanged. We report unconditional familywise coverage together with the raw and gated decision rates.

Regime	Reps	Eligible	FW cover	Raw decide	Gated decide
Randomized reference	1000	1.000	0.945	0.064	0.064
Observational reference	1000	1.000	0.949	0.079	0.079
Both nuisances wrong	200	1.000	0.910	0.230	0.230
Weak overlap	200	0.000	0.965	0.007	0.000
Rare group	200	0.000	0.915	0.043	0.000
Heavy tail	200	0.790	0.960	0.083	0.068
Larger family	200	0.155	0.905	0.056	0.008
Randomized, 200-run grid	200	1.000	0.965	0.145	0.145
Randomized external nuisance	200	1.000	0.960	0.142	0.142
Calibration constant	200	1.000	0.940	0.052	0.052
Observational bounded	200	1.000	0.975	0.072	0.072
Larger n	200	1.000	0.935	0.568	0.568
Propensity wrong	200	1.000	0.925	0.155	0.155
Outcome model wrong	200	0.915	0.980	0.087	0.080

Table 18: Complete operational eligibility audit for direct max-t. Coverage is unconditional; raw and gated decision rates are reported.

The diagnostic flags all weak-overlap and rare-group runs, 21% of heavy-tail runs, and 84.5% of enlarged-family runs.

H.4 Same-respondent HB preference evidence for tuberculosis preventive treatment

Having characterized the causal audit, we next analyze a public discrete-choice experiment (DCE) of 400 adults receiving HIV care in Kampala. Of these respondents, 394 passed the study's quality checks and have complete vectors of hierarchical-Bayes part-worth posterior means over duration, frequency, pills per dose, ART adjustment, mild adverse-event risk, moderate/severe adverse-event risk, and no treatment (Aschmann et al., 2026a;b). The raw response vectors, design, codebook, and part-worth file are distributed under CC0. Our loader rejects a file unless its SHA-256 matches the frozen Dryad release.

We compare the published profiles: 3HP is three months, weekly, four pills, no ART adjustment, and 11.5%/6.0% mild/moderate-or-severe adverse-event risk; 6H is six months, daily, two pills, no adjustment, and 36.1%/8.2% risk. Pill counts and both adverse-event probabilities fall between elicited levels, so we use prespecified piecewise-linear interpolation in the primary analysis and map to the nearest elicited level in sensitivity analysis. We resample complete respondent vectors in the bootstrap, preserving empirical cross-attribute dependence, and condition intervals on the deposited posterior means. A marginal-preserving permutation ablation isolates the contribution of joint coefficient dependence.

Nested empirical set	n	3HP share	$P(U_{3HP} > U_{6H})$	Contrast support
Ω_0: all rows	394	0.947	0.982	[-1.94,7.78]
Ω_1: selected 3HP-like	338	0.954	0.985	[-1.94,7.78]
Ω_2: selected and accepted	335	0.960	0.985	[-1.94,7.78]

Table 19: Retrospective nested refinement within the same DCE sample. Each row reports the empirical contrast support for its retained respondents.

Published WHIP3TB outcome	3HP	6H	Published effect [95% CI]
Completion	90.40	50.50	adjusted risk ratio 1.78 [1.61,1.96]
TB incidence, 0–12 months	1.74	1.12	adjusted hazard ratio 1.60 [0.58,4.42]
Mortality, 0–12 months	0.57	0.28	unadjusted hazard ratio 2.07 [0.28,15.50]
AE discontinuation	0.70	1.20	adjusted odds ratio 0.58 [0.22,1.94]

Table 20: Frozen aggregate randomized evidence from WHIP3TB Table 2. The 3HP column pools the two first-year 3HP groups exactly as reported there; units differ by row.

For clinical context, WHIP3TB provides arm-level randomized estimates for completion, tuberculosis incidence, mortality, and adverse-event discontinuation (Churchyard et al., 2021). We report these disease- and regimen-matched clinical estimates alongside the DCE analysis; under the prespecified four-way audit, this evidence configuration returns evidence uncertainty.

Within the DCE sample, we finally describe two nested respondent subsets, first retaining respondents who selected a 3HP-like profile and then those who both selected and accepted it. The predicted 3HP share increases from .947 to .954 and .960, while each nested support continues to contain both positive and negative 3HP–6H utility contrasts.

H.5 Artifact contract

We store each analysis configuration, raw run-level CSVs, derived summaries, and a manifest of byte sizes and SHA-256 digests. Our submission-table generator reads those summaries and rejects missing columns or nonformal repetition counts. We include the verified public DCE release and participant-free IST and ACTG summaries in the archive; controlled-access WHIP3TB participant data are not included.

ADDITIONAL DISCUSSION AND LIMITATIONS

Our results place the structured basis and distributional approaches in different operating regimes. Distributional learners achieve lower predictive RMSE in our benchmarks, whereas the basis representation provides low-order attribution, smaller high-dimensional memory requirements, and fast post-fit utility queries. Turning from representation to certification, strong, well-separated effects yield high certification power, and direct max-t bands improve substantially on rectangle propagation. Near the boundary, the four-way audit converts unresolved signs to evidence uncertainty rather than a treatment conclusion.

Several boundaries define the current guarantees. We certify prespecified target-group averages; profile-level surfaces remain exploratory unless target construction uses an additional independent split. Within this target scope, the results require causal identification, a suitable outcome basis, and a valid utility family. Without a valid approximation radius, the target is the basis projection. The eligibility gate diagnoses empirical fragility rather than nuisance-model correctness: dual nuisance misspecification can pass, and inference conditional on passing would require a selective-inference analysis. Separately, our direct max-t theory covers a fixed finite group–utility family. Growing or continuously indexed families require additional Gaussian-approximation, covariance, and anti-concentration arguments. Finally, the 0.923 coverage in the separated score-level experiment marks a finite-sample regime that may need larger samples or stronger calibration; the observational reference covers only the bounded-overlap design studied here.

The empirical applications validate complementary parts of the framework rather than a single joint outcome–preference analysis. In the tuberculosis component, the intervals condition on the deposited hierarchical-Bayes posterior means because posterior draws are unavailable. Moreover, the DCE attributes and WHIP3TB aggregate outcomes do not form a common patient-level object, and transport would require covariates unavailable in the public DCE release. In the IST component, the fixed sensitivity rows combine separately elicited anchors through an additive utility model and therefore do not represent a sampled joint preference distribution. Existing stroke-preference studies also do not match the IST death, recurrence, and major-bleeding vector (Rusin et al., 2023; Slot & Berge, 2009; Firth et al., 2026; Zhang et al., 2026).

A study measuring compatible outcomes and preferences in the same population would allow us to propagate both uncertainty sources in one causal audit. Further work can also develop growing-family theory, confirmatory inference for data-adaptive target groups, and cluster-robust versions for multi-center studies.