ROBUST ACTIVE STATISTICAL INFERENCE: SEPARATE GEOMETRY FROM UNCERTAINTY

Anonymous authors
Paper under double-blind review

ABSTRACT

Active statistical inference reduces labeling costs by combining inexpensive proxy labels with selectively queried ground-truth labels for a downstream inferential target. The sampling policy determines which labels to acquire based on the estimated error of the proxy labels. However, such policies can be sensitive to misspecification. When the error of proxy labels is severely underestimated in a region where it is in fact large, the resulting query probabilities can be small, substantially inflating the estimator's variance and confidence-interval width. To mitigate this issue, we introduce **R**esidual-robust **A**ctive **S**tatistical **I**nference (RASI), which exploits a structural decomposition of the optimal sampling policy into a target-specific loading determined by the geometry of the estimand and an unobserved residual-risk component approximated by an uncertainty proxy. RASI preserves the former and robustifies only the latter. We formulate policy design as a distributionally robust optimization problem over plausible residual risks and derive minimax sampling policies over an ambiguity set. Our analysis shows that mixing with uniform sampling, a common robustness strategy in active statistical inference, is generally not the appropriate robust fallback because it discards target-specific geometry. We extend RASI to smooth functions of parameters defined by outcome-affine estimating equations, covering linear and logistic regression coefficients and log-odds ratios. Across synthetic and real-data experiments, RASI protects against misspecification of the uncertainty proxy while retaining the efficiency gains of active sampling when the proxy is informative.

1 INTRODUCTION

Many scientific studies observe inexpensive features or model predictions for a large population, while ground-truth outcomes or labels are costly to acquire. Many fields use predictive models as proxies for true outcomes. Examples include radiologist-assessed findings in chest radiographs (Irvin et al., 2019), gene expression measured by laboratory assays (Vaishnav et al., 2022), human coding of political text (Egami et al., 2023), and human evaluation of generated summaries and question-answer outputs (Fisch et al., 2024). Prediction-powered inference (PPI) is a framework for valid statistical inference based on predictive models and a small number of true labels (Angelopoulos et al., 2023a).

Active statistical inference (ASI) goes one step further by asking which outcomes are most informative for the inferential target (Zrnic & Candès, 2024). It assigns larger query probabilities to observations whose proxy-label errors are expected to matter more for the inferential target, while inverse-probability weighting preserves validity. Uniform sampling gives a PPI-type estimator in this sampling formulation. The sampling policy relies on an *uncertainty proxy*, a score intended to approximate the predictive model's error. When the uncertainty proxy accurately reflects conditional proxy-label error, active sampling can substantially reduce variance at the same labeling budget compared with uniform sampling.

However, these gains can be fragile. If proxy-label error is severely underestimated in an important region, a plug-in policy may assign that region a very small query probability, amplifying the inverse-probability-weighted correction and, consequently, the variance of the resulting estimator. Such misspecification need not invalidate asymptotic coverage, but it can make the resulting confidence interval much wider. The central robustness problem is therefore: how can we protect an active

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.

sampling policy against misspecification of its uncertainty proxy without discarding the information that makes active sampling useful?

A common safeguard is to interpolate between the active policy and uniform sampling (Zrnic & Candès, 2024). More recently, Li et al. (2025) use robust optimization to select a point along a prescribed interpolation path between these two policies. Such approaches provide protection against misspecification, but restrict the robust policy to a path specified in advance and treat uniform sampling as the natural fallback. More related work on PPI, active inference, and distributionally robust optimization is discussed in Appendix C.

The use of uniform sampling as a robust endpoint seems natural. However, for a general inferential target, prediction errors at different covariate values can have substantially different effects on the estimator. For example, in regression, an error at a high-leverage covariate value may be far more consequential than an equally large error elsewhere. Shrinking toward uniform sampling ignores known target-specific structure. This suggests a different principle: *robustify only what is uncertain, but preserve what is known from the inferential target*.

We introduce **R**esidual-robust **A**ctive **S**tatistical **I**nference (RASI), which implements this principle through robust optimization (Ben-Tal & Nemirovski, 1998). The key observation is that the policy-dependent component of the asymptotic variance separates into two conceptually distinct quantities: a target-specific loading determined by the geometry of the estimand, and an unobserved residual-risk component that is accessible only through an uncertainty proxy. RASI preserves the former and places an ambiguity set only around the latter. We then optimize directly over the full admissible class of sampling policies instead of restricting the solution to a pre-specified interpolation path.

This formulation leads to a qualitatively different robust fallback. At zero ambiguity, RASI recovers the usual plug-in policy. As the ambiguity radius increases, the policy increasingly protects covariate regions in which under-sampling would be most consequential. Under an L_2 ambiguity set, the large-ambiguity limit allocates labels proportionally to $\Lambda^{2/3}$, up to clipping induced by the labeling budget, where Λ denotes the target-specific loading. As a result, the robust limit is determined by the geometry of the inferential target.

The practical difference can be substantial. Figure 1 shows a linear-regression experiment on U.S. Census data (Ding et al., 2021) in which the uncertainty proxy is fit using only 500 labeled observations. We define relative efficiency as the squared ratio of the CI width under uniform sampling to that under a given method, which can be interpreted as an approximate effective-sample-size gain. RASI maintains a relative efficiency of approximately 1.7 across the full range of evaluation labeling budgets. In contrast, the misspecified plug-in policy deteriorates rapidly and eventually attains less than half the efficiency of uniform sampling. RASI also consistently outperforms oracle-tuned linear and geometric mixtures, whose mixing weights are selected in hindsight to minimize CI width. Meanwhile, all methods achieve approximately nominal or conservative coverage. This contrast highlights the role of RASI: it protects statistical efficiency against a misspecified uncertainty proxy without sacrificing the target-specific structure that makes active sampling effective.

Our contributions are threefold. First, we build on the separation between target-specific geometry and residual risk in active inference (Zrnic & Candès, 2024; Li et al., 2025), and propose RASI to robustify only the latter. Second, we derive the resulting sampling policies from a robust optimization problem and characterize their behavior across ambiguity levels, showing in particular that the robust limit is generally geometry-aware rather than uniform. Third, we extend RASI to smooth functions of parameters defined by outcome-affine estimating equations, covering logistic regression coefficients and log-odds ratios, and demonstrate its benefits in synthetic and real-data experiments.

2 SAMPLING GEOMETRY: LINEAR REGRESSION

We start from the example of linear regression problem to demonstrate our ideas. Let $(X_i, Y_i)_{i=1}^n$ be i.i.d. from $\mathbb{P}_X \times \mathbb{P}_{Y|X}$, with all features $X_i \in \mathbb{R}^d$ observed and labels Y_i acquired selectively. The estimand is the linear regression coefficient $\theta^* = e_1^\top \beta^* \in \mathbb{R}$, where β^* solves

$$\mathbb{E}[X\{Y - X^\top \beta^*\}] = 0$$

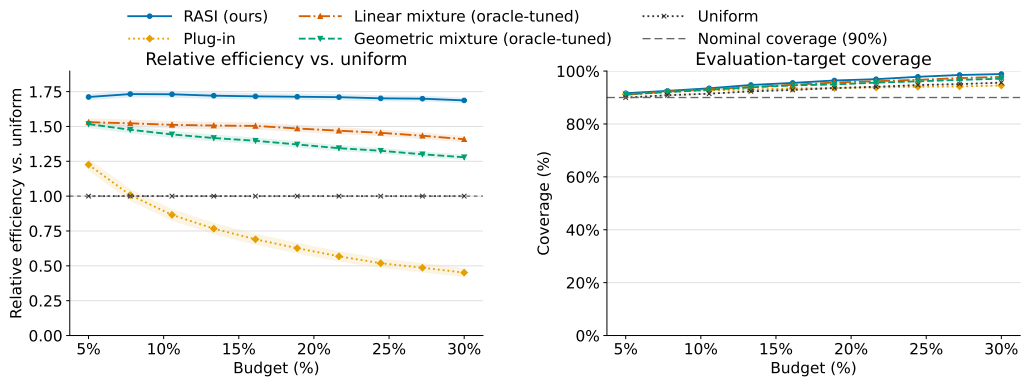

Figure 1: **Linear regression on U.S. Census data with a misspecified uncertainty proxy.** The uncertainty proxy is fit using 500 labeled observations. *Left:* relative efficiency against uniform sampling as a function of the evaluation labeling budget. *Right:* empirical confidence-interval coverage for the same methods and budgets. Linear and geometric mixtures are oracle-tuned in hindsight to minimize CI width.

We assume $Q := \mathbb{E}(XX^\top)$ is nonsingular. This population linear projection does not require a correctly specified linear model. We first treat the predictor f and policy π as fixed. In practice, they can be learned from another independent dataset.

Active inference estimator. Each label Y_i is queried with probability $\pi(X_i) \in [\underline{\pi}, 1]$, where $\underline{\pi}$ is a small positive number to avoid overflow. With budget fraction $\gamma \in (\underline{\pi}, 1)$, the expected-budget policy class is

$$\Pi_\gamma = \{\pi : \underline{\pi} \le \pi(X) \le 1, \quad \mathbb{E}\{\pi(X)\} = \gamma\}.$$

Let $\xi_i \in \{0, 1\}$ be the indicator of being queries. The corrected pseudo-outcome is

$$\widetilde{Y}_i = f(X_i) + \frac{\xi_i}{\pi(X_i)}\{Y_i - f(X_i)\}. \tag{1}$$

Then the active inference estimator can be defined as $\hat{\theta}_\pi = e_1^\top \widehat{Q}^{-1} \frac{1}{n} \sum_{i=1}^n X_i \widetilde{Y}_i$, where $\widehat{Q} = n^{-1} \sum_i X_i X_i^\top$. Because $\mathbb{E}_\xi(\widetilde{Y}_i \mid X_i, Y_i) = Y_i$, $\hat{\theta}_\pi$ is still a consistent estimator.

Variance geometry. Let $\Lambda(x) = (e_1^\top Q^{-1} x)^2$, and $v_f(x) = \mathbb{E}[\{Y - f(X)\}^2 \mid X = x]$. Under the regularity conditions in Appendix A.1, $\sqrt{n}(\widehat{\theta}_\pi - \theta^*)$ is asymptotically normal with variance

$$V(\pi) = V_{\text{full}} + \mathbb{E}\left[\left\{\frac{1}{\pi(X)} - 1\right\} \Lambda(X) v_f(X)\right], \tag{2}$$

where $V_{\text{full}} = \mathbb{E}[\Lambda(X)\{Y - X^\top \beta^*\}^2]$ is the usual variance for linear coefficient estimator. The loading Λ depends only on the coefficient direction and P_X; the residual risk v_f must be approximated using labeled data. A consistent variance estimator gives confidence intervals for the superpopulation target. Appendix A.1 also gives the design-based formulation for a fixed-pool target.

Oracle allocation. For fixed v_f with $\Lambda v_f > 0$ almost surely, the oracle policy minimizing $V(\pi)$ is

$$\pi^*(x) = \text{clip}_{[\underline{\pi}, 1]}\left\{\sqrt{\frac{\Lambda(x) v_f(x)}{\lambda}}\right\}, \tag{3}$$

where $\lambda > 0$ enforces the budget condition $\mathbb{E}[\pi^*(X)] = \gamma$ and $\text{clip}_{[a,b]}(t) = \min\{b, \max(a, t)\}$. It is suggested in Zrnic & Candès (2024) to use an estimate $\hat{v}_f(x)$ to obtain a plug-in estimate of $\pi^*(x)$. Thus an underestimated residual risk can cause under-sampling where errors have a large effect on the target. In the next section, we develop RASI to protect this residual component without discarding the loading.

3 RESIDUAL-ROBUST ACTIVE SAMPLING

Following notations in the last section, we further assume a nonnegative estimate/proxy $\hat{v}$ of the model's residual risk. RASI makes two distinct design choices. First, it treats the target loading Λ as part of the inferential geometry and places ambiguity only on the residual-risk component:

$$\mathcal{V}_\delta(\widehat{v}) = \left\{v \geq 0 : \|v - \widehat{v}\|_{L_2(P_X)} \leq \delta\right\}, \qquad \delta \geq 0.$$

Second, given this structured ambiguity set, it optimizes the resulting worst-case criterion over the full admissible policy class Π_γ, rather than restricting the policy to a prescribed interpolation path. Specifically, we solve the minimax problem

$$\pi_\delta \in \arg\min_{\pi\in\Pi_\gamma} \sup_{v\in\mathcal{V}_\delta(\widehat{v})} \mathbb{E}\left[\frac{\Lambda(X)v(X)}{\pi(X)}\right] = \arg\min_{\pi\in\Pi_\gamma} \mathbb{E}\left[\frac{\Lambda(X)\widehat{v}(X)}{\pi(X)}\right] + \delta\left\|\frac{\Lambda}{\pi}\right\|_{L_2(P_X)}. \tag{4}$$

The equality follows from Cauchy–Schwarz. We use the inverse-propensity criterion $\mathbb{E}\left[\frac{\Lambda(X)v(X)}{\pi(X)}\right]$ because it is the only term depends on policy π in $V(\pi)$. For fixed v, subtracting $\mathbb{E}\{\Lambda(X)v(X)\} - V_{\text{full}}$ recovers the asymptotic variance.

The L_2 ambiguity set provides a simple global measure of residual misspecification while preserving heterogeneity across covariate values. It also yields the explicit regularized form in (4), making transparent how residual uncertainty interacts with the target loading. Other ambiguity geometries like $L_p, p > 1$ are possible. For alternative loss design, see Appendix A.2.

Geometry-aware fallback At $\delta = 0$, (4) reduces to the plug-in sampling policy based on $\widehat{v}$. As δ grows, the robust penalty increasingly determines the allocation. In the large-ambiguity limit, the optimization reduces to minimizing $\|\Lambda/\pi\|_2$ over Π_γ. Under the usual integrability conditions, and when $\Lambda > 0$ almost surely, the limit sampling policy is

$$\pi_\infty(x) = \text{clip}_{[\underline{\pi},1]}\{c_\infty \Lambda(x)^{2/3}\}, \qquad \mathbb{E}\{\pi_\infty(X)\} = \gamma. \tag{5}$$

Thus, even when the residual-risk proxy becomes completely uninformative, RASI does not generally revert to uniform sampling. Instead, it discards uncertain information about residual risk while retaining the geometry of the inferential target. For an interior budget, uniform sampling is recovered only when Λ is constant. This distinction is a consequence of both ingredients above: the ambiguity set preserves target geometry, and the policy is free to respond to that ambiguity over the full feasible class.

Geometry gain The large-ambiguity limit above suggests a direct way to quantify the value of retaining target geometry. Let $\pi_{\text{unif}}(x) \equiv \gamma$ and define

$$Q_\delta(\pi) := \sup_{v\in\mathcal{V}_\delta(\widehat{v})} \mathbb{E}\left[\frac{\Lambda(X)v(X)}{\pi(X)}\right].$$

Then we have the following proposition.

Proposition 3.1 (Geometry gain of target-aware fallback)**.** *Suppose $\Lambda > 0$ almost surely, $\Lambda \in L_2(P_X)$, and $\mathbb{E}\{\Lambda(X)\widehat{v}(X)\} < \infty$, and suppose the propensity bounds are inactive. Then*

$$\lim_{\delta\to\infty} \frac{Q_\delta(\pi_{\text{unif}}) - \inf_{\pi\in\Pi_\gamma} Q_\delta(\pi)}{\delta} = \Delta_\Lambda := \frac{1}{\gamma}\left[\|\Lambda\|_2 - \{\mathbb{E}(\Lambda^{2/3})\}^{3/2}\right] \geq 0.$$

Moreover, $\Delta_\Lambda = 0$ if and only if Λ is constant almost surely.

Thus, we can define *geometry gain* as Δ_Λ. Consequently, $\delta\Delta_\Lambda$ is the leading-order cost of reverting to uniform sampling rather than retaining target geometry under the robust criterion. This is not literally a difference in worst-case asymptotic variances, since the ambiguity set is applied to the inverse-propensity criterion. However, for any fixed residual risk v, differences in this criterion across policies coincide exactly with differences in asymptotic variance. Hence Δ_Λ can be viewed as a large-ambiguity proxy for the variance benefit of preserving target geometry.

Uniform sampling is therefore geometry-optimal for constant-loading targets, such as mean estimation. For heterogeneous loadings, including regression coefficients, even after discarding all information in the residual-risk proxy, target geometry alone still supports a better robust fallback than uniform sampling.

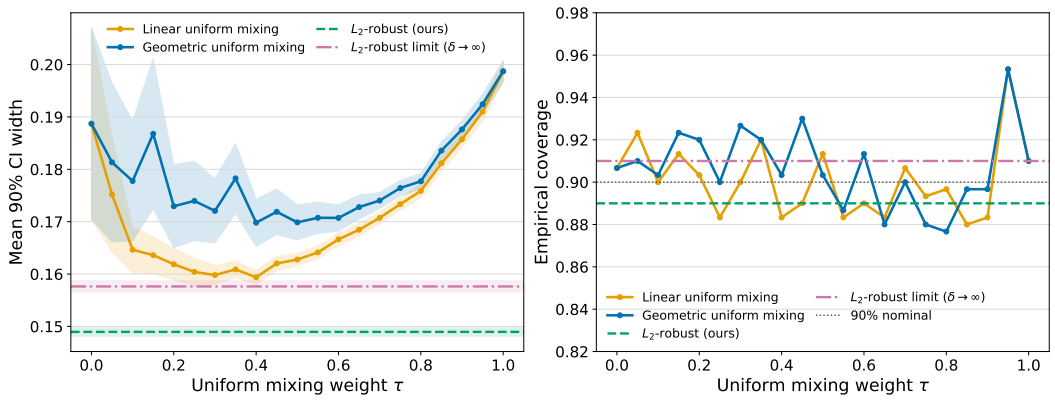

(a) CI width (b) Coverage

Figure 2: Synthetic linear regression. Mean 90% confidence-interval width (left) and empirical coverage (right) as a function of the uniform-mixing weight τ. Arithmetic and geometric mixtures follow prescribed paths between the plug-in and uniform policies. Horizontal lines show RASI and its $\delta \to \infty$ geometry-only limit. Finite-radius RASI achieves the smallest interval width while maintaining approximately nominal coverage.

Relation to uniform mixing and synthetic illustration. Existing robustification strategies often interpolate between a plug-in active policy π_0 and uniform sampling (Zrnic & Candès, 2024; Li et al., 2025), for example

$$\pi_\tau^A = (1-\tau)\pi_0 + \tau\gamma, \qquad \pi_\tau^G = \mathrm{clip}_{[\underline{\pi},1]}\{c_\tau \pi_0^{1-\tau}\gamma^\tau\}, \qquad \tau \in [0,1],$$

where c_τ enforces the budget. These approaches search over a prescribed one-dimensional path, whereas RASI optimizes the robust criterion over the full feasible class Π_γ.

Figure 2 illustrates the resulting distinction in a heteroskedastic linear-regression problem; Appendix B.1 gives the full protocol. Even the $\delta \to \infty$ RASI limit achieves width 0.158, slightly improving on both oracle-tuned mixing paths despite discarding all information in the residual-risk proxy. RASI further reduces the width to 0.149 by retaining useful residual-risk information. The experiment therefore separates the benefit of the geometry-aware fallback from the additional gain obtained from an informative residual-risk proxy.

4 Beyond Linear Regression

The linear-regression development in Section 2 extends to a broad class of targets defined through estimating equations. Let $\beta^\star \in \mathbb{R}^d$ solve

$$\mathbb{E}\{\psi(X, Y; \beta^\star)\} = 0, \qquad \theta^\star = g(\beta^\star), \tag{6}$$

where g is a smooth scalar-valued function.

Outcome-affine assumption. We assume that the estimating equation depends affinely on the outcome:

$$\psi(x, y; \beta) = r(x; \beta) + s(x; \beta)y. \tag{7}$$

The assumption covers all targets considered here, including means, linear and logistic regression coefficients, and marginal log odds ratios. It also allows the same prediction-assisted correction to be used across targets. Replacing Y_i by the corrected pseudo-outcome $\widetilde{Y}_i$ from (1) leads us to

$$\frac{1}{n}\sum_{i=1}^{n} \psi(X_i, \widetilde{Y}_i; \widehat{\beta}) = 0, \qquad \widehat{\theta} = g(\widehat{\beta}). \tag{8}$$

Target loading. To identify the analogue of the regression leverage term, consider the usual first-order linearization of (6). Let $M = \mathbb{E}\left[\frac{\partial}{\partial \beta^\top}\psi(X, Y; \beta^\star)\right]$ be nonsingular. Standard M-estimation theory gives the full-data influence function

$$\varphi_{\text{full}}(X, Y) = -\nabla g(\beta^\star)^\top M^{-1}\psi(X, Y; \beta^\star).$$

Under (7), its dependence on the outcome is linear, with coefficient

$$a(x) = -\nabla g(\beta^\star)^\top M^{-1}s(x; \beta^\star), \qquad \Lambda(x) := a(x)^2. \tag{9}$$

Thus $a(x)$ measures the first-order effect of an outcome perturbation at covariate value x on the target, while $\Lambda(x)$ measures its squared contribution to variance. A detailed derivation is given in Appendix A.4; see also standard M-estimation and influence-function theory (van der Vaart, 1998). By definition of $\widetilde{Y}$, for a fixed sampling policy π, the active influence function becomes

$$\varphi_\pi(O) = \varphi_{\text{full}}(X, Y) + \left\{\frac{\xi}{\pi(X)} - 1\right\} a(X)\{Y - f(X)\}.$$

The correction has conditional mean zero given (X, Y), and therefore

$$V(\pi) = V_{\text{full}} + \mathbb{E}\left[\left\{\frac{1}{\pi(X)} - 1\right\}\Lambda(X)v_f(X)\right]. \tag{10}$$

Here $V_{\text{full}} = \mathbb{E}\{\varphi_{\text{full}}(X, Y)^2\}$ is the full-data asymptotic variance. The same target–residual factorization therefore reappears: Λ describes where prediction errors matter for the target, whereas v_f describes where the predictor is inaccurate. Consequently, the RASI construction in Section 3 applies without modification after replacing the linear-regression loading by (9).

Examples. For mean estimation, take $\psi(y; \beta) = y - \beta$ and $g(\beta) = \beta$. Then $a(x) = 1$ and $\Lambda(x) = 1$, so the target geometry is constant.

For coefficient j of the population logistic projection,

$$\mathbb{E}[X\{Y - \mu_{\beta^\star}(X)\}] = 0, \qquad \mu_\beta(x) = \text{expit}(x^\top\beta).$$

Let $H = \mathbb{E}[\mu_{\beta^\star}(X)\{1 - \mu_{\beta^\star}(X)\}XX^\top]$, then $\Lambda_j(x) = \{e_j^\top H^{-1}x\}^2$.

For a marginal log odds ratio, let Y be binary and let $G \in \{0, 1\}$ be an observed group indicator included in X. Write $p_g = \mathbb{P}(G = g) > 0$ and $\mu_g = \mathbb{E}(Y \mid G = g) \in (0, 1)$. Then log odds ratio is

$$\theta^\star = \text{logit}(\mu_1) - \text{logit}(\mu_0).$$

The resulting loading is $\Lambda(X) = \left(\frac{G}{p_1\mu_1(1-\mu_1)} - \frac{1-G}{p_0\mu_0(1-\mu_0)}\right)^2$. The corresponding corrected estimators and derivations are given in Appendix A.4.

In the above examples, some Λ depend on β^*, in practice, we need to first obtain a rough estimation $\hat{\beta}_0$ and then use plug in estimation to obtain $\widehat{\Lambda}$.

5 RASI IMPLEMENTATION AND INFERENCE

We consider a batch setting in which all evaluation features $\{X_i\}_{i=1}^n$ are observed before any evaluation labels are queried. We also have access to a synthetic label $f(X)$. An independent labeled pilot sample is used to fit the residual-risk proxy $\widehat{v}$ and any pilot-dependent loading components $\widehat{\Lambda}$. The robustness radius is also selected using pilot data only.

With these quantities fixed, the evaluation features are used jointly to estimate covariate-dependent loading components and construct the final sampling probabilities. Thus the final policy may depend on the entire evaluation covariate sample, whereas pilot learning and radius selection use no evaluation data. The full algorithm is given in Algorithm 1. In the following, we explain each component and give a theoretical guarantee.

Empirical policy construction. $\widehat{\Lambda}_r = \widehat{\Lambda}(X_r)$ and $\widehat{v}_r = \widehat{v}(X_r)$, solve

$$\min_{\underline{\pi} \le p_r \le 1, n^{-1}\sum_{r=1}^n p_r = \gamma} \left[\frac{1}{n}\sum_{r=1}^n \frac{\widehat{\Lambda}_r \widehat{v}_r}{p_r} + \delta \left\{ \frac{1}{n}\sum_{r=1}^n \left(\frac{\widehat{\Lambda}_r}{p_r} \right)^2 \right\}^{1/2} \right]. \tag{11}$$

At the pilot-selected radius, this program directly yields the query probabilities $p_1, \ldots, p_n$ for the evaluation batch. They are frozen before labels are queried. It is a convex problem. The same construction is applied separately within the pilot for radius selection.

Consistent variance estimation. Inference can be implemented with the usual sandwich variance estimator applied to the corrected estimating equation. Define

$$\widehat{M} = \frac{1}{n}\sum_{i=1}^n \frac{\partial}{\partial \beta^\top}\psi(X_i, \widetilde{Y}_i; \widehat{\beta}), \qquad \widehat{\Omega} = \frac{1}{n}\sum_{i=1}^n \psi(X_i, \widetilde{Y}_i; \widehat{\beta})\psi(X_i, \widetilde{Y}_i; \widehat{\beta})^\top.$$

Then

$$\widehat{V} = \nabla g(\widehat{\beta})^\top \widehat{M}^{-1} \widehat{\Omega} \widehat{M}^{-\top} \nabla g(\widehat{\beta}) \tag{12}$$

is consistent for the asymptotic variance in (25) under the conditions of Theorem A.4. Equivalently, defining the estimated influence values $\widehat{\varphi}_i = -\nabla g(\widehat{\beta})^\top \widehat{M}^{-1}\psi(X_i, \widetilde{Y}_i; \widehat{\beta})$, we have $\widehat{V} = \frac{1}{n}\sum_{i=1}^n \widehat{\varphi}_i^2$. Thus a $(1-\alpha)$ asymptotically valid confidence interval is $\widehat{\theta} \pm z_{1-\alpha/2}\sqrt{\widehat{V}/n}$. All quantities in (12) are observable: for an unqueried unit, $\widetilde{Y}_i = f(X_i)$, while queried units use the recorded propensity p_i in the correction.

The inferential target remains the superpopulation parameter $\theta^\star$. We first analyze the procedure with its training inputs and radius fixed, and then account for independent pilot learning by conditioning.

Theorem 5.1 (Batch RASI inference, informal)**.** *Fix f, $\hat{v}$ and radius $\delta \ge 0$. Under suitable design-stability, moment, and estimating-equation regularity conditions, the empirical RASI allocation converges to the unique population counterpart $\pi_\delta^\star$ of the fitted design. With a fixed positive propensity floor and conditionally independent Bernoulli queries using the learned probabilities, we have*

$$\sqrt{n}(\widehat{\theta} - \theta^\star)/\sqrt{\widehat{V}} \rightsquigarrow \mathcal{N}(0,1). \tag{13}$$

The formal statement and proof are in Appendix A.5. The proof first establishes empirical allocation consistency, then couples the batch estimating equation to its limiting fixed-policy counterpart.

Independent pilot learning. For an independent pilot of fixed size, condition on the pilot and any independent external training information. The predictor, residual-risk proxy, and pilot-fitted loading components are then fixed, so Theorem 5.1 applies whenever its assumptions hold. If these assumptions hold for almost every training realization, the confidence interval has asymptotic coverage $1-\alpha$ both conditionally and unconditionally.

Pilot-only radius selection. We use pooled out-of-fold pilot risk to select the radius. For pilot observation i, write $r_i = \{Y_i^{\mathrm{pil}} - f(X_i^{\mathrm{pil}})\}^2$. Cross-fit the residual model to obtain nonnegative predictions $\widehat{v}_i^{\mathrm{OOF}}$, and set $\widehat{s} = \{m^{-1}\sum_i (r_i - \widehat{v}_i^{\mathrm{OOF}})^2\}^{1/2}$. For each multiplier t in a fixed finite set $\mathcal{T}$, solve (11) jointly on the full pilot, using full-pilot loadings $\widehat{\Lambda}_i^{\mathrm{pil}}$, residual predictions $\widehat{v}_i^{\mathrm{OOF}}$, and radius $t\widehat{s}$. Denote the resulting probabilities by $\widehat{p}_{t,i}$, and select

$$\begin{aligned} \widehat{J}_{\mathrm{pool}}(t) &= \frac{1}{m}\sum_{i=1}^m \frac{\widehat{\Lambda}_i^{\mathrm{pil}} r_i}{\widehat{p}_{t,i}}, \\ \widehat{t} &\in \arg\min_{t \in \mathcal{T}} \widehat{J}_{\mathrm{pool}}(t), \qquad \widehat{\delta} = \widehat{t}\widehat{s}. \end{aligned} \tag{14}$$

Only the residual predictions are cross-fitted; the scale, loadings, and budget normalization use the full pilot. After selection, refit the residual model on the full pilot and solve (11) on the evaluation covariates at $\widehat{\delta}$. This empirical tuning criterion does not guarantee unbiased risk estimation, selection consistency, or ambiguity-set containment. Corollary A.5 permits arbitrary pilot-measurable radius selection under its assumptions.

Algorithm 1 Residual-robust active statistical inference (RASI)

Require: Independent labeled pilot; evaluation covariates $\{X_i\}_{i=1}^n$; predictor f; target; budget fraction γ; propensity floor $\underline{\pi}$; radius multipliers $\mathcal{T}$; confidence level $1-\alpha$.

1: **Construct pilot candidates.** Cross-fit the residual model, compute full-pilot loadings and scale $\widehat{s}$, and jointly solve the pilot program at radii $t\widehat{s}$, $t \in \mathcal{T}$.
2: **Select and refit.** Minimize the pooled criterion (14), set $\widehat{\delta} = \widehat{t}\widehat{s}$, and refit the residual model on the full pilot.
3: **Construct the evaluation-batch policy.** Using the evaluation covariates, estimate any covariate-dependent loading components and solve (11) at the fixed radius $\widehat{\delta}$ to obtain $p_1, \ldots, p_n$.
4: **Acquire labels.** Freeze the propensities, draw $\xi_i \sim \mathrm{Bernoulli}(p_i)$ conditionally independently, and query Y_i when $\xi_i = 1$.
5: **Estimate and infer.** Form $\widetilde{Y}_i$ using the recorded propensities, solve (8), and report the confidence interval using (12).

6 EXPERIMENTS

We evaluate RASI in two real-data applications: a linear-regression coefficient using U.S. Census data and a marginal log odds ratio using chest X-rays.

In both applications, the evaluation population contains $n = 10{,}000$ observations and the expected evaluation-labeling budget ranges from 5% to 30%. Labels are queried independently with propensity floor $\underline{\pi} = 0.01$. We form nominal 90% confidence intervals for each application. Each configuration is evaluated over 100 outer replications and 300 active-sampling realizations per outer replication.

Methods and metrics. We compare RASI with uniform sampling, the plug-in policy based on $\widehat{v}$, and arithmetic and geometric mixtures between the plug-in and uniform policies. For each mixture family, $\tau \in \{0, 0.1, \ldots, 1\}$; to provide a deliberately favorable benchmark, we select τ in hindsight separately within each outer replication and budget to minimize mean interval width. RASI selects its radius using pooled out-of-fold pilot risk (Appendix B.2), over multipliers $\{0.03, 0.1, 0.3, 1, 3\}$ applied to the out-of-fold residual-model RMSE. RASI uses no evaluation outcomes for policy construction or radius selection. We additionally compare with OPAL (Ma & Candès, 2026) in Appendix B.4 using the same pilot-trained squared-residual model.

Our primary metric is relative efficiency against uniform sampling $\mathrm{RE}_m = \left(W_{\mathrm{unif}}/W_m\right)^2$, where W_m is the mean confidence-interval width for method m. Thus $\mathrm{RE} > 1$ indicates narrower intervals than uniform sampling and can be interpreted as an approximate effective-sample-size ratio. Since the superpopulation target is unknown, we report coverage of the replication-specific full-evaluation target, which can be conservative. For comparison, Appendix B.5 reports desired coverage using nominal 90% design-based confidence intervals constructed directly for the full-evaluation target.

6.1 CENSUS INCOME COEFFICIENT

We use the 2019 one-year California ACS person-level PUMS (Ding et al., 2021). After requiring observed annual personal income, age, and sex, $318{,}215$ observations remain. In each outer replication, we draw disjoint proxy-training, pilot, and evaluation samples. The headline experiment uses $3{,}000$ observations to train the outcome proxy, 500 pilot labels for policy construction, and $10{,}000$ evaluation observations.

We estimate the age coefficient in the unweighted least-squares projection of annual personal income on age and the numeric ACS sex code. Its full-evaluation benchmark is $\theta_{\mathrm{eval}} = e_1^\top \left(\frac{1}{n}\sum_i Z_i Z_i^\top\right)^{-1} \left(\frac{1}{n}\sum_i Z_i Y_i\right)$, where $Z_i = (\mathrm{AGEP}_i, \mathrm{SEX}_i)^\top$. The executed regression contains no intercept. A histogram gradient-boosting model trained on 17 demographic and socioeconomic variables provides the proxy f. On the pilot sample, we estimate $\mathbb{E}[(Y - f(X))^2 \mid X]$ using five-fold cross-fitted XGBoost (Chen & Guestrin, 2016); the final residual model is refit on the complete pilot before constructing the sampling policy.

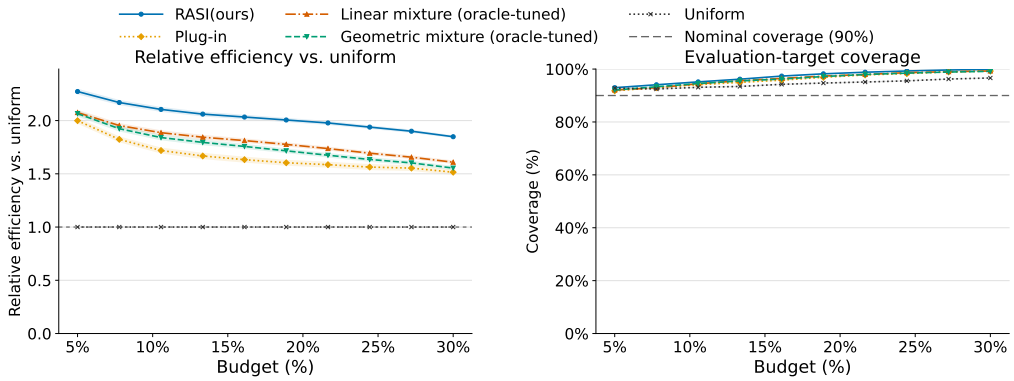

Figure 3: **Chest X-ray marginal log odds ratio.** Relative efficiency against uniform sampling (left) and empirical evaluation-target coverage of nominal 90% intervals (right), using 500 pilot labels. Mixture baselines are oracle-tuned in hindsight.

Figure 1 shows the main result. RASI maintains a relative efficiency of roughly 1.7 throughout the reported budget range, whereas the plug-in policy degrades sharply as the budget increases and eventually becomes substantially less efficient than uniform sampling. RASI also improves on the oracle-tuned arithmetic and geometric mixtures across the reported budgets. All methods retain approximately nominal or conservative coverage.

6.2 Chest X-ray log odds ratio

We estimate the cardiomegaly log odds ratio for images of patients aged at least 40 versus younger patients (Irvin et al., 2019). After requiring a binary cardiomegaly label, a valid model score, and observed age, the analysis contains $19{,}596$ image observations. We define $G_i = \mathbf{1}\{\mathrm{Age}_i \geq 40\}$; the full-evaluation benchmark is $\theta_{\mathrm{eval}} = \mathrm{logit}(\mu_1) - \mathrm{logit}(\mu_0)$, where $\mu_g = \frac{\sum_i \mathbf{1}(G_i=g)Y_i}{\sum_i \mathbf{1}(G_i=g)}$. This is a marginal association rather than a conditional logistic-regression coefficient or causal effect.

The frozen cardiomegaly probability is used directly as f. The headline experiment again uses 500 pilot labels and a disjoint evaluation population of $10{,}000$ observations. A pilot-trained model averages squared proxy residuals within up to ten score-quantile bins per age group. Five-fold out-of-fold predictions guide radius selection; refitting on the full pilot supplies the shared residual-risk estimates for all adaptive methods, including OPAL.

As shown in Figure 3, RASI attains the highest relative efficiency across the reported evaluation budgets, with approximately nominal or conservative evaluation-target coverage. The gain over oracle-tuned mixtures extends the Census finding to a nonlinear target with pilot-estimated loadings.

Sensitivity analyses. Appendix B.3 examines Census pilot sizes $m \in \{100, 300, 500, 700, 900\}$, compares pilot-selected and fixed radius multipliers on CheXpert, and reports a fixed-multiplier sensitivity diagnostic on one Census split. Appendix B.4 gives the matched-proxy OPAL comparison for CheXpert.

AI use statement

We used generative AI tools to assist with revising the manuscript, developing and checking mathematical arguments, providing feedback on experimental design and interpretation, and implementing and debugging experimental code. The authors reviewed and revised all AI-assisted material, verified the mathematical claims and proofs, and checked the implementation and experimental results. The authors take full responsibility for the final content of this work, including all text, theoretical claims, code, and empirical findings.

REFERENCES

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

Anastasios N. Angelopoulos, John C. Duchi, and Tijana Zrnic. PPI++: Efficient prediction-powered inference. *arXiv preprint arXiv:2311.01453*, 2023b. doi: 10.48550/arXiv.2311.01453.

Aharon Ben-Tal and Arkadi Nemirovski. Robust convex optimization. *Mathematics of Operations Research*, 23(4):769–805, 1998. doi: 10.1287/moor.23.4.769.

Claes M. Cassel, Carl E. Särndal, and Jan H. Wretman. Some results on generalized difference estimation and generalized regression estimation for finite populations. *Biometrika*, 63(3):615–620, 1976. doi: 10.1093/biomet/63.3.615.

Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pp. 785–794, 2016.

Tong Chen and Thomas Lumley. Optimal sampling for design-based estimators of regression models. *Statistics in Medicine*, 41(8):1482–1497, 2022. doi: 10.1002/sim.9300.

Frances Ding, Moritz Hardt, John Miller, and Ludwig Schmidt. Retiring adult: New datasets for fair machine learning. *Advances in neural information processing systems*, 34:6478–6490, 2021.

Naoki Egami, Musashi Hinck, Brandon M. Stewart, and Hanying Wei. Using imperfect surrogates for downstream inference: Design-based supervised learning for social science applications of large language models. In *Advances in Neural Information Processing Systems*, volume 36, pp. 68589–68601, 2023. doi: 10.52202/075280-3000.

Laurent El Ghaoui and Hervé Lebret. Robust solutions to least-squares problems with uncertain data. *SIAM Journal on Matrix Analysis and Applications*, 18(4):1035–1064, 1997. doi: 10.1137/S0895479896298130.

Adam Fisch, Joshua Maynez, R. Alex Hofer, Bhuwan Dhingra, Amir Globerson, and William W. Cohen. Stratified prediction-powered inference for effective hybrid evaluation of language models. In *Advances in Neural Information Processing Systems*, volume 37, pp. 111489–111514, 2024. doi: 10.52202/079017-3541.

D. G. Horvitz and D. J. Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association*, 47(260):663–685, 1952. doi: 10.1080/01621459.1952.10483446.

Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silviana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghgoo, Robyn Ball, Katie Shpanskaya, Jayne Seekins, David A. Mong, Safwan S. Halabi, Jesse K. Sandberg, Ricky Jones, David B. Larson, Curtis P. Langlotz, Bhavik N. Patel, Matthew P. Lungren, and Andrew Y. Ng. CheXpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 590–597, 2019. doi: 10.1609/aaai.v33i01.3301590.

Wenlong Ji, Lihua Lei, and Tijana Zrnic. Predictions as surrogates: Revisiting surrogate outcomes in the age of AI. *Biometrika*, 2026. doi: 10.1093/biomet/asag053. Advance article.

Daniel Kuhn, Soroosh Shafiee, and Wolfram Wiesemann. Distributionally robust optimization. *Acta Numerica*, 34:579–804, 2025. doi: 10.1017/S0962492924000084.

Puheng Li, Tijana Zrnic, and Emmanuel J. Candès. Robust sampling for active statistical inference. In *Advances in Neural Information Processing Systems*, volume 38, 2025. doi: 10.52202/085713-2312.

Jesús López-Fidalgo and Douglas P. Wiens. Robust active learning with binary responses. *Journal of Statistical Planning and Inference*, 220:1–14, 2022. doi: 10.1016/j.jspi.2022.01.004.

Virginia L. Ma and Emmanuel J. Candès. Optimized labeling resource allocation for prediction-assisted inference via OPAL. *arXiv preprint arXiv:2606.03211*, 2026. doi: 10.48550/arXiv.2606.03211.

Jerzy Neyman. On the two different aspects of the representative method: The method of stratified sampling and the method of purposive selection. *Journal of the Royal Statistical Society*, 97(4): 558–606, 1934. doi: 10.1111/j.2397-2335.1934.tb04184.x.

Rui Nie, Douglas P. Wiens, and Zhichun Zhai. Minimax robust active learning for approximately specified regression models. *Canadian Journal of Statistics*, 46(1):104–122, 2018. doi: 10.1002/cjs.11327.

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

Daniel Ting and Eric Brochu. Optimal subsampling with influence functions. In *Advances in Neural Information Processing Systems*, volume 31, 2018. URL https://proceedings.neurips.cc/paper/2018/hash/57c0531e13f40b91b3b0f1a30b529a1d-Abstract.html.

Eeshit Dhaval Vaishnav, Carl G. de Boer, Jennifer Molinet, Moran Yassour, Lin Fan, Xian Adiconis, Dawn A. Thompson, Joshua Z. Levin, Francisco A. Cubillos, and Aviv Regev. The evolution, evolvability and engineering of gene regulatory DNA. *Nature*, 603(7901):455–463, 2022. doi: 10.1038/s41586-022-04506-6.

A. W. van der Vaart. *Asymptotic Statistics*, volume 3 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, 1998. doi: 10.1017/CBO9780511802256.

Douglas P. Wiens. Bias constrained minimax robust designs for misspecified regression models. In N. Balakrishnan, V. B. Melas, and S. Ermakov (eds.), *Advances in Stochastic Simulation Methods*, Statistics for Industry and Technology, pp. 117–133. Birkhäuser, Boston, MA, 2000. doi: 10.1007/978-1-4612-1318-5_7.

Tijana Zrnic and Emmanuel J. Candès. Active statistical inference. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 62993–63010, 2024. URL https://proceedings.mlr.press/v235/zrnic24a.html.

A ADDITIONAL THEORY AND IMPLEMENTATION

We first give the inference results for a fixed regression policy, then study the robust allocation and its implementation. We next treat outcome-affine estimating equations, distinguishing a fixed policy from one fitted jointly to the evaluation covariates, and finally account for independent pilot learning. These results accompany Sections 2–5.

Unless stated otherwise, expectations and L_2 norms refer to the population distribution, and queries are independent Bernoulli draws conditional on the feature–outcome pairs. We take $0 < \underline{\pi} < \gamma < 1$; at $\gamma = \underline{\pi}$ or $\gamma = 1$, only the corresponding constant policy is feasible. For regression coefficient j, write $a_j(x) = e_j^\top Q^{-1} x$ and $\Lambda_j = a_j^2$. Section 2 specializes this notation to $j = 1$.

A.1 LINEAR-REGRESSION INFERENCE

Proposition A.1 (Active linear-regression inference)**.** *Let (X_i, Y_i) be i.i.d., and fix a measurable predictor f and policy $\pi \in \Pi_\gamma$. Conditional on these pairs, draw $\xi_i \mid X_i, Y_i \sim \mathrm{Bernoulli}\{\pi(X_i)\}$ independently across i. Suppose $Q = \mathbb{E}(XX^\top)$ is finite and positive definite, $\mathbb{E}[X(Y - X^\top\beta^*)] = 0$, and*

$$\mathbb{E}\big[\|X\|^2\{(Y - X^\top\beta^*)^2 + (Y - f(X))^2\}\big] < \infty.$$

Define $\widehat{\beta}_\pi = \widehat{Q}^{-1} n^{-1} \sum_i X_i \widetilde{Y}_i$ and $\widehat{\theta}_\pi = e_j^\top \widehat{\beta}_\pi$. Then $\widehat{\beta}_\pi \to_P \beta^*$ and $\sqrt{n}(\widehat{\theta}_\pi - e_j^\top\beta^*)$ converges to $\mathcal{N}\{0, V(\pi)\}$, with $V(\pi)$ given by (2) after setting $\Lambda = \Lambda_j$. If additionally $\mathbb{E}\|X\|^4 < \infty$, then

$$\widehat{V}(\pi) = \frac{1}{n}\sum_{i=1}^{n}\Big[e_j^\top \widehat{Q}^{-1} X_i\{\widetilde{Y}_i - X_i^\top\widehat{\beta}_\pi\}\Big]^2 \tag{15}$$

converges in probability to $V(\pi)$. For $V(\pi) > 0$, the interval $\widehat{\theta}_\pi \pm z_{1-\alpha/2}\sqrt{\widehat{V}(\pi)/n}$ has asymptotic coverage $1 - \alpha$, where z_u is the u quantile of the standard normal distribution.

Proof. Set $D_i = \xi_i/\pi(X_i) - 1$. Conditional on (X_i, Y_i), $\mathbb{E}(D_i \mid X_i, Y_i) = 0$ and $\mathbb{E}(D_i^2 \mid X_i, Y_i) = 1/\pi(X_i) - 1$. To separate full-data variation from query variation, use the identity

$$\widehat{\beta}_\pi - \beta^* = \widehat{Q}^{-1}\frac{1}{n}\sum_i X_i\{Y_i - X_i^\top\beta^* + D_i(Y_i - f(X_i))\}$$

whose summands have mean zero and finite second moments under the stated moment condition and propensity floor. The law of large numbers gives consistency. Applying the multivariate central limit theorem and using $\widehat{Q} \to_P Q$ then gives

$$\sqrt{n}(\widehat{\theta}_\pi - e_j^\top\beta^*) = \frac{1}{\sqrt{n}}\sum_i a_j(X_i)\{Y_i - X_i^\top\beta^* + D_i(Y_i - f(X_i))\} + o_P(1).$$

The query contribution has conditional mean zero given (X_i, Y_i), so it is uncorrelated with the full-data contribution. The two variances are V_{full} and $\mathbb{E}[\{1/\pi(X) - 1\}\Lambda_j(X)v_f(X)]$, respectively.

For variance consistency, let $u_i = \widetilde{Y}_i - X_i^\top\beta^*$ and $\widehat{u}_i = \widetilde{Y}_i - X_i^\top\widehat{\beta}_\pi$. The law of large numbers applies to $X_i X_i^\top u_i^2$. Moreover,

$$\frac{1}{n}\sum_i \|X_i\|^2(\widehat{u}_i - u_i)^2 \le \|\widehat{\beta}_\pi - \beta^*\|^2\frac{1}{n}\sum_i\|X_i\|^4 = o_P(1).$$

Together with Cauchy–Schwarz for the cross term, this bound shows that replacing u_i by $\widehat{u}_i$ preserves the limiting second-moment matrix. Consistency of $\widehat{Q}^{-1}$ therefore gives consistency of (15), and studentization yields the stated coverage. □

Finite-population target and design variance. For design-based inference, condition instead on the pool $(X_i, Y_i)_{i=1}^n$, the predictor, and the recorded propensities $\pi_i \ge \underline{\pi}$. Define $Q_n =$

$n^{-1}\sum_i X_i X_i^\top$, $\beta_n = Q_n^{-1} n^{-1}\sum_i X_i Y_i$, and $b_i = e_j^\top Q_n^{-1} X_i\{Y_i - f(X_i)\}$, assuming Q_n is nonsingular. Under independent Bernoulli queries, $\widehat{\theta}_\pi$ is design-unbiased for $e_j^\top \beta_n$ and

$$V_{n,\mathrm{des}} := \mathrm{Var}_\xi\{\sqrt{n}(\widehat{\theta}_\pi - e_j^\top \beta_n)\} = \frac{1}{n}\sum_i \left(\frac{1}{\pi_i} - 1\right) b_i^2. \tag{16}$$

An observable, design-unbiased estimator of this variance is

$$\widehat{V}_{n,\mathrm{des}} = \frac{1}{n}\sum_i \frac{\xi_i(1-\pi_i)}{\pi_i^2}\{e_j^\top Q_n^{-1} X_i\}^2\{Y_i - f(X_i)\}^2. \tag{17}$$

Unqueried summands are evaluated as zero and do not require Y_i. For a sequence of pools, sufficient conditions for a design-based studentized central limit theorem are $V_{n,\mathrm{des}} \geq c > 0$ and $n^{-2}\sum_i b_i^4 \to 0$, with a common propensity floor. Indeed, the fourth moments of $(\xi_i/\pi_i - 1)b_i$ are bounded by a constant times b_i^4, giving Lyapunov's condition; the variance of (17) is bounded by a constant times $n^{-2}\sum_i b_i^4$. Conditioning on the pool removes the full-data sampling term. Thus (15) estimates variation for a superpopulation target, whereas (17) estimates query variation around the fixed-pool target.

A.2 Other Robust objective

Throughout this subsection, assume $\Lambda \geq 0$, $\widehat{v} \geq 0$, $\Lambda, \widehat{v} \in L_2(P_X)$, and $\Lambda \not\equiv 0$. To distinguish a policy's cost at a fixed residual risk from its worst-case cost, write

$$\mathcal{R}(\pi; v) := \mathbb{E}\{\Lambda(X)v(X)/\pi(X)\}. \tag{18}$$

We use $\mathcal{Q}_\delta$ for the robust criterion denoted by Q_δ in Section 3. For the fixed true residual risk, $V(\pi) = V_{\mathrm{full}} + \mathcal{R}(\pi; v_f) - \mathbb{E}(\Lambda v_f)$. For fixed v_f, the two terms outside $\mathcal{R}$ do not depend on the policy. This equivalence of policy comparisons does not in general extend to an adversary that varies v.

Robustifying the added variance instead. For the alternative loss $\mathbb{E}[\{1/\pi(X) - 1\}\Lambda(X)v(X)]$, nonnegativity and Cauchy–Schwarz yield

$$\begin{aligned}\sup_{v\in\mathcal{V}_\delta(\widehat{v})} &\mathbb{E}\left[\left(\frac{1}{\pi} - 1\right)\Lambda v\right] \\ &= \mathbb{E}\left[\left(\frac{1}{\pi} - 1\right)\Lambda\widehat{v}\right] + \delta\left\|\left(\frac{1}{\pi} - 1\right)\Lambda\right\|_2.\end{aligned}$$

When the norm is positive, a perturbation in its norming direction attains the supremum. For fixed v, subtracting $\mathbb{E}[\Lambda v]$ does not affect the oracle policy. However, this term varies over the ambiguity set, so the resulting minimax problem generally differs from (4). We adopt the inverse-propensity criterion for its simpler robust penalty and explicit large-radius limiting policy.

Other L_p ambiguity sets. The extension mentioned in Section 3 follows from the same support-function calculation. For $1 < p < \infty$, let $q = p/(p-1)$. In this paragraph, replace the L_2 moment assumptions by $\widehat{v} \in L_p$ and $\Lambda \in L_q$, retaining nonnegativity and $\Lambda \not\equiv 0$, and replace the ambiguity constraint by $\|v - \widehat{v}\|_p \leq \delta$. Hölder's inequality gives

$$\sup_{\substack{v\geq 0 \\ \|v-\widehat{v}\|_p \leq \delta}} \mathcal{R}(\pi; v) = \mathcal{R}(\pi; \widehat{v}) + \delta\|\Lambda/\pi\|_q.$$

The nonnegative perturbation $\delta(\Lambda/\pi)^{q-1}/\|\Lambda/\pi\|_q^{q-1}$ has L_p norm δ and attains this bound. When $\Lambda > 0$ almost surely, the geometry-only minimizer is $\mathrm{clip}_{[\underline{\pi},1]}\{c\Lambda^{q/(q+1)}\}$, with c enforcing the budget. Indeed, minimizing the qth power of the norm reduces pointwise to minimizing $\Lambda(x)^q/u^q + \lambda u$. Thus the L_2 exponent $2/3$ is the case $q = 2$.

A.3 Proof and interpretation of geometry gain

Proof of Proposition 3.1. For the nonnegative residual proxy of Section 3, separate the plug-in term from the robust penalty by writing $A(\pi) = \mathbb{E}(\Lambda\widehat{v}/\pi)$, $S(\pi) = \|\Lambda/\pi\|_2$, and $S_* = \inf_{\Pi_\gamma} S$. The assumptions of the proposition and the propensity floor imply $0 \le A(\pi) \le C := \mathbb{E}(\Lambda\widehat{v})/\underline{\pi}$. Consequently, for every $\delta > 0$,

$$S_* \le \frac{\inf_{\pi\in\Pi_\gamma} \mathcal{Q}_\delta(\pi)}{\delta} \le S_* + \frac{C}{\delta}.$$

The upper bound follows by using policies arbitrarily close to the infimum of S; no convergence of optimizing policies is needed. For the uniform policy, $\mathcal{Q}_\delta(\pi_{\mathrm{unif}})/\delta \to \|\Lambda\|_2/\gamma$. Subtracting the optimal-value limit from this limit gives the large-radius difference in Proposition 3.1.

Substituting $\pi_\infty = \gamma\Lambda^{2/3}/\mathbb{E}(\Lambda^{2/3})$ gives $S_* = \frac{\{\mathbb{E}(\Lambda^{2/3})\}^{3/2}}{\gamma}$ when the propensity bounds are inactive. Jensen's inequality for the strictly convex map $t \mapsto t^3$ applied to $\Lambda^{2/3}$ gives $\mathbb{E}(\Lambda^2) \ge \{\mathbb{E}(\Lambda^{2/3})\}^3$, with equality exactly when Λ is constant almost surely. Substituting this value of S_* identifies Δ_Λ in Proposition 3.1; Jensen's inequality gives its nonnegativity and equality condition. □

Convergence of the policies. Retain the preceding moment assumptions and $\Lambda > 0$ almost surely, now allowing active propensity bounds. Let π_δ minimize $\mathcal{Q}_\delta$; existence follows from weak compactness of Π_γ in L_2 and convex lower semicontinuity of the objective. Pointwise minimization of $\Lambda^2/u^2 + \lambda u$ gives (5); continuity of its mean from $\underline{\pi}$ to 1 gives a budget-enforcing scale. Strict convexity makes this policy the unique minimizer of $F(\pi) := S(\pi)^2$. Using A, S, S_*, C from the preceding proof, optimality gives $0 \le S(\pi_\delta) - S_* \le C/\delta$ for $\delta > 0$. The second derivative of Λ^2/u^2 is at least $6\Lambda^2$ on the feasible interval, so the variational inequality at π_∞ yields

$$3\mathbb{E}\{\Lambda^2(\pi_\delta - \pi_\infty)^2\} \le F(\pi_\delta) - F(\pi_\infty) \le \frac{2S_*C}{\delta} + \frac{C^2}{\delta^2}.$$

For every $\eta > 0$, bounded propensities also give

$$\mathbb{E}(\pi_\delta - \pi_\infty)^2 \le P_X(\Lambda < \eta) + \eta^{-2}\mathbb{E}\{\Lambda^2(\pi_\delta - \pi_\infty)^2\}.$$

Taking $\delta \to \infty$ and then $\eta \downarrow 0$ proves $\pi_\delta \to \pi_\infty$ in $L_2(P_X)$.

Equality with clipping. The preceding optimal-value argument also applies when clipping is active: the limiting difference is $\|\Lambda\|_2/\gamma - S_*$, where $S_* = \inf_{\pi\in\Pi_\gamma} \|\Lambda/\pi\|_2$. The closed form Δ_Λ in Proposition 3.1 applies when the power allocation satisfies the propensity bounds. For $\underline{\pi} < \gamma < 1$, uniform sampling also minimizes the clipped problem only for constant Λ. To see necessity, uniform sampling is an interior feasible policy, so the directional derivative of $F(\pi) = \mathbb{E}(\Lambda^2/\pi^2)$ in every bounded mean-zero direction h must vanish: $-2\gamma^{-3}\mathbb{E}(\Lambda^2 h) = 0$. Taking $h = \mathbf{1}_B - P_X(B)$ for each measurable set B forces Λ^2 to be constant almost surely. Conversely, Jensen's inequality shows that a constant loading makes uniform sampling optimal. At either endpoint budget, every feasible policy is constant irrespective of the loading; the strict-gain conclusion therefore presumes an interior budget.

A.4 Outcome-affine inference for fixed policies

For a fixed policy, outcome affinity transfers the regression variance decomposition to the estimating equation. We first give the fixed-policy result and three examples. Evaluation-batch policies and independent pilot learning are treated in Appendix A.5.

Fixed-policy regularity conditions. For fixed f and π, write $O = (X, \xi, \xi Y)$ for the observed data, with $\xi Y = 0$ when $\xi = 0$, and put $U(O;\beta) = \psi(X, \widetilde{Y}; \beta)$, where $\widetilde{Y}$ is defined by (1). Assume the following:

(i) The evaluation pairs are i.i.d.; conditional on these pairs, queries are independent Bernoulli draws with measurable propensities $\pi(X) \ge \underline{\pi} > 0$.

(ii) The population score is integrable and has an interior, locally unique zero β^*. The score is affine in Y as in (7). A measurable root $\widehat{\beta}_\pi \to_P \beta^*$ of the corrected equation exists.

(iii) On an open set containing a compact convex neighborhood B of β^*, $U(O;\beta)$ is continuously differentiable almost surely, with

$$\mathbb{E}\sup_{\beta\in B}\|\partial_\beta U(O;\beta)\| < \infty, \qquad \mathbb{E}\sup_{\beta\in B}\|U(O;\beta)\|^2 < \infty.$$

The full-data score has a finite second moment at β^*.

(iv) Differentiation and expectation can be interchanged for the full-data score, and $M = \mathbb{E}\{\partial_\beta \psi(X,Y;\beta^*)\}$ is nonsingular. The map g is continuously differentiable near β^*.

The local envelopes supply the uniform laws of large numbers used in linearization and variance estimation. Existence of a consistent root is a separate assumption; outcome affinity alone does not supply it.

Theorem A.2 (Inference for outcome-affine estimating equations)**.** *Under conditions (i)–(iv),* $\widehat{\theta}_\pi = g(\widehat{\beta}_\pi)$ *is asymptotically linear with the active influence function derived below and variance* (10)*. Define*

$$\widehat{M} = \frac{1}{n}\sum_i \partial_\beta U(O_i;\widehat{\beta}_\pi), \qquad \widehat{\Omega} = \frac{1}{n}\sum_i U(O_i;\widehat{\beta}_\pi)U(O_i;\widehat{\beta}_\pi)^\top.$$

Then

$$\widehat{V} = \nabla g(\widehat{\beta}_\pi)^\top \widehat{M}^{-1}\widehat{\Omega}\widehat{M}^{-\top}\nabla g(\widehat{\beta}_\pi) \tag{19}$$

is consistent. If $V(\pi) > 0$*, the confidence interval* $\widehat{\theta}_\pi \pm z_{1-\alpha/2}\sqrt{\widehat{V}/n}$ *has asymptotic coverage* $1-\alpha$*.*

Proof. Affinity gives the exact identity, for every β,

$$U(O;\beta) = \psi\{X, f(X);\beta\} + \frac{\xi}{\pi(X)}s(X;\beta)\{Y - f(X)\}.$$

Conditional expectation over ξ recovers the full-data score and, under the derivative envelope, its Jacobian. Hence $\mathbb{E}U(O;\beta^*) = 0$, and the expected corrected-score Jacobian remains M. An integral Taylor expansion along the segment from β^* to the consistent root gives

$$\sqrt{n}(\widehat{\beta}_\pi - \beta^*) = -M^{-1}\frac{1}{\sqrt{n}}\sum_i U(O_i;\beta^*) + o_P(1).$$

The delta method identifies the full-data influence function and its signed outcome coefficient as

$$\varphi_{\text{full}}(X,Y) = -\nabla g(\beta^*)^\top M^{-1}\psi(X,Y;\beta^*),$$
$$a(X) = -\nabla g(\beta^*)^\top M^{-1}s(X;\beta^*).$$

The active influence function is consequently

$$\varphi_\pi(O) = \varphi_{\text{full}}(X,Y) + \left\{\frac{\xi}{\pi(X)} - 1\right\}a(X)\{Y - f(X)\},$$

as in Section 4. The query term has conditional mean zero given (X,Y) and conditional variance $\{1/\pi(X)-1\}\Lambda(X)\{Y-f(X)\}^2$. It is therefore uncorrelated with the full-data influence function, which gives (10). The i.i.d. central limit theorem then gives asymptotic normality. The envelope assumptions, consistency of the root, and the continuous mapping theorem give $\widehat{M} \to_P M$ and $\widehat{\Omega} \to_P \mathbb{E}\{U(O;\beta^*)U(O;\beta^*)^\top\}$. Consistency of (19) follows, and studentization gives the stated coverage when $V(\pi) > 0$. □

Mean estimation. For $\psi(X,Y;\beta) = Y - \beta$ and $g(\beta) = \beta$, the corrected root is $\widehat{\theta} = n^{-1}\sum_i \widetilde{Y}_i$. Here $M = -1$, $a = 1$, and the sandwich variance is $n^{-1}\sum_i(\widetilde{Y}_i - \widehat{\theta})^2$.

Logistic regression. The logistic score has derivative

$$\partial_\beta[X\{Y-\mu_\beta(X)\}] = -\mu_\beta(X)\{1-\mu_\beta(X)\}XX^\top.$$

Thus $M = -H$, giving the logistic loading in Section 4. Corrected pseudo-outcomes need not lie in $[0,1]$; existence of a finite consistent score root remains an assumption.

Marginal log odds ratio. Take $\beta = (\mu_0, \mu_1)^\top$ and

$$\psi(X, Y; \beta) = \begin{pmatrix} (1-G)(Y-\mu_0) \\ G(Y-\mu_1) \end{pmatrix}.$$

The corrected estimating equation gives

$$\widehat{\mu}_g = \frac{\sum_i \mathbf{1}\{G_i = g\}\widetilde{Y}_i}{\sum_i \mathbf{1}\{G_i = g\}}, \qquad g \in \{0,1\},$$

whenever both groups are nonempty. When both corrected means lie in $(0,1)$, the target estimator is $\widehat{\theta} = \mathrm{logit}(\widehat{\mu}_1) - \mathrm{logit}(\widehat{\mu}_0)$. For this score, $M = -\mathrm{diag}(p_0, p_1)$ and $\nabla g = (-1/\{\mu_0(1-\mu_0)\}, 1/\{\mu_1(1-\mu_1)\})^\top$, which gives the marginal log-odds loading in Section 4. Here $p_g > 0$ and $0 < \mu_g < 1$ are required. The fitted influence contribution is

$$\widehat{\varphi}_i = \left\{ \frac{G_i}{\widehat{p}_1\widehat{\mu}_1(1-\widehat{\mu}_1)} - \frac{1-G_i}{\widehat{p}_0\widehat{\mu}_0(1-\widehat{\mu}_0)} \right\} (\widetilde{Y}_i - \widehat{\mu}_{G_i}),$$

where $\widehat{p}_g = n^{-1}\sum_i \mathbf{1}\{G_i = g\}$. The sandwich variance equals $n^{-1}\sum_i \widehat{\varphi}_i^2$. Consistency places $\widehat{\mu}_g$ in $(0,1)$ with probability tending to one. Finite-sample boundary failures nevertheless require a prespecified reporting rule; they must not be silently discarded.

A.5 Evaluation-batch policies and independent pilot learning

We now allow the sampling probabilities to depend jointly on the entire evaluation covariate sample. Throughout this subsection, f, the nonnegative residual-risk proxy $\widehat{v}$, the pilot-fitted loading components, and $0 \le \delta < \infty$ are fixed. Let $\bar{\Lambda}$ denote the limiting design loading, obtained by replacing quantities estimated from evaluation covariates by their population limits while holding pilot-fitted components fixed. This design loading need not equal the true target loading Λ. We first establish allocation convergence, then prove inference using the recorded propensities, and finally condition on an independent pilot.

For an array $h = (h_i)_{i=1}^n$, write $P_n h = n^{-1}\sum_i h_i$ and $\|h\|_n = (P_n h^2)^{1/2}$. For a function, evaluation at the sample points is implicit.

Proposition A.3 (Consistency of the empirical RASI allocation)**.** *Let $X_1, \ldots, X_n$ be i.i.d. and $0 < \underline{\pi} < \gamma < 1$. Suppose $\bar{\Lambda}, \widehat{v} \ge 0$ belong to $L_2(P_X)$ and*

$$w(X) := \bar{\Lambda}(X)\widehat{v}(X) + \delta\bar{\Lambda}(X)^2 > 0 \quad \textit{almost surely}. \tag{20}$$

Let the nonnegative fitted loadings $\widehat{\Lambda}_{i,n}$ depend only on the fixed inputs and $X_{1:n}$, and assume

$$\frac{1}{n}\sum_{i=1}^n \{\widehat{\Lambda}_{i,n} - \bar{\Lambda}(X_i)\}^2 \xrightarrow{P} 0. \tag{21}$$

The population program (4)*, with Λ replaced by $\bar{\Lambda}$, then has a unique minimizer $\pi_\delta^\star$ up to P_X-null sets. If $p_n = (p_{1,n}, \ldots, p_{n,n})$ is a measurable minimizer of* (11)*, then*

$$\frac{1}{n}\sum_{i=1}^n \{p_{i,n} - \pi_\delta^\star(X_i)\}^2 \xrightarrow{P} 0. \tag{22}$$

The conclusion also holds for feasible solutions whose objective values exceed the empirical minimum by $o_P(1)$.

Proof. Existence follows from weak compactness of the feasible subset of $L_2(P_X)$ and convexity and lower semicontinuity of the objective. Condition (20) gives strict convexity: the reciprocal term is strictly convex wherever $\bar{\Lambda}\widehat{v} > 0$, and, for $\delta > 0$, the norm penalty is strictly convex along any two policies differing on a set where $\bar{\Lambda} > 0$. Thus the minimizer is unique.

Write $L_i = \bar{\Lambda}(X_i)$, $A_i = L_i\widehat{v}(X_i)$, and $B_i = L_i^2$. Let $\widehat{Q}_n$ be the objective in (11) and define its fixed-loading counterpart

$$\bar{Q}_n(p) = P_n\frac{A_i}{p_i} + \delta\left(P_n\frac{B_i}{p_i^2}\right)^{1/2}.$$

The propensity floor, Cauchy–Schwarz, and the reverse triangle inequality give the uniform perturbation bound

$$\sup_{p\in[\underline{\pi},1]^n}|\widehat{Q}_n(p) - \bar{Q}_n(p)| \le \frac{\|\widehat{\Lambda}_n - L\|_n}{\underline{\pi}}\big(\|\widehat{v}\|_n + \delta\big) = o_P(1).$$

Here $\|\widehat{v}\|_n = O_P(1)$ by the law of large numbers.

Set $q_i = \pi^\star_\delta(X_i)$ and $\mu_n = P_n q$. Since $\mu_n \to_P \gamma$, there is a feasible vector $\widetilde{q}$ with $\max_i|\widetilde{q}_i - q_i| = o_P(1)$: mix q with the constant vector 1 when $\mu_n < \gamma$, and with the constant vector $\underline{\pi}$ when $\mu_n > \gamma$. The mixing weights are respectively $(\gamma - \mu_n)/(1 - \mu_n)$ and $(\mu_n - \gamma)/(\mu_n - \underline{\pi})$. Both vanish in probability. Moreover,

$$|\bar{Q}_n(\widetilde{q}) - \bar{Q}_n(q)| \le \frac{\max_i|\widetilde{q}_i - q_i|}{\underline{\pi}^2}\big(P_nA + \delta\|L\|_n\big) = o_P(1).$$

Optimality of p_n and the uniform perturbation bound therefore imply

$$\bar{Q}_n(p_n) - \bar{Q}_n(q) \le o_P(1). \tag{23}$$

To turn this objective comparison into allocation convergence, put

$$S_\star = \|\bar{\Lambda}/\pi^\star_\delta\|_2, \qquad S_n = \left(P_n\frac{B_i}{q_i^2}\right)^{1/2}.$$

Then $S_n \to_P S_\star \in (0,\infty)$. The optimal value of the population program is a finite convex function of the budget on $[\underline{\pi},1]$. A supporting line at the interior budget γ therefore gives a finite multiplier $\lambda_\star$. Differentiating the resulting unconstrained Lagrangian over the box of policies and localizing the variational inequality yields

$$\{d_\star(x) + \lambda_\star\}\{u - \pi^\star_\delta(x)\} \ge 0, \qquad u \in [\underline{\pi},1], \quad P_X\text{-a.s.},$$
$$d_\star(x) = -\frac{\bar{\Lambda}(x)\widehat{v}(x)}{\pi^\star_\delta(x)^2} - \frac{\delta\bar{\Lambda}(x)^2}{S_\star\pi^\star_\delta(x)^3}.$$

The derivatives are integrable by the moment assumptions and the propensity floor. Define

$$d_{i,n} = -\frac{A_i}{q_i^2} - \frac{\delta B_i}{S_n q_i^3}.$$

Because $P_nB = O_P(1)$ and $S_n \to_P S_\star$,

$$P_n|d_{i,n} - d_\star(X_i)| = o_P(1).$$

The variational inequality and $P_np_n = \gamma$ imply

$$P_nd_{i,n}(p_{i,n} - q_i) \ge -\lambda_\star(\gamma - \mu_n) - o_P(1) = -o_P(1). \tag{24}$$

The norm support inequality gives, for every p in the box,

$$\left(P_n\frac{B_i}{p_i^2}\right)^{1/2} \ge \frac{1}{S_n}P_n\frac{B_i}{q_ip_i}.$$

Using this bound and the exact reciprocal identity

$$\frac{1}{p} - \frac{1}{q} = -\frac{p-q}{q^2} + \frac{(p-q)^2}{pq^2},$$

we obtain

$$\begin{aligned}\bar{Q}_n(p) - \bar{Q}_n(q) &\geq P_n d_{i,n}(p_i - q_i) \\ &\quad + P_n\left(A_i + \frac{\delta B_i}{S_n q_i}\right)\frac{(p_i - q_i)^2}{p_i q_i^2} \\ &\geq P_n d_{i,n}(p_i - q_i) + P_n\left(A_i + \frac{\delta B_i}{S_n}\right)(p_i - q_i)^2.\end{aligned}$$

The last inequality uses $p_i, q_i \leq 1$. Combining this with (23) and (24) gives

$$P_n w(X_i)(p_{i,n} - q_i)^2 = o_P(1),$$

since S_n converges to a finite positive constant. For every $\eta > 0$,

$$P_n(p_{i,n} - q_i)^2 \leq P_n \mathbf{1}\{w(X_i) \leq \eta\} + \eta^{-1} P_n w(X_i)(p_{i,n} - q_i)^2.$$

Letting $n \to \infty$ and then $\eta \downarrow 0$ proves (22) by (20). An $o_P(1)$ optimization error only adds an $o_P(1)$ term to (23), so the same proof applies. □

Scope of the design conditions. These are sufficient conditions, not a claim that arbitrary loading estimators satisfy (21). For $\delta > 0$, positive exposure requires only $\bar{\Lambda} > 0$ almost surely; for $\delta = 0$, it requires $\bar{\Lambda}\widehat{v} > 0$ almost surely. If exposure vanishes on a set of positive probability, the proof controls weighted allocation error but does not identify allocations on that set. Unweighted policy convergence then requires an additional argument, including appropriate control of any tie-breaking rule. No correctness assumption on $\widehat{v}$ is imposed.

Checking loading consistency. A useful sufficient condition is $\widehat{\Lambda}_{i,n} = \{\widehat{b}_n^\top z(X_i)\}^2$, $\bar{\Lambda}(x) = \{b_0^\top z(x)\}^2$, $\widehat{b}_n \to_P b_0$, and $\mathbb{E}\|z(X)\|^4 < \infty$. Indeed,

$$P_n(\widehat{\Lambda}_{i,n} - \bar{\Lambda}(X_i))^2 \leq \|\widehat{b}_n - b_0\|^2(\|\widehat{b}_n\| + \|b_0\|)^2 P_n\|z(X_i)\|^4 = o_P(1).$$

For linear regression, take $\widehat{b}_n = \widehat{Q}^{-1}e_j$; consistency follows from the law of large numbers and nonsingularity of the limiting second-moment matrix. For a pilot-fitted logistic loading, the same argument applies to the empirical weighted second-moment matrix, with the pilot coefficient held fixed, provided its population limit is nonsingular. For marginal log odds ratios, take $z(X) = (1 - G, G)^\top$: positive population group probabilities, consistent evaluation group frequencies, and fixed pilot means strictly between zero and one make the two loading coefficients converge. These are convergence statements for the fitted design geometry, not assertions that the pilot-fitted loading equals the true loading.

Batch estimating-equation conditions. Assume the outcome-affine form (7) and define the residual-score contribution

$$C(X, Y; \beta) = s(X; \beta)\{Y - f(X)\}.$$

Assume the population score has an interior, locally unique zero β^*, differentiation and expectation can be interchanged, $M = \mathbb{E}\{\partial_\beta \psi(X, Y; \beta^*)\}$ is nonsingular, and g is continuously differentiable near β^*. On an open set containing a compact convex neighborhood B of β^*, suppose $\psi(X, Y; \beta)$ and $C(X, Y; \beta)$ are continuously differentiable almost surely. Assume

$$\begin{aligned}H_0(X, Y) &= \sup_{\beta \in B}\{\|\psi(X, Y; \beta)\| + \|C(X, Y; \beta)\|\}, \\ H_1(X, Y) &= \sup_{\beta \in B}\{\|\partial_\beta \psi(X, Y; \beta)\| + \|\partial_\beta C(X, Y; \beta)\|\}, \\ \mathbb{E}H_0^2 &< \infty, \qquad \mathbb{E}H_1 < \infty.\end{aligned}$$

Finally, assume a measurable consistent root $\widehat{\beta} \to_P \beta^*$ of the corrected batch equation exists. These componentwise envelopes control every propensity above the common floor. Root existence and consistency remain separate assumptions; they do not follow from outcome affinity alone.

Theorem A.4 (Batch RASI inference)**.** *Let the evaluation pairs be i.i.d., and let p_n be the empirical RASI allocation in Proposition A.3. Suppose that proposition's assumptions and the batch estimating-equation conditions above hold. Conditional on the complete evaluation sample, draw independent $\xi_i \sim \mathrm{Bernoulli}(p_{i,n})$, where the probabilities are constructed without evaluation outcomes. Form the corrected estimator using these recorded probabilities. Then*

$$\sqrt{n}(\widehat{\theta} - \theta^*) \rightsquigarrow \mathcal{N}\{0, V(\pi_\delta^\star)\}, \qquad \theta^* = g(\beta^*),$$

where

$$V(\pi_\delta^\star) = V_{\mathrm{full}} + \mathbb{E}\left[\left\{\frac{1}{\pi_\delta^\star(X)} - 1\right\}\Lambda(X)v_f(X)\right]. \tag{25}$$

The sandwich estimator in (12) *converges in probability to this variance. If $V(\pi_\delta^\star) > 0$, the confidence interval has asymptotic coverage $1-\alpha$. The variance uses the true target loading Λ, not the possibly different design loading $\bar{\Lambda}$.*

Proof. Write $q_i = \pi_\delta^\star(X_i)$, $\psi_i(\beta) = \psi(X_i, Y_i; \beta)$, and $C_i(\beta) = C(X_i, Y_i; \beta)$. By outcome affinity, the observed corrected score equals

$$U_{i,n}(\beta) = \psi_i(\beta) + \left(\frac{\xi_i}{p_{i,n}} - 1\right) C_i(\beta).$$

We compare this batch score with a fixed-policy score on a common probability space. Without changing the joint law of the data and queries, take independent $T_i \sim \mathrm{Unif}(0,1)$, independent of the evaluation pairs, and set

$$\xi_i = \mathbf{1}\{T_i \le p_{i,n}\}, \qquad \xi_i^\circ = \mathbf{1}\{T_i \le q_i\}, \qquad R_{i,n} = \frac{\xi_i}{p_{i,n}} - \frac{\xi_i^\circ}{q_i}.$$

Let $\mathcal{F}_n = \sigma\{(X_i, Y_i) : 1 \le i \le n\}$. Conditional on $\mathcal{F}_n$, the $R_{i,n}$ are independent and

$$\begin{aligned}
\mathbb{E}(R_{i,n} \mid \mathcal{F}_n) &= 0, \\
\mathbb{E}(R_{i,n}^2 \mid \mathcal{F}_n) &= \frac{|p_{i,n} - q_i|}{p_{i,n} q_i} \le \underline{\pi}^{-2} |p_{i,n} - q_i|, \\
\mathbb{E}(|R_{i,n}| \mid \mathcal{F}_n) &\le 2\underline{\pi}^{-1} |p_{i,n} - q_i|.
\end{aligned} \tag{26}$$

Define the i.i.d. fixed-policy scores

$$U_i^\circ(\beta) = \psi_i(\beta) + \left(\frac{\xi_i^\circ}{q_i} - 1\right) C_i(\beta).$$

Then $U_{i,n}(\beta) - U_i^\circ(\beta) = R_{i,n} C_i(\beta)$.

First note a weighted consequence of allocation convergence. For any nonnegative integrable $W(X,Y)$, let $e_{i,n} = |p_{i,n} - q_i| \le 1$. For every $K > 0$,

$$P_n W_i e_{i,n} \le K(P_n e_{i,n}^2)^{1/2} + P_n W_i \mathbf{1}\{W_i > K\}.$$

The first term is $o_P(1)$ for fixed K by (22); the law of large numbers and then $K \to \infty$ control the second term. Consequently,

$$P_n W_i |p_{i,n} - q_i| = o_P(1). \tag{27}$$

Using (26), conditional independence, and (27),

$$\begin{aligned}
\mathbb{E}\left[\left\| \frac{1}{\sqrt{n}} \sum_i R_{i,n} C_i(\beta^*) \right\|^2 \middle| \mathcal{F}_n \right] &\le \underline{\pi}^{-2} P_n \|C_i(\beta^*)\|^2 |p_{i,n} - q_i| \\
&= o_P(1).
\end{aligned}$$

Conditional Chebyshev therefore gives

$$\frac{1}{\sqrt{n}} \sum_i U_{i,n}(\beta^*) = \frac{1}{\sqrt{n}} \sum_i U_i^\circ(\beta^*) + o_P(1). \tag{28}$$

For clarity, a conditional bound that converges in probability is sufficient here: the unconditional error probability is bounded by the expectation of its truncation at one, which tends to zero.

The same coupling controls the Jacobian and second-moment matrix. Indeed,

$$\sup_{\beta\in B}\|P_n\partial_\beta U_{i,n}(\beta)-P_n\partial_\beta U_i^\circ(\beta)\|\le P_n|R_{i,n}|H_{1,i},$$

$$\sup_{\beta\in B}\|P_nU_{i,n}(\beta)U_{i,n}(\beta)^\top-P_nU_i^\circ(\beta)U_i^\circ(\beta)^\top\|\le c_{\underline{\pi}}P_n|R_{i,n}|H_{0,i}^2,$$

where $c_{\underline{\pi}}<\infty$ depends only on the propensity floor. Their conditional expected upper bounds are $o_P(1)$ by (26) and (27), so both left sides are $o_P(1)$ by conditional Markov. The i.i.d. counterpart classes satisfy uniform laws of large numbers by continuity, compactness of B, and the integrable envelopes H_1 and H_0^2.

Since $\mathbb{E}U_i^\circ(\beta)=\mathbb{E}\psi(X,Y;\beta)$, the limiting Jacobian at β^* is M. An integral Taylor expansion between β^* and the consistent batch root, together with (28), gives

$$\sqrt{n}(\widehat{\beta}-\beta^*)=-M^{-1}\frac{1}{\sqrt{n}}\sum_i U_i^\circ(\beta^*)+o_P(1).$$

The i.i.d. central limit theorem and delta method apply. Their scalar influence function is

$$-\nabla g(\beta^*)^\top M^{-1}U_i^\circ(\beta^*)=\varphi_{\mathrm{full}}(X_i,Y_i)+\left(\frac{\xi_i^\circ}{q_i}-1\right)a(X_i)\{Y_i-f(X_i)\}.$$

The second term has conditional mean zero given (X_i,Y_i), so its covariance with the full-data term is zero. Its conditional variance is $(1/q_i-1)\Lambda(X_i)\{Y_i-f(X_i)\}^2$. This proves (25) and asymptotic normality.

Finally, the uniform comparisons above and consistency of the root imply

$$\widehat{M}\to_P M,\qquad \widehat{\Omega}\to_P \mathbb{E}\{U_i^\circ(\beta^*)U_i^\circ(\beta^*)^\top\}.$$

The continuous mapping theorem yields consistency of (12). Studentization gives the stated coverage when the limiting variance is positive. □

What the batch argument requires. The inference proof uses the empirical RASI program only through (22); it also applies to any covariate-only allocation converging in that sense to a fixed measurable policy above the propensity floor. No particular rate for allocation convergence is required. Crucially, the correction uses the actual probabilities that generated the queries, rather than estimates of a different sampling design. The result concerns a superpopulation target and independent Bernoulli queries. It does not automatically cover sequential outcome-dependent policies or fixed-size sampling without replacement.

Regression conditions for the batch results. For $\psi(X,Y;\beta)=X(Y-X^\top\beta)$, suppose Q is positive definite. The moment assumptions of Proposition A.1 together with $\mathbb{E}\|X\|^4<\infty$ imply the batch envelope conditions. Indeed, on a compact neighborhood of β^* the score is bounded by $\|X\||Y-X^\top\beta^*|+c\|X\|^2$, its derivative is $-XX^\top$, and $C=X\{Y-f(X)\}$ does not depend on β. The explicit corrected root exists whenever $\widehat{Q}$ is nonsingular, an event of probability tending to one. Conditional on the complete sample, the centered query correction to the averaged score, $n^{-1}\sum_i(\xi_i/p_{i,n}-1)X_i\{Y_i-f(X_i)\}$, has second moment at most $\underline{\pi}^{-1}n^{-1}P_n\|X_i\{Y_i-f(X_i)\}\|^2=O_P(n^{-1})$. Together with the full-data law of large numbers and $\widehat{Q}\to_P Q$, this proves root consistency. Under the assumptions of Proposition A.3, these regression moments and the stated query design suffice to apply the batch theorem to regression.

Corollary A.5 (Independent pilot learning)**.** *Fix an independent pilot and any independent external training data as the evaluation sample size tends to infinity, and let $\mathcal{G}$ contain this training information. The predictor, residual-risk proxy, selected radius $\widehat{\delta}$, and pilot-fitted loading components may be arbitrary measurable functions of $\mathcal{G}$. Suppose that for almost every training realization, Theorem A.4 applies conditionally with positive limiting variance. Then*

$$\sup_{t\in\mathbb{R}}\left|\mathbb{P}\left(\frac{\sqrt{n}(\widehat{\theta}-\theta^*)}{\sqrt{\widehat{V}}}\le t\,\middle|\,\mathcal{G}\right)-\Phi(t)\right|\longrightarrow 0\quad \textit{almost surely}. \tag{29}$$

The confidence interval has conditional coverage converging almost surely to $1-\alpha$ *and unconditional coverage converging to* $1-\alpha$*. The selected radius need not converge to a deterministic value. The same conclusions hold for a policy constructed entirely from independent data if Theorem A.2 applies conditionally instead.*

Proof. Conditioning on $\mathcal{G}$ fixes the training inputs and selected radius, but leaves the evaluation pairs i.i.d. from the original population distribution. The final batch allocation may still depend jointly on $X_{1:n}$; that dependence is handled by Theorem A.4, not removed by conditioning. The conditional studentized central limit theorem applies for almost every training realization. Continuity of the normal distribution function makes the convergence uniform in t. Conditional coverage follows, and bounded convergence after averaging over $\mathcal{G}$ gives unconditional coverage. For an entirely independently learned policy, use Theorem A.2 in the same argument. □

Fixed versus growing pilots. The corollary does not require consistency of pilot fitting, risk consistency of radius selection, or containment of v_f in the ambiguity set. Its moment, design, root-consistency, and positive-variance conditions must nevertheless hold for almost every realized pilot. The unstudentized limiting variance may remain random across pilots. A growing-pilot extension requires additional uniform conditional control of the relevant errors and moment envelopes; it is not implied by the fixed-pilot corollary.

B Additional experimental details

We give the synthetic protocol, the real-data reporting conventions, and additional experiments examining pilot size, robustness-radius selection, the comparison with OPAL, and design-based coverage. Throughout this appendix, RASI is also labeled L_2-robust in the figures, and the linear-mixture baseline refers to arithmetic mixing with uniform sampling.

B.1 Synthetic linear-regression separation example

Data-generating process and target. Let $X_1, U, \varepsilon \overset{\text{ind}}{\sim} \mathcal{N}(0,1)$ and set

$$X_2 = 0.45X_1 + \sqrt{1-0.45^2}\,U,$$
$$Y = 0.5 + X_1 - 0.6X_2 + \sigma(X_1)\varepsilon, \qquad \sigma(x) = 0.35 + \frac{0.85}{1+\exp(-1.5x)}.$$

The target is the slope on X_1, whose population value is one. For $X = (1, X_1, X_2)^\top$ and $e_j = (0,1,0)^\top$,

$$Q = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0.45 \\ 0 & 0.45 & 1 \end{pmatrix}, \qquad a_j(X) = \frac{X_1 - 0.45X_2}{1-0.45^2}, \qquad \Lambda_j(X) = a_j(X)^2.$$

The frozen outcome proxy is

$$f(X) = 0.5 + X_1 - 0.6X_2 + 0.65\sin(1.4X_1) + 0.20X_2^2,$$

so its conditional residual second moment is

$$v_f(X) = \sigma(X_1)^2 + \{0.65\sin(1.4X_1) + 0.20X_2^2\}^2.$$

This nonlinear, heteroskedastic residual risk lies outside the quadratic working model used for policy construction.

Policy construction and inference. Each of the 300 Monte Carlo replications uses 300 fully labeled pilot observations and 1,700 evaluation observations. Quadratic ridge regression estimates squared proxy residuals, using five-fold out-of-fold predictions, a full-pilot refit, and truncation at zero. Loadings use empirical pilot and evaluation matrices; the displayed Q is their population counterpart. Evaluation allocations are frozen before querying. All methods use expected query fraction $\gamma = 0.2$ and propensity floor 0.005. Inference uses the active regression estimator of Proposition A.1 and nominal 90% confidence intervals based on (15). Empirical coverage is evaluated against the known population slope.

Table 1: Real-data settings for the headline experiments in Section 6. Labeling budgets refer to evaluation queries and exclude pilot and proxy-training labels.

Setting	Census	CheXpert
Analysis cohort	318,215 observations	19,596 images
Target	Age coefficient in a no-intercept income regression	Marginal log odds ratio, age ≥ 40 versus age < 40
Outcome proxy	Histogram gradient boosting; 3,000 training labels	Frozen cardiomegaly probability
Residual model	Pilot-trained XGBoost	Up to ten score-quantile bins per age group
Headline pilot	500 labels	500 labels
Evaluation pool	10,000 observations	10,000 observations
Query design	Independent Bernoulli; floor 0.01	Independent Bernoulli; floor 0.01
Replication	100 outer, 300 query draws each	100 outer, 300 query draws each
Intervals	90% confidence interval	90% confidence interval from delta-method

Mixing paths and robustness radius. The arithmetic and geometric paths use $\tau \in \{0, 0.05, \ldots, 1\}$, interpolating from the plug-in policy at $\tau = 0$ to uniform sampling at $\tau = 1$. Geometric probabilities are normalized to the common budget, and all policies satisfy the propensity bounds. The finite RASI radius is selected once using 300 extra independent labels, then fixed at $\delta \approx 0.178$ throughout. We also evaluate the geometry-only $\delta \to \infty$ policy described in (5). For each mixing family, the ex-post curve-wise oracle selects the single grid value of τ minimizing mean interval width across all replications.

Every method uses the same 300 pilot labels and an expected 340 evaluation labels per replication, for an expected total of 640. These counts exclude the one-time external calibration of the RASI radius. Conversely, the oracle mixing weights use the complete simulation rather than pilot-only information. The experiment therefore compares the allocation geometries under these stated tuning conventions; it is not a fully resource-matched comparison of implementable tuning procedures. The main-text figure separates the gain from the geometry-aware fallback, with mean interval width approximately 0.158, from the additional gain of finite-radius RASI, whose mean width is approximately 0.149.

B.2 Real-data protocol and reporting conventions

Data splits and label accounting. Table 1 summarizes the headline experiments in Section 6. The Census experiment draws disjoint proxy-training, pilot, and evaluation samples; CheXpert uses a frozen pretrained predictor and disjoint pilot and evaluation samples. Pilots are excluded from inference. Each eligible CheXpert image has a distinct patient identifier. The horizontal axis counts only expected evaluation queries. With $n = 10{,}000$ and $\gamma \in [0.05, 0.30]$, these counts range from 500 to 3,000. At the headline pilot size $m = 500$, the expected pilot-plus-evaluation counts are therefore 1,000–3,500. Census uses a further 3,000 labels to train the outcome proxy; the pretraining cost of the frozen CheXpert predictor is not included.

Targets and design loadings. For Census, the regression design is $Z_i = (\mathrm{AGEP}_i, \mathrm{SEX}_i)^\top$, with no intercept. Writing $Q_n = n^{-1} \sum_i Z_i Z_i^\top$, the full-evaluation benchmark and design loading are

$$\theta_{\mathrm{eval}} = e_1^\top Q_n^{-1} \frac{1}{n} \sum_i Z_i Y_i, \qquad \widehat{\Lambda}_i = (e_1^\top Q_n^{-1} Z_i)^2.$$

The proxy predictor uses the larger demographic and socioeconomic feature vector described in Section 6. For CheXpert, $G_i = \mathbf{1}\{\mathrm{Age}_i \geq 40\}$ and the benchmark is $\theta_{\mathrm{eval}} = \mathrm{logit}(\mu_{1,n}) - \mathrm{logit}(\mu_{0,n})$, where $\mu_{g,n}$ is the full-evaluation outcome mean in group g. The design uses the log-odds loading in Section 4, with group proportions and outcome means estimated from the pilot; the latter are clipped to $[10^{-3}, 1 - 10^{-3}]$ for policy construction only. The loading is constant within each age group.

Shared residual proxy and radius selection. Adaptive policies share the outcome proxy and pilot-trained residual-risk estimates. At each budget, RASI uses (14) to select $t \in \{0.03, 0.1, 0.3, 1, 3\}$ from pooled five-fold out-of-fold residual predictions. The radius is t times their full-pilot RMSE against observed squared residuals. Each candidate uses full-pilot loadings and joint budget normalization; we select the one with the smallest pilot average of $\widehat{\Lambda}_i(Y_i - f(X_i))^2/\pi_i$. Only the residual predictions are cross-fitted: the scale, loadings, and normalization reuse the full pilot. The resulting score is not an unbiased held-out risk estimate, and no selection-consistency guarantee is asserted. The residual model refit on the full pilot and the selected radius determine the evaluation allocation. Fitting and tuning use no evaluation information and do not certify ambiguity-set containment.

Replication and performance summaries. Unless stated otherwise, the real-data experiments use $B = 100$ outer replications and $S = 300$ query realizations within each replication. The fixed-split multiplier diagnostic and the design-based experiment in Appendix B.5 are exceptions. Let W_{bsk} be the confidence-interval width for outer replication b, query realization s, and method k. For each budget, we first compute

$$\overline{W}_{bk} = \frac{1}{S}\sum_{s=1}^{S} W_{bsk}, \qquad \mathrm{RE}_{bk} = \left(\frac{\overline{W}_{b,\mathrm{unif}}}{\overline{W}_{bk}}\right)^2.$$

The reported relative efficiency averages these replication-level ratios. Thus, widths are averaged before taking their squared ratio, and $\mathrm{RE} > 1$ indicates narrower intervals than uniform sampling. Coverage is likewise averaged first over query realizations within each outer replication. The displayed Monte Carlo bands use 1.96 times the standard error across outer-replication summaries.

Oracle mixture benchmarks. For each mixture family, outer replication, and budget, the mixing weight is chosen in hindsight from $\{0, 0.1, \ldots, 1\}$ to minimize $\overline{W}_{bk}(\tau)$. This differs from the synthetic experiment, which chooses one curve-wise oracle across all replications. Both use the same query draws for selection and reporting, allowing selection optimism; they are hindsight benchmarks rather than pilot-tuned policies.

Interpreting evaluation-target coverage. Outside Appendix B.5, Census uses HC0 intervals and CheXpert uses the group-mean delta method with within-group sample variances of corrected outcomes. The evaluation outcomes are used after policy construction to compute the full-evaluation benchmark, not to fit or tune the implementable policies. The plotted coverage is the frequency with which the intervals contain this replication-specific benchmark; it is not a direct empirical measurement of coverage for the unknown superpopulation parameter. Conditioning on the evaluation pool removes the full-data sampling variation retained by the superpopulation sandwich. For example, the Census design variance is (16), whereas the reported HC0 intervals use (15). The same distinction applies to the CheXpert delta-method intervals. Consequently, evaluation-target coverage can be conservative, particularly at larger query fractions; the reported intervals should not be interpreted as exact design-based intervals.

B.3 Pilot-size and radius sensitivity

These experiments examine three different questions: sensitivity to pilot size in Census, the effect of selecting rather than fixing the radius multiplier in CheXpert, and sensitivity to fixed multipliers on one Census split.

Census: pilot-size sensitivity. Figures 4 and 5 vary the pilot size over $m \in \{100, 300, 500, 700, 900\}$, keeping the proxy-training and evaluation sizes at 3,000 and 10,000, respectively. RASI uses pilot-selected radii throughout. Its relative efficiency remains approximately 1.6–1.8 over the displayed pilot sizes and budgets, and it exceeds both oracle-tuned mixture baselines throughout the displayed grid. In contrast, the plug-in policy becomes less efficient than uniform sampling as the evaluation budget increases at every pilot size. Increasing the pilot size over this range does not remove that pattern. The corresponding evaluation-target coverage is approximately nominal or conservative for all methods. These results show that the Census efficiency advantage is not specific to the headline choice $m = 500$.

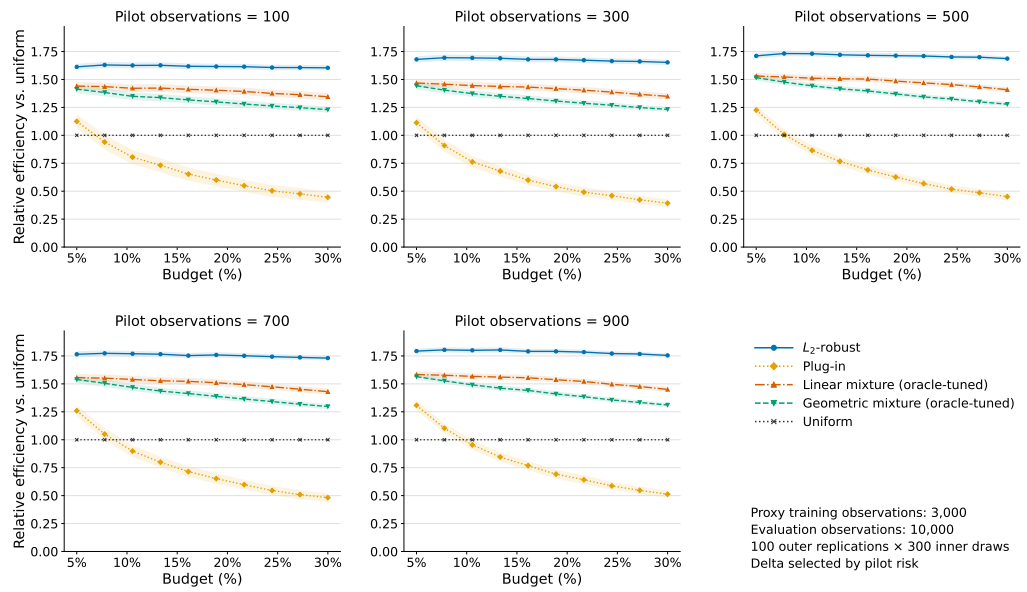

Figure 4: **Census pilot-size sensitivity: relative efficiency.** Each panel uses the indicated pilot size, 3,000 proxy-training observations, and 10,000 evaluation observations. RASI selects its radius using pilot risk; arithmetic and geometric mixtures are oracle-tuned within each outer replication and budget. Curves summarize 100 outer replications with 300 query realizations each, and bands are 95% Monte Carlo intervals across outer replications. Budgets count evaluation queries only.

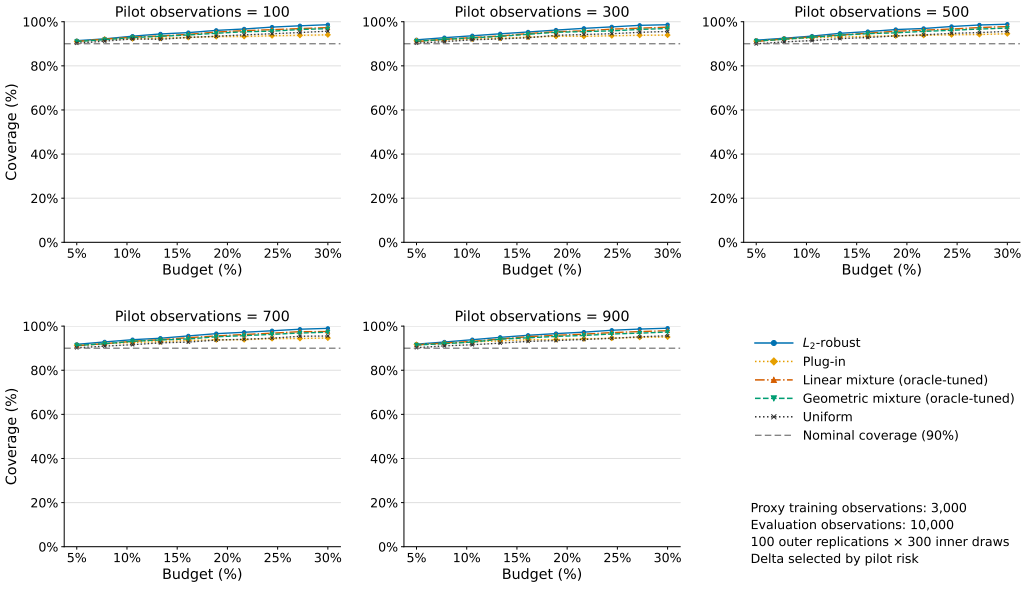

Figure 5: **Census pilot-size sensitivity: evaluation-target coverage.** Coverage of nominal 90% HC0 intervals against the replication-specific full-evaluation coefficient for the same configurations as Figure 4. The dashed horizontal line marks 90%. Bands summarize Monte Carlo uncertainty across the 100 outer replications, not just repeated queries of one fixed pool.

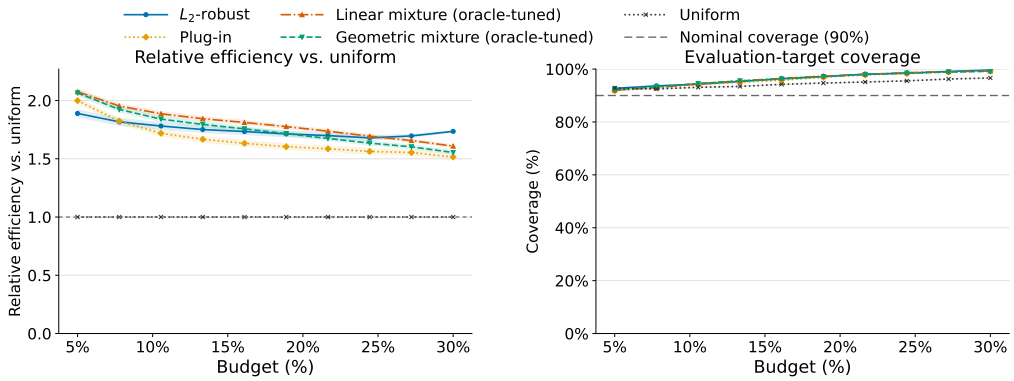

Figure 6: **CheXpert with fixed radius multiplier $t = 1$.** The experiment uses 500 pilot labels and an evaluation pool of 10,000 observations. The radius equals the pilot out-of-fold residual-model RMSE. The left panel reports relative efficiency against uniform sampling; the right panel reports evaluation-target coverage of nominal 90% intervals. Mixture baselines are oracle-tuned. Figure 3 gives the corresponding pilot-selected-radius results.

CheXpert: selected versus fixed radius. Figure 6 shows the fixed multiplier $t = 1$, using $m = 500$, for comparison with pilot-based radius selection in Figure 3. The fixed-multiplier policy remains more efficient than uniform sampling but does not uniformly outperform the oracle-tuned mixtures. In contrast, the selected-radius curve has higher relative efficiency than the fixed-multiplier curve throughout the displayed budget range and exceeds both mixture baselines. Evaluation-target coverage remains approximately nominal or conservative under both choices.

Census: fixed-multiplier sensitivity on one split. Figure 7 holds a single proxy-training, pilot, and evaluation split fixed, with sizes 3,000, 500, and 10,000, respectively. We compare the five fixed multipliers $t \in \{0.03, 0.1, 0.3, 1, 3\}$, with $\delta = t\widehat{s}$ and $\widehat{s}$ the pilot out-of-fold residual-model RMSE. All policies share the frozen proxy, fitted residual moments, and propensity floor 0.01. At each budget, 500 query realizations use common random seeds across policies.

All five choices produce narrower intervals than uniform sampling on this split, with relative efficiencies approximately 1.5–1.75 across the displayed settings. The smaller multipliers have similar efficiency, while the largest multiplier incurs a modest efficiency cost. Thus the gain over uniform sampling on this split does not depend on selecting a single narrowly optimal multiplier. However, unlike the preceding repeated-split experiments, this diagnostic measures sensitivity conditional on one fitted pilot; its Monte Carlo intervals do not quantify variability from refitting the predictor or residual model on new samples.

B.4 OPAL WITH PILOT-TRAINED RESIDUALS

Controlled comparison. For the CheXpert log-odds-ratio target, we evaluate a matched-proxy version of OPAL (Ma & Candès, 2026), using the same full-pilot fitted squared-residual model, target loadings, labeling budget, and propensity bounds as the other methods. This experiment uses 500 pilot labels and 10,000 evaluation observations. OPAL and plug-in sampling optimize the same nominal criterion,

$$\sum_i \frac{\widehat{\Lambda}_i \widehat{v}_i}{\pi_i},$$

but OPAL restricts log-inverse query probabilities to group-specific spline functions of $\widehat{v}_i$. The pilot-trained moment $\widehat{v}_i$, rather than $f(X_i)\{1 - f(X_i)\}$, enters both the spline coordinate and nominal variance weights; allocation is joint across groups.

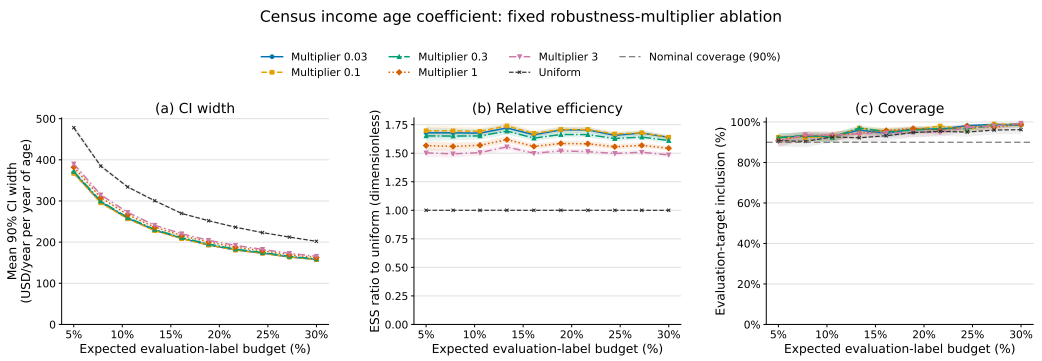

Figure 7: **Census fixed-multiplier diagnostic on one data split.** Mean 90% interval width (left), relative efficiency against uniform sampling (middle), and evaluation-target coverage (right) for five fixed radius multipliers. The split uses $3{,}000$ proxy-training observations, 500 pilot observations, and $10{,}000$ evaluation observations. Results use 500 query realizations per budget with common random seeds across policies. Shading shows pointwise approximate 95% Monte Carlo intervals conditional on this split; relative-efficiency intervals use a paired delta method. These are not uncertainty intervals across outer replications.

Results and interpretation. Figure 8 shows nearly overlapping relative-efficiency curves for OPAL and plug-in sampling. Both improve on uniform sampling, while RASI achieves higher relative efficiency throughout the displayed budget range. Evaluation-target coverage is approximately nominal or conservative for all displayed methods.

OPAL and plug-in sampling optimize the same nominal target-variance criterion, $\sum_i \widehat{\Lambda}_i \widehat{v}_i / \pi_i$. Over all feasible query probabilities, the plug-in solution is

$$\pi_i^{\mathrm{plug}} = \mathrm{clip}_{[\underline{\pi},1]} \left\{ c\sqrt{\widehat{\Lambda}_i \widehat{v}_i} \right\},$$

where c enforces the labeling budget. OPAL optimizes the same criterion within a spline policy class, whose restriction acts as structural regularization. When this class closely approximates the plug-in allocation, the two methods can achieve nearly the same efficiency. The observed near-overlap is therefore unsurprising, without requiring the policies to be exactly identical.

This does not contradict the gains reported by Ma & Candès (2026). Our plug-in baseline already incorporates target-specific variance optimization. The matched-proxy comparison therefore isolates the spline restriction, while RASI changes the objective to protect against residual-risk misspecification.

B.5 Census design-based coverage

We also evaluate intervals for the full-data age coefficient of a fixed Census evaluation pool. For each pilot-fitted policy, conditional on this pool, only the Bernoulli queries are random. We use $\widehat{\theta}_\pi \pm z_{0.95}\sqrt{\widehat{V}_{n,\mathrm{des}}/n}$, with the queried-residual variance estimator in (17), so the intervals account for query variation around this finite-population target. Figure 9 shows that RASI achieves approximately nominal 90% coverage while retaining smaller mean interval widths than the displayed baselines. Across the ten budgets, its empirical coverage ranges from 88.0% to 91.8%. These results support design-based calibration on this split; the normal intervals have asymptotic validity under the conditions in Appendix A.1, rather than an exact finite-sample coverage guarantee.

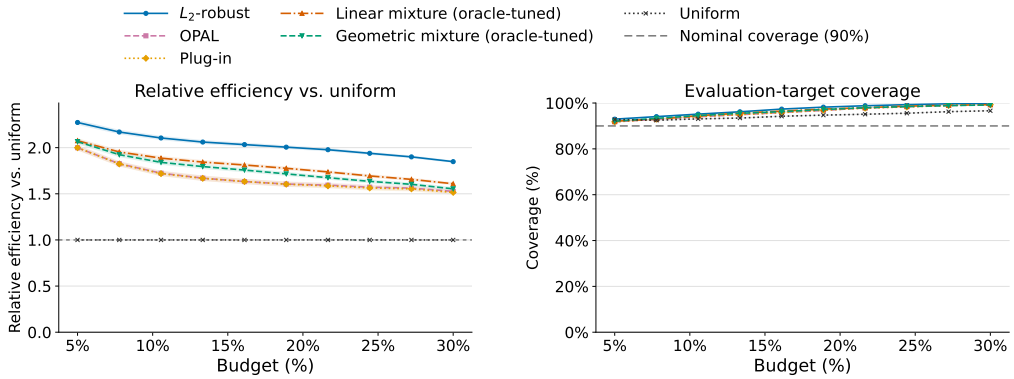

Figure 8: **Matched-proxy OPAL comparison on CheXpert.** Relative efficiency against uniform sampling (left) and evaluation-target coverage of nominal 90% intervals (right), using 500 pilot labels and 10,000 evaluation observations. Adaptive methods share the full-pilot fitted squared-residual model, target loadings, evaluation budget, and propensity bounds. RASI uses pilot-selected radii, and mixture weights are oracle-tuned. OPAL and plug-in have nearly overlapping relative-efficiency curves in this comparison.

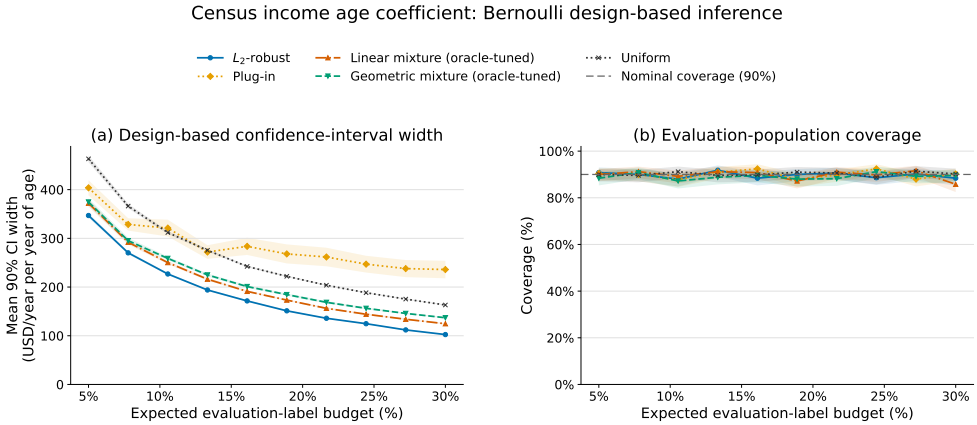

Figure 9: **Census design-based confidence intervals.** Mean width (a) and empirical coverage (b) of nominal 90% intervals for the fixed-pool age coefficient. The split uses 3,000 proxy-training observations, 500 pilot observations, and 10,000 evaluation observations (seed 2026), with 500 paired Bernoulli-query repetitions per budget. RASI selects its radius from the pilot using the multiplier grid $\{0.03, 0.1, 0.3, 1, 3\}$; all policies use propensity floor 0.01. Budgets exclude proxy-training and pilot labels. Shading shows pointwise 95% Monte Carlo intervals conditional on this split; the dashed line marks 90%. Oracle mixture weights minimize mean width on the same repetitions and may exhibit selection optimism.

C ADDITIONAL RELATED WORK

Prediction-powered and model-assisted inference. Prediction-powered inference combines predictions with labeled-data correction to enable valid inference despite systematic prediction error (Angelopoulos et al., 2023a). This principle has precedents in inverse-probability weighting, generalized-difference estimation, and semiparametric augmentation (Horvitz & Thompson, 1952; Cassel et al., 1976; Robins et al., 1994). PPI++ and RePPI improve how predictions enter the estimator (Angelopoulos et al., 2023b; Ji et al., 2026), while stratified PPI adapts prediction weights and label allocations to heterogeneous prediction quality (Fisch et al., 2024). RASI instead fixes a prediction-assisted estimating equation and optimizes acquisition when the residual-risk proxy guiding allocation is unreliable.

Target-specific and robust active inference. Classical Neyman allocation and influence-function-based subsampling tailor sampling to the estimand and estimator (Neyman, 1934; Ting & Brochu, 2018; Chen & Lumley, 2022). ASI combines prediction-assisted correction with budgeted acquisition; its generalized-linear-model allocation already separates residual uncertainty from a target-specific loading (Zrnic & Candès, 2024). RASI builds on this factorization. The closest explicit robustification is Li et al. (2025), who optimize a worst-case criterion along prescribed paths between an active policy and uniform sampling. Their formulation incorporates target geometry but places ambiguity on the combined error contribution. OPAL instead optimizes nominal target-specific variance over spline-parameterized policies, using policy smoothness to address noisy uncertainty scores (Ma & Candès, 2026). RASI differs in both the ambiguity set and the policy class: it perturbs only residual risk, retains the design loading, and optimizes over the full feasible class Π_γ. Minimax regression design also addresses model misspecification (Nie et al., 2018; López-Fidalgo & Wiens, 2022). For unbiased weighted regression, Wiens (2000) considers L_2-bounded variance functions and obtains design densities proportional to a regression-geometry term raised to the power $2/3$. RASI uses this related robust-design structure in prediction-assisted inference, with ambiguity centered on a learned residual second moment and Bernoulli query probabilities constrained by propensity bounds and a labeling budget.

Robust optimization and distributional ambiguity. RASI uses the classical support-function connection between uncertainty sets and regularization (Ben-Tal & Nemirovski, 1998; El Ghaoui & Lebret, 1997): its L_2 residual-risk class adds $\delta\|\Lambda/\pi\|_2$ to the nominal objective. Whereas distributionally robust optimization generally varies probability laws (Kuhn et al., 2025), RASI holds P_X fixed and varies a nonnegative conditional residual-second-moment function. Its ambiguity set is therefore a structured residual-function relaxation, not a model of arbitrary covariate shift or a complete family of joint feature–outcome distributions.