When Is Coarse Supervision Worth It? Cost-Aware Learning under Unknown Aggregation

Anonymous authors
Paper under double-blind review

Abstract

Modern learning systems often acquire supervision at multiple resolutions, trading annotation cost against information content. We study cost-aware two-resolution learning, where expensive fine labels reveal a vector response and cheaper coarse labels reveal a scalar aggregate formed with unknown weights, while the target remains the full response. The challenge is that unknown aggregation changes which directions coarse data can identify, so the value of coarse supervision depends jointly on cost, noise, and identification. We characterize this information geometry and develop an estimate-and-track policy that learns the aggregation rule and tracks the optimal resolution mix. We derive a closed-form break-even condition for coarse supervision and prove that the online policy attains the optimal leading cumulative-risk coefficient, with a matching local asymptotic minimax lower bound. Synthetic experiments support the predicted all-fine/mixed transition, show the online learner approaching the oracle-share benchmark, and demonstrate a finite-budget gain over all-fine acquisition when coarse supervision is sufficiently favorable. Our results provide a principled way to balance information and annotation cost across supervision resolutions.

1 Introduction

Modern learning systems increasingly acquire supervision at different resolutions. For the same example one may purchase a detailed record of many attributes or a cheaper summary of them: pixel-level versus image-level annotation in vision, token-level versus sentence-level labels in language, a detailed diagnostic profile versus an overall outcome, or a multi-criterion evaluation rubric versus one overall score. A growing body of work in active learning and weak supervision accordingly treats the *level* of supervision, and not only the choice of example, as part of the acquisition decision (Tejero et al., 2023; Hu et al., 2019; Rotman & Reichart, 2022; Matsuo et al., 2025; Nguyen et al., 2026). This raises a basic question: when supervision is available at several resolutions and prices, how should a finite annotation budget be divided across them?

Consider evaluating model-generated responses. A detailed annotation scores a response on several criteria such as factuality, relevance, safety, clarity, and style, while a cheaper annotation asks only for an overall judgment. The quantity we wish to predict may nevertheless be the full criterion vector. The cheap label is therefore not simply a noisier copy of the target: it is a scalar projection of a K-dimensional response, carrying information only along one direction. The relevant tradeoff is whether the information gained from this inexpensive direction compensates for the fine information forgone elsewhere.

A further difficulty is that the aggregation rule is often not known to the learner. An overall score may combine fine criteria according to weights implicit in an evaluation protocol or annotator population rather than specified in advance. If the rule were known, the cheap channel would supply information along a known response direction. When it is unknown, that direction must itself be identified from fine and coarse observations before that information can be fully exploited. Cheap supervision therefore creates an additional identification cost: its value depends not only on price and noise, but also on uncertainty about what is being aggregated.

We study this tradeoff in a stylized linear two-resolution model. Each example carries a d-dimensional covariate; a fine query returns a noisy K-dimensional response generated by a linear map, and a cheaper but coarse query returns a noisy scalar aggregation under unknown simplex

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.

weights. The inferential target is the entire fine map, measured in prediction risk. Linearity makes the nuisance-efficient information geometry exactly computable, yielding closed-form break-even conditions, optimal resolution mixes, and leading-order annotation savings.

Our approach starts from a simple geometric question: which parts of the fine map can each annotation type distinguish? If the aggregation rule were known, coarse labels would inform a fixed response direction, while fine labels identify the remaining directions. When the weights are unknown, however, some changes in the fine map can be matched by changes in the aggregation weights and therefore leave the coarse conditional mean unchanged. In the linear model, this ambiguity removes exactly $K-1$ directions from the part of the target that coarse data can help estimate. The remaining information still separates into fine-only and shared blocks, so the design problem reduces to their dimensions and a single scalar measuring cost-adjusted precision. This yields a one-dimensional objective and a sharp break-even boundary between all-fine and mixed acquisition.

Our problem connects multiresponse optimal design, multifidelity estimation, and cost-sensitive supervision (Sagnol, 2011; Peherstorfer et al., 2016; Xu et al., 2022; Matsuo et al., 2025; Fotakis et al., 2021). These literatures separately address heterogeneous observation operators, budget allocation, and supervision granularity. What is specific here is that the cheap channel is a lower-dimensional projection of the prediction target whose direction is itself unknown, and that we seek exact first-order answers for the resulting cost–information tradeoff.

Summary of Contributions. Under the canonical linear model, we develop a first-order theory of cost-aware two-resolution supervision with unknown aggregation. Specifically, we separate two sources of difficulty. First, not knowing the aggregation direction changes the information available from coarse labels and creates an intrinsic statistical price. Second, the optimal resolution mix depends on that unknown direction, so the learner must estimate and track the allocation online. We show that this second difficulty need not add another first-order penalty.

- **A sharp static cost–information frontier.** For fixed d and K, a single effective ratio λ, combining relative annotation costs, noise variances, and the aggregation strength $\|w_\star\|_2^2$, determines whether coarse supervision is worth purchasing. With unknown simplex aggregation weights, coarse labels enter the optimal design if and only if $\lambda > dK/(d-K+1)$, compared with threshold K when the aggregation rule is known. Above the boundary, both the optimal coarse spending share and its leading fractional gain are available in closed form. The gain measures the leading reduction in prediction risk at fixed budget, equivalently the leading reduction in annotation cost at fixed target risk, and remains structurally bounded even when coarse feedback becomes arbitrarily cheap or precise.
- **First-order optimal online adaptation with a matching lower bound.** When the aggregation rule is unknown, so is the optimal resolution mix. We give an epochal estimate-and-track policy that reserves a vanishing design stream to estimate the aggregation rule and target share, while the remaining budget tracks that share. A cross-fitted estimator uses the estimation stream without feeding its labels back into future allocation decisions. The cumulative risk matches the optimal leading $\log T$ coefficient for exogenous static designs under unknown aggregation, with an $O(1)$ remainder. A matching local asymptotic minimax lower bound shows that learning the allocation adds no first-order penalty beyond the intrinsic nuisance price.

Synthetic experiments support both parts of the theory. A structural comparison illustrates the gap between known- and unknown-aggregation acquisition decisions, while the online policy approaches the oracle-share benchmark and, in a strongly favorable coarse-information regime, yields a clear finite-budget improvement over all-fine acquisition.

2 Related Works

We discuss the closest connections to optimal design, multifidelity estimation, and cost-sensitive learning. Additional comparisons with the broader literature appear in Appendix A.1.

Optimal design and cost-aware partial measurement. Multiresponse optimal design allows experiment types to contribute different vector observations and information matrices. Sagnol (2011) develops A-optimal design under linear resource constraints, while later work treats multivariate

responses and variable response dimensions (Soumaya et al., 2015; Somogyi et al., 2026; Filová et al., 2026). With known aggregation, our fine and coarse channels are two structured observation types in this framework; related work also optimizes complete versus partial measurement under collection costs (Mareis & Drton, 2026). Our focus is the closed-form two-resolution frontier and how it changes when the coarse observation operator is itself unknown.

Multifidelity estimation and adaptive design. Multifidelity methods allocate budgets across expensive and cheap information sources, including analytic allocations and multi-output estimators (Peherstorfer et al., 2016; Schaden & Ullmann, 2020; Croci et al., 2023). Cross-fidelity relationships can also be learned adaptively (Xu et al., 2022; Dixon et al., 2026), while Fontaine et al. (2021) study online A-optimal scalar regression under unknown noise. These methods motivate our adaptive question, but our cheap channel is a lower-dimensional projection of the same vector target, and the unknown projection is a nuisance parameter in the conditional-mean law itself.

Cost-sensitive learning and supervision granularity. Active-learning work has long selected among annotation sources or forms with different costs and information content (Donmez & Carbonell, 2008; Wallace et al., 2011; Hu et al., 2019; Rotman & Reichart, 2022; Tejero et al., 2023; Matsuo et al., 2025). Related work studies learning from intrinsically coarse labels or cheap coarse machine annotations (Fotakis et al., 2021; Nguyen et al., 2026). These settings use different observation structures; here the cheaper action is an unknown scalar projection of the same vector-valued response whose full conditional mean remains the target.

3 Problem Setup

We formulate the two-resolution acquisition problem, define the performance criterion and reference regimes, and state the regularity assumptions used throughout.

3.1 Problem Formulation

We study a sequential information-acquisition problem with two annotation resolutions. At round t, the learner chooses the resolution $A_t \in \{\mathrm{F}, \mathrm{C}\}$ as a measurable function of the history before observing the current covariate $x_t \in \mathbb{R}^d$ or current response noise. The prediction target is the full K-dimensional conditional response $z_\star(x) = \Theta_\star^\top x$, where $\Theta_\star \in \mathbb{R}^{d\times K}$ is unknown.

A fine query returns

$$Y_t^{\mathrm{F}} = \Theta_\star^\top x_t + \varepsilon_t^{\mathrm{F}}, \qquad \varepsilon_t^{\mathrm{F}} \sim \mathcal{N}(0, \sigma_{\mathrm{F}}^2 I_K), \tag{1}$$

and costs $c_{\mathrm{F}} > 0$. A coarse query costs $c_{\mathrm{C}} > 0$ and returns only

$$Y_t^{\mathrm{C}} = w_\star^\top \Theta_\star^\top x_t + \varepsilon_t^{\mathrm{C}} = x_t^\top u_\star + \varepsilon_t^{\mathrm{C}}, \qquad u_\star := \Theta_\star w_\star, \qquad \varepsilon_t^{\mathrm{C}} \sim \mathcal{N}(0, \sigma_{\mathrm{C}}^2), \tag{2}$$

where $w_\star$ lies in the probability simplex $\Delta_K := \{w \in \mathbb{R}^K : \mathbf{1}^\top w = 1,\ w_k \geq 0\}$. Hence coarse supervision reveals a scalar projection of the same fine conditional mean rather than a separate prediction target.

The round index counts completed queries, while the budget clock records their cumulative annotation cost. For a budget level b, let $N_{\mathrm{F}}(b)$ and $N_{\mathrm{C}}(b)$ be the numbers of completed queries whose cumulative annotation cost does not exceed b. Bounded action costs can leave a bounded unused residue, so $c_{\mathrm{F}} N_{\mathrm{F}}(b) + c_{\mathrm{C}} N_{\mathrm{C}}(b) \leq b$. The learner may use the entire observed history to choose future resolutions but never observes the unqueried response. We consider both known and unknown aggregation. Known $w_\star$ provides a static reference problem; our main setting treats $w_\star$ as unknown and learns the allocation sequentially.

3.2 Key Quantities

Definition 1 (Fine-prediction risk). *For an estimator $\widehat{\Theta}$, define $\ell(\widehat{\Theta}, \Theta_\star) := \mathbb{E}_x[\|\widehat{\Theta}^\top x - \Theta_\star^\top x\|_2^2]$. Let B_0 denote the end of a fixed finite initialization and the start of the reported budget clock. For budget $b \geq B_0$, let $L_b := \mathbb{E}[\ell(\widehat{\Theta}_b, \Theta_\star)]$, with the current estimator held fixed between completed estimation queries. The cumulative fine-prediction risk through budget T is*

$$\mathfrak{R}_T := \sum_{b=B_0}^{T} L_b. \tag{3}$$

Therefore, T is an annotation-cost horizon rather than a query count. The budget clock in Equation (3) uses integer cost units for convenience; rational costs can be rescaled, and a continuous budget integral gives the same first-order conclusions. This criterion measures the anytime quality of a predictor updated throughout an annotation campaign, rather than only its terminal risk.

Definition 2 (Reference regimes)**.** *We use three reference regimes.*

- B_1*:* known-$w_\star$ static oracle*: the best exogenous static allocation when $w_\star$ is known;*
- B_2*:* unknown-aggregation oracle-share benchmark*: the true-parameter optimal share is given exogenously, but $\Theta_\star$ and $w_\star$ remain unknown and the schedule carries no inferential information;*
- B_3*:* online learned-allocation setting*: $w_\star$ and the optimal share are learned sequentially; Theorem 6 addresses it.*

3.3 Assumptions

We first impose standard regularity on the covariates.

Assumption 1 (Covariates)**.** *The covariates are i.i.d. with $\mathbb{E}[x_t] = 0$ and covariance $\Sigma_x \succ 0$. Unlimited unlabeled covariates are available at zero annotation cost, so Σ_x is treated as known. After whitening, $\mathbb{E}[x_t x_t^\top] = I_d$ and $\|x_t\|_2 \le L_x$ almost surely.*

Whitening only changes coordinates and is without loss for prediction risk, and bounded covariates are a standard regularity condition.

Assumption 2 (Noise and costs)**.** *Fine-query noises are i.i.d. $\mathcal{N}(0, \sigma_{\mathrm{F}}^2 I_K)$ and coarse-query noises are i.i.d. $\mathcal{N}(0, \sigma_{\mathrm{C}}^2)$. The queried noises are independent across items, independent of covariates, and independent across resolutions conditional on the covariate. The positive costs $c_{\mathrm{F}}, c_{\mathrm{C}}$ and noise variances $\sigma_{\mathrm{F}}^2, \sigma_{\mathrm{C}}^2$ are known.*

The Gaussian independent-channel model is the canonical setting in which the two resolutions have separate measurement noise; correlated or shared item-level noise would lead to a different pairing problem (see Section A.2). We next require the aggregation parameter to be regular and identifiable.

Assumption 3 (Aggregation regularity and identifiability)**.** *The dimensions satisfy $d \ge K \ge 2$. For some fixed $\tau \in (0, 1/(2K))$, $w_\star$ lies in the simplex interior with $\min_k w_{\star,k} \ge 2\tau$. Let $Q \in \mathbb{R}^{K\times(K-1)}$ have orthonormal columns spanning $\mathbf{1}^\perp$. For fixed constants $\kappa_\Theta, M_\Theta > 0$, the whitened coefficient matrix satisfies*

$$\sigma_{\min}(\Theta_\star Q) \ge \kappa_\Theta, \qquad \|\Theta_\star\|_{\mathrm{op}} \le M_\Theta. \tag{4}$$

The interior condition fixes the local simplex dimension at $K-1$, while the singular-value condition rules out observationally indistinguishable aggregation directions. Finally, we specify the sequential policy class and asymptotic regime.

Assumption 4 (Sequential protocol and asymptotics)**.** *Each action A_t is measurable with respect to the history before round t and is therefore conditionally independent of the current x_t and current noises. The dimensions, costs, noise variances, and regularity constants are fixed as $T \to \infty$.*

Choosing the resolution before seeing the current covariate isolates the resolution-allocation problem; allowing current-item leverage would introduce an additional optimal-design dimension.

4 Static Allocation with Known Aggregation

We first isolate the cost–information tradeoff when the aggregation vector $w_\star$ is known. This benchmark removes the need to learn the coarse direction and reveals, in closed form, when coarse supervision is worth purchasing, how much budget should be assigned to it, and how much it can improve over all-fine acquisition.

4.1 Fine-Prediction Risk under a Static Allocation

Fix fine and coarse sample counts N_{F} and N_{C}. Let $\beta = \mathrm{vec}(\Theta)$ and define $\alpha := N_{\mathrm{F}}/\sigma_{\mathrm{F}}^2$ and $\gamma := N_{\mathrm{C}}/\sigma_{\mathrm{C}}^2$. Since $w_\star$ is known, the expected Fisher information for β is

$$\mathcal{I}_{\mathrm{kn}} = \alpha I_{dK} + \gamma\big(w_\star w_\star^\top \otimes I_d\big). \tag{5}$$

Fine labels contribute in every response direction, while the coarse term acts only along $w_\star$.

Proposition 1 (Known-aggregation information trace)**.** *Under Assumptions 1 and 2, $\mathcal{I}_{\mathrm{kn}}$ has eigenvalue α with multiplicity $d(K-1)$ and eigenvalue $\alpha+\gamma\|w_\star\|_2^2$ with multiplicity d. Hence*

$$\operatorname{tr}(\mathcal{I}_{\mathrm{kn}}^{-1}) = d(K-1)\frac{\sigma_{\mathrm{F}}^2}{N_{\mathrm{F}}} + d\left(\frac{N_{\mathrm{F}}}{\sigma_{\mathrm{F}}^2} + \frac{N_{\mathrm{C}}\|w_\star\|_2^2}{\sigma_{\mathrm{C}}^2}\right)^{-1}. \tag{6}$$

For exogenous static allocations, this inverse-information trace is the first-order local minimax fine-prediction risk and is attained by the Gaussian efficient estimator described in Appendix C.1.

The $d(K-1)$ directions orthogonal to $w_\star$ can only be learned from fine labels, whereas the remaining d directions can use both resolutions. This separation makes the cost-allocation problem one-dimensional.

4.2 Cost-Optimal Resolution Mix

Let $\eta \in [0,1)$ denote the fraction of annotation budget spent on coarse labels. At total budget B, ignoring bounded integer effects, $N_{\mathrm{F}} = (1-\eta)B/c_{\mathrm{F}}$ and $N_{\mathrm{C}} = \eta B/c_{\mathrm{C}}$. Substituting into Equation (6) introduces the effective coarse-information ratio

$$\lambda := \frac{c_{\mathrm{F}}\sigma_{\mathrm{F}}^2\|w_\star\|_2^2}{c_{\mathrm{C}}\sigma_{\mathrm{C}}^2}, \tag{7}$$

which combines relative annotation cost, relative noise, and aggregation strength. For simplex weights, $\|w_\star\|_2^2 \in [1/K, 1]$: concentrated aggregation makes each coarse label more informative in this sense, while uniform averaging is least informative. The resulting leading risk is

$$\mathcal{R}_{\mathrm{K}}(B,\eta) = \frac{c_{\mathrm{F}}\sigma_{\mathrm{F}}^2 d}{B}\psi_{\mathrm{K}}(\eta,\lambda), \qquad \psi_{\mathrm{K}}(\eta,\lambda) := \frac{K-1}{1-\eta} + \frac{1}{1+(\lambda-1)\eta}. \tag{8}$$

Theorem 2 (Static break-even boundary)**.** *Suppose $K \geq 2$ and $\lambda > 0$. The function $\psi_{\mathrm{K}}(\eta,\lambda)$ is strictly convex in $\eta \in [0,1)$. Its unique minimizer is*

$$\eta_{\mathrm{K}}^\star(\lambda) = \begin{cases} 0, & \lambda \leq K, \\ \dfrac{q_{\mathrm{K}}(\lambda)-1}{\lambda-1+q_{\mathrm{K}}(\lambda)}, & \lambda > K, \end{cases} \qquad q_{\mathrm{K}}(\lambda) := \sqrt{\frac{\lambda-1}{K-1}}. \tag{9}$$

Therefore coarse supervision enters the optimal static design if and only if $\lambda > K$.

Proof. Differentiating Equation (8) gives $\partial_\eta\psi_{\mathrm{K}}(0,\lambda) = K-\lambda$, while the second derivative is strictly positive on $[0,1)$. Hence the all-fine boundary is optimal exactly when $\lambda \leq K$. Solving the first-order condition above the boundary gives Equation (9). □

The threshold has a direct economic interpretation: coarse labels must be sufficiently precise per unit cost along the aggregation direction to compensate for the $K-1$ response directions that still require fine labels. Define $\Phi_{\mathrm{K}}^\star := c_{\mathrm{F}}\sigma_{\mathrm{F}}^2 d \min_{\eta\in[0,1)}\psi_{\mathrm{K}}(\eta,\lambda)$; its cumulative-risk interpretation is formalized together with the unknown-aggregation benchmark in Corollary 5.

4.3 Value of Coarse Supervision

Let $\psi_{\mathrm{K}}^\star(\lambda) := \psi_{\mathrm{K}}(\eta_{\mathrm{K}}^\star(\lambda),\lambda)$. Below the break-even boundary, $\psi_{\mathrm{K}}^\star(\lambda) = K$. Above it,

$$\psi_{\mathrm{K}}^\star(\lambda) = \frac{\left(\sqrt{(K-1)(\lambda-1)}+1\right)^2}{\lambda}, \qquad \lambda > K. \tag{10}$$

We measure the leading value of optimal coarse acquisition relative to all-fine by

$$G_{\mathrm{K}}(\lambda) := 1 - \frac{\psi_{\mathrm{K}}^\star(\lambda)}{K}. \tag{11}$$

Thus $G_{\mathrm{K}} = 0.04$, for example, means a 4% leading reduction in prediction risk at fixed budget, equivalently a 4% leading reduction in annotation cost at fixed target risk. The gain is zero for $\lambda \le K$, strictly positive above the boundary, and satisfies

$$0 \le G_{\mathrm{K}}(\lambda) < \frac{1}{K}, \qquad \lim_{\lambda\to\infty} G_{\mathrm{K}}(\lambda) = \frac{1}{K}. \tag{12}$$

Even arbitrarily cheap or precise coarse labels therefore cannot replace the fine information needed in the remaining $K-1$ response directions.

5 Unknown Aggregation and Online Adaptation

We now study the main setting in which the aggregation vector $w_\star$ is unknown. The section first identifies the statistical price of this nuisance parameter and then shows that the resulting optimal resolution mix can be learned online without an additional first-order penalty.

5.1 Unknown Aggregation: Static Benchmark

When $w_\star$ is unknown, a coarse observation depends jointly on the fine map and aggregation direction. Profiling this nuisance changes the two information-block dimensions. Using the same η and λ as in Section 4, let A_{U} and D_{U} denote the fine-only and fine-plus-coarse dimensions, and define

$$A_{\mathrm{U}} := (K-1)(d+1), \qquad D_{\mathrm{U}} := d-K+1, \qquad \psi_{\mathrm{U}}(\eta,\lambda) := \frac{A_{\mathrm{U}}/d}{1-\eta} + \frac{D_{\mathrm{U}}/d}{1+(\lambda-1)\eta}. \tag{13}$$

Let $\eta^\star_{\mathrm{U}}(\lambda)$ minimize $\psi_{\mathrm{U}}(\eta,\lambda)$ over $\eta \in [0,1)$.

Theorem 3 (Unknown-aggregation frontier)**.** *Under Assumptions 1 to 3, the first-order local minimax fine-prediction risk of an exogenous static allocation with budget B and coarse spending share η satisfies*

$$\mathcal{R}_{\mathrm{U}}(B,\eta) = \frac{c_{\mathrm{F}}\sigma_{\mathrm{F}}^2 d}{B}\psi_{\mathrm{U}}(\eta,\lambda) + o(B^{-1}). \tag{14}$$

A projected cross-fitted one-step estimator attains the same leading term. Coarse supervision enters the optimal design if and only if

$$\lambda > \lambda_{\mathrm{U}}^{\mathrm{crit}} := \frac{dK}{d-K+1}. \tag{15}$$

Compared with the known-$w_\star$ threshold K, unknown aggregation therefore converts exactly $K-1$ formerly coarse-informative directions into fine-only directions.

Define $\Phi^\star_{\mathrm{U}} := c_{\mathrm{F}}\sigma_{\mathrm{F}}^2 d\,\psi_{\mathrm{U}}(\eta^\star_{\mathrm{U}},\lambda)$. Above the threshold, $\eta^\star_{\mathrm{U}} = (q_{\mathrm{U}}-1)/(\lambda-1+q_{\mathrm{U}})$ with $q_{\mathrm{U}} := \sqrt{D_{\mathrm{U}}(\lambda-1)/A_{\mathrm{U}}}$. The leading gain relative to all-fine is $G_{\mathrm{U}}(\lambda) := 1-\psi_{\mathrm{U}}(\eta^\star_{\mathrm{U}},\lambda)/K$. It is zero at and below the boundary, increases above it, and approaches the structural ceiling $(d-K+1)/(dK)$, compared with $1/K$ when $w_\star$ is known.

Proof roadmap. The key step is Lemma 4: profiling out the $(K-1)$-dimensional simplex tangent of $w_\star$ moves exactly $K-1$ directions from the coarse-informative block into the fine-only block, giving the multiplicities in Equation (13). Taking the inverse-information trace yields Equation (14). Strict convexity then reduces the design problem to one dimension: the derivative at $\eta = 0$ gives the break-even boundary, while the first-order condition above the boundary gives the closed-form optimizer. Appendix C.3 gives the full information calculation, lower bound, and matching attainment.

Lemma 4 (Nuisance information split)**.** *After profiling out local perturbations of $w_\star$ in the simplex tangent space, the efficient information for $\Theta_\star$ has two eigenspaces: A_{U} directions receive only fine-label information, while D_{U} directions receive both fine and coarse information.*

Proof sketch. A local perturbation (H,a) of $(\Theta_\star, w_\star)$ changes the coarse mean through $Hw_\star + \Theta_\star a$. The tangent space of an interior simplex point has dimension $K-1$, and Assumption 3 makes its image under $\Theta_\star$ $K-1$ dimensional. Profiling over a removes exactly those $K-1$ directions from the coarse-informative block. The full Schur-complement calculation appears in Appendix C.4. □

Since first-order pointwise risk scales as $1/b$, these coefficients become $\log T$ cumulative-risk benchmarks on the budget clock.

Corollary 5 (Static benchmark attainment)**.** *Regular exogenous static procedures achieve* $\mathfrak{R}_T^{\mathsf{B}_1} \le \Phi_{\mathrm{K}}^\star \log T + O(1)$ *and* $\mathfrak{R}_T^{\mathsf{B}_2} \le \Phi_{\mathrm{U}}^\star \log T + O(1)$*; the two leading coefficients are local asymptotic minimax for their respective static experiments.*

5.2 Learning the Optimal Mix Online

The static optimum depends on the unknown $w_\star$ through λ. Algorithm 1 therefore separates allocation learning from prediction: at the start of epoch m, the accumulated vanishing design stream is used to estimate $w_\star$ and the target coarse share, while the estimation stream tracks that share and fits the predictor without feeding its labels back into future allocation decisions. Within the estimation stream, Algorithm 1 uses two deterministic folds; each fold takes one Fisher-scoring step from an opposite-fold pilot, and the two fine-map estimates are averaged and projected. Cross-fitting preserves pilot–score independence conditional on the design history while using estimation data at the full $b^{-1/2}$ scale.

Algorithm 1 Epochal Estimate-and-Track with Cross-Fitted Estimation (EET)

Require: B_0, $\zeta_0 \in (0, 1/4)$, fixed design split $\eta_{\mathrm{des}} \in (0,1)$, costs/noise variances, and known regularity bounds $\tau, \kappa_\Theta, M_\Theta$

Run a fixed predictable initialization through budget B_0 containing both resolutions in the design stream and both estimation folds; initialize signed tracking imbalance $s \leftarrow 0$
for $m = 0, 1, 2, \ldots$ **do**
  Set $B_m = 2^m B_0$, $\Delta B_m = B_{m+1} - B_m$, and $\zeta_m = \zeta_0/(m+2)^2$
  From accumulated design data available at B_m, fit guarded projected regressions $\widetilde{\Theta}_m, \widetilde{u}_m$ and set

$$\widetilde{w}_m \in \operatorname*{arg\,min}_{w \in \Delta_K,\, w_k \ge \tau} \|\widetilde{u}_m - \widetilde{\Theta}_m w\|_2^2, \qquad \widetilde{\lambda}_m = \frac{c_{\mathrm{F}}\sigma_{\mathrm{F}}^2}{c_{\mathrm{C}}\sigma_{\mathrm{C}}^2}\|\widetilde{w}_m\|_2^2.$$

  Set target coarse-budget share $\widehat{\eta}_m \leftarrow \eta_{\mathrm{U}}^\star(\widetilde{\lambda}_m)$
  Reserve $\zeta_m \Delta B_m$ for the design stream and query it using the fixed split η_{des}
  while the remaining estimation budget permits a query **do**
    $s_{\mathrm{C}} \leftarrow s + (1 - \widehat{\eta}_m)c_{\mathrm{C}}, \quad s_{\mathrm{F}} \leftarrow s - \widehat{\eta}_m c_{\mathrm{F}}$
    Choose the affordable action whose updated imbalance has smaller absolute value (fine on ties); stop the block if the only affordable action increases $|s|$
    Assign the query to the next deterministic parity fold of that resolution, then observe $(x_t, Y_t^{A_t})$ and update s
    Recompute the projected two-fold cross-fitted one-step estimator from all estimation data collected so far
  end while
  Hold the current predictor fixed through the next design block until another estimation query completes
end for

The bounded-imbalance rule tracks $\widehat{\eta}_m$ to $O(1)$ cost. More importantly, conditional on the accumulated design history, the entire estimation resolution schedule and fold assignment are fixed before the associated covariates and labels are observed. Estimation labels therefore never feed back into future allocation decisions. Let $w_0 := K^{-1}\mathbf{1}$ and $\vartheta = (\mathrm{vec}(\Theta), v)$ be a local coordinate with $w = w_0 + Qv$, and for the local statements let $\vartheta_\star$ be an interior point of the regular class with a sufficiently small radius-δ ball $\mathcal{U}_\delta$ around it.

Theorem 6 (First-order optimal online adaptation)**.** *Under Assumptions 1 to 4, EET satisfies*

$$\mathfrak{R}_T^{\mathrm{alg}} \le \Phi_{\mathrm{U}}^\star \log T + O(1). \tag{16}$$

The bound is locally uniform: for sufficiently small δ, $\sup_{\vartheta \in \mathcal{U}_\delta} \mathfrak{R}_T^{\mathrm{alg}}(\vartheta) \le (\Phi_{\mathrm{U}}^\star(\vartheta_\star) + o_\delta(1)) \log T + O_\delta(1)$. *Moreover,*

$$\lim_{\delta \downarrow 0} \liminf_{T \to \infty} \frac{1}{\log T} \inf_{\pi, \widehat{\Theta}} \sup_{\vartheta \in \mathcal{U}_\delta} \mathfrak{R}_T(\pi, \widehat{\Theta}; \vartheta) \ge \Phi_{\mathrm{U}}^\star(\vartheta_\star), \tag{17}$$

where the infimum ranges over predictable resolution policies and estimators. Hence EET is first-order locally asymptotically minimax.

Relative to the known-$w_\star$ oracle, the leading coefficient decomposes as $\Phi_{\mathrm{U}}^\star = \Phi_{\mathrm{K}}^\star + (\Phi_{\mathrm{U}}^\star - \Phi_{\mathrm{K}}^\star)$: the nonnegative difference is the nuisance price of unknown aggregation, while learning the allocation introduces no additional first-order term.

Proof roadmap. For the upper bound, Lemma 7 shows that the design stream uses only $O(b/(\log b)^2)$ budget and that the realized estimation-stream objective approaches the unknown-$w_\star$ static optimum fast enough that both the design opportunity cost and the induced allocation error are summable after the leading $1/b$ risk scaling. Conditional on this design history, Lemma 8 fixes the resolution schedule and cross-fitting folds, so the predictor attains the efficient risk of the realized fine/coarse mix up to an $O(b_{\mathrm{E}}^{-3/2})$ remainder. Combining the two lemmas gives $L_b \le \Phi_{\mathrm{U}}^\star/b$ plus summable terms at estimation endpoints. Together with the held-predictor accounting in Appendix C.6, summing over the full budget clock proves Equation (16); the same uniform bounds on a small regular neighborhood give the locally uniform statement.

For the lower bound, Lemma 9 applies van Trees to a local prior on the joint parameter (Θ, w). Predictability makes the policy kernel parameter-free conditional on the past, so the same nuisance-profiled information frontier lower-bounds the per-budget Bayes risk by $\Phi_{\mathrm{U}}^\star/b$ up to lower-order terms. Summing yields Equation (17); hence the upper and lower bounds meet in their leading $\log T$ coefficient.

Lemma 7 (Pilot and tracking error)**.** *Let $\Xi(b)$ be design-stream spending by budget b, let $b_{\mathrm{E}} := b - \Xi(b)$, and let η_b^{E} be the realized coarse-budget share in the estimation stream. Uniformly over the regular parameter class,*

$$\Xi(b) = O\left(\frac{b}{(\log b)^2}\right), \qquad \mathbb{E}\big[\psi_{\mathrm{U}}(\eta_b^{\mathrm{E}}, \lambda) - \psi_{\mathrm{U}}(\eta_{\mathrm{U}}^\star, \lambda)\big] = O\left(\frac{(\log b)^3}{b}\right). \tag{18}$$

No separation from the break-even boundary is required.

Proof sketch. The fixed design split gives $\Theta(b/(\log b)^2)$ design observations of each resolution. Guarded projected regression therefore estimates λ with mean-square error $O((\log b)^2/b)$. Uniform strong convexity of ψ_{U} makes the plug-in objective loss quadratic, including at the all-fine boundary. Geometric weighting across dyadic epochs and bounded-imbalance tracking give Equation (18). Appendix C.7 gives the full argument. □

Lemma 8 (Cross-fitted efficient estimation)**.** *Let $\mathcal{G}_m$ be the design history determining the current epoch schedule. For sufficiently large estimation endpoints b,*

$$L_b(\mathcal{G}_m) := \mathbb{E}\Big[\ell(\widehat{\Theta}_b, \Theta_\star) \mid \mathcal{G}_m\Big] \le \frac{c_{\mathrm{F}}\sigma_{\mathrm{F}}^2 d}{b_{\mathrm{E}}}\psi_{\mathrm{U}}(\eta_b^{\mathrm{E}}, \lambda) + O(b_{\mathrm{E}}^{-3/2}), \tag{19}$$

with a constant uniform over the regular parameter class and the realized predictable schedules generated by EET.

At estimation endpoints, combining Lemmas 7 and 8 and $b_{\mathrm{E}}^{-1} - b^{-1} = O(1/[b(\log b)^2])$ gives

$$L_b \le \frac{\Phi_{\mathrm{U}}^\star}{b} + O\left(\frac{1}{b(\log b)^2} + \frac{(\log b)^3}{b^2} + \frac{1}{b^{3/2}}\right), \tag{20}$$

whose remainder is summable over estimation endpoints. Appendix C.6 shows that the intervening held-predictor budget levels preserve summability on the full budget clock.

Lemma 9 (Adaptive information lower bound)**.** *For any predictable resolution policy, a smooth compactly supported prior on the joint local parameter (Θ, w) inside $\mathcal{U}_\delta$ gives per-budget Bayes risk at least*

$$\frac{\Phi_{\mathrm{U}}^\star(\vartheta_\star) - o_\delta(1)}{b} - O_\delta(b^{-2}). \tag{21}$$

Proof sketch. Predictability makes the policy kernel parameter free, so adaptive Fisher information is the expected sum of queried-label information. Joint van Trees and nuisance profiling as in Lemma 4 yield Equation (21); Appendix C.9 gives the details. □

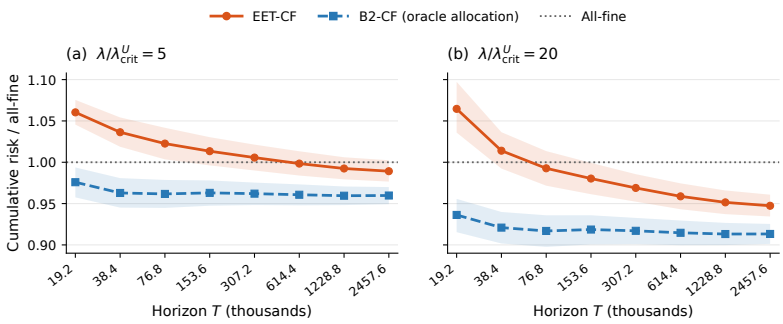

Figure 1: **Finite-budget value of online adaptation.** Cumulative fine-prediction risk is normalized by all-fine, so values below 1 indicate an improvement. EET-CF is EET with the cross-fitted estimator in Algorithm 1; B2-CF uses the same estimator with the oracle unknown-aggregation static share and still treats $w_\star$ as unknown. Shaded bands are paired-bootstrap 95% intervals over 20 seeds. In the $20\times$ regime EET-CF clearly improves on all-fine; in the $5\times$ regime the mean improvement appears later and remains smaller.

6 Numerical Experiments

We use two synthetic instances to test complementary consequences of the theory: $d = 6, K = 5$ widens the known/unknown break-even gap, while $d = 20, K = 5$ is the canonical online instance. All losses are exact population fine-prediction risks under the whitened covariance; full protocols, estimators, seeds, and diagnostics appear in Appendix B.

Structural effect of unknown aggregation. We first isolate the price of not knowing the aggregation direction. For $d = 6$ and $K = 5$, the known-aggregation break-even point is $\lambda = 5$, whereas the unknown-aggregation boundary is $\lambda = 15$. At budget $B = 76{,}800$, the known-$w_\star$ design achieves fixed-budget gains relative to all-fine of 2.87% (95% CI $[1.35, 4.40]\%$) at $\lambda = 10$ and 4.66% ($[3.03, 6.30]\%$) at $\lambda = 14$, compared with the closed-form targets 2.00% and 3.68%; both targets lie inside the corresponding intervals. The unknown-$w_\star$ optimum remains all-fine throughout this intermediate region. Appendix B.3 reports the full sweep and both tested budgets.

Online adaptation. For the online experiment we use $d = 20$, $K = 5$, $w_\star = (0.30, 0.25, 0.20, 0.15, 0.10)$, $c_{\mathrm{F}} = 5$, $c_{\mathrm{C}} = 1$, $\sigma_{\mathrm{F}} = 1$, independent Rademacher covariates, and independent Gaussian query noise. We vary σ_{C} and report $\lambda/\lambda^{\mathrm{crit}}_{\mathrm{U}} \in \{5, 20\}$ over 20 seeds. At $T = 2{,}457{,}600$, $TL_T/\Phi^\star_{\mathrm{U}}$ is 1.031 versus 1.028 at $5\times$ and 1.026 versus 1.023 at $20\times$ for EET-CF versus B2-CF, showing that the learned allocation is numerically close to the oracle-share benchmark. Figure 1 reports cumulative fine-prediction risk relative to all-fine across the tested horizons. EET-CF improves the final mean by 1.08% at $5\times$, although the 95% interval includes zero, and by 5.26% at $20\times$ (95% CI $[3.92, 6.56]\%$). The slower cumulative approach is consistent with a bounded additive lower-order term, whose relative effect decays only logarithmically; Appendix B.5 reports the direct remainder diagnostic. Below and at the break-even boundary, all-fine is the static optimum and the online learner retains a visible transient exploration cost; Appendix B.6 reports that diagnostic.

7 Discussion and Conclusion

In this work, we study cost-aware two-resolution supervision when expensive fine labels reveal a vector response and cheaper coarse labels reveal an unknown scalar aggregation. The theory gives a closed-form break-even rule, optimal resolution mixes, and an online estimate-and-track policy that attains the optimal first-order cumulative-risk coefficient with a matching local asymptotic minimax lower bound. Experiments illustrate the structural cost of unknown aggregation and show EET approaching the oracle-share benchmark, with finite-budget gains when coarse supervision is favorable. The analysis isolates a linear-Gaussian setting with a fixed aggregation rule and pre-covariate resolution choice. We discuss richer extensions in Appendix A.2.

AI USE STATEMENT

Generative AI tools were used as research assistants for theoretical hypothesis refinement, proof development of mathematical claims, experimental design and implementation including synthetic-data generation, literature search, figure preparation, and drafting, polishing and editing portions of the manuscript. Dataset cleaning/reformatting and qualitative or thematic data analysis are not applicable to this work. The authors directed the research, made the final decisions on the problem formulation, assumptions, theorem statements, proof arguments, experimental protocols, interpretation, and presentation, and reviewed all AI-assisted outputs. Literature claims were manually checked against primary sources, and AI-assisted code and numerical results were tested and independently checked for consistency with the stated theory. The authors take full responsibility for the final content of this work, including any text, claims, code, or artifacts produced with AI assistance.

REPRODUCIBILITY STATEMENT

The statistical model, assumptions, performance metric, and acquisition policy are specified in Sections 3 to 5. The full synthetic-data protocol, estimator definitions, seed sets, uncertainty calculations, and additional numerical diagnostics are recorded in Section B.1 and the remainder of Appendix B. Complete proofs and auxiliary calculations for all formal results appear in Appendix C. The anonymized supplementary code contains the scripts and fixed configurations used to reproduce the reported experiments and figures.

REFERENCES

Matteo Croci, Karen E Willcox, and Stephen J Wright. Multi-output multilevel best linear unbiased estimators via semidefinite programming. *Computer Methods in Applied Mechanics and Engineering*, 413:116130, 2023.

Thomas Dixon, Alex A Gorodetsky, John Jakeman, Akil Narayan, and Yiming Xu. Optimally balancing exploration and exploitation to automate multifidelity statistical estimation. *SIAM/ASA Journal on Uncertainty Quantification*, 14(3):1104–1133, 2026.

Pinar Donmez and Jaime G Carbonell. Proactive learning: cost-sensitive active learning with multiple imperfect oracles. In *Proceedings of the 17th ACM conference on Information and knowledge management*, pp. 619–628, 2008.

Lenka Filová, Pál Somogyi, and Radoslav Harman. Optimal exact designs of multiresponse experiments under linear and sparsity constraints. *Applied Stochastic Models in Business and Industry*, 42(2):e70072, 2026.

Xavier Fontaine, Pierre Perrault, Michal Valko, and Vianney Perchet. Online a-optimal design and active linear regression. In *International Conference on Machine Learning*, pp. 3374–3383. PMLR, 2021.

Dimitris Fotakis, Alkis Kalavasis, Vasilis Kontonis, and Christos Tzamos. Efficient algorithms for learning from coarse labels. In *Conference on Learning Theory*, pp. 2060–2079. PMLR, 2021.

Richard D Gill and Boris Y Levit. Applications of the van trees inequality: a bayesian cramér-rao bound. *Bernoulli*, 1995.

Peiyun Hu, Zachary C Lipton, Anima Anandkumar, and Deva Ramanan. Active learning with partial feedback. In *International Conference on Learning Representations*, 2019.

Leopold Mareis and Mathias Drton. Cost-aware optimized front-door experimental design. *arXiv preprint arXiv:2603.22024*, 2026.

Shinnosuke Matsuo, Riku Togashi, Ryoma Bise, Seiichi Uchida, and Masahiro Nomura. Instance-wise supervision-level optimization in active learning. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4939–4947. IEEE, 2025.

Phuong Ngoc Nguyen, Kaito Shiku, Ryoma Bise, Seiichi Uchida, and Shinnosuke Matsuo. Leveraging vision-language models as weak annotators in active learning. *arXiv preprint arXiv:2605.00480*, 2026.

Benjamin Peherstorfer. Multifidelity monte carlo estimation with adaptive low-fidelity models. *SIAM/ASA Journal on Uncertainty Quantification*, 7(2):579–603, 2019.

Benjamin Peherstorfer, Karen Willcox, and Max Gunzburger. Optimal model management for multifidelity monte carlo estimation. *SIAM Journal on Scientific Computing*, 38(5):A3163–A3194, 2016.

Thomas P Prescott, David J Warne, and Ruth E Baker. Efficient multifidelity likelihood-free bayesian inference with adaptive computational resource allocation. *Journal of Computational Physics*, 496:112577, 2024.

Guy Rotman and Roi Reichart. Multi-task active learning for pre-trained transformer-based models. *Transactions of the Association for Computational Linguistics*, 10:1209–1228, 2022.

Guillaume Sagnol. Computing optimal designs of multiresponse experiments reduces to second-order cone programming. *Journal of Statistical Planning and Inference*, 141(5):1684–1708, 2011.

Guillaume Sagnol and Radoslav Harman. Computing exact d-optimal designs by mixed integer second-order cone programming. *The Annals of Statistics*, 2015.

Daniel Schaden and Elisabeth Ullmann. On multilevel best linear unbiased estimators. *SIAM/ASA Journal on Uncertainty Quantification*, 8(2):601–635, 2020.

Samuel D Silvey and DM Titterington. A geometric approach to optimal design theory. *Biometrika*, 1973.

Pál Somogyi, Samuel Rosa, and Radoslav Harman. A randomized exchange algorithm for optimal design of multi-response experiments. *Metrika*, 89(2):217–242, 2026.

Moudar Soumaya, Norbert Gaffke, and Rainer Schwabe. Optimal design for multivariate observations in seemingly unrelated linear models. *Journal of Multivariate Analysis*, 142:48–56, 2015.

Javier Gamazo Tejero, Martin S Zinkernagel, Sebastian Wolf, Raphael Sznitman, and Pablo Márquez-Neila. Full or weak annotations? an adaptive strategy for budget-constrained annotation campaigns. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11381–11391, 2023.

Aad W Van der Vaart. *Asymptotic statistics*, volume 3. Cambridge university press, 2000.

Byron C Wallace, Kevin Small, Carla E Brodley, and Thomas A Trikalinos. Who should label what? instance allocation in multiple expert active learning. In *Proceedings of the 2011 SIAM international conference on data mining*, pp. 176–187. SIAM, 2011.

Yiming Xu, Vahid Keshavarzzadeh, Robert M Kirby, and Akil Narayan. A bandit-learning approach to multifidelity approximation. *SIAM Journal on Scientific Computing*, 44(1):A150–A175, 2022.

Yifan Yan and Sheng-Jun Huang. Cost-effective active learning for hierarchical multi-label classification. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*, pp. 2962–2968, 2018.

Appendix

APPENDIX ROADMAP

Appendix A expands the literature comparison and discusses limitations and extensions. Appendix B gives the complete synthetic protocols, uncertainty calculations, structural comparisons, and finite-budget diagnostics. Appendix C contains the proofs: Proposition 1 in Appendix C.1, Theorem 2 in Appendix C.2, Theorem 3 in Appendix C.3, Lemma 4 in Appendix C.4, Corollary 5 in Appendix C.5, and the online theorem and its three supporting lemmas in Appendices C.6–C.9.

Frequently used notation.

Symbol	Meaning
B, b, T	static budget, running budget level, and cumulative-risk horizon
B_0	end of finite initialization and start of the reported budget clock
$\mathsf{B}_1, \mathsf{B}_2, \mathsf{B}_3$	known-$w_\star$ oracle, unknown-aggregation oracle-share, online learned allocation
$N_{\mathrm{F}}, N_{\mathrm{C}}$	fine and coarse query counts
$\eta, \widehat{\eta}_m, \eta_b^{\mathrm{E}}$	static, epoch-target, and realized estimation-stream coarse shares
λ	effective coarse-information ratio in Equation (7)
$\psi_{\mathrm{K}}, \psi_{\mathrm{U}}$	known/unknown static risk objectives
$G_{\mathrm{K}}, G_{\mathrm{U}}$	known/unknown leading gains relative to all-fine
$A_{\mathrm{U}}, D_{\mathrm{U}}$	fine-only and fine-plus-coarse multiplicities under unknown aggregation
$\Phi_{\mathrm{K}}^\star, \Phi_{\mathrm{U}}^\star$	optimal known/unknown first-order cumulative-risk coefficients
Q, w_0, v, ϑ	simplex tangent basis, center, coordinate, and joint local parameter
$\mathcal{U}_\delta$	local regular neighborhood around $\vartheta_\star$
B_m, ζ_m	dyadic epoch boundary and design-stream budget fraction
$\Xi(b), b_{\mathrm{E}}$	design-stream spending and remaining estimation budget
$\mathcal{G}_m$	design-stream history determining epoch-m allocation
L_x	almost-sure whitened covariate-norm bound

A EXTENDED DISCUSSION

This section expands the literature comparison and the limitations discussed briefly in the main text, focusing on aspects that interact directly with the two-resolution information structure.

A.1 EXTENDED RELATED WORK

Multiresponse optimal design and heterogeneous measurements. Classical and modern optimal-design theory provides the broad statistical framework for combining heterogeneous measurements. Silvey & Titterington (1973) develops a geometric approach to optimal design, while Sagnol (2011) shows that several multiresponse criteria, including A-optimality, can be handled under linear resource constraints. Related work considers multivariate observations and increasingly general response structures (Soumaya et al., 2015; Somogyi et al., 2026; Filová et al., 2026; Sagnol & Harman, 2015). In the known-aggregation benchmark of our model, vectorizing Θ turns a fine query and a coarse query into two known information matrices, so the static problem is a structured two-experiment instance of this general framework. The special rank-one response structure reduces the cost-constrained design to a scalar budget share and yields the explicit threshold in Section 4. Our main unknown-aggregation setting differs because the coarse observation operator itself contains an unknown nuisance parameter.

A related measurement-design viewpoint appears in Mareis & Drton (2026), who optimize complete and partial measurements under collection costs for a causal estimand. This work reinforces the point that a full-versus-partial acquisition decision can be statistically central without being specific to active learning. Our target and observation geometry are different: the target here is the full multivariate regression map, while the cheaper observation is a scalar projection of that same map.

Multifidelity estimation and adaptive resource allocation. Multifidelity methods also allocate finite budgets across expensive and cheap information sources. Peherstorfer et al. (2016) derives

analytic cost-optimal model-management rules for multifidelity Monte Carlo, while Schaden & Ullmann (2020) and Croci et al. (2023) study budget-optimal linear unbiased estimation and multi-output extensions. The relationship between fidelities need not be known in advance: Peherstorfer (2019) adapts a low-fidelity model, Prescott et al. (2024) estimates allocation-relevant quantities during multifidelity inference, Xu et al. (2022) learns cross-model statistical relationships while allocating a budget, and Dixon et al. (2026) balances pilot estimation against the final multifidelity estimator.

These papers are close to our online problem at the level of resource allocation: all must learn quantities that determine whether a cheap source is useful. The inferential distinction is that $w_\star$ is part of the conditional-mean observation law $u_\star = \Theta_\star w_\star$, so uncertainty about the cross-resolution relationship changes the efficient information for the target itself. Profiling out $w_\star$ removes exactly $K - 1$ coarse-informative directions and changes the static break-even boundary. Online optimal design gives another neighboring perspective: Fontaine et al. (2021) learns an A-optimal scalar-regression design under unknown heteroscedastic noise, whereas our acquisition action chooses response resolution before observing the current covariate.

Cost-sensitive active learning and annotation granularity. Cost-sensitive active learning has long treated the annotation source as part of the acquisition decision. Donmez & Carbonell (2008) select among imperfect oracles while accounting for cost and reliability, and Wallace et al. (2011) allocate instances across experts of different quality and cost. Other work varies the information requested from the same example: Hu et al. (2019) studies partial-feedback queries, Yan & Huang (2018) considers hierarchy-dependent query costs, and Rotman & Reichart (2022) allows task-specific annotations with different granularities and costs. More recent work makes supervision level especially explicit, including full versus weak annotations (Tejero et al., 2023) and instance-wise supervision-level optimization (Matsuo et al., 2025).

A related literature studies inference from intrinsically coarse labels. Fotakis et al. (2021) formalizes categorical coarsening through partitions of a latent fine-label space and derives learning guarantees that depend on the information preserved by coarsening. Nguyen et al. (2026) considers expensive fine-grained human labels together with cheaper coarse VLM labels in active learning. These settings motivate resolution-aware acquisition but use different observation structures. Here the learner purchases either a K-dimensional response or a scalar projection of that same response, while the target remains the full vector-valued conditional mean.

A.2 Limitations and Extensions

Joint item and resolution selection. The canonical model chooses the resolution before observing the current covariate. This isolates the global value of each resolution while retaining covariates in the prediction problem itself. If x_t were available before the action, the learner would need to decide both *where* information is most valuable and *which resolution* to purchase. The resulting problem is a matrix-valued optimal-design problem rather than the one-dimensional frontier derived here.

Richer cross-resolution structure. Several extensions change the cross-resolution information geometry, including shared item-level effects, multiple or heterogeneous coarse channels, nonlinear aggregation, and unknown noise levels; we discuss these separately below. Near the asymptotic break-even boundary, finite-budget benefits can also be delayed by estimation and adaptation costs.

Shared item-level noise and repeated measurements. The action space in the present paper requests only one resolution for each arriving item, so no joint fine/coarse noise law is needed for the main likelihood. If an extension allowed both resolutions to be collected on the same item, a shared latent item effect or correlated measurement noise could make pairing informative beyond two separately sampled observations. Such a model would also make the effective coarse variance depend on the form of the shared component and could require learning additional noise parameters.

Multiple resolutions and heterogeneous aggregation. We analyze one fine channel and one scalar coarse channel with a single fixed aggregation vector. Multiple coarse channels, several levels of granularity, or annotator-dependent aggregation vectors would replace the single rank-one direction by a collection of partially overlapping information subspaces. The same efficient-information

viewpoint may still apply, but the phase diagram could contain several entry thresholds rather than one scalar break-even point.

Scale normalization and unconstrained aggregation. The simplex assumption fixes the scale of the coarse aggregation and gives a $(K-1)$-dimensional nuisance tangent. If $w_\star$ were instead locally unconstrained in $\mathbb{R}^K$ and $\Theta_\star$ had full column rank with $K < d$, the same tangent-space calculation gives

$$A = d(K-1) + K, \qquad D = d - K, \qquad \lambda_{\mathrm{crit}} = \frac{dK}{d-K}.$$

Thus the efficient-information argument and scalar allocation optimization carry over, but with a larger nuisance price. We do not make a corresponding intercept claim because an additional offset interacts differently with the centered covariate model and requires a separate derivation.

Nonlinear aggregation. The exact decomposition relies on the coarse conditional mean being linear in the fine response. With a smooth nonlinear aggregator, the locally informative direction can depend on the current parameter and potentially on the item itself. The coarse-information subspace would then vary across observations, and the static design would generally no longer collapse to a single scalar allocation share.

High-dimensional structure. The theory keeps d and K fixed as the annotation budget grows. If either dimension grows, or if $\Theta_\star$ is sparse or low rank, projected least squares and classical efficiency must be replaced by regularized estimators and a corresponding high-dimensional analysis. Structural assumptions may also change the number of directions that coarse feedback can usefully inform.

Estimated covariance and noise levels. The main theorem treats Σ_x and the two noise variances as known; exact whitening is justified by free unlabeled covariates. The allocation depends on noise and cost through their cost-adjusted ratio. If an independent or cross-fitted estimator of that ratio has mean-square error no larger than the aggregation-weight contribution, the same quadratic plug-in argument suggests that the first-order allocation result should persist. We leave a complete theorem for estimated variances to future work.

Expected rather than time-uniform guarantees. Our main performance criterion is ex-ante expected cumulative fine-prediction risk. A time-uniform or high-probability guarantee would require simultaneous control of the sequential score, the pilot, and the allocation error over all budget levels. Such a result is stronger and may introduce additional logarithmic factors even if the first-order expected-risk coefficient is unchanged.

B Extended Numerical Experiments

This section records the full synthetic protocol, uncertainty calculations, and additional diagnostics behind Section 6. The experiments test quantitative consequences of the linear theory; they are not evidence that the canonical model is an accurate description of a particular real annotation system.

B.1 Experimental Setup and Reproducibility

Canonical $d = 20, K = 5$ instance. The static-convergence and online experiments use $w_\star = (0.30, 0.25, 0.20, 0.15, 0.10)$, $c_{\mathrm{F}} = 5$, $c_{\mathrm{C}} = 1$, and $\sigma_{\mathrm{F}} = 1$. The columns of $\Theta_\star$ are the first five columns of a deterministic orthonormal DCT-II basis in $\mathbb{R}^{20}$, so $\Theta_\star^\top \Theta_\star = I_5$. Covariates have independent Rademacher coordinates, and queried observations use independent Gaussian noise. For this instance, $A_{\mathrm{U}} = 84$, $D_{\mathrm{U}} = 16$, and $\lambda_{\mathrm{U}}^{\mathrm{crit}} = 6.25$. We vary only σ_{C} to obtain the desired effective ratios.

The static experiments use 200 seeds, 42001–42200, common random numbers across allocations, and exact population fine-prediction risk. The larger-budget confirmation uses $B \in \{19200, 38400, 76800\}$ together with the earlier budgets $2400, 4800, 9600$; the exact theory-optimal allocations are evaluated in the $2\times$, $5\times$, and $20\times$ regimes.

The online experiment uses exactly 20 seeds, 44001–44020, initialization budget $B_0 = 9600$, dyadic epochs, design share $\zeta_m = 0.20/(m+2)^2$, and a fixed 50% coarse budget share inside each design block. The reported regimes are $\lambda/\lambda_{\mathrm{U}}^{\mathrm{crit}} \in \{0.75, 1, 5, 20\}$. All four use horizons through 614400; the two mixed regimes additionally use 1228800 and 2457600.

Initialization and cumulative-risk clock. For the canonical online instance, initialization through B_0 is fixed in advance: 40 design-fine labels, 40 design-coarse labels, 40 mandatory estimation-fine labels, and 40 mandatory estimation-coarse labels are collected first, after which the remaining pre-B_0 budget purchases 1,824 additional estimation-fine labels. Thus the design stream contains 40 fine and 40 coarse observations (cost 240), the estimation stream contains 1,864 fine and 40 coarse observations (cost 9,360), and there is no unused residue. The design observations are used only to form the first allocation target; the predictor at B_0 is fit from the estimation folds only. The reported cumulative risk starts inclusively at B_0: $\mathfrak{R}_T = \sum_{b=B_0}^{T} L_b$.

Cross-fitted EET estimator. The design stream estimates $w_\star$ and the target allocation. Estimation observations are permanently assigned by deterministic parity within each resolution stream to two folds before the current label is observed. At every completed estimation query, each score fold takes one expected-Fisher scoring step from a guarded projected pilot fitted on the opposite fold; the two fine-map estimates are averaged and projected onto the common operator-norm ball. The predictor is held fixed between estimation-query completions, including throughout design blocks. The static B_2 comparator uses the same cross-fitted estimator but is given only the locally optimal unknown-aggregation static share; its estimator still treats $w_\star$ as unknown. At the all-fine boundary the efficient estimator reduces to pooled guarded projected fine regression, so we use that estimator rather than artificially splitting the fine-only data.

Uncertainty. For online curves, each point is a ratio of aggregate mean cumulative risks. The shaded intervals are pointwise 95% percentile intervals from 20,000 paired bootstrap resamples of the 20 common seeds, using fixed bootstrap seed 48001. Paired method differences use seedwise differences with mean ± 1.96 standard errors. For static gains, the primary quantity is the ratio-of-means gain $1 - \bar{L}_{\mathrm{method}}/\bar{L}_{\mathrm{AF}}$; its 95% interval uses a paired delta-method calculation. The cumulative risk is the exact integer-budget sum of held population risk, not a checkpoint interpolation or test-set estimate.

No scoring guard, fallback, or final-projection event occurred in the online estimator evaluations. The anonymized supplementary code provides the fixed configurations and scripts needed to regenerate the reported experiments, tables, and figures from the documented random seeds.

B.2 Static Convergence Check

The first static experiment asks whether the finite-budget unknown-aggregation risk approaches the closed-form coefficient from Theorem 3. All six budgets use the same two-fold cross-fitted projected one-step estimator as the static B_2 construction. Figure 2 reports paired uncertainty over the 200 seeds. At $B = 76800$, the empirical gains in the $2\times$, $5\times$, and $20\times$ regimes are 1.226%, 4.848%, and 9.617%, compared with theoretical values 1.548%, 5.271%, and 10.012%. The last-step changes in coefficient excess are much smaller than their paired standard errors, so the slight visual flattening or increase at the largest budgets is not resolved statistically.

B.3 Known versus Unknown Aggregation

To make the structural cost of unknown aggregation visible, we use a second instance with $d = 6$ and $K = 5$. The first five DCT-II basis vectors define $\Theta_\star$; $w_\star$, costs, and fine-noise variance are unchanged. Here the known-aggregation break-even point is 5, whereas the unknown-aggregation threshold is 15. We evaluate $\lambda \in \{4, 5, 7.5, 10, 14, 15, 20, 40\}$ at budgets 19200 and 76800 using 200 seeds. B_1 receives the known-$w_\star$ theory-optimal share and uses efficient known-weight Gaussian regression; B_2 receives the unknown-$w_\star$ theory-optimal share and uses the cross-fitted estimator.

Figure 3 shows both budgets. The vertical axis is the fixed-budget gain relative to all-fine, $1 - L_{\mathrm{method}}/L_{\mathrm{AF}}$: a value of 4% means 4% lower fine-prediction risk at the same annotation budget. In the leading theory, the same percentage is also the annotation-cost saving at a fixed target risk.

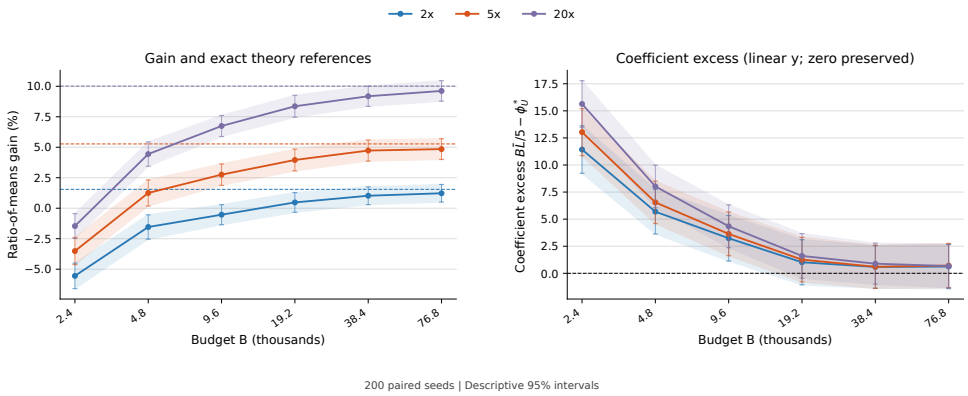

Figure 2: **Static convergence in the canonical $d = 20, K = 5$ instance.** Left: ratio-of-means annotation gain with paired 95% intervals and exact theory references. Right: first-order coefficient excess with descriptive 95% intervals. All points use the same 200 paired seeds.

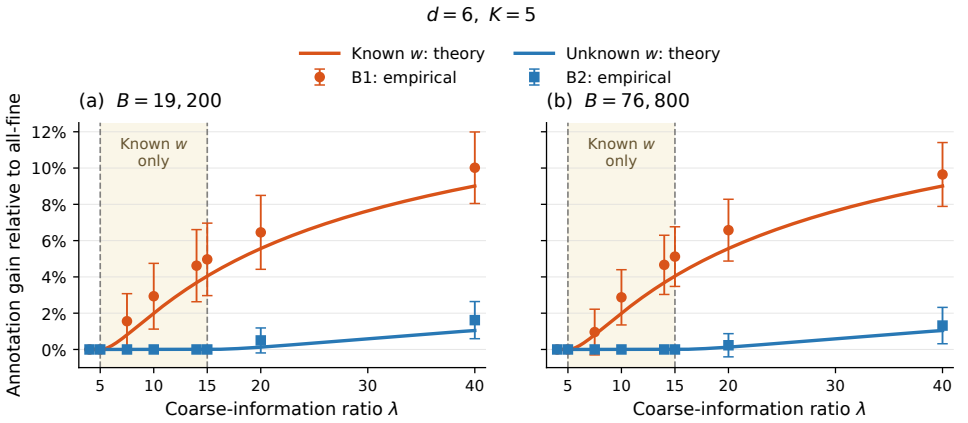

Figure 3: **Known versus unknown aggregation, $d = 6$, $K = 5$.** Theory curves and empirical gains are shown at two budgets. The shaded region lies between the known-aggregation threshold 5 and the unknown-aggregation threshold 15. Error bars are 95% intervals over 200 seeds.

The shaded interval $5 < \lambda \leq 15$ is the regime in which known aggregation recommends a mixed design while unknown aggregation recommends all-fine acquisition. The exact known-$w_\star$ gains are 2.000% at $\lambda = 10$ and 3.683% at $\lambda = 14$. At $B = 76800$, the corresponding empirical gains are 2.874% (95% CI $[1.352, 4.397]\%$) and 4.664% ($[3.032, 6.296]\%$); both theory targets lie inside the empirical intervals. B_2 is exactly all-fine throughout this intermediate region by the pre-specified theory-optimal allocation. The experiment illustrates the consequence of the analytic thresholds; it does not independently estimate them from an allocation grid.

The size of this structural gap depends strongly on d/K. For fixed $K = 5$, the unknown threshold $5d/(d-4)$ equals 25 at $d = 5$, 15 at $d = 6$, 10 at $d = 8$, 6.25 at $d = 20$, and approaches the known threshold 5 as d grows. Equivalently, the relative loss in the maximal gain ceiling is $(K-1)/d$.

B.4 Online Allocation Learning and Finite-Budget Value

The main text reports cumulative risk relative to all-fine. The allocation path itself depends only on the independent design stream, not on estimation labels. Figure 4 shows that the learned and realized coarse shares approach the distinct theoretical targets 0.042891 and 0.029955 in the $5\times$ and $20\times$ regimes. The smaller spending share at $20\times$ is not paradoxical: when each coarse dollar is much more informative, less coarse spending is needed to complement the fine-only directions.

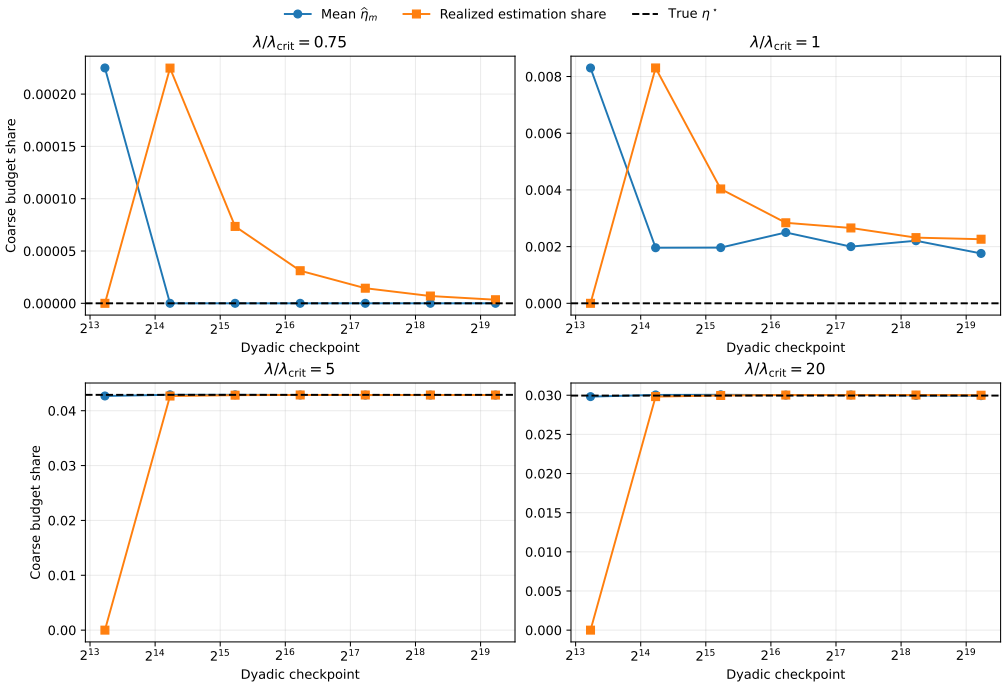

Figure 4: **Allocation learning.** The design-stream plug-in target and the realized estimation-stream coarse-budget share track the theoretical optimal share. Estimation labels never feed back into the allocation rule.

Regime	Final EET coeff.	Final B2 coeff.	EET/all-fine	Mean crossover	95% crossover
5×	1.031	1.028	0.989	614,400	none
20×	1.026	1.023	0.947	76,800	153,600

Table 1: **Online summary at $T = 2{,}457{,}600$.** The coefficient columns report $TL_T/\Phi^\star_{\mathrm{U}}$. "Mean crossover" is the first tested horizon at which mean cumulative EET risk is below all-fine; "95% crossover" additionally requires the paired pointwise 95% interval for EET minus all-fine to be entirely below zero.

At the largest tested horizon, EET-CF is numerically close to B_2 in *instantaneous* first-order coefficient even though a positive cumulative gap remains from earlier adaptation. Table 1 summarizes the two mixed regimes. The B_2 column is an oracle-share reference that EET is intended to approach, not beat.

At the final horizon, the paired cumulative EET-minus-B_2 gaps are 80.42 (95% CI $[59.94, 100.90]$) and 93.26 ($[79.38, 107.15]$) in the $5\times$ and $20\times$ regimes. A nonzero finite cumulative gap is compatible with equal leading $\log T$ coefficients; the theorem does not require the difference to vanish.

B.5 First-Order Remainder Diagnostic

To inspect the lower-order term directly, we subtract the discrete first-order benchmark $H_T := \sum_{b=B_0}^{T} b^{-1}$. Figure 5 plots $\mathfrak{R}_T - \Phi^\star H_T$, using $\Phi^\star_{\mathrm{U}}$ for EET-CF and B_2 and the all-fine coefficient $dKc_{\mathrm{F}}\sigma^2_{\mathrm{F}}$ for all-fine. Over the last three horizons, all consecutive-increment 95% intervals contain zero. The displayed tails are therefore visually consistent with flattening, but with 20 paths they do not distinguish a bounded plateau from modest continuing drift. We use this figure only as a finite-horizon diagnostic, not as empirical proof of the $O(1)$ remainder.

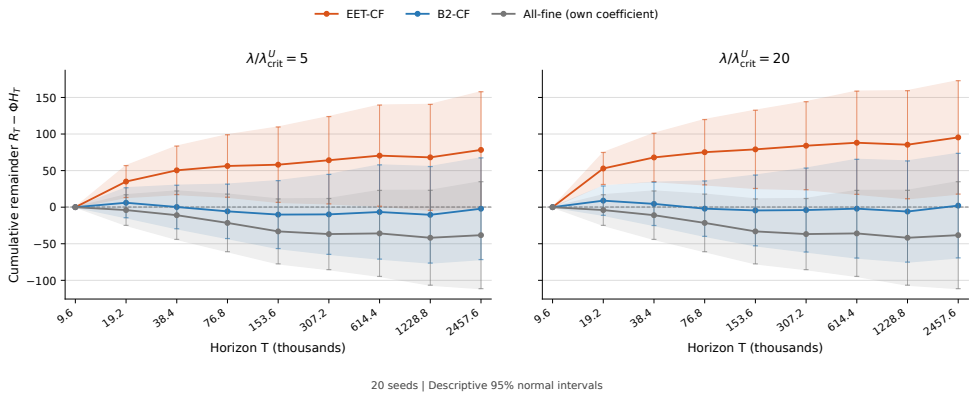

Figure 5: **First-order remainder diagnostic.** Cumulative risk after subtracting the exact discrete harmonic first-order benchmark. Bands are descriptive pointwise 95% intervals over the 20 saved paths. The late trajectories are consistent with a bounded lower-order term but remain too noisy to establish a plateau.

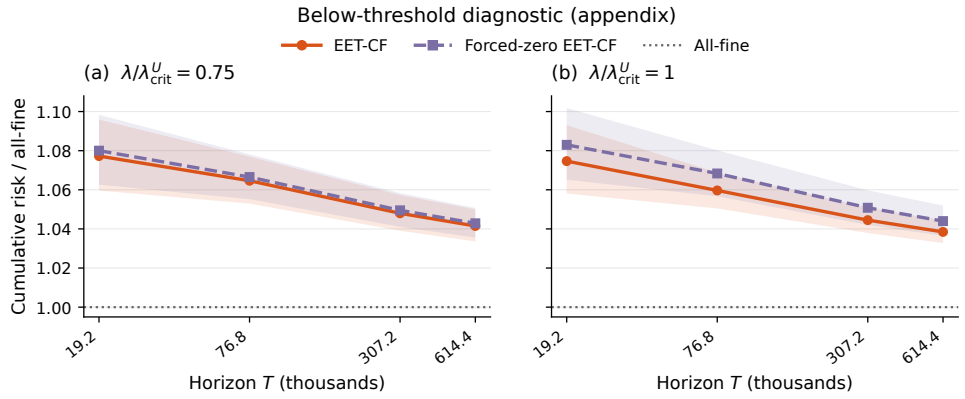

Figure 6: **Below-threshold finite-budget diagnostic.** Cumulative risk is normalized by all-fine. Forced-zero acquisition preserves design spending but removes positive post-initialization target shares. Its failure to reduce the gap indicates that erroneous positive coarse targets are not the main source of the remaining transient overhead. Bands are paired-bootstrap 95% intervals.

B.6 Below-Threshold Finite-Budget Diagnostic

At $0.75\lambda_{\mathrm{U}}^{\mathrm{crit}}$ and $\lambda_{\mathrm{U}}^{\mathrm{crit}}$, all-fine is the static optimum. We therefore compare normal EET-CF with a diagnostic variant that forces the post-initialization estimation target to zero while preserving the same design-stream spending rule. Figure 6 shows that forcing zero does not remove the roughly 4% cumulative overhead at $T = 614400$. In the $0.75\times$ regime, 18 of 20 normal and forced-zero post-initialization acquisition paths are identical; at the exact boundary they differ more often because finite-sample plug-in estimates cross the threshold.

At $T = 614400$, the forced-zero design stream spends 4346 cost units and retains the fixed 40-unit estimation-coarse initialization. The corresponding endpoint fine-count inflation proxy is about 0.72%, while cumulative overhead is about 4.3–4.4%. Earlier opportunity costs were larger and remain integrated into cumulative risk, and finite-sample estimation effects can persist. These accounting quantities are diagnostics, not a causal decomposition.

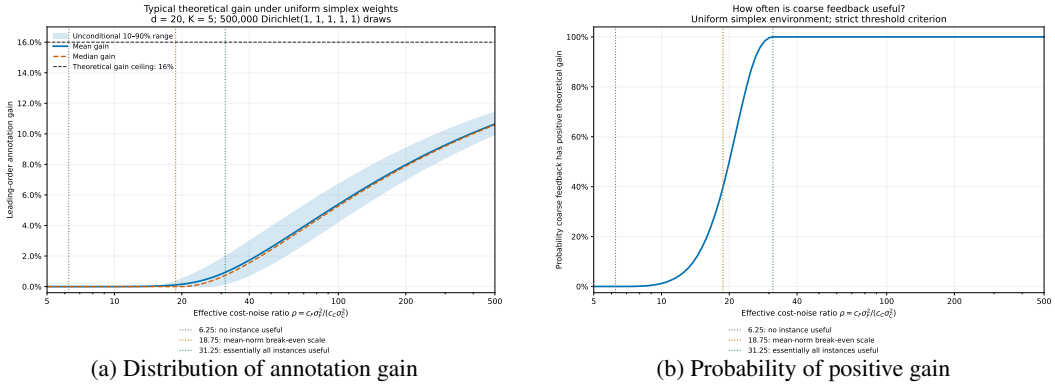

(a) Distribution of annotation gain (b) Probability of positive gain

Figure 7: **Theorem-magnitude illustration under random simplex weights.** Positive gain can become common before its magnitude becomes substantial. The variation across sampled weights is heterogeneity under the illustrative Dirichlet environment, not estimation uncertainty.

B.7 Distributional Illustration of Annotation Gain

Finally, we illustrate the magnitude of the static gain over heterogeneous aggregation vectors. Fix $d = 20$, $K = 5$ and draw 500,000 independent weights from $\mathrm{Dirichlet}(1, \dots, 1)$. For each draw, set $\rho = c_{\mathrm{F}}\sigma_{\mathrm{F}}^2/(c_{\mathrm{C}}\sigma_{\mathrm{C}}^2)$ and evaluate the closed-form gain at $\lambda = \rho\|w\|_2^2$. This is Monte Carlo integration of a theorem quantity: no covariates or labels are generated and no estimator is fitted. At $\rho = 20$, approximately half of the sampled weights have positive gain but the unconditional mean gain is only 0.15%; the mean rises to approximately 1.73%, 5.39%, and 7.97% at $\rho = 40, 100, 200$.

C Proof Details

This section gives complete proofs of the formal results in Sections 4 and 5. We work in whitened coordinates throughout. Let $w_0 := K^{-1}\mathbf{1}$, let $Q \in \mathbb{R}^{K\times(K-1)}$ have orthonormal columns spanning $\mathbf{1}^\perp$, and write $w(v) = w_0 + Qv$. We use column stacking $\beta := \mathrm{vec}(\Theta)$. Constants may depend on the fixed dimensions, costs, variances, L_x, τ, κ_Θ, and M_Θ, but not on the running budget. Unless otherwise stated, these constants are uniform over the regular parameter class in Assumptions 1 to 3.

C.1 Proof of Proposition 1

For one fine observation, the derivative of the conditional mean with respect to β is the usual multivariate-regression design. Since $\mathbb{E}[xx^\top] = I_d$, its expected information is $\sigma_{\mathrm{F}}^{-2} I_{dK}$. A coarse observation has mean $x^\top\Theta w_\star = (w_\star^\top \otimes x^\top)\beta$, so its expected information is

$$\sigma_{\mathrm{C}}^{-2}(w_\star w_\star^\top \otimes I_d). \tag{22}$$

Summing N_{F} fine and N_{C} coarse observations gives Equation (5).

Let $\bar{w}_\star = w_\star/\|w_\star\|_2$. Any perturbation $H \in \mathbb{R}^{d\times K}$ has the unique orthogonal decomposition

$$H = h\bar{w}_\star^\top + H_0, \qquad H_0\bar{w}_\star = 0, \tag{23}$$

with $\|H\|_F^2 = \|h\|_2^2 + \|H_0\|_F^2$. The H_0 subspace has dimension $d(K-1)$ and receives no coarse information because $H_0 w_\star = 0$. Its information eigenvalue is therefore $\alpha := N_{\mathrm{F}}/\sigma_{\mathrm{F}}^2$. On the complementary d-dimensional subspace, the eigenvalue is $\alpha+\gamma\|w_\star\|_2^2$, where $\gamma := N_{\mathrm{C}}/\sigma_{\mathrm{C}}^2$. Taking the inverse trace gives Equation (6).

For completeness, consider an exogenous static sequence with $N_{\mathrm{F}} \asymp B$ and either $N_{\mathrm{C}} = 0$ or $N_{\mathrm{C}} \asymp B$, which covers the cost-optimal static designs below. Conditional on the sampled covariates, the Gaussian weighted least-squares estimator has information

$$J_B = \frac{1}{\sigma_{\mathrm{F}}^2}(I_K \otimes S_{\mathrm{F}}) + \frac{1}{\sigma_{\mathrm{C}}^2}(w_\star w_\star^\top \otimes S_{\mathrm{C}}), \qquad S_a := \sum_{i:A_i=a} x_i x_i^\top, \tag{24}$$

and conditional covariance J_B^{-1}. Its expectation is exactly $\mathcal{I}_{\mathrm{kn}}$. Bounded isotropic covariates give exponentially small tails for $\|S_a/N_a - I_d\|_{\mathrm{op}}$. On the event where both relevant Gram matrices lie within one half of their expectations, write $J_B = \mathcal{I}_{\mathrm{kn}} + \Delta_B$. Since $\|\mathcal{I}_{\mathrm{kn}}^{-1}\|_{\mathrm{op}} = O(B^{-1})$, $\mathbb{E}\Delta_B = 0$, and $\mathbb{E}\|\Delta_B\|_{\mathrm{op}}^2 = O(B)$, the resolvent identity

$$J_B^{-1} = \mathcal{I}_{\mathrm{kn}}^{-1} - \mathcal{I}_{\mathrm{kn}}^{-1}\Delta_B\mathcal{I}_{\mathrm{kn}}^{-1} + \mathcal{I}_{\mathrm{kn}}^{-1}\Delta_B\mathcal{I}_{\mathrm{kn}}^{-1}\Delta_B J_B^{-1} \tag{25}$$

shows after taking traces and expectations that $\mathbb{E}\,\mathrm{tr}(J_B^{-1}) = \mathrm{tr}(\mathcal{I}_{\mathrm{kn}}^{-1}) + O(B^{-2})$; the bad-Gram event contributes exponentially little after bounded projection. Thus the projected weighted estimator has risk at most the displayed trace plus $O(B^{-2})$. At the all-fine boundary the same argument reduces to projected multivariate least squares. The fixed-dimensional Gaussian experiment is differentiable in quadratic mean and locally asymptotically normal. The local asymptotic minimax theorem for bowl-shaped loss therefore gives the inverse-information trace as the local lower bound (Van der Vaart, 2000, Theorem 8.11). Projection onto the fixed operator-norm ball bounds squared loss, so the corresponding truncated-loss argument and uniform integrability justify the expected-risk statement. Together these facts prove the first-order attainment statement in Proposition 1.

C.2 Proof of Theorem 2

Substituting $N_{\mathrm{F}} = (1-\eta)B/c_{\mathrm{F}}$ and $N_{\mathrm{C}} = \eta B/c_{\mathrm{C}}$ into Equation (6) gives Equation (8). Its first two derivatives are

$$\partial_\eta \psi_{\mathrm{K}}(\eta, \lambda) = \frac{K-1}{(1-\eta)^2} - \frac{\lambda - 1}{\{1 + (\lambda-1)\eta\}^2}, \tag{26}$$

$$\partial^2_{\eta\eta} \psi_{\mathrm{K}}(\eta, \lambda) = \frac{2(K-1)}{(1-\eta)^3} + \frac{2(\lambda-1)^2}{\{1 + (\lambda-1)\eta\}^3} > 0. \tag{27}$$

Hence the objective is strictly convex. At zero, $\partial_\eta \psi_{\mathrm{K}}(0, \lambda) = K - \lambda$. If $\lambda \le K$, strict convexity makes $\eta = 0$ the unique minimizer. If $\lambda > K$, the minimizer is interior and solves

$$\frac{\sqrt{K-1}}{1-\eta} = \frac{\sqrt{\lambda - 1}}{1 + (\lambda-1)\eta}. \tag{28}$$

Setting $q_{\mathrm{K}} = \sqrt{(\lambda-1)/(K-1)}$ and rearranging gives Equation (9).

Substitution yields

$$\psi^\star_{\mathrm{K}}(\lambda) = \frac{\left(\sqrt{(K-1)(\lambda-1)} + 1\right)^2}{\lambda} \tag{29}$$

above threshold. Since this tends to $K-1$ as $\lambda \to \infty$, Equation (12) follows directly.

C.3 Proof of Theorem 3

We first invoke Lemma 4, proved in Section C.4. It implies that for fixed counts the efficient inverse-information trace for $\Theta_\star$ is

$$\mathcal{R}_{\mathrm{U}}(N_{\mathrm{F}}, N_{\mathrm{C}}) := \frac{A_{\mathrm{U}}\sigma_{\mathrm{F}}^2}{N_{\mathrm{F}}} + \frac{D_{\mathrm{U}}}{N_{\mathrm{F}}/\sigma_{\mathrm{F}}^2 + N_{\mathrm{C}}\|w_\star\|_2^2/\sigma_{\mathrm{C}}^2}. \tag{30}$$

For fixed positive proportions, the joint Gaussian experiment in the local coordinate (β, v) is differentiable in quadratic mean and locally asymptotically normal. The local asymptotic minimax theorem (Van der Vaart, 2000, Theorem 8.11) gives the trace of the inverse efficient information as the local lower bound for the β target, yielding Equation (30). The projected two-fold one-step construction analyzed in Section C.8 attains the same leading covariance under an exogenous static schedule; when $N_{\mathrm{C}} = 0$, projected fine OLS gives the boundary case. The final projection onto the fixed operator-norm ball makes the squared loss uniformly bounded, so the asymptotic linear expansion also yields convergence of expected squared loss, not only distributional convergence.

At budget B and coarse spending share η, substitute $N_{\mathrm{F}} = (1-\eta)B/c_{\mathrm{F}}$ and $N_{\mathrm{C}} = \eta B/c_{\mathrm{C}}$ into Equation (30). With Equation (7),

$$\mathcal{R}_{\mathrm{U}}(B, \eta) = \frac{c_{\mathrm{F}}\sigma_{\mathrm{F}}^2 d}{B}\left[\frac{A_{\mathrm{U}}/d}{1-\eta} + \frac{D_{\mathrm{U}}/d}{1 + (\lambda-1)\eta}\right] + o(B^{-1}), \tag{31}$$

which is Equation (14).

To optimize, set $A = A_{\mathrm{U}}$ and $D = D_{\mathrm{U}}$ and ignore the common factor $1/d$. Define

$$\phi(\eta, \lambda) = \frac{A}{1-\eta} + \frac{D}{1+(\lambda-1)\eta}. \tag{32}$$

Its derivatives are

$$\partial_\eta \phi = \frac{A}{(1-\eta)^2} - \frac{D(\lambda-1)}{\{1+(\lambda-1)\eta\}^2}, \tag{33}$$

$$\partial^2_{\eta\eta} \phi = \frac{2A}{(1-\eta)^3} + \frac{2D(\lambda-1)^2}{\{1+(\lambda-1)\eta\}^3} > 0. \tag{34}$$

At zero, $\partial_\eta \phi(0, \lambda) = A - D(\lambda - 1)$, so the boundary changes sign at

$$1 + \frac{A}{D} = \frac{A+D}{D} = \frac{dK}{d-K+1}. \tag{35}$$

Above this boundary, solving the first-order condition gives

$$\eta^\star_{\mathrm{U}}(\lambda) = \frac{q_{\mathrm{U}} - 1}{\lambda - 1 + q_{\mathrm{U}}}, \qquad q_{\mathrm{U}} := \sqrt{\frac{D_{\mathrm{U}}(\lambda-1)}{A_{\mathrm{U}}}}. \tag{36}$$

This proves the threshold and optimizer.

The all-fine coefficient in the unnormalized objective is $A + D = dK$. Above threshold the exact gap is

$$dK - \phi^\star(\lambda) = \frac{A_{\mathrm{U}}(q_{\mathrm{U}} - 1)^2}{\lambda}, \tag{37}$$

so

$$G_{\mathrm{U}}(\lambda) = \frac{A_{\mathrm{U}}(q_{\mathrm{U}} - 1)^2}{\lambda dK}. \tag{38}$$

The envelope theorem gives

$$\frac{d}{d\lambda}\phi^\star(\lambda) = -\frac{D_{\mathrm{U}} \eta^\star_{\mathrm{U}}(\lambda)}{\{1+(\lambda-1)\eta^\star_{\mathrm{U}}(\lambda)\}^2} < 0 \tag{39}$$

above threshold, proving that the gain is strictly increasing there. Finally $\phi(\eta, \lambda) \geq A_{\mathrm{U}}$ for every feasible η, while $\phi^\star(\lambda) \to A_{\mathrm{U}}$ as $\lambda \to \infty$. Hence the gain ceiling is $1 - A_{\mathrm{U}}/(dK) = D_{\mathrm{U}}/(dK) = (d-K+1)/(dK)$.

C.4 Proof of Lemma 4

The simplex tangent at the interior point $w_\star$ is $\mathbf{1}^\perp$. Consider a local perturbation (H, h_w) with $H \in \mathbb{R}^{d \times K}$ and $h_w \in \mathbf{1}^\perp$. A fine observation has local mean perturbation $H^\top x$, while a coarse observation has local mean perturbation $x^\top(Hw_\star + \Theta_\star h_w)$. Therefore the expected Fisher-information quadratic form for fixed counts is

$$\mathcal{I}(H, h_w) = \alpha \|H\|_F^2 + \gamma \|Hw_\star + \Theta_\star h_w\|_2^2, \qquad \alpha := \frac{N_{\mathrm{F}}}{\sigma^2_{\mathrm{F}}}, \quad \gamma := \frac{N_{\mathrm{C}}}{\sigma^2_{\mathrm{C}}}. \tag{40}$$

Let $\mathcal{S}_\star := \mathrm{col}(\Theta_\star Q) = \Theta_\star \mathbf{1}^\perp$. By Assumption 3, $\Theta_\star Q$ has full column rank, so $\dim(\mathcal{S}_\star) = K - 1$. Profiling Equation (40) over h_w gives

$$\mathcal{I}_{\mathrm{eff}}(H) = \alpha \|H\|_F^2 + \gamma \left\| P_{\mathcal{S}_\star^\perp} H w_\star \right\|_2^2. \tag{41}$$

Indeed, the second term is the squared distance from $-Hw_\star$ to the nuisance image $\mathcal{S}_\star$.

Let $\bar{w}_\star = w_\star / \|w_\star\|_2$ and write

$$H = h\bar{w}_\star^\top + H_0, \qquad H_0 \bar{w}_\star = 0. \tag{42}$$

The H_0 subspace has dimension $d(K-1)$ and receives only fine information α. Decompose $h = h_\parallel + h_\perp$ with $h_\parallel \in \mathcal{S}_\star$ and $h_\perp \in \mathcal{S}_\star^\perp$. The $K - 1$ directions in $\mathcal{S}_\star$ can be canceled by a nuisance

perturbation h_w, so they also receive only α. The remaining $d-K+1$ directions receive $\alpha+\gamma\|w_\star\|_2^2$. Thus the fine-only multiplicity is $d(K-1)+(K-1)=(K-1)(d+1)=A_{\mathrm{U}}$ and the fine-plus-coarse multiplicity is $d-K+1=D_{\mathrm{U}}$.

For later use, write $B_\Theta := \Theta Q$. In the joint coordinate (β, v) the expected information blocks are

$$I_{\beta\beta} = \alpha I_{dK} + \gamma(ww^\top \otimes I_d), \tag{43}$$

$$I_{\beta v} = \gamma(w \otimes B_\Theta), \tag{44}$$

$$I_{vv} = \gamma B_\Theta^\top B_\Theta. \tag{45}$$

Their β Schur complement is

$$\Sigma_\beta(\Theta, w) = \alpha I_{dK} + \gamma\big(ww^\top \otimes (I_d - P_{\Theta Q})\big), \tag{46}$$

where $P_{\Theta Q} := \Theta Q(Q^\top \Theta^\top \Theta Q)^{-1} Q^\top \Theta^\top$. At the truth, the inverse trace of Equation (46) is exactly Equation (30).

C.5 Proof of Corollary 5

For B_1, choose the static fraction $\eta_{\mathrm{K}}^\star$ from Theorem 2 and use known-$w_\star$ weighted Gaussian least squares. By Section C.1, at budget b its expected risk is

$$L_b^{\mathsf{B}_1} \le \frac{\Phi_{\mathrm{K}}^\star}{b} + O(b^{-2}), \tag{47}$$

including bounded count rounding. Summing gives $\mathfrak{R}_T^{\mathsf{B}_1} \le \Phi_{\mathrm{K}}^\star \log T + O(1)$.

For B_2, choose the exogenous unknown-aggregation fraction $\eta_{\mathrm{U}}^\star$. If $\eta_{\mathrm{U}}^\star = 0$, projected fine OLS again has an $O(b^{-2})$ remainder. If $\eta_{\mathrm{U}}^\star > 0$, both fine and coarse counts are linear in b and the static two-fold specialization of Lemma 8 gives

$$L_b^{\mathsf{B}_2} \le \frac{\Phi_{\mathrm{U}}^\star}{b} + O(b^{-3/2}). \tag{48}$$

The remainder is summable, yielding the stated B_2 cumulative bound. The local minimax interpretations follow from the known-weight inverse-information argument and the unknown-weight local lower bound in Section C.3.

C.6 Proof of Theorem 6

Let $\Xi(b)$ be design-stream spending and $b_{\mathrm{E}} = b - \Xi(b)$. By Lemma 7,

$$\Xi(b) = O\left(\frac{b}{(\log b)^2}\right), \tag{49}$$

so $b_{\mathrm{E}} \ge b/2$ for all sufficiently large b and

$$0 \le \frac{1}{b_{\mathrm{E}}} - \frac{1}{b} = \frac{\Xi(b)}{b b_{\mathrm{E}}} = O\left(\frac{1}{b(\log b)^2}\right). \tag{50}$$

The conditional-efficiency lemma is stated for all sufficiently large estimation endpoints. Let $b_\star$ be a fixed threshold after which it applies. The finitely many levels $B_0 \le b < b_\star$ contribute only $O(1)$ to cumulative risk because projection onto the compact parameter set uniformly bounds L_b.

At a sufficiently large estimation endpoint, taking total expectation in Equation (19) gives

$$L_b \le \frac{c_{\mathrm{F}} \sigma_{\mathrm{F}}^2 d}{b_{\mathrm{E}}} \mathbb{E}\psi_{\mathrm{U}}(\eta_b^{\mathrm{E}}, \lambda) + O(b_{\mathrm{E}}^{-3/2}). \tag{51}$$

Subtract $\Phi_{\mathrm{U}}^\star/b = (c_{\mathrm{F}} \sigma_{\mathrm{F}}^2 d/b)\psi_{\mathrm{U}}(\eta_{\mathrm{U}}^\star, \lambda)$, apply Lemma 7, and use boundedness of ψ_{U} on the compact range of realized shares. This gives Equation (20) at estimation endpoints.

It remains to cover the budget levels at which the predictor is held fixed. Within an estimation block, consecutive estimation endpoints are separated by at most the fixed maximum action cost, so repeating an endpoint bound across these bounded gaps preserves summability. For a design

block in epoch m, let s_m be the last estimation endpoint before the block and let D_m be its budget length. Bounded endpoint residue gives $s_m \asymp B_m$, while $D_m = O(B_m/m^2)$. The predictor is fixed throughout the block, so its leading contribution is $D_m \Phi^\star_{\mathrm{U}}/s_m$. Relative to the harmonic benchmark on the same budget levels,

$$D_m \frac{\Phi^\star_{\mathrm{U}}}{s_m} - \Phi^\star_{\mathrm{U}} \sum_{j=1}^{D_m} \frac{1}{s_m + j} = O\left(\frac{D_m^2}{s_m^2}\right) = O(m^{-4}). \tag{52}$$

Repeating the three remainder terms in Equation (20) over the same block contributes respectively $O(m^{-4})$, $O(m2^{-m})$, and $O(2^{-m/2}/m^2)$. These sequences are summable over epochs. Thus the non-endpoint budget levels contribute only $O(1)$ additional cumulative excess, and

$$\mathfrak{R}_T^{\mathrm{alg}} \le \Phi^\star_{\mathrm{U}} \sum_{b=B_0}^{T} \frac{1}{b} + O(1) = \Phi^\star_{\mathrm{U}} \log T + O(1). \tag{53}$$

All constants in Lemmas 7 and 8 depend only on the fixed regularity bounds. Therefore the same per-budget inequality holds uniformly on any sufficiently small $\mathcal{U}_\delta$, with $\Phi^\star_{\mathrm{U}}$ evaluated at the local parameter. The map $\vartheta \mapsto \Phi^\star_{\mathrm{U}}(\vartheta)$ is continuous on the regular class: $w \mapsto \|w\|_2^2$ is continuous, the minimizer remains in a compact interval away from one, and the minimum of the continuous objective is continuous. Consequently

$$\sup_{\vartheta \in \mathcal{U}_\delta} \mathfrak{R}_T^{\mathrm{alg}}(\vartheta) \le \big(\Phi^\star_{\mathrm{U}}(\vartheta_\star) + o_\delta(1)\big) \log T + O_\delta(1), \tag{54}$$

proving the locally uniform statement in Theorem 6.

The lower bound follows from Lemma 9: sum Equation (21) over b, use that the Bayes risk is at most the supremum risk on $\mathcal{U}_\delta$, divide by $\log T$, and let $T \to \infty$ and then $\delta \downarrow 0$. This proves Equation (17).

Finally, the pointwise objectives satisfy

$$\psi_{\mathrm{U}}(\eta, \lambda) - \psi_{\mathrm{K}}(\eta, \lambda) = \frac{K-1}{d} \left[\frac{1}{1-\eta} - \frac{1}{1+(\lambda-1)\eta}\right] \ge 0, \tag{55}$$

because $1 + (\lambda - 1)\eta - (1 - \eta) = \lambda\eta \ge 0$. Hence $\Phi^\star_{\mathrm{U}} \ge \Phi^\star_{\mathrm{K}}$, and the stated known-versus-unknown decomposition is the identity $\Phi^\star_{\mathrm{U}} = \Phi^\star_{\mathrm{K}} + (\Phi^\star_{\mathrm{U}} - \Phi^\star_{\mathrm{K}})$ inserted into the upper bound.

C.7 Proof of Lemma 7

We use a guarded projected regression primitive repeatedly. Given n whitened observations from a linear regression $Y = B_\star^\top x + \varepsilon$, let $S_n = \sum_{i=1}^n x_i x_i^\top$. If $\lambda_{\min}(S_n) \ge n/2$, use ordinary least squares and project the coefficient onto a relevant fixed closed convex compact set containing the truth; otherwise return any fixed point in that set. Bounded isotropic covariates imply

$$\Pr(\|S_n/n - I_d\|_{\mathrm{op}} > 1/2) \le Ce^{-cn}. \tag{56}$$

Conditionally on the covariates and on the good event, the Gaussian regression error has second and fourth moments of orders n^{-1} and n^{-2}. Projection is nonexpansive and bounds the error on the bad event. Thus, for $p = 1, 2$,

$$\mathbb{E}\|\widetilde{B} - B_\star\|_F^{2p} \le C_p (n \vee 1)^{-p}. \tag{57}$$

This proves the moment statement used below and also makes every pilot well defined on every sample path.

Design-stream size and aggregation pilot. At the start of epoch m, the installed pilot is based on the fixed initialization and design blocks $j < m$; the design block collected in epoch m first enters the pilot used in epoch $m+1$. Let $B_m = 2^m B_0$ and $\Delta B_m = B_{m+1} - B_m$. Since $\zeta_m = \zeta_0/(m+2)^2$,

$$\Xi(B_m) \asymp \sum_{j<m} \frac{2^j}{(j+2)^2} = \Theta\left(\frac{2^m}{m^2}\right). \tag{58}$$

The upper bound follows by splitting the sum at $m/2$; the lower bound follows from the last fixed number of terms. The fixed design split is bounded away from zero and one, so both fine and coarse design counts are $\Theta(2^m/m^2)$. Applying Equation (57) gives, for $p = 1, 2$,

$$\mathbb{E}\|\widetilde{\Theta}_m - \Theta_\star\|_F^{2p} \le C\left(\frac{m^2}{2^m}\right)^p, \tag{59}$$

$$\mathbb{E}\|\widetilde{u}_m - u_\star\|_2^{2p} \le C\left(\frac{m^2}{2^m}\right)^p. \tag{60}$$

The constrained pilot satisfies

$$\widetilde{w}_m \in \operatorname*{arg\,min}_{w\in\Delta_K,\, w_k\ge\tau} \|\widetilde{u}_m - \widetilde{\Theta}_m w\|_2^2. \tag{61}$$

The truth is feasible because $w_{\star,k} \ge 2\tau$. On the event $\|(\widetilde{\Theta}_m - \Theta_\star)Q\|_{\mathrm{op}} \le \kappa_\Theta/2$, Weyl's inequality gives $\sigma_{\min}(\widetilde{\Theta}_m Q) \ge \kappa_\Theta/2$. Because $\widetilde{w}_m - w_\star \in \mathbf{1}^\perp$, write $\widetilde{w}_m - w_\star = Qh_m$. Optimality and feasibility imply

$$\frac{\kappa_\Theta}{2}\|\widetilde{w}_m - w_\star\|_2 \le \|\widetilde{\Theta}_m(\widetilde{w}_m - w_\star)\|_2 \tag{62}$$

$$\le 2\|\widetilde{u}_m - u_\star\|_2 + 2\|\widetilde{\Theta}_m - \Theta_\star\|_{\mathrm{op}}\|w_\star\|_2. \tag{63}$$

The complementary event has exponentially small probability by Equation (56) together with the conditional Gaussian tail for the guarded regression, while compact simplex constraints bound the error there. Hence

$$\mathbb{E}\|\widetilde{w}_m - w_\star\|_2^{2p} \le C\left(\frac{m^2}{2^m}\right)^p. \tag{64}$$

Let $\rho = c_{\mathrm{F}}\sigma_{\mathrm{F}}^2/(c_{\mathrm{C}}\sigma_{\mathrm{C}}^2)$. Since simplex vectors have norm at most one,

$$|\widetilde{\lambda}_m - \lambda| = \rho\left|\|\widetilde{w}_m\|_2^2 - \|w_\star\|_2^2\right| \le 2\rho\|\widetilde{w}_m - w_\star\|_2, \tag{65}$$

so

$$\mathbb{E}|\widetilde{\lambda}_m - \lambda|^2 \le C\frac{m^2}{2^m}. \tag{66}$$

Because both $w_\star$ and $\widetilde{w}_m$ lie in the simplex, $\lambda, \widetilde{\lambda}_m \in [\rho/K, \rho]$.

Quadratic stability of the plug-in optimum. Every minimizer of ψ_{U} lies in a fixed compact interval away from one: comparison with $\eta = 0$ gives $(A_{\mathrm{U}}/d)/(1 - \eta_{\mathrm{U}}^\star) \le K$, hence $\eta_{\mathrm{U}}^\star \le D_{\mathrm{U}}/(dK) < 1$. On a slightly larger compact interval and for $\lambda \in [\rho/K, \rho]$, there are constants $\mu > 0$ and $C_{\eta\lambda} < \infty$ such that

$$\partial_{\eta\eta}^2\psi_{\mathrm{U}} \ge \mu, \qquad |\partial_{\eta\lambda}^2\psi_{\mathrm{U}}| \le C_{\eta\lambda}. \tag{67}$$

Let $\eta_i = \eta_{\mathrm{U}}^\star(\lambda_i)$. The variational inequalities for constrained minimizers are

$$\partial_\eta\psi_{\mathrm{U}}(\eta_1, \lambda_1)(\eta_2 - \eta_1) \ge 0, \qquad \partial_\eta\psi_{\mathrm{U}}(\eta_2, \lambda_2)(\eta_1 - \eta_2) \ge 0. \tag{68}$$

Adding them and using strong monotonicity in η gives

$$|\eta_2 - \eta_1| \le \frac{C_{\eta\lambda}}{\mu}|\lambda_2 - \lambda_1|. \tag{69}$$

This argument also covers a minimizer at the boundary zero. Next, optimality of η_2 under λ_2 implies $\psi_{\mathrm{U}}(\eta_2, \lambda_2) \le \psi_{\mathrm{U}}(\eta_1, \lambda_2)$. Therefore

$$\psi_{\mathrm{U}}(\eta_2, \lambda_1) - \psi_{\mathrm{U}}(\eta_1, \lambda_1) \tag{70}$$

$$\le \big[\psi_{\mathrm{U}}(\eta_2, \lambda_1) - \psi_{\mathrm{U}}(\eta_2, \lambda_2)\big] - \big[\psi_{\mathrm{U}}(\eta_1, \lambda_1) - \psi_{\mathrm{U}}(\eta_1, \lambda_2)\big] \tag{71}$$

$$\le C_{\eta\lambda}|\eta_2 - \eta_1|\,|\lambda_2 - \lambda_1| \le \frac{C_{\eta\lambda}^2}{\mu}|\lambda_2 - \lambda_1|^2. \tag{72}$$

Applying this with $\lambda_1 = \lambda$ and $\lambda_2 = \widetilde{\lambda}_m$, then using Equation (66), yields

$$\mathbb{E}\big[\psi_{\mathrm{U}}(\widehat{\eta}_m, \lambda) - \psi_{\mathrm{U}}(\eta_{\mathrm{U}}^\star, \lambda)\big] \le C\frac{m^2}{2^m}. \tag{73}$$

No margin from the break-even boundary is used.

Accumulated tracking. Ignoring bounded endpoint residues, the cumulative estimation-stream coarse fraction at a budget $b \in [B_m, B_{m+1})$ is a cost-weighted convex combination of the installed epoch targets,

$$\bar{\eta}_b^{\mathrm{E}} = \sum_{j=0}^{m} \omega_{j,b} \widehat{\eta}_j, \qquad \omega_{j,b} \geq 0, \quad \sum_j \omega_{j,b} = 1, \tag{74}$$

with $\omega_{j,b} \leq C 2^{j-m}$ for $j < m$. Convexity and Equation (73) give

$$\mathbb{E}\big[\psi_{\mathrm{U}}(\bar{\eta}_b^{\mathrm{E}}, \lambda) - \psi_{\mathrm{U}}(\eta_{\mathrm{U}}^{\star}, \lambda)\big] \leq C 2^{-m} + C \sum_{j=1}^{m} 2^{j-m} \frac{j^2}{2^j} \tag{75}$$

$$\leq C \frac{m^3}{2^m}. \tag{76}$$

Writing s as coarse expenditure minus its target expenditure, a fine action changes s by $-\widehat{\eta}_m c_{\mathrm{F}}$ and a coarse action by $(1-\widehat{\eta}_m)c_{\mathrm{C}}$. The smaller-absolute-value rule therefore keeps $|s| \leq \max\{c_{\mathrm{F}}, c_{\mathrm{C}}\}$ by induction, up to the bounded terminal affordability residue. The signed-imbalance tracker changes cumulative coarse expenditure by only $O(1)$ within each epoch; bounded unused residues contribute $O(m)$ through epoch m. Since $\partial_\eta \psi_{\mathrm{U}}$ is bounded on the compact interval containing all realized shares, these effects add only $O(m/2^m)$. Using $2^m \asymp b$ and $m \asymp \log b$ gives the second statement of Equation (18); Equation (58) gives the first.

C.8 Proof of Lemma 8

Fix an epoch and condition on its design history $\mathcal{G}_m$. The target share and the entire estimation action schedule through the epoch are then deterministic, as are the parity-fold assignments within each resolution. All estimation covariates and noises remain independent of $\mathcal{G}_m$. Let $N_{a,r}$ be the number of estimation observations of resolution $a \in \{\mathrm{F}, \mathrm{C}\}$ in fold $r \in \{1, 2\}$ at budget b. The deterministic parity rule gives $|N_{a,1} - N_{a,2}| \leq 1$.

Opposite-fold pilots. For score fold r, fit guarded projected fine and coarse regressions on the opposite fold $-r$ using the primitive from Equation (57), obtaining $\widetilde{\Theta}^{(-r)}$ and $\widetilde{u}^{(-r)}$. Set

$$\widetilde{w}^{(-r)} \in \operatorname*{arg\,min}_{w \in \Delta_K,\ w_k \geq \tau} \|\widetilde{u}^{(-r)} - \widetilde{\Theta}^{(-r)} w\|_2^2, \qquad \widetilde{v}^{(-r)} = Q^\top (\widetilde{w}^{(-r)} - w_0). \tag{77}$$

The finite initialization guarantees a fixed positive number of coarse observations in each fold; hence the nuisance block I_{vv} is well defined even when the post-initialization coarse count remains bounded. If a coarse Gram matrix is ill conditioned, the guarded regression returns a bounded projected fallback. Because the fine-budget share is uniformly bounded away from zero, $N_{\mathrm{F},-r} = \Theta(b_{\mathrm{E}})$. The same tangent-stability argument as in Section C.7 gives, for $q = 1, 2$,

$$\mathbb{E}\|\widetilde{\Theta}^{(-r)} - \Theta_\star\|_F^{2q} \leq C b_{\mathrm{E}}^{-q}, \tag{78}$$

$$\mathbb{E}\|\widetilde{v}^{(-r)} - v_\star\|_2^{2q} \leq C \left(b_{\mathrm{E}}^{-q} + (N_{\mathrm{C},-r} \vee 1)^{-q}\right). \tag{79}$$

The second bound is allowed to remain $O(1)$ when the coarse count is bounded; this is sufficient because the nuisance contribution to the target row of the inverse information is then downweighted by the coarse information itself.

One-step map and inverse-information bounds. Let $S_r(\vartheta)$ be the Gaussian score on fold r and let $I_r(\vartheta)$ be its expected Fisher information conditional on the deterministic fold counts. Define

$$\check{\vartheta}^{(r)} = \widetilde{\vartheta}^{(-r)} + I_r(\widetilde{\vartheta}^{(-r)})^{-1} S_r(\widetilde{\vartheta}^{(-r)}). \tag{80}$$

If the fine-pilot tangent guard $\sigma_{\min}(\widetilde{\Theta}^{(-r)} Q) \geq \kappa_\Theta / 2$ fails, use the bounded projected fine estimator for that fold. The guard failure probability is $O(e^{-c b_{\mathrm{E}}})$ by Equation (56) and the conditional Gaussian tail of the fine-regression error. The final estimator averages the two β components and projects the reshaped matrix onto the compact operator-norm ball.

For one score fold let $\alpha_r = N_{\mathrm{F},r}/\sigma_{\mathrm{F}}^2$ and $\gamma_r = N_{\mathrm{C},r}/\sigma_{\mathrm{C}}^2$. The exact information blocks are Equations (43) to (45). Standard block inversion gives

$$[I_r^{-1}]_{\beta v} = -\frac{1}{\alpha_r}(w \otimes \Theta Q)\{(\Theta Q)^\top(\Theta Q)\}^{-1}, \tag{81}$$

$$[I_r^{-1}]_{\beta\beta} = I_{\beta\beta}^{-1} + \frac{\gamma_r}{\alpha_r(\alpha_r + \gamma_r\|w\|_2^2)}\big(ww^\top \otimes P_{\Theta Q}\big). \tag{82}$$

On the regular neighborhood, $\alpha_r \asymp b_{\mathrm{E}}$ and $0 < \gamma_r \le C b_{\mathrm{E}}$. Hence the two blocks in the β row are $O(b_{\mathrm{E}}^{-1})$. Moreover, for two parameters ϑ_1, ϑ_2 in that neighborhood,

$$\|[I_r(\vartheta_1)^{-1}]_{\beta\beta} - [I_r(\vartheta_2)^{-1}]_{\beta\beta}\|_{\mathrm{op}} \le C\frac{\gamma_r}{b_{\mathrm{E}}^2}\|\vartheta_1 - \vartheta_2\|, \tag{83}$$

$$\|[I_r(\vartheta_1)^{-1}]_{\beta v} - [I_r(\vartheta_2)^{-1}]_{\beta v}\|_{\mathrm{op}} \le \frac{C}{b_{\mathrm{E}}}\|\vartheta_1 - \vartheta_2\|. \tag{84}$$

These follow directly from Equations (81) and (82) because $\Theta \mapsto P_{\Theta Q}$ and $\Theta \mapsto \{(\Theta Q)^\top(\Theta Q)\}^{-1}$ are Lipschitz when the tangent singular value is bounded away from zero. The factor $\gamma_r/b_{\mathrm{E}}^2$ is important when coarse counts are sublinear.

Exact score decomposition. Write $\Delta_\Theta = \widetilde{\Theta}^{(-r)} - \Theta_\star$, $\Delta_v = \widetilde{v}^{(-r)} - v_\star$, and $B_\star = \Theta_\star Q$. For a coarse score observation, the residual at the pilot is exactly

$$Y_i^{\mathrm{C}} - x_i^\top \widetilde{\Theta}^{(-r)} w(\widetilde{v}^{(-r)}) = \varepsilon_i^{\mathrm{C}} - p_i - q_i, \tag{85}$$

where

$$p_i = x_i^\top \Delta_\Theta w_\star + x_i^\top B_\star \Delta_v, \qquad q_i = x_i^\top \Delta_\Theta Q \Delta_v. \tag{86}$$

The pilot score directions are $(w_\star + Q\Delta_v) \otimes x_i$ for β and $(B_\star + \Delta_\Theta Q)^\top x_i$ for v. Expanding Equation (80) around $\vartheta_\star$ gives

$$\check{\beta}^{(r)} - \beta_\star = [I_r(\vartheta_\star)^{-1} S_r(\vartheta_\star)]_\beta + R_{r,1} + R_{r,2} + R_{r,3} + R_{r,4} + R_{r,5}, \tag{87}$$

where the five remainder classes are:

1. $R_{r,1}$: the quadratic mean term q_i multiplied by the truth score directions;
2. $R_{r,2}$: the perturbed score directions $Q\Delta_v \otimes x_i$ and $(\Delta_\Theta Q)^\top x_i$, multiplied by $\varepsilon_i^{\mathrm{C}} - p_i - q_i$;
3. $R_{r,3}$: the empirical-score Jacobian minus its expected Fisher information, multiplied by $(\Delta_\Theta, \Delta_v)$;
4. $R_{r,4}$: replacing the truth inverse information by the inverse information evaluated at the pilot;
5. $R_{r,5}$: the exponentially rare fine-pilot guard failure, handled by the bounded projected fallback.

This list is exhaustive because the fine mean is linear and the coarse mean is exactly bilinear in (Θ, v). To make the bookkeeping explicit, abbreviate $\widetilde{\vartheta} = \widetilde{\vartheta}^{(-r)}$, let $M_r(\vartheta) = I_r(\vartheta)^{-1}$, and write $M_{r,\beta\bullet}$ for its β block row. For a coarse observation define the truth joint score direction

$$g_i^\star = \begin{pmatrix} w_\star \otimes x_i \\ B_\star^\top x_i \end{pmatrix}, \qquad \delta g_i = \begin{pmatrix} Q\Delta_v \otimes x_i \\ (\Delta_\Theta Q)^\top x_i \end{pmatrix}.$$

Let $\widehat{J}_r(\widetilde{\vartheta})$ denote the negative empirical score Jacobian on fold r. After the linear score term cancels against its expected Jacobian, the first four classes can be represented as

$$R_{r,1} = -\big[M_r(\widetilde{\vartheta})\big]_{\beta\bullet}\frac{1}{\sigma_{\mathrm{C}}^2}\sum_{i\in\mathcal{C}_r} g_i^\star q_i, \tag{88}$$

$$R_{r,2} = \big[M_r(\widetilde{\vartheta})\big]_{\beta\bullet}\frac{1}{\sigma_{\mathrm{C}}^2}\sum_{i\in\mathcal{C}_r} \delta g_i(\varepsilon_i^{\mathrm{C}} - p_i - q_i), \tag{89}$$

$$R_{r,3} = \big[M_r(\widetilde{\vartheta})\big]_{\beta\bullet}\{I_r(\widetilde{\vartheta}) - \widehat{J}_r(\widetilde{\vartheta})\}(\widetilde{\vartheta} - \vartheta_\star), \tag{90}$$

$$R_{r,4} = \big[M_r(\widetilde{\vartheta}) - M_r(\vartheta_\star)\big]_{\beta\bullet} S_r(\vartheta_\star), \tag{91}$$

with terms that differ only by truth-versus-pilot score directions grouped into the corresponding class. $R_{r,5}$ is supported on the guard-failure event.

We now bound each class conditionally on the opposite-fold pilot. Bounded covariates and Gaussian fourth moments imply the usual L_2 bounds for centered sums. Using Equations (78) and (79), $\gamma_r \le C b_{\mathrm{E}}$, and the target-row inverse bounds above gives

$$\|R_{r,1}\|_{L_2} \le C\frac{\gamma_r}{b_{\mathrm{E}}}\|\Delta_\Theta\|_{L_4}\|\Delta_v\|_{L_4} \le \frac{C}{b_{\mathrm{E}}}, \tag{92}$$

$$\|R_{r,2}\|_{L_2} \le C\left[\frac{\gamma_r}{b_{\mathrm{E}}}(\|\Delta_v\|_{L_4}^2 + \|\Delta_\Theta\|_{L_4}\|\Delta_v\|_{L_4}) + \frac{\sqrt{\gamma_r}}{b_{\mathrm{E}}}(\|\Delta_v\|_{L_4} + \|\Delta_\Theta\|_{L_4})\right] \le \frac{C}{b_{\mathrm{E}}}, \tag{93}$$

$$\|R_{r,3}\|_{L_2} \le \frac{C}{b_{\mathrm{E}}}\left(\sqrt{N_{\mathrm{F},r}}\|\Delta_\Theta\|_{L_4} + \sqrt{N_{\mathrm{C},r}}\|\Delta_v\|_{L_4}\right) \le \frac{C}{b_{\mathrm{E}}}, \tag{94}$$

$$\|R_{r,4}\|_{L_2} \le \frac{C}{b_{\mathrm{E}}}, \tag{95}$$

where the last inequality uses Equations (83) and (84) together with truth-score sizes $O_{L_4}(\sqrt{b_{\mathrm{E}}})$ and $O_{L_4}(\sqrt{\gamma_r})$. The guard contribution satisfies $\mathbb{E}\|R_{r,5}\|_2^2 \le Ce^{-cb_{\mathrm{E}}}$ because the fallback is bounded. Thus

$$\mathbb{E}\|R_{r,\beta}\|_2^2 \le \frac{C}{b_{\mathrm{E}}^2}. \tag{96}$$

The bounds remain valid when $N_{\mathrm{C},r}$ is bounded: then Δ_v need not converge, but every occurrence is multiplied by either γ_r/b_{E} or $\sqrt{\gamma_r}/b_{\mathrm{E}}$, which is enough for Equation (95).

Leading variance and risk. Conditional on $\mathcal{G}_m$, the score folds are independent and centered, and $\mathrm{Cov}(S_r(\vartheta_\star) \mid \mathcal{G}_m) = I_r(\vartheta_\star)$. Hence the two leading influence functions in Equation (87) have covariances $[I_r(\vartheta_\star)^{-1}]_{\beta\beta}$. Their average has mean-square

$$\frac{1}{4}\sum_{r=1}^{2}\mathrm{tr}[I_r(\vartheta_\star)^{-1}]_{\beta\beta} = \mathcal{R}_{\mathrm{U}}(N_{\mathrm{F}}^{\mathrm{E}}, N_{\mathrm{C}}^{\mathrm{E}}) + O(b_{\mathrm{E}}^{-2}), \tag{97}$$

because each fold has half of the full counts up to $O(1)$, so each fold inverse information is twice the full-count inverse up to $O(b_E^{-2})$; the factor $1/4$ from averaging the two independent fold estimates restores the full-sample efficient trace. The leading term has L_2 norm $O(b_{\mathrm{E}}^{-1/2})$. Cauchy–Schwarz with Equation (96) therefore makes the influence–remainder cross term $O(b_{\mathrm{E}}^{-3/2})$, and the squared remainder is $O(b_{\mathrm{E}}^{-2})$. Final projection cannot increase squared prediction loss. We obtain

$$L_b(\mathcal{G}_m) \le \mathcal{R}_{\mathrm{U}}(N_{\mathrm{F}}^{\mathrm{E}}, N_{\mathrm{C}}^{\mathrm{E}}) + O(b_{\mathrm{E}}^{-3/2}). \tag{98}$$

Writing the count-based efficient trace in terms of b_{E} and η_b^{E} gives Equation (19). All constants used above are uniform over the compact regular class and over the predictable schedules generated by EET, proving the lemma.

C.9 Proof of Lemma 9

Let $\vartheta = (\beta, v)$ and choose a smooth prior Π_δ with compact support strictly inside the regular coordinate ball $\mathcal{U}_\delta$. Its density is C^1, vanishes at the boundary, and has finite Fisher-information matrix $J(\Pi_\delta)$. The Gaussian observation model is differentiable in quadratic mean on this neighborhood, so the multivariate van Trees inequality applies (Gill & Levit, 1995). It is essential that the prior vary both β and v; fixing v would remove the nuisance price.

Fix an arbitrary predictable resolution policy and let τ_b be the last completed query whose known cumulative cost does not exceed b. If H_{t-1} is the observed history before query t, the stopped-history density factorizes as

$$p_\vartheta(H_{\tau_b}) = \prod_{t=1}^{\tau_b} q_t(A_t \mid H_{t-1})\, p(x_t)\, p_\vartheta(Y_t^{A_t} \mid x_t, A_t), \tag{99}$$

where each predictable policy kernel q_t is parameter free. The stopping rule is also parameter free conditional on the action sequence because costs are known. Therefore the history score is the sum of queried-observation scores. These are martingale differences conditional on the past, so their cross terms have zero expectation and the Fisher information is the expected sum of the queried-label information matrices. Let its prior average be $\overline{I}_b(\Pi_\delta)$.

The multivariate van Trees inequality for the target map $\vartheta \mapsto \beta$ gives

$$\int \mathbb{E}_\vartheta[(\widehat{\beta}_b - \beta)(\widehat{\beta}_b - \beta)^\top]\,\Pi_\delta(d\vartheta) \succeq [\overline{I}_b(\Pi_\delta) + J(\Pi_\delta)]^{-1}_{\beta\beta}. \tag{100}$$

Taking traces converts the left side to Bayes fine-prediction risk.

Let

$$\bar{N}_{\mathrm{F}}(b) := \int \mathbb{E}_\vartheta N_{\mathrm{F}}(b)\,\Pi_\delta(d\vartheta), \qquad \bar{N}_{\mathrm{C}}(b) := \int \mathbb{E}_\vartheta N_{\mathrm{C}}(b)\,\Pi_\delta(d\vartheta). \tag{101}$$

If a policy leaves usable budget idle, adding hypothetical observations can only increase information and decrease the inverse-information trace, so it is enough for the lower bound to complete the design up to a bounded residue $r_b \le \max\{c_{\mathrm{F}}, c_{\mathrm{C}}\}$. Replacing $b - r_b$ by b changes each $1/b$ leading term by $O(b^{-2})$, which is absorbed below. Let $\bar{\eta}_b = c_{\mathrm{C}}\bar{N}_{\mathrm{C}}(b)/b$. At the center $\vartheta_\star$, Equation (46) gives

$$\mathrm{tr}\left[\Sigma_{\beta,\star}(\bar{N}_{\mathrm{F}}(b), \bar{N}_{\mathrm{C}}(b))^{-1}\right] = \frac{c_{\mathrm{F}}\sigma_{\mathrm{F}}^2 d}{b}\psi_{\mathrm{U}}(\bar{\eta}_b, \lambda) + O(b^{-2}) \ge \frac{\Phi^\star_{\mathrm{U}}}{b} - O(b^{-2}). \tag{102}$$

It remains to pass from center information to the prior-averaged information uniformly over policies.

Uniform center-information comparison. For $\vartheta \in \mathcal{U}_\delta$, write $n_a(\vartheta) = \mathbb{E}_\vartheta N_a(b)$ and let $\mathcal{J}_{\mathrm{C}}(\vartheta)$ denote the one-coarse-observation joint information matrix in the local coordinate (β, v). The fine information is parameter independent, while

$$\overline{I}_b(\Pi_\delta) = \frac{\bar{N}_{\mathrm{F}}(b)}{\sigma_{\mathrm{F}}^2}\begin{pmatrix} I_{dK} & 0 \\ 0 & 0 \end{pmatrix} + \int n_{\mathrm{C}}(\vartheta)\,\mathcal{J}_{\mathrm{C}}(\vartheta)\,\Pi_\delta(d\vartheta). \tag{103}$$

If $\bar{N}_{\mathrm{C}}(b) = 0$, set $\bar{\Sigma}_\beta = (\bar{N}_{\mathrm{F}}/\sigma_{\mathrm{F}}^2)I_{dK}$; then the comparison below holds with $\gamma = 0$. Otherwise normalize the second term by $\bar{N}_{\mathrm{C}}(b)$. Because every parameter in the prior support lies within $O(\delta)$ of $\vartheta_\star$, the maps $w \mapsto ww^\top$, $\Theta \mapsto \Theta Q$, and $\Theta \mapsto P_{\Theta Q}$ are uniformly Lipschitz, and the normalized coarse-information matrix differs from $\mathcal{J}_{\mathrm{C}}(\vartheta_\star)$ by at most $C\delta$ in operator norm. For δ small enough, its nuisance block remains uniformly positive definite by Assumption 3. The Schur-complement map is Lipschitz on this compact set, hence the efficient β information $\bar{\Sigma}_\beta$ of the prior-averaged likelihood satisfies

$$\left\|\bar{\Sigma}_\beta - \Sigma_{\beta,\star}(\bar{N}_{\mathrm{F}}(b), \bar{N}_{\mathrm{C}}(b))\right\|_{\mathrm{op}} \le C\delta\gamma, \qquad \gamma := \frac{\bar{N}_{\mathrm{C}}(b)}{\sigma_{\mathrm{C}}^2}. \tag{104}$$

If $\bar{\eta}_b < 1/K$, then

$$\alpha := \frac{\bar{N}_{\mathrm{F}}(b)}{\sigma_{\mathrm{F}}^2} \ge \frac{(1 - 1/K)b}{c_{\mathrm{F}}\sigma_{\mathrm{F}}^2} + O(1) \asymp b. \tag{105}$$

Every β Schur complement is bounded below by αI. Equation 104 changes the efficient β information by at most $O(\delta b)$, while the fixed prior information $J(\Pi_\delta)$ contributes only $O_\delta(1)$. The resolvent identity therefore gives

$$\mathrm{tr}[\overline{I}_b(\Pi_\delta) + J(\Pi_\delta)]^{-1}_{\beta\beta} \ge \mathrm{tr}\left[\Sigma_{\beta,\star}(\bar{N}_{\mathrm{F}}, \bar{N}_{\mathrm{C}})^{-1}\right] - \frac{C\delta}{b} - \frac{C_\delta}{b^2}. \tag{106}$$

Combining with Equation (102) gives the desired lower bound in this case.

If $\bar{\eta}_b \ge 1/K$, use the fixed center fine-only subspace $V_{w_\star} = \{H : Hw_\star = 0\}$ of dimension $d(K-1)$. For a positive-definite operator, the trace of the full inverse is at least the trace of the inverse of its compression to any fixed subspace. By Equation (104), coarse information contributes

at most $C\delta\gamma$ on this subspace. Adding $J(\Pi_\delta)$ can increase the compressed information by at most its fixed operator norm C_δ. Hence

$$\mathrm{tr}[\overline{I}_b(\Pi_\delta) + J(\Pi_\delta)]^{-1}_{\beta\beta} \geq \frac{d(K-1)}{\alpha + C\delta\gamma + C_\delta} \tag{107}$$

$$\geq \frac{dKc_{\mathrm{F}}\sigma_{\mathrm{F}}^2}{b} - \frac{C\delta}{b} - \frac{C_\delta}{b^2} \tag{108}$$

$$\geq \frac{\Phi_{\mathrm{U}}^\star - C\delta}{b} - \frac{C_\delta}{b^2}, \tag{109}$$

where the last step uses $\Phi_{\mathrm{U}}^\star \leq dKc_{\mathrm{F}}\sigma_{\mathrm{F}}^2$. Because $\bar{\eta}_b \geq 1/K$, for sufficiently small fixed δ we have $\alpha + C\delta\gamma \leq [(K-1)/(Kc_{\mathrm{F}}\sigma_{\mathrm{F}}^2) + C\delta]b + O(1)$; expanding its reciprocal gives the second line above. The constants in the final $O_\delta(b^{-2})$ term may depend on δ.

Thus both cases yield

$$\mathrm{tr}[\overline{I}_b(\Pi_\delta) + J(\Pi_\delta)]^{-1}_{\beta\beta} \geq \frac{\Phi_{\mathrm{U}}^\star - C\delta}{b} - \frac{C_\delta}{b^2}. \tag{110}$$

Together with Equation (100), this is Equation (21), with $o_\delta(1) = C\delta$. Summation gives the lower half of Theorem 6.