PERFECT MEMORY IS NOT ENOUGH FOR RELIABLE LIFELONG PREDICTION

Anonymous authors
Paper under double-blind review

ABSTRACT

Lifelong and continual learning have long been organized around a single adversary: catastrophic forgetting. Retrieval systems, exact key-value caches, and persistent memory push that defense to its limit by storing historical interactions verbatim. Yet, even when retention error is eliminated by construction, a fundamental question remains: does perfect memory guarantee reliable prediction? We show that it does not, and we isolate the irreducible statistical bottleneck: the missing mass of unobserved tasks. In heterogeneous, heavy-tailed request streams, standard split conformal prediction offers only marginal coverage, dangerously masking catastrophic coverage collapse on rare, user-specific tasks. We formalize *taskwise honest prediction*, requiring each recurring task to satisfy a distribution-free coverage guarantee at level $1-\alpha$. Through an indistinguishable-worlds construction, we prove sharp minimax lower bounds on prediction set efficiency (Lebesgue measure and cardinality), showing that taskwise validity necessitates set inflation on unobserved tasks. To reduce this cost, we introduce *active query allocation*: proactively spending a hard query budget B over a submodular polymatroid before test requests arrive. We develop a scalable stochastic top-B dual Hedge algorithm circumventing exponential subset constraints. To bridge theory and practice, we introduce the finite-sample *taskwise coverage guard*, certifying validity around any nominal set predictor. Instantiating this guard around pooled conformal prediction (*Guarded CP*) separates certification from prediction, which eliminates tail undercoverage while preserving small empirical set sizes on frequent tasks. Furthermore, we establish *lifetime coverage* across extended deployment horizons, proving that a single greedy query policy is universally optimal across all horizons.

1 INTRODUCTION

Continual and lifelong learning have long been defined by the struggle against catastrophic forgetting (McCloskey and Cohen, 1989; French, 1999) and the stability-plasticity dilemma (Parisi et al., 2019). Modern agent architectures fundamentally alter that struggle: retrieval-augmented generation (RAG) (Lewis et al., 2020; Borgeaud et al., 2022; Packer et al., 2023), vector stores, persistent scratchpads, and exact key-value caches store historical interactions verbatim, so past records no longer decay on disk. Yet, even when historical retention is made exact by construction, a critical question remains: *does perfect memory make a lifelong system reliable?*

In high-stakes lifelong deployments—such as clinical diagnosis, legal consultation, and autonomous coding agents—the primary failure mode is no longer merely losing past weights, but uncalibrated hallucination on recurring and newly arriving requests. In such domains, returning a fragile point prediction is unacceptable; safety demands *distribution-free uncertainty quantification* with rigorous statistical coverage guarantees. While standard conformal prediction (CP) (Vovk et al., 2005; Lei et al., 2018; Angelopoulos and Bates, 2023) provides finite-sample prediction sets, it guarantees only *marginal coverage*: $\mathbb{P}(Y \in C(X)) \geq 1-\alpha$ averaged across the entire test distribution. In lifelong interaction streams, incoming requests follow heavy-tailed, multi-task distributions (Zipf, 1949): a handful of popular head tasks recur frequently, while thousands of rare, user-specific, or safety-critical tasks populate the long tail. In this heterogeneous regime, marginal coverage is profoundly deceptive. A pooled predictor can achieve a flawless 99.5% stream-level coverage rate by guaranteeing coverage on high-frequency head tasks, while suffering total coverage collapse (0% coverage) on rare or newly requested tasks. For an individual user or a specialized domain, global average safety is meaningless.

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.

Reliable lifelong prediction instead demands *taskwise coverage*:

$$\forall j \in \mathcal{J}, \qquad \mathbb{P}(a_j \in C(\mathcal{T}, j)) \geq 1 - \alpha. \tag{1}$$

Achieving conditional validity, however, faces a celebrated theoretical barrier: Barber et al. (2021) proved that distribution-free exact conditional coverage is impossible without inflating prediction sets to trivial size. In a lifelong multi-task stream, if the learner is equipped with a "perfect memory" (retaining all historical encounters without noise or decay), what prevents taskwise validity? We show that the fundamental obstacle is not memory capacity or representational drift, but *missing mass* (Good, 1953; McAllester and Schapire, 2000): tasks whose responses have never appeared in the historical stream. Through an indistinguishable-worlds construction, we prove sharp distribution-free minimax lower bounds on prediction set efficiency (Lebesgue measure and cardinality) for any unobserved task, and we show that this unavoidable price is attained by certified coverage guards.

How can a lifelong learner overcome this fundamental barrier? We propose two complementary solutions:

- **Active Query Allocation**: Given a hard query budget B spent before test requests arrive, how should queries be allocated among unobserved tasks to minimize prediction set size? We show that feasible query allocations form a submodular polymatroid (Edmonds, 1970) governed by joint absence probabilities, and we design a scalable stochastic dual Hedge algorithm with a conservative calibration audit.
- **Lifetime Coverage Across Deployments**: Real agents serve requests across an extended deployment stream. Independent checkpoint predictions leak error across time, causing the lifetime failure probability to reach $\alpha(1 + pH)$ under moving-threshold checkpoint rules (Proposition 29). We formalize lifetime coverage and show that a single universal query policy is simultaneously optimal across all evaluation horizons.

Contributions.

1. **Fundamental Limits of Taskwise Honest Prediction**: We establish sharp, distribution-free minimax lower bounds on prediction set efficiency (Lebesgue measure and cardinality) for taskwise coverage, and we determine the exact reduction enabled by budgeted proactive queries over submodular polymatroids (Section 3).
2. **Scalable Stochastic Dual Optimization**: We develop a stochastic top-B Hedge algorithm that circumvents exponential subset constraints via dual prices, paired with a finite-sample conservative calibration audit that rigorously guarantees validity under empirical estimation (Section 4).
3. **Lifetime Coverage Across Deployments**: We solve the problem of time-uniform taskwise coverage across extended deployment streams, proving that a single universal greedy query policy is simultaneously optimal for every horizon and discount factor (Section 5).
4. **Taskwise Coverage Guard and Guarded Prediction**: We establish the *taskwise coverage guard*, a finite-sample wrapper certifying taskwise validity around any nominal set predictor. Instantiating it around pooled conformal prediction (*Guarded CP*) separates certification from prediction: the guard alone guarantees tail safety, while conformal prediction exploits learned representations on the non-guard branch. On natural digit tasks, Guarded CP eliminates detected violations and elevates coverage from 95.73% to 99.84% over the prediction-agnostic certified rule with mean set size increasing only from 2.34 to 2.38 (Section 6).

Section 7 examines how inductive biases and task relatedness break discrete indistinguishability, and Section 8 positions the work. All proofs and auxiliary lemmas are collected in the appendices; Appendix A provides a complete proof map and Table 2 defines all symbols.

2 Problem Formulation: Taskwise Honest Prediction in Interaction Streams

A canonical model isolating statistical limits. To isolate the statistical cost of taskwise validity from confounding optimization failures, we study a canonical lifelong setting: tasks recur in a

stationary stream, every encounter reveals the true response without noise, and past observations are cached with zero retention decay. Granting perfect retention by construction strips away catastrophic forgetting, plasticity-stability trade-offs, and optimization errors, proving that the remaining failure is fundamentally *inferential*—even when encounters are exchangeable across time, marginal coverage averages over task frequencies and fails to guarantee reliability on rare or unobserved tasks.

Tasks, observations, and queries. A task is an identifiable recurring prediction problem indexed by a finite set $\mathcal{J} = \{1, \ldots, K\}$ (countable extensions are addressed via finite-support approximations; Corollary 17). The request distribution π ($\pi_j > 0$, $\sum_j \pi_j = 1$) is known to the learner for planning and query allocation (with misspecification and unknown π analyzed in Proposition 25 and Theorem 10), and task j has a fixed response $a_j \in \mathcal{A}$. We write $a = (a_j)_{j \in \mathcal{J}}$ for the complete response vector. The initial history consists of n independent requests together with their responses:

$$J_1, \ldots, J_n \overset{\text{iid}}{\sim} \pi, \qquad H_n = ((J_t, a_{J_t}))_{t=1}^n, \qquad U = \mathcal{J} \setminus \{J_1, \ldots, J_n\}. \tag{2}$$

The set U collects the tasks that never appeared in the history; we call them *unobserved*. Every encounter reveals the exact response, and the learner retains everything verbatim.

Before an independent test request $J \sim \pi$ arrives, the learner may query a set $Q \subseteq U$ of unobserved tasks, each revealing its exact response. The query count is restricted by a hard budget: $|Q| \leq B$ with probability one. Queries are spent before the request arrives; Section 7 contrasts this proactive preparation with verifying an answer on demand after the request is known. The learner's complete record—history, queries, revealed responses, and random coins—is its transcript $\mathcal{T}$, and on request J it outputs a prediction set $C(\mathcal{T}, J) \subseteq \mathcal{A}$.

Taskwise validity versus marginal coverage. Throughout, $\mathbb{P}_a$ denotes probability under response vector a, averaging over history, queries, and algorithm randomization.

Definition 1 (Taskwise coverage)**.** Let $\mathcal{A}$ be the response domain and $0 < \alpha < 1$ the allowed miscoverage. A procedure is *taskwise honest* (or *honest*) at level α if

$$\forall a \in \mathcal{A}^{\mathcal{J}},\ \forall j \in \mathcal{J}, \qquad \mathbb{P}_a\{a_j \in C(\mathcal{T}, j)\} \geq 1 - \alpha. \tag{3}$$

In contrast to marginal coverage ($\sum_j \pi_j \mathbb{P}_a\{a_j \in C(\mathcal{T}, j)\} \geq 1 - \alpha$), taskwise validity demands that every individual task satisfies the $1 - \alpha$ standard, preventing common tasks from subsidizing errors on rare ones.

Prediction cost metrics. Coverage alone is trivially achieved by returning $\mathcal{A}$, so every guarantee must be priced by set efficiency. We consider three domains: the real line $\mathcal{A} = \mathbb{R}$ (unbounded regression), a bounded interval $\mathcal{A} = [-A, A]$ (bounded regression), and a finite label set $\mathcal{A} = [M]$ (classification). We score a procedure at an independent request $J \sim \pi$:

$$\begin{aligned} V_a &= \mathbb{P}_a\{\mathrm{Leb}(C(\mathcal{T}, J)) = \infty\} && (\mathcal{A} = \mathbb{R}), \\ W_a &= \mathbb{E}_a\, \mathrm{Leb}(C(\mathcal{T}, J)) && (\mathcal{A} = [-A, A]), \\ S_a &= \mathbb{E}_a |C(\mathcal{T}, J)| && (\mathcal{A} = [M]). \end{aligned} \tag{4}$$

Dividing W_a by $2A$ yields the *normalized regression cost* $\widetilde{V}_a$, equal to $W_a/(2A)$ on $[-A, A]$ and V_a on $\mathbb{R}$, obeying identical bounds throughout. The subscript a records that bounds hold at every configuration, constraining the worst-case response vector.

3 Fundamental limits of taskwise honest prediction and active allocation

The limits of taskwise reliability hinge on a single event: whether a task's response has ever appeared in the transcript. Section 3.1 establishes the passive baseline; Section 3.2 introduces active queries; Section 3.3 identifies subset bottlenecks.

3.1 PASSIVE COSTS: INDISTINGUISHABILITY AND MISSING MASS

Let $q_j = (1-\pi_j)^n$ be the probability that task j is absent from history H_n. On histories where j never appears, the learner sees identical information whether the true response is a_j or an alternative y. Exact retention preserves this information perfectly, but cannot distinguish between alternative responses. Consequently, for every candidate $y \in \mathcal{A}$ and task j, indistinguishability forces $\mathbb{P}_a\{y \in C(\mathcal{T}, j),\ j \in U\} \geq (q_j - \alpha)_+$ (Lemma 15, Appendix C.1). Averaging over requested tasks yields the exact passive cost:

Theorem 2 (Exact passive costs)**.** *Let $B = 0$, $0 < \alpha < 1$, and $n < \infty$. Define the* passive cost

$$\Psi_n(\pi, \alpha) = \sum_j \pi_j \big((1-\pi_j)^n - \alpha\big)_+. \tag{5}$$

Every honest procedure satisfies, at every response vector a,

$$V_a \geq \Psi_n, \qquad W_a \geq 2A\Psi_n, \qquad S_a \geq 1 - \alpha + (M-1)\Psi_n, \tag{6}$$

and for each criterion one honest procedure attains equality at every a simultaneously. If some $q_j > \alpha$ in $\mathcal{A} = \mathbb{R}$, then $\mathbb{E}_a \operatorname{Leb}(C(\mathcal{T}, j)) = \infty$ for every a.

Proved in Appendix C.2, p. 20.

The passive cost Ψ_n is not the standard missing mass $\sum_j \pi_j q_j$: each task first receives its own error allowance α, and only the positive residual is averaged. On unobserved requests, the minimax optimal procedure returns the full domain with probability $(1 - \alpha/q_j)_+$ and an empty set otherwise. On infinite task vocabularies, $\sup_j q_j = 1$ at every finite n, forcing $\Psi_n > 0$ and infinite expected length on $\mathbb{R}$ (Corollary 17, Appendix C.2).

3.2 ACTIVE ALLOCATION OVER SUBMODULAR POLYMATROIDS

A natural remedy for missing mass is to resolve unobserved tasks proactively before test requests arrive. Let $r_j(a) = \mathbb{P}_a(j \in Q)$ be the query probability and $u_j(a) = q_j - r_j(a)$ the unresolved probability of task j. A total constraint $\sum_j r_j \leq B$ is insufficient: a query can only be spent on a task that is actually unobserved in the realized history, and tasks may be absent together.

For any subset $S \subseteq \mathcal{J}$, a realized history offers at most $\min\{B, |U \cap S|\}$ query opportunities in S. Taking expectations defines the *capacity* $F_{n,B}(S) = \mathbb{E}\min\{B, |U \cap S|\}$ and the feasible region

$$\mathcal{P}_{n,B} = \{r \geq 0 : \sum_{j \in S} r_j \leq F_{n,B}(S) \text{ for every } S \subseteq \mathcal{J}\}. \tag{7}$$

The capacity $F_{n,B}$ is a submodular set function, and $\mathcal{P}_{n,B}$ is the polymatroid it defines (Edmonds, 1970). By Edmonds' theorem, every hard-budget query policy satisfies $r(a) \in \mathcal{P}_{n,B}$, and every point in $\mathcal{P}_{n,B}$ is achievable by a response-independent randomized query kernel via random permutations and sub-budget truncation (Lemma 18, Appendix C.4).

For threshold $b \geq 0$, the optimal prediction cost under active querying is

$$\Phi_{n,B}(\pi, b) = \min_{r \in \mathcal{P}_{n,B}} \sum_j \pi_j (q_j - r_j - b)_+. \tag{8}$$

Theorem 3 (Exact active costs)**.** *Let $0 < \alpha < 1$, $n < \infty$, and $B \geq 0$. Every honest procedure with a hard-budget query policy satisfies, for every response vector a, task j, and candidate $y \in \mathcal{A}$,*

$$\mathbb{P}_a\{y \in C(\mathcal{T}, j),\ j \in U \setminus Q\} \geq (u_j(a) - \alpha)_+, \tag{9}$$

and consequently the cost bounds

$$V_a \geq \Phi_{n,B}(\pi, \alpha), \quad W_a \geq 2A\Phi_{n,B}(\pi, \alpha), \quad S_a \geq 1 - \alpha + (M-1)\Phi_{n,B}(\pi, \alpha) \tag{10}$$

hold at every a. For each criterion, one response-independent policy–predictor pair attains equality at every a simultaneously. If classification sets must be nonempty, the sharp bound is instead

$$S_a \geq 1 + (M-1)\Phi_{n,B}(\pi, \alpha_M), \qquad \alpha_M = \frac{M\alpha}{M-1}. \tag{11}$$

Proved in Appendix C.4, p. 21; nonempty bound proved in Appendix D.1.

For nonempty classification, the fallback on unresolved requests is a uniform random label, correct with probability $1/M$; this expands the miscoverage threshold from α to α_M. When $B = 0$, $\Phi_{n,0} = \Psi_n$, recovering Theorem 2.

3.3 Subset bottlenecks and deficit characterization

When does budget B suffice to eliminate prediction set inflation entirely? Defining task demand as $c_j = (q_j - \alpha)_+$, the cost $\Phi_{n,B}(\pi, \alpha) = 0$ if and only if every subset's demand is within its capacity: $\sum_{j \in S} c_j \le F_{n,B}(S)$ for all $S \subseteq \mathcal{J}$ (Corollary 20, Appendix C.4). A total budget B alone can mask subset bottlenecks: two rare tasks that tend to be absent together can exceed their pairwise capacity even when total demand is within total capacity (Example 19). Theorem 11 in Appendix C.5 establishes an exact closed-form expression for $\Phi_{n,B}$ in terms of weighted subset deficits.

4 From an exact bound to an implementable procedure

Evaluating $\Phi_{n,B}$ via (8) requires checking 2^K subset constraints. We overcome this computational barrier via dual optimization (Section 4.1) and introduce a finite-sample calibration audit ensuring validity under empirical estimates (Section 4.2).

4.1 Stochastic dual optimization without enumerating subsets

For dual prices $\lambda \in \mathbb{R}^K$ and unobserved set U, let $\mathrm{Top}_B(\lambda, U)$ denote the sum of the B largest weights in U. Passing to the dual representation eliminates subset enumeration:

Theorem 4 (Dual representation and stochastic optimization)**.** *(i) The exact cost admits the dual representation*

$$\Phi_{n,B} = \max_{0 \le \lambda_j \le \pi_j} \left\{ \lambda^\top (q - \alpha \mathbf{1}) - \mathbb{E}\, \mathrm{Top}_B(\lambda, U) \right\}. \tag{12}$$

(ii) The mixture policy returned by Algorithm 1 obeys $|Q| \le B$ on every run, and its true unresolved probabilities u_j satisfy

$$\mathbb{E} \sum_j \pi_j (u_j - \alpha)_+ \le \Phi_{n,B} + e_{T,R}, \qquad e_{T,R} = \frac{\sqrt{\log 2/2} + 1/2}{\sqrt{T}} + \frac{1}{2\sqrt{R}}, \tag{13}$$

where the expectation is over the algorithm's training randomness.

Proved in Appendix D.2, p. 29.

In (12), $\lambda_j \in [0, \pi_j]$ acts as a price for leaving task j unresolved. Algorithm 1 ascends this dual: each iteration simulates an unobserved set U_t using task identities only (no true responses needed), queries the B highest-priced unobserved tasks, and increases the price of tasks left unresolved.

Algorithm 1 Budgeted query policy by stochastic top-B Hedge

Require: task frequencies π, history length n, budget B, level α, iterations T, snapshots R
Ensure: randomized query policy obeying $|Q| \le B$ on *every* run

$z_1 \leftarrow \mathbf{0} \in \mathbb{R}^K$; $\eta \leftarrow \sqrt{8 \log 2/T}$; draw $I_1, \dots, I_R \overset{\text{iid}}{\sim} \mathrm{Unif}[T]$
for $t = 1, \dots, T$ **do**
  simulate $J_1, \dots, J_n \overset{\text{iid}}{\sim} \pi$; $U_t \leftarrow \mathcal{J} \setminus \{J_1, \dots, J_n\}$ ▷ task identities only
  $\lambda_{t,j} \leftarrow \pi_j \,\mathrm{sigmoid}(z_{t,j})$ for all j ▷ dual prices in $[0, \pi_j]$
  $Q_t \leftarrow$ the $\min(B, |U_t|)$ tasks of U_t with largest $\lambda_{t,j}$ ▷ Top_B selection
  $z_{t+1,j} \leftarrow z_{t,j} + \eta\big(\mathbf{1}\{j \in U_t \setminus Q_t\} - \alpha\big)$ for all j ▷ unresolved ⇒ price up
  if $t \in \{I_1, \dots, I_R\}$ **then** store λ_t
return mixture policy g: draw $r \sim \mathrm{Unif}[R]$, query top-B tasks of real unobserved set U ranked by λ_{I_r}

4.2 Conservative calibration and certified prediction

The attaining rule requires the true unresolved probabilities u_j of the learned policy g. Underestimating u_j violates coverage; we therefore audit g on m independent identity-only episodes, computing

empirical rates $\widehat{u}_j$ and Hoeffding upper bounds $\overline{u}_j$:

$$\widehat{u}_j = \frac{1}{m}\sum_{i=1}^{m} \mathbf{1}\{j \text{ unresolved in audit } i\}, \quad \epsilon_m = \sqrt{\frac{\log(2K/\delta)}{2m}}, \quad \overline{u}_j = \min\{1, \widehat{u}_j + \epsilon_m\}. \tag{14}$$

With probability at least $1-\delta$, $\overline{u}_j \geq u_j$ holds simultaneously for all K tasks.

Certified rules and taskwise coverage guard. On request j: (1) if j is resolved, return cached $\{a_j\}$; (2) in regression, set $t_j = (1-(\alpha-\delta)/\overline{u}_j)_+$ and return $\mathcal{A}$ with probability t_j (and $\emptyset$ otherwise); (3) in nonempty M-class classification with $\kappa = M/(M-1)$, set $t_j = (1-\kappa(\alpha-\delta)/\overline{u}_j)_+$ and return $[M]$ with probability t_j (and a uniform singleton otherwise). Replacing the fallback with any nominal predictor $C^{\text{nom}}(j)$ defines a generic **taskwise coverage guard**: validity stems entirely from the guard wrapper ($u_j(1-t_j) \leq \alpha - \delta$), holding even if C^{nom} is empty; **Guarded CP** (Algorithm 3, Appendix D.3) decouples certification from prediction, exploiting learned structure on the non-guard branch.

Theorem 5 (Coverage after independent calibration)**.** *The certified prediction rule is taskwise honest at level α, with audit randomness included in the guarantee, and its normalized regression cost satisfies*

$$\widetilde{V}_a \leq \Phi_{n,B}(\pi, \alpha) + e_{T,R} + 2\epsilon_m + 2\delta. \tag{15}$$

For nonempty classification at threshold α_M, $(S_a - 1)/(M-1) \leq \Phi_{n,B}(\pi, \alpha_M) + e_{T,R} + 2\epsilon_m + (1+\kappa)\delta$.

Proved in Appendix D.3, p. 31.

Even if Algorithm 1 is trained on a misspecified simulator $\widetilde{\pi}$, auditing under the true law π rigorously preserves taskwise coverage; its efficiency degradation is bounded by $2\big(\min\{1, n\,\mathrm{TV}(\pi, \widetilde{\pi})\} + \|\pi - \widetilde{\pi}\|_1\big)$ when objective weights are also replaced by $\widetilde{\pi}$ (Proposition 25, Appendix D.4).

5 Lifetime coverage across extended deployments

In real deployments, an agent serves recurring tasks across an extended stream. An instantaneous guarantee that holds at each checkpoint separately does not bound the probability of making at least one error over the agent's operating lifetime.

5.1 Repeated guarantees and error accumulation

Let requests continue arriving independently from π, with every encounter revealing its exact response. Let $C_t(j)$ denote the prediction set for task j after t new encounters (checkpoint t). While checkpoint validity requires $\mathbb{P}_a\{a_j \in C_t(j)\} \geq 1-\alpha$ at each t separately, *lifetime coverage* requires

$$\forall a \in \mathcal{A}^{\mathcal{J}},\ \forall j \in \mathcal{J}, \qquad \mathbb{P}_a\{a_j \in C_t(j) \text{ for every } t \geq 0\} \geq 1-\alpha. \tag{16}$$

The gap between these criteria is severe: a rule satisfying α miscoverage at each individual checkpoint can accumulate error across time, causing the probability of at least one error over H checkpoints to reach $\alpha(1+pH)$ for a task of frequency p under a shifting-threshold construction with $(1-p)^H \geq \alpha$ ($(1-p)^H \geq \alpha_M = \frac{M\alpha}{M-1}$ for nonempty classification; Proposition 29, Appendix D.6).

5.2 Cumulative lifetime cost and universal query policy

A task initially unresolved remains so at checkpoint t with probability $(1-\pi_j)^t$. Summing over checkpoints $t = 0, \ldots, H$ defines expected unresolved encounters $v_j(H) = \sum_{t=0}^{H}(1-\pi_j)^t$, cumulative frequency weights $w_j(H) = \pi_j v_j(H) = 1-(1-\pi_j)^{H+1}$, and cumulative cost:

$$\Gamma_{n,B,H}(b) = \min_{r \in \mathcal{P}_{n,B}} \sum_j w_j(H)(q_j - r_j - b)_+. \tag{17}$$

Theorem 6 (Exact costs throughout deployment)**.** *Under lifetime coverage (16), the cumulative costs over checkpoints $0, \ldots, H$ satisfy:*

$$\begin{aligned} \textit{Vacuity probability:}\ & \textstyle\sum_{t=0}^{H} V_{a,t} \geq \Gamma_{n,B,H}(\alpha), \\ \textit{Bounded interval length:}\ & \textstyle\sum_{t=0}^{H} W_{a,t} \geq 2A\Gamma_{n,B,H}(\alpha), \\ \textit{Classification cardinality:}\ & \textstyle\sum_{t=0}^{H} S_{a,t} \geq (H+1)(1-\alpha) + (M-1)\Gamma_{n,B,H}(\alpha), \\ \textit{Nonempty classification:}\ & \textstyle\sum_{t=0}^{H} S_{a,t} \geq H + 1 + (M-1)\Gamma_{n,B,H}(\alpha_M). \end{aligned} \tag{18}$$

For each criterion, one query policy and one lifetime prediction rule attain equality at every response vector and every $H \geq 0$ simultaneously, without prior knowledge of the deployment horizon.

Proved in Appendix D.5, p. 33.

To attain the lower bound, the predictor draws a single full-domain-or-fallback decision at time zero for each unresolved task and preserves it until that task is first observed, spending its error budget exactly once.

Corollary 7 (Universal query policy)**.** *Order tasks by decreasing π_j, and let $S_k = \{1, \ldots, k\}$. Let G be the rank function of the capped polymatroid $\{r \in \mathcal{P}_{n,B} : r_j \leq (q_j - b)_+\}$ (Appendix C.5). The greedy increments $r_k^* = G(S_k) - G(S_{k-1})$ are feasible and minimize cumulative cost simultaneously for single-checkpoint weights π and lifetime weights $w(H)$ for every horizon $H \geq 0$. Furthermore, r^* is optimal for any nonincreasing time weights $\omega_0 \geq \omega_1 \geq \cdots \geq 0$, including geometric discounting $\omega_t = \gamma^t$ (Corollary 28, Appendix D.5).*

Proved in Appendix D.5, p. 33.

From a cold start ($n = B = 0$), cumulative regression cost equals $(1-\alpha)\mathbb{E}D_{H+1}$, where D_{H+1} is the number of distinct tasks observed; on infinite vocabularies this diverges as $H \to \infty$ even as per-step cost vanishes (Corollary 27, Appendix D.5).

6 Experiments: What Remains After Memory Is Perfect?

We evaluate our theoretical predictions on neural classifiers, scalable query optimization, and lifetime error control. Complete protocols, architectures, and regression logs are in Appendix E.

6.1 Taskwise Coverage Guard on Neural Classifiers

We train ten independent MLPs on disjoint splits of the digit vision benchmark (scikit-learn developers, 2026), holding out 600 tasks requested under a power-law distribution. Every encountered or queried label is cached verbatim, ensuring perfect retention. We compare: (1) *Cached pooled CP*: standard split conformal calibration on the historical stream to ensure marginal $1-\alpha$ coverage (Vovk et al., 2005). (2) *Certified nonempty*: the distribution-free rule of Theorem 5 with uniform singleton fallback. (3) *Guarded CP*: our instantiation of the taskwise coverage guard wrapping cached pooled CP; the guard guarantees tail safety, while conformal sets provide feature-conditioned informativeness on the non-guard branch (the guard equally accommodates Guarded Top-1 or Guarded APS/RAPS; Appendix D.3). A task is counted as *detected* if a one-sided Clopper–Pearson upper bound on its coverage falls below $1-\alpha$ after Bonferroni correction over the entire evaluation grid.

Table 1 and Figure 1 confirm our theoretical predictions:

- **Marginal Coverage Deceives**: Pooled CP attains 99.42% average coverage with set size 1.00, but leaves an average of 14.9 tasks significantly undercovered. Under a permuted-label control, undercoverage surges to 74.5 tasks.
- **Decoupling Certification from Prediction**: Guarded CP eliminates detected violations across all conditions. On natural tasks, it raises coverage from 95.73% to 99.84% over the prediction-agnostic certified rule with mean set size increasing only from 2.34 to 2.38.
- **Active Queries Compress Prediction Sets**: Under passive operation ($B = 0$), Guarded CP incurs an average set size of 3.24; allocating a proactive budget of $B = 32$ compresses the set size to 2.38, and $B = 96$ further compresses it to 1.87 without retraining the neural backbone (Appendix E, Table 4, Figure 11).

Table 1: Digit classifiers with perfect caches ($n = 128, B = 32, \alpha = 0.1$): high marginal coverage masks tail failure, while the coverage guard (Guarded CP) eliminates detected undercoverage. Means and 95% intervals over ten models (600 tasks each). "Detected" counts tasks with statistically significant undercoverage via simultaneous Clopper–Pearson bounds.

Labels	Prediction rule	Coverage (%)	Detected	Set size
Natural	Cached pooled CP	99.42 ± 0.12	14.9 ± 3.1	1.00 ± 0.00
Natural	Certified nonempty	95.73 ± 0.02	0.0 ± 0.0	2.34 ± 0.00
Natural	Guarded CP	99.84 ± 0.06	0.0 ± 0.0	2.38 ± 0.00
Permuted	Cached pooled CP	97.89 ± 0.76	74.5 ± 26.3	2.55 ± 0.08
Permuted	Certified nonempty	95.73 ± 0.02	0.0 ± 0.0	2.34 ± 0.00
Permuted	Guarded CP	99.53 ± 0.15	0.0 ± 0.0	2.72 ± 0.02

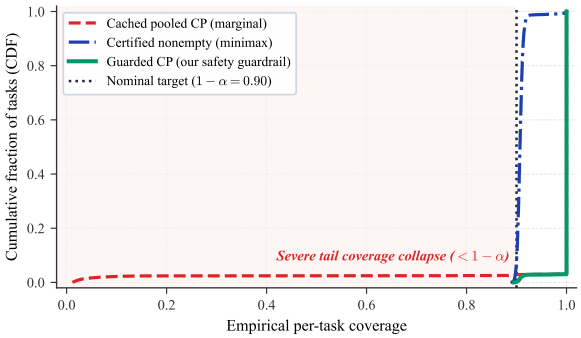

Figure 1: Taskwise coverage empirical CDF on digit tasks with perfect caches ($B = 32$). While pooled CP exhibits severe tail collapse despite high marginal coverage, the coverage guard (Guarded CP) eliminates detected undercoverage across all tasks.

6.2 Optimization, Calibration, and Lifetime Control

On small vocabularies where (8) is solved by linear programming, Algorithm 1 converges rapidly to the exact optimum (Appendix E, Figure 2). On large streams ($K = 4{,}096$), top-B Hedge trains in seconds and outperforms naive baselines. In calibration audits, certified rules using $\overline{u}_j$ rigorously maintain coverage, whereas plug-in estimates $\widehat{u}_j$ undercover (Figure 6). Auditing under the true law restores coverage under misspecified simulators (Figure 8). Under lifetime deployment, moving-threshold checkpoint rules accumulate error ($\alpha(1 + pH)$), while our lifetime rule strictly respects the target (Figure 9). Across hundreds of random instances, the universal greedy policy of Corollary 7 matches independent horizon-specific LP optima in every case.

7 Discussion: Inductive Biases, Task Relatedness, and Prior Knowledge

The minimax bounds of Sections 3–5 represent the worst-case baseline where task identities are discrete and independent. In practical systems, tasks often share underlying structure:

Task groups as latent equivalence classes.

Proposition 8 (Known response groups)**.** *If a known map $g : \mathcal{J} \to \{1, \ldots, G\}$ partitions tasks into G groups such that $a_j = \theta_{g(j)}$, then every exact cost in Sections 3 and 5 holds with π replaced by group frequencies $\rho_h = \sum_{j:g(j)=h} \pi_j$, and one query resolves an entire group.*

Proved in Appendix C.7, p. 25 (formal statement Proposition 21).

Proposition 8 collapses the effective vocabulary from K to G groups, enabling zero-inflation taskwise validity once $\max_h (1-\rho_h)^n \le \alpha$. Crucially, observing empirical agreement without a prior grouping does not suffice: honest inference must guard against unseen exceptions, paying the full discrete cost Ψ_n (Cai and Low, 2004).

Continuous geometry and non-parametric minimax rates. When task descriptors reside in a metric space $(\mathcal{X}, d)$ with Hölder continuous responses $|f(x) - f(x')| \leq L\, d(x, x')^{\beta}$, unobserved tasks are no longer completely indistinguishable. Uncertainty is governed by continuous *geometric visibility*: nearest-neighbor honest intervals achieve the minimax optimal expected width of order $\Theta(Ln^{-\beta/d})$, which robustly extends to matching rates under observation noise (Theorems 14, 23, Appendix C.8, C.10).

Heavy-tailed dynamics and information timing. In open-ended streams with Zipfian demand ($\pi_j \propto j^{-s}$, $s > 1$), the scale $B \asymp n^{1/s}$ marks the crossover budget where proactive queries begin to compress the passive tail cost $\Phi_{n,B} \asymp ((n+1)^{1/s} + B)^{1-s}$ (Theorem 30, Appendix D.7). Conversely, if an agent verifies responses *on demand* after a request is revealed, the necessary and sufficient query rate is precisely Ψ_n (Theorem 13, Appendix C.7). Reliable lifelong prediction is fundamentally an information-acquisition problem, not merely a storage problem.

8 Related work

Continual learning and memory-augmented architectures. Continual learning has traditionally focused on catastrophic forgetting (McCloskey and Cohen, 1989; French, 1999) via regularization (Kirkpatrick et al., 2017), replay (Lopez-Paz and Ranzato, 2017; Rebuffi et al., 2017), and parameter isolation (Rusu et al., 2016; Parisi et al., 2019). Modern agent architectures (Khandelwal et al., 2020; Lewis et al., 2020; Borgeaud et al., 2022; Packer et al., 2023) leverage persistent retrieval caches where historical interactions do not decay. We analyze the limiting regime of perfect memory, proving that zero retention error cannot solve reliable prediction on unobserved tasks. While Shihab et al. (2026) observe empirical coverage collapse during fine-tuning, we characterize the underlying information-theoretic bottleneck that persists with an exact cache and frozen representations.

Distribution-free uncertainty quantification and conformal prediction. Conformal prediction provides finite-sample marginal coverage (Vovk et al., 2005; Lei et al., 2018; Angelopoulos and Bates, 2023). Since exact conditional validity requires set inflation (Barber et al., 2021), relaxed conditional targets (Gibbs et al., 2023; Lee and Barber, 2021; Xie et al., 2025) and multi-task calibration (Fisch et al., 2021; Park et al., 2022; 2024) have been explored. We characterize the minimax cost of taskwise coverage under perfect memory and introduce a finite-sample *taskwise coverage guard* certifying validity independently of the nominal model; instantiating it around pooled conformal prediction (*Guarded CP*) decouples certification from prediction, exploiting learned structure on the non-guard branch. While time-uniform conformal prediction (Scharfstein and Kuchibhotla, 2026) tracks expanding sample sizes, our lifetime framework controls error across recurring deployments of fixed latent responses. Proactive query budgets define a submodular polymatroid (Edmonds, 1970), connecting missing-mass acquisition (Good, 1953; McAllester and Schapire, 2000; Noorani et al., 2025) with reduced-form combinatorial allocation (Border, 1991; Cai et al., 2011; Dughmi, 2021).

9 Conclusion

Perfect memory solves historical retention, but retention alone does not ensure reliable lifelong prediction. In open-ended streams, marginal coverage masks severe failures on long-tail requests. We formalized taskwise honest prediction in lifelong streams and established exact minimax lower bounds on prediction set efficiency, showing that missing mass forces set inflation on unobserved tasks under distribution-free conditions. We demonstrated that proactive query allocation over submodular polymatroids, solved via stochastic top-B dual Hedge with conservative calibration, significantly compresses certified sets. Furthermore, under lifetime deployment, a single greedy query policy is universally optimal across all horizons. Finally, the taskwise coverage guard bridges theory and practice, certifying validity independently of the nominal predictor, while Guarded CP couples certification with learned informativeness on the non-guard branch. Reliable lifelong intelligence requires principled certification and proactive acquisition for missing tasks.

AI USE STATEMENT

Generative AI tools were used for grammatical refinement, language polishing, and literature search assistance during the preparation of this manuscript. Generative AI was not used to synthesize the underlying mathematical theorems or replace human conceptual analysis. The authors have independently derived and verified all mathematical proofs, designed and executed all empirical evaluations. The authors take full responsibility for the technical accuracy and integrity of the entire manuscript.

REPRODUCIBILITY STATEMENT

All mathematical claims are rigorously and completely proved in the appendices (Appendix C and D). All algorithmic procedures are fully specified with explicit pseudocode (Algorithms 1 and 3).

ETHICS STATEMENT

This work presents foundational theoretical research, and we foresee no direct negative ethical or societal implications.

REFERENCES

Anastasios N. Angelopoulos and Stephen Bates. Conformal prediction: A gentle introduction. *Foundations and Trends in Machine Learning*, 16(4):494–591, 2023. Preprint: arXiv:2107.07511.

Rina Foygel Barber, Emmanuel J. Candès, Aaditya Ramdas, and Ryan J. Tibshirani. The limits of distribution-free conditional predictive inference. *Information and Inference: A Journal of the IMA*, 2021. Preprint: arXiv:1903.04684.

Kim C. Border. Implementation of reduced form auctions: A geometric approach. *Econometrica*, 59 (4):1175–1187, 1991. doi: 10.2307/2938181.

Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, et al. Improving language models by retrieving from trillions of tokens. In *Proceedings of the 39th International Conference on Machine Learning*, 2022. arXiv:2112.04426.

T. Tony Cai and Mark G. Low. An adaptation theory for nonparametric confidence intervals. *The Annals of Statistics*, 32(5):1805–1840, 2004. Preprint: arXiv:math/0503662.

Yang Cai, Constantinos Daskalakis, and S. Matthew Weinberg. A constructive approach to reduced-form auctions with applications to multi-item mechanism design. *arXiv preprint arXiv:1112.4572*, 2011. Revised 2017.

Shaddin Dughmi. Matroid secretary is equivalent to contention resolution. *arXiv preprint arXiv:2103.04205*, 2021.

Jack Edmonds. Submodular functions, matroids, and certain polyhedra. In *Combinatorial Structures and Their Applications*, pages 69–87. Gordon and Breach, 1970.

Adam Fisch, Tal Schuster, Tommi Jaakkola, and Regina Barzilay. Few-shot conformal prediction with auxiliary tasks. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 3329–3339, 2021.

Robert M. French. Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*, 3 (4):128–135, 1999. doi: 10.1016/S1364-6613(99)01294-2.

Isaac Gibbs, John J. Cherian, and Emmanuel J. Candès. Conformal prediction with conditional guarantees. *arXiv preprint arXiv:2305.12616*, 2023.

I. J. Good. The population frequencies of species and the estimation of population parameters. *Biometrika*, 40(3–4):237–264, 1953. doi: 10.1093/biomet/40.3-4.237.

Samuel Karlin. Central limit theorems for certain infinite urn schemes. *Journal of Mathematics and Mechanics*, 17(4):373–401, 1967.

Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. Generalization through memorization: Nearest neighbor language models. In *International Conference on Learning Representations*, 2020. arXiv:1911.00172.

James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13):3521–3526, 2017. doi: 10.1073/pnas.1611835114.

Samory Kpotufe. k-NN regression adapts to local intrinsic dimension. In *Advances in Neural Information Processing Systems*, 2011. arXiv:1110.4300.

Yonghoon Lee and Rina Foygel Barber. Distribution-free inference for regression: Discrete, continuous, and in between. In *Advances in Neural Information Processing Systems*, 2021. arXiv:2105.14075.

Jing Lei, Max G'Sell, Alessandro Rinaldo, Ryan J. Tibshirani, and Larry Wasserman. Distribution-free predictive inference for regression. *Journal of the American Statistical Association*, 113(523): 1094–1111, 2018. doi: 10.1080/01621459.2017.1307116.

Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. Retrieval-augmented generation for knowledge-intensive NLP tasks. In *Advances in Neural Information Processing Systems*, 2020. arXiv:2005.11401.

David Lopez-Paz and Marc'Aurelio Ranzato. Gradient episodic memory for continual learning. In *Advances in Neural Information Processing Systems*, 2017. arXiv:1706.08840.

David McAllester and Robert E. Schapire. On the convergence rate of good-turing estimators. In *Proceedings of the 13th Annual Conference on Computational Learning Theory*, pages 1–6, 2000.

Michael McCloskey and Neal J. Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of Learning and Motivation*, volume 24, pages 109–165. Academic Press, 1989.

Sima Noorani, Shayan Kiyani, George Pappas, and Hamed Hassani. Conformal prediction beyond the seen: A missing mass perspective for uncertainty quantification in generative models. *arXiv preprint arXiv:2506.05497*, 2025.

Charles Packer, Sarah Wooders, Kevin Lin, Vivian Fang, Shishir G. Patil, Ion Stoica, and Joseph E. Gonzalez. MemGPT: Towards LLMs as operating systems. *arXiv preprint arXiv:2310.08560*, 2023.

German I. Parisi, Ronald Kemker, Jose L. Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113:54–71, 2019. doi: 10.1016/j.neunet.2019.01.012.

Sangdon Park, Edgar Dobriban, Insup Lee, and Osbert Bastani. PAC prediction sets for meta-learning. In *Advances in Neural Information Processing Systems*, 2022. arXiv:2207.02440.

Sangwoo Park, Kfir M. Cohen, and Osvaldo Simeone. Few-shot calibration of set predictors via meta-learned cross-validation-based conformal prediction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. Preprint: arXiv:2210.03067.

Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H. Lampert. iCaRL: Incremental classifier and representation learning. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017. arXiv:1611.07725.

Andrei A. Rusu, Neil C. Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *arXiv preprint arXiv:1606.04671*, 2016.

Kayla E. Scharfstein and Arun Kumar Kuchibhotla. Time-uniform conformal and PAC prediction. *arXiv preprint arXiv:2602.06297*, 2026.

scikit-learn developers. load_digits: Optical recognition of handwritten digits dataset. Official scikit-learn dataset documentation, 2026. URL https://scikit-learn.org/stable/modules/generated/sklearn.datasets.load_digits.html. Accessed September 17, 2026; experiments use scikit-learn 1.8.0.

Burr Settles. Active learning literature survey. Technical Report Computer Sciences Technical Report 1648, University of Wisconsin–Madison, 2009.

Ibne Farabi Shihab, Sanjeda Akter, and Anuj Sharma. Continual calibration: Coverage can collapse before accuracy in lifelong LLM fine-tuning. *arXiv preprint arXiv:2604.23987*, 2026.

Vladimir Vovk, Alexander Gammerman, and Glenn Shafer. *Algorithmic Learning in a Random World*. Springer, 2005.

Tianmin Xie, Yanfei Zhou, Ziyi Liang, Stefano Favaro, and Matteo Sesia. Conformal inference for open-set and imbalanced classification. *arXiv preprint arXiv:2510.13037*, 2025.

George Kingsley Zipf. *Human Behavior and the Principle of Least Effort*. Addison-Wesley, 1949.

TABLE OF CONTENTS

1 Introduction **1**

2 Problem formulation: taskwise honest prediction in interaction streams **2**

3 Fundamental limits of taskwise honest prediction and active allocation **3**

3.1 Passive costs: indistinguishability and missing mass 4

3.2 Active allocation over submodular polymatroids 4

3.3 Subset bottlenecks and deficit characterization 5

4 From an exact bound to an implementable procedure **5**

4.1 Stochastic dual optimization without enumerating subsets 5

4.2 Conservative calibration and certified prediction 5

5 Lifetime coverage across extended deployments **6**

5.1 Repeated guarantees and error accumulation 6

5.2 Cumulative lifetime cost and universal query policy 6

6 Experiments: what remains after memory is perfect? **7**

6.1 Taskwise coverage guard on neural classifiers 7

6.2 Optimization, calibration, and lifetime control 8

7 Discussion: inductive biases, task relatedness, and prior knowledge **8**

8 Related work **9**

9 Conclusion **9**

A Notation, terminology, and where each result is proved **15**

B Additional foundational statements **16**

B.1 Passive costs, infinite support, and unknown environments 16

B.2 An explicit formula for the exact cost with queries 17

B.3 Verification and structured transfer: formal statements 18

C Proofs of foundational results **19**

C.1 Avoidance and indistinguishability 19

C.2 Passive costs and the infinite-support separation 20

C.3 Unknown task probabilities 21

C.4 The query polymatroid and the exact costs with queries 21

C.5 An exact weighted subset-deficit formula 23

C.6 Exact algorithms and limiting cases 23

C.7 Verification and known task groups 25
C.8 Geometric visibility and exact-response rates 25
C.9 Power-law asymptotics 26
C.10 Gaussian observation noise: a complete interpolation theorem 26

D Proofs of algorithmic and lifetime results **28**
D.1 Nonempty prediction sets and oracle reductions 28
D.2 The stochastic optimization guarantee 29
D.3 Independent audit: coverage and efficiency 31
D.4 Sensitivity to a misspecified simulator 32
D.5 Lifetime honesty: complete proof and universal query policy 33
D.6 A sharp checkpoint-versus-lifetime counterexample 36
D.7 Active heavy tails and cumulative lifetime growth 37

E Experimental settings and numerical results **38**
E.1 Experimental environment and evaluation protocols 38
E.2 Exact optimization, duality, and nonempty-set checks 38
E.3 Vocabulary scaling and query-budget ablations 39
E.4 Calibration size and simulator misspecification 42
E.5 Lifetime trajectories and common optimal query policy 43
E.6 A learned classifier with exact memory 45
E.7 Noisy geometric recovery 47

A NOTATION, TERMINOLOGY, AND WHERE EACH RESULT IS PROVED

Terminology map. The proofs occasionally use the following synonyms, all of which refer to objects defined in the main text.

Used in the appendices	Meaning in the main text
response, label	the fixed correct answer a_j of task j ("label" in the classification case)
configuration	a response vector $a = (a_j)_{j \in \mathcal{J}}$
honest, taskwise honest	satisfies taskwise coverage (3) at level α
acquisition, query	revealing an unobserved task's response before the request arrives
kernel, query kernel	a randomized query rule $Q \mid U$, i.e. a response-independent policy
hard-B policy, hard budget	$\lvert Q \rvert \le B$ on every run
resolved / unresolved	observed in the history or queried / neither
unseen set	U, the set of unobserved tasks after the history
transcript $\mathcal{T}$	the learner's complete record (history, queries, responses, random seed)
frontier, exact cost	the sharp lower bound of Theorem 2, 3, or 6
residual, vacuity	the cost that remains after queries; the event that a set has infinite length
rank function, polymatroid	$F_{n,B}$ in (7); the region $\mathcal{P}_{n,B}$ it defines

Symbols are listed in Table 2. Two symbols are deliberately reused with local meaning inside proofs and never leak out: T denotes a subset of tasks inside Appendix C.5, and L_t denotes a per-checkpoint cost inside Corollary 28.

Table 2: Symbols used throughout the manuscript.

Symbol	Meaning	Where
π_j, K, n, B	request frequency; number of tasks; history length; hard query budget	§2
U, Q	unobserved tasks; queried tasks ($Q \subseteq U$, $\lvert Q \rvert \le B$)	(2)
α; $\mathbb{P}_a$	allowed miscoverage per task; probability at response vector a	Def. 1
$V_a, W_a, S_a, \widetilde{V}_a$	vacuity probability; expected length; expected cardinality; normalized cost	(4)
$q_j = (1-\pi_j)^n$	probability that task j is unobserved	§3.1
Ψ_n	passive cost (no queries)	(5)
$r_j, u_j = q_j - r_j$	probability that j is queried; that j remains unresolved	§3.2
$F_{n,B}(S), \mathcal{P}_{n,B}$	capacity of subset S; feasible region of query probabilities	(7)
$\Phi_{n,B}(\pi, b)$	exact cost with queries at threshold b	(8)
$\alpha_M = M\alpha/(M-1)$	threshold for nonempty M-class prediction	§3.2
$c_j = (q_j - \alpha)_+$	demand of task j	Cor. 20
λ, Top_B, $e_{T,R}$	dual prices; sum of B largest prices; optimization error	§4.1
$\widehat{u}_j, \epsilon_m, \overline{u}_j, \delta$	audit estimate; its margin; upper bound; audit failure budget	(14)
$C_t(j), v_j(H), w_j(H), \Gamma_{n,B,H}$	prediction at checkpoint t; unresolved checkpoints; lifetime weight; lifetime cost	§5

Proof map.

Main-text statement	What it says	Proved in
Lemma 15	unobserved task forces inclusion	App. C.1 (Lemma 15)
Theorem 2	exact passive costs	App. C.2
Lemma 18	feasible query probabilities form $\mathcal{P}_{n,B}$	App. C.4 (Lemma 18)
Theorem 3	exact costs with queries	App. C.4; nonempty case App. D.1
Corollary 17	$B = 0$ limit; infinite support	App. C.2; App. C.6
Corollary 20	zero cost iff demand within capacity	App. C.4
Theorem 4	dual form; Algorithm 1 error	App. D.2
Theorem 5	calibrated coverage and cost	App. D.3
Proposition 25	misspecified simulator	App. D.4 (Proposition 25)
Proposition 29	checkpoint rule fails lifetime target	App. D.6
Theorem 6	exact lifetime costs	App. D.5
Corollary 7	one policy for every horizon and discount	App. D.5 (and Corollary 28)
Corollary 27	cold-start cumulative cost	App. D.5
Proposition 8	known groups	App. C.7 (Proposition 21)
Theorems 14, 23	geometric visibility & noisy recovery	App. C.8, C.10
Theorem 30	power-law query scaling	App. D.7
Theorem 13	on-demand verification cost	App. C.7

B Additional foundational statements

The main text states the results needed for its argument. This appendix states, and briefly motivates, the additional results that the proofs and experiments rely on: the infinite-support separation, the cost of an unknown task distribution, an explicit formula for the exact cost $\Phi_{n,B}$ in terms of subset deficits, the cold-start case, and the formal versions of the structure results of Section 7. Their proofs are in Appendix C; the stochastic algorithm and the lifetime results are proved in Appendix D.

B.1 Passive costs, infinite support, and unknown environments

Recall the passive cost from (5),

$$\Psi_n(\pi, \alpha) = \sum_j \pi_j (q_j - \alpha)_+, \qquad q_j = (1 - \pi_j)^n. \tag{19}$$

(Earlier versions of this manuscript call Ψ_n the *residual visibility profile*.) It is not the ordinary missing mass $\sum_j \pi_j q_j$: coverage permits a separate error budget α for *each* task before the residuals are averaged. Theorem 2 in the main text states the exact passive costs; its proof is in Appendix C.2.

Proof mechanism. On the event that task j is unobserved, the entire transcript law is unchanged by replacing a_j with any y. Honesty therefore implies the inclusion bound (29). Integrating that bound over $y \in [-A, A]$ gives the bounded-width bound. Integrating over $[-L, L]$, dividing by $2L$, and letting $L \to \infty$ yields the infinite-length probability bound. For classification, count the true-label inclusion probability $1 - \alpha$ and the $M - 1$ other labels using (29). An unobserved task attains the regression bounds by returning the full domain with probability $(1 - \alpha/q_j)_+$ and the empty set otherwise; an observed task returns its exact response. A slight thinning of observed singletons gives the classification equality.

Corollary 9 (Infinite task support and the exact recovery threshold)**.** *With infinite positive support, $\Psi_n > 0$ at every finite n, and honest unbounded-response prediction has infinite expected length at every configuration, including $a_j = 0$ for all j. With finite support, an honest rule of finite expected length exists exactly when*

$$\max_j (1 - \pi_j)^n \le \alpha. \tag{20}$$

When (20) holds, zero-length sets suffice in the exact-response experiment.

Proved in Appendix C.2, p. 20.

Remark. This theorem does not preclude pointwise lifelong success. Memorizing encountered responses and returning an empty set otherwise gives task-j error $q_j \to 0$ and marginal error $\sum_j \pi_j q_j \to 0$ as $n \to \infty$. Yet $\sup_j q_j = 1$ at every finite n for infinite support. Thus pointwise learning, environmental average performance, and uniform taskwise reliability are three different objectives.

Theorem 10 (The cost of an unknown task environment)**.** *Suppose $B = 0$, there are at least two tasks, and the procedure is not supplied π. Require taskwise coverage (3) uniformly over all positive task distributions as well as all response configurations. Then the three exact costs in (6) hold with Ψ_n replaced by*

$$\widetilde{\Psi}_n = (1-\alpha)\sum_j \pi_j(1-\pi_j)^n. \tag{21}$$

One distribution-independent rule attains each bound at every π, a. In the unbounded-response model, finite expected length is impossible at any finite horizon, even for a finite task vocabulary.

Proved in Appendix C.3, p. 21.

Remark. The distinction is inferential: an empirical frequency cannot be substituted for a supplied frequency while retaining the same unrestricted uniform guarantee. Conditional on not seeing j, shrinking its population frequency toward zero while keeping the relative frequencies of other tasks fixed. The conditional transcript law is unchanged, forcing conditional coverage $1-\alpha$ on that unseen event, with no error allowance left to trade against. The remaining results in this appendix return to the supplied-π setting of the main text.

B.2 An explicit formula for the exact cost with queries

The main text defines $\Phi_{n,B}$ in (8) as a minimum over the polymatroid $\mathcal{P}_{n,B}$, and Example 19 in Appendix C.4 computes it by hand from the deficit of one pair of tasks. The following theorem shows that this is the general pattern: the exact cost is a weighted sum of subset deficits, where the deficit of a set is the amount by which the demand of some subset exceeds its capacity.

Theorem 11 (Weighted subset-deficit formula)**.** *Order $\pi_1 \geq \cdots \geq \pi_K$, set $\pi_{K+1} = 0$, and let $S_k = \{1, \ldots, k\}$ be the k most frequent tasks. With the demands $c_j = (q_j - \alpha)_+$ of Corollary 20, write $c(S) = \sum_{j\in S} c_j$ and define the* deficit

$$D(S) = \max_{T\subseteq S}\{c(T) - F_{n,B}(T)\} \ \ (\geq 0).$$

Then

$$\Phi_{n,B} = \sum_{k=1}^{K}(\pi_k - \pi_{k+1})D(S_k). \tag{22}$$

Writing $G(S) = c(S) - D(S)$, the increments $r^*_k = G(S_k) - G(S_{k-1})$ are optimal query probabilities, and they are realizable by a hard-budget response-independent policy.

Proved in Appendix C.5, p. 23.

The formula converts subset infeasibility into an exact weighted prediction cost. Its proof shows that G is the rank function of the capped polymatroid $\{r \in \mathcal{P}_{n,B} : r \leq c\}$, so that the greedy vector r^* is the polymatroid greedy solution for the weights π; the same vector reappears in Corollary 7 because the lifetime weights $w_j(H)$ have the same order as π_j.

Constructive solution by linear programming. Write $p_u = \mathbb{P}(U = u)$ for the law of the unobserved set. A finite linear program uses joint masses $z_{u,Q} \geq 0$, one for each unobserved set u and each admissible query set $Q \subseteq u$ with $|Q| \leq B$, satisfying

$$\sum_Q z_{u,Q} = p_u, \quad r_j = \sum_{u,Q:j\in Q} z_{u,Q}, \quad d_j \geq q_j - r_j - \alpha, \quad d_j \geq 0, \qquad \min \sum_j \pi_j d_j. \tag{23}$$

Given $U = u$, sample Q with probabilities $z_{u,Q}/p_u$. With $u_j = q_j - r_j$, unresolved tasks use the full response domain with probability $(1-\alpha/u_j)_+$; resolved tasks use their exact labels. This implements the regression frontier. Appendix C.6 gives the classification modification and exact pseudocode. The enumerative LP is exponential in K, so we use it as an exact benchmark rather than as a scalable solver; Algorithm 1 is the scalable substitute.

Proposition 12 (Cold start and finite approximation)**.** *At $n = 0$, order $\pi_1 \geq \cdots \geq \pi_K$ and write $c = 1-\alpha$. If $B < Kc$, let $m = \lfloor B/c \rfloor$ and $\eta = B - mc$. Then*

$$\Phi_{0,B} = c\sum_{j=m+1}^{K} \pi_j - \eta\pi_{m+1}; \tag{24}$$

if $B \geq Kc$, $\Phi_{0,B} = 0$. On countable supports, optimizing queries only over a finite set G of tasks approximates the optimal cost within $\sum_{j \notin G} \pi_j (q_j - \alpha)_+$, which is at most the discarded task mass.

Proved in Appendix C.6, p. 23.

B.3 Verification and structured transfer: formal statements

Verification after seeing the requested task. Proactive queries and target-specific verification are different resources. Suppose no proactive queries are made, but after observing the request J the learner may verify its response once before predicting. Section 7 in the main text discusses this setting; the following is the formal statement.

Theorem 13 (Exact on-demand verification cost)**.** *In the unbounded-response model, require taskwise honesty and finite-length output almost surely at every configuration. Every such rule must verify the requested task with probability at least Ψ_n, at every configuration. This bound is attained by verifying an unobserved task j with probability $(1 - \alpha/q_j)_+$ and returning its known response, or an empty set when it remains unresolved.*

Proved in Appendix C.7, p. 25.

Thus the passive optimal vacuity probability is exactly the necessary verification frequency. A finite proactive budget cannot cover infinitely many unrestricted tasks (Corollary 17), whereas on-demand verification yields an honest nonvacuous response at a finite average cost per request. This is an information-timing distinction, not a memory advantage.

Known groups. If a known map $g : \mathcal{J} \to \{1, \ldots, G\}$ satisfies $a_j = \theta_{g(j)}$, one encounter reveals an entire task group. Proposition 21 in Appendix C.7 states and proves that every exact cost above applies with group masses $\rho_h = \sum_{j:g(j)=h} \pi_j$ in place of π. In particular, infinitely many task identities can admit zero-width prediction once $\max_h (1 - \rho_h)^n \leq \alpha$. This is the formal content of Proposition 8.

Known smoothness. Suppose task descriptors form a metric probability space $(\mathcal{X}, d, \mu)$ of diameter at most one, and responses obey a *known* Hölder condition

$$|f(x) - f(x')| \leq L d(x, x')^\beta, \qquad 0 < \beta \leq 1. \tag{25}$$

The $n \geq 1$ historical descriptors are i.i.d. from μ, with exact responses. Require coverage at every fixed test descriptor x uniformly over this function class. Let $R_n(x) = \min_i d(X_i, x)$ be the distance to the nearest observed descriptor, and let $q_x(h) = (1 - \mu\{z : d(z, x) < h\})^n$ be the probability that no observation falls within distance h of x; this is the continuous analogue of q_j.

Theorem 14 (Geometric visibility and matching rates)**.** *Every honest rule obeys, at the constant function $f = 0$,*

$$\mathbb{E}_0 \operatorname{Leb}(C_n(x)) \geq 2L \sup_{0 < h \leq 1} h^\beta (q_x(h) - \alpha)_+. \tag{26}$$

The nearest-neighbor interval centered at $f(X_{i^*})$ with radius $LR_n(x)^\beta$ has coverage one and expected length $2L\mathbb{E}R_n(x)^\beta$. If, uniformly in x and $0 < h \leq 1$,

$$c_0 h^d \leq \mu(B(x, h)) \leq C_0 h^d, \tag{27}$$

then the optimal worst-case honest expected length is $\Theta(Ln^{-\beta/d})$, with constants depending only on $\alpha, \beta, d, c_0, C_0$.

Proved in Appendix C.8, p. 25.

The lower bound uses a continuum of localized Hölder bumps that are invisible when their support is missed. This puts discrete identity visibility and geometric transfer in the same proof framework. The rate and the nearest-neighbor mechanism have established nonparametric precedents (Kpotufe, 2011); what is new here is that they supply the construction matching an honest taskwise lower bound. Theorem 23 in Appendix C.10 proves the interpolation rate (37) with Gaussian observation noise.

Remark. If all responses happen to be equal, an algorithm honest over unrestricted configurations still pays (10). If equality is known as a model restriction, one encounter suffices. This exact separation is

an instance of the broader distinction between estimation adaptivity and honest inference adaptivity (Cai and Low, 2004). It locates the extra assumption needed to turn apparent transfer into a uniform reliability guarantee.

Passive power-law asymptotics. For $\pi_j = c_s j^{-s}$, $s > 1$, $c_s = (\sum_{j\geq 1} j^{-s})^{-1}$, define

$$K_{s,\alpha} = \frac{c_s^{1/s}}{s} \int_0^{\log(1/\alpha)} (e^{-u} - \alpha) u^{-1/s}\, du. \tag{28}$$

Proposition 22 in Appendix C.9 proves $\Psi_n \sim K_{s,\alpha} n^{1/s-1}$. Theorem 30 in Appendix D.7 adds the uniform active-budget rate discussed in Section 7 and a lifetime occupancy consequence.

C Proofs of foundational results

This appendix proves the foundational statements of Sections 3 and 7 and the deferred results of Appendix B. The stochastic algorithm, calibration, and lifetime proofs are in Appendix D. Throughout, the randomness used by the learner is independent of the task responses, and $\mathcal{T}$ denotes the learner's complete transcript.

C.1 Avoidance and indistinguishability

A random prediction set is understood through a jointly measurable indicator $(\mathcal{T}, y) \mapsto \mathbf{1}\{y \in C(\mathcal{T}, j)\}$; thus all integrals below are well defined by Tonelli's theorem. In the finite-label experiment, intersecting the output with $[M]$ never decreases coverage or increases its size. The main-text Lemma 15 is the special case $B = 0$ of the following lemma, in which the unresolved event E_j reduces to $\{j \in U\}$ and u_j to q_j.

Lemma 15. *Fix j and all other responses a_{-j}. Let E_j be the event that j is neither encountered passively nor queried. The law of the final transcript restricted to E_j, including predictor randomization for task j, is independent of a_j. In particular, $u_j = \mathbb{P}_a(E_j)$ is independent of a_j. Every taskwise honest procedure satisfies*

$$\mathbb{P}_a\{y \in C(\mathcal{T}, j), E_j\} \geq (u_j - \alpha)_+ \quad \textit{for every allowed response } y. \tag{29}$$

The same conclusion holds for verification after revealing the test identity, conditional on that identity being j.

Proof. Couple two configurations differing only at j by the same passive task identities and the same complete random seed. If j is passively observed, E_j is false in both experiments. Otherwise their passive transcripts agree. Inductively, as long as neither has queried j, their available labels, random seed, stopping decisions, and next-query choices agree. Therefore their first query of j, if any, occurs at the same step. On the event that no such query occurs, the entire transcripts and predictions agree. This proves the restricted-law assertion, without assuming the query sequence is nonadaptive. Write $a^{(y)}$ for the configuration with coordinate j replaced by y. Honesty and the event's common probability give

$$1 - \alpha \leq \mathbb{P}_{a^{(y)}}\{y \in C(\mathcal{T}, j)\} \leq 1 - u_j + \mathbb{P}_{a^{(y)}}\{y \in C(\mathcal{T}, j), E_j\}.$$

Transfer the last term by the restricted-law identity. Its nonnegativity gives the positive part. Revealing j before the query changes none of this coupling for that fixed identity. □

The next lemma converts forced inclusion of each candidate into forced size, one domain at a time.

Lemma 16. *Let $v_j = (u_j - \alpha)_+$. At every configuration, an honest rule satisfies*

$$\mathbb{P}_a\{E_j, \mathrm{Leb}(C(\mathcal{T}, j)) = \infty\} \geq v_j \qquad (\mathcal{A} = \mathbb{R}), \tag{30}$$

$$\mathbb{E}_a \mathrm{Leb}(C(\mathcal{T}, j)) \geq 2Av_j \qquad (\mathcal{A} = [-A, A]), \tag{31}$$

$$\mathbb{E}_a |C(\mathcal{T}, j)| \geq 1 - \alpha + (M - 1)v_j \qquad (\mathcal{A} = [M]). \tag{32}$$

Proof. For bounded responses, integrate (29) over $[-A, A]$. For real responses, for positive integers m define

$$Z_m = \mathbf{1}_{E_j} \frac{\mathrm{Leb}(C(\mathcal{T}, j) \cap [-m, m])}{2m}.$$

Tonelli gives $\mathbb{E}Z_m \geq v_j$. Also $0 \leq Z_m \leq 1$, and $Z_m \to 0$ whenever E_j is false or the output has finite Lebesgue measure. Consequently

$$v_j \leq \limsup_m \mathbb{E}Z_m \leq \mathbb{E} \limsup_m Z_m \leq \mathbb{P}_a\{E_j, \mathrm{Leb}(C(\mathcal{T}, j)) = \infty\}.$$

The middle inequality is reverse Fatou for variables bounded above by one. This argument is stronger than merely proving infinite expected length; the latter alone would not imply a positive probability of infinite length. For classification, separate the true label from the other $M - 1$ labels:

$$\mathbb{E}_a|C(\mathcal{T}, j)| = \sum_{y=1}^{M} \mathbb{P}_a\{y \in C(\mathcal{T}, j)\} \geq 1 - \alpha + \sum_{y \neq a_j} \mathbb{P}_a\{y \in C(\mathcal{T}, j), E_j\}.$$

Apply Lemma 15 to each term. □

C.2 Passive costs and the infinite-support separation

Recall that without queries $u_j = q_j = (1 - \pi_j)^n$ and that $\Psi_n = \sum_j \pi_j (q_j - \alpha)_+$.

Proof of Theorem 2. Weight the three conclusions of Lemma 16 by π_j and sum, using nonnegative summation for a countable task set. This proves all lower bounds at each fixed configuration, not only after a supremum over configurations.

For regression, on a resolved task output the singleton containing its response. On an unresolved task with $q_j > \alpha$, output the full domain with probability $1 - \alpha/q_j$, and the empty set otherwise. If $q_j \leq \alpha$, output the empty set on unresolved histories. Thus coverage is

$$1 - q_j + (q_j - \alpha)_+ = 1 - \min(q_j, \alpha) \geq 1 - \alpha.$$

The probability of the full domain is $(q_j - \alpha)_+$, giving both equalities. This construction works for the real domain and the bounded interval, with different meanings of "full domain."

For classification, when $q_j > \alpha$ use exactly the same rule. Its expected size is

$$1 - q_j + M(q_j - \alpha) = 1 - \alpha + (M - 1)(q_j - \alpha).$$

When $q_j \leq \alpha$, return an empty set on unresolved histories; on resolved histories return the true singleton with probability $(1 - \alpha)/(1 - q_j)$, otherwise the empty set. This probability is in $[0, 1]$ and has a nonzero denominator because $q_j \leq \alpha < 1$. Coverage and expected size both equal $1 - \alpha$. All randomization probabilities depend only on known task frequencies, not on responses. If $q_j > \alpha$ in the real domain, (30) proves infinite expected length at every configuration. □

Corollary 17. *(i) At $B = 0$, $\Phi_{n,0} = \Psi_n$, and Theorem 3 reduces to Theorem 2. (ii) If $\mathcal{J}$ is infinite, then for every finite n and B and every query policy, $\Phi_{n,B} > 0$. Moreover, for every response configuration a, there exists a task j (which may depend on a) such that $u_j(a) > \alpha$, implying $\mathbb{E}_a \mathrm{Leb}(C(\mathcal{T}, j)) = \infty$ in the unbounded domain, and hence $\sup_{j' \in \mathcal{J}} \mathbb{E}_a \mathrm{Leb}(C(\mathcal{T}, j')) = \infty$.*

Proof of Corollary 9, of Corollary 17(i), and of the consistency statements. Part (i) of Corollary 17 is immediate: at $B = 0$ every capacity $F_{n,0}(S)$ is zero, so $\mathcal{P}_{n,0} = \{0\}$ and $\Phi_{n,0} = \Psi_n$.

On an infinite positive support, $\pi_j \to 0$ along some sequence of distinct indices. For any fixed finite n, $q_j = (1 - \pi_j)^n \to 1$ along this sequence, including $q_j = 1$ at $n = 0$. Some $q_j > \alpha$, so its positively weighted contribution makes $\Psi_n > 0$. With finitely many tasks, the absence of such a term is exactly (20), and the attaining rule then has length zero almost surely. Conversely, any positive term forces infinite expected length, even at the all-zero configuration.

For each fixed j, $\pi_j > 0$ implies $q_j \to 0$ as $n \to \infty$. Since $0 \leq q_j \leq 1$, dominated convergence under the probability measure π gives $\sum_j \pi_j q_j \to 0$. Infinite support gives $\sup_j q_j = 1$ for every finite n as just shown. The memorize-or-abstain predictor has precisely these taskwise and marginal error probabilities; it need not satisfy the finite-sample uniform honesty requirement. □

C.3 UNKNOWN TASK PROBABILITIES

Proof of Theorem 10. Fix any positive task distribution π, a task j, and the other responses a_{-j}. Because the support has at least two tasks, $\pi_j < 1$. Define the conditional distribution of other identities by $\eta_i = \pi_i/(1-\pi_j)$ for $i \neq j$. For every $\varepsilon \in (0,1)$ form the positive task law

$$\pi_j^{(\varepsilon)} = \varepsilon, \qquad \pi_i^{(\varepsilon)} = (1-\varepsilon)\eta_i \quad (i \neq j).$$

Let E_j denote no passive encounter of j. Conditional on E_j, all n identities are i.i.d. from η under every law in this family, including the original π. Their labels involve only a_{-j}. Since the algorithm is not given π, its conditional output law is the same for all ε and all replacements $a_j = y$. If $v(y)$ is the conditional probability of including y on this common unseen-event experiment, uniform honesty gives

$$1 - \alpha \leq 1 - (1-\varepsilon)^n + (1-\varepsilon)^n v(y).$$

Letting $\varepsilon \downarrow 0$ yields $v(y) \geq 1-\alpha$ for every allowed y. This also holds at $n = 0$ without taking a limit. Returning to the actual distribution π,

$$\mathbb{P}_{\pi,a}\{y \in C(\mathcal{T}, j), E_j\} \geq (1-\alpha)q_j.$$

Apply the normalized-volume argument of Lemma 16 for real responses, integrate over $[-A, A]$ for bounded responses, and use true-label coverage plus the $M-1$ other-label inequalities for classification. Averaging by π_j gives exactly the asserted three lower bounds.

A universal attaining regression procedure returns the true singleton if observed, and otherwise returns the full response domain with probability $1-\alpha$, the empty set with probability α. Its coverage for task j is $1 - \alpha q_j \geq 1-\alpha$, and its cost is the claimed value without knowledge of any frequency. For classification, independently thin an observed singleton with retention probability $1-\alpha$, while on unobserved tasks use the same full-set/empty-set randomization. This has coverage exactly $1-\alpha$ for every task and expected size $(1-\alpha)(1+(M-1)q_j)$. Finally every task in a positive distribution with at least two tasks has $q_j > 0$ at every finite n, and thus the real-response output has positive infinite-length probability at every π, a. This proves the final assertion. □

C.4 THE QUERY POLYMATROID AND THE EXACT COSTS WITH QUERIES

For compactness write $F = F_{n,B}$ and $\mathcal{P} = \mathcal{P}_{n,B}$, the capacity function and feasible region of (7). All arguments in this section concern finite K unless stated otherwise; the countable case is treated in Appendix C.6. The main-text Lemma 18 is Lemma 18 below.

Lemma 18. *The set of query-probability vectors achievable by a response-independent randomized kernel $Q \mid U$ with $Q \subseteq U$, $|Q| \leq B$, is exactly $\mathcal{P}$. Every response-adaptive procedure has its query-probability vector in this same set at every fixed response configuration.*

Proof. For a fixed unseen set u, the function $S \mapsto \min(B, |u \cap S|)$ is normalized and monotone. Its increment when an element is added is either zero or one and can only decrease as the original set grows. It is therefore submodular. Averaging over U proves these properties for F.

For any policy, pathwise

$$|Q \cap S| \leq \min(B, |U \cap S|).$$

Taking expectations gives all constraints of $\mathcal{P}$, irrespective of how labels influenced the policy. Let $\mathcal{Z}$ be the region achieved by response-independent kernels. It is the linear image of a finite product of probability simplexes, so it is compact and convex. It contains zero and is downward closed: independently retaining queried coordinates with suitable probabilities lowers any chosen marginal without violating the hard budget. We already have $\mathcal{Z} \subseteq \mathcal{P}$, and $\mathcal{P}$ is also compact because $0 \leq r_j \leq F(\{j\}) \leq 1$.

Fix any real weight vector w. Coordinates with nonpositive weight can be set to zero in either region. Relabel the positive coordinates so that $w_1 \geq \cdots \geq w_m > 0$, put $w_{m+1} = 0$, and let $S_k = \{1, \ldots, k\}$. For every $r \in \mathcal{P}$,

$$w^\top r \leq \sum_{k=1}^{m} (w_k - w_{k+1}) r(S_k) \leq \sum_{k=1}^{m} (w_k - w_{k+1}) F(S_k).$$

The first inequality discards nonpositive coordinates. Query, for each realized U, the first B available positive-weight coordinates in this order (or all if there are fewer). For each prefix S_k, its queried count is exactly $\min(B, |U \cap S_k|)$. Its expected marginal vector lies in $\mathcal{Z}$ and attains the displayed upper bound. Thus $\mathcal{Z}$ and $\mathcal{P}$ have identical support functions for every w. If a point of $\mathcal{P}$ lay outside the compact convex set $\mathcal{Z}$, finite-dimensional strict separation would give a weight with different support values, a contradiction. Hence the regions are equal. □

Proof of Theorem 3. Fix any configuration a and any honest adaptive policy. The passive unseen probability is independent of responses, and a query is made only when a task is unseen. Thus

$$u_j(a) = \mathbb{P}_a(E_j) = q_j - r_j(a).$$

Lemma 15 gives (9). By Lemma 18, $r(a) \in \mathcal{P}$. Apply Lemma 16, average over j, and minimize over $\mathcal{P}$ to obtain (10). This argument holds at every configuration even though the actual vector $r(a)$ may differ between configurations.

The objective in (8) is continuous on a nonempty compact region, so it has a minimizer r^*. Lemma 18 implements r^* by one response-independent query kernel. Now $u_j = q_j - r_j^*$ is independent of all responses and known to the designer. Repeat the attaining rules in Appendix C.2, replacing q_j by u_j and "observed" by "resolved." In particular, if $u_j = 0$ the task is always resolved; regression uses its singleton, and classification independently retains that singleton with probability $1 - \alpha$. For $u_j > 0$ all other randomization probabilities are exactly those in the passive proof. This proves equality simultaneously at every configuration for each criterion. Taking a supremum over configurations and then an infimum over honest rules leaves the same constants, proving the minimax interpretation. The nonempty bound (11) is proved in Appendix D.1. □

Example 19. Take $n = 10$, $B = 1$, $\alpha = 0.17$, and $\pi = (0.05, 0.05, 0.225, 0.225, 0.225, 0.225)$. The four frequent tasks are unobserved with probability $0.775^{10} \approx 0.078 < \alpha$, so they have zero demand. The two rare tasks each have demand $0.95^{10} - 0.17 \approx 0.429$, for a total demand of 0.857, below the total capacity $F_{n,1}(\mathcal{J}) \approx 0.914$; a total-budget calculation would predict zero cost. But the two rare tasks are unobserved together with probability $0.9^{10} \approx 0.349$, and the single query can serve only one of them on that event, so the capacity of the pair is $2(0.95^{10}) - 0.9^{10} \approx 0.849 < 0.857$ and the exact minimum cost is positive.

Corollary 20. *Call $c_j = (q_j - \alpha)_+$ the* demand *of task j: the query probability it needs before its unresolved probability drops to α. Then $\Phi_{n,B}(\pi, \alpha) = 0$ if and only if*

$$\sum_{j \in S} c_j \le F_{n,B}(S) \qquad \text{for every } S \subseteq \mathcal{J}, \tag{33}$$

that is, if and only if every subset's total demand is within its capacity.

Proof of Corollary 20. Let $c_j = (q_j - \alpha)_+$. Since all $\pi_j > 0$, the minimum is zero exactly when some $r \in \mathcal{P}$ satisfies $r_j \ge c_j$ for every j. The region is downward closed, so this is equivalent to $c \in \mathcal{P}$, precisely (33). If the condition fails, $\Phi_{n,B} > 0$; Theorem 3 gives positive infinite-length probability at every configuration, ruling out finite expected length. If it holds, the attaining regression construction never outputs the full real line, and uses only singletons or empty sets.

For Example 19, only the first two coordinates have positive demand. Their individual unseen probabilities equal 0.95^{10}, and the probability both are unseen is 0.9^{10}. Thus the subset capacity at $B = 1$ and its demand are

$$F_{n,B}(\{1,2\}) = 2(0.95^{10}) - 0.9^{10}, \qquad c(\{1,2\}) = 2(0.95^{10} - 0.17),$$

and the difference is $0.9^{10} - 0.34 > 0$. Both tasks have weight 0.05, so the exact residual is

$$\Phi_{n,B} = 0.05\,(0.9^{10} - 0.34) \approx 4.34 \times 10^{-4}.$$

Feasibility of allocating the whole subset capacity to these two coordinates follows from the query-first-among-these-two policy and mixtures; no other coordinate has useful positive demand. Total capacity is directly evaluated via exact finite-state calculation. □

C.5 AN EXACT WEIGHTED SUBSET-DEFICIT FORMULA

Recall from Theorem 11 the demands $c_j = (q_j - \alpha)_+$, the deficit $D(S) = \max_{T \subseteq S}\{c(T) - F(T)\}$, and $G(S) = c(S) - D(S)$. In this subsection T denotes a subset of tasks. The following formula identifies which subset obstructions determine a nonzero cost, and constructs optimal marginal query rates without optimizing a hinge-loss LP.

Proof of Theorem 11. Any feasible r can be replaced coordinatewise by $\min(r, c)$: feasibility is preserved, and the objective is unchanged. Therefore

$$\Phi_{n,B} = \pi^\top c - \max\{\pi^\top r : r \in \mathcal{P},\ 0 \le r \le c\}.$$

Observe that

$$G(S) = \min_{T \subseteq S}\{F(T) + c(S \setminus T)\} = c(S) - D(S).$$

The choice $T = \varnothing$ shows $G(S) \le c(S)$, and $T = S$ shows $G(S) \le F(S)$. A nonnegative vector obeys the constraints of $\mathcal{P}$ and $r \le c$ exactly when $r(S) \le G(S)$ for every S. Indeed, the forward direction follows by decomposing S into T and $S \setminus T$, and the reverse uses $G \le F$ and $G(\{j\}) \le c_j$.

We verify that G is a normalized, monotone submodular function. Normalization is immediate. For monotonicity consider $S \cup \{j\}$ and a minimizing set T. If $j \notin T$, its value is at least $G(S) + c_j$. If $j \in T$, remove j from T and use monotonicity of F to bound the value below by $G(S)$. For submodularity, choose minimizing sets $T_A \subseteq A$ and $T_B \subseteq B$. Submodularity of F and modularity of c imply

$$\begin{aligned} G(A) + G(B) &\ge F(T_A \cup T_B) + F(T_A \cap T_B) \\ &\quad + c((A \cup B) \setminus (T_A \cup T_B)) + c((A \cap B) \setminus (T_A \cap T_B)) \\ &\ge G(A \cup B) + G(A \cap B). \end{aligned}$$

Order tasks so that $\pi_1 \ge \cdots \ge \pi_K$ and put $S_k = \{1, \dots, k\}$. Define $r^*_k = G(S_k) - G(S_{k-1})$. Monotonicity makes these increments nonnegative. For any A, add its elements in this order. By diminishing returns, each increment of G along the full prefix is no greater than its increment when only previously added elements of A are present. Summing gives $r^*(A) \le G(A)$. Therefore r^* is feasible for the capped region. It saturates every prefix. With $\pi_{K+1} = 0$, summation by parts gives

$$\max \pi^\top r = \sum_{k=1}^{K} (\pi_k - \pi_{k+1}) G(S_k).$$

Likewise $\pi^\top c = \sum_k (\pi_k - \pi_{k+1}) c(S_k)$. Subtracting proves (22). Equal weights cause zero coefficients, so any ordering within a tie is valid. Lemma 18 turns these optimal marginals into a hard-budget kernel. □

C.6 EXACT ALGORITHMS AND LIMITING CASES

Unseen-set distribution and query kernel. Represent a seen subset by a bit mask $s \in \{0, \dots, 2^K - 1\}$. Start with $v_0(0) = 1$. For $t = 0, \dots, n-1$, update

$$v_{t+1}(s \cup \{j\}) \mathrel{+}= v_t(s)\pi_j \quad \text{for every } s, j.$$

Then $p_u = v_n(\mathcal{J} \setminus u)$. This recursion is exact in real arithmetic and sums all repeated visits. It takes $O(nK2^K)$ arithmetic operations. Compute $F(S) = \sum_u p_u \min(B, |u \cap S|)$ and solve either (8) with all rank inequalities, or the joint-mass LP (23). Both are finite but generally exponential in K. Formula (22) is another exact evaluation by subset enumeration. We make no efficiency claim for large task vocabularies; these routines exist to certify the frontier, not to compute it at scale.

Algorithm 2 is the complete prediction rule given an optimal kernel, with independent uniforms available for every random choice.

The predictor's random coins need not be shared between tasks. This is fixed-horizon per-task coverage, not a simultaneous guarantee. Querying a task whose response is already known is unnecessary and can always be omitted.

Algorithm 2 Exact frontier-attaining rule from an optimal query kernel

Require: optimal joint masses $z_{u,Q}$ solving (23); an independent uniform for every random choice
precompute $r_j \leftarrow \sum_{u,Q:j\in Q} z_{u,Q}$ and $u_j \leftarrow q_j - r_j$ for every j
observe H_n and determine U
draw Q with probability $z_{U,Q}/p_U$ and query its tasks ▷ zero-probability unseen sets are ignored
on request of the independent test task j:
if $u_j > \alpha$ **then** ▷ identical for regression and classification
  return $\{a_j\}$ if j is resolved
  otherwise **return** the full response domain with probability $1 - \alpha/u_j$, else $\varnothing$
else if regression **then** ▷ $u_j \le \alpha$
  return $\{a_j\}$ if j is resolved, else $\varnothing$
else ▷ $u_j \le \alpha$, classification; well defined since $u_j \le \alpha < 1$
  return $\varnothing$ if j is unresolved
  otherwise **return** $\{a_j\}$ with probability $(1-\alpha)/(1-u_j)$, else $\varnothing$

Proof of Proposition 12: cold start. At cold start $U = \mathcal{J}$ and $F(S) = \min(B, |S|)$. Hence $\mathcal{P} = \{r : 0 \le r_j \le 1,\ \sum_j r_j \le B\}$. Let $c = 1 - \alpha$. Capping r_j at c does not change the objective, so it is optimal to maximize $\sum_j \pi_j r_j$ under $0 \le r_j \le c$ and $\sum_j r_j \le B$. An exchange argument fills higher-weight coordinates before lower-weight ones. If $B < Kc$, this gives $r_j = c$ for $j \le m$, $r_{m+1} = \eta$, and zero thereafter; subtracting from c proves (24). If $B \ge Kc$, all coordinates can be filled. Although these are fractional marginal values, Lemma 18 implements them using a randomized subset of hard size at most the integer B. □

Countably many tasks: finite approximation and Corollary 17(ii). Define $\Phi_{\infty,B}$ as the infimum of $\sum_j \pi_j (q_j - r_j - \alpha)_+$ over all hard-B query kernels on the countable task set. This definition does not presume infinite-dimensional attainment. For finite $G \subset \mathcal{J}$, let Φ_G be the optimum of the finite problem restricted to querying tasks in G, with objective restricted to G. Passive identities outside G are a dummy category of mass $1 - \pi(G)$: probabilities in G must *not* be renormalized. The resulting unseen-set distribution on G is the one induced by the original experiment.

For any countable query policy, its marginal vector on G obeys all finite rank constraints on G. Lemma 18 therefore gives

$$\Phi_G \le \Phi_{\infty,B} \le \Phi_G + \sum_{j \notin G} \pi_j (q_j - \alpha)_+. \tag{34}$$

The upper inequality follows by using an optimal kernel on G and never querying outside it. The tail is at most $1 - \pi(G)$ and tends to zero along an exhaustion. Thus finite kernels attain the infimum arbitrarily closely. The three prediction costs have the corresponding exact infima, by the same avoidance lower bound and these approaching constructions. This completes the finite-approximation assertion in Proposition 12.

For infinite positive support, every finite n, B has $\Phi_{\infty,B} > 0$, which is part (ii) of Corollary 17. To see this, $c_j = (q_j - \alpha)_+$ is bounded below by $(1-\alpha)/2$ for infinitely many j. Choose a finite set S of these tasks with $c(S) > B$, and let $\pi_{\min,S} = \min_{j\in S} \pi_j > 0$. For any policy,

$$\sum_j \pi_j (c_j - r_j)_+ \ge \pi_{\min,S} \sum_{j\in S} (c_j - r_j)_+ \ge \pi_{\min,S} (c(S) - B) > 0.$$

The lower bound holds uniformly over all hard-B policies. Furthermore, under any adaptive query policy and for any fixed response configuration a, the marginals satisfy $\sum_{j\in S} r_j(a) \le B$. Since $c(S) > B$, there must exist at least one task $j \in S$ (which may depend on the configuration a) such that $r_j(a) < c_j$, meaning $u_j(a) = q_j - r_j(a) > \alpha$. Lemma 16 then guarantees that $\mathbb{E}_a \operatorname{Leb}(C(\mathcal{T}, j)) = \infty$ at this configuration in the unbounded domain, so $\sup_{j'\in\mathcal{J}} \mathbb{E}_a \operatorname{Leb}(C(\mathcal{T}, j')) = \infty$ for every configuration a. We emphasize the quantifier ordering $\forall a,\ \exists j :\ u_j(a) > \alpha$: one cannot swap these quantifiers to assert that a single fixed task remains unresolved across all configurations, because a response-adaptive policy could observe the first response and query that specific task index. Hence finite proactive acquisition does not eliminate infinite expected length on unrestricted infinite supports.

C.7 VERIFICATION AND KNOWN TASK GROUPS

Recall that verification means revealing the requested task's response after the request is known, and that Ψ_n is the passive cost.

Proof of Theorem 13. Condition on the requested identity being j, and let r_j be the probability that an as-yet unseen task j is verified. Then $u_j = q_j - r_j$ is the probability its value remains unresolved. Lemmas 15 and 16, applied after fixing the request, imply $\mathbb{P}_a\{\mathrm{Leb}(C(\mathcal{T}, j)) = \infty\} \geq (u_j - \alpha)_+$. Finite length almost surely for every task forces $r_j \geq (q_j - \alpha)_+$. If finite length is instead stated for the random request, positivity of every π_j implies the same taskwise conclusion. Total verification probability is at least $\sum_j \pi_j r_j \geq \Psi_n$; verifying already resolved tasks only increases cost.

For attainability, never verify resolved tasks. For an unresolved request j, verify with probability $(q_j - \alpha)_+/q_j$, defining this probability as zero if $q_j = 0$. Output the exact singleton whenever resolved and an empty set otherwise. Its taskwise coverage is $1 - q_j + (q_j - \alpha)_+ \geq 1 - \alpha$, all outputs have zero length, and its mean verification cost is Ψ_n. □

Proposition 21 (Known-group reduction)**.** *Suppose a known partition has $G < \infty$ positive-mass groups and responses $a_j = \theta_{g(j)}$ with unrestricted group responses. All exact passive, proactive, and on-demand frontiers apply with π replaced by the group masses ρ and with one query revealing one group response.*

Proof. The group identities $g(J_t)$ are i.i.d. with probabilities ρ_h. Observing any member gives the exact group response. A query to an already resolved group provides no new information and can be omitted. Any policy hence gives a set of at most B newly resolved groups; its group inclusion marginals obey the same rank constraints. Avoidance now concerns the whole group. For each requested member j of group h, varying θ_h leaves the transcript unchanged until that group is resolved, so Lemmas 15 and 16 apply to that member with the group's unresolved probability. Summing over members gives the weight ρ_h. Conversely a group-level policy is implemented by querying a fixed representative of each selected group. Use the group prediction for all its members. The conditional identity of a member inside a group is independent of its response; any extra randomization it induces cannot evade the lower bound. The same reasoning applies when the requested group is known before verification. □

C.8 GEOMETRIC VISIBILITY AND EXACT-RESPONSE RATES

We use open metric balls throughout. Let $\mathcal{H}(L, \beta)$ consist of real functions satisfying (25); their overall additive level is not bounded. The response domain is real. At every fixed x, the guarantee is

$$\inf_{f \in \mathcal{H}(L,\beta)} \mathbb{P}_f\{f(x) \in C_n(x)\} \geq 1 - \alpha.$$

Recall $R_n(x)$, the distance from x to the nearest observed descriptor, and $q_x(h)$, the probability that no observation lies within distance h of x.

Proof of Theorem 14. Fix x, h and, for $|t| \leq Lh^\beta$, set

$$f_t(z) = t\big(1 - (d(z, x)/h)^\beta\big)_+.$$

For $0 < \beta \leq 1$, $|u^\beta - v^\beta| \leq |u - v|^\beta$ for nonnegative u, v. The reverse triangle inequality and the fact that $v \mapsto (1 - v)_+$ is one-Lipschitz show

$$|f_t(z) - f_t(z')| \leq |t| h^{-\beta} d(z, z')^\beta \leq L d(z, z')^\beta.$$

Thus every f_t belongs to the class, with $f_t(x) = t$, and vanishes outside $B(x, h)$. On the event E_h that no historical descriptor falls in that ball, its observations are identical to those for $f = 0$. This event has probability $q_x(h)$ for all t. The same coupling and honesty argument as Lemma 15 gives

$$\mathbb{P}_0\{t \in C_n(x), E_h\} \geq (q_x(h) - \alpha)_+.$$

Integrating t over $[-Lh^\beta, Lh^\beta]$ proves (26) after taking a supremum over h.

Choose a nearest descriptor by the smallest-index tie rule. The Hölder bound implies $|f(x) - f(X_{i^*})| \le L R_n(x)^\beta$ for every sample, proving the stated upper construction. Under (27), the layer-cake identity and $\mathbb{P}(R_n(x) > h) \le \exp(-c_0 n h^d)$ yield

$$\mathbb{E}R_n(x)^\beta = \int_0^1 \beta h^{\beta-1} \mathbb{P}(R_n(x) > h)\, dh \le \Gamma(1+\beta/d)(c_0 n)^{-\beta/d}. \tag{35}$$

The distinction between strict and weak distance inequalities causes no problem: avoiding the closed ball is no more probable than avoiding the open ball. Enlarge C_0 to at least one and reduce c_0 to at most one if necessary. Take

$$h_n = \left(\frac{1-\alpha}{2C_0 n}\right)^{1/d} \le 1.$$

By the union bound, $q_x(h_n) \ge 1 - n\mu(B(x, h_n)) \ge (1+\alpha)/2$. Hence

$$\mathbb{E}_0 \operatorname{Leb}(C_n(x)) \ge L(1-\alpha)\left(\frac{1-\alpha}{2C_0 n}\right)^{\beta/d}. \tag{36}$$

This lower bound holds at every fixed x at the same constant function, and the upper bound holds uniformly over x, f. In particular, both the supremum-in-x minimax expected length and the μ-averaged minimax expected length have rate $Ln^{-\beta/d}$. This is the precise risk interpretation of the rate quoted in Section 7 (Theorem 14). □

C.9 Power-law asymptotics

Proposition 22 (Exact fixed-sample asymptotic)**.** *For $\pi_j = cj^{-s}$, $s > 1$, $c = 1/\zeta(s)$, fixed $\alpha \in (0,1)$, and integer $n \to \infty$, $\Psi_n \sim K_{s,\alpha} n^{1/s-1}$, with $K_{s,\alpha}$ defined in (28).*

Proof. Let $b_n = (nc)^{1/s}$ and $x_j = j/b_n$. Then $\pi_j = x_j^{-s}/n$, and

$$\Psi_n = \frac{b_n}{n}\frac{1}{b_n}\sum_{j\ge 1} x_j^{-s}\left((1 - x_j^{-s}/n)^n - \alpha\right)_+.$$

On grid points the base in parentheses is nonnegative. Put $a = \log(1/\alpha) > 0$ and $x_0 = a^{-1/s}$. Since $(1 - v/n)^n \le e^{-v}$ for $0 \le v \le n$, all summands with $x_j \le x_0$ vanish. On every compact subinterval of (x_0, ∞), the summand function converges uniformly to

$$g(x) = x^{-s}(e^{-x^{-s}} - \alpha)_+.$$

The same is true on any compact interval beginning at x_0, by continuity of positive part and uniform exponential convergence there. The Riemann sums on $[x_0, R]$ converge to the corresponding integral as $b_n \to \infty$. Uniformly for large n, the tail beyond $R > x_0$ is bounded by a constant times R^{1-s}: dominate the summands by x_j^{-s} and use the integral bound for the decreasing power function. The limiting integral has the same tail bound. Letting $R \to \infty$ proves

$$\frac{n}{b_n}\Psi_n \longrightarrow \int_{x_0}^\infty x^{-s}(e^{-x^{-s}} - \alpha)\, dx = \frac{1}{s}\int_0^a (e^{-u} - \alpha) u^{-1/s}\, du.$$

The integral is finite since $1/s < 1$, and strictly positive. Substituting $b_n/n = c^{1/s} n^{1/s-1}$ gives the claimed equivalent. This derivation uses the binomial absence probability throughout and does not replace the sample size by a Poisson variable. □

C.10 Gaussian observation noise: a complete interpolation theorem

This extension estimates the *latent task response* $f(x)$, not a new noisy outcome. The exact-response minimax equalities do not automatically apply to noisy observations. We now give a separate matching rate theorem with full assumptions; it establishes the noisy case (Theorem 23) discussed in Section 7.

Let $X_i \overset{\text{iid}}{\sim} \mu$ and

$$Y_i = f(X_i) + \sigma\varepsilon_i, \qquad \varepsilon_i \overset{\text{iid}}{\sim} N(0,1), \qquad 0 \le \sigma \le L,$$

where the noise is independent of the descriptors, $L > 0$, $f \in \mathcal{H}(L, \beta)$, and σ is known. Assume (27) with $0 < c_0 \le 1 \le C_0$, $d > 0$, and diameter at most one. Let $\mathcal{C}_\alpha$ be the class of procedures satisfying taskwise honesty for $f(x)$ at every fixed x. For a fixed x, define

$$\mathcal{W}_n(x) = \inf_{C \in \mathcal{C}_\alpha} \sup_{f \in \mathcal{H}(L,\beta)} \mathbb{E}_f \operatorname{Leb}(C_n(x)).$$

The same theorem applies if the risk is a supremum over x, or an average under μ.

Theorem 23 (Visibility–noise interpolation)**.** *For $n \ge 1$, $0 < \beta \le 1$, and fixed $0 < \alpha < 1$, there exist finite positive constants c_1, c_2 depending only on $\alpha, \beta, d, c_0, C_0$ such that*

$$c_1 A_n \le \mathcal{W}_n(x) \le c_2 A_n, \qquad A_n = L n^{-\beta/d} + L^{d/(2\beta+d)} \left(\frac{\sigma^2}{n}\right)^{\beta/(2\beta+d)}. \tag{37}$$

The lower bound holds at the constant function $f = 0$ for every honest rule. A k-nearest-neighbor interval attains the upper bound uniformly over the function class.

Proof. Visibility lower bound. Use the bumps from Appendix C.8. On the event of no sampled descriptor in their support, both the noiseless signals and the noisy observations (coupled by the same Gaussian noise) coincide with those at zero. The same proof gives (36), valid for every σ. This supplies a constant times $L n^{-\beta/d}$.

Noise lower bound. Suppose $\sigma > 0$. Choose a number $\delta \in (0, L]$, set $h = (\delta/L)^{1/\beta}$, and use the bumps

$$f_t(z) = t\big(1 - (d(z,x)/h)^\beta\big)_+, \qquad |t| \le \delta.$$

Let P_t and P_0 denote the full n-sample laws. Conditioning on descriptors and using the normal density ratio gives exactly

$$\mathrm{KL}(P_t \| P_0) = \frac{n}{2\sigma^2} \int f_t(z)^2 \, d\mu(z) \le \frac{n C_0 \delta^{2+d/\beta}}{2\sigma^2 L^{d/\beta}}.$$

Independent algorithmic randomization does not change this divergence. Choose

$$\delta = \left(\frac{(1-\alpha)^2 \sigma^2 L^{d/\beta}}{4 C_0 n}\right)^{\beta/(2\beta+d)}.$$

Since $\sigma \le L$, $n \ge 1$, and $C_0 \ge 1$, this is at most L. The KL bound is at most $(1-\alpha)^2/8$. Pinsker's inequality then gives $\mathrm{TV}(P_t, P_0) \le (1-\alpha)/4$. Hence honesty at f_t implies

$$\mathbb{P}_0\{t \in C_n(x)\} \ge \mathbb{P}_t\{t \in C_n(x)\} - \mathrm{TV}(P_t, P_0) \ge \tfrac{3}{4}(1-\alpha).$$

Integrating over $[-\delta, \delta]$ yields $\mathbb{E}_0 \operatorname{Leb}(C_n(x)) \ge \frac{3}{2}(1-\alpha)\delta$, a constant times the second term in (37). For completeness, the form of Pinsker used here follows by mapping to an event attaining total variation (or an approximating sequence), applying the log-sum inequality to the event and its complement, and using binary relative entropy $\mathrm{kl}(p, q) \ge 2(p-q)^2$. The log-sum inequality on an event follows by applying Jensen's inequality to $t \log t$ under the normalized restricted reference measure. The last inequality follows because its second derivative in p is $1/[p(1-p)] \ge 4$, and its value and first derivative at $p = q$ are zero. These arguments also prove the needed data-processing inequality directly for the event. Combining the two lower bounds uses $\max(a, b) \ge (a+b)/2$. At $\sigma = 0$ only the first term is nonzero.

A moment bound for the kth nearest neighbor. Let $R_{n,k}(x)$ be the kth smallest sample distance. If $h^d \ge 2k/(c_0 n)$, the expected number of descriptors in $B(x, h)$ is at least $2k$. For a binomial variable N of mean m, Markov's inequality on e^{-tN} and $1 + p(e^{-t} - 1) \le \exp(p(e^{-t} - 1))$ give

$$\mathbb{P}(N \le m/2) \le \exp\{m(t/2 + e^{-t} - 1)\} \le e^{-m/8}$$

by $t = \log 2$, since $(\log 2)/2 - 1/2 < -1/8$. Thus

$$\mathbb{P}(R_{n,k}(x) > h) \le e^{-c_0 n h^d/8} \quad \text{when } h \ge h_0 = (2k/(c_0 n))^{1/d}.$$

If $h_0 \ge 1$, simply use $R_{n,k} \le 1 \le h_0$. Otherwise integrate the tail below and above h_0, bounding the latter by its integral on $(0, \infty)$. In both cases,

$$\mathbb{E} R_{n,k}(x)^\beta \le \left[(2/c_0)^{\beta/d} + \Gamma(1 + \beta/d)(8/c_0)^{\beta/d}\right](k/n)^{\beta/d}. \tag{38}$$

Honest upper interval. Average the Y_i at the k nearest descriptors, calling the result $\overline{Y}_k$. Conditional on the descriptors, its noise is $N(0, \sigma^2/k)$ and its bias relative to $f(x)$ has magnitude at most $LR_{n,k}(x)^\beta$. Therefore

$$C_n(x) = \left[\overline{Y}_k - LR_{n,k}(x)^\beta - z_{1-\alpha/2}\frac{\sigma}{\sqrt{k}},\ \overline{Y}_k + LR_{n,k}(x)^\beta + z_{1-\alpha/2}\frac{\sigma}{\sqrt{k}}\right]$$

has coverage at least $1-\alpha$, even conditional on descriptors. Here $z_{1-\alpha/2}$ is the standard normal quantile. Its expected length is at most a constant times

$$L(k/n)^{\beta/d} + \sigma k^{-1/2}.$$

Put $t = (\sigma/L)^{2d/(2\beta+d)} n^{2\beta/(2\beta+d)}$ and take $k = \min\{n, \max\{1, \lceil t \rceil\}\}$. For $\sigma = 0$ take $k = 1$. When $t < 1$, $\sigma < Ln^{-\beta/d}$, so $k = 1$ gives a constant times the first term. When $t \geq 1$, our assumptions give $t \leq n$ and $t \leq k \leq 2t$. Substitution bounds both terms by a constant times $L^{d/(2\beta+d)}(\sigma^2/n)^{\beta/(2\beta+d)}$. This proves the upper bound. The constants are uniform in x, f; the lower construction at $f = 0$ holds for every x, so integration or taking the supremum preserves the conclusions. □

D Proofs of algorithmic and lifetime results

This appendix proves the stochastic optimization guarantee, the calibration and robustness bounds, and the lifetime results of Sections 4–5. Foundational proofs are in Appendix C. Throughout, "kernel" means a response-independent randomized query rule $Q \mid U$, and u_j denotes a task's unresolved probability under the kernel in question.

D.1 Nonempty prediction sets and oracle reductions

Recall that the nonempty classification bound (11) uses the threshold $\alpha_M = M\alpha/(M-1)$ in place of α.

Nonempty part of Theorem 3. Fix a response configuration and a task j, and let E_j be its unresolved event, with probability $\mathbb{P}_a(E_j) = u_j$, so that $\mathbb{P}_a(E_j^c) = 1 - u_j$. Every output is required to be a nonempty subset of $[M]$, which implies $|C| \geq 1$ almost surely.

On the resolved event E_j^c, this nonemptiness property immediately gives

$$\mathbb{E}_a[|C|\mathbf{1}_{E_j^c}] \geq 1 \cdot \mathbb{P}_a(E_j^c) = 1 - u_j.$$

On the unresolved event E_j, we obtain two distinct lower bounds:

(i) By nonemptiness, $|C|\mathbf{1}_{E_j} \geq \mathbf{1}_{E_j}$, which yields $\mathbb{E}_a[|C|\mathbf{1}_{E_j}] \geq \mathbb{P}_a(E_j) = u_j$.

(ii) Decomposing the set size into label indicators, $|C| = \sum_{y=1}^M \mathbf{1}\{y \in C\}$. Taking expectation on E_j yields $\mathbb{E}_a[|C|\mathbf{1}_{E_j}] = \sum_{y=1}^M \mathbb{P}_a(y \in C, E_j)$. By Lemma 15, to maintain $1-\alpha$ coverage under every possible counterfactual configuration without knowing a_j, the predictor must satisfy $\mathbb{P}_a(y \in C, E_j) \geq (u_j - \alpha)_+$ for each of the M candidates. Summing over all $y \in [M]$ gives $\mathbb{E}_a[|C|\mathbf{1}_{E_j}] \geq M(u_j - \alpha)_+$.

Combining both lower bounds gives

$$\mathbb{E}_a[|C|\mathbf{1}_{E_j}] \geq \max\{u_j, M(u_j - \alpha)_+\}.$$

Adding the expectations over the disjoint events E_j and E_j^c, we have

$$\begin{aligned}\mathbb{E}_a|C| &= \mathbb{E}_a[|C|\mathbf{1}_{E_j}] + \mathbb{E}_a[|C|\mathbf{1}_{E_j^c}] \\ &\geq \max\{u_j, M(u_j - \alpha)_+\} + 1 - u_j \\ &= 1 + \max\{0, M(u_j - \alpha)_+ - u_j\}.\end{aligned}$$

If $u_j \le \alpha$, $(u_j - \alpha)_+ = 0$, so the second term is $\max\{0, -u_j\} = 0$. If $u_j > \alpha$, we have $M(u_j - \alpha) - u_j = (M-1)u_j - M\alpha = (M-1)(u_j - \alpha_M)$ with $\alpha_M = M\alpha/(M-1)$. Thus in both cases,

$$\mathbb{E}_a|C| \ge 1 + [(M-1)u_j - M\alpha]_+ = 1 + (M-1)(u_j - \alpha_M)_+.$$

Lemma 18 and minimization over r prove (11). This lower bound holds separately at every a, including for label-adaptive queries.

For attainability, take a response-independent minimizing kernel. On resolved histories return the true singleton. On unresolved histories with $u_j > 0$, return $[M]$ with probability $t_j = (1 - \alpha_M/u_j)_+$; otherwise draw one label uniformly from $[M]$. For every response, miscoverage is exactly

$$u_j(1 - t_j)(1 - 1/M) = \min\{u_j, \alpha_M\}(1 - 1/M) \le \alpha.$$

Expected size is $1 + (M-1)u_j t_j = 1 + (M-1)(u_j - \alpha_M)_+$. At $u_j = 0$ no unresolved branch occurs. If $\alpha_M \ge 1$, $t_j = 0$ everywhere and uniform singleton guessing already meets the target. The construction neither assumes correct pretrained guesses nor discards empty predictions from the coverage calculation. □

Proposition 24 (Exact-oracle reduction)**.** *Suppose a weaker task experiment supplies observations from known probability kernels $K_j(a_j)$ whenever task j is encountered or queried. Arbitrarily many finite within-task observations and response-dependent stopping are allowed, provided they can be simulated from the revealed a_j and independent randomness. Task identities and the hard budget of newly queried tasks agree with the main model. Every lower bound in the exact-response experiment applies to honest inference for the latent responses a_j in this weaker experiment.*

Proof. Given any weaker-experiment algorithm, an exact-oracle learner reveals a_j on each encounter or query and generates the weaker observations with independent random seeds from $K_j(a_j)$. It follows the algorithm's decisions, including its stopping decisions, without disclosing the exact value to that algorithm. The simulated transcript, output, coverage, and cost have exactly the weaker experiment's law at every a. The query budget is unchanged. An algorithm violating the exact-oracle bound would therefore violate it after this simulation. This is a direct comparison of experiments, not an approximation. □

For example, $X \sim \nu$ and $Y = g_0(X) + a_j + \sigma\varepsilon$ with known ν, g_0, σ and noise law define such a kernel. Thus unrestricted task-specific intercepts reproduce the same information obstruction inside an ordinary statistical model. The reduction only transfers lower bounds: it does not claim that exact-response attaining sets stay optimal with noisy measurements, nor that a latent-response interval covers a fresh noisy Y.

D.2 The stochastic optimization guarantee

We prove a weighted version of Theorem 4, which is also used for the lifetime objective. Let $w_j \ge 0$, $\sum_j w_j = 1$, and $0 \le b \le 1$; the main text is the case $w = \pi$, $b = \alpha$. For a kernel g, write $u_j(g) = \mathbb{P}(j \in U \setminus Q_g)$ and define the objective and its optimum

$$\mathcal{L}_{w,b}(g) = \sum_j w_j(u_j(g) - b)_+, \qquad \Phi_{w,b} = \min_g \mathcal{L}_{w,b}(g).$$

(The symbol $\mathcal{L}$ is used so as not to collide with the capacity function $F_{n,B}$.) Zero weights may simply be omitted from the coordinate updates; tie-breaking still selects a feasible subset. All expectations here use the true initial unseen-set law.

Finite-dimensional minimax and the box dual. For $u \in [0,1]^K$,

$$\sum_j w_j(u_j - b)_+ = \max_{0 \le \lambda_j \le w_j} \lambda^\top(u - b\mathbf{1}).$$

The response-independent kernels form a finite product of probability simplices, indexed by the finitely many unseen sets and feasible query subsets. The box and this product are compact convex

sets, and the last display is bilinear in the box coordinate and kernel. Finite-dimensional minimax, equivalently linear-programming duality, gives

$$\Phi_{w,b} = \max_{0\le\lambda\le w} \min_{g} \left[\lambda^\top (q - b\mathbf{1}) - \mathbb{E}\sum_{j\in Q_g} \lambda_j\right] = \max_{0\le\lambda\le w} \left[\lambda^\top (q - b\mathbf{1}) - \mathbb{E}\,\mathrm{Top}_B(\lambda, U)\right].$$

One can also obtain the first equality by writing joint kernel masses as in (23), introducing positive-part slack variables, and taking its ordinary finite LP dual. The inner optimization is solved on every unseen set by its highest B nonnegative weights. This proves (12), part (i) of Theorem 4, without any assumption of independent coordinate availabilities.

A self-contained weighted Hedge bound. Run Algorithm 1 with w, b instead of π, α. Put $X_{t,j} = \mathbf{1}\{j \in U_t \setminus Q_t\}$ (task j left unresolved at iteration t), $d_{t,j} = X_{t,j} - b$, $p_{t,j} = \mathrm{sigmoid}(z_{t,j})$, and $\lambda_{t,j} = w_j p_{t,j}$. We first prove, for every realized sequence,

$$\max_{0\le\lambda\le w} \sum_{t=1}^{T} \lambda^\top d_t - \sum_{t=1}^{T} \lambda_t^\top d_t \le \frac{\log 2}{\eta} + \frac{\eta T}{8}. \tag{39}$$

For each coordinate there are two actions, receiving rewards 0 and $d_{t,j}$; their cumulative weights start at $1/2$ and update exponentially. If a random variable takes two values in an interval of length at most one, its log moment generating function minus its mean-linear term is at most $\eta^2/8$. Indeed, the second derivative of that log moment generating function is the variance under an exponential tilt, which is at most $1/4$; integrate twice from zero. Thus the log growth of each coordinate's weight normalizer is at most $\eta p_{t,j} d_{t,j} + \eta^2/8$. Its terminal value is at least $\eta \max\{0, \sum_t d_{t,j}\} - \log 2$. Combine these bounds, divide by η, multiply by w_j, and sum over j. Since $\sum_j w_j = 1$, this proves (39). The same proof covers negative $d_{t,j}$ because $|d_{t,j}| \le 1$.

Conditional on the past, the current weights are fixed before U_t is drawn. If g_t is their top-B kernel, then

$$\mathbb{E}[\lambda_t^\top d_t \mid \mathcal{F}_{t-1}] = \lambda_t^\top (q - b\mathbf{1}) - \mathbb{E}\,\mathrm{Top}_B(\lambda_t, U) \le \Phi_{w,b}.$$

Writing $\widehat{u} = T^{-1}\sum_t X_t$, taking expectation in (39), and optimizing $\eta = \sqrt{8\log 2/T}$ gives

$$\mathbb{E}\sum_j w_j(\widehat{u}_j - b)_+ \le \Phi_{w,b} + \sqrt{\frac{\log 2}{2T}}. \tag{40}$$

This is not yet a deployment bound, because $\widehat{u}$ is computed on the training episodes.

From the training path to a deployable mixture. Let $\bar{u}_j = T^{-1}\sum_{t=1}^{T} u_j(g_t)$. Denoting the error increment $\Delta_{t,j} = X_{t,j} - u_j(g_t)$, the sequence $(\Delta_{t,j})_{t=1}^{T}$ forms a martingale difference sequence with respect to the algorithmic filtration $\mathcal{F}_{t-1}$. Specifically, since $\mathbb{E}[X_{t,j} \mid \mathcal{F}_{t-1}] = u_j(g_t)$, we have $\mathbb{E}[\Delta_{t,j} \mid \mathcal{F}_{t-1}] = 0$. By the tower property of conditional expectation, the cross terms vanish for all $s < t$:

$$\mathbb{E}[\Delta_{s,j}\Delta_{t,j}] = \mathbb{E}\big[\mathbb{E}[\Delta_{s,j}\Delta_{t,j} \mid \mathcal{F}_{t-1}]\big] = \mathbb{E}\big[\Delta_{s,j}\underbrace{\mathbb{E}[\Delta_{t,j} \mid \mathcal{F}_{t-1}]}_{=0}\big] = 0.$$

Furthermore, conditional on $\mathcal{F}_{t-1}$, $X_{t,j} \in \{0,1\}$ is a Bernoulli indicator with parameter $p = u_j(g_t)$, so its conditional variance satisfies $\mathbb{E}[\Delta_{t,j}^2 \mid \mathcal{F}_{t-1}] = p(1-p) \le 1/4$ (maximized at $p = 1/2$). Expanding the mean squared error and applying Jensen's inequality yields

$$\mathbb{E}|\widehat{u}_j - \bar{u}_j| \le \sqrt{\mathbb{E}(\widehat{u}_j - \bar{u}_j)^2} = \sqrt{\frac{1}{T^2}\sum_{t=1}^{T}\mathbb{E}[\Delta_{t,j}^2]} \le \sqrt{\frac{T\cdot(1/4)}{T^2}} = \frac{1}{2\sqrt{T}}.$$

The positive-part function is 1-Lipschitz. Weighting and combining with (40) gives

$$\mathbb{E}\sum_j w_j(\bar{u}_j - b)_+ \le \Phi_{w,b} + \frac{\sqrt{\log 2/2} + 1/2}{\sqrt{T}}.$$

Drawing a uniform training iterate at deployment implements exactly the marginals $\bar{u}$. To store only R iterates, independently draw $I_1, \dots, I_R$ uniformly with replacement from $[T]$, independently of the complete training path. These indices can be drawn before training and the corresponding weights saved when visited. Conditional on that path, each $u_j(g_{I_r})$ lies in $[0,1]$, so the same variance argument gives expected absolute error at most $1/(2\sqrt{R})$. The uniform mixture on these R snapshots therefore proves (13), part (ii) of Theorem 4. Every component, and hence every mixture realization, obeys the hard budget.

For unnormalized nonnegative weights w, divide them by $W = \sum_j w_j$ during training; the objective error bound is then $We_{T,R}$. In particular this applies to the lifetime weights $w_j(H)$ of (17). Changing the weights changes this finite-iteration error, even though the exact optimum is common to all horizons, as proved in Appendix D.5. The result uses independent identity episodes from the true law, not repeated uncorrected reuse of a finite dataset.

D.3 Independent audit: coverage and efficiency

Recall the audit quantities of (14): $\widehat{u}_j$ is the empirical unresolved frequency over m audit episodes, $\epsilon_m = \sqrt{\log(2K/\delta)/(2m)}$, and $\overline{u}_j = \min\{1, \widehat{u}_j + \epsilon_m\}$. The complete procedure for auditing and deploying the certified predictor is summarized in Algorithm 3.

Algorithm 3 Independent calibration and taskwise coverage guard

Require: frozen policy g from Algorithm 1, audit size m, level α, slack $0 < \delta < \alpha$

Audit phase (offline; m identity-only episodes from the true π, independent of g)

for $i = 1, \dots, m$ **do**
  simulate a fresh history, run the frozen g, and record which tasks remain unresolved
form $\widehat{u}_j$, ϵ_m, and $\overline{u}_j$ as in (14) ▷ $\overline{u}_j \geq u_j$ for all j w.p. $\geq 1-\delta$

Deployment phase

run g on the real history H_n; cache every observed and every queried response
on request j:
if j is resolved **then**
  return $\{a_j\}$ ▷ exact cached response
else if taskwise coverage guard (nominal set predictor $C^{\mathrm{nom}}(j)$) **then**
  $t_j \leftarrow \left(1 - (\alpha - \delta)/\overline{u}_j\right)_+$
  return full domain $\mathcal{A}$ w.p. t_j; otherwise nominal set $C^{\mathrm{nom}}(j)$ ▷ yields Guarded CP, Guarded Top-1, etc.
else if regression (minimal length fallback) **then**
  $t_j \leftarrow \left(1 - (\alpha - \delta)/\overline{u}_j\right)_+$
  return the full domain $\mathcal{A}$ with probability t_j; otherwise $\emptyset$
else ▷ nonempty M-class, $\kappa = M/(M-1)$, uniform singleton fallback
  $t_j \leftarrow \left(1 - \kappa(\alpha - \delta)/\overline{u}_j\right)_+$
  return $[M]$ with probability t_j; otherwise one uniformly drawn singleton

Proof of Theorem 5. Condition on the trained response-independent mixture, including snapshot randomization. It is independent of the m subsequent audit episodes. Each coordinate's m unresolved indicators are i.i.d. Bernoulli with mean u_j. Hoeffding's inequality and a union bound give an event G of probability at least $1-\delta$ on which

$$|\widehat{u}_j - u_j| \leq \epsilon_m \quad \text{for every } j, \qquad u_j \leq \overline{u}_j \leq \min\{1, u_j + 2\epsilon_m\}.$$

For completeness, the required Hoeffding inequality follows by applying the log-mgf bound proved in Appendix D.2 to the sum of centered Bernoulli variables, using Markov's inequality, and optimizing the exponential parameter; the resulting one-sided tail is $e^{-2m\epsilon_m^2}$. There are $2K$ tails.

Put $a_0 = \alpha - \delta$. On G, the regression rule has miscoverage no larger than

$$u_j(1 - t_j) = u_j \min\{1, a_0/\overline{u}_j\} \leq a_0.$$

The branch $\overline{u}_j = 0$ has $u_j = 0$ on G and is immaterial. On G^c miscoverage is at most one. Averaging over the audit yields at most $a_0 + \delta = \alpha$ for every configuration and task. The proof conditions on the trained kernel but not on the realized audit in the final guarantee.

On G, its normalized regression cost $\widetilde{V}_a$ of (4) equals $\sum_j \pi_j u_j t_j$ in both the bounded and the unbounded domain. Because $u_j \le \overline{u}_j$,

$$u_j t_j \le (\overline{u}_j - \alpha + \delta)_+ \le (u_j - \alpha)_+ + 2\epsilon_m + \delta.$$

On G^c the normalized cost is at most one. Consequently its unconditional cost is at most $\sum_j \pi_j (u_j - \alpha)_+ + 2\epsilon_m + 2\delta$. Apply Theorem 4 to complete (15).

For nonempty classification, put $\kappa = M/(M-1)$ and use $t_j = (1 - \kappa a_0/\overline{u}_j)_+$. A uniform singleton fallback gives error $u_j(1-t_j)/\kappa \le a_0$ on G, so the same argument proves honesty. Its normalized excess size is again $\sum_j \pi_j u_j t_j$. On G it is at most

$$\sum_j \pi_j (u_j - \kappa\alpha)_+ + 2\epsilon_m + \kappa\delta.$$

Allowing G^c adds δ and proves the stated bound. When $\kappa\alpha \ge 1$, simply use uniform singletons on unresolved tasks, without training or auditing. Resolved tasks use exact singleton labels. □

The taskwise coverage guard: wrapping arbitrary nominal predictors. The construction of Theorem 5 is not limited to empty-set or uniform-singleton fallbacks; it constitutes a generic *taskwise coverage guard*. Let $C^{\mathrm{nom}}(j)$ be *any* nominal set predictor—such as cached top-1, pooled conformal prediction, adaptive prediction sets (APS), regularized APS (RAPS), or a heuristic neural set. When task j is unresolved, returning the full response domain $\mathcal{A}$ with probability $t_j = (1 - (\alpha - \delta)/\overline{u}_j)_+$ and falling back to $C^{\mathrm{nom}}(j)$ with probability $1 - t_j$ guarantees taskwise validity:

$$\mathbb{P}_a(a_j \in C(j)) \ge \mathbb{P}(\text{return } \mathcal{A}) + \mathbb{P}(a_j \in C^{\mathrm{nom}}(j) \mid \text{fallback}) \cdot \mathbb{P}(\text{fallback}) \ge t_j \cdot 1 + 0 = 1 - \frac{\alpha - \delta}{\overline{u}_j}.$$

On the audit event $G = \{\forall j,\ \overline{u}_j \ge u_j\}$, taskwise miscoverage is bounded by $u_j(1-t_j) \le \alpha - \delta$. Including the audit failure probability δ, taskwise miscoverage is at most $(\alpha - \delta) + \delta = \alpha$, establishing finite-sample distribution-free honesty *regardless* of whether C^{nom} satisfies any marginal or conditional coverage guarantees. Instantiating C^{nom} with different nominal predictors defines an entire family of guarded architectures:

$$\text{Guarded Top-1}, \quad \text{Guarded CP}, \quad \text{Guarded APS/RAPS}, \quad \text{Guarded Neural Sets}, \quad \ldots$$

In particular, instantiating the guard with pooled conformal prediction yields *Guarded CP*: validity is certified entirely by the guard's fallback t_j, while conformal sets minimize empirical set size on predictable head tasks.

For lifetime regression, draw a single guarded full/fallback decision at time zero for each initially unresolved task and keep it until revelation; then output the exact response. The probability of ever excluding its response is bounded by the same $u_j(1-t_j)$. With uniform-singleton classification, also retain the same guessed label until revelation. Thus the coverage proof extends directly to lifetime honesty. The cumulative efficiency bound is weighted by $w_j(H)$ and has additional audit error at most $(2\epsilon_m + 2\delta)\sum_j w_j(H)$ for regression, with the corresponding κ correction for nonempty classification. Independent audit data do not turn an estimated task law into an exact known law or promise zero vacuity at finite audit size.

D.4 Sensitivity to a misspecified simulator

This subsection is the formal version of Proposition 25. Recall $\mathcal{L}^P_{w,b}(g)$ denotes the objective of Appendix D.2 when the unseen set is drawn from the law P.

Proposition 25 (Efficiency transfer and true-law certification)**.** *Let P, Q be two laws of the initial unseen set, and let $w, \widetilde{w}$ be normalized nonnegative objective weights on the same finite vocabulary. Fix $0 \le b \le 1$. For every response-independent hard-budget kernel g,*

$$|\mathcal{L}^P_{w,b}(g) - \mathcal{L}^Q_{\widetilde{w},b}(g)| \le d(P, Q, w, \widetilde{w}) := \mathrm{TV}(P, Q) + \|w - \widetilde{w}\|_1.$$

An expected-e approximate minimizer under $(Q, \widetilde{w})$ has true expected objective at most $\Phi^P_{w,b} + e + 2d(P, Q, w, \widetilde{w})$. Independently auditing this fixed kernel under P gives the honesty conclusions of Theorem 5 regardless of the simulator used in training. Auditing only under Q does not generally do so. If P, Q are unseen-set laws from n i.i.d. draws from $\pi, \widetilde{\pi}$, then $\mathrm{TV}(P, Q) \le \min\{1, n\,\mathrm{TV}(\pi, \widetilde{\pi})\}$.

Proof. For each j, the kernel's conditional unresolved probability given U is a function in $[0,1]$. Its expectations under P, Q differ by at most $\mathrm{TV}(P,Q)$. Apply the 1-Lipschitz property of positive part, sum against w, and then change the weights, noting that every positive part is in $[0,1]$. Apply this uniform objective comparison first to the learned kernel and then to a true optimal kernel to obtain the excess-cost bound. The independent-audit proof involves only its own actual unresolved probabilities, so it is unchanged by training misspecification. Finally couple each pair of identity draws so that they disagree with probability $\mathrm{TV}(\pi,\widetilde{\pi})$; by a union bound the n-draw sequences, and hence their unseen sets, disagree with probability at most $n\,\mathrm{TV}(\pi,\widetilde{\pi})$. Equivalently this follows from product-law total variation and deterministic data processing. □

The proposition distinguishes simulator mismatch from lack of a true-law audit. It does not promise useful efficiency when the discrepancy bound is large. It also does not replace the unknown-environment impossibility in Theorem 10, which assumes only the specified passive history rather than an additional audit resource.

D.5 Lifetime honesty: complete proof and universal query policy

This subsection proves Theorem 6, Corollary 7, and Corollary 27. All probabilities in this section include the initial history, queries, future identities, and one complete independent algorithmic random seed. Predictions are adapted: at time t they may depend on the first t future encounters but not on future identities. Including the full independent algorithmic seed in the filtration is harmless. Responses are fixed, not redrawn each time a task recurs. Recall $v_j(H) = \sum_{t=0}^{H}(1-\pi_j)^t$ and $w_j(H) = \pi_j v_j(H)$ from (17).

A first-exclusion stopping argument.

Lemma 26. *Fix a, a task j, and an initially unresolved event E_j of probability u_j. Let $N_j(t)$ be the event that task j has not appeared among the first t future encounters. Under lifetime honesty, for every allowed candidate response y and finite H,*

$$\mathbb{E}_a \sum_{t=0}^{H} \mathbf{1}\{E_j, N_j(t), y \in C_t(j)\} \geq (u_j - \alpha)_+ v_j(H). \tag{41}$$

Proof. Define the stopping time

$$\tau_y = \inf\{t \geq 0 : E_j \text{ occurs},\ N_j(t) \text{ occurs},\ y \notin C_t(j)\},$$

with value ∞ when the set is empty. Couple a and $a^{(y)}$ by their histories, future identities, and random seeds. Until j is read their transcripts agree, exactly as in Lemma 15. Thus each event $\{\tau_y = t\}$, and their countable union, has the same probability under the two configurations. Under $a^{(y)}$, finite τ_y implies a lifetime coverage failure. Therefore

$$\mathbb{P}_a(\tau_y < \infty) \leq \alpha, \qquad \mathbb{P}_a(\tau_y < \infty) \leq u_j.$$

If exclusion occurs at a time with $E_j, N_j(t)$, some first exclusion has already occurred. Consequently

$$\sum_{t=0}^{H} \mathbf{1}\{E_j, N_j(t), y \notin C_t(j)\} \leq \sum_{s=0}^{H} \mathbf{1}\{\tau_y = s\} \sum_{t=s}^{H} \mathbf{1}\{N_j(t)\}.$$

Conditional on $\{\tau_y = s\}$ and the available filtration, no future identity has been seen and $N_j(s)$ holds. Independent geometric survival gives the conditional expected inner sum $v_j(H-s) \leq v_j(H)$. Hence the last display's expectation is at most $v_j(H)\min\{u_j, \alpha\}$. On the other hand, independence of future identities from initial acquisition gives $\mathbb{E}\sum_{t=0}^{H}\mathbf{1}\{E_j, N_j(t)\} = u_j v_j(H)$. Subtract the exclusion bound from this total to prove (41). Allowing later re-inclusion only makes the lower bound stronger; no monotonicity of the procedure's sets is assumed. □

All four lower bounds. For bounded regression, integrate (41) over $y \in [-A, A]$ and use Tonelli. For unbounded regression, put

$$Z_L = \sum_{t=0}^{H} \mathbf{1}\{E_j, N_j(t)\} \frac{\mathrm{Leb}(C_t(j) \cap [-L, L])}{2L}.$$

This is between zero and $H + 1$, and its expectation is at least $(u_j - \alpha)_+ v_j(H)$. On a realization where a particular $C_t(j)$ has finite length its summand tends to zero. All other normalized summands are at most one. Reverse Fatou, applicable to these uniformly bounded variables, gives

$$\sum_{t=0}^{H} \mathbb{P}_a\{E_j, N_j(t), \mathrm{Leb}(C_t(j)) = \infty\} \geq (u_j - \alpha)_+ v_j(H).$$

This is an event-probability conclusion, not merely an infinite expected length argument.

For M-class prediction with empty sets allowed, lifetime honesty implies $\mathbb{P}_a\{a_j \in C_t(j)\} \geq 1 - \alpha$ for each t. Sum these $H + 1$ inclusions and apply (41) to the other $M - 1$ labels. This yields

$$\mathbb{E}_a \sum_{t=0}^{H} |C_t(j)| \geq (H + 1)(1 - \alpha) + (M - 1)(u_j - \alpha)_+ v_j(H).$$

For nonempty sets, the expected unresolved cardinality sum is at least both $u_j v_j(H)$ and $M(u_j - \alpha)_+ v_j(H)$. At all other times it is at least $(H + 1) - u_j v_j(H)$. Adding gives

$$\mathbb{E}_a \sum_{t=0}^{H} |C_t(j)| \geq H + 1 + (M - 1)(u_j - \alpha_M)_+ v_j(H).$$

Multiply each task bound by π_j and sum. Lemma 18 gives $u_j = q_j - r_j(a)$ for a feasible $r(a)$, even under label-adaptive queries. Minimizing the appropriate weighted residual over that region gives all lower bounds in (18). They hold at every configuration, not just under a favorable prior on responses.

Attainment with one lifetime randomization. Fix any response-independent kernel and its unresolved probabilities u_j. The following policies are specified for all $t \geq 0$, not just for a reporting horizon.

For regression, initially resolved tasks always return their known singleton. If $u_j > \alpha$, an initially unresolved task independently chooses a good branch with probability $1 - \alpha/u_j$. On this branch return the whole response domain until its first revelation, then its true singleton forever. On the other branch return the empty set until revelation, then the true singleton. If $u_j \leq \alpha$, all initially unresolved histories use the latter branch. The probability of a lifetime failure is at most $\min\{u_j, \alpha\}$ for every response. The expected number of full outputs through H is exactly $(u_j - \alpha)_+ v_j(H)$. Regression empty sets can be replaced by any fixed singleton without changing the cost; their coverage can only increase.

For classification allowing empty sets, exact cardinality additionally spends the error budget after revelation. If $u_j > \alpha$, initially resolved histories always return their true singleton. Initially unresolved good histories (probability $u_j - \alpha$ unconditionally) return $[M]$ until revelation and the true singleton thereafter; the remaining histories (unconditional probability α) return the empty set forever, including after revelation. The true label is then included with probability exactly $1 - \alpha$ at every time, and each other label is included with probability $(u_j - \alpha)(1 - \pi_j)^t$. If $u_j \leq \alpha$, initially unresolved histories return the empty set forever; initially resolved histories choose the true singleton forever with probability $(1 - \alpha)/(1 - u_j)$, and otherwise the empty set forever. This is well defined because $u_j \leq \alpha < 1$ and again has true inclusion probability $1 - \alpha$. These constructions attain the empty-allowed formula. Their persistent abstention is an artifact of scoring by expected cardinality alone, not a recommended way to use a label that is already known.

For nonempty classification, initially resolved histories always output their true singleton. An initially unresolved task draws a good branch with probability $(1 - \alpha_M/u_j)_+$. It returns $[M]$ on this branch until revelation, or one *fixed uniformly chosen label* on the other branch; after revelation it returns the true singleton. Its lifetime error is exactly $\min\{u_j, \alpha_M\}(1 - 1/M) \leq \alpha$. Its excess expected size at time t is $(M - 1)(u_j - \alpha_M)_+(1 - \pi_j)^t$. This proves attainment for any kernel, once a common optimal kernel has been identified. In particular, no repeated checkpoint error spending is needed.

Why the same query kernel solves every horizon (Corollary 7). Fix a threshold b and write $c_j = (q_j - b)_+$. Let $F = F_{n,B}$ and, as in Theorem 11,

$$G(S) = \min_{T \subseteq S}\{F(T) + c(S \setminus T)\}, \qquad D(S) = c(S) - G(S).$$

Appendix C.5 proves that G is the rank of the capped polymatroid $\{r \in \mathcal{P}_{n,B} : r \le c\}$. Capping a feasible r at c preserves feasibility and its residual objective. On this capped set that objective is linear: $\sum_j w_j(H)(c_j - r_j)$.

Order tasks by decreasing π_j and put $S_k = \{1, \dots, k\}$. Since $x \mapsto 1 - (1-x)^{H+1}$ is increasing on $[0, 1]$, this same ordering sorts $w_j(H)$ for every integer $H \ge 0$. The greedy vector

$$r^*_k = G(S_k) - G(S_{k-1}) \tag{42}$$

is feasible and maximizes every linear objective whose nonnegative weights have that order. Indeed, summation by parts bounds such an objective by $\sum_k (w_k - w_{k+1})G(S_k)$, and the greedy vector saturates every prefix. Lemma 18 implements this one vector by one response-independent hard-budget kernel. The prediction rules just constructed depend only on $u_j = q_j - r^*_j$ and not on H. Together they attain every reported horizon, proving Theorem 6 and the horizon part of Corollary 7. Ties in π cause ties in all $w(H)$, so any consistent ordering suffices.

The resulting explicit value is

$$\Gamma_{n,B,H}(b) = \sum_{k=1}^{K} \big(w_k(H) - w_{k+1}(H)\big) D(S_k), \qquad w_{K+1}(H) = 0. \tag{43}$$

Thus the cumulative cost changes with the horizon even though a common exact optimizing kernel exists. Some other one-step optimal kernels may also work; the theorem specifically provides a capped greedy one. A stochastic approximation need not be an exact optimizer, so reweighting its finite training loss can improve lifetime approximation without contradicting horizon universality.

Corollary 27. *With no initial history and no queries ($n = B = 0$), the minimum cumulative regression cost of Theorem 6 equals*

$$(1-\alpha) \sum_j \big[1 - (1-\pi_j)^{H+1}\big] = (1-\alpha)\, \mathbb{E} D_{H+1}, \tag{44}$$

where D_{H+1} is the number of distinct tasks among $H+1$ independent requests. On an infinite support it diverges as $H \to \infty$, although the cost per checkpoint tends to zero.

Proof. At $n = B = 0$, $u_j = 1$ and no optimization remains, proving (44). The identity $\mathbb{E} D_N = \sum_j [1 - (1-\pi_j)^N]$ follows by writing distinct-task count as a sum of appearance indicators. It holds for countably infinite vocabularies by Tonelli. The lifetime construction above is also defined on that countable vocabulary, giving the same exact passive cumulative values for every finite H.

For finite K, $w_j(H) \uparrow 1$ and the optimal passive cumulative vacuity tends to $(1-\alpha)K$. For infinite positive support it diverges to infinity, by monotone convergence, even though the per-time passive cost tends to zero by dominated convergence. This is not an assertion that all tasks must be simultaneously covered with probability $1 - \alpha$: the guarantee remains separate for each task, simultaneous only over its times. □

Discounting and other nonincreasing time weights.

Corollary 28. *Fix a finite task vocabulary. Evaluate the weighted sum $\sum_{t=0}^{H} \omega_t L_t$ of the corresponding time-t costs L_t, with $\omega_0 \ge \cdots \ge \omega_H \ge 0$. The same kernel and lifetime prediction rule in Theorem 6 are optimal, with $w_j(H)$ replaced by*

$$w^{\omega}_j = \pi_j \sum_{t=0}^{H} \omega_t (1-\pi_j)^t$$

and each baseline $H+1$ replaced by $\sum_{t=0}^{H} \omega_t$. This also holds for a nonnegative summable nonincreasing infinite sequence. For geometric discount $\omega_t = \gamma^t$, $0 < \gamma < 1$, the exact regression residual is

$$\min_{r \in \mathcal{P}_{n,B}} \sum_j \frac{\pi_j}{1 - \gamma(1-\pi_j)} (q_j - r_j - \alpha)_+.$$

Proof. For any sequence of nonnegative time costs L_t, let $A_h = \sum_{t=0}^{h} L_t$. Finite summation by parts gives

$$\sum_{t=0}^{H} \omega_t L_t = \sum_{h=0}^{H-1} (\omega_h - \omega_{h+1}) A_h + \omega_H A_H.$$

Every coefficient is nonnegative. The theorem lower-bounds each A_h for every honest procedure, and a single procedure attains all these bounds. It therefore attains the weighted minimum as well. Substituting its explicit per-time costs gives w_j^{ω} and the stated baselines. In the infinite summable case $\omega_h \to 0$ and $\omega_t = \sum_{h \geq t} (\omega_h - \omega_{h+1})$. Tonelli gives the same decomposition with an infinite nonnegative sum, so the argument is unchanged. The geometric formula follows by summing its convergent geometric series. □

This extension requires nonincreasing time prices, and it completes the proof of Corollary 7. If arbitrary weights concentrate on a late checkpoint, the summation-by-parts argument and the task-weight ordering it relies on need not apply. The theorem does not claim a query policy universally optimal for every possible temporal objective.

D.6 A SHARP CHECKPOINT-VERSUS-LIFETIME COUNTEREXAMPLE

Proposition 29 (Checkpoint coverage can fail the lifetime target)**.** *Consider one task of frequency p, no initial history or queries, and checkpoints $t = 0, \ldots, H$. In regression (or unrestricted prediction), under the horizon condition $(1-p)^H \geq \alpha$, there is a rule that draws a single uniform random coin and, at every checkpoint, returns either the whole domain or the empty set according to that coin and current unresolved probability $(1-p)^t$, whose miscoverage at every single checkpoint is exactly α, but whose probability of at least one miscoverage over the $H+1$ checkpoints equals $\alpha(1+pH)$. For nonempty M-class classification, under the strictly stronger condition $(1-p)^H \geq \alpha_M := \frac{M\alpha}{M-1}$, the identical lifetime error $\alpha(1+pH)$ is achieved by the corresponding moving-threshold rule with fallback to a single fixed random label. The lifetime rule of Theorem 6 below miscovers at some checkpoint with probability at most α.*

Proof. Consider a single monitored task of probability $p \in (0,1)$ among the recurring identities, start without history or queries, and write $q_t = (1-p)^t$ for its unresolved probability at checkpoint t. Predictions are made after t further encounters and before the next one. Assume $q_H \geq \alpha$.

Choose one $Z \sim \mathrm{Unif}[0,1]$, independent of identities. At checkpoint t, an unresolved task outputs the empty set if $Z \leq \alpha/q_t$ and the whole response domain otherwise; resolved tasks use their exact response. This is the single-checkpoint optimal rule of Theorem 2 applied afresh at each checkpoint with the current q_t, sharing one coin. Its error probability at every fixed $t \leq H$ is exactly $q_t \alpha / q_t = \alpha$. Yet its probability of any error through H is

$$\alpha + \sum_{t=1}^{H} q_t \left(\frac{\alpha}{q_t} - \frac{\alpha}{q_{t-1}} \right) = \alpha(1+pH). \tag{45}$$

To prove the first equality, partition the uniform coin into the initial interval $[0, \alpha]$ and the successive disjoint intervals $(\alpha/q_{t-1}, \alpha/q_t]$. On the latter, the first possible error is exactly at t if the task has remained unresolved; before then the same coin always selects the full set. The intervals are within $[0,1]$ by the assumed horizon condition $(1-p)^H \geq \alpha$. Independence gives each summand. This also shows the expression never exceeds one under that condition.

For nonempty M-class sets use the same construction with threshold α_M and a single independent uniformly chosen fallback label shared at all times. On the event this fallback label is incorrect (probability $1 - 1/M$), the first-error calculation proceeds identically to the regression case, yielding cumulative error $\alpha_M(1+pH)$. Multiplying by $1 - 1/M$ recovers exactly $\alpha(1+pH)$. Crucially, this requires the intervals $(\alpha_M/q_{t-1}, \alpha_M/q_t]$ to remain within $[0,1]$, which is guaranteed if and only if $q_H \geq \alpha_M = \frac{M\alpha}{M-1}$. If $q_H \geq \alpha$ but $q_H < \alpha_M$ (for example, $M = 10, \alpha = 0.1, p = 0.1, H = 21$ yields $0.9^{21} \approx 0.1094 \geq 0.1$ but $0.9^{21} < 1/9$), the threshold exceeds 1 before horizon H and the exact formula cannot be directly applied. Thus empty sets are not responsible for the separation, but nonempty classification requires the tighter horizon condition. Lifetime-optimal sets instead choose

their branch once using the *initial* unresolved probability, without increasing the fallback probability as q_t shrinks. Their ever-error is at most α, by the attainment paragraph of Appendix D.5. This proves Proposition 29.

The experiment uses $p = 0.1$, $H = 20$, $\alpha = 0.1$, where $q_H \approx 0.12158$ exceeds both α and $\alpha_{10} = 1/9$. Equation (45) is exactly 0.3 in both versions. Independently refreshed checkpoint coins give an even higher ever-error in the executed simulation; no universal closed formula or optimality claim is needed for that comparison. □

D.7 Active heavy tails and cumulative lifetime growth

This subsection proves Theorem 30 (discussed in Section 7).

Theorem 30 (Uniform power-law query rate)**.** *Let $\pi_j = c_s j^{-s}$, $s > 1$, $c_s = 1/\zeta(s)$, and fix $0 < \alpha < 1$. On this countable vocabulary let $\Phi_{n,B}$ be the infimum defined in Appendix C.6. There are positive finite constants depending only on s, α such that, for all integers $n, B \geq 0$,*

$$\Phi_{n,B} \asymp \Big((n+1)^{1/s} + B\Big)^{1-s}. \tag{46}$$

For cold start without proactive queries, lifetime cumulative vacuity has the asymptotic

$$\Gamma_{0,0,H}(\alpha) \sim (1-\alpha)\Gamma(1-1/s)\big(c_s(H+1)\big)^{1/s}. \tag{47}$$

Proof. Set $a = (n+1)^{1/s}$. First ignore queries. For $n \geq 1$,

$$\Psi_n \leq \sum_j \pi_j e^{-n\pi_j} \leq C_s n^{1/s-1}.$$

Here the function $h(x) = c_s x^{-s} e^{-nc_s x^{-s}}$ on $(0,\infty)$ is nonnegative and unimodal. Its sum on positive integers is at most its integral plus twice its supremum: split the grid on either side of a mode and compare each monotone part to its adjacent integral. The integral equals $(c_s^{1/s}/s)\Gamma(1-1/s)n^{1/s-1}$ by substitution $u = nc_s x^{-s}$; the supremum is at most $1/(en)$. This proves the displayed bound, and $n = 0$ is immediate. Alternatively query every unseen task among the first B identities. This obeys the hard budget and leaves cost at most $(1-\alpha)\sum_{j>B}\pi_j \leq C_s(B+1)^{1-s}$. The minimum of these two upper bounds is at most a constant times $(a+B)^{1-s}$.

For the lower bound choose $m = \lceil C(a+B)\rceil$, where C is large enough, depending only on s, α, that $nc_s m^{-s} \leq (1-\alpha)/2$ and $m \geq 4B/(1-\alpha)$. For every $j \in \{m+1, \dots, 2m\}$, Bernoulli's inequality yields $q_j \geq 1 - n\pi_j \geq (1+\alpha)/2$ and hence $q_j - \alpha \geq (1-\alpha)/2$. Any hard-$B$ policy has $\sum_j r_j \leq B$, including at each fixed response configuration under adaptivity. Therefore

$$\begin{aligned}\sum_j \pi_j (q_j - r_j - \alpha)_+ &\geq c_s(2m)^{-s} \sum_{j=m+1}^{2m} (q_j - r_j - \alpha)_+ \\ &\geq c_s(2m)^{-s}\,(m(1-\alpha)/2 - B) \geq c_s 2^{-s}(1-\alpha)m^{1-s}/4.\end{aligned}$$

Since $a + B \geq 1$ and $m \leq (C+1)(a+B)$, this is a constant times the desired lower rate. The bound is uniform over policies and needs no claim of infinite-dimensional attainment.

For the second assertion put $N = H+1$ and $b_N = (c_s N)^{1/s}$. The occupancy identity gives

$$\frac{\mathbb{E}D_N}{b_N} = \frac{1}{b_N}\sum_{j\geq 1}\left[1 - \left(1 - \frac{(j/b_N)^{-s}}{N}\right)^N\right].$$

On every compact interval inside $(0,\infty)$ the summands converge uniformly to $1 - e^{-x^{-s}}$. For $x \leq \varepsilon$ they are at most one, giving normalized contribution at most $\varepsilon + 1/b_N$. For $x \geq R$ Bernoulli's inequality bounds them by x^{-s}, whose Riemann-sum tail is at most a constant times R^{1-s}. The middle Riemann sums therefore converge and the two tails can be made arbitrarily small. Substitution and integration by parts give

$$\int_0^\infty (1 - e^{-x^{-s}})\,dx = \frac{1}{s}\int_0^\infty (1-e^{-u})u^{-1-1/s}\,du = \Gamma(1-1/s).$$

Multiply by the exact lifetime factor $1-\alpha$ from (44) to prove (47). Occupancy asymptotics themselves are classical (Karlin, 1967); their role here is to evaluate the exact lifetime reliability price. □

The scale $B \asymp n^{1/s}$ marks when proactive labels change the order of the passive tail cost. This is an order statement for unrestricted responses, not a claim that a frequent-first policy attains the exact finite-sample optimum.

E EXPERIMENTAL SETTINGS AND NUMERICAL RESULTS

This section reports the full experimental protocols and results underlying the main-text comparisons of Section 6. There are three different evidential levels: exact finite-model optimization in floating-point arithmetic; independent simulated deployments; and experiments with trained classifiers on a public dataset. None substitutes for the proofs above. All reported confidence intervals are descriptive across the explicitly stated independent outer seeds; they do not convert a pointwise empirical minimum into a simultaneous statistical guarantee.

E.1 EXPERIMENTAL ENVIRONMENT AND EVALUATION PROTOCOLS

The numerical experiments were executed using Python 3.13.5, NumPy 2.3.5, SciPy 1.17.0, and scikit-learn 1.8.0 on Linux x86-64, with one OpenMP/OpenBLAS thread requested. The subsections below detail the task distributions, query budgets, neural architectures, and optimization hyperparameters for each numerical evaluation.

Unless otherwise stated, an interval for a reported mean is $\overline{x} \pm t_{0.975,S-1} s_x/\sqrt{S}$, computed across S outer seeds. This is not an exact finite-sample interval for arbitrary seed-output distributions; it is a standard seed-level summary. We never flatten the dependent tasks within a single history and treat them as independent draws when forming these intervals. Main-table intervals rounded to 0.00 are small but nonzero.

Optimization, audit, and deployment use separately specified random-seed families. The large-vocabulary baselines share the same distributions, budgets, and seed-indexed replication counts, but not an exact common-random-number coupling: different policy draws consume different amounts of randomness. The digit prediction rules, in contrast, are evaluated on the exact same history/query realization within each episode. No audit episode is used to train the query policy.

E.2 EXACT OPTIMIZATION, DUALITY, AND NONEMPTY-SET CHECKS

This subsection supports Section 6.2. The main small problem uses $n = 6$, $\alpha = 0.1$, and

$$\pi = (0.32,\ 0.23,\ 0.16,\ 0.11,\ 0.075,\ 0.055,\ 0.03,\ 0.02).$$

The unseen-set probabilities are calculated by the exact finite recursion from Appendix C.6. SciPy's HiGHS solves the finite LP (23) in double precision; this is an exact formulation with numerical solutions, not interval-certified arithmetic. Across budgets $B \in \{1, 2, 3, 4\}$, optimization lengths $T \in \{256, 1024, 4096, 16384\}$, each configuration uses 20 independent runs with $R = 1024$ randomly sampled iterates. Every deployed mixture's true residual marginals are evaluated by enumeration, so the plotted objective gap has no deployment Monte Carlo error. Figure 2 reports $B = 2$.

An independently formulated dual LP enumerates every feasible query subset inside each unseen set and constrains its support epigraph. Its value agrees with the rank-primal LP on 30 random instances ($K = 2, \ldots, 6$) to maximum observed discrepancy 1.67×10^{-16}. A further 90 LPs over all nonempty class subsets check the nonempty cardinality formula (11), using $M = 2, \ldots, 6$, $\alpha \in \{0.03, 0.1, 0.3\}$, and $u \in \{0, 0.05, 0.1, 0.2, 0.5, 1\}$.

Numerical LP comparisons verify the agreement between the primal rank formulation and the dual subset LP across random instances ($K = 2, \ldots, 6$), with maximum observed numerical discrepancy below 10^{-15}. An additional 90 LPs over nonempty class subsets check the cardinality formula (11) across $M = 2, \ldots, 6$, $\alpha \in \{0.03, 0.1, 0.3\}$, and $u \in \{0, 0.05, 0.1, 0.2, 0.5, 1\}$. A further 125 weighted LP checks confirm the nonincreasing-discount extension of Corollary 28 (maximum discrepancy 1.95×10^{-16}). These finite checks confirm numerical consistency; the formal theoretical guarantees are established by the mathematical proofs in the preceding appendices.

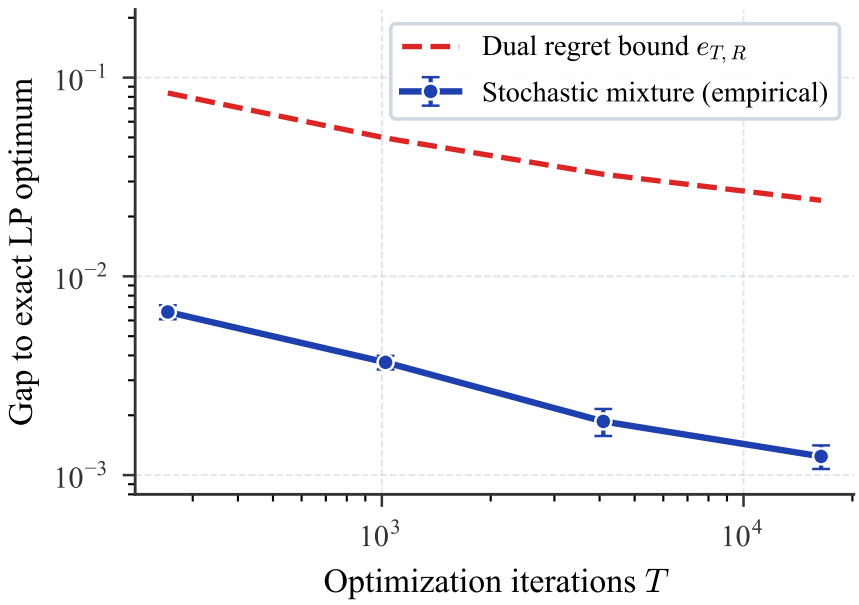

Figure 2: The sampled query policy approaches the exact finite optimum at $K = 8$, $n = 6$, and $B = 2$. The plotted cost is evaluated by enumeration, so only policy-training randomness contributes to the 95% seed-level intervals.

E.3 Vocabulary scaling and query-budget ablations

For $K \in \{64, 256, 1024, 4096\}$ use normalized finite Zipf probabilities $\pi_j \propto j^{-1.2}$, $n = K/2$, $B = K/16$, $\alpha = 0.1$, $T = 8192$, and $R = 1024$. The first three sizes have ten independent optimization/deployment seeds and 20,000 evaluation episodes per seed; $K = 4096$ has five seeds and 10,000 episodes per seed. The large-problem objective is calculated from independent deployment estimates of u, not an exact LP optimum. Training times measure query-policy optimization only: they exclude audits, baseline deployment evaluation, network training, and manuscript production.

Five query policies receive the same information and obey the same hard budget: the learned mixture; uniform subsets of unseen tasks; largest π_j first; smallest π_j first; and a fixed demand score $\pi_j((1-\pi_j)^n - \alpha)_+$. Scores are applied only to currently unseen tasks. When fewer than B tasks are available all available tasks can be queried. The core routine asserts the budget on every batch. The fixed-K curve uses $K = 256$, $n = 128$, budgets $0, 4, 8, 16, 32, 64, 128$, twelve seeds, $T = 4096$, $R = 1024$, and 15,000 independent evaluation episodes per seed.

Table 3: Scaling experiment: residual cost $\sum_j \pi_j(u_j - \alpha)_+$ after queries, and optimization-only time. Intervals are over independent seeds (ten, or five for $K = 4096$).

K	Stochastic mixture	Frequent-first	Uniform	Train seconds
64	0.16439 ± 0.00013	0.17350 ± 0.00012	0.22038 ± 0.00022	0.192 ± 0.028
256	0.12870 ± 0.00004	0.13609 ± 0.00004	0.18007 ± 0.00010	0.223 ± 0.007
1024	0.10092 ± 0.00004	0.10672 ± 0.00001	0.14732 ± 0.00002	0.503 ± 0.007
4096	0.07913 ± 0.00004	0.08365 ± 0.00001	0.12063 ± 0.00002	1.657 ± 0.050

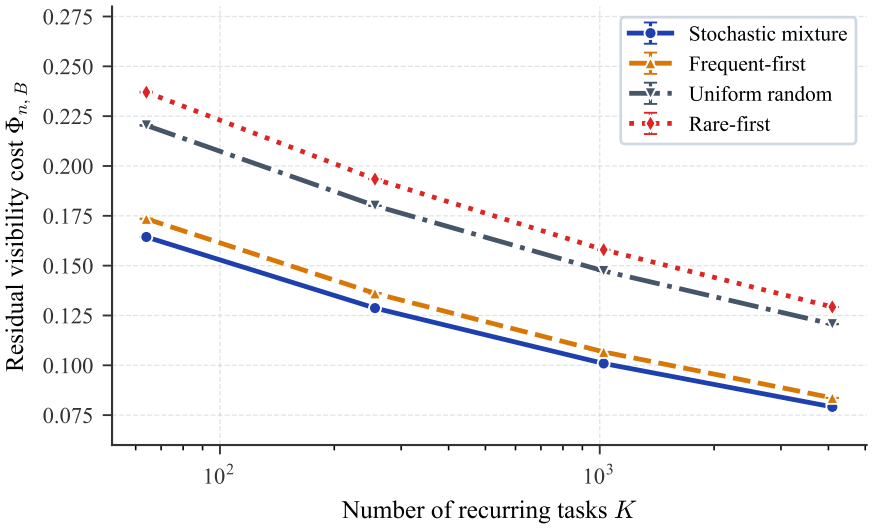

Figure 3: Independent large-vocabulary deployment. The lower cost of the learned mixture is an empirical efficiency result under the same hard budget, not a comparison of different coverage targets.

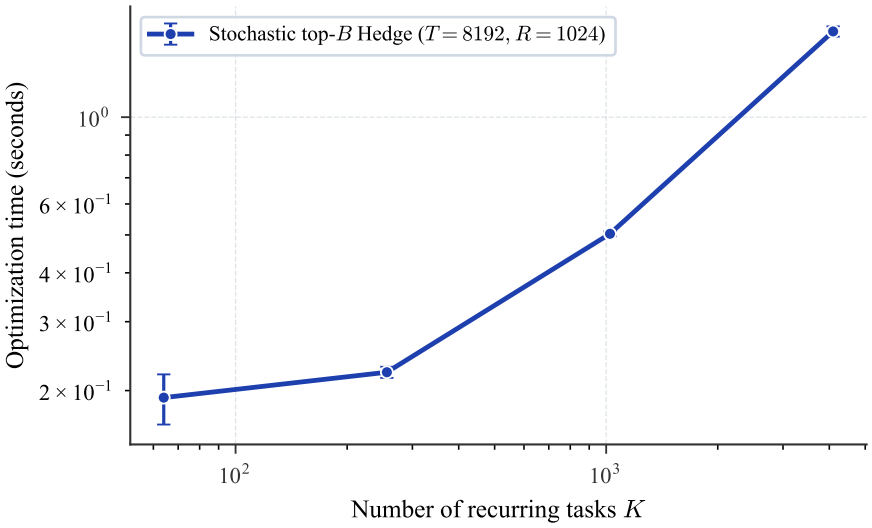

Figure 4: Optimization-only runtime for the scalable solver. Storage uses 1,024 snapshots; identity simulation and top-B selection are included.

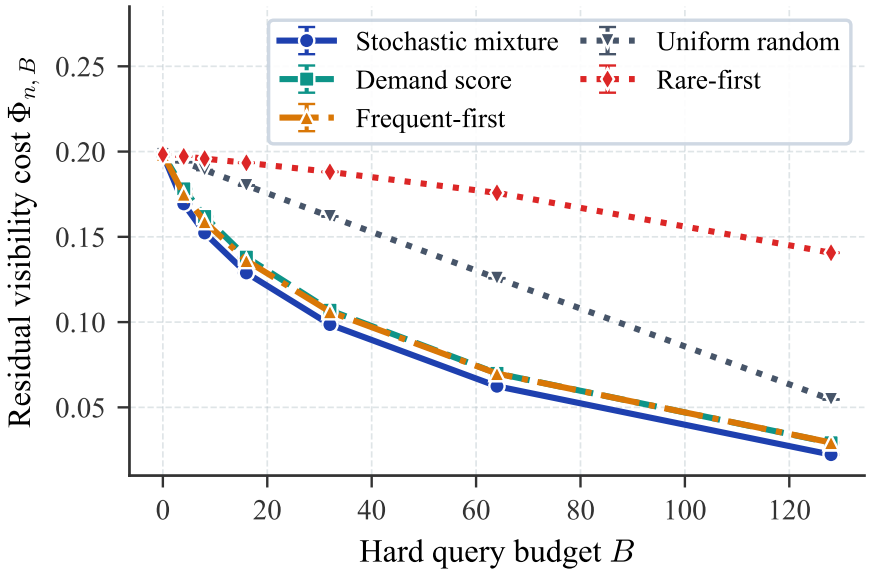

Figure 5: Independent deployment at $K = 256$, $n = 128$, across query budgets. All query methods obey the same hard budget; intervals summarize 12 independent seeds.

E.4 Calibration size and simulator misspecification

This subsection supports Section 6.2. For the $K = 8$, $n = 6$, $B = 2$ problem, train one fixed policy with $T = 8192$, $R = 2048$. Its true u is evaluated by enumeration. For each $m \in \{16, 64, 256, 1024, 4096, 16384\}$, generate 200 independent audits with $\delta = 0.005$. Compare the uncorrected plug-in $\widehat{u}$ and the certified upper bound $\overline{u}$ of (14). For each realized audit, evaluate subsequent coverage and normalized regression width using the *true* u. Thus Figure 6 isolates calibration-estimation error without further Monte Carlo error in the measured coverage.

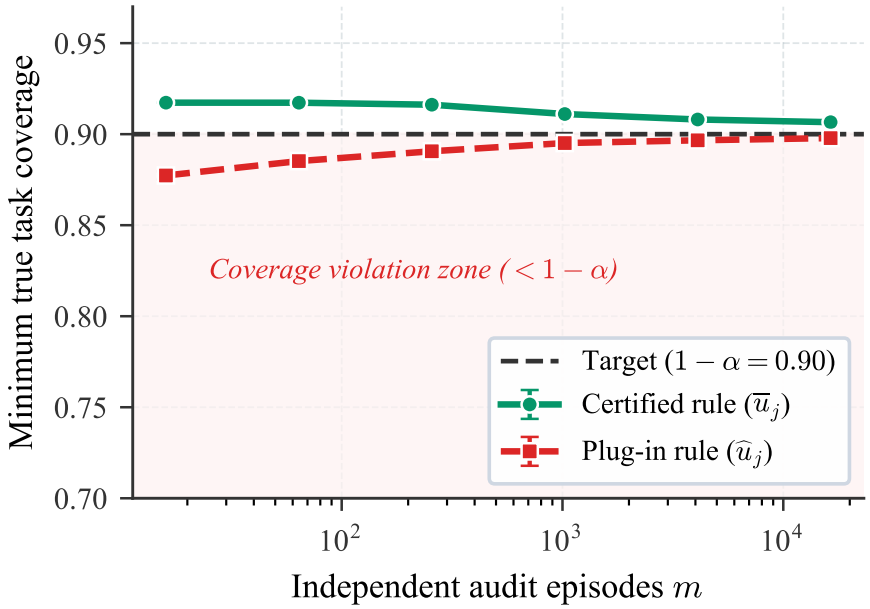

Figure 6: Minimum true task coverage for the independently calibrated and plug-in rules, conditional on each audit and then averaged across 200 audits. More data reduce calibration conservatism; the certified rule accounts for the audit failure event.

The plug-in rule's minimum coverage approaches 0.9 from below in this experiment. Its frequency of having at least one coordinate below 0.9 need not vanish at the same rate as the *magnitude* of the violation; we do not read the frequency alone as an effect size. The certified rule is wider, as shown in Figure 7. That extra width is not a disadvantage relative to the plug-in rule, which is invalid and therefore never competes at the same honest guarantee.

For misspecification, train using $\widetilde{\pi}_j \propto \pi_j^e$ with $e \in \{0.3, 1, 3\}$, ten seeds, $T = 8192$, $R = 1024$, and the same nominal $B = 2$, $n = 6$, $\alpha = 0.1$. Audit each policy either under this assumed law or the true law using 50,000 episodes and $\delta = 0.005$. Evaluate coverage exactly under the true law. At $e = 0.3$ the assumed-law audit yields minimum coverage 0.85534, compared with 0.90595 under true-law auditing. At $e = 3$ the assumed-law rule is instead conservative; mismatch does not always cause undercoverage. Figure 8 shows all three settings; Proposition 25 is the corresponding statement.

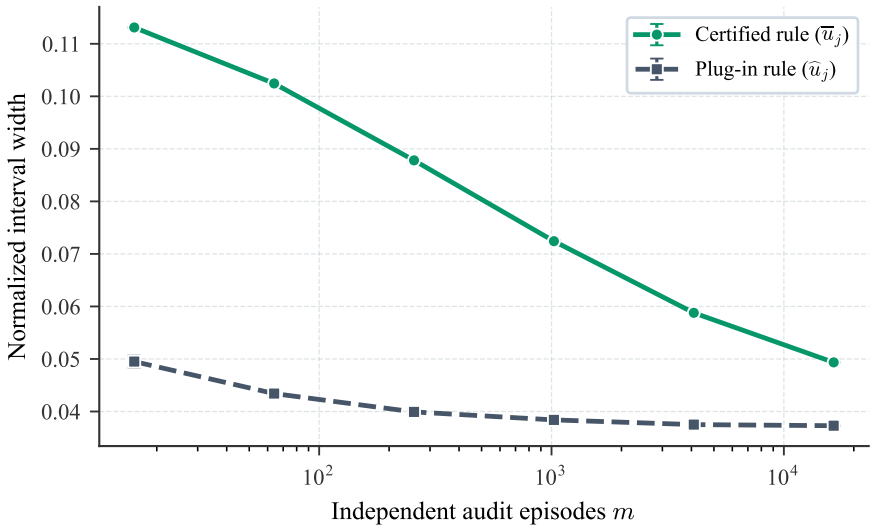

Figure 7: Efficiency price of the independent audit on the fixed finite model. Plug-in width is smaller but accompanies the undercoverage shown in the main figure.

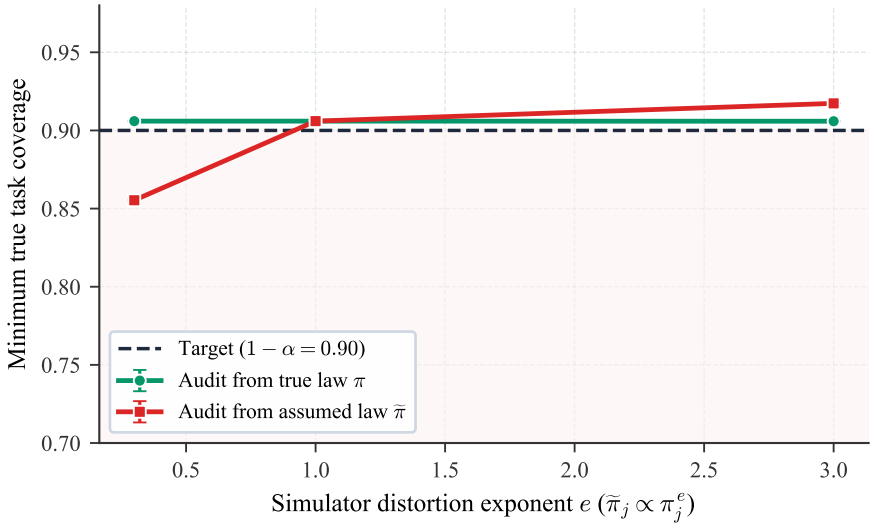

Figure 8: Auditing the training simulator is not the same as auditing the deployment law. True-law audit episodes preserve validity despite misspecified query-policy training.

E.5 Lifetime trajectories and common optimal query policy

This subsection supports Section 6.2. For the checkpoint experiment take $p = 0.1$, $H = 20$, $\alpha = 0.1$. Each of twenty outer seeds generates 100,000 first-revelation times from a geometric law. On these shared trajectories compare checkpoint thresholds with a shared coin (the rule of Proposition 29), checkpoint thresholds with fresh coins, and the lifetime threshold with a shared decision (the rule of Theorem 6). Record probability of at least one error, maximum checkpoint error frequency, and cumulative vacuity. Thus two million trajectories are generated and used across the three rules, not six million independent datasets. The respective mean ever-errors are 0.3003575, 0.622502, and 0.100183 (Figure 9).

A separate twenty-seed nonempty ten-class experiment also uses 100,000 trajectories per seed, with a fixed uniformly chosen fallback label. The shared checkpoint rule gives 0.2998725 ever-error; the lifetime rule gives 0.0998900. Mean cumulative set sizes are 80.20245 and 92.29038, respectively.

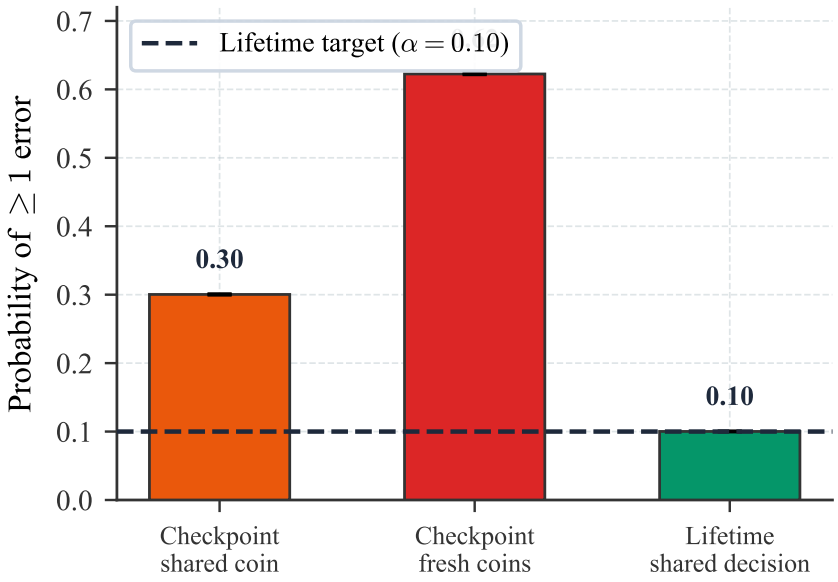

Figure 9: Probability of at least one error across 21 predictions, with nominal error 0.10. Policies that satisfy the checkpoint target can exceed the lifetime target substantially; sharing one initial full-set/fallback decision retains the correct lifetime budget. Intervals summarize 20 independent replications.

The larger lifetime size pays for a genuinely stronger guarantee; the checkpoint rule is not a valid cheaper lifetime competitor. The construction is exactly the one analyzed in Appendix D.6.

We check 48 flow LPs with initial unresolved probability one, and 180 further LPs with $u \in \{0, 0.03, 0.1, 0.4, 1\}$. These verify the first-exclusion budget formula of Lemma 26 across several probabilities, horizons, and error levels. For the general test, LP variables b_t represent unresolved probability mass first spending the error budget at time t. Constraints are

$$\sum_{s \leq t} b_s (1-p)^{t-s} \leq u(1-p)^t, \qquad \sum_t b_t \leq \alpha, \qquad b_t \geq 0.$$

Maximizing $\sum_t b_t v(H-t)$ subtracts the saved inclusions from $uv(H)$. The optimum agrees with $(u-\alpha)_+ v(H)$ in all tested cases. These LPs independently optimize stopping-budget allocation, rather than merely evaluating the proposed policy.

For the query policy, train mixtures under normalized $w(H)$ and separately under π, with the $K = 8$ model, $B = 1, 2, 3$, $H = 0, 1, 4, 16, 64, 256, 1024$, twenty seeds, $T = 8192$, $R = 1024$. Deployment is evaluated exactly. Figure 10 shows $B = 3$. Horizon-reweighting improves the finite-iteration approximation for large H. Separately, 350 random-instance LP comparisons test the same capped greedy marginal vector (42) at five different horizons; maximum observed discrepancy is 1.36×10^{-10}, below the 10^{-8} numerical tolerance. The exact optimum is therefore common across horizons, as Corollary 7 predicts, and what horizon reweighting changes is the finite-iteration approximation rather than the allocation being approximated.

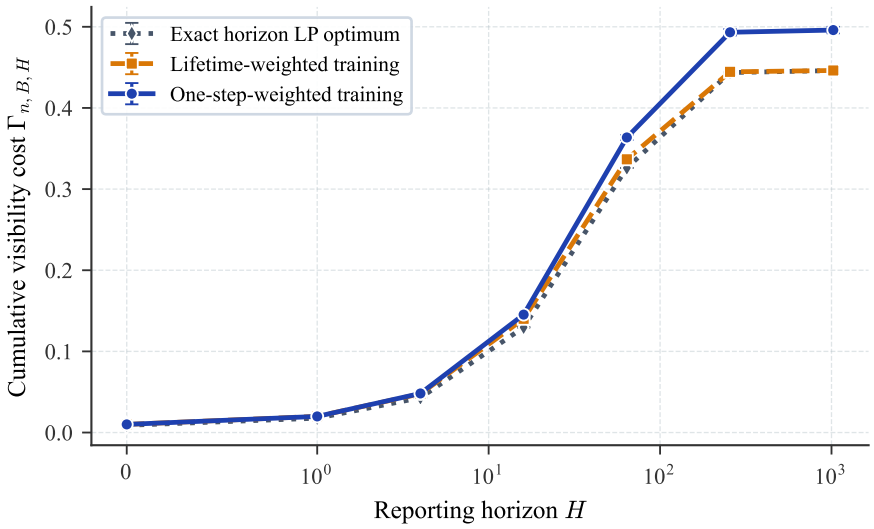

Figure 10: Cumulative objective for $B = 3$. The exact common-kernel optimum is compared with two finite-iteration training weightings. Different approximation errors do not imply different exact optimal query policies.

E.6 A LEARNED CLASSIFIER WITH EXACT MEMORY

This subsection supports Section 6.1. The built-in digit dataset contains 1,797 images, represented by 64 pixel values, with ten labels (scikit-learn developers, 2026). Pixels are divided by 16. For each of ten independent split seeds, use a stratified split with 1,197 training images and a 600-image task bank. Training and bank indices are disjoint *within* each split; the ten splits are not mutually disjoint datasets. Independently permute the bank's identity order before assigning Zipf rank probabilities. The network has architecture $64 \to 64 \to 32 \to 10$, ReLU hidden layers, Adam, penalty 10^{-4}, batch size 64, initial learning rate 0.001, at most 200 epochs, and early stopping on 15% of the assigned training set with patience 15. Executed fits stopped between 39 and 72 epochs, had no recorded convergence warnings, and achieved clean-bank top-1 accuracy between 0.95667 and 0.98000.

For each bank, tasks are its recurring image identities with fixed labels. Arrival weights are normalized $j^{-1.2}$, initial history has $n = 128$, and hard query budgets are $0, 32, 96$. A cache records all encountered or queried labels exactly. The backbone is not subsequently fine-tuned. A second condition independently permutes the bank labels while leaving all pixels and pretrained predictions unchanged. This is a controlled removal of transfer structure, not a new natural-data classification task.

The nonempty query policy is trained for threshold $\alpha_{10} = 1/9$, with $T = 8192$, $R = 2048$ and one independent 50,000-episode audit per budget, $\delta = 0.005$. It is shared across all prediction rules and network seeds; confidence intervals for this experiment therefore condition on these particular policy/audit realizations. The mathematical coverage guarantee of Theorem 5 separately averages over audit randomness. Each network, condition, and budget receives 4,000 independent history episodes. Within each episode the four rules below receive exactly the same cached labels.

Cached top-1. Return the known label when cached, otherwise the network's top-1 label. No taskwise guarantee is asserted.

Cached pooled conformal. Form scores $s_j(y) = 1 - \widehat{p}_j(y)$. The threshold is the $\lceil (n+1)(1-\alpha) \rceil$th historical true-label score, or infinity when that rank exceeds n. Only the original n i.i.d. history draws enter this calibration sample; adaptively queried labels are not added. Return all labels below the threshold, adding top-1 if this set is empty. Override by the true singleton on cached tasks. Conditional on the fixed bank and pretrained model, calibration and independent test identities are

exchangeable; ties only make the unrandomized quantile conservative. Thus the pooled baseline has its usual *marginal* guarantee. Cache overrides cannot remove the true label.

Certified nonempty. Use the independent-audit rule of Algorithm 3 with full-set probability $(1 - \alpha_{10}(1 - \delta/\alpha)/\overline{u}_j)_+$ and uniform singleton fallback, overridden by cached truth.

Guarded conformal (Guarded CP). Instantiate the taskwise coverage guard of Appendix D.3 using cached pooled conformal prediction as the nominal set predictor C^{nom}. On unresolved requests, return the full label set with probability $(1 - (\alpha - \delta)/\overline{u}_j)_+$ and the pooled conformal set otherwise. This guard guarantees finite-sample taskwise honesty for arbitrary nominal predictors; because the nominal set is not an unconditioned uniform singleton, the general threshold $\alpha - \delta$ is used without the $\kappa = M/(M-1)$ guessing relaxation. Consequently, Guarded CP delivers feature-conditioned predictions rather than unconditioned noise, raising empirical coverage from 95.73% to 99.84% at the cost of a marginal size difference ($2.34 \to 2.38$). This construction exemplifies the broader family of guarded architectures (such as Guarded Top-1 or Guarded APS/RAPS; Appendix D.3).

Each episode evaluates the prediction rule on every possible requested identity and then averages against π; this is an exact integration over the test identity conditional on that episode's predictions. Historical episodes remain independent. Evaluated per-task success counts give binomial sampling over the 4,000 episodes for each task, conditional on the fixed model and audit. To declare a task undercovered, form a one-sided Clopper–Pearson *upper* confidence bound on its true coverage and require that bound to be below 0.9. The per-test tail is $0.05/(600 \cdot 4 \cdot 2 \cdot 3 \cdot 10)$, covering all tasks, methods, conditions, budgets, and seeds by Bonferroni. No independence across tasks is needed for that union bound. A raw empirical minimum below 0.9 is not by itself evidence of a violation, and conversely zero detected violations are evidence of coverage rather than proof of it.

Table 4: Full digit-task results. Marginal coverage is a percentage; detected tasks are counted out of 600. All sets are nonempty. Intervals summarize ten independent network/deployment seeds, conditional on the evaluated query policies and audits.

Labels	B	Rule	Coverage (%)	Detected	Set size
Natural	0	Cached top-1	98.91 ± 0.26	17.9 ± 3.5	1.00 ± 0.00
Natural	0	Cached pooled CP	98.96 ± 0.30	17.6 ± 3.7	1.00 ± 0.01
Natural	0	Certified nonempty	95.53 ± 0.02	0.0 ± 0.0	3.20 ± 0.00
Natural	0	Guarded CP	99.83 ± 0.07	0.0 ± 0.0	3.24 ± 0.00
Natural	32	Cached top-1	99.39 ± 0.10	15.2 ± 3.0	1.00 ± 0.00
Natural	32	Cached pooled CP	99.42 ± 0.12	14.9 ± 3.1	1.00 ± 0.00
Natural	32	Certified nonempty	95.73 ± 0.02	0.0 ± 0.0	2.34 ± 0.00
Natural	32	Guarded CP	99.84 ± 0.06	0.0 ± 0.0	2.38 ± 0.00
Natural	96	Cached top-1	99.56 ± 0.10	13.3 ± 2.9	1.00 ± 0.00
Natural	96	Cached pooled CP	99.57 ± 0.12	13.1 ± 3.0	1.00 ± 0.00
Natural	96	Certified nonempty	95.80 ± 0.01	0.0 ± 0.0	1.83 ± 0.00
Natural	96	Guarded CP	99.84 ± 0.07	0.0 ± 0.0	1.87 ± 0.00
Permuted	0	Cached top-1	73.31 ± 0.47	534.8 ± 5.1	1.00 ± 0.00
Permuted	0	Cached pooled CP	96.87 ± 1.07	82.0 ± 28.2	3.33 ± 0.11
Permuted	0	Certified nonempty	95.54 ± 0.01	0.0 ± 0.0	3.20 ± 0.00
Permuted	0	Guarded CP	99.51 ± 0.15	0.0 ± 0.0	3.59 ± 0.02
Permuted	32	Cached top-1	82.27 ± 0.22	484.4 ± 4.7	1.00 ± 0.00
Permuted	32	Cached pooled CP	97.89 ± 0.76	74.5 ± 26.3	2.55 ± 0.08
Permuted	32	Certified nonempty	95.73 ± 0.02	0.0 ± 0.0	2.34 ± 0.00
Permuted	32	Guarded CP	99.53 ± 0.15	0.0 ± 0.0	2.72 ± 0.02
Permuted	96	Cached top-1	87.37 ± 0.12	422.3 ± 3.2	1.00 ± 0.00
Permuted	96	Cached pooled CP	98.49 ± 0.52	65.2 ± 22.5	2.10 ± 0.06
Permuted	96	Certified nonempty	95.80 ± 0.01	0.0 ± 0.0	1.83 ± 0.00
Permuted	96	Guarded CP	99.54 ± 0.15	0.0 ± 0.0	2.21 ± 0.02

The cache has zero retention errors by construction; the experiment isolates what remains after retention is ideal. The trained network is useful side information in the natural-label condition, while the guard's honesty remains valid when that structure is destroyed.

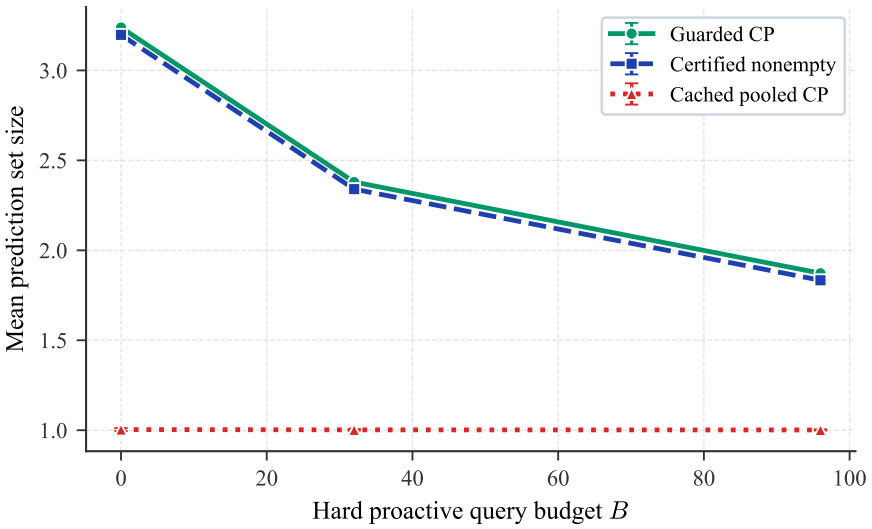

Figure 11: Additional truth queries reduce the certified prediction-set size. Pooled conformal sets are often smaller, but provide the weaker marginal guarantee rather than the taskwise guarantee studied here.

E.7 Noisy geometric recovery

This subsection supports the experiment quoted in Section 7. For actual nonparametric observations use $X \sim \mathrm{Unif}[0,1]$, $f(x) = 0.25\sin(2\pi x)$, fixed test descriptor $x_0 = 0.37$, $L = \pi/2$, $\beta = d = 1$, and $\alpha = 0.1$. Use $n = 16, 64, 256, 1024, 4096$ and $\sigma/L \in \{0, 0.01, 0.1, 1\}$. Each combination has ten seeds and 2,000 independently generated datasets per seed, with fresh Gaussian observation noise. Choose $k = \min\{n, \max\{1, \lceil(\sigma/L)^{2/3}n^{2/3}\rceil\}\}$ and form the nearest-neighbor interval proved in Theorem 23. Its noise quantile is $z_{0.95}$ because the two-sided error target is 0.1. Across the reported mean coverages, the smallest is 0.95245 and the noiseless rule covers every simulated case. The noiseless width decreases at the n^{-1} scale while the noise-dominated regime is consistent with $n^{-1/3}$; constants are not fitted or declared minimax from the plots.

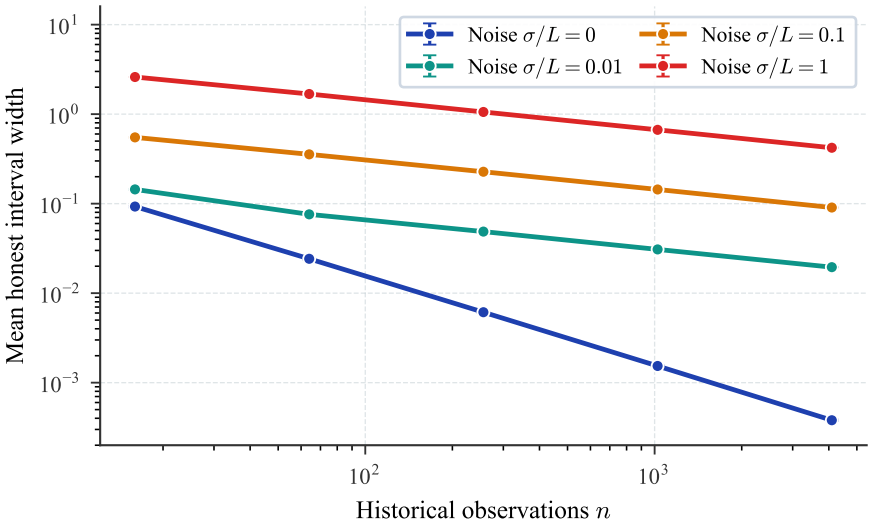

Figure 12: Executed noisy-regression experiment, not a plot of the theoretical rate formula alone. The known smoothness and noise assumptions justify finite intervals; their coverage and rates are proved separately.