NEAR-OPTIMAL SAMPLE COMPLEXITY FOR PAC SUBMODULAR MAXIMIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

How many noisy observations are needed to select a high-quality set with high probability? We study monotone submodular maximization, where an item's benefit can decrease as other items are selected. Given n items and a selection budget k, we provide a polynomial-time algorithm that returns a set of value at least $(1-1/e)\,\mathrm{OPT}_k - \varepsilon$ with probability $1-\delta$, where OPT_k is the optimal value. Under independent unit-variance Gaussian noise on marginal gains, the algorithm uses $\widetilde{O}(nk^2/\varepsilon^2)$ observations, without knowing OPT_k or an upper bound on function values. A matching worst-case lower bound, up to logarithmic factors, holds even for additive functions, for fixed $\delta \le 0.05$ and $1 \le k \le n/8$. The algorithm begins by evaluating each item and then selects from random candidate subsets. An exponential-moment bound for sample maxima controls the total loss from these subsets; the initial gain and the slack in the greedy guarantee together offset that loss. With exact marginal queries, the same analysis gives the strict $1-1/e$ ratio with high probability using $O(n\log(k/\delta))$ queries. We further show that, at comparable optimal values, submodular maximization can require polynomially more observations than maximizing an additive function.

1 INTRODUCTION

Monotone submodular maximization models tasks such as data summarization and sensor selection. Recent work studies faster optimization of structured objectives (Lv et al., 2026) and finite-confidence guarantees for identifying one of several valid answers (Poiani et al., 2026). In our setting, marginal gains must be estimated from fresh noisy observations and change as the selected set grows. We ask how many observations are needed over this trajectory to ensure that the returned set meets a prescribed quality target.

For a normalized monotone submodular function f, write $O = \mathrm{OPT}_k(f) = \max_{|S|\le k} f(S)$ and $\rho = 1 - e^{-1}$. We require the returned set C to satisfy

$$\mathbb{P}_f\{f(C) \ge \rho O - \varepsilon\} \ge 1-\delta, \tag{1}$$

including both candidate randomness and observation noise. Classical greedy attains $1-(1-1/k)^k \ge \rho$ with exact marginals (Nemhauser et al., 1978). Estimating all available marginals to accuracy $O(\varepsilon/k)$ at every round extends this argument to (1), but uses $\widetilde{O}(nk^3/\varepsilon^2)$ scalar observations: about nk estimates, each costing about k^2/ε^2. The question is whether the same final-output guarantee can be obtained with quadratic dependence on k, uniformly over unknown objective amplitudes.

Random candidate reduction gives efficient expected-quality algorithms (Mirzasoleiman et al., 2015), and noisy variants can have quadratic k dependence (Jawanpuria et al., 2024). An expectation over candidate draws, however, does not establish (1) for the returned set. Moreover, a guarantee $(\rho-\zeta)O - \varepsilon'$ meets the target only when $\zeta O + \varepsilon' \le \varepsilon$; a fixed coefficient loss cannot be absorbed uniformly over an unknown, unbounded objective range.

Near-optimal minimax observation cost. An initial singleton scan followed by random-candidate selection achieves (1) with $\widetilde{O}(nk^2/\varepsilon^2)$ observations on *every* normalized monotone submodular function. Its parameters use only $n, k, \varepsilon, \delta$, so the bound is uniform over unknown, arbitrarily large objective amplitudes at fixed noise variance. A direct partial-recovery bound for modular (additive)

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.

functions with $O \leq 4\varepsilon$ matches the joint polynomial dependence for $1 \leq k \leq n/8$, fixed $\delta \leq 0.05$ and $0 < \varepsilon \leq 1$. Thus, in that regime, handling general diminishing returns adds no polynomial worst-case observation cost to the modular subclass. The lower bound applies established coordinatewise change-of-measure ideas (Castro, 2014; Kaufmann et al., 2016); the upper bound controls marginals that change along the selection trajectory. Table 1 compares sufficient observation budgets for Gaussian implementations of existing methods; Theorem 2 gives the matching polynomial-order lower bound for the PAC task.

Table 1: Output guarantees and observation costs under independent unit-Gaussian marginal feedback, without a supplied bound on function values. PAC denotes probability at least $1 - \delta$.

Method	Output guarantee	Scalar observations
Bandit stochastic greedy (2024)†	Expected $(\rho - \zeta)O - \varepsilon$	$\widetilde{O}(nk^2\varepsilon^{-2}\log(1/\zeta))$
Fixed-batch threshold greedy (2025)†	PAC $\rho O - \varepsilon$	$\widetilde{O}(nk^3/\varepsilon^2)$
This work (Theorem 1)	**PAC** $\rho O - \varepsilon$	$\boldsymbol{\widetilde{O}(nk^2/\varepsilon^2)}$

†Sufficient budgets derived here for the Gaussian constructions in Appendix D. Here $0 < \varepsilon \leq 1$, $0 < \zeta \leq 1/2$; $\widetilde{O}$ suppresses logarithms in n, k, δ^{-1}.

Controlling the cumulative candidate deficit. The initial gain and finite-k greedy surplus provide a reserve proportional to the unknown largest singleton value. A conditional exponential-moment inequality for sample maxima bounds the aggregate signed candidate deficit at that same scale, with no explicit trajectory-length factor. Their cancellation leaves only the additive estimation error. The algorithm never identifies the top-k marginal sets used in this analysis (Section 4). With exact feedback, the same argument gives the strict $1 - 1/e$ ratio with probability $1 - \delta$ in $O(n\log(k/\delta))$ queries (Corollary 5); Appendix D.8 compares related exact-query guarantees. Separating candidate deficits from estimation error also permits unit sub-Gaussian noise and adaptive local sampling.

Dependence on objective amplitude. Restricting objective amplitude reveals a distinction between the function classes: on $k \leq O \leq 2k$, general submodular functions require $\Omega(nk)$ observations while modular functions admit $\widetilde{O}(n)$, for $n \geq 64k^2$, $k \geq 4$ and fixed $0 < \varepsilon \leq 1/10$, $0 < \delta \leq 0.05$. This polynomial separation emerges when objective amplitude is controlled. Section 5 and Table 2 develop this refinement.

2 Model and Sample-Complexity Results

Let $\mathcal{F}_n$ be the class of functions $f : 2^{[n]} \to \mathbb{R}_{\geq 0}$ with $f(\varnothing) = 0$, monotonicity, and diminishing marginal returns: $\Delta_f(S, i) \geq \Delta_f(R, i)$ when $S \subseteq R$ and $i \notin R$, where $\Delta_f(S, i) = f(S \cup \{i\}) - f(S)$. Normalization fixes $f(\varnothing) = 0$, while the range of f is unrestricted. A query (S, i) with $i \notin S$ returns

$$Y = \Delta_f(S, i) + Z, \qquad Z \sim N(0, 1), \tag{2}$$

with independent fresh noise of known unit variance. Each scalar observation costs one unit. The probability in (1) includes the observation noise, candidate draws, and other algorithmic randomness. Queries may depend on the complete past and use any context S not containing i. Neither the function representation nor a bound on its singleton values is supplied. A mean computed from b observations costs b. The unit-noise model fixes an absolute accuracy scale: rescaling f by $1/A$ also rescales noise variance to $1/A^2$ and tolerance to ε/A.

An algorithm is (ε, δ)-correct if it terminates, returns $|C| \leq k$ almost surely, and satisfies (1) for every $f \in \mathcal{F}_n$. Write T for its total observation count and define

$$C^*(n, k, \varepsilon, \delta) = \inf_{\mathcal{A} \text{ correct}} \sup_{f \in \mathcal{F}_n} \mathbb{E}_f T_{\mathcal{A}}. \tag{3}$$

This definition places no computational restriction on the lower bound. The algorithm for the upper bound uses only polynomial computation. Its observation cap is deterministic, hence also upper-bounds this expected-cost criterion. All logarithms are natural.

Algorithm 1 Greedy selection with an initial singleton scan

Require: $n, k, \varepsilon, \delta$ and the scalar marginal oracle (2)
if $k = n$ **then return** $[n]$.
Set r, L, s, Q, τ, b by (4); set $S \leftarrow \varnothing$.
for $j = 0, \ldots, k-1$ **do**
  if $j = 0$ or $s \geq n$ **then**
    $D \leftarrow [n] \setminus S$.
  else
    Draw s labels independently and uniformly from $[n]$.
    $D \leftarrow$ the distinct drawn labels not in S.
  end if
  if $D \neq \varnothing$ **then**
    For each $i \in D$, average b fresh observations of (S, i).
    Let i maximize these empirical means, breaking ties by label.
  else
    Let i be the smallest label in $[n] \setminus S$.
  end if
  $S \leftarrow S \cup \{i\}$.
end for
return S

Theorem 1 (Uniform upper bound)**.** *For $1 \leq k \leq n$, $0 < \varepsilon \leq 1$, and $0 < \delta \leq 1/2$, Algorithm 1 is (ε, δ)-correct on $\mathcal{F}_n$. If $k < n$, set*

$$r = \lceil 4e \log(3/\delta) \rceil, \qquad L = \max\{2r, 8\log(3k/\delta)\},$$
$$s = \lceil nL/k \rceil, \qquad \tau = \varepsilon/(2k),$$
$$Q = \begin{cases} nk - k(k-1)/2, & s \geq n, \\ n + (k-1)s, & s < n, \end{cases} \qquad b = \begin{cases} \lceil 2\tau^{-2} \log(2Q/\delta) \rceil, & s \geq n, \\ \lceil 2\tau^{-2} \log(6Q/\delta) \rceil, & s < n. \end{cases} \tag{4}$$

The algorithm uses at most Qb scalar observations. The full-scan branch allocates its entire failure budget to mean estimation. In particular,

$$Qb = O\left(\frac{nk^2}{\varepsilon^2} \log \frac{k}{\delta} \log \frac{n \log(k/\delta)}{\delta} \right). \tag{5}$$

If $k = n$, it returns $[n]$ without querying the oracle.

Theorem 2 (Lower bound on a modular subclass)**.** *For $0 < \varepsilon \leq 1$, $0 < \delta \leq 0.05$, and $1 \leq k \leq n/8$, every correct algorithm has an instance $f \in \mathcal{F}_n$ on which*

$$\mathbb{E}_f T \geq \frac{nk^2}{128\varepsilon^2}. \tag{6}$$

The instance can be chosen modular, and the bound permits arbitrary adaptive marginal contexts and arbitrary computation.

Together the theorems determine the joint polynomial dependence of $C^*(n, k, \varepsilon, \delta)$ on n, k, ε^{-1}, up to logarithmic factors, for fixed $\delta \leq 0.05$ throughout $1 \leq k \leq n/8$. The unmatched factor is $\log(k/\delta) \log(n \log(k/\delta)/\delta)$; the lower bound does not characterize the dependence on δ.

3 An initial scan followed by random candidates

Algorithm 1 (*candidate uniform*) first estimates all singletons, then draws candidates independently with replacement from $[n]$. It discards duplicates and selected labels, and charges a fresh batch for each remaining candidate.

For $k < n$, the output has exactly k labels and at most Q means are estimated. Sampling, deduplication and maximum selection use polynomial computation. The auxiliary top-k sets and samples below are used only in the analysis.

4 CONTROLLING THE COMPLETE ADAPTIVE TRAJECTORY

The proof bounds the cumulative loss from candidate sampling and noisy estimation along the selected prefixes. For $s \geq n$, Algorithm 1 is full-scan greedy with exactly the parameters of Appendix D.2, which proves its PAC guarantee directly. Consider the random-candidate branch. Fix $f \in \mathcal{F}_n$ and write $a = \max_i f(\{i\})$, $O = \mathrm{OPT}_k(f)$. Submodularity and monotonicity give

$$0 \leq \Delta_f(S, i) \leq a, \qquad a \leq O \leq ka. \tag{7}$$

If $a = 0$, the function is identically zero. Otherwise $a > 0$ is a fixed analysis parameter, even though the algorithm does not know it. The bounded variables in the next lemma are true marginal values sampled through random labels, not the unbounded Gaussian observations.

Proof mechanism. The singleton scan improves on the initial gap-based guarantee O/k. Together with the finite-k surplus over ρ, it pays for the signed candidate deficits over all later rounds:

$$f(S_k) - \rho O \;\geq\; \underbrace{\frac{a}{2e}}_{\text{scan + surplus}} \;-\; \underbrace{\frac{a}{2e}}_{\text{candidate deficit}} \;-\; \underbrace{2k\tau}_{\text{estimation}} \;=\; -\varepsilon.$$

This holds jointly with probability at least $1 - \delta$ by (16) and (12). A candidate maximum can exceed its reference mean: retaining these negative deficits controls their sum without certifying each round. The scale a cancels, leaving $\widetilde{O}(n)$ estimates at accuracy $O(\varepsilon/k)$ each.

4.1 AN EXPONENTIAL MOMENT FOR A SAMPLE MAXIMUM

Lemma 3 (Maximum exponential moment)**.** *Let $a > 0$ and let V take values in $[0, a]$ with mean μ. For independent copies $V_1, \ldots, V_r$, $r \geq 2$, put $M = \max_{t \leq r} V_t$. Then $\mathbb{E}e^{h(\mu - M)} \leq 1$ for $0 \leq h \leq r/(2a)$.*

The proof uses the classical extreme-point reduction for moment constraints (Pinelis, 2016): fixing mass and mean leaves at most two atoms. The required maximum bound then follows from a two-point inequality and discretization.

Proof. First assume finite support. For $h > 0$, writing F for the distribution function gives

$$\mathbb{E}e^{-hM} = e^{-ha} + h \int_0^a e^{-ht} F(t)^r \, dt. \tag{8}$$

Fix the support and mean, and vary the probabilities. The expression is convex in that probability vector, since F is linear in it and z^r is convex for $z \geq 0$. The polytope specified by nonnegative probabilities, total mass one, and a fixed mean has extreme points with at most two positive entries. Convexity reduces the claim to these distributions, preserving both the mean and the support bound.

Let the two values be $u < v$, with upper mass p, and write $x = h(v - u) \leq r/2$. The desired inequality reduces to

$$e^{-x} + (1 - e^{-x})(1 - p)^r \leq e^{-xp}. \tag{9}$$

The cases $x = 0$ and $p \in \{0, 1\}$ are immediate. Otherwise put $R(p) = (e^{-xp} - e^{-x})/(1 - e^{-x})$. The function $t/(1 - e^{-t})$ is increasing, and

$$\frac{t}{1 - e^{-t}} \leq \frac{r/2}{1 - e^{-r/2}} \leq r \quad (0 < t \leq x \leq r/2),$$

where the last inequality uses $r \geq 2$. It follows that

$$\frac{d}{dp} \log R(p) = -\frac{x}{1 - e^{-x(1-p)}} \geq -\frac{r}{1 - p} = \frac{d}{dp} \log(1 - p)^r.$$

Integrating from zero proves $R(p) \geq (1 - p)^r$, hence (9). Substitution into (8) proves the finite-support case. For a general bounded V, apply that case to $a\lfloor dV/a \rfloor/d$ and let $d \to \infty$. The means and sample maxima converge; bounded convergence passes the exponential moment inequality to the limit. □

The maximum's upward shift absorbs the convexity cost: no positive variance term is needed. For every $h \geq r/a$, a sufficiently rare two-point distribution on $\{0, a\}$ violates the moment bound (Appendix A.3). Thus the lemma's distribution-uniform exponential range is sharp within a factor of two.

Corollary 4 (Conditional trajectory bound)**.** *Let $a > 0$, $r \geq 2$, and let $(\mathcal{H}_t)_{t=0}^T$ be a filtration. Given $\mathcal{H}_{t-1}$, draw r independent samples from an $\mathcal{H}_{t-1}$-measurable distribution on $[0, a]$ with mean μ_t, and let their maximum M_t be $\mathcal{H}_t$-measurable. For predictable weights $w_t \in [0, 1]$ and $0 < \alpha < 1$,*

$$\mathbb{P}\left\{\sum_{t=1}^{T} w_t(\mu_t - M_t) > \frac{2a}{r}\log\frac{1}{\alpha}\right\} \leq \alpha. \tag{10}$$

Proof. With $\lambda = r/(2a)$, Lemma 3 conditionally gives $\mathbb{E}[e^{\lambda w_t(\mu_t - M_t)} \mid \mathcal{H}_{t-1}] \leq 1$. Iterating these inequalities bounds the exponential moment of the sum by one. Markov's inequality proves (10). □

The conditional distributions may depend on the entire past; the corollary itself requires no submodularity. Here, submodularity supplies the support bound and relates the conditional mean to the remaining optimization gap. Observation error is controlled separately.

4.2 Sampling useful marginals without identifying their labels

At each round, the analysis uses the largest k marginals as a reference population whose mean measures the remaining optimization gap. Let S_j be the set after j selections. Extend the marginal vector to all labels by putting $v_{j,i} = 0$ when $i \in S_j$; otherwise set $v_{j,i} = \Delta_f(S_j, i)$. Let A_j be its largest k entries, with ties broken deterministically, and put $\mu_j = k^{-1}\sum_{i \in A_j} v_{j,i}$. For any optimum S^*, submodularity gives

$$k\mu_j \geq \sum_{i \in S^* \setminus S_j} \Delta_f(S_j, i) \geq O - f(S_j). \tag{11}$$

Conditionally on the past, a random-candidate round hits A_j a binomial number B_j of times with mean $sk/n \geq L \geq 2r$. Thus $\mathbb{P}(B_j < r \mid \text{past}) \leq e^{-L/8} \leq \delta/(3k)$. For the analysis, generate each draw by an independent hit indicator of probability k/n, then take the next uniform label from an independent sequence inside or outside A_j. This produces exactly the algorithm's uniform label draws. The first r inside labels define an auxiliary ("ghost") maximum W_j, even if the actual round uses fewer than r of them. Given the past, W_j always has the distribution in Lemma 3; on $B_j \geq r$, the actual candidate maximum is at least W_j. This construction separates the moment bound from the hit event, so no conditioning on successful candidate draws is needed.

Let $q = 1 - 1/k$ and $w_j = q^{k-1-j}$, the weight of round j's deficit after unfolding the greedy recurrence. The filtration includes the past observations, selections and auxiliary variables, but no future randomness. Corollary 4, with $\alpha = \delta/3$, yields

$$\mathbb{P}\left\{\sum_{j=1}^{k-1} w_j(\mu_j - W_j) > \frac{2a}{r}\log\frac{3}{\delta}\right\} \leq \delta/3. \tag{12}$$

By the choice of r, the threshold is at most $a/(2e)$. The learner observes only the original draws, not the auxiliary labels.

4.3 The initial gain pays for the remaining trajectory deficit

Every fresh batch mean has conditional error probability at most $2e^{-b\tau^2/2} \leq \delta/(3Q)$ outside radius τ. A union bound over the at most Q actual estimates gives simultaneous accuracy with probability at least $1 - \delta/3$. Intersect this event with all the candidate-hit events and (12). Their total failure probability is at most δ; they need not be independent. On the intersection, $f(S_1) \geq a - 2\tau$ and, for $j \geq 1$,

$$f(S_{j+1}) - f(S_j) \geq \frac{O - f(S_j)}{k} - (\mu_j - W_j) - 2\tau. \tag{13}$$

If all sampled labels were already selected, the comparison uses $W_j = 0$ on the hit event and the nonnegativity of the added marginal. Writing $\rho_k = 1 - q^k$ and unfolding the recurrence gives

$$f(S_k) \geq \rho_k O + q^{k-1}(a - O/k) - a/(2e) - 2k\tau. \tag{14}$$

For $k \geq 2$, $q^{k-1} \geq e^{-1}$ and

$$k(\rho_k - \rho) \geq \frac{1}{2e}. \tag{15}$$

Indeed, with $x = 1/k$, the power series and $\ell 2^\ell \geq \ell + 1$ give

$$k \log(1 - 1/k) + 1 = -\sum_{\ell \geq 1} \frac{x^\ell}{\ell + 1} \leq -\sum_{\ell \geq 1} \frac{x^\ell}{\ell 2^\ell} = \log(1 - x/2),$$

proving (15). Consequently, using $O \leq ka$ with its nonpositive coefficient,

$$\begin{aligned}(\rho_k - \rho)O + q^{k-1}(a - O/k) &= q^{k-1}a + (e^{-1} - q^{k-1})O \\ &\geq k(\rho_k - \rho)a \geq a/(2e).\end{aligned} \tag{16}$$

Together with $2k\tau = \varepsilon$, (14) proves (1). For $k = 1$, the initial scan itself returns value at least $O - \varepsilon$ on the accuracy event. The case $k = n$ is immediate. Finally, $Q \leq n + (k-1)\min\{n, s\} \leq n(1 + L) + k = O(n \log(k/\delta))$, which proves (5) and completes Theorem 1.

The separation of candidate randomness from estimation error gives two immediate consequences for other feedback models.

Corollary 5 (Exact marginal queries)**.** *With exact marginals in place of batch estimates, Algorithm 1 satisfies* $\mathbb{P}\{f(C) \geq \rho O\} \geq 1 - \delta$ *using* $O(n \log(k/\delta))$ *marginal queries.*

Proof. Replace each batch estimate by an exact marginal and set the estimation losses to zero in the trajectory argument, keeping the candidate parameters. Each exact marginal needs at most two exact value queries. □

Corollary 6 (Conditionally unit sub-Gaussian noise)**.** *The upper bound also holds when, conditional on the past and current query, each fresh error Z has mean zero and known unit sub-Gaussian proxy:* $\mathbb{E}[e^{tZ} \mid \textit{past, query}] \leq e^{t^2/2}$ *for all t.*

Proof. Iterating the conditional exponential-moment bound gives $2e^{-b\tau^2/2}$ for each batch's mean error, conditional on its starting history and query. The remaining argument concerns true marginals and is unchanged. The lower bound uses the Gaussian subclass. □

4.4 Replacing uniform estimation inside a candidate set

The same trajectory proof accepts a local procedure that, conditionally on the history and D_j, returns a ν_j-optimal candidate except with probability η_j. It suffices that

$$\sum_{j=0}^{k-1} q^{k-1-j}\nu_j \leq \varepsilon, \qquad \sum_{j=0}^{k-1} \eta_j \leq \delta/3. \tag{17}$$

The first term includes the initial singleton scan. Appendix C establishes this interface for predictable budgets, including empty candidates and boundary cases.

The *candidate adaptive* implementation uses standard successive elimination (Even-Dar et al., 2006), using $\nu = \varepsilon/k$ and $\eta = \delta/(3k)$. For m candidates, cap the pilot at mb_m observations, where $b_m = \lceil 8\nu^{-2} \log(4m/\eta) \rceil$. If it has not certified an answer, obtain b_m fresh samples for every candidate and select their empirical maximum. For $k < n$, the total cost is at most $2Q\lceil 8k^2\varepsilon^{-2} \log(12nk/\delta) \rceil$ on every path, preserving the joint polynomial order. The appendix gives the confidence sequence on geometric stages and gap-dependent cost bounds, including an unconditional expectation bound. Section 7 isolates candidate reduction with a *full adaptive* control that applies this same selector to all remaining labels.

5 A LOWER BOUND FOR THE FINAL-OUTPUT TASK

The final output need only recover a fraction of the useful labels. We bound this task directly by comparing supports differing in one coordinate, following the approach used for adaptive support estimation (Castro, 2014). Put $d = 4\varepsilon/k$ and consider

$$f_G(S) = d|S \cap G|, \tag{18}$$

for every support of size $k-1$ or k. These functions belong to $\mathcal{F}_n$, and a marginal query for label i has mean $d\mathbf{1}\{i \in G\}$ in every context. Each query therefore gives exactly one unit-variance Gaussian arm observation.

On each k-element support H, correctness implies $\mathbb{E}_H|C \cap H| \geq \beta k$, with $\beta = (1-\delta)(\rho - 1/4) > 1/3$. Averaged over $(k-1)$-element supports and absent labels, output capacity bounds the reference inclusion probability by $k/(n-k+1)$. The stopped Gaussian change-of-measure inequality (Kaufmann et al., 2016, Lemma 1), binary Pinsker and averaging give

$$\frac{1}{\binom{n}{k-1}} \sum_{|G|=k-1} \mathbb{E}_G T \geq \frac{4(n-k+1)}{d^2} \left(\beta - \frac{k}{n-k+1}\right)^2 \geq \frac{nk^2}{128\varepsilon^2}. \tag{19}$$

Appendix B proves this calculation for arbitrary adaptive algorithms, without a symmetry assumption.

The support size is not promised; correctness includes both sizes and the empty reference when $k = 1$. All comparison instances have $O \leq 4\varepsilon$; at $O = 4\varepsilon$ the target is $(\rho - 1/4)O$, a positive constant fraction of the optimum. This establishes the full-class minimax lower bound.

Refinement by objective amplitude. The global minimax comparison takes the worst case over all amplitudes. Estimating the singleton scale refines the upper bound to $\widetilde{O}(n + nk^2/(\varepsilon + O/k)^2)$ (Appendix B.2), allowing smaller budgets as O grows. Table 2 compares the minimax expected costs C^*_A and $C^*_{A,\mathrm{mod}}$ on $A \leq O \leq 2A$, with correctness on each full function class; A is not an algorithm input. At $A = k$, the nonmodular lower bound exceeds the modular upper bound by a polynomial factor, isolating the statistical cost of diminishing returns at comparable objective values.

Table 2: Minimax observation costs by amplitude for $n \geq 64k^2$, $k \geq 4$, fixed $0 < \varepsilon \leq 1/10$, $0 < \delta \leq 0.05$, and fixed $M \geq 4$. Orders suppress logarithmic factors. Proofs: Appendices B.1–B.4.

Amplitude	Modular class	General class
$A = M\varepsilon$	$\widetilde{\Theta}(nk^2/\varepsilon^2)$	$\widetilde{\Theta}(nk^2/\varepsilon^2)$
$A = k$	$\widetilde{\Theta}(n)$	$\Omega(nk) \leq C^*_k \leq \widetilde{O}(nk^2)$
$A \geq k^2$	$\widetilde{\Theta}(n)$	$\widetilde{\Theta}(n)$

The nonmodular construction forces recovery of $\Omega(\sqrt{k})$ hidden labels, and its information bound covers every adaptive marginal context. For $A \geq 10k\varepsilon$, $k \geq 4$, $n \geq 64k^{3/2}$ and $\delta \leq 0.05$, the bound is $\Omega(nk^3/[A^2(1 + k^2/n)^2])$ at $O = A$, even when the learner knows $A \leq O \leq 2A$ (Appendix B.5). This preserves polynomial separations below the k^2 ground-set scale. For $A = k$ and $n \geq 64k^2$, aggregate certification attains $\widetilde{O}(nk)$ expected cost on the nonmodular reference instance, matching its lower bound up to logarithmic factors while retaining correctness on $\mathcal{F}_n$ (Appendix B.6). Every k-set is optimal on this reference instance. Its observation cost arises because the algorithm must also be correct on statistically similar alternatives with hidden useful labels.

6 RELATED GUARANTEES UNDER MARGINAL FEEDBACK

Noisy greedy and the feedback interface. Adaptive sampling of candidate marginals is established (Singla et al., 2016; Jawanpuria et al., 2024; Chen et al., 2025). Bandit greedy (BG) guarantees the quality of its returned set; bandit stochastic greedy (BSG) averages over candidate draws. Appendix D.1 implements their naive elimination (NE) selector with marginal observations, using the common prefix-value offset. A fixed-batch version of Confident Threshold Greedy (CTG), with threshold parameter $1/(4k)$, gives (1) at sufficient cost $\widetilde{O}(nk^3/\varepsilon^2)$. The experiments use a batched implementation of its adaptive Confident Sample subroutine (CTG-CS; Appendix D.3.1).

Repetition, boosting and other guarantees. Appendix D.5 treats repetition and scale-calibrated stochastic greedy. A finite-k conversion of boosted stochastic greedy (Castillo J. et al., 2023), including Gaussian estimation and selection, gives (1) without a value bound at sufficient cost $\widetilde{O}(nk^3/\varepsilon^2)$; repetition contributes the extra k (Appendix D.6). Other high-probability analyses use local-ratio assumptions (Hassidim & Singer, 2017a; Hashemi et al., 2018; Kaya et al., 2025); exact-oracle methods allow a multiplicative coefficient loss (Li et al., 2022; Breuer et al., 2020). Appendix D.7 accounts for the gradient and objective-scale costs of continuous-oracle methods (Hassani et al., 2019; Becker et al., 2023).

Noise persistence, structure and online objectives. Hassidim & Singer (2017b) primarily study persistent value noise; repeated queries return the same corrupted value. Chen & Crawford (2025) learn unknown weights of known basis functions using value-oracle access. Tajdini et al. (2024) establish minimax regret bounds with noisy set values. Our target is (1) under fresh marginal observations of a general submodular function.

Relation to identification. Average-additive top-k selection (Zhou et al., 2014) targets $O - k\eta$, which at $\eta = \varepsilon/k$ is stronger than (1), so its lower bound does not directly imply ours. Average-relative selection (Cao et al., 2015) uses Bernoulli feedback. Appendix B compares these and other identification targets.

7 Finite-scale costs and the trajectory mechanism

The experiments examine the finite-scale costs and trajectory mechanism behind the PAC upper bound: candidate reduction, local adaptation and signed deficits for Algorithm 1 and its adaptive variant. Gaussian sufficient-statistic simulations charge all observations, including initialization and fallback (Appendix E).

Cost calibration and a shared-selector comparison. Figure 1a calibrates the exact expected-cost formula (Appendix F.1) on modular and block-coverage instances at $n = 1024$, $\varepsilon = \delta = 0.05$ and $O/\varepsilon \in \{4, 40\}$. Candidate/full uniform ratios are 0.303 at $k = 256$ and 0.160 at $k = 512$; full scans for $k \leq 90$ explain equality at $k = 64$. Constructions and exact optima appear in Appendix E.1. The same reduction persists with adaptive estimation (Figure 1b): at $n = 16k$ and $O/\varepsilon = 40$, candidate/full adaptive cost ratios reach 0.328–0.335 at $k = 256$ and 0.180–0.182 at $k = 512$. Both methods use the same local selector and confidence allocation.

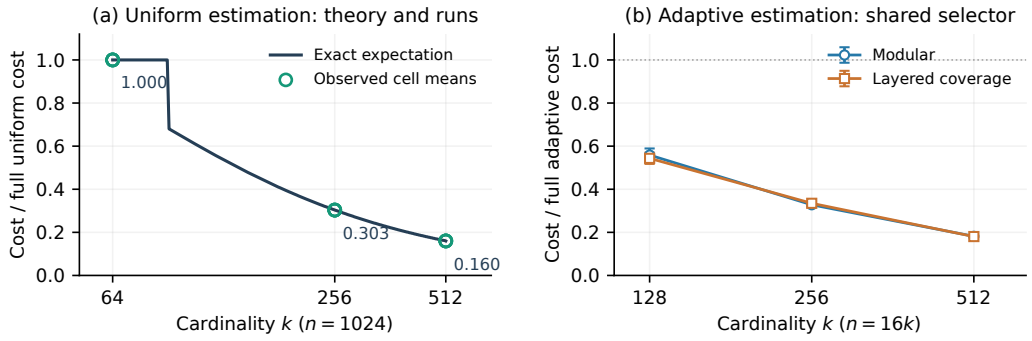

Figure 1: Complete observation costs under proved PAC parameters. (a) Cost calibration: exact expectation and all four cell means per k (overlapping markers), 100 trials per method and cell. (b) Shared-selector comparison at $O/\varepsilon = 40$: 10 trials at $k = 128$, four at $k = 256, 512$; bars are pointwise bootstrap 95% intervals. Each panel uses its stated full-scan control.

Candidate reduction and local adaptation. Table 3 crosses full/candidate selection with uniform/adaptive estimation at $n = 16k$, $k = 128$ and $\varepsilon = \delta = 0.05$. Candidate reduction lowers mean cost by 44.1–47.5% relative to full adaptive greedy in all four conditions. Local adaptation

further reduces candidate cost by 18.9–82.6% relative to uniform estimation. All 160 runs meet the output-quality target.

Table 3: Mean scalar observations (units of 10^{13}), 10 trials per method and condition. Ratio: candidate adaptive / full adaptive. Bold marks the lowest mean in each row. All methods use proved PAC parameters.

Family	O/ε	Full unif.	Full adapt.	Cand. unif.	Cand. adapt.	Ratio
Modular	4	21.49	5.97	11.47	**3.23**	0.541
Modular	40	21.49	3.57	11.46	**2.00**	0.559
Layered	4	21.49	17.69	11.46	**9.29**	0.525
Layered	40	21.49	3.91	11.47	**2.13**	0.543

With exact queries at $k = 128$, candidate selection uses 0.509 of full-scan queries, and all 80 runs meet $f(C) \geq \rho O$ (Appendices E.2 and F.2). Figure 4 illustrates the signed-deficit cancellation underlying the exponential-moment argument.

Comparison under marginal-value dispersion. We compare local implementations of CTG-CS (Chen et al., 2025) and BG-NE (Jawanpuria et al., 2024), calibrated to the same Gaussian feedback and final PAC target (Appendix D). Weighted partial-layer coverage varies column-weight spread $h \in \{4, 16\}$, coverage fraction $\omega \in \{1/4, 1/2, 1\}$ and layer strength $\gamma \in \{1.001, 2\}$. Here $n = 16k$, $k = 128$, $O/\varepsilon = 40$ and $\varepsilon = \delta = 0.05$, with six trials per method and condition. The optimum is known only to the evaluator. Across these 12 conditions, candidate adaptive uses 45.8–82.8% fewer observations than CTG-CS and 57.7–94.6% fewer than BG-NE (Figure 2). Appendices E.3 and F.3 give the construction, controls and query-count decomposition. Equal-weight controls with $\gamma = 1.001$ favor CTG-CS through common-threshold acceptance (Table 7); Table 6 reports absolute costs and output quality.

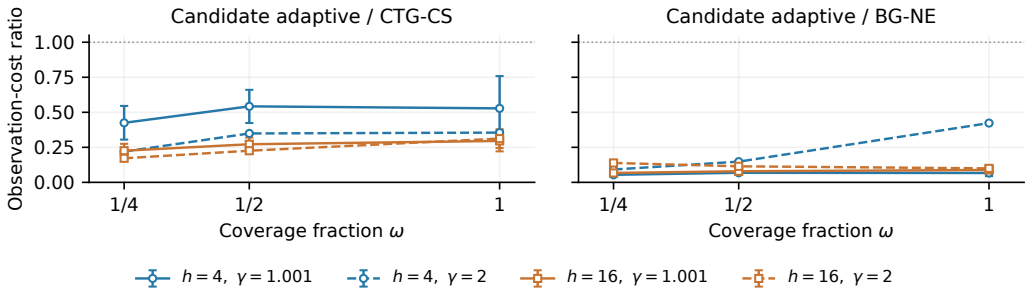

Figure 2: Observation-cost ratios for all 12 heterogeneous-coverage conditions. Ratios below one favor candidate adaptive; the dotted line denotes equal cost. Bars are pointwise paired-bootstrap 95% intervals from six trials.

Discussion. Controlling the complete trajectory gives a final-set PAC guarantee at near-optimal polynomial observation cost, uniformly over unknown objective amplitudes. The initial gain and finite-k surplus absorb candidate deficits; the same argument yields a near-linear exact-query guarantee. Amplitude-restricted bounds refine this global result by identifying a cost of diminishing returns beyond the modular subclass. Open questions concern sharp confidence dependence, matching minimax bounds at intermediate amplitudes (Table 2), and amplitude-layer bounds for $n = O(k)$.

REPRODUCIBILITY STATEMENT

Assumptions and complete proofs are provided in the main text and Appendices A–D. Experimental details are given in Appendices E–F. Anonymous code and instructions for reproducing the experiments are included in the supplementary material.

AI USE STATEMENT

In this work, we used generative AI tools for literature retrieval and discovery, to aid and polish the writing of the manuscript, and to assist with implementing research methods and experimental code. All AI-assisted code was reviewed and tested for correctness by the authors, and all AI-assisted literature retrieval and writing edits were verified. We take responsibility for the final content of this work, including its code, claims, results, and citations.

ETHICS STATEMENT

This work studies a mathematical oracle model using synthetic experiments. It involves no human subjects or personal data.

REFERENCES

Evan Becker, Jingdong Gao, Ted Zadouri, and Baharan Mirzasoleiman. High probability bounds for stochastic continuous submodular maximization. In *Proceedings of the 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, 2023.

Adam Breuer, Eric Balkanski, and Yaron Singer. The FAST algorithm for submodular maximization. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 1134–1143, 2020.

Wei Cao, Jian Li, Yufei Tao, and Zhize Li. On top-k selection in multi-armed bandits and hidden bipartite graphs. In *Advances in Neural Information Processing Systems*, volume 28, 2015.

Andrés C. Castillo J., Ege C. Kaya, and Abolfazl Hashemi. High probability guarantees for submodular maximization via boosted stochastic greedy. In *57th Asilomar Conference on Signals, Systems, and Computers*, pp. 602–606, 2023. doi: 10.1109/IEEECONF59524.2023.10476958.

Rui M. Castro. Adaptive sensing performance lower bounds for sparse signal detection and support estimation. *Bernoulli*, 20(4):2217–2246, 2014. doi: 10.3150/13-BEJ555.

Lijie Chen, Jian Li, and Mingda Qiao. Nearly instance optimal sample complexity bounds for top-k arm selection. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pp. 101–110, 2017.

Shouyuan Chen, Tian Lin, Irwin King, Michael R. Lyu, and Wei Chen. Combinatorial pure exploration of multi-armed bandits. In *Advances in Neural Information Processing Systems*, volume 27, 2014.

Wenjing Chen and Victoria G. Crawford. Linear submodular maximization with bandit feedback. In *Proceedings of the 28th International Conference on Artificial Intelligence and Statistics*, volume 258 of *Proceedings of Machine Learning Research*, pp. 4249–4257, 2025.

Wenjing Chen, Shuo Xing, and Victoria G. Crawford. Adaptive threshold sampling for pure exploration in submodular bandits. In *Proceedings of the 41st Conference on Uncertainty in Artificial Intelligence*, volume 286 of *Proceedings of Machine Learning Research*, pp. 612–646, 2025.

Yixin Chen, Tonmoy Dey, and Alan Kuhnle. Best of both worlds: Practical and theoretically optimal submodular maximization in parallel. In *Advances in Neural Information Processing Systems*, volume 34, 2021.

Rémy Degenne and Wouter M. Koolen. Pure exploration with multiple correct answers. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

Eyal Even-Dar, Shie Mannor, and Yishay Mansour. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *Journal of Machine Learning Research*, 7:1079–1105, 2006.

Abolfazl Hashemi, Mahsa Ghasemi, Haris Vikalo, and Ufuk Topcu. Randomized greedy sensor selection: Leveraging weak submodularity. arXiv:1807.08627v1, 2018. Version used for Theorem 3; subsequently published in IEEE Transactions on Automatic Control, 66(1):199–212, 2021.

Hamed Hassani, Amin Karbasi, Aryan Mokhtari, and Zebang Shen. Stochastic continuous greedy++: When upper and lower bounds match. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

Avinatan Hassidim and Yaron Singer. Robust guarantees of stochastic greedy algorithms. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, 2017a.

Avinatan Hassidim and Yaron Singer. Submodular optimization under noise. In *Proceedings of the 2017 Conference on Learning Theory*, volume 65 of *Proceedings of Machine Learning Research*, pp. 1069–1122, 2017b.

Pratik Jawanpuria, Bamdev Mishra, and Karthik S. Gurumoorthy. Revisiting stochastic submodular maximization with cardinality constraint: A bandit perspective. *Transactions on Machine Learning Research*, 2024.

Shivaram Kalyanakrishnan, Ambuj Tewari, Peter Auer, and Peter Stone. PAC subset selection in stochastic multi-armed bandits. In *Proceedings of the 29th International Conference on Machine Learning*, pp. 655–662, 2012.

Julian Katz-Samuels and Kevin Jamieson. The true sample complexity of identifying good arms. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 1781–1791, 2020.

Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the complexity of best-arm identification in multi-armed bandit models. *Journal of Machine Learning Research*, 17(1):1–42, 2016.

Ege C. Kaya, Michael Hibbard, Takashi Tanaka, Ufuk Topcu, and Abolfazl Hashemi. Randomized greedy methods for weak submodular sensor selection with robustness considerations. *Automatica*, 171:111984, 2025. doi: 10.1016/j.automatica.2024.111984. Theorem comparison uses the author manuscript arXiv:2404.03740v2.

Wenxin Li, Moran Feldman, Ehsan Kazemi, and Amin Karbasi. Submodular maximization in clean linear time. In *Advances in Neural Information Processing Systems*, volume 35, 2022.

Yang Lv, Guihao Wang, Dachuan Xu, and Ruiqi Yang. Efficient submodular maximization for sums of concave over modular functions. In *International Conference on Learning Representations*, 2026.

Baharan Mirzasoleiman, Ashwinkumar Badanidiyuru, Amin Karbasi, Jan Vondrák, and Andreas Krause. Lazier than lazy greedy. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2015.

George L. Nemhauser, Laurence A. Wolsey, and Marshall L. Fisher. An analysis of approximations for maximizing submodular set functions—i. *Mathematical Programming*, 14:265–294, 1978. doi: 10.1007/BF01588971.

Iosif Pinelis. On the extreme points of moments sets. *Mathematical Methods of Operations Research*, 83(3):325–349, 2016. doi: 10.1007/s00186-015-0530-0.

Riccardo Poiani, Martino Bernasconi, and Andrea Celli. Non-asymptotic analysis of (sticky) track-and-stop. In *International Conference on Learning Representations*, 2026.

Max Simchowitz, Kevin Jamieson, and Benjamin Recht. The simulator: Understanding adaptive sampling in the moderate-confidence regime. In *Proceedings of the 2017 Conference on Learning Theory*, volume 65 of *Proceedings of Machine Learning Research*, pp. 1794–1834, 2017.

Adish Singla, Sebastian Tschiatschek, and Andreas Krause. Noisy submodular maximization via adaptive sampling with applications to crowdsourced image collection summarization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2016.

Artin Tajdini, Lalit Jain, and Kevin Jamieson. Nearly minimax optimal submodular maximization with bandit feedback. In *Advances in Neural Information Processing Systems*, volume 37, 2024.

Yuan Zhou, Xi Chen, and Jian Li. Optimal PAC multiple arm identification with applications to crowdsourcing. In *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pp. 217–225, 2014.

A Additional details for the upper bound

A.1 Filtration and candidate coupling

For a fixed instance, include all completed oracle observations, chosen sets, label draws, and auxiliary variables from previous rounds in the filtration $\mathcal{H}_j$ before round j. The set A_j of largest k current marginals is then measurable with respect to $\mathcal{H}_j$. When $k < n$, generate fresh, mutually independent sequences: Bernoulli indicators with parameter k/n, uniform labels inside A_j, and uniform labels in $[n] \setminus A_j$. At each of the s draws, take the next label from the sequence named by its indicator. Conditional on $\mathcal{H}_j$, these are exactly independent uniform draws from $[n]$.

Extend the inside sequence to at least r entries even if fewer were used by the algorithm. The corresponding maximum W_j has its required law irrespective of B_j, and is dominated by the actual candidate maximum when $B_j \geq r$. The unused entries remain auxiliary; the algorithm is not supplied with them. For the full-scan branch, draw only the auxiliary inside sequence. Add this round's auxiliary entries to the next filtration after completing the actual round. Future sequences and oracle noise remain independent conditional on that enlarged past.

The exponential product formed using $q^{k-1-j}(\mu_j - W_j)$ is therefore a nonnegative supermartingale. Its argument does not condition on the event that all candidate hits or all Gaussian confidence intervals succeed. This matters because such conditioning could alter the distribution of later choices. Only after deriving (12) do we intersect the three events. Gaussian accuracy is established separately, conditionally at each actual batch. Randomly missing estimates can be padded by dummy failure indicators equal to zero when applying the union bound with the deterministic cap Q.

If the candidate set contains no unselected label, all its extended true marginal values are zero. On the hit event this also forces $W_j = 0$, so the arbitrary feasible addition is sufficient by monotonicity. If $a = 0$, every marginal is zero by (7). The algorithm continues to be well-defined without recognizing this case.

A.2 Finite-cardinality surplus

For $k \geq 2$, the inequality $\log(1-x) \geq -x/(1-x)$ at $x = 1/k$ implies $(1-1/k)^{k-1} \geq e^{-1}$. The other bound used in the main proof is

$$(1-1/k)^k \leq e^{-1}(1-1/(2k)).$$

The convergent power-series comparison in (15) proves this for every $k \geq 2$. Since $e^{-1} - q^{k-1} \leq 0$, the smallest value of $q^{k-1}a + (e^{-1} - q^{k-1})O$ over $a \leq O \leq ka$ occurs at $O = ka$. This establishes the uniform margin $a/(2e)$.

The argument uses the initial gain to absorb the trajectory deficit. The exponential parameter is proportional to $1/a$ for each fixed f, while every parameter chosen by Algorithm 1 depends only on $n, k, \varepsilon, \delta$.

A.3 Scale of the exponential range

Lemma 3 allows $0 \leq h \leq r/(2a)$ for every $[0,a]$-valued distribution and $r \geq 2$. Its order in r/a cannot be enlarged uniformly. Fix any $h \geq r/a$ and let $V = a$ with probability p and $V = 0$ otherwise. With $x = ha$,

$$F(p) := \mathbb{E}e^{h(\mu - M)} = e^{xp}\left[e^{-x} + (1-e^{-x})(1-p)^r\right], \qquad F(0) = 1.$$

The right derivative at zero is

$$F'(0) = x - r(1-e^{-x}) = (x-r) + re^{-x} > 0 \quad (x \geq r).$$

Hence $F(p) > 1$ for sufficiently small positive p. Together with the lemma's sufficient interval, this locates the largest distribution-uniform exponential range within a factor of two. The distribution may depend on h and r; no uniform numerical lower bound on p is needed.

B A DIRECT PARTIAL-RECOVERY LOWER BOUND

The following form separates the output-overlap requirement from the information calculation. It uses the coordinatewise comparison method of Castro (2014, Section 4) and the stopped bandit inequality of Kaufmann et al. (2016, Lemma 1). Here the overlap requirement comes directly from the final-output quality target.

Proposition 7 (Bounded-list partial recovery). *Fix $1 \le k \le n$, $d > 0$, and $\beta \in [0,1]$. For every $G \subseteq [n]$ of size $k-1$ or k, observations of label i are independent $N(d\mathbf{1}\{i \in G\}, 1)$ draws. An adaptive algorithm terminates almost surely on each such instance and outputs $|C| \le k$. Suppose $\mathbb{E}_H|C \cap H| \ge \beta k$ for every k-element support H. With $[x]_+ = \max\{x, 0\}$, its mean cost over the $(k-1)$-element supports satisfies*

$$\frac{1}{\binom{n}{k-1}} \sum_{|G|=k-1} \mathbb{E}_G T \ge \frac{4(n-k+1)}{d^2} \left[\beta - \frac{k}{n-k+1}\right]_+^2. \tag{20}$$

No symmetry of the algorithm or exact identification of its correct output labels is required.

Proof. If any expectation on the left is infinite, the result is immediate. Otherwise write $N = n-k+1$, $B = \binom{n}{k-1}$, and index comparisons by the BN pairs (G, i) with $|G| = k-1$ and $i \notin G$. Put

$$p_{G,i} = \mathbb{P}_G\{i \in C\}, \qquad q_{G,i} = \mathbb{P}_{G \cup \{i\}}\{i \in C\}.$$

Let N_i count observations of label i. Only this label changes distribution between G and $G \cup \{i\}$. Thus

$$\frac{d^2}{2} \mathbb{E}_G N_i \ge \mathrm{kl}(p_{G,i}, q_{G,i}) \ge 2(p_{G,i} - q_{G,i})^2, \tag{21}$$

where $\mathrm{kl}(p,q) = p\log(p/q) + (1-p)\log((1-p)/(1-q))$ with the usual endpoint conventions. The first inequality applies to random stopping: at the truncation $T \wedge m$, the transcript KL is $d^2 \mathbb{E}_G N_i(T \wedge m)/2$. All action kernels, internal randomness, and stopping rules are the same under a common history. Apply data processing to $\{T \le m, i \in C\}$ and let $m \to \infty$. Almost-sure termination, monotone convergence of the counts, and lower semicontinuity of binary KL give the first inequality in (21). The second is binary Pinsker with natural logarithms.

Let $\bar{p}$ and $\bar{q}$ be the averages of $p_{G,i}$ and $q_{G,i}$ over these pairs. The output-capacity constraint gives

$$\bar{p} = \frac{1}{BN} \sum_{|G|=k-1} \mathbb{E}_G|C \setminus G| \le \frac{k}{N}.$$

Each pair (H, i) with $|H| = k$ and $i \in H$ appears exactly once, as $(G, i) = (H \setminus \{i\}, i)$. Since $BN = k\binom{n}{k}$, we obtain

$$\bar{q} = \frac{1}{k\binom{n}{k}} \sum_{|H|=k} \mathbb{E}_H|C \cap H| \ge \beta.$$

Sum (21) over $i \notin G$ and average over G. Using $T \ge \sum_{i \notin G} N_i$ and Jensen's inequality gives

$$\frac{d^2}{2B} \sum_{|G|=k-1} \mathbb{E}_G T \ge 2N \frac{1}{BN} \sum_{G, i \notin G} (p_{G,i} - q_{G,i})^2 \ge 2N(\bar{p} - \bar{q})^2.$$

If $\beta > k/N$, the last square is at least $(\beta - k/N)^2$; otherwise the claimed lower bound is zero. This proves the proposition. □

For Theorem 2, take $d = 4\varepsilon/k$ and $\beta = (1-\delta)(\rho - 1/4)$. On the success event for a support H of size k, the overlap is at least $(\rho - 1/4)k$; outside it the overlap is nonnegative. Since $e > 8/3$ and $\delta \le 1/20$, we have $\beta > 57/160 > 1/3$. For $n \ge 8k$, $k/(n-k+1) \le 1/7$ and $n - k + 1 \ge 7n/8$. Consequently (20) is at least

$$\frac{4}{d^2} \frac{7n}{8} \left(\frac{1}{3} - \frac{1}{7}\right)^2 = \frac{8n}{63d^2} = \frac{nk^2}{126\varepsilon^2} \ge \frac{nk^2}{128\varepsilon^2}.$$

All adaptive marginal contexts are covered: for modular f_G, querying another label whose context contains i reveals no additional information about i. Each scalar observation is charged once. The lower bound uses correctness on supports of size k and a well-defined, capacity-limited output on supports of size $k-1$. These requirements follow from correctness on $\mathcal{F}_n$, which includes both support sizes without revealing their cardinality to the learner. The condition $n \geq 8k$ supplies a sufficient separation between the two inclusion probabilities.

Relation to average top-k identification. Other fixed-mean settings include PAC subset selection and combinatorial pure exploration (Kalyanakrishnan et al., 2012; Chen et al., 2017; 2014), and multiple-answer identification and discovery (Degenne & Koolen, 2019; Poiani et al., 2026; Katz-Samuels & Jamieson, 2020; Simchowitz et al., 2017). Zhou et al. (2014) study the average-additive target $f(C) \geq O - k\eta$ for bounded arm rewards, with Bernoulli lower-bound constructions. Their Theorems 5.1 and 5.5 give $\Omega(n\eta^{-2}(1+\log(1/\delta)/k))$ in their parameter range. Their argument also requires recovering a sufficient fraction of good arms. Cao et al. (2015) instead study average-relative quality with Bernoulli feedback; their Theorem 5 additionally depends on the mean reward scale. Matching our quality threshold would require $\eta = ((1-\rho)O + \varepsilon)/k$, rather than $\eta = \varepsilon/k$. Neither a lower bound for the stronger target $O - \varepsilon$ nor a Bernoulli noise calculation directly establishes our final-output Gaussian lower bound. Proposition 7 supplies that calculation with arbitrary contexts and no exact-identification requirement.

B.1 Lower bounds on amplitude strata

For $A > 0$, let $\mathcal{F}_n(A) = \{f \in \mathcal{F}_n : A \leq \mathrm{OPT}_k(f) \leq 2A\}$ and define

$$C_A^* = \inf_{\mathcal{A} \text{ correct on } \mathcal{F}_n} \sup_{f \in \mathcal{F}_n(A)} \mathbb{E}_f T_{\mathcal{A}}.$$

Correctness is required on the full function class. The amplitude A indexes the evaluation class; it is not an algorithm input.

Proposition 8 (A lower bound at a prescribed optimum)**.** *Let $2 \leq k \leq n/8$, $0 < \delta \leq 0.05$, and $A \geq 4\varepsilon > 0$. Every (ε, δ)-correct algorithm has a modular instance with $\mathrm{OPT}_k = A$ on which*

$$\mathbb{E}_f T \geq \max\left\{\frac{n}{6}, \frac{nk^2}{32A^2}\right\}.$$

The same lower bound holds even if correctness is required only on $\mathcal{F}_n(A)$ and A is supplied to the learner.

Proof. In Proposition 7, take $d = A/(k-1)$. The comparison instances have optima A and $Ak/(k-1)$, both in $[A, 2A]$. Correctness on each k-element support implies $\beta = (1-\delta)(\rho - \varepsilon/(kd)) \geq (1-\delta)(\rho - 1/4) > 1/3$. The argument following that proposition lower-bounds the average cost on the $(k-1)$-element supports by

$$\frac{8n}{63d^2} = \frac{8n(k-1)^2}{63A^2} \geq \frac{2nk^2}{63A^2} \geq \frac{nk^2}{32A^2}.$$

There is also a discovery cost independent of d. Couple the runs on G and $G \cup \{i\}$ using the same internal randomness and the same observations until the first query to i. Every other marginal has the same law in both instances, even if its context contains i. Thus the event of ever querying i has the same probability in both runs, and their output indicators can differ only on that event. With the notation of Proposition 7,

$$q_{G,i} - p_{G,i} \leq \mathbb{P}_G\{N_i \geq 1\}.$$

Summing and averaging, then using the same output-capacity calculation, gives mean cost at least $(n-k+1)(\bar{q} - \bar{p}) \geq (n-k+1)/3 - k \geq n/6$. Both lower bounds concern the same average over supports, so one support has cost at least their maximum. All cost-bearing instances have optimum A. □

For fixed $M = A/\varepsilon \geq 4$, this gives the same polynomial order in n, k, ε^{-1}, with a constant depending on M. The statistical lower-bound term scales as M^{-2} when M varies.

B.2 USING PART OF THE RESERVE TO ESTIMATE LESS PRECISELY

The following variant of Algorithm 1 estimates a singleton scale and allocates part of the reserve to coarser marginal estimation under the Gaussian interface (2).

Write $a = \max_i f(\{i\})$, $u = \varepsilon + a$, and

$$H(u) = \log \frac{16n}{\delta} + 2\log\big(2 + \log_+(1/u)\big), \qquad \log_+ x = \max\{0, \log x\}.$$

Let Q_+ denote the estimate cap in (4), with

$$r_+ = \lceil 8e \log(4/\delta) \rceil, \quad L_+ = \max\{2r_+, 8\log(4k/\delta)\}, \quad s_+ = \lceil nL_+/k \rceil.$$

Thus $Q_+ = O(n \log(k/\delta))$, including the full-scan branch.

Proposition 9 (Upper bound from a scale probe)**.** *For $1 \le k < n$, $0 < \varepsilon \le 1$ and $0 < \delta \le 1/2$, there is an (ε, δ)-correct algorithm on $\mathcal{F}_n$, using only $n, k, \varepsilon, \delta$ as inputs, whose expected observation count satisfies*

$$\mathbb{E}_f T = O\left(n(1 + u^{-2}) H(u) + Q_+ \left[1 + \frac{k^2}{u^2} \log \frac{8Q_+}{\delta} \right] \right). \tag{22}$$

The constants are universal. Computation is polynomial in the number of observations and n, k. When $k = n$, no observations are needed.

Proof. Scale probe. Put $J = \lceil \log_2(1/\varepsilon) \rceil$, and for $m = 0, \dots, J$ set

$$R_m = 2^{-m}, \qquad h_m = \log \frac{16n(m+1)^2}{\delta}, \qquad b_m = \lceil 2R_m^{-2} h_m \rceil.$$

At stage m, take b_m fresh observations of every singleton and let $\widehat{a}_m$ be the largest batch mean. If $\widehat{a}_m \ge 3R_m$, stop the probe and set $\ell = \widehat{a}_m - R_m$. If no stage triggers, set $\ell = 0$ after stage J. Every probe observation is charged. Gaussian tails and a union bound give simultaneous singleton accuracy within R_m at every executed stage with failure probability at most

$$\sum_{m \ge 0} \frac{\delta}{8(m+1)^2} = \frac{\pi^2 \delta}{48} < \frac{\delta}{4}.$$

Call this event E_0. On E_0, $0 \le \ell \le a$.

Selection and correctness. Run a fresh copy of the initial scan and candidate selection, using r_+, L_+, s_+ and

$$\tau_\ell = \frac{\varepsilon + \ell/(4e)}{2k}, \qquad b_\ell = \left\lceil 2\tau_\ell^{-2} \log \frac{8Q_+}{\delta} \right\rceil.$$

Selection uses fresh observations for both the initial scan and subsequent candidates. Conditional on the probe, the estimates fail with probability at most $\delta/4$. In the candidate branch, the hit and trajectory events each fail with probability at most $\delta/4$. The choice of r_+ makes the trajectory deficit at most $a/(4e)$, while (16) supplies $a/(2e)$. Consequently, with probability at least $1 - 3\delta/4$ conditional on the probe,

$$f(C) \ge \rho O + \frac{a}{4e} - 2k\tau_\ell. \tag{23}$$

The full-scan branch has no candidate deficit and hence also satisfies this inequality. For $k = 1$, it follows from $f(C) \ge a - 2\tau_\ell$ and $1 - \rho = 1/e$. On E_0, substituting τ_ℓ proves $f(C) \ge \rho O - \varepsilon$. The total failure probability is at most δ. The probe stops by J, and $\tau_\ell \ge \varepsilon/(2k)$, so termination and a finite deterministic observation cap hold on every sample path.

Expected observation cost. Let $N \le J$ be the last probe stage, whether or not it triggers. All the stage means may be generated in advance. For an index attaining singleton value a, its mean at stage m is $N(a, 1/b_m)$. Thus

$$\mathbb{P}\{N \ge m+1\} \le \mathbb{P}\{\widehat{a}_m < 3R_m\} \le \exp[-b_m(a - 3R_m)^2/2] \quad \text{if } a > 3R_m. \tag{24}$$

These tail bounds hold unconditionally. If a stage triggers, $\ell \ge 2R_N$. If no stage triggers, $R_J \le \varepsilon$, so $(\varepsilon + \ell)^{-2} \le R_J^{-2}$ still holds.

The bounds needed below are

$$\mathbb{E}\sum_{m=0}^{N} b_m = O((1+u^{-2})H(u)), \qquad \mathbb{E}(\varepsilon+\ell)^{-2} = O(u^{-2}). \tag{25}$$

Here are details covering small and large singleton scales. If $a \le 8\varepsilon$, the deterministic bounds $\sum_{m=0}^{J} b_m = O(\varepsilon^{-2}h_J)$ and $(\varepsilon+\ell)^{-2} \le \varepsilon^{-2}$ prove (25), since $u \in [\varepsilon, 9\varepsilon]$ and $h_J = O(H(u))$. If $8\varepsilon < a \le 8$, let m_* be the first stage with $R_{m_*} \le a/8$. It exists by stage J, and $a/16 < R_{m_*} \le a/8$. The total cost through m_* is $O(a^{-2}h_{m_*})$; stopping by that stage gives $(\varepsilon+\ell)^{-2} = O(a^{-2})$. For $j \ge 1$, (24) gives

$$\mathbb{P}\{N \ge m_* + j\} \le \exp[-c4^{j-1}h_{m_*}]$$

for a universal $c > 0$, because $a - 3R_m \ge 5a/8$ for $m \ge m_*$. Stage costs grow by at most a constant times $4^j(1+\log(j+1))$, and inverse squared precision by at most a constant times 4^j. Their products with this tail are summable, proving (25) in this case.

Finally, suppose $a > 8$. The first stage costs $O(h_0)$. On $\widehat{a}_0 \ge a/2+1$, it triggers with $\ell \ge a/2$. The complementary event has probability at most $\exp(-ca^2h_0)$ by the tail of the best singleton mean. If it still triggers, $\ell \ge 2$; if it does not, the later-stage arrival probabilities are at most $\exp(-ca^2 4^{m-1}h_0)$ for $m \ge 1$, by (24). Summing the corresponding costs and inverse squared precisions yields $O(h_0)$ and $O(a^{-2})$, respectively. This proves (25) also for $a > 8$.

The probe cost is $n\sum_{m=0}^{N} b_m$. The selection cost is at most $Q_+ b_\ell$, and

$$b_\ell \le 1 + 128e^2k^2(\varepsilon+\ell)^{-2}\log\frac{8Q_+}{\delta}.$$

Taking expectations and applying (25) proves (22). □

Amplitude regimes. Since $a \ge O/k$, the same algorithm, without A as an input, gives

$$C_A^* \le O\left(n(1+v^{-2})H(v) + Q_+\left[1 + \frac{k^2}{v^2}\log\frac{8Q_+}{\delta}\right]\right), \qquad v = \varepsilon + A/k.$$

For fixed $n, k, \delta, A > 0$, this bound stays finite as $\varepsilon \downarrow 0$: positive objective amplitude allows the observation cost to saturate with respect to additive accuracy. For $A \ge 10k\varepsilon$, $k \ge 4$, $n \ge 64k^2$ and $\delta \le 0.05$, combining this upper bound with Proposition 10 gives inverse-square terms $\Omega(nk^3/A^2)$ and $\widetilde{O}(nk^4/A^2)$. At the high-amplitude endpoint $A \ge k^2$, the upper bound is $\widetilde{O}(n)$; the discovery lower bound $n/6$ matches its polynomial order for $2 \le k \le n/8$ and fixed $\delta \le 0.05$. Thus the two characterized endpoints are $A = M\varepsilon$ for fixed $M \ge 4$, with polynomial order nk^2/ε^2, and $A \ge k^2$, with order n. The scale k^2 is in the fixed unit-noise normalization.

B.3 A NONMODULAR OBSTRUCTION AT LARGER AMPLITUDES

Nonmodular instances yield an additional amplitude-dependent observation cost. The construction controls the information in every marginal query, including contexts larger than the output budget.

Proposition 10 (A nonmodular amplitude lower bound). *Let $k \ge 4$, $n \ge 64k^2$, $0 < \delta \le 0.05$, and $A \ge 40\varepsilon > 0$. Put $\theta = \max\{1/k, 10\varepsilon/A\}$. Every correct algorithm has a normalized monotone submodular instance with $\mathrm{OPT}_k = A$ on which*

$$\mathbb{E}_f T \ge \frac{nk}{8A^2\theta^2} \ge \frac{nk}{800(\varepsilon + A/k)^2}. \tag{26}$$

In particular, if $A \ge 10k\varepsilon$, the first bound is $nk^3/(8A^2)$. The claim allows arbitrary adaptive marginal contexts and computation. It holds even if the learner is given A and correctness is required only on $\mathcal{F}_n(A)$, but does not supply the exact optimum of each instance.

Proof. Instances and the output event. Set $\gamma = 1-\theta$, $q = 1-\gamma/k$, $\beta = 1-q^k$, and $B = A/\beta$. Define a reference function and an alternative for each k-set G by

$$f_0(S) = B(1-q^{|S|}), \qquad f_G(S) = B\left[1 - \left(1 - \frac{|S\cap G|}{k}\right)q^{|S\setminus G|}\right]. \tag{27}$$

Write $a = B/k$, $g = |S \cap G|$, and $b = |S \setminus G|$. For $i \notin S$, the alternative marginal is aq^b if $i \in G$, and $a\gamma(1 - g/k)q^b$ otherwise. These nonnegative marginals decrease under every feasible addition to the context, proving monotonicity and submodularity on all subsets. The reference function has the same properties.

Since $1/k \le \theta \le 1/4$, $\gamma \ge 3/4$ and $\beta \ge 1 - e^{-3/4} > 1/2$. The bound $\log(1 - x) \ge -x - x^2/(2(1 - x))$ gives

$$k \log q \ge -1 + \theta - \frac{1}{2(k-1)} \ge -1 + \theta/3, \qquad \rho - \beta \ge \frac{\theta}{3e} > 0.$$

Consequently $\mathrm{OPT}_k(f_0) = A$ and $\mathrm{OPT}_k(f_G) = B \in [A, 2A]$. If $C \cap G = \varnothing$ and $|C| \le k$, then $f_G(C) \le A < \rho B - \varepsilon$, since $B(\rho - \beta) \ge A\theta/(3e) \ge 10\varepsilon/(3e) > \varepsilon$. Thus correctness on f_G implies $\mathbb{P}_G\{C \cap G \ne \varnothing\} \ge 1 - \delta$.

Information in an arbitrary query. Let G be uniform over the k-sets, only for averaging the alternatives. For every fixed (S, i), $i \notin S$, we show

$$\mathbb{E}_G\big[(\Delta_{f_G}(S, i) - \Delta_{f_0}(S, i))^2\big] \le \frac{4k}{n} a^2\theta^2. \tag{28}$$

Here S can have any size $s \le n - 1$. Conditional on $i \in G$ or $i \notin G$, $X = |S \cap G|$ is hypergeometric with population $n - 1$ and respectively $k - 1$ or k successes. Put $r = q^{-1}$ and $z = r^2 < 2$. Since $r - 1 \le 1/k \le \theta$, the absolute marginal difference when $i \in G$ is at most $a\theta q^s r^X(1 + X)$. For $i \notin G$, use $\phi(t) = (1 - t/k)r^t$ and the identity

$$\phi(t) - \phi(t + 1) = \frac{r^{t+1}}{k}\left(\theta + \frac{\gamma t}{k}\right).$$

Summing over $t < X$ bounds the absolute difference by $a\theta q^s r^X X(X + 1)/(2k)$.

The required weighted hypergeometric moments admit elementary bounds. If X has population N, m successes and s draws, then for $z \ge 1$ and integers $j \ge 0$,

$$\mathbb{E}[(X)_j z^X] \le \lambda^j z^j e^{\lambda(z-1)}, \qquad \lambda = sm/N, \tag{29}$$

where $(X)_j = X(X - 1)\cdots(X - j + 1)$ and $(X)_0 = 1$. Indeed, $\mathbb{E}(X)_h = (s)_h(m)_h/(N)_h \le (sm/N)^h$; expanding $z^X = (1 + (z - 1))^X$ proves the exponential bound for $j = 0$. For general j, condition on j designated draws being successes; the remaining success proportion $(m - j)/(N - j)$ is at most m/N.

In both conditional laws, set $\alpha = k^2/(n - 1) \le 1/32$ and $x = s/k$. Then $\lambda \le \alpha x$, and

$$q^{2s} e^{\lambda(z-1)} \le e^{-x}.$$

To see the latter, $\log q \le -3/(4k)$, $z - 1 \le 4/k$, and $k/(n - 1) \le 1/16$ make its logarithm at most $-s/k$. Expanding $(1+X)^2 = 1+3(X)_1+(X)_2$ and $X^2(X+1)^2 = 4(X)_1+14(X)_2+8(X)_3+(X)_4$, and using $\sup_{x \ge 0} x^j e^{-x} \le j!$, gives

$$q^{2s}\mathbb{E}[(1 + X)^2 z^X] \le 1 + 6\alpha + 8\alpha^2 \le 2,$$
$$q^{2s}\mathbb{E}[X^2(X + 1)^2 z^X] \le 8\alpha + 112\alpha^2 + 384\alpha^3 + 384\alpha^4 \le 12\alpha.$$

Weighting the two cases by k/n and $1 - k/n$ bounds the squared difference by $a^2\theta^2[2k/n + 3/(n - 1)] \le 4ka^2\theta^2/n$, proving (28).

Adaptive stopping and change of measure. If $\mathbb{E}_0 T = \infty$ there is nothing to prove. Otherwise let $p_G = \mathbb{P}_0(C \cap G \ne \varnothing)$ and $v_G = \mathbb{P}_G(C \cap G \ne \varnothing)$. Under the reference law the transcript and queries are independent of the averaging variable G. The stopped Gaussian change-of-measure argument of Proposition 7, now using the full marginal mean difference at each query, yields

$$\mathbb{E}_G\,\mathrm{kl}(p_G, v_G) \le \frac{1}{2}\mathbb{E}_0 \sum_{t=1}^{T} \mathbb{E}_G\big[(\Delta_{f_G}(S_t, i_t) - \Delta_{f_0}(S_t, i_t))^2\big] \le \frac{2B^2\theta^2}{nk}\mathbb{E}_0 T.$$

Random stopping can equivalently be handled first at $T \wedge m$ with the event $\{T \le m, C \cap G \ne \varnothing\}$, then by taking $m \to \infty$. Under the reference, output capacity and a union bound give $\mathbb{E}_G p_G \le k^2/n \le 1/64$, whereas $v_G \ge 1 - \delta$. Binary Pinsker and Jensen therefore give

$$\mathbb{E}_G\,\mathrm{kl}(p_G, v_G) \ge 2(1 - \delta - k^2/n)^2 \ge 1.$$

It follows that $\mathbb{E}_0 T \ge nk/(2B^2\theta^2) \ge nk/(8A^2\theta^2)$. Finally $A\theta = \max\{A/k, 10\varepsilon\} \le 10(\varepsilon + A/k)$. □

B.4 Separation from the modular subclass

Let $C^*_{A,\mathrm{mod}}$ be the analogous amplitude-layer minimax cost when both the correctness domain and the evaluation class are restricted to nonnegative modular functions.

Proposition 11 (Modular amplitude upper bound)**.** *There is an algorithm, correct on every nonnegative modular function and using no amplitude input, with*

$$\mathbb{E}_f T = O\left(n\left[1+\frac{k^2}{(\varepsilon+O)^2}\right]H\left(\frac{\varepsilon+O}{k}\right)\right), \qquad 0<\varepsilon\le 1, \quad 0<\delta\le 1/2,$$

where H is defined in Appendix B.2.

Proof. The case $k=n$ returns all labels without sampling. Otherwise use $R_m = 2^{-m}$, $h_m = \log(16n(m+1)^2/\delta)$ and $b_m = \lceil 2R_m^{-2}h_m\rceil$ for $m=0,\dots,J=\lceil\log_2(2k/\varepsilon)\rceil$. At each stage sample every singleton b_m times and let C_m be the k largest empirical means. Write $\widehat{v}_m$ for their average. Stop and return C_m if $\widehat{v}_m \ge 8R_m$, or at stage J. Every observation is charged.

With probability at least $1-\delta$, all singleton errors are within R_m at every stage. On this event, writing $v = O/k$, $|\widehat{v}_m - v| \le R_m$ and $f(C_m) \ge O - 2kR_m$. A trigger implies $v \ge 7R_m$, hence $f(C_m) \ge (5/7)O > \rho O$. Without a trigger, $R_J \le \varepsilon/(2k)$ gives $f(C_J) \ge O-\varepsilon$.

For expected cost, the average empirical mean on a fixed optimal k-set is $N(v, 1/(kb_m))$ and is no larger than $\widehat{v}_m$. Thus, for the last executed stage N,

$$\mathbb{P}\{N \ge m+1\} \le e^{-kb_m(v-8R_m)^2/2} \quad (v>8R_m).$$

This is unconditional. Put $w=\varepsilon/(2k)$. If $v\le 16w$, the deterministic geometric sum to J is $O((w+v)^{-2}H(w+v))$ per label. If $16w < v \le 16$, let m_* be the first stage with $R_{m_*} \le v/16$. The cost through m_* is $O(v^{-2}h_{m_*})$; the displayed tail beyond m_* decays as $\exp[-c4^{m-m_*}h_{m_*}]$, which is summable against the geometric stage costs. If $v>16$, stage zero costs $O(h_0)$ and all later arrivals have probability at most $\exp[-cv^2 4^{m-1}h_0]$. These cases give expected cost $O(n[1+(w+v)^{-2}]H(w+v))$. Since $w+v$ and $(\varepsilon+O)/k$ differ by at most a factor two, the claimed bound follows. □

For $A \ge 40\varepsilon$, $k\ge 4$, $n\ge 64k^2$ and fixed $\delta\le 0.05$, the general-class bounds give

$$C^*_A \ge \Omega\left(n + \frac{nk^2}{A^2} + \frac{nk}{(\varepsilon+A/k)^2}\right).$$

When $A \gg k\varepsilon$, the inverse-square lower and upper terms are $\Omega(nk^3/A^2)$ and $\widetilde{O}(nk^4/A^2)$, respectively.

There is also an explicit separation. Fix $0<\varepsilon\le 1/10$ and $0<\delta\le 0.05$, and take $A=k\ge 4$, $n\ge 64k^2$. Then

$$C^*_{k,\mathrm{mod}} = \widetilde{\Theta}(n), \qquad C^*_k = \Omega(nk). \tag{30}$$

The modular lower bound follows from Proposition 8; its proof requires correctness only on modular functions. Thus equality of the full-class worst-case polynomial orders does not extend to every amplitude layer. For example, with $n=64k^2$ the separation grows polynomially, up to logarithms, as $k\to\infty$. The separation holds for the final-output PAC target under arbitrary adaptive marginal algorithms. Appendix B.5 extends the lower bound to smaller ground sets with an explicit density correction.

B.5 Recovery requirements and the cardinality range

The same construction extends beyond $n\ge 64k^2$ if the information bound retains its dependence on the hidden-label density.

Proposition 12 (An amplitude lower bound with explicit density)**.** *Let $k\ge 4$, $A\ge 40\varepsilon>0$, $0<\delta\le 0.05$, and $\theta = \max\{1/k, 10\varepsilon/A\}$. If $n\ge 64k/\sqrt{\theta}$, every correct algorithm has an instance with $\mathrm{OPT}_k = A$ satisfying*

$$\mathbb{E}_f T \ge \frac{nk}{128A^2(\theta+k/n)^2}. \tag{31}$$

As in Proposition 10, the learner may know the bracket $[A, 2A]$, correctness can be confined to that bracket, and all adaptive marginal contexts are allowed. In particular, if $A \geq 10k\varepsilon$ and $n \geq 64k^{3/2}$, then

$$\mathbb{E}_f T \geq \frac{nk^3}{128A^2(1+k^2/n)^2}. \tag{32}$$

Proof. Use $f_0, f_G, B, a, q, \gamma$ from (27). We strengthen the required output event and refine the per-query bound.

Number of hidden labels in a successful output. Since $\gamma \leq 1 - 1/k$, $\gamma^2/(k-\gamma) \leq 1/k \leq \theta$. The logarithmic bound used earlier therefore gives $k \log q \geq -1 + \theta/2$, so $\rho - \beta \geq \theta/(2e)$. For an output containing g hidden labels, monotonicity implies

$$f_G(C) \leq B[1 - (1 - g/k)q^{k-g}].$$

With $\phi(t) = (1 - t/k)q^{-t}$, summing the identity for $\phi(t) - \phi(t+1)$ in Appendix B.3 gives

$$f_G(C) - B\beta \leq B\left(\theta\frac{g}{k} + \frac{g(g-1)}{2k^2}\right).$$

If $g \leq k\sqrt{\theta}/8$, this is at most $9B\theta/128$, since $\theta \leq 1/4$. However,

$$\rho B - \varepsilon - B\beta \geq B\theta\left(\frac{1}{2e} - \frac{1}{10}\right) > \frac{9B\theta}{128}.$$

Every successful output thus has $|C \cap G| > k\sqrt{\theta}/8$. Under f_0, averaging over the independent uniform G and applying Markov's inequality bounds the probability of this event by $8k/(n\sqrt{\theta}) \leq 1/8$.

Information in an arbitrary marginal query. Write $p = k/n$, $\xi = k/(n-1)$, $s = |S|$, $x = s/k$, $r = q^{-1}$ and $z = r^2$. Our range implies $\xi \leq 1/64$ and $\xi \leq 4p/3$. In either conditional hypergeometric law for $X = |S \cap G|$, its mean is at most $k\xi x$. Thus (29) and $q^{2s}e^{\lambda(z-1)} \leq e^{-x}$ imply

$$\begin{aligned} q^{2s}\mathbb{E}[X^2 z^X] &\leq 2k\xi + 8k^2\xi^2, \\ q^{2s}\mathbb{E}[X^4 z^X] &\leq 2k\xi + 56k^2\xi^2 + 288k^3\xi^3 + 384k^4\xi^4. \end{aligned}$$

Here $X^4 = (X)_1 + 7(X)_2 + 6(X)_3 + (X)_4$; the bounds again use $z < 2$ and $\sup_{x \geq 0} x^j e^{-x} \leq j!$.

Conditional on $i \in G$, the absolute marginal difference is bounded by $aq^s r^X(\theta + X/k)$. Its mean square is consequently at most $a^2(2\theta^2 + 4\xi/k + 16\xi^2) \leq 16a^2(\theta+\xi)^2$. Conditional on $i \notin G$, the telescoping identity bounds that difference by

$$aq^s r^X\left(\theta\frac{X}{k} + \frac{X^2}{2k^2}\right).$$

Its mean square, divided by a^2, is at most

$$\frac{4\theta^2\xi}{k} + 16\theta^2\xi^2 + \frac{\xi}{k^3} + \frac{28\xi^2}{k^2} + \frac{144\xi^3}{k} + 192\xi^4 \leq 3\xi(\theta+\xi)^2.$$

The final inequality follows from $1/k \leq \theta$, $k \geq 4$ and $\xi \leq 1/64$. Weighting the cases by p and $1-p$ gives

$$\mathbb{E}_G\left[(\Delta_{f_G}(S,i) - \Delta_{f_0}(S,i))^2\right] \leq 19a^2\xi(\theta+\xi)^2 \leq 64a^2p(\theta+p)^2. \tag{33}$$

This bound holds for every context size $s \leq n-1$.

Change of measure. Use the event $|C \cap G| > k\sqrt{\theta}/8$ in the stopped argument of Proposition 10. Its probability is at least $1-\delta$ under each alternative and at most $1/8$ on average under the reference. Pinsker and Jensen give average binary KL at least $2(1 - \delta - 1/8)^2 \geq 1$. By (33), the average transcript KL is at most $32B^2(\theta+p)^2\mathbb{E}_0 T/(nk)$. Since $B \leq 2A$, this proves (31). □

A separation on denser ground sets. At $A = k$, fixed $0 < \varepsilon \le 1/10$ and $0 < \delta \le 0.05$, Proposition 12 and the modular upper bound give

$$C_k^* \ge \Omega\left(\frac{nk}{(1+k^2/n)^2}\right), \qquad C_{k,\mathrm{mod}}^* = \widetilde{\Theta}(n), \quad n \ge 64k^{3/2}.$$

In particular, for any fixed $0 < \eta \le 1/2$, the scale sequence $n = \lceil 64k^{3/2+\eta} \rceil$ has a general-class lower bound $\Omega(nk^{2\eta})$ and a modular upper bound $\widetilde{O}(n)$. This yields a polynomial separation for ground sets growing subquadratically in k. The density factor accounts for both the greater overlap of an uninformed output with a denser hidden set and the increased information in individual queries.

B.6 Matching cost on the nonmodular reference instance

The reference instance in the amplitude lower bound admits a matching upper bound. The construction certifies the final set through an upper bound on the sum of the k largest singleton means. A multiscale coordinate test implements this certificate using the same scalar Gaussian oracle; its sampling principle is related to sparse-signal detection (Castro, 2014).

Proposition 13 (Reference-instance sharpness)**.** *Let $k \ge 4$, $n \ge 64k^2$, $0 < \varepsilon \le 1/10$ and $0 < \delta \le 0.05$. Let f_0 be the reference in (27) with $A = k$, so $\theta = 1/k$ and $\mathrm{OPT}_k(f_0) = k$. Then*

$$\frac{nk}{8} \le I(f_0) := \inf_{\mathcal{A} \text{ correct on } \mathcal{F}_n} \mathbb{E}_{f_0} T_{\mathcal{A}} \le \widetilde{O}(nk). \tag{34}$$

The upper bound is achieved with polynomial computation and inputs $n, k, \varepsilon, \delta$ only. Neither the function formula nor its amplitude is supplied to the algorithm. The suppressed factors are logarithmic in $n, k, \varepsilon^{-1}, \delta^{-1}$.

Proof. A one-sided aggregate certificate. For a vector of means μ, let $W_k(\mu)$ be the sum of its k largest entries. Given $u \in \mathbb{R}$, $t > 0$ and $0 < \eta \le 1/4$, the following test accepts with probability at least $1 - \eta$ if $\max_i \mu_i \le u$, and at most η if $W_k(\mu) > ku + t$. Set $J = 1 + \lfloor \log_2 k \rfloor$, and for $j = 0, \dots, J-1$ put

$$m_j = 2^j, \qquad d_j = \frac{t}{Jm_j}, \qquad M_j = \left\lceil \frac{n}{m_j} \log \frac{2}{\eta} \right\rceil, \qquad M = \sum_j M_j.$$

At scale j, draw M_j independent uniform labels with replacement. For each draw, take a fresh batch of $b_j = \lceil 32 d_j^{-2} \log(4M/\eta) \rceil$ observations. Reject if any batch mean exceeds $u + d_j/2$; otherwise accept. Every batch, including repeated labels, is charged.

With probability at least $1 - \eta/2$, every batch error is at most $d_j/4$. On this event the test accepts whenever $\max_i \mu_i \le u$. For the other claim, order $x_i = (\mu_i - u)_+$ decreasingly. If $W_k(\mu) > ku + t$, then $\sum_{i=1}^k x_{(i)} > t$. Partition the ranks into blocks $[m_j, \min\{2m_j - 1, k\}]$. If all block leaders satisfied $x_{(m_j)} \le t/(Jm_j)$, each block would sum to at most t/J, a contradiction. Thus some scale has at least m_j labels with $\mu_i > u + d_j$. The probability of missing all of them is at most $e^{-M_j m_j/n} \le \eta/2$; hitting one with an accurate batch forces rejection. This proves the two probability claims, also conditionally on any history fixing u, when the test uses fresh observations.

For $H = \log(2/\eta)$, $M \le 3nH$ and $\sum_j M_j m_j^2 \le 2nkH + 4k^2/3 \le 4nkH$. Hence its deterministic observation cap is

$$T_{\mathrm{cert}} = \sum_j M_j b_j \le 3nH + \frac{128nkJ^2H}{t^2} \log \frac{12nH}{\eta}. \tag{35}$$

A uniformly correct algorithm. Fix $C_0 = \{1, \dots, k\}$, $t = (16e)^{-1}$ and $\eta = \delta\varepsilon^2/(64k^2)$. Query each of the k marginal pairs along its fixed insertion path $b_{\mathrm{val}} = \lceil 2kt^{-2} \log(2/\eta) \rceil$ times. The sum $\widehat{V}$ of their batch means estimates $V = f(C_0)$ with variance k/b_{val}, so $\mathbb{P}(|\widehat{V} - V| > t) \le \eta$. All kb_{val} scalar observations are counted. Set $V_L = \widehat{V} - t$ and

$$u = \frac{V_L + \varepsilon - \rho t}{\rho k}.$$

Apply the aggregate test to the singleton means $\mu_i = f(\{i\})$ with new observations. If it accepts, return C_0; if it rejects, run a fresh copy of Algorithm 1 with parameters $(\varepsilon, \delta/2)$, including its initial scan.

For every $f \in \mathcal{F}_n$, submodularity gives $\mathrm{OPT}_k(f) \le W_k(\mu)$. On accurate value estimation and a valid acceptance, $\rho\,\mathrm{OPT}_k(f) \le \rho(ku+t) = V_L + \varepsilon \le V + \varepsilon$. An incorrect return from this branch has probability at most 2η; the fallback has conditional failure probability at most $\delta/2$. Thus the full algorithm is correct on $\mathcal{F}_n$, since $2\eta + \delta/2 \le \delta$. All branches terminate within a deterministic observation cap.

Cost on the reference instance. Under f_0, $V = k$ and every singleton mean equals $a_0 = B\gamma/k$, with $\gamma = 1 - 1/k$ and $B = k/[1-(1-\gamma/k)^k]$. The function $h(x) = 1 - (1-x/k)^k - \rho x$ is concave on $[0,1]$, $h(0) = 0$, and (15) gives $h(1) \ge 1/(2ek)$. Consequently

$$k - \rho k a_0 = Bh(\gamma) \ge \frac{B\gamma}{2ek} \ge \frac{3}{8e}.$$

If value estimation is accurate, then $V_L \ge k - 2t$ and

$$\rho k(u - a_0) \ge \frac{3}{8e} + \varepsilon - (2+\rho)t > 0.$$

The test therefore accepts with conditional probability at least $1-\eta$, so the fallback probability is at most 2η. Writing $B_0 = \widetilde{O}(nk^2/\varepsilon^2)$ for its deterministic cap, the unconditional expected cost is bounded by

$$\mathbb{E}_{f_0} T \le k b_{\mathrm{val}} + T_{\mathrm{cert}} + 2\eta B_0 = \widetilde{O}(k^2 + nk + n) = \widetilde{O}(nk).$$

The smaller internal failure probability adds only logarithmic cost. Finally, the proof of Proposition 10 gives $\mathbb{E}_{f_0} T \ge nk/8$ for every correct algorithm when $A = k$ and $\theta = 1/k$, proving (34). □

C A REPLACEABLE CANDIDATE SELECTOR

The trajectory argument separates candidate generation from local estimation. The selector may use any procedure that identifies an approximately best candidate with the required conditional probability.

Proposition 14 (Conditional local errors). *For $k \ge 2$, put $q = 1 - 1/k$ and $w_j = q^{k-1-j}$ for $0 \le j < k$. Keep the candidate law of Algorithm 1. Given the history and the realized nonempty D_j, choose $\nu_j > 0$ and $\eta_j > 0$ before obtaining new observations. Suppose the local procedure returns $i_j \in D_j$ with conditional probability at least $1 - \eta_j$ of*

$$\Delta_f(S_j, i_j) \ge \max_{i \in D_j} \Delta_f(S_j, i) - \nu_j.$$

If, on every path, $\sum_{j=0}^{k-1} w_j \nu_j \le \varepsilon$ and $\sum_{j=0}^{k-1} \eta_j \le \delta/3$, the final set satisfies (1). An empty D_j uses the original zero-cost fallback. For full candidate scans, the failure-budget condition can be relaxed to $\sum_j \eta_j \le \delta$. For $k = 1$, one local call with $\nu_0 \le \varepsilon$ and $\eta_0 \le \delta$ suffices; for $k = n$, return $[n]$.

Proof. Conditional correctness and the tower rule give a total local failure probability at most $\mathbb{E}\sum_j \eta_j$. This remains true with predictable, history-dependent budgets. The ghost-label construction and the conditional maximum inequality are unchanged: new label draws remain independent of the enlarged past, and the selector need not observe any ghost variables. On the intersection of the local, hit, and trajectory events, the initial value is at least $a - \nu_0$ and the recurrence (13) has ν_j in place of 2τ. Consequently,

$$f(S_k) \ge \rho_k O + q^{k-1}(a - O/k) - a/(2e) - \sum_{j=0}^{k-1} w_j \nu_j \ge \rho O - \varepsilon.$$

An empty candidate set is treated by nonnegativity exactly as in the main proof. For full scans the ordinary greedy recurrence directly gives $f(S_k) \ge \rho_k O - \sum_j w_j \nu_j$ without candidate or trajectory failures. The stated boundary cases follow directly. □

C.1 Capped successive elimination with fresh fallback

Fix one nonempty candidate set of size m, local tolerance ν, and local failure budget η. If $m = 1$, return its label at zero cost. Otherwise set

$$b_m = \left\lceil 8\nu^{-2} \log(4m/\eta) \right\rceil, \qquad P_m = m b_m. \tag{36}$$

The pilot is an ordinary confidence-based elimination procedure (Even-Dar et al., 2006). At stage $h \geq 1$, bring every surviving arm to $t_h = 4^{h-1}$ observations using independent increments. Define

$$c_h = \sqrt{\frac{2}{t_h} \log \frac{8mh^2}{\eta}}. \tag{37}$$

Remove arms whose empirical upper endpoint is below the largest lower endpoint. Return the empirical best if only one arm remains or $2c_h \leq \nu$. Stop the pilot before its observations exceed P_m. The pilot spends any remainder on one equal increment for the surviving arms; it makes no confidence decision from that final, non-geometric increment. If the pilot has not returned a certified answer, draw b_m *fresh* observations from each of the original m candidates and return their empirical maximum. This includes previously eliminated arms and pays for every failed pilot observation.

Proposition 15 (Local validity and full cost)**.** *The procedure terminates and returns a ν-optimal candidate with conditional probability at least $1 - \eta$. Its observation count is at most $2mb_m$ on every path. Thus, for $k < n$, using it with $\nu = \varepsilon/k$, $\eta = \delta/(3k)$ at all rounds, including the initial scan, gives a final-output PAC algorithm with deterministic cost at most*

$$2Q \left\lceil \frac{8k^2}{\varepsilon^2} \log \frac{12nk}{\delta} \right\rceil. \tag{38}$$

This preserves the joint polynomial order of Theorem 1. All computations have polynomial overhead.

Proof. For each arm, expose an infinite independent stream in the fixed current context, including samples not actually used. A union bound at the deterministic stage sizes gives

$$\mathbb{P}\{\exists i, h : \ |\widehat{\theta}_{i,h} - \theta_i| > c_h\} \leq \sum_{h \geq 1} \frac{\eta}{4h^2} < \eta/2.$$

Elimination and stopping are therefore valid despite adaptive removal and random termination. A true maximizer is never removed, so a sole survivor is a true maximizer. In the other stopping case, the returned empirical maximum loses at most $2c_h \leq \nu$. Conditionally on the entire pilot, the fresh fallback means each have error at most $\nu/2$ except with union probability $2m \exp(-b_m \nu^2/8) \leq \eta/2$. Its empirical maximum loses at most ν. Taking the union of these two errors proves local correctness. The pilot and fallback each cost at most P_m, on both successful and unsuccessful noise paths. Summing over the at most Q candidate-arm occurrences and using $m \leq n$ proves (38). Proposition 14 applies since $w_j \leq 1$ and there are k calls. Stage counts are logarithmic in the finite sample cap; candidate generation, elimination and averaging have polynomial cost. □

Gap-dependent stopping costs. The same proof quantifies how local gaps affect the observation cost. For fixed local means, let $\Delta_i = \max_\ell \theta_\ell - \theta_i$. Let h_i be the first stage with $c_{h_i} < \max\{\nu, \Delta_i\}/4$, and define $G(\theta, \nu, \eta) = \sum_i 4^{h_i - 1}$. On the simultaneous pilot accuracy event, the uncapped pilot samples arm i at most $4^{h_i - 1}$ times: when $\Delta_i \geq \nu$ it is eliminated by that stage, and otherwise the tolerance stopping rule has already applied. If the cap prevents completing a stage, this uncapped total exceeds P_m; the pilot and fallback together then cost at most $2P_m$. Consequently the actual local cost T_j satisfies

$$T_j \leq 2 \min\{P_m, G(\theta, \nu, \eta)\} \quad \text{with conditional probability at least } 1 - \eta/2,$$

and, without conditioning on correctness,

$$\mathbb{E}[T_j \mid \text{call history}] \leq 2 \min\{P_m, G(\theta, \nu, \eta)\} + \eta P_m.$$

The deterministic cap controls the contribution of inaccurate pilot paths to the unconditional expectation. Nearly tied useful arms can require precision ν and trigger fallback, as illustrated by the cost decomposition in Appendix F.2.

Full-candidate adaptive control. The standard adaptive-greedy construction solves an approximate best-arm problem at each full set of remaining labels (Singla et al., 2016, Theorem 1, $k' = 1$). Our full-candidate control uses the identical capped selector, ν, and η just described. It has the same guarantee for the returned set, with sequential, gap-dependent sampling. Full greedy could use $\eta = \delta/k$ rather than $\delta/(3k)$, improving a logarithmic factor. Sharing the allocation isolates candidate reduction in our comparison. The experimental label *Full adaptive* refers to this selector with the stated observation cap.

D Related guarantees and baseline implementations

D.1 Local selection and output-value estimation

Every estimate below is charged its scalar marginal-observation cost. The parameter choices below convert published guarantees to the feedback model and final PAC target in Section 2. Fix a prefix S and a candidate set R of size d_R before collecting fresh observations. Section 3.1 of Jawanpuria et al. (2024) uses the arm mean $f(S \cup \{i\})$ in place of the marginal mean, since $f(S)$ is common to all candidates. Its stochastic sample functions need not be submodular. The paper uses bounded observations for its displayed constants and explicitly permits sub-Gaussian extensions. We instantiate its uniform-sampling NE rule with fresh independent marginal samples.

For independent Gaussian marginal samples of known proxy variance σ^2, average

$$b_R = \lceil 8\sigma^2 \nu^{-2} \log(2d_R/\beta) \rceil$$

observations per candidate. A Gaussian tail bound and a union bound give simultaneous error at most $\nu/2$ with conditional probability $1 - \beta$. The empirical maximizer therefore has marginal at least $\max_{i \in R} \Delta_f(S, i) - \nu$. The proof uses the arm gaps and centered concentration, which are unchanged by a common shift of the means. Thus $d_R b_R$ fresh marginal observations realize the same per-round PAC selection property. The learner need not estimate $f(S)$ or know an amplitude bound. All statements are conditional on the chosen prefix and candidates, with fresh independent observations for each subsequent estimate.

For an unconditional expected-quality Gaussian version of bandit stochastic greedy, the empirical maximizer's conditional expected loss is at most $2\sqrt{2\log(2d_R)/b}$ at unit variance. Taking $b = \lceil 8k^2\varepsilon^{-2} \log(2n) \rceil$ makes this at most ε/k. Sampling at most $\lceil (n/k) \log(1/\zeta) \rceil$ candidates per round as in stochastic greedy gives the unconditional recursion

$$\mathbb{E} f(S_{j+1}) \geq \left(1 - \frac{1-\zeta}{k}\right) \mathbb{E} f(S_j) + \frac{1-\zeta}{k} O - \varepsilon/k.$$

Unfolding it yields $\mathbb{E} f(S_k) \geq (1 - e^{-1} - \zeta)O - \varepsilon$, at sufficient cost $\widetilde{O}(nk^2\varepsilon^{-2} \log(1/\zeta))$ for the usual $0 < \zeta \leq 1/2$ range. The ceiling contributes at most $k \leq n$. The expectation averages over both candidate draws and observation noise.

Comparison of returned sets. For R returned sets of size at most k, fix an ordering of each set and estimate its value by summing its telescoping marginal increments. With

$$b_{\mathrm{val}}(R) = \lceil 32k\varepsilon^{-2} \log(4R/\delta) \rceil$$

fresh observations per increment, each sum has Gaussian variance at most $k/b_{\mathrm{val}}(R)$. A union bound gives simultaneous value error at most $\varepsilon/4$ with probability $1 - \delta/2$. Choosing the largest estimated value loses at most $\varepsilon/2$ relative to the best returned set and costs at most $Rkb_{\mathrm{val}}(R)$ observations. This applies conditionally on all previously returned sets. Two independent unit-variance value observations instead give one marginal difference with variance two and cost two; shared component-function observations are a different interface.

D.2 Uniform full-scan greedy

For $k < n$, a full scan at every round uses $Q_{\mathrm{g}} = nk - k(k-1)/2$ marginal estimates. Average

$$b_{\mathrm{g}} = \lceil 8k^2\varepsilon^{-2} \log(2Q_{\mathrm{g}}/\delta) \rceil \tag{39}$$

fresh observations per estimate. All errors are at most $\varepsilon/(2k)$ with probability $1-\delta$. At each round the chosen marginal is then at least the maximum minus ε/k, giving

$$f(S_{j+1}) \geq (1-1/k)f(S_j) + O/k - \varepsilon/k.$$

Iteration gives $f(S_k) \geq [1-(1-1/k)^k]O - \varepsilon$. The sufficient budget is $Q_{\mathrm{g}} b_{\mathrm{g}}$.

D.3 Threshold greedy with the same final target

The CTG analysis of Chen et al. (2025) permits an additive threshold-test tolerance and a geometric threshold decay. We give a fixed-batch realization of that schedule so its complete observation cost and exact final coefficient can be read directly. Let $\alpha = 1/(4k)$ and $\xi = \varepsilon/(2k)$. Estimate all singleton values and let d_0 be the largest estimate. If $d_0 \leq 0$, return the empty set. Otherwise scan thresholds $d_0, (1-\alpha)d_0, \ldots$ while $w > \alpha d_0/k$, adding an unselected element whenever its estimated current marginal is at least w, until k elements have been selected. All singleton and threshold estimates have simultaneous error at most ξ. The number of threshold scans is at most

$$H = \left\lceil 4k \log(4k^2) \right\rceil,$$

so $n(H+1)$ estimates suffice. Using $\lceil 2\xi^{-2} \log(2n(H+1)/\delta) \rceil$ samples per estimate ensures the accuracy event. Its cost is $O(nk^3\varepsilon^{-2} \log(2k) \log(nk \log(2k)/\delta))$.

If the algorithm fills the budget, an accepted marginal g_j at threshold w satisfies $g_j \geq w - \xi$. A remaining element that was rejected at the preceding threshold has current marginal at most $w/(1-\alpha) + \xi$ by diminishing returns. At the initial threshold the singleton estimates provide the same bound. Therefore

$$g_j \geq \frac{1-\alpha}{k}(O - f(S_j)) - 2\xi.$$

Put $q_0 = 1-(1-\alpha)/k$. For $u = 1/k \in (0,1]$,

$$q_0 = 1 - u + u^2/4 \leq 1 - u + u^2/2 - u^3/6 \leq e^{-u}.$$

Unfolding the recurrence thus yields $f(S_k) \geq (1-q_0^k)O - 2k\xi \geq \rho O - \varepsilon$.

If it stops early, use the last threshold actually scanned, $w_{\mathrm{last}} \leq \alpha d_0/[k(1-\alpha)]$. Each remaining optimal element was rejected at that scan or has since lost marginal value. Thus

$$O - f(S) \leq k(w_{\mathrm{last}} + \xi) \leq \frac{\alpha}{1-\alpha} O + \left(k + \frac{\alpha}{1-\alpha}\right)\xi,$$

where $d_0 \leq a + \xi \leq O + \xi$. Since $\alpha/(1-\alpha) \leq 1/3$, the retained coefficient is at least $2/3 > \rho$ and the additive loss is at most $2k\xi = \varepsilon$. If $d_0 \leq 0$ initially, $O \leq k\xi$ and the empty set qualifies. The displayed cost uses CTG's threshold logic with uniform Gaussian batches. The adaptive Confident Sample implementation below can stop earlier when a marginal is well separated from the threshold.

D.3.1 Adaptive Confident Sample implementation

Use the threshold schedule of Algorithm 2 in Chen et al. (2025), with $\alpha = 1/(4k)$, $\xi = \varepsilon/(2k)$ and $H = \lceil \log(k/\alpha)/\alpha \rceil$. The singleton initialization uses $N_2 = \lceil 2\xi^{-2} \log(6n/\delta) \rceil$ observations per label. Each subsequent threshold decision uses their Confident Sample (CS) rule with local failure budget $\eta = 2\delta/(3nH)$, radius and cap

$$C_t = \sqrt{\frac{2}{t} \log \frac{8t^2}{\eta}}, \qquad N_1 = \left\lceil \frac{2}{\xi^2} \log \frac{4}{\eta} \right\rceil.$$

It accepts when $\overline{Y}_t - C_t \geq w - \xi$, rejects when $\overline{Y}_t + C_t \leq w + \xi$, and otherwise uses $\overline{Y}_{N_1} \geq w$ at the terminal cap. Accepted labels are removed from subsequent scans; their zero current marginal is known from the query definition. The initial scan and every CS call are included in the observation count. If the largest singleton estimate is nonpositive, the procedure returns the empty set, as in Appendix D.3.

Batched implementation. The original CS checks after each scalar observation. We check at $t_1 = 1$, $t_{h+1} = \min\{N_1, \lceil 1.1 t_h \rceil\}$, retaining the original radius and terminal decision. We report observation counts for this *batched CTG-CS* variant.

Same final guarantee. Restricting inspection to deterministic endpoints preserves the CS guarantee. Indeed, summing the two-sided Gaussian tail probabilities over all integer times gives at most $\eta \sum_{t\geq 1}(4t^2)^{-1} < \eta/2$; the fixed terminal mean has error at most ξ except with probability $\eta/2$. Thus an acceptance certifies $\Delta_f(S,i) \geq w - \xi$ and a rejection certifies $\Delta_f(S,i) \leq w + \xi$, except with probability η. These statements hold conditionally on each chosen prefix. The singleton event fails with probability at most $\delta/3$, and at most nH CS calls consume at most $2\delta/3$. On their intersection, exactly the recurrence and early-termination argument in Appendix D.3 apply, yielding $f(C) \geq \rho O - \varepsilon$. The choice $\alpha = 1/(4k)$ recovers the strict ρ coefficient in this threshold recurrence.

D.4 BG-NE WITH ITS OPTIMIZED ERROR SCHEDULE

BG-NE uses all remaining labels at every step. We use the published uniform-confidence, optimized-tolerance policy (Jawanpuria et al., 2024, Theorem C.4). With $j = 0, \ldots, k-1$, put

$$w_j = (1 - 1/k)^{k-1-j}, \quad Z = \sum_{j=0}^{k-1} w_j^{2/3}, \quad \nu_j = \frac{\varepsilon w_j^{-1/3}}{Z}, \quad \delta_j = \delta/k.$$

For $k = 1$, take $w_0 = Z = 1$. Thus $\sum_j w_j \nu_j = \varepsilon$. We calibrate its NE batch to the independent unit-Gaussian feedback:

$$b_j = \left\lceil 4\nu_j^{-2} \log(nk/\delta) \right\rceil.$$

For any arm with gap greater than ν_j, the difference between its sample mean and a best arm's sample mean has variance $2/b_j$. Its probability of beating the best arm is at most $\exp(-b_j\nu_j^2/4)$. A union bound over remaining arms bounds the local failure probability by δ/k; a second union bound and the weighted greedy recurrence give (1). This Gaussian calibration uses BG-NE's full confidence budget and optimized nonuniform tolerances. BG supplies the final-output guarantee used in the comparison; the BSG guarantee averages over random candidate sets.

D.5 A KNOWN-BOUND EXPECTATION ALGORITHM AND OUTPUT SELECTION

An alternative to a supplied bound is a singleton scan giving a valid upper confidence bound $\bar{a} \geq a$. Conditional on its validity, $U = k\bar{a}$ bounds O, and choosing stochastic-greedy precision $\zeta = \varepsilon/(8U)$ limits its expected multiplicative loss to $\varepsilon/8$. The candidate count then includes $\log(U/\varepsilon)$, and noisy estimation receives its own error budget. Attaining the final PAC event also requires an expectation-to-tail conversion. The construction below includes the cost of that conversion; using an estimated bound additionally incurs the scale-scan cost. Proposition 9 uses a lower scale estimate to spend a retained trajectory reserve.

Proposition 16 (Amplification with output selection)**.** *Given $U \geq O$, ordinary random-candidate selection followed by independent repetitions and fresh candidate-value estimation satisfies* (1) *with $\widetilde{O}(nk^2U/\varepsilon^3)$ observations, unless $\rho U \leq \varepsilon$, when the empty set suffices.*

Assume $U \geq O$ is supplied. If $\rho U \leq \varepsilon$, return the empty set. Otherwise $U/\varepsilon > 1/\rho$. Set

$$\zeta = \frac{\varepsilon}{8U}, \qquad \tau_0 = \frac{\varepsilon}{32k}, \qquad \eta = \frac{\varepsilon}{16kU}, \qquad s_0 = \left\lceil \frac{n}{k} \log(1/\zeta) \right\rceil.$$

At each step, sample a uniform candidate subset of size at most s_0 from the remaining labels, using the full remaining set if smaller. Estimate each candidate to radius τ_0 with joint conditional failure probability at most η, and add the largest empirical mean. A fixed batch of $b_0 = \lceil 2\tau_0^{-2} \log(2\min(n, s_0)/\eta) \rceil$ is sufficient. Conditional on the past, the stochastic-greedy expectation lemma (Mirzasoleiman et al., 2015, Lemma 2) and nonnegative marginals imply

$$\mathbb{E}[g_j \mid \text{past}] \geq \frac{1-\zeta}{k}(O - f(S_j)) - 2\tau_0 - U\eta. \tag{40}$$

On a failed estimation event the true chosen marginal is still nonnegative and the candidate maximum is at most $a \leq U$. Consequently, (40) is an unconditional expectation recurrence, not a recurrence conditional on simultaneous success of all noise intervals. It gives

$$\mathbb{E}f(C) \geq (1 - e^{-(1-\zeta)})O - k(2\tau_0 + U\eta) \geq \rho O - \varepsilon/4.$$

There are at most $k\min(n, s_0)$ estimates, so the scalar cost $T_A = k\min(n, s_0)b_0$ is $\widetilde{O}(nk^2/\varepsilon^2)$ with logarithms depending also on U/ε.

Put $t = \rho O - \varepsilon/2$. If $t > 0$, then $0 \le f(C) \le U$ and

$$p := \mathbb{P}\{f(C) \ge t\} \ge \frac{\mathbb{E}f(C) - t}{U - t} \ge \frac{\varepsilon}{4U}.$$

When $t \le 0$, all candidates exceed the threshold. Run the base method independently $R = \lceil (4U/\varepsilon)\log(2/\delta)\rceil$ times. The chance that none exceeds t is at most $\delta/2$. Apply the output-value comparison in Appendix D.1 to these R outputs. The selected set satisfies (1), at total cost

$$R\big(T_A + kb_{\mathrm{val}}(R)\big) = \widetilde{O}(nk^2U/\varepsilon^3),$$

proving Proposition 16. This construction has quadratic cardinality dependence when U/ε is fixed, with its repetition count determined by the expectation-to-tail conversion.

D.6 Boosted stochastic greedy at the strict approximation ratio

The original guarantee and its full repetition cost. Let $X_m = f(C_m) \in [0, O]$ be the values of M independent stochastic-greedy outputs. In the notation of Castillo J. et al. (2023, Theorem 2 and equations (12)–(13)), denote their candidate parameter by ζ, to distinguish it from our additive error. The expectation guarantee gives $\mathbb{E}(O - X_m) \le (e^{-1} + \zeta)O$. For $C > 1$, Markov's inequality and independence imply

$$\mathbb{P}\Big\{\max_{m\le M} X_m < [1 - C(e^{-1} + \zeta)]O\Big\} \le C^{-M}.$$

This re-derives their stated high-probability guarantee with $C = \delta^{-1/M}$. Their Algorithm 1 uses exact evaluations and does not input O or an amplitude bound; O is used only in the analysis. The repetition cost is $O(Mn\log(1/\zeta) + n)$. Expressing it as $O(n\log(1/\delta)\log(1/\zeta))$ keeps $C > 1$ fixed and absorbs the factor $1/\log C$, which diverges as C approaches one.

The stronger finite-round calculation. The ordinary stochastic-greedy recurrence gives

$$\mathbb{E}X_m \ge [1 - \beta_k(\zeta)]O, \qquad \beta_k(\zeta) = \left(1 - \frac{1-\zeta}{k}\right)^k.$$

Using this deficit in the same amplification step gives coefficient $1 - \delta^{-1/M}\beta_k(\zeta)$. For $k \ge 2$, this reaches ρ only if $-1 - \log\beta_k(\zeta) > 0$ and

$$M \ge \frac{\log(1/\delta)}{-1 - \log\beta_k(\zeta)} \ge \frac{\log(1/\delta)}{-1 - k\log(1 - 1/k)} = \Omega(k\log(1/\delta)). \tag{41}$$

The last denominator is $1/(2k) + O(k^{-2})$. Equation (41) specifies the repetition requirement for the strict coefficient in this expectation-based guarantee.

An amplitude-uniform Gaussian realization. We implement the repetition method under (2), including the cost of comparing its outputs. For $2 \le k < n$, set

$$\zeta = \frac{1}{4k}, \qquad s_B = \left\lceil \frac{n}{k}\log(4k)\right\rceil, \qquad b_B = \left\lceil 128k^2\varepsilon^{-2}\log(2n)\right\rceil.$$

In each round of a base run, sample a uniform subset of size $\min\{s_B, n - |S|\}$ from the remaining labels, estimate each marginal with b_B fresh observations, and choose the largest mean. Conditional on the history and candidates, the expected loss from using noisy means is at most

$$2\mathbb{E}\max_{i\in D}|\widehat{\Delta}_i - \Delta_i| \le 2\sqrt{2\log(2n)/b_B} \le \varepsilon/(4k).$$

The middle inequality follows by applying the Gaussian exponential moment bound to the at most $2n$ signed errors and optimizing its parameter. It is uniform in the true marginal means. The usual candidate expectation bound therefore yields the unconditional recurrence

$$\mathbb{E}[g_j \mid \text{past}] \ge \frac{1-\zeta}{k}(O - f(S_j)) - \frac{\varepsilon}{4k}, \qquad \mathbb{E}f(C_m) \ge [1 - \beta_k(\zeta)]O - \varepsilon/4.$$

Writing $d_k = -1 - \log \beta_k(1/(4k))$, the inequality $\log(1-x) \le -x - x^2/2$ gives, for $k \ge 2$,

$$d_k \ge -\frac{1}{4k} + \frac{(1-1/(4k))^2}{2k} \ge \frac{17}{128k} \ge \frac{1}{8k}.$$

Since $O - f(C_m) \ge 0$, Markov's inequality now implies

$$\mathbb{P}\{f(C_m) < \rho O - \varepsilon/2\} \le \frac{\beta_k(\zeta)O + \varepsilon/4}{e^{-1}O + \varepsilon/2} \le e^{-1/(8k)}.$$

The last step uses $\beta_k(\zeta) \le e^{-1-1/(8k)}$ and $1/2 \le e^{-1/(8k)}$. Consequently $M = \lceil 8k \log(2/\delta) \rceil$ independent base runs contain a candidate of value at least $\rho O - \varepsilon/2$ except with probability $\delta/2$.

Apply the same output-value comparison to the M outputs. The selected set satisfies (1), with complete scalar cost

$$M\big[k \min\{n, s_B\} b_B + k b_{\mathrm{val}}(M)\big] = \widetilde{O}(nk^3/\varepsilon^2).$$

For $k = n$ return $[n]$, and for $k = 1$ use the full-scan construction. With exact marginals and exact output-value comparison, the same finite-k amplification uses $\widetilde{O}(nk)$ queries.

Expected value and success probability. A finite-k expected-value surplus can coexist with a success probability of order $1/k$ at the target. For $k \ge 16$, put $V = k^2\varepsilon$, $\rho_k = 1 - (1-1/k)^k$, $z = \rho V - 2\varepsilon$, and let a random variable X take values z and V, with upper probability

$$p_k = \frac{(\rho_k - \rho)V + 3\varepsilon/2}{(1-\rho)V + 2\varepsilon}.$$

Then $\mathbb{E}X = \rho_k V - \varepsilon/2$, yet $\mathbb{P}\{X \ge \rho V - \varepsilon\} = p_k = \Theta(1/k)$. This two-point random-variable example concerns inference from an output's expectation and range.

D.7 Continuous-oracle conversions and objective scale

For a stochastic set function $f = \mathbb{E}_\gamma f_\gamma$, the expected-quality result of Hassani et al. (2019, Theorem 2) has stochastic-oracle cost $O(r^3 d D_f^2/\varepsilon^2)$: r is rank, d is dimension, and $D_f^2 = \mathbb{E}_\gamma[(\max_i f_\gamma(\{i\}))^2]$. Its gradient construction costs $O(d)$ component-value queries and its Hessian estimator reuses the same sampled component function on several sets. Implementing these operations with independent marginal feedback requires charging the observations used to construct them.

The high-probability bound of Becker et al. (2023, Theorem 4) contains terms of order $D\sigma\sqrt{\log(1/\delta)/T} + LD^2/T$, where D is the feasible-set diameter, L is gradient smoothness, and σ controls gradient noise. Their dependence on objective scale persists under a direct marginal-oracle substitution. For $n > k \ge 2$, let

$$f_A(S) = A\mathbf{1}\{S \cap \{1,2\} \ne \varnothing\}, \qquad A > 0.$$

Its multilinear extension $A(x_1 + x_2 - x_1x_2)$ has Hessian norm $L = A$. Sampling a random context with inclusion probabilities x and taking one noisy marginal gives the unbiased coordinate estimate

$$G_1 = A\mathbf{1}\{2 \notin R\} + Z, \quad Z \sim N(0,1), \qquad \mathrm{Var}(G_1) = A^2 x_2(1-x_2) + 1.$$

At $x_2 = 1/2$ the variance is $A^2/4 + 1$; constructing a full vector by this route costs n scalar observations. Thus this substitution retains both objective-scale dependence and a per-vector observation cost.

D.8 Exact-query guarantees and approximation loss

Corollary 5 removes estimation error from the same trajectory argument. Table 4 compares representative guarantees under exact queries. Approximation loss is denoted by ζ to distinguish it from our additive observation tolerance ε.

Setting $\zeta = c/k$ with $c > 0$ in a $(\rho - \zeta)$ guarantee still leaves its coefficient below ρ. A conversion to the strict target requires a separate finite-k argument, supplied for boosted SG in Appendix D.6. Corollary 5 attains this target in $O(n \log(k/\delta))$ queries.

Table 4: Representative exact-oracle guarantees with their accuracy and confidence parameters.

Result	Guarantee and query cost	Strict ρ target
Stochastic greedy (Mirzasoleiman et al., 2015)	Expected value at least $(\rho-\zeta)O$, $O(n\log(1/\zeta))$ queries	An expected-quality guarantee.
Boosted SG (Castillo J. et al., 2023)	Independent SG runs and exact final comparison; Appendix D.6 retains the finite-k surplus	The conversion uses $\widetilde{O}(nk)$ exact queries.
FAST (Breuer et al., 2020, Theorem 1)	Ratio $\rho-4\zeta$ with probability $1-\delta$, explicit inverse powers of ζ and a lower condition on k	The query bound requires $\zeta>0$.
LS+PGB (Chen et al., 2021, Theorem 3)	Ratio $\rho-\zeta$ with probability $1-1/n-p_\alpha$; conditional expected $O(n\log(1/\alpha)/\zeta^2)$ queries after preprocessing	Retains accuracy and confidence dependence.
Clean linear time (Li et al., 2022, Theorems 3.2, 3.6)	Deterministic ratio $\rho-\zeta$ in $O(n/\zeta)$ time	Retains the positive accuracy parameter ζ.
This corollary	Strict ρ, probability $1-\delta$, deterministic $O(n\log(k/\delta))$ marginal queries	At most twice as many exact value queries.

Local-ratio concentration. The high-probability local-ratio bound of Hassidim & Singer (2017a, Theorem 2) has coefficient $1-e^{-(1-t)\mu}$ and condition $k\geq 2/(t^2\mu)$. Even retaining a finite-k product would require $(1-t)\mu\geq k(1-e^{-1/k})=1-1/(2k)+O(k^{-2})$ to reach the exact coefficient ρ. For $\mu\leq 1$, this needs $t=O(1/k)$, incompatible with the displayed size condition along growing k.

E EXPERIMENTAL SETUP AND OBSERVATION ACCOUNTING

The experiments study Algorithm 1 and its local selectors in Section 7. All instances are normalized monotone submodular functions with exact optima available to the evaluator. The learner receives only $n, k, \varepsilon, \delta$ and the fresh marginal-observation oracle. Objective values, optima, feature representations, scales and gaps are not supplied to the algorithm. Gaussian experiments fix $\varepsilon = \delta = 0.05$.

Sampling and complete costs. An increment of b unit-Gaussian observations is generated as its exact sufficient statistic $N(\Delta_f(S, i), 1/b)$ and charged b observations. Independent increments are used at every sampled context. Adaptive decisions occur between increments; failed pilots, fresh fallback and singleton initialization are included in the observation count. For batched CTG-CS, the simulator generates the Gaussian endpoint path and exposes only its first stopping decision. Any unused suffix is discarded and cannot affect subsequent decisions. Elapsed simulation time measures computation on batch statistics rather than the acquisition of individual oracle observations.

E.1 FIXED-GROUND-SET COMPARISON WITH EXACT EXPECTED COST

This comparison fixes $n = 1024$, $\varepsilon = \delta = 0.05$, $k \in \{64, 256, 512\}$, $O/\varepsilon \in \{4, 40\}$, and 100 trials per function/cardinality/amplitude/method combination. The two methods are Algorithm 1 and the full-scan construction in Appendix D.2. The full-scan method averages b_{g} samples per marginal and pays $(nk - k(k-1)/2)b_{\mathrm{g}}$ observations. This comparison isolates candidate reduction under uniform local estimation.

Instances and exact optima. For the modular family, draw 1,024 weights independently and uniformly on $[0.05, 1]$. Sorting yields its exact cardinality optima. For coverage, use 128 disjoint feature blocks, each with eight elements and sixteen features. Element i covers its own private feature and three uniformly chosen distinct shared features among the other eight. Shared-feature weights are uniform on $[0.5, 1.5]$; private-feature weights are independently drawn from that interval and multiplied by 0.3. Blocks are independent. Enumerating the 2^8 subsets in each block gives the best value for each local cardinality, and a dynamic program allocates the global budget. This computes the combinatorial optimum, up to floating-point arithmetic, without assuming greedy optimality. Private and overlapping shared features make the coverage instances nonmodular.

The base instance for each family stays fixed across k. For each cell it is rescaled so that $O = 4\varepsilon$ or 40ε. Both targets are positive. This rescaling is an experimental construction using the external exact optimum; the selector receives neither the scale nor the optimum, weights, features, or block structure. Each trial draws a new label permutation shared by the two methods. Candidate draws and oracle noise have separate streams, also separate between methods. All twelve conditions and their output-quality summaries appear in Table 5.

E.2 SHARED-SELECTOR COMPARISON AND EXACT-QUERY RUNS

For Table 3, the ground set has $n = 16k$ labels. The modular family has k nonzero weights equal to O/k and $15k$ zeros. The second family is layered coverage with $c = 15$, rescaled to the prescribed optimum. Both use $k = 128$ and $O/\varepsilon \in \{4, 40\}$.

Layered coverage. Let $t = 1.001/k$ and $q_t = 1 - t$. The feature universe has k columns, each of total weight one. In each column, layer $\ell = 0, \ldots, k-1$ has weight tq_t^ℓ; the remaining feature has weight q_t^k. There are k column elements covering one entire column, and $c = 15$ duplicate elements for each layer, each covering that layer in all columns, giving $n = 16k$. The optimal k columns cover the entire universe, so $O = k$ and $a = kt = 1.001$. After selecting the first j layers, the next layer has marginal ktq_t^j, strictly exceeding the column marginal q_t^j. Thus exact full greedy selects all layers and has ratio $1 - q_t^k$, equal to 0.63393 at $k = 128$ and 0.63321 at $k = 256$. The perturbation removes reliance on ties between columns and layers. Copies increase the chance of drawing a layer without changing this full-greedy path or the optimum. All these functions are weighted coverage functions and hence normalized monotone submodular.

The full and candidate adaptive methods share the capped selector of Appendix C, with $\nu = \varepsilon/k$ and $\eta = \delta/(3k)$ including initialization. The uniform controls use their stated fixed batches. There

are 10 proved-parameter trials per method and main condition, plus four trials per method at each scale $k = 256, 512$, $O/\varepsilon = 40$. Each trial uses independent label permutations, candidate draws and observation noise. The candidate and full adaptive methods are simulated independently; the pointwise percentile bootstrap intervals in Figure 1b therefore resample them independently (10,000 bootstrap samples). The separate exact-query comparison uses 20 trials per family and method at $k = 128$, judged against $f(C) \geq \rho O$ without additive tolerance.

E.3 Weighted partial-layer coverage

There are k column elements and fifteen copies of each of k layer elements. Column weights a_i form a geometric progression with $\max_i a_i / \min_i a_i = h$ and $\sum_i a_i = k$. Layer ℓ covers a cyclic window W_ℓ of $d = \lceil \omega k \rceil$ columns. Put $t = \gamma/d$ and $p_\ell = t(1-t)^\ell$, $\ell = 0, \ldots, k-1$. In column i, layer ℓ has feature weight $a_i p_\ell$ if $i \in W_\ell$; a residual feature has weight $a_i(1 - \sum_{\ell: i \in W_\ell} p_\ell) \geq 0$. Each column element covers all features in its column. A copy of layer ℓ covers its feature in every column of W_ℓ. This is a normalized monotone weighted-coverage function. The k columns cover the entire universe, proving $O = k$ before scaling. All features are then multiplied by the prescribed O/k.

If A is the selected column set and B is the set of distinct selected layers, write $c_i = \sum_{\ell \in B: i \in W_\ell} p_\ell$. Before the common scaling, the value is $\sum_{i \in A} a_i + \sum_{i \notin A} a_i c_i$. An unselected column has marginal $a_i(1 - c_i)$; an unselected layer has marginal $p_\ell \sum_{i \notin A,\, i \in W_\ell} a_i$. After selecting one copy of a layer, the remaining copies have zero marginal. Hidden label permutations are shared by competing methods; the learner sees neither the feature representation nor these analytical quantities.

To examine marginal-value dispersion, the grid fixes $k = 128$, $O/\varepsilon = 40$ and crosses $h \in \{1, 4, 16\}$, $\omega \in \{1/4, 1/2, 1\}$ and $\gamma \in \{1.001, 2\}$. With uniform column weights, the choice $t = \gamma/d$ keeps the leading layer's initial marginal equal to γ across coverage fractions, while changing the layer decay. With heterogeneous weights, window composition also matters. The ω comparison therefore captures the joint effect of coverage fraction, layer decay and window composition within this family. The uniform full-overlap case with $\gamma = 1.001$ is exactly the layered family in Appendix E.2. Two weighted modular controls have $h = 4, 16$, k nonzero weights and $15k$ zero weights.

Layer strength and estimation tolerance. For uniform columns and full overlap, the first layer exceeds a column by $(\gamma - 1)O/k$. At $\gamma = 1.001$, this is only 0.004 or 0.04 times the candidate selector's tolerance ε/k for $O/\varepsilon \in \{4, 40\}$, respectively. Consequently, a locally approximate selector can choose columns instead of following exact greedy's near-ρ trajectory. The strength level $\gamma = 2$ makes this difference visible relative to the local tolerance and changes the exact-greedy trajectory accordingly.

Each of the 18 coverage conditions and two weighted modular controls uses six independent trials per method, for 360 proved-parameter runs. Hidden label permutations are paired across the three methods; their noise streams are independent. Methods use the proved parameters of the candidate selector in Appendix C and the baselines in Appendices D.3.1 and D.4. Pointwise percentile bootstrap cost-ratio intervals resample the six paired instances 10,000 times. Figure 2 presents the 12 heterogeneous coverage conditions; Appendix F.3 supplies the remaining controls and the query-pair decomposition.

F Supplementary cost, quality and mechanism evidence

F.1 Uniform sampling: predicted costs and attained quality

Proposition 17 (Exact expected observation count)**.** *Put* $Q_{\mathrm{g}} = nk - k(k-1)/2$ *and* $p_s = 1 - (1 - 1/n)^s$. *For* $k < n$ *and* $s < n$, *Algorithm 1 satisfies*

$$\mathbb{E}_f T = b\big[n + p_s(Q_{\mathrm{g}} - n)\big] \qquad (f \in \mathcal{F}_n). \tag{42}$$

For $s \geq n$, *its cost is* $Q_{\mathrm{g}} b$ *deterministically; for* $k = n$, *it is zero.*

Proof of Proposition 17. After j selections each unselected label is drawn with conditional probability p_s, so $\mathbb{E}[|D_j| \mid \mathcal{H}_j] = (n - j)p_s$. Sum over $j = 1, \ldots, k-1$ and add the initial n estimates. The full-scan and $k = n$ cases follow directly from the algorithm. □

Table 5 pairs the cost calibration in Figure 1a with output quality for all twelve conditions. The full-scan costs at $k = 64, 256, 512$ are respectively $1.228 \cdot 10^{13}$, $7.717 \cdot 10^{14}$ and $5.470 \cdot 10^{15}$ scalar observations. All 2,400 runs meet their quality target; the pointwise one-sided 95% upper failure endpoint for each 100-run method-cell is $1 - 0.05^{1/100} \simeq 0.0295$.

Table 5: Uniform-estimation comparison at $n = 1024$, 100 trials per method and cell. Ratios are candidate/full-scan observation costs; quality is the minimum $f(C)/O$. Expected ratios are from Proposition 17.

Family	k	O/ε	Expected ratio	Observed ratio	Min. quality (cand./full)
Modular	64	4	1.0000	1.0000	0.998 / 0.997
Modular	64	40	1.0000	1.0000	1.000 / 1.000
Modular	256	4	0.3033	0.3033	0.999 / 0.999
Modular	256	40	0.3033	0.3034	1.000 / 1.000
Modular	512	4	0.1601	0.1600	0.999 / 1.000
Modular	512	40	0.1601	0.1600	1.000 / 1.000
Block coverage	64	4	1.0000	1.0000	0.995 / 0.995
Block coverage	64	40	1.0000	1.0000	1.000 / 1.000
Block coverage	256	4	0.3033	0.3033	0.974 / 0.974
Block coverage	256	40	0.3033	0.3034	0.976 / 0.980
Block coverage	512	4	0.1601	0.1601	0.990 / 0.992
Block coverage	512	40	0.1601	0.1601	0.992 / 0.993

F.2 Local adaptation, scale and exact queries

The methods and families are specified in Appendix E.2. Figure 3 decomposes the costs in Table 3 into initialization, subsequent sampling and subsequent fallback.

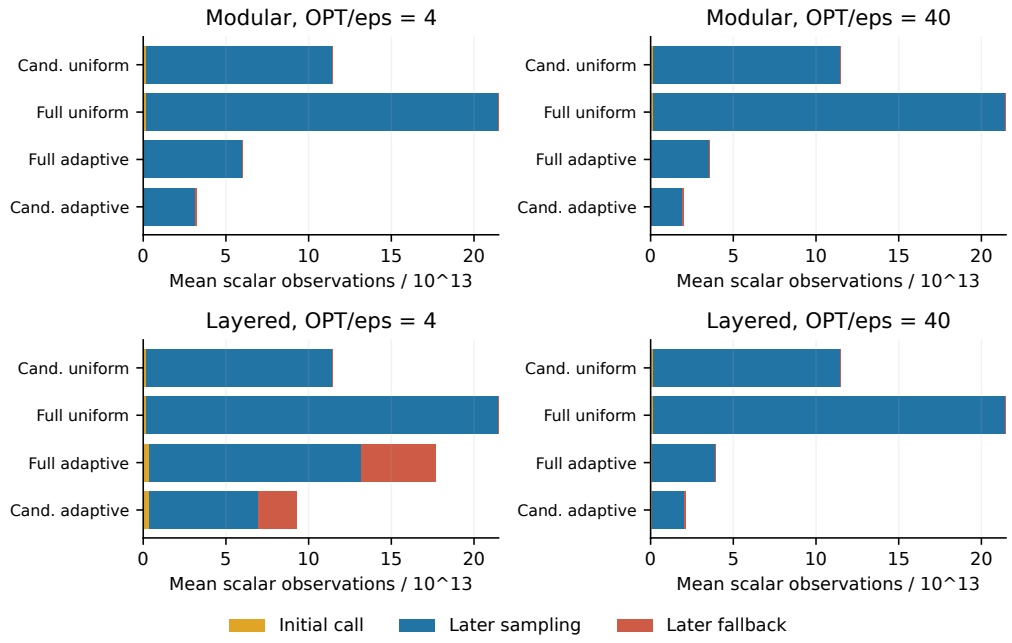

Figure 3: Proved-parameter cost decomposition, 10 trials per method and condition. The initial call includes its possible fallback; the other two parts cover subsequent calls. These three parts are disjoint and sum to the complete cost.

All 160 main runs, 64 scale runs and 80 exact-query runs meet their respective quality targets. For layered coverage, exact full greedy follows the slightly better layers and attains $1 - (1 - 1.001/k)^k$

of the optimum. Noisy approximate choices can instead select columns and attain larger final value. Thus estimation precision can change the greedy path as well as its observation cost.

F.3 Complete costs, output quality and structural controls

Appendix E.3 defines the weighted coverage family and its exact optimum. Table 6 pairs the absolute costs in Figure 2 with output quality for all twelve heterogeneous conditions. Table 7 adds the six equal-weight coverage and two weighted-modular controls to the heterogeneous coverage conditions in Figure 2. All 360 returned sets in this structural study meet the quality target.

Table 6: Absolute costs and output quality for all twelve conditions in Figure 2. Each method has six trials per condition, with $n = 2048$, $k = 128$, $O/\varepsilon = 40$ and $\varepsilon = \delta = 0.05$. Cost is the mean number of scalar observations in units of 10^{13}; quality is the minimum observed $f(C)/O$. The normalized PAC threshold is $\rho - 1/40 \simeq 0.6071$.

h	ω	γ	Candidate adaptive		CTG-CS		BG-NE	
			Cost	Min. quality	Cost	Min. quality	Cost	Min. quality
4	1/4	1.001	0.216	0.996	0.509	1.000	4.009	1.000
4	1/4	2	0.367	0.974	1.675	0.974	4.009	0.974
4	1/2	1.001	0.274	0.996	0.505	1.000	4.009	1.000
4	1/2	2	0.592	0.941	1.697	0.941	4.009	0.939
4	1	1.001	0.268	0.993	0.507	1.000	4.009	1.000
4	1	2	1.696	0.867	4.777	0.867	4.009	0.867
16	1/4	1.001	0.269	0.998	1.191	1.000	4.009	1.000
16	1/4	2	0.555	0.988	3.230	0.988	4.009	0.988
16	1/2	1.001	0.317	0.998	1.166	1.000	4.009	1.000
16	1/2	2	0.459	0.983	2.032	0.984	4.009	0.983
16	1	1.001	0.354	0.997	1.196	1.000	4.009	1.000
16	1	2	0.400	0.997	1.279	1.000	4.009	1.000

Table 7: Controls for the heterogeneous-coverage comparison, with six trials per method and condition. Ratios compare candidate adaptive with each baseline; brackets give pointwise paired-bootstrap 95% intervals. Values below one favor candidate adaptive. h is weight spread and (ω, γ) gives coverage fraction and layer strength.

Family	h	(ω, γ)	O/ε	Cand./CTG [95% CI]	Cand./BG [95% CI]
Coverage	1	$(0.25, 1.001)$	40	6.300 [5.805, 6.867]	0.367 [0.358, 0.376]
Coverage	1	$(0.25, 2)$	40	0.979 [0.952, 1.007]	0.274 [0.266, 0.281]
Coverage	1	$(0.5, 1.001)$	40	8.690 [8.059, 9.587]	0.535 [0.506, 0.559]
Coverage	1	$(0.5, 2)$	40	0.476 [0.469, 0.484]	0.207 [0.204, 0.211]
Coverage	1	$(1, 1.001)$	40	8.460 [7.590, 9.542]	0.503 [0.493, 0.515]
Coverage	1	$(1, 2)$	40	0.356 [0.355, 0.358]	0.416 [0.414, 0.417]
Modular	4	–	40	0.615 [0.491, 0.681]	0.072 [0.058, 0.080]
Modular	16	–	40	0.248 [0.196, 0.302]	0.068 [0.053, 0.082]

For each realized trajectory let Q_{obs} count distinct sampled pairs (S, i) and let T count all scalar observations. Repeated tests or elimination stages at the same pair increase T without increasing Q_{obs}. Ratios of group means satisfy

$$\frac{\overline{T}_{\mathrm{cand}}}{\overline{T}_{\mathrm{CTG}}} = \frac{\overline{Q}_{\mathrm{cand}}}{\overline{Q}_{\mathrm{CTG}}} \frac{\overline{T}_{\mathrm{cand}}/\overline{Q}_{\mathrm{cand}}}{\overline{T}_{\mathrm{CTG}}/\overline{Q}_{\mathrm{CTG}}}.$$

Table 8 reports these two factors. CTG can accept many equal-weight columns at a common threshold. Unequal weights spread acceptances over more thresholds and increase its queried pairs. The schedule traverses a weight factor h in approximately $\log(h)/[-\log(1 - 1/(4k))] \simeq 4k \log(h)$ levels. The samples-per-pair factor also reflects repeated tests and local elimination along each method's own trajectory.

Table 8: Cost decomposition at full overlap, layer strength 1.001, $k = 128$ and $O/\varepsilon = 40$. The pair ratio times the samples-per-pair ratio equals the complete candidate/CTG cost ratio.

Weight spread	Cand./CTG cost	Pair ratio	Samples/pair ratio	CTG levels
1	8.460	11.063	0.765	7.7
4	0.528	0.510	1.035	709.5
16	0.296	0.509	0.581	1414.3

F.4 Signed deficits along the adaptive trajectory

The layered construction of Appendix E.2 also illustrates the signed recurrence. For $j \geq 1$, let M_j be the actual candidate maximum (zero for an empty set), g_{j+1} the selected true marginal, μ_j the top-k mean, and $w_j = q^{k-1-j}$. Define

$$R = (\rho_k - \rho)O + q^{k-1}(f(S_1) - O/k),$$
$$D_+ = \sum_{j=1}^{k-1} w_j(\mu_j - M_j)_+, \qquad D_- = \sum_{j=1}^{k-1} w_j(M_j - \mu_j)_+,$$
$$E = \sum_{j=1}^{k-1} w_j(M_j - g_{j+1}), \qquad H = \sum_{j=1}^{k-1} w_j \left(\mu_j - \frac{O - f(S_j)}{k}\right) \geq 0.$$

Unfolding the recurrence gives the exact identity

$$f(S_k) - \rho O = R - D_+ + D_- - E + H. \tag{43}$$

The term E retains its sign, since a feasible addition after an empty candidate set can have positive marginal. The actual maximum M_j is distinct from the auxiliary maximum W_j used in the tail proof. Across all 1,120 simulated trajectories, the decomposition agrees with the final values to within $2 \cdot 10^{-13}$. Corollary 4 controls the proof's auxiliary maxima; Figure 4 illustrates the signed cancellation on realized paths.

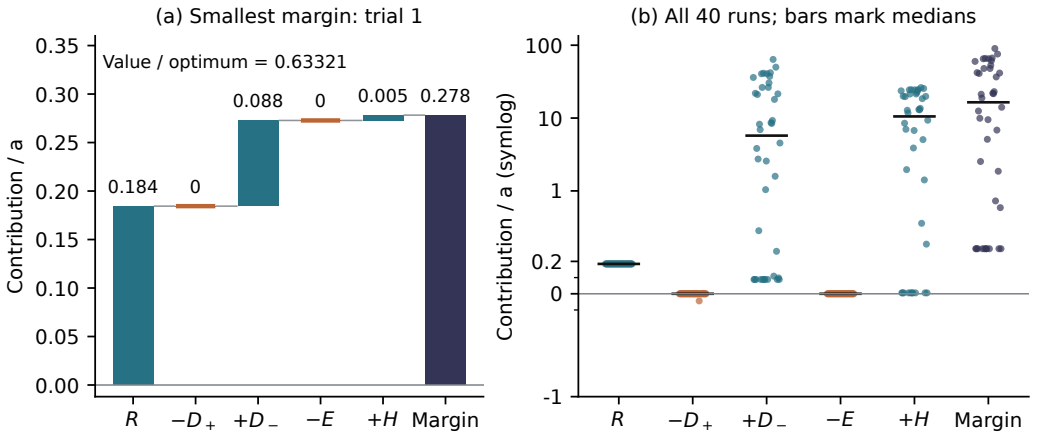

Figure 4: Signed trajectory decomposition on layered coverage, $c = 15$, $k = 256$, exact queries, 40 runs of the initial-scan method. Left: the run with the smallest final margin (ties by trial index). Right: all 40 runs and component medians. All quantities are divided by a; the symmetric log axis is linear for magnitudes at most 0.5.

In the forty exact-query runs of Figure 4, positive local deficits occur in 30 runs, while every final signed deficit $D_+ - D_-$ is negative. This decomposition shows how negative increments offset positive ones within a trajectory.