Identifying Approximate Minimizers Under Stochastic Uncertainity

Hessa Al-Thani
Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor, MI, USA

Viswanath Nagarajan
Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor, MI, USA

Abstract

We study a fundamental stochastic selection problem involving n independent random variables, each of which can be queried at some cost. Given a tolerance level δ, the goal is to find a δ-approximately minimum (or maximum) value over all the random variables, at minimum expected cost. A solution to this problem is an adaptive sequence of queries, where the choice of the next query may depend on previously-observed values. Two variants arise, depending on whether the goal is to find a δ-minimum value or a δ-minimizer. When all query costs are uniform, we provide a 4-approximation algorithm for both variants. When query costs are non-uniform, we provide a 5.83-approximation algorithm for the δ-minimum value and a 7.47-approximation for the δ-minimizer. All our algorithms rely on non-adaptive policies (that perform a fixed sequence of queries), so we also upper bound the corresponding "adaptivity" gaps. Our analysis relates the stopping probabilities in the algorithm and optimal policies, where a key step is in proving and using certain stochastic dominance properties.

2012 ACM Subject Classification Theory of computation → Stochastic approximation; Theory of computation → Discrete optimization

Keywords and phrases Approximation algorithms, stochastic optimization, selection problem

Digital Object Identifier 10.4230/LIPIcs.ICALP.2025.8

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: on arXiv [2]

Funding *Hessa Al-Thani*: This publication was made possible by the Graduate Sponsorship Research Award from the Qatar Research and Development Institute. The findings herein reflect the work, and are solely the responsibility, of the authors.
Viswanath Nagarajan: Research supported in part by NSF grant CCF-2418495.

1 Introduction

We study a natural stochastic selection problem that involves querying a set of random variables so as to identify their minimum (or maximum) value within a desired precision. Consider a car manufacturer who wants to chose one design from n options so as to optimize some attribute (e.g., maximum velocity or energy efficiency). Each option i corresponds to an attribute value X_i which is uncertain and drawn from a known probability distribution. It is possible to determine the *exact* value of X_i by further testing – but this incurs some cost c_i. Identifying the exact minimum (or maximum) value among the X_is might be too expensive. Instead, our goal is to identify an *approximately* minimum (or maximum) value, within a prescribed tolerance level. For example, we might be satisfied with a value (and corresponding option) that is within 10% of the true minimum. The objective is to minimize the expected cost. In this paper, we provide the first constant-factor approximation algorithm for this problem.

© Hessa Al-Thani and Viswanath Nagarajan;
licensed under Creative Commons License CC-BY 4.0

Our problem is related to two lines of work: *stochastic combinatorial optimization* and *optimization under explorable uncertainty*. In stochastic combinatorial optimization, a solution makes selections incrementally and adaptively (i.e., the next selection can depend on previously observed random outcomes). An optimal solution here may even require exponential space to describe. Nevertheless, there has been much recent success in obtaining (efficient) approximation algorithms for such problems, see e.g., [7, 14, 5, 15, 12, 19, 20, 17, 18]. Optimization problems under explorable uncertainty involve querying values drawn from known intervals in order to identify a minimizer. Typically, these results focus on the *competitive ratio*, which relates the algorithm's (expected) query cost to the optimum query-cost in hindsight, see e.g., [21, 6, 10, 9, 22, 8, 4, 23]. In particular, for the problem of finding an *exact* minimizer among n intervals, [21] obtained a 2-competitive algorithm in the adversarial setting and [6] obtained a 1.45-approximation algorithm in the stochastic setting. The problem we study is a significant generalization of the stochastic exact minimizer problem [6].

1.1 Problem Definition

In the *stochastic minimum query* (SMQ) problem, there are n independent discrete random variables $X_1, ..., X_n$ that lie in intervals $I_1, ..., I_n$ respectively. The random variables (r.v.s) may be negative. We assume that each interval is bounded and closed, i.e., $I_j = [\ell_j, r_j]$ for each $j \in [n]$. We also assume (without loss of generality) that each r.v. has non-zero probability at the endpoints of its interval, i.e., $\Pr[X_j = \ell_j] > 0$ and $\Pr[X_j = r_j] > 0$ for each $j \in [n]$.[1] We will use the terms random variable (r.v.) and interval interchangeably. The exact value of any r.v. X_j can only be determined by querying it, which incurs some cost $c_j \geq 0$. Additionally, we are given a "precision" value $\delta \geq 0$, where the goal is to identify the minimum value over *all* r.v.s up to an additive precision of δ. Formally, if $\mathsf{MIN} = \min_{j=1}^{n} X_j$ then we want to find a deterministic value VAL such that $\mathsf{MIN} \leq \mathsf{VAL} \leq \mathsf{MIN} + \delta$. Such a value VAL is called a *δ-minimum* value. The objective in SMQ is to minimize the expected cost of the queried intervals. Note that it may be sufficient to probe only a (small) subset of intervals before stopping.

We also consider a related, but harder, problem where the goal is to *identify* some δ-minimizer $i^* \in [n]$, i.e., an interval that satisfies $X_{i^*} \leq \mathsf{MIN} + \delta$. We refer to this problem as *stochastic minimum query for identification* (SMQI). If a δ-minimum value is found then it also provides a δ-minimizer (see §1.4). However, the converse is not true. So, an SMQI solution may return an un-queried a δ-minimizer i^* without determining a δ-minimum value.

Although our formulation above uses *additive* precision (we aim to find a value that is at most $\mathsf{MIN} + \delta$), we can also handle *multiplicative* precision where the goal is to find a value that is at most $\alpha \cdot \mathsf{MIN}$. This just requires a simple logarithmic transformation; see Appendix A. We can also handle the goal of finding the *maximum* value by working with negated r.v.s $\{-X_i\}_{i=1}^{n}$.

Throughout, we use $N := [n] = \{1, 2, \ldots, n\}$ to denote the index set of the r.v.s.

Adaptive and Non-adaptive policies. Any solution to SMQ involves querying r.v.s sequentially until a δ-minimum value is found. In general, the sequence of queries may depend on the realizations of previously queried r.v.s. We refer to such solutions as *adaptive* policies. Formally, such a solution can be described as a decision tree where each node corresponds to

[1] Otherwise, we can just work with a smaller interval representing the same r.v.

the next r.v. to query and the branches out of a node represent the realization of the queried r.v. *Non-adaptive* policies are a special class of solutions where the sequence of queries is fixed upfront: the policy then performs queries in this order until a δ-minimum value is found. A central notion in stochastic optimization is the *adaptivity gap* [7], which is the worst-case ratio between the optimal non-adaptive value and the optimal adaptive value. All our algorithms produce non-adaptive policies and hence also bound the adaptivity gap.

1.2 Results

Our first result is on the SMQ problem with unit costs, for which we provide a 4-approximation algorithm. Moreover, we achieve this result via a non-adaptive policy, which also proves an upper bound of 4 on the adaptivity gap. This algorithm relies on combining two natural policies. The first policy simply queries the r.v. with the smallest left-endpoint. The second policy queries the r.v. that maximizes the probability of stopping in the very next step. When used in isolation, both these policies have unbounded approximation ratios. However, interleaving the two policies leads to a constant-factor approximation algorithm.

We also consider the (harder) unit-cost SMQI problem and show that the same policy leads to a 4-approximation algorithm: the only change is in the criterion to stop, which is now more relaxed. While the algorithm is the same as SMQ, the analysis for SMQI is significantly more complex due to the new stopping criterion, which allows us to infer a δ-minimizer i^* even when it has not been queried. Specifically, we prove a stochastic dominance property between r.v.s in our algorithm and the optimum (conditioned on the SMQ stopping criterion not occurring), and use this in relating the SMQI stopping-probability in the algorithm and the optimum.

Our next result is for the SMQ problem with non-uniform costs. We obtain a constant-factor approximation again, with a slightly worse ratio of 5.83. This is based on combining ideas from the unit-cost algorithm with a "power-of-two" approach. In particular, the algorithm proceeds in several iterations, where the i^{th} iteration incurs cost roughly 2^i. In each iteration i, the algorithm selects a subset of r.v.s with cost $O(2^i)$ based on the following two criteria (i) smallest left-endpoint and (ii) maximum probability of stopping in one step. In order to select the r.v.s for criterion (ii) we need to use a PTAS for an appropriate version of the knapsack problem.

Finally, we consider the SMQI problem with non-uniform costs. Directly using the SMQ algorithm for SMQI (as in the unit-cost case) does not work here: it leads to a poor approximation ratio. However, a modification of the SMQ algorithm works. Specifically, we modify step (i) above: instead of just selecting a prefix of intervals with the smallest left-endpoints, we select an "almost prefix" set by skipping some expensive intervals. We show that this approach leads to an approximation ratio of 7.47, which is slightly worse than what we obtain for SMQ. The analysis combines aspects of unit-cost SMQI and SMQ with non-uniform costs.

1.3 Related Work

Computing an approximately minimum or maximum value by querying a set of random variables is a central question in stochastic optimization. Most of the prior works on this topic have focused on *budgeted* variants. Here, one wants to select a subset of queries of total cost within some budget so as to maximize or minimize the value *among the queried* r.v.s. The results for the minimization and maximization versions are drastically different. A $1 - \frac{1}{e}$ approximation algorithm for the budgeted max-value problem follows from results

on stochastic submodular maximization [3]; more complex "budget" constraints can also be handled in this setting [1, 16]. These results also bound the adaptivity gap. In addition, PTASes are known for non-adaptive and adaptive versions of budgeted max-value [11, 24]. For the budgeted min-value problem, it is known that the adaptivity gap is unbounded and results for the non-adaptive and adaptive versions are based on entirely different techniques. [13] obtained a bi-criteria approximation algorithm for the non-adaptive problem (the queried subset must be fixed upfront) that achieves a $1+\epsilon$ approximation to the optimal value while exceeding the budget by at most an $O(\log\log m)$ factor, where each r.v. takes an integer value in the range $\{0, 1, \ldots, m\}$. Subsequently, [26] studied the adaptive setting (the queried subset may depend on observed realizations) and obtained a 4-approximation while exceeding the budget by at most an $O(\log\log m)$ factor. In contrast to these results, the goal in SMQ is to achieve a value close to the true minimum/maximum taken over *all* random variables $X_1, X_2, \ldots, X_n$ (not just the queried ones). Moreover, we want to find an approximately min/max value with probability one, as opposed to optimizing the expected min/max value.

A different formulation of the minimum-element problem is studied in [25]: this combines the query-cost and the value of the minimum-queried element into a single objective. They obtain an exact algorithm for this setting, which also extends to a wider class of constrained problems.

Closely related to our work, [6] studied the SMQI problem with exact precision, i.e., $\delta = 0$. In particular, their goal is to identify an *interval* that is an exact minimizer. [6] obtained a 1.45-approximation ratio for general query costs. The SMQI problem that we study allows for arbitrary precision δ, and is significantly more complex than the setting in [6]. One indication of the difficulty of handling arbitrary δ is that the simpler SMQ problem with $\delta = 0$ (where we want to find the exact minimum value) admits a straightforward exact algorithm that queries by increasing left-endpoint; however, this algorithm has an unbounded ratio for SMQ with arbitrary δ (see §2 for an example).

As mentioned earlier, the SMQ problem is also related to optimization problems under explorable uncertainty. Apart from the minimum-value problem [21], various other problems like computing the median [10], minimum spanning tree [9, 22] and set selection [8, 4, 23] have been studied in this setting. The key difference from our work is that these results focus on the competitive ratio. In contrast, we compare to the optimal policy that is limited in the same manner as the algorithm. We note that there is an $\tilde{\Omega}(n)$ lower bound on the competitive ratio for SMQ and SMQI; see Appendix B. Our results show that much better (constant) approximation ratios are achievable for SMQ and SMQI in the stochastic setting, relative to an optimal policy.

1.4 Preliminaries

Stopping rule for SMQ. Even without querying any interval, we know that the minimum value is at most $R := \min_{i\in N}\{r_i\}$, the minimum right-endpoint. In order to simplify notation, we incorporate this information using a dummy r.v. $X_0 = [R, R]$ that is queried at the start of any policy and incurs no cost. We now formally define the condition under which a policy for SMQ is allowed to stop. We will refer to the partial observations at any point in a policy (i.e., values of r.v.s queried so far) as the *state*. Consider any state, given by a subset $S \subseteq N$ of queried r.v.s along with their observations $\{x_i\}_{i\in S}$. The minimum observed value is $\min_{i\in S} x_i$ and the minimum possible value among the un-queried r.v.s is $\min_{j\in N\setminus S} \ell_j$. The stopping criterion is:

$$\min_{i\in S} x_i \quad \leq \quad \min_{j\in N\setminus S} \ell_j \; + \; \delta. \tag{1}$$

If this criterion is met then $\mathsf{VAL} = \min_{i \in S} x_i$ is guaranteed to satisfy $\mathsf{MIN} \leq \mathsf{VAL} \leq \mathsf{MIN}+\delta$, where $\mathsf{MIN} = \min_{j \in N} X_j$. Also, $\arg\min_{i \in S} x_i$ is a δ-minimizer. On the other hand, if this criterion is not met then there is no value v that guarantees $\mathsf{MIN} \leq v \leq \mathsf{MIN} + \delta$: there is a non-zero probability that the minimum value is $\min_{j \in N \setminus S} \ell_j$ or $\min_{i \in S} x_i$ (and these values are more than δ apart). So,

▶ **Proposition 1.1.** *A policy for* SMQ *can stop if and only if criterion* (1) *holds.*

The stopping rule for SMQI is described in §2.2. An SMQI policy can stop either due to the SMQ stopping rule (above) or by inferring an un-queried interval i^* as a δ-minimizer.

Adaptivity gap. We show that the adaptivity gap for the SMQ problem is more than one: so adaptive policies may indeed perform better. This example also builds some intuition for the problem. Consider an instance $\mathcal{I}$ with three intervals as shown in Figure 1. In particular, $X_1 \in \{0, 3, \infty\}$, $X_2 \in \{1, \infty\}$, $X_3 \in \{2, \infty\}$ and $\delta = 1$. Let $\Pr(X_1 = 0) = \frac{1}{3}, \Pr(X_1 = 3) = \frac{1}{3}, \Pr(X_1 = \infty) = \frac{1}{3}, \Pr(X_2 = 1) = \epsilon, \Pr(X_2 = \infty) = 1 - \epsilon, \Pr(X_3 = 2) = 1 - \epsilon, \Pr(X_3 = \infty) = \epsilon$. An adaptive policy is shown in Figure 2, which has cost at most $1 + \frac{2}{3} + \frac{\epsilon}{3} = \frac{5+\epsilon}{3}$. We present a case analysis in [2] and show that the best non-adaptive cost is $\min\left\{\frac{6-\epsilon}{3}, \frac{5+2\epsilon}{3}\right\}$. Setting $\epsilon = \frac{1}{3}$, we obtain an adaptivity gap of $\frac{17}{16}$. We can also modify this instance slightly to get a worse adaptivity gap of $\frac{12}{11}$.

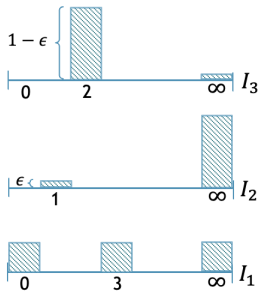

Figure 1 Adaptivity gap instance.

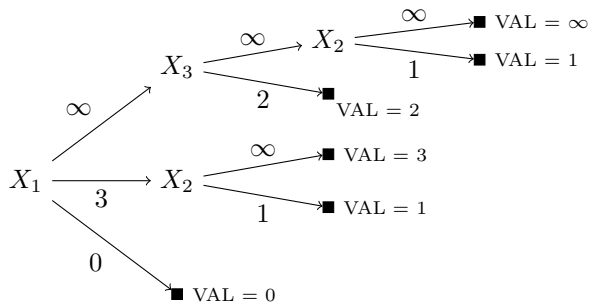

Figure 2 Optimal adaptive policy.

Fixed threshold problem. In our analysis, we relate SMQ to the following simpler problem. Given n r.v.s $\{X_i : i \in N\}$ with costs as before, a *fixed* threshold θ and budget k, find a policy having query-cost at most k that maximizes the probability of observing a realization less than θ. A useful property of this fixed threshold problem is that it has adaptivity gap one [2].

▶ **Proposition 1.2.** *Consider any instance of the fixed threshold problem. Let V^* and F^* denote the maximum success probabilities over adaptive and non-adaptive policies respectively. Then, $V^* = F^*$*

2 Algorithm for Unit Costs

Before presenting our algorithm, we discuss two simple greedy policies and show why they fail to achieve a good approximation.

1. A natural approach is to select intervals by increasing left-endpoint. Indeed, [21] shows that this algorithm is optimal when $\delta = 0$, even in an online setting (with open intervals). Consider the instance with two types of intervals as shown in Figure 3. The r.v.s

$X_1, \cdots X_{n/2}$ are identically distributed with $X_i = 0$ w.p. $\frac{1}{n}$ and $X_i = n$ otherwise. The remaining r.v.s $X_{n/2+1}, \cdots X_n$ are identically distributed with $X_i = \frac{\delta}{2}$ w.p. $\frac{1}{2}$ and $X_i = n$ otherwise. The greedy policy queries r.v.s in the order $1, 2, \cdots n$, resulting in an expected cost of $\Omega(n)$ as it can stop only when it observes a "low" realization for some r.v. However, the policy that probes in the reverse order $n, n-1, \cdots 1$ has constant expected cost: the policy can stop upon observing *any* "low" realization (even if a value of $\delta/2$ is observed, it is guaranteed to be within δ of the true minimum). So the approximation ratio of this greedy policy is $\Omega(n)$.

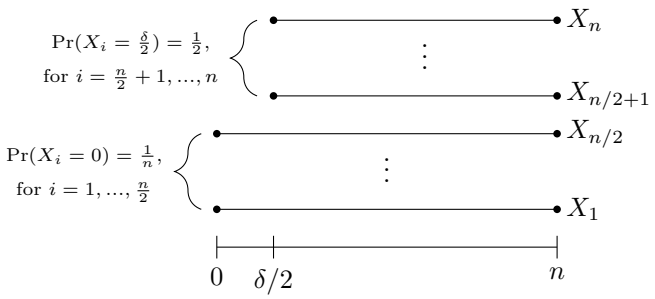

Figure 3 Bad example for greedy by left-endpoint.

2. A different greedy policy (based on the instance in Figure 3) is to always select the interval that maximizes the likelihood of stopping in one step. Now consider another instance with three types of intervals; see Figure 4. The r.v. X_n is always 1.4δ. The r.v. X_1 takes value 0 w.p. $\frac{1}{2n}$ and has value n otherwise. The remaining r.v.s $X_2, \cdots X_{n-1}$ are identically distributed with $X_i = \frac{\delta}{2}$ w.p. $\frac{1}{n}$ and $X_i = n$ otherwise. As long as X_1 is not queried, the probability of stopping (in one step) is as follows: $\frac{1}{2n}$ for X_1, $\frac{1}{n}$ for $X_2, \cdots X_{n-1}$ and zero for X_n. So this greedy policy will query in the order $2, 3, \cdots n-1, 1, n$ resulting in an $\Omega(n)$ expected cost. On the other hand, querying the r.v.s X_1 and X_n guarantees that the policy can stop. So the optimal cost is at most 2, implying an $\Omega(n)$ approximation ratio.

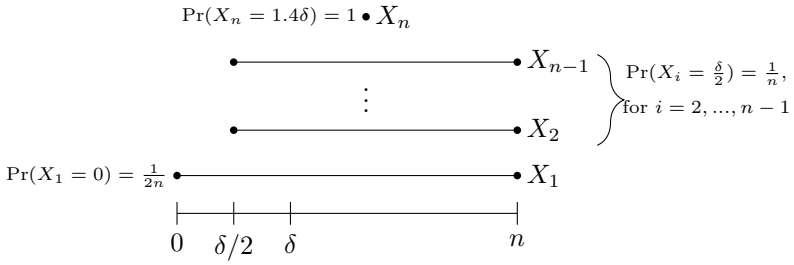

Figure 4 Bad example for greedy by stopping probability.

Our approach is to interleave the above two greedy criteria. In particular, each iteration of our algorithm makes two queries: the interval with the smallest left-endpoint and the interval that maximizes the probability of stopping in one step. We will show that this leads

to a constant-factor approximation. We first re-number intervals by increasing order of their left-endpoint, i.e., $\ell_1 \leq \ell_2 \leq \cdots \ell_n$. For each $k \in N$, let $\theta_k := \ell_{k+1} + \delta$. Algorithm 1 describes our algorithm formally.

Algorithm 1 Non-Adaptive Double Greedy.

Let $\ell^* = \min_{i \in N} \ell_i$, $m^* = R := \min_{i=1}^n r_i$, and $\pi \leftarrow \emptyset$.
for $j = 1, \ldots, n$ **do** ▷ iterations
  Query interval j (if not already in π).
  Query interval $b(j) = \text{argmax}_{i \in N \setminus (\pi \circ j)} \Pr[X_i \leq \theta_j]$.
  Update list $\pi \leftarrow \pi \circ j \circ b(j)$. ▷ skip j if it was already in π
  Update $m^* = \min\{m^*, X_j, X_{b(j)}\}$ and $\ell^* = \min_{i \in N \setminus \pi}\{\ell_i\}$.
  if $m^* - \ell^* \leq \delta$ **then** stop.

Equivalently, we can view Algorithm 1 as first computing the permutation π (without querying) and then performing queries in the order given by π until the stopping criterion is met. Note that Algorithm 1 is non-adaptive because it uses observations only to determine when to stop. So, our analysis also upper bounds the adaptivity gap.

We overload notation slightly and use π to also denote the non-adaptive policy given in Algorithm 1. Note that each *iteration* in this policy involves *two* queries. We use σ to denote the optimal (adaptive) policy. Let $c_{exp}(\pi)$ and $c_{exp}(\sigma)$ denote the expected number of queries in policies π and σ, respectively. The key step in the analysis is to relate the termination probabilities in these two policies, formalized below.

▶ **Lemma 2.1.** *For any $k \geq 1$, we have*

$$\Pr[\sigma \text{ finishes in } k \text{ queries}] \leq \Pr[\pi \text{ finishes in } 2k \text{ iterations}].$$

We will prove this lemma in the next subsection. First, we complete the analysis using this.

▶ **Theorem 2.2.** *We have $c_{exp}(\pi) \leq 4 \cdot c_{exp}(\sigma)$.*

Proof. Let C_σ denote the random variable that captures the number of queries made by the optimal policy σ. Similarly, let C_π denote the number of queries made by our policy. Using Lemma 2.1 and the fact that policy π makes two queries in each iteration, for any $k \geq 1$ we have

$$\Pr[C_\sigma \leq k] = \Pr[\sigma \text{ finishes in } k \text{ queries}] \leq \Pr[\pi \text{ finishes in } 2k \text{ iterations}] \leq \Pr[C_\pi \leq 4k] \quad (2)$$

Hence,

$$\begin{aligned} c_{exp}(\sigma) &= \int_0^\infty \Pr[C_\sigma > t]dt = \int_0^\infty (1 - \Pr[C_\sigma \leq t])\, dt \geq \int_0^\infty (1 - \Pr[C_\pi \leq 4t])dt \\ &= \frac{1}{4}\int_0^\infty (1 - \Pr[C_\pi \leq y])dy = \frac{1}{4}\int_0^\infty \Pr[C_\pi > y]dy = \frac{1}{4}c_{exp}(\pi) \end{aligned} \quad (3)$$

The first equality in (3) is by a change of variables $y = 4t$. ◀

2.1 Proof of Key Lemma

We now prove Lemma 2.1. Fix any $k \geq 1$ and define threshold $\theta := \theta_k = \ell_{k+1} + \delta$.

Let $T^* \subseteq N$ denote the optimal solution to the non-adaptive "fixed threshold" problem:

$$\max_{T \subseteq N, |T| \leq k} \quad \Pr\left[\min_{i \in T} X_i \leq \theta\right]. \quad (4)$$

We then proceed in two steps, as follows.

$$\Pr[\sigma \text{ finishes in } k \text{ queries}] \leq \Pr\left[\min_{i \in T^*} X_i \leq \theta\right] \leq \Pr[\pi \text{ finishes in } 2k \text{ iterations}]$$

The first inequality is shown in Lemma 2.3: this uses the fact that the fixed-threshold problem has adaptivity gap one (Proposition 1.2). The second inequality is shown in Lemma 2.4: this relies on the greedy criteria used in our algorithm.

▶ **Lemma 2.3.** $\Pr[\sigma$ *finishes in* k *queries*$] \leq \Pr\left[\min_{i \in T^*} X_i \leq \theta\right]$.

Proof. Let σ^k denote the optimal policy truncated after k queries: so the cost of σ^k is always at most k. Let $L(\sigma^k) = \min_{i \in N \setminus \sigma^k}\{\ell_i\}$ denote the smallest un-queried left-endpoint at the end of σ^k; this is a random value because σ^k is an adaptive policy. Then,

$$\Pr[\sigma \text{ finishes in } k \text{ queries}] = \Pr\left[\min_{i \in \sigma^k} X_i \;\leq\; L(\sigma^k) + \delta\right] \tag{5}$$

$$\leq \Pr\left[\min_{i \in \sigma^k} X_i \;\leq\; \ell_{k+1} + \delta\right] = \Pr\left[\min_{i \in \sigma^k} X_i \;\leq\; \theta\right] \tag{6}$$

$$\leq \Pr\left[\min_{i \in T^*} X_i \leq \theta\right] \tag{7}$$

Equality (5) is by the stopping criterion for SMQ. The inequality in (6) uses the observation that after *any* k queries, the smallest un-queried left-endpoint must be at most ℓ_{k+1}: so $L(\sigma^k) \leq \ell_{k+1}$ always. The equality in (6) is by definition of the threshold θ. Inequality (7) follows from Proposition 1.2: we view σ^k is a feasible adaptive policy for the fixed-threshold problem and T^* is the optimal non-adaptive policy. ◀

▶ **Lemma 2.4.** $\Pr\left[\min_{i \in T^*} X_i \leq \theta\right] \leq \Pr[\pi$ *finishes in* $2k$ *iterations*$]$.

Proof. Recall that each iteration j of Algorithm 1 selects two intervals: j in Step 3 and $b(j)$ in Step 4. Let $B = \{b(1), \cdots b(2k)\}$ be the set of intervals chosen by our policy π in Step 4 of the first $2k$ iterations. We partition B into $B' = \{b(1), \cdots b(k)\}$ and $B'' = \{b(k+1), \cdots b(2k)\}$. Let $d^* = \operatorname{argmin}_{d \in T^* \setminus B} \Pr(X_d > \theta_k) = \operatorname{argmax}_{d \in T^* \setminus B} \Pr(X_d \leq \theta_k)$.

$$\begin{aligned}
\Pr\left[\min_{i \in T^*} X_i \;>\; \theta_k\right] &= \prod_{i \in T^*} \Pr[X_i > \theta_k] \\
&= \prod_{i \in T^* \cap B} \Pr[X_i > \theta_k] \cdot \prod_{i \in T^* \setminus B} \Pr[X_i > \theta_k] \\
&\geq \prod_{i \in T^* \cap B} \Pr[X_i > \theta_k] \;\cdot\; (\Pr[X_{d^*} > \theta_k])^{|T^* \setminus B|} && (8) \\
&\geq \prod_{i \in T^* \cap B} \Pr[X_i > \theta_{2k}] \cdot \prod_{i \in B'' \setminus T^*} \Pr[X_i > \theta_{2k}] && (9) \\
&\geq \prod_{i \in B} \Pr[X_i > \theta_{2k}] = \Pr\left[\min_{i \in B} X_i \;>\; \theta_{2k}\right] && (10)
\end{aligned}$$

(8) follows from the definition of d^*. (10) just uses that $T^* \cap B$ and $B'' \setminus T^*$ are disjoint subsets of B. The key step above is (9), which we prove using two cases:

- Suppose that $T^* \setminus B = \emptyset$. Then, using $\theta_k \leq \theta_{2k}$ we obtain $\Pr[X_i > \theta_k] \geq \Pr[X_i > \theta_{2k}]$, which proves (9) for this case.

- Suppose that $T^* \setminus B \neq \emptyset$. In this case, d^* is well-defined. We now claim that:

$$\text{For each } j = k+1, \cdots 2k, \text{ either } b(j) \in T^* \text{ or } \Pr[X_{b(j)} > \theta_{2k}] \leq \Pr[X_{d^*} > \theta_k]. \quad (11)$$

Indeed, consider any such j and suppose that $b(j) \notin T^*$. As d^* is a valid choice for $b(j)$, the greedy rule implies:

$$\Pr\left[X_{d^*} > \theta_j\right] \geq \Pr\left[X_{b(j)} > \theta_j\right].$$

Further, using the fact that $\theta_k \leq \theta_j \leq \theta_{2k}$, we get

$$\Pr\left[X_{d^*} > \theta_k\right] \geq \Pr\left[X_{d^*} > \theta_j\right] \geq \Pr\left[X_{b(j)} > \theta_j\right] \geq \Pr\left[X_{b(j)} > \theta_{2k}\right],$$

which proves (11). Let h denote the number of iterations $j \in \{k+1, \cdots 2k\}$ where $b(j) \notin T^*$. Note that $h = |B''| - |T^* \cap B''| = k - |T^* \cap B''| \geq |T^* \setminus B|$, where we used $|T^*| = k$. Using (11), it follows that $\Pr[X_i > \theta_{2k}] \leq \Pr[X_{d^*} > \theta_k]$ for all $i \in B'' \setminus T^*$. Hence,

$$\Pr[X_{d^*} > \theta_k]^{|T^* \setminus B|} \geq \Pr[X_{d^*} > \theta_k]^h \geq \prod_{i \in B'' \setminus T^*} \Pr[X_i > \theta_{2k}],$$

Combined with the fact that $\theta_k \leq \theta_{2k}$ (as before), we obtain (9).

We are now ready to complete the proof. Using the SMQ stopping criterion and the fact that π queries all the intervals in B within $2k$ iterations,

$$\Pr[\pi \text{ finishes in } 2k \text{ iterations}] \geq \Pr\left[\min_{i \in B} X_i \leq \theta_{2k}\right] = 1 - \Pr\left[\min_{i \in B} X_i > \theta_{2k}\right].$$

Combined with (10),

$$\Pr[\pi \text{ finishes in } 2k \text{ iterations}] \geq 1 - \Pr\left[\min_{i \in T^*} X_i > \theta_k\right] = \Pr\left[\min_{i \in T^*} X_i \leq \theta\right],$$

where we use the definition $\theta = \theta_k$. ◀

2.2 Finding the minimum interval

In this section, we consider the SMQI problem, where the goal is to identify *an interval* that is guaranteed to be a δ-minimizer. Unlike the previous SMQ setting (where we find a δ-minimum value), for SMQI we just want to identify some interval $i^* \in N$ such that $X_{i^*} \leq \mathsf{MIN} + \delta$. Recall that $\mathsf{MIN} = \min_{i \in N} X_i$. It is important to note that the interval i^* may not have been queried. It is easy to see that any SMQ policy is also feasible to SMQI. Indeed, by the stopping rule (1) for SMQ, the δ-minimum value returned is always the minimum value of a queried interval: so we also identify i^*. However, an SMQI policy may return an interval i^* without querying it. So the optimal value of SMQI may be smaller than SMQ.

Remark. We note that the optimal values of SMQ and SMQI differ by at most the maximum query cost c_{max}. As noted above, the optimal SMQI value is at most that of SMQ. On the other hand, the optimal SMQ value is at most the optimal SMQI value plus the cost to query i^*. In the unit-cost setting, $c_{max} = 1$ and any policy has expected cost at least 1: so the optimal values of SMQ and SMQI are within a factor two of each other. This immediately

implies that Algorithm 1 is also an 8-approximation for unit-cost SMQI. In the rest of this subsection, we will prove a stronger result, that Algorithm 1 is a 4-approximation for SMQI. Apart from the improved constant factor, these ideas will also be helpful for SMQI with general costs. We note that under general costs, the optimal SMQ and SMQI values may differ by an arbitrarily large factor because c_{max} is not a lower bound on the optimal value.

Stopping criteria for SMQI. Consider any state, given by a subset $S \subseteq N$ of queried r.v.s along with their observations $\{x_i\}_{i \in S}$. There are two conditions under which the SMQI policy can stop.

- The first stopping rule is just the one for SMQ (1), which corresponds to the situation that interval i^* is queried. We restate this rule below for easy reference:

$$\min_{i \in S} x_i \quad \leq \quad \min_{j \in N \setminus S} \ell_j + \delta. \tag{12}$$

  In this case, we return $i^* = \arg\min_{i \in S} x_i$. We refer to this as the old stopping rule.
- The second stopping rule handles the situation where an un-queried interval i^* is returned. For any $i \in N$, define the "almost prefix" set $P_i := \{j \in N \setminus i : \ell_j < r_i - \delta\}$. Note that either P_i or $P_i \cup i$ is a *prefix* of $[n]$. (As before, we assume that intervals are indexed by increasing order of their left-endpoint, i.e., $\ell_1 \leq \ell_2 \leq \cdots \ell_n$.) The new rule is:

$$\exists\, i \in N \text{ such that } P_i \subseteq S \text{ and } \min_{j \in P_i} x_j \geq r_i - \delta. \tag{13}$$

  In other words, there is some interval i where (1) all intervals $j \neq i$ with left-endpoint $\ell_j < r_i - \delta$ have been queried, and (2) the minimum value of these r.v.s is at least $r_i - \delta$. In this case, we return $i^* = i$ (we may not know a δ-minimum value). We refer to this as the new stopping rule. See Figure 5 for an example.

▶ **Proposition 2.5.** *A policy for* SMQI *can stop if and only if either criterion* (12) *or* (13) *holds.*

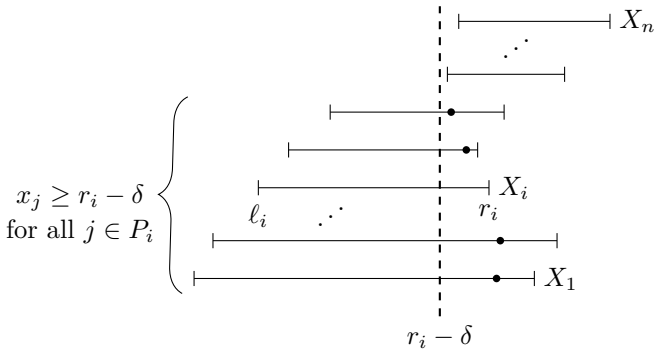

Figure 5 Illustration of new SMQI stopping criterion.

Our algorithm for SMQI with unit costs remains the same as for SMQ (Algorithm 1). The only difference is in the new stopping criterion (described above). Recall that π is the permutation used by our non-adaptive policy. When it is clear from the context, we will also use π to denote our SMQI policy that performs queries in the order of π until stopping criteria (12) or (13) applies.

▶ **Theorem 2.6.** *The non-adaptive policy π is a 4-approximation algorithm for* SMQI.

We now prove this result. Let σ denote an optimal adaptive policy for SMQI. For any $k \geq 1$, we will show:

$$\Pr[\sigma \text{ finishes in } k \text{ queries}] \leq \Pr[\pi \text{ finishes in } 2k \text{ iterations}]. \tag{14}$$

This would suffice to prove the 4-approximation, exactly as in Theorem 2.2.

In order to prove (14), we fix some $k \geq 1$. As in the previous proof, let $\theta = \theta_k = \ell_{k+1} + \delta$ and let T^* be defined as in (4). To reduce notation, define the following events.

$\mathcal{A}_1$: our policy π finishes within $2k$ iterations due to (12).

$\mathcal{A}_2$: our policy π finishes within $2k$ iterations due to (13).

$\mathcal{O}_1$: optimal policy σ finishes within k queries due to (12).

$\mathcal{O}_2$: optimal policy σ finishes within k queries due to (13).

Handling the old stopping criterion. Let L denote the smallest un-queried left-endpoint at the end of iteration $2k$ in π. Note that L is a deterministic value as π is a non-adaptive policy. Moreover, $L \geq \ell_{2k+1}$ as π would have queried the first $2k$ r.v.s. Let $\mathcal{G}$ be the event that $X_i > L + \delta$ for all intervals i queried by π in its first $2k$ iterations. In other words, $\mathcal{G}$ is precisely the event that stopping criterion (12) *does not* apply at the end of iteration $2k$ in π, i.e., $\mathcal{G} = \neg\mathcal{A}_1$. By Lemma 2.4,

$$\Pr[\neg\mathcal{G}] = \Pr[\mathcal{A}_1] \geq \Pr\left[\min_{i \in T^*} X_i \leq \theta\right].$$

Similarly, let $\mathcal{G}^*$ be that event that $X_i > \theta$ for all intervals i in the first k queries of σ. From the proof of Lemma 2.3, we obtain $\mathcal{O}_1 \subseteq \neg\mathcal{G}^*$ and

$$\Pr[\neg\mathcal{G}^*] \leq \Pr\left[\min_{i \in T^*} X_i \leq \theta\right].$$

Combining the above two inequalities, we have

$$\Pr[\neg\mathcal{G}^*] \leq \Pr[\neg\mathcal{G}]. \tag{15}$$

Handling the new stopping criterion. Let $\mathcal{G}_A$ be the event that $X_j > L + \delta$ for *all* r.v.s $j \in N$. Similarly, let $\mathcal{G}_A^*$ be the event that $X_j > \theta$ for *all* $j \in N$. Clearly,

$$\Pr[\mathcal{A}_2 \,|\, \mathcal{G}] = \Pr[\mathcal{A}_2 \,|\, \mathcal{G}_A] \quad \text{and} \quad \Pr[\mathcal{O}_2 \,|\, \mathcal{G}^*] = \Pr[\mathcal{O}_2 \,|\, \mathcal{G}_A^*]. \tag{16}$$

We will now prove that

$$\Pr[\mathcal{A}_2 \,|\, \mathcal{G}_A] \geq \Pr[\mathcal{O}_2 \,|\, \mathcal{G}_A^*]. \tag{17}$$

If σ finishes due to (13) in k queries then $P_{i^*} \subseteq [k+1]$: otherwise $|P_{i^*}| > k$ which contradicts with the fact that all r.v.s in P_{i^*} must be queried. Let $R = \{i \in N : P_i \subseteq [k+1]\}$ be all such intervals. It now follows that the event $\mathcal{O}_2$ (which corresponds to policy σ) is a subset of the event

$$\mathcal{E} := \bigvee_{i \in R} \left(\wedge_{j \in P_i}(X_j \geq r_i - \delta)\right). \tag{18}$$

Note that $\mathcal{E}$ is independent of the policy: it only depends on the realizations of the r.v.s (and doesn't depend on whether/not an interval has been queried).

Moreover, our policy π queries all the r.v.s in $[2k] \supseteq [k+1]$ within $2k$ iterations. So, for all $i \in R$, the r.v.s in $P_i \subseteq [k+1]$ are queried by π in $2k$ iterations. Hence, event $\mathcal{A}_2$ (which corresponds to policy π) contains event $\mathcal{E}$.

Recall that the event $\mathcal{G}_A$ (resp. $\mathcal{G}_A^*$) in policy π (resp. σ) means that every r.v. is more than $L+\delta$ (resp. θ). Also, $\theta \leq L+\delta$, which means

$$\Pr[X_j \geq u | X_j > L+\delta] \geq \Pr[X_j \geq u | X_j > \theta], \quad \forall u \in \mathbb{R}, \forall j \in N.$$

In other words, for any $j \in N$, if Y_j (resp. Z_j) is the r.v. X_j conditioned on $\mathcal{G}_A$ (resp. $\mathcal{G}_A^*$) then Y_j *stochastically dominates* Z_j.[2] Note also that the r.v.s Y_js (resp. Z_js) are independent. Using the fact that event $\mathcal{E}$ corresponds to a *monotone* function, we obtain:

▶ **Lemma 2.7.** *Let $\{Y_j : j \in N\}$ and $\{Z_j : j \in N\}$ be independent r.v.s such that Y_j stochastically dominates Z_j for each $j \in N$. Then,* $\Pr[\mathcal{E}(Y_1, ..., Y_n)] \geq \Pr[\mathcal{E}(Z_1, ..., Z_n)]$ *where event $\mathcal{E}$ is a function of independent r.v.s as defined in* (18)*.*

Proof. It suffices to prove the following.

$$\Pr[\mathcal{E}(Y_1, ..., Y_k, Z_{k+1}, ..., Z_n)] \geq \Pr[\mathcal{E}(Y_1, ..., Y_{k-1}, Z_k, ..., Z_n)], \quad \forall k \in [n].$$

Note that the r.v.s above only differ at position k. To keep notation simple, for any $j \in [n] \setminus k$ let $X'_j = Y_j$ if $j < k$ and $X'_j = Z_j$ if $j > k$. So, we need to show $\Pr[\mathcal{E}(X', Y_k)] \geq \Pr[\mathcal{E}(X', Z_k)]$. We *condition* on the realizations of the X' r.v.s. For each $j \in [n] \setminus k$ let t_j denote the realization of the r.v. X'_j. Having conditioned on these r.v.s, the only randomness is in Y_k and Z_k. We will show:

$$\Pr\left[\mathcal{E}(X', Y_k) | X' = t\right] \geq \Pr\left[\mathcal{E}(X', Z_k) | X' = t\right]. \tag{19}$$

Using the definition of the event $\mathcal{E}$ from (18), let $R(t) = \{i \in R : k \in P_i \text{ and } t_j > r_i - \delta \text{ for all } j \in P_i \setminus k\}$. In other words, $R(t) \subseteq R$ corresponds to those "clauses" in (18) that have not evaluated to true or false based on the realizations $\{X'_j = t_j : j \in [n] \setminus k\}$. If there is some clause in (18) that already evaluates to true (based on t) then $\mathcal{E}$ holds regardless of Y_k or Z_k. So, (19) holds in this case (both terms are one). Now, we assume that no clause in (18) already evaluates to true. We can write

$$\{\mathcal{E}(X', Y_k) | X' = t\} \quad = \quad \bigvee_{i \in R(t)} (Y_k \geq r_i - \delta) \quad = \quad \{Y_k \geq f\},$$

where $f = \min_{i \in R(t)} r_i - \delta$ is a deterministic value.[3] Similarly, we have

$$\{\mathcal{E}(X', Z_k) | X' = t\} \quad = \quad \{Z_k \geq f\}.$$

Using the fact that Y_k (resp. Z_k) is independent of X' and that Y_k stochastically dominates Z_k,

$$\Pr\left[\mathcal{E}(X', Y_k) | X' = t\right] = \Pr[Y_k \geq f] \geq \Pr[Z_k \geq f] = \Pr\left[\mathcal{E}(X', Z_k) | X' = t\right].$$

This completes the proof of (19). De-conditioning the X' r.v.s, we obtain $\Pr[\mathcal{E}(X', Y_k)] \geq \Pr[\mathcal{E}(X', Z_k)]$ as desired. ◀

Using Lemma 2.7, we obtain $\Pr[\mathcal{E}|\mathcal{G}_A] \geq \Pr[\mathcal{E}|\mathcal{G}_A^*]$, which proves (17). Combined with (16),

$$\Pr\left[\mathcal{A}_2 \,|\, \mathcal{G}\right] \geq \Pr\left[\mathcal{O}_2 \,|\, \mathcal{G}^*\right]. \tag{20}$$

[2] We say that r.v. Y stochastically dominates Z if $\Pr[Y \geq u] \geq \Pr[Z \geq u]$ for all $u \in \mathbb{R}$.
[3] If $R(t) = \emptyset$ then we set $f = \infty$.

Wrapping up. We have

$$\begin{aligned}
\Pr[\mathcal{A}_1 \vee \mathcal{A}_2] &= \Pr[\mathcal{A}_1] + \Pr[\mathcal{A}_2 \wedge \neg\mathcal{A}_1] = \Pr[\neg\mathcal{G}] + \Pr[\mathcal{A}_2 \wedge \mathcal{G}] \\
&= \Pr[\neg\mathcal{G}] + \Pr[\mathcal{A}_2|\mathcal{G}] \cdot \Pr[\mathcal{G}] = 1 - (1 - \Pr[\mathcal{A}_2|\mathcal{G}]) \cdot \Pr[\mathcal{G}] \\
&\geq 1 - (1 - \Pr[\mathcal{O}_2|\mathcal{G}^*]) \cdot \Pr[\mathcal{G}^*] \qquad \text{by (15) and (20)} \\
&= \Pr[\neg\mathcal{G}^*] + \Pr[\mathcal{O}_2 \wedge \mathcal{G}^*] \\
&\geq \Pr[\mathcal{O}_1] + \Pr[\mathcal{O}_2 \wedge \neg\mathcal{O}_1] \qquad \text{using } \mathcal{O}_1 \subseteq \neg\mathcal{G}^* \\
&= \Pr[\mathcal{O}_1 \vee \mathcal{O}_2].
\end{aligned}$$

This completes the proof of (14) and the theorem.

3 Algorithm for General Costs

We now consider the SMQ problem with non-uniform query costs. The high-level idea is similar to the unit-cost case: interleaving the two greedy criteria of smallest left-endpoint and highest probability of stopping. However, we need to incorporate the costs carefully. To this end, we use an iterative algorithm that in every iteration g, makes a *batch* of queries having total cost about 2^g. (In order to optimize the approximation ratio, we use a generic base y for the exponential costs.)

We extend the algorithm and analysis in this section to get a 7.47 approximation for the SMQI with non-uniform costs in [2].

For any subset $S \subseteq N$, let $c(S) := \sum_{j \in S} c_j$ denote the cost of querying all intervals in S. Again, we renumber intervals so that $\ell_1 \leq \ell_2 \leq \cdots \ell_n$.

▶ **Definition 3.1.** *For any $g \geq 0$, let T_g be the maximal prefix of intervals having cost at most y^g.*

Algorithm 2 Double Greedy for General Cost.

1: Let $\ell^* = \min_{j \in N} \ell_j$, $m^* = R := \min_{j \in N} r_j$, and $\pi \leftarrow \emptyset$.
2: **for** $g = 0, 1, 2, \cdots$ **do** ▷ iteration
3: Query intervals $T_g \setminus \pi$ and update list $\pi \leftarrow \pi \circ T_g$.
4: Update $\ell^* = \min_{j \in N \setminus \pi}\{\ell_j\}$ and let threshold $\theta_g = \ell^* + \delta$.
5: Compute a $(1, 1+\epsilon)$ bicriteria approximate solution U_g for:

$$p_g^* = \min_{T \subseteq N \setminus \pi} \left\{ \Pr\left[\min_{j \in T} X_j > \theta_g\right] : c(T) \leq y^g \right\}. \tag{KP}$$

6: Query intervals U_g and update list $\pi \leftarrow \pi \circ U_g$.
7: Update $\ell^* = \min_{j \in N \setminus \pi}\{\ell_j\}$ and $m^* = \min\left\{m^*, \min_{j \in T_g \cup U_g} X_j\right\}$.
8: **if** $m^* - \ell^* \leq \delta$ **then** stop.

The complete algorithm is given in Algorithm 2. The optimization problem (KP) solved in Step 5 is a variant of the classic knapsack problem: in Theorem 3.2 we provide a $(1, 1+\epsilon)$ bicriteria approximation algorithm for (KP) for any constant $\epsilon > 0$. In particular, this ensures that $c(U_g) \leq y^g(1+\epsilon)$ and

$$\Pr\left[\min_{j \in U_g} X_j > \theta_g\right] \leq p_g^*.$$

Note that the left-hand-side above equals $\prod_{j \in U_g} \Pr[X_j > \theta_g]$ as all r.v.s are independent.

▶ **Theorem 3.2.** *Given discrete random variables* $\{X_i\}_{i=1}^n$ *with costs* $\{c_i\}_{i=1}^n$*, budget* d *and threshold* $\theta \in \mathbb{R}$*, there is an* $n^{O(1/\epsilon)}$ *time algorithm that finds* $T \subseteq N$ *such that* $\Pr[\min_{j\in T} X_j > \theta] \le p^*$ *and* $c(T) \le (1+\epsilon)d$*, for any* $\epsilon > 0$*. Here,*

$$p^* = \min_{T\subseteq N} \left\{ \Pr\left[\min_{j\in T} X_j > \theta\right] : c(T) \le d \right\}. \tag{*}$$

The proof of Theorem 3.2 is presented in [2].

Furthermore, just like Algorithm 1, we can view Algorithm 2 as first computing the permutation π (without querying) and then performing queries in that order until the stopping criterion. So, our algorithm is a non-adaptive policy and our analysis also upper-bounds the adaptivity gap.

3.1 Analysis

We use σ to denote the optimal (adaptive) policy and π to denote our non-adaptive policy.

▶ **Definition 3.3.** *For any* $g \ge 0$*, let*

$$o_g := \Pr[\sigma \text{ does not finish by cost } y^g].$$

Similarly, for our policy we define

$$v_g := \Pr[\pi \text{ does not finish by iteration } g].$$

We also define σ_g *to be the optimal policy truncated at cost* y^g*, i.e., the total cost of queried intervals is always at most* y^g*. Similarly, we define* π_g *to be our policy truncated at the end of iteration* g*.*

The key part of the analysis lies in relating the non-stopping probabilities o_g and a_g in the optimal and algorithmic policies: see Lemma 3.5. Our first lemma bounds the (worst-case) cost incurred in g iterations of our policy.

▶ **Lemma 3.4.** *The cost of our policy until the end of iteration* g *is*

$$c(\pi_g) \le (1+\epsilon)\left(1 + \frac{y}{y-1}\right) y^g.$$

Proof. We handle separately the costs of intervals queried in Steps 3 and 6. The total cost incurred in Step 3 of the first g iterations is $c(T_g) \le y^g$: this uses $\cup_{k=0}^g T_k = T_g$ because T_g are prefixes. The total cost due to Step 6 can be bounded using a geometric series:

$$\sum_{k=0}^{g} c(U_k) \le (1+\epsilon)\sum_{k=0}^{g} y^k = (1+\epsilon)\cdot\frac{y^{g+1}-1}{y-1}.$$

The inequality above is by the cost guarantee for (KP). The lemma now follows. ◀

▶ **Lemma 3.5.** *For all* $g \ge 0$*, we have* $a_g \le o_g$*.*

Proof. Recall that σ_g denotes the optimal policy truncated at cost y^g. We let $L(\sigma_g) = \min_{j\in N\setminus\sigma_g}\{\ell_j\}$ be the smallest un-queried left-endpoint: this is a random value as σ_g is adaptive. In the algorithm, consider iteration g and let $L(T_g) = \min_{j\in N\setminus T_g}\{\ell_j\}$; note that the threshold $\theta_g \ge L(T_g) + \delta$ in Step 4. Let $\pi' = \pi_{g-1} \circ T_g$ denote the list after Step 3 in

iteration g. Note that the optimization in (KP) of iteration g is over $T \subseteq N \setminus \pi'$, which yields U_g. Also, $\pi_g = \pi' \cup U_g$.

$$
\begin{aligned}
o_g &= \Pr\left[\text{OPT does not finish within cost } y^g\right] \\
&= \Pr\left[\min_{j\in\sigma_g} X_j > L(\sigma_g) + \delta\right] \geq \Pr\left[\min_{j\in\sigma_g} X_j > L(T_g) + \delta\right] && (21)\\
&\geq \Pr\left[\min_{j\in\sigma_g} X_j > \theta_g\right] = 1 - \Pr\left[\min_{j\in\sigma_g} X_j \leq \theta_g\right] && (22)\\
&\geq 1 - \max_{T\subseteq N, c(T)\leq y^g} \Pr\left[\min_{j\in T} X_j \leq \theta_g\right] = \min_{T\subseteq N, c(T)\leq y^g} \Pr\left[\min_{j\in T} X_j > \theta_g\right] && (23)\\
&= \min_{T\subseteq N, c(T)\leq y^g} \prod_{j\in T} \Pr\left[X_j > \theta_g\right] \geq \prod_{j\in\pi'} \Pr[X_j > \theta_g] \cdot \min_{T\subseteq N\setminus\pi', c(T)\leq y^g} \prod_{j\in T} \Pr[X_j > \theta_g] && (24)\\
&= \prod_{j\in\pi'} \Pr[X_j > \theta_g] \cdot p^*_g = \Pr\left[\min_{j\in\pi'} X_j > \theta_g\right] \cdot p^*_g && (25)\\
&\geq \Pr\left[\min_{j\in\pi'} X_j > \theta_g\right] \cdot \Pr\left[\min_{j\in U_g} X_j > \theta_g\right] = \Pr\left[\min_{j\in\pi_g} X_j > \theta_g\right] \geq v_g && (26)
\end{aligned}
$$

The equality in (21) is given by the definition of $L(\sigma_g)$ and the stopping rule. The inequality in (21) uses the fact that $L(\sigma_g) \leq L(T_g)$ always, which in turn is because σ_g has cost at most y^g and T_g is the maximal prefix within this cost. The inequality in (22) uses $\theta_g \geq L(T_g) + \delta$. The inequality in (23) is by Proposition 1.2: we view σ_g as a feasible adaptive policy for the fixed-threshold problem with threshold θ_g and budget y^g. The equality in (24) follows from independence of the random variables. The first equality in (25) uses the definition of p^*_g from (KP) and independence. The first inequality in (26) uses the choice of U_g and Theorem 3.2. The equality in (26) is by $\pi_g = \pi' \cup U_g$. To see the last inequality in (26), note that if $\min_{j\in\pi_i}\{X_j\} \leq \theta_g$ then π finishes by iteration g. ◀

In Lemma 3.6 we lower bound the expected cost of the optimal policy. Let $c_{exp}(\pi)$ and $c_{exp}(\sigma)$ denote the expected cost of our greedy policy and the optimal policy, respectively.

▶ **Lemma 3.6.** *For any base $y \geq 1$, we have $\sum_{g\geq 0} y^g \cdot o_g \leq \frac{y}{y-1} c_{exp}(\sigma) - \frac{1}{y-1}$.*

Proof. Let Z denote the random variable that represents the cost of the optimal policy σ: so $c_{exp}(\sigma) = \mathbb{E}[Z]$. Let $\mathbf{1}(Z > y^g)$ be the indicator variable for when $Z > y^g$; so $\mathbb{E}[I(Z > y^g)] = o_g$. We now show that:

$$\sum_{g\geq 0} y^g \cdot \mathbf{1}(Z > y^g) \leq \frac{y}{y-1} Z - \frac{1}{y-1} \tag{27}$$

To see this, suppose that $y^k < Z \leq y^{k+1}$ for some integer $k \geq 0$. Then the left-hand-side of (27) equals

$$\sum_{g=0}^{k} y^g = \frac{y^{k+1}-1}{y-1} \leq Z\frac{y}{y-1} - \frac{1}{y-1},$$

which proves (27). Taking the expectation of (27) proves the lemma. ◀

▶ **Theorem 3.7.** *There is a $(3 + 2\sqrt{2} + \epsilon)$-approximation for the* SMQ *problem with general costs.*

Proof. By Lemma 3.4, we have $c_{exp}(\pi) \leq (1+\epsilon)\left(1+\frac{y}{y-1}\right)\sum_{g\geq 1} y^g(v_{g-1}-v_g)$. Now,

$$\sum_{g\geq 1} y^g(v_{g-1}-v_g) = v_0 + (y-1)\sum_{g\geq 0} y^g v_g \leq 1 + (y-1)\sum_{g\geq 0} y^g o_g \tag{28}$$

$$\leq 1 + (y-1)\sum_{g\geq 0} y^g o_g \leq 1 + y\cdot c_{exp}(\sigma) - 1 = y\cdot c_{exp}(\sigma) \tag{29}$$

The inequality in (28) is by Lemma 3.4 and $v_0 = 1$. The first inequality in (29) uses $y \geq 1$ and the second inequality is by Lemma 3.6.

Hence, we obtain $c_{exp}(\pi) \leq (1+\epsilon)y \cdot \left(1+\frac{y}{y-1}\right)\cdot c_{exp}(\sigma)$. Now, optimizing for y, we obtain the stated approximation ratio. ◀

References

1 Marek Adamczyk, Maxim Sviridenko, and Justin Ward. Submodular stochastic probing on matroids. *Math. Oper. Res.*, 41(3):1022–1038, 2016. doi:10.1287/MOOR.2015.0766.

2 Hessa Al-Thani and Viswanath Nagarajan. Identifying approximate minimizers under stochastic uncertainty, 2025. arXiv:2504.17019.

3 Arash Asadpour and Hamid Nazerzadeh. Maximizing stochastic monotone submodular functions. *Manag. Sci.*, 62(8):2374–2391, 2016. doi:10.1287/MNSC.2015.2254.

4 Evripidis Bampis, Christoph Dürr, Thomas Erlebach, Murilo Santos de Lima, Nicole Megow, and Jens Schlöter. Orienting (hyper)graphs under explorable stochastic uncertainty. In *29th Annual European Symposium on Algorithms (ESA)*, volume 204 of *LIPIcs*, pages 10:1–10:18, 2021. doi:10.4230/LIPICS.ESA.2021.10.

5 N. Bansal, A. Gupta, J. Li, J. Mestre, V. Nagarajan, and A. Rudra. When LP is the cure for your matching woes: Improved bounds for stochastic matchings. *Algorithmica*, 63(4):733–762, 2012. doi:10.1007/S00453-011-9511-8.

6 Steven Chaplick, Magnús M. Halldórsson, Murilo S. de Lima, and Tigran Tonoyan. Query minimization under stochastic uncertainty. *Theoretical Computer Science*, 895:75–95, 2021. doi:10.1016/J.TCS.2021.09.032.

7 B. C. Dean, M. X. Goemans, and J. Vondrák. Approximating the stochastic knapsack problem: The benefit of adaptivity. *Math. Oper. Res.*, 33(4):945–964, 2008. doi:10.1287/MOOR.1080.0330.

8 Thomas Erlebach, Michael Hoffmann, and Frank Kammer. Query-competitive algorithms for cheapest set problems under uncertainty. *Theoretical Computer Science*, 613:51–64, 2016. doi:10.1016/J.TCS.2015.11.025.

9 Thomas Erlebach, Michael Hoffmann, Danny Krizanc, Matús Mihal'Ák, and Rajeev Raman. Computing minimum spanning trees with uncertainty. In *STACS*, pages 277–288, 2008.

10 Tomás Feder, Rajeev Motwani, Rina Panigrahy, Chris Olston, and Jennifer Widom. Computing the median with uncertainty. In *Proceedings of the thirty-second annual ACM symposium on Theory of computing*, pages 602–607, 2000. doi:10.1145/335305.335386.

11 Hao Fu, Jian Li, and Pan Xu. A PTAS for a class of stochastic dynamic programs. In *45th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 107 of *LIPIcs*, pages 56:1–56:14, 2018. doi:10.4230/LIPICS.ICALP.2018.56.

12 Dimitrios Gkenosis, Nathaniel Grammel, Lisa Hellerstein, and Devorah Kletenik. The Stochastic Score Classification Problem. In *26th Annual European Symposium on Algorithms (ESA)*, pages 36:1–36:14, 2018. doi:10.4230/LIPICS.ESA.2018.36.

13 Ashish Goel, Sudipto Guha, and Kamesh Munagala. How to probe for an extreme value. *ACM Transactions on Algorithms (TALG)*, 7(1):1–20, 2010. doi:10.1145/1868237.1868250.

14 Sudipto Guha and Kamesh Munagala. Approximation algorithms for budgeted learning problems. In *39th Annual ACM Symposium on Theory of Computing (STOC)*, pages 104–113. ACM, 2007. doi:10.1145/1250790.1250807.

15 Anupam Gupta, Ravishankar Krishnaswamy, Viswanath Nagarajan, and R. Ravi. Running errands in time: Approximation algorithms for stochastic orienteering. *Math. Oper. Res.*, 40(1):56–79, 2015. doi:10.1287/MOOR.2014.0656.

16 Anupam Gupta, Viswanath Nagarajan, and Sahil Singla. Adaptivity gaps for stochastic probing: Submodular and XOS functions. In *Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1688–1702. SIAM, 2017. doi:10.1137/1.9781611974782.111.

17 Lisa Hellerstein, Devorah Kletenik, and Srinivasan Parthasarathy. A tight bound for stochastic submodular cover. *J. Artif. Intell. Res.*, 71:347–370, 2021. doi:10.1613/JAIR.1.12368.

18 Lisa Hellerstein, Naifeng Liu, and Kevin Schewior. Quickly determining who won an election. In *15th Innovations in Theoretical Computer Science Conference (ITCS)*, LIPIcs, pages 61:1–61:14, 2024. doi:10.4230/LIPICS.ITCS.2024.61.

19 Sungjin Im, Viswanath Nagarajan, and Ruben van der Zwaan. Minimum latency submodular cover. *ACM Trans. Algorithms*, 13(1):13:1–13:28, 2016. doi:10.1145/2987751.

20 Haotian Jiang, Jian Li, Daogao Liu, and Sahil Singla. Algorithms and adaptivity gaps for stochastic k-tsp. In *11th Innovations in Theoretical Computer Science Conference (ITCS)*, volume 151 of *LIPIcs*, pages 45:1–45:25, 2020. doi:10.4230/LIPICS.ITCS.2020.45.

21 Simon Kahan. A model for data in motion. In *Proceedings of the Twenty-third Annual ACM Symposium on Theory of computing*, pages 265–277, 1991.

22 Nicole Megow, Julie Meißner, and Martin Skutella. Randomization helps computing a minimum spanning tree under uncertainty. *SIAM Journal on Computing*, 46(4):1217–1240, 2017. doi:10.1137/16M1088375.

23 Nicole Megow and Jens Schlöter. Set selection under explorable stochastic uncertainty via covering techniques. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 319–333. Springer, 2023. doi:10.1007/978-3-031-32726-1_23.

24 Danny Segev and Sahil Singla. Efficient approximation schemes for stochastic probing and prophet problems. In *22nd ACM Conference on Economics and Computation (EC)*, pages 793–794. ACM, 2021. doi:10.1145/3465456.3467614.

25 Sahil Singla. The price of information in combinatorial optimization. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2523–2532. SIAM, 2018. doi:10.1137/1.9781611975031.161.

26 Weina Wang, Anupam Gupta, and Jalani K Williams. Probing to minimize. In *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*, volume 215, page 120, 2022.

A Multiplicative Precision

Given an instance with non-negative r.v.s $\{X_i\}_{i=1}^n$ and multiplicative precision $\alpha \geq 1$, consider a new instance of SMQ with r.v.s $\{X_i' := \ln(X_i)\}_{i=1}^n$ and additive precision $\delta := \ln \alpha$. Note that

$$\mathsf{MIN}' = \min_{i=1}^n X_i' = \min_{i=1}^n \ln(X_i) = \ln\left(\min_{i=1}^n X_i\right) = \ln(\mathsf{MIN}).$$

An α-approximately minimum value W for the original instance satisfies $\mathsf{MIN} \leq W \leq \alpha \cdot \mathsf{MIN}$, where $\mathsf{MIN} = \min_{i=1}^n X_i$. Then, $\mathsf{VAL} = \ln(W)$ satisfies $\mathsf{MIN}' = \ln(\mathsf{MIN}) \leq \mathsf{VAL} \leq \ln(\mathsf{MIN}) + \ln \alpha = \mathsf{MIN}' + \delta$, i.e., VAL is a δ-minimum value for the new instance. Similarly, if VAL is a δ-minimum value for the new instance then $W := e^{\mathsf{VAL}}$ is an α-approximately minimum value for the original instance.

B Bad Example for Competitive Ratio

We provide an example that rules out any reasonable *competitive ratio* bound for SMQ and SMQI with precision $\delta > 0$. This is in sharp contrast to the corresponding problem with exact precision ($\delta = 0$) for which a constant competitive ratio is known [21]. We note that

results in the online setting assume open intervals, which in our setting (with discrete r.v.s) corresponds to all left-endpoints being distinct.[4] The benchmark in the online setting is the *hindsight optimum*, which is the minimum number (or cost) of queries that are needed to verify a δ-minimum value *conditioned* on the realizations $\{x_i\}_{i=1}^n$ of the r.v.s.

Consider an instance with n r.v.s with $\Pr[X_i = i] = \frac{\ln n}{n}$ and $\Pr[X_i = n^2] = 1 - \frac{\ln n}{n}$ for all $i \in [n]$. All costs are unit and the precision $\delta = n$. We refer to the values $\{1, 2, \cdots n\}$ as *low* values: note that any low value is a δ-minimum value for this instance.

We first consider the hindsight optimum. If any of the n r.v.s (say k) realizes to a low value then verifying the δ-minimum value just requires querying k, which has cost 1. On the other hand, the probability that none of the n r.v.s realizes to a low value is $\left(1 - \frac{\ln n}{n}\right)^n \leq \frac{1}{n}$: in this case the optimal verification cost is n (querying all r.v.s). So the expected optimal cost is at most 2.

Now, consider an SMQ policy: this does not know the realizations. It is easy to see that the only way to stop querying is when some low value is observed (or all n r.v.s are queried). So, the expected cost of any policy is at least $\frac{n}{\ln n}$. Hence the competitive ratio for SMQ is $\Omega(\frac{n}{\ln n})$.

[4] Alternatively, our example can be modified into one with open intervals where the competivity ratio is still $\tilde{\Omega}(n)$.