Set Selection under Explorable Stochastic Uncertainty via Covering Techniques

Nicole Megow*

Jens Schlöter*

June 16, 2023

Abstract

Given subsets of uncertain values, we study the problem of identifying the subset of minimum total value (sum of the uncertain values) by querying as few values as possible. This *set selection problem* falls into the field of *explorable uncertainty* and is of intrinsic importance therein as it implies strong adversarial lower bounds for a wide range of interesting combinatorial problems such as knapsack and matchings. We consider a stochastic problem variant and give algorithms that, in expectation, improve upon these adversarial lower bounds. The key to our results is to prove a strong structural connection to a seemingly unrelated covering problem with uncertainty in the constraints via a linear programming formulation. We exploit this connection to derive an algorithmic framework that can be used to solve both problems under uncertainty, obtaining nearly tight bounds on the competitive ratio. This is the first non-trivial stochastic result concerning the sum of unknown values without further structure known for the set. With our novel methods, we lay the foundations for solving more general problems in the area of explorable uncertainty.

1 Introduction

In the setting of *explorable uncertainty*, we consider optimization problems with uncertainty in the numeric input parameters. Instead of having access to the numeric values, we are given *uncertainty intervals* that contain the precise values. Each uncertainty interval can be queried, which reveals the corresponding precise value. The goal is to adaptively query intervals until we have sufficient information to optimally (or approximately) solve the underlying optimization problem, while minimizing the number of queries.

This paper mainly considers the set selection problem (MINSET) under explorable uncertainty. In this problem, we are given a set of n uncertain values represented by uncertainty intervals $\mathcal{I} = \{I_1, \ldots, I_n\}$ and a family of m sets $\mathcal{S} = \{S_1, \ldots, S_m\}$ with $S \subseteq \mathcal{I}$ for all $S \in \mathcal{S}$. A value w_i lies in its uncertainty interval I_i , is initially unknown, and can be revealed via a query. The value of a subset S is $w(S) = \sum_{I_i \in S} w_i$. Our goal is to determine a subset of minimum value as well as the corresponding value by using a minimal number of queries. It can be seen as an optimization problem with uncertainty in the coefficients of the objective function:

$$\min \sum_{\substack{j=1 \ x_j}}^m x_j \sum_{I_i \in S_j} w_i$$
s.t.
$$\sum_{\substack{j=1 \ x_j}}^m x_j = 1$$

$$x_j \in \{0,1\} \quad \forall j \in \{1,\ldots,m\}.$$
(SETSELIP)

Since the precise w_i 's are uncertain, we do not always have sufficient information to just compute an optimal solution to (SETSELIP) and instead might have to execute queries in order to determine such a solution. An algorithm for MINSET under uncertainty can adaptively query intervals until it has sufficient information to

^{*}University of Bremen, Faculty of Mathematics and Computer Science, {nicole.megow, jschloet}@uni-bremen.de

determine an optimal solution to (SETSELIP). Adaptivity in this context means that the algorithm can take previous query results into account to decide upon the next query.

In this paper, we consider the stochastic problem variant, where we assume that all values w_i are drawn independently at random from their intervals I_i according to unknown distributions d_i . Since there are instances that cannot be solved without querying the entire input, we analyze an algorithm ALG in terms of its *competitive ratio*: for the set of problem instances \mathcal{J} , it is defined as $\max_{J \in \mathcal{J}} \mathbb{E}[ALG(J)]/\mathbb{E}[OPT(J)]$, where ALG(J) is the number of queries needed by ALG to solve instance J, and OPT(J) is the minimum number of queries necessary to solve the instance.

MINSET is a fundamental problem and of intrinsic importance within the field of explorable uncertainty. The majority of existing works considers the *adversarial setting*, where query outcomes are not stochastic but returned in a worst-case manner. Selection type problems have been studied in the adversarial setting and constant (matching) upper and lower bounds are known, e.g., for selecting the minimum [26], the *k*-th smallest element [16,26], a minimum spanning tree [13,14,24,32], sorting [23] and geometric problems [6]. However, these problems essentially boil down to comparing *single* uncertainty intervals and identifying the minimum of two unknown values. Once we have to compare two (even disjoint) sets and the corresponding *sums* of unknown values, no deterministic algorithm can have a better adversarial competitive ratio than *n*, the number of uncertainty intervals. This has been shown by Erlebach et al. [12] for MINSET, and it implies adversarial lower bounds for classical combinatorial problems, such as, knapsack [33] and matchings [33], and solving ILPs with uncertainty in the cost coefficients [33] as in (SETSELIP) above. Thus, solving MINSET under stochastic uncertainty is an important step towards obtaining improved results for this range of problems. As a main result, we provide substantially better algorithms for MINSET under stochastic uncertainty adversarial lower bounds for a wide range of problems.

For the *stochastic setting*, the only related results we are aware of concern sorting [7] and the problem of finding the minimum in each set of a given collection of sets [3]. Asking for the *sum* of unknown values is substantially different.

1.1 The Covering Point of View

Our key observation is that we can view MINSET as a covering problem with uncertainty in the constraints. To see this, we focus on the structure of the uncertainty intervals and how a query affects it. We assume that each interval $I_i \in \mathcal{I}$ is either open (*non-trivial*) or *trivial*, i.e., $I_i = (L_i, U_i)$ or $I_i = \{w_i\}$; a standard technical assumption in explorable uncertainty. In the latter case, $L_i = U_i = w_i$. We call L_i and U_i lower and upper limit. If S contains only trivial uncertainty intervals, then we define $I_S = [L_S, U_S] = \{w(S)\}$ and call I_S trivial. Otherwise, we define $I_S = (L_S, U_S)$. Clearly, the value w(S) of a set $S \in S$ is contained in the interval I_S , i.e., $w(S) \in I_S$. We call I_S the uncertainty interval of set S. See Figure 1 for an example.

Since the intervals (L_S, U_S) of the sets $S \in S$ can overlap, we might have to execute queries to determine the set of minimum value. A query to an interval I_i reveals the precise value w_i and, thus, replaces both, L_i and U_i , with w_i . In a sense, a query to an $I_i \in S$ reduces the range (L_S, U_S) in which w(S) might lie by increasing L_S by $w_i - L_i$ and decreasing U_S by $U_i - w_i$. See Figure 2 for an example. We use L_S and U_S to refer to the initial limits and $L_S(Q)$ and $U_S(Q)$ to denote the limits of a set $S \in S$ after querying a set of intervals $Q \subseteq I$.

Let $w^* = \min_{S \in S} w(S)$ be the initially uncertain minimum value. To solve the problem, we have to adaptively query a set of intervals Q until $U_{S^*}(Q) = L_{S^*}(Q) = w^*$ holds for some $S^* \in S$ and $L_S(Q) \ge w^*$ holds for all $S \in S$. Only then, we know for sure that w^* is indeed the minimum set value and that S^* achieves this value. Figure 3 shows the structure of an instance that has been solved. For an instance $(\mathcal{I}, \mathcal{S})$ of MINSET, the following integer linear program (ILP) with $a_i = w_i - L_i$ for all $I_i \in \mathcal{I}$ and



Figure 1: Instance for set selection under explorable uncertainty with intervals $\mathcal{I} = \{I_1, I_2, \ldots, I_8\}$ and sets $\mathcal{S} = \{S_1, S_2, S_3, S_4\}$ with $S_1 = \{I_1, I_2\}$, $S_2 = \{I_3, I_4, I_5\}$, $S_3 = \{I_4, I_5, I_6\}$ and $S_4 = \{I_7, I_8\}$ (left) and the corresponding uncertainty intervals I_{S_j} for the sets S_j with $j \in \{1, \ldots, 4\}$ (right). Green circles illustrate the precise values.

 $b_S = w^* - L_S$ for all $S \in S$ formulates this problem:

$$\begin{array}{ll} \min & \sum_{I_i \in \mathcal{I}} x_i \\ \text{s.t.} & \sum_{I_i \in S} x_i \cdot a_i \ge b_S \quad \forall S \in \mathcal{S} \\ & x_i \in \{0, 1\} \qquad \quad \forall I_i \in \mathcal{I} \end{array}$$
(MINSETIP)

Here, the variable x_i , $I_i \in \mathcal{I}$, indicates whether interval I_i is selected to be queried ($x_i = 1$) or not ($x_i = 0$) and our objective is to minimize the number of queries.

Observe that this ILP is a special case of the *multiset multicover problem* (see, e.g., [35]). If $a_i = w_i - L_i = 1$ for all $I_i \in \mathcal{I}$ and $b_S = w^* - L_S = 1$ for all $S \in S$, then the problem is exactly the classical SETCOVER problem with \mathcal{I} corresponding to the SETCOVER sets and S corresponding to the SETCOVER elements.

The optimal solution to (MINSETIP) is the optimal query set for the corresponding MINSET instance; this is not hard to see but we also formally prove it in Appendix B. Under uncertainty however, the coefficients $a_i = w_i - L_i$ and right-hand sides $b_S = w^* - L_S$ are unknown to us. We only know that $a_i \in (L_i - L_i, U_i - L_i) = (0, U_i - L_i)$ as $a_i = (w_i - L_i)$ and $w_i \in (L_i, U_i)$. Only once we query an interval I_i , the precise value w_i and, thus, the coefficient a_i is revealed to us. In a sense, to solve MINSET under uncertainty, we have to solve (MINSETIP) with uncertainty in the coefficients and with irrevocable decisions. For the rest of the paper, we interpret MINSET under uncertainty in exactly that way: We have to solve (MINSETIP) without knowing the coefficients in the constraints. Whenever we *irrevocably* add an interval I_i to our solution (i.e., set x_i to 1), the information on the coefficients (in form of w_i) is revealed to us. Our goal is to add elements to our solution until it becomes feasible for (MINSETIP), and to minimize the number of added elements. In this interpretation, the terms "querying an element" and "adding it to the solution" are interchangeable, and we use them as such.

$$I_{i} \underbrace{\underbrace{I_{S_{1}}}_{w_{i}-L_{i}} U_{i}-w_{i}} I_{S_{2}} \underbrace{\underbrace{I_{S_{2}}}_{w_{i}-L_{i}} U_{i}-w_{i}} U_{i}-w_{i}$$

Figure 2: Example of how a query to an interval I_i changes the intervals of two sets S_1, S_2 with $I_i \in S_1 \cap S_2$ in the set selection problem under explorable uncertainty.



Figure 3: Instance of Figure 1 after querying $Q = \{I_1, I_2, I_5, I_7\}$: Updated uncertainty intervals \mathcal{I} (left) and updated set uncertainty intervals (right).

Our main contribution is an algorithmic framework that exploits techniques for classical covering problems and adapts them to handle uncertainty in the coefficients a_i and the right-hand sides b_S . This framework not only allows us to obtain improved results for MINSET under stochastic uncertainty but also to solve other covering problems with uncertainty in the constraints.

1.2 Our Results

We design a polynomial-time algorithm for MINSET under stochastic uncertainty with competitive ratio $O(\frac{1}{\tau} \cdot \log^2 m)$, where *m* is the number of sets (number of constraints in (MINSETIP)) and parameter τ characterizes how "balanced" the distributions of values within the given intervals are. More precisely, $\tau = \min_{I_i \in \mathcal{I}} \tau_i$ and τ_i is the probability that w_i is larger than the center of I_i (e.g., for uniform distributions $\tau = \frac{1}{2}$). All our results assume $\tau > 0$. This is the first stochastic result in explorable uncertainty concerning the sum of unknown values and it builds on new methods that shall be useful for solving more general problems in this field. The ratio is independent of the number of elements, *n*. In particular for a small number of sets, *m*, this is a significant improvement upon the adversarial lower bound of *n* [12]. Dependencies on parameters such as τ are quite standard and necessary [4, 5, 18, 30, 38]. For example, in [30] the upper bounds depend on the probability to draw the largest value of the uncertainty interval, which is an even stricter assumption that does not translate to open intervals.

We remark that the hidden constants in the performance bounds depend on the upper limits of the given intervals. Assuming those to be constant is also a common assumption; see, e.g., [31]. Even greedy algorithms for covering problems similar to (MINSETIP) *without* uncertainty have such dependencies [11, 35, 37]. While there exist non-greedy algorithms for covering problems without such dependencies [28, 29], it remains open whether they can be adjusted to the setting with uncertainty and, in particular, irrevocable decisions.

As MINSET contains the classical SETCOVER problem, an approximation factor better than $O(\log m)$ is unlikely, unless P=NP [10]. We show that this holds also in the stochastic setting, even with uniform distributions. We further show that $\frac{2}{\tau}$ is a lower bound for both problems under stochastic explorable uncertainty, even if the sets are pairwise disjoint. Hence, the dependencies on $\log m$ and $\frac{2}{\tau}$ in our upper bounds are necessary.

In the special case that all given sets are disjoint, we provide a simpler algorithm with competitive ratio $\frac{2}{\tau}$, which matches the lower bound. This is a gigantic improvement compared to the adversarial setting, where the lower bound of *n* holds even for disjoint sets [12].

We remark that all our results for MINSET translate to the maximization variant of the problem, where

we have to determine the set of maximum value (cf. Appendix D).

Algorithmically, we exploit the covering point of view to introduce a class of greedy algorithms that use the same basic strategy as the classical SETCOVER greedy algorithm [8]. However, we do not have sufficient information to compute and query an exact greedy choice under uncertainty as this choice depends on uncertain parameters. Instead, we show that it is sufficient to query a small number of elements that together achieve a similar greedy value to the exact greedy choice. If we do this repeatedly and the number of queries per iteration is small in expectation, then we achieve guarantees comparable to the approximation factor of a greedy algorithm with full information. It is worth noting that this way of comparing an algorithm to the optimal solution is a novelty in explorable uncertainty as all previous algorithms for adversarial explorable uncertainty (MINSET and other problems) exploit witness sets. A witness set is a set of queries Q such that each feasible solution has to query at least one element of Q, which allows to compare an algorithm with an optimal solution.

Our results translate to different covering problems under uncertainty. In particular, we consider the variant of (MINSETIP) under uncertainty with *deterministic right-hand sides*. We give a simplified algorithm with improved competitive ratio $\mathcal{O}(\frac{1}{\tau} \cdot \log m)$. For a different balancing parameter, this holds even for the more general variant, where a variable can have different coefficients for the different constraints each with an individual uncertainty interval and distribution. For this problem, adding an element to the solution reveals *all* corresponding coefficients.

1.3 Further Previous Work

Since MINSET under uncertainty can be interpreted as both, a query minimization problem and a covering problem with uncertainty, we in the following summarize previous work from both fields.

1.3.1 Previous Work on Query Problems

For adversarial MINSET under uncertainty, Erlebach et al. [15] show a (best possible) competitive ratio of 2d, where d is the cardinality of the largest set. In the lower bound instances, $d \in \Omega(n)$. The algorithm repeatedly queries disjoint witness sets of size at most 2d. This result was stated for the setting, in which it is not necessary to determine the value of the minimal set; if the value has to be determined, the bounds change to d.

Further related work on MINSET includes the result by Maehara and Yamaguchi [31], who consider packing ILPs with (stochastic) uncertainty in the cost coefficients, which can be queried. They present a framework for solving several problems and bound the absolute number of iterations that it requires to solve them, instead of the competitive ratio. However, we show in Appendix A that their algorithm has competitive ratio $\Omega(n)$ for MINSET under uncertainty, even for uniform distributions. Thus, it does not improve upon the adversarial lower bound.

Also Wang et al. [39] consider selection-type problems in a somewhat related model. In contrast to our setting, they consider different constraints on the set of queries that, in a way, imply a budget on the number of queries. They solve optimization problems with respect to this budget, which has a very different flavor than our setting of minimizing the number of queries.

Furthermore, there is related work in a setting, where a query reveals the *existence* of entities instead of numeric values, e.g., the existence of edges in a graph, c.f. [5, 18, 38]. For example, Behnezhad et al. [4] showed that vertex cover can be approximated within a factor of $(2 + \epsilon)$ with only a constant number of queried edges per vertex. As edges define constraints, the result considers uncertainty only in the right-hand sides.

1.3.2 Previous Work on Covering Problems with Uncertainty

We continue by summarizing previous work on covering problems in different adversarial and stochastic settings.

In the *online* version of SETCOVER [2], we are given a ground set of elements and a family of subsets of these elements. In contrast to offline SETCOVER, we do not necessarily have to cover all elements of the ground set. Instead, the members of the ground set that we do actually have to cover arrive online in an adversarial manner. Whenever an element arrives, we have to cover it by irrevocably adding a set containing the element to our solution, unless a previously added set already contains the element. In a sense, online SETCOVER is a variant of (MINSETIP) under uncertainty, where only the right-hand sides are uncertain in $\{0, 1\}$ and all left-hand side coefficients are known and either one or zero. In contrast to MINSET under uncertainty, the adversary for online SETCOVER is in a sense more powerful when selecting the right-hand sides as they do not depend on a common value w^* . Because of these differences, online SETCOVER has a very different flavor to MINSET under uncertainty. The same holds for the stochastic version of online SETCOVER [21, 22], where the subset of elements to be covered is drawn from a probability distribution.

A different stochastic variant of SETCOVER considers a *two-stage* version of the problem [36]. In the first stage, we do not yet know which members of the ground set actually need to be covered. After the first stage, the elements to be covered are drawn from a probability distribution and in the second stage we have full knowledge of the elements to be covered. The crux of this two-stage variant is that adding sets to the solution in the first stage can be cheaper than adding them to the solution during the second stage. This again leads to a very different flavor than our setting.

While these SETCOVER variants consider uncertainty in the set of elements that need to be covered, Goemans and Vondrák [19] consider a variant where the elements to be covered are certain but there is uncertainty in which elements are covered by the sets. For each set, a vector describing the elements that are covered by the set is drawn according to a probability distribution. This corresponds to a variant of (MINSETIP), where all right-hand sides are one but the left-hand side coefficients are uncertain in $\{0,1\}$. Even in comparison to MINSET under uncertainty with deterministic right-hand sides, there are several further difference besides the restriction of the coefficients to values in $\{0, 1\}$. For one, [19] assumes access to the probability distributions. In particular, their algorithms are able to compute certain expected values. For our stochastic setting, we do not have sufficient information to compute expected values and the adversary still has some power in the selection of the unknown distributions as long as it respects the balancing parameter. On the other hand, their SETCOVER variant allows some distributions that are not possible in MINSET. In particular, an interval I_i in MINSET has the same coefficient $a_i = (w_i - L_i)$ in each constraint for a set S with $I_i \in S$. Such a restriction does not exist in the problem considered in [19]. This in a sense makes their problem incomparable to MINSET under uncertainty. Furthermore, [19] analyzes the approximation ratio instead of the competitive ratio. That is, they compare the expected objective value of an algorithm against the expected objective of the best possible algorithm instead of the expected optimum. To that end, they give an *m*-approximation for the stochastic SETCOVER variant. If sets can be added to the solution multiple times while each time drawing a new realization from the same distribution, they give a $\mathcal{O}(\log m)$ -approximation.

Besides related work on stochastic SETCOVER variants, there is previous related work on the more general (stochastic) submodular covering problem (cf., e.g., [1,17,40]). In the submodular covering problem, we are given a ground set of elements E and a submodular function $f: 2^E \to \mathbb{N}_+$. The goal is to find a subset $S \subseteq E$ of minimum cardinality such that f(S) = f(E). This non-stochastic submodular covering problem contains offline MINSET [40], i.e., (MINSETIP) with full knowledge of all coefficients and right-hand sides. To see this, consider an instance $(\mathcal{I}, \mathcal{S})$ of MINSET. We can interpret the intervals as the ground set of elements, i.e., $E = \mathcal{I}$ and use the submodular function $f(Q) = \sum_{S \in \mathcal{S}} \min\{\sum_{I_i \in S \cap Q} w_i - L_i, w^* - L_S\}$ for $Q \subseteq E$. Then, f(E) is the sum of right-hand sides of (MINSETIP) and f(Q) = f(E) holds for a subset

 $Q \subseteq \mathcal{I} = E$ if and only if Q is feasible for (MINSETIP). The best-known algorithm for the submodular covering problem achieves an approximation ratio of $\mathcal{O}(\log(f(E)))$ [40] and no polynomial-time algorithm can be better unless P=NP [10].

In the stochastic submodular covering problem, we are given random variables X_1, \ldots, X_n that independently realize to subsets of E according to known probability distributions. The task is to sequentially and irrevocably add random variables X_i to the solution \mathcal{X} until $f(\bigcup_{X_i \in \mathcal{X}} X_i) = f(E)$. Whenever a random variable X_i is added to the solution, the realization of the variable is revealed. While this general setting is similar to MINSET under uncertainty, there are some differences. In the stochastic submodular covering problem, the value $f(\bigcup_{X_i \in \mathcal{X}} X_i)$ only depends on the realizations of the random variables in \mathcal{X} . For the submodular function defined above for a MINSET instance, the function value f(Q) depends also on the uncertain w^* and, therefore, on elements outside of Q. Furthermore, as the intervals \mathcal{I} in a MINSET instance are continuous, modeling them as a stochastic submodular covering instance would require some form of discretization. Independent of these differences, all results on the stochastic submodular covering problem (to our knowledge) assume known distributions and actively use them, which is in contrast to our stochastic setting. Furthermore, all these results analyze the approximation ratio instead of the competitive ratio. Thus, existing results for the stochastic submodular covering problem cannot directly be applied to our stochastic setting. This also holds for a range of problem variants that have been considered in the literature (see, e.g., [9, 20, 25, 27, 34]).

1.4 Outline

To start the paper, we, in Section 2, consider the special case of MINSET under uncertainty with pairwise disjoint sets. We give a lower bound of $\frac{2}{\tau}$ and a matching upper bound on the competitive ratio for this special case. These bounds nicely illustrate the challenges caused by the uncertainty and the techniques that we use to tackle them, also later on for the general problem.

Afterwards, in Section 3, we move on to the general MINSET and discuss the hardness of approximation as well as approximations of the offline problem variant. Based on observations for the offline problem, we introduce an algorithmic framework that can be used to solve MINSET under uncertainty and other covering problems with uncertainty in the constraints.

For the remaining paper, we show how to implement the framework for MINSET under uncertainty with deterministic right-hand sides (Section 4), for more general covering problems under uncertainty with deterministic right-hand sides (Section 4) and, finally, for the general MINSET (Section 5). Using these implementations and our observations for the special case of disjoint sets, we prove our algorithmic results.

2 Disjoint MINSET

Consider the special case of MINSET where all sets are pairwise disjoint, i.e., $S \cap S' = \emptyset$ for all $S, S' \in S$ with $S \neq S'$. We call this special case *disjoint* MINSET. Disjoint MINSET is of particular interest as it gives lower bounds for several problems under adversarial explorable uncertainty, cf. [12, 33]. To illustrate the challenges posed by having stochastic uncertainty in the input, we give the following lower bound. Recall that the balancing parameter is defined as $\tau = \min_{I_i \in \mathcal{I}} \tau_i$, where τ_i is the probability that w_i is larger than the center of I_i . We use ALG and OPT to refer to an algorithm and an optimal solution, respectively. Slightly abusing the notation, we use the same terms to refer also to the corresponding numbers of queries.

First, we show the following lower bound that even holds for known probability distributions. Afterwards, we prove a slightly stronger bound exploiting unknown distributions.

Theorem 2.1. No deterministic algorithm for MINSET under uncertainty has a competitive ratio better than $\frac{1}{\tau}$, even if all given sets are pairwise disjoint and the distributions are known.



Figure 4: Lower bound example for the set selection problem under explorable uncertainty consisting of the intervals $\mathcal{I} = \{I_0, \ldots, I_n\}$ and the sets $\mathcal{S} = \{S_1, S_2\}$ with $S_1 = \{I_1, \ldots, I_n\}$, $S_2 = \{I_0\}$, $I_0 = \{0.65\}$ and $I_i = (0, 1)$ for $i \in \{1, \ldots, d\}$.

Proof. Consider an instance with the set of uncertainty intervals $\mathcal{I} = \{I_0, I_1, \ldots, I_n\}$ with $I_0 = \{0.65\}$ and $I_i = (0, 1)$ for all $i \neq 0$, and sets $\mathcal{S} = \{S_1, S_2\}$ with $S_1 = \{I_0\}$ and $S_2 = \mathcal{I} \setminus \{I_0\}$. See Figure 4 for an illustration. Define the distributions d_i with $i \neq 1$ as $d_i(a) = (1 - \tau)$ if $a = \epsilon$, $d_i(a) = \tau$ if a = 0.7 and $d_i(a) = 0$ otherwise, for some infinitesimally small $\epsilon > 0$.

If there exists some $I_i \in S_2$ with $w_i = 0.7$, then OPT = 1 as a query to that interval already proves that S_1 is the set of minimum value since $w(S_1) = 0.65 < 0.7$. Otherwise, OPT = n. Therefore, $\mathbb{E}[OPT] = (1 - (1 - \tau)^n) + (1 - \tau)^n \cdot n$ and $\lim_{n \to \infty} \mathbb{E}[OPT] = 1$.

Since I_0 is trivial and all I_i with $i \neq 0$ are identical with the same distribution, each deterministic algorithm ALG will just query the elements of S_2 in some order until it either reaches an I_i with $w_i = 0.7$ or has queried all intervals. This implies that ALG is a geometrical distribution with success probability τ and, therefore, $\mathbb{E}[ALG] = \min\{n, \frac{1}{\tau}\}$. For *n* towards infinity, we get

$$\lim_{n \to \infty} \frac{\mathbb{E}[\text{ALG}]}{\mathbb{E}[\text{OPT}]} = \frac{1}{\tau}$$

Next, we show a slightly stronger bound for unknown distributions. The lower bound instance heavily exploits that, for unknown distributions, the adversary still has some power when selecting the probability distributions.

Theorem 2.2. No deterministic algorithm for MINSET under uncertainty has a competitive ratio better than $\frac{2}{\pi}$, even if all given sets are pairwise disjoint.

Proof. Consider the same instance as in the proof of Theorem 2.1 but with different, now unknown distributions. Since the distributions are unknown, an algorithm cannot distinguish the intervals I_1, \ldots, I_n even if they have different distributions. This means that the adversary still has some power and can set the distributions in a worst-case manner for the algorithm, as long as the distributions respect balancing parameter τ .

To that end, consider a fixed value τ and an arbitrary deterministic algorithm ALG. As ALG cannot distinguish the intervals I_1, \ldots, I_n , we can assume w.l.o.g. that it queries the intervals in order of their indices until the instance is solved. For all 0 < i < n, the adversary sets the distribution to $d_i(a) = \tau$ for a = 0.51, $d_i(a) = (1 - \tau)$ for $a = \epsilon$ and $d_i = 0$ otherwise, for some infinitesimally small $\epsilon > 0$. Finally,

Algorithm 1: Algorithm for disjoint MINSET under uncertainty.

Input: Instance of MINSET under uncertainty with pairwise disjoint sets.

1 $Q \leftarrow \emptyset;$

- 2 while the problem is not solved do
- 3 $S_{\min} \leftarrow \arg\min_{S \in \mathcal{S}} L_S(Q);$
- 4 repeat
- 5 $I_i \leftarrow \arg \max_{I_j \in S_{\min} \setminus Q} U_j L_j;$ Query $I_i; Q \leftarrow Q \cup \{I_i\};$
- 6 **until** $w_i L_i \ge \frac{1}{2} \cdot (U_i L_i)$ or S_{\min} has been completely queried;

the adversary sets distribution d_n to $d_n(a) = 1$ for a = 0.7 and $d_n(a) = 0$ otherwise. These distributions clearly respect the balancing parameter τ .

For these distributions, we always have $w(S_2) > w(S_1)$ as $w_n > w(S_1)$ holds with a probability of one. Thus, every algorithm has to query until the lower limit of set S_2 increases to a value of at least $w(S_1)$. For ALG, this is the case once it found two intervals I_i with i < n and $w_i = 0.51$ or once it queries interval I_n in case no two such intervals exist. Thus, the expected query cost is $\mathbb{E}[ALG] = \min\{\frac{2}{\tau}, n\}$. The optimal solution on the other hand only queries I_n and is done after a single query. Therefore, $\mathbb{E}[OPT] = 1$ and the competitive ratio of ALG is at least $\min\{\frac{2}{\tau}, n\}$. We can conclude the theorem by picking a sufficiently large value for n.

We continue by giving a quite simple algorithm for disjoint MINSET that matches the lower bound of Theorem 2.2.

In disjoint MINSET, each I_i occurs in exactly one constraint for one set S in the corresponding (MINSETIP). Thus, each set S defines a disjoint subproblem and the optimal solution OPT of the instance is the union of optimal solutions for the subproblems. The optimal solution for a subproblem S is to query the elements of $I_i \in S$ in order of non-decreasing $(w_i - L_i)$ until the sum of those coefficients is at least $(w^* - L_S)$.

Under uncertainty, we adapt this strategy and query in order of non-decreasing $(U_i - L_i)$. While this does not guarantee that we query the interval with maximum $(w_j - L_j)$ in S, it gives us a probability of τ to query an interval I_i such that $(w_i - L_i)$ is at least half the maximum $(w_j - L_j)$. We will prove that this is sufficient to achieve the guarantee. Since we do not know w^* , we do not know when to stop querying in a subproblem. We handle this by only querying in the set S of minimum current lower limit as the subproblem for this set is clearly not yet solved. Algorithm 1 formalizes this approach.

Theorem 2.3. There is an algorithm for disjoint MINSET under uncertainty with competitive ratio at most $\frac{2}{\tau}$.

Proof. Consider a fixed realization of values w_i and the corresponding (MINSETIP) instance. For each $S \in S$, a feasible solution Q must satisfy $\sum_{I_i \in S \cap Q} (w_i - L_i) \ge (w^* - L_S)$. This implies $|Q \cap S| \ge |P_S^*|$ for the minimum cardinality prefix P_S^* of I_1, \ldots, I_k with $\sum_{I_i \in P_S^*} (w_i - L_i) \ge w^* - L_S$, where $S = \{I_1, \ldots, I_k\}$ and all I_i are indexed by non-increasing $w_i - L_i$. As the sets are disjoint, we get $OPT = \sum_{S \in S} |P_S^*|$.

Using this, we show that Algorithm 1 satisfies the theorem. To that end, let X_j be a random variable denoting the number of queries in iteration j of the outer while-loop of Algorithm 1 and let Y_j be an indicator variable indicating whether iteration Y_j is actually executed ($Y_j = 1$) or not ($Y_j = 0$).

We prove the theorem by separately showing $\sum_{j} \mathbb{P}[Y_j = 1] \leq 2 \cdot \mathbb{E}[OPT]$ and $\mathbb{E}[X_i \mid Y_j = 1] \leq \frac{1}{\tau}$.

Since

$$\mathbb{E}[ALG] = \sum_{j} \mathbb{E}[X_j] = \sum_{j} \mathbb{P}[Y_j = 0] \mathbb{E}[X_j \mid Y_j = 0] + \sum_{j} \mathbb{P}[Y_j = 1] \mathbb{E}[X_j \mid Y_j = 1]$$
$$= \sum_{j} \mathbb{P}[Y_j = 1] \mathbb{E}[X_j \mid Y_j = 1]$$

follows from $\mathbb{E}[X_j \mid Y_j = 0] = 0$ and the law of total expectations, the two inequalities imply the theorem.

Note that $\sum_{j} \mathbb{P}[Y_j = 1]$ is just the expected number of iterations of the algorithm. Thus, if we show for each realization of precise values that the number of iterations is at most 2 · OPT, we directly get $\sum_{j} \mathbb{P}[Y_j = 1] \leq 2 \cdot \mathbb{E}[\text{OPT}].$

Consider a fixed realization. For each S, let h_S denote the number of iterations with $S_{\min} = S$. We claim that $h_S \leq 2 \cdot |P_S^*|$. Then, $OPT \geq \sum_{S_\ell \in S} |P_S^*|$ implies $\sum_j \mathbb{P}[Y_j = 1] \leq 2 \cdot \mathbb{E}[OPT]$.

Let j be an iteration with $S_{\min} = S$, let G_j denote the queries of this iteration and let Q_j denote the set of all previous queries. Observe that $P_S^* \setminus Q_j \neq \emptyset$. Otherwise, the definition of P_S^* would either imply that the lower limit of S after querying Q_j is larger than w^* , which contradicts $S_{\min} = S$, or that the lower limit is equal to w^* , which implies that the problem is already solved.

We argue that $\sum_{I_i \in G_j} w_i - L_i \ge \frac{1}{2} \cdot \max_{I_i \in S \setminus Q_j} w_i - L_i$. In case that interval $I_i = \arg \max_{I_i \in S \setminus Q_j} w_i - L_i$ is contained in G_j , the inequality clearly holds. Otherwise, let $I_{i'}$ denote the last element that is queried in the iteration. Then, $w_{i'} - L_{i'} \ge \frac{1}{2} \cdot (U_{i'} - L_{i'}) \ge \frac{1}{2} \cdot \max_{I_i \in S \setminus Q_j} w_i - L_i$, where the first inequality holds as $I_{i'}$ is the last query of the iteration and the second inequality holds by the order in which the elements of S are queried by the algorithm. Thus, this last interval $I_{i'}$ alone satisfies the inequality.

The inequality suffices to conclude that $\sum_{I_i \in Q_j \cap S} w_i - L_i \ge \sum_{I_i \in P_S^*} w_i - L_i$ holds after at most $2 \cdot |P_S^*|$ iterations with $S_{\min} = S$. As $S_{\min} \neq S$ holds for all following iterations, the claim $h_S \le 2 \cdot |P_S^*|$ follows. Next, we show the second inequality, $\mathbb{E}[X_j \mid Y_j = 1] \le \frac{1}{\tau}$. If an iteration j of the outer while-loop is executed $(Y_j = 1)$, the repeat statement queries intervals I_i until either $(w_i - L_i) \ge \frac{1}{2} \cdot (U_i - L_i)$ or $S_{\min} \subseteq Q$. Thus, it terminates at the latest when it finds an I_i with $(w_i - L_i) \ge \frac{1}{2} \cdot (U_i - L_i)$. The number of queries until such an interval occurs is described by a geometric distribution with success probability at least τ . So, in expectation, this number is at most $\frac{1}{\tau}$ and we can conclude $\mathbb{E}[X_j \mid Y_j = 1] \le \frac{1}{\tau}$.

We remark that Theorems 2.2 and 2.3 imply that, even with full knowledge of the distributions, the competitive ratio for disjoint MINSET cannot be improved by more than a factor of two compared to the ratio with unknown distributions.

3 Algorithmic framework

In the previous section, we have seen an algorithm for disjoint MINSET under uncertainty with a tight competitive ratio. The key observation that allowed us to achieve that ratio was the simple characterization of an (offline) optimal solution. In this section, we consider the offline variant of the general MINSET and give inapproximability results that prevent such simple characterizations for optimal solutions of the general problem. Thus, we need alternative algorithms and, based on observations for the offline problem, present an algorithmic framework that can be used to solve MINSET and other covering problems under uncertainty.

3.1 Offline Problems and Hardness of Approximation

We refer to the problem of solving (MINSETIP) with full knowledge of the precise values w_i (and w^*) as *offline* problem. This means that we have full knowledge of all coefficients of the ILP. For MINSET under uncertainty, we say that a solution is optimal, if it is an optimal solution for the corresponding offline

problem. We use OPT to refer to an optimal solution and, slightly abusing the notation, to its objective value.

Offline MINSET contains the classical SETCOVER problem and, thus, it is as hard to approximate. This result transfers to the stochastic setting, even for uniform distributions. Results by Dinur and Steurer [10] imply the following, as we formally prove in Appendix C.

Theorem 3.1. For any fixed $\alpha > 0$, it is NP-hard to compute a query strategy that is $(1 - \alpha) \cdot \ln m$ competitive for MINSET under uncertainty even if the precise value w_i of each I_i is drawn independently and uniformly at random from (L_i, U_i) . The same inapproximability holds also for offline MINSET.

On the positive side, we can approximate offline MINSET by adapting covering results (see, e.g., [8, 11, 28, 29, 35]). In particular, we want to use greedy algorithms that iteratively and irrevocably add elements to the solution that are selected by a certain greedy criterion. Recall that "adding an element to the solution" corresponds to both, setting the variable x_i of an interval $I_i \in \mathcal{I}$ in (MINSETIP) to one and querying I_i . While we are technically not restricted to greedy algorithms when solving offline MINSET, our goal is to later on generalize the offline algorithm to the setting with uncertainty and irrevocable decisions. Hence, greedy algorithms seem to be a suitable choice.

Since the greedy criterion for adding an element depends on previously added elements, we define a version of (MINSETIP) that is parametrized by the set $Q \subseteq \mathcal{I}$ of elements that have already been added to the solution and adjust the right-hand sides to the remaining covering requirement after adding Q. Recall that $a_i = w_i - L_i$ and $b_S = w^* - L_S$. Here, $b_S(Q) = \max\{b_S - \sum_{I_i \in Q \cap S} a_i, 0\}$ and $b(Q) = \sum_{S \in S} b_S(Q)$.

$$\begin{array}{ll} \min & \sum_{I_i \in \mathcal{I} \setminus Q} x_i \\ \text{s.t.} & \sum_{I_i \in S \setminus Q} x_i \cdot a_i & \geq b_S(Q) \quad \forall S \in \mathcal{S} \\ & x_i & \in \{0, 1\} \quad \forall I_i \in \mathcal{I} \setminus Q \end{array}$$
 (MINSETIP-Q)

Based on this ILP, we adjust the algorithm by Dobson [11] for the multiset multicover problem to our setting (cf. Algorithm 2). The algorithm scales the coefficients such that all non-zero left-hand side coefficients are at least 1. We refer to such instances as *scaled*. Then it greedily adds the element to the solution that reduces the right-hand sides the most, i.e., the interval $I_i \in \mathcal{I} \setminus Q$ that maximizes $g_c(Q, I_i) = b'(Q) - b'(Q \cup \{I_i\})$ (a' and b' indicate scaled coefficients). For a subset $G \subseteq \mathcal{I}$, we define $g_c(Q, G) = b'(Q) - b'(Q \cup \{I_i\})$

After $b'_S(Q) < 1$ for all $S \in S$, we can exploit that all scaled non-zero coefficients a'_i are at least one. This means that adding an element $I_i \in \mathcal{I} \setminus Q$ satisfies all remaining constraints of sets S with $I_i \in S$. Thus, the remaining problem reduces to a SETCOVER instance, which can be solved by using the classical greedy algorithm by Chvatal [8]. This algorithm greedily adds the element $I_i \in \mathcal{I} \setminus Q$ that maximizes $g_s(Q, I_i) = A(Q) - A(Q \cup \{I_i\})$ with $A(Q) = |\{S \in S \mid b'_S(Q) > 0\}|$, i.e., the element that satisfies the largest number of constraints that are not already satisfied by Q. For a subset $G \subseteq \mathcal{I}$, we define $g_s(Q, G) = A(Q) - A(Q \cup G)$.

During the course of this paper, we refer to $g_c(Q, I_i)$, $g_s(Q, I_i)$, $g_c(Q, G)$ and $g_s(Q, G)$ as the greedy values of I_i and G, respectively.

Theorem 3.2 (Follows from Dobson [11]). Algorithm 2 is a polynomial-time $\mathcal{O}(\log m)$ -approximation for offline MINSET. The precise approximation factor is $\rho(\gamma) = \lceil \ln(\gamma \cdot m \cdot \max_S(w^* - L_S)) \rceil + \lceil \ln(m) \rceil$ with $s_{\min} = \min_{I_i \in \mathcal{I}: (w_i - L_i) > 0} (w_i - L_i), \ \gamma = 1/s_{\min} \text{ and } m = |\mathcal{S}|.$

During the remaining course of the paper, we will state the competitive ratios of our algorithms in terms of ρ . To that end, define $\bar{\rho}(\gamma) = \lceil \ln(\gamma \cdot m \cdot \max_{S,S'}(U_S - L_{S'}) \rceil + \lceil \ln(m) \rceil$, which is an upper bound on $\rho(\gamma)$. Under uncertainty, we compare against $\bar{\rho}$ to avoid the random variable w^* . For constant U_i 's, $\bar{\rho}$ and ρ are asymptotically the same.

Algorithm 2: Greedy algorithm by Dobson [11] for offline MINSET.
Input: An instance of offline MINSET, i.e., an instance of (MINSETIP)
$1 \ s_{\min} = \min_{I_i \in \mathcal{I}: \ a_i > 0} a_i; \forall S \in \mathcal{S}: b'_S = \frac{b_S}{s_{\min}}; \forall I_i \in \mathcal{I}: \ a'_i = \frac{a_i}{s_{\min}};$
2 while $\exists S \in \mathcal{S} \colon b_S'(Q) \ge 1$ do
3 $I_i \leftarrow \arg \max_{I_j \in \mathcal{I} \setminus Q} g_c(Q, I_j); $ Query $I_i; Q \leftarrow Q \cup \{I_i\};$
4 while the problem is not solved do
5 $I_i \leftarrow \arg \max_{I_i \in \mathcal{I} \setminus Q} g_s(Q, I_j);$ Query $I_i; Q \leftarrow Q \cup \{I_j\};$

We remark again that the approximation ratio of Algorithm 2 has dependencies on the numerical input parameter s_{\min} and $\max_S(w^* - L_S)$. While there exist algorithms that achieve an approximation ratio of $\mathcal{O}(\log m)$ for the offline problem without such dependencies [28, 29], these algorithms are not greedy and it remains open whether there exist algorithms with this improved ratio that execute irrevocable decisions, even with full knowledge of the coefficients. Thus, we consider Algorithm 2 and aim at extending it for the setting under uncertainty.

3.2 Algorithmic framework

We introduce our algorithmic framework that we use to solve MINSET under uncertainty. Ideally, we would like to apply the offline greedy algorithm. However, since the coefficients $a_i = w_i - L_i$ and $b_S = w^* - L_S$ are unknown, we cannot apply Algorithm 2 to solve MINSET under uncertainty as we cannot compute the element that maximizes the greedy value g_c or g_s .

While we cannot precisely compute the greedy choice, our strategy is to *approximate* it and to show that approximating the greedy choice is sufficient to obtain the desired guarantees. To make this more precise, consider an *iterative* algorithm for (MINSETIP), i.e., an algorithm that iteratively adds pairwise disjoint subsets G_1, \ldots, G_h of \mathcal{I} to the solution. For each j, let $Q_j = \bigcup_{1 \le j' \le j-1} G_{j'}$, i.e., Q_j contains the elements that have been added to the solution before G_j . If the combined greedy value of G_j is within a factor of α to the best greedy value for the problem instance *after* adding Q_j , then we say that $G_j \alpha$ -approximates the greedy choice. The following technical definition makes this more precise and the subsequent lemma connects the definition to the actual greedy values while taking into account that there are two different greedy values g_c and g_s (cf. Algorithm 2).

Definition 3.3. For a (MINSETIP) instance with scaled coefficients and optimal solution OPT, let $\alpha \in \mathbb{R}_{\geq 1}$ and consider the corresponding instance of (MINSETIP-Q) for some $Q \subseteq \mathcal{I}$. A set $G \subseteq \mathcal{I} \setminus Q$ α -approximates the current greedy choice after adding Q if either

1.
$$A(Q \cup G) \leq (1 - \frac{1}{\alpha \cdot \text{OPT}}) \cdot A(Q) \text{ or }$$

2. $b'(Q) \ge 1$ and $b'(Q \cup G) \le (1 - \frac{1}{\alpha \cdot QPT}) \cdot b'(Q)$.

Intuitively, the two conditions of the following lemma seem like a more appropriate definition of approximating a greedy choice. While the conditions of the lemma imply that the definition above is satisfied, in our proofs it will sometimes be easier to directly show that the definition is satisfied, without using the lemma. Therefore, we use the more technical Definition 3.3 but the lemma captures the intuition behind the definition.

Lemma 3.4. For a scaled instance of (MINSETIP), $Q \subseteq \mathcal{I}$, $\alpha \geq 1$ and $G \subseteq \mathcal{I} \setminus Q$:

1. If $b'_S(Q) < 1$ for all $S \in S$ and $g_s(Q,G) \ge \frac{1}{\alpha} \cdot \max_{I_i \in \mathcal{I} \setminus Q} g_s(Q,I_i)$, then G satisfies the first condition of Definition 3.3 and, thus, α -approximates the greedy choice.

2. If $b'(Q) \ge 1$ and $g_c(Q, G) \ge \frac{1}{\alpha} \cdot \max_{I_i \in \mathcal{I} \setminus Q} g_c(Q, I_i)$, then G satisfies the second condition of Definition 3.3 and, thus, α -approximates the greedy choice.

Proof. First, assume that $b'_S(Q) < 1$ for all $S \in S$ and consider a set $G \subseteq \mathcal{I} \setminus Q$ with $g_s(Q,G) \ge \frac{1}{\alpha} \cdot \max_{I_i \in \mathcal{I} \setminus Q} g_s(Q,I_i)$.

Let $I^* = \arg \max_{I_i \in \mathcal{I} \setminus Q} A(Q) - A(Q \cup \{I_i\}) = \arg \max_{I_i \in \mathcal{I} \setminus Q} g_s(Q, I_i)$. By assumption $b'_S(Q) < 1$ for all $S \in S$ and, as we consider a scaled instance, $a'_i \ge 1$ for all $I_i \in \mathcal{I}$. Thus, the remaining instance is a set cover instance as adding an interval I_i to the solution satisfies all constraints S with $I_i \in S$ that have not already been satisfied by Q.

Let OPT_Q denote the optimal solution for the remaining instance after adding Q to the solution, i.e., the optimal solution to (MINSETIP-Q). Using a standard set cover argument, we can observe that $\frac{A(Q)}{OPT_Q} \leq A(Q) - A(Q \cup \{I^*\})$ as the optimal solution satisfies the remaining constraints at cost OPT_Q , but a single interval can satisfy at most $A(Q) - A(Q \cup \{I^*\})$ constraints. Note that this argument only holds because all left-hand side coefficients are at least as large as the right-hand sides. Otherwise, adding an interval I_i later, i.e., after $Q' \supset Q$ has already been added to the solution, could satisfy more constraints, i.e., $A(Q) - A(Q \cup \{I_i\}) < A(Q') - A(Q' \cup \{I_i\})$. This is one of the reasons why the offline greedy algorithm uses two greedy criteria.

By assumption and definition of g_s , we have $\alpha \cdot (A(Q) - A(Q \cup G)) \ge \max_{I_i \in \mathcal{I} \setminus Q} A(Q) - A(Q \cup \{I_i\})$ and, therefore, $\frac{A(Q)}{OPT_Q} \le \alpha \cdot (A(Q) - A(Q \cup G))$. Rearranging the latter inequality, we obtain $A(Q \cup G) \le A(Q) \cdot \left(1 - \frac{1}{\alpha OPT_Q}\right)$. Since $OPT \ge OPT_Q$ for the optimal solution OPT of the complete instance, we get $A(Q \cup G) \le A(Q) \cdot \left(1 - \frac{1}{\alpha OPT}\right)$. This implies that G satisfies the first condition of Definition 3.3. Next, assume $b'(Q) \ge 1$ and consider a set $G \subseteq \mathcal{I} \setminus Q$ with $g_c(Q, G) \ge \frac{1}{\alpha} \cdot \max_{I_i \in \mathcal{I} \setminus Q} g_c(Q, I_i)$.

Let $I^* = \arg \max_{I_i \in \mathcal{I} \setminus Q} b'(Q) - b'(Q \cup \{I_i\}) = \arg \max_{I_i \in \mathcal{I} \setminus Q} g_c(Q, I_i)$. Observe that $\frac{b'(Q)}{OPT_Q} \leq b'(Q) - b'(Q \cup \{I^*\})$ as the optimal solution covers the remaining constraints at cost OPT_Q , but a single interval can decrease the total slack between left-hand and right-hand sides of (MINSETIP-Q) by at most $b'(Q) - b'(Q \cup \{I^*\})$. By assumption and definition of g_c , we have $\alpha \cdot (b'(Q) - b'(Q \cup G)) \geq \max_{I_i \in \mathcal{I} \setminus Q} b'(Q) - b'(Q \cup \{I_i\})$ and, therefore, $\frac{b'(Q)}{OPT_Q} \leq \alpha \cdot (b'(Q) - b'(Q \cup G))$. Rearranging the latter inequality, we obtain $b'(Q \cup G) \leq b'(Q) \cdot \left(1 - \frac{1}{\alpha OPT_Q}\right)$. Since $OPT \geq OPT_Q$ for the optimal solution OPT of the complete instance, we get $b'(Q \cup G) \leq b'(Q) \cdot \left(1 - \frac{1}{\alpha OPT_Q}\right)$. This and the assumption $b'(Q) \geq 1$ imply that G satisfies the second condition of Definition 3.3. Note that the argument for the second case does not use that all non-zero coefficients are at least one. Thus, the statement also holds if there are coefficients $0 < a'_i < 1$.

With the following lemma, we bound the number of iterations j in which $G_j \alpha$ -approximates the current greedy choice via an adjusted set cover greedy analysis.

Lemma 3.5. Consider an arbitrary algorithm for (MINSETIP) that scales the coefficients by factor γ and iteratively adds disjoint subsets G_1, \ldots, G_h of \mathcal{I} to the solution until the instance is solved. The number of groups G_j that α -approximate the current greedy choice (after adding $Q_j = \bigcup_{1 \le j' \le j-1} G_{j'}$) is at most $\alpha \cdot \rho(\gamma) \cdot \text{OPT}$.

Proof. We first show that the number of iterations j with $b'(Q_j) \ge 1$ and $b'(Q_j \cup G_j) \le (1 - \frac{1}{\alpha \cdot \text{OPT}}) \cdot b'(Q_j)$, i.e., the number of iterations that satisfy the second condition of Definition 3.3, is at most $\alpha \lceil \ln(\gamma \cdot m \cdot \max_S(w^* - L_S)) \rceil \cdot \text{OPT}$.

Let $G_1, \ldots, G_k \subseteq \mathcal{I}$ denote the sets that are added to the solution by the algorithm *and* satisfy the second condition of Definition 3.3. Assume that the sets are indexed in the order they are added. For each $j \in \{1, \ldots, k\}$, let $\bar{Q}_j \subseteq \mathcal{I}$ denote the set of intervals that are added to the solution *before* \bar{G}_j . Note that

 $\{\bar{G}_1, \ldots, \bar{G}_{j-1}\} \subseteq \bar{Q}_j$, but \bar{Q}_j might contain additional added groups that just do not satisfy the second condition of Definition 3.3.

By assumption, $b'(\bar{Q}_j \cup \bar{G}_j) \leq (1 - \frac{1}{\alpha \cdot \text{OPT}}) \cdot b'(\bar{Q}_j)$. A recursive application of this inequality and the fact that $(1 - x) < e^{-x}$ for all $x \in \mathbb{R} \setminus \{0\}$ implies:

$$b'(\bar{Q}_j \cup \bar{G}_j) \le b'(\emptyset) \cdot \left(1 - \frac{1}{\alpha \text{OPT}}\right)^j < b'(\emptyset) \cdot e^{-\frac{j}{\alpha \text{OPT}}}.$$

Thus, after $j = \alpha \cdot \text{OPT} \cdot \lceil \ln b'(\emptyset) \rceil$ iterations that satisfy the second condition of Definition 3.3, we have $b'(\bar{Q}_j \cup \bar{G}_j) < b'(\emptyset) \cdot e^{-\ln b'(\emptyset)} = 1$. But if $b'(\bar{Q}_j \cup \bar{G}_j) < 1$, then there can be no further iteration that satisfies the second condition of Definition 3.3. Thus, the number of such iterations is at most $\alpha \cdot \text{OPT} \cdot \lceil \ln b'(\emptyset) \rceil$. Since $b'(\emptyset)$ is upper bounded by $\gamma \cdot m \cdot \max_S(w^* - L_S)$ as we have m constraints with scaled right-hand side values of at most $\gamma \cdot \max_S(w^* - L_S)$, the number of such iterations is at most $\alpha \cdot \lceil \ln(\gamma \cdot m \cdot \max_S(w^* - L_S) \rceil \cdot \text{OPT}$.

Next, we show that the number of iterations j with $A(Q_j \cup G_j) \le (1 - \frac{1}{\alpha \cdot \text{OPT}}) \cdot A(Q_j)$, i.e., the number of iterations that satisfy the first condition of Definition 3.3, is at most $\alpha \lceil \ln(m) \rceil \cdot \text{OPT}$.

Let $G_1, \ldots, G_k \subseteq \mathcal{I}$ denote the sets that are added to the solution by the algorithm and satisfy the first condition of Definition 3.3. Assume that the sets are indexed in the order they are added. For each $j \in \{1, \ldots, k\}$, let $\bar{Q}_j \subseteq \mathcal{I}$ again denote the set of intervals that are added to the solution before \bar{G}_j . Note that $\{\bar{G}_1, \ldots, \bar{G}_{j-1}\} \subseteq \bar{Q}_j$, but \bar{Q}_j might contain additional sets that just do not satisfy the first condition of Definition 3.3.

By assumption, $A(\bar{Q}_j \cup \bar{G}_j) \leq (1 - \frac{1}{\alpha \cdot \text{OPT}}) \cdot A(\bar{Q}_j)$. A recursive application of this inequality and the fact that $(1 - x) < e^{-x}$ for all $x \in \mathbb{R} \setminus \{0\}$ implies:

$$A(\bar{Q}_j \cup \bar{G}_j) \le A(\emptyset) \cdot \left(1 - \frac{1}{\alpha \text{OPT}}\right)^j < A(\emptyset) \cdot e^{-\frac{j}{\alpha \text{OPT}}}.$$

Thus, after $j = \alpha \cdot \text{OPT} \cdot \lceil \ln A(\emptyset) \rceil$ such iterations, we have $A(\bar{Q}_j \cup \bar{G}_j) < A(\emptyset) \cdot e^{-\ln A(\emptyset)} = 1$. But if $A(\bar{Q}_j \cup \bar{G}_j) < 1$, then $A(\bar{Q}_j \cup \bar{G}_j) = 0$ and the instance is solved and no further iteration is executed. Since $A(\emptyset)$ is upper bounded by the number of constraints m, the number of iterations that satisfy the first condition of Definition 3.3 is at most $\alpha \lceil \ln(m) \rceil \cdot \text{OPT}$.

In total, at most $\alpha \lceil \ln(m) \rceil \cdot \text{OPT}$ iterations satisfy the first condition of Definition 3.3 and at most $\alpha \cdot \lceil \ln(\gamma \cdot m \cdot \max_S(w^* - L_S) \rceil \cdot \text{OPT}$ iterations satisfy the second condition of Definition 3.3. In summation, there are at most $\alpha \cdot (\lceil \ln(\gamma \cdot n \cdot \max_{e \in E} b_e) \rceil + \lceil \ln(n) \rceil) \cdot \text{OPT} = \alpha \cdot \rho(\gamma) \cdot \text{OPT}$ iterations that satisfy Definition 3.3.

The lemma states that the number of groups G_j that α -approximate their greedy choice is within a factor of α of the performance guarantee $\rho(\gamma)$ of the offline greedy algorithm. If each $G_j \alpha$ -approximates its greedy choice, the iterative algorithm achieves an approximation factor of $\max_j |G_j| \cdot \alpha \cdot \rho(\gamma)$. Thus, approximating the greedy choices by a constant factor using a constant group size is sufficient to only lose a constant factor compared to the offline greedy algorithm.

This insight gives us a framework to solve MINSET under uncertainty. Recall that the w_i 's (and by extension the a_i 's and b_s 's) are uncertain and only revealed once we irrevocably add an $I_i \in \mathcal{I}$ to the solution. We refer to a revealed w_i as a *query result*, and to a fixed set of revealed w_i 's for all $I_i \in \mathcal{I}$ as a *realization of query results*. Consider an iterative algorithm. The sets G_j can be computed and queried adaptively and are allowed to depend on (random) query results from previous iterations. Hence, $X_j = |G_j|$ is a random variable. Let Y_j be an indicator variable denoting whether the algorithm executes iteration j $(Y_j = 1)$ or terminates beforehand $(Y_j = 0)$. We define the following class of iterative algorithms and show that algorithms from this class achieve certain guarantees.

Definition 3.6. An iterative algorithm is (α, β, γ) -GREEDY if it satisfies:

- 1. For every realization of query results; each $G_j \alpha$ -approximates the greedy choice as characterized by Q_j on the instance with coefficients scaled by γ .
- 2. $\mathbb{E}[X_j \mid Y_j = 1] \leq \beta$ holds for all iterations j.

Theorem 3.7. Each (α, β, γ) -GREEDY algorithm for MINSET under uncertainty achieves a competitive ratio of $\alpha \cdot \beta \cdot \bar{\rho}(\gamma) \in \mathcal{O}(\alpha \cdot \beta \cdot \log(m))$.

Proof. Consider an (α, β, γ) -GREEDY algorithm ALG for MINSET. The expected cost of ALG is $\mathbb{E}[ALG] = \sum_{j} \mathbb{E}[X_{j}]$. Using the total law of expectations, we get

$$\mathbb{E}[ALG] = \sum_{j} \mathbb{P}[Y_j = 1] \mathbb{E}[X_j \mid Y_j = 1] + \mathbb{P}[Y_j = 0] \mathbb{E}[X_j \mid Y_j = 0]$$
$$= \sum_{j} \mathbb{P}[Y_j = 1] \mathbb{E}[X_j \mid Y_j = 1],$$

where the last inequality holds as $\mathbb{E}[X_j | Y_j = 0] = 0$ by definition (if the algorithm terminates before iteration *j*, then it adds no more elements to the solution and, thus, $X_j = 0$). By the second property of Definition 3.6, this implies $\mathbb{E}[ALG] \leq \beta \cdot \sum_j \mathbb{P}[Y_j = 1]$.

Thus, it remains to bound $\sum_{j} \mathbb{P}[Y_j = 1]$, which corresponds to the expected number of iterations of ALG. Consider a fixed realization of query results, then, by the first property of (α, β, γ) -GREEDY, each $G_j \alpha$ -approximates its greedy choice for the (MINSETIP) instance of the realization scaled by factor γ . Then, Lemma 3.5 implies that the number of iterations is at most $\alpha \cdot \rho(\gamma) \cdot \text{OPT}$, which is upper bounded by $\alpha \cdot \bar{\rho}(\gamma) \cdot \text{OPT}$. As this upper bound on the number of iterations holds for every realization and OPT is the only random variable of that term (since we substituted ρ with $\bar{\rho}$), we can conclude $\sum_{j} \mathbb{P}[Y_j = 1] \leq \alpha \cdot \bar{\rho}(\gamma) \cdot \mathbb{E}[\text{OPT}]$, which implies $\mathbb{E}[\text{ALG}] \leq \alpha \cdot \beta \cdot \bar{\rho}(\gamma) \cdot \mathbb{E}[\text{OPT}]$.

4 MINSET with Deterministic Right-Hand Sides and More Covering Problems

We first introduce an algorithm for a variant of MINSET under uncertainty, where the right-hand sides b_S of the ILP representation (MINSETIP) are deterministic and explicit part of the input. Afterwards, we generalize the algorithm for more general covering problems and a different balancing parameter.

4.1 MINSET With Deterministic Right-Hand Sides

We consider a variant of MINSET under uncertainty, where the right-hand sides b_S of the ILP representation (MINSETIP) are deterministic and explicit part of the input. Thus, only the coefficients $a_i = (w_i - L_i)$ remain uncertain within the interval $(0, U_i - L_i)$. For this problem variant, it can happen that the instance has no feasible solution. In that case, we require every algorithm (including OPT) to reduce the covering requirements as much as possible. As we consider the stochastic problem variant, recall that the balancing parameter is defined as $\tau = \min_{I_i \in \mathcal{I}} \tau_i$ for $\tau_i = \mathbb{P}[w_i \ge \frac{U_i + L_i}{2}]$.

Theorem 4.1. There is an algorithm for MINSET under uncertainty with deterministic right-hand sides and a competitive ratio of $\frac{2}{\tau} \cdot \rho(\gamma) \in \mathcal{O}(\frac{1}{\tau} \cdot \log m)$ with $\gamma = 2/s_{\min}$ for $s_{\min} = \min_{I_i \in \mathcal{I}: U_i - L_i > 0} U_i - L_i$.

Algorithm 3: MINSET with deterministic right-hand sides.

Input: Instance of MINSET with deterministic right-hand sides. $Q = \emptyset$; Scale *a* and *b* by $\frac{2}{s_{\min}}$ to *a'* and *b'* for $s_{\min} = \min_{I_i \in \mathcal{I}: |U_i - L_i| > 0} |U_i - L_i|$; 2 while the problem is not solved do **if** $b'(Q) \ge 1$ then $g = \bar{g}_c$ else $g = \bar{g}_s$; **if** $b'(Q) \ge 1$ then $g = \bar{g}_c$ else $g = \bar{g}_s$; **if** $I_i \leftarrow \arg \max_{I_j \in \mathcal{I} \setminus Q} g(Q, I_j)$; Query $I_i; Q \leftarrow Q \cup \{I_i\}$; **until** the problem is solved or $w_i - L_i \ge \frac{1}{2} \cdot (U_i - L_i)$;

The algorithm of the theorem loses only a factor $\frac{2}{\tau}$ compared to the greedy approximation factor $\rho(\gamma)$ on the corresponding offline problem. We show the theorem by proving that Algorithm 3 is an (α, β, γ) -GREEDY algorithm for $\alpha = 2$, $\beta = \frac{1}{\tau}$ and $\gamma = \frac{2}{s_{\min}}$ with $s_{\min} = \min_{I_i \in \mathcal{I}: U_i - L_i > 0} U_i - L_i$. Then, Theorem 3.7 implies the theorem. We remark that we scale by $\frac{2}{s_{\min}}$ instead of $\frac{1}{s_{\min}}$ because of technical reasons that will become clear in the proof of the theorem.

The algorithm scales the coefficients by factor γ ; we use a' and b' to refer to the scaled coefficients. The idea of Algorithm 3 is to execute the greedy Algorithm 2 under the assumption that $a_i = U_i - L_i$ (and $a'_i = \gamma a_i$) for all $I_i \in \mathcal{I}$ that were not yet added to the solution. As $a_i = (w_i - L_i) \in (0, U_i - L_i)$, this means that we assume a_i to have the largest possible value. Consequently, s_{\min} is the smallest (non-zero) coefficient a_i under this assumption. The algorithm computes the greedy choice based on the *optimistic greedy values*

$$\bar{g}_c(Q, I_i) = \sum_{S \in \mathcal{S}: \ I_i \in S} b'_S(Q) - \max\{0, b'_S(Q) - \gamma(U_i - L_i)\}$$

(if $b'(Q) \ge 1$) and

$$\bar{g}_s(Q, I_i) = |\{S \in \mathcal{S} : I_i \in S \mid b'_S(Q) > 0 \land b'_S(Q) - \gamma(U_i - L_i) \le 0\}|$$

(otherwise). That is, the greedy values under the assumption $a_i = U_i - L_i$. We call these values optimistic as they might overestimate but never underestimate the actual greedy values. For subsets $G \subseteq \mathcal{I}$, we define $\bar{g}_s(Q,G)$ and $\bar{g}_c(Q,G)$ analogously.

In contrast to g_s and g_c , Algorithm 3 has sufficient information to compute \bar{g}_s and \bar{g}_c , and, therefore, the best greedy choice based on the optimistic greedy values. The algorithm is designed to find, in each iteration, an element I_i with $g_c(Q, I_i) \ge \frac{1}{2} \cdot \bar{g}_c(Q, I_i)$ for the current Q (or analogously for \bar{g}_s and g_s). We show that (i) this ensures that each iteration 2-approximates the greedy choice and (ii) that finding such an element takes only $\frac{1}{\tau}$ tries in expectation.

Proof of Theorem 4.1. Let j be an arbitrary iteration of the outer while-loop, X_j denote the number of queries during the iteration, and Y_j indicate whether the algorithm executes iteration j ($Y_j = 1$) or not ($Y_j = 0$).

Assuming $Y_j = 1$, the algorithm during iteration j executes queries to elements I_i until either $w_i - L_i \geq \frac{1}{2}(U_i - L_i)$ or the problem is solved. Since $w_i \geq \frac{(U_i + L_i)}{2}$ implies $w_i - L_i \geq \frac{1}{2}(U_i - L_i)$ and $\mathbb{P}[w_i \geq \frac{(U_i + L_i)}{2}] \geq \tau$ holds by assumption, the number of attempts until the current I_i satisfies the inequality follows a geometric distribution with success probability at least τ . Hence, $\mathbb{E}[X_j | Y_j = 1] \leq \frac{1}{\tau}$; proving Property 2 of Definition 3.6.

We continue by proving Property 1 of Definition 3.6. Consider a fixed realization. Let \bar{G}_j denote the queries of iteration *j* except the last one and let $I_{\bar{j}}$ denote the last query of iteration *j*. Then $G_j = \bar{G}_j \cup \{I_{\bar{j}}\}$

is the set of queries during the iteration. Finally, let Q_j denote the set of queries before iteration j. We show that G_j 2-approximates the greedy choice of the scaled instance, which implies Property 1 of Definition 3.6.

If the iteration solves the problem, then G_j clearly 1-approximates the greedy choice and we are done. Thus, assume otherwise. We distinguish between the two cases (1) $b'(Q_j) \ge 1$ and (2) $b'(Q_j) < 1$.

Case (1): We show first that G_j 2-approximates the greedy choice if $b'(Q_j) \ge 1$. In this case, we have $g = \overline{g}_c$ (cf. Line 3). By choice of $I_{\overline{j}}$, we have $\overline{g}_c(Q_j \cup \overline{G}_j, I_{\overline{j}}) = \max_{I_i \in \mathcal{I} \setminus (Q_j \cup \overline{G}_j)} \overline{g}_c(Q_j \cup \overline{G}_j, I_i)$, i.e., $I_{\overline{j}}$ has the best optimistic greedy value when it is chosen.

As the iteration does not solve the instance, we have $(w_{\bar{j}} - L_{\bar{j}}) \ge \frac{1}{2}(U_{\bar{j}} - L_{\bar{j}})$ by Line 6. This directly implies that the real greedy value of $I_{\bar{j}}$ is at least half the optimistic greedy value, i.e., $g_c(Q_j \cup \bar{G}_j, I_{\bar{j}}) \ge \frac{1}{2} \cdot \bar{g}_c(Q_j \cup \bar{G}_j, I_{\bar{j}})$.

Since the best optimistic greedy value is never smaller than the best real greedy value, we get $g_c(Q_j \cup \bar{G}_j, I_{\bar{j}}) \geq \frac{1}{2} \cdot \max_{I_i \in \mathcal{I} \setminus (Q_j \cup \bar{G}_j)} g_c(Q_j \cup \bar{G}_j, I_i)$. This allows us to apply Lemma 3.4 to get $b'(Q_j \cup \bar{G}_j \cup I_{\bar{j}}) \leq (1 - \frac{1}{2\text{OPT}}) \cdot b'(Q_j \cup \bar{G}_j)$. Using $b'(Q_j \cup \bar{G}_j) \leq b'(Q_j)$ and $G_j = \bar{G}_j \cup \{I_{\bar{j}}\}$, we can conclude $b'(Q_j \cup G_j) \leq (1 - \frac{1}{2\text{OPT}}) \cdot b'(Q_j)$, which shows that G_j satisfies Condition 2 of Definition 3.3.

Case (2): Next, we show that G_j 1-approximates the greedy choice if $b'(Q_j) < 1$. In this case, we have $g = \bar{g}_s$ (cf. Line 3). Similar to the previous case, we have $\bar{g}_s(Q_j \cup \bar{G}_j, I_{\bar{j}}) = \max_{I_i \in \mathcal{I} \setminus (Q_j \cup \bar{G}_j)} \bar{g}_s(Q_j \cup \bar{G}_j, I_{\bar{j}})$.

From $b'(\dot{Q}_j) = \sum_{S \in S} b'_S(Q_j) < 1$ follows $b'_S(Q_j) < 1$ for all $S \in S$. Furthermore, every element I_i with $w_i - L_i \ge \frac{1}{2}U_i - L_i$ satisfies $a_i = w_i - L_i \ge \frac{s_{\min}}{2}$ and, therefore $a'_i = \gamma a_i = \frac{2}{s_{\min}} \cdot a_i \ge 1$. This means that adding I_i to the solution satisfies *all* constraints for sets S with $I_i \in S$ that are previously not satisfied. Thus, the optimistic greedy value $\bar{g}_s(Q_j, I_i)$ and the real greedy value $g_s(Q_j, I_i)$ are the same, i.e., $\bar{g}_s(Q_j, I_i) = g_s(Q_j, I_i)$, as adding I_i cannot satisfy more constraints even if the coefficient a_i was $U_i - L_i$. This observation is crucial for the remaining proof and the reason we scale with $\gamma = \frac{2}{s_{\min}}$ instead of $\frac{1}{s_{\min}}$.

Since the iteration does not solve the instance by assumption, we have $(w_{\bar{j}} - L_{\bar{j}}) \geq \frac{1}{2}(U_{\bar{j}} - L_{\bar{j}})$ by Line 6. As argued above, this implies that $I_{\bar{j}}$ has the best optimistic *and* actual greedy value when it is added to the solution, i.e., $g_s(Q_j \cup \bar{G}_j, I_{\bar{j}}) = \bar{g}_s(Q_j \cup \bar{G}_j, I_{\bar{j}}) = \max_{I_i \in \mathcal{I} \setminus (Q_j \cup \bar{G}_j)} \bar{g}_s(Q_j \cup \bar{G}_j, I_i)$. Thus, even under the assumption that all elements I_i of $\mathcal{I} \setminus (Q_j \cup \bar{G}_j)$ have coefficients $a_i = (U_i - L_i)$, interval $I_{\bar{j}}$ achieves the best greedy value.

Let LB denote the optimal solution value for the remaining instance after querying $Q_j \cup \overline{G}_j$ under exactly this assumption that $a_i = (U_i - L_i)$ and $a'_i = \gamma(U_i - L_i)$ for all $I_i \in \mathcal{I} \setminus (Q_j \cup \overline{G}_j)$. Clearly $LB \leq \text{OPT}$.

Under the assumption that $a'_i = \gamma(U_i - L_i)$ for all $I_i \in \mathcal{I} \setminus (Q_j \cup \bar{G}_j)$, the instance is scaled (i.e., it satisfies that all non-zero coefficients are at least one). Thus, we can apply Lemma 3.4 under the assumption to get $A(Q_j \cup \bar{G}_j \cup \{I_{\bar{j}}\}) \leq (1 - \frac{1}{LB}) \cdot A(Q_j \cup \bar{G}_j)$. Since $LB \leq \text{OPT}$, this gives us $A(Q_j \cup \bar{G}_j \cup \{I_{\bar{j}}\}) \leq (1 - \frac{1}{DPT}) \cdot A(Q_j \cup \bar{G}_j)$. Using $A(Q_j \cup \bar{G}_j) \leq A(Q_j)$ and $G_j = \bar{G}_j \cup \{I_{\bar{j}}\}$, we can conclude $A(Q_j \cup G_j) \leq (1 - \frac{1}{OPT}) \cdot A(Q_j)$, which shows that G_j satisfies Condition 2 of Definition 3.3 for $\alpha = 1$.

4.2 Multiset Multicover under Stochastic Explorable Uncertainty

We extend the result of the previous section to a stochastic variant of the *multiset multicover problem* (see, e.g., [35]), which generalizes the classical set cover problem. We are given a universe \mathcal{U} of n elements and a family \mathcal{M} of m multi-sets with $M \subseteq \mathcal{U}$ for all $M \in \mathcal{M}$. For each $e \in \mathcal{U}$, we are given a deterministic covering requirement $b_e \in \mathbb{R}_+$. Each multi-set $M \in \mathcal{M}$ contains a number of (fractional) copies of element $e \in \mathcal{U}$, denoted by $a_{M,e} \in \mathbb{R}_+$. The goal is to select a subset $C \subseteq \mathcal{M}$ of minimum cardinality that satisfies

all covering requirements:

$$\min_{\mathbf{x}_{M} \in \mathcal{M}} \sum_{M \in \mathcal{M}} x_{M} x_{M} \cdot a_{M,e} \geq b_{e} \quad \forall e \in \mathcal{U}$$

$$x_{M} \in \{0,1\} \quad \forall M \in \mathcal{M}.$$
(MINCOVERIP)

In our stochastic variant, MINCOVER under uncertainty, the constraints are uncertain; more precisely, the coefficients $a_{M,e}$ are initially unknown. We are given uncertainty intervals $I_{M,e} = (L_{M,e}, U_{M,e})$ with $a_{M,e} \in I_{M,e}$. We may query a set $M \in \mathcal{M}$, which reveals the precise values $a_{M,e}$ for set M and all $e \in \mathcal{U}$. In contrast to MINSET, we study MINCOVER in a *query-commit model*, which means that we can only add a set $M \in \mathcal{M}$ to the solution if we query it, and we *have* to add all queried sets to the solution. In a sense, we solve (MINCOVERIP) with uncertainty in the coefficients $a_{M,e}$ and irrevocable decisions: Once we add a set M to the solution ($x_M = 1$), all $a_{M,e}$ are revealed and we can never remove M from the solution. Irrevocably adding a set corresponds to querying it. The goal is to find a feasible solution to the ILP with a minimal number of queries.

From the problem description alone, the problem of solving MINCOVER under uncertainty is quite similar to solving (MINSETIP) under uncertainty. The difference is that in (MINSETIP), each interval I_i has the same coefficient a_i in all constraints where the variable occurs, i.e., where the variable has a non-zero coefficient. This value is drawn from a single distribution over I_i . For MINCOVER under uncertainty on the other hand, each coefficient $a_{M,e}$ of a multiset M can be different and has its own uncertainty interval $I_{M,e}$. In contrast to MINSET under uncertainty, we still have deterministic right-hand sides.

4.2.1 Offline MINCOVER.

All our results for offline MINSET directly translate to offline MINCOVER, where we assume full knowledge of all coefficients $a_{M,e}$. To see this, we again define an ILP parametrized by the set $Q \subseteq \mathcal{M}$ of multisets that have already been added to the solution and adjust the right-hand sides to the remaining covering requirement after adding Q. Here, $b_e(Q) = \max\{b_e - \sum_{M \in Q} a_{M,e}, 0\}$ and $b(Q) = \sum_{e \in \mathcal{U}} b_e(Q)$.

$$\begin{array}{ll} \min & \sum_{M \in \mathcal{M} \setminus Q} x_M \\ \text{s.t.} & \sum_{M \in \mathcal{M} \setminus Q} x_M \cdot a_{M,e} & \geq b_e(Q) \quad \forall e \in \mathcal{U} \\ & x_M & \in \{0,1\} \quad \forall M \in \mathcal{M} \setminus Q \end{array}$$
 (MINCOVERIP-Q)

Based on this ILP, we define an offline greedy algorithm for MINCOVER in the same way as for offline MINSET, only with slightly different greedy value definitions. The algorithm scales the coefficients such that all non-zero left-hand side coefficients are at least 1, and then greedily adds the set to the solution that reduces the right-hand sides the most, i.e., the set $M \in \mathcal{M}$ that maximizes $g_c(Q, M) = b'(Q) - b'(Q \cup \{M\})$ (a' and b' indicate scaled coefficients). For a set $G \subseteq \mathcal{M}$, we define $g_c(Q, G) = b'(Q) - b'(Q \cup \{M\})$

After $b'_e(Q) < 1$ for all $e \in \mathcal{U}$, we can exploit that all scaled non-zero coefficients $a'_{M,e}$ are at least one. This means that adding a set M satisfies all remaining constraints of elements e with $a'_{M,e} > 0$. Thus, the remaining problem reduces to a SETCOVER instance, which can be solved by using the greedy algorithm by Chvatal [8], which greedily adds the set M that maximizes $g_s(Q, M) = A(Q) - A(Q \cup \{M\})$ with $A(Q) = |\{e \in \mathcal{U} \mid b_e(Q) > 0\}|$, i.e., the set that satisfies the largest number of constraints that are not already satisfied by Q. For a set $G \subseteq \mathcal{M}$, we define $g_s(Q, G) = A(Q) - A(Q \cup G)$. We call $g_c(Q, M)$ and $g_s(Q, M)$ the greedy values of M after adding Q.

This offline algorithm achieves an approximation factor of $\rho'(\gamma) = \lceil \ln(\gamma \cdot n \cdot \max_{e \in \mathcal{U}} b_e) \rceil + \lceil \ln(n) \rceil$ with $s_{\min} = \min_{e \in \mathcal{U}, M \in \mathcal{M}: a_{M,e} > 0} a_{M,e}$, $\gamma = 1/s_{\min}$ and $n = |\mathcal{U}|$ (follows from [11]). It is easy to see that Definition 3.3, Lemma 3.4, Lemma 3.5, Definition 3.6 and Theorem 3.7 all transfer to MINCOVER using ρ' instead of ρ (or $\bar{\rho}$) and using the adjusted greedy value definitions. This is because none of the Algorithm 4: Algorithm MINCOVER under stochastic explorable uncertainty.

Input: Instance if MINCOVER under stochastic explorable uncertainty.

1 $Q \leftarrow \emptyset$; Scale coefficients by $\gamma = (\frac{2}{s_{\min}})$ to a' and b' with $s_{\min} = \min_{e \in \mathcal{U}, M \in \mathcal{M}: U_{M,e} > 0} U_{M,e}$;

2 while the problem is not solved do

- 3 if $b'(Q) \ge 1$ then $g = \bar{g}_c$ else $g = \bar{g}_s$;
- 4 repeat
- 5 $M \leftarrow \arg \max_{M' \in \mathcal{M} \setminus Q} g(Q, M');$ Query $M; Q \leftarrow Q \cup \{M\};$
- 6 **until** the problem is solved or $g_c(Q, M) \ge \frac{1}{2}\bar{g}_c(Q, M)$ (if $g = \bar{g}_c$) or $g_s(Q, M) \ge \frac{1}{2}\bar{g}_s(Q, M)$ (if $g = \bar{g}_s$);

proofs actually uses the special case properties of (MINSETIP) that (MINCOVERIP) does not have. In the following, we will use these definitions, lemmas, and theorems to prove algorithmic results for stochastic MINCOVER.

4.2.2 MINCOVER under uncertainty.

Similar to MINSET with deterministic right-hand sides, it can happen that the instance has no feasible solution. In that case, we require every algorithm (including OPT) to reduce the covering requirements as much as possible.

Our algorithm again relies on the optimistic greedy values $\bar{g}_c(Q, M) = b'(Q) - \sum_{e \in \mathcal{U}} \max\{0, b'_e(Q) - \gamma U_{M,e}\}$ (if $b'(Q) \ge 1$) and $\bar{g}_s(Q, M) = A(Q) - |\{e \in \mathcal{U} \mid b'_e(Q) - \gamma U_{M,e} > 0\}|$ (otherwise). For sets $G \subseteq \mathcal{M}$, we define $\bar{g}_s(Q, G)$ and $\bar{g}_c(Q, G)$ analogously. Furthermore, the algorithm again scales the left-hand side coefficients by $\gamma = 2/s_{\min}$ for $s_{\min} = \min_{e \in \mathcal{U}, M \in \mathcal{M} : U_{M,e} > 0} U_{M,e}$.

In contrast to MINSET under uncertainty, we need a different balancing parameter to take the more general nature of the coefficients into account. For MINSET under uncertainty, the special structure of the coefficients for an interval I_i allowed us to relate the event $w_i \ge \frac{U_i+L_i}{2}$ to the actual and optimistic greedy value of I_i after some set Q has already been queried. For MINCOVER the greedy values of a multiset M depend on the outcomes of several random variables, one for each $e \in E$ with $U_{M,e} > 0$. In particular, the greedy value $g_s(Q, M)$ now depends on the number of random variables $a_{M,e}$ with a sufficiently large value that satisfies the corresponding constraint. Similar, the greedy value $g_c(Q, M)$ now depends on a sum of random variables. Even if $a_{M,e} \ge \frac{U_{M,e}+L_{M,e}}{2}$ holds with probability τ for each M and e, this property does not necessarily translate to the greedy values.

For these reasons, we define the balancing parameter τ' based on the greedy values for a scaled instance. We characterize the distributions for an $M \in \mathcal{M}$ by $\tau'_M = \min\{\tau'_{M,1}, \tau'_{M,2}\}$, where $\tau'_{M,1}$ is the minimum probability $\Pr[g_c(M,Q) \ge \frac{1}{2}\bar{g}_c(M,Q)]$ over all $Q \subseteq \mathcal{M} \setminus \{M\}$ and all realizations of the elements of Q. Similar, $\tau'_{M,2}$ is the minimum probability $\Pr[g_s(M,Q) \ge \frac{1}{2}\bar{g}_s(M,Q)]$ over all $Q \subseteq \mathcal{M} \setminus \{M\}$ and all realizations of the elements of Q similar, $\tau'_{M,2}$ is the minimum probability $\Pr[g_s(M,Q) \ge \frac{1}{2}\bar{g}_s(M,Q)]$ over all $Q \subseteq \mathcal{M} \setminus \{M\}$ and all realizations of the elements of Q with $b'_e(Q) \le 1$ for all $e \in \mathcal{U}$. For the complete instance, we use parameter $\tau' = \min_{M \in \mathcal{M}} \tau'_M$.

In contrast to the parameter for MINSET, the parameter τ' is defined in a more artificial and restrictive way. Nevertheless, we argue that small values of $\frac{1}{\tau'}$ still capture interesting distributions. For example, if the distribution of each $a_{M,e}$ is symmetric over the center of $I_{M,e}$, then $\tau' = 0.5$. Our main result is the following theorem.

Theorem 4.2. There exists an algorithm for MINCOVER under uncertainty with a competitive ratio of $\frac{2}{\tau'} \cdot \rho'(\gamma) \in \mathcal{O}(\frac{1}{\tau'} \cdot \log n)$ with $\gamma = 2/s_{\min}$ for $s_{\min} = \min_{e \in \mathcal{U}, M \in \mathcal{M}: U_{M,e} > 0} U_{M,e}$.

Thus, the algorithm of the theorem loses only a factor $\frac{2}{\tau}$ compared to the greedy approximation fac-

tor $\rho'(\gamma)$ on the corresponding offline problem. We show the theorem by proving that Algorithm 4 is an (α, β, γ) -GREEDY algorithm for $\alpha = 2$, $\beta = \frac{1}{\tau'}$ and $\gamma = \frac{2}{s_{\min}}$. Then, Theorem 3.7 implies the theorem. The algorithm scales the coefficients by factor γ ; we use a' and b' to refer to the scaled coefficients. Algorithm 4 uses the exact same idea as Algorithm 3 for MINSET with deterministic right-hand sides. That is, it executes the greedy algorithm by Dobson [11] under the assumption that $a_{M,e} = U_{M,e}$ (and $a'_{M,e} = \gamma U_{M,e}$) for all $e \in \mathcal{U}$ and all not yet queried $M \in \mathcal{M}$.

In contrast to g_s and g_c , Algorithm 4 has sufficient information to compute the optimistic greedy values \bar{g}_s and \bar{g}_c and, therefore, can compute the best greedy choice based on the optimistic greedy values. The algorithm is designed to find, in each iteration, one set M with $g_c(Q, M) \ge \frac{1}{2} \cdot \bar{g}_c(Q, M)$ for the current Q (or analogously for \bar{g}_s and g_s). We show that this ensures that each iteration 2-approximates the greedy choice for the scaled instance.

Proof of Theorem 4.2. Let *i* be an arbitrary iteration of the outer while-loop, X_i denote the number of queries during the iteration, and Y_i indicate whether the algorithm executes iteration i ($Y_i = 1$) or not. Assuming $Y_i = 1$, the algorithm, during *i*, executes queries to sets *M* until either $g_s(Q, M) \ge \frac{1}{2}\bar{g}_s(Q, M)$ (if $g = \bar{g}_s$) or $g_c(Q, M) \ge \frac{1}{2}\bar{g}_c(Q, M)$ (if $g = \bar{g}_c$) or the problem is solved. Thus, the iteration terminates at the latest when it finds an *M* with $g_s(Q, M) \ge \frac{1}{2}\bar{g}_s(Q, M)$ or $g_c(Q, M) \ge \frac{1}{2}\bar{g}_c(Q, M)$ depending on *g*. By definition of τ' , the number of attempts until this happens follows a geometric distribution with success probability at least τ' . Hence, $\mathbb{E}[X_i | Y_i = 1] \le \frac{1}{\tau'}$; proving Property 2 of Definition 3.6.

We continue by proving the first property of (α, β, γ) -GREEDY. Consider a fixed realization. Let \overline{G}_i denote the queries of iteration *i* except the last one and let \overline{M} denote the last query of iteration *i*. Then $G_i = \overline{G}_i \cup {\overline{M}}$ is the set of queries during the iteration. Finally, let Q_i denote the set of queries before iteration *i*. We show that G_i 2-approximates the greedy choice of the scaled instance, which implies Property 1 of Definition 3.6. If the iteration solves the problem, then G_i clearly 1-approximates the greedy choice and we are done. Thus, assume otherwise.

We first show that G_i 2-approximates the greedy choice if $b'(Q_i) \ge 1$. If $b'(Q_i) \ge 1$, we have $g = \bar{g}_c$ (cf. Line 3). By choice of \bar{M} , we have $\bar{g}_c(Q_i \cup \bar{G}_i, \bar{M}) = \max_{M \in \mathcal{M} \setminus (Q_i \cup \bar{G}_i)} \bar{g}_c(Q_i \cup \bar{G}_i, M)$. Then, $g_c(Q_i \cup \bar{G}_i, \bar{M}) \ge \frac{1}{2} \cdot \bar{g}_c(Q_i \cup \bar{G}_i, \bar{M})$ holds by definition of the algorithm as $g = \bar{g}_c$ and iteration i does not solve the instance. This implies $g_c(Q_i \cup \bar{G}_i, \bar{M}) \ge \frac{1}{2} \cdot \max_{M \in \mathcal{M} \setminus (Q_i \cup \bar{G}_i)} g_c(Q_i \cup \bar{G}_i, M)$. By Lemma 3.4, $b'(Q_i \cup \bar{G}_i \cup \{\bar{M}\}) \le (1 - \frac{1}{2\text{OPT}}) \cdot b'(Q_i \cup \bar{G}_i)$ and, thus, $b'(Q_i \cup G_i) \le (1 - \frac{1}{2\text{OPT}}) \cdot b'(Q_i)$, which implies the property.

Next, we show that G_i 2-approximates the greedy choice if $b'(Q_i) < 1$. By assumption, we have $b'_e(Q_i) < 1$ for all $e \in \mathcal{U}$. Recall that we scale by factor $\gamma = \frac{2}{s_{\min}}$ with $s_{\min} = \min_{e \in \mathcal{U}, M \in \mathcal{M}: U_{M,e} > 0} U_{M,e}$. If a set M satisfies $a_{M,e} \ge \frac{1}{2} \cdot U_{M,e}$ for some $e \in \mathcal{U}$ with $U_{M,e} > 0$, then $a'_{M,e} = \gamma a_{M,e} \ge \frac{2}{s_{\min}} \cdot \frac{1}{2} \cdot U_{M,e} \ge 1$. Thus, adding M to the solution satisfies the constraint of element e (if it was not already satisfied).

By choice of \overline{M} , we have $\overline{g}_s(Q_i \cup \overline{G}_i, \overline{M}) = \max_{M \in \mathcal{M} \setminus (Q \cup \overline{G}_i)} \overline{g}_s(Q_i \cup \overline{G}_i, M)$. Let *LB* denote the optimal solution value for the remaining instance under the assumption that $a_{M,e} = U_{M,e}$ (and $a'_{M,e} = \gamma U_{M,e}$) for all $e \in \mathcal{U}$ and all not yet queried $M \in \mathcal{M} \setminus (Q_i \cup \overline{G}_i)$. Clearly $LB \leq \text{OPT}$.

Under this assumption, the remaining instance is again just a SETCOVER instance. Similar to the proof of Lemma 3.4, we can argue that $\frac{A(Q_i \cup \bar{G}_i)}{LB} \leq \bar{g}_s(Q_i \cup \bar{G}_i, \bar{M})$ as the optimal solution satisfies the remaining constraints at cost LB, but a single multiset can satisfy at most $\bar{g}_s(Q_i \cup \bar{G}_i, \bar{M})$ constraints.

Observe that $g_s(Q_i \cup \bar{G}_i, \bar{M}) \geq \frac{1}{2}\bar{g}_s(Q_i \cup \bar{G}_i, \bar{M})$ holds by definition of the algorithm as \bar{M} is the last added set in the iteration with $g = \bar{g}_s$ and iteration i does not solve the instance. By definition of g_s , this implies $A(Q_i \cup \bar{G}_i) - A(Q_i \cup \bar{G}_i \cup \bar{M}) = A(Q_i \cup \bar{G}_i) - A(Q_i \cup G_i) \geq \frac{1}{2} \cdot \bar{g}_s(Q_i \cup \bar{G}_i, \bar{M})$. In combination with $\frac{A(Q_i \cup \bar{G}_i)}{LB} \leq \bar{g}_s(Q_i \cup \bar{G}_i, \bar{M})$, we get $A(Q_i \cup \bar{G}_i) - A(Q_i \cup G_i) \geq \frac{A(Q_i \cup \bar{G}_i)}{2 \cdot LB}$. This implies $A(Q_i \cup G_i) \leq (1 - \frac{1}{2 \cdot LB}) \cdot A(Q_i \cup \bar{G}_i)$. As OPT $\geq LB$ and $A(Q_i) \geq A(Q_i \cup \bar{G}_i)$, we can conclude $A(Q_i \cup G_i) \leq (1 - \frac{1}{2 \cdot OPT}) \cdot A(Q_i)$, which implies that G_i 2-approximates the greedy choice and, thus,

satisfies the first property of (α, β, γ) -GREEDY.

5 MINSET under uncertainty

We consider the general MINSET under uncertainty. In contrast to the previous section, we now also have uncertainty in the right-hand sides of (MINSETIP). Since we consider the stochastic problem variant, recall that the balancing parameter is $\tau = \min_{I_i \in \mathcal{I}} \tau_i$ with $\tau_i = \mathbb{P}[w_i \geq \frac{U_i + L_i}{2}]$. Our goal is to iteratively add intervals from \mathcal{I} to the solution until it becomes feasible for (MINSETIP). To that end, we prove the following main result.

Theorem 5.1. For $\tau > 0$. There is an algorithm for MINSET under uncertainty with a competitive ratio of $\mathcal{O}(\frac{1}{\tau} \cdot \log m \cdot \bar{\rho}(\gamma)) \subseteq \mathcal{O}(\frac{1}{\tau} \cdot \log^2 m)$ with $\gamma = 2/s_{\min}$ for $s_{\min} = \min_{I_i \in \mathcal{I}: (U_i - L_i) > 0} (U_i - L_i)$.

Exploiting Theorem 3.7, we prove the statement by providing Algorithm 5 and showing that it is an (α, β, γ) -GREEDY algorithm for $\alpha = 2$, $\gamma = 2/s_{\min}$ and $\beta = \frac{1}{\tau}(\lceil \log_{1.5}(m \cdot (2/s_{\min}) \cdot \max_{I_i \in \mathcal{I}}(U_i - L_i)) \rceil + \lceil \log_2(m) \rceil)$. Note that α and γ are defined as in the previous section for MINSET with deterministic right-hand sides and will be used analogously. For β on the other hand, we require a larger value to adjust for the additional uncertainty in the right-hand sides $b_S = w^* - L_S$ for the uncertain w^* . Notice that we do not have sufficient information to just execute Algorithm 3 for MINSET with deterministic right-hand sides as we need the right-hand side values to compute even the optimistic greedy values.

To handle this additional uncertainty, we want to ensure that each iteration of our algorithm α -approximates the greedy choice for *each* possible value of w^* . To do so, we compute and query the best optimistic greedy choice for several carefully selected possible values of w^* .

To state our algorithm, we define a parametrized variant of (MINSETIP) that states the problem under the assumptions that $w^* = w$ for some w and that the set $Q \subseteq \mathcal{I}$ has already been queried. The coefficients are scaled to $a'_i = (2/s_{\min})(w_i - L_i)$ and $b'_S(Q, w) = \max\{(2/s_{\min})(w - L_S) - \sum_{I_i \in Q \cap S} a'_i, 0\}$. As before, let $b'(Q, w) = \sum_{S \in S} b'_S(Q, w)$ denote the sum of right-hand sides.

$$\begin{array}{ll} \min & \sum_{I_i \in \mathcal{I} \setminus Q} x_i \\ \text{s.t.} & \sum_{I_i \in S \setminus Q} x_i \cdot a'_i & \geq b'_S(Q, w) \quad \forall S \in \mathcal{S} \\ & x_i & \in \{0, 1\} \quad \forall I_i \in \mathcal{I} \end{array}$$
 (MINSETIP-QW)

As the right-hand sides are unknown, we define the greedy values for every possible value w for w^* . To that end, let $g_c(Q, I_i, w) = b'(Q, w) - b'(Q \cup \{I_i\}, w)$ and $g_s(Q, I_i, w) = A(Q, w) - A(Q \cup \{I_i\}, w)$, where $A(Q, w) = |\{S \in S \mid b'_S(Q, w) > 0\}|$ denotes the number of constraints in (MINSETIP-QW) that are not yet satisfied. As before, $g_c(Q, I_i, w)$ and $g_s(Q, I_i, w)$ describe how much adding I_i to the solution reduces the sum of right-hand sides and the number of non-satisfied constraints, respectively; now under the assumption that $w^* = w$. For subsets $G \subseteq \mathcal{I} \setminus Q$, we define the greedy values in the same way, i.e., $g_s(Q, G, w) = A(Q, w) - A(Q \cup G, w)$ and $g_c(Q, G, w) = b'(Q, w) - b'(Q \cup G, w)$.

Since our algorithm again does not have sufficient information to compute the precise greedy values $g_s(Q, I_i, w)$ and $g_c(Q, I_i, w)$ even for a fixed w, we again use the optimistic greedy values defined in the same way as in the previous section. That is

$$\bar{g}_c(Q, I_i, w) = \sum_{S \in \mathcal{S}: \ I_i \in S} b'_S(Q, w) - \max\{0, b'_S(Q, w) - \gamma(U_i - L_i)\}$$

and

$$\bar{g}_s(Q, I_i, w) = |\{S \in \mathcal{S} \colon I_i \in S \mid b'_S(Q, w) > 0 \land b'_S(Q, w) - \gamma(U_i - L_i) \le 0\}|.$$

For subsets $G \subseteq \mathcal{I} \setminus Q$, the optimistic greedy values are defined analogously.

Algorithm 5: Algorithm for MINSET under uncertainty.

Input: Instance of MINSET under uncertainty. 1 Scale all coefficients with $\gamma = 2/s_{\min}$ for $s_{\min} = \min_{I_i \in \mathcal{I}: (U_i - L_i) > 0} (U_i - L_i);$ 2 $Q \leftarrow \emptyset, w_{\min} \leftarrow \text{minimum possible value } w^*$ (keep up-to-date); 3 while the problem is not solved do **foreach** g from the ordered list \bar{g}_c, \bar{g}_s **do** 4 $d \leftarrow 1; Q' \leftarrow Q;$ 5 if $g = \bar{g}_c$ then $w_{\max} \leftarrow \max$ possible value w^* ; 6 else $w_{\max} \leftarrow \max w$ s.t. $b'_S(Q, w) < 1$ for all $S \in S$; 7 while $\exists w_{\min} \leq w \leq w_{\max}$ such that $\max_{I_h \in \mathcal{I} \setminus Q} g(Q, I_h, w) \geq d$ do 8 9 repeat $w \leftarrow \min w_{\min} \le w \le w_{\max} \text{ s.t. } \max_{I_h \in \mathcal{I} \setminus Q} g(Q, w, I_h) \ge d;$ 10 $\begin{bmatrix} I_i \leftarrow \arg \max_{I_h \in I \setminus Q} g(Q, I_h, w); \operatorname{Query} I_i; Q \leftarrow Q \cup \{I_i\}; \\ Q_{1/2} \leftarrow \{I_j \in Q \setminus Q' \mid w_j - L_j \geq \frac{U_j - L_j}{2}\}; \\ \text{if } g = \bar{g}_c \text{ then } d \leftarrow g_c(Q', Q_{1/2}, w) \text{ else } d \leftarrow g_s(Q', Q_{1/2}, w); \end{bmatrix}$ 11 12 13 until $w_i - L_i \ge \frac{U_i - L_i}{2}$ or $\nexists w \le w_{\max}$: $\max_{I_h \in \mathcal{I} \setminus Q} g(Q, w, I_h) \ge d$; 14

Similar to Algorithm 3, we would like to repeatedly compute and query the best optimistic greedy choice until the queried I_i satisfies $w_i - L_i \ge \frac{U_i - L_i}{2}$ (cf. the repeat-statement). However, we cannot decide which greedy value, \bar{g}_c or \bar{g}_s , to use as deciding whether $b'_S(Q, w^*) < 1$ depends on the unknown w^* . Instead, we compute and query the best optimistic greedy choice for both greedy values (cf. the for-loop). Even then, the best greedy choice still depends on the unknown right-hand sides. Thus, we compute and query the best optimistic greedy choice for several carefully selected values w (cf. the inner while-loop) to make sure that the queries of the iteration approximate the greedy choice for every possible w^* . Additionally, we want to ensure that we use at most β queries in expectation within an iteration of the outer while-loop.

To illustrate the ideas of the algorithm, consider an iteration of the outer while-loop. In particular, consider the for-loop iteration with $g = \bar{g}_s$ within this iteration. Let Q' denote the set of queries that were executed before the start of the iteration. Since we only care about the greedy value g_c if there exists some $S \in S$ with $b'_S(Q') \ge 1$ (otherwise we use \bar{g}_s and g_s instead), we assume that this is the case. If not, we use a separate analysis for the for-loop iteration with $g = \bar{g}_s$.

Our goal for the iteration is to query a set of intervals \bar{Q} that 2-approximates the best greedy choice I^* after querying Q', i.e., it has a greedy value $g_c(Q', \bar{Q}, w^*) \geq \frac{1}{2}g_c(Q', I^*, w^*)$ and, thus, satisfies Lemma 3.4. To achieve this for the unknown w^* , the algorithm uses the parameter d, which is initialized with 1 (cf. Line 5), the minimum possible value for $\bar{g}_c(Q', I^*, w^*)$ under the assumption that there exists some $S \in S$ with $b'_S(Q') \geq 1$. In an iteration of the inner while-loop, the algorithm repeatedly picks the minimal value wsuch that the best current optimistic greedy choice has an optimistic greedy value of at least d (cf. Line 10). If no such value exists, then the loop terminates (cf. Lines 8, 14). Afterwards, it queries the corresponding best optimistic greedy choice I_i for the selected value w (cf. Line 11). Similar to the algorithms of the previous section, this is done repeatedly until $w_i - L_i \geq (U_i - L_i)/2$.

The key idea to achieve the 2-approximation with an expected number of queries that does not exceed β , is to always reset the value d to $g_c(Q', Q_{1/2}, w)$, where $Q_{1/2}$ is the subset of all intervals I_j that have already been queried in the current iteration of the outer while-loop and satisfy $w_j - L_j \ge (U_j - L_j)/2$ (cf. Lines 12, 13). This can be seen as an implicit doubling strategy to search for an unknown value. It leads to an exponential increase of d over the iterations of the inner while-loop, which will allow us to bound their number.

With the following lemma, we prove that this choice of d also ensures that the queries of the iteration indeed 2-approximate the best greedy choice for w^* if there exists a $S \in S$ with $b'_S(Q', w^*) \ge 1$. If there is no such set, we can use a similar proof w.r.t. greedy value g_s . For an iteration j of the outer while-loop, let G_j be the set of queries during the iteration and let $Q_j = \bigcup_{j' < j} G_j$ denote the queries before the iteration (cf. Q' in the algorithm).

Lemma 5.2. If there is an $S \in S$ with $b'_S(Q_j, w^*) \ge 1$, then G_j 2-approximates the greedy choice for the scaled instance with $w = w^*$ after querying Q_j .

Proof. For an arbitrary but fixed realization, consider an iteration j of the outer while-loop such that there exists a set $S \in S$ with $b'_S(Q_j, w^*) \ge 1$.

Consider the subset $\bar{G}_j \subseteq G_j$ of queries that were executed with $g = \bar{g}_c$ before the increasing value w (cf. Line 10) surpasses w^* . That is, \bar{G}_j only contains intervals that were queried for a current value $w \leq w^*$. Let \bar{I}_i be the element of \bar{G}_j that is queried last. Finally, let \bar{d}_j denote the value d computed by the algorithm in Line 13 directly after querying \bar{I}_i . We continue to show that G_j 2-approximates the greedy choice of (MINSETIP-QW) for $Q = Q_j$ and $w = w^*$.

Observe that $g_c(Q_j, \bar{G}_j, w^*) \geq \bar{d}_j$. To see this, recall that \bar{d}_j was computed in Line 13 after \bar{I}_i was queried. Thus, $\bar{d}_j = g_c(Q', Q_{1/2}, w)$ for $Q' = Q_j$, $Q_{1/2} = \{I_j \in \bar{G}_j \mid w_j - L_j \geq \frac{U_j - L_j}{2}\}$ and some value w with $w \leq w^*$ by assumption. Since $w^* \geq w$ and $Q_{1/2} \subseteq \bar{G}_j$, the greedy value $g_c(Q_j, \bar{G}_j, w^*)$ can never be smaller than $\bar{d}_j = g_c(Q', Q_{1/2}, w)$. This implies $g_c(Q_j, \bar{G}_j, w^*) \geq \bar{d}_j$.

We continue by showing that $d^* \leq 2 \cdot g_c(Q_j, \bar{G}_j, w^*)$ holds for the best greedy value d^* at the start of the iteration, i.e., $d^* = \max_{I_i \in \mathcal{I} \setminus Q_j} g_c(Q_j, I_i, w^*)$. As $\bar{G}_j \subseteq G_j$, this implies $d^* \leq 2 \cdot g_c(Q_j, G_j, w^*)$ and, thus, that G_j satisfies Definition 3.3.

To upper bound d^* , first observe that the best optimistic greedy value d' after querying $G_j \cup Q_j$ is smaller than \bar{d}_j , i.e., $d' = \max_{I_i \in \mathcal{I} \setminus (Q_j \cup \bar{G}_j)} \bar{g}_c(Q_j \cup \bar{G}_j, I_i, w^*) < \bar{d}_j$. This follows directly from Line 10 as \bar{I}_i is the last query for a value $w \leq w^*$ by assumption. As $g_c(Q_j, \bar{G}_j, w^*) \geq \bar{d}_j$, we get $g_c(Q_j, \bar{G}_j, w^*) \geq d'$.

By definition of g_c , the best greedy value after querying Q_j can never be larger than the sum of the greedy value of \bar{G}_j after querying Q_j and the best optimistic greedy value after querying $\bar{G}_j \cup Q_j$. Thus, we have $d^* \leq g_c(Q_j, \bar{G}_j, w^*) + d' \leq 2 \cdot g_c(Q_j, \bar{G}_j, w^*)$. This proves that \bar{G}_j satisfies Lemma 3.4 and, thus, concludes this proof.

Using a similar proof, we show the following lemma for the case where $b'_S(Q', w^*) < 1$ for all $S \in S$, which together with Lemma 5.2 implies Property 1 of Definition 3.6. While the main arguments remain the same as for the previous lemma, the more discrete nature of the greedy values g_s and \bar{g}_s poses several additional technical challenges that need to be taken care of. For an iteration j of the outer while-loop, let G_j again be the set of queries during the iteration and let $Q_j = \bigcup_{j' < j} G_j$ denote the queries before the iteration (cf. set Q' in the algorithm).

Lemma 5.3. If $b'_S(Q_j, w^*) < 1$ for all $S \in S$, then G_j 2-approximates the greedy choice for the scaled instance with $w = w^*$ after querying Q_j .

Proof. For an arbitrary but fixed realization, consider an iteration j of the outer while-loop such that $b'_S(Q_j, w^*) < 1$ for all $S \in S$. Our goal is to prove that G_j approximates the greedy choice for the scaled instance with $w = w^*$ after querying Q_j within a factor of two. That is, we have to prove $A(Q_j \cup G_j, w^*) \le (1 - \frac{1}{2 \cdot \text{OPT}}) \cdot A(Q_j, w^*)$. Recall that $A(Q_j, w^*)$ denotes the number of constraints that are not yet satisfied in the (MINSETIP-QW) instance for $Q = Q_j$ and $w = w^*$.

Consider the subset $G_j \subseteq G_j$ of queries to intervals I_i that were executed with $g = \bar{g}_s$ for a current value $w \leq w^*$ during iteration j of the outer while-loop. Let $P_j \subseteq G_j$ denote the queries of the iteration that were executed before \bar{G}_j , i.e., that were executed during the iteration of the for-loop with $g = \bar{g}_c$. Note

that $Q' = Q_j \cup P_j$ is the set of intervals queried before the beginning of the for-loop iteration with $g = \bar{g}_s$ during iteration j of the outer while-loop.

Proof outline. We start the proof by making some preliminary observations regarding greedy value g_s and the scaling factor γ that will be crucial for the remainder of the proof. Then we proceed by proving that \bar{G}_j approximates the greedy choice of (MINSETIP-QW) for $Q = Q_j \cup P_j$ and $w = w^*$ within a factor of 2. To that end, we first derive a lower bound on the greedy value $g_s(Q_j \cup P_j, \bar{G}_j, w^*)$ and afterwards compare this lower bound with OPT. Finally, we use the fact that \bar{G}_j 2-approximates the greedy choice after querying $Q_j \cup P_j$ to show that G_j approximates the greedy choice of (MINSETIP-QW) for $Q = Q_j$ and $w = w^*$ within a factor of 2.

Preliminary observations. Before we start with the proof, recall that an interval I_i with $(w_i - L_i) \geq \frac{1}{2} \cdot (U_i - L_i)$ satisfies $a'_i = \frac{2 \cdot (w_i - L_i)}{s_{\min}} \geq \frac{U_i - L_i}{s_{\min}} \geq 1$ by choice of the scaling parameter $\gamma = \frac{2}{s_{\min}}$. This implies that adding I_i to the solution satisfies all constraints for sets S with $I_i \in S$ as long as we are considering values w with $b'_S(Q_j, w) \leq 1$ for all $S \in S$. Thus, for such intervals and values w, the greedy value g_s of I_i is then equal to the optimistic greedy value \bar{g}_s as even under the assumption $a'_i = \gamma(U_i - L_i)$ adding interval I_i cannot satisfy more constraints. By assumption, this in particular holds for all values $w \leq w^*$. This also means that the greedy values g_s and \bar{g}_s of such intervals I_i only increase with an increasing value w, as long as $b'_S(Q_j, w) \leq 1$ still holds for all $S \in S$.

Lower bound on $g_s(Q_j \cup P_j, G_j, w^*)$. We continue by deriving a lower bound on $g_s(Q_j \cup P_j, G_j, w^*)$. Let \bar{I}_i be the element of \bar{G}_j that is queried last and let \bar{d}_j denote the value d computed by the algorithm in Line 13 directly after querying \bar{I}_i . We first observe that $g_s(Q_j \cup P_j, \bar{G}_j, w^*) \ge \bar{d}_j$. To see this, recall that \bar{d}_j was computed in Line 13 after \bar{I}_i was queried. Thus, $\bar{d}_j = g_s(Q', Q_{1/2}, w)$ for $Q' = Q_j \cup P_j$, $Q_{1/2} = \{I_j \in \bar{G}_j \mid w_j - L_j \ge \frac{U_j - L_j}{2}\}$ and some value w with $w \le w^*$ by assumption. Since $w^* \ge w$ and $Q_{1/2} \subseteq \bar{G}_j$, the greedy value $g_s(Q_j \cup P_j, \bar{G}_j, w^*)$ can never be smaller than $\bar{d}_j = g_s(Q', Q_{1/2}, w)$ according to the observations stated at the beginning of the proof. This implies $g_s(Q_j \cup P_j, \bar{G}_j, w^*) \ge \bar{d}_j$.

Assume for now that the algorithm queried $Q_s = \bar{G}_j \setminus Q_{1/2}$ before $Q_{1/2}$. We consider the best optimistic greedy value d^* after $Q_j \cup P_j \cup Q_s$ has already been queried, i.e, $d^* = \max_{I_i \in \mathcal{I} \setminus (Q_j \cup P_j \cup Q_s)} \bar{g}_s(Q_j \cup P_j, I_i, w^*)$. To bound d^* , first observe that the best optimistic greedy value d' after querying $\bar{G}_j \cup Q_j \cup P_j$ is smaller than \bar{d}_j , i.e., $d' = \max_{I_i \in \mathcal{I} \setminus (Q_j \cup P_j \cup \bar{G}_j)} \bar{g}_s(Q_j \cup P_j \cup \bar{G}_j, I_i, w^*) < \bar{d}_j$. This follows directly from Line 10 as \bar{I}_i is the last query for a value $w \leq w^*$ by assumption. Since we already showed $g_s(Q_j \cup P_j, \bar{G}_j, w^*) \geq \bar{d}_j$, we get $g_s(Q_j \cup P_j, \bar{G}_j, w^*) \geq d'$.

By definition of \bar{g}_s , definition of $Q_{1/2}$, and the assumption that we only consider values w with $b'_S(Q', w) < 1$ for all $S \in S$, the best optimistic greedy value after querying $Q_j \cup P_j \cup Q_s$ can never be larger than the sum of the optimistic greedy value of $Q_{1/2}$ after querying $Q_j \cup P_j \cup Q_s$ and the best optimistic greedy value after querying $Q_j \cup P_j \cup Q_s$ and the best optimistic greedy value after querying $Q_j \cup P_j \cup Q_s$ and the best optimistic greedy value after querying $Q_j \cup P_j \cup Q_s \cup Q_{1/2} = Q_j \cup P_j \cup \bar{G}_j$. By assumption that we only consider values w with $b'_S(Q', w) < 1$ for all $S \in S$, we have $\bar{g}_s(Q_j \cup P_j \cup Q_s, Q_{1/2}, w^*) = g_s(Q_j \cup P_j \cup Q_s, Q_{1/2}, w^*)$ (as argued at the beginning of the proof). Putting it together, we get

$$\begin{aligned} d^* &\leq \bar{g}_s(Q_j \cup P_j \cup Q_s, Q_{1/2}, w^*) + d' \\ &= g_s(Q_j \cup P_j \cup Q_s, Q_{1/2}, w^*) + d' \\ &\leq g_s(Q_j \cup P_j \cup Q_s, Q_{1/2}, w^*) + g_s(Q_j \cup P_j, \bar{G}_j, w^*) \end{aligned}$$

Proving that \bar{G}_j **approximates its greedy choice.** We continue the proof by using the inequality for d^* in order to show that \bar{G}_j approximates its greedy choice.

To that end, we consider a relaxed instance R of (MINSETIP-QW) with $Q = Q_j \cup P_j$ and $w = w^*$, i.e., we assume that $Q_j \cup P_j$ has already been queried and consider the right-hand sides $b'_S(Q_j \cup P_j, w^*)$ for all $S \in S$. We relax the instance by increasing the left-hand side coefficients and decreasing the cost coefficients. Let \hat{a} denote the relaxed coefficients. First consider the intervals in \bar{G}_j . For these intervals, we use the original scaled coefficients. That is, we set $\hat{a}_i = a'_i = \gamma(w_i - L_i)$. Recall that $Q_{1/2} = \{I_i \in \bar{G}_j \mid w_i - L_i \geq \frac{1}{2}(U_i - L_i)\}$ and $Q_s = \bar{G}_j \setminus Q_{1/2}$. For intervals $I_i \in Q_s$, we set the cost coefficients to zero. Thus, these intervals can be added to any solution without increasing the objective value. All other intervals keep their cost coefficients of one. For all other intervals $I_i \in \mathcal{I} \setminus (Q_j \cup P_j \cup \bar{G}_j)$, we set $\hat{a}_i = \gamma(U_i - L_i)$. That is, we assume that the coefficients of these intervals are slightly larger than their largest possible value. Since instance R compared to the original instance only increases left-hand side coefficients and decreases cost coefficients, it clearly is a relaxation and, thus, $LB \leq OPT$ for the optimal objective value LB of the relaxed instance R.

Since we assume $b'_S(Q_j \cup P_j, w^*) < 1$ for all $S \in S$, the right-hand sides of instance R are all strictly smaller than one. Furthermore, by choice of the scaling parameter γ , all intervals I_i satisfy $\gamma(U_i - L_i) \ge 1$. Thus, the intervals I_i in $Q \setminus (Q_j \cup P_j \cup \overline{G}_j)$ have coefficients $\hat{a}_i \ge 1$. By the observations at the beginning of the proof, the same holds for the intervals in $Q_{1/2}$. Only the intervals in Q_s can potentially have coefficients smaller than the right-hand sides. This also means that the greedy value g_s of an interval $I_i \in Q \setminus (Q_j \cup P_j \cup Q_s)$ for instance R can never be increased by previously adding intervals of Q_s to the solution, since adding I_i already satisfies all constraints for sets S with $I_i \in S$. For intervals in Q_s , it might be the case that previously adding further elements of Q_s to the solution increases the greedy value. This is, because the coefficient of an $I_i \in Q_s$ might be too small to satisfy a constraint S with $I_i \in S$ but previously adding further intervals from Q_s might decrease the right-hand side of S enough such that adding I_i can satisfy the constraint and, thus, increase the greedy value of I_i .

Now, consider the number of constraints that are initially not satisfied in R. This number is exactly the same as in the non-relaxed original instance after querying $Q_j \cup P_j$ since R uses the exact same right-hand sides. Thus, the number of not yet satisfied constraints in R is $A(Q_j \cup P_j, w^*)$. Let $\hat{d} = \max_{I_i \in \mathcal{I} \setminus (Q_j \cup P_j \cup Q_s)} |\{S \in S \mid b'_S(Q_j \cup P_j, w^*) > 0 \land b'_S(Q_j \cup P_j, w^*) - \hat{a}_i - \sum_{I_h \in Q_s} \hat{a}_h \leq 0\}|$ denote the maximum number of constraints that can be satisfied by adding Q_s and one interval I_i in $Q_j \cup P_j \cup Q_s$ to the solution for R. In a sense, \hat{d} is the best possible greedy value g_s for instance R that can be achieve by adding such a set $Q_s \cup \{I_i\}$ to the solution. Remember that the elements of Q_s do not incur any costs, so \hat{d} constraints can be satisfied at cost one. Following the arguments in the proof of Lemma 3.4, this implies $\hat{d} \geq \frac{A(Q_j \cup P_j, w^*)}{LB}$ as the optimal solution satisfies $A(Q_j \cup P_j, w^*)$ constraints at cost LB but it is impossible to satisfy more than \hat{d} constraints with cost one. This last implication crucially uses the observation that the greedy values of intervals in $Q \setminus (Q_j \cup P_j \cup Q_s)$ cannot increase by previously adding other intervals to the solution. Because this is the case, even adding an interval later cannot satisfy more than \hat{d} constraints, which gives us $\hat{d} \geq \frac{A(Q_j \cup P_j, w^*)}{LB}$. We remark that this also is the reason why the offline Algorithm 3 starts with greedy value g_c and only switches to greedy value g_s once the instance reduced to a SETCOVER instance.

Taking a closer look at \hat{d} , we can observe that this greedy value is exactly the sum of the greedy values $g_s(Q_j \cup P_j, Q_s, w^*)$ and $d^* = \max_{I_i \in Q \setminus (Q_j \cup P_j \cup Q_s)} \bar{g}_s(Q_j \cup P_j \cup Q_s, I_i, w^*)$ for the non-relaxed instance. This is, because the original instance uses the same right-hand sides as R, the coefficients of intervals in Q_s are the same in both instances, and the optimistic greedy value $\bar{g}_s(Q_j \cup P_j \cup Q_s, I_i, w^*)$ like instance R already assumes that all intervals in $Q \setminus (Q_j \cup P_j \cup Q_s)$ have coefficients larger than all right-hand sides. Thus, these three arguments imply $\hat{d} = g_s(Q_j \cup P_j, Q_s, w^*) + d^*$. Combining this equality with

 $\hat{d} \ge \frac{A(Q_j \cup P_j, w^*)}{LB}$ and the previously derived upper bound on d^* gives us:

$$\begin{aligned} \frac{A(Q_j \cup P_j, w^*)}{LB} &\leq \hat{d} \\ &= g_s(Q_j \cup P_j, Q_s, w^*) + d^* \\ &\leq g_s(Q_j \cup P_j, Q_s, w^*) + g_s(Q_j \cup P_j \cup Q_s, Q_{1/2}, w^*) + g_s(Q_j \cup P_j, \bar{G}_j, w^*) \\ &= 2 \cdot g_s(Q_j \cup P_j, \bar{G}_j, w^*). \end{aligned}$$

Here the second equality uses that

$$g_{s}(Q_{j} \cup P_{j}, Q_{s}, w^{*}) + g_{s}(Q_{j} \cup P_{j} \cup Q_{s}, Q_{1/2}, w^{*})$$

$$= (A(Q_{j} \cup P_{j}, w^{*}) - A(Q_{j} \cup P_{j} \cup Q_{s}, w^{*}))$$

$$+ (A(Q_{j} \cup P_{j} \cup Q_{s}, w^{*}) - A(Q_{j} \cup P_{j} \cup Q_{s} \cup Q_{1/2}, w^{*}))$$

$$= A(Q_{j} \cup P_{j}, w^{*}) - A(Q_{j} \cup P_{j} \cup Q_{s} \cup Q_{1/2}, w^{*})$$

$$= A(Q_{j} \cup P_{j}, w^{*}) - A(Q_{j} \cup P_{j} \cup \bar{G}_{j}, w^{*})$$

$$= g_{s}(Q_{j} \cup P_{j}, \bar{G}_{j}, w^{*}).$$

Plugging in the definition of g_s and the inequality $LB \leq OPT$ yields

$$\frac{A(Q_j \cup P_j, w^*)}{LB} \le 2 \cdot (A(Q_j \cup P_j, w^*) - A(Q_j \cup P_j \cup \bar{G}_j, w^*))$$
$$\Rightarrow A(Q_j \cup P_j \cup \bar{G}_j, w^*) \le A(Q_j \cup P_j, w^*) \cdot (1 - \frac{1}{2 \cdot LB})$$
$$\Rightarrow A(Q_j \cup P_j \cup \bar{G}_j, w^*) \le A(Q_j \cup P_j, w^*) \cdot (1 - \frac{1}{2 \cdot OPT}).$$

Thus, \bar{G}_j approximates the greedy choice of (MINSETIP-QW) for $Q = Q_j \cup P_j$ and $w = w^*$ within a factor of 2.

Concluding the proof. Remember that our goal was to show that G_j approximates the greedy choice for $w = w^*$ after querying Q_j within a factor of 2. To that end, observe that $P_j \cup \bar{G}_j \subseteq G_j$ implies $A(Q_j \cup P_j \cup \bar{G}_j, w^*) \ge A(Q_j \cup G_j, w^*)$ and that $Q_j \subseteq Q_j \cup P_j$ implies $A(Q_j, w^*) \ge A(Q_j \cup P_j, w^*)$. Plugging these inequalities into the previously derived inequality for $A(Q_j \cup P_j \cup \bar{G}_j, w^*)$ yields $A(Q_j \cup G_j, w^*) \le A(Q_j, w^*) \le A(Q_j, w^*) \le A(Q_j, w^*) \le A(Q_j, w^*)$. This means that G_j also 2-approximates the greedy choice after querying Q_j for $w = w^*$, which concludes the proof.

Since the Lemmas 5.2 and 5.3 imply Condition 1 of Definition 3.6, it remains to show Condition 2 in order to apply Theorem 3.7. The proof idea is to show that parameter d increases by a factor of at least 1.5 in each iteration of the inner while-loop with $g = \bar{g}_c$. As $\bar{g}_c(Q, I_i, w)$ is upper bounded by $m(2/s_{\min}) \max_{I_i \in \mathcal{I}} (U_i - L_i)$, this means the inner loop executes at most $\lceil \log_{1.5}(m(2/s_{\min}) \max_{I_i \in \mathcal{I}} (U_i - L_i)) \rceil$ iterations for $g = \bar{g}_c$. For $g = \bar{g}_s$, we can argue in a similar way that at most $\lceil \log_2(m) \rceil$ iterations are executed. Similar to the previous section on MINSET with deterministic right-hand sides, we can also show that each iteration of the inner while-loop executes at most $\frac{1}{\tau}$ queries in expectation. Combining these insights, we can bound the expected number of queries during an execution of the outer while-loop by β . Formally proving the increase of d by 1.5 requires to take care of several technical challenges. The basic idea is to exploit that the interval I_i queried in Line 11 has an optimistic greedy value of at least d. If it satisfies $w_i - L_i \ge \frac{1}{2}(U_i - L_i)$, we show that the actual greedy value is at least d/2. When d is recomputed in Line 13, then I_i is a new member of the set $Q_{1/2}$ and leads to the increase of d/2.

We conclude the section by formalizing this proof idea. For each iteration j of the outer while-loop, let X_j be a random variable denoting the number of queries executed during iteration j and let Y_j be a variable indicating whether the iteration is executed ($Y_j = 1$) or not ($Y_j = 0$). We prove the following lemma, which implies that the algorithm also satisfies the second condition of Definition 3.6 and, thus, satisfies Theorem 5.1.

Lemma 5.4.
$$\mathbb{E}[X_j \mid Y_j = 1] \leq \beta = \frac{1}{\tau} \cdot (\lceil \log_{1.5}(m \cdot \frac{2 \cdot \max_{I_i \in \mathcal{I}}(U_i - L_i)}{s_{\min}}) \rceil + \lceil \log_2(m) \rceil).$$

Proof. Consider an iteration j of the outer while-loop of Algorithm 5. In the following, all probabilities and expected values are under the condition $Y_j = 1$.

For iteration j, let A_{jk} and B_{jk} be random variables denoting the number of queries in iteration k of the inner while-loop for $g = \bar{g}_c$ and $g = \bar{g}_s$, respectively. Then, $\mathbb{E}[X_j] = \sum_k \mathbb{E}[B_{jk}] + \sum_k \mathbb{E}[A_{jk}]$.

Let A_{jk} and B_{jk} be random variables indicating whether iteration k of the inner while-loop with $g = \bar{g}_c$ and $g = \bar{g}_s$, respectively, is executed $(\bar{A}_{jk} = 1, \bar{B}_{jk} = 1)$ or not. Given that such an iteration is executed, we can exploit that the termination criterion $w_i \ge \frac{1}{2} \cdot (U_i - L_i)$ occurs with probability $\mathbb{P}[(w_i - L_i) \ge (U_i - L_i)/2] \ge \mathbb{P}[w_i \ge (U_i + L_i)/2] \ge \tau$ to show $\mathbb{E}[A_{jk} \mid \bar{A}_{jk} = 1] \le \frac{1}{\tau}$ and $\mathbb{E}[B_{jk} \mid \bar{B}_{jk} = 1] \le \frac{1}{\tau}$. Using these bounds and exploiting that $\mathbb{E}[A_{jk} \mid \bar{A}_{jk} = 0] = \mathbb{E}[B_{jk} \mid \bar{B}_{jk} = 0] = 0$, we get

$$\begin{split} \mathbb{E}[X_{j}] &= \sum_{k} (\mathbb{E}[A_{jk}] + \mathbb{E}[B_{jk}]) \\ &= \sum_{k} \mathbb{P}[\bar{A}_{jk} = 1] \cdot \mathbb{E}[A_{jk} \mid \bar{A}_{jk} = 1] + \mathbb{P}[\bar{B}_{jk} = 0] \cdot \mathbb{E}[B_{jk} \mid \bar{A}_{jk} = 0] \\ &+ \sum_{k} \mathbb{P}[\bar{B}_{jk} = 1] \cdot \mathbb{E}[B_{jk} \mid \bar{B}_{jk} = 1] + \mathbb{P}[\bar{B}_{jk} = 0] \cdot \mathbb{E}[B_{jk} \mid \bar{B}_{jk} = 0] \\ &= \sum_{k} \mathbb{P}[\bar{A}_{jk} = 1] \cdot \mathbb{E}[A_{jk} \mid \bar{A}_{jk} = 1] + \sum_{k} \mathbb{P}[\bar{B}_{jk} = 1] \cdot \mathbb{E}[B_{jk} \mid \bar{B}_{jk} = 1] \\ &\leq \frac{1}{\tau} \cdot \left(\sum_{k} \mathbb{P}[\bar{A}_{jk} = 1] + \sum_{k} \mathbb{P}[\bar{B}_{jk} = 1] \right). \end{split}$$

Thus, it only remains to bound $\sum_k \mathbb{P}[\bar{A}_{jk} = 1] + \sum_k \mathbb{P}[\bar{B}_{jk} = 1]$. We do this by separately proving $\sum_k \mathbb{P}[\bar{A}_{jk} = 1] \leq \lceil \log_{1.5} (m \cdot (2/s_{\min}) \cdot \max_{I_i \in \mathcal{I}} (U_i - L_i)) \rceil$ and $\sum_k \mathbb{P}[\bar{B}_{jk} = 1] \leq \lceil \log_2(m) \rceil$. Putting all bounds together, we then get $\mathbb{E}[X_j] \leq \frac{1}{\tau} \cdot (\sum_k \mathbb{P}[\bar{A}_{jk} = 1] + \mathbb{P}[\bar{B}_{jk} = 1]) \leq \frac{1}{\tau} \cdot (\lceil \log_{1.5}(m \cdot D_{jk}) - D_{jk} - D_{jk}) \rceil$.

Putting all bounds together, we then get $\mathbb{E}[X_j] \leq \frac{1}{\tau} \cdot (\sum_k \mathbb{P}[A_{jk} = 1] + \mathbb{P}[B_{jk} = 1]) \leq \frac{1}{\tau} \cdot (|\log_{1.5}(m \cdot \frac{2 \cdot \max_{I_i \in \mathcal{I}}(U_i - L_i)}{s_{\min}})| + \lceil \log_2(m) \rceil).$

Upper bound for $\sum_k \mathbb{P}[\bar{A}_{jk} = 1]$. Note that $\sum_k \mathbb{P}[\bar{A}_{jk} = 1]$ is just the expected number of iterations of the inner while-loop with $g = \bar{g}_c$ during iteration j of the outer while-loop. We bound this number by showing that the value d increases by a factor of at least 1.5 in each iteration except possibly the first and last ones. As d is upper bounded by $a = (2/s_{\min}) \cdot m \cdot \max_{I_i \in \mathcal{I}} (U_i - L_i)$, the number of iterations then is at most $\lceil \log_{1.5} (a) \rceil$ for every realization of values.

To prove that d increases by a factor of 1.5 in each iteration of the inner while-loop, consider the last execution of the repeat-statement within an iteration of the inner while-loop with $g = \bar{g}_c$ (that is not the first or last one). Let I_i be the interval queried in that last iteration of the repeat-statement, let Q denote the set of I_i and all intervals that were queried before I_i , let Q' denote the set of all intervals queried before the current execution of the inner while-loop was started, and let $Q_{1/2} = \{I_j \in Q \setminus Q' \mid w_j - L_j \ge \frac{U_j - L_j}{2}\}$. Furthermore, let w denote the value of Line 10 computed before I_i was queried in the following line. Note that this notation matches the one used in the algorithm at the execution of Line 12 directly after I_i was queried.

Let \bar{d} denote the value d computed by the algorithm before I_i was queried. That is, $\bar{d} = g_c(Q', Q_{1/2} \setminus \{I_i\}, w')$ for the value w' that was computed in the previous execution of Line 10. By choice of I_i , we have $\bar{g}_c(Q \setminus \{I_i\}, I_i, w) \ge \bar{d}$ and $g_c(Q \setminus \{I_i\}, I_i, w) \ge \bar{d}/2$ (since $w_i - L_i \ge (U_i - L_i)/2$).

The value d computed after querying I_i is defined as

$$d = g_c(Q', Q_{1/2}, w) \ge g_c(Q', Q_{1/2} \setminus \{I_i\}, w) + g_c(Q' \cup Q_{1/2} \setminus \{I_i\}, I_i, w)$$

$$\ge g_c(Q', Q_{1/2} \setminus \{I_i\}, w) + g_c(Q \setminus \{I_i\}, I_i, w),$$

where the first inequality holds by definition of g_c and the second inequality holds because $Q \setminus \{I_i\} \supseteq Q' \cup Q_{1/2} \setminus \{I_i\}$ implies $g_c(Q' \cup Q_{1/2} \setminus \{I_i\}, I_i, w) \ge g_c(Q \setminus \{I_i\}, I_i, w)$.

In an iteration of the for-loop with $g = g_c$, the current value w only increases, which implies $w' \leq w$. Since the greedy value g_c only increases for increasing values w, we get $g_c(Q', Q_{1/2} \setminus \{I_i\}, w) \geq g_c(Q', Q_{1/2} \setminus \{I_i\}, w')$ and, therefore $g_c(Q', Q_{1/2}, w) \geq g_c(Q', Q_{1/2} \setminus \{I_i\}, w') + g_c(Q \setminus \{I_i\}, I_i, w)$. Plugging in $\overline{d} = g_c(Q', Q_{1/2} \setminus \{I_i\}, w')$ and $g_c(Q \setminus \{I_i\}, I_i, w) \geq \overline{d}/2$ yields $d \geq 1.5 \cdot \overline{d}$.

Note that this only shows that d increases by a factor of 1.5 compared to \bar{d} , the old value computed in the previous execution of the repeat statement. Our goal was to show that d increases by factor 1.5 compared to the previous iteration of the inner-while loop. However, since the value w only increases, the greedy values g_c only increase. This implies that \bar{d} can only be larger than the corresponding value at the end of the previous iteration of the inner while-loop. Thus, we have shown that d increases by a factor of 1.5 compared to its value at the end of the previous iteration of the inner-while loop. This implies $\sum_k \mathbb{P}[\bar{A}_{jk} = 1] \leq \lceil \log_{1.5} (m \cdot (2 \cdot \max_{I_i \in \mathcal{I}} (U_i - L_i))/s_{\min}) \rceil$.

Upper bound for $\sum_k \mathbb{P}[\bar{B}_{jk} = 1]$. Next, we use similar arguments to show $\sum_k \mathbb{P}[\bar{B}_{jk} = 1] \leq \lceil \log_2(m) \rceil$. Consider an iteration j of the outer while-loop.

We first show that, assuming $g = \bar{g}_s$, the value *d* increases by a factor of at least 2 in each iteration of the inner while-loop (except possibly the first and last ones) during iteration *j* of the outer while-loop.

By Line 6, iterations with $g = \overline{g}_s$ only consider values w with $b'_S(Q', w) < 1$ for all $S \in S$, where Q' is the set of intervals queried before the start of the current execution of the inner while-loop. For an interval I_i with $(w_i - L_i) \ge \frac{1}{2} \cdot (U_i - L_i)$, this means $a'_i = \frac{2 \cdot (w_i - L_i)}{s_{s_{\min}}} \ge \frac{U_i - L_i}{s_{\min}} \ge 1$. This implies that adding I_i to the solution satisfies all constraints for sets S with $I_i \in S$ (at least for all values w with $b'_S(Q', w) < 1$ for all $S \in S$). Thus, the greedy value g_s of I_i is then equal to the optimistic greedy value \overline{g}_s as even if $a_i = U_i - L_i$ the interval cannot satisfy more constraints.

Furthermore, as long as we only consider values w with $b'_S(Q', w) < 1$ for all $S \in S$, the greedy values g_s and \bar{g}_s of intervals I_i with $(w_i - L_i) \ge \frac{1}{2} \cdot (U_i - L_i)$ can only increase with increasing w and never increase for decreasing w. The reason for this is that such an increase in w can only lead to *more* constraints becoming *not* satisfied for the right-hand sides $b'_S(Q', w)$. Thus, the number of constraints that adding I_i can satisfy only increases. This also means, that the value w as used by the algorithm will never decrease during the current iteration of the outer while-loop with $g = \bar{g}_s$.

Using these observations, we can essentially repeat the proof of the previous case to show that the value d increases by a factor of at least 2 in each iteration of the inner while-loop with $g = \bar{g}_s$ (except possibly the first and last ones).

Let I_i be the interval queried in the last iteration of the repeat-statement within such an iteration of the inner while-loop, let Q denote the set of I_i and all intervals that were queried before I_i , let Q' denote the set of all intervals queried before the current execution of the inner while-loop was started, and let $Q_{1/2} = \{I_j \in Q \setminus Q' \mid w_j - L_j \geq \frac{U_j - L_j}{2}\}$. Furthermore, let w denote the value of Line 10 computed before I_i was queried in the following line. Note that this notation matches the one used in the algorithm at the execution of Line 12 directly after I_i was queried.

Let \bar{d} denote the value d computed by the algorithm *before* I_i was queried. That is, $\bar{d} = g_s(Q', Q_{1/2} \setminus \{I_i\}, w')$ for the value w' that was computed in the previous execution of Line 10. By choice of I_i , we have

 $\bar{g}_s(Q \setminus \{I_i\}, I_i, w) \ge \bar{d}$ and $g_s(Q \setminus \{I_i\}, I_i, w) = \bar{g}_s(Q \setminus \{I_i\}, I_i, w) \ge \bar{d}$ (since $w_i - L_i \ge (U_i - L_i)/2$ and as argued above).

The value d computed after querying I_i is defined as

$$d = g_s(Q', Q_{1/2}, w) \ge g_s(Q', Q_{1/2} \setminus \{I_i\}, w) + g_s(Q' \cup Q_{1/2} \setminus \{I_i\}, I_i, w)$$

$$\ge g_s(Q', Q_{1/2} \setminus \{I_i\}, w) + g_s(Q \setminus \{I_i\}, I_i, w).$$

Here, the first inequality holds by definition of g_s , definition of $Q_{1/2}$, and by the assumption that we only consider values w with $b'_S(Q', w) < 1$ for all $S \in S$. The second inequality holds because $Q \setminus \{I_i\} \supseteq Q' \cup Q_{1/2} \setminus \{I_i\}$ implies $g_s(Q' \cup Q_{1/2} \setminus \{I_i\}, I_i, w) \ge g_s(Q \setminus \{I_i\}, I_i, w)$. Note that this implication only holds under the assumption that $b'_S(Q', w) < 1$ for all $S \in S$ for intervals I_i with $a'_i \ge 1$.

As argued above, the value w used by the algorithm only increases, i.e., $w' \leq w$. Furthermore, again as argued above, the greedy value g_s of intervals I_j with $a'_j \geq 1$ only increases for increasing values w. Thus, we get $g_s(Q', Q_{1/2} \setminus \{I_i\}, w) \geq g_s(Q', Q_{1/2} \setminus \{I_i\}, w')$ and, therefore, $g_s(Q', Q_{1/2}, w) \geq g_s(Q', Q_{1/2} \setminus \{I_i\}, w') + g_s(Q \setminus \{I_i\}, I_i, w)$. Plugging in $\overline{d} = g_s(Q', Q_{1/2} \setminus \{I_i\}, w')$ and $g_s(Q \setminus \{I_i\}, I_i, w) \geq \overline{d}$ yields $d \geq 2 \cdot \overline{d}$. This implies that d increased by a factor of at least 2 compared to its old value at the end of the previous iteration of the inner while-loop. As the greedy value g_s is upper bounded by m, we get $\sum_k \mathbb{P}[\overline{B}_{jk} = 1] \leq \lceil \log_2(m) \rceil$.

The Lemmas 5.2 to 5.4 imply that Algorithm 5 satisfies Definition 3.6 for $\alpha = 2$, $\gamma = 2/s_{\min}$ and $\beta = \frac{1}{\tau} (\lceil \log_{1.5}(m \cdot 2(\max_{I_i \in \mathcal{I}}(U_i - L_i))/s_{\min}) \rceil + \lceil \log_2(m) \rceil)$. Thus, Theorem 3.7 implies Theorem 5.1.

6 Final remarks

In this paper, we provide the first results for MINSET under stochastic explorable uncertainty by exploiting a connection to a covering problem and extending techniques for solving covering problems to our setting with uncertainty.

Since our results, in expectation, break adversarial lower bound instances for a number of interesting combinatorial problems under explorable uncertainty, e.g., matching, knapsack, solving ILPs [33], we hope that our techniques lay the foundation for solving more general problems. In particular if we consider the mentioned problems with uncertainty in the cost coefficients and our goal is to query elements until we can identify an optimal solution to the underlying problem, then all these problems admit the same connection to covering problems as MINSET and can also be written as covering ILPs with uncertain coefficients and right-hand sides. The difference to MINSET is that the number of constraints in the covering representation for these problems might be exponential in the input size of the underlying optimization problem. In a sense, each feasible solution for the underlying optimization problem would define a constraint in the covering representation. This means that using the results of this paper as a blackbox to solve such optimization problems under explorable uncertainty does not yield sublinear competitive ratios as the number of constraints becomes too large. For future research, we suggest to investigate whether one can exploit the additional structure in these constraints to still achieve improved competitive ratios.

With respect to the results of this paper, we leave open whether the second log factor in our main result, Theorem 5.1, is necessary. Furthermore, the best competitive ratio achievable in exponential running time also remains open. Finally, we remark that we expect our algorithms to be parameterizable by the choice of the balancing parameter. We defined the parameter as the probability that the precise value is larger than the center of the corresponding interval. Alternatively, one could pick any fraction of the interval, say one third, as a threshold and define the parameter as the probability that the precise value lies outside the first third of its interval. As the only parts of our algorithm that use the definition of the balancing parameter are the termination criteria of the repeat-statements, the definition of the scaling factor γ and the definition of the set $Q_{1/2}$ in Algorithm 5, we expect that adjusting these parts to a different balancing parameter suffices to make our algorithms work for such parameters. The choice of the threshold for the balancing parameter would then influence the constant factor within the competitive ratios.

References

- [1] Arpit Agarwal, Sepehr Assadi, and Sanjeev Khanna. Stochastic submodular cover with limited adaptivity. In *SODA*, pages 323–342. SIAM, 2019.
- [2] Noga Alon, Baruch Awerbuch, and Yossi Azar. The online set cover problem. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 100–105, 2003.
- [3] Evripidis Bampis, Christoph Dürr, Thomas Erlebach, Murilo Santos de Lima, Nicole Megow, and Jens Schlöter. Orienting (hyper)graphs under explorable stochastic uncertainty. In ESA, volume 204 of LIPIcs, pages 10:1–10:18, 2021.
- [4] Soheil Behnezhad, Avrim Blum, and Mahsa Derakhshan. Stochastic vertex cover with few queries. In SODA, pages 1808–1846. SIAM, 2022.
- [5] Avrim Blum, John P. Dickerson, Nika Haghtalab, Ariel D. Procaccia, Tuomas Sandholm, and Ankit Sharma. Ignorance is almost bliss: Near-optimal stochastic matching with few queries. *Oper. Res.*, 68(1):16–34, 2020.
- [6] R. Bruce, M. Hoffmann, D. Krizanc, and R. Raman. Efficient update strategies for geometric computing with uncertainty. *Theory of Computing Systems*, 38(4):411–423, 2005.
- [7] Steven Chaplick, Magnús M. Halldórsson, Murilo S. de Lima, and Tigran Tonoyan. Query minimization under stochastic uncertainty. *Theor. Comput. Sci.*, 895:75–95, 2021.
- [8] Vasek Chvatal. A greedy heuristic for the set-covering problem. *Mathematics of Operations Research*, 4(3):233–235, 1979.
- [9] Amol Deshpande, Lisa Hellerstein, and Devorah Kletenik. Approximation algorithms for stochastic submodular set cover with applications to boolean function evaluation and min-knapsack. ACM Trans. Algorithms, 12(3):42:1–42:28, 2016.
- [10] Irit Dinur and David Steurer. Analytical approach to parallel repetition. In *STOC*, pages 624–633. ACM, 2014.
- [11] Gregory Dobson. Worst-case analysis of greedy heuristics for integer programming with nonnegative data. *Mathematics of Operations Research*, 7(4):515–531, 1982.
- [12] T. Erlebach, M. Hoffmann, and F. Kammer. Query-competitive algorithms for cheapest set problems under uncertainty. *Theoretical Computer Science*, 613:51–64, 2016.
- [13] Thomas Erlebach, Murilo Santos de Lima, Nicole Megow, and Jens Schlöter. Learning-augmented query policies for minimum spanning tree with uncertainty. In *ESA*, volume 244 of *LIPIcs*, pages 49:1–49:18. Schloss Dagstuhl Leibniz-Zentrum für Informatik, 2022.
- [14] Thomas Erlebach and Michael Hoffmann. Minimum spanning tree verification under uncertainty. In *WG*, volume 8747 of *Lecture Notes in Computer Science*, pages 164–175. Springer, 2014.

- [15] Thomas Erlebach, Michael Hoffmann, and Frank Kammer. Query-competitive algorithms for cheapest set problems under uncertainty. *Theoretical Computer Science*, 613:51–64, 2016.
- [16] T. Feder, R. Motwani, R. Panigrahy, C. Olston, and J. Widom. Computing the median with uncertainty. *SIAM Journal on Computing*, 32(2):538–547, 2003.
- [17] Rohan Ghuge, Anupam Gupta, and Viswanath Nagarajan. The power of adaptivity for stochastic submodular cover. In *ICML*, volume 139 of *Proceedings of Machine Learning Research*, pages 3702– 3712. PMLR, 2021.
- [18] Michel X. Goemans and Jan Vondrák. Covering minimum spanning trees of random subgraphs. Random Struct. Algorithms, 29(3):257–276, 2006.
- [19] Michel X. Goemans and Jan Vondrák. Stochastic covering and adaptivity. In LATIN, volume 3887 of Lecture Notes in Computer Science, pages 532–543. Springer, 2006.
- [20] Daniel Golovin and Andreas Krause. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. J. Artif. Intell. Res., 42:427–486, 2011.
- [21] Fabrizio Grandoni, Anupam Gupta, Stefano Leonardi, Pauli Miettinen, Piotr Sankowski, and Mohit Singh. Set covering with our eyes closed. SIAM J. Comput., 42(3):808–830, 2013.
- [22] Anupam Gupta, Gregory Kehne, and Roie Levin. Set covering with our eyes wide shut. *CoRR*, abs/2304.02063, 2023.
- [23] M. M. HalldÄ³rsson and M. S. de Lima. Query-competitive sorting with uncertainty. In *MFCS*, volume 138 of *LIPIcs*, pages 7:1–7:15, 2019.
- [24] Michael Hoffmann, Thomas Erlebach, Danny Krizanc, Matús Mihalák, and Rajeev Raman. Computing minimum spanning trees with uncertainty. In STACS, volume 1 of LIPIcs, pages 277–288. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Germany, 2008.
- [25] Sungjin Im, Viswanath Nagarajan, and Ruben van der Zwaan. Minimum latency submodular cover. *ACM Trans. Algorithms*, 13(1):13:1–13:28, 2016.
- [26] Simon Kahan. A model for data in motion. In STOC, pages 267-277. ACM, 1991.
- [27] Prabhanjan Kambadur, Viswanath Nagarajan, and Fatemeh Navidi. Adaptive submodular ranking. In IPCO, volume 10328 of Lecture Notes in Computer Science, pages 317–329. Springer, 2017.
- [28] Stavros G. Kolliopoulos and Neal E. Young. Tight approximation results for general covering integer programs. In FOCS, pages 522–528. IEEE Computer Society, 2001.
- [29] Stavros G. Kolliopoulos and Neal E. Young. Approximation algorithms for covering/packing integer programs. J. Comput. Syst. Sci., 71(4):495–505, 2005.
- [30] Takanori Maehara and Yutaro Yamaguchi. Stochastic packing integer programs with few queries. *Math. Program.*, 182(1):141–174, 2020.
- [31] Takanori Maehara and Yutaro Yamaguchi. Stochastic packing integer programs with few queries. *Mathematical Programming*, 182(1):141–174, 2020.
- [32] N. Megow, J. Meißner, and M. Skutella. Randomization helps computing a minimum spanning tree under uncertainty. SIAM Journal on Computing, 46(4):1217–1240, 2017.

- [33] J. Meißner. Uncertainty Exploration: Algorithms, Competitive Analysis, and Computational Experiments. PhD thesis, Technischen Universität Berlin, 2018.
- [34] Fatemeh Navidi, Prabhanjan Kambadur, and Viswanath Nagarajan. Adaptive submodular ranking and routing. Oper. Res., 68(3):856–877, 2020.
- [35] Sridhar Rajagopalan and Vijay V Vazirani. Primal-dual rnc approximation algorithms for set cover and covering integer programs. SIAM Journal on Computing, 28(2):525–540, 1998.
- [36] David B. Shmoys and Chaitanya Swamy. Stochastic optimization is (almost) as easy as deterministic optimization. In FOCS, pages 228–237. IEEE Computer Society, 2004.
- [37] Vijay V Vazirani. Approximation algorithms, volume 1. Springer, 2001.
- [38] Jan Vondrák. Shortest-path metric approximation for random subgraphs. *Random Struct. Algorithms*, 30(1-2):95–104, 2007.
- [39] Weina Wang, Anupam Gupta, and Jalani Williams. Probing to minimize. In *ITCS*, volume 215 of *LIPIcs*, pages 120:1–120:23. Schloss Dagstuhl Leibniz-Zentrum für Informatik, 2022.
- [40] Laurence A. Wolsey. An analysis of the greedy algorithm for the submodular set covering problem. *Comb.*, 2(4):385–393, 1982.

A Comparison with Maehara and Yamaguchi [31]

We consider the framework by Maehara and Yamaguchi [31] on the set selection problem. Since their algorithm is designed for maximization problems, we consider the maximization variant of MINSET, i.e., we have to find the set $S \in S$ of maximum w(S) and determine the corresponding value. We remark that all our results also translate to the maximization variant (cf. Appendix D).

The algorithm by Maehara and Yamaguchi, in each iteration, solves the LP-relaxation of the *optimistic* version of the given ILP, which assumes $w_i = U_i$ for all $I_i \in \mathcal{I}$. In this case, the ILP under consideration formalizes the set selection problem (and *not* the query minimization problem as formalized by (MINSETIP)). The following LP-relaxation formulates the optimistic LP for a set selection instance $(\mathcal{I}, \mathcal{S})$:

$$\begin{array}{ll} \max & \sum_{I_i \in \mathcal{I}} x_i \cdot U_i \\ \text{s.t.} & \sum_{I_i \in S} x_i \geq y_S \cdot |S| & \forall S \in \mathcal{S} \\ & \sum_{I_i \in E} x_i \leq \sum_{S \in \mathcal{S}} y_S \cdot |S| \\ & \sum_{S \in \mathcal{S}} y_S = 1 \\ & 0 \leq x_i \leq 1 & \forall I_i \in \mathcal{I} \\ & 0 \leq y_S \leq 1 & \forall S \in \mathcal{S} \end{array}$$

Here variable y_S models whether set S is selected as the set of maximum value ($y_S = 1$) or not ($y_S = 0$) and the third constraint makes sure that, at least integrally, exactly one set is selected. The variables x_i model whether an interval I_i is part of the selected set or not, and the first two constraints ensure that (integrally) exactly the members of the selected set S (with $y_S = 1$) are selected. Note that we use this ILP instead of (SETSELIP) because the algorithm by Maehara and Yamaguchi requires variables that correspond to elements that can be queried. The algorithm in each iteration solves the LP-relaxation to obtain an optimal fractional solution (x, y), and queries each I_i with probability x_i .

In the following, we give an instance of the set selection problem under stochastic explorable uncertainty for which the algorithm has a competitive ratio of $\Omega(n)$. Let $\mathcal{I} = \{I_0, \ldots, I_n\}$ with $I_0 = (d \cdot n - \epsilon, d \cdot n + \epsilon)$ and $I_i = (0, d + \epsilon)$, i > 0, for some large d and some small $\epsilon > 0$. Let $\mathcal{S} = \{S_1, S_2\}$ with $S_1 = \{I_0\}$ and $S_2 = \{I_1, \ldots, I_n\}$. Assume uniform distributions and consider the algorithm that starts by querying I_0 and afterwards queries the intervals of S_2 in an arbitrary order. In expectation, this algorithm only needs a constant number of queries to solve the instance.

The algorithm by Maehara and Yamaguchi on the other hand, in the first iteration, solves the LPrelaxation and obtains the optimal solution (x, y) with $x_0 = 0$ and $x_i = 1$ for all i > 1. This means that all elements of S_2 are queried with a probability of 1. Thus, the algorithm queries at least n elements, which implies a competitive ratio of $\Omega(n)$. This means that applying the algorithm by Maehara and Yamaguchi [31] to the set selection problem does not improve upon the adversarial lower bound.

B Equivalence of MINSET and Solving (MINSETIP)

In this paper, we heavily exploit that solving MINSET is equivalent to solving (MINSETIP). With the following lemma, we show that this is indeed the case.

Lemma B.1. Solving MINSET is equivalent to solving (MINSETIP).

Proof. We show the lemma by proving the following claim: A query set $Q \subseteq \mathcal{I}$ is feasible for MINSET if and only if vector x, with $x_i = 1$ for all $I_i \in Q$ and $x_i = 0$ otherwise, is a feasible solution for the corresponding (MINSETIP).

Let Q be feasible for MINSET, and let x be a vector with $x_i = 1$ for all $I_i \in Q$ and $x_i = 0$ otherwise. By definition, each feasible solution for MINSET must query all non-trivial elements of some cheapest set S^* with $w(S^*) = w^*$ as this is the only way to determine the value w^* . Let $N(S^*)$ denote those non-trivial elements, then we can rewrite the initial lower limit of S^* as $L_{S^*} = w^* + \sum_{I_i \in N(S^*)} (L_i - w_i)$ (the lower limits of trivial elements are covered by w^*). This implies $w^* - L_{S^*} = \sum_{I_i \in N(S^*)} (w_i - L_i)$. As $N(S^*) \subseteq Q$, we get $\sum_{I_i \in S^*} x_i \cdot (w_i - L_i) = \sum_{I_i \in S^* \cap Q} (w_i - L_i) \ge w^* - L_{S^*}$. Thus, x satisfies the constraint for S^* . Let $S \neq S^*$. Since Q is feasible, $L_S(Q)$ has to be at least w^* as otherwise querying Q would not prove that w^* is the minimum set value. After querying Q, the lower limit of set S is

$$L_S(Q) = \sum_{I_i \in S \setminus Q} L_i + \sum_{I_i \in S \cap Q} w_i$$

= $\sum_{I_i \in S} L_i - \sum_{I_i \in S \cap Q} L_i + \sum_{I_i \in S \cap Q} w_i$
= $L_S + \sum_{I_i \in S \cap Q} (w_i - L_i).$

Thus, $L_S(Q) = L_S + \sum_{I_i \in S \cap Q} (w_i - L_i) \ge w^*$ must hold, which implies $\sum_{I_i \in S \cap Q} (w_i - L_i) \ge w^* - L_S$ and $\sum_{I_i \in S} x_i \cdot (w_i - L_i) \ge w^* - L_S$. We can conclude that x is feasible.

For the other direction consider a feasible solution x for (MINSETIP) and the corresponding set $Q = \{I_i \mid x_i = 1\}$. Consider some cheapest set S^* . As $\sum_{I_i \in N(S^*)} (w_i - L_i) = w^* - L_{S^*}$ and $\sum_{I_i \in S^* \setminus N(S^*)} (w_i - L_i) = \sum_{I_i \in S^* \setminus N(S^*)} (w_i - w_i) = 0$, for x to be feasible it must hold $x_i = 1$ for all $I_i \in N(S^*)$. Thus, Q contains all non-trivial elements of some cheapest set S^* . To show that Q is a feasible solution for MINSET, it remains to show that $L_S(Q) \ge w^*$ for all $S \ne S^*$. Consider an arbitrary $S \ne S^*$. As $\sum_{I_i \in S} x_i \cdot (w_i - L_i) \ge (w^* - L_S)$, we have $\sum_{I_i \in S \cap Q} (w_i - L_i) \ge w^* - L_S$, which implies $L_S(Q) = L_S + \sum_{I_i \in S \cap Q} (w_i - L_i) \ge w^*$.

C Hardness of Approximation

In the following, we show that (MINSETIP) is not only a special case of the multiset multicover problem but also contains the hard instances of this problem.

Erlebach et al. [15] showed that offline MINSET, for the problem variant where it is not necessary to compute the value w^* , is NP-hard via reduction from vertex cover. In the reduction by [15] all intervals of the set S^* with $w(S^*) = w^*$ are trivial and, thus, the result translates to the problem variant where one has to compute w^* . In the following, we strengthen this result by showing that the offline problem is as hard to approximate as SETCOVER.

In SETCOVER, we are given a set of elements $U = \{1, ..., n\}$ and a family of sets $\overline{S} = \{\overline{S}_1, ..., \overline{S}_m\}$ with $\overline{S}_j \subseteq U$. The goal is to find a subset $H \subseteq \overline{S}$ of minimum cardinality such that $\bigcup_{\overline{S}_j \in H} \overline{S}_j = U$.

Theorem C.1. There is an approximation-factor preserving reduction from SETCOVER to offline MINSET.

Proof. Given an instance (U, \overline{S}) of SETCOVER, we construct an offline MINSET instance as follows:

- 1. Add a trivial interval $I_r = \{w_r\}$.
- 2. Add a single set $C = \{I_r\}$.
- 3. For each $j \in U$, add a set S_j .
- 4. For each $\bar{S}_i \in \bar{S}$:
 - (a) Add an interval $I_i = (L_i, U_i)$ with $L_i = 0$, $U_i = w_r + \delta$ and $w_i = w_r + \epsilon$ for a common $\delta > \epsilon > 0$ and some infinitesimally small $\epsilon > 0$.
 - (b) For each $j \in \overline{S}_i$, add interval I_i to set S_j .

This reduction clearly runs in polynomial time. To finish the proof, we show the following claim: *There* is a SETCOVER solution H of cardinality k if and only if there is a feasible query set Q for the constructed offline MINSET instance with |Q| = k.

By definition of the constructed instance, set $C = \{I_r\}$ is the set of minimum value $w^* = w_r$. Each feasible query set Q for the offline MINSET instance must prove that $L_S(Q) \ge w_r$ holds for each $S \in S \setminus \{C\}$. Recall that $L_S(Q)$ is the lower limit of S after querying Q. By definition of the constructed intervals and sets, a query set Q is feasible if and only if $|Q \cap S| \ge 1$ for each $S \in S \setminus \{C\}$, i.e., Q has to contain at least one element of each $S \in S \setminus \{C\}$.

For the first direction, consider an arbitrary set cover H for the given SETCOVER instance and construct $Q = \{I_i \mid \overline{S}_i \in H\}$. Clearly |Q| = |H|. Since H is a set cover, each $j \in U$ is contained in at least one $\overline{S}_i \in H$. If $j \in U$ is contained in \overline{S}_i , then, by construction, I_i is contained in S_j . Thus, as H covers all elements $j \in U$, set Q contains at least one member of each $S_j \in S \setminus \{C\}$ and, therefore, is a feasible query set.

For the second direction, consider an arbitrary feasible query set Q of the constructed instance and construct $H = \{\overline{S}_i \mid I_i \in Q\}$. Clearly, |Q| = |H|. Since Q is feasible, it contains at least one member of each $S_j \in S \setminus \{C\}$. If $I_i \in S_j$ is contained in Q, then, by construction, set $\overline{S}_i \in H$ covers element j. As Q contains at least one member of each $S_j \in S \setminus \{C\}$, it follows that H covers U.

Dinur and Steurer [10] showed that it is NP-hard to approximate SETCOVER within a factor of $(1 - \alpha) \cdot \ln n$ for any $\alpha > 0$, where n is the number of elements in the instance, via a reduction running in time $n^{1/\alpha}$. Consider the construction of Theorem C.1. Since the sets in the constructed offline MINSET instance correspond to the elements in the input SETCOVER instance, the construction implies the following corollary.

Corollary C.2. For every $\alpha > 0$, it is NP-hard to approximate offline MINSET within a factor of $(1 - \alpha) \cdot \ln m$, where m = |S| is the number of sets. The reduction runs in time $m^{1/\alpha}$.

We show that Theorem C.1 and Corollary C.2 apply also to MINSET under stochastic explorable uncertainty, even if the precise value w_i of each I_i is drawn independently and uniformly at random from (L_i, U_i) .

Theorem 3.1. For any fixed $\alpha > 0$, it is NP-hard to compute a query strategy that is $(1 - \alpha) \cdot \ln m$ competitive for MINSET under uncertainty even if the precise value w_i of each I_i is drawn independently and uniformly at random from (L_i, U_i) . The same inapproximability holds also for offline MINSET.

Proof. The hardness of approximation for offline MINSET follows directly from Corollary C.2.

We continue to show the statement on MINSET under uncertainty with uniform distributions. Consider the reduction of Theorem C.1 with w_r towards 0 and/or δ towards ∞ . With w_r running towards 0 and/or δ running towards ∞ , the probability that it is sufficient to query one $I_i \in S_j$ to show that $w_r \leq L_{S_j}(Q)$, for some query set Q, goes towards 1. Thus, the probability that any set Q that contains at least one member of each $S \in S$ is feasible goes towards one as well. Thus, $\lim_{w_r \to 0} \mathbb{E}[\text{OPT}] = \lim_{\delta \to \infty} \mathbb{E}[\text{OPT}] = |H^*|$, where H^* is the optimal solution for the input SETCOVER instance. Therefore, by Theorem C.1, in order to be $((1 - \alpha) \cdot \ln m)$ -competitive, the query strategy has to compute an $((1 - \alpha) \cdot \ln n)$ -approximation for set cover. This implies NP-hardness.

D The Maximization Variant of MINSET

Consider the set selection problem MAXSET, which has the same input as MINSET, but now the goal is to determine a set of *maximum value* and to determine the corresponding value.

In MAXSET, a set $Q \subseteq \mathcal{I}$ is *feasible* if all non-trivial elements of a set S^* of maximum value w^* are in Q and $U_S(Q) \leq w^*$ for all sets $S \neq S^*$. If Q would not contain all non-trivial elements of some set S^* of maximum value, then the maximum value w^* would still be unknown and the problem would not be solved yet. If $U_S(Q) \geq w^*$ for some S, then both S^* and S could still be of maximum value or not and we have not yet determined the set of maximum value. Thus, the problem would not be solved yet. Our goal is again to adaptively query a feasible query set and minimize the number of queries.

In the following, we briefly sketch why all our results on MINSET transfer to MAXSET. Analogous to MINSET, one can show that the following ILP characterizes the problem. That is, a query set Q is feasible if and only if vector x with $x_i = 1$ if $I_i \in Q$ and $x_i = 0$ otherwise is feasible for the ILP.

$$\begin{array}{ll} \min & \sum_{I_i \in \mathcal{I}} x_i \\ \text{s.t.} & \sum_{I_i \in S} x_i \cdot (U_i - w_i) \ge (U_S - w^*) \quad \forall S \in \mathcal{S} \\ & x_i \in \{0, 1\} \qquad \qquad \forall I_i \in \mathcal{I} \end{array}$$
 (MAXSETIP)

In the offline setting, the ILP is the exact same special case of the multiset multicover problem as (MINSETIP), which suffices to observe that all our observations on offline MINSET translate to offline MAXSET. Under uncertainty, the only difference to MINSET is, that, in contrast to (MINSETIP), a small value w_i leads to a larger coefficient $(U_i - w_i)$ and a small value w^* leads to larger right-hand sides. Using the inverse balancing parameter $\bar{\tau} = \min_{I_i \in \mathcal{I}} \bar{\tau}_i$ with $\bar{\tau}_i = \mathbb{P}[w_i \leq \frac{U_i + L_i}{2}]$, it is not hard to see, that, even under uncertainty, MINSET and MAXSET are essentially the same problem. Thus, all our results translate.