

A Tight Bound for Stochastic Submodular Cover

Lisa Hellerstein

*Dept. of Computer Science and Engineering
NYU Tandon School of Engineering
Brooklyn, NY 11201 USA*

LISA.HELLERSTEIN@NYU.EDU

Devorah Kletenik

*Dept. of Computer and Information Science
Brooklyn College, City University of New York
Brooklyn, NY 11210 USA*

KLETENIK@SCI.BROOKLYN.CUNY.EDU

Srinivasan Parthasarathy

*IBM T. J. Watson Research Center
Yorktown Heights, NY 10598 USA*

SPARTHA@US.IBM.COM

Abstract

We show that the Adaptive Greedy algorithm of Golovin and Krause achieves an approximation bound of $(\ln(Q/\eta)+1)$ for Stochastic Submodular Cover: here Q is the “goal value” and η is the minimum gap between Q and any attainable utility value $Q' < Q$. Although this bound was claimed by Golovin and Krause in the original version of their paper, the proof was later shown to be incorrect by Nan and Saligrama. The subsequent corrected proof of Golovin and Krause gives a quadratic bound of $(\ln(Q/\eta)+1)^2$. A bound of $56(\ln(Q/\eta)+1)$ is implied by work of Im et al. Other bounds for the problem depend on quantities other than Q and η . Our bound restores the original bound claimed by Golovin and Krause, generalizing the well-known $(\ln m+1)$ approximation bound on the greedy algorithm for the classical Set Cover problem, where m is the size of the ground set.

1. Introduction

The *Stochastic Submodular Cover* problem was introduced by Golovin and Krause (2011). Their interest in this problem was motivated by applications in a variety of areas, including sensor placement, viral marketing, and active learning. The problem combines and generalizes two previously studied generalizations of the classical NP-complete Set Cover problem: *Submodular Cover* and *Stochastic Set Cover*.

In the Submodular Cover problem, introduced by Wolsey (1982), the input consists of a finite set of items E , a utility function $f: 2^E \rightarrow \mathbb{R}_{\geq 0}$ (given by an oracle), and a cost function $c: E \rightarrow \mathbb{R}_{> 0}$. The function f is a *polymatroid function*, that is, $f(\emptyset) = 0$ and f is both *monotone* and *submodular*. Monotonicity means that $\forall A \subseteq B \subseteq E, f(A) \leq f(B)$. Submodularity means that

$$\forall A \subseteq B \subseteq E, \forall e \in E \setminus B$$

$$f(A \cup \{e\}) - f(A) \geq f(B \cup \{e\}) - f(B)$$

The goal in the Submodular Cover problem is to find a subset $A \subseteq E$ such that $f(A) = f(E)$ and $\sum_{e \in A} c(e)$ is minimized.

The fact that the Submodular Cover problem is a generalization of the classical Set Cover problem can be seen as follows. In the classical Set Cover problem, the input consists of a ground set of elements $D = \{d_1, \dots, d_m\}$, a family of subsets of the ground elements, $\mathcal{S} = \{S_1, \dots, S_n\}$ whose union is equal to D , and a cost function $c: \mathcal{S} \rightarrow \mathbb{R}_{> 0}$. View each $S_i \in \mathcal{S}$ as an item, and define the polymatroid utility function $f: 2^{\mathcal{S}} \rightarrow \mathbb{R}_{\geq 0}$ such that for $A \subseteq \mathcal{S}$, $f(A) = |\bigcup_{S_i \in A} S_i|$.




Figure 1: Relationship between cover problems

Finding a min-cost subset of items A such that $f(A) = f(\mathcal{S})$ is equivalent to finding a min-cost cover of the ground elements.

Wolsey’s greedy algorithm for Submodular Set Cover repeatedly selects the item yielding the largest increase in utility, per unit cost. Wolsey’s analysis of this algorithm shows that it achieves an approximation bound of $(\ln(Q/\eta) + 1)$, where $Q = f(E)$ and η is the minimum gap between Q and any attainable utility value $Q' < Q$.

The Stochastic Set Cover problem is a different generalization of the classical Set Cover problem, which captures situations in which there is uncertainty about which ground elements will be covered by each “subset”. A motivating example for this problem is a sensor placement problem, where prior to placing a sensor in a given location, it is unclear which subset of locations will be monitored by the sensor (Golovin & Krause, 2011; Goemans & Vondrák, 2006). The input to the Stochastic Set Cover problem consists of a ground set $D = \{d_1, \dots, d_m\}$, a family $\mathcal{S} = \{S_1, \dots, S_n\}$ of stochastic subsets of the ground set, and a cost function $c: \mathcal{S} \rightarrow \mathbb{R}_{>0}$. There is an integer $k > 0$, such that each S_i is an independent random variable, whose value can be any one of k different subsets of ground set D (the k subsets can be different for each i). For each of the k^n possible combinations of values of the n variables S_i , it holds that $\bigcup_{i=1}^n S_i = D$. Choosing S_i incurs cost $c(S_i)$, the value of S_i is not revealed until S_i is chosen, and each S_i can be chosen at most once. The problem is to determine how to sequentially and adaptively select stochastic subsets S_i , until the union of the (values of the) chosen S_i is equal to D , so as to minimize the expected cost incurred. Selecting “adaptively” means that the choice of the next subset at each step can depend on the revealed values of the S_i ’s that were already chosen (but not on the values of the other S_i ’s).

A natural greedy policy for this problem is to repeatedly choose the stochastic subset which is expected to cover the largest number of uncovered items, per unit cost. Liu et al. (2008) presented and analyzed this policy for a special case of Stochastic Set Cover. They claimed that the expected cost incurred by this policy is at most a factor of $(\ln m + 1)$ greater than the optimal expected cost of any other sequential, adaptive policy for the problem. However, as we describe in Appendix B, the analysis of Liu et al. was not correct.

The problem addressed in this paper, the Stochastic Submodular Cover problem, combines and generalizes the Stochastic Set Cover problem and the Submodular Cover problem.¹ We

1. When Wolsey introduced the Submodular Cover problem, he called it Submodular Set Cover. Similarly, Stochastic Submodular Cover is sometimes called Stochastic Submodular Set Cover.

give an informal definition of the problem here, and a more formal definition in Section 1.1. In this problem, there is a set of stochastic items E , and a cost function $c: E \rightarrow \mathbb{R}_{>0}$. Each item can be in one of k possible “states”. The states of the different items are independent. There is a utility function f that assigns a real value to each subset A of items from E . This value depends not only on which items are in A , but also on the states of those items. The value of f on the empty set is 0. For any fixed assignment of states to items, f is monotone and submodular. On the set of all items in E , the value of f is guaranteed to be equal to a fixed value Q , called the *goal value*, regardless of the states of the items. The state of an item $e \in E$ is not revealed until item e is chosen and its cost $c(e)$ is incurred. Each item can be chosen at most once. The problem is to find a policy for sequentially selecting items from E until the value of f on the set of chosen items (in their revealed states) is equal to Q , so as to minimize the expected cost incurred. The policy can be adaptive, meaning that the choice of the next item to select can depend on the states of the items already chosen (but not on the states of the other items).

We say that an algorithm achieves a γ -approximation bound for this problem, where $\gamma \geq 1$, if it produces a sequential, adaptive policy whose expected cost is at most a factor of γ larger than the minimum expected cost of any sequential, adaptive policy.

Golovin and Krause (2011) proposed and analyzed a greedy algorithm for Stochastic Submodular Cover, called Adaptive Greedy. It repeatedly selects the item which would yield the largest expected increase in utility, per unit cost. This generalizes the greedy algorithms described above. Golovin and Krause also considered a setting where the Adaptive Greedy algorithm does not make an exactly optimal choice in each greedy step, but instead makes an α -approximate greedy choice. They claimed that Adaptive Greedy achieves an approximation bound of $\alpha(\ln(Q/\eta) + 1)$, but subsequently Nan and Saligrama (2017) showed that their proof of the bound was incorrect. (The error in Golovin and Krause’s proof is not the same as the one in the paper of Liu et al., and the proofs use different approaches.) In a later version of their paper, Golovin and Krause (2017) presented a corrected version of their proof, but with a weaker quadratic bound of $\alpha(\ln(Q/\eta) + 1)^2$.

Im et al. (2016) studied a problem called Weighted Stochastic Submodular Ranking, which generalizes Stochastic Submodular Cover. Their work implies that Adaptive Greedy, with $\alpha = 1$, achieves a bound of $56(\ln(Q/\eta) + 1)$ for Stochastic Submodular Cover. The proof technique used by Im et al. relies on a latency argument with geometrically increasing time intervals. They did not attempt to minimize the constant of 56 in their bound. While it is possible that their analysis could be tightened up to reduce the constant somewhat, a latency argument based approach is unlikely to yield optimal constant factors. (See the encyclopedia article on Min-Sum Set Cover by Im, 2016, for a related discussion).

Other bounds for Stochastic Submodular Cover have a dependence on quantities other than Q and η . In two related papers, Deshpande et al. (2016) and Hellerstein and Kletenik (2018) showed that Adaptive Greedy, with $\alpha = 1$, achieves a bound of $k(\ln(Q/\eta) + 1)$, where k is the number of states. Their analysis used an LP formulation of the problem. Recent independent work of Esfandiari et al. (2019) on related problems shows that Adaptive Greedy with $\alpha = 1$ achieves a bound close to $\ln(Qn/\eta)$. In particular, $c_{\text{greedy}} \leq (c_{\text{opt}} + 1)\ln(Qn/\eta) + 1$ where c_{greedy} is the expected cost of Adaptive Greedy, and c_{opt} is the expected cost of the optimal policy. The analyses for these bounds, and for the $56(\ln(Q/\eta) + 1)$ bound of Im et al., can easily be extended to $\alpha \geq 1$, leading to an extra multiplicative factor of α in the bounds.

In this paper, we prove that the Adaptive Greedy algorithm yields an $\alpha(\ln(Q/\eta) + 1)$ approximation bound for Stochastic Submodular Cover; the bound is $\alpha H(Q)$, where $H(Q)$ is the Q^{th} Harmonic number, when the utility function is integer valued. The tightness of the

bound follows from the fact that it generalizes the $(\ln m + 1)$ bound for classical Set Cover. Unless $P = NP$, the classical Set Cover problem cannot be approximated to within a factor of $(1 - o(1))(\ln m)$ (Dinur & Steurer, 2014; Feige, 1998).

Our bound restores the bound originally claimed by Golovin and Krause for Stochastic Submodular Cover. It also generalizes the bound claimed by Liu et al. (2008) for their restricted version of Stochastic Set Cover.

1.1 Related Results

In this paper, we use the definition of the Stochastic Submodular Cover problem previously used by Deshpande et al. (2016) and by Hellerstein and Kletenik (2018). This definition is slightly more general than the original definition of Stochastic Submodular Cover given by Golovin and Krause (2011). In Appendix A, we describe the differences in the definitions, and its relevance to the definition of the Weighted Stochastic Submodular Ranking problem studied by Im et al. (2016). The analyses of Golovin and Krause, and of Im et al., still hold for the definition of Stochastic Submodular Cover used in this paper.

In addition to defining Stochastic Submodular Cover, Golovin and Krause also defined a more general covering problem where item states are not necessarily independent (2011, 2017). In this problem, the function f , together with the joint distribution on item states, are assumed to satisfy the properties of “adaptive submodularity” and “adaptive monotonicity”. These properties are defined in terms of the *expected* value of f on random assignments of states to items. In their original paper, Golovin and Krause claimed the $\alpha(\ln(Q/\eta) + 1)$ bound for this more general problem, and then applied it to Stochastic Submodular Cover. They later proved their corrected bound of $\alpha(\ln(Q/\eta) + 1)^2$ for a version of their original problem with somewhat stronger assumptions, and again applied the bound to Stochastic Submodular Cover.

In the previous section, we stated a bound for Stochastic Submodular Cover that is linear in $\ln(Qn/\eta)$, implied by the work of Esfandiari et al. (2019). In fact, Esfandiari et al. proved this bound for the original version of the more general problem of Golovin and Krause, assuming adaptive submodularity and adaptive monotonicity.

Our analysis of Stochastic Submodular Cover relies heavily on the assumption of independent item states, as do the analyses that led to the previous $56(\ln(Q/\eta) + 1)$ and $k(\ln(Q/\eta) + 1)$ bounds. We leave open the question of whether an $\alpha(\ln(Q/\eta) + 1)$ bound can be achieved for the more general problem of Golovin and Krause.

In their work on Stochastic Submodular Cover, Deshpande et al. (2016) and Hellerstein and Kletenik (2018) actually showed that Adaptive Greedy with $\alpha = 1$ achieves a bound of $k(\ln(R/\eta) + 1)$. Here R is the largest increase in utility attainable from a single item, and thus is at most Q . This is a generalization of the $(\ln(R/\eta) + 1)$ bound of Wolsey for deterministic Submodular Cover. Although the bound we prove in this paper is tight in terms of Q/η , we leave open the question of whether the Q in our bound can be replaced by R . For some problems, R is much less than Q .

Hellerstein and Kletenik also proved an alternative bound of $(\ln(R_E/\eta_E) + 1)$. Here R_E is the largest expected marginal increase in utility from a single item, and η_E is the smallest expected non-zero marginal increase in utility from that same item. This latter bound is incomparable to the $(\ln(Q/\eta) + 1)$ bound shown in this paper, if items can have non-zero utility in one state, and zero utility in another.

To achieve the bound in this paper, we use an amortized approach. Liu et al. (2008) also used an amortized approach in their incorrect proof, but the approach here is different. A major difference is that Liu et al. only considered the Stochastic Set Cover problem, and their approach

relied on charging costs to elements of the ground set as they were covered. In our problem, Stochastic Submodular Cover, costs cannot be charged in this way, since there is no ground set, but only an abstract submodular utility function. This makes amortization more challenging.

An analysis due to Wan et al. (2010) for a deterministic generalization of Submodular Cover used an amortized approach with an abstract submodular utility function, but it did not need to address the challenges arising from the stochastic nature of our problem and the branching in the adaptive ordering.

We note that Deshpande et al. also introduced a dual greedy algorithm for Stochastic Submodular Cover, and proved that it achieves a bound that is qualitatively different and incomparable to the bounds given above. The bound generalizes the dual greedy approximation bound of Hochbaum (1982) for the classical Set Cover problem, and Fujito’s extension of it to Submodular Cover (cf. Fujito, 2000). These bounds all assume that the dual greedy algorithm can find the exactly optimal dual greedy choice at every step. It is not clear how to extend these results to approximately optimal dual greedy choices.

2. Preliminaries

For the aid of the reader, a table of frequently used notation is presented in Appendix C.

We now present a formal definition of the Stochastic Submodular Cover problem and introduce associated notation.

Let $E = \{e_1, \dots, e_n\}$ be a finite set of *items*. Let $O = \{o_1, \dots, o_k\}$ be a finite set of *states*. A *realization* is a function $\varphi: E \rightarrow O$, assigning state $\varphi(e)$ to item e . A *subrealization* is a function $\psi: E \rightarrow O \cup \{*\}$ which is a partial assignment of states to items (we assume $* \notin O$). If $\psi(e) \in O$ then $\psi(e)$ is the state assigned to e , while if $\psi(e) = *$, then the state of e is unassigned. We define

$$\text{dom}(\psi) \triangleq \{e \in E \mid \psi(e) \in O\}.$$

Realization φ has a corresponding relation

$$\text{rel}(\varphi) \triangleq \{(e, o) \in E \times O \mid \varphi(e) = o\}.$$

Subrealization ψ has a corresponding relation $\text{rel}(\psi) \triangleq \{(e, o) \in E \times O \mid e \in \text{dom}(\psi) \wedge \psi(e) = o\}$.

Given subrealizations ψ and ψ' , we write $\psi \subset \psi'$ to denote that $\text{rel}(\psi) \subset \text{rel}(\psi')$ and $\psi \subseteq \psi'$ to denote that $\text{rel}(\psi) \subseteq \text{rel}(\psi')$. When $\psi \subseteq \psi'$ we call ψ' an *extension* of ψ .

Given a subrealization ψ and a pair $(e, o) \in E \times O$ where $e \notin \text{dom}(\psi)$, we use $\psi \cup \{(e, o)\}$ to denote the extension ψ' of ψ such that $\text{rel}(\psi') = \text{rel}(\psi) \cup \{(e, o)\}$.

Let \mathbb{D} be a distribution whose domain is the set of all realizations $\varphi: E \rightarrow O$. Let Φ be a random variable distributed according to distribution \mathbb{D} . Thus Φ is a random assignment of states to items.

In what follows, all probabilities and expectations are with respect to distribution \mathbb{D} for Φ .

Assumption 1 (Independence). *We assume that in distribution \mathbb{D} , item states are independent, i.e., for all realizations $\varphi: E \rightarrow O$,*

$$P[\Phi = \varphi] = \prod_{e \in E} P[\Phi(e) = \varphi(e)].$$

Let $\mathcal{F} = \{f^\varphi \mid \varphi: E \rightarrow O\}$ be a family of utility functions $f^\varphi: 2^E \rightarrow \mathbb{R}_{\geq 0}$, whose elements f^φ have a one-to-one correspondence with the set of realizations φ . We assume that the functions f^φ in \mathcal{F} have the following two properties:

Assumption 2 (Pointwise Polymatroid). *Every function in \mathcal{F} is a polymatroid function.*

Assumption 3 (Sufficiency). *Given $E' \subseteq E$ and realizations φ_1, φ_2 such that $\varphi_1(e) = \varphi_2(e)$ for all $e \in E'$, we have $f^{\varphi_1}(E') = f^{\varphi_2}(E')$.*

Assumption 3 (Sufficiency) says that the value of f^φ on any subset of items $E' \subset E$ depends only on the states that φ assigns to items in E' , and not to the states it assigns to items outside of E' .

Define f to be the mapping from subrealizations ψ to $\mathbb{R}_{\geq 0}$ such that $f(\psi) = f^\varphi(\text{dom}(\psi))$, for all realizations φ such that $\psi \subseteq \varphi$. By the Sufficiency Assumption, f is well-defined.

Assumption 4 (Coverability). *There is a value $Q \in \mathbb{R}_{>0}$ such that for all realizations φ , $f(\varphi) = Q$.*

We call Q the *goal value* of f . We say that subrealization ψ is a *cover* of f if $f(\psi) = Q$.

Relative to ψ , the marginal increase in utility due to observing $e \in E \setminus \text{dom}(\psi)$ in state o , is

$$f_\psi(e, o) \triangleq f(\psi \cup \{(e, o)\}) - f(\psi)$$

We use $F_\psi(e)$ to denote the random marginal increase in utility from $e \in E \setminus \text{dom}(\psi)$, relative to ψ , that will result from observing the state $\Phi(e)$ of e . That is,

$$F_\psi(e) \triangleq f_\psi(e, \Phi(e))$$

An *adaptive covering policy* π for f is a function which maps every subrealization ψ such that $f(\psi) < Q$ to an item $\pi(\psi) \in E \setminus \text{dom}(\psi)$.²

Policy π specifies a sequential order in which to select the items of a cover for f . The order is adaptive, meaning that the choice of the next item depends on the states of the already selected items.

The selection of items proceeds in *iterations*. In each iteration, an item e is selected and its state is observed. Each item can only be selected once. As the iterations proceed, the set of selected items, and their states, can be represented by a subrealization ψ : $\text{dom}(\psi)$ is the set of selected items, and for each e in that set, $\psi(e)$ is its observed state. Initially ψ is such that $\text{rel}(\psi) = \emptyset$. Policy π specifies the item to select in each iteration, namely item $e = \pi(\psi)$. Selection of items, and observations of their states (with associated updating of ψ), continues until ψ is a cover for f . Executing policy π on realization φ means following this procedure when the observed state of each selected item e is $\varphi(e)$.

The adaptive sequential order specified by policy π can be represented by a decision tree $T(\pi)$, whose nodes are identified with subrealizations ψ . The root of the tree corresponds to the initial iteration, identified with realization ψ where $\text{rel}(\psi) = \emptyset$. The leaves are identified with the possible covers of f constructed by the policy. An example tree is shown in Figure 2.

Each non-leaf node ψ of the tree corresponds to an iteration where ψ represents the states of the items selected so far. The children of node ψ correspond to the possible updated values for ψ at the end of the iteration, after the state of selected item $\pi(\psi)$ is observed. Letting $e = \pi(\psi)$, for each $o \in O$ there is a child node identified with $\psi' = \psi \cup \{(e, o)\}$, if there is a non-zero probability of observing state o for item e under distribution \mathbb{D} (i.e., if $P[\Phi(e) = o] > 0$). In Figure 2, the edge from ψ to child ψ' is labeled with the selected item e , and with the associated increase in utility $f(\psi') - f(\psi)$.

2. We consider only deterministic policies, but our approximation bound for Adaptive Greedy also holds relative to a randomized policy r with minimum expected cost. Under the standard assumptions on the sources of randomness, executing r is equivalent to first choosing a policy π at random from a distribution D over deterministic policies, such that π is independent of Φ , and then executing π . Since r is optimal, every policy in the support of D has the same expected cost as r . Since our approximation bound holds with respect to each π in the support of D , it also holds with respect to r .

By the Independence Assumption, for subrealization ψ we have

$$\forall e \in E \setminus \text{dom}(\psi): \mathbf{E}[F_\psi(e) | \psi \subset \Phi] = \mathbf{E}[F_\psi(e)]$$

Let $L(\pi)$ be the set of leaf nodes of $T(\pi)$ and let $N(\pi)$ be the set of non-leaf nodes. If subrealization ψ is a non-leaf node then $f(\psi) < Q$, and if it is a leaf node then $f(\psi) = Q$.

The execution of π on random realization Φ corresponds to a random walk on $T(\pi)$ which starts at the root and proceeds downwards towards a random leaf of the tree. We regard the walk as visiting not just a sequence of nodes, but also the associated sequence of subrealizations. Formally, the associated sequence of subrealizations is $\psi_1, \psi_2, \dots, \psi_k$ for some $k > 0$ where $f(\psi_1) = \emptyset$, $f(\psi_k) = Q$, and for $i = 1, \dots, k-1$, $f(\psi_i) < Q$ and $\psi_{i+1} = \psi_i \cup \{(\pi(\psi_i), \Phi(\pi(\psi_i)))\}$. The final subrealization in this sequence, ψ_k , is identified with the leaf at the end of the path in $T(\pi)$ corresponding to the sequence. We refer to the final subrealization of the sequence as the *cover constructed by executing π on Φ* and define the random variable

$$\Psi(\pi) \triangleq \text{the cover constructed by executing } \pi \text{ on random realization } \Phi.$$




Figure 2: Tree representation of a policy π with $Q = 10$. Annotation (e, δ) on edge (ψ_i, ψ_j) indicates that e is selected at node ψ_i and $\delta = f(\psi_j) - f(\psi_i)$.

For $\psi \in N(\pi)$ and $e = \pi(\psi)$, $\mathbf{E}[F_\psi(e)]$ is the expected increase in utility at node ψ of $T(\pi)$, immediately prior to observing the state of e .

Consider the execution of π on random realization Φ . If $\psi \in N(\pi)$ is visited during this execution, then clearly $\psi \subset \Psi(\pi)$ and $\psi \subset \Phi$. If ψ is not visited, then there must be an ancestor ψ' of ψ such that the item $\pi(\psi)$ chosen at node ψ is assigned to different states by ψ and by $\Psi(\pi)$. Thus the following conditions are equivalent for $\psi \in N(\pi)$: (1) ψ is visited when executing π on Φ , (2) $\psi \subset \Phi$, and (3) $\psi \subset \Psi(\pi)$.

We define $S(\pi)$ to be the random set of items selected by π when executed on random realization Φ . That is,

$$S(\pi) \triangleq \text{dom}(\Psi(\pi))$$

We define

$$C(\pi) \triangleq \sum_{e \in S(\pi)} c(e)$$

Thus $C(\pi)$ is a random variable whose value is the sum of the costs incurred by executing π on Φ .

Formally, the inputs to the Stochastic Submodular Cover problem, consistent with the definitions and assumptions above, are as follows: (1) the set $E = \{e_1, \dots, e_n\}$ of items, (2) the value of $P[\Phi(e) = o]$ under distribution \mathbb{D} , for each $e \in E$ and $o \in O$, (3) a cost function $c: E \rightarrow \mathbb{R}_{>0}$, and (4) an oracle that takes as input any subrealization ψ , and outputs the value $f(\psi)$. The problem is to find an adaptive covering policy π for f with minimum expected cost, where the expectation is with respect to the distribution \mathbb{D} of Φ :

$$\min_{\pi} \mathbf{E}[C(\pi)]$$

The approximation bound we prove in this paper depends on the quantity η , which is the minimum gap between Q and $f(\psi)$ for any subrealization ψ where $f(\psi) < Q$. Formally:

$$\eta \triangleq \min\{Q - f(\psi) \mid \psi: E \rightarrow O \cup \{*\} \text{ and } f(\psi) < Q\}$$

When f is integer valued, $\eta \geq 1$.

In our analysis, we will sometimes use the following intuitive perspective. Think of the execution of an adaptive covering policy π on Φ in terms of driving a car along a road of length Q , with stops along the way. The start of the execution corresponds to the start of the drive. During the execution, when an item e is selected and its state is observed, there is a marginal increase δ in utility, and cost $c(e)$ is incurred. This corresponds to driving a further distance of δ along the road, at a cost of $c(e)$, before stopping again. At the end of execution, the sum of the increases in utility is Q , which corresponds to having traveled the entire length Q of the road.

3. Adaptive Greedy Policy

The Adaptive Greedy algorithm of Golovin and Krause (2011, 2017), when used to solve Stochastic Submodular Cover, executes the *adaptive greedy policy* defined as follows.

Define the *unit price* of an item $e \in E \setminus \text{dom}(\psi)$ with respect to subrealization ψ , denoted by $v(e|\psi)$, as follows:

$$v(e|\psi) \triangleq \frac{c(e)}{\mathbf{E}[F_{\psi}(e) \mid \psi \subset \Phi]} = \frac{c(e)}{\mathbf{E}[F_{\psi}(e)]} \tag{1}$$

The unit price is considered to be infinite if $\mathbf{E}[f_{\psi}(e, \Phi(e))] = 0$.

We denote the adaptive greedy policy by σ . The adaptive greedy policy with $\alpha = 1$ is the adaptive covering policy that, for subrealization ψ such that $f(\psi) < Q$, selects the item e with smallest unit price, so $\sigma(\psi) = \operatorname{argmin}_{e \in E \setminus \text{dom}(\psi)} v(e|\psi)$ with ties broken deterministically.

In some applications, though, finding the item with smallest unit price can be a computationally hard problem. We may instead solve the problem only approximately, with an approximation ratio α .

More generally, then, for fixed $\alpha \geq 1$, the adaptive greedy policy σ selects some element $\sigma(\psi)$ such that

$$v(\sigma(\psi)|\psi) \leq \alpha \left(\min_{e \in E \setminus \text{dom}(\psi)} v(e|\psi) \right) \tag{2}$$

We use e_{ψ} to denote the item $\sigma(\psi)$. We use v_{ψ} to denote the quantity $v(e_{\psi}|\psi)$. Thus

$$v_{\psi} \triangleq v(e_{\psi}|\psi) \leq \alpha \left(\min_{e \in E \setminus \text{dom}(\psi)} v(e|\psi) \right) \tag{3}$$

4. Comparison of π^* and σ

Let π^* be an optimal adaptive covering policy for an instance of Stochastic Submodular Cover. We will prove an approximation bound that compares $\mathbf{E}[C(\sigma)]$, the expected cost of the adaptive greedy policy, to $\mathbf{E}[C(\pi^*)]$, the expected cost of the optimal policy. To accomplish this, we introduce different ways of accounting for the costs incurred while executing policies π^* and σ on the same random realization Φ .

Given item e and policy π , let $\pi_{-1}[e]$ denote the *inverse image* of π under $\{e\}$: this is the set of all subrealizations corresponding to nodes in $N(\pi)$ at which item e is selected.

Consider the execution of π^* and σ on the same random realization Φ . The two executions give rise to values for random variables $S(\pi^*)$ (the items selected by policy π^*), $S(\sigma)$ (the items selected by policy σ), and $\Psi(\sigma)$ (the cover constructed by σ). We define the *revenue* collected by π^* and σ , in their executions on realization Φ , in terms of these random variables.

Revenue collection by π^*

Policy π^* collects its revenue from the items $e \in E$. For all $e \in E$, we define the revenue $\mu(e)$ that π^* collects from e as follows:

$$\mu(e) \triangleq \begin{cases} 0 & \text{if } e \notin S(\pi^*) \\ c(e) & \text{if } e \in S(\pi^*) \setminus S(\sigma) \\ v_\psi F_\psi(e) & \text{if } e \in S(\pi^*) \cap S(\sigma), \text{ where } \psi \text{ is the} \\ & \text{subrealization in } \sigma_{-1}[e] \text{ s.t. } \psi \subset \Psi(\sigma) \end{cases} \quad (4)$$

That is, if e was not selected by π^* , it does not contribute any revenue. If e was selected by π^* but not by σ , then its revenue equals $c(e)$. Otherwise, let ψ be the subrealization visited by σ where e was selected: the revenue in this case is defined to be the product of the unit price v_ψ and the marginal increase in utility obtained by σ when it observed e in state $\Phi(e)$.

The total revenue collected by π^* is

$$\mu^* \triangleq \sum_{e \in E} \mu(e) \quad (5)$$

Revenue collection by σ

Policy σ also collects revenue when executed on Φ . The revenue in this case is defined as follows. Let $\rho_1, \rho_2, \dots, \rho_{|N(\sigma)|}$ be all the subrealizations in $N(\sigma)$ indexed in non-decreasing order of their utilities so that $f(\rho_1) \leq f(\rho_2) \leq \dots \leq f(\rho_{|N(\sigma)|})$ (ties are broken arbitrarily). We treat each ρ_i as a *marker* located at point $f(\rho_i)$ on the interval $[0, Q]$ (corresponding to the road of length Q in our driving formulation). Let $\rho_{|N(\sigma)|+1}$ be an extra marker placed at point Q . Define $f(\rho_{|N(\sigma)|+1}) \triangleq Q$, and for $i \in \{1, \dots, |N(\sigma)|\}$ define $\epsilon_i \triangleq f(\rho_{i+1}) - f(\rho_i)$.

Observe that ϵ_i is the length of the i^{th} subinterval, which lies between ρ_i and ρ_{i+1} . Also, if ψ is the subrealization corresponding to a node in $T(\sigma)$, then $f(\psi)$ is on the boundary of one of these subintervals. We have

$$\sum_{i=1}^{|N(\sigma)|} \epsilon_i = Q \quad (6)$$

$$\epsilon_{|N(\sigma)|} \geq \eta \quad (7)$$

where (7) follows because $f(\rho_{|N(\sigma)|}) < Q$ and $f(\rho_{|N(\sigma)|+1}) = Q$.

Policy σ collects its revenue from the markers ρ_i . Viewing the execution of σ in terms of the drive along the road of length Q , revenue is collected from marker ρ_i when the car reaches a distance of $f(\rho_i)$ from the start.

We now describe the details of how σ collects revenue from ρ_i . Let R_i be the set of markers in the interval $[0, f(\rho_i)]$. Define:

$\psi \rightsquigarrow i \triangleq$ the event that, in the execution of σ on Φ , ψ was the last visited subrealization in R_i

That is, $\psi \rightsquigarrow i$ means that among all subrealizations that were visited when executing σ on Φ , ψ is the last one visited whose utility did not exceed $f(\rho_i)$.

Note that if $f(\rho_i) = f(\rho_{i+1})$, then R_i and R_{i+1} are equal and contain both ρ_i and ρ_{i+1} . For each realization Φ and $i \in \{1, \dots, |N(\sigma)|\}$ there is a unique $\psi \in R_i$ satisfying $\psi \rightsquigarrow i$.

Before completing the description of the revenue collection process for σ , we give an example to illustrate the above definitions.

Example: Suppose that σ is the policy π whose tree is shown in Figure 2. Here $Q = 10$, $f(\psi_1) = 0$, $f(\psi_2) = 3$, and $f(\psi_3) = f(\psi_4) = 1$. Breaking ties by increasing index, we have markers $\rho_1 = \psi_1$, $\rho_2 = \psi_3$, $\rho_3 = \psi_4$, and $\rho_4 = \psi_2$, and we place them at the points $f(\rho_1) = 0$, $f(\rho_2) = 1$, $f(\rho_3) = 1$, and $f(\rho_4) = 3$ along the interval $[0, 10]$. This is shown graphically in Figure 3, along with the associated values of the ϵ_i . The additional marker ρ_5 , which is not shown in the figure, is located at the right end of the interval, at $f(\rho_5) = 10$. Suppose that the execution of σ on Φ terminates with subrealization ψ_5 , so $\Psi(\sigma) = \psi_5$. In this case the subrealizations in $N(\sigma)$ visited during the execution are ψ_1 and ψ_2 ; there is an increase of 3 units of utility after observing the state of e_1 , followed by an increase of 7 units after observing the state of e_2 . Further, $\psi_1 \rightsquigarrow i$ for $i = 1, \dots, 3$, while $\psi_2 \rightsquigarrow i$ for $i = 4$. The marginal increases in utility for this case are shown at the bottom of Figure 3.

If, however, $\Psi(\sigma) = \psi_6$, then after ψ_1 , policy σ visits ψ_3 and ψ_4 before reaching a leaf. In this case, $\psi_1 \rightsquigarrow i$ for $i = 1$, and $\psi_4 \rightsquigarrow i$ for $i = 2, 3, 4$.

Policy σ collects a random amount of revenue λ_i from marker ρ_i as follows:

$$\lambda_i \triangleq v_\psi \epsilon_i, \text{ where } \psi \in R_i \text{ is such that } \psi \rightsquigarrow i \quad (8)$$

By the monotonicity property of the Pointwise Polymatroid Assumption, the event $\psi \rightsquigarrow i$ implies that there is no subrealization $\psi' \in N(\sigma)$ such that $\psi' \subset \Psi(\sigma)$ and $f(\psi') \in (f(\psi), f(\rho_i)]$. This implies the following fact.

Fact 1.

$$\forall \psi \in N(\sigma) \text{ such that } \psi \subset \Psi(\sigma): \sum_{i \text{ s.t. } \psi \rightsquigarrow i} \lambda_i = v_\psi \sum_{i \text{ s.t. } \psi \rightsquigarrow i} \epsilon_i = v_\psi F_\psi(e_\psi)$$




Figure 3: Revenue collection when σ is the policy π in Figure 2.

Hybrid policy

We also define a hybrid policy that combines σ and π^* in a careful manner. In particular, let ψ be a subrealization such that $\psi \in N(\sigma)$. We define the hybrid policy π^ψ whose execution on a random realization Φ consists of two stages. In the first stage, the execution of π^ψ on Φ follows the same sequence as σ until it either visits ψ or terminates. As soon as it visits ψ , π^ψ enters its second stage, during which it executes policy π^* , starting from the root of $T(\pi^*)$ (i.e., as if it has not yet selected any items). However, during the execution of π^* , π^ψ does not re-select an item e if it was already selected during the first stage. Instead, π^ψ uses the previously observed state of e (which is $\psi(e)$) and proceeds with the execution of π^* as if π^* had observed that state for e .³

Note that when π^ψ is executed on a random realization Φ such that $\psi \not\subset \Psi(\sigma)$ (i.e., σ does not visit ψ when executed on Φ), then π^ψ selects the same items that would have been selected by executing σ on Φ . When π^ψ is executed on a realization Φ such that $\psi \subset \Psi(\sigma)$, then the items it selects include all items that would have been selected by executing π^* on Φ . It follows that π^ψ is an adaptive covering policy.

Because σ selects item e_ψ at node ψ of $T(\sigma)$, it follows that $e_\psi \notin \text{dom}(\psi)$.

3. Technically, our description of π^ψ should also specify the value of $\pi^\psi(\psi')$ for subrealizations ψ' that do not appear as nodes in its decision tree. These values do not affect execution of π^ψ and can be assigned arbitrarily.

Further, by the definition of the hybrid policy, we have the following fact.

Fact 2.

$$\forall \psi \in N(\sigma) \text{ such that } \psi \subset \Psi(\sigma), \forall e \in E \setminus \text{dom}(\psi): e \in S(\pi^*) \text{ iff } e \in S(\pi^\psi) \quad (9)$$

We define $\delta_\psi^+(e)$ to be the marginal utility obtained by π^ψ when e is selected, if e is selected in the second stage; if e is not selected by π^ψ in the second stage, this value is 0.

Formally:

$$\delta_\psi^+(e) \triangleq \begin{cases} F_{\psi'}(e) & \text{if } \exists \psi' \in \pi_{-1}^\psi[e] \text{ s.t. } \psi \subseteq \psi', \psi' \subset \Psi(\pi^\psi) \\ 0 & \text{otherwise} \end{cases} \quad (10)$$

5. Approximation Analysis: Overview

Recall that η is equal to the minimum value of $Q - f(\psi)$ over all subrealizations ψ such that $f(\psi) < Q$. The goal of our approximation analysis is to prove the following theorem:

Theorem 1. *The adaptive greedy policy σ achieves an approximation bound of $\alpha\kappa$ for Stochastic Submodular Cover, where $\kappa = \ln(Q/\eta) + 1$ when the utility function f is real valued and $\kappa = H(Q)$ when the utility function is integer valued. That is, $\mathbf{E}[C(\sigma)] \leq \alpha\kappa \mathbf{E}[C(\pi^*)]$.*

We present the proof of the theorem in the following section. It is based on a sequence of lemmas. We first prove that the revenue collecting procedures defined above correctly account for the expected cost incurred by their associated algorithms. Specifically, in Lemma 1, we prove $\mathbf{E}[C(\sigma)] = \sum_{i=1}^{N(\sigma)} \mathbf{E}[\lambda_i]$, meaning that the expected cost of the greedy policy σ is equal to the expected revenue it collects from the markers λ_i . In Lemma 2 we prove $\mathbf{E}[C(\pi^*)] = \mathbf{E}[\mu^*]$, meaning that the expected cost of the optimal policy π^* is equal to the revenue it collects from the items e .

We then focus on comparing $\mathbf{E}[\mu^*]$ to $\sum_{i=1}^{N(\sigma)} \mathbf{E}[\lambda_i]$. To make this comparison, consider λ_i , the revenue collected by σ at marker ρ_i . By the definition of λ_i , its value depends on the sequence of subrealizations ψ visited by σ when executed on Φ . In particular, it depends on which ψ in that sequence satisfies $\psi \rightsquigarrow i$ (i.e., which is the last subrealization in the sequence that satisfies $f(\psi) \leq f(\rho_i)$). For each i , we partition the sample space according to which ψ satisfies $\psi \rightsquigarrow i$. In Lemmas 3 and 4, we prove bounds relating the quantities $\mathbf{E}[\lambda_i | \psi \rightsquigarrow i]$, $\mathbf{E}[\mu(e) | \psi \rightsquigarrow i]$, and $\mathbf{E}[\delta_\psi^+(e) | \psi \rightsquigarrow i]$.

Lemma 3 shows that for all items $e \notin \text{dom}(\psi)$, $\alpha \mathbf{E}[\mu(e) | \psi \rightsquigarrow i] \geq \mathbf{E}[\delta_\psi^+(e) | \psi \rightsquigarrow i] v_\psi$. Recall that $\delta_\psi^+(e)$ is the expected increase in utility from e in the second stage of π^ψ (after the switch to π^*), and v_ψ is the unit price paid by σ when it makes its greedy choice at node ψ of tree $T(\sigma)$.

Lemma 3 is used to prove Lemma 4, which states that $\forall i \in \{1, \dots, N(\sigma)\}$, $\mathbf{E}[\lambda_i | \psi \rightsquigarrow i] \leq \frac{\alpha \mathbf{E}[\mu^* | \psi \rightsquigarrow i]}{Q - f(\rho_i)} \epsilon_i$. For some intuition behind this statement, consider the execution of the greedy policy, and the corresponding drive along the road of length Q , conditional on $\psi \rightsquigarrow i$. After the car reaches marker ρ_i (has traveled a distance of $f(\rho_i)$), it travels an additional distance ϵ_i until it reaches the next marker ρ_{i+1} . If we view λ_i as the revenue paid to finance this travel, then the expected revenue collected by σ , per unit distance, in traveling from ρ_i to ρ_{i+1} , is $\mathbf{E}[\lambda_i | \psi \rightsquigarrow i] / \epsilon_i$. Consider this ratio to be the rate paid by σ in traveling from ρ_i to ρ_{i+1} . Then consider what happens in the drive corresponding to policy π^* when it reaches marker ρ_i . At marker ρ_i , the remaining distance to be traveled is $Q - f(\rho_i)$. Since the total revenue collected by π^* is μ^* , it will collect at most that much revenue in traveling the remaining distance $Q - f(\rho_i)$. Thus $\mathbf{E}[\mu^* | \psi \rightsquigarrow i]$ is an upper bound on the expected revenue collected by π^* while traveling the

remaining distance, and the ratio $\mathbf{E}[\mu^*|\psi \rightsquigarrow i]/(Q-f(\rho))$ is an upper bound on the rate paid by π^* during that travel. Lemma 4 says that this ratio (times α) is an upper bound on the rate paid by σ in traveling from marker ρ_i to ρ_{i+1} .

The next lemma, Lemma 5, states that $\forall i \in \{1, \dots, |N(\sigma)|\}$, $\mathbf{E}[\lambda_i] \leq \alpha \mathbf{E}[C(\pi^*)] \frac{\epsilon_i}{Q-f(\rho_i)}$. It is the same as Lemma 4, except that the conditioning on $\psi \rightsquigarrow i$ is removed, and we use the fact that $\mathbf{E}[\mu^*] = \mathbf{E}[C(\pi^*)]$. Having removed the conditioning, we no longer need to worry about the branching in the decision tree for σ , and can simply think of σ as deterministically traveling step by step from each marker ρ_i to the next one, at a cost of $\mathbf{E}[\lambda_i]$ for the step.

The rest of the analysis is done in a standard way and leads to the desired approximation bound for the adaptive greedy policy: $\mathbf{E}[C(\sigma)] \leq \alpha \mathbf{E}[C(\pi^*)](\ln(Q/\eta) + 1)$ (and $\mathbf{E}[C(\sigma)] \leq \alpha \mathbf{E}[C(\pi^*)]H(Q)$ when f is integer-valued).

6. Approximation Analysis

In this section, we will be considering the execution of the adaptive greedy policy σ , the optimal policy π^* , and the hybrid policy π^ψ on the same random realization Φ . To make our expressions more compact, we will use the following notation to express, for each of these policies, the event that a subrealization ψ' is visited:

- $\mathcal{V}^\sigma(\psi')$ denotes the event $\psi' \subset \Psi(\sigma)$
- $\mathcal{V}^*(\psi')$ denotes the event $\psi' \subset \Psi(\pi^*)$
- $\mathcal{V}^\psi(\psi')$ denotes the event $\psi' \subset \Psi(\pi^\psi)$

Consider running the greedy policy and the optimal policy on the same random realization Φ . We will use the following notation to express that item e is selected by the optimal policy but not the greedy policy, or that it is selected by both policies:

- $\mathcal{A}(e)$ denotes the event $e \in S(\pi^*) \setminus S(\sigma)$
- $\mathcal{B}(e)$ denotes the event $e \in S(\sigma) \cap S(\pi^*)$

The following lemma equates the expected cost of greedy policy σ to its expected revenue.

Lemma 1.

$$\mathbf{E}[C(\sigma)] = \sum_{i=1}^{|N(\sigma)|} \mathbf{E}[\lambda_i] \tag{11}$$

Proof. Recall that $e_\psi = \sigma(\psi)$. We have

$$\begin{aligned}
 \sum_{i=1}^{|\mathcal{N}(\sigma)|} \mathbf{E}[\lambda_i] &= \mathbf{E} \left[\sum_{\psi \in \mathcal{N}(\sigma) \text{ s.t. } \mathcal{V}^\sigma(\psi)} v_\psi \sum_{i \text{ s.t. } \psi \rightsquigarrow i} \epsilon_i \right] && \text{by the definition of } \lambda_i \text{ in (8)} \\
 &= \mathbf{E} \left[\sum_{\psi \in \mathcal{N}(\sigma) \text{ s.t. } \mathcal{V}^\sigma(\psi)} v_\psi F_\psi(e_\psi) \right] && \text{by Fact 1} \\
 &= \sum_{\psi \in \mathcal{N}(\sigma)} \Pr[\mathcal{V}^\sigma(\psi)] v_\psi \mathbf{E}[F_\psi(e_\psi) | \mathcal{V}^\sigma(\psi)] \\
 &= \sum_{\psi \in \mathcal{N}(\sigma)} \Pr[\mathcal{V}^\sigma(\psi)] v_\psi \mathbf{E}[F_\psi(e_\psi)] && \text{by the Independence Assumption} \\
 &= \sum_{\psi \in \mathcal{N}(\sigma)} \Pr[\mathcal{V}^\sigma(\psi)] c(e_\psi) && \text{by the definition of } v_\psi \text{ in (3)} \\
 &= \sum_{e \in E} \sum_{\psi \in \sigma_{-1}[e]} \Pr[\mathcal{V}^\sigma(\psi)] c(e) \\
 &= \sum_{e \in E} \Pr[e \in S(\sigma)] c(e) \\
 &= \mathbf{E}[C(\sigma)]. \quad \square
 \end{aligned}$$

The following lemma equates the expected cost of π^* to its expected revenue.

Lemma 2.

$$\mathbf{E}[C(\pi^*)] = \mathbf{E}[\mu^*]$$

Proof. Recalling that $\mathcal{A}(e)$ denotes the event $e \in S(\pi^*) \setminus S(\sigma)$, define

$$J \triangleq \left(\sum_{e \in E} \Pr[\mathcal{A}(e)] c(e) \right).$$

Recalling that $\mathcal{B}(e)$ denotes the event $e \in S(\pi^*) \cap S(\sigma)$, and using the definitions of μ and μ^* in (4) and (5), we have

$$\begin{aligned}
 \mathbf{E}[\mu^*] &= \sum_{e \in E} \mathbf{E}[\mu(e)] \\
 &= J + \left(\sum_{e \in E} \Pr[\mathcal{B}(e)] \sum_{\psi \in \sigma_{-1}[e]} \Pr[\mathcal{V}^\sigma(\psi) | \mathcal{B}(e)] v_\psi \mathbf{E}[F_\psi(e) | \mathcal{B}(e) \wedge \mathcal{V}^\sigma(\psi)] \right) \quad (12)
 \end{aligned}$$

To simplify the above expression, we first prove the following:

$$\forall e \in E, \psi \in \sigma_{-1}[e]: \mathbf{E}[F_\psi(e) | \mathcal{B}(e) \wedge \mathcal{V}^\sigma(\psi)] = \mathbf{E}[F_\psi(e)] \quad (13)$$

To prove (13), assume $e \in E$ and $\psi \in \sigma_{-1}[e]$. Since $\psi \in \sigma_{-1}[e]$, $\mathcal{V}^\sigma(\psi)$ implies $e \in S(\sigma)$. Therefore, the condition $\mathcal{B}(e) \wedge \mathcal{V}^\sigma(\psi)$ can be replaced by $e \in S(\pi^*) \wedge \mathcal{V}^\sigma(\psi)$. We can refine this

conditioning further by taking into account not just the fact that $e \in S(\pi^*)$, but by considering which subrealization ψ^* in π^* satisfies $\psi^* \in \pi_{-1}^*[e] \wedge \mathcal{V}^*(\psi^*)$. So:

$$\begin{aligned}
 & \mathbf{E}[F_\psi(e) | \mathcal{B}(e) \wedge \mathcal{V}^\sigma(\psi)] \\
 &= \mathbf{E}[F_\psi(e) | e \in S(\pi^*) \wedge \mathcal{V}^\sigma(\psi)] \\
 &= \sum_{\psi^* \in \pi_{-1}^*[e]} \Pr[\mathcal{V}^*(\psi^*) | e \in S(\pi^*) \wedge \mathcal{V}^\sigma(\psi)] \mathbf{E}[F_\psi(e) | \mathcal{V}^*(\psi^*) \wedge \mathcal{V}^\sigma(\psi)] \\
 &= \sum_{\psi^* \in \pi_{-1}^*[e]} \Pr[\mathcal{V}^*(\psi^*) | e \in S(\pi^*) \wedge \mathcal{V}^\sigma(\psi)] \mathbf{E}[F_\psi(e)] \\
 &= \mathbf{E}[F_\psi(e)].
 \end{aligned} \tag{14}$$

$$= \mathbf{E}[F_\psi(e)]. \tag{15}$$

Line (14) holds due to the fact that $\psi \in \sigma_{-1}[e]$ by assumption, and $\psi^* \in \pi_{-1}^*[e]$, so e is not an element of either $\text{dom}(\psi)$ or $\text{dom}(\psi^*)$. Whether σ visits node ψ of $T(\sigma)$ depends only on the states of the items in $\text{dom}(\psi)$, and whether π^* visits node ψ^* of $T(\pi^*)$ depends only on the states of the items in $\text{dom}(\psi^*)$. Thus by the Independence Assumption, the value of $\Phi(e)$, and hence the value of $\mathbf{E}[F_\psi(e)]$, is independent of the conditions $\mathcal{V}^\sigma(\psi)$ and $\mathcal{V}^*(\psi^*)$. Line (15) holds because if π^* selects e when executed on Φ , it visits exactly one realization $\psi^* \in \pi_{-1}^*[e]$. Hence $\sum_{\psi^* \in \pi_{-1}^*[e]} \Pr[\mathcal{V}^*(\psi^*) | e \in S(\pi^*) \wedge \mathcal{V}^\sigma(\psi)] = 1$. This completes the proof of (13).

We can now use (13) to rewrite the expression for $\mathbf{E}[\mu^*]$ in (12), yielding

$$\begin{aligned}
 \mathbf{E}[\mu^*] &= J + \sum_{e \in E} \Pr[\mathcal{B}(e)] \sum_{\psi \in \sigma_{-1}[e]} \Pr[\mathcal{V}^\sigma(\psi) | \mathcal{B}(e)] v_\psi \mathbf{E}[F_\psi(e)] \\
 &= J + \sum_{e \in E} \Pr[\mathcal{B}(e)] \sum_{\psi \in \sigma_{-1}[e]} \Pr[\mathcal{V}^\sigma(\psi) | \mathcal{B}(e)] c(e) && \text{by the definition of } v_\psi \\
 &= J + \left(\sum_{e \in E} \Pr[\mathcal{B}(e)] c(e) \right) && \text{because } \sum_{\psi \in \sigma_{-1}[e]} \Pr[\mathcal{V}^\sigma(\psi) | \mathcal{B}(e)] = 1 \\
 &= \left(\sum_{e \in E} \Pr[e \in S(\pi^*) \setminus S(\sigma)] c(e) \right) + \left(\sum_{e \in E} \Pr[e \in S(\pi^*) \cap S(\sigma)] c(e) \right) \\
 &= \sum_{e \in E} \Pr[e \in S(\pi^*)] c(e) \\
 &= \mathbf{E}[C(\pi^*)]. \quad \square
 \end{aligned}$$

Recall the description of the hybrid policy π^ψ , where ψ is a non-leaf node of $T(\sigma)$. At node ψ , σ selects item e_ψ and $e_\psi \notin \text{dom}(\psi)$. The tree $T(\pi^\psi)$ for the hybrid policy also contains node ψ , but the item selected at that node may be different from e_ψ (since the hybrid policy switches to π^* at this point).

With that background, we are ready to prove the following technical lemma.

Lemma 3. *Let $\psi \in N(\sigma)$ and $i \in \{1, \dots, |N(\sigma)|\}$ such that $\Pr[\psi \rightsquigarrow i] > 0$. Then*

$$\forall e \in E \setminus \text{dom}(\psi): \alpha \mathbf{E}[\mu(e) | \psi \rightsquigarrow i] \geq \mathbf{E}[\delta_\psi^+(e) | \psi \rightsquigarrow i] v_\psi \tag{16}$$

Proof. Let $e \in E \setminus \text{dom}(\psi)$ and let ψ and i be as given in the statement of the lemma.

Define $H \triangleq \{\psi_h \mid \psi_h \in \pi_{-1}^\psi[e] \wedge \psi \subseteq \psi_h\}$. Thus H is the set of nodes of $T(\pi^\psi)$ where e could be selected during the second stage of π^ψ .

We divide our analysis according to whether e is selected by both σ and π^* when they are executed on Φ , or just by π^* . We first prove the inequality in the statement of the lemma with respect to the additional condition $\mathcal{A}(e)$ (e is selected by π^* but not by σ), and then with respect to the additional condition $\mathcal{B}(e)$ (e is selected by both σ and π^*). We do not need to consider conditions where e is not selected by π^* , because $\mu(e)$ and $\delta_\psi^+(e)$ are equal to 0 in these cases.

- Additional condition $\mathcal{A}(e)$

We will show

$$\alpha \mathbf{E}[\mu(e) \mid \mathcal{A}(e) \wedge \psi \rightsquigarrow i] \geq v_\psi \mathbf{E}[\delta_\psi^+(e) \mid \mathcal{A}(e) \wedge \psi \rightsquigarrow i] \quad (17)$$

Recall that $\mathcal{A}(e)$ denotes the event that e is in $S(\pi^*)$ but not in $S(\sigma)$. The event $\psi \rightsquigarrow i$ implies that $\psi \subseteq \Psi(\sigma)$. Thus by Fact 2, $\mathcal{A}(e) \wedge \psi \rightsquigarrow i$ implies that $e \in S(\pi^\psi)$. Further, since $e \notin \text{dom}(\psi)$, e must be selected in the second stage of π^ψ .

Let $X \subseteq L(\sigma)$ denote the set of terminal subrealizations observable (covers constructed) by σ given that $\mathcal{A}(e) \wedge \psi \rightsquigarrow i$. Thus X is a subset of the possible values of $\Psi(\sigma)$. We have

$$\begin{aligned} & \alpha \mathbf{E}[\mu(e) \mid \mathcal{A}(e) \wedge \psi \rightsquigarrow i] \\ &= \sum_{\psi_h \in \pi_{-1}^\psi[e]} \sum_{\psi_x \in X} \Pr[\mathcal{V}^\psi(\psi_h) \wedge \Psi(\sigma) = \psi_x \mid \mathcal{A}(e) \wedge \psi \rightsquigarrow i] \alpha c(e) \\ & \quad \text{since } \mu(e) = c(e) \text{ when } \mathcal{A}(e) \text{ holds, and partitioning event } (\mathcal{A}(e) \wedge \psi \rightsquigarrow i) \text{ w.r.t. which} \\ & \quad \psi_h \in \pi_{-1}^\psi[e] \text{ is visited by } \pi^\psi \text{ and which cover } \psi_x \in X \text{ is constructed by } \sigma \\ & \geq \sum_{\psi_h \in \pi_{-1}^\psi[e]} \sum_{\psi_x \in X} \Pr[\mathcal{V}^\psi(\psi_h) \wedge \Psi(\sigma) = \psi_x \mid \mathcal{A}(e) \wedge \psi \rightsquigarrow i] v_\psi \mathbf{E}[F_\psi(e)] \quad \text{by the greedy choice (3)} \\ &= \sum_{\psi_h \in \pi_{-1}^\psi[e]} \sum_{\psi_x \in X} \Pr[\mathcal{V}^\psi(\psi_h) \wedge \Psi(\sigma) = \psi_x \mid \mathcal{A}(e) \wedge \psi \rightsquigarrow i] v_\psi \mathbf{E}[F_\psi(e) \mid \mathcal{V}^\psi(\psi_h) \wedge \Psi(\sigma) = \psi_x] \\ & \quad \text{by the Independence Assumption since } e \notin \text{dom}(\psi_h) \cup \text{dom}(\psi_x) \\ & \geq \sum_{\psi_h \in \pi_{-1}^\psi[e]} \sum_{\psi_x \in X} \Pr[\mathcal{V}^\psi(\psi_h) \wedge \Psi(\sigma) = \psi_x \mid \mathcal{A}(e) \wedge \psi \rightsquigarrow i] v_\psi \mathbf{E}[F_{\psi_h}(e) \mid \mathcal{V}^\psi(\psi_h) \wedge \Psi(\sigma) = \psi_x] \\ & \quad \text{replacing } F_\psi \text{ by } F_{\psi_h}, \text{ by submodularity since } \psi \subseteq \psi_h \\ &= v_\psi \sum_{\psi_h \in \pi_{-1}^\psi[e]} \sum_{\psi_x \in X} \Pr[\mathcal{V}^\psi(\psi_h) \wedge \Psi(\sigma) = \psi_x \mid \mathcal{A}(e) \wedge \psi \rightsquigarrow i] \mathbf{E}[\delta_\psi^+(e) \mid \mathcal{V}^\psi(\psi_h) \wedge \Psi(\sigma) = \psi_x] \\ & \quad \text{because } \mathcal{V}^\psi(\psi_h) \text{ implies } e \text{ was selected by } \pi^\psi \text{ at node } \psi_h \text{ during its second stage} \\ &= v_\psi \mathbf{E}[\delta_\psi^+(e) \mid \mathcal{A}(e) \wedge \psi \rightsquigarrow i]. \end{aligned}$$

- Additional condition $\mathcal{B}(e)$

Because $e \notin \text{dom}(\psi)$, condition $\mathcal{B}(e) \wedge \psi \rightsquigarrow e$ implies, by Fact 2, that π^ψ visits some $\psi_h \in H$.

We will show

$$\alpha \mathbf{E}[\mu(e) \mid \mathcal{B}(e) \wedge \psi \rightsquigarrow i] \geq v_\psi \mathbf{E}[\delta_\psi^+(e) \mid \mathcal{B}(e) \wedge \psi \rightsquigarrow i] \quad (18)$$

We first give the proof of (18) for the case where $e = e_\psi$. If $e = e_\psi$, we have

$$\begin{aligned} & \mathbf{E}[\mu(e_\psi)|\psi \rightsquigarrow i] \\ &= \Pr[e_\psi \in S(\pi^*)|\psi \rightsquigarrow i] \mathbf{E}[\mu(e_\psi)|\psi \rightsquigarrow i \wedge e_\psi \in S(\pi^*)] && \text{by the definition of } \mu \\ &= \Pr[e_\psi \in S(\pi^\psi)|\psi \rightsquigarrow i] \mathbf{E}[\mu(e_\psi)|\psi \rightsquigarrow i \wedge e_\psi \in S(\pi^\psi)] && \text{by Fact 2} \\ &= \sum_{\psi_h \in H} \Pr[\mathcal{V}^\psi(\psi_h)|\psi \rightsquigarrow i] \mathbf{E}[\mu(e_\psi)|\psi \rightsquigarrow i \wedge \mathcal{V}^\psi(\psi_h)] && \text{partitioning w.r.t. which } \psi_h \in H \text{ is visited} \\ &= \sum_{\psi_h \in H} \Pr[\mathcal{V}^\psi(\psi_h)|\psi \rightsquigarrow i] v_\psi \mathbf{E}[F_{\psi_h}(e_\psi)|\psi \rightsquigarrow i \wedge \mathcal{V}^\psi(\psi_h)] && \text{by the definition of } \mu \text{ since } \psi \rightsquigarrow i \Rightarrow \mathcal{V}^\sigma(\psi) \\ &\geq \sum_{\psi_h \in H} \Pr[\mathcal{V}^\psi(\psi_h)|\psi \rightsquigarrow i] v_\psi \mathbf{E}[F_{\psi_h}(e_\psi)|\psi \rightsquigarrow i \wedge \mathcal{V}^\psi(\psi_h)] && \text{replacing } f_\psi \text{ by } f_{\psi_h}, \text{ by submodularity} \\ &= v_\psi \mathbf{E}[\delta_\psi^+(e_\psi)|\psi \rightsquigarrow i] && \text{by definition of } \delta_\psi^+. \end{aligned}$$

Thus (18) holds when $e = e_\psi$.

So suppose $e \neq e_\psi$. In this case, define G to be the set of subrealizations ψ_g satisfying the following three properties:

$$\begin{aligned} \psi_g &\in \sigma_{-1}[e] \\ \psi_g &\supset \psi \\ f(\psi \cup (e_\psi, \psi_g(e_\psi))) &> f(\rho_i) \end{aligned}$$

The second property says that in the greedy tree, node ψ is an ancestor of node ψ_g . The third property concerns the child of node ψ that is on the path from ψ down to ψ_g : this is the node identified with subrealization $\psi \cup (e_\psi, \psi_g(e_\psi))$. It states that the value of f on this subrealization is strictly greater than $f(\rho_i)$. Therefore,

$$\forall \psi_g \in G: \mathcal{V}^\sigma(\psi_g) \Rightarrow \psi \rightsquigarrow i \quad (19)$$

Conversely, if $\psi \rightsquigarrow i$, then because $e \notin \text{dom}(\psi)$, condition $\mathcal{B}(e)$ implies that σ must visit some $\psi_g \in G$. As before, condition $\mathcal{B}(e) \wedge \psi \rightsquigarrow i$ implies that π^ψ must visit some $\psi_h \in H$. Therefore,

$$\mathcal{B}(e) \wedge \psi \rightsquigarrow i \Rightarrow \exists! \psi_g \in G, \exists! \psi_h \in H \text{ such that } \mathcal{V}^\sigma(\psi_g) \wedge \mathcal{V}^\psi(\psi_h)$$

where $\exists!$ means the quantified element exists and is unique.

Note that by the Independence Assumption, for $\psi_g \in G$ and $\psi_h \in H$, since $e \notin \text{dom}(\psi_g) \cup \text{dom}(\psi_h)$,

$$\mathbf{E}[F_{\psi_g}(e)|\mathcal{V}^\sigma(\psi_g) \wedge \mathcal{V}^\psi(\psi_h)] = \mathbf{E}[F_{\psi_g}(e)] \quad (20)$$

Before proceeding with the proof of (18), we show that the following equation holds for $\psi_g \in G$:

$$\mathbf{E}[\mu(e)|\mathcal{B}(e) \wedge \mathcal{V}^\sigma(\psi_g)] = c(e) \quad (21)$$

The proof of (21) is as follows. Recalling that $e \notin \text{dom}(\psi)$,

$$\begin{aligned} & \mathbf{E}[\mu(e)|\mathcal{B}(e) \wedge \mathcal{V}^\sigma(\psi_g)] \\ &= v_{\psi_g} \mathbf{E}[F_{\psi_g}(e)|\mathcal{B}(e) \wedge \mathcal{V}^\sigma(\psi_g)] && \text{by the definition of } \mu(e), \text{ since } \psi_g \in \sigma_{-1}[e] \\ &= \sum_{\psi_h \in H} P[\mathcal{V}^\psi(\psi_h)|\mathcal{B}(e) \wedge \mathcal{V}^\sigma(\psi_g)] v_{\psi_g} \mathbf{E}[F_{\psi_g}(e)|\mathcal{V}^\sigma(\psi_g) \wedge \mathcal{V}^\psi(\psi_h)] \\ & && \text{partitioning w.r.t. which } \psi_h \in H \text{ is visited by } \pi^\psi \\ &= \sum_{\psi_h \in H} P[\mathcal{V}^\psi(\psi_h)|\mathcal{B}(e) \wedge \mathcal{V}^\sigma(\psi_g)] v_{\psi_g} \mathbf{E}[F_{\psi_g}(e_{\psi_g})] \\ & && \text{by (20) and because } \psi_g \in \sigma_{-1}[e], \text{ so } e = e_{\psi_g} \\ &= c(e) \end{aligned}$$

because $\sum_{\psi_h \in H} P[\mathcal{V}^\psi(\psi_h) | \mathcal{B}(e) \wedge \mathcal{V}^\sigma(\psi_g)] = 1$ and $v_{\psi_g} = c(e_{\psi_g}) / \mathbf{E}[F_{\psi_g}(e_{\psi_g})]$.

Having proved (21), we now prove (18), the desired inequality for condition $\mathcal{B}(e)$. For $\psi_g \in G$, define

$$K_g \triangleq P[\mathcal{V}^\sigma(\psi_g) | \mathcal{B}(e) \wedge \psi \rightsquigarrow i]$$

We have

$$\begin{aligned} & \alpha \mathbf{E}[\mu(e) | \mathcal{B}(e) \wedge \psi \rightsquigarrow i] \\ &= \alpha \sum_{\psi_g \in G} P[\mathcal{V}^\sigma(\psi_g) | \mathcal{B}(e) \wedge \psi \rightsquigarrow i] \mathbf{E}[\mu(e) | \mathcal{B}(e) \wedge \mathcal{V}^\sigma(\psi_g)] \\ & \quad \text{partitioning w.r.t. which } \psi_g \in G \text{ is visited, and because } \mathcal{V}^\sigma(\psi_g) \Rightarrow \psi \rightsquigarrow i \\ &= \alpha \sum_{\psi_g \in G} K_g c(e) \quad \text{by (21) and the definition of } K_g \\ &= \alpha \sum_{\psi_g \in G} K_g \sum_{\psi_h \in H} P[\mathcal{V}^\psi(\psi_h) | \mathcal{B}(e) \wedge \mathcal{V}^\sigma(\psi_g)] v(e | \psi_h) \mathbf{E}[F_{\psi_h}(e)] \\ & \quad \text{partitioning w.r.t. which } \psi_h \in H \text{ is visited, and because } v(e | \psi_h) = c(e) / \mathbf{E}[F_{\psi_h}(e)] \\ &= \alpha \sum_{\psi_g \in G} K_g \sum_{\psi_h \in H} P[\mathcal{V}^\psi(\psi_h) | \mathcal{B}(e) \wedge \mathcal{V}^\sigma(\psi_g)] v(e | \psi_h) \mathbf{E}[F_{\psi_h}(e) | \mathcal{V}^\sigma(\psi_g) \wedge \mathcal{V}^\psi(\psi_h)] \\ & \quad \text{by the Independence Assumption} \\ &= \alpha \sum_{\psi_g \in G} K_g \sum_{\psi_h \in H} P[\mathcal{V}^\psi(\psi_h) | \mathcal{B}(e) \wedge \mathcal{V}^\sigma(\psi_g)] v(e | \psi_h) \mathbf{E}[\delta_\psi^+(e) | \mathcal{V}^\sigma(\psi_g) \wedge \mathcal{V}^\psi(\psi_h)] \\ & \quad \text{by the definition of } \delta_\psi^+(e), \text{ since } \mathcal{V}^\psi(\psi_h) \text{ implies } e \text{ was selected by} \\ & \quad \quad \pi^\psi \text{ at node } \psi_h \text{ during its second stage} \\ &\geq \alpha v(e | \psi) \sum_{\psi_g \in G} K_g \sum_{\psi_h \in H} P[\mathcal{V}^\psi(\psi_h) | \mathcal{B}(e) \wedge \mathcal{V}^\sigma(\psi_g)] \mathbf{E}[\delta_\psi^+(e) | \mathcal{V}^\sigma(\psi_g) \wedge \mathcal{V}^\psi(\psi_h)] \\ & \quad \text{since } \psi \subseteq \psi_h \text{ so by the submodularity property, } F_\psi(e) \geq F_{\psi_h}(e) \\ & \quad \quad \text{and hence } v(e | \psi_h) \geq v(e | \psi) \\ &\geq v_\psi \sum_{\psi_g \in G} K_g \sum_{\psi_h \in H} P[\mathcal{V}^\psi(\psi_h) | \mathcal{B}(e) \wedge \mathcal{V}^\sigma(\psi_g)] \mathbf{E}[\delta_\psi^+(e) | \mathcal{V}^\sigma(\psi_g) \wedge \mathcal{V}^\psi(\psi_h)] \\ & \quad \text{since } v_\psi = v(e_\psi | \psi) \leq \alpha v(e | \psi) \text{ by the greedy choice (3)} \\ &= v_\psi \sum_{\psi_g \in G} K_g \mathbf{E}[\delta_\psi^+(e) | \mathcal{B}(e) \wedge \mathcal{V}^\sigma(\psi_g)] \\ &= v_\psi \sum_{\psi_g \in G} P[\mathcal{V}^\sigma(\psi_g) | \mathcal{B}(e) \wedge \psi \rightsquigarrow i] \mathbf{E}[\delta_\psi^+(e) | \mathcal{B}(e) \wedge \mathcal{V}^\sigma(\psi_g)] \\ & \quad \text{by the definition of } K_g \\ &= v_\psi \mathbf{E}[\delta_\psi^+(e) | \mathcal{B}(e) \wedge \psi \rightsquigarrow i]. \end{aligned}$$

Finally, because $\mu(e) = 0$ and $\delta_\psi^+(e) = 0$ whenever $e \notin S(\pi^*)$, the lemma follows from the inequalities (18) and (17) proved for the two conditions, $\mathcal{A}(e)$ and $\mathcal{B}(e)$. \square

Fact 3. For non-negative numbers $\alpha_1, \dots, \alpha_n$, and β_1, \dots, β_n such that $\sum_{i=1}^n \beta_i > 0$:

$$\min_{i \text{ s.t. } \beta_i > 0} \frac{\alpha_i}{\beta_i} \leq \frac{\sum_{i=1}^n \alpha_i}{\sum_{i=1}^n \beta_i} \quad (22)$$

We combine the above fact with Lemma 3 to prove the following lemma.

Lemma 4.

$$\forall i \in \{1, \dots, |N(\sigma)|\}: \mathbf{E}[\lambda_i | \psi \rightsquigarrow i] \leq \frac{\alpha \mathbf{E}[\mu^* | \psi \rightsquigarrow i]}{Q - f(\rho_i)} \epsilon_i \quad (23)$$

Proof. Let $E_1 = \{e | e \in E \setminus \text{dom}(\psi) \wedge \mathbf{E}[\delta_\psi^+(e) | \psi \rightsquigarrow i] > 0\}$.

$$\begin{aligned} \mathbf{E}[\lambda_i | \psi \rightsquigarrow i] &= v_\psi \epsilon_i \leq \min_{e \in E_1} \frac{\alpha \mathbf{E}[\mu(e) | \psi \rightsquigarrow i]}{\mathbf{E}[\delta_\psi^+(e) | \psi \rightsquigarrow i]} \epsilon_i && \text{by Lemma 3} \\ &\leq \frac{\sum_{e \in E \setminus \text{dom}(\psi)} \alpha \mathbf{E}[\mu(e) | \psi \rightsquigarrow i]}{\sum_{e \in E \setminus \text{dom}(\psi)} \mathbf{E}[\delta_\psi^+(e) | \psi \rightsquigarrow i]} \epsilon_i && \text{by Fact 3} \\ &\leq \frac{\alpha \mathbf{E}[\mu^* | \psi \rightsquigarrow i]}{Q - f(\psi)} \epsilon_i && \text{by definitions of } \mu^* \text{ and } \delta_\psi^+ \quad \square \end{aligned}$$

Lemma 5.

$$\forall i \in \{1, \dots, |N(\sigma)|\}: \mathbf{E}[\lambda_i] \leq \alpha \mathbf{E}[C(\pi^*)] \frac{\epsilon_i}{Q - f(\rho_i)} \quad (24)$$

Proof. Fix $i \in \{1, \dots, |N(\sigma)|\}$. Let $R = \{\psi | \psi \in N(\sigma) \wedge \Pr[\psi \rightsquigarrow i] > 0\}$. The events $\{\psi \rightsquigarrow i | \psi \in R\}$ partition the sample space. Hence:

$$\begin{aligned} \mathbf{E}[\lambda_i] &= \sum_{\psi \in R} \Pr[\psi \rightsquigarrow i] \mathbf{E}[\lambda_i | \psi \rightsquigarrow i] \\ &\leq \sum_{\psi \in R} \frac{\Pr[\psi \rightsquigarrow i] \alpha \mathbf{E}[\mu^* | \psi \rightsquigarrow i] \epsilon_i}{Q - f(\rho_i)} && \text{by Lemma 4} \\ &= \frac{\alpha \mathbf{E}[\mu^*] \epsilon_i}{Q - f(\rho_i)} \\ &= \frac{\alpha \mathbf{E}[C(\pi^*)] \epsilon_i}{Q - f(\rho_i)} && \text{by Lemma 2} \quad \square \end{aligned}$$

Lemma 6.

$$\mathbf{E}[C(\sigma)] \leq \alpha \mathbf{E}[C(\pi^*)] \sum_{i=1}^{|N(\sigma)|} \frac{\epsilon_i}{Q - f(\rho_i)} \quad (25)$$

Proof. The lemma follows by combining Lemmas 1 and 5. \square

With the above lemmas, we can now prove Theorem 1, which states that σ achieves an approximation bound of $\alpha \kappa$ for Stochastic Submodular Cover, where $\kappa = \ln(Q/\eta) + 1$ when the utility function is real valued and $\kappa = H(Q)$ when the utility function is integer valued.

Proof of Theorem 1: The theorem follows from Lemma 6 by bounding the summation

$$\sum_{i=1}^{|\mathcal{N}(\sigma)|} \frac{\epsilon_i}{Q - f(\rho_i)}$$

using an argument of Azar and Gamzu (2011), which we include for completeness.

Let t denote $|\mathcal{N}(\sigma)|$. Recall that $\epsilon_i = f(\rho_{i+1}) - f(\rho_i)$. For $i \in \{1, \dots, t-1\}$,

$$\frac{\epsilon_i}{Q - f(\rho_i)} = \int_{f(\rho_i)}^{f(\rho_{i+1})} \frac{1}{Q - f(\rho_i)} dx \leq \int_{f(\rho_i)}^{f(\rho_{i+1})} \frac{1}{Q - x} dx.$$

Therefore,

$$\begin{aligned} \sum_{i=1}^t \frac{\epsilon_i}{Q - f(\rho_i)} &\leq \frac{\epsilon_t}{Q - f(\rho_t)} + \int_0^{f(\rho_t)} \frac{1}{Q - x} dx \\ &= 1 + \ln Q/\epsilon_t \quad \text{because } \epsilon_t = Q - f(\rho_t) \\ &\leq 1 + \ln Q/\eta \quad \text{by (7)} \end{aligned} \tag{26}$$

If f is integer-valued, then $\epsilon_t \geq 1$ and similarly,

$$\begin{aligned} \sum_{i=1}^t \frac{\epsilon_i}{Q - f(\rho_i)} &\leq 1 + \sum_{j=0}^{f(\rho_t)-1} \frac{1}{Q - j} \\ &= 1 + H(Q) - H(\epsilon_t) \\ &\leq H(Q). \end{aligned} \tag{27}$$

□

Acknowledgments

Partial support for this work came from NSF Award IIS-1909335 (for L. Hellerstein) and from a PSC-CUNY Award, jointly funded by The Professional Staff Congress and The City University of New York (for D. Kletenik). We thank the anonymous referees for their comprehensive and useful comments. We thank Naifeng Liu for proofreading help.

Appendix A. The Definition of the Stochastic Submodular Cover Problem

In the definition of Stochastic Submodular Cover given by Golovin and Krause (who called the problem Stochastic Submodular *Coverage*), the input utility function is $\hat{f}: 2^{E \times O} \rightarrow \mathbb{R}_{\geq 0}$ where E is the set of items and O is the set of states, such that \hat{f} is a monotone and submodular set function. Since \hat{f} is defined on all subsets of $E \times O$, its domain includes “inconsistent” subsets containing pairs (e, o) and (e, o') , where $o \neq o'$. Intuitively, this means that the input utility function must assign values to collections of items in which a single item can simultaneously be in more than one state; further, these values must be compatible with the properties of monotonicity and submodularity. In contrast, our input utility function need only assign values to collections of items where each item has exactly one state. In fact, though, the value of \hat{f} on the “inconsistent” subsets is irrelevant both to the remainder of the Golovin and Krause problem definition, and to the proof of their bound for Stochastic Submodular Cover. While in their definition they require \hat{f} to have values on these subsets, they effectively ignore these values.

The definition of Stochastic Submodular Cover given by Golovin and Krause, with item set $E = \{e_1, \dots, e_n\}$ and state set O , is equivalent to a special case of the Weighted Stochastic

Submodular Ranking (WSSR) problem, studied by Im et al. (2016). In particular, it is equivalent to the special case of the WSSR problem where there is a single function $f_1 : 2^D \rightarrow [0,1]$ with $D = E \times O$, and such that for each X_i , $P[X_i = (e,o)] = 0$ if $e \neq e_i$, otherwise $P[X_i = (e,o)]$ equals the probability $P[\Phi(e_i) = o]$ in Stochastic Submodular Cover. Although the problem definition for the WSSR problem requires that f_1 be defined on all subsets of $D = E \times O$ in this case, including the “incompatible” subsets, the analysis of Im et al. applied to this case makes no use of the value of f_1 on such subsets.

Appendix B. The Error in the Analysis of Liu et al., 2008

The paper of Liu et al. claims to show an $(\ln m + 1)$ approximation bound for the adaptive greedy policy, applied to Stochastic Set Cover. Consider the execution of this policy, which always selects the stochastic subset covering the largest expected number of items, per unit cost. Let S_{i_1}, \dots, S_{i_k} denote the sequence of chosen subsets. Each S_{i_ℓ} covers some number h of ground elements in S_{i_ℓ} not already covered by $S_{i_1}, \dots, S_{i_{\ell-1}}$. Consider the h elements to be covered one by one by S_{i_ℓ} . The *price* paid to cover each one is $\frac{c_{i_\ell}}{h}$.

Consider a run of the greedy policy, and of an optimal policy OPT, on a random assignment of values to the S_i . Order the elements of the ground set in the order they are covered by the greedy policy. Lemma 3 of Liu et al. gives an upper bound on $price(j)$, the price paid when the j th item in this order is covered. The bound is conditioned on ψ and Υ , where ψ is the sequence of S_i chosen by the greedy policy (together with the observed values of those S_i), and Υ is the subset of S_i chosen by both the greedy policy and OPT.

For example, suppose the set of ground elements is $\{e_1, \dots, e_6\}$ and there are only three stochastic subsets, S_1, S_2, S_3 . Consider the following realization of these subsets: $S_1 = \{e_1, e_2, e_3, e_6\}$, $S_2 = \{e_3, e_4, e_6\}$, and $S_3 = \{e_1, e_2, e_5\}$. Suppose that under this realization, the greedy policy covers the ground elements using all three subsets, in the following order: S_1, S_3, S_2 . Suppose that the optimal policy, under the above realization, covers the elements using just S_2 and S_3 . Figure 4 shows the $m = 6$ elements of the ground set from left to right in the order they were covered by the S_i in the greedy policy (elements covered during the same step are listed in order of their indices). Here ψ is equal to S_1, S_3, S_2 with the realizations listed above, and $\Upsilon = \{S_2, S_3\}$.




Figure 4: Greedy cover example

Right before the greedy policy covers its j th element, there remain $m - j + 1$ uncovered elements. Fix j and let M_{last} denote that set of elements. Let $M_{com} \subseteq M_{last}$ denote the subset of elements in M_{last} that are first covered, during execution of the greedy policy, by subsets in Υ . A crucial step of the proof of Lemma 3 is to show that

$$m-j+1 = \sum_{S_i \in \Upsilon} cov(S_i) + \sum_{S_i \in \mathcal{S} \setminus \psi} optcov(S_i)$$

where $cov(S_i)$ is the number of elements in M_{last} that are first covered by S_i during execution of the greedy policy, and $optcov(S_i)$ is the expected number of elements in $M_{last} \setminus M_{com}$ first covered by S_i during execution of the optimal algorithm. However, this equation does not hold. (It would be sufficient to show that the left hand side of the equation is bounded above by the right hand side, but that does not hold either.)

Figure 4 with $j=1$ is a counterexample. In this case, M_{last} consists of all 6 ground elements. In this ψ , S_1 covers 4 new elements, S_3 covers 1 (namely e_5), and S_2 covers 1 (namely e_4). Therefore, $cov(S_2)=1$ and $cov(S_3)=1$. However, $\mathcal{S} \setminus \psi$ is empty. So in this case, the left hand side of the equation equals 6, while the right hand side equals 2. The problem is that in the optimal policy, S_2 and S_3 are used to cover all 6 ground elements, but the greedy policy uses them to cover only 2 elements, because the others were already covered by an S_i not in Υ . The conditioning on ψ and Υ does not imply the equality. While the condition can be strengthened, it is not clear how to do so in a way that doesn't invalidate other parts of the proof.

Appendix C. Table of Frequently Used Symbols

Symbol	Description
E, e	Set of items, an individual item
$c(e)$	Cost of item e
O	Set of possible states for an item
Φ	Random realization of all item states
φ	A fixed realization; a possible outcome for Φ
ψ	A fixed subrealization
$dom(\psi)$	Items whose states are recorded in ψ
f	Utility function mapping from set of all subrealizations to $\mathbb{R}_{\geq 0}$
$F_\psi(e)$	Random marginal increase in f from observing state of e , after observing states in ψ
Q	Value of $f(\varphi)$ for all realizations φ (cover value)
η	Minimum value of $Q - f(\psi)$ for any $f(\psi) < Q$
π, π^*, σ	A fixed policy; optimal policy; greedy policy
π^ψ	Hybrid policy, runs σ and switches to π^* at ψ
$\pi(\psi)$	Item selected by π after observing states in ψ
$T(\pi)$	Decision tree for policy π
$L(\pi), N(\pi)$	Sets of subrealizations for leaf and non-leaf nodes of $T(\pi)$,
$\Psi(\pi)$	Random cover constructed by π
$C(\pi)$	Random cost of policy π
$\mu(e), \mu^*$	Revenue from e in π^* ; total revenue in π^*
λ_i	Revenue from marker ρ_i in σ
e_ψ	Item chosen at node ψ of greedy tree $T(\sigma)$
$v(e \psi), v_\psi$	unit price $c(e)/F_\psi(e)$ of e w.r.t. ψ ; unit price paid at node ψ of $T(\sigma)$
ρ_i, ϵ_i	Subrealization that is i th marker of σ ; difference in utility between ρ_i and ρ_{i+1}
$\delta_\psi^+(e)$	Equal to marginal utility from e in hybrid policy, if e selected during second stage, else equals 0
$\pi_{-1}[e]$	Subrealizations for nodes of greedy tree where e is chosen
$\psi \rightsquigarrow i$	ψ is last realization visited by σ such that $f(\psi) \leq f(\rho_i)$
$\mathcal{V}^\sigma(e), \mathcal{V}^\psi(e), \mathcal{V}^*(e)$	e is chosen by: policy σ , policy π^ψ , policy π^*
$\mathcal{A}(e), \mathcal{B}(e)$	e is chosen by π^* but not by σ , e is chosen by both

Table 1: Frequently used symbols

References

- Azar, Y., & Gamzu, I. (2011). Ranking with submodular valuations. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms, SODA*, pp. 1070–1079. SIAM.
- Deshpande, A., Hellerstein, L., & Kletenik, D. (2016). 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.
- Dinur, I., & Steurer, D. (2014). Analytical approach to parallel repetition. In *Proceedings of the ACM Symposium on Theory of Computing, STOC*, pp. 624–633.
- Esfandiari, H., Karbasi, A., & Mirrokni, V. S. (2019). Adaptivity in adaptive submodularity. *CoRR*, abs/1911.03620.
- Feige, U. (1998). A threshold of $\ln n$ for approximating set cover. *J. ACM*, 45(4), 634–652.
- Fujito, T. (2000). Approximation algorithms for submodular set cover with applications. *IEICE Transactions on Informations and Systems*, 83.
- Goemans, M., & Vondrák, J. (2006). Stochastic covering and adaptivity. In *Proceedings of LATIN 2006: Theoretical Informatics*, pp. 532–543. Springer Berlin Heidelberg.
- Golovin, D., & Krause, A. (2011). Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *J. Artif. Int. Res.*, 42(1), 427–486.
- Golovin, D., & Krause, A. (2017). Adaptive submodularity: A new approach to active learning and stochastic optimization (version 5). *CoRR*, abs/1003.3967.
- Hellerstein, L., & Kletenik, D. (2018). Revisiting the approximation bound for stochastic submodular cover. *J. Artif. Intell. Res.*, 63, 265–279.
- Hochbaum, D. S. (1982). Approximation algorithms for the set covering and vertex cover problems. *SIAM Journal on Computing*, 11(3), 555–556.
- Im, S. (2016). Min-sum set cover and its generalizations. In *Encyclopedia of Algorithms*, pp. 1331–1334. Springer.
- Im, S., Nagarajan, V., & van der Zwaan, R. (2016). Minimum latency submodular cover. *ACM Trans. Algorithms*, 13(1), 13:1–13:28.
- Liu, Z., Parthasarathy, S., Ranganathan, A., & Yang, H. (2008). Near-optimal algorithms for shared filter evaluation in data stream systems. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pp. 133–146.
- Nan, F., & Saligrama, V. (2017). Comments on the proof of adaptive stochastic set cover based on adaptive submodularity and its implications for the group identification problem in "Group-based active query selection for rapid diagnosis in time-critical situations". *IEEE Trans. Information Theory*, 63(11), 7612–7614.
- Wan, P., Du, D., Pardalos, P. M., & Wu, W. (2010). Greedy approximations for minimum submodular cover with submodular cost. *Comp. Opt. and Appl.*, 45(2), 463–474.
- Wolsey, L. A. (1982). An analysis of the greedy algorithm for the submodular set covering problem. *Combinatorica*, 2(4), 385–393.