

Submodular Stochastic Probing with Prices

Ben Chugg* Takanori Maehara†

Abstract

We introduce *Stochastic Probing with Prices (SPP)*, a variant of the *Stochastic Probing (SP)* model in which we must pay a price to probe an element. A SPP problem involves two set systems $(\mathcal{N}, \mathcal{I}_{\text{in}})$ and $(\mathcal{N}, \mathcal{I}_{\text{out}})$ where each $e \in \mathcal{N}$ is *active* with probability p_e . To discover whether an element e is active, it must be *probed* by paying the *price* Δ_e . If an element is probed and is active, then it is irrevocably added to the solution. Moreover, at all times, the set of probed elements must lie in \mathcal{I}_{out} , and the solution (the set of probed and active elements) must lie in \mathcal{I}_{in} . The goal is to maximize a submodular set function f minus the cost of the probes. We give a bi-criteria approximation algorithm to the online version of this problem, in which the elements are shown to the algorithm in a possibly adversarial order. Our results translate to state-of-the-art approximations for the traditional (online) stochastic probing problem.

1 Introduction

Life is hard: There is lots of uncertainty, it is rarely the case that one can gather all the information before making a decision, and information isn't free. Whether we're buying a plane ticket for a conference, or a pair of socks for a significant other¹, it is impossible to be aware of every item in the market and gathering the information to make a good decision costs time and resources. The job of computer scientists is to try and model such scenarios. In this paper, we propose the *Stochastic Probing with Prices (SPP)* model, which combines two previous models—Stochastic Probing [16] and the Price of Information [20]—in order to study decision making under uncertainty where ascertaining new information has a cost.

1.1 Problem Formulation

Let $\mathcal{N} = \{e_1, \dots, e_n\}$ be a set of n elements, and $(\mathcal{N}, \mathcal{I}_{\text{in}})$, $(\mathcal{N}, \mathcal{I}_{\text{out}})$ be two downward-closed set systems². Every element $e \in \mathcal{N}$ is *active* with probability p_e , independently of all other elements. In order to determine whether e is active, we can *probe* e . However, if e is probed and is active, then it must be added to our current solution (which thus consists of all previously probed elements which are also active). Moreover, at all times the set of probed elements must lie in \mathcal{I}_{out} and the solution must lie in \mathcal{I}_{in} . Thus, if $P \subset \mathcal{N}$ is the set probed elements so far and $S \subset P$ is the solution, we can only query an element e that satisfies $P \cup \{e\} \in \mathcal{I}_{\text{out}}$ and $S \cup \{e\} \in \mathcal{I}_{\text{in}}$. The set of active elements is determined a priori, i.e., before the execution of the algorithm. Given an instance of the problem, let \mathcal{A} refer to the

*benjamin.chugg@maths.ox.ac.uk. Supported by a research internship at RIKEN.

†takanori.maehara@riken.jp

¹We leave the question of whether a pair of socks is the optimal gift for another paper.

² $\mathcal{I} \subset 2^{\mathcal{N}}$ is *downward-closed* if $A \in \mathcal{I}$ and $B \subset A$ implies $B \in \mathcal{I}$.

set of active elements. We suppose we are given two set-functions, $f, \text{cost} : 2^{\mathcal{N}} \rightarrow \mathbb{R}$, where f is a *utility* or *objective* function, and cost is (naturally) a *cost* or *payment* function. The goal is to query the elements in such a way as to maximize f while minimizing the cost paid for the queries, i.e., to maximize $\mathbf{E}_{\mathcal{A}}[f(P \cap \mathcal{A}) - \text{cost}(P)]$. In this paper we assume that the cost function is linear: for each $e \in \mathcal{N}$ there exists a price $\Delta_e \in \mathbb{R}$ drawn from some distribution F_e , and $\text{cost}(P) = \sum_{e \in P} \Delta_e$. For simplicity, we will assume that Δ_e is deterministic (i.e., F_e is a point mass). Otherwise, since cost is linear, we can simply replace Δ_e with $\mathbf{E}[\Delta_e]$ to obtain our results for more general distributions. We call this model *Stochastic Probing with Prices (SPP)*. If we can choose the order in which the data is revealed to us is, the problem is *offline*, otherwise it is *online*. We draw a distinction between two kinds of online settings: an *adversarial setting*, in which the order may be chosen by an adversary, and a *random-order setting*, in which the elements are chosen uniformly at random. In this paper, all of our results pertaining to online SPP apply in the adversarial online setting. Surprisingly, when f is submodular, our algorithms do not seem to yield better approximations in the offline or random-order settings than in the adversarial setting.

SPP is a generalization of the *Stochastic Probing (SP)* model of Gupta and Nagaranjan [16] and borrows the pricing idea from the recently developed *Price of Information (PoI)* model of Singla [20]—see Sections 1.2 and 1.3. It falls under the general category of what we might call a *query-and-commit* model of computation, in which irrevocable decisions must be made on the fly. Stochastic matching, packing and the secretary problem are examples of problems studied under such a model [12, 17, 9, 7]. Since SPP is a generalization of SP (as we will see), it is useful in the same application domains as the latter model. However, we believe few problems encountered in practice do not have a cost associated with obtaining information (even if it is only a time cost).

We now formally introduce the models upon which SPP is built, and relevant known results.

1.2 Stochastic Probing

The SP model is precisely the SPP model without the cost function. That is, the goal is to maximize f only. It has been a valuable abstraction of stochastic matching and packing, and has been useful in modeling such practical problems as kidney exchange, online dating, and posted price mechanisms [16, 2, 1].

In order to present many of the known results for SP, we need to introduce *contention resolution schemes (CRSs)*, first introduced by Chekuri et al. [6]. Suppose we are given a fractional solution \mathbf{x} to an LP-relaxation of some constrained maximization problem. The goal is now to round \mathbf{x} to an integral solution \mathbf{x}_I which respects the constraints, without losing too much of the value given by the fractional solution. Intuitively, a c -CRS gives a guarantee that if $x_e > 0$ then x_e will be in the solution with probability at least cx_e . Thus, if L is the linear (or concave) objective function, then $L(\mathbf{x}_I) \geq cL(\mathbf{x})$, with high probability (w.h.p.). A (b, c) -CRS gives the same guarantee assuming that $x \in b\mathcal{P}$, where \mathcal{P} is a polytope relaxation of the constraints of the problem. (Thus, a $(1, c)$ -CRS is a c -CRS.) Feldman et al. [15] recently extended the notion of CRSs to online settings. We will provide the formal definition of both offline and online CRSs in Section 2.2.

As mentioned above, the general framework of offline SP problems was first introduced by [16] who restrict their attention to modular (linear) objective functions. When \mathcal{I}_{in} and \mathcal{I}_{out} admit offline (b, c_{in}) and (b, c_{out}) schemes, they give a $b(c_{\text{out}} + c_{\text{in}} - 1)$ -approximation to modular SP. This implies a $\Omega(\frac{1}{(k+\ell)^2})$ -approximation when \mathcal{I}_{in} and \mathcal{I}_{out} are k and ℓ systems

[16]. If \mathcal{I}_{in} and \mathcal{I}_{out} are the intersection of k and ℓ matroids, respectively, then they give a $\frac{1}{4(k+\ell)}$ -approximation. Under these same constraints, Adamczyk et al. [2] improve this approximation to $\frac{1}{k+\ell}$ via a randomized rounding technique, and give a $\frac{1-1/e}{k+\ell+1}$ -approximation in the more general case when f is monotone submodular. Online SP was first introduced and studied by Feldman et al. [15]. When \mathcal{I}_{in} and \mathcal{I}_{out} admit (b, c_{in}) and (b, c_{out}) online CRSs respectively, they give a $bc_{\text{out}}c_{\text{in}}$ approximation to online modular SP, and a $(1 - e^{-b})c_{\text{in}}c_{\text{out}}$ -approximation to online, submodular SP for monotone objective functions. These results hold in the adversarial setting. More recently, Adamczyk and Włodarczyk [3] initiated the study of random order CRSs, and use such schemes to achieve the first approximations to non-monotone submodular SP in the random-order setting, giving a $\frac{1}{(k+\ell+1)e}$ -approximation.

1.3 Price of Information

Perhaps the first, well-known problem to incorporate costs for querying elements is Weitzman’s Pandora’s Box [22]. Generalizing this problem, Singla [20] recently developed the PoI model.

In this setting, we are given a single set system³ $(\mathcal{N}, \mathcal{I})$ and each element e has a value v_e which is drawn from some distribution D_e . The distribution is not limited to a Bernoulli distribution as in the SP model. Each element e also a *price*, Δ_e . The goal is to query a set of elements P in order to maximize $\mathbf{E}_{\{D_e\}}[\max_{I \subset P, I \in \mathcal{I}} \{\sum_{e \in I} v_e + h(I)\} - \sum_{e \in P} \Delta_e]$, where h is some function that depends only on the chosen set, not the values of the elements. The function $\sum_{e \in I} v_e + h(I)$ is called a *semiadditive* function for this reason. Besides the fact that one must pay for making queries in this model, the key distinction is that one gets to choose the maximizing argument $I \subset P$. That is, this model is not a query-and-commit model. The major result of Singla [20] was to prove that if there is an α -approximation to a given problem using what he calls a “frugal” algorithm (essentially, but slightly more general than, a greedy algorithm) when all prices are zero, then there is an α -approximation to the problem with arbitrary prices. Intriguingly, in our work we will obtain a similar result for modular SPP—this will be described in detail later.

While SPP is a generalization of SP, there is no strict inclusion relation between SPP and the PoI model. The PoI model is not a query-and-commit model and as such it does not capture SPP. Furthermore, it is unclear how one would expand a PoI problem to use more general classes of objective functions—submodular functions, for example—while this is easily modeled in SPP. Conversely, SPP does not capture PoI because it is more limited in its assumption that an element is either active or inactive, and cannot model more complicated distributions.

In his Ph.D. thesis, Singla does examine the PoI setting with commitment constraints [19]. In this scenario, the values of the elements change in time as a Markov process, and at each step, we can decide to advance the Markov chain and pay a cost, or stop it and obtain its current value. The setting of activation probabilities can be viewed as a two-stage Markov chain; therefore this model can be viewed as a generalization in that sense. However, there is no outer-constraint, i.e., every element can be queried.

³Although it is possible to generalize the model, as discussed in [20].

1.4 Our Results

In this paper, we are concerned with two classes of objective functions: modular and submodular. We call the corresponding problems modular and submodular SPP. Note that if f is modular then there exists weights $(w_e)_{e \in \mathcal{N}}$ such that $f(E) = \sum_{e \in E} w_e$. Consequently, when all prices are zero, this problem has also been called weighted stochastic probing.

We begin by studying modular SPP. We demonstrate that if there exists an α -approximation for SP, then there exists an α -approximation for SPP. This result applies in all SPP settings (e.g., adversarial, random order, offline).

Theorem 1. *If there exists an α -approximation to modular SP then there exists an α -approximation to modular SPP.*

For submodular SPP, the problem becomes significantly more difficult. It is known that submodular maximum facility location, which is a special case of SPP, is inapproximable to within any non-trivial factor [10]. It is therefore highly unlikely that there exists an approximation of SPP, even one that holds with high probability. Thus, we relax our requirements and search for constant-factor bi-criteria approximations:

Definition 1. Let Opt be an optimal query strategy. A query strategy Alg is an (α, β) -approximation to SPP if $\mathbf{E}[\text{Alg}] \geq \alpha \mathbf{E}[f(\text{Opt} \cap \mathcal{A})] - \beta \text{cost}[\text{Opt}]$, where the expectation is taken over \mathcal{A} and any random choices of the algorithm.

We demonstrate that slight variations of the algorithm for modular SP of Gupta and Naragajan [16] (using offline CRSs) and the algorithm of Feldman et al. [15] (using online CRSs) give good bi-criteria approximations to adversarial submodular SPP if \mathcal{I}_{in} and \mathcal{I}_{out} admit (b, c_{in}) and (b, c_{out}) CRSs. Interestingly, our variation of Gupta and Naragajan’s algorithm increases its scope in two significant ways: we do not require the inner CRS to be ordered, nor the ability to choose the order in which the elements are presented. Thus, our version of their algorithm (which is called **Offline-Rounding** because it uses offline CRSs—Section 4.2) applies to online SPP.

Theorem 2. *Suppose $\mathcal{P}(\mathcal{I}_{\text{in}})$ and $\mathcal{P}(\mathcal{I}_{\text{out}})$ admit (b, c_{in}) and (b, c_{out}) CRSs and let $\gamma = \max\{c_{\text{out}} + c_{\text{in}} - 1, c_{\text{out}}c_{\text{in}}\}$. There exists a $(\gamma \cdot \alpha(b), b)$ -approximation to adversarial, submodular SPP where $\alpha(b) = 1 - e^{-b}$ if the objective function is monotone, and be^{-b} otherwise.*

A note now on how our results translate to SP. When all prices are zero, Theorem 2 gives a $\gamma \cdot \alpha(b)$ -approximation for online, submodular SP. In this case, if \mathcal{I}_{in} and \mathcal{I}_{out} are the intersection of k and ℓ matroids, we demonstrate that the optimal value of b is $\frac{1}{2}(k + \ell + 2 - \sqrt{(k + \ell)(k + \ell + 4)})$ when f is non-monotone, and $k + \ell + 1 - W((k + \ell)e^{k + \ell + 1})$ when f is monotone, where W is the Lambert W function [8].⁴ This gives a $(W(\rho) - k - \ell)^{k + \ell + 1} / W(\rho)$ -approximation in the monotone case, where $\rho = (k + \ell)e^{k + \ell + 1}$, and something significantly uglier in the non-monotone case. These results are given in Section 4.3. These results outperform those of Feldman et al. [15] (who study the adversarial setting) and Adamczyk and Włodarczyk [3] (who study the random-order setting). This is also shown in Section 4.3. In the offline case, the approximation ratio of Adamczyk et al. [2] beats this approximation ratio, however, we remark that our results are the first pertaining to a non-monotone objective function in the adversarial setting.

We hope our results provide a unified view of SP and SPP problems. The analyses use existing techniques which should be familiar to those working in submodular maximization which we hope makes the results easily accessible.

⁴ W is defined by $x = W(x)e^{W(x)}$. It is also called the product log function.

1.5 Proof Techniques

For a modular objective function, there is little daylight between SP and SPP. The former is reduced to the latter by modifying the weights. No such transformation can be made in the more general case of a submodular objective function however, which thus increases the difficulty of the problem.

Similar to the algorithms presented in [16, 2, 3, 15] for SP, our algorithms proceed by obtaining a fractional to a suitably chosen LP-relaxation of the problem, and then rounding according to both CRSs. The first major difficulty is that while $\mathbf{E}[f(P \cap A)] - \text{cost}(P)$ is still a submodular function, it may not be non-negative. Since traditional algorithms for submodular maximization require that the objective function be non-negative, this is a hurdle which needs to be overcome. To do so we use the recent technique of Sviridenko et al. [21] in order to maximize a function of the form $f^+(\mathbf{x} \circ \mathbf{p}) - C(\mathbf{x})$, where f^+ is concave and C is linear. Lemma 3 demonstrates that this objective function (subject to the appropriate constraints) provides an upper bound on the value achieved by the optimal policy.

2 Preliminaries

Let $[n] = \{1, \dots, n\}$. Given a vector $\mathbf{x} \in [0, 1]^{\mathcal{N}}$, let $R(\mathbf{x})$ be a random set where each $e \in \mathcal{N}$ is selected with probability x_e . We will sometimes write $R \sim \mathbf{x}$ to refer to choosing a random set $R(\mathbf{x})$. We call p_e the *activation probability of $e \in \mathcal{N}$* . We emphasize that it is independent of the activation probabilities of other elements. For a set $E \subset \mathcal{N}$ and element $e \in \mathcal{N}$, we will often write $E + e$ in place of the more cumbersome $E \cup \{e\}$. Given a (stochastic) probing algorithm Alg, we will abuse notation somewhat and write $\mathbf{E}[\text{Alg}]$ to mean the expected value of the solution obtained by querying according to Q , i.e., $\mathbf{E}[f(\text{Alg} \cap \mathcal{A}) - \text{cost}(\text{Alg})]$. Note that the expectation is over the joint distribution $\{p_e\}$ and any randomness in the choices made by the algorithm.

2.1 Modular and Submodular Functions

A function $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}$ is *submodular* if for all $A, B \subset \mathcal{N}$ $f(A) + f(B) \geq f(A \cap B) + f(A \cup B)$. If this inequality holds with equality, then f is *modular*. Given $E \subset \mathcal{N}$, write $f_E(e)$ for $f(E + e) - f(E)$. We will consider two extensions of a set function $f : 2^{\mathcal{N}} \rightarrow \mathbb{R}$ to $[0, 1]^{\mathcal{N}}$. The first is the *multilinear extension* [4, 5]:

$$F(\mathbf{y}) = \sum_{A \subset \mathcal{N}} f(A) \prod_{e \in A} y_e \prod_{e \in A^c} (1 - y_e) = \mathbf{E}_{R \sim \mathbf{y}}[f(R)]$$

. The second is the concave closure of f [6]:

$$f^+(\mathbf{y}) = \max_{S \subset \mathcal{N}} \left\{ \sum_{S \subset \mathcal{N}} p_S f(S) : p_S \geq 0 \forall S \subset \mathcal{N}, \sum_{S \subset \mathcal{N}} p_S = 1, \sum_{S: e \in S} p_S = y_e \forall e \in \mathcal{N} \right\}.$$

As indicated by the wording f^+ is concave. It is well known that $F(\mathbf{y}) \leq f^+(\mathbf{y})$. Indeed, $f^+(\mathbf{y})$ can be viewed as selecting the probability distribution over $2^{\mathcal{N}}$ maximizing $\mathbf{E}[f(E)]$ subject to the constraint that $\mathbb{P}[e \in E] \leq y_e$, while $F(\mathbf{y})$ is the value of $\mathbf{E}[f(E)]$ under a particular such distribution.

2.2 Contention Resolution Schemes

Let $\mathcal{I} \subset 2^{\mathcal{N}}$ be downward-closed. The *polytope relaxation* of \mathcal{I} is the set $\mathcal{P}(\mathcal{I}) \subset [0, 1]^{\mathcal{N}}$ defined as the convex hull of all characteristic vectors of \mathcal{I} .

For the rest of this section fix a downward closed set system \mathcal{I} , and let $\mathcal{P}(\mathcal{I})$ be its convex relaxation. We now introduce the formal definitions of offline and online CRSs.

Definition 2 (Offline CRS [6]). For $b, c \in [0, 1]$, a (b, c) *offline Contention Resolution scheme* for $\mathcal{P}(\mathcal{I})$ is family of (possibly randomized) functions $\{\pi_{\mathbf{x}} : 2^U \rightarrow 2^U : \mathbf{x} \in b\mathcal{P}(\mathcal{I})\}$ such that for all $\mathbf{x} \in b\mathcal{P}(\mathcal{I})$ and $E \subset \mathcal{N}$: (1) $\pi_{\mathbf{x}}(E) \subset E \cap \text{supp}(\mathbf{x})$; (2) $\pi_{\mathbf{x}}(E) \in \mathcal{I}$ with probability 1; and (3) for all $e \in \text{supp}(\mathbf{x})$, $\mathbb{P}_{R(\mathbf{x}), \pi}[e \in \pi_{\mathbf{x}}(R(\mathbf{x})) | e \in R(\mathbf{x})] \geq c$.

The CRS π is *monotone* if for all $e \in E_1 \subset E_2$, $\mathbb{P}_{\pi}[e \in \pi(E_1)] \geq \mathbb{P}_{\pi}[e \in \pi(E_2)]$. We emphasize that in condition (3) of Definition 2, the probability is over both the random set $R(\mathbf{x})$ and the CRS, whereas in the definition of monotonicity of a CRS, the probability is taken only over the (possibly) random choices of scheme itself.

Definition 3 (Online CRS [15]). For $b, c \in [0, 1]$, a (b, c) *online contention resolution scheme* of $\mathcal{P}(\mathcal{I})$ is a procedure which defines for any $\mathbf{x} \in \mathcal{P}(\mathcal{I})$ a family $\mathcal{I}_{\mathbf{x}} \subset \mathcal{I}$ such that for all $e \in \mathcal{N}$, $\mathbb{P}[I + e \in \mathcal{I}_{\mathbf{x}} \text{ for all } I \subset R(\mathbf{x}), I \in \mathcal{I}_{\mathbf{x}}] \geq c$.

Given an online CRS of $\mathcal{P}(\mathcal{I})$ the authors of [15] define a related offline CRS, called the *characteristic* CRS $\pi(E) = \{e \in E : I + e \in \mathcal{I}_{\pi, \mathbf{x}}, \text{ for all } I \subset E, I \in \mathcal{I}_{\pi, \mathbf{x}}\}$. They verify that the characteristic CRS meets the conditions of an offline CRS and demonstrate that the characteristic CRS of a (b, c) online CRS is a monotone, (b, c) CRS. If we refer to a CRS without specifying whether it is online or offline, the discussion should apply to both. Unless otherwise stated, we will assume that any CRS discussed in this paper is *efficient*: that it can be computed in polynomial time.

2.3 A general probing strategy

A *permutation* or *ordering* on a subset $E \subset \mathcal{N}$ is a bijection $\sigma : [n] \rightarrow E$. To say *traverse E in the order of σ* should be taken to mean iterate over E in the order $\sigma(1), \sigma(2), \dots, \sigma(|E|)$. The following probing procedure will be repeated often enough throughout the paper to warrant a name.

GreedyProbing $(\sigma, \mathcal{I}_1, \mathcal{I}_2)$ **on** $E \subset \mathcal{N}$. Let $P \leftarrow \emptyset$ and $S \leftarrow \emptyset$. For all $e \in E$ in the order σ if $P + e \in \mathcal{I}_1$ and $S + e \in \mathcal{I}_2$, then

1. if $f_S(e) \geq 0$ then probe e and set $P \leftarrow P + e$. If e was active, set $S \leftarrow S + e$.
2. if $f_S(e) < 0$, then add e to S with probability p_e .

Return S, P .

We observe that this strategy does in fact produce a valid set of probed elements and a valid solution.

Observation 1. Let S and P be the sets returned by **GreedyProbing**. Let \overline{S} be the set of elements queried by **GreedyProbing** which were active. Then $P \cap \mathcal{A} = \overline{S} \subset S$. Moreover, $S \in \mathcal{I}_2$ hence $\overline{S} \in \mathcal{I}_2$ and $P \in \mathcal{I}_1$ (assuming \mathcal{I}_1 and \mathcal{I}_2 are downward closed).

3 Modular SPP

The main result of this section is to verify the intuitive result that when the objective function, f , is modular, we can obtain the same approximation ratio to the optimal as in SP (i.e., with no prices). Recall that if f is modular, there exist weights $(w_e)_{e \in \mathcal{N}}$ such that $f(E) = \sum_{e \in E} w_e$. Here, we will not draw a distinction between online and offline SPP. This is because the results apply in both settings. We will reduce an instance of SPP to an instance of SP. The result requires only a change of variables, and thus only knowledge of the weights, prices, and activation probabilities in the SPP instance. Thus, if the algorithm for SP is online or applies in any other setting⁵, so too does the corresponding SPP algorithm.

We begin with an observation that no algorithm worth its salt will query elements whose expected weight is upper bounded by their price.

Lemma 1. *If Q is a query strategy which queries e where $w_e p_e \leq \Delta_e$, there exists a query strategy Q' which does not query e such that $\mathbf{E}[Q'] \geq \mathbf{E}[Q]$.*

Proof. Given Q , define Q' to be the strategy obtained by mimicking Q but refraining from querying any element e such that $w_e p_e \leq \Delta_e$. It follows that Q' is still a valid strategy because \mathcal{I}_{in} and \mathcal{I}_{out} are downward-closed. Moreover, any element queried by Q which is not queried by Q' adds a non-positive expected value to the solution. \square

Henceforth we will apply Lemma 1 and assume that for all elements e , $w_e p_e > \Delta_e$. In this section we will allow prices to be negative; therefore, it is not necessarily the case that $p_e > 0$. This section requires that we make comparisons between solutions of SSP and SP, and hence need to introduce the relevant notation. For a probing strategy Q , let $\mathbf{E}[Q^{\text{SP}}(\mathbf{w}, \mathbf{p})]$ denote the expected value in the SP setting with weights $\mathbf{w} = (w_e)_{e \in \mathcal{N}}$ and probabilities $\mathbf{p} = (p_e)_{e \in \mathcal{N}}$. Similarly, let $\mathbf{E}[Q^{\text{SPP}}(\mathbf{w}, \mathbf{p}, \Delta)]$ denote the expected value of the solution in the SSP setting with weights (w_e) , probabilities (p_e) and prices (Δ_e) .

Given an instance of SSP, we define new weights and activation probabilities as follows. For all e , let $z_e = w_e - \Delta_e/p_e$ if $p_e > 0$ and $-\Delta_e$ otherwise and let $\hat{p}_e = p_e$ for all e with $p_e > 0$, and $\hat{p}_e = 1$ for e with $p_e = 0$.

Lemma 2. *For any querying strategy Q and any weights \mathbf{w} , activation probabilities \mathbf{p} and prices Δ , $\mathbf{E}[Q^{\text{SPP}}(\mathbf{w}, \mathbf{p}, \Delta)] = \mathbf{E}[Q^{\text{SP}}(\mathbf{z}, \hat{\mathbf{p}})]$.*

Proof. Let $\hat{\mathcal{A}}$ be the random variable denoting the set of active elements according to the probabilities $\hat{\mathbf{p}}$. Note that $p_e > 0$ for all e , hence $\mathbb{P}[e \in Q \cap \hat{\mathcal{A}}]/\hat{p}_e = \mathbb{P}[e \in Q]$. Unwinding definitions now gives

$$\begin{aligned} \mathbf{E}[Q^{\text{SPP}}(\mathbf{w}, \mathbf{p}, \Delta)] &= \sum_{e \in \mathcal{N}} w_e \mathbb{P}[e \in Q \cap \mathcal{A}] - \Delta_e \mathbb{P}[e \in Q] \\ &= \sum_{e: p_e > 0} (w_e p_e - \Delta_e) \mathbb{P}[e \in Q] - \sum_{e: p_e = 0} \Delta_e \mathbb{P}[e \in Q] \\ &= \sum_{e \in \mathcal{N}} z_e \mathbb{P}[e \in Q \cap \hat{\mathcal{A}}] = \mathbf{E}[Q^{\text{SP}}(\mathbf{z}, \hat{\mathbf{p}})]. \end{aligned} \quad \square$$

We can now prove Theorem 1.

⁵For example, stochastic probing with deadlines has been examined [16, 15].

Proof. Let Alg be an algorithm for linear SP which obtains an α -approximation. Let $\mathbf{w}, \mathbf{p}, \Delta$ be an instance of linear SPP. Running Alg on weights \mathbf{z} and probabilities $\hat{\mathbf{p}}$ and applying Lemma 2 gives

$$\mathbf{E}[\text{Alg}^{\text{SPP}}(\mathbf{w}, \mathbf{p}, \Delta)] = \mathbf{E}[\text{Alg}^{\text{SP}}(\mathbf{z}, \hat{\mathbf{p}})] \geq \alpha \mathbf{E}[\text{Opt}^{\text{SP}}(\mathbf{z}, \hat{\mathbf{p}})] = \alpha \mathbf{E}[\text{Opt}^{\text{SPP}}(\mathbf{w}, \mathbf{p}, \Delta)]. \quad \square$$

4 Submodular SPP

We now proceed to the more general problem of submodular SPP. We assume in this section that the prices are non-negative. We will employ the common approach of solving a relaxed linear program (Section 4.1), and then rounding the solution to obtain a probing policy (Section 4.2). First we observe a natural upper bound on the value of the optimal solution, against which we can gauge the quality of approximations. It will be notationally convenient to work with a single polytope instead of both $\mathcal{P}(\mathcal{I}_{\text{in}})$ and $\mathcal{P}(\mathcal{I}_{\text{out}})$. Accordingly, we will henceforth let \mathcal{P} refer to the polytope

$$\left\{ \mathbf{x} \in [0, 1]^{\mathcal{N}} : \mathbf{x} \in \mathcal{P}(\mathcal{I}_{\text{out}}), \mathbf{x} \circ \mathbf{p} \in \mathcal{P}(\mathcal{I}_{\text{in}}) \right\}.$$

Let C be the (multi)linear extension of cost . Thus $C(\mathbf{x}) = \sum_{e \in \mathcal{N}} \Delta_e x_e$. A natural relaxation of submodular SPP is the following program:

$$\max_{\mathbf{x}} \left\{ f^+(\mathbf{x} \circ \mathbf{p}) - C(\mathbf{x}) : \mathbf{x} \in \mathcal{P} \right\}. \quad (\text{LP})$$

The relationship between SPP algorithms and LP is given by the following lemma.

Lemma 3. *If Alg be any (stochastic) probing strategy then there exists a point $\mathbf{x} \in \mathcal{P}$ such that $\mathbf{E}[f(\text{Alg} \cap \mathcal{A})] \leq f^+(\mathbf{x} \circ \mathbf{p})$ and $\mathbf{E}[\text{cost}(\text{Alg})] = \sum_{e \in \mathcal{N}} \Delta_e x_e$.*

Proof. Define \mathbf{x} and \mathbf{y} by $x_e = \mathbb{P}[e \in \text{Alg}]$ and $y_e = \mathbb{P}[e \in \text{Alg} \cap \mathcal{A}]$. It is immediate that $\mathbf{x} \in \mathcal{P}(\mathcal{I}_{\text{in}})$ and $\mathbf{y} \in \mathcal{P}(\mathcal{I}_{\text{out}})$ if Alg is a valid strategy. First, we claim that \mathbf{x} is feasible solution. For this it suffices to show that $y_e = x_e p_e$. Recall that by the constraints of the problem, if an element is queried, then its addition to the current set of probed elements and to the solution must be allowed by the constraints of \mathcal{I}_{out} and \mathcal{I}_{in} respectively. Thus, $y_e = \mathbb{P}[e \in \text{Alg} \cap \mathcal{A}] = \mathbb{P}[e \in \text{Alg}] \mathbb{P}[e \in \mathcal{A}] = x_e p_e$, where we've used the fact that whether a particular element is active or not is fixed a priori. Now, notice that $\mathbf{E}[f(\text{Alg} \cap \mathcal{A})] = \mathbf{E}_{S \sim D}[f(S)]$ where D is a particular distribution such that $\mathbb{P}_{S \sim D}[e \in S] \leq p_e x_e$. Conversely, $f^+(\mathbf{x} \circ \mathbf{p})$ is the maximum over all such distributions, i.e., $f^+(\mathbf{x} \circ \mathbf{p}) = \max_D \mathbf{E}_{S \sim D}[f(S)]$. Therefore, $f^+(\mathbf{x} \circ \mathbf{p}) \geq \mathbf{E}[f(\text{Opt} \cap \mathcal{A})]$. Moreover, by the linearity of cost it's easy to see that $\mathbf{E}[\text{cost}(\text{Alg})] = \sum_{e \in \mathcal{N}} \Delta_e \mathbb{P}[e \in \text{Alg}] = \sum_{e \in \mathcal{N}} \Delta_e x_e$. \square

4.1 Obtaining a fractional solution

While the problem of maximizing a (non-monotone) submodular function subject to various constraints has been the subject of intense study (e.g., [6, 14, 18, 11]), less is known about combinations of submodular functions. In our case, the difficulty in solving (LP) efficiently arises because the function $f^+(\mathbf{x} \circ \mathbf{p}) - C(\mathbf{x})$ is not necessarily non-negative. Removing the non-negativity condition in a non-monotone submodular maximization problem makes the problem intractable in general, since, as noted in [6], it may take an exponential number

of queries to determine whether the optimum is greater than zero. We must therefore take advantage of the special form of our problem; namely the fact that C is linear.

Recently, Sviridenko et al. gave an approximation algorithm for maximizing the sum of a non-negative, normalized, monotone, submodular function f and a linear function over a matroid constraint [21, Theorem 3.1]. More precisely, given a submodular function g , a linear function ℓ and a matroid \mathcal{I} , with high probability they obtain a set E such that $g(E) + \ell(E) \geq (1 - 1/e)g(I) + \ell(I)$ minus an arbitrarily small constant term, for any base $I \in \mathcal{I}$. The idea is elegant and straightforward, and involves using the traditional continuous greedy algorithm but over the polytope $\mathcal{P} \cap \{\mathbf{x} : \ell(\mathbf{x}) \geq \lambda\}$ (rather than simply \mathcal{P}) where λ is a guess for the value of $\ell(\text{Opt})$. Intuitively, this guarantees that the fractional solution \mathbf{x}^* satisfies $L(\mathbf{x}^*) \geq \ell(\text{Opt})$ (where L is the linear extension of ℓ). Somewhat surprisingly, restricting the polytope in this way does not damage the approximation to f .

Using Measured Continuous Greedy [14] instead of continuous greedy and the modification of its proof used in [1] and [3], we are able to extend the approach of Sviridenko et al. to non-monotone submodular functions and arbitrary constraints. The result is summarized as Theorem 3. We say a function $G : [0, 1]^{\mathcal{N}} \rightarrow \mathbb{R}$ can be *efficiently estimated with high probability* if for any $\mathbf{x} \in [0, 1]^{\mathcal{N}}$, $G(\mathbf{x})$ can be determined to within exponentially small error with a polynomial number of queries. For example, if we are given oracle access to a submodular function f , then its multilinear extension can be efficiently estimated with high probability (see, e.g., [6]).

Theorem 3. *Let $g : 2^{\mathcal{N}} \rightarrow \mathbb{R}_{\geq 0}$ be a normalized submodular function with multilinear extension G which can be efficiently estimated with high probability. Let $\ell : 2^{\mathcal{N}} \rightarrow \mathbb{R}$ be a non-decreasing modular function with (multi)linear extension L . For any $\epsilon > 0$, $T \geq 0$, and downward-closed system $\mathcal{I} \subset 2^{\mathcal{N}}$, there exists a polynomial time algorithm which produces a point $\mathbf{x}/T \in \mathcal{P}(\mathcal{I})$ such that*

$$G(\mathbf{x}) - L(\mathbf{x}) \geq \alpha(T)g^+(\mathbf{y}) - TL(\mathbf{y}) - O(\epsilon)R,$$

for any $\mathbf{y} \in \mathcal{P}$ with high probability, where $R = \max\{\max_{e \in \mathcal{N}} g(e), \max_{e \in \mathcal{N}} |\text{cost}(e)|\}$ and $\alpha(T) = 1 - e^{-T}$ if g is monotone, and Te^{-T} otherwise.

The proof of Theorem 3 can be found in the Appendix.

4.2 Obtaining a policy

While Theorem 3 guarantees the existence of a *fractional* solution which gives a good approximation to the value of the optimal policy, it does not tell us how to query the elements. This section focuses on extracting a probing policy from this fractional solution. We give two rounding techniques, which apply depending on what kind of CRSs to which one has access (i.e., online or offline). Let $\mathbf{x} \in b\mathcal{P}$.

Offline-Rounding. Let π^{in} and π^{out} be offline CRSs for \mathcal{I}_{in} and \mathcal{I}_{out} respectively. Draw $R \sim \mathbf{x}$ and compute $Q = \pi^{\text{out}}(R)$. Run GreedyProbing($\sigma, 2^{\mathcal{N}}$) on $\pi^{\text{in}}(R) \cap \pi^{\text{out}}(R)$ where σ is any ordering.

Online-Rounding. Let $\mathcal{I}_{\text{in}}^{\mathbf{x}} \subset \mathcal{I}_{\text{in}}$ and $\mathcal{I}_{\text{out}}^{\mathbf{x}} \subset \mathcal{I}_{\text{out}}$ be the (random) subsets given by the inner and outer CRSs. If π^{out} is not online, take $\mathcal{I}_{\text{out}}^{\mathbf{x}} = \mathcal{I}_{\text{out}}$. Draw $R \sim \mathbf{x}$. Run

GreedyProbing($\sigma, \mathcal{I}_{\text{out}}^{\mathbf{x}}, \mathcal{I}_{\text{in}}^{\mathbf{x}}$) on R if π^{out} is online, and on Q otherwise where σ is any ordering (even adaptively and adversarially chosen).

Analysis. For either rounding technique, let P , S , and \bar{S} be as in Observation 1. Given any run of the algorithm, let \mathcal{E} be defined as those elements which were either probed and active, or else had negative marginals but were nonetheless added to S . Note that \mathcal{E} is distributed as $R(\mathbf{p})$ and thus as \mathcal{A} . However, it might not be the case that $\mathcal{E} = \mathcal{A}$ because not all elements in \mathcal{E} were actually probed. Note that P , S and \bar{S} are actually functions of R and \mathcal{E} , and it will oftentimes be helpful to write them as such (i.e., $S = S(R, \mathcal{E})$). Finally, we define $J(R, \mathcal{E}) = \pi^{\text{in}}(R) \cap \pi^{\text{out}}(R) \cap \mathcal{E}$, where π^{in} and π^{out} were the inner and outer CRSs used in the rounding technique (the given schemes in the case of offline rounding, and the characteristic schemes in the case of online rounding). The following lemma uses properties of the CRSs to obtain the inequality which is crucial to the main result.

Lemma 4. *Let π^{in} and π^{out} be monotone (b, c_{in}) and (b, c_{out}) CRSs respectively. Let $\mathbf{x} \in b\mathcal{P}$ and set $R = R(\mathbf{x})$. Then $\mathbb{P}[e \in \pi^{\text{in}}(R) \cap \pi^{\text{out}}(R) | e \in R] \geq \max\{c_{\text{out}} + c_{\text{in}} - 1, c_{\text{out}} \cdot c_{\text{in}}\}$. Moreover, $\mathbb{P}[e \in P] \leq x_e$ where P is the set of elements probed by the algorithm.*

Proof. When conditioning on the choice of the random set R , the randomness in $\pi^{\text{in}}(R)$ and $\pi^{\text{out}}(R)$ stems purely from the CRSs themselves. Hence, the events $\{e \in \pi^{\text{in}}(R) | R\}$ and $\{e \in \pi^{\text{out}}(R) | R\}$ are independent. Additionally, notice that given R and $e \in R$, the quantity $\mathbb{P}[e \in \pi^{\text{in}}(R)]$ is monotonically decreasing in R . Combining these two facts and conditioning on $e \in R$ gives

$$\begin{aligned}
& \mathbb{P}_{R, \pi^{\text{in}}, \pi^{\text{out}}}[e \in \pi^{\text{in}}(R) \cap \pi^{\text{out}}(R)] & (1) \\
&= \mathbf{E}_R [\mathbb{P}_{\pi^{\text{in}}, \pi^{\text{out}}}[e \in \pi^{\text{in}}(R) \cap \pi^{\text{out}}(R) | R]] \\
&= \mathbf{E}_R [\mathbb{P}_{\pi^{\text{in}}}[e \in \pi^{\text{in}}(R) | R] \mathbb{P}_{\pi^{\text{out}}}[e \in \pi^{\text{out}}(R) | R]] \\
&\geq \mathbf{E}_R[\mathbb{P}_{\pi^{\text{in}}}[e \in \pi^{\text{in}}(R) | R]] \cdot \mathbf{E}_R[\mathbb{P}_{\pi^{\text{out}}}[e \in \pi^{\text{out}}(R) | R]] \\
&= \mathbb{P}_{R, \pi^{\text{in}}}[e \in \pi^{\text{in}}(R)] \mathbb{P}_{R, \pi^{\text{out}}}[e \in \pi^{\text{out}}(R)], & (2)
\end{aligned}$$

using the FKG inequality. This implies that $\mathbb{P}[e \in \pi^{\text{in}}(R) \cap \pi^{\text{out}}(R) | e \in R] \geq c_{\text{out}} \cdot c_{\text{in}}$. To obtain the other bound, we compute

$$\begin{aligned}
& \mathbb{P}[e \in \pi^{\text{in}}(R) \cap \pi^{\text{out}}(R)] \\
&= \mathbb{P}[e \in \pi^{\text{out}}(R)] - \mathbb{P}[e \in \pi^{\text{in}}(R)^c \cap \pi^{\text{out}}(R)] \\
&\geq c_{\text{out}} - \mathbb{P}[e \in \pi^{\text{in}}(R)^c] = c_{\text{out}} - (1 - \mathbb{P}[e \in \pi^{\text{in}}(R)]) \\
&\geq c_{\text{out}} - (1 - c_{\text{in}}).
\end{aligned}$$

Rearranging gives the desired result. The final statement follows from noticing that $P \subset R$ in either rounding scheme, and $\mathbb{P}[e \in R] = x_e$. \square

Instead of reasoning directly about the set of elements in the solution, we will reason about the set $J(R, \mathcal{E})$. The following two lemmas prove useful properties about this set.

Lemma 5. *Given $R_1 \subset R_2$ and $\mathcal{E}_1 \subset \mathcal{E}_2$ with $e \in R_1 \cap \mathcal{E}_1$, $\mathbb{P}[e \in J(R_1, \mathcal{E}_1)] \geq \mathbb{P}[e \in J(R_2, \mathcal{E}_2)]$.*

Proof. The randomness comes only from the CRSs, since R_1, R_2, \mathcal{E}_1 and \mathcal{E}_2 are given. Therefore, $\mathbb{P}_{\pi^{\text{in}}, \pi^{\text{out}}}[e \in \pi^{\text{in}}(R_1) \cap \pi^{\text{out}}(R_1) \cap \mathcal{E}_1] = \mathbb{P}_{\pi^{\text{in}}}[e \in \pi^{\text{in}}(R_1)] \mathbb{P}_{\pi^{\text{out}}}[e \in \pi^{\text{out}}(R_1)] \geq \mathbb{P}_{\pi^{\text{in}}}[e \in \pi^{\text{in}}(R_2)] \mathbb{P}_{\pi^{\text{out}}}[e \in \pi^{\text{out}}(R_2)] = \mathbb{P}_{\pi^{\text{in}}, \pi^{\text{out}}}[e \in \pi^{\text{in}}(R_2) \cap \pi^{\text{out}}(R_2) \cap \mathcal{E}_2]$. \square

Lemma 6. *For any R and \mathcal{E} , $J(R, \mathcal{E}) \subset S(R, \mathcal{E})$.*

Proof. Let $e \in J(R, \mathcal{E})$. We first observe that regardless of the rounding technique, if e advances to having its marginal considered in **GreedyProbing** then, by definition of \mathcal{E} , it will be added to $S = S(R, \mathcal{E})$. Therefore, it remains only to show that e meets the condition $P_e + e \in \mathcal{I}_1$ and $S_e + e \in \mathcal{I}_2$ in **GreedyProbing**, where $P_e \subset P$ and $S_e \subset S$ are the respective intermediary solutions immediately before e is considered by the algorithm. In the case of offline rounding, this is immediate since **GreedyProbing** is run on $\pi^{\text{out}}(R) \cap \pi^{\text{in}}(R)$ which is in $\mathcal{I}_{\text{out}} = \mathcal{I}_1$ and $\mathcal{I}_{\text{in}} = \mathcal{I}_2$ (since $\pi^{\text{out}}(R) \subset \mathcal{I}_{\text{out}}$ and $\pi^{\text{in}}(R) \subset \mathcal{I}_{\text{in}}$ w.p. 1). Now consider online rounding with two online schemes. Here, recall that $\pi^{\text{in}}(R) = \{e \in R : I + e \in \mathcal{I}_{\text{in}}, \forall I \subset R, I \in \mathcal{I}_{\text{in}}\}$. Taking $I = S_e$ we see that $S_e + e \in \mathcal{I}_{\text{in}}$. Similarly, $P_e + e \in \mathcal{I}_{\text{out}}$. The argument for online rounding when π^{out} is offline is similar. \square

This final technical lemma derives a lower bound on the marginal of our solution with respect to that of the optimal's. The main approximation guarantee will then result from decomposing the objective function into the sum of its marginals and applying the following lemma.

Lemma 7. *Let $\mathbf{x} \in b\mathcal{P}$ and let $\{e_1, e_2, \dots, e_n\}$ be any ordering on \mathcal{N} . For every $e_i \in \mathcal{N}$, $\mathbf{E}[f_{\overline{S}^{i-1}}(e_i)] \geq \gamma \mathbf{E}[f_{R^{i-1}(\mathbf{x}) \cap \mathcal{A}}(e_i)]$ where $\gamma = \max\{c_{\text{out}} + c_{\text{in}} - 1, c_{\text{out}}c_{\text{in}}\}$.*

Proof. Fix $i \in [n]$ and let $e = e_i$. Set $R = R(\mathbf{x})$. We have:

$$\begin{aligned} \mathbf{E}_{\mathcal{E}, \pi}[f_{\overline{S}^{i-1}}(e)] &= \mathbb{P}[e \in R \cap \mathcal{E}] \mathbf{E}_{\mathcal{E}, \pi}[\mathbf{1}(e \in \overline{S}) f_{\overline{S}^{i-1}}(e) | e \in R \cap \mathcal{E}] \\ &= \mathbb{P}[e \in R \cap \mathcal{E}] \mathbf{E}_{\mathcal{E}, \pi}[\mathbf{1}(e \in \overline{S}) \max\{0, f_{\overline{S}^{i-1}}(e)\} | e \in R \cap \mathcal{E}] \\ &= \mathbb{P}[e \in R \cap \mathcal{E}] \mathbf{E}_{\mathcal{E}, \pi}[\mathbf{1}(e \in S) \max\{0, f_{\overline{S}^{i-1}}(e)\} | e \in R \cap \mathcal{E}] \\ &\geq \mathbb{P}[e \in R \cap \mathcal{E}] \mathbf{E}_{\mathcal{E}, \pi}[\mathbf{1}(e \in S) \max\{0, f_{R^{i-1} \cap \mathcal{E}}(e)\} | e \in R \cap \mathcal{E}], \end{aligned}$$

where the third equality follows from the fact that if $e \in S \setminus \overline{S}$ then $\max\{0, f_{\overline{S}^{i-1}}(e)\} = 0$, and the final inequality follows from submodularity. Let $\phi(R, \mathcal{E}) = \max\{0, f_{R^{i-1} \cap \mathcal{E}}(e)\}$. Now, condition on $e \in R \cap \mathcal{E}$ and write

$$\begin{aligned} \mathbf{E}_{R, \mathcal{E}, \pi}[\mathbf{1}(e \in \overline{S}(R, \mathcal{E})) \cdot \phi(R, \mathcal{E})] &= \mathbf{E}_{R, \mathcal{E}}[\mathbf{E}_{\pi}[\mathbf{1}(e \in \overline{S}(R, \mathcal{E})) \cdot \phi(R, \mathcal{E}) | R, \mathcal{E}]] \\ &= \mathbf{E}_{R, \mathcal{E}}[\mathbf{E}_{\pi}[\mathbf{1}(e \in \overline{S}(R, \mathcal{E})) | R, \mathcal{E}] \cdot \phi(R, \mathcal{E})] \\ &\geq \mathbf{E}_{R, \mathcal{E}}[\mathbf{E}_{\pi}[\mathbf{1}(e \in J(R, \mathcal{E})) | R, \mathcal{E}] \cdot \phi(R, \mathcal{E})], \end{aligned}$$

where the final inequality follows from Lemma 6. Both $\phi(R, \mathcal{A})$ and $\mathbf{E}[\mathbf{1}(e \in J(R, \mathcal{E})) | R \cap \mathcal{E}]$ are decreasing functions of R, \mathcal{E} (the former from submodularity and the latter from Lemma 5). Therefore, by the FKG inequality, the above is at least

$$\begin{aligned} &\mathbf{E}_{R, \mathcal{E}}[\mathbf{E}_{\pi}[\mathbf{1}(e \in J(R, \mathcal{E})) | R, \mathcal{E}] | e \in R \cap \mathcal{E}] \cdot \mathbf{E}_{R, \mathcal{E}}[\phi(R, \mathcal{E}) | e \in R \cap \mathcal{E}] \\ &= \mathbf{E}_{R, \mathcal{E}, \pi}[\mathbf{1}(e \in J(R, \mathcal{E})) | e \in R \cap \mathcal{E}] \cdot \mathbf{E}_{R, \mathcal{E}}[\phi(R, \mathcal{E}) | e \in R \cap \mathcal{E}] \\ &= \mathbf{E}_{R, \pi}[\mathbf{1}(e \in \pi^{\text{in}}(R) \cap \pi^{\text{out}}(R)) | e \in R] \cdot \mathbf{E}_{R, \mathcal{E}}[\phi(R, \mathcal{E}) | e \in R \cap \mathcal{E}] \\ &\geq \gamma \mathbf{E}_{R, \mathcal{E}}[\phi(R, \mathcal{E}) | e \in R \cap \mathcal{E}], \end{aligned}$$

by Lemma 4, where the final equality uses the fact that the events $\{e \in \mathcal{E}\}$ and $\{e \in \pi^{\text{in}}(R) \cap \pi^{\text{out}}(R)\}$ are independent. Combining everything and keeping in mind that $R \cap \mathcal{E}$ is distributed as $R \cap \mathcal{A}$, we obtain

$$\mathbf{E}[f_{S^{i-1}}(e)] \geq \gamma \mathbb{P}[e \in R \cap \mathcal{E}] \mathbf{E}[\phi(R, \mathcal{E}) | e \in R \cap \mathcal{E}] \geq \gamma \mathbf{E}[f_{R^{i-1}(\mathbf{x}) \cap \mathcal{A}}(e)],$$

as desired. \square

4.3 Approximation Guarantees

Given the rounding policies presented in the previous section, we are now ready to prove Theorem 2, and explore the guarantees given in the case where \mathcal{I}_{in} and \mathcal{I}_{out} are the intersection of matroids.

4.3.1 Proof of Theorem 2

Proof. Apply Theorem 3 with $G(\mathbf{x}) = F(\mathbf{x} \circ \mathbf{p})$ and $L(\mathbf{x}) = C(\mathbf{x})$ and run CRS-Rounding on the resulting point $\mathbf{x} \in b \cdot \mathcal{P}$. Then, applying Lemmas 4 and 7 gives

$$\begin{aligned} \mathbf{E}[\text{Alg}] &= \mathbf{E}[f(\text{Alg} \cap \mathcal{A})] - \mathbf{E}[\text{cost}(\text{Alg})] \\ &= f(\emptyset) + \sum_{i=1}^n \mathbf{E}[f_{\text{Alg}^{i-1} \cap \mathcal{A}}(e_i)] - \sum_{e \in \mathcal{N}} \Delta_e \mathbb{P}[e \in \text{Alg}] \\ &\geq f(\emptyset) + \gamma \sum_{i=1}^n \mathbf{E}[f_{R^{i-1} \cap \mathcal{A}}(e_i)] - \sum_{e \in \mathcal{N}} \Delta_e x_e \\ &= \gamma \mathbf{E}[f(R(\mathbf{x}) \cap \mathcal{A})] - C(\mathbf{x}). \end{aligned}$$

Applying Lemma 3, let \mathbf{x}^* satisfy $\mathbf{E}[f(\text{Opt} \cap \mathcal{A})] \leq f^+(\mathbf{x}^* \circ \mathbf{p})$ and $\mathbf{E}[\text{cost}(\text{Opt})] = \sum_{e \in \mathcal{N}} \Delta_e x_e^*$, where Opt is an optimal probing strategy. Noticing that $\mathbf{E}[f(R(\mathbf{x}) \cap \mathcal{A})] = \mathbf{E}[f(R(\mathbf{x}) \circ \mathbf{p})]$, we have

$$\gamma \mathbf{E}[f(R(\mathbf{x}) \cap \mathcal{A})] - C(\mathbf{x}) \geq \gamma \alpha(b) f^+(\mathbf{x}^* \circ \mathbf{p}) - bL(\mathbf{x}^*) - \gamma O(\epsilon)R,$$

with high probability by Theorem 3. Lemma 3 implies that Alg is a $(\gamma \alpha(b), b)$ -approximation. \square

4.3.2 Intersection of matroids for SP

Here we focus on traditional Stochastic Probing when \mathcal{I}_{in} and \mathcal{I}_{out} are the intersection of k and ℓ matroids, respectively, we have the existence of $(b, 1 - b)$ offline CRSs and online CRSs, for any $b \in [0, 1)$ [15, 16]. For any fixed value of b , Theorem 2 gives a $\max\{(1 - b)^{k+\ell}, (1 - b)^k + (1 - b)^\ell - 1\}(1 - e^{-b})$ approximation. By an easy induction, we see that for all $k, \ell \in \mathbb{N}$, $(1 - b)^{k+\ell} \geq (1 - b)^k + (1 - b)^\ell - 1$. Therefore, in the case of matroids, it's always more optimal to use the online CRSs than the offline CRSs and we obtain a $\max_{b \in (0, 1)}\{(1 - b)^{k+\ell}(1 - e^{-b})\}$ approximation. Next, we determine the optimal value of b when the objective function is monotone.

Lemma 8. *For any $k, \ell \in \mathbb{N}$,*

$$(k + \ell + 1) - W((k + \ell)e^{k+\ell+1}) = \arg\max_{b \in (0, 1)}\{(1 - b)^{k+\ell}(1 - e^{-b})\}.$$

Proof. Let $z = k + \ell$ and set $\xi(b) = (1 - b)^z(1 - e^{-b})$. Solving $\xi'(b) = 0$ yields $(1 - b)^z e^{-b} = z(1 - b)^{z-1}(1 - e^{-b})$, i.e., $(1 - b + z)e^{-b} = z$ which is solved by $b = z + 1 - W(ze^{z+1}) \equiv b^*$. The only other critical points of ξ in the region $0 \leq b \leq 1$ are $b = 0, 1$, which yield $\xi(b) = 0$. Thus, b^* is the unique maximizer of $\xi(b)$. \square

Corollary 1. *If \mathcal{I}_{in} and \mathcal{I}_{out} are the intersection of k and ℓ matroids then there exists a $(W(ze^{z+1}) - z)^{z+1}/W(ze^{z+1})$ -approximation to non-monotone, adversarial SP where $z = k + \ell$.*

Proof. Evaluate the approximation ratio at $b = k + \ell - W((k + \ell)e^{k+\ell+1})$. \square

We now observe that the above approximation is state-of-the-art compared to known results for online, monotone SP. Recall that Feldman et al. give a $c_{\text{in}}c_{\text{out}} \cdot (1 - e^{-b})$ -approximation in the adversarial setting, and Adamczyk and Włodarczyk [3] give a $\frac{1}{(k+\ell+1)e}$ -approximation. The former is less than $\max\{c_{\text{in}}c_{\text{out}}, c_{\text{in}} + c_{\text{out}} - 1\}(1 - e^{-b})$. For the latter we use the following lemma.

Lemma 9. $\max_{b \in [0,1]} (1 - e^{-b})(1 - b)^{k+\ell} \geq 1/(k + \ell + 1)e$.

Proof. We evaluate the left-hand side at $b = t = 1/(k + \ell + 1)$. Let $\phi(t) = (1 - t)^{(1/t)-1}$. Then, our goal is to show that $(1 - e^{-t})\phi(t) \geq t/e$ for $t = 1/3, 1/4, \dots$. This is easily verified at $t = 1/3$ by direct computation. For $t \leq 1/4$, we use

$$1 - e^{-t} \geq t - \frac{t^2}{2}. \quad (3)$$

Also, we can show that $\phi(0) = 1/e$, $\phi'(0) = 1/2e$, and $\phi''(t) \geq 7/12e$. Therefore, by Taylor's theorem,

$$\phi(t) \geq \frac{1}{e} + \frac{t}{2e} + \frac{7t^2}{24e}.$$

Thus,

$$(1 - e^{-t})\phi(t) \geq \frac{t}{e} \left(1 + \frac{t^2}{24} - \frac{7t^3}{48}\right) \geq \frac{t}{e} \square$$

Finally, we give the optimal value of b when the objective function is non-monotone.

Lemma 10. For any $k, \ell \in \mathbb{N}$,

$$\frac{1}{2}(k + \ell + 2 - \sqrt{(k + \ell)(k + \ell + 4)}) = \operatorname{argmax}_{b \in (0,1]} \{(1 - b)^{k+\ell} \cdot be^{-b}\}.$$

Proof. Let $z = k + \ell$ and set $\xi(b) = (1 - b)^z be^{-b}$. Then $\xi'(b) = 0$ iff $e^{-b}(1 - b)^{z-1}(1 - (2 + z)b + b^2) = 0$. We require that $b > 0$, implying that $b = \frac{1}{2}(z + 2 - \sqrt{z(z + 4)}) \equiv b^*$ (it could not have been the other root since $b \leq 1$). It is easy to verify that $0 < b \leq 1$. Now,

$$\xi''(b) = -\xi'(b) - (z - 1)(1 - b)^{-1}\xi'(b) + e^{-b}(1 - b)^{z-1}(2b - (2 + z)),$$

and so $\xi'(b^*) = e^{-b^*}(1 - b^*)^{z-1}(2b^* - (2 + z)) < 0$ since $e^{-b^*}, (1 - b^*)^{z-1} > 0$ and $2b^* - (2 + z) < 0$ because $2b^* \leq 2 < 2 + z$. \square

Acknowledgements We would like to thank Moran Feldman for pointing out an error in an earlier version of the paper.

References

- [1] Marek Adamczyk. Non-negative submodular stochastic probing via stochastic contention resolution schemes. *arXiv preprint arXiv:1508.07771*, 2015.

- [2] Marek Adamczyk, Maxim Sviridenko, and Justin Ward. Submodular stochastic probing on matroids. *Mathematics of Operations Research*, 41(3):1022–1038, 2016.
- [3] Marek Adamczyk and Michał Włodarczyk. Random order contention resolution schemes. *arXiv preprint arXiv:1804.02584*, 2018.
- [4] Alexander A Ageev and Maxim I Sviridenko. Pipage rounding: A new method of constructing algorithms with proven performance guarantee. *Journal of Combinatorial Optimization*, 8(3):307–328, 2004.
- [5] Gruia Calinescu, Chandra Chekuri, Martin Pál, and Jan Vondrák. Maximizing a submodular set function subject to a matroid constraint. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 182–196. Springer, 2007.
- [6] Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 783–792. ACM, 2011.
- [7] YS Chow, Sigaiti Moriguti, Herbert Robbins, and SM Samuels. Optimal selection based on relative rank (the secretary problem). *Israel Journal of mathematics*, 2(2):81–90, 1964.
- [8] Robert M Corless, Gaston H Gonnet, David EG Hare, David J Jeffrey, and Donald E Knuth. On the lambertw function. *Advances in Computational mathematics*, 5(1):329–359, 1996.
- [9] Brian C Dean, Michel X Goemans, and Jan Vondrák. Adaptivity and approximation for stochastic packing problems. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 395–404. Society for Industrial and Applied Mathematics, 2005.
- [10] Uriel Feige, Nicole Immorlica, Vahab S Mirrokni, and Hamid Nazerzadeh. Pass approximation: A framework for analyzing and designing heuristics. *Algorithmica*, 66(2):450–478, 2013.
- [11] Uriel Feige, Vahab S Mirrokni, and Jan Vondrak. Maximizing non-monotone submodular functions. *SIAM Journal on Computing*, 40(4):1133–1153, 2011.
- [12] Jon Feldman, Aranyak Mehta, Vahab Mirrokni, and S Muthukrishnan. Online stochastic matching: Beating $1-1/e$. In *Foundations of Computer Science, 2009. FOCS'09. 50th Annual IEEE Symposium on*, pages 117–126. IEEE, 2009.
- [13] Moran Feldman. *Maximization Problems with Submodular Objective Functions*. PhD thesis, Technion - Israel Institute of Technology, July 2013.
- [14] Moran Feldman, Joseph Naor, and Roy Schwartz. A unified continuous greedy algorithm for submodular maximization. In *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on*, pages 570–579. IEEE, 2011.
- [15] Moran Feldman, Ola Svensson, and Rico Zenklusen. Online contention resolution schemes. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 1014–1033. Society for Industrial and Applied Mathematics, 2016.

- [16] Anupam Gupta and Viswanath Nagarajan. A stochastic probing problem with applications. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 205–216. Springer, 2013.
- [17] Richard M Karp, Umesh V Vazirani, and Vijay V Vazirani. An optimal algorithm for on-line bipartite matching. In *Proceedings of the twenty-second annual ACM symposium on Theory of computing*, pages 352–358. ACM, 1990.
- [18] Jon Lee, Vahab S Mirrokni, Viswanath Nagarajan, and Maxim Sviridenko. Non-monotone submodular maximization under matroid and knapsack constraints. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, pages 323–332. ACM, 2009.
- [19] Sahil Singla. *Combinatorial Optimization Under Uncertainty: Probing and Stopping-Time Algorithms*. PhD thesis, Carnegie Mellon University, 2018.
- [20] Sahil Singla. The price of information in combinatorial optimization. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '18*, pages 2523–2532, Philadelphia, PA, USA, 2018. Society for Industrial and Applied Mathematics.
- [21] Maxim Sviridenko, Jan Vondrák, and Justin Ward. Optimal approximation for submodular and supermodular optimization with bounded curvature. *Mathematics of Operations Research*, 42(4):1197–1218, 2017.
- [22] Martin L Weitzman. Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, pages 641–654, 1979.

A Proof of Theorem 3

Let g be a normalized submodular function and c a non-decreasing modular function. Let G and C be their respective multilinear extensions, and let $T \geq 0$ be given. Throughout the proof we will assume that we have an oracle to evaluate G . Without such an oracle, we may evaluate G with high probability with polynomially many queries. This is standard practice (e.g., [6]).

Recall that $R = \max\{\max_{e \in \mathcal{N}} g_\emptyset(e), \max_{e \in \mathcal{N}} |\text{cost}(e)|\}$ and note that $g(OPT)$ and $|\text{cost}(OPT)|$ are both upper bounded by nR .

Let $\mathbf{y}^* \in \text{argmax}_{\mathbf{y}} \{\alpha(T)g^+(\mathbf{y}) - TC(\mathbf{y})\}$. Recall that our goal is to find a point $\mathbf{x} \in T \cdot \mathcal{P}$ such that $G(\mathbf{x}) - C(\mathbf{x}) \geq \alpha(T)g^+(\mathbf{y}^*) - TC(\mathbf{y}^*)$. Our first goal is to estimate the value of $C(\mathbf{y}^*)$. As in [21], we do this by sampling $O(\epsilon^{-1}n \log n)$ points from the interval $[-nR, nR]$. For each point θ , we will essentially run Measured Continuous Greedy over the polytope $\mathcal{P} \cap \{\mathbf{x} : C(\mathbf{x}) \leq \theta\}$. The algorithm is described formally below.

Modified Measured Continuous Greedy. Assume that $1/\epsilon \in \mathbb{N}$; otherwise decrease ϵ sufficiently. Similarly to [21], fill $[0, TR]$ with $O(\epsilon^{-1})$ points of the form $i\epsilon TR$ for $i \in \{0, 1, \dots, \epsilon^{-1}\}$, and $[TR, TnR]$ with $O(\epsilon^{-1} \log n)$ points of the form $(1 + \epsilon/n)^i \log(Tn)$ for $i \in \{0, 1, \dots, \lceil \log_{1+\epsilon/n} n \rceil\}$. For each point θ , perform the following.

Set $\mathbf{x}_\theta^0 = \mathbf{1}_\emptyset$ and assume δ is sufficiently small. Let $\mathcal{P}(\theta) = \mathcal{P} \cap \{\mathbf{x} : C(\mathbf{x}) \leq \theta\}$. For $i = k\delta$, $k = 0, 1, \dots, T/\delta - 1$ let $\mathbf{v}^i = \text{argmax}_{\mathbf{v}} \{\mathbf{v} \cdot (\nabla G(\mathbf{x}_\theta^i) \circ (1 - \mathbf{x}_\theta^i)) : \mathbf{v} \in \mathcal{P}(\theta)\}$. Define $\mathbf{x}_\theta^{i+\delta}$ by $x_{\theta,e}^{i+\delta} = x_{\theta,e}^i + \delta v_e^i (1 - x_{\theta,e}^i)$. As the final solution we return $\text{argmax}_\theta G(\mathbf{x}_\theta^T) - C(\mathbf{x}_\theta^T)$.

We emphasize the similarity of this procedure to Measured Continuous Greedy [14]. The only differences are the estimation of $C(\mathbf{y}^*)$ and restriction of the given polytope \mathcal{P} to $\mathcal{P}(\theta)$.

Analysis. For some point θ we have $\theta \leq C(\mathbf{y}^*) \leq \theta + \epsilon R$; see [21] for more details. We will perform the rest of the analysis for this value of θ . For notational simplicity let $\mathbf{x} = \mathbf{x}_\theta^T$. We want to demonstrate that

$$G(\mathbf{x}) - C(\mathbf{x}) \geq \alpha(T)g^+(\mathbf{y}^*) - TC(\mathbf{y}^*) - O(\epsilon)R$$

with high probability, and that $\mathbf{x}/T \in \mathcal{P}$. The argument of the latter fact does not change from the analysis of Measured Continuous Greedy in [14], thus we proceed to prove the former.

We begin by upper bounding C . Noting that $\mathbf{x} = \sum_{i=0}^{T/\delta-1} \delta \mathbf{v}^{i\delta} \circ (1 - \mathbf{x}^{i\delta})$ and that C is linear and non-decreasing, we have

$$C(\mathbf{x}) = \delta \sum_{i=0}^{T/\delta-1} C(\mathbf{v}^{i\delta} \circ (1 - \mathbf{x}^{i\delta})) \leq \delta \sum_{i=0}^{T/\delta-1} C(\mathbf{v}^{i\delta}) \leq T\theta,$$

since $C(\mathbf{v}^{i\delta}) \leq \theta$ for each i by construction. Therefore, $C(\mathbf{x}) = T\theta \leq TC(\mathbf{y}^*)$.

Now we proceed to lower bounding G . Here we follow the analysis of Measured Continuous Greedy as found in [13] and the modification of its proof as in [1] and [3]. We begin by assuming that g is non-monotone. We need the following lemmas from Feldman et al.

Lemma 11 ([14]). *If $|x'_e - x_e| \leq \delta$ for all $e \in \mathcal{N}$, then*

$$G(\mathbf{x}') - G(\mathbf{x}) \geq \sum_{e \in \mathcal{N}} (x'_e - x_e) \partial_e G(\mathbf{x}) - O(n^3 \delta^2) \max_e f(e). \quad (4)$$

Lemma 12 ([14]). *For all i , and $e \in \mathcal{N}$, $x_e^i \leq 1 - (1 - \delta)^i / \delta \leq 1 - e^{-i} + O(\delta)$. Moreover, if $x_e \leq a$ for every $e \in \mathcal{N}$ then for all $E \subset \mathcal{N}$, $G(\mathbf{x} \vee \mathbf{1}_E) \geq (1 - a)g(E)$.*

Let $H_e(\mathbf{x}^i) = G(\mathbf{x}_i \vee \mathbf{1}_e) - G(\mathbf{x}_i)$ and define $H_E(\mathbf{x}_i)$ similarly. Applying (4) to $\mathbf{x}^{i+\delta} - \mathbf{x}^i$ and noting that $\partial_e G(\mathbf{x}) = \frac{G(\mathbf{x} \vee \mathbf{1}_e) - G(\mathbf{x})}{1 - x_e}$ yields

$$G(\mathbf{x}^{i+\delta}) - G(\mathbf{x}^i) \geq \delta \sum_{e \in \mathcal{N}} v_e^i H_e(\mathbf{x}^i) - O(n^3 \delta^2) \max_e g(e) \geq \delta \sum_{e \in \mathcal{N}} y_e^* H_e(\mathbf{x}^i) - O(n^3 \delta^2) \max_e g(e), \quad (5)$$

where the second inequality uses the definition of \mathbf{v}^i . Let $(p_E)_{E \subset \mathcal{N}}$ be the maximizing argument of $g^+(\mathbf{y}^*)$, i.e., $g^+(\mathbf{y}^*) = \sum_{E \subset \mathcal{N}} p_E g(E)$. Then,

$$\begin{aligned} \sum_{e \in \mathcal{N}} y_e^* H_e(\mathbf{x}_i) &= \sum_{E \subset \mathcal{N}} p_E \sum_{e \in E} H_e(\mathbf{x}_i) \geq \sum_{E \subset \mathcal{N}} p_E H_E(\mathbf{x}_i) \\ &\geq \left(\sum_{E \subset \mathcal{N}} p_E g(E) (e^{-i} - O(\delta)) \right) - G(\mathbf{x}^i) = (e^{-i} - O(\delta))g^+(\mathbf{y}^*) - G(\mathbf{x}^i), \end{aligned} \quad (6)$$

where the first inequality is due to monotonicity. Now, notice that $\max_e g(e) \leq \max_{E \subset \mathcal{N}} g(E) = \max_{E \subset \mathcal{N}} g^+(\mathbf{1}_E) \leq g^+(\mathbf{z}^*)$ where $\mathbf{z}^* \in \operatorname{argmax}_{\mathbf{z}} \{g^+(\mathbf{z}) : \mathbf{z} \in \mathcal{P}\}$ (because $\mathbf{1}_E \in \mathcal{P}$). By

definition of \mathbf{y}^* we have $\alpha(T)g^+(\mathbf{y}^*) - T\beta C(\mathbf{y}^*) \geq \alpha(T)g^+(\mathbf{z}^*) - T\beta C(\mathbf{z}^*)$. Moreover, $C(\mathbf{z}^*) = \sum_{e \in \mathcal{N}} \Delta_e z_e^* \leq \Delta$ where $\Delta = \sum_{e \in \mathcal{N}} \Delta_e$. Hence,

$$\max_e g(e) \leq g^+(\mathbf{z}^*) \leq g^+(\mathbf{y}^*) - T\beta g^+(\mathbf{y}^*) + T\beta\Delta \leq g^+(\mathbf{y}^*) + T\beta\Delta$$

Setting $\kappa = T\beta\Delta$, and combining the above with (5) and (6) gives

$$G(\mathbf{x}^{i+\delta}) - G(\mathbf{x}^i) \geq \delta(e^{-i}g^+(\mathbf{y}^*) - G(\mathbf{x}^i)) - O(n^3\delta^2)g^+(\mathbf{y}^*) - O(n^3\delta^2)\kappa. \quad (7)$$

As noted in [1], the analysis of [13] relies on demonstrating that for all i

$$G(\mathbf{x}^{i+\delta}) - G(\mathbf{x}^i) \geq \delta[(e^{-i} \cdot g(OPT)) - G(\mathbf{x}^i)] - O(n^3\delta^2)g(OPT). \quad (8)$$

The analysis after this point does not use any properties of $g(OPT)$, and can be replaced by any constant. In our case, we will replace it with $g^+(\mathbf{y}^*)$. Additionally (and crucially in our case) the analysis uses no properties of the vector $v_e(t)$, meaning that it is not affected by working over the polytope $\mathcal{P}(\theta)$. Therefore, continuing with their analysis, but tacking on the term $O(n^3\delta^2)\kappa$ onto the end, eventually yields that

$$G(\mathbf{x}^T) \geq [Te^{-T} - O(n^3\delta^2)]g^+(\mathbf{y}^*) - O(n^3\delta^2)\kappa = (Te^{-T} - o(1))g^+(\mathbf{y}^*) - o(1),$$

if δ is chosen sufficiently small. Combining this inequality with (8) gives the desired result, and completes the proof if g is non-monotone. If g is monotone, then we may strengthen the bound with precisely the same techniques used in [13].