Online Contention Resolution Schemes with Applications to Bayesian Selection Problems

Journal Article**Author(s):**

Feldman, Moran; Svensson, Ola; Zenklusen, Rico

Publication date:

2021

Permanent link:

<https://doi.org/10.3929/ethz-b-000489010>

Rights / license:

In Copyright - Non-Commercial Use Permitted

Originally published in:

SIAM Journal on Computing 50(2), <https://doi.org/10.1137/18m1226130>

Funding acknowledgement:

165866 - New Approaches to Constrained Submodular Maximization (SNF)

184622 - Toward Stronger Approximation Algorithms for Fundamental Network Design and Optimization Problems (SNF)

817750 - Fundamental Problems at the Interface of Combinatorial Optimization with Integer Programming and Online Optimization (EC)

Online Contention Resolution Schemes*

Moran Feldman[†]

Ola Svensson[‡]

Rico Zenklusen[§]

Abstract

We introduce a new rounding technique designed for online optimization problems, which is related to contention resolution schemes, a technique initially introduced in the context of submodular function maximization. Our rounding technique, which we call *online contention resolution schemes* (OCRSs), is applicable to many online selection problems, including Bayesian online selection, oblivious posted pricing mechanisms, and stochastic probing models. It allows for handling a wide set of constraints, and shares many strong properties of offline contention resolution schemes. In particular, OCRSs for different constraint families can be combined to obtain an OCRS for their intersection. Moreover, we can approximately maximize submodular functions in the online settings we consider.

We, thus, get a broadly applicable framework for several online selection problems, which improves on previous approaches in terms of the types of constraints that can be handled, the objective functions that can be dealt with, and the assumptions on the strength of the adversary. Furthermore, we resolve two open problems from the literature; namely, we present the first constant-factor constrained oblivious posted price mechanism for matroid constraints, and the first constant-factor algorithm for weighted stochastic probing with deadlines.

Keywords: contention resolution schemes, online algorithms, matroids, prophet inequalities, oblivious posted pricing, stochastic probing

*This work appeared in SIAM Journal on Computing (2021) 50(2):255–300, <https://doi.org/10.1137/18M1226130>. Copyright © by SIAM.

[†]School of Computer and Communication Sciences, EPFL. Email: moranfe3@gmail.com. Supported by ERC Starting Grant 335288-OptApprox.

[‡]School of Computer and Communication Sciences, EPFL. Email: ola.svensson@epfl.ch. Supported by ERC Starting Grant 335288-OptApprox.

[§]Department of Mathematics, ETH Zurich. Email: ricoz@ethz.ch. Supported by ERC Consolidator grant 817750-ICOPT under the European Union’s Horizon 2020 research and innovation program.

1 Introduction

Recently, interest has surged in Bayesian and stochastic online optimization problems. These are problems where we can use limited a priori information to select elements arriving online, often subject to classical combinatorial constraints such as matroids, matchings and knapsack constraints. Examples include posted pricing mechanisms [8, 24, 18], prophet inequalities [18], probing models [17, 1], stochastic matchings [3], and secretary problems [2, 20, 15].¹ Simultaneous with this development, interest has arose also in generalizing the optimization of linear objective functions to relevant nonlinear objective functions. A particular focus was set on submodular functions, which is a function class that captures the property of diminishing returns, a very natural property in many of the above-mentioned settings [21, 1, 4, 16].

A very successful approach for these problems is based on first using the a priori information to formulate an (often linear) relaxation whose optimal fractional solution x^* upper bounds the performance of any online (or even offline) algorithm. Then, x^* is used to devise an online algorithm whose goal is to recover a solution of a similar objective value as x^* . Such an online algorithm can also be interpreted as an online rounding procedure for rounding x^* . In particular, online rounding approaches have recently been used to obtain nearly optimal and surprisingly elegant results for stochastic matchings (see Bansal et al. [3]), and for a very general probing model introduced by Gupta and Nagarajan [17] with applications in posted pricing mechanisms, online matching problems and beyond.

A key ingredient in the general rounding algorithms presented in [17] are so-called contention resolution schemes (CRSs), a rounding technique introduced by Chekuri et al. [10] in the context of (offline) submodular function maximization. CRSs are defined with respect to a constraint family, like matroids, matching, or knapsack constraints. Interestingly, the existence of a so-called *ordered* CRS for the given constraint family is all that is needed to apply the techniques of [17]. Whereas this generality is very appealing, there are some inherent barriers in current CRSs that hinder a broader applicability to online settings beyond the probing model defined in [17]. More precisely, most settings considered in [17] require that the *algorithm can choose the order* in which to obtain new online information about an underlying ground set over which the objective is optimized. This is due to the fact that most CRSs need to round the components of a fractional point x^* step by step in a particular order.

In this paper we introduce a stronger notion of contention resolution schemes that overcomes this restriction and allows for the online information to be presented *adversarially*. We show that such schemes exist for many interesting constraint families, including matroid constraints and knapsack constraints. As we discuss in Section 1.3, this leads to a broadly applicable online rounding framework that works in considerably more general settings than previous approaches. Furthermore, our techniques answer two open problems from the literature: we show the existence of a constrained oblivious posted-pricing mechanism (COPM) for matroids, intersection of matroids, and further constraints families (a question first raised in [8]); and we get an $O(1)$ -competitive algorithm for the weighted probing problem with deadlines introduced in [17]. Additionally, our rounding approach yields optimal guarantees (up to moderate constant factors) for a class of online submodular function maximization problems.

Before we formally define our rounding framework (Section 1.1), state our results (Section 1.2) and describe the aforementioned applications (Section 1.3), it is helpful to introduce our rounding framework in the light of a concrete example. Consider the following Bayesian online selection problem studied by Kleinberg and Weinberg [18] in the context of prophet inequalities. There is a finite set N of items or elements, and a nonnegative random variable Z_e for each $e \in N$, where all $\{Z_e\}_{e \in N}$ are independent. The distributions of all Z_e are known and for simplicity we assume they are continuous. Furthermore, a matroid

¹Strictly speaking, secretary problems have no a priori information. However, as items arrive in a random order, most algorithms first observe a fraction of the elements (serving as the a priori information), and then devise an online strategy based on this information.

$M = (N, \mathcal{F})$ on N is given.² Let $\{z_e\}_{e \in N}$ be realizations of the random variables $\{Z_e\}_{e \in N}$. The goal is to select a subset $I \subseteq N$ of the elements that is independent, i.e., $I \in \mathcal{F}$, and whose value $z(I) := \sum_{e \in I} z_e$ is as large as possible. The way how elements can be selected works as follows. Elements $e \in N$ reveal one by one their realization z_e , in a fixed prespecified order that is unknown to the algorithm. Whenever a value z_e is revealed, one has to choose whether to select e or discard it, before the next element reveals its realization.

A natural way to approach this problem is to define a threshold $t_e \geq 0$ for each $e \in N$ and only accept elements $e \in N$ whose realization is at least the threshold, i.e., $z_e \geq t_e$; we call such elements *active*. Let x_e be the probability of e being active, i.e., $x_e = \Pr[Z_e \geq t_e]$. Notice that the set of all active elements is distributed like a random set that contains each element e independently with probability x_e . We denote such a set by $R(x)$. As we show in Section 4, using a convex relaxation one can find thresholds t_e such that: (i) $x = (x_e)_{e \in N}$ is in the matroid polytope $P_{\mathcal{F}}$,³ and (ii) an algorithm that disregards the matroid constraint and accepts any active element would have an expected return at least as good as the one of an optimal offline algorithm.

Our goal is to design an online algorithm that only selects active elements, such that an independent set I is obtained where $\Pr[e \in I] \geq c \cdot x_e$ for all $e \in E$, where $c \in (0, 1]$ is a constant as large as possible. It is not hard to check that such a procedure would lead to an objective value of at least c times the offline optimum. The guarantee we are seeking closely resembles the notion of c -balanced CRSs as defined in [10], which is an *offline* algorithm that depends on x and returns for any set $S \subseteq N$ a (potentially random) subset $\pi(S) \subseteq S$ with $\pi(S) \in \mathcal{F}$ such that $\Pr[e \in \pi(R(x))] \geq c \cdot x_e$. The only reason why this procedure is not applicable in the above context is that, in general, π needs to know the realization of the full set $R(x)$ in advance to determine $\pi(R(x))$. However, $R(x)$ is revealed element by element in the above selection problem. A key observation in [17] is that some CRSs do not need to know the full set $R(x)$ upfront, but can round step by step if the elements come in some prescribed order chosen by the algorithm. However, in the above setting, as in many other combinatorial online problems, the order cannot be chosen freely.

We overcome this limitation through a considerably stronger notion of CRSs, which we call *online contention resolution schemes* (OCRSs).

1.1 Online contention resolution schemes

OCRSs, like classical contention resolution schemes, are defined with respect to a relaxation of the feasible sets of a combinatorial optimization problem. Consider a finite ground set $N = \{e_1, \dots, e_n\}$, and a family of *feasible* subsets $\mathcal{F} \subseteq 2^N$, which is down-closed, i.e., if $I \in \mathcal{F}$ and $J \subseteq I$ then $J \in \mathcal{F}$. Let $P_{\mathcal{F}} \subseteq [0, 1]^N$ be the polytope corresponding to the feasible sets $\mathcal{F}$, i.e., $P_{\mathcal{F}}$ is the convex hull of all characteristic vectors of feasible sets:

$$P_{\mathcal{F}} = \text{conv}(\{\mathbf{1}_I \mid I \in \mathcal{F}\}) .$$

We highlight that throughout this paper we focus on down-closed feasibility constraints.

Definition 1.1 (relaxation). *We say that a polytope $P \subseteq [0, 1]^N$ is a relaxation of $P_{\mathcal{F}}$ if it contains the same $\{0, 1\}$ -points, i.e., $P \cap \{0, 1\}^N = P_{\mathcal{F}} \cap \{0, 1\}^N$.*

We start by defining online contention resolution schemes (OCRS) simply as algorithms that can be applied to the online selection problem highlighted in the introduction. The performance of an OCRS is then characterized by additional properties that we define later.

²We recall that a matroid $M = (N, \mathcal{F})$ consists of a finite ground set N and a nonempty family $\mathcal{F} \subseteq 2^N$ of subsets of N satisfying: (i) If $I \in \mathcal{F}$, $J \subseteq I$, then $J \in \mathcal{F}$, and (ii) if $I, J \in \mathcal{F}$, $|J| > |I|$, then $\exists e \in J \setminus I$ with $I \cup \{e\} \in \mathcal{F}$. If not stated otherwise, we assume that matroids are given by an independence oracle that, for every $I \subseteq N$, returns whether $I \in \mathcal{F}$.

³The matroid polytope $P_{\mathcal{F}}$ is the convex hull of all characteristic vectors of independent sets. In particular it can be described by $P_{\mathcal{F}} = \{x \in \mathbb{R}_{\geq 0}^N \mid x(S) \leq \text{rank}(S) \forall S \subseteq N\}$, where $\text{rank}(S) = \max\{|I| \mid I \subseteq S, I \in \mathcal{F}\}$ is the *rank function* of M .

Definition 1.2 (Online contention resolutions scheme (OCRS)). *Let us consider the following online selection setting. A point $x \in P$ is given and let $R(x)$ be a random subset of active elements. The elements $e \in N$ reveal one by one whether they are active, i.e., $e \in R(x)$, and the decision whether to select an active element is taken irrevocably before the next element is revealed. An OCRS for P is an online algorithm that selects a subset $I \subseteq R(x)$ such that $\mathbf{1}_I \in P$.*

Most of the OCRSs that we present follow a common algorithmic theme, which leads us to the definition of *greedy OCRS*.

Definition 1.3 (Greedy OCRS). *Let $P \subseteq [0, 1]^N$ be a relaxation for the feasible sets $\mathcal{F} \subseteq 2^N$. A greedy OCRS π for P is an OCRS that for any $x \in P$ defines a down-closed subfamily of feasible sets $\mathcal{F}_{\pi, x} \subseteq \mathcal{F}$, and an element e is selected when it arrives if, together with the already selected elements, the obtained set is in $\mathcal{F}_{\pi, x}$.*

If the choice of $\mathcal{F}_{\pi, x}$ given x is randomized, we talk about a randomized greedy OCRS; otherwise, we talk about a deterministic greedy OCRS. We also simplify notation and abbreviate $\mathcal{F}_{\pi, x}$ by $\mathcal{F}_x$ when the OCRS π is clear from the context.

For simplicity of presentation, and because all our main results are based on greedy OCRSs, we restrict our attention to this class of OCRSs, and focus on greedy OCRSs when defining and analyzing properties.

As mentioned in the example shown in the introduction, a desirable property of OCRSs would be that every element $e \in N$ gets selected with probability at least $c \cdot x_e$ for a constant $c > 0$ as large as possible. This property is called *c-balancedness* in the context of classical contention resolution schemes. However, to be precise about such properties in the online context that we consider, one has to specify the power of the adversary who chooses the order of the elements. Adversaries of different strengths have been considered in various online settings. For example, one arguably weak type of adversary is an *offline adversary*, who has to choose the order of the elements upfront, before any elements get revealed. On the other end, the most powerful adversary that can be considered is what we call the *almighty adversary*; an almighty adversary knows upfront the outcomes of all random events, which includes the realization of $R(x)$ and the outcome of the random bits that the algorithm may query. An almighty adversary can thus calculate exactly how the algorithm will behave and reveal the elements in a worst case order. A typical adversary type that is in between these two extremes is an *online adversary*, who can choose the next element to reveal online depending on what happened so far; thus, it has the same information as the online algorithm. Throughout this paper, when not indicated otherwise, we assume to play against the almighty adversary.

In the context of greedy OCRSs, we define a considerably stronger notion than *c-balancedness*, which we call *c-selectability*, and which leads to results even against the almighty adversary. In words, a greedy OCRS is *c-selectable* if with probability at least c , the random set $R(x)$ is such that an element e is selected no matter what other elements I of $R(x)$ have been selected so far as long as $I \in \mathcal{F}_x$. Thus, it guarantees that an element is selected with probability at least c against any (even the almighty) adversary.

Definition 1.4 (*c-selectability*). *Let $c \in [0, 1]$. A greedy OCRS for P is *c-selectable* if for any $x \in P$ we have*

$$\Pr[I \cup \{e\} \in \mathcal{F}_x \mid \forall I \subseteq R(x), I \in \mathcal{F}_x] \geq c \quad \forall e \in N.$$

We highlight that the probability in Definition 1.4 is over the random outcomes of $R(x)$ when dealing with a deterministic greedy OCRS; when the greedy OCRS is randomized, then the probability is over $R(x)$ and the random choice of $\mathcal{F}_x$. We call an element $e \in N$ *selectable* for a particular realization of $R(x)$ and random choice of $\mathcal{F}_x$ if $I \cup \{e\} \in \mathcal{F}_x$ for all $I \subseteq R(x)$ with $I \in \mathcal{F}_x$.

As aforementioned, the *c-selectability* is a very strong property that implies guarantees against any adversary. Despite this strong definition, we show that $\Omega(1)$ -selectable greedy OCRSs exist for many natural constraints.

Often, a larger factor c can be achieved when x is supposed to be in a down-scaled version of P . This is similar to the situation in classical contention resolution schemes.

Definition 1.5 ((b, c) -selectability). *Let $b, c \in [0, 1]$. A greedy OCRS for P is (b, c) -selectable if for any $x \in b \cdot P$ we have*

$$\Pr[I \cup \{e\} \in \mathcal{F}_x \mid \forall I \subseteq R(x), I \in \mathcal{F}_x] \geq c \quad \forall e \in N.$$

Notice that a (b, c) -selectable greedy OCRS implies a randomized bc -selectable greedy OCRS because we can “scale down” x online by only considering each element e with probability b independent of the other elements.

Observation 1.6. *A (b, c) -selectable greedy OCRS for P implies a (randomized) bc -selectable greedy OCRS for P .*

The existence of OCRSs is interesting even regardless of efficiency issues. Still, in many applications it is important to have efficient OCRSs.

Definition 1.7 (efficiency). *A greedy OCRS π is efficient if there exists a polynomial time algorithm that, for a given input x , samples an efficient independence oracle for the set $\mathcal{F}_{\pi, x}$. That is, an oracle that answers in polynomial time queries of the form: is a set $S \subseteq N$ in $\mathcal{F}_{\pi, x}$?*

We next summarize our technical results before highlighting the implications of our results to various online settings.

1.2 Our results

Our first technical result proves the existence of greedy OCRSs with constant selectability for relaxations of several interesting families of constraints. All the greedy OCRSs described by Theorem 1.8 are either efficient, or can be made efficient at the cost of an arbitrarily small constant $\varepsilon > 0$ loss in the selectability guarantee.

Theorem 1.8. *There exist:*

- *For every $b \in [0, 1]$, a $(b, 1 - b)$ -selectable deterministic greedy OCRS for matroid polytopes.*
- *For every $b \in [0, 1]$, a (b, e^{-2b}) -selectable randomized greedy OCRS for matching polytopes.⁴*
- *For every $b \in [0, 1/2]$, a $(b, \frac{1-2b}{2-2b})$ -selectable randomized greedy OCRS for the natural relaxation of a knapsack constraint.*

Interestingly, it turns out that there is no (b, c) -selectable *deterministic* greedy OCRS for the natural relaxation of a knapsack constraint for any constants b and c . This stands in contrast to the case of the matching polytope, for which the randomized greedy OCRS given by Theorem 1.8 can be made deterministic at the cost of only a small loss in the selectability.

Like offline CRSs, greedy OCRSs can be combined to form greedy OCRSs for more involved constraints.

Theorem 1.9. *If π^1 is a (b, c_1) -selectable greedy OCRS for a polytope P_1 , and π^2 is a (b, c_2) -selectable greedy OCRS for a polytope P_2 , then there exists a $(b, c_1 \cdot c_2)$ -selectable greedy OCRS for the polytope $P_1 \cap P_2$. Moreover, the last greedy OCRS is efficient if π^1 and π^2 are.*

⁴Our greedy OCRS works also for a weaker relaxation of matching which only bounds the degree of each node.

Notice that Theorem 1.9 can be applied repeatedly to combine several OCRSs. Thus, Theorems 1.8 and 1.9 prove together the existence of constant selectability greedy OCRSs for any constant intersection of matroid, matching and knapsack constraints.

It is easy to see that, given a non-negative increasing linear objective function, a (b, c) -selectable greedy OCRS for a polytope P can be used to round online a vector $x \in bP$ while losing only a factor of c in the objective. Theorem 1.10 proves this result formally, and extends it to nonnegative submodular⁵ functions. To state this theorem, we need to define some notation. Given a function $f: 2^N \rightarrow \mathbb{R}$, the *multilinear* extension of f is a function $F: [0, 1]^N \rightarrow \mathbb{R}$ whose value for a vector $x \in [0, 1]^N$ is $F(x) = \mathbb{E}[f(R(x))]$. Informally, $F(x)$ is the expected value of f over a set obtained by randomly rounding every coordinate of x independently.

Theorem 1.10. *Given a nonnegative monotone⁶ submodular function $f: 2^N \rightarrow \mathbb{R}_{\geq 0}$ and a (b, c) -selectable greedy OCRS for a polytope P , applying the greedy OCRS to an input $x \in bP$ results in a random set S satisfying $\mathbb{E}[f(S)] \geq c \cdot F(x)$, where F is the multilinear extension of f . Moreover, even if f is not monotone, $\mathbb{E}[f(R(1/2 \cdot \mathbf{1}_S))] \geq (c/4) \cdot F(x)$, where the random decisions used to calculate $R(1/2 \cdot \mathbf{1}_S)$ are considered part of the algorithm, and thus, known to the almighty adversary.*

In many applications the use of Theorem 1.10 requires finding *offline*, using the available a priori information, a vector x (approximately) maximizing the multilinear extension F . This can often be done using known algorithms. For example, Călinescu et al. [7] proved that given a non-negative monotone submodular function $f: 2^N \rightarrow \mathbb{R}_{\geq 0}$ and a solvable⁷ polytope $P \subseteq [0, 1]^N$, one can efficiently find a fractional point $x \in P$ for which $F(x) \geq (1 - e^{-1}) \cdot \max\{f(S) \mid \mathbf{1}_S \in P\}$. Chekuri et al. [10] showed that even when f is not monotone, an analogous result can be obtained with a worse constant factor of 0.325 instead of $1 - e^{-1}$ when P is solvable and down-closed. A simpler procedure with a stronger constant factor was later presented by Feldman et al. [14], implying that one can efficiently find a fractional point $x \in P$ for which $F(x) \geq (e^{-1} - o(1)) \cdot \max\{f(S) \mid \mathbf{1}_S \in P\}$ as long as the polytope P is solvable and down-closed.

The result of Theorem 1.10 for a non-monotone submodular objective can sometimes be improved when assuming an online adversary (instead of an almighty one). The class of OCRSs for which this can be done is a bit involved to define, and we defer its definition to Section 3. We state here only the following special case of the result we prove.

Theorem 1.11. *Let π be a (b, c) -selectable greedy OCRS π for a polytope P that was obtained by using Theorem 1.9 to combine the OCRS of Theorem 1.8. Then, for every given non-negative submodular function $f: 2^N \rightarrow \mathbb{R}_{\geq 0}$ there exists an OCRS π' for P that for every input vector $x \in bP$ and online adversary selects a random set S such that $\mathbb{E}[f(S)] \geq c \cdot F(x)$.*

The OCRS π' guaranteed by Theorem 1.11 is not efficient. However, if π is efficient then π' can be made efficient at the cost of an additive loss of $|N|^{-d} \cdot \max\{f(\{e\}) \mid x_e > 0\}$ in the guarantee (where d is any positive constant).

1.3 Applications

In this section we present a few applications for our technical results. All these applications were previously studied in the literature, and connections have been found between them. In this work we show that all three applications can be reduced to finding appropriate OCRSs. In addition to proving new results, we believe that these reductions into one common setting clarify the connections between the three applications.

⁵A set function $f: 2^N \rightarrow \mathbb{R}$ is submodular if $f(A) + f(B) \geq f(A \cup B) + f(A \cap B)$ for every two sets $A, B \subseteq N$.

⁶A set function $f: 2^N \rightarrow \mathbb{R}$ is monotone if $f(A) \leq f(B)$ for every two sets $A \subseteq B \subseteq N$.

⁷A polytope is *solvable* if one can optimize linear functions over it.

Prophet inequalities for Bayesian online selection problems

Consider again the Bayesian online selection problem we sketched earlier in the introduction. We recall that the setting in this problem consists of a matroid $M = (N, \mathcal{F})$ and independent non-negative random variables Z_e for every $e \in N$ with known distributions. Moreover, the random variables Z_e satisfy $\max_{e \in N} \mathbb{E}[Z_e] < \infty$. An offline adversary chooses upfront the order in which the elements $e \in N$ reveal a realization z_e of Z_e . The task is to select online an independent set of elements $I \in \mathcal{F}$ with total weight $z(I) = \sum_{e \in I} z_e$ as high as possible.

A fundamental result about the relative power of offline and online algorithms in a Bayesian setting was obtained by Krengel, Sucheston and Garling (see [19]) for the special case when M is the uniform matroid of rank one, *i.e.*, precisely one element can be selected. They showed that there exists a selection algorithm returning a single element of expected weight at least $\frac{1}{2} \mathbb{E}[\max_{e \in N} Z_e]$, *i.e.*, half of the weight of the best offline solution, which is the best solution obtainable by an algorithm that knows all realizations upfront. Recently, Kleinberg and Weinberg [18] extended this result considerably by showing that the same guarantee can be obtained when selecting multiple elements that have to be independent in the matroid M , *i.e.*, there exists an online algorithm returning a set $I \in \mathcal{F}$ satisfying

$$\mathbb{E} \left[\sum_{e \in I} Z_e \right] \geq \frac{1}{2} \mathbb{E} \left[\max \left\{ \sum_{e \in S} Z_e \mid S \in \mathcal{F} \right\} \right]. \quad (1)$$

Inequalities of type (1) are often called *prophet inequalities* due to the interpretation of the offline adversary as a prophet. Moreover, Kleinberg and Weinberg generalized their result to the setting where $\mathcal{F}$ are the common independent sets in the intersection of p matroids. For this setting, they present an online algorithm whose expected profit is at least $\frac{1}{4p-2}$ times the expected maximum weight of a feasible set. Kleinberg and Weinberg's algorithms work not just against an offline adversary, which is the adversary type typically assumed in Bayesian online selection, but also against an online adversary.

Using a simple, yet very general link between greedy OCRSs and prophet inequalities we can generate prophet inequalities from greedy OCRSs.

Theorem 1.12. *Let $\mathcal{F} \subseteq 2^N$ be a down-closed set family and P be a relaxation of $\mathcal{F}$. If there exists a c -selectable greedy OCRS for P then there is an online algorithm for the Bayesian online selection problem with almighty adversary that returns a set $I \in \mathcal{F}$ satisfying*

$$\mathbb{E} \left[\sum_{e \in I} Z_e \right] \geq c \cdot \mathbb{E} \left[\max \left\{ \sum_{e \in S} Z_e \mid S \in \mathcal{F} \right\} \right].$$

As we discuss in Section 4, the above theorem can be made constructive in many cases, assuming that the OCRS is efficient and some natural optimization problems involving the distributions of the random weights Z_e can be solved efficiently.

Our results show that constant-factor prophet inequalities are often possible even against an almighty adversary. Moreover, we get $\Theta(1)$ -factor prophet inequalities for a wide set of new constraint families.

Corollary 1.13. *There are $\Theta(1)$ -factor prophet inequalities for the Bayesian online selection problem against the almighty adversary for any constraint family that is an intersection of a constant number of matroid, knapsack, and matching constraints.*

In contrast, so far, the most general prophet inequality was the $\frac{1}{4p-2}$ -factor prophet inequality of Kleinberg and Weinberg for the intersection of p matroids. Interestingly, even for this specific setting of the intersection of p matroids, considered by Kleinberg and Weinberg, our general approach allows for obtaining a better constant for $p \geq 4$ (against a stronger adversary).

Corollary 1.14. *There is an $\frac{1}{e(p+1)}$ -factor prophet inequality for the Bayesian online selection problem against the almighty adversary when the feasible sets are described by the intersection of p matroids.*

That our results hold against an almighty adversary is in particular of importance for applications of prophet inequalities to posted pricing mechanisms. Indeed, one of the main technical difficulties that Kleinberg and Weinberg [18] had to overcome to apply their results to posted pricing mechanism, was the fact that their results were only with respect to an online adversary.

Oblivious posted pricing mechanisms

We start by introducing the Bayesian single-parameter mechanism design setting (short BSMD), largely following [8]. There is a single seller providing a set N of services, and for each service $e \in N$ there is one agent interested in e , whose valuation is drawn from a nonnegative random variable Z_e . The Z_e are independent and have known distributions. Furthermore, there is a down-closed family $\mathcal{F} \subseteq 2^N$ representing feasibility constraints faced by the seller, *i.e.*, the seller can provide any set of services $S \in \mathcal{F}$. The setting is called *single-parameter* because every agent is interested in precisely one service. The goal in this setting is to find truthful mechanisms maximizing the expected revenue.

From a theoretical point of view, this setting is well understood and optimally solved by Myerson's mechanism [22]. Unfortunately, the resulting mechanism is impractical, and thus, rarely employed. Furthermore, it does not extend to multi-parameter settings where an agent may, for example, be interested in buying one out of several items, a setting known as *Bayesian multi-parameter unit-demand mechanism design* (BMUMD). Therefore, Chawla et al. [8] suggested considerably simpler and more robust alternatives having many advantages while maintaining an almost optimal performance [8, 24, 18]. The idea is to offer to the agents sequentially take-it-or-leave-it prices as follows. Agents are considered one by one, in an order chosen by the algorithm. Whenever an agent $e \in N$ is considered, the algorithm either makes no offer to e —and thus e does not get served—or, if e can be added to the elements selected so far without violating feasibility, an offer $p_e \in \mathbb{R}_{\geq 0}$ is made to e . Agent e will then accept the offer if $Z_e \geq p_e$ and decline if $Z_e < p_e$.

This type of mechanism, with the additional freedom that the algorithm can choose the order in which to consider the agents, is called a *sequential posted price mechanism*. A natural stronger version of sequential posted price mechanisms, called *constrained oblivious posted price mechanisms* (COPM), suggested by Chawla et al. [8], allows for dealing with the multi-parameter setting BMUMD, and has many further interesting properties. Formally, a COPM is defined by a tuple $(p \in \mathbb{R}_{\geq 0}^E, \mathcal{F}')$, where p are the take-it-or-leave-it prices, and $\mathcal{F}' \subseteq \mathcal{F}$. A COPM defined by $(p, \mathcal{F}')$ works as follows. Consider the moment when a new agent e arrives and let S be the set of agents served so far. If $S \cup \{e\} \notin \mathcal{F}'$, then e is skipped; otherwise, e is offered the price p_e . In short the COPM maintains a feasible set in the more restricted family $\mathcal{F}'$, and greedily selects any agent e that does not destroy feasibility in $\mathcal{F}'$ and has a valuation of at least p_e . Furthermore, the order of the agents is chosen by an adversary at the beginning of the procedure, knowing all valuations z_e , the prices p_e and the family $\mathcal{F}'$. A COPM can also be randomized, in which case the tuple $(p, \mathcal{F}')$ is chosen at random at the beginning of the algorithm.

So far, COPMs with an $O(1)$ gap with respect to the optimal mechanism were only known for very restricted types of matroids, and the intersection of two partition matroids [8]. For general matroids, the best previously known COPM was a non-efficient procedure with an optimality gap of $O(\log(\text{rank}))$, where rank is the size of a largest feasible set of the matroid [8]. In particular, the existence of an COPM for general matroids with constant optimality gap remained an open problem.

Exploiting a link between greedy OCRSs and COPMs, we resolve the open question about $O(1)$ -optimal COPMs for matroids raised in [8], and show that even much more general constraint families admit $O(1)$ -optimal COPMs.

Theorem 1.15. *Let $\mathcal{F} \subseteq 2^N$ be a down-closed family and P be a relaxation of $\mathcal{F}$. If there is a c -selectable greedy OCRS for P , then there is a COPM for $\mathcal{F}$ that, even against an almighty adversary, is at most a factor of c worse than the optimal truthful mechanism.*

Using the reduction from the multi-parameter setting to the single-parameter setting presented in [8] we obtain results for BMUMD under very general feasibility constraints.

Corollary 1.16. *Let $\mathcal{F}$ be the intersection of a constant number of matroid, knapsack, and matching constraints. Then there is a posted price mechanism for BMUMD on $\mathcal{F}$ whose optimality gap with respect to the optimal truthful mechanism is at most a constant.*

Moreover, as we highlight in Section 4, the mechanisms obtained through our greedy OCRSs can be implemented efficiently under mild assumptions.

Stochastic probing

Recently, Gupta and Nagarajan [17] introduced the following stochastic probing model. Given is a finite ground set N and each element $e \in N$ is *active* with a given probability $p_e \in [0, 1]$, independently of the other elements. Furthermore, there is a weight function $w : N \rightarrow \mathbb{Z}$, and two down-closed constraint families $\mathcal{F}_{in}, \mathcal{F}_{out} \subseteq 2^N$ on N , which are called the *inner* and *outer constraints*, respectively. The goal is to select a subset of active elements of high weight according to the following rules. An algorithm must first *probe* an element $e \in N$ to select it. If a probed element e is active, then e is selected, otherwise it is not. The algorithm can choose the order in which elements are probed. The set Q of all probed elements must satisfy $Q \in \mathcal{F}_{out}$ and the set $S \subseteq Q$ of all selected elements must satisfy $S \in \mathcal{F}_{in}$. Hence, at any step of the algorithm an element can only be probed if adding it to the currently probed elements does not violate $\mathcal{F}_{out}$, and adding it to the elements selected so far does not violate $\mathcal{F}_{in}$. Gupta and Nagarajan [17] showed that this model captures numerous applications. Furthermore, they show how so-called *ordered CRSs* can be used to get approximately optimal approximation factors for many constraint families including the intersection of a constant number of matroids. However, due to their use of ordered CRSs, the presented algorithms crucially rely on the fact that the order in which elements are probed can be chosen freely.

Replacing their use of ordered CRSs with OCRSs we can drop this requirement.

Theorem 1.17. *Let $\mathcal{F}_{in}, \mathcal{F}_{out} \subseteq 2^N$ be two down-closed families. If there are a (b, c_{in}) -selectable greedy OCRS π_{in} for a relaxation P_{in} of $\mathcal{F}_{in}$ and a (b, c_{out}) -selectable greedy OCRS π_{out} for a relaxation P_{out} of $\mathcal{F}_{out}$, then there is a $(b \cdot c_{in} \cdot c_{out})$ -approximation for the weighted stochastic probing problem where the order in which elements can be probed is chosen by an almighty adversary and the inner and outer constraints are given by $\mathcal{F}_{in}$ and $\mathcal{F}_{out}$, respectively.⁸ Moreover, if π_{in} and π_{out} are efficient and there are separation oracles for P_{in} and P_{out} , then the above algorithm has a polynomial time complexity.*

It turns out that the extension to arbitrary probing orders resolves an open question of [17] about stochastic probing with deadline. In a probing problem with deadlines there is a deadline $d_e \in \mathbb{Z}_{\geq 1}$ for each element $e \in N$, indicating that e can only be probed as one of the first d_e elements that get probed. Using a clever technique, Gupta and Nagarajan [17] presented an $O(1)$ -approximation for this problem setting for the unweighted case, i.e., w is the all-ones vector, when $\mathcal{F}_{in}, \mathcal{F}_{out}$ are k -systems⁹, for $k = O(1)$. They left it as an open question how to approach the weighted version of stochastic probing with deadlines. Using OCRSs

⁸Similar to [17] one can strengthen the theorem, and only assume an offline CRS for $\mathcal{F}_{out}$ and an OCRS for $\mathcal{F}_{in}$. The fact that an offline CRS suffices for $\mathcal{F}_{out}$ can sometimes be used to get better approximation factors. For simplicity of presentation, we do not go into these details here, and also in Theorem 1.18 and 1.19.

⁹A k -system $\mathcal{F} \subseteq 2^N$ is a down-closed family such that, for any $S \subseteq N$, the ratio of the sizes of any two maximal sets of $\mathcal{F}$ that are contained in S is at most k . In particular, k -systems generalize the intersection of k matroids.

we can leverage Theorem 1.17 to consider elements in increasing order of their deadlines, which allows for addressing this open question.

Theorem 1.18. *Let $\mathcal{F}_{in}, \mathcal{F}_{out} \subseteq 2^N$ be two down-closed families. If there are a (b, c_{in}) -selectable greedy OCRS π_{in} for a relaxation P_{in} of $\mathcal{F}_{in}$ and a (b, c_{out}) -selectable greedy OCRS π_{out} for a relaxation P_{out} of $\mathcal{F}_{out}$, then there is a $(b(1-b) \cdot c_{in} \cdot c_{out})$ -approximation for the weighted stochastic probing with deadlines problem where the inner and outer constraints are given by $\mathcal{F}_{in}$ and $\mathcal{F}_{out}$, respectively. Moreover, if π_{in} and π_{out} are efficient and there are separation oracles for P_{in} and P_{out} , then the above algorithm has a polynomial time complexity.*

We highlight that the stochastic probing problem with monotone submodular objectives but without deadlines—in short, *submodular stochastic probing*—was considered by Adamczyk et al. [1], who presented for this setting a $(1 - 1/e)/(k_{in} + k_{out} + 1)$ -approximation when $\mathcal{F}_{in}$ and $\mathcal{F}_{out}$ are the intersection of k_{in} and k_{out} matroids, respectively. Using our techniques we obtain $O(1)$ -approximations for considerably more general settings of submodular stochastic probing. More precisely, we can handle a very broad set of constraints, with a probing order chosen by an almighty adversary (instead of being choosable by the algorithm).

Theorem 1.19. *Let $\mathcal{F}_{in}, \mathcal{F}_{out} \subseteq 2^N$ be two down-closed families. If there are a (b, c_{in}) -selectable greedy OCRS π_{in} for a relaxation P_{in} of $\mathcal{F}_{in}$ and a (b, c_{out}) -selectable greedy OCRS π_{out} for a relaxation P_{out} of $\mathcal{F}_{out}$, then there is a $((1 - e^{-b} - o(1)) \cdot c_{in} \cdot c_{out})$ -approximation for the submodular stochastic probing problem where the order in which elements can be probed is chosen by an almighty adversary and the inner and outer constraints are given by $\mathcal{F}_{in}$ and $\mathcal{F}_{out}$, respectively. Moreover, if π_{in} and π_{out} are efficient and there are separation oracles for P_{in} and P_{out} , then the above algorithm has a polynomial time complexity.*

We remark that the same idea used to derive Theorem 1.18 from Theorem 1.17 can also be used to derive from Theorem 1.19 a result for the submodular stochastic probing with deadlines problem.

2 Constructing Online Contention Resolution Schemes

In this section we prove the existence (or non-existence) of OCRSs for various polytopes. Sections 2.1, 2.2 and 2.3 study OCRSs for matroid polytopes, matching polytopes and the natural relaxation of knapsack constraints, respectively. The results proved in these sections prove together Theorem 1.8. Theorem 1.9, which shows that greedy OCRSs for different polytopes can be combined to create greedy OCRSs for the intersection of these polytopes, is proved in Section 2.4.

2.1 OCRS for matroids

In this section we give a greedy OCRS for matroid polytopes. For standard matroidal concepts such as span, rank, contraction and restriction, we refer the reader to Appendix A. Also recall that, for a given matroid $M = (N, \mathcal{F})$, the matroid polytope $P_{\mathcal{F}}$ is defined by $\{x \in \mathbb{R}_{\geq 0}^N \mid \forall S \subseteq N \sum_{e \in S} x_e \leq \text{rank}(S)\}$. The main result of this section can be stated as follows.

Theorem 2.1. *Let $b \in [0, 1]$. There exists a $(b, 1-b)$ -selectable deterministic greedy OCRS for any matroid polytope $P_{\mathcal{F}} \subseteq [0, 1]^N$ on ground set N .*

Combining Theorem 1.9 with Theorem 2.1, we also get a greedy OCRS for the intersection of k matroids.

Corollary 2.2. *Let $b \in [0, 1]$, and let $P_1, \dots, P_k \subseteq [0, 1]^N$ be k matroid polytopes over a common ground set N . Then there exists a $(b, (1-b)^k)$ -selectable deterministic greedy OCRS for $P = \cap_{i=1}^k P_i$.*

Note that the above corollary implies (by Observation 1.6) that there is a $b(1 - b)^k$ -selectable greedy OCRS for P . Moreover, choosing $b = \frac{1}{1+k}$ in Corollary 2.2, we obtain the following.

Corollary 2.3. *Let $P_1, \dots, P_k \subseteq [0, 1]^N$ be matroid polytopes over a common ground set N , let $P = \bigcap_{i=1}^k P_i$, and let $c = \frac{1}{k+1}(1 - \frac{1}{1+k})^k \geq \frac{1}{e(k+1)}$. Then there exists a c -selectable greedy OCRS for P .*

The rest of this section is devoted to the proof of Theorem 2.1. Consider a matroid $M = (N, \mathcal{F})$ and let $P_{\mathcal{F}}$ be the associated polytope. Our greedy OCRS is based on using $x \in P_{\mathcal{F}}$ to find a *chain decomposition* of the elements

$$\emptyset = N_{\ell} \subsetneq N_{\ell-1} \subsetneq \dots \subsetneq N_1 \subsetneq N_0 = N .$$

It then accepts an active element $e \in N_i \setminus N_{i+1}$ if e together with the already accepted elements in $N_i \setminus N_{i+1}$ forms an independent set in the matroid $(M/N_{i+1})|_{N_i}$, i.e., the matroid obtained from M by contracting N_{i+1} and then restricting to N_i . To see that this OCRS is a greedy OCRS, note that the above algorithm is equivalent to defining the family $\mathcal{F}_x = \{I \subseteq N : \forall i I \cap (N_i \setminus N_{i+1}) \text{ is independent in } (M/N_{i+1})|_{N_i}\}$. The family $\mathcal{F}_x$ is clearly a down-closed family of sets (since each $(M/N_{i+1})|_{N_i}$ is a matroid and its independent sets are, thus, down-closed). Moreover, $\mathcal{F}_x$ is a subset of feasible sets because (see, e.g., Theorem 5.1 in [23]) if I_i is an independent set of $(M/N_{i+1})|_{N_i}$ for every i , then the set $\bigcup_i I_i$ is independent in M . Even though we do not need this fact, we highlight that $\mathcal{F}_x$ itself describes a family of independent sets of a matroid. This follows from the fact that $\mathcal{F}_x$ is the family of all (disjoint) unions of independent sets in the matroids $(M/N_{i+1})|_{N_i}$ for $i \in \{0, \dots, \ell - 1\}$, implying that $\mathcal{F}_x$ are the independent sets of the union matroid obtained by taking the union of all matroids $(M/N_{i+1})|_{N_i}$ for $i \in \{0, \dots, \ell - 1\}$.

Having shown that the above algorithm is a greedy OCRS for any chain decomposition of the elements, we turn our attention to the task of defining a chain that maximizes the selectability of our greedy OCRS. For a fixed chain, we have that the selectability of an element $e \in N_i \setminus N_{i+1}$ is

$$\Pr[I \cup \{e\} \in \mathcal{F}_x \mid I \subseteq R(x), I \in \mathcal{F}_x] = \Pr[e \notin \text{span}_i((R(x) \cap (N_i \setminus N_{i+1})) \setminus \{e\})] ,$$

where $\text{span}_i(\cdot)$ denotes the span function of matroid $(M/N_{i+1})|_{N_i}$. Our objective is therefore to construct a chain decomposition that maximizes $\Pr[e \notin \text{span}_i((R(x) \cap (N_i \setminus N_{i+1})) \setminus \{e\})]$ or equivalently minimizes

$$\Pr[e \in \text{span}_i((R(x) \cap (N_i \setminus N_{i+1})) \setminus \{e\})] . \quad (2)$$

We now describe an iterative procedure for constructing a chain decomposition so that (2) is at most b for each element. Initially, the chain only consists of the ground set $N_0 = N$. We need to refine the chain if there exists an element $e \in N$ such that $\Pr[e \in \text{span}(R(x) \setminus \{e\})] > b$. We do that in a “minimal” way as follows:

- Let $S = \emptyset$.
- While there exists $e \in N_0 \setminus S$ such that $\Pr[e \in \text{span}((R(x) \cup S) \setminus \{e\})] > b$, add e to S .

We then let $N_1 = S$. Note that if N_1 is a strict subset of N_0 then we have made progress and we repeat the above procedure (on the matroid induced by N_1) to obtain $N_2 \subsetneq N_1$, and so on, until we obtain $N_{\ell} = \emptyset$. Thus, if we assume that the procedure terminates, then, by construction, the chain satisfies

$$\Pr[\text{span}_i((R(x) \cap (N_i \setminus N_{i+1})) \setminus \{e\})] \leq b \quad \text{for all } i \text{ and } e \in N_i \setminus N_{i+1} .$$

The rest of this section is devoted to show that the chain construction always terminates. As described above, this implies a $(b, 1 - b)$ -selectable greedy OCRS, and thus proves Theorem 2.1.

2.1.1 Proof of termination of chain construction

To prove that the chain decomposition terminates, it is sufficient to show that $N_1 = S \subsetneq N_0 = N$ for any (non-empty) matroid as the chain decomposition then recurses on the matroid induced by N_1 (so the same argument implies that $N_2 \subsetneq N_1$ assuming $N_1 \neq \emptyset$, and so on). Notice that the definition of S implies that S can only increase as coordinates of x are increased. Hence, it is safe to assume that $x \in b \cdot P_{\mathcal{B}}$, where $P_{\mathcal{B}} = \{x \in P_{\mathcal{F}} \mid x(N) = \text{rank}(N)\}$ is the base polytope of the matroid M , which is the set of all maximal vectors in $P_{\mathcal{F}}$.

For proofs in the rest of the section it will be convenient to have the following equivalent view of the construction of $S = N_1$ in the refinement procedure.

- Let $S_0 = \emptyset$, and let $S_1 = \{e \in N \mid \Pr[e \in \text{span}(R(x) \setminus \{e\})] > b\}$ be the set of elements that are “likely” to be spanned.
- Assuming we have defined $S_0, S_1, \dots, S_{i-1}$, let

$$S_i = \{e \in N \mid \Pr[e \in \text{span}((R(x) \cup S_{i-1}) \setminus \{e\})] > b\} ,$$

that is, S_i contains those elements that are likely to be spanned assuming that the elements of S_{i-1} are contracted (or equivalently appear with probability 1).

Notice that $S_{i-1} \subseteq S_i$ for every $i \geq 1$, and $S_i = S_{i-1}$ implies $S_i = S_j$ for every $j > i$. Thus, we must have $N_1 = S = S_{|N|}$. The key technical part of the termination analysis is the following lemma.

Lemma 2.4. *It always holds that:*

$$\sum_{e \in N} x_e \Pr[e \in \text{span}(R(x) \cup S)] \leq b \cdot (x(N) + (1 - b) \text{rank}(S)) .$$

Moreover, the inequality is strict if $S \neq \emptyset$.

Before proving the lemma, let us see that it implies what we want:

Corollary 2.5. *If $N \neq \emptyset$ then $N_1 = S \subsetneq N$.*

Proof. The corollary is clearly true if $S = \emptyset$. Otherwise, by Lemma 2.4,

$$\begin{aligned} x(S) &\leq x(\text{span}(R(x) \cup S)) = \sum_{e \in N} x_e \Pr[e \in \text{span}(R(x) \cup S)] \\ &< b(x(N) + (1 - b) \text{rank}(S)) \leq b \text{rank}(N) . \end{aligned}$$

As $x(S) < b \cdot \text{rank}(N)$ and $x(N) = b \cdot \text{rank}(N)$ by our assumption that $x \in bP_{\mathcal{B}}$, we get $x(S) < x(N)$, which implies $N \setminus S \neq \emptyset$. $\square$

Let us now continue with the proof of the our technical lemma.

Proof of Lemma 2.4. Let $S' = \{e_1, e_2, \dots, e_k\}$ be a basis of the matroid $M|_S$ obtained by first greedily selecting elements from S_0 , then greedily adding elements from S_1 , and so on. Consider the distribution μ over 2^N defined according to the following sampling procedure:

1. Let A be a random set originally distributed like $R(x)$.
2. For $j = 1, 2, \dots, k$, if $e_j \notin \text{span}(A)$, add e_j to A .
3. Output A .

Observe that this sampling procedure guarantees that the distributions of $\text{span}(S \cup R(x))$, $\text{span}(S' \cup R(x))$ and $\text{span}(A)$ are identical. Therefore,

$$\Pr[e \in \text{span}(R(x) \cup S)] = \Pr_{A \sim \mu}[e \in \text{span}(A)] .$$

Now simple calculations yield

$$\begin{aligned} \sum_{e \in N} x_e \Pr_{A \sim \mu}[e \in \text{span}(A)] &= \mathbb{E}_{A \sim \mu}[x(\text{span}(A))] \\ &\leq b \cdot \mathbb{E}_{A \sim \mu}[\text{rank}(\text{span}(A))] && \text{(using } x \in b \cdot P_T) \\ &= b \cdot \mathbb{E}_{A \sim \mu}[\text{rank}(A)] \leq b \cdot \mathbb{E}_{A \sim \mu}[|A|] . \end{aligned}$$

We complete the proof by showing that $\mathbb{E}_{A \sim \mu}[|A|] \leq x(N) + (1 - b) \text{rank}(S)$. To see this, note that

$$\begin{aligned} \mathbb{E}_{A \sim \mu}[|A|] &= x(N) + \sum_{j=1}^k \Pr[e_j \notin \text{span}(R(x) \cup \{e_1, \dots, e_{j-1}\})] \\ &\leq x(N) + (1 - b)k = x(N) + (1 - b) \cdot \text{rank}(S) , \end{aligned}$$

where the inequality follows from the following argument. If we let i be the smallest index so that $e_j \in S_i$, then by the construction of S' we have $\Pr[e_j \in \text{span}(R(x) \cup \{e_1, \dots, e_{j-1}\})] \geq \Pr[e_j \in \text{span}(R(x) \cup S_{i-1})] > b$. Finally, we remark that the strict inequality $\Pr[e_j \in \text{span}(R(x) \cup S_{i-1})] > b$ implies that the inequality in the statement is strict if $S \neq \emptyset$. $\square$

Efficient implementation

The only step that is not constructive in the description of our OCRS for matroids is the computation of probabilities of the type $\Pr[e \in \text{span}((R(x) \cup S_{i-1}) \setminus \{e\})]$. One can easily get around this issue by using good estimates through Monte-Carlo sampling, leading to the following.

Lemma 2.6. *For any $\epsilon > 0$ and $\alpha > 0$, there is a randomized construction of a chain $\emptyset = N_\ell \subsetneq N_{\ell-1} \subsetneq \dots \subsetneq N_1 \subsetneq N_0 = N$ such that with probability at least $1 - |N|^{-\alpha}$ the greedy OCRS defined by the chain is $(b, 1 - b - \epsilon)$ -selectable. Furthermore, the time needed for this construction is $O(\alpha \cdot \frac{1}{\epsilon^2} \cdot \text{poly}(|N|) \cdot T)$, where T is the time for a single call to the independence oracle of the matroid, and we assume that we can sample a Bernoulli random variable in $O(1)$ time.*

Proof. For simplicity we define $p_{e,S} = \Pr[e \in \text{span}((R(x) \cup S) \setminus \{e\})]$ for $e \in N$ and $S \subseteq N$. As before, we focus on the construction of N_1 ; the algorithm is then applied recursively. We recall that in the above-mentioned construction we set $N_1 = S$, where S was constructed from $S = \emptyset$ by adding elements $e \in N$ satisfying $p_{e,S} > b$. To perform this step constructively we will, whenever we need a probability $p_{e,S}$, use an estimate $\hat{p}_{e,S}$ obtained through Monte-Carlo sampling. By standard results on Monte-Carlo sampling (see, e.g., [11]) it suffices to use $O(\alpha \cdot \frac{1}{\epsilon^2} \cdot \log |N|)$ many samples to obtain a value $\hat{p}_{e,S}$ such that

$$\Pr[\hat{p}_{e,S} \in [p_{e,S} - \epsilon, p_{e,S}]] \geq 1 - |N|^{-3-\alpha} .$$

Hence, to construct S , we start with $S = \emptyset$ and successively add elements $e \in N \setminus S$ with $\hat{p}_{e,S} > b$. There are at most $|N|$ elements we add to S , and to add one element to S we may have to check the values $\hat{p}_{e,S}$ of all elements in $N \setminus S$. Hence, to construct S we use at most $O(|N|^2)$ estimates of the type $\hat{p}_{e,S}$. At the end of this procedure we set $N_1 = S$ and repeat. Since there are at most $O(|N|)$ sets in the final chain $\emptyset = N_\ell \subsetneq \dots \subsetneq N_0 = N$, the total number of estimates we need is bounded by $O(|N|^3)$, which implies the claimed running time of our algorithm. Moreover, with probability at least $1 - |N|^{-\alpha}$, all our estimates

$\hat{p}_{e,S}$ satisfy $\hat{p}_{e,S} \in [p_{e,S} - \epsilon, p_{e,S}]$. To prove the lemma, we assume from now on that all our estimates fulfil this property and show that this implies that the OCRS we obtain is $(b, 1 - b - \epsilon)$ -selectable.

Notice that during our construction of S we only add elements e to S satisfying $b < \hat{p}_{e,S} \leq p_{e,S}$. Hence, the elements we add would also have been added to S in the construction that uses the true probabilities $p_{e,S}$. Therefore, for the same reasons showing that $S \subsetneq N$ when using the true probabilities $p_{e,S}$, we also have $S \subsetneq N$. It remains to observe that at the end of the construction of S , i.e., when we set $N_1 = S$, the probability $1 - p_{e,S} = \Pr[e \notin \text{span}((S \cup R(x)) \setminus \{e\})]$ of an element $e \in N \setminus S$ being selectable is at least $1 - b - \epsilon$. This indeed holds since

$$1 - p_{e,S} \geq 1 - \hat{p}_{e,S} - \epsilon \geq 1 - b - \epsilon . \quad \square$$

2.2 OCRSs for matchings in general graphs

In this section we describe an OCRS that works for a relaxation P_G of matching in a general graph $G = (V, E)$. Specifically, the relaxation polytope P_G is defined as:

$$\begin{aligned} \sum_{g \in \delta(u)} x_g &\leq 1 \quad \forall u \in V \\ x_g &\geq 0 \quad \forall g \in E , \end{aligned}$$

where $\delta(u)$ is the set of edges incident to the node u . Observe that this relaxation is weaker than the matching polytope, hence, our results hold also for the matching polytope. We would like to stress that the ground set in this section is the set of edges, and thus, unlike in the rest of the paper, we denote it by E . Some ideas from the proof of Theorem 2.7 can be traced to offline CRSs given by [14].¹⁰

Theorem 2.7. *For every $b \in [0, 1]$, there exists a (b, e^{-2b}) -selectable randomized greedy OCRS for the relaxation $P_G \subseteq [0, 1]^E$ of matching in a graph $G = (V, E)$.*

Proof. Let $x \in bP_G$ be the input point to the OCRS, and let $A \sim R(x)$ be the set of active elements. Our OCRS begins by selecting a subset K of potential edges, where every edge $g \in E$ belongs to K with probability $(1 - e^{-x_g})/x_g$, independently (observe that this probability is indeed always within the range $[0, 1]$). Whenever an edge g reveals whether it is active, the OCRS selects it if $g \in A \cap K$ and the addition of g to the set of already selected edges does not make this set an illegal matching. Observe that for any fixed choice of K this OCRS is a deterministic greedy OCRS, and thus, for a random K it is a randomized greedy OCRS.

Next, let us show that our OCRS is (b, e^{-2b}) -selectable. Consider an arbitrary edge $g' = uv \in E$. We need to prove that with probability at least e^{-2b} the edge g' is in K , and can be added to any matching which is a subset of $A \cap K$. Formally, we need to prove:

$$\Pr[g' \in K, A \cap K \cap (\delta(u) \cup \delta(v) \setminus \{g'\}) = \emptyset] \geq e^{-2b} .$$

Clearly, every edge $g \in E$ belongs to $A \cap K$ with probability $x_g \cdot (1 - e^{-x_g})/x_g = 1 - e^{-x_g}$. Since the membership of every edge in K and A is independent from the membership of other edges in these sets, we get:

$$\begin{aligned} \Pr[g' \in K, A \cap K \cap (\delta(u) \cup \delta(v) \setminus \{g'\}) = \emptyset] &= \Pr[g' \in K] \cdot \prod_{g \in \delta(u) \cup \delta(v) \setminus \{g'\}} \Pr[g \notin A \cap K] \\ &= \frac{(1 - e^{-x_{g'}})}{x_{g'}} \cdot \prod_{g \in \delta(u) \cup \delta(v) \setminus \{g'\}} e^{-x_g} = \frac{(1 - e^{-x_{g'}})}{x_{g'}} \cdot e^{-\sum_{g \in \delta(u) \cup \delta(v) \setminus \{g'\}} x_g} \\ &\geq \frac{(1 - e^{-x_{g'}})}{x_{g'}} \cdot e^{-2(b - x_{g'})} = \frac{e^{x_{g'}}(e^{x_{g'}} - 1)}{x_{g'}} \cdot e^{-2b} \geq e^{-2b} , \end{aligned}$$

¹⁰The details of these CRSs are omitted in [14], but can be found in [13].

where the first inequality holds since the membership of x in bP_G guarantees that the total x -values of the edges in $\delta(u) \setminus \{g'\}$ (or $\delta(v) \setminus \{g'\}$) is at most $b - x_{g'}$. $\square$

Remark: If one is interested in a *deterministic* greedy OCRS, it is possible to set K equal to E deterministically in the greedy OCRS described by the last proof. A simple modification of the proof shows that the resulting deterministic greedy OCRS is $(b, (1 - b)^2)$ -selectable.

2.3 OCRSs for knapsack constraint

In this section we consider the problem of defining an OCRS for a polytope $P \subseteq [0, 1]^N$ defined by a single knapsack constraint. That is, each element $e \in N$ has an associated size $s_e \in [0, 1]$ and P is defined by

$$\sum_{e \in N} s_e x_e \leq 1 \quad x_e \in [0, 1] \quad \text{for } e \in N .$$

We begin with an interesting simple observation.

Proposition 2.8. *For every $n \geq 1$, there exists a knapsack constraint over a ground set of n elements such that no deterministic greedy OCRS for the polytope defined by this constraint is (b, c) -selectable for any pair of $b \in [0, 1]$ and $c > (1 - b)^{n-1}$.*

Proof. Consider the knapsack constraint defined by the ground set $N = \{1, 2, \dots, n\}$ and sizes $s_1 = s_2 = \dots = s_{n-1} = 1/n$, $s_n = 1$. Assume towards a contradiction that there exists a deterministic greedy OCRS that is (b, c) -selectable for some $b \in [0, 1]$ and $c > (1 - b)^{n-1}$, and consider the family $\mathcal{F}_x$ of feasible sets used by this OCRS for the possible input $x_1 = x_2 = \dots = x_{n-1} = b$, $x_n = b/n$.

Since the OCRS is (b, c) -selectable for $c > 0$, each element $e \in N$ must be included in at least one set of the family $\mathcal{F}_x$. Thus, by the down-monotonicity of $\mathcal{F}_x$ it must contain the set $\{e\}$ for every $e \in N$. On the other hand, for every $e \in N \setminus \{n\}$ we have $\{e, n\} \notin \mathcal{F}$, and thus, also $\{e, n\} \notin \mathcal{F}_x$. Combining all these observations, we get:

$$\Pr[I \cup \{n\} \in \mathcal{F}_x \quad \forall I \subseteq R(x), I \in \mathcal{F}_x] \leq \Pr[e \notin R(x) \quad \forall e \in N \setminus \{n\}] = (1 - b)^{n-1} < c ,$$

which contradicts the (b, c) -selectability of the assumed OCRS. $\square$

The next theorem shows that randomized greedy OCRSs can do much better. Some of the ideas used by this theorem can be traced back to an offline CRS presented by [10] for knapsack constraints.

Theorem 2.9. *For every $b \in [0, 1/2]$, there exists a $(b, \frac{1-2b}{2-2b})$ -selectable randomized greedy OCRS for any polytope P defined by a knapsack constraint.*

Proof. Let $x \in bP$ be the input point to the OCRS, and let $N_{\text{big}} = \{e \in N \mid s_e > 1/2\}$ be the subset of elements that are big. We use b_{big} to denote the total part of the knapsack occupied by big elements in the fractional solution x . Formally,

$$b_{\text{big}} = \sum_{e \in N_{\text{big}}} s_e x_e .$$

Observe that b_{big} is always within the range $[0, b]$. The randomized greedy OCRS we use is defined as follows. With probability p_{big} accept greedily the elements of N_{big} while respecting the knapsack inequality, where p_{big} is the probability

$$p_{\text{big}} = \frac{1 - 2b + 2b_{\text{big}}}{2 - 2b} .$$

With the remaining probability accept greedily the small elements of $N \setminus N_{\text{big}}$ while respecting the knapsack inequality. It is easy to see that this OCRS is indeed a randomized greedy OCRS.

We continue to analyze the selectability of the above OCRS. Observe that for any big element $e' \in N_{\text{big}}$:

$$\Pr[I \cup \{e'\} \in \mathcal{F}_x \quad \forall I \subseteq R(x), I \in \mathcal{F}_x] \geq \Pr\left[\{e'\} \in \mathcal{F}_x, \sum_{e \in R(x) \cap N_{\text{big}}} s_e \leq 1 - s_{e'}\right] . \quad (3)$$

The event $\{e'\} \in \mathcal{F}_x$ is simply the event that the OCRS decides to accept big elements. Given that this event occurs, the condition $\sum_{e \in R(x) \cap N_{\text{big}}} s_e \leq 1 - s_{e'}$ guarantees that for every subset I of $R(x) \cap N_{\text{big}}$ one can add the element e' without violating the knapsack constraint. We can lower bound (3) as follows.

$$\begin{aligned} \Pr\left[\{e'\} \in \mathcal{F}_x, \sum_{e \in R(x) \cap N_{\text{big}}} s_e \leq 1 - s_{e'}\right] &= p_{\text{big}} \cdot \Pr[R(x) \cap (N_{\text{big}} \setminus \{e'\}) = \emptyset] \\ &\geq \frac{1 - 2b + 2b_{\text{big}}}{2 - 2b} \cdot (1 - x(N_{\text{big}})) \geq \frac{1 - 2b + 2b_{\text{big}}}{2 - 2b} \cdot (1 - 2b_{\text{big}}) \geq \frac{1 - 2b}{2 - 2b} , \end{aligned}$$

where the first inequality follows from the union bound and the second inequality uses the fact that $s_e > 1/2$ for every $e \in N_{\text{big}}$. Similarly, for an element $e \notin N_{\text{big}}$ we have:

$$\begin{aligned} \Pr[I \cup \{e'\} \in \mathcal{F}_x \quad \forall I \subseteq R(x), I \in \mathcal{F}_x] &\geq \Pr\left[\{e'\} \in \mathcal{F}_x, \sum_{e \in R(x) \setminus N_{\text{big}}} s_e \leq 1 - s_{e'}\right] \\ &\geq (1 - p_{\text{big}}) \cdot \Pr[s(R(x) \setminus (N_{\text{big}} \cup \{e'\})) \leq 1/2] \geq \left(1 - \frac{1 - 2b + 2b_{\text{big}}}{2 - 2b}\right) \cdot \left(1 - \frac{b - b_{\text{big}}}{1/2}\right) \\ &= \frac{1 - 2b_{\text{big}}}{2 - 2b} \cdot (1 - 2b + 2b_{\text{big}}) \geq \frac{1 - 2b}{2 - 2b} , \end{aligned}$$

where the third inequality follows from Markov's inequality. $\square$

2.4 Combining constraints

Like offline CRSs (see [10]), greedy OCRSs can be combined to form greedy OCRSs for more involved constraints.

Definition 2.10. *Given two greedy OCRSs π^1 and π^2 for polytopes P_1 and P_2 , the combination of π^1 and π^2 is a greedy OCRS π for the polytope $P = P_1 \cap P_2$. For every input $x \in P$, π is defined by the down-closed family $\mathcal{F}_{\pi,x} = \mathcal{F}_{\pi^1,x} \cap \mathcal{F}_{\pi^2,x}$.*

Theorem 1.9 is an immediate implication of the next lemma.

Lemma 2.11. *Let π^1 and π^2 be (b, c_1) -selectable and (b, c_2) -selectable greedy OCRSs, respectively. The combination of π^1 and π^2 is $(b, c_1 \cdot c_2)$ -selectable.*

Proof. Let P_1 and P_2 be the polytopes of π^1 and π^2 , respectively. Additionally, let $x \in b(P_1 \cap P_2)$, and let e be an arbitrary element of N . We need to prove that:

$$\Pr[I \cup \{e\} \in \mathcal{F}_{\pi^1,x} \cap \mathcal{F}_{\pi^2,x} \quad \forall I \subseteq R(x), I \in \mathcal{F}_{\pi^1,x} \cap \mathcal{F}_{\pi^2,x}] \geq c_1 \cdot c_2 .$$

For ease of notation, let us denote by $\chi_e(A, \mathcal{F}, \mathcal{F}')$ an indicator for the event that $I \cup \{e\} \in \mathcal{F}$ for every set $I \subseteq A$ obeying $I \in \mathcal{F}'$. Using this notation, the inequality that we need to prove becomes: $\Pr[\chi_e(R(x), \mathcal{F}_{\pi^1,x} \cap \mathcal{F}_{\pi^2,x}, \mathcal{F}_{\pi^1,x} \cap \mathcal{F}_{\pi^2,x})] \geq c_1 \cdot c_2$.

On the other hand, observe that $x \in b(P_1 \cap P_2) \subseteq bP_1$. Hence, the (b, c_1) -selectability of π^1 implies that:

$$\begin{aligned} \Pr[\chi_e(R(x), \mathcal{F}_{\pi^1,x}, \mathcal{F}_{\pi^1,x} \cap \mathcal{F}_{\pi^2,x})] &\geq \Pr[\chi_e(R(x), \mathcal{F}_{\pi^1,x}, \mathcal{F}_{\pi^1,x})] \\ &= \Pr[I \cup \{e\} \in \mathcal{F}_{\pi^1,x} \quad \forall I \subseteq R(x), I \in \mathcal{F}_{\pi^1,x}] \geq c_1 . \end{aligned}$$

where the first inequality follows since χ_e is a non-increasing function of its third argument (when the other two arguments are fixed). Similarly, we also get: $\Pr[\chi_e(R(x), \mathcal{F}_{\pi^2, x}, \mathcal{F}_{\pi^1, x} \cap \mathcal{F}_{\pi^2, x})] \geq c_2$. Next, observe that χ_e is also non-increasing in its first argument (when the other two arguments are fixed). Hence, if we let R be a random set distributed like $R(x)$, then, by the FKG inequality:

$$\begin{aligned} & \Pr[\chi_e(R(x), \mathcal{F}_{\pi^1, x} \cap \mathcal{F}_{\pi^2, x}, \mathcal{F}_{\pi^1, x} \cap \mathcal{F}_{\pi^2, x})] \\ &= \Pr[\chi_e(R, \mathcal{F}_{\pi^1, x}, \mathcal{F}_{\pi^1, x} \cap \mathcal{F}_{\pi^2, x}) \cdot \chi_e(R, \mathcal{F}_{\pi^2, x}, \mathcal{F}_{\pi^1, x} \cap \mathcal{F}_{\pi^2, x})] \\ &\geq \Pr[\chi_e(R(x), \mathcal{F}_{\pi^1, x}, \mathcal{F}_{\pi^1, x} \cap \mathcal{F}_{\pi^2, x})] \cdot \Pr[\chi_e(R(x), \mathcal{F}_{\pi^2, x}, \mathcal{F}_{\pi^1, x} \cap \mathcal{F}_{\pi^2, x})] \geq c_1 \cdot c_2. \quad \square \end{aligned}$$

3 From Selectability to Approximation

In this section we prove Theorems 1.10 and 1.11 which study the relation between the value of a vector x and the expected value of the output produced by an OCRS given this vector as input. Following is a restatement of Theorem 1.10. Unlike the original statement of the theorem, this restatement uses the notation $S(p)$, where S is a set and p is a probability, to denote a random set containing every element $e \in S$ with probability p , independently. Observe that $S(p)$ has the same distribution as $R(p \cdot \mathbf{1}_S)$.

Theorem 1.10. *Given a non-negative monotone submodular function $f: 2^N \rightarrow \mathbb{R}_{\geq 0}$ and a (b, c) -selectable greedy OCRS for a polytope P , applying the greedy OCRS to an input $x \in bP$ results in a random set S satisfying $\mathbb{E}[f(S)] \geq c \cdot F(x)$, where F is the multilinear extension of f . Moreover, even if f is not monotone, $\mathbb{E}[f(S(1/2))] \geq (c/4) \cdot F(x)$, where the random decisions used to calculate $S(1/2)$ are considered part of the algorithm, and thus, known to the almighty adversary.*

To prove Theorem 1.10 we need to define some concepts regarding *offline* CRSs. Recall that a CRS for a polytope P is a (possibly random) function $\pi: 2^N \rightarrow 2^N$ that depends on an input vector $x \in P$, and for every set $S \subseteq N$ returns a subset $\pi(S) \subseteq S$ that obeys the polytope P (i.e., $\mathbf{1}_{\pi(S)} \in P$).

Definition 3.1. A CRS π for a polytope P is:

- (b, c) -balanced *some* for $b, c \in [0, 1]$ if $\Pr[e \in \pi(R(x)) \mid e \in R(x)] \geq c \cdot x_e$ whenever the input vector x belongs to bP and $x_e > 0$.
- monotone if $\Pr[e \in \pi(S_1)] \geq \Pr[e \in \pi(S_2)]$ whenever $e \in S_1 \subseteq S_2$.

We now define and analyze an interesting notion that allows us to connect the world of OCRSs with that of CRSs.

Definition 3.2 (characteristic CRS). *The characteristic CRS $\bar{\pi}$ of a greedy OCRS π for a polytope P is a CRS for the same polytope P . It is defined for an input $x \in P$ and a set $A \subseteq N$ by $\bar{\pi}(A) = \{e \in A \mid I \cup \{e\} \in \mathcal{F}_{\pi, x} \ \forall I \subseteq A, I \in \mathcal{F}_{\pi, x}\}$.*

The following observation shows that $\bar{\pi}(A)$ obeys the polytope P and $\bar{\pi}(A) \subseteq A$. Thus, the characteristic CRS $\bar{\pi}$ is a true CRS for the polytope P .

Observation 3.3. *For every set $A \subseteq N$ and a characteristic CRS $\bar{\pi}$ of a greedy OCRS π , the set $\bar{\pi}(A)$ is always a subset of the elements selected by π when the active elements are the elements of A .*

Proof. Fix an element $e \in \bar{\pi}(A)$, and let T be the set of elements selected by the greedy OCRS π immediately before e reveals whether it is active. The definition of $\bar{\pi}(A)$ guarantees that $T \cup \{e\} \in \mathcal{F}_{\pi, x}$ since $T \in \mathcal{F}_{\pi, x}$, and thus, e is accepted by π . $\square$

Lemma 3.4. *The characteristic CRS $\bar{\pi}$ of a (b, c) -selectable greedy OCRS π is (b, c) -balanced and monotone.*

Proof. Let P denote the polytope of the greedy OCRS π (and its characteristic CRS $\bar{\pi}$), and fix a vector $x \in bP$. Since π is (b, c) -selectable, we get for a set A distributed like $R(x)$ and an arbitrary element $e \in N$:

$$\begin{aligned} \Pr[e \in \bar{\pi}(A) \mid e \in A] &= \Pr[I \cup \{e\} \in \mathcal{F}_{\pi, x} \mid \forall I \subseteq A, I \in \mathcal{F}_{\pi, x} \mid e \in A] \\ &= \Pr[I \cup \{e\} \in \mathcal{F}_{\pi, x} \mid \forall I \subseteq A, I \in \mathcal{F}_{\pi, x}] \geq c . \end{aligned}$$

The last inequality implies, by definition, that $\bar{\pi}$ is a (b, c) -balanced CRS.

Next, let us prove that $\bar{\pi}$ is monotone. Fix an instantiation of $\mathcal{F}_{\pi, x}$, an element $e \in N$ and two sets $e \in A_1 \subseteq A_2 \subseteq N$. If $e \in \bar{\pi}(A_2)$, then we know that $I \cup \{e\} \in \mathcal{F}_{\pi, x}$ for every $I \subseteq A_2, I \in \mathcal{F}_{\pi, x}$. Thus, clearly, $I \cup \{e\} \in \mathcal{F}_{\pi, x}$ for every $I \subseteq A_1 \subseteq A_2, I \in \mathcal{F}_{\pi, x}$, which implies $e \in \bar{\pi}(A_1)$. In summary, we got that:

$$e \in \bar{\pi}(A_2) \Rightarrow e \in \bar{\pi}(A_1) ,$$

and thus, even when we do not fix the instantiation of $\mathcal{F}_{\pi, x}$:

$$\Pr[e \in \bar{\pi}(A_2)] \leq \Pr[e \in \bar{\pi}(A_1)] ,$$

which completes the proof that $\bar{\pi}$ is monotone. □

To continue the proof of Theorem 1.10 we need to state some known results. The following lemma follows from the work of [10] (mostly from Theorem 1.3 in that work).

Lemma 3.5. *For every given non-negative submodular function $f: 2^N \rightarrow \mathbb{R}_{\geq 0}$, there exists a function $\eta_f: 2^N \rightarrow 2^N$ that always returns a subset of its argument (i.e., $\eta_f(S) \subseteq S$ for every $S \subseteq N$) having the following property. For every monotone (b, c) -balanced CRS π for a polytope P and input vector $x \in bP$:*

$$\mathbb{E}[f(\eta_f(\pi(R(x))))] \geq c \cdot F(x) ,$$

where $F(x)$ is the multilinear extension of f .

The following two known lemmata about submodular functions have been rephrased to make them fit our notation better.

Lemma 3.6 (Lemma 2.2 of [12]). *Let $g: 2^N \rightarrow \mathbb{R}$ be a submodular function, and let T be an arbitrary set $T \subseteq N$. For every random set $T_p \subseteq T$ which contains every element of T with probability p (not necessarily independently):*

$$\mathbb{E}[g(T_p)] \geq (1 - p)g(\emptyset) + p \cdot g(T) .$$

Lemma 3.7 (Lemma 2.2 of [5]). *Let $g: 2^N \rightarrow \mathbb{R}_{\geq 0}$ be a non-negative submodular function. For every random set $N_p \subseteq N$ which contains every element of N with probability at most p (not necessarily independently):*

$$\mathbb{E}[g(N_p)] \geq (1 - p)g(\emptyset) .$$

We are now ready to prove Theorem 1.10.

Proof of Theorem 1.10. Let $\bar{\pi}$ be the characteristic CRS of the OCRS π we consider, and let $A \sim R(x)$ be the set of active elements. For notational convenience, let us also denote by S' the set $\eta_f(\bar{\pi}(A))$. Notice that S' is a subset of the set S of accepted elements since $\bar{\pi}(A)$ is a subset of S by Observation 3.3. Lemmata 3.4 and 3.5 imply together the inequality $\mathbb{E}[f(S')] = \mathbb{E}[f(\eta_f(\bar{\pi}(A)))] \geq c \cdot F(x)$. To complete the proof of the

first part of the theorem, it is now enough to observe that by the monotonicity of f : $\mathbb{E}[f(S)] \geq \mathbb{E}[f(S')] \geq c \cdot F(x)$.

To prove the second part of the theorem, let us fix the set A and the family $\mathcal{F}_{\pi,x}$ which, for brevity, we denote by $\mathcal{F}_x$ in the rest of the proof. Observe that the set S' is deterministic once A and $\mathcal{F}_x$ are fixed, and let us denote this set by $S'_{A,\mathcal{F}_x}$. Hence, we can think of $S(1/2)$ as obtained by first calculating a set $S'_{A,\mathcal{F}_x}(1/2)$ containing every element of $S'_{A,\mathcal{F}_x}$ with probability $1/2$, independently, and then adding to it a random set $\Delta \subseteq N \setminus S'_{A,\mathcal{F}_x}$. By controlling the order in which elements reveal whether they are active, the adversary can make the distribution of Δ depend on $S'_{A,\mathcal{F}_x}(1/2)$; however, Δ is guaranteed to contain every element with probability at most $1/2$ for every given choice of $S'_{A,\mathcal{F}_x}(1/2)$. Using this observation we get:

$$\begin{aligned} \mathbb{E}[f(S(1/2)) \mid A, \mathcal{F}_x] &= \mathbb{E}[f(S'(1/2) \cup \Delta) \mid A, \mathcal{F}_x] \\ &= \sum_{B \subseteq S'_{A,\mathcal{F}_x}} \Pr[S'_{A,\mathcal{F}_x}(1/2) = B \mid A, \mathcal{F}_x] \cdot \mathbb{E}[f(B \cup \Delta) \mid S'_{A,\mathcal{F}_x}(1/2) = B, A, \mathcal{F}_x] \\ &\geq \sum_{B \subseteq S'_{A,\mathcal{F}_x}} \Pr[S'_{A,\mathcal{F}_x}(1/2) = B \mid A, \mathcal{F}_x] \cdot \frac{\mathbb{E}[f(B) \mid A, \mathcal{F}_x]}{2} = \frac{\mathbb{E}[f(S'(1/2)) \mid A, \mathcal{F}_x]}{2} \\ &\geq \frac{\mathbb{E}[f(S') \mid A, \mathcal{F}_x]}{4} , \end{aligned}$$

where the first inequality follows from Lemma 3.7 since the function $h_B(T) = f(B \cup T)$ is non-negative and submodular for every set $B \subseteq N$, and the second inequality follows from Lemma 3.6. Taking an expectation over the possible values of the set A and the family $\mathcal{F}_x$, we get:

$$\begin{aligned} \mathbb{E}[f(S(1/2))] &= \mathbb{E}_{A,\mathcal{F}_x}[\mathbb{E}[f(S(1/2)) \mid A, \mathcal{F}_x]] \\ &\geq \mathbb{E}_{A,\mathcal{F}_x} \left[\frac{\mathbb{E}[f(S') \mid A, \mathcal{F}_x]}{4} \right] = \frac{\mathbb{E}[f(S')]}{4} \geq (c/4) \cdot F(x) . \quad \square \end{aligned}$$

Notice that the result proved by Theorem 1.10 for non-monotone functions loses a factor of 4 in the guarantee. To avoid that, we also consider online adversaries. Unfortunately, we do not have an improved result for greedy OCRSs against online adversaries. Instead, we study the performance of a different family of OCRSs against such adversaries.

Definition 3.8 (element-monotone OCRS). *An element-monotone OCRS is an OCRS characterized by a down-closed family $\mathcal{F}_e \subseteq 2^{N \setminus \{e\}}$ for every element $e \in N$, where the families $\{\mathcal{F}_e\}_{e \in N}$ can be either deterministic (a deterministic element-monotone OCRS) or taken from some joint distribution (a randomized element-monotone OCRS). Such an OCRS accepts an active element e if and only if $A^{<e} \in \mathcal{F}_e$, where $A^{<e}$ is the set of active elements that revealed that they are active before e does so.*

One can observe that all the results we prove for greedy OCRSs in this paper can be easily applied also to element-monotone OCRSs with the following modified definition of (b, c) -selectability.

Definition 3.9 $((b, c)$ -selectability for element-monotone OCRSs). *Let $b, c \in [0, 1]$. An element-monotone OCRS for P is (b, c) -selectable if for any $x \in b \cdot P$ we have*

$$\Pr[R(x) \setminus \{e\} \in \mathcal{F}_e] \geq c \quad \forall e \in N .$$

Moreover, the greedy OCRSs we describe for specific polytopes can be converted into similar element-monotone OCRSs with the same (b, c) -selectability guarantee (in fact the greedy OCRS we describe for matroids is already an element-monotone OCRS without any modifications). Our objective in the rest of this section is to prove the next theorem. Notice that Theorem 1.11 is a special case of this theorem.

Theorem 3.10. *Given an element-monotone (b, c) -selectable OCRS π and a non-negative submodular function $f: 2^N \rightarrow \mathbb{R}_{\geq 0}$, there exists an OCRS π' that for every input vector $x \in bP$ and online adversary selects a random set S such that $\mathbb{E}[f(S)] \geq c \cdot F(x)$.*

To prove Theorem 3.10, we need some notation. Let σ_a denote an arbitrary fixed absolute order over the elements of N . Given element $e \in N$ and a vector $y \in [0, 1]^N$, let $F(e \mid y)$ denote the marginal contribution obtained by increasing the coordinate of e in y to 1. Formally, $F(e \mid y) = F(y \vee \mathbf{1}_{\{e\}}) - F(y)$. Finally, we say that x is *non-reducible* if $F(e \mid x \wedge \mathbf{1}_{\sigma_a^{<e}}) \geq 0$ for every element $e \in N$ having $x_e > 0$, where $\sigma_a^{<e}$ is the set of elements that appear before e in the order σ_a .

Observation 3.11. *If x is reducible, then there exists a non-reducible vector $x' \leq x$ obeying $F(x') \geq F(x)$.*

Proof. Consider the vector x' obtained from x by the following process. Start with $x' \leftarrow x$. Scan the elements in the order σ_a . For every element e , if $F(e \mid x' \wedge \mathbf{1}_{\sigma_a^{<e}}) < 0$, then reduce x'_e to 0. Clearly this process ends up with a non-reducible vector x' . Moreover, every step of the process only increases the value of $F(x')$, and thus, this value ends up at least as large as its initial value $F(x)$. To see why this is the case, consider a step in which the value of x'_e is reduced to 0 for some element $e \in N$, and let x^1 and x^2 be the vector x' before and after the reduction. Then,

$$\begin{aligned} F(x^1) &= f(\emptyset) + \sum_{e' \in N} x_{e'}^1 \cdot F(e' \mid x^1 \wedge \mathbf{1}_{\sigma_a^{<e'}}) \leq f(\emptyset) + \sum_{e' \in N \setminus \{e\}} x_{e'}^1 \cdot F(e' \mid x^1 \wedge \mathbf{1}_{\sigma_a^{<e'}}) \\ &\leq f(\emptyset) + \sum_{e' \in N \setminus \{e\}} x_{e'}^1 \cdot F(e' \mid x^2 \wedge \mathbf{1}_{\sigma_a^{<e'}}) = F(x^2), \end{aligned}$$

where the first inequality holds since $F(e \mid x^1 \wedge \mathbf{1}_{\sigma_a^{<e}}) < 0$, and the second inequality holds by submodularity. $\square$

Observation 3.11 shows that it is enough to prove Theorem 3.10 for non-reducible x . If x is reducible, the OCRS π' can calculate a non-reducible x' having $F(x') \geq F(x)$, and then “pretend” that some active elements are in fact inactive in a way that makes every element e active with probability x'_e , independently. Determining x' requires exponential time, but even if one wants a polynomial time OCRS π' it is possible to use sampling to get, for every constant $d > 0$, a non-reducible vector x' having $F(x') \geq F(x) - |N|^{-d} \cdot \max\{f(\{e\}) \mid x_e > 0\}$.

The OCRS π' we use in the proof of Theorem 3.10 works as follows: whenever π' learns that an element e is active it complements the set $A^{<e}$ to have the distribution $R(x \wedge \mathbf{1}_{N \setminus \{e\}})$, and then checks whether the resulting random set is in $\mathcal{F}_e$. More formally, let σ be the order in which the elements reveal whether they are active, and let $\sigma^{>e}$ be the set of elements that appears after e in the order σ . Then, when e reveals that it is active, the OCRS π' decides to accept it if $A^{<e} \cup R(x \wedge \mathbf{1}_{\sigma^{>e}}) \in \mathcal{F}_e$ (for a random realization of the random set $R(x \wedge \mathbf{1}_{\sigma^{>e}})$).

Observation 3.12. *The OCRS π' always selects a set $S \in \mathcal{F}$.*

Proof. Fix the set A of active elements, the families $\{\mathcal{F}_e\}_{e \in N}$ and the order σ in which elements reveal whether they are active. We prove that the observation holds for any possible choice for these fixed values.

Let S' be the set of elements selected by π . The fact that π' selects some element $e \in N$ (i.e., $e \in S$) means that $(S \cap \sigma^{<e}) \cup R(x \wedge \mathbf{1}_{\sigma^{>e}}) \in \mathcal{F}_e$ for some realization of the random set $R(x \wedge \mathbf{1}_{\sigma^{>e}})$. Hence, $S \cap \sigma^{<e} \in \mathcal{F}_e$, which implies that π also selects e (i.e., $e \in S'$). This argument implies that $S \subseteq S'$. The observation now follows by the down-monotonicity of $\mathcal{F}$ since π always returns a set in $\mathcal{F}$. $\square$

Next, we need to prove the following technical lemma.

Lemma 3.13. *Let X_1, X_2 be two random sets that never contain elements $e, e' \in N$. Additionally, let Y_1 and Y_2 be two random sets distributed like $R(x \wedge \mathbf{1}_{\{e'\}})$ which are independent from each other and from X_1, X_2 and $\mathcal{F}_e$. Then, assuming $\Pr[X_2 \in \mathcal{F}_e] > 0$,*

$$\mathbb{E}[f(e \mid X_1 \cup Y_1) \mid X_2 \cup Y_2 \in \mathcal{F}_e] \leq \mathbb{E}[f(e \mid X_1 \cup Y_1) \mid X_2 \cup Y_1 \in \mathcal{F}_e] .$$

Proof. It is enough to prove the lemma for fixed values of X_1, X_2 and $\mathcal{F}_e$ obeying $X_2 \in \mathcal{F}_e$. If $X_2 \cup \{e'\} \in \mathcal{F}_e$, then the conditions $X_2 \cup Y_2 \in \mathcal{F}_e$ and $X_2 \cup Y_1 \in \mathcal{F}_e$ always hold and the lemma is trivial, thus, we also assume $X_2 \cup \{e'\} \notin \mathcal{F}_e$. Then, conditioned on the values of X_1, X_2 and $\mathcal{F}_e$:

$$\begin{aligned} \mathbb{E}[f(e \mid X_1 \cup Y_1) \mid X_2 \cup Y_2 \in \mathcal{F}_e] &\leq \mathbb{E}[f(e \mid X_1) \mid X_2 \cup Y_2 \in \mathcal{F}_e] = \mathbb{E}[f(e \mid X_1) \mid Y_2 = \emptyset] \\ &= \mathbb{E}[f(e \mid X_1 \cup Y_1) \mid Y_1 = \emptyset] \\ &= \mathbb{E}[f(e \mid X_1 \cup Y_1) \mid X_2 \cup Y_1 \in \mathcal{F}_e] , \end{aligned}$$

where the inequality holds by submodularity. $\square$

Using the last lemma we can now prove the following one, which lower bounds the contribution of an element e to the value of the solution selected by π' .

Lemma 3.14. *Let $A \sim R(x)$ be the set of active elements, and let S be the output of π' given A and an arbitrary online adversary. Then, for every element $e \in N$ having $x_e > 0$ and $\Pr[\emptyset \in \mathcal{F}_e] > 0$, $\Pr[e \in S \mid e \in A] = \Pr[R(x) \setminus \{e\} \in \mathcal{F}_e]$ and $\mathbb{E}[f(e \mid A \cap \sigma_a^{\leq e}) \mid e \in S] \geq F(e \mid x \wedge \mathbf{1}_{\sigma_a^{\leq e}})$.*

Proof. Observe that the conditions $x_e > 0$ and $\Pr[\emptyset \in \mathcal{F}_e] > 0$ are equivalent to $\Pr[e \in S] > 0$, and thus, the second expectation we want to bound is well defined for every element e having $x_e > 0$ and $\Pr[\emptyset \in \mathcal{F}_e] > 0$.

The strategy of an online adversary for determining the order in which elements reveal whether they are active, up to the point when the element e does so, can be characterized by a binary tree with the following properties:

- Every leaf is marked by e .
- Every internal node is marked by some other element $e' \neq e$ and has two children called the “accepted child” and the “rejected child”.
- The path from the root of the tree to every leaf does not contain a single element more than once.

The semantics of understanding the tree as a strategy are as follows. The adversary starts at the root of the tree, and moves slowly down the tree till reaching a leaf. When the adversary is at some node e' it makes e' the next element that reveals whether it is active. If e' is accepted by the OCRS, then the adversary moves to its “accepted child”, otherwise it moves to its “rejected child”. Notice that strategies defined by such trees do not allow the adversary to use the information whether e' was active or not when the OCRS rejects it, however, since the activation of elements is independent and our OCRS behaves in the same way under both cases, this information cannot help the adversary.

In the rest of this proof, we show by induction on the number of leafs in the strategy of the adversary that the lemma holds for every online adversary. Let us start with the case of the single strategy having a single leaf. In this strategy the adversary makes e the first element to reveal whether it is active. This means that:

$$\Pr[e \in S \mid e \in A] = \Pr[R(x) \setminus \{e\} \in \mathcal{F}_e] ,$$

and

$$\begin{aligned} \mathbb{E}[f(e \mid A \cap \sigma_a^{\leq e}) \mid e \in S] &= \mathbb{E}[f(e \mid A \cap \sigma_a^{\leq e}) \mid R(x) \setminus \{e\} \in \mathcal{F}_e, e \in A] \\ &= \mathbb{E}[f(e \mid A \cap \sigma_a^{\leq e})] = F(e \mid x \wedge \mathbf{1}_{\sigma_a^{\leq e}}) , \end{aligned}$$

where the second equality holds since $A \cap \sigma_a^{<e}$ is independent of the conditions $R(x) \setminus \{e\} \in \mathcal{F}_e$ and $e \in A$.

Next, assume the lemma holds for strategies having $\ell - 1 \geq 1$ leaves, and let us prove it for a strategy $\mathcal{T}$ having ℓ leaves. Since every internal node of $\mathcal{T}$ has two children, $\mathcal{T}$ must contain a node w (marked by an element $e' \neq e$) having two leaves as children. Let $\mathcal{T}'$ be the strategy resulting from $\mathcal{T}$ by removing the two children of w and making w itself a leaf (by marking it with e). Since $\mathcal{T}'$ has $\ell - 1$ leaves, it obeys the lemma by the induction hypothesis. Hence, it is enough to show that $\Pr[e \in S \mid e \in A]$ is identical under both $\mathcal{T}$ and $\mathcal{T}'$ and $\mathbb{E}[f(e \mid A \cap \sigma_a^{<e}) \mid e \in S]$ is at least as large under $\mathcal{T}$ as under $\mathcal{T}'$.

Let χ_w be the event that the adversary reaches the node w in its strategy. Clearly $\Pr[\chi_w \mid \mathcal{T}] = \Pr[\chi_w \mid \mathcal{T}']$, where by conditioning a probability on a strategy of the adversary we mean that the probability is calculated for the case that the adversary uses this strategy. If the last two probabilities are strictly smaller than 1 we also have:

$$\Pr[e \in S \mid e \in A, \neg\chi_w, \mathcal{T}] = \Pr[e \in S \mid e \in A, \neg\chi_w, \mathcal{T}'] .$$

On the other hand, observe that for both strategies $\mathcal{T}$ and $\mathcal{T}'$ the event χ_w implies the same (deterministic) fixed values for the set B_w of the elements that revealed whether they are active up to the point that the adversary reached w , and the set $A_w = A \cap B_w$. Thus, whenever $\Pr[\chi_w \mid \mathcal{T}] = \Pr[\chi_w \mid \mathcal{T}'] > 0$:

$$\begin{aligned} \Pr[e \in S \mid e \in A, \chi_w, \mathcal{T}] &= \Pr[A_w \cup (R(x) \setminus (B_w \cup \{e\})) \in \mathcal{F}_e] \\ &= \Pr[A_w \cup (A \cap \{e'\}) \cup (R(x) \setminus (B_w \cup \{e, e'\})) \in \mathcal{F}_e] = \Pr[e \in S \mid e \in A, \chi_w, \mathcal{T}'] . \end{aligned}$$

In conclusion, we got: $\Pr[e \in S \mid e \in A, \mathcal{T}] = \Pr[e \in S \mid e \in A, \mathcal{T}']$. Notice that both probabilities must be strictly positive by the induction hypothesis.

By Bayes' law we now get:

$$\begin{aligned} \Pr[\chi_w \mid \mathcal{T}, e \in S] &= \Pr[\chi_w \mid e \in A, \mathcal{T}, e \in S] = \Pr[e \in S \mid e \in A, \mathcal{T}, \chi_w] \cdot \frac{\Pr[\chi_w \mid e \in A, \mathcal{T}]}{\Pr[e \in S \mid e \in A, \mathcal{T}]} \\ &= \Pr[e \in S \mid e \in A, \mathcal{T}', \chi_w] \cdot \frac{\Pr[\chi_w \mid e \in A, \mathcal{T}']}{\Pr[e \in S \mid e \in A, \mathcal{T}']} = \Pr[\chi_w \mid e \in A, \mathcal{T}', e \in S] = \Pr[\chi_w \mid \mathcal{T}', e \in S] , \end{aligned}$$

and when the probabilities $\Pr[\chi_w \mid \mathcal{T}, e \in S] = \Pr[\chi_w \mid \mathcal{T}', e \in S]$ are strictly smaller than 1 we also have:

$$\mathbb{E}[f(e \mid A \cap \sigma_a^{<e}) \mid e \in S, \neg\chi_w, \mathcal{T}] = \mathbb{E}[f(e \mid A \cap \sigma_a^{<e}) \mid e \in S, \neg\chi_w, \mathcal{T}'] .$$

Hence, to prove that $\mathbb{E}[f(e \mid A \cap \sigma_a^{<e}) \mid e \in S]$ is at least as large under $\mathcal{T}$ as under $\mathcal{T}'$ we are only left to show the inequality:

$$\mathbb{E}[f(e \mid A \cap \sigma_a^{<e}) \mid e \in S, \chi_w, \mathcal{T}] \geq \mathbb{E}[f(e \mid A \cap \sigma_a^{<e}) \mid e \in S, \chi_w, \mathcal{T}']$$

whenever $\Pr[\chi_w \mid \mathcal{T}, e \in S] > 0$. The last inequality holds since:

$$\begin{aligned} &\mathbb{E}[f(e \mid A \cap \sigma_a^{<e}) \mid e \in S, \chi_w, \mathcal{T}'] \\ &= \mathbb{E}\left[f(e \mid (A_w \cup (A \setminus B_w)) \cap \sigma_a^{<e}) \mid A_w \cup ((R(x) \setminus B_w) \setminus \{e\}) \in \mathcal{F}_e\right] \\ &\leq \mathbb{E}\left[f(e \mid (A_w \cup (A \setminus B_w)) \cap \sigma_a^{<e}) \mid A_w \cup (A \cap \{e'\}) \cup ((R(x) \setminus B_w) \setminus \{e', e\}) \in \mathcal{F}_e\right] \\ &= \mathbb{E}[f(e \mid A \cap \sigma_a^{<e}) \mid e \in S, \chi_w, \mathcal{T}] , \end{aligned}$$

where the inequality holds by Lemma 3.13. □

We are now ready to prove Theorem 3.10.

Proof of Theorem 3.10. If $c = 0$, then the theorem is trivial. Thus, we may assume $c > 0$. Recall that S is the set produced by the OCRS π' , and let $A \sim R(x)$ be the set of active elements. Since, π is a (b, c) -selectable element-monotone contention resolution scheme, $A \setminus \{e\}$ must belong to $\mathcal{F}_e$, for every element $e \in N$, with a positive probability. By the down-monotonicity of $\mathcal{F}_e$, we get that $\emptyset \in \mathcal{F}_e$ for every $e \in N$.

Next, observe that:

$$\begin{aligned} \mathbb{E}[f(S)] &= f(\emptyset) + \sum_{e \in N, x_e > 0} \Pr[e \in S] \cdot \mathbb{E}[f(e \mid S \cap \sigma_a^{\leq e}) \mid e \in S] \\ &\geq f(\emptyset) + \sum_{e \in N, x_e > 0} \Pr[e \in S] \cdot \mathbb{E}[f(e \mid A \cap \sigma_a^{\leq e}) \mid e \in S] \\ &\geq f(\emptyset) + \sum_{e \in N, x_e > 0} \Pr[R(x) \setminus \{e\} \in \mathcal{F}_e, e \in A] \cdot F(e \mid x \wedge \mathbf{1}_{\sigma_a^{\leq e}}) \\ &= f(\emptyset) + \sum_{e \in N, x_e > 0} x_e \cdot \Pr[R(x) \setminus \{e\} \in \mathcal{F}_e] \cdot F(e \mid x \wedge \mathbf{1}_{\sigma_a^{\leq e}}) . \end{aligned}$$

where the first inequality holds by submodularity and the second by Lemma 3.14. Since π is a (b, c) -selectable contention resolution scheme, the probability $\Pr[R(x) \setminus \{e\} \in \mathcal{F}_e]$ must be at least c for every element $e \in N$. Additionally, since we assumed that x is non-reducible, we also have $F(e \mid x \wedge \mathbf{1}_{\sigma_a^{\leq e}}) \geq 0$ for every element $e \in N$ obeying $x_e > 0$. Plugging both observations into the previous inequality gives:

$$\mathbb{E}[f(S)] \geq f(\emptyset) + \sum_{e \in N, x_e > 0} c \cdot x_e \cdot F(e \mid x \wedge \mathbf{1}_{\sigma_a^{\leq e}}) = c \cdot F(x) + (1 - c) \cdot f(\emptyset) \geq c \cdot F(x) . \quad \square$$

4 Details on applications

In this section we prove the theorems stated in Section 1.3 and provide some additional information. We reuse notation introduced in Section 1.3.

4.1 Prophet inequalities for Bayesian online selection problems

Proof of Theorem 1.12

We start by introducing a relaxation for the expected value of the prophet, *i.e.*, $\mathbb{E}[\max\{\sum_{e \in S} Z_e \mid S \in \mathcal{F}\}]$. Similar relaxation techniques to the one we use here have been used previously for different related problems (see, *e.g.*, [24]).

Let $I^* \in \operatorname{argmax}\{\sum_{e \in I} Z_i \mid I \in \mathcal{F}\}$ be an optimal (random) set for the prophet, and we define $p_e^* = \Pr[e \in I^*]$ for $e \in N$. Our relaxation seeks probabilities p_e that have the role of p_e^* . Observe first that since $I^* \in \mathcal{F}$ with probability 1, we have that p^* is a convex combination of characteristic vectors of feasible sets, and hence, $p^* \in P$ (we recall that P is a relaxation for $\mathcal{F}$).

Our relaxation assigns an optimistic objective value to each probability vector $p \in P$. More precisely, for any $e \in N$ we assume in the relaxation that element e gets selected when Z_e takes one of its p_e -fraction of highest values. In particular, if Z_e follows a continuous distribution with cumulative distribution function F_e , then we assume that Z_e gets selected whenever $Z_e \geq q_e := F_e^{-1}(1 - p_e)$ and, consequently, the contribution of e to the objective of the prophet is $g_e(p_e) = \int_{q_e}^{\infty} x \cdot dF_e(x)$. More generally, when Z_e does not follow a continuous distribution, the expected value of Z_e on the highest p_e -fraction of realizations can

be described as follows:

$$g_e(p_e) = (p_e - (1 - F_e(q_e))) \cdot q_e + \int_{x \in (q_e, \infty)} x \cdot dF_e(x), \quad \text{where}$$

$$q_e = q_e(p_e) = \min\{\alpha \mid F_e(\alpha) \geq 1 - p_e\}.$$

In words, we assume that e gets selected whenever $Z_e > q_e$. Moreover, if $Z_e = q_e$, then e gets selected with probability $p_e - 1 + F_e(q_e)$.

Putting things together, the relaxation we consider is the following.

$$\max_{p \in P} \sum_{e \in N} g_e(p_e) \tag{4}$$

By the above discussion this is indeed a relaxation. Moreover, one can easily observe that g_e is a concave functions, and thus, in many settings one can efficiently obtain a near-optimal solution to this relaxation using convex optimization techniques. However, if we are only interested in proving the existence of a prophet inequality as stated in Theorem 1.12, we do not need an efficient procedure to solve (4).

Let $x \in P$ be an optimal solution to (4). We now create an algorithm for the Bayesian online selection problem based on the point $x \in P$ and the c -selectable OCRS for P which exists by assumption. Whenever an element $e \in N$ reveals in the Bayesian online selection problem we say that e is active if its random variable Z_e realizes within the largest x_e -fraction of realizations. More formally, e is active if either:

- (i) $Z_e > q_e(x)$, or
- (ii) if $Z_e = q_e(x)$ (assuming $\Pr[Z_e = q_e(x)] > 0$), we toss a coin and declare e to be active with probability $\frac{x_e - 1 + F_e(q_e(x))}{\Pr[Z_e = q_e(x)]}$.

Let $A \subseteq N$ be the random set of active elements. Observe that A is distributed like $R(x)$, the random subset of N that contains each element $e \in N$ with probability x_e independently of the others. Also, by definition of active elements we have

$$\mathbb{E}[Z_e \cdot \mathbf{1}_{e \in A}] = g_e(x_e) \quad \forall e \in N. \tag{5}$$

Our algorithm for the Bayesian online selection problem applies a c -selectable OCRS to the set A to obtain a random set $I \subseteq A$, $I \in \mathcal{F}$. To prove the theorem, we show that the expected value of I is at least a c -fraction of the optimal value of (4), i.e., $\mathbb{E}[\sum_{e \in I} Z_e] \geq c \cdot \sum_{e \in N} g_e(x_e)$.

Since the OCRS is c -selectable, we have

$$\Pr[e \in I] \geq c \cdot x_e \quad \forall e \in N. \tag{6}$$

A key observation is that the distribution of Z_e conditioned on $e \in I$ is the same as the distribution of Z_e conditioned on $e \in A$. This follows from the fact that the OCRS does not consider the precise value of Z_e , but only knows whether $e \in A$ or not. In particular, this implies

$$\mathbb{E}[Z_e \mid e \in I] = \mathbb{E}[Z_e \mid e \in A] \quad \forall e \in N. \tag{7}$$

Combining the above observations we get

$$\begin{aligned}
\mathbb{E} \left[\sum_{e \in I} Z_e \right] &= \sum_{e \in N} \Pr[e \in I] \cdot \mathbb{E}[Z_e \mid e \in I] \\
&\geq c \cdot \sum_{e \in N} x_e \cdot \mathbb{E}[Z_e \mid e \in A] && \text{(by (6) and (7))} \\
&= c \cdot \sum_{e \in N} \Pr[e \in A] \cdot \mathbb{E}[Z_e \cdot \mathbf{1}_{e \in A} \mid e \in A] && \text{(since } \Pr[e \in A] = x_e \text{)} \\
&= c \cdot \sum_{e \in N} g_e(x_e) && \text{(by (5)) .}
\end{aligned}$$

The last inequality shows that our procedure is worse by at most a factor of c compared to the value of the relaxation (4), which completes the proof.

4.2 Oblivious posted pricing mechanisms

Proof of Theorem 1.15

The proof for Theorem 1.15 goes along the same lines as the proof of Theorem 1.12 presented in Section 4.1, but uses a different relaxation. The relaxation we employ is the same as the one used by Yan [24]. For completeness and ease of understanding we replicate some of the arguments in [24] and refer to the excellent discussion of this relaxation in [24] for more details about it.

Consider the random set of agents $I^* \subseteq N$, $I^* \in \mathcal{F}$ served by Myerson's mechanism, *i.e.*, an optimal truthful mechanism, and let $q_e^* = \Pr[e \in I^*]$ be the probability that e gets served. Since only feasible subsets of agents can be served, we have $q^* \in P$, because P is a relaxation of $\mathcal{F}$.

For this fixed q^* we can now define independent mechanism design problems for the different agents as follows. For each $e \in N$, we are interested in finding a price distribution (for the price offered to e), that maximizes the expected revenue under the constraint that e gets served with probability equal to q_e^* . Based on results by Myerson [22], it follows that the optimal price distribution can be chosen to be a two-price distribution, which can be determined through a well-known technique in mechanism design known as *ironing* (we refer to [24] for details). We denote by $R_e(q_e^*)$ the expected revenue of this optimal distribution, which can be shown to be concave in q_e^* , and by $\mathcal{D}_e(q_e^*)$ the distribution itself. Since the family of these independent mechanism design problems for the agents is less constrained than the original BSMD, in which we also had to make sure that the set of all served agents is in $\mathcal{F}$, we have that $\sum_{e \in N} R_e(q_e^*)$ is an upper bound to the expected revenue of an optimal mechanism for the original BSMD (we refer to [24] for a formal proof). Hence, the following is a convex relaxation of the original BSMD.

$$\max \left\{ \sum_{e \in N} R_e(q_e) \mid q \in P \right\} \tag{8}$$

The COPM we construct needs probabilities $x \in P$ such that $\sum_{e \in N} R_e(x_e)$ is an upper bound on the revenue of Myerson's mechanism. By the above discussion, this holds for $x = q^*$, or for x being an optimal solution to (8). To make this step constructive one can follow, for example, the sampling-based approach of Chawla et al. [9]; they estimate the probabilities q^* by running Myerson's mechanism sufficiently many times. Alternatively, one could use convex optimization techniques to optimize the relaxation (8).¹¹ Our

¹¹We highlight that even if only an approximately optimal $x \in P$ can be obtained, *i.e.*, one such that $\alpha \cdot \sum_{e \in N} R_e(x_e)$ upper bounds the optimal revenue for some $\alpha > 1$, then all of what follows still goes through simply with an additional loss of a factor α . This will lead to a COPM that is at most a factor of $\alpha \cdot c$ worse than Myerson's mechanism.

COPM is randomized and defined by the following randomization over tuples $(p, \mathcal{F}')$. The price vector $p \in \mathbb{R}_{\geq 0}^N$ is drawn according to the product distribution $\times_{e \in N} \mathcal{D}_e(x_e)$, where each p_e for $e \in N$ is drawn independently according to the two-price distribution $\mathcal{D}_e(x_e)$. The family $\mathcal{F}'$ is chosen to be equal to the family $\mathcal{F}_x$ of the c -selectable OCSR for the point $x \in P$. Hence, if our OCSR is deterministic, then also $\mathcal{F}' = \mathcal{F}_x$ is deterministic, in which case the randomization of our COPM is solely on the price vector p .

To prove Theorem 1.15 we show that our COPM has an expected revenue of at least $c \cdot \sum_{e \in N} R_e(x_e)$. We call an agent $e \in N$ *active* if its personal (random) valuation Z_e is at least as large as the (random) price p_e . In other words, we say that e is active if it would accept the offer presented by our COPM. By the definition of the distributions $\mathcal{D}_e(x_e)$, we have that each agent e is active with probability x_e , independently of all other agents. Notice that an agent $e \in N$ being active does not imply that e gets selected, because a COPM is allowed to reject an agent if feasibility in $\mathcal{F}'$ is not maintained. However, because our OCSR is c -selectable, we have that for any $e \in N$ with probability at least c , the agent e can be added to the agents served so far no matter which subset of the active agents has been already served; furthermore, this event is independent of whether e is active itself. The expected revenue that our COPM gets from agent e is therefore at least $c \cdot R_e(x_e)$, which completes the proof.

Extension to other objectives

Finally, we notice that, as highlighted by Yan [24], the used relaxation can easily be extended to a much larger class of objective functions that are decomposable with respect to the agents. More precisely, these are objectives of the form $\mathbb{E}[\sum_{e \in N} \mathbf{1}_{Z_e \geq p_e} \cdot g_e(Z_e, p_e)]$, where for $e \in N$, g_e is a function of the (random) valuation Z_e of e and the price p_e offered to e . In particular, the maximization of revenue which we discussed above corresponds to $g_e(Z_e, p_e) = p_e$. Similarly, one can deal with welfare maximization or surplus maximization by defining $g_e(Z_e, p_e) = Z_e$ and $g_e(Z_e, p_e) = Z_e - p_e$, respectively. The above reasoning why our COPM is at most a factor of c worse than the optimal truthful mechanism extends to such objectives without modifications.

4.3 Stochastic probing

We begin with the proof of Theorem 1.17. Consider the following relaxation. We use w in the relaxation to denote the natural extension of w to vectors (formally, $w(x) = \sum_{e \in N} w(e) \cdot x_e$). Additionally, the relaxation uses the expression $p \circ x$ to denote the element-wise multiplication of the probabilities vector p and the variables vector x . Clearly this relaxation can be solved efficiently when there is a separation oracle for the polytopes P_{in} and P_{out} .

$$(R1) \quad \begin{aligned} \max \quad & w(p \circ x) \\ & p \circ x \in P_{in} \\ & x \in P_{out} \\ & x \in [0, 1]^N \end{aligned}$$

The following lemma proves an important property of (R1). The proof of this lemma is based on the observation that one feasible solution for (R1) is the vector $x \in [0, 1]^N$ in which x_e is equal to the marginal probability that the optimal algorithm probes element e .

Lemma 4.1 (Claim 3.1 of [17]). *The optimal value of (R1) upper bounds the the expected performance of the optimal algorithm for the weighted stochastic probing problem.*

Let x^* be an optimal solution for (R1). By Lemma 4.1, to prove Theorem 1.17 we only need to show an algorithm for the weighted stochastic probing problem which finds a solution of expected weight at least $b(c_{in} \cdot c_{out}) \cdot w(p \circ x^*)$. Our algorithm for the weighted stochastic probing problem is given as Algorithm 1.

Algorithm 1: Probing Algorithm

```

1 Let  $A_{out}$  be a random set distributed like  $R(bx^*)$ .
2 Let  $\hat{\mathcal{F}}_{out} \subseteq \mathcal{F}_{out}$  be an instantiation of the random family  $\mathcal{F}_{\pi_{out}, bx^*}$ .
3 Let  $\hat{\mathcal{F}}_{in} \subseteq \mathcal{F}_{in}$  be an instantiation of the random family  $\mathcal{F}_{\pi_{in}, p \circ (bx^*)}$ .
4 Let  $Q, S \leftarrow \emptyset$ .
5 for every element  $e \in N$  in the order chosen by the adversary do
6   if  $e \in A_{out}$ ,  $S \cup \{e\} \in \hat{\mathcal{F}}_{in}$  and  $Q \cup \{e\} \in \hat{\mathcal{F}}_{out}$  then
7     Add  $e$  to  $Q$  and probe  $e$ .
8   if  $e$  is active then Add  $e$  to  $S$ .
9 return  $S$ .
```

The next observation shows that Algorithm 1 is a legal algorithm for the weighted stochastic probing problem.

Observation 4.2. *The following always hold when Algorithm 1 terminates:*

- The set Q of probed elements is in $\hat{\mathcal{F}}_{out} \subseteq \mathcal{F}_{out}$.
- The set S of selected elements is in $\hat{\mathcal{F}}_{in} \subseteq \mathcal{F}_{in}$.
- The set S contains exactly the active elements of Q .

In the rest of the section we use notation and results introduced in Section 3. Let $\bar{\pi}_{in}$ and $\bar{\pi}_{out}$ denote the characteristic CRSs of the greedy OCRSs π_{in} and π_{out} , respectively. We would like to use $\bar{\pi}_{in}$ and $\bar{\pi}_{out}$ to lower bound the value of a subset of S , and through that subset also the value of S . To achieve that objective we first need to describe the said subset of S as an expression of $\bar{\pi}_{in}$ and $\bar{\pi}_{out}$.

Let A_{in} be the set of elements that belong to A_{out} and are also active. One can observe that A_{in} is distributed like $R(p \circ x)$. Additionally, let us couple the randomness of Algorithm 1 and the CRSs $\bar{\pi}_{in}$ and $\bar{\pi}_{out}$ as follows: we use the same instantiations for the random families $\mathcal{F}_{\pi_{in}, p \circ (bx^*)}$ and $\mathcal{F}_{\pi_{out}, bx^*}$ in both Algorithm 1 and the definitions of $\bar{\pi}_{in}$ and $\bar{\pi}_{out}$, respectively. Notice that this coupling implies:

$$\bar{\pi}_{in}(A_{in}) = \{e \in A_{in} \mid I \cup \{e\} \in \hat{\mathcal{F}}_{in} \ \forall I \subseteq A_{in}, I \in \hat{\mathcal{F}}_{in}\}$$

and

$$\bar{\pi}_{out}(A_{out}) = \{e \in A_{out} \mid I \cup \{e\} \in \hat{\mathcal{F}}_{out} \ \forall I \subseteq A_{out}, I \in \hat{\mathcal{F}}_{out}\}.$$

Observation 4.3. $\bar{\pi}_{in}(A_{in}) \cap \bar{\pi}_{out}(A_{out}) \subseteq S$.

Proof. Consider an element $e \in \bar{\pi}_{in}(A_{in}) \cap \bar{\pi}_{out}(A_{out})$. Since $e \in \bar{\pi}_{in}(A_{in}) \subseteq A_{in}$, e must be both active and in A_{out} . Let S_e and Q_e denote the sets S and Q immediately before e is processed by Algorithm 1. Clearly $S_e \subseteq A_{in}$ and $S_e \in \hat{\mathcal{F}}_{in}$. Together with the fact that $e \in \bar{\pi}_{in}(A_{in})$, these observations imply that $S_e \cup \{e\} \in \hat{\mathcal{F}}_{in}$. An analogous arguments shows also that $Q_e \cup \{e\} \in \hat{\mathcal{F}}_{out}$. Hence, we proved that e is an active element obeying all the conditions on Line 6 when processed by Algorithm 1, and therefore, e is added to S by the algorithm. $\square$

We defined A_{out} as a random set distributed like $R(bx^*)$ and A_{in} as the intersection of A_{out} and the set of active elements. However, for the purpose of analyzing the value of the set $\bar{\pi}_{in}(A_{in}) \cap \bar{\pi}_{out}(A_{out})$ we can assume any construction procedure that results in the same joint distribution of A_{in} and A_{out} . The following is a convenient construction that we are going to use from this point on: the set A_{in} is a random set distributed like $R(p \circ (bx^*))$. The set A_{out} is calculated by starting with A_{in} and adding to it every

element $e \notin A_{in}$ with probability $bx_e^*(1 - p_e)/(1 - bp_ex_e^*)$, independently. Notice that this construction indeed produces the same joint distribution of the sets A_{in} and A_{out} as the original construction. For ease of notation, let us denote by z a vector in $[0, 1]^N$ defined by: $z_e = bx_e^*(1 - p_e)/(1 - bp_ex_e^*)$ for every $e \in N$. Using this notation we get that A_{out} has the same distribution as $A_{in} \cup R(z)$.

The new construction of A_{out} implies that A_{out} is a random function of A_{in} , and so is the expression $\bar{\pi}_{in}(A_{in}) \cap \bar{\pi}_{out}(A_{out})$. Thus, we can define a new CRS $\bar{\pi}$ for P_{in} by the equality $\bar{\pi}(A_{in}) = \bar{\pi}_{in}(A_{in}) \cap \bar{\pi}_{out}(A_{out})$. Notice that $\bar{\pi}$ is a true CRS for P_{in} in the sense that it always outputs a set in $\mathcal{F}_{in}$ since $\bar{\pi}_{in}(A_{in})$ is guaranteed to be in $\mathcal{F}_{in}$. Let us now study the properties of $\bar{\pi}$.

Lemma 4.4. *The CRS $\bar{\pi}$ is monotone.*

Proof. We need to show that every element $e \in N$ and two sets $T_1 \subseteq T_2 \subseteq N$ obey the inequality:

$$\begin{aligned} \Pr[e \in \bar{\pi}(A_{in}) \mid A_{in} = T_1] &= \Pr[e \in \bar{\pi}_{in}(A_{in}) \cap \bar{\pi}_{out}(A_{out}) \mid A_{in} = T_1] \\ &\geq \Pr[e \in \bar{\pi}_{in}(A_{in}) \cap \bar{\pi}_{out}(A_{out}) \mid A_{in} = T_2] = \Pr[e \in \bar{\pi}(A_{in}) \mid A_{in} = T_2] . \end{aligned}$$

This is true since:

$$\begin{aligned} \Pr[e \in \bar{\pi}_{in}(A_{in}) \cap \bar{\pi}_{out}(A_{out}) \mid A_{in} = T_1] &= \Pr[e \in \bar{\pi}_{in}(T_1) \cap \bar{\pi}_{out}(T_1 \cup R(z))] \\ &= \Pr[e \in \bar{\pi}_{in}(T_1)] \cdot \Pr[e \in \bar{\pi}_{out}(T_1 \cup R(z))] \geq \Pr[e \in \bar{\pi}_{in}(T_2)] \cdot \Pr[e \in \bar{\pi}_{out}(T_2 \cup R(z))] \\ &= \Pr[e \in \bar{\pi}_{in}(T_2) \cap \bar{\pi}_{out}(T_2 \cup R(z))] = \Pr[e \in \bar{\pi}_{in}(A_{in}) \cap \bar{\pi}_{out}(A_{out}) \mid A_{in} = T_2] , \end{aligned}$$

where the inequality follows since both $\bar{\pi}_{in}$ and $\bar{\pi}_{out}$ are monotone by Lemma 3.4. $\square$

The following lemma shows that $\bar{\pi}$ obeys a weak variant of balanceness.

Definition 4.5. *A CRS π for a polytope P is (x, c) -balanced for a vector x and $c \in [0, 1]$ if $\Pr[e \in \pi(R(x)) \mid e \in R(x)] \geq c$ for every element $e \in N$ having $x_e > 0$.*

Lemma 4.6. *The CRS $\bar{\pi}$ is $(p \circ (bx^*), c_{in} \cdot c_{out})$ -balanced.*

Proof. Since A_{in} is distributed like $R(p \circ (bx^*))$, we need to show that $\Pr[e \in \bar{\pi}_{in}(A_{in}) \cap \bar{\pi}_{out}(A_{out}) \mid e \in A_{in}] \geq c_1 \cdot c_2$ holds for every element $e \in N$ having $bp_ex_e^* > 0$.

Since $\bar{\pi}_{in}$ and $\bar{\pi}_{out}$ are monotone by Lemma 3.4, we get that both $\Pr[e \in \bar{\pi}_{in}(T)]$ and $\Pr[e \in \bar{\pi}_{out}(T \cup R(z))]$ are decreasing functions of T as long as e is in T . Thus, by the FKG inequality:

$$\begin{aligned} \Pr[e \in \bar{\pi}_{in}(A_{in}) \cap \bar{\pi}_{out}(A_{out}) \mid e \in A_{in}] &= \Pr[e \in \bar{\pi}_{in}(A_{in}) \cap \bar{\pi}_{out}(A_{in} \cup R(z)) \mid e \in A_{in}] \\ &\geq \Pr[e \in \bar{\pi}_{in}(A_{in}) \mid e \in A_{in}] \cdot \Pr[e \in \bar{\pi}_{out}(A_{in} \cup R(z)) \mid e \in A_{in}] \\ &= \Pr[e \in \bar{\pi}_{in}(A_{in}) \mid e \in A_{in}] \cdot \Pr[e \in \bar{\pi}_{out}(A_{out}) \mid e \in A_{out}] , \end{aligned}$$

where the last equality uses the fact that the membership of every element in the sets A_{in} and $R(z)$ is independent from the membership of other elements in these sets.

The set A_{in} is distributed like $R(p \circ (bx^*))$ and the vector $p \circ (bx^*)$ is inside the polytope bP_{in} . Thus, since $\bar{\pi}_{in}$ is (b, c_{in}) -balanced by Lemma 3.4, we get: $\Pr[e \in \bar{\pi}_{in}(A_{in}) \mid e \in A_{in}] \geq c_{in}$. Similarly, the set A_{out} is distributed like $R(bx^*)$ and the vector bx^* is inside the polytope bP_{out} . Thus, since $\bar{\pi}_{out}$ is (b, c_{out}) -balanced by Lemma 3.4, we get: $\Pr[e \in \bar{\pi}_{out}(A_{out}) \mid e \in A_{out}] \geq c_{out}$. The lemma now follows by combining the above inequalities. $\square$

In the following corollary we use w to denote its natural extension to sets, i.e., $w(T) = \sum_{e \in T} w(e)$.

Corollary 4.7. $\mathbb{E}[w(S)] \geq b(c_{in} \cdot c_{out}) \cdot w(p \circ x^*)$.

Proof. By Lemma 4.6, linearity of the expectation and the observation that A_{in} is distributed like $R(p \circ (bx^*))$, we get:

$$\begin{aligned} \mathbb{E}[w(\pi(A_{in}))] &= \sum_{e \in N} w(e) \cdot \Pr[e \in \pi(R(p \circ (bx^*)))] \\ &= \sum_{e \in N} w(e) \cdot \Pr[e \in R(p \circ (bx^*))] \cdot \Pr[e \in \pi(R(p \circ (bx^*))) \mid e \in R(p \circ (bx^*))] \\ &\geq \sum_{e \in N} w(e) \cdot (p_e \circ (bx_e^*)) \cdot (c_{in} \cdot c_{out}) = b(c_{in} \cdot c_{out}) \cdot w(p \circ x^*) . \end{aligned}$$

Since all the weights are non-negative and $\pi(A_{in}) = \pi_{in}(A_{in}) \cap \pi_{out}(A_{out}) \subseteq S$ by Observation 4.3, the last inequality implies $\mathbb{E}[w(S)] \geq \mathbb{E}[w(\pi(A_{in}))] \geq b(c_{in} \cdot c_{out}) \cdot w(p \circ x^*)$. $\square$

Following the above discussion, Theorem 1.17 is implied by the last corollary. We can now prove Theorem 1.18 as a direct consequence of Theorem 1.17.

Proof of Theorem 1.18. Consider a down-closed set $\mathcal{F} \subseteq 2^N$ containing every set $T \subseteq N$ if and only if all the elements of T can be queried while respecting the deadlines. Formally,

$$\mathcal{F} = \left\{ T \subseteq 2^N \mid \forall_{1 \leq d \leq |N|} |\{e \in T \mid d_e \leq d\}| \leq d \right\} .$$

By definition $(N, \mathcal{F})$ is a laminar matroid, thus, by Theorem 1.8 there exists a $(b, 1 - b)$ -selectable greedy OCRS for its matroid polytope $P_{\mathcal{F}}$. Together with the existence of π_{out} we get, by Theorem 1.9, a $(b, (1 - b)c_{out})$ -selectable greedy OCRS for the polytope $P'_{out} = P_{out} \cap P_{\mathcal{F}}$. Notice that this polytope is a relaxation of the down-closed family $\mathcal{F}'_{out} = \mathcal{F}_{out} \cap \mathcal{F}$. Moreover, P'_{out} has a separation oracle whenever P_{out} has such an oracle.

For convenience, let us use $(\mathcal{F}_1, \mathcal{F}_2)$ -probing as a shorthand for the weighted stochastic probing problem with $\mathcal{F}_1$ and $\mathcal{F}_2$ as the inner and outer constraints, respectively. Consider the best algorithm for $(\mathcal{F}_{in}, \mathcal{F}_{out})$ -probing with deadlines. Since this algorithm probes with respect to the deadlines, the set of elements it probes must be in $\mathcal{F}$. Hence, the same algorithm is also an algorithm for $(\mathcal{F}_{in}, \mathcal{F}'_{out})$ -probing. Thus, by Theorem 1.17 we have an algorithm ALG for $(\mathcal{F}_{in}, \mathcal{F}'_{out})$ -probing whose approximation ratio is $b(1 - b) \cdot c_{in} \cdot c_{out}$ compared to the best algorithm for $(\mathcal{F}_{in}, \mathcal{F}_{out})$ -probing with deadlines. Moreover, the approximation ratio of ALG holds regardless of the order in which ALG can probe elements.

The algorithm we suggest for $(\mathcal{F}_{in}, \mathcal{F}_{out})$ -probing with deadlines is ALG when we allow it to probe elements in increasing deadlines order. We have already proved that ALG has the approximation ratio guaranteed by the theorem, so we only need to explain why does it respect the deadlines. Assume towards a contradiction that ALG probes element e after time d_e . This means that ALG probes a set T of at least d_e elements before it probes e . However, since ALG can probe elements only in increasing deadlines order, all the elements of $T \cup \{e\}$ have a deadline of at most d_e . The last observation implies that $T \cup \{e\} \notin \mathcal{F}$, which contradicts the fact that the set of elements probed by ALG is always in $\mathcal{F}'_{out} \subseteq \mathcal{F}$. $\square$

The rest of this section is devoted to proving Theorem 1.19. Let $f: 2^N \rightarrow \mathbb{R}_{\geq 0}$ be the non-negative monotone submodular objective function of the problem, i.e., the value of an output S of the probing algorithm is $f(S)$. We need to introduce an extension of f to $[0, 1]^N$ studied by [6].

$$f^+(x) = \max \left\{ \sum_{T \subseteq N} \alpha_T \cdot f(T) \mid \sum_{T \subseteq N} \alpha_T \leq 1, \alpha_T \geq 0 \text{ and } \forall_{e \in N} \sum_{e \in T \subseteq N} \alpha_T \leq x_e \right\} .$$

Intuitively, f^+ is equal to the largest possible expected value of f over a distribution of sets in which every element e appears with a marginal probability at most x_e . Using this extension, we can now introduce a variant of the relaxation (R1) that works for monotone submodular objectives.

$$(R2) \quad \max_{\substack{p \circ x \in P_{in} \\ x \in P_{out} \\ x \in [0, 1]^N}} f^+(p \circ x)$$

Let us explain why we use the extension f^+ in (R2) instead of the simpler multilinear extension. An algorithm for the submodular stochastic probing may choose the next element to probe based on the set of elements previously probed and the results of these probes. Thus, the membership of elements in the solution produced by the algorithm is not independent, and this is captured by f^+ . Using this intuition, the work of [1] implies the following counterpart of Lemma 4.1.

Lemma 4.8. *The optimal value of (R2) upper bounds the expected performance of the optimal algorithm for the submodular stochastic probing problem.*

Additionally, [1] also shows that a variant of the continuous greedy algorithm of [7] can be used to find a point x obeying all the constraints of (R2) and also the inequality $F(p \circ x) \geq (1 - e^{-1}) \cdot f^+(p \circ x^*)$, where F is the multilinear extension of f and x^* is the optimal solution for (R2). The same argument can also be used to show that by stopping the continuous greedy algorithm at time b , instead of letting it reach time 1, one gets a point $\tilde{x} \in [0, 1]^N$ obeying $\tilde{p} \circ x \in bP_{in}$, $\tilde{x} \in bP_{out}$ and $F(p \circ \tilde{x}) \geq (1 - e^{-b} - o(1)) \cdot f^+(p \circ x^*)$.¹²

The algorithm we use to prove Theorem 1.19 is Algorithm 1 modified by replacing the vector bx^* with $\tilde{x}$. Additionally, we use the same definition of the CRS $\bar{\pi}$ given above, and observe that Lemmata 4.4 and 4.6 still apply to this CRS (with $\tilde{x}$ replacing bx^* in the appropriate places). To analyze the output of this CRS we observe that the proof of [10] for Lemma 3.5 in fact proves the following stronger version of the lemma. Note that this version strictly generalizes Lemma 3.5 since being (b, c) -balanced is equivalent to being (x, c) -balanced for every vector $x \in bP$.

Lemma 4.9. *For every given non-negative submodular function $f: 2^N \rightarrow \mathbb{R}_{\geq 0}$, there exists a function $\eta_f: 2^N \rightarrow 2^N$ that always returns a subset of its argument (i.e., $\eta_f(S) \subseteq S$ for every $S \subseteq N$) having the following property. For every input vector $x \in [0, 1]^N$ and monotone (x, c) -balanced CRS π :*

$$\mathbb{E}[f(\eta_f(\pi(R(x))))] \geq c \cdot F(x) ,$$

where $F(x)$ is the multilinear extension of f .

We are now ready to prove the next lemma, which together with Lemma 4.8, proves Theorem 1.19.

Lemma 4.10. $\mathbb{E}[f(S)] \geq (c_{in} \cdot c_{out}) \cdot F(p \circ \tilde{x})$, where S is the output set of the modified Algorithm 1.

Proof. Let S' be the set $\eta_f(\pi(A_{in})) \subseteq \pi(A_{in}) = \bar{\pi}_{in}(A_{in}) \cap \bar{\pi}_{out}(A_{out}) \subseteq S$. By combining Lemmata 4.4, 4.6 and 4.9 and using the observation that A_{in} is distributed like $R(p \circ \tilde{x})$, we get:

$$\mathbb{E}[f(S')] \geq (c_{in} \cdot c_{out}) \cdot F(p \circ \tilde{x}) .$$

Using the monotonicity of f we now get: $\mathbb{E}[f(S)] \geq \mathbb{E}[f(S')] \geq (c_{in} \cdot c_{out}) \cdot F(p \circ \tilde{x})$. □

¹²The idea of running the continuous greedy for less time instead of scaling its output was first introduced by [14].

References

- [1] M. Adamczyk, M. Sviridenko, and J. Ward. Submodular stochastic probing on matroids. In *31st Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 29–40, 2014.
- [2] M. Babaioff, N. Immorlica, and R. Kleinberg. Matroids, secretary problems, and online mechanisms. In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 434–443, 2007.
- [3] N. Bansal, A. Gupta, J. Li, J. Mestre, V. Nagarajan, and A. Rudra. When LP is the cure for your matching woes: Improved bounds for stochastic matchings. *Algorithmica*, 63(4):733–762, 2012.
- [4] M. Bateni, M. Hajiaghayi, and M. Zadimoghaddam. Submodular secretary problem and extensions. *ACM Transactions on Algorithms*, 9(4):32:1–32:23, October 2013.
- [5] N. Buchbinder, M. Feldman, J. Naor, and R. Schwartz. Submodular maximization with cardinality constraints. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1433–1452, 2014.
- [6] G. Călinescu, C. Chekuri, M. Pál, and J. Vondrák. Maximizing a submodular set function subject to a matroid constraint. In *Proceedings of the 12th International Integer Programming and Combinatorial Optimization Conference (IPCO)*, pages 182–196, 2007.
- [7] G. Călinescu, C. Chekuri, M. Pál, and J. Vondrák. Maximizing a monotone submodular function subject to a matroid constraint. *SIAM Journal on Computing*, 40(6):1740–1766, 2011.
- [8] S. Chawla, J. Hartline, D. Malec, and B. Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC)*, pages 311–320, New York, NY, USA, 2010. ACM.
- [9] S. Chawla, J. Hartline, D. Malec, and B. Sivan. Multi-parameter mechanism design and sequential posted pricing. <http://arxiv.org/abs/0907.2435v2>, 2010.
- [10] C. Chekuri, J. Vondrák, and R. Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. *SIAM Journal on Computing*, 43(6):1831–1879, 2014.
- [11] P. Dagum, R. Karp, M. Luby, and S. Ross. An optimal algorithm for Monte Carlo estimation. *SIAM Journal on Computing*, 29(5):1484–1496, 2000.
- [12] U. Feige, V. S. Mirrokni, and J. Vondrák. Maximizing non-monotone submodular functions. *SIAM Journal on Computing*, 40(4):1133–1153, 2011.
- [13] M. Feldman. *Maximization Problems with Submodular Objective Functions*. PhD thesis, Computer Science Department, Technion - Israel Institute of Technology, July 2013.
- [14] M. Feldman, J. Naor, and R. Schwartz. A unified continuous greedy algorithm for submodular maximization. In *Proceedings of the 52nd IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 570–579, 2011.
- [15] M. Feldman, O. Svensson, and R. Zenklusen. A simple $O(\log \log(\text{rank}))$ -competitive algorithm for the matroid secretary problem. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1189–1201, 2015.

- [16] M. Feldman and R. Zenklusen. The submodular secretary problem goes linear. In *Proceedings of 56th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2015.
- [17] A. Gupta and V. Nagarajan. A stochastic probing problem with applications. In *Proceedings of Integer Programming and Combinatorial Optimization (IPCO)*, pages 205–216, 2013.
- [18] R. Kleinberg and S. M. Weinberg. Matroid prophet inequalities. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC)*, pages 123–136, 2012.
- [19] U. Krengel and L. Sucheston. On semiamarts, amarts, and processes with finite value. In J. Kuelbs, editor, *Probability on Banach Spaces*, volume 4 of *Advances in probability and related topics*, pages 197–266. Marcel Dekker Inc., 1978.
- [20] O. Lachish. $O(\log \log \text{Rank})$ competitive ratio for the matroid secretary problem. In *Proceedings of the 55th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 326–335. IEEE, 2014.
- [21] B. Lehmann, D. Lehmann, and N. Nisan. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, 55(2):270 – 296, 2006. Mini Special Issue: Electronic Market Design.
- [22] R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- [23] J. A. Soto. Matroid secretary problem in the random-assignment model. *SIAM Journal on Computing*, 42(1):178–211, 2013.
- [24] Q. Yan. Mechanism design via correlation gap. In *Proceedings of the Twenty-second Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 710–719. SIAM, 2011.

A Definitions of Matroidal Concepts

Recall that a matroid $M = (N, \mathcal{F})$ is a tuple consisting of a finite ground set N , and a nonempty family $\mathcal{F} \subseteq 2^N$ of subsets of the ground set, called *independent* sets, which satisfy:

- (i) $I \subseteq J \in \mathcal{F} \Rightarrow I \in \mathcal{F}$, and
- (ii) $I, J \in \mathcal{F}, |I| > |J| \Rightarrow \exists e \in I \setminus J \text{ s.t. } J \cup \{e\} \in \mathcal{F}$.

The *rank* of a set $S \subseteq 2^N$ is the size of a maximum cardinality independent subset of S . The rank function of the matroid M is a function $\text{rank}: 2^N \rightarrow \mathbb{Z}_{\geq 0}$ (where $\mathbb{Z}_{\geq 0}$ is the set of all non-negative integers) assigning each set its rank. More formally,

$$\text{rank}(S) = \max\{|I| \mid I \in \mathcal{F}, I \subseteq S\} .$$

Observe that the rank of an independent set is equal to its size. The rank of the matroid M itself is defined as $\text{rank}(N)$. Notice that $\text{rank}(S) \leq \text{rank}(N)$ for every subset $S \subseteq N$. A set $S \subseteq N$ is called a *base* of M if it is independent and has maximum rank, i.e., $|S| = \text{rank}(S) = \text{rank}(N)$.

We say that an element $e \in N$ is *spanned* by a set $S \subseteq N$ if adding e to S does not increase the rank of S . On the other hand, the *span* of a subset $S \subseteq N$ is the set of elements that are spanned by it. More formally, the span of a set S is $\text{span}(S) = \{e \in N \mid \text{rank}(S + e) = \text{rank}(S)\}$.

Given a subset $N' \subseteq N$, the restriction of M to N' , denoted by $M|_{N'}$, is the matroid obtained from M by keeping only the elements of N' . Formally, $M|_{N'}$ is the matroid $(N', \mathcal{F} \cap 2^{N'})$. On the other hand, contracting N' in M results in another matroid, denoted by M/N' , over the ground set $N \setminus N'$. A set is independent in M/N' if and only if adding a base of N' to it results in an independent set of M . It turns

out that this definition is independent of the base that is chosen for N' . Formally, M/N' is the matroid $(N \setminus N', \mathcal{F}')$, where:

$$\mathcal{F}' = \{S \subseteq N \setminus N' \mid \text{rank}(S \cup N') = |S| + \text{rank}(N')\} \ .$$