

# Budget-Feasible Mechanism Design for Non-Monotone Submodular Objectives: Offline and Online

GEORGIOS AMANATIDIS\*, Centrum Wiskunde & Informatica (CWI), The Netherlands

PIETER KLEER\*, Centrum Wiskunde & Informatica (CWI), The Netherlands

GUIDO SCHÄFER, Centrum Wiskunde & Informatica (CWI) and Vrije Universiteit Amsterdam, The Netherlands

The framework of budget-feasible mechanism design studies procurement auctions where the auctioneer (buyer) aims to maximize his valuation function subject to a hard budget constraint. We study the problem of designing truthful mechanisms that have good approximation guarantees and never pay the participating agents (sellers) more than the budget. We focus on the case of general (non-monotone) submodular valuation functions and derive the first truthful, budget-feasible and  $O(1)$ -approximation mechanisms that run in polynomial time in the value query model, for both offline and online auctions. Since the introduction of the problem by Singer [40], obtaining efficient mechanisms for objectives that go beyond the class of monotone submodular functions has been elusive. Prior to our work, the only  $O(1)$ -approximation mechanism known for non-monotone submodular objectives required an exponential number of value queries.

At the heart of our approach lies a novel greedy algorithm for non-monotone submodular maximization under a knapsack constraint. Our algorithm builds two candidate solutions simultaneously (to achieve a good approximation), yet ensures that agents cannot jump from one solution to the other (to implicitly enforce truthfulness). Ours is the first mechanism for the problem where—crucially—the agents are not ordered according to their marginal value per cost. This allows us to appropriately adapt these ideas to the online setting as well.

To further illustrate the applicability of our approach, we also consider the case where additional feasibility constraints are present, e.g., at most  $k$  agents can be selected. We obtain  $O(p)$ -approximation mechanisms for both monotone and non-monotone submodular objectives, when the feasible solutions are independent sets of a  $p$ -system. With the exception of additive valuation functions, no mechanisms were known for this setting prior to our work. Finally, we provide lower bounds suggesting that, when one cares about non-trivial approximation guarantees in polynomial time, our results are asymptotically best possible.

CCS Concepts: • **Theory of computation** → **Algorithmic mechanism design**; *Algorithmic game theory*; Approximation algorithms analysis; Online algorithms.

Additional Key Words and Phrases: procurement auctions; budget-feasible mechanism design; non-monotone submodular maximization

## ACM Reference Format:

Georgios Amanatidis, Pieter Kleer, and Guido Schäfer. 2019. Budget-Feasible Mechanism Design for Non-Monotone Submodular Objectives: Offline and Online. In *ACM EC '19: ACM Conference on Economics and Computation (EC '19)*, June 24–28, 2019, Phoenix, AZ, USA. ACM, New York, NY, USA, 19 pages. <https://doi.org/10.1145/3328526.3329622>

\*G. Amanatidis and P. Kleer are supported by the NWO Gravitation Project NETWORKS, Grant Number 024.002.003.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

EC '19, June 24–28, 2019, Phoenix, AZ, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6792-9/19/06...\$15.00

<https://doi.org/10.1145/3328526.3329622>

## 1 INTRODUCTION

We consider the problem of designing *budget-feasible mechanisms* for a natural model of procurement auctions. In this model, an auctioneer is interested in buying services (or goods) from a set of agents  $A$ . Each agent  $i \in A$  specifies a cost  $c_i$  to be paid by the buyer for using his service; crucially, these costs are assumed to be private information. The auctioneer has a budget  $B$  and a valuation function  $v(\cdot)$ , where  $v(S)$  specifies the value derived from the services of the agents in  $S \subseteq A$ . Given the (reported) costs of the agents, the goal of the auctioneer is to choose a *budget-feasible* subset  $S \subseteq A$  of the agents, such that the valuation  $v(S)$  is maximized. Budget-feasibility here means that  $\sum_{i \in S} p_i \leq B$ , where  $p_i$  is the payment issued from the mechanism to agent  $i$ .

Note that the agents might try to extract larger payments from the mechanism by misreporting their actual costs—which of course is undesirable from the auctioneer's perspective. The goal, therefore, is to design budget-feasible mechanisms that (i) elicit truthful reporting of the costs by all agents, and (ii) achieve a good approximation with respect to the optimal value for the auctioneer. What makes the problem so intriguing is the fact that truthfulness and budget-feasibility are two directly conflicting goals, since the former is achieved by paying as much as needed to make agents indifferent to lying (see Lemma 2.4). Indicatively, the use of the celebrated truthful VCG mechanism in this setting completely fails with respect to keeping the payments bounded [40].

The problem of designing budget-feasible mechanisms was introduced by Singer [40] and has received a lot of attention, both because of its theoretical appeal and of its relevance to several emerging application domains. A prominent such application is in crowdsourcing marketplaces (such as Mechanical Turk, Figure Eight and Clickworker) which provide online platforms to procure workforce (see [4, 26, 31]). Another application is in the context of influence maximization in social networks, where one seeks to select influential users (see [1, 41]).

We focus on the design of budget-feasible mechanisms for the general class of *non-monotone submodular* valuation functions. Submodular objectives constitute an important class of valuation functions as they satisfy the property of diminishing returns, which naturally arises in many settings. Most existing works make the assumption that the valuation functions are *monotone* (non-decreasing), i.e.,  $v(S) \leq v(T)$  for  $S \subseteq T$ . Although the monotonicity assumption makes sense in certain applications, there are several examples where it is violated. For example, in the context of influence maximization in social networks, adding more users to the selected set may sometimes result in negative influence (see [12]). The most prominent example of a non-monotone submodular objective studied in our setting is the *budgeted max-cut problem* [2, 18], where  $v(\cdot)$  is determined by the cuts of a given graph.

A natural generalization of this framework is to assume that the space of feasible sets has some structure, e.g., the feasible sets form a matroid. This variant has been studied only for additive valuation functions [1, 35], despite its wide range of applications varying from team formation to spectrum markets (see [35]). Here we study the problem for monotone and non-monotone submodular objectives under  $p$ -system constraints.

The purely algorithmic versions of these mechanism design problems ask for the maximization of a (non-monotone) submodular function subject to the constraint that the total cost of the selected agents does not exceed the budget; often referred to as a *knapsack constraint*. These problems are typically NP-hard, hence our focus is on approximation algorithms that compute a close to optimal solution in polynomial time. From an algorithmic point of view, most of these problems are well-understood and admit good approximations. However, it is not clear how to appropriately convert these algorithms into truthful, budget-feasible mechanisms and, up to this work, this goal had been elusive. Our results illustrate that for the mechanism design problems it is possible to achieve the same asymptotic guarantees that are known for their algorithmic counterparts.

**Our Contributions.** We derive the first budget-feasible and  $O(1)$ -approximate mechanisms for non-monotone submodular objectives, both for the offline and the online setting. Our results for the online setting hold for the well-studied *secretary model*, where the agents arrive in a uniformly random order. Our mechanisms run in polynomial time in the value query model. The highlights of this work are as follows:

- We obtain the first universally truthful, budget-feasible  $O(1)$ -approximation mechanism for non-monotone submodular objectives in the value query model.
- We derive the first universally truthful, budget-feasible  $O(1)$ -approximation *online* mechanism for non-monotone submodular objectives. As a consequence, we obtain the first  $O(1)$ -approximation algorithm for the non-monotone *Submodular Knapsack Secretary Problem* (see also Remark 4.3), a budget constrained variant of the infamous Secretary Problem.
- We give universally truthful, budget-feasible  $O(p)$ -approximation mechanisms for both monotone and non-monotone submodular objectives, when the feasible solutions are independent sets of a  $p$ -system. Beyond the additive case, nothing was known for this constrained setting.
- We provide lower bounds illustrating that asymptotically our results are *best possible*. On a high level, only trivial guarantees can be achieved in polynomial time if one goes beyond the class of general submodular functions or imposes constraints beyond downward closed systems.

**Technical Challenges.** It should be noted that for monotone submodular objectives all known mechanisms essentially use the same greedy subroutine introduced by Singer [40]: Sort all agents in decreasing order of marginal value per cost and pick as many agents as possible before hitting some carefully selected threshold. This is a simplified version of the optimal greedy algorithm of Sviridenko [43] and indeed gives non-trivial approximation guarantees. Further, due to its simplicity it also has the other desired properties of truthfulness, individual rationality, and budget-feasibility.

Unfortunately, it is easy to construct examples where running such a greedy algorithm for a non-monotone objective results in a solution of arbitrarily poor quality. The algorithmic state-of-the-art for non-monotone submodular maximization under a knapsack constraint, e.g., [16, 24, 32], provides us with quite involved algorithms on continuous relaxations of the problem that seem very unlikely to yield monotone allocation rules, and thus truthful mechanisms. The only simple (and deterministic) exception is the two-pass greedy algorithm of Gupta et al. [28], where it is shown that running Sviridenko's greedy algorithm twice and then maximizing without the knapsack constraint is sufficient to get a deterministic 6-approximation algorithm. Despite being significantly simpler, however, this two-pass greedy algorithm still suffers with respect to monotonicity. More recently, several simple randomized greedy approaches for maximizing non-monotone submodular objectives subject to other (i.e., non-knapsack) constraints were proposed [14, 15, 22, 25]. However, these approaches are also not applicable here; we refer the interested reader to the discussion in the full version of our paper [3].

At the heart of our approach lies a novel deterministic greedy algorithm for non-monotone submodular maximization under a knapsack constraint. Our algorithm builds two candidate solutions *simultaneously*, yet prevents agents to jump from one solution to the other by changing their cost. To do the latter we offer each agent a take-it-or-leave-it price based on an estimate of the optimal value which we obtain by sampling. Moreover, this is the first mechanism for the problem where—crucially—the agents are not ordered with respect to their marginal value per cost. This further allows us to appropriately modify the algorithm and adapt it to the online secretary setting and to settings with additional feasibility constraints, while maintaining all its desired properties.

All of our mechanisms are randomized and, in fact, random sampling is an essential building block in our approach, as in previous related works [2, 6, 10, 34]. Designing *deterministic* mechanisms

seems very challenging and, beyond additive functions [17, 40], no deterministic, polynomial-time, budget-feasible  $O(1)$ -approximation mechanisms are known, except for some special cases [1, 2, 18, 29, 41]. Obtaining deterministic, budget-feasible,  $O(1)$ -approximation mechanisms—or showing that they do not exist—is an intriguing open problem.

**Related Work.** As mentioned above, the study of budget-feasible mechanisms was initiated by Singer [40], who gave a randomized  $O(1)$ -approximation mechanism for monotone submodular functions. Later, Chen et al. [17] significantly improved the approximation ratio and also suggested a deterministic  $O(1)$ -approximation mechanism, albeit with superpolynomial running time. Several follow-up results modified this deterministic mechanism so that it runs in polynomial time for special cases, including coverage functions [1, 41] and information gain functions [29]. For subadditive functions, Dobzinski et al. [18] suggested a  $O(\log^2 n)$ -approximation mechanism, and gave the first constant factor mechanisms for a special case of non-monotone objectives, namely cut functions. The factor for subadditive functions was later improved to  $O(\log n / \log \log n)$  by Bei et al. [10], who also gave a randomized  $O(1)$ -approximation mechanism for XOS functions, albeit in exponential time in the value query model, and further initiated the Bayesian analysis in this setting.<sup>1</sup> Amanatidis et al. [2] suggested  $O(1)$ -approximation mechanisms for a subclass of non-monotone submodular objectives, namely symmetric submodular objectives, however their approach does not seem to generalize beyond this subclass. For settings with additional combinatorial constraints, Amanatidis et al. [1] and Leonardi et al. [35] gave  $O(1)$ -approximation mechanisms for additive valuation functions subject to independent system constraints. There is also a line of related work under the *large market* assumption (where no participant can significantly affect the market outcome), which allows for mechanisms with improved performance (see, e.g., [4, 8, 26, 31, 42]).

The online version of the problem was introduced and studied by Badanidiyuru et al. [6] who give an  $O(1)$ -approximation mechanism for monotone submodular functions. This is closely related to the purely algorithmic version of the problem (i.e., without the incentives), namely the Submodular Knapsack Secretary Problem introduced by Bateni et al. [9] as a generalization of the Knapsack Secretary Problem [5]. Bateni et al. studied the problem for monotone and non-monotone submodular objectives, although their argument for the latter case is not sound (see Remark 4.3). While the monotone submodular case has been improved [23] and generalized [30], there is no follow-up work on the non-monotone case to the best of our knowledge.

On maximization of submodular functions subject to knapsack or other type of constraints, there is a vast literature, going back several decades (see, e.g., [39, 44]). Focusing on knapsack constraints, there is a rich line of recent work on developing algorithms on continuous relaxations of the problem (see, e.g., [16, 20, 24, 32] and references therein) achieving an  $\epsilon$ -approximation for non-monotone objectives. However, the most relevant recent work to ours is that of Gupta et al. [28] who proposed a deterministic 6-approximation algorithm for the non-monotone case, related on a high level to our main approach. Gupta et al. also gave algorithms for certain constrained secretary problems, although not with knapsack constraints. When  $\ell$  knapsack constraints and a  $p$ -system constraint are both present, the algorithmic state-of-the-art is a  $(p + 2\ell + 1)$ -approximation algorithm for the monotone submodular case due to Badanidiyuru and Vondrák [7] and a  $(2p + 2\ell + 1)(p + 1)/p$ -approximation algorithm for the non-monotone submodular case due to Mirzasoleiman et al. [37].

<sup>1</sup>Bei et al. [10] propose an  $O(1)$ -approximation mechanism for *non-decreasing* XOS objectives in the much stronger *demand query* model. However, they discuss how to extend their result to general XOS functions via the use of  $\hat{v}(S) = \max_{T \subseteq S} v(T)$ . It is easy to see that  $\hat{v}$  is non-decreasing and that  $S$  is an optimal solution of  $v$  if and only if it is a minimal optimal solution for  $\hat{v}$ . Moreover, Gupta et al. [27] proved that if  $v$  is general XOS then  $\hat{v}$  is monotone XOS. It should be noted that this transformation does not work for submodular functions [2]. Therefore, known results for monotone submodular functions do not extend to the non-monotone case, even in the demand query model.

## 2 PRELIMINARIES

We use  $A = [n] = \{1, 2, \dots, n\}$  to denote a set of  $n$  agents. Each agent  $i$  is associated with a private cost  $c_i$ , denoting the cost for participating in the solution. We consider a procurement auction setting, where the auctioneer is equipped with a valuation function  $v : 2^A \rightarrow \mathbb{Q}_{\geq 0}$  and a budget  $B > 0$ . For  $S \subseteq A$ ,  $v(S)$  is the value derived by the auctioneer if the set  $S$  is selected (for singletons, we will often write  $v(i)$  instead of  $v(\{i\})$ ). Therefore, the algorithmic goal in all the problems we study is to select a set  $S$  that maximizes  $v(S)$  subject to the constraint  $\sum_{i \in S} c_i \leq B$ . We assume oracle access to  $v$  via *value queries*, i.e., we assume the existence of a polynomial time value oracle that returns  $v(S)$  when given as input a set  $S$ .

A function  $v$  is *non-decreasing* (often referred to as *monotone*), if  $v(S) \leq v(T)$  for any  $S \subseteq T \subseteq A$ . We consider general (i.e., not necessarily monotone), normalized (i.e.,  $v(\emptyset) = 0$ ), non-negative submodular valuation functions. Since marginal values are extensively used, we adopt the shortcut  $v(i | S)$  for the marginal value of agent  $i$  with respect to the set  $S$ , i.e.,  $v(i | S) = v(S \cup \{i\}) - v(S)$ . The following three definitions of submodularity are equivalent. While definition (i) is the most standard, the other two alternative definitions will be useful later on.

*Definition 2.1.* A function  $v$ , defined on  $2^A$  for some set  $A$ , is *submodular* if and only if

- (i)  $v(i | S) \geq v(i | T)$  for all  $S \subseteq T \subseteq A$ , and  $i \notin T$ .
- (ii)  $v(S) + v(T) \geq v(S \cup T) + v(S \cap T)$  for all  $S, T \subseteq A$ .
- (iii)  $v(T) \leq v(S) + \sum_{i \in T \setminus S} v(i | S) - \sum_{i \in S \cap T} v(i | S \cup T \setminus \{i\})$  for all  $S, T \subseteq A$ .

In Section 6 we also deal with valuation functions that come from a superclass of submodular functions, namely *XOS* or *fractionally subadditive* functions. In particular, it is known that non-negative (monotone) submodular functions are a strict subset of (monotone) XOS functions [27, 33].

*Definition 2.2.* A function  $v$ , defined on  $2^A$  for some set  $A$ , is *XOS* or *fractionally subadditive*, if there exist additive functions  $\alpha_1, \dots, \alpha_r$ , for some finite  $r$ , such that  $v(S) = \max_{i \in [r]} \alpha_i(S)$ .

We often need to argue about optimal solutions of sub-instances of the original instance  $(A, v, \mathbf{c}, B)$ . Given a cost vector  $\mathbf{c}$ , and a subset  $X \subseteq A$ , we denote by  $\mathbf{c}_X$  the projection of  $\mathbf{c}$  on  $X$ , and by  $\mathbf{c}_{-X}$  the projection of  $\mathbf{c}$  on  $A \setminus X$ . By  $\text{OPT}(X, v, \mathbf{c}_X, B)$  we denote the value of an optimal solution to the problem restricted on  $X$ . Similarly,  $\text{OPT}(X, v, \infty)$  denotes the value of an optimal solution to the unconstrained version of the problem restricted on  $X$ . For the sake of readability, we usually drop the valuation function and the cost vector, and write  $\text{OPT}(X, B)$  and  $\text{OPT}(X, \infty)$ , respectively.

**Mechanism Design.** In the strategic version that we consider here, every agent  $i \in A$  only has his true cost  $c_i$  as private information. Hence, this is a *single-parameter environment*. A mechanism  $\mathcal{M} = (f, p)$  in our context consists of an outcome rule  $f$  and a payment rule  $p$ . Given a vector of cost declarations,  $\mathbf{b} = (b_i)_{i \in A}$ , where  $b_i$  denotes the cost reported by agent  $i$ , the outcome rule of the mechanism selects the set  $f(\mathbf{b}) \subseteq A$ . At the same time, it computes payments  $p(\mathbf{b}) = (p_i(\mathbf{b}))_{i \in A}$  where  $p_i(\mathbf{b})$  denotes the payment issued to agent  $i$ . Hence, the final utility of agent  $i$  is  $p_i(\mathbf{b}) - c_i$ .

Unless stated otherwise, our mechanisms run in polynomial time in the value query model. Further properties we want to enforce in our mechanism design problem are the following.

*Definition 2.3.* A mechanism  $\mathcal{M} = (f, p)$  is

- *truthful*, if reporting  $c_i$  is a dominant strategy for every agent  $i$ .
- *individually rational*, if  $p_i(\mathbf{b}) \geq 0$  for every  $i \in A$ , and  $p_i(\mathbf{b}) \geq c_i$ , for every  $i \in f(\mathbf{b})$ .
- *budget-feasible*, if  $\sum_{i \in A} p_i(\mathbf{b}) \leq B$  for every  $\mathbf{b}$ .

For our randomized mechanisms we use the strong notion of *universal truthfulness*, which means that the mechanism is a probability distribution over deterministic truthful mechanisms. As all the

mechanisms we suggest are universally truthful, we will consistently use  $\mathbf{c} = (c_i)_{i \in A}$  rather than  $\mathbf{b} = (b_i)_{i \in A}$  for the declared costs in their description and analysis.

To design truthful mechanisms for single-parameter environments, we use a characterization by Myerson [38]. We say that an outcome rule  $f$  is *monotone*, if for every agent  $i \in A$ , and any vector of cost declarations  $\mathbf{b}$ , if  $i \in f(\mathbf{b})$ , then  $i \in f(b'_i, \mathbf{b}_{-i})$  for  $b'_i \leq b_i$ . That is, if an agent  $i$  is selected by declaring cost  $b_i$ , then he should still be selected by declaring a lower cost. Myerson's lemma, below, implies that monotone algorithms admit truthful payment schemes (often referred to as *threshold payments*). This greatly simplifies the design of truthful mechanisms, as one may focus on constructing monotone algorithms rather than having to worry about the payment scheme. For all of our mechanisms, we assume that the underlying payment scheme is given by Myerson's lemma.

**LEMMA 2.4** (Myerson [38]). *Given a monotone algorithm  $f$ , there is a unique payment scheme  $p$ , such that  $(f, p)$  is a truthful and individually rational mechanism, given by*

$$p_i(\mathbf{b}) = \begin{cases} \sup_{b'_i \in [c_i, \infty)} \{b'_i : i \in f(b'_i, \mathbf{b}_{-i})\}, & \text{if } i \in f(\mathbf{b}), \\ 0, & \text{otherwise.} \end{cases}$$

**Remark 2.5.** We may assume, without loss of generality, that in any given instance all the costs are upper bounded by the budget. To see this notice that neither our mechanisms nor the optimal offline solution will ever consider any agent with cost higher than  $B$ . Furthermore, no agent has an incentive to misreport a very high true cost. Indeed, due to budget-feasibility, if agent  $i$  reports a cost  $b_i \leq B$  instead of his true cost  $c_i > B$  and is selected, then he has utility  $p_i(\mathbf{b}) - c_i < B - B = 0$ . Thus, in all of our mechanisms we implicitly assume a preprocessing step that removes all the agents with declared costs exceeding  $B$ . The resulting instance (given as input to the corresponding mechanism) has the same set of optimal solutions subject to the budget constraint as the original one. Note that in the case of the online mechanism GENSM-ONLINE rejecting such agents as they arrive suffices.

**Remark 2.6.** We should stress that wherever tie-breaking is needed (e.g., in lines 3 and 10 of SIMULTANEOUS GREEDY, during the execution of the auxiliary algorithms ALG<sub>1</sub>, ALG<sub>2</sub> and ALG<sub>3</sub>, etc.), we assume the consistent use of a tie-breaking rule that is *independent of the declared costs*. An obvious such choice would be a deterministic lexicographic tie-breaking rule.

Due to space restrictions, several of our proofs are deferred to the full version of this paper [3].

### 3 AN EFFICIENT MECHANISM FOR SUBMODULAR OBJECTIVES

The main result of this section is the first  $O(1)$ -approximation mechanism (termed GENSM-MAIN below) for non-monotone submodular valuation functions.

**THEOREM 3.1.** *GENSM-MAIN is a universally truthful, individually rational, budget-feasible,  $O(1)$ -approximation mechanism.*

At the heart of our approach lies a novel greedy algorithm for non-monotone submodular maximization under a knapsack constraint (SIMULTANEOUS GREEDY below). As we mentioned in the Introduction, all known mechanisms use the same greedy subroutine: sort all agents in decreasing order of marginal value per cost and pick as many agents as possible before hitting some threshold. While for monotone submodular objectives this gives a non-trivial approximation guarantee, for non-monotone objectives may result in arbitrarily bad solutions. Moreover, continuous algorithmic approaches for non-monotone submodular maximization under a knapsack constraint [24, 32] seem very unlikely to yield monotone allocation rules, and thus truthful mechanisms. The only algorithm that is conceptually close to our approach is the two-pass greedy algorithm of Gupta et al. [28], that runs Sviridenko's greedy algorithm twice and then maximizes without the knapsack

constraint to get a deterministic 6-approximate solution. The intuition behind this approach is that submodularity prevents the greedy algorithm from getting stuck in consecutive “bad” local maxima. Despite being significantly simpler, however, this two-pass greedy algorithm still suffers irreparably with respect to monotonicity, as it allows agents to jump from one solution to the other by changing their cost.

Here we introduce **SIMULTANEOUS GREEDY**, a greedy mechanism that builds two candidate solutions *simultaneously*. While the analysis of Gupta et al. [28] does not apply here (our solutions are neither built sequentially nor according to the standard greedy algorithm), the way we obtain our approximation guarantee is of the same flavor: at least one of the solutions will contain an approximately optimal set. At the same time **SIMULTANEOUS GREEDY** prevents agents to choose their favorite candidate solution by misreporting their cost. To achieve that, we offer each agent a take-it-or-leave-it price based on an estimate  $x$  of the optimal value which we obtain by sampling. This is the first mechanism for the problem where it is crucial that the agents are *not* ordered with respect to their marginal value per cost. This will further allow us to appropriately modify **SIMULTANEOUS GREEDY** for the *online* setting of Section 4 while maintaining all its desired properties.

The parameter  $\beta$  is later set to 9.185 in order to get the approximation factor of Corollary 3.8 but, otherwise, our analysis is independent of its value.  $\text{ALG}_2$  in line 9 can be any approximation algorithm for *unconstrained* non-monotone submodular maximization. In particular, here we may use the deterministic 2-approximation algorithm of Buchbinder and Feldman [13].

---

**SIMULTANEOUS GREEDY**( $D, v, c_D, B, x$ )

---

```

1  $S_1 = S_2 = \emptyset; B_1 = B_2 = B; U = D$                                 /* each  $S_j$  has its own budget  $B_j$  */
2 while  $\max_{i \in U, j \in \{1,2\}} v(i|S_j) > 0$  do
3   Let  $(i, j) \in \arg \max_{i \in U, j \in \{1,2\}} v(i|S_j)$ 
4   if  $c_i \leq \frac{\beta B}{x} v(i|S_j) \leq B_j$  then
5      $S_j = S_j \cup \{i\}$ 
6      $B_j = B_j - \frac{\beta B}{x} v(i|S_j)$ 
7    $U = U \setminus \{i\}$ 
8 for  $j \in \{1, 2\}$  do
9    $T_j = \text{ALG}_2(S_j)$                                 /* a 2-approximate solution with respect to  $\text{OPT}(S_j, v, c_{S_j}, \infty)$  */
10 Let  $S$  be the best solution among  $S_1, S_2, T_1, T_2$ 
11 return  $S$ 

```

---

Ideally, we would like the rate parameter  $x$  to be close to  $\text{OPT}(A, B)$  and also to be robust in the sense that no single agent can significantly affect its value. To achieve that, **SAMPLE-THEN-GREEDY** randomly partitions the set of agents into two sets  $A_1$  and  $A_2$ , then approximately solves the problem on  $A_1$  to obtain an estimate of  $\text{OPT}(A_1, B)$ , and finally uses this  $x$  to set the threshold rate for **SIMULTANEOUS GREEDY** on  $A_2$ .

$\text{ALG}_1$  in line 2 can be any approximation algorithm for non-monotone submodular maximization subject to a knapsack constraint. In particular, here we may use the  $e$ -approximation algorithm of Kulik et al. [32] (also see Remark 3.9).

---

**SAMPLE-THEN-GREEDY**( $A, v, c, B$ )

---

```

1 Put each agent of  $A$  in either  $A_1$  or  $A_2$  independently at random with probability  $\frac{1}{2}$ 
2  $x = v(\text{ALG}_1(A_1))$                                 /* an  $e$ -approximation of  $\text{OPT}(A_1, v, c_{A_1}, B)$  */
3 return SIMULTANEOUS GREEDY( $A_2, v, c_{A_2}, B, x$ )

```

---

Lemma 3.6 in Subsection 3.1, due to Bei et al. [10] and Leonardi et al. [34], guarantees that with high probability both  $A_1$  and  $A_2$  contain enough value subject to the budget constraint for things to work, as long as no agent is too valuable. The latter leads to the final mechanism GENSM-MAIN (GENERAL SUBMODULAR-MAIN) that randomizes between all the above and just returning a best singleton.

---

GENSM-MAIN( $A, v, c, B$ )

---

- 1 With probability  $p = 0.201$  : **return**  $i^* \in \arg \max_{i \in A} v(i)$
  - 2 With probability  $1 - p$  : **return** SAMPLE-THEN-GREEDY( $A, v, c, B$ )
- 

### 3.1 Proving the Properties of GENSM-MAIN

We fix some additional notation to facilitate the presentation of the proofs. We use  $(D, v, c, B, x)$  for a generic instance given to SIMULTANEOUS GREEDY and  $S$  for the set returned. By  $i_1, i_2, \dots, i_t$  we denote the sequence of agents of  $D$  examined during this execution of the algorithm in this exact order. All the agents of  $S$  clearly appear within this sequence, so for any particular  $\ell \in S$  we have that  $\ell = i_k$  for some  $k$ . By  $j_k$  we denote the index  $\hat{j}$  picked during the  $k$ th execution of line 3 of SIMULTANEOUS GREEDY, while by  $S_{j_k}^{(k)}$  and  $B_{j_k}^{(k)}$  we denote the set  $S_{j_k}$  and its remaining budget, respectively, at that time. Conventionally, we use notation like  $S_{j_k}^{(k+1)}$  to denote  $S_{j_k}$  right after the  $k$ th execution of line 7, even if line 3 is never executed more than  $k$  times. Recall that we use a tie-breaking rule that is independent of the costs, as mentioned in Remark 2.6.

LEMMA 3.2. *The allocation rule defined by SIMULTANEOUS GREEDY is monotone. Thus, using the threshold payments of Myerson's lemma, the resulting mechanism is truthful and individually rational.*

In all the following statements, when we refer to mechanisms, we always assume threshold payments. Before we study the total payment, we should point out that enforcing budget-feasibility has been the main source of technical difficulties in the budget-feasible mechanism design literature. A significant advantage of the take-it-or-leave-it approach used in threshold mechanisms like SIMULTANEOUS GREEDY is that the budget-feasibility becomes much more manageable. To some extent this comes at the expense of the approximation guarantee and its analysis, but also offers some additional flexibility that will be explored in Sections 4 and 5.

LEMMA 3.3. *The mechanism SIMULTANEOUS GREEDY is budget-feasible.*

PROOF. Let  $S$  be the set returned given the instance  $(D, v, c, B, x)$  and fix  $i_k \in S$ . We claim that the payment  $p_{i_k}(c)$  is exactly  $\pi_k = \frac{\beta B}{x} v(i_k | S_{j_k}^{(k)})$ , i.e.,  $i_k \in S$  if and only if he bids  $c'_{i_k} \leq \pi_k$ . First note that  $i_k$  cannot affect the time when he is examined by the mechanism or which agents come before him. So, since  $c_{-i_k}$  is fixed, during the  $k$ th execution of line 3, he is always "offered"  $\pi_k$ ; either he accepts, i.e.,  $c'_{i_k} \leq \pi_k$ , and the algorithm proceeds in the exact same way as with  $c'_{i_k} = c_{i_k}$  (see also the proof of Lemma 3.2) or he rejects, i.e.,  $c'_{i_k} > \pi_k$ , and he is removed from the active set of agents. Once an agent is removed, however, he is never reexamined and thus, if  $c'_{i_k} > \pi_k$  then  $i_k$  is not in the winning set.

Recall that  $S$  can be any of  $S_1, S_2, T_1, T_2$ . We will show that all four sets are budget-feasible. Let  $T_1 = \{i_{a_1}, i_{a_2}, \dots, i_{a_{|T_1|}}\}$  and  $S_1 = \{i_{b_1}, i_{b_2}, \dots, i_{b_{|S_1|}}\}$ , where  $(a_i)_{i=1}^{|T_1|}$  is a subsequence of  $(b_i)_{i=1}^{|S_1|}$  which is a subsequence of  $1, 2, \dots, t$ . Recall that the budget  $B_1$  for  $S_1$  is never exhausted. We have

$$\sum_{\tau=1}^{|T_1|} \pi_{a_\tau} \leq \sum_{\tau=1}^{|S_1|} \pi_{b_\tau} = \sum_{\tau=1}^{|S_1|} \frac{\beta B}{x} v(i_{b_\tau} | S_{j_{b_\tau}}^{(b_\tau)}) = B - B_1^{(|S_1|+1)} \leq B.$$

The first and the second sum represent the total payment when  $S = T_1$  and when  $S = S_1$ , respectively. The budget-feasibility of  $T_2$  and  $S_2$  is proved in the exact same way.  $\square$



**COROLLARY 3.4.** *The mechanism GENSM-MAIN is universally truthful, individually rational and budget-feasible.*

**LEMMA 3.5.** *If there is a positive integer  $\ell$  such that  $\max_{i \in D} v(i) < \frac{x}{\ell \cdot \beta}$ , then SIMULTANEOUS GREEDY( $D, v, c, B, x$ ) outputs a set  $S$  such that*

$$v(S) \geq \min \left\{ \frac{\ell x}{(\ell + 1)\beta}, \frac{1}{6} \left( \text{OPT}(D, B) - \frac{2x}{\beta} \right) \right\}.$$

**PROOF.** Let  $t$  be the number of times line 3 was executed. At the end of the  $t$ th iteration,  $U$  is the set of agents never examined. That is,  $U$  only contains agents that have non-positive marginal utilities with respect to  $S_1^{(t+1)}$  and  $S_2^{(t+1)}$ . For the sake of readability, we henceforth use  $S_1$  and  $S_2$  to denote  $S_1^{(t+1)}$  and  $S_2^{(t+1)}$ , respectively. Let  $R = D \setminus (U \cup S_1 \cup S_2)$  be the agents  $i_k$  that were considered at some point by the mechanism but were *rejected*, i.e., not added to either  $S_1^{(k)}$  or  $S_2^{(k)}$ . We first partition  $R$  into two sets depending on *why* the corresponding agents were rejected. The set  $R_c = \{i_k \mid \frac{\beta B}{x} v(i_k | S_{j_k}^{(k)}) < c_i\}$  contains the agents rejected because the first inequality in line 4 was violated during the corresponding iteration. Similarly, the set  $R_B = \{i_k \mid B_{j_k}^{(k)} < \frac{\beta B}{x} v(i_k | S_{j_k}^{(k)})\}$  contains the agents rejected because the second inequality in line 4 was violated. Clearly,  $R = R_c \cup R_B$ . We consider two cases, depending on whether  $R_B$  is empty or not.

**Case 1.** Assume that  $R_B \neq \emptyset$  and let  $i_k \in R_B$ . That is, during the  $k$ th execution of line 3,  $(i, j) = (i_k, j_k)$ , but  $\frac{\beta B}{x} v(i_k | S_{j_k}^{(k)}) > B_{j_k}^{(k)}$ . Let  $S_{j_k}^{(k)} = \{i_{a_1}, i_{a_2}, \dots, i_{a_s}\}$ , where  $(a_i)_{i=1}^s$  is a subsequence of  $1, 2, \dots, t$ . Further, notice that, by its definition,  $B_{j_k}^{(k)} = B - \sum_{\tau=1}^s \frac{\beta B}{x} v(i_{a_\tau} | S_{j_k}^{(a_\tau)})$ . We have

$$\begin{aligned} v(S_{j_k}^{(k)}) &= \sum_{i=1}^s v(i_{a_i} | S_{j_k}^{(a_i)}) = \frac{x}{\beta B} (B - B_{j_k}^{(k)}) \\ &> \frac{x}{\beta B} (B - B_{j_k}^{(k)}) + \frac{x}{\beta B} B_{j_k}^{(k)} - v(i_k | S_{j_k}^{(k)}) = \frac{x}{\beta} - v(i_k | S_{j_k}^{(k)}). \end{aligned} \quad (1)$$

By submodularity and the way the agents in  $S_{j_k}^{(k)}$  are chosen, we have

$$v(i_{a_1}) = v(i_{a_1} | S_{j_k}^{(a_1)}) \geq v(i_{a_2} | S_{j_k}^{(a_2)}) \geq \dots \geq v(i_{a_s} | S_{j_k}^{(a_s)}) \geq v(i_k | S_{j_k}^{(k)}).$$

Yet, each one of these values is at most  $\max_{i \in D} v(i) < \frac{x}{\ell \cdot \beta}$ . Combining with (1), we have

$$\ell \cdot \max_{i \in D} v(i) < \frac{x}{\beta} \leq v(S_{j_k}^{(k)}) + v(i_k | S_{j_k}^{(k)}) \leq \sum_{\tau=1}^s v(i_{a_\tau} | S_{j_k}^{(a_\tau)}) + v(i_k | S_{j_k}^{(k)}) \leq (s+1) \cdot v(i_{a_1}),$$

and therefore, we conclude that  $|S_{j_k}^{(k)}| = s \geq \ell$ . Now we repeat the same argument for the average marginal value in the sum  $\sum_{\tau=1}^s v(i_{a_\tau} | S_{j_k}^{(a_\tau)})$ . Using the simple observation that the smallest term of a sum cannot exceed the average of the remaining terms, we get

$$\frac{x}{\beta} \leq v(S_{j_k}^{(k)}) + v(i_k | S_{j_k}^{(k)}) \leq v(S_{j_k}^{(k)}) + \frac{1}{s} \sum_{\tau=1}^s v(i_{a_\tau} | S_{j_k}^{(a_\tau)}) \leq \frac{s+1}{s} v(S_{j_k}^{(k)}) \leq \frac{\ell+1}{\ell} v(S_{j_k}^{(k)}), \quad (2)$$

where the last inequality follows from the fact that  $f(z) = \frac{z+1}{z}$  is decreasing.

Finally, to get the approximation guarantee for this case, we combine (2) with the fact that  $S$  is at least as good as each greedy solution:

$$v(S) \geq v(S_{j_k}) \geq v(S_{j_k}^{(k)}) \geq \frac{\ell}{\ell+1} \cdot \frac{x}{\beta}.$$

**Case 2.** Now assume that  $R_B = \emptyset$ , i.e.,  $R = R_c$ . Let  $C^*$  be an optimal solution for the given instance and define  $C_1 = C^* \cap S_1$ ,  $C_2 = C^* \cap S_2$  and  $C_3 = C^* \setminus (C_1 \cup C_2)$ . By subadditivity, we have

$$\text{OPT}(D, B) = v(C^*) \leq v(C_1) + v(C_2) + v(C_3). \quad (3)$$

Recall that  $T_j = \text{ALG}_2(S_j)$ ,  $j \in \{1, 2\}$ , is a 2-approximate solution with respect to  $\text{OPT}(S_j, \infty)$ . Thus,  $v(C_j) \leq \text{OPT}(S_j, B) \leq 2 \cdot v(T_j)$ , for  $j \in \{1, 2\}$ , and inequality (3) gives

$$\text{OPT}(D, B) \leq 2v(T_1) + 2v(T_2) + v(C_3). \quad (4)$$

Upper bounding  $v(C_3)$  in terms of  $S_1, S_2, T_1, T_2$  and  $r$  is somewhat more involved. We begin by invoking the non-negativity of  $v$ , as well as its submodularity (as defined in of Definition 2.1(ii)) on  $S_1 \cup C_3$  and  $S_2 \cup C_3$ . We have

$$v(C_3) \leq v(C_3) + v(S_1 \cup C_3 \cup S_2) \leq v(S_1 \cup C_3) + v(S_2 \cup C_3). \quad (5)$$

In order to upper bound  $v(S_1 \cup C_3)$  we again use the submodularity of  $v$ , together with a couple of facts about the marginal utilities of agents outside of  $S_1$ . Since the mechanism stopped after  $t$  iterations,  $\max_{i \in D \setminus (S_1 \cup S_2 \cup R)} v(i|S_1) \leq 0$ . Also, given that  $R = R_c$ , for all agents that got rejected at some point, we know that they had very low marginal value per cost ratio with respect to both  $S_1$  and  $S_2$ . In particular, if  $i_k \in R$ , then  $c_{i_k} > \frac{\beta B}{x} v(i_k|S_j^{(k)})$ , for both  $j \in \{1, 2\}$ . We may now rely on Definition 2.1(iii) to get

$$\begin{aligned} v(S_1 \cup C_3) &\leq v(S_1) + \sum_{i_k \in C_3} v(i_k|S_1) \\ &\leq v(S_1) + \sum_{i_k \in C_3 \cap R} v(i_k|S_1) \quad (v(i_k|S_1) \leq 0 \text{ for } i_k \in C_3 \setminus R) \\ &\leq v(S_1) + \sum_{i_k \in C_3 \cap R} v(i_k|S_1^{(k)}) \quad (\text{by submodularity, } v(i_k|S_1) \leq v(i_k|S_1^{(k)}) \text{ for } i_k \in D) \\ &\leq v(S_1) + \sum_{i_k \in C_3 \cap R} \frac{x}{\beta B} c_{i_k} \quad (\frac{\beta B}{x} v(i_k|S_1^{(k)}) < c_{i_k} \text{ for } i_k \in R) \end{aligned}$$

Similarly,  $v(S_2 \cup C_3) \leq v(S_2) + \sum_{i_k \in C_3 \cap R} \frac{x}{\beta B} c_{i_k}$ . Also, recall that  $\sum_{i \in C^*} c_i \leq B$  to get

$$v(S_j \cup C_3) \leq v(S_j) + \frac{x}{\beta}, \quad \text{for } j \in \{1, 2\}. \quad (6)$$

Finally, we may combine (4), (5) and (6) to get

$$\begin{aligned} \text{OPT}(D, B) &\leq 2v(T_1) + 2v(T_2) + v(S_1 \cup C_3) + v(S_2 \cup C_3) \\ &\leq 2v(T_1) + 2v(T_2) + v(S_1) + v(S_2) + \frac{2x}{\beta} \leq 6 \cdot v(S) + \frac{2x}{\beta}, \end{aligned}$$

or, equivalently,  $v(S) \geq \frac{1}{6} (\text{OPT}(D, B) - \frac{2x}{\beta})$ .

Combining Case 1 and Case 2, we obtain the claimed guarantee.  $\square$

So far, unless  $x = \Theta(\text{OPT}(D, B))$ , the approximation guarantee seems to be rather weak. In fact, the way **SIMULTANEOUS GREEDY** is used within **SAMPLE-THEN-GREEDY** requires that both  $x = v(\text{ALG}_1(A_1))$  and  $\text{OPT}(A_2, B)$  are  $\Theta(\text{OPT}(A, B))$ . The next technical lemma guarantees that this happens with high probability, unless there is an extremely valuable agent.

**LEMMA 3.6** (Follows from Bei et al. [10] and Leonardi et al. [34]). *Consider any submodular function  $v(\cdot)$ . For any given subset  $T \subseteq A$  and a positive integer  $k$  assume that  $v(T) \geq k \cdot \max_{i \in T} v(i)$ . Further, suppose that  $T$  is divided uniformly at random into two subsets  $T_1$  and  $T_2$ . Then with probability at least  $\frac{1}{2}$ , we have that  $v(T_1) \geq \frac{k-1}{4k} v(T)$  and  $v(T_2) \geq \frac{k-1}{4k} v(T)$ .*

We are now ready to lower bound the approximation guarantee of SAMPLE-THEN-GREEDY.

LEMMA 3.7. *Assume that for some positive integer  $k$ ,  $\text{OPT}(A, B) > k \cdot \max_{i \in A} v(i)$ . Then with probability at least  $\frac{1}{2}$ , SAMPLE-THEN-GREEDY( $A, v, c, B$ ) outputs a set  $S$  such that*

$$v(S) \geq \min \left\{ \frac{\lfloor \frac{k-1}{4e\beta} \rfloor (k-1)}{e \left( \lfloor \frac{k-1}{4e\beta} \rfloor + 1 \right)}, \frac{\beta(k-1) - 8k}{6} \right\} \cdot \frac{1}{4\beta k} \cdot \text{OPT}(A, B).$$

PROOF. Let  $C^*$  be an optimal solution for the given instance. By applying Lemma 3.6 with  $T = C^*$  we have that with probability at least  $1/2$  it holds that  $v(A_i \cap C^*) \geq \frac{k-1}{4k} v(C^*)$  for both  $i \in \{1, 2\}$ . In what follows we assume that this is indeed the case. Thus,

$$\text{OPT}(A, B) \geq x = v(\text{ALG}_1(A_1)) \geq \frac{1}{e} \text{OPT}(A_1, B) \geq \frac{k-1}{4ek} \text{OPT}(A, B)$$

and also  $\text{OPT}(A_2, B) \geq \frac{k-1}{4k} \text{OPT}(A, B)$ .<sup>2</sup>

The lower bound on  $x$  paired with the upper bound on  $\max_{i \in A} v(i)$ , imply that

$$\max_{i \in A} v(i) < \frac{1}{k} \cdot \text{OPT}(A, B) \leq \frac{1}{k} \cdot \frac{4ek\beta}{k-1} \cdot \frac{x}{\beta} \leq \frac{1}{\lfloor \frac{k-1}{4e\beta} \rfloor} \cdot \frac{x}{\beta}.$$

Thus, we can use Lemma 3.5 with  $D = A_2$ ,  $x = v(\text{ALG}_1(A_1))$  and  $\ell = \lfloor \frac{k-1}{4e\beta} \rfloor$ . Therefore, SIMULTANEOUS GREEDY( $A_2, v, c_{A_2}, B, x$ ) outputs an  $S$  such that

$$\begin{aligned} v(S) &\geq \min \left\{ \frac{\lfloor \frac{k-1}{4e\beta} \rfloor x}{\left( \lfloor \frac{k-1}{4e\beta} \rfloor + 1 \right) \beta}, \frac{1}{6} \left( \text{OPT}(A_2, B) - \frac{2x}{\beta} \right) \right\} \\ &\geq \min \left\{ \frac{\lfloor \frac{k-1}{4e\beta} \rfloor (k-1)}{4ek \left( \lfloor \frac{k-1}{4e\beta} \rfloor + 1 \right) \beta} \text{OPT}(A, B), \frac{1}{6} \left( \frac{k-1}{4k} \text{OPT}(A, B) - \frac{2}{\beta} \text{OPT}(A, B) \right) \right\} \\ &\geq \min \left\{ \frac{\lfloor \frac{k-1}{4e\beta} \rfloor (k-1)}{e \left( \lfloor \frac{k-1}{4e\beta} \rfloor + 1 \right)}, \frac{\beta(k-1) - 8k}{6} \right\} \cdot \frac{1}{4\beta k} \cdot \text{OPT}(A, B). \quad \square \end{aligned}$$

COROLLARY 3.8. *The set  $S$  returned by GENSM-MAIN( $A, v, c, B$ ) satisfies*

$$505 \cdot \mathbb{E}(v(S)) \geq \text{OPT}(A, B).$$

PROOF. Suppose that  $\max_{i \in A} v(i) \geq \frac{1}{101} \cdot \text{OPT}(A, B)$ . Then, with probability  $p$  at least  $1/101$  of the optimal value is returned. Hence,

$$\mathbb{E}(v(S)) \geq p \cdot \max_{i \in A} v(i) \geq \frac{0.201}{101} \cdot \text{OPT}(A, B) \geq \frac{1}{505} \cdot \text{OPT}(A, B).$$

Next suppose that  $\max_{i \in A} v(i) < \frac{1}{101} \cdot \text{OPT}(A, B)$ . We may apply Lemma 3.7 with  $k = 101$ . As discussed before the description of mechanism SIMULTANEOUS GREEDY, the parameter  $\beta$  is equal to 9.185. This implies that  $\lfloor \frac{k-1}{4e\beta} \rfloor = 1$ . By substituting the values of  $k$  and  $\beta$  to the bound of Lemma 3.7, we get that with probability at least  $(1-p)/2$

$$v(S) \geq \min \left\{ \frac{50}{e}, \frac{110.5}{6} \right\} \cdot \frac{1}{3710.74} \cdot \text{OPT}(A, B) \geq \frac{1}{201.7367} \cdot \text{OPT}(A, B),$$

<sup>2</sup>For the sake of presentation, here we write  $x = v(\text{ALG}_1(A_1)) \geq \frac{1}{e} \text{OPT}(A_1, B)$  rather than the technically correct  $\mathbb{E}(x) \geq \frac{1}{e} \text{OPT}(A_1, B)$ . However, as discussed in Remark 3.9, one can formally deal with this issue with a negligible effect on the expected approximation guarantee while keeping the running time polynomial.

and thus,

$$\mathbb{E}(v(S)) \geq \frac{1 - 0.201}{2} \cdot \frac{1}{201.7367} \cdot \text{OPT}(A, B) \geq \frac{1}{505} \cdot \text{OPT}(A, B). \quad \square$$

*Remark 3.9.* In our mechanisms we often use randomized approximation algorithms as sub-routines. In particular,  $\text{ALG}_1$  in  $\text{SAMPLE-THEN-GREEDY}$  and  $\text{GENSM-ONLINE}$  and  $\text{ALG}_3$  in  $\text{MONSM-CONSTRAINED}$  are randomized. Yet, in the description of our mechanisms—and more importantly in our analyses—we treat them as if there were deterministic, e.g., we assume that  $x$  in line 2 of  $\text{SAMPLE-THEN-GREEDY}$  is at least  $\frac{1}{e} \cdot \text{OPT}(A_1, v, c_{A_1}, B)$ . While this is not technically accurate, we do so for the sake of presentation. However, we may instead use the fact that for any constants  $\delta, \eta > 0$ ,  $\text{ALG}_1$  can be modified so that *with probability at least  $1 - \delta$  it returns a solution of value at least  $(\frac{1}{e} - \eta) \cdot \text{OPT}(A_1, v, c_{A_1}, B)$  in polynomial time*, using standard arguments. Thus, for any constant  $\varepsilon > 0$  the analysis of  $\text{GENSM-MAIN}$  can be adjusted to hold for an approximation factor  $505 + \varepsilon$  instead. Similarly, the analyses of our other mechanisms can be adjusted accordingly.

#### 4 ONLINE PROCUREMENT

Note that the mechanism presented in the last section already bares some resemblance to online algorithms for the *secretary model* (although truthfulness is rarely a requirement there). Namely, a part of the input is only used to estimate the quality of the optimal solution and then, based on that estimation, some threshold is set for the remaining instance. On a high level, this is straightforward to adjust for the secretary model; we use the first (roughly) half of the stream of agents to find an estimate of  $\text{OPT}(A, B)$  and then set a threshold similar to the one in  $\text{SIMULTANEOUS GREEDY}$ . However, there are a few issues one has to deal with.

First,  $\text{SIMULTANEOUS GREEDY}$  goes through the agents in a specific order (in decreasing order of the maximum marginal value with respect to either one of the two constructed sets). Even though this fact is indeed used in the proof of Lemma 3.5, we show that even examining agents in arbitrary order works well, albeit with a somewhat worse approximation factor. Note that this is not true when there are other constraints on top of the budget-feasibility requirement, as in Section 5.

Second, towards the end, in line 9,  $\text{SIMULTANEOUS GREEDY}$  runs an unconstrained submodular maximization algorithm on  $S_1$  and  $S_2$  to possibly reveal a subset of them with much higher value. While this is a critical step, we rely on a very elegant result of Feige et al. [21]: a uniformly random set gives a 4-approximation for the unconstrained problem. Thus, every agent that passes the threshold and is added to  $S_j$  is only accepted to  $T_j$  with probability  $1/2$ . The actual output of the mechanism is a random choice  $S$  between  $S_1, S_2, T_1$  and  $T_2$ , made *before* the arrival of the first agent. So, while the four sets are built obliviously with respect to the choice of  $S$ , the agents added to  $S$  are irrevocably chosen while everyone else is irrevocably discarded.

One last issue is that we want the mechanism to occasionally return the single most valuable agent. This, however, is easily resolved by running *Dynkin's algorithm* [19] with constant probability instead. This mechanism samples the first  $n/e$  agents and then it picks the first agent  $i'$ , among the remaining agents, who is at least as good as the best agent in the sample, i.e.,  $v(i') \geq \max_{k \leq n/e} v(i_k)$ . This guarantees that  $\mathbb{E}(v(i')) \geq \frac{1}{e} \max_{i \in A} v(i)$ , where the expectation is over the order of the agents.

The mechanism  $\text{GENSM-ONLINE}$  below incorporates all these adjustments, yet maintains all the good properties of  $\text{GENSM-MAIN}$ . We assume a secretary setting, where the agents arrive *uniformly at random*. In particular, agents have no control over their arrival time, so this is still a single-parameter environment and truthfulness still means *universal truthfulness*, i.e., if we fix the random bits of the mechanism, then *for any arrival order* no agent has an incentive to lie.

Again,  $\text{ALG}_1$  is the  $e$ -approximation algorithm of Kulik et al. [32]. The parameter  $\beta$  is set to 8.725 and, like the parameter in  $\text{SIMULTANEOUS GREEDY}$ , is only relevant for the approximation factor.

**THEOREM 4.1.** *GENSM-ONLINE is a universally truthful, individually rational, budget-feasible online mechanism and achieves an  $O(1)$ -approximation in the secretary model.*

One immediate consequence of Theorem 4.1 is the existence of an  $O(1)$ -approximation algorithm for the non-monotone *Submodular Knapsack Secretary Problem* (SKS). To the best of our knowledge GENSM-ONLINE is the first such algorithm (see also Remark 4.3).

Formally, an instance of SKS consists of a ground set  $A = [n]$ , a non-negative submodular objective  $v : A \rightarrow \mathbb{R}_+$  and a given budget  $B$ . The elements of  $A$  arrive in a uniformly random order and each element must be accepted or rejected immediately upon arrival. An algorithm for SKS has access to  $n = |A|$ , to the costs of items that have arrived (i.e., each cost is revealed upon arrival) and to a value oracle that, given a subset  $S \subseteq A$  of elements that have already arrived, returns  $v(S)$ . The objective is to accept a set of elements maximizing  $v$  without exceeding the budget.

It is straightforward to see that the only difference of SKS with the online procurement problem studied in this section is the information about the costs. In SKS there is no notion of misreporting a cost and thus it can be seen as a special case of our online problem where agents are guaranteed to always reveal their true costs.

**COROLLARY 4.2.** *There is an  $O(1)$ -approximation algorithm for the non-monotone SKS.*

*Remark 4.3.* Bateni et al.[9] give an  $O(1)$ -approximation algorithm for the *monotone* SKS. They claim, without a proof, that their result extends to the non-monotone SKS as well, using the same ideas that work for the Submodular Secretary problem with a cardinality or a matroid constraint (rather than a knapsack constraint). However, this does not seem to be the case. It is indeed true that these ideas do pair well with cardinality and, more generally, (intersection of) matroid constraints and they have been recently used to obtain fast randomized algorithms with good approximation guarantees [14, 22]. Unfortunately, to the best of our knowledge, there are no examples in the literature where this *random greedy* approach does work for non-monotone submodular maximization with knapsack constraints.

---

GENSM-ONLINE( $A, v, c, B$ )

---

```

1  With probability  $q = 0.4$  :
2  | Run Dynkin's algorithm and return the winner
3  With probability  $1 - q$  :
4  |  $S_1 = S_2 = T_1 = T_2 = \emptyset$ ;  $B_1 = B_2 = B$ 
5  |  $S = \begin{cases} S_j, & \text{with probability } 1/10, \text{ for each } j \in \{1, 2\} \\ T_j, & \text{with probability } 2/5, \text{ for each } j \in \{1, 2\} \end{cases}$ 
6  | Draw  $\xi$  from the binomial distribution  $\mathcal{B}(n, 0.5)$ 
7  | Let  $A_1$  be the set of the first  $\xi$  agents, and  $A_2 = A \setminus A_1$ 
8  | Reject all the agents in  $A_1$  and calculate  $x = v(\text{ALG}_1(A_1))$ 
9  | for each  $i \in A_2$  as he arrives do
10 | | Let  $\hat{j} \in \arg \max_{j \in \{1, 2\}} v(i|S_j)$ 
11 | | if  $c_i \leq \frac{\beta B}{x} v(i|S_{\hat{j}}) \leq B_{\hat{j}}$  then
12 | | |  $S_{\hat{j}} = S_{\hat{j}} \cup \{i\}$ 
13 | | |  $B_{\hat{j}} = B_{\hat{j}} - \frac{\beta B}{x} v(i|S_{\hat{j}})$ 
14 | | | With probability  $1/2$ ,  $T_{\hat{j}} = T_{\hat{j}} \cup \{i\}$ ; otherwise,  $T_{\hat{j}} = T_{\hat{j}}$ 
15 | | | Update  $S$  /* the update is consistent to the choice made in line 5 */
16 | return  $S$ 
```

---

## 5 ADDING COMBINATORIAL CONSTRAINTS

To illustrate the applicability of our approach, we turn to the case where the solution has to satisfy some additional combinatorial constraint. With the exception of additive valuation functions [1, 35], even for *monotone* submodular objectives no polynomial-time mechanisms using only value queries are known. Here we show that the general approach of GENSM-MAIN can be utilized to achieve an  $O(p)$ -approximation for  $p$ -systems, i.e., for independence systems with *rank quotient* at most  $p$ . In particular, as stated in Corollary 5.4, this implies constant factor approximation for cardinality, matroid and matching constraints. As it is shown in Section 6, going beyond independence systems (e.g., require that the solution forms a spanning tree) is hindered by strong impossibility results.

**Definition 5.1.** An *independence system* is a pair  $(U, \mathcal{I})$ , where  $U$  is a finite set and  $\mathcal{I} \subseteq 2^U$  is a family of subsets, whose members are called the *independent sets* of  $U$  and satisfy: (i)  $\emptyset \in \mathcal{I}$ , and (ii) if  $B \in \mathcal{I}$  and  $A \subseteq B$ , then  $A \in \mathcal{I}$ .

Given a set  $S \subseteq U$ , a maximal independent set contained in  $S$  is called a *basis* of  $S$ . The *upper rank*  $ur(S)$  (resp. the *lower rank*  $lr(S)$ ) is defined as the cardinality of a largest (resp. smallest) basis of  $S$ . A  $p$ -system  $(U, \mathcal{I})$  is an independence system such that  $\max_{S \subseteq U} ur(S)/lr(S) \leq p$ .

For the sake of readability, we present the case of *monotone* submodular objectives here; the non-monotone case is deferred to the full version. A technical highlight of our analysis, later used for the non-monotone case as well, is Claim 5.3. The claim crucially depends on the order we consider the agents, in order to bound the value lost because of the  $p$ -system constraint.

As usual, we assume the existence of an independence oracle. In particular, when we write that  $\mathcal{I}$  is part of the input of the mechanism, we mean that the mechanism has access to a membership oracle for  $\mathcal{I}$ . The parameter  $\beta$  is later set to  $13/3$ .  $ALG_3$  in line 5 can be any approximation algorithm for monotone submodular maximization subject to a knapsack and a  $p$ -system constraint. Here we assume the  $(p+3)$ -approximation algorithm of Badanidiyuru and Vondrák [7].

---

MONSM-CONSTRAINED( $A, \mathcal{I}, v, c, B$ )

---

```

1  With probability  $q = 0.2$  :
2  | return  $i^* \in \arg \max_{i \in A} v(i)$ 
3  With probability  $1 - q$  :
4  | Put each agent of  $A$  in either  $A_1$  or  $A_2$  independently at random with probability  $\frac{1}{2}$ 
5  |  $x = v(ALG_3(A_1))$  /* a  $(p+3)$ -approximation of  $OPT(A_1, v, c_{A_1}, B)$  */
6  |  $S = \emptyset; B_R = B; U = A_2$ 
7  | while  $U \neq \emptyset$  do
8  | | Let  $i \in \arg \max_{i \in U} v(i|S)$ 
9  | | if  $c_i \leq \frac{\beta B}{x} v(i|S) \leq B_R$  and  $S \cup \{i\} \in \mathcal{I}$  then
10 | | |  $S = S \cup \{i\}$ 
11 | | |  $B_R = B_R - \frac{\beta B}{x} v(i|S)$ 
12 | |  $U = U \setminus \{i\}$ 
13 | return  $S$ 

```

---

**THEOREM 5.2.** Assuming that the solution has to be an independent set of a  $p$ -system, there is a universally truthful, individually rational, budget-feasible,  $O(p)$ -approximation mechanism that runs in polynomial time for (non-monotone) submodular objectives.

**PROOF.** Here we prove the theorem for the non-monotone case. First, we observe that  $S$  starts as an independent set, namely the empty set, and it is expanded only if it remains an independent set. Hence, at the end MONSM-CONSTRAINED does return a feasible solution, i.e.,  $S$  is in  $\mathcal{I}$ .

At this point, following the same reasoning used for GENSM-MAIN and GENSM-ONLINE, it should be easy to see that MONSM-CONSTRAINED is universally truthful, individually rational, budget-feasible, and runs in polynomial time.

Next we show that the solution returned by the mechanism is an  $O(p)$ -approximation of the optimum. First, suppose that  $\max_{i \in A} v(i) \geq \text{OPT}(A, B)/(26(p+10))$ . Then, for the set  $S$  returned by MONSM-CONSTRAINED,  $\mathbb{E}(v(S)) \geq q \cdot \max_{i \in A} v(i) \geq \frac{1}{5} \cdot \text{OPT}(A, B)/(26(p+10)) \geq \text{OPT}(A, B)/(138(p+10))$ .

For the case where  $\max_{i \in A} v(i) < \text{OPT}(A, B)/(26(p+10))$ , we follow the same notation and the same high level approach as with the approximation guarantees of GENSM-MAIN and GENSM-ONLINE. So,  $i_1, i_2, \dots, i_{|A_2|}$  are the agents of  $A_2$  in the order considered by the mechanism. By  $S^{(k)}$  and  $B_R^{(k)}$  we denote  $S$  and  $B_R$ , respectively, at the time  $i_k$  arrives, and we only use  $S$  for the final set returned. The set  $R = A_2 \setminus S$  contains the agents  $i_k$  that were not added to  $S^{(k)}$  and it is further partitioned to  $R_c = \{i_k \mid \frac{\beta B}{x} v(i_k | S^{(k)}) < c_i\}$ ,  $R_B = \{i_k \mid B_R^{(k)} < \frac{\beta B}{x} v(i_k | S^{(k)})\}$ , and  $R_I = R \setminus (R_c \cup R_B)$ .

Assume that  $\text{OPT}(A_i, B) \geq \frac{k-1}{4k} \text{OPT}(A, B)$  for  $i \in \{1, 2\}$ , where  $k = 26(p+10)$ . Thus,  $x = v(\text{ALG}_1(A_1)) \geq \frac{k-1}{4(p+3)k} \text{OPT}(A, B)$ . Recall that this does happen with probability at least  $\frac{1}{2}$ , as discussed in the beginning of the proof of Lemma 3.7.

**Case 1.** Assume that  $R_B \neq \emptyset$ . Let  $i_k \in R_B$ , i.e.,  $\frac{\beta B}{x} v(i_k | S^{(k)}) > B_R^{(k)}$ . Using the same argument as in the proof of Lemma 3.5, we get  $v(S^{(k)}) \geq \frac{x}{\beta} - \max_{i \in A} v(i)$  and, given the known bounds on  $x$  and  $\max_{i \in A} v(i)$ , this leads to  $v(S) \geq (\frac{k-1}{4(p+3)k\beta} - \frac{1}{k}) \cdot \text{OPT}(A, B)$ .

By substituting  $k = 26(p+10)$  and  $\beta = \frac{13}{3}$ , it is a matter of simple calculations to get

$$v(S) \geq \frac{5}{276(p+10)} \cdot \text{OPT}(A, B). \quad (7)$$

**Case 2.** Assume that  $R_B = \emptyset$  and let  $C^*$  be an optimal solution for the instance  $(A_2, v, c_{A_2}, B)$ . By monotonicity, we have

$$\text{OPT}(A_2, B) = v(C^*) \leq v(S \cup C^*). \quad (8)$$

Because of the  $p$ -system constraint, however, deriving the analog of inequality (6) needs some extra work. By Definition 2.1(iii), we have

$$v(S \cup C^*) \leq v(S) + \sum_{i_k \in C^* \cap S} v(i_k | S) \leq v(S) + \sum_{i_k \in C^* \cap R_c} v(i_k | S) + \sum_{i_k \in C^* \cap R_I} v(i_k | S). \quad (9)$$

We may upper bound the first sum using the fact that all agents involved got rejected because they had very low marginal value per cost ratio. That is,

$$\sum_{i_k \in C^* \cap R_c} v(i_k | S) \leq \sum_{i_k \in C^* \cap R_c} v(i_k | S^{(k)}) < \frac{x}{\beta B} \sum_{i_k \in C^* \cap R_c} c_{i_k} \leq \frac{x}{\beta} \leq \frac{\text{OPT}(A, B)}{\beta}. \quad (10)$$

For the second sum we prove the following result that crucially relies on the fact that agents are examined in decreasing marginal value.

**CLAIM 5.3.**  $\sum_{i_k \in C^* \cap R_I} v(i_k | S) \leq p \cdot v(S)$ .

**PROOF OF CLAIM 5.3.** Recall that when we index agents we follow the ordering imposed by the mechanism, i.e.,  $i_k$  is always the agent picked at the  $k$ th execution of line 8 of MONSM-CONSTRAINED.

Suppose that there is a mapping  $f : C^* \cap R_I \rightarrow S$  such that

- (i) if  $f(i_k) = i_\ell$ , then  $v(i_k | S^{(k)}) \leq v(i_\ell | S^{(\ell)})$  for all  $i_k \in C^* \cap R_I$ , and
- (ii)  $|f^{-1}(i_\ell)| \leq p$  for all  $i_\ell \in S$ .

We slightly abuse the notation and write  $S^{f(i_k)}$  instead of  $S^{(\ell)}$  when  $f(i_k) = i_\ell$ . The existence of  $f$  implies that

$$\sum_{i_k \in C^* \cap R_I} v(i_k | S) \leq \sum_{i_k \in C^* \cap R_I} v(i_k | S^{(k)}) \leq \sum_{i_k \in C^* \cap R_I} v(f(i_k) | S^{f(i_k)}) \leq p \cdot \sum_{i_\ell \in S} v(i_\ell | S^{(\ell)}) = p \cdot v(S).$$

The first inequality follows from the submodularity of  $v$ , while the second and third inequalities follow from (i) and (ii), respectively.

Next, we are going to construct such an  $f$ . Let  $S = \{i_{a_1}, i_{a_2}, \dots, i_{a_s}\}$  and  $C^* \cap R_I = \{i_{b_1}, i_{b_2}, \dots, i_{b_t}\}$ , where both  $(a_i)_{i=1}^s$  and  $(b_i)_{i=1}^t$  are subsequences of  $1, 2, \dots, |A_2|$ . We are going to map the first  $p$  elements of  $C^* \cap R_I$ ,  $i_{b_1}, \dots, i_{b_p}$ , to  $i_{a_1}$ , the next  $p$  elements  $i_{b_{p+1}}, \dots, i_{b_{2p}}$ , to  $i_{a_2}$ , and so on. That is,  $f(i_{b_j}) = i_{a_{\lceil j/p \rceil}}$ .

It is straightforward that  $f$  satisfies property (ii). In order to prove property (i), it suffices to show that for all  $j \in \{1, 2, \dots, t\}$ , agent  $i_{b_j}$  is considered by MONSM-CONSTRAINED after agent  $f(i_{b_j})$ . Indeed, if that was the case, by the definition of  $\hat{I}$  in line 8 and submodularity, we would get

$$v(i_{a_{\lceil j/p \rceil}} | S^{(a_{\lceil j/p \rceil})}) \geq v(i_{b_j} | S^{(a_{\lceil j/p \rceil})}) \geq v(i_{b_j} | S^{(b_j)}),$$

for all  $i_{b_j} \in C^* \cap R_I$ , as desired. Suppose, towards a contradiction, that there is some  $k \in \{1, 2, \dots, t\}$ , such that  $b_k < a_{\lceil k/p \rceil}$ ; in fact, suppose  $k$  is the smallest such index. Consider the sets  $T = \{i_{a_1}, i_{a_2}, \dots, i_{a_{\lceil k/p \rceil - 1}}\} \subseteq S$  and  $Q = \{i_{b_1}, i_{b_2}, \dots, i_{b_k}\} \subseteq C^* \cap R_I$ . By construction,  $T \in \mathcal{I}$ . Moreover, we claim that  $T$  is maximally independent in  $T \cup Q$ . Indeed, each  $i_{b_\tau} \in Q$  was rejected because  $S^{(b_\tau)} \cup \{i_{b_\tau}\} \notin \mathcal{I}$ , and since  $S^{(b_\tau)} \subseteq T$  we get  $T \cup \{i_{b_\tau}\} \notin \mathcal{I}$ . This implies that  $\text{lr}(T \cup Q) \leq |T|$ . On the other hand,  $Q \in \mathcal{I}$  because  $Q \subseteq C^* \in \mathcal{I}$ . As a result  $\text{ur}(T \cup Q) \geq |Q|$ . However, notice that

$$p \cdot |T| = p(\lceil k/p \rceil - 1) < p(k/p + 1 - 1) = k = |Q|.$$

Thus,  $\frac{\text{ur}(T \cup Q)}{\text{lr}(T \cup Q)} \geq \frac{|Q|}{|T|} > p$ , contradicting the fact that  $(A, \mathcal{I})$  is a  $p$ -system. We conclude that  $f$  satisfies both (i) and (ii), and therefore,  $\sum_{i_k \in C^* \cap R_I} v(i_k | S) \leq p \cdot v(S)$ .  $\blacktriangleleft$

Now, combining (8), (9), (10), and Claim 5.3, we have  $\text{OPT}(A_2, B) \leq (p+1) \cdot v(S) + \frac{\text{OPT}(A, B)}{\beta}$ , and using the lower bound on  $\text{OPT}(A_2, B)$ ,  $v(S) \geq \frac{1}{p+1} \cdot \left(\frac{k-1}{4k} - \frac{1}{\beta}\right) \text{OPT}(A, B)$ . Again, by substituting  $k$  and  $\beta$ , it is a matter of calculations to get

$$v(S) \geq \frac{5}{276(p+10)} \cdot \text{OPT}(A, B). \quad (11)$$

By Lemma 3.6, both (7) and (11) hold with probability at least  $1/2$ . Hence,

$$\mathbb{E}(v(S)) \geq (1-q) \cdot \frac{1}{2} \cdot \frac{5}{276(p+10)} \cdot \text{OPT}(A, B) = \frac{1}{138(p+10)} \cdot \text{OPT}(A, B). \quad \square$$

For matroid constraints we have  $p = 1$  and for matching constraints  $p = 2$ . Since cardinality constraints are a special case of matroid constraint, we directly get the following.

**COROLLARY 5.4.** *For cardinality, matroid and matching constraints, there is a universally truthful, budget-feasible  $O(1)$ -approximation mechanism for (non-monotone) submodular objectives.*

## 6 LOWER BOUNDS

In the value query model there is a strong lower bound on the number of queries for deterministic algorithms for monotone XOS objectives due to Singer [40]. This result is based on a lower bound of Mirrokni et al. [36] on welfare maximization in combinatorial auctions. As the latter also holds for randomized algorithms, so does Singer's result as well, essentially with the same proof. We restate it here for completeness. Note that it holds even when the costs are public knowledge.



**THEOREM 6.1** (Singer [40]). *For any fixed  $\varepsilon > 0$ , any (randomized)  $n^{\frac{1}{2}-\varepsilon}$ -approximation algorithm for monotone XOS function maximization subject to a budget constraint requires exponentially many value queries (in expectation).*

When one moves to non-monotone objectives, as it is the case in this work, it is possible to prove even stronger lower bounds. Below we show that for general XOS objectives, exponentially many value queries are needed for *any* non-trivial approximation even without the budget constraint. As this result applies to the purely algorithmic setting, it is of independent interest.

It is known that in many settings there is a separation between the power of value and demand queries of polynomial size, see, e.g., [11]. To stress this difference in our setting, recall that in the demand query model, the class of XOS objectives admits a truthful  $O(1)$ -approximation mechanism with a polynomial number of queries.

**THEOREM 6.2.** *For any fixed  $\varepsilon > 0$ , any (randomized)  $n^{1-\varepsilon}$ -approximation algorithm for XOS function maximization requires exponentially many value queries (in expectation).*

One immediate consequence of Theorem 6.2 is that when we care for constant approximation ratios, the result of Theorem 3.1 is (asymptotically) the *best possible* for budget-feasible mechanism design. General submodular objectives is the broadest class of well studied non-monotone functions one could hope for, even for randomized mechanisms.

**Combinatorial Constraints.** We now turn to the problem of maximizing subject to additional constraints on top of the budget constraint. To further motivate our restriction to  $p$ -system constraints, we restate here a lower bound of Badanidiyuru and Vondrák [7]: for independence system constraints one cannot achieve an approximation factor better than  $\max_{S \subseteq U} \frac{\text{ur}(S)}{\text{lr}(S)}$  with a polynomial number of queries. Thus, the result of Theorem 5.2 is asymptotically optimal.

**THEOREM 6.3** (Badanidiyuru and Vondrák [7]). *For any fixed  $\varepsilon > 0$ , any (randomized)  $(p + \varepsilon)$ -approximation algorithm for additive function maximization subject to  $p$ -system constraints requires exponentially many independence oracle queries (in expectation).*

As we mentioned in the beginning of Section 5, we cannot really go beyond independence systems and have any non-trivial approximation guarantee in polynomial time. This is illustrated in Theorem 6.4 and Corollary 6.5 below. Theorem 6.4 generalizes Singer's [40] strong impossibility result for deterministically “hiring a team of agents” to *any* constraint that is not downward closed below. Note that it holds even for super-constant approximation ratios, even for the special case of additive objectives, irrespectively of any complexity assumptions.

**THEOREM 6.4.** *Let  $\mathcal{F} \subseteq 2^A$  be any collection of feasible sets that is not downward closed. Then there is no deterministic, truthful, individually rational, budget-feasible mechanism achieving a bounded approximation when restricted on  $\mathcal{F}$ , even for additive objectives.*

The next corollary of Theorem 6.3 states that under general combinatorial constraints it is not possible to achieve any non-trivial approximation with polynomially many queries. While it is not hard to prove it directly, given Theorem 6.3 it suffices to notice that such a lower bound holds even for general independence systems. Indeed, there are cases where  $\frac{\text{ur}(U)}{\text{lr}(U)}$  is  $\Theta(n)$  like the  $(n - 1)$ -systems of independent sets of star graphs.

**COROLLARY 6.5.** *For any fixed  $\varepsilon > 0$ , any (randomized)  $n^{1-\varepsilon}$ -approximation algorithm for additive function maximization subject to general feasibility constraints requires exponentially many queries (in expectation).*

## REFERENCES

- [1] Georgios Amanatidis, Georgios Birmpas, and Evangelos Markakis. 2016. Coverage, Matching, and Beyond: New Results on Budgeted Mechanism Design. In *Web and Internet Economics - 12th International Conference, WINE 2016, Montreal, Canada, December 11-14, 2016, Proceedings (Lecture Notes in Computer Science)*, Vol. 10123. Springer, 414–428.
- [2] Georgios Amanatidis, Georgios Birmpas, and Evangelos Markakis. 2017. On Budget-Feasible Mechanism Design for Symmetric Submodular Objectives. In *Web and Internet Economics - 13th International Conference, WINE 2017, Bangalore, India, December 17-20, 2017, Proceedings (Lecture Notes in Computer Science)*, Vol. 10660. Springer, 1–15.
- [3] Georgios Amanatidis, Pieter Kleer, and Guido Schäfer. 2019. Budget-Feasible Mechanism Design for Non-Monotone Submodular Objectives: Offline and Online. *CoRR* abs/1905.00848 (2019). arXiv:1905.00848 <http://arxiv.org/abs/1905.00848>
- [4] Nima Anari, Gagan Goel, and Afshin Nikzad. 2014. Mechanism Design for Crowdsourcing: An Optimal  $1-1/e$  Competitive Budget-Feasible Mechanism for Large Markets. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*. 266–275.
- [5] Moshe Babaioff, Nicole Immorlica, David Kempe, and Robert Kleinberg. 2007. A Knapsack Secretary Problem with Applications. In *APPROX-RANDOM (Lecture Notes in Computer Science)*, Vol. 4627. Springer, 16–28.
- [6] Ashwinkumar Badanidiyuru, Robert Kleinberg, and Yaron Singer. 2012. Learning on a budget: posted price mechanisms for online procurement. In *EC*. ACM, 128–145.
- [7] Ashwinkumar Badanidiyuru and Jan Vondrák. 2014. Fast algorithms for maximizing submodular functions. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*. SIAM, 1497–1514.
- [8] Eric Balkanski and Jason D. Hartline. 2016. Bayesian Budget Feasibility with Posted Pricing. In *Proceedings of the 25th International Conference on World Wide Web, WWW 2016, Montreal, Canada, April 11 - 15, 2016*. ACM, 189–203.
- [9] MohammadHossein Bateni, Mohammad Taghi Hajiaghayi, and Morteza Zadimoghaddam. 2013. Submodular secretary problem and extensions. *ACM Trans. Algorithms* 9, 4 (2013), 32:1–32:23.
- [10] Xiaohui Bei, Ning Chen, Nick Gravin, and Pinyan Lu. 2017. Worst-Case Mechanism Design via Bayesian Analysis. *SIAM J. Comput.* 46, 4 (2017), 1428–1448.
- [11] Liad Blumrosen and Noam Nisan. 2009. On the Computational Power of Demand Queries. *SIAM J. Comput.* 39, 4 (2009), 1372–1391.
- [12] Allan Borodin, Yuval Filmus, and Joel Oren. 2010. Threshold Models for Competitive Influence in Social Networks. In *Proceedings of the 6th International Workshop on Internet and Network Economics, WINE 2010*. 539–550.
- [13] Niv Buchbinder and Moran Feldman. 2018. Deterministic Algorithms for Submodular Maximization Problems. *ACM Trans. Algorithms* 14, 3 (2018), 32:1–32:20.
- [14] Niv Buchbinder, Moran Feldman, Joseph Naor, and Roy Schwartz. 2014. Submodular Maximization with Cardinality Constraints. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*. SIAM, 1433–1452.
- [15] Chandra Chekuri, Shalmoli Gupta, and Kent Quanrud. 2015. Streaming Algorithms for Submodular Function Maximization. In *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part I (Lecture Notes in Computer Science)*, Vol. 9134. Springer, 318–330.
- [16] Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. 2014. Submodular Function Maximization via the Multilinear Relaxation and Contention Resolution Schemes. *SIAM J. Comput.* 43, 6 (2014), 1831–1879.
- [17] Ning Chen, Nick Gravin, and Pinyan Lu. 2011. On the Approximability of Budget Feasible Mechanisms. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*. 685–699.
- [18] Shahar Dobzinski, Christos H. Papadimitriou, and Yaron Singer. 2011. Mechanisms for complement-free procurement. In *Proceedings 12th ACM Conference on Electronic Commerce (EC-2011), San Jose, CA, USA, June 5-9, 2011*. 273–282.
- [19] Evgenii Borisovich Dynkin. 1963. Optimal choice of the stopping moment of a Markov process. In *Doklady Akademii Nauk*, Vol. 150. Russian Academy of Sciences, 238–240.
- [20] Alina Ene and Huy L. Nguyen. 2017. A Nearly-linear Time Algorithm for Submodular Maximization with a Knapsack Constraint. *CoRR* abs/1709.09767 (2017). arXiv:1709.09767 <http://arxiv.org/abs/1709.09767>
- [21] Uriel Feige, Vahab S. Mirrokni, and Jan Vondrák. 2011. Maximizing Non-monotone Submodular Functions. *SIAM J. Comput.* 40, 4 (2011), 1133–1153.
- [22] Moran Feldman, Christopher Harshaw, and Amin Karbasi. 2017. Greed Is Good: Near-Optimal Submodular Maximization via Greedy Optimization. In *Proceedings of the 30th Conference on Learning Theory, COLT 2017, Amsterdam, The Netherlands, 7-10 July 2017 (Proceedings of Machine Learning Research)*, Vol. 65. PMLR, 758–784.
- [23] Moran Feldman, Joseph Naor, and Roy Schwartz. 2011. Improved Competitive Ratios for Submodular Secretary Problems (Extended Abstract). In *APPROX-RANDOM (Lecture Notes in Computer Science)*, Vol. 6845. Springer, 218–229.

- [24] Moran Feldman, Joseph Naor, and Roy Schwartz. 2011. A Unified Continuous Greedy Algorithm for Submodular Maximization. In *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*. IEEE Computer Society, 570–579.
- [25] Moran Feldman and Rico Zenklusen. 2018. The Submodular Secretary Problem Goes Linear. *SIAM J. Comput.* 47, 2 (2018), 330–366.
- [26] Gagan Goel, Afshin Nikzad, and Adish Singla. 2014. Mechanism Design for Crowdsourcing Markets with Heterogeneous Tasks. In *Proceedings of the Seconf AAAI Conference on Human Computation and Crowdsourcing, HCOMP 2014, November 2-4, 2014, Pittsburgh, Pennsylvania, USA*.
- [27] Anupam Gupta, Viswanath Nagarajan, and Sahil Singla. 2017. Adaptivity Gaps for Stochastic Probing: Submodular and XOS Functions. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*. SIAM, 1688–1702.
- [28] Anupam Gupta, Aaron Roth, Grant Schoenebeck, and Kunal Talwar. 2010. Constrained Non-monotone Submodular Maximization: Offline and Secretary Algorithms. In *Internet and Network Economics - 6th International Workshop, WINE 2010, Stanford, CA, USA, December 13-17, 2010. Proceedings (LNCS)*, Vol. 6484. Springer, 246–257.
- [29] Thibaut Horel, Stratis Ioannidis, and S. Muthukrishnan. 2014. Budget Feasible Mechanisms for Experimental Design. In *LATIN 2014: Theoretical Informatics - 11th Latin American Symposium, Montevideo, Uruguay, March 31 - April 4, 2014. Proceedings*. 719–730.
- [30] Thomas Kesselheim and Andreas Tönnis. 2017. Submodular Secretary Problems: Cardinality, Matching, and Linear Constraints. In *APPROX-RANDOM (LIPIcs)*, Vol. 81. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 16:1–16:22.
- [31] Pooya Jalaly Khalilabadi and Éva Tardos. 2018. Simple and Efficient Budget Feasible Mechanisms for Monotone Submodular Valuations. In *Web and Internet Economics - 14th International Conference, WINE 2018, Oxford, UK, December 15-17, 2018, Proceedings (Lecture Notes in Computer Science)*, Vol. 11316. Springer, 246–263.
- [32] Ariel Kulik, Hadas Shachnai, and Tami Tamir. 2013. Approximations for Monotone and Nonmonotone Submodular Maximization with Knapsack Constraints. *Math. Oper. Res.* 38, 4 (2013), 729–739.
- [33] Benny Lehmann, Daniel J. Lehmann, and Noam Nisan. 2006. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior* 55, 2 (2006), 270–296.
- [34] Stefano Leonardi, Gianpiero Monaco, Piotr Sankowski, and Qiang Zhang. 2016. Budget Feasible Mechanisms on Matroids. *CoRR* abs/1612.03150 (2016).
- [35] Stefano Leonardi, Gianpiero Monaco, Piotr Sankowski, and Qiang Zhang. 2017. Budget Feasible Mechanisms on Matroids. In *Integer Programming and Combinatorial Optimization - 19th International Conference, IPCO 2017, Waterloo, ON, Canada, June 26-28, 2017, Proceedings (Lecture Notes in Computer Science)*, Vol. 10328. Springer, 368–379.
- [36] Vahab S. Mirrokni, Michael Schapira, and Jan Vondrák. 2008. Tight information-theoretic lower bounds for welfare maximization in combinatorial auctions. In *Proceedings 9th ACM Conference on Electronic Commerce (EC-2008), Chicago, IL, USA, June 8-12, 2008*. ACM, 70–77.
- [37] Baharan Mirzasoleiman, Ashwinkumar Badanidiyuru, and Amin Karbasi. 2016. Fast Constrained Submodular Maximization: Personalized Data Summarization. In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016 (JMLR)*, Vol. 48. JMLR.org, 1358–1367.
- [38] Roger Myerson. 1981. Optimal Auction Design. *Mathematics of Operations Research* 6, 1 (1981).
- [39] George L. Nemhauser, Laurence A. Wolsey, and Marshall L. Fisher. 1978. An analysis of approximations for maximizing submodular set functions - I. *Math. Program.* 14, 1 (1978), 265–294.
- [40] Yaron Singer. 2010. Budget Feasible Mechanisms. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*. IEEE Computer Society, 765–774.
- [41] Yaron Singer. 2012. How to win friends and influence people, truthfully: influence maximization mechanisms for social networks. In *Proceedings of the Fifth International Conference on Web Search and Web Data Mining, WSDM 2012, Seattle, WA, USA, February 8-12, 2012*. 733–742.
- [42] Adish Singla and Andreas Krause. 2013. Incentives for Privacy Tradeoff in Community Sensing. In *Proceedings of the First AAAI Conference on Human Computation and Crowdsourcing, HCOMP 2013, November 7-9, 2013, Palm Springs, CA, USA*. AAAI.
- [43] Maxim Sviridenko. 2004. A note on maximizing a submodular set function subject to a knapsack constraint. *Oper. Res. Lett.* 32, 1 (2004), 41–43.
- [44] Laurence A. Wolsey. 1982. Maximising Real-Valued Submodular Functions: Primal and Dual Heuristics for Location Problems. *Math. Oper. Res.* 7, 3 (1982), 410–425.