

Budget Feasible Mechanisms

Yaron Singer
Computer Science Division
University of California at Berkeley
Berkeley, CA 94720
yaron@cs.berkeley.edu

Abstract—We study a novel class of mechanism design problems in which the outcomes are constrained by the payments. This basic class of mechanism design problems captures many common economic situations, and yet it has not been studied, to our knowledge, in the past. We focus on the case of procurement auctions in which sellers have private costs, and the auctioneer aims to maximize a utility function on subsets of items, under the constraint that the sum of the payments provided by the mechanism does not exceed a given budget. Standard mechanism design ideas such as the VCG mechanism and its variants are not applicable here. We show that, for general functions, the budget constraint can render mechanisms arbitrarily bad in terms of the utility of the buyer. However, our main result shows that for the important class of submodular functions, a bounded approximation ratio is achievable. Better approximation results are obtained for subclasses of the submodular functions. We explore the space of budget feasible mechanisms in other domains and give a characterization under more restricted conditions.

I. INTRODUCTION

Consider the following familiar problems:

- **Knapsack:** Given a budget B and a set of items $\mathcal{N} = \{1, \dots, n\}$, each with cost c_i and value v_i , find a subset of items S which maximizes $\sum_{i \in S} v_i$ under the budget constraint.
- **Matching:** Given a budget B and a bipartite graph, with set of edges $\mathcal{N} = \{e_1, \dots, e_n\}$ each with cost c_e and value v_e , find a legal matching S which maximizes $\sum_{e \in S} v_e$ under the budget constraint.
- **Coverage:** Given a budget B and subsets $\mathcal{N} = \{T_1, \dots, T_n\}$ of some ground set, each with a cost c_i find a subset S which maximizes $|\cup_{i \in S} T_i|$ under the budget constraint.

Three much studied, and much solved, optimization problems. However, suppose that the elements of \mathcal{N} are not combinatorial objects, but *strategic agents* with private costs. The above problems then capture natural economic interactions: Knapsack, for example, models a simple procurement auction, while Coverage may model the problem of maximizing exposure of an advertising campaign under a budget. These are precisely the kinds of economic interactions we wish to study here: reverse

auctions with private costs, with the goal of optimizing the auctioneer’s value.

At first glance it may seem that the problem we describe falls within the scope of familiar domains. However, closer inspection reveals a new dimension of difficulty: the budget constraint applies not to the costs but to the payments the mechanism uses to support truthfulness. We need mechanisms whose sum of payments never exceeds the given budget.

Can we design mechanisms that implement these intended economic interactions in the most favorable way to the auctioneer without their payments exceeding the budget?

Mechanism design is by now a very mature discipline and the recent injection of computational thought has helped develop it even further, and in new and forward-looking directions [24]. Procurement auctions, introduced to computer scientists already in [24], were at first studied under *utilitarian* objectives, seeking to optimize social welfare [24], [12]. More recently procurement auctions have been studied under the non-utilitarian framework of *frugality* [8], [3], [15], [11], [26] — essentially, payment optimization in reverse auctions.

These situations still fall within classical mechanism design theory, where *the set of possible outcomes is a priori fixed and publicly known*. By “set of all possible outcomes” here we mean the set of all possible allocations, with payments projected out. In other words, there is a rich class of allocations, independent of payments, each of which is realizable by a truthful mechanism. In the three introductory examples, however, the set of possible outcomes is *not* fixed or publicly known: It depends crucially on the participants’ private information, ultimately on the mechanism’s payments. It is this peculiarity that makes these three problems difficult, and places them at a blind spot of mechanism design.

Budget Feasible Mechanisms

We say a truthful mechanism is *budget feasible* if its payments do not exceed a given budget. In single

parameter domains, where each agent’s private information is a single number, designing truthful mechanisms often reduces to designing monotone allocation rules, since payments can be computed via binary search [21]. This no longer holds when the payments are restricted by a budget: Designing a budget feasible allocation rule requires understanding its payments, which in-turn depend on the allocation rule itself. Not surprisingly, it seems that budget feasible mechanisms are very tricky to find.

The VCG mechanism does not work: Consider a simple Knapsack instance where all items have identical values, and except for one item whose cost equals the budget, all items have small costs. The VCG mechanism will choose the $n - 1$ small-cost agents, paying the budget to each. Thus, while this mechanism returns the optimal solution with total *cost* within the budget, the total *payment* will be way over budget (in fact, $(n - 1)$ times the budget).

In general, nothing can work: Consider a slight variation of the above problem, in which all items have small costs, and identical values *as long as a particular item i is in the solution*, and otherwise all have value 0 (for example, think of i as a corkscrew and the rest of the items as bottles of wine). How well can a budget feasible mechanism do here? If the mechanism has a bounded approximation ratio it must always guarantee to include i in its solution. This however implies that as long as i declares a cost that is less than the mechanism’s budget, the mechanism includes her in the solution. A truthful mechanism must therefore surrender its entire budget to i . This of course results in an unbounded approximation ratio.

Our Results

The question, then, is: Which classes of functions have budget feasible mechanisms with good approximation properties? *Our main result is a randomized constant factor budget feasible mechanism that is universally truthful for the quite general, and important, class of nondecreasing submodular functions* (Theorem 4.8).

For a slightly broader class, that of fractionally subadditive functions, we show that computational constraints dictate a lower bound. As shown in the simple example above, superadditive functions bring out the clash between truthfulness and the budget constraint. On a positive note, the three problems in the beginning of the section correspond to subclasses (additive, OXS, and coverage) of submodular maximization problems. We show improved approximations for these problems and other special cases. We further explore the space of budget feasible mechanisms, showing several impossibilities as well as a characterization under more restricted conditions.

Related Work

Budgets in auctions: Budgets came under scrutiny in auction theory [9], [7], [13], [6] after observing behavior of bidders in online automated auctions [2], as well as in spectrum auctions where bidding is performed by groups of strategic experts [7]. While these pioneering works highlight the significance and challenges that budgets introduce to mechanism design, they relate to an entirely different concept than the one we study here. While these works study the impact of budgets on strategic *bidders*, our interest is to explore the budget’s effect on the *mechanism*. These papers, however, do point out the complexity induced by budget constraints in mechanism design, and the need for approximations.

Frugality: In recent years a theory of *frugality* has been developed with the goal of providing mechanisms for procurement auctions that admit minimal payments [8], [3], [15], [11], [26]. Budget feasibility and frugality are complementary concepts. Frugality is about buying a feasible solution at minimum cost — there are no preferences among the solutions, and the goal is to minimize payments. In our setting we have no preferences among payments — as long as they are below the budget — but we do care about the value of the solutions. The two approaches are complementary also in another important sense: in our last section we show that for all the problems studied in the frugality literature there are no budget feasible mechanisms.

Cost Sharing: Somewhat conceptually closer to our work is the subject of cost sharing, in which agents have private values for a service, there is a nondecreasing cost for allocating the service to agents, and the goal is to maximize the agents’ valuations under the cost (see [14] for a survey). The proportional share mechanism we study in this paper is inspired by [20] and [25]. The relationship between cost sharing and our setting, however, is quite limited: we are not aiming to optimize a function of the agents private information under a public cost function, but rather optimize a public function under constraints dictated by agents’ private information and a fixed budget. Here, our goal is non-utilitarian — we aim to maximize the *buyer’s* demand, which is independent of the agents’ utilities.

Submodular Maximization: From a pure algorithmic perspective, even under a cardinality constraint, maximizing a submodular function is well known to be NP-hard, and an $1 - 1/e$ approximation ratio can be achieved by greedily taking items based on their marginal contribution [23]. When items have costs, variations of greedy on marginal contribution normalized by cost can achieve constant factor approximations, and even the optimal $1 - 1/e$ ratio [16], [17]. For submodular maximization problems that can be expressed as integer programs, rounding solutions of linear and nonlinear pro-

grams can, in some cases, achieve the optimal constant approximation ratio [1].

Paper Organization

After the necessary definitions in Section II, we present a mechanism for the class of symmetric submodular functions (Section III); this special case simplifies the problem enormously and facilitates the introduction of ideas and intuition for the general submodular case. Our main result for submodular functions is developed in Section IV. Finally, in Section V, we further discuss the space of budget feasibility, improved approximations, impossibility results and characterization.

II. THE MODEL

In a *budget-limited reverse auction* we have a set of items $[n] = \{1, \dots, n\}$, and a single buyer. Each item $i \in [n]$ is associated with a cost $c_i \in \mathcal{R}_+$, while the buyer has a budget $B \in \mathcal{R}_+$ and a demand valuation function $V : 2^{[n]} \rightarrow \mathcal{R}_+$. In the *full information* case, costs are common knowledge, and the objective is to maximize the demand function under the budget, i.e. find the subset $S \in \{T \mid \sum_{i \in T} c_i \leq B\}$ for which $V(S)$ is maximized.

We focus on the *strategic case*, in which each item is held by a unique agent and costs are *private*. The budget and demand function of the buyer are common knowledge. A solution is a subset and a payment vector, and the objective is to maximize the demand function while the *payments* (not costs) are within the budget. More formally, a mechanism $\mathcal{M} = (f, p)$ consists of an allocation function $f : \mathcal{R}_+^n \rightarrow 2^{[n]}$ and a payment function $p : \mathcal{R}_+^n \rightarrow \mathcal{R}_+^n$. The allocation function f maps a set of n bids to a subset $S = f(c_1, \dots, c_n) \subseteq [n]$. The payment function p returns a vector p_1, \dots, p_n of payments to the agents. We shall often omit the arguments c_1, \dots, c_n when writing f and p . We shall denote by s_1, \dots, s_n the characteristic vector of S , that is, $s_i = 1$ iff $i \in S$. As usual, we seek normalized ($s_i = 0$ implies $p_i = 0$), individually rational ($p_i \geq s_i \cdot c_i$) mechanisms with no positive transfers ($p_i \geq 0$). As it is common in algorithmic mechanism design, our goal is manifold. We seek mechanisms that are:

- 1) **Truthful**, that is, reporting the true costs is a dominant strategy for sellers. Formally, a mechanism $\mathcal{M} = (f, p)$ is *truthful (incentive compatible)* if for every $i \in [n]$ with cost c_i and bid c'_i , and every set of bids by $[n] \setminus \{i\}$ we have $p_i - s_i \cdot c_i \geq p'_i - s'_i \cdot c_i$, where (s_i, p_i) and (s'_i, p'_i) are the allocations and payments when the bidding is c_i and c'_i , respectively. A mechanism that is a randomization over truthful mechanisms is *universally truthful*.
- 2) **Computationally Efficient**. The functions f and p can be computed in polynomial time. In cases

where the demand function requires exponential data to be represented (as in the general submodular case), we take the common “black-box” approach and assume the buyer has access to an oracle which allows evaluating any subset $S \subseteq [n]$, with polynomially many queries. Such queries are known as value queries. This is a weaker model than ones allowing demand or general queries (see [5] for a definition). Since our main interest here is algorithmic, this strengthens our results.

- 3) **Budget Feasible**. Importantly, we require that a mechanism’s allocation rule and payments do not exceed the budget: $\sum_i p_i s_i \leq B$. We call such mechanisms *budget feasible*.
- 4) **Approximation**. We want the allocated subset to yield the highest possible value for the buyer. For $\alpha \geq 1$ we say that a mechanism is α -approximate if the mechanism allocates to a set S such that $OPT(c, \mathcal{N}, B) \leq \alpha V(S)$, where $OPT(c, \mathcal{N}, B)$ denotes the value of full information optimal solution over a set of agents \mathcal{N} with cost vector c and budget B . As usual, when dealing with randomization we seek mechanisms that yield constant factor approximations in expectation.

This is a *single parameter* mechanism design problem, in that each bidder has only one private value. We shall repeatedly rely on Myerson’s well-known characterization ¹:

Theorem 2.1 ([22]): In single parameter domains a normalized mechanism $\mathcal{M} = (f, p)$ is truthful iff:

- (i) **f is monotone**: $\forall i \in [n]$, if $c'_i \leq c_i$ then $i \in f(c_i, c_{-i})$ implies $i \in f(c'_i, c_{-i})$ for every c_{-i} ;
- (ii) **winners are paid threshold payments**: payment to each winning bidder is $\inf \{c_i : i \notin f(c_i, c_{-i})\}$.

III. SYMMETRIC SUBMODULAR FUNCTIONS

We now introduce a subclass of submodular functions which is devoid of many of the intricacies of the general case. It will serve as an exposition of the basic ideas, and will help understand the difficulties in the general case.

We say a set function is *symmetric* if it only depends on the cardinality of the set, rather than the identity of the items it holds. Symmetric submodular functions (also called *downward sloping*) were used by Vickrey in his seminal work on multi-unit auctions [27]. They have a very simple structure:

¹Note that although there is a budget constraint on the payments, Myerson’s characterization applies to our setting as well. Due to the characterization, we know that the allocation rule determines the payment function. The budget constraint can therefore be viewed as a property of the allocation rule alone.

Definition 3.1: A function $V : 2^{[n]} \rightarrow \mathcal{R}_+$ is symmetric submodular if there exist $r_1 \geq \dots \geq r_n \geq 0$, such that $V(S) = \sum_{i=1}^{|S|} r_i$.

Consider the following allocation rule $f_{\mathcal{M}}$: Sort the n bids so that $c_1 \leq c_2 \leq \dots \leq c_n$, and consider the largest k such that $c_k \leq B/k$. That is, k is the place where the curve of the increasing costs intersects the hyperbola B/k . The set allocated here is $\{1, 2, \dots, k\}$. That is, $f_{\mathcal{M}} = \{1, 2, \dots, k\}$. This is obviously a monotone allocation rule: an agent cannot be excluded when decreasing her bid. In the Appendix we show that paying each agent $\theta_i = \min\{B/k, c_{k+1}\}$ results in a truthful mechanism.²

Observe that this allocation rule has the property we seek: summing over the payments that support truthfulness satisfies the budget constraint. Hence this gives us a budget feasible mechanism. Importantly, this is also a good approximation of the optimum solution:

Theorem 3.2: The above mechanism has approximation ratio of two.

Proof: Observe that the optimal solution is obtained by greedily choosing the lowest-priced items until the budget is exhausted. By the downward sloping property, to prove the result it suffices to show that the mechanism returns at least half of the items in the greedy solution. Assume for purpose of contradiction that the optimum solution has ℓ items, and the mechanism returns less than $\ell/2$. It follows that $c_{\lceil \ell/2 \rceil} > 2B/\ell$. Note however, that this is impossible since we assume that $c_{\lceil \ell/2 \rceil} \leq \dots \leq c_\ell$, and $\sum_{i=\lceil \ell/2 \rceil}^{\ell} c_i \leq B$ which implies that $c_{\lceil \ell/2 \rceil} \leq 2B/\ell$, a contradiction. ■

In Section V we show that no better approximation ratio is possible. This is rather surprising, given the simplicity of the full-information problem, and illustrates the intricacies of budget feasibility.

IV. GENERAL SUBMODULAR FUNCTIONS

We now turn to the general case of nondecreasing submodular functions. A demand function V is nondecreasing if $S \subseteq T$ implies $V(S) \leq V(T)$.

Definition 4.1: $V : 2^{[n]} \rightarrow \mathcal{R}_+$ is submodular if $V(S \cup \{i\}) - V(S) \geq V(T \cup \{i\}) - V(T) \quad \forall S \subseteq T$.

In general, submodular functions may require exponential data to be represented. We therefore assume the buyer has access to a *value oracle* which given a query $S \subseteq [n]$ returns $V(S)$ (see related work section for more

²It is rather interesting that the second term is needed; we show in the Appendix that the mechanism breaks down in its absence.

discussion on submodular maximization). In designing truthful mechanisms for submodular maximization problems, the greedy approach is a natural fit, since it is monotone when agents are sorted according to their increasing marginal contributions relative to cost: the marginal contribution of an agent i given a subset S is $V_{i|S} := V(S \cup \{i\}) - V(S)$. In the marginal contribution-per-cost sorting the $i+1$ agent is the agent j for which $V_{j|S_i}/c_j$ is maximized over all agents \mathcal{N} where $S_i = \{1, 2, \dots, i\}$, and $S_0 = \emptyset$. To simplify notation we will write V_i instead of $V_{i|S_{i-1}}$. This sorting, in the presence of submodularity, implies:

$$V_1/c_1 \geq V_2/c_2 \geq \dots \geq V_n/c_n. \quad (1)$$

Notice that $V(S_k) = \sum_{i \leq k} V_i$ for all k .

A. The Proportional Share Allocation Rule

The mechanism from the previous section for the limited symmetric case can be generalized appropriately to work for various classes in the submodular family of functions.

Definition 4.2: For a budget B and set of agents \mathcal{N} with cost vector c , the generalized proportional share allocation rule, denoted $f_{\mathcal{M}}(c, B, \mathcal{N})$ sorts agents according to (1) with costs vector c and budget B and allocates to agents $\{1, \dots, k\}$ that respect $c_i \leq B \cdot V_i/V(S_i)$. Observe that this condition is met for every $\{1, \dots, i\}$ when $i \leq k$.

For concreteness consider the case of additive valuations (Knapsack from the Introduction): each agent is associated with a fixed value v_i and $V(S) = \sum_{i \in S} v_i$. Here the marginal contribution of each agent is independent of their place in the sorting, and we simply have that $V_i = v_i$ for all agents $i \in [n]$. In this case $f_{\mathcal{M}}$ produces a budget-feasible mechanism. The reason is, it assures us that for each agent i , the threshold payments of $f_{\mathcal{M}}$, denoted θ_i do not exceed the agent's proportional share:

$$\theta'_i = \min \left\{ \frac{V_i \cdot B}{\sum_{i \in S} V_i}, \frac{V_i \cdot c_{k+1}}{V_{k+1}} \right\}.$$

which allows budget feasibility, as well as individually rationality: $\theta'_i \leq c_i$. This seems to make the proportional share allocation rule an ideal candidate to obtain budget feasible mechanisms. Indeed, with some minor adjustments, for many problems with functions in the submodular class (e.g. symmetric, Knapsack, Matching,) this general approach works well and produces budget feasible mechanisms with good approximation guarantees (see the following section for more details). Furthermore, as we discussed above, the proportional share mechanism is optimal in some cases, and in some restricted environments our characterizations show that

this is essentially the *only* budget feasible mechanism (see the following section). It seems however that this natural approach completely fails as soon as we encounter more involved cases, such as Coverage.

B. The Difficulties

The Coverage problem captures many of the difficulties that are associated with designing budget feasible mechanisms for the general submodular case. In Coverage the marginal contribution of an agent is not fixed, but depends on the subset allocated by the algorithm in the previous stages. An agent’s marginal contribution therefore depends on its position in the sorting, which introduces several difficulties.

Marginal Contributions are Affected by Costs:

When applying the proportional share mechanism in Coverage, paying agents θ_i^j as above will be under the budget, as we desire. However, observe that for each agent i the payment depends on the marginal contribution V_i , which is determined by i ’s position in the sorting. Thus, in such a case the payments will depend on the agent’s declared cost, and therefore cannot induce truthfulness, making the proportional share mechanism hopeless here.

Simple allocation and payment schemes that are independent of the agent’s position in the sorting also fail. An approach that may seem natural is to replace marginal contributions with Shapley values [14] since they make the proportional contribution of an allocated agent independent of her position in the sorting. Unfortunately, such an approach cannot approximate better than a factor of \sqrt{n} , as we show in the full version of this paper. While it is tempting to get rid of the marginal contribution sorting, it is the only known means for obtaining good approximation guarantees for the general submodular case.

Non-monotonicity of the Maximum Operator:

Bounded approximation ratios for submodular maximization under a knapsack constraint depend crucially on the ability to take the maximum between a greedy solution and the item with highest value. In the general case, as well in the case of Coverage, taking this maximum does not preserve monotonicity: simple examples show that for allocation rules that depend on marginal contribution sorting, by decreasing her cost an agent can *decrease* the value of the allocation.

C. Overview of Our Approach

Our approach is based in three ideas:

- First, we derive an alternative characterization of the threshold payments of the proportional share allocation rule. Since we know that this rule does not work, this may seem futile. We’ll show that this characterization plays a significant role in our design.

- Using the above characterization, we show that for any (nondecreasing) submodular function, the threshold payments of a slightly modified version of the proportional share allocation rule are “not too far” from the agents’ proportional contributions. This enables us to guarantee that when running the modified version of the proportional share allocation rule *with a constant fraction of the budget*, the threshold payments will be budget feasible.
- Finally, to obtain the approximation guarantee we partition the agents in a manner that allows us to include the variation of the proportional share rule over a *subset of agents* and obtain good approximation guarantees.

D. Characterizing Threshold Payments

The following definition is key in our characterization.

Definition 4.3: The marginal contribution of agent i at point j is $V_{i(j)} := V(T_{j-1} \cup \{i\}) - V(T_{j-1})$ where T_j denotes the subset of the first j agents in the marginal-contribution-per-cost sorting (as in (1)) over the subset $\mathcal{N} \setminus \{i\}$.

The intuition behind the payments characterization can be described as follows. Consider running the proportional share mechanism without agent i . For the first j agents in the marginal contribution sorting, using the marginal contribution of i at point j we can find the maximal cost that agent i can declare in order to be allocated instead of the agent in the j th place in the sorting. While these costs may have arbitrary behavior as a function of j , we will show that taking the maximum of these values guarantees payments that support truthfulness. To avoid confusion we use T_j to denote the first j agents according to this sorting, V_j' to denote the marginal contribution of the j th agent in this case, and k' to denote the index of the last agent $j \in \mathcal{N} \setminus \{i\}$ that respects $c_j \leq V_j' \cdot B/V(T_j)$. For brevity we will write $c_{i(j)} := V_{i(j)} \cdot c_j/V_j'$ and $\rho_{i(j)} := V_{i(j)} \cdot B/V(T_{j-1} \cup \{i\})$.

Lemma 4.4 (Payments Characterization): The threshold payment for $f_{\mathcal{M}}$ is

$$\theta_i = \max_{j \in [k'+1]} \left\{ \min\{c_{i(j)}, \rho_{i(j)}\} \right\}.$$

Proof: To characterize the threshold payment for agent i , relabel the agents according to the marginal-contribution-per-cost sorting over the subset $\mathcal{N} \setminus \{i\}$. Consider $f_{\mathcal{M}}$ as a sequential allocation rule: at each stage j , the mechanism resorts the remaining agents that have not yet been allocated according to marginal contribution-per-cost sorting, and allocates to the first

agent in the sorting if she meets the condition $c_j \leq V'_j \cdot B/V(T_{j-1} \cup \{j\})$.

For a given stage j in this sequential allocation, we can find the maximal cost agent i would have been able declare and be allocated, if she had been considered by the mechanism at this stage: The value $c_{i(j)} = V_{i(j)} \cdot c_j/V'_j$ is the maximal cost i can declare which would place her ahead of j in the sorting, and if this cost does not exceed $\rho_{i(j)} = B \cdot V'_{i(j)}/V(T_{j-1} \cup \{i\})$, the mechanism would have allocated to agent i . Therefore, had i appeared at stage j , the minimum between these values is the maximal cost she can declare and be allocated at this stage. Since $V_{i(j)}$ monotonically decreases with j while c_j/V'_j increases, $c_{i(j)}$ may have arbitrary behavior as a function of j . However, as we now show, taking the maximum of these values result in threshold payments.

Let r be the index in $[k' + 1]$ for which $\min\{c_{i(j)}, \rho_{i(j)}\}$ is maximal. Declaring a cost below $\theta_i \leq c_{i(r)}$ guarantees i to be within the first $r \leq k' + 1$ elements in the sorting stage of the mechanism, with $r-1$ items allocated. Since $\theta_i \leq \rho_{i(r)}$, i will be allocated.

To see that declaring a higher cost prevents i from being allocated, consider first the case where $c_{i(r)} \leq \rho_{i(r)}$. A higher cost places i after r in the sorting stage of the mechanism. If the maximum of $c_{i(j)}$ over all $j \in [k' + 1]$ is $c_{i(r)}$, reporting a higher cost places i after an element which is not allocated and therefore it will not be allocated. Otherwise, if $c_{i(r)} < c_{i(j)}$, for some $j \leq k' + 1$, by the maximality of r it must be the case that:

$$\frac{B \cdot V_{i(j)}}{V(T_{j-1} \cup \{i\})} = \rho_{i(j)} < c_{i(r)} < c_{i(j)}$$

and i will not be allocated as a cost above $\rho_{i(j)}$ will not meet the allocation condition.

In the second case when $c_{i(r)} > \rho_{i(r)}$, if r is the index which maximizes $\rho_{i(j)}$ over all indices in $[k' + 1]$, reporting a higher cost will not meet the mechanism's allocation condition at each index in $[k' + 1]$. Otherwise, if there is some other index $j \in [k' + 1]$ for which this maximum is achieved, then:

$$\frac{V_{i(j)} \cdot c_j}{V'_j} = c_{i(j)} < \rho_{i(r)} < \rho_{i(j)}$$

and thus declaring a higher cost in this case places i after j in the sorting, and the mechanism will not consider i . ■

Lemma 4.5 (Individual Rationality): The mechanism $f_{\mathcal{M}}$ is individually rational, i.e., $c_i \leq \theta_i$.

Proof: Observe that:

- (a) $V_{i(j)} \geq V_{i(j+1)} \quad \forall j \in \mathcal{N}$;
- (b) $T_j = S_j \quad \forall j < i$;

(c) $V_{i|T_{i-1}} = V_i$.

Since the threshold payment is the maximum over all $\min\{c_{i(j)}, \rho_{i(j)}\}$ in $[k' + 1]$, it is enough to show that $c_i \leq \min\{c_{i(j)}, \rho_{i(j)}\}$ for a certain $j \leq k' + 1$. Since (b) implies that $i \leq k' + 1$ we can consider i 's replacement j which appears in the i th place in the marginal-contribution-per-cost sorting over $\mathcal{N} \setminus \{i\}$. Since $i \in [k]$, and due to (b) and (c) above, we have that

$$c_i \leq \frac{V_i \cdot B}{V((S_{i-1} \cup \{i\}))} = \frac{V_{i|T_{i-1}} \cdot B}{V(T_{i-1} \cup \{i\})} = c_{i(j)}.$$

In the original sorting, i appears ahead of j (as implied from (b)), and therefore its relative marginal contribution is greater. Thus:

$$c_i \leq \frac{V_{i|S_{i-1}} \cdot c_j}{V_{j|S_{i-1}}} = \frac{V_{i|T_{i-1}} \cdot c_j}{V'_{j|T_{j-1}}} = \rho_{i(j)}.$$

It therefore follows that $c_i \leq \min\{c_{i(j)}, \rho_{i(j)}\} \leq \theta_i$. ■

E. Payment Bounds

The characterization above allows us to include a slightly modified version of the proportional share allocation rule in our mechanism, with threshold payments that are guaranteed to be no more than a constant factor away from agents' proportional contribution. This is a key property which guides the design of our mechanism.

We will run the modified proportional share allocation rule over a subset of the agents \mathcal{N}_s , with a constant fraction of the budget, denoted B' . We describe \mathcal{N}_s and B' explicitly in the following section, but for the purpose of showing the payment bounds, we can think of these as any subset of agents and any budget. In this modified version of the proportional share allocation rule, for $i_s^* := \operatorname{argmax}_{j \in \mathcal{N}_s} V(\{j\})$, we sort the agents of $\mathcal{N}_s \setminus \{i_s^*\}$ according to the marginal-contribution-per-cost order, and allocate to $W = S_k \cup \{i_s^*\}$, where S_k are all k agents in \mathcal{N}_s that respect the condition $c_i \leq V_i \cdot B/V(S_i \cup \{i_s^*\})$. The characterization from above easily extends to this case using $\rho_{i(j)} = V_{i(j)} \cdot B/V(T_{j-1} \cup \{i, i_s^*\})$. Under this modification we can show the following desirable bound on the threshold payments:

Lemma 4.6 (Payment Bounds): For $i \in W \setminus \{i_s^*\}$:

$$\theta_i \leq \left(\frac{6e-2}{e-1}\right) \frac{V_i \cdot B'}{V(W)}.$$

Proof: For $T_{k'}$ as above, let $W' = T_{k'} \cup \{i_s^*\}$ and let r be the index for which $\theta_i = \min\{c_{i(r)}, \rho_{i(r)}\}$. If $r \leq k'$, observe that the sorting implies $c_r/V'_r \leq c_{k'}/V'_{k'}$ and therefore:

$$\theta_i \leq \frac{V_{i(r)} \cdot c_r}{V'_r} \leq \frac{V_{i(r)} \cdot c_{k'}}{V'_{k'}} \leq \frac{V_{i(r)} \cdot B'}{V(W')} \leq \frac{V_i \cdot B'}{V(W')}$$

where the last inequality relies on the observation that $V_{i(j)} \leq V_i$ for every $j \in \mathcal{N}_s \setminus \{i^*\}$, which is due to the fact that $T_j = S_j$ for $j \leq i$ and the decreasing marginal utility property of V . In the case where if $r = k' + 1$:

$$\theta_i \leq \rho_{i(r)} = \frac{V_{i(k'+1)} \cdot B'}{V(W' \cup \{i\})} \leq \frac{V_i \cdot B'}{V(W')}.$$

We therefore see that in both cases $\theta_i \leq V_i \cdot B'/V(W')$. To complete our proof, we will show that $V(W) \leq ((5e - 1)/(e - 1))V(W')$.

Consider the marginal contribution-per-cost sorting on $\mathcal{N}_s \setminus \{i, i^*\}$ and let $\ell' := \max\{t \mid \sum_{j \leq t} c_j \leq B'\}$. Due to the marginal contribution-per-cost sorting we have:

$$\frac{c_{k'+1}}{V'_{k'+1}} \sum_{j=k'+1}^{\ell'} V'_j \leq \sum_{j=k'+1}^{\ell'} \left(\frac{c_j}{V'_j} \right) V'_j \leq B'.$$

Since $c_{k'+1}/V'_{k'+1} > B/V(W' \cup \{k' + 1\})$, the above inequality implies:

$$V(T_{\ell'}) - V(T_{k'}) < V(W' \cup \{k' + 1\})$$

and since $T_{k'} \subset W'$, and $i^* \in W'$ this implies that $3V(W') \geq V(T_{\ell'})$. From submodularity, it can be shown that $V(T_{\ell'+1})$ is a $e/(e - 1)$ -approximation of the optimal solution (over $\mathcal{N}_s \setminus \{i, i^*\}$) [16], [17]. Since $i^* \in W'$ we have that:

$$\begin{aligned} V(W) &\leq OPT(c, B', \mathcal{N}_s) \\ &\leq OPT(B', \mathcal{N}_s \setminus \{i, i^*\}) + V(\{i\}) + V(\{i^*\}) \\ &< \left(\frac{4e}{e-1} + 2 \right) V(W') \end{aligned}$$

which concludes our proof. \blacksquare

F. Approximation Guarantee

To obtain the approximation guarantee we partition the set of agents based on their value and declared cost. Let $i^* := \operatorname{argmax}_{i \in \mathcal{N}} V(\{i\})$, $\mathcal{N}_s := \{i \neq i^* : c_i \leq B/2\}$, and $\mathcal{N}_\ell := \mathcal{N} \setminus (\mathcal{N}_s \cup \{i^*\})$. This partition, as we argue later, does not break monotonicity. Our mechanism will run the modified proportional share allocation rule over \mathcal{N}_s , with a constant fraction of the budget B/α .

The rationale behind this approach is the following. Due to their large cost, any feasible solution includes at most one agent in \mathcal{N}_ℓ , and since each such agent has value lower than that of i^* , we will be able to discard \mathcal{N}_ℓ without sacrificing too much from the quality of our solution. Since all the costs in \mathcal{N}_s are below $B/2$, including i^* in our solution W is always

feasible. Importantly, we can show this allocation rule is a constant factor approximation over \mathcal{N}_s :

$$\text{Lemma 4.7: } \left(\frac{(3+2\alpha)e-1}{e-1} \right) V(W) \geq OPT(c, B, \mathcal{N}_s).$$

Sketch: Similarly to what we have shown in lemma 4.6, one can show that

$$V(S_\ell) - V(S_k) < \alpha V(W \cup \{k + 1\})$$

where ℓ is the maximal index for which $\sum_{j \leq \ell} c_j \leq B$ when items are taken according to their marginal contribution relative to cost on $\mathcal{N}_s \setminus \{i^*\}$. Since W includes i^* , one can verify that $V(S_{\ell+1}) \leq (2+2\alpha)V(W)$. Similar to before, we know that $(1 - 1/e)OPT(c, B, \mathcal{N}_s \setminus \{i^*\}) \leq V(S_{\ell+1})$, and thus:

$$\begin{aligned} OPT(c, B, \mathcal{N}_s) &\leq OPT(c, B, \mathcal{N}_s \setminus \{i^*\}) + V(\{i^*\}) \\ &\leq \left(\frac{(2+2\alpha)e}{e-1} + 1 \right) V(W) \end{aligned}$$

which implies our desired bound. \blacksquare

G. Main Result

We can now prove our main theorem.

Theorem 4.8: For any submodular maximization problem there exists a constant factor approximation randomized mechanism in the value query model which is budget feasible and universally truthful. Furthermore, no budget feasible mechanism can do better than $2 - \epsilon$, for any fixed $\epsilon > 0$.

Our analysis shows that in expectation our mechanism guarantees an approximation ratio of $\left(\frac{58e^2 - 32e + 6}{(e-1)^2} \right) \approx 117.7$. It is possible that tighter analysis can show the mechanism does better. We further discuss this points in Section VI.

Proof: Consider the following mechanism:

A Budget Feasible Approximation Mechanism for Submodular Functions	
Initialize:	$i^* \leftarrow \operatorname{argmax}_{i \in \mathcal{N}} V(\{i\}),$ $\mathcal{N}_s \leftarrow \{i \neq i^* : c_i \leq B/2\},$ $i_s^* \leftarrow \operatorname{argmax}_{i \in \mathcal{N}_s} V(\{i\}),$ $W \leftarrow \{i_s^*\},$ $B' \leftarrow \left((e-1)/(12e-4) \right) B,$ $i \leftarrow \operatorname{argmax}_{j \in \mathcal{N}_s \setminus \{i_s^*\}} V(\{j\})/c_j$
While $c_i \leq V_i \cdot B'/V(W \cup \{i\})$	
Do:	$W \leftarrow W \cup \{i\}$ $i \leftarrow \operatorname{argmax}_{j \in \mathcal{N} \setminus \{i_s^*\}} V_j _{S_i} / c_j$
Output:	Choose u.a.r. from $\{(W, \hat{\theta}), (\{i^*\}, B)\}$

The payment $\hat{\theta}$ here is $\hat{\theta}_{i_s^*} = B/2$ and $\hat{\theta}_i = \min\{\theta_i, B/2\}$ for $i \neq i_s^*$, where θ_i are the payments as described for the modified proportional share allocation rule, using the budget $B' = B/\alpha$, for $\alpha = (12e-4)/(e-1)$. Observe that α is exactly twice the constant from the bound in Lemma 4.6.

In case i^* is allocated, B is clearly her threshold payment. If W is allocated, from the characterization lemma and the fact that \mathcal{N}_s consists only of agents with cost less than $B/2$, $\hat{\theta}$ as described above are clearly the threshold payments. One can verify the partition of agents is monotone and since payments are bounded by $B/2$ agents in \mathcal{N}_ℓ cannot benefit by misreporting their cost.

Since the modified proportional share allocation rule uses $B' = B/\alpha$ as its budget, from lemma 4.6, we can conclude that:

$$\sum_i \hat{\theta}_i \leq B/2 + \alpha \sum_{i \neq i_s^*} \frac{V_i \cdot B'}{V(W)} \leq B$$

and the mechanism is therefore truthful and budget feasible. Individual rationality and monotonicity were discussed above, and the lower bound as described in the following section applies here as well.

Finally, let $\beta = ((3 + 2\alpha)/(e - 1))/(e - 1)$. From Lemma 4.7 and the partition of agents we have:

$$\begin{aligned} OPT(c, B, \mathcal{N}) &\leq OPT(c, B, \mathcal{N} \setminus \mathcal{N}_\ell) + V(\{i^*\}) \\ &\leq OPT(c, B, \mathcal{N}_s) + 2V(\{i^*\}) \\ &\leq \beta V(W) + 2V(\{i^*\}) \\ &\leq (2 + \beta) \max\{V(W), V(\{i^*\})\} \end{aligned}$$

which gives us our desired approximation ratio. \blacksquare

V. THE SPACE OF BUDGET FEASIBLE MECHANISMS

In this section we address some natural questions that arise when exploring budget feasible mechanisms. We limit our discussion to deterministic mechanisms.

A. Lower Bounds and Improved Approximations

For special cases of submodular functions, specialized techniques based on the main result yield better approximation ratios:

Theorem 5.1: For Knapsack there is a budget feasible 5-approximation mechanism. For Matching there is a budget feasible $(\frac{5e-1}{e-1})$ -approximation mechanism. For either problem, no budget feasible mechanism can approximate within a factor better than $2 - \epsilon$, for any fixed $\epsilon > 0$.

We defer the details of upper bound proofs to the full version of the paper. We will now show a lower bound that is independent of computational assumptions and shows that no approximation ratio better than two is possible. For our lower bound we'll consider a very simple function, $V(S) = |S|$. Observe that all the functions in the submodular class we have discussed, including symmetric submodular include this demand valuation.

Proposition 5.2: For $V(S) = |S|$, no budget feasible mechanism can guarantee an approximation of $2 - \epsilon$, for any $\epsilon > 0$.

Proof: Suppose we have n items with costs $c_1 = c_2 = \dots = c_n = B/2 + \epsilon$, for some positive $\epsilon < B/2$. Assume, for purpose of contradiction that f is a budget feasible mechanism with approximation ratio better than 2. In particular, f has a finite approximation ratio and must therefore allocate to at least one agent in this case. W.l.o.g., assume f allocates to agent 1.

By monotonicity, agent 1 can reduce her cost to $c'_1 = \epsilon' < B/2 - \epsilon$ and remain allocated. For this cost vector, (c_1, c_{-1}) , Myerson's characterization implies that the threshold payment for agent 1 should be at least $B/2 + \epsilon$, by individual rationality and budget feasibility, f cannot allocate to any other agent. Observe however that the optimal full information solution in this case allocates to two agents which contradicts f 's approximation ratio guarantee. \blacksquare

B. Lower Bound on Fractionally Subadditive Functions

In light of the positive results for submodular functions, our natural desire would be to extend the main result to more general classes of problems. We now show there is little hope in that, at least in the value query model. Let us consider fractionally subadditive functions:

Definition 5.3: A function $V : 2^{[n]} \rightarrow \mathcal{R}$ is called fractionally subadditive if there exists a finite set of additive valuations $\{a_1, \dots, a_t\}$ s.t. $V(S) = \max_{i \in [t]} a_i(S)$.

It is known that every submodular function can be represented as a fractionally subadditive function, and that all fractionally subadditive functions are subadditive [18]. Using a simple reduction from [19] we show that for mechanisms which use value query oracles, obtaining reasonable approximations in the case of fractionally subadditive demands is hard, regardless of incentive considerations.

Theorem 5.4: In the case of fractionally subadditive demands, any algorithm which approximates within a

factor better than $n^{\frac{1}{2}-\epsilon}$, for any fixed $\epsilon > 0$, requires exponentially many value queries. This is true even in the setting where all costs are public knowledge.

Proof: Our proof relies on the following lemma shown in [19]. We will use v_S to denote the additive function that assigns the value 1 to an item $j \in S$ and 0 to all other items, and \bar{v} to denote the function that assigns a value $(1+\epsilon/2)/(n^{\frac{1}{2}})$ to every item $j \in [n]$.

Lemma 5.5 ([19]): There exists a set T , $|T| = \sqrt{n}$, for which distinguishing between the functions $V = \max\{v_S: |S| \leq (1+\epsilon/2)n^\epsilon, \bar{v}\}$, and $V' = \max\{V, v_T\}$ requires exponentially (in n^ϵ) many value queries.

Let $c_1 = c_2, \dots, c_n = 1$ and $B = \sqrt{n}$. For such T as in the lemma, we have that $\sum_{i \in T} c_i = B$, and $V'(T) = \sqrt{n}$ while $V(T) = (1+\epsilon/2)n^\epsilon$. Thus, an algorithm that approximates better than the desired approximation ratio must be able to distinguish between these two valuations, and requires exponentially many value queries. ■

C. Impossibility of Hiring a Team of Agents

Looking at problems previously studied in procurement, one could also ask whether “hiring a team of agents” problems [3], [26], [15], [8] have good budget feasible mechanisms. In these problems there is a set of feasible outcomes (e.g. all possible spanning trees or all paths from s to t) and the goal is to design a mechanism that yields a feasible solution. In the literature these problems have been studied with the goal of designing mechanisms which yield minimal payments (frugal), according to various benchmarks.

In our notation, such problems can be written as having a function $V(S) = 1$ if $S \in \mathcal{F}$, and 0 otherwise, where \mathcal{F} is the set of all feasible outcomes. Call such a problem *nontrivial* if all solutions in \mathcal{F} contain more than one element.

Theorem 5.6: There is no budget feasible mechanism with a bounded approximation ratio for any nontrivial “hiring a team of agents” problem.

Proof: Assume for purpose of contradiction, there exists a budget feasible mechanism f that guarantees a bounded approximation ratio. Let S be a feasible outcome, and consider the bid profile in which all agents in S declare positive cost $\epsilon < B/|S|$, and all other agents declare B . Since the problem is nontrivial, the minimal cost of a feasible outcome different than S (if it exists) exceeds the budget. Since f guarantees a bounded approximation ratio, it must allocate to S . For agent $i \in S$, consider the cost vector (c'_i, c_{-i}) with

$c'_i = B - \epsilon(|S| - 1)$. Observe that S remains a cost-feasible solution, and the only one in \mathcal{F} , and therefore $f(c'_i, c_{-i}) = S$, which implies the threshold payment of agent i is at least c'_i . Since this holds for all agents in S , budget feasibility is contradicted. ■

D. Characterizing Budget Feasible Mechanisms in Restricted Settings

The characterization of budget feasible mechanisms with good approximation ratios naturally depends on the environment in which the mechanisms are implemented, or more concretely, the function we aim to optimize. In characterizations, we often introduce additional restrictions on our mechanisms which we then use as guidelines in their design or for complete characterization. We now consider mechanisms that respect the two additional conditions of anonymity [4] and weak stability (similar to [10]). Informally, a mechanism is weakly stable if an agent doesn’t hurt the rest when reducing her cost, and anonymous if its allocation rule does not depend on the agents’ identities.

Definition 5.7: An allocation rule f satisfies anonymity if $i \in f(c_i, c_j, c_{-ij})$ implies $j \in f(c'_i, c'_j, c_{-ij})$ when $c'_i = c_j, c'_j = c_i$.

Definition 5.8: An allocation rule f satisfies weak stability if for every $i, j \in f(c_i, c_{-i}), c'_i \leq c_i$ implies $j \in f(c'_i, c_{-i})$.

Note that for the class of symmetric submodular functions, the proportional share mechanism respects these conditions. Furthermore, for symmetric submodular functions it seems quite reasonable that mechanisms with good approximation properties satisfy these conditions. The obvious question here is whether all budget feasible mechanisms for the symmetric submodular case take the form of the proportional share allocation rule.

Theorem 5.9: Let f be a budget feasible mechanism that is anonymous and weakly stable, and let $S = f(c)$ for some bid profile c . Then, for all $i \in S$ it must be that $c_i \leq B/|S|$.

Proof: Assume for purpose of contradiction that there is a bid profile $c = (c_1, \dots, c_n)$, s.t. $f(c) = S$ and there is some $i \in S$ for which $c_i > B/|S|$. Let c' be the bid profile in which all agents in $S \setminus \{i\}$ bid $c_{min} = \min_{j \in S} c_j$, and the rest bid as in c . Since f is weakly stable, we have that $S \subseteq f(c')$. Let c'' the bid profile where i bids c_{min} as well, and the rest of the agents bid as in c' . From monotonicity we have that i is allocated, and again $S \subseteq f(c'')$. We now claim that

under the profile c'' , the threshold price for each agent in S is at least $c_i > B/|S|$. To see this, observe that i 's threshold price must be at least c_i , since $i \in f(c')$. Since f is anonymous, and all agents in S declare the same price, the threshold price for each agent in S must also be at least c_i . Thus, payments to agents in S exceed the budget, contradicting budget feasibility. ■

VI. DISCUSSION

The space of budget feasible mechanisms seems quite broad and invites for further investigation. The richness of the submodular class implies there are many problems for which better approximation ratios are achievable. We have briefly discussed some specific examples which include Knapsack, Matching, Coverage and the symmetric submodular case, and believe there are many more interesting problems to study. Furthermore, we believe a better approximation ratio is achievable for the general case, and it may be possible to improve the analysis of the upper bound by showing tighter payment bounds.

While we have made a first step towards characterization by considering more restricted mechanisms, we believe a much more general characterization for budget feasible mechanisms awaits to be revealed. Finally, it would be interesting to further explore the lower bounds dictated by budget feasibility. Here we've shown several lower bound which are independent of computational assumptions, and it would be interesting to extend these techniques.

ACKNOWLEDGEMENTS

The author wishes to thank Christos Papadimitriou for endless discussions and help. To Dave Buchfuhrer, Iftah Gamzu, Arpita Ghosh, Mohammad Mahdian, George Pierrakos, Amin Saberi, Michael Schapira, Meromit Singer and Mukund Sundararajan the author wishes to express gratitude for meaningful discussions and valuable advice.

REFERENCES

- [1] Alexander A. Ageev and Maxim Sviridenko. Pipage rounding: A new method of constructing algorithms with proven performance guarantee. *J. Comb. Optim.*, 8(3):307–328, 2004.
- [2] Gagan Aggarwal, Nir Ailon, Florin Constantin, Eyal Even-Dar, Jon Feldman, Gereon Frahling, Monika Rauch Henzinger, S. Muthukrishnan, Noam Nisan, Martin Pál, Mark Sandler, and Anastasios Sidiropoulos. Theory research at google. *SIGACT News*, 39(2):10–28, 2008.
- [3] Aaron Archer and Éva Tardos. Frugal path mechanisms. *ACM Transactions on Algorithms*, 3(1), 2007.
- [4] Itai Ashlagi, Shahar Dobzinski, and Ron Lavi. An optimal lower bound for anonymous scheduling mechanisms. In *ACM Conference on Electronic Commerce*, pages 169–176, 2009.
- [5] Liad Blumrosen and Noam Nisan. On the computational power of iterative auctions. In *ACM Conference on Electronic Commerce*, pages 29–43, 2005.
- [6] Christian Borgs, Jennifer T. Chayes, Nicole Immorlica, Mohammad Mahdian, and Amin Saberi. Multi-unit auctions with budget-constrained bidders. In *ACM Conference on Electronic Commerce*, pages 44–51, 2005.
- [7] Jeremy Bulow, Jonathan Levin, and Paul Milgrom. Winning play in spectrum auctions. *Working Paper*.
- [8] Matthew Cary, Abraham D. Flaxman, Jason D. Hartline, and Anna R. Karlin. Auctions for structured procurement. In *SODA*, pages 304–313, 2008.
- [9] Shahar Dobzinski, Ron Lavi, and Noam Nisan. Multi-unit auctions with budget limits. In *FOCS*, pages 260–269, 2008.
- [10] Shahar Dobzinski and Mukund Sundararajan. On characterizations of truthful mechanisms for combinatorial auctions and scheduling. In *ACM Conference on Electronic Commerce*, pages 38–47, 2008.
- [11] Edith Elkind, Amit Sahai, and Kenneth Steiglitz. Frugality in path auctions. In *SODA*, pages 701–709, 2004.
- [12] Joan Feigenbaum, Christos H. Papadimitriou, Rahul Sami, and Scott Shenker. A bgp-based mechanism for lowest-cost routing. In *PODC*, pages 173–182, 2002.
- [13] Jon Feldman, S. Muthukrishnan, Martin Pál, and Clifford Stein. Budget optimization in search-based advertising auctions. In *ACM Conference on Electronic Commerce*, pages 40–49, 2007.
- [14] Kamal Jain and Mohammad Mahdian. Cost sharing. In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani, editors, *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [15] Anna R. Karlin, David Kempe, and Tami Tamir. Beyond VCG: Frugality of truthful mechanisms. In *FOCS*, pages 615–626, 2005.
- [16] Samir Khuller, Anna Moss, and Joseph (Seffi) Naor. The budgeted maximum coverage problem. *Inf. Process. Lett.*, 70(1):39–45, 1999.
- [17] Andreas Krause and Carlos Guestrin. A note on the budgeted maximization of submodular functions. In *CMU Technical Report*, pages CMU–CALD–05–103, 2005.
- [18] Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. In *ACM conference on electronic commerce*, 2001.
- [19] Vahab S. Mirrokni, Michael Schapira, and Jan Vondrák. Tight information-theoretic lower bounds for welfare maximization in combinatorial auctions. In *ACM Conference on Electronic Commerce*, pages 70–77, 2008.
- [20] Hervé Moulin and Scott Shenker. Strategyproof sharing of submodular costs: budget balance versus efficiency. *Economic Theory*, 18(3), pages 511–533, 2001.
- [21] Ahuva Mu'alem and Noam Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. *Games and Economic Behavior*, 64(2):612–631, 2008.
- [22] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1), 1981.
- [23] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis of approximations for maximizing submodular set functions ii. *Math. Programming Study* 8, pages 73–87, 1978.
- [24] Noam Nisan and Amir Ronen. Algorithmic mechanism design. *Games and Economic Behaviour*, 35:166 – 196, 2001. A preliminary version appeared in *STOC* 1999.
- [25] Tim Roughgarden and Mukund Sundararajan. New trade-offs in cost-sharing mechanisms. In *STOC*, pages 79–88, 2006.
- [26] Kunal Talwar. The price of truth: Frugality in truthful mechanisms. In *STACS*, pages 608–619, 2003.
- [27] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961.