

Budget Feasible Mechanism Design

YARON SINGER

Harvard University

In this letter we sketch a brief introduction to budget feasible mechanism design. This framework captures scenarios where the goal is to buy items or services from strategic agents under budget. The setting introduces interesting challenges that arise from the tension between incentive compatibility and the budget constraint, and leaves many interesting open questions. We will discuss several application domains that include crowdsourcing, information dissemination in social networks, and privacy-preserving recommendation systems.

Categories and Subject Descriptors: J.4 [Social and Behavioral Science]: Economics

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: mechanism design, approximation algorithms

1. INTRODUCTION

Consider the following model. There are n agents $\{a_1, \dots, a_n\}$, each with a *private* cost $c_i \in \mathbb{R}_+$ for selling a unique item or performing a service. There is a buyer with a budget $B \in \mathbb{R}_+$ and a combinatorial utility function $f : 2^{[n]} \rightarrow \mathbb{R}_+$ defined over subsets of agents.¹ We say that a mechanism is *budget feasible* if the *payments* it makes to agents do not exceed the budget. The goal is to design an *incentive compatible* budget feasible mechanism which yields the largest value possible to the buyer. This model was introduced in [Singer 2010] and gives us a simple framework to study procurement in strategic environments. We will start with a few simple examples that introduce the main ideas in the budget feasibility framework. In the sections that follow we will describe interesting generalizations and applications.

VCG is inapplicable. We will begin by investigating the celebrated Vickrey-Clarke-Groves (VCG) framework, adapted to our setting. One way to interpret the class of VCG mechanisms is that they allocate resources optimally assuming agents bid truthfully, and enforce prices that support truthful reporting. Now, consider an instance in which all agents have some small cost $\varepsilon > 0$, the budget is at least εn and the buyer's objective is simply to maximize the number of items purchased (i.e. $f(S) = |S|$). Our adaptation of VCG finds the optimal solution, which in this case selects all agents since their costs are small in comparison to the budget. When considering payments, each agent should be paid $B - \varepsilon(n - 1)$, and for a

¹We assume agents have quasi-linear utilities and are strategic and rational (i.e. aim to maximize the difference between the payment they receive and their true costs and are willing to lie about their costs if and only if it increases their utilities). The budget and utility function of the buyer are common knowledge, as is the mechanism being implemented. We first discuss *direct revelation* mechanisms, i.e. protocols where agents submit bids that represent their declared costs. A solution is a (possibly randomization over) subset of agents and payment vector.

Author's address: yarons@seas.harvard.edu.

sufficiently small ε the total payments exceed the budget by a factor proportional to the number of agents! Since we require mechanisms whose sum of payments does not exceed the budget this kind of mechanism is inapplicable.

Approximation is necessary. The above example implies something stronger than the fact that VCG is generally inapplicable. Due to uniqueness of prices, the example above suggests that the optimal solution – the allocation which maximizes the utility of the buyer when all agents' costs are known – cannot be implemented truthfully under budget. This is true even for utility functions that are as simple as the one above and *is independent of any computational assumptions*. We will therefore aim to *approximate* the optimal solution. We will quantify the quality of a mechanism in terms of its approximation ratio against the full-information optimum. For $\alpha \geq 1$ we say that a mechanism is α -approximate if for any input of bids it allocates to a set S s.t. $\alpha f(S) \geq f(S^*)$ where $S^* \in \operatorname{argmax}_{\{T: \sum_{i \in T} c_i \leq B\}} f(T)$.²

In general, nothing works. Consider an instance almost identical to the one from the previous example where again all agents have small costs $\varepsilon > 0$, the budget is sufficiently large, i.e. $B > \varepsilon n$, but now the utility function the buyer aims to optimize is the following variant of the cardinality function:

$$f(S) := \begin{cases} |S| & \text{if agent } a_1 \text{ is in } S \\ 0 & \text{otherwise} \end{cases}$$

A buyer with the above utility function only benefits when a particular agent a_1 is allocated (think of a buyer who wishes to consume as many bottles of wine as possible, and only a_1 sells her bottle with a complementary bottle opener). It is not hard to be convinced that any truthful mechanism which has a positive value to the buyer must surrender its entire budget to a_1 in this case. So, for general utility functions no truthful mechanism has a reasonable approximation ratio. Importantly, note that here too this is due to the conflict between truthfulness and the budget constraint and does not depend on any computational assumptions.

The question is then:

which utility functions have budget feasible mechanisms with reasonable approximation guarantees?

The first main result we will discuss shows that for any monotone submodular function there exists a randomized truthful budget feasible mechanism that has a constant factor approximation ratio. As we will soon see, this result is developed by a careful study of what we call *proportional share mechanisms*.

2. THE PROPORTIONAL SHARE MECHANISM

Consider the following mechanism: Sort the n bids so that $b_1 \leq b_2 \leq \dots \leq b_n$, and let k be the largest index for which $b_k \leq B/k$. That is, k is the place where

²We note that similar impossibilities to the ones presented in this letter apply to Bayesian mechanisms, implying that positive results for worst-case analysis would be particularly strong.

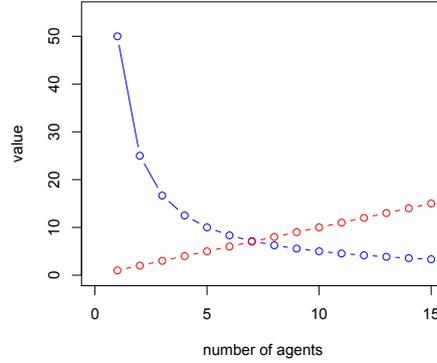


Fig. 1. An illustration of the allocation rule of the proportional share mechanism for $f(S) = |S|$. The blue line is the curve B/k which is the border of budget feasibility when sharing the budget equality between agents. The red line is the value of the function we aim to optimize $f(S) = |S|$.

the curve of the increasing costs intersects the hyperbola B/k (see Figure 1). The set allocated here is $S = \{1, 2, \dots, k\}$ and for every agent $a_i \in S$ the payment that supports truthful reporting is $\min\{B/k, b_{k+1}\}$. It is not difficult to verify that this mechanism is truthful and importantly that it has the property we seek: summing over the payments that support truthfulness satisfies the budget constraint, and we therefore have a budget feasible mechanism. Importantly, this mechanism has a good approximation ratio.

PROPOSITION 2.1. *For $f(S) = |S|$ the mechanism is a two-approximation.*

PROOF. Observe that the optimal solution is obtained by greedily choosing the lowest-priced items until the budget is exhausted. Assume for purpose of contradiction that the optimal solution has ℓ items, and the mechanism returns less than $\ell/2$ items. 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. \square

Somewhat surprisingly, this simple mechanism is in fact *optimal*.

PROPOSITION 2.2. *For $f(S) = |S|$, no budget feasible mechanism can guarantee an approximation ratio strictly better than two.*

PROOF. Suppose $c_1 = c_2 = \dots = c_n = B/2 + \delta$, for some positive $\delta < B/2$. Let \mathcal{M} be a mechanism with a finite approximation ratio and w.l.o.g. assume \mathcal{M} allocates to agent a_1 . By monotonicity, a_1 can reduce her bid to $b'_1 < B/2 - \delta$ and remain allocated. For this bid vector, (b'_1, c_{-1}) , Myerson's characterization [Myerson 1981] implies that the payment for agent a_1 should be at least $B/2 + \delta$. By individual rationality and budget feasibility \mathcal{M} cannot allocate to any other agent. Observe that the optimal solution in this case includes two agents. \square

We call the mechanism above a *proportional share mechanism* since it shares the budget among agents proportionally to their contribution. We will now develop this idea further and show how it can be applied to monotone submodular utility functions ($f(S) = |S|$ is a very special case of a submodular function). As we will briefly discuss, this requires overcoming some fundamental challenges that emphasize the clash between incentive compatibility and the budget constraint.

Submodular functions. Recall that a function $f : 2^{[n]} \rightarrow \mathcal{R}_+$ is *submodular* if $f(S \cup T) \leq f(S) + f(T) - f(S \cap T)$, and *monotone* if $S \subseteq T$ implies $f(S) \leq f(T)$. For monotone submodular functions, the marginal contribution of an agent a_i given a subset S , is $f_S(a_i) = f(S \cup \{a_i\}) - f(S)$. In an analogous manner to the way in which we sorted agents according to their bids, when the buyer has a general monotone submodular utility function we can consider bids by a marginal-contribution-per-cost order. That is, a sorting where the agent that appears in position $i + 1$ is the agent a_j for which $f_{S_i}(a_j)/c_j$ is maximal, where $S_i = \{a_1, a_2, \dots, a_i\}$ and $S_0 = \emptyset$. To simplify notation, we will write f_i instead of $f_{S_{i-1}}(a_i)$. Note that $f(S_k) = \sum_{i \leq k} f_i$. In the presence of submodularity this sorting implies:

$$\frac{f_1}{c_1} \geq \frac{f_2}{c_2} \geq \dots \geq \frac{f_n}{c_n} \quad (1)$$

The proportional share allocation rule. The *proportional share allocation rule* sorts agents according to (1) and allocates to agents $i \in \{1, \dots, k\}$ that respect $c_i \leq B \cdot f_i / f(S_i)$. For concreteness consider the special case when the utility function is additive, i.e. each agent a_i is associated with a fixed value v_i and $f(S) = \sum_{i \in S} v_i$. Here the marginal contribution of each agent is independent of their place in the sorting, and we simply have that $f_i = v_i$ for all agents. In this case the proportional share allocation rule produces a budget-feasible mechanism. The reason is, it assures us that for each agent a_i , the payments that support truthfulness θ_i do not exceed the agent's proportional share:

$$\theta_i = \min \left\{ \frac{v_i \cdot B}{\sum_{i \in S} v_i}, \frac{v_i \cdot b_{k+1}}{v_{k+1}} \right\}$$

which is budget feasible since $\sum_i \theta_i \leq B$, and $b_i \leq \theta_i$.

For many special cases of submodular functions, minor adjustments to the proportional share allocation rule produce budget feasible mechanisms with good approximation guarantees, which in some cases are even optimal (e.g. when $f(S) = |S|$). For the more intricate cases of submodular functions however, the proportional share allocation rule does not guarantee budget feasibility. For coverage functions for example, where each agent is associated with some set of elements T_i and $f(S) = |\cup_{i \in S} T_i|$, the marginal contribution of an agent is not fixed, but depends on the subset allocated by the algorithm in the previous stages. In this case, paying θ_i as above with f_i instead of v_i will be under budget but will not support truthfulness. This is because the payment depends on the marginal contribution f_i , which is determined by a_i 's position in the sorting. Since the position in the sorting depends on the agent's bid paying $f_i \cdot \min\{B/f(S_k), b_{k+1}/f_{k+1}\}$ is budget feasible but cannot be truthful (see Figure 2).

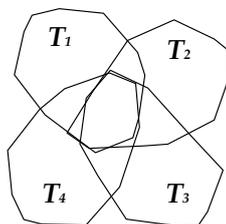


Fig. 2. An illustration of a coverage function. Assume each agent a_i contributes the set T_i , and the order in which the proportional share allocation rule selects agents is (a_1, a_2, a_3, a_4) . If the mechanism paid agents proportionally to their marginal contribution, agent a_4 may wish to bid 0, which would maximize her marginal contribution since this bid places her first in the marginal-contribution-per-cost sorting.

For utility functions like coverage the fact that agents' contributions heavily depend on their bids make it seem like there is little hope in exploring proportional share mechanisms. Fortunately, it turns out that the proportional share allocation rule can be modified into a budget feasible mechanism with constant factor approximation guarantee for monotone submodular functions. This is the main technical result in [Singer 2010] and can be summarized as follows:

- (1) **Payment characterization.** First, we derive a characterization of the truthful payments of the proportional share allocation rule. The characterization reveals an underlying structure which plays an important role in our design.
- (2) **Bounding payments.** Using the above characterization, we show that for any monotone submodular function, we can slightly modify the proportional share allocation rule so that its truthful payments are only a constant factor away from the agents' proportional contributions $B \cdot f_i/f(S_k)$. We can therefore run this modified version of the proportional share allocation rule *with a constant fraction of the budget*, and the payments will be budget feasible.
- (3) **Approximation guarantee.** Finally, we apply the variation of the proportional share rule over a particular *subset of agents* and use randomization to provide a good approximation guarantee without breaking monotonicity of the allocation rule.³

THEOREM 2.3 [SINGER 2010]. *For any monotone submodular function there exists a randomized universally truthful budget feasible mechanism with a constant factor approximation ratio.*

3. BEYOND SUBMODULARITY

The positive result for submodular functions opens the possibility for broader classes of functions to have budget feasible mechanisms. In our introduction we saw a simple example of a superadditive utility function where no budget feasible with

³The approximation ratio shown in [Singer 2010] was 117.7. Chen, Gravin, and Lu later gave a much tighter bound of 7.91 using clever analysis [Chen et al. 2011].

approximation ratio better than n can be obtained (and a similar function can be constructed to show such a lower bound for randomized mechanisms). It therefore seems like the natural class of functions that can potentially allow for budget feasibility is that of *subadditive* functions, i.e. $f(S \cup T) \leq f(S) + f(T)$.

For subadditive utility functions budget feasible mechanisms with polylogarithmic approximation ratios were presented in [Dobzinski et al. 2011]. Bei et al. give a mechanism with a sublogarithmic approximation ratio, and a mechanism with a constant factor approximation ratio for a class known as *fractionally subadditive* functions [Bei et al. 2012]. Interestingly, they also show a budget feasible mechanism for subadditive functions with a constant factor approximation ratio when agents' costs are drawn from a *known distribution*. This result opens a new and exciting direction for *Bayesian* budget feasible mechanism design.⁴

It is in fact an open question whether there exists a budget feasible mechanism with a constant factor approximation guarantee for subadditive utility functions. It is important to note that without incentive constraints, a constant factor approximation algorithm (which relies on access to demand oracles) for maximizing subadditive functions under a budget constraint exists [Bei et al. 2012]. Since subadditive functions are arguably the most natural class where budget feasibility seems possible, whether a constant factor budget feasible mechanism for this class exists is a fundamental open question.

4. ONLINE LEARNING AND POSTED PRICE MECHANISMS

Throughout our discussion we considered *direct revelation* mechanisms where the agents bid their costs. An alternative is *posted price* mechanisms where the mechanism makes agents “take-it-or-leave-it” offers. In this model we assume agents are drawn sequentially from an *unknown distribution* that describes their costs, and for each agent a_i the mechanism posts a price p_i . If $p_i \geq c_i$ the agent accepts and the buyer receives the item or service a_i sells at the expense of losing p_i from the remaining budget. Otherwise the agent rejects the offer, and the budget and utility of the buyer remain as they were before making the offer.

The main open question here is whether posted price mechanisms can do asymptotically as well as (i.e. a constant factor away from) direct revelation mechanisms. The technical challenge is to learn enough about distribution under the budget: high offers quickly exhaust the budget, and low offers can quickly exhaust the pool of agents. Surprisingly, it turns out that for simple utility functions such as $f(S) = |S|$, posted prices are as asymptotically powerful as direct revelation [Badanidiyuru et al. 2012]. Whether this is also true for broader classes of utility functions is open.

⁴Throughout our discussion we ignored computational and information theoretic considerations. To give a quick overview, submodular functions and their superclasses may have exponential representation and thus an efficient algorithm requires an oracle to evaluate these functions. The results discussed above for submodular functions are obtained using *value oracles*. For slightly broader classes there are known impossibility results for optimization with value oracles [Mirrokni et al. 2008] which apply to our setting as well. The mechanisms described in this section use the stronger *demand oracle* model. For more details see [Blumrosen and Nisan 2009].

5. APPLICATIONS

Influence in Social Networks. In this setting, formalized in [Kempe et al. 2003], the goal is to select a small set of individuals to become initial adopters or recommend a technology so that the word-of-mouth effect is maximized. In their seminal work Kempe, Kleinberg, and Tardos showed that for many natural word-of-mouth processes the task of maximizing influence reduces to maximizing a submodular function under a cardinality constraint. The underlying assumption in influence maximization is that every individual has some inherent cost for being a recommender or an initial adopter. Despite the growing availability of personalized data and the emergence of sophisticated machine learning techniques, inferring personal costs is a difficult task. The results for submodular functions here imply that individuals' costs can be elicited through a truthful mechanism with good performance guarantees. Budget feasible mechanisms for particular classes of influence functions as well as experiments on large-scale data are shown in [Singer 2012].

Crowdsourcing. In crowdsourcing markets like Amazon's Mechanical Turk, *requesters* typically seek to outsource simple human computation tasks such as image labeling, in a cost-effective manner. Requesters often face task completion deadlines and must account for dramatic elasticity in the workforce supply. In addition, requesters must accommodate the large variance in effort required to complete different tasks, which largely depends workers' heterogeneous skill levels. To automate the process of pricing crowdsourcing tasks the MECHANICAL PERK™ platform implements budget feasible mechanisms tailored for crowdsourcing markets that maximize requesters' utilities and minimize payments [Singer and Mittal 2011; 2013]. Singla and Krause develop direct revelation as well as posted price mechanisms that minimize regret [Singla and Krause 2013b], and applications of budget feasible mechanisms for crowdsourcing with mobile devices are shown in [Yang et al. 2012].

Data Analysis. There are several interesting applications where data is procured from strategic agents and the goal is to optimize estimations and predictions. Roth and Schoenebeck assume costs are drawn from a known distribution and use a posted price budget feasible mechanism for minimizing variance of an estimator [Roth and Schoenebeck 2012]. Recent work by Horel et al. shows budget feasible mechanisms for experimental design, and manage to elegantly optimize objectives such as entropy minimization [Horel et al. 2013].

Privacy-Preserving Systems. In a data-rich era, recommendation systems aim to strike a fine balance between making useful recommendations while preserving privacy. Budget feasible mechanisms that procure data from users and maximize prediction in a privacy-preserving manner were developed in [Dandekar et al. 2013]. For a different set of problems known as *community sensing* Singla and Krause develop budget feasible mechanisms for *adaptive submodular* utilities [Singla and Krause 2013a].

6. CONCLUSION

Our goal in this letter is to provide a taste of the challenges and applications of budget feasible mechanism design. The model fits scenarios in which the buyer is limited by a budget, or ones where budget feasible mechanism design can serve other objectives, such as minimizing costs as in the crowdsourcing example. Of course, our focus in this brief survey was intentionally narrow and one can certainly consider procurement models with different objectives and different constraints.

REFERENCES

- BADANIDIYURU, A., KLEINBERG, R., AND SINGER, Y. 2012. Learning on a budget: posted price mechanisms for online procurement. In *ACM Conference on Electronic Commerce*. 128–145.
- BEI, X., CHEN, N., GRAVIN, N., AND LU, P. 2012. Budget feasible mechanism design: from prior-free to bayesian. In *STOC*. 449–458.
- BLUMROSEN, L. AND NISAN, N. 2009. On the computational power of demand queries. *SIAM J. Comput.* 39, 4, 1372–1391.
- CHEN, N., GRAVIN, N., AND LU, P. 2011. On the approximability of budget feasible mechanisms. In *SODA*. 685–699.
- DANDEKAR, P., FAWAZ, N., AND IOANNIDIS, S. 2013. Privacy auctions for recommender systems. In *ACM Transactions on Economics and Computation (TEAC)*.
- DOBZINSKI, S., PAPADIMITRIOU, C., AND SINGER, Y. 2011. Mechanisms for complement-free procurement. In *EC*.
- HOREL, T., IOANNIDIS, S., AND MUTHUKRISHNAN, S. 2013. Budget feasible mechanisms for experimental design. *CoRR abs/1302.5724*.
- KEMPE, D., KLEINBERG, J. M., AND TARDOS, É. 2003. Maximizing the spread of influence through a social network. In *KDD*. 137–146.
- MIRROKNI, V. S., SCHAPIRA, M., AND VONDRÁK, J. 2008. Tight information-theoretic lower bounds for welfare maximization in combinatorial auctions. In *ACM Conference on Electronic Commerce*. 70–77.
- MYERSON, R. 1981. Optimal auction design. *Mathematics of Operations Research* 6, 1.
- ROTH, A. AND SCHOENEBECK, G. 2012. Conducting truthful surveys, cheaply. In *ACM Conference on Electronic Commerce*. 826–843.
- SINGER, Y. 2010. Budget feasible mechanisms. In *FOCS*. 765–774.
- SINGER, Y. 2012. How to win friends and influence people, truthfully: influence maximization mechanisms for social networks. In *WSDM*. 733–742.
- SINGER, Y. AND MITTAL, M. 2011. Mechanisms for online labor markets. In *Workshop on Human Computation*.
- SINGER, Y. AND MITTAL, M. 2013. Pricing mechanisms for crowdsourcing markets. In *WWW*. 1157–1166.
- SINGLA, A. AND KRAUSE, A. 2013a. Incentives for privacy tradeoff in community sensing. In *HCOMP*.
- SINGLA, A. AND KRAUSE, A. 2013b. Truthful incentives in crowdsourcing tasks using regret minimization mechanisms. In *WWW*. 1167–1178.
- YANG, D., XUE, G., FANG, X., AND TANG, J. 2012. Crowdsourcing to smartphones: Incentive mechanism design for mobile phone sensing. In *Proceedings of the 18th Annual International Conference on Mobile Computing and Networking*. Mobicom '12. ACM, New York, NY, USA, 173–184.