Competitive Auctions for Markets with Positive Externalities

Nick Gravin[1] and Pinyan Lu[2]

[1] Division of Mathematical Sciences, School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore
ngravin@pmail.ntu.edu.sg
[2] Microsoft Research Asia, China
pinyanl@microsoft.com

Abstract. In digital goods auctions, the auctioneer sells an item in unlimited supply to a set of potential buyers. The objective is to design a truthful auction that maximizes the auctioneer's total profit. Motivated by the observation that the buyers' valuation of the good might be interconnected through a social network, we study digital goods auctions with positive externalities among buyers. This defines a multi-parameter auction design problem where the private valuation of every buyer is a function of the set of other winning buyers. The main contribution of this paper is a truthful competitive mechanism for subadditive valuations. Our competitive result is with respect to a new solution benchmark $\mathcal{F}^{(3)}$. On the other hand, we show a surprising impossibility result if comparing to the stronger benchmark $\mathcal{F}^{(2)}$, where the latter has been used quite successfully in digital goods auctions without externalities [16].

1 Introduction

In economics, the term externality is used to describe situations in which private costs or benefits to the producers or purchasers of a good or service differ from the total social costs or benefits entailed in its production and consumption. In this context a benefit is called a positive externality, while a cost is referred to as a negative one. One needs not to go far to find examples of positive external influence in digital and communications markets, when a customer's decision to buy a good or purchase a service strongly relies on its popularity among his/her friends or generally among other customers, e.g. instant messenger and cell phone users will want a product that allows them to talk easily and cheaply with their friends. Another good example is social network, where a user is more likely to appreciate membership in a network if many of his/her friends are already using it. There exist a number of applications, like the very popular Farm Ville in online social network Facebook, where a user would have more fun when participating with friends. In fact, quite a few such applications explicitly reward players with a large number of friends.

On the other hand, negative external effects occur when a potential buyer, e.g. a big company, incurs a great loss if a subject it fights for, like a small firm or

F.V. Fomin et al. (Eds.): ICALP 2013, Part II, LNCS 7966, pp. 569–580, 2013.

company, goes to its direct competitor. Another well-studied example related to computer science is the allocation of advertisement slots [1, 13–15, 17, 23], where every customer would like to see a smaller number of competitors' advertisements on a web page that contains his/her own advert. One may also face mixed externalities as in the case of selling nuclear weapons [21], where countries would like to see their allies win the auction rather than their foes.

We investigate the problem of *mechanism design* for auctions with positive externalities. We study a scenario where an auctioneer sells the good, of no more than a single copy in the hands of each customer. We define a model for externalities among the buyers in the sealed-bid auction with an unlimited supply of the good. This types of auctions arise naturally in digital markets, where making a copy of the good (e.g. cd with songs or games, or extra copy of online application) has a negligible cost compared to the final price and can be done at any time the seller chooses.

A similar agenda has been introduced in the paper [18], where authors consider a Bayesian framework and study positive externalities in the social networks with single-parameter bidders and submodular valuations. The model in the most general form can be described by a number of bidders n, each with a non-negative private valuation function $v_i(S)$ depending on the possible winning set S. This is a natural multi-parameter mechanism design model that may be considered a generalization of the classical auctions with unlimited supply, i.e. auctions where the amount of items being sold is greater than the number of buyers.

Traditionally the main question arising in such situations is how to maximize the seller's revenue. In literature on the classical auctions without any externalities many diverse approaches to this question have been developed. In the current work we pick a classical approach and benchmark (cf. [16]), namely the best-uniform-price benchmark called $\mathcal{F}$, which is different from Bayesian framework. There one seeks to maximize the ratio of the mechanism's revenue to the revenue of $\mathcal{F}$ taken in the worst case over all possible bids. In particular a mechanism is called competitive if such a ratio is bounded by some uniform constant for each possible bid. However, it was shown that there is no competitive truthful mechanism w.r.t. $\mathcal{F}$, and therefore to get around this problem, a slightly modified benchmark $\mathcal{F}^{(2)}$ [16] was proposed. The only difference of $\mathcal{F}^{(2)}$ to $\mathcal{F}$ is in one additional requirement that at least two buyers should be in a winning set. Thus $\mathcal{F}^{(2)}$ becomes a standard benchmark in analyzing digital auctions [11, 12, 16, 20]. Similarly to $\mathcal{F}^{(2)}$ one may define benchmark $\mathcal{F}^{(k)}$ for any fixed constant k. It turns out that the same benchmarks can be naturally adopted to the case of positive externalities. Surprisingly $\mathcal{F}^{(2)}$ fails to serve as a benchmark in social networks with positive externalities, i.e. no competitive mechanism exists w.r.t. $\mathcal{F}^{(2)}$. Therefore, we go further and consider the next natural candidate for the benchmark, which is $\mathcal{F}^{(3)}$.

The main contribution of this paper is an universally truthful competitive mechanism for the general multi-parameter model with subadditive valuations (substantially broader class than submodular) w.r.t. $\mathcal{F}^{(3)}$ benchmark. We complement this result with a proof that no truthful mechanism can achieve constant ratio w.r.t. $\mathcal{F}^{(2)}$. In order to do so we introduce a restricted model with a single private parameter which in some respects resembles the one considered in [18]; further for this restricted model we give a simple geometric characterization of all truthful mechanisms and based on the characterization then show that there is no competitive truthful mechanism w.r.t. $\mathcal{F}^{(2)}$.

Our model is the so-called multi-parameter or multi-dimensional model (see [25]), as utility of every agent may not be described by a single real number for all possible outcomes of the mechanism. Mechanism design in this case is known to be harder than in the single-parameter domains.

1.1 Related Work

Many studies on externalities in the direction of pricing and marketing strategies over social networks have been conducted over the past few years. In many ways, they have been caused by the development of social-networks on the Internet, which has allowed companies to collect information about each user and user relationships.

Earlier works have generally been focused on the influence maximization problems (see Chapter 24 of [24]). For instance, Kempe *et al.* [22] study the algorithmic question of searching a set of nodes in a social network of highest influence. From the economic literature one could name such papers as [26], which studies the effect of network topology on a monopolist's profits and [10], which studies a multi-round pricing game, where a seller may lower his price in an attempt to attract low value buyers. These works take no heed of algorithmic motivation.

There are several more recent papers [2, 7, 9, 19] studying the question of revenue maximization as well as work studying the posted price mechanisms [3,5,8,19].

We could not continue without mentioning a beautiful line of research on revenue maximization for classical auctions, where the objective is to maximize the seller's revenue compared to a benchmark in the worst case. We cite here only some papers that are most relevant to our setting [4,11,12,16,20]. With respect to the refined best-uniform-price benchmark $\mathcal{F}^{(2)}$ a number of mechanisms with constant competitive ratio were obtained; each subsequent paper improving the competitive ratio of the previous one [11, 12, 16, 20]. The best known current mechanism is due to Hartline and McGrew [20] and has a competitive ratio of 3.25. On the other hand a lower bound of 2.42 has been proven in [16] by Goldberg *et.al.*. The question of closing the gap still remains open.

2 Preliminaries

We suppose that in a marketplace n agents are present, the set of which we denote by $[n]$. Each agent i has a private valuation function v_i, which is a non-negative real number for each possible winner set $S \subset [n]$. The seller organizes a single round sealed bid auction, where agents submit their valuations $b_i(S)$ to an auctioneer for all possible winner sets S, and the auctioneer then chooses agents who will obtain the good and vector of prices to charge each of them. The auctioneer is interested in maximizing his/her revenue.

For every $i \in [n]$ we impose the following mild requirements on v_i.

1. $v_i(S) \geq 0$.
2. $v_i(S) = 0$ if $i \notin S$.
3. $v_i(S)$ is a monotone sub-additive function of S, i.e.
   (a) $v_i(S) \leq v_i(R)$ if $S \subseteq R \subseteq [n]$.
   (b) $v_i(S \cup R) \leq v_i(S) + v_i(R)$, for each $i \in S, R \subseteq [n]$

We should note here that the sub-additivity requirement is only for those subsets that include the agent i. This is a natural assumption since $v_i(S) = 0$ if $i \notin S$.

2.1 Mechanism Design

Each agent in turn would like to get a positive utility that is as high as possible and may lie strategically about his/her valuation. The utility $u_i(S)$ of an agent i for a winning set S is simply the difference of his valuation $v_i(S)$ and the price p_i the auctioneer charges i. Thus one of the desired properties for the auction is the well known concept of truthfulness or incentive compatibility, i.e. the condition that every agent maximizes his utility by truth telling.

It is worth mentioning that our model is that of multi-parameter mechanism design and, moreover, that collecting the whole bunch of values $v_i(S)$ for every $i \in [n]$ and $S \subset [n]$ would require an exponential amount of bits in n and thus is inefficient. However, in the field of mechanism design there is a way to get around such a problem of exponential input size with the broadly recognized concept of black box value queries. The latter simply means that the auctioneer, instead of getting the whole collection of bids instantly, may ask during the mechanism execution every agent i only for a small part of his input, i.e. a number of questions about valuation of i for certain sets. We note that the agent still may lie in the response to each such query. We denote the bid of i by $b_i(S)$ to distinguish it from the actual valuation $v_i(S)$. Thus if we are interested in designing a computationally efficient mechanism, we can only ask in total a polynomial in n number of queries.

Throughout the paper, with $\mathcal{M}$ we denote a mechanism with allocation rule $\mathcal{A}$ and payment rule $\mathcal{P}$. Allocation algorithm $\mathcal{A}$ may ask queries about valuations of any agent for any possible set of winners. Thus $\mathcal{A}$ has an oracle black box access to the collection of bid functions $b_i(S)$. For each agent i in the winning set S the payment algorithm decides a price p_i to charge. The utility of agent i is then $u_i = v_i(S) - p_i$ if $i \in S$ and 0 otherwise. To emphasize the fact that agents may report untruthfully we will use $u_i(b_i)$ notation for the utility function in the general case and $u_i(v_i)$ in the case of truth telling. We assume voluntary participation for the agents, that is $u_i \geq 0$ for each i who reports the truth.

2.2 Revenue Maximization and Possible Benchmarks

We discuss here the problem of revenue maximization from the seller's point of view. The revenue of the auctioneer is simply the total payment $\sum_{i \in S} p_i$ of all buyers in the winning set. We assume that the seller incurs no additional cost for making a copy of the good. This assumption is essential for our model, since unlike the classical digital auction case there is no simple reduction of the settings with a positive price per issuing the item to the settings with zero price.

The best revenue the seller can hope for is $\sum_{i\in[n]} v_i([n])$. However, it is not realistic when the seller does not know agents' valuation functions. We follow the tradition of previous literature [11, 12, 16, 20] of algorithmic mechanism design on competitive auctions with limited or unlimited supply and consider the best revenue uniform price benchmark, which is defined as maximal revenue that the auctioneer can get for a fixed uniform price for the good. In the literature on classical competitive auctions this benchmark was called $\mathcal{F}$ and is formally defined as follows.

Definition 1 ($\mathcal{F}$ without Externalities). *For the vector of agent's bids* $\mathbf{b}$

$$\mathcal{F}(\mathbf{b}) = \max_{c\geq 0, S\subset[n]} \left(c\cdot |S| \middle| \forall i \in S \quad b_i \geq c\right).$$

This definition generalizes naturally to our model with externalities and is defined rigorously as follows.

Definition 2 ($\mathcal{F}$ with Externalities). *For the collection of agents' bid functions* $\mathbf{b}$.

$$\mathcal{F}(\mathbf{b}) = \max_{c\geq 0, S\subset[n]} \left(c\cdot |S| \middle| \forall i \in S \quad b_i(S) \geq c\right).$$

The important point in considering $\mathcal{F}$ in the setting of classical auctions is that the auctioneer, when he/she is given in advance the best uniform price, can run a truthful mechanism with corresponding revenue. It turns out that the same mechanism works neatly for our model. Specifically, a seller who is given the price c in advance can begin with the set of all agents and drop one by one those agents with negative utility ($b_i(S) - c < 0$); once there are left no agents to delete, the auctioneer sells the item to all surviving buyers at the given price c.

In these circumstances, a natural problem arising for the auctioneer is to devise a truthful mechanism which has a good approximation ratio of the mechanism's revenue to the revenue of the benchmark at any possible bid vector $\mathbf{b}$. Such a ratio is usually called the *competitive* ratio of a mechanism. However, it was shown (cf. [16]) that no truthful mechanism can guarantee any constant competitive ratio w.r.t. $\mathcal{F}$. Specifically, the unbounded ratio appears in the instances where the benchmark buys only one item at the highest price. To overcome this obstacle, a slightly modified benchmark $\mathcal{F}^{(2)}$ has been proposed and a number of competitive mechanisms w.r.t. $\mathcal{F}^{(2)}$ were obtained [11, 12, 16, 20]. The only difference of $\mathcal{F}^{(2)}$ from $\mathcal{F}$ is one additional requirement that at least two buyers should be in the winning set. Similarly, for any $k \geq 2$ we may define $\mathcal{F}^{(k)}$.

Definition 3.

$$\mathcal{F}^{(k)}(\mathbf{b}) = \max_{c\geq 0, S\subset[n]} \left(c\cdot |S| \middle| |S| \geq k, \quad \forall i \in S \quad b_i(S) \geq c\right).$$

However, in case of our model the benchmark $\mathcal{F}^{(2)}$ does not imply the existence of a constant approximation truthful mechanism. In order to illustrate that later in Section 4 we will introduce a couple of new models which differ from the original one in certain additional restrictions on the domain of agent's bids. We further give a complete characterization of truthful mechanisms for these new restricted settings substantially exploiting the fact that every agent's bidding language is single-parameter. Later, we use that characterization to argue that no truthful mechanism can achieve constant approximation with respect to $\mathcal{F}^{(2)}$ benchmark even for these cases. On the positive side, and quite surprisingly, we can furnish our work in the next section with the truthful mechanism which has a constant approximation ratio w.r.t. $\mathcal{F}^{(3)}$ benchmark for the general case of multi-parameter bidding.

3 Competitive Mechanism

Here we give a competitive truthful mechanism, that is a mechanism which guarantees that the auctioneer gets a constant fraction of the revenue he could get for the best fixed price benchmark assuming that all agents bid truthfully. We call it PROMOTION-TESTING-SELLING MECHANISM. In the mechanism we give the good to certain agents *for free*, that is without requiring any payment. The general scheme of the mechanism is as follows.

PROMOTION-TESTING-SELLING MECHANISM

1. Put every agent at random into one of the sets A, B, C.
2. Denote $r_A(C)$ and $r_B(C)$ the largest fixed price revenues one can extract from C given that, respectfully, either A, or B got the good for free.
3. Let $r(C) = \max\{r_A(C), r_B(C)\}$.
4. Sell items to agents in A for free.
5. Apply COST SHARING MECHANISM(r(C), B, A) to extract revenue $r(C)$ from set B given that A got the good for free.

Bidders in A receive items for free and increase the demand of agents from B. One may say that they "advertise" the goods and resemble the promotion that occurs when selling to participants. The agents in C play the role of the "testing" group, the only service of which is to determine the right price. Note that we take no agents of the testing group into the winning set, therefore, they have nothing to gain for bidding untruthfully. The agents of B appear to be the source of the mechanism's revenue, which is being extracted from B by a cost sharing mechanism as follows.

We note here that a more "natural" mechanism is simply to set that $r(C) = r_A(C)$ rather than $\max\{r_A(C), r_B(C)\}$. But unfortunately, we have a counter example to show that this simpler mechanism cannot guarantee a constant approximation ratio compared to our benchmark.

```
COST SHARING MECHANISM(r,X,Y)

1. S ← X.
2. Repeat until T = ∅:
     – T ← {i|i ∈ S and b_i(S ∪ Y) < r/|S|}.
     – S ← S \ T.
3. If S ≠ ∅ sell items to everyone in S at r/|S| price.
```

Lemma 4. Promotion-Testing-Selling Mechanism *is universally truthful.*

Proof. The partitioning of the set $[n]$ into A, B, C does not depend on the agent bids. When the partition is fixed, our mechanism becomes deterministic. Therefore, we are only left to prove the truthfulness for that deterministic part. Let us do so by going through the proof separately for each set A, B and C.

– Bids of agents in A do not affect the outcome of the mechanism. Therefore, they have no incentive to lie.
– No agent from C could profit from bidding untruthfully, since her utility will be zero regardless of the bid.
– Let us note that the Cost Sharing Mechanism is applied to the agents in B and the value of r does not depend on their bids, since both $r_A(C)$ and $r_B(C)$ are retracted from C irrespectively of bids from A and B. Also let us note that at each step of the cost sharing mechanism the possible payment $\frac{r}{|S|}$ is rising, and meanwhile the valuation function, because of monotonicity condition, is going down. Hence, manipulation of a bid does not help any agent to survive in the winning set and receive positive utility, if by bidding truthfully he/she has been dropped from the winning set. Mis-reporting a bid could not help an agent to alter the surviving set and at the same time remain a winner. These two observations conclude the proof of truthfulness for B.

Therefore, from now on we may assume that $b_i(S) = v_i(S)$.

Theorem 5. Promotion-Testing-Selling Mechanism *is universally truthful and has an expected revenue of at least* $\frac{\mathcal{F}^{(3)}}{324}$.

Proof. We are left to prove the lower bound on the competitive ratio of our mechanism, as we have shown the truthfulness in Lemma 4.

For the purpose of analysis, we separate the random part of our mechanism into two phases. In the first phase, we divide agents randomly into three groups S_1, S_2, S_3 and in the second one, we label the groups at random by A, B and C. Note that the combination of these two phases produces exactly the same distribution over partitions as in the mechanism.

Let S be the set of winners in the optimal $\mathcal{F}^{(3)}$ solution and the best fixed price be p^*. For $1 \le i \ne j \le 3$ we may compute r_{ij} the largest revenue for a fixed price that one can extract from set S_i given S_j is "advertising" the good, that is agents in S_j get the good for free and thus increase the valuations of agents from S_i though contribute nothing directly to the revenue.

First, let us note that the cost-sharing part of our mechanism will extract one of these r_{ij} from at least one of the six possible labels for every sample of the dividing phase (in general cost-sharing mechanism may extract 0 revenue, e.g. if the target revenue is set too high). Indeed, let i_0 and j_0 be the indexes for which $r_{i_0 j_0}$ achieves maximum over all r_{ij} and let $k_0 = \{1, 2, 3\} \setminus \{i_0, j_0\}$. Then the cost-sharing mechanism will retract the revenue $r(C) = max(r_{_A}(C), r_{_B}(C))$ on the labeling with $S_{j_0} = A$, $S_{i_0} = B$ and $S_{k_0} = C$. It turns out, as we will prove in the following lemma, that one can get a lower bound on this revenue within a constant factor of $r_{_\mathcal{F}}(C)$; the revenue we got from the agents of C in the benchmark $\mathcal{F}^{(3)}$.

Lemma 6. $r(C) \geq \frac{r_{\mathcal{F}}(C)}{4}$.

Proof. Let $S_c = S \cap C$. Thus, by the definition of $\mathcal{F}^{(3)}$, we have $r_{_\mathcal{F}}(C) = |S_c| \cdot p^*$ and for all $i \in S_c$, $v_i(S) \geq p^*$.

We define a subset T of S_c as a final result of the following procedure.

```
1. T ← ∅ and X ← {i|i ∈ S_c and v_i(A ∪ {i}) ≥ p*/2}.
2. While X ≠ ∅
     – T ← T ∪ X,
     – X ← {i|i ∈ S_c and v_i(A ∪ T ∪ {i}) ≥ p*/2}
```

For any agent of T we have $v_i(A \cup T) \geq \frac{p^*}{2}$ because the valuation function is monotone. Now if $|T| \geq \frac{|S_c|}{2}$, we get the desired lower bound. Indeed,

$$r(C) \geq r_{_A}(C) \geq \frac{|S_c|}{2} \cdot \frac{p^*}{2} = \frac{|S_c| \cdot p^*}{4} = \frac{r_{_\mathcal{F}}(C)}{4}.$$

Otherwise, let $W = S_c \setminus T$. Then we have $|W| \geq \frac{|S_c|}{2}$. For an agent $i \in W$ it holds true that $v_i(A \cup T \cup \{i\}) < \frac{p^*}{2}$, since otherwise we should include i into T. However, since i wins in the optimal $\mathcal{F}^{(3)}$ solution, we have $v_i(S) \geq p^*$. The former two inequalities together with the subadditivity of $v_i(\cdot)$ $(v_i(S \setminus (A \cup T)) + v_i(A \cup T \cup \{i\}) \geq v_i(S))$ allow us to conclude that $v_i(S \setminus (A \cup T)) \geq \frac{p^*}{2}$ for each $i \in W$. Hence, we get $v_i(B \cup W) \geq \frac{p^*}{2}$ for each $i \in W$, since $S \setminus (A \cup T) \subseteq B \cup W$. Therefore, we are done with the lemma's proof, since

$$r(C) \geq r_{_B}(C) \geq |W| \cdot \frac{p^*}{2} \geq \frac{|S_c| \cdot p^*}{4} = \frac{r_{_\mathcal{F}}(C)}{4}.$$

Let k_1, k_2, k_3 be the number of winners of the optimal $\mathcal{F}^{(3)}$ solution, respectively, in S_1, S_2, S_3.

For any fixed partition S_1, S_2, S_3 of the dividing phase by applying Lemma 6, we get that the expected revenue of our mechanism over a distribution of six permutations in the second phase should be at least

$$\frac{1}{6} \cdot \frac{1}{4} \min\{k_1, k_2, k_3\} \cdot p^*.$$

In order to conclude the proof of the theorem we are only left to estimate the expected value of $\min\{k_1, k_2, k_3\}$ from below by some constant factor of $|S|$. The next lemma will do this for us.

Lemma 7. *Let $m \geq 3$ items independently at random be put in one of the three boxes and let a, b and c be the random variables denoting the number of items in these boxes. Then $\mathbb{E}[\min\{a, b, c\}] \geq \frac{2}{27}m$.*

By definition of the benchmark $F^{(3)}$ we have $m = k_1 + k_2 + k_3 \geq 3$ and thus we can apply Lemma 7. Combining every bound we have so far on the expected revenue of our mechanism we conclude the proof with the following lower bound.

$$\frac{1}{6} \cdot \frac{1}{4} \mathbb{E}\left[\min\{k_1, k_2, k_3\}\right] \cdot p^* \geq \frac{1}{24} \cdot \frac{2}{27} \cdot p^* \cdot m = \frac{F^{(3)}}{324}.$$

4 Restricted Single-Parameter Valuations

Here we introduce a couple of special restricted cases of the general setting with a single parameter bidding language. For these models we only specify restrictions on the valuation functions. In each case we assume that t_i is a single private parameter for agent i that he submits as a bid and $w_i(S)$ and $w'_i(S)$ are fixed publicly known functions for each possible winning set S. The models then are described as follows.

- *Additive* valuation $v_i(t_i, S) = t_i + w_i(S)$.
- *Scalar* valuation $v_i(t_i, S) = t_i \cdot w_i(S)$.
- *Linear* valuation $v_i(t_i, S) = t_i w_i(S) + w'_i(S)$, i.e. combination of previous two.

Note that we still require $w_i(S) = w'_i(S) = 0$ if $i \notin S$. These settings are now single parameter domains, which is the most well studied and understood case in mechanism design.

4.1 A Characterization

The basic question of mechanism design is to describe truthful mechanisms in terms of simple geometric conditions. Given a vector of n bids, $\mathbf{b} = (b_1, \ldots, b_n)$, let b_{-i} denote the vector, where b_i is replaced with a '?'. It is well known that truthfulness implies a *monotonicity* condition stating that if an agent i wins for the bid vector $\mathbf{b} = (b_{-i}, b_i)$ then she should win for any bid vector (b_{-i}, b'_i) with $b'_i \geq b_i$. In single-dimensional domains monotonicity turns out to be a sufficient condition for truthfulness [6], where prices are determined by the threshold functions.

In our model, valuation of an agent may vary for different winning sets and thus may depend on his/her bid. Nevertheless, any truthful mechanism still has to have a bid-independent allocation rule, although now it is not sufficient for the truthfulness. However, in the case of linear valuation functions we are capable of giving a complete characterization.

Theorem 8. *In the model with linear valuation functions $v_i(t_i, S) = t_i \cdot w_i(S) + w'_i(S)$ an allocation rule $\mathcal{A}$ may be truthfully implemented if and only if it satisfies the following conditions:*

1. *$\mathcal{A}$ is bid-independent, that is for each agent i, bid vector $\mathbf{b} = (b_{-i}, b_i)$ with $i \in \mathcal{A}(\mathbf{b})$ and any $b'_i \geq b_i$, it holds that $i \in \mathcal{A}(b_{-i}, b'_i)$.*
2. *$\mathcal{A}$ encourages asymptotically higher bids, i.e. for any fixed b_{-i} and $b'_i \geq b_i$, it holds that $w_i(\mathcal{A}(b_{-i}, b'_i)) \geq w_i(\mathcal{A}(b_{-i}, b_i))$.*

Here we prove that these conditions are indeed necessary. The sufficiency part of the theorem is deferred to the full paper version, where we prove the characterization for a slightly more general family of single parameter valuation functions.

Proof. The necessity of the first monotonicity condition was known, so we prove here that the second condition is also necessary. In the truthful mechanism, an agent's payment should not depend on his/her bid, if by changing it the mechanism does not shift the allocated set. We denote by p the payment of agent i for winner set $\mathcal{A}(b_{-i}, b_i)$ and by p' the payment of agent i for winner set $\mathcal{A}(b_{-i}, b'_i)$. If the agent's true value is b_i, by truthfulness, we have

$$b_i \cdot w_i(\mathcal{A}(b_{-i}, b_i)) + w'_i(\mathcal{A}(b_{-i}, b_i)) - p \geq b_i \cdot w_i(\mathcal{A}(b_{-i}, b'_i)) + w'_i(\mathcal{A}(b_{-i}, b'_i)) - p'.$$

And if the agent's true value is b'_i, we have

$$b'_i \cdot w_i(\mathcal{A}(b_{-i}, b'_i)) + w'_i(\mathcal{A}(b_{-i}, b'_i)) - p' \geq b'_i \cdot w_i(\mathcal{A}(b_{-i}, b_i)) + w'_i(\mathcal{A}(b_{-i}, b_i)) - p.$$

Adding these two inequalities and using the fact that $b'_i \geq b_i$, we have

$$w_i(\mathcal{A}(b_{-i}, b'_i)) \geq w_i(\mathcal{A}(b_{-i}, b_i)).$$

4.2 From $\mathcal{F}^{(2)}$ to $\mathcal{F}^{(3)}$

Here we show that the usage of $\mathcal{F}^{(2)}$ as a benchmark may lead to an unbounded approximation ratio even for the restricted single parameter scalar valuations. This justifies why we used a slightly modified benchmark $\mathcal{F}^{(3)}$ in Section 3.

Theorem 9. *There is no universally truthful mechanism that can achieve a constant approximation ratio w.r.t. $\mathcal{F}^{(2)}$.*

Proof. Consider the example of two people, in which every bidder evaluates the outcome, where both agents get items much higher than the outcome, where only one agent gets the item. That is $v_1(x, \{1\}) = x, v_2(y, \{2\}) = y$ and $v_1(x, \{1, 2\}) = Mx, v_2(y, \{1, 2\}) = My$ for a large constant M. We note that these are single parameter scalar valuations. We also note that these valuation functions are indeed subadditive according to our definition. The subadditive requirement is only for the subsets that includes the current agent and, in fact, any valuation function for two agents is subadditive by our definition.

We will show that any universally truthful mechanism $\mathcal{M}_{\mathcal{D}}$ with a distribution $\mathcal{D}$ over truthful mechanisms cannot achieve an approximation ratio better than

M. Each truthful mechanism $\mathcal{M}$ in $\mathcal{D}$ either sells items to both bidders for some pair of bids (b_1, b_2), or for all pairs of bids sells not more than one item. In the first case, by our characterization of truthful mechanisms (see theorem 8), $\mathcal{M}$ should also sell two items for the bids (x, b_2) and (b_1, y), where $x \geq b_1$ and $y \geq b_2$. Therefore, $\mathcal{M}$ has to sell two items for any bid (x, y) with $x \geq b_1$ and $y \geq b_2$. Let us denote the first and the second group of mechanisms in $\mathcal{D}$ by $\mathcal{G}_1$ and $\mathcal{G}_2$ respectively.

For any small ϵ we may pick sufficiently large x_0, such that at least $1-\epsilon$ fraction of $\mathcal{G}_1$ mechanisms in $\mathcal{D}$ are selling two items for the bids $(x = \frac{x_0}{2M}, y = \frac{y_0}{2M})$. Note that

- revenue of $\mathcal{F}^{(2)}$ for the bids (x_0, x_0) is $2Mx_0$,
- revenue of any $\mathcal{M}$ in $\mathcal{G}_2$ for the bids (x_0, x_0) is not greater than x_0,
- revenue of more than $1 - \epsilon$ fraction of $\mathcal{G}_1$ mechanisms in $\mathcal{D}$ is not greater than $2M\frac{x_0}{2M} = x_0$.
- revenue of the remaining ϵ fraction of $\mathcal{G}_1$ mechanisms is not greater than $2Mx_0$.

Thus we can upper bound the revenue of $\mathcal{M}_{\mathcal{D}}$ by $x_0(1-\epsilon) + 2Mx_0\epsilon$ while the revenue of $\mathcal{F}^{(2)}$ is $2Mx_0$. By choosing sufficiently large M and small ϵ we get an arbitrarily large approximation ratio.

Remark 10. In fact, the same inapproximability results w.r.t. $\mathcal{F}^{(2)}$ holds for a weaker notion of truthfulness, namely truthfulness in expectation.

References

1. Aggarwal, G., Feldman, J., Muthukrishnan, S.M., Pál, M.: Sponsored search auctions with markovian users. In: Papadimitriou, C., Zhang, S. (eds.) WINE 2008. LNCS, vol. 5385, pp. 621–628. Springer, Heidelberg (2008)
2. Akhlaghpour, H., Ghodsi, M., Haghpanah, N., Mahini, H., Mirrokni, V.S., Nikzad, A.: Optimal Iterative Pricing over Social Networks. In: Proceedings of the Fifth Workshop on Ad Auctions (2009)
3. Akhlaghpour, H., Ghodsi, M., Haghpanah, N., Mirrokni, V.S., Mahini, H., Nikzad, A.: Optimal iterative pricing over social networks (Extended abstract). In: Saberi, A. (ed.) WINE 2010. LNCS, vol. 6484, pp. 415–423. Springer, Heidelberg (2010)
4. Alaei, S., Malekian, A., Srinivasan, A.: On random sampling auctions for digital goods. In: EC, pp. 187–196 (2009)
5. Anari, N., Ehsani, S., Ghodsi, M., Haghpanah, N., Immorlica, N., Mahini, H., Mirrokni, V.S.: Equilibrium pricing with positive externalities (Extended abstract). In: Saberi, A. (ed.) WINE 2010. LNCS, vol. 6484, pp. 424–431. Springer, Heidelberg (2010)
6. Archer, A., Tardos, É.: Truthful mechanisms for one-parameter agents. In: FOCS, pp. 482–491 (2001)
7. Arthur, D., Motwani, R., Sharma, A., Xu, Y.: Pricing strategies for viral marketing on Social Networks, pp. 101–112. Springer (2009)
8. Candogan, O., Bimpikis, K., Ozdaglar, A.: Optimal pricing in the presence of local network effects. In: Saberi, A. (ed.) WINE 2010. LNCS, vol. 6484, pp. 118–132. Springer, Heidelberg (2010)

9. Chen, W., Lu, P., Sun, X., Tang, B., Wang, Y., Zhu, Z.A.: Optimal pricing in social networks with incomplete information. In: Chen, N., Elkind, E., Koutsoupias, E. (eds.) WINE 2011. LNCS, vol. 7090, pp. 49–60. Springer, Heidelberg (2011)
10. Domingos, P., Richardson, M.: Mining the network value of customers. In: ACM SIGKDD, pp. 57–66. ACM Press, New York (2001)
11. Feige, U., Flaxman, A.D., Hartline, J.D., Kleinberg, R.D.: On the competitive ratio of the random sampling auction. In: Deng, X., Ye, Y. (eds.) WINE 2005. LNCS, vol. 3828, pp. 878–886. Springer, Heidelberg (2005)
12. Fiat, A., Goldberg, A.V., Hartline, J.D., Karlin, A.R.: Competitive generalized auctions. In: STOC, pp. 72–81 (2002)
13. Ghosh, A., Mahdian, M.: Externalities in online advertising. In: WWW, pp. 161–168. ACM (2008)
14. Ghosh, A., Sayedi, A.: Expressive auctions for externalities in online advertising. In: WWW, pp. 371–380. ACM (2010)
15. Giotis, I., Karlin, A.R.: On the equilibria and efficiency of the GSP mechanism in keyword auctions with externalities. In: Papadimitriou, C., Zhang, S. (eds.) WINE 2008. LNCS, vol. 5385, pp. 629–638. Springer, Heidelberg (2008)
16. Goldberg, A.V., Hartline, J.D., Karlin, A.R., Saks, M., Wright, A.: Competitive auctions. Games and Economic Behavior 55(2), 242–269 (2006)
17. Gomes, R., Immorlica, N., Markakis, E.: Externalities in keyword auctions: An empirical and theoretical assessment. In: Leonardi, S. (ed.) WINE 2009. LNCS, vol. 5929, pp. 172–183. Springer, Heidelberg (2009)
18. Haghpanah, N., Immorlica, N., Mirrokni, V.S., Munagala, K.: Optimal auctions with positive network externalities. In: EC, pp. 11–20 (2011)
19. Hartline, J., Mirrokni, V., Sundararajan, M.: Optimal marketing strategies over social networks. In: WWW, pp. 189–198. ACM (2008)
20. Hartline, J.D., McGrew, R.: From optimal limited to unlimited supply auctions. In: EC, pp. 175–182. ACM (2005)
21. Jehiel, P., Moldovanu, B., Stacchetti, E.: How (not) to sell nuclear weapons. American Economic Review 86(4), 814–829 (1996)
22. Kempe, D., Kleinberg, J.M., Tardos, É.: Influential nodes in a diffusion model for social networks. In: Caires, L., Italiano, G.F., Monteiro, L., Palamidessi, C., Yung, M. (eds.) ICALP 2005. LNCS, vol. 3580, pp. 1127–1138. Springer, Heidelberg (2005)
23. Kempe, D., Mahdian, M.: A cascade model for externalities in sponsored search. In: Papadimitriou, C., Zhang, S. (eds.) WINE 2008. LNCS, vol. 5385, pp. 585–596. Springer, Heidelberg (2008)
24. Kleinberg, J.: Cascading behavior in networks: algorithmic and economic issues. Cambridge University Press (2007)
25. Nisan, N., Roughgarden, T., Tardos, É., Vazirani, V.V.: Algorithmic game theory. Cambridge University Press (2007)
26. Sääskilahti, P.: Monopoly pricing of social goods, vol. 3526. University Library of Munich, Germany (2007)