Connecticut College Digital Commons @ Connecticut College

Computer Science Faculty Publications

Computer Science Department

11-1-2011

The Power of Fair Pricing Mechanisms

Christine Chung Connecticut College, cchung@conncoll.edu

Katrina Ligett California Institute of Technology

Aaron Roth University of Pennsylvania

Kirk Pruhs University of Pittsburgh

Follow this and additional works at: http://digitalcommons.conncoll.edu/comscifacpub



Part of the Computer Sciences Commons

Recommended Citation

Chung, C., Ligett, K., Pruhs, K., & Roth, A. (2011). The power of fair pricing mechanisms. Algorithmica 63,3 (July, 2012), 634-644.

This Article is brought to you for free and open access by the Computer Science Department at Digital Commons @ Connecticut College. It has been accepted for inclusion in Computer Science Faculty Publications by an authorized administrator of Digital Commons @ Connecticut College. For more information, please contact bpancier@conncoll.edu.

The views expressed in this paper are solely those of the author.

The Power of Fair Pricing Mechanisms

Keywords

fair pricing mechanisms

Comments

Published in Algorithmica, Nov 2011. (Earlier version appeared in LATIN 2010.)

Published version may be found at http://www.springerlink.com/content/r50l5h1708307598.

DOI 10.1007/s00453-011-9587-1

Noname manuscript No. (will be inserted by the editor)

The Power of Fair Pricing Mechanisms*

Christine Chung \cdot Katrina Ligett \cdot Kirk Pruhs \cdot Aaron Roth

Received: date / Accepted: date

Abstract We explore the revenue capabilities of truthful, monotone ("fair") allocation and pricing functions for resource-constrained auction mechanisms within a general framework that encompasses unlimited supply auctions, knapsack auctions, and auctions with general non-decreasing convex production cost functions. We study and compare the revenue obtainable in each fair pricing scheme to the profit obtained by the ideal omniscient multi-price auction. We show that for capacitated knapsack auctions, no constant pricing scheme can achieve any approximation to the optimal profit, but proportional pricing is as powerful as general monotone pricing. In addition, for auction settings with arbitrary bounded non-decreasing convex production cost functions, we present a proportional pricing mechanism which achieves a poly-logarithmic approximation. Unlike existing approaches, all of our mechanisms have fair (monotone) prices, and all of our competitive analysis is with respect to the optimal profit extraction.

C. Chung

Department of Computer Science, Connecticut College, New London, CT E-mail: cchung@conncoll.edu

K. Ligett

Department of Computing and Mathematical Sciences and Division of the Humanities and Social Sciences, California Institute of Technology, Pasadena, CA E-mail: katrina@cs.cornell.edu

K. Pruhs

Department of Computer Science, University of Pittsburgh, PA E-mail: kirk@cs.pitt.edu

A. Roth

Department of Computer and Information Sciences, University of Pennsylvania, Philadelphia, PA

 $\hbox{E-mail: aaroth@cis.upenn.edu}$

 $^{^{\}star}$ A preliminary version of this paper appeared in the proceedings of LATIN 2010

1 Introduction

Practical experience [9,10,8] demonstrates that any store charging nonmonotone prices (that is, charging some buyer i more than buyer j despite the fact that buyer i receives strictly less of the good than j) risks public outrage and accusations of unfair practices. There are of course very simple auction pricing schemes that are monotone: for example, constant pricing, in which each bidder is quoted the same price regardless of the quantity of the good she receives, and proportional pricing in which each bidder is quoted a price proportional to her demand. Given that fairness may thus in many situations be considered a first-order mechanism design constraint, even at the expense of short-term profit maximization, it is natural to ask, "are clever implementations of these simple monotone pricing schemes capable of maximizing profit?"

We answer this question in the affirmative in a broad class of auctions in which bidders demand different quantities of a given resource, for example, server capacity, bandwidth, or electricity. We consider natural subclasses of this class of auctions: unlimited supply, limited supply ("knapsack auctions"), and a more general setting in which the cost to the mechanism may be some arbitrary non-decreasing convex function of the supply sold. This last model we propose generalizes the first two, and models any way in which the auctioneer may incur decreasing marginal utility as the production of the good being sold increases (for example if increased demand for raw materials increases the producer's per unit cost for these materials).

In general, if h is the highest bidder's true valuation, no truthful auction can acquire the value h from the highest bidder, and at best can hope to compete with $\mathbf{OPT} - h$ (see, for example, [1]). In the unlimited supply setting, the constant pricing mechanism of Goldberg and Hartline [5] is $O(\log n)$ competitive with $\mathbf{OPT} - h$, where n is the number of bidders, and \mathbf{OPT} is the sum of all bidders' valuations, not just the optimal profit obtainable by any constant price mechanism. Our findings are as follows:

- In the limited supply (knapsack) setting, we show that no constant pricing mechanism can achieve an approximation factor of o(S) to $\mathbf{OPT} o(n)h$, where S is knapsack size. However, we give a mechanism that uses proportional pricing that, aside from an extra profit loss of h, achieves an $O(\log S)$ approximation to $\mathbf{OPT} 2h$.
- In a general setting in which the mechanism incurs some non-decreasing convex cost as a function C of the supply it sells, we give a proportional pricing mechanism that achieves (again aside from an extra loss of h) a polylogarithmic approximation to $\mathbf{REV} 3h 2C(S^*)$, where \mathbf{REV} is the revenue obtained and $C(S^*)$ the cost incurred by the auctioneer in the optimal solution maximizing $\mathbf{REV} C(S)$. (Here we assume the cost function is polynomially bounded.)

In each of these settings, there is essentially a log lower bound on the profit competitiveness for any monotone pricing mechanism (implicit in [5]). Additionally, in the generalized auction setting, we show that proportional pricing is strictly weaker than monotone pricing (independent of truthfulness). We give an instance that shows that no proportional pricing scheme can achieve profit within any finite factor of $\mathbf{REV} - O(1)C(S^*)$. This perhaps makes it more surprising that a truthful proportional pricing mechanism can be close to optimally competitive with $\mathbf{OPT} - h$. Overall, our results show that there exist proportional pricing schemes that compete with full profit extraction essentially as effectively as the best possible monotone pricing scheme.

1.1 Related Work

The framework of competitive analysis in the setting of auction design was introduced by Goldberg et al. [6]. In the digital goods setting, where each bidder demands one unit of the resource, and the supply of the resource is unlimited, Goldberg et al. [6] and Fiat et al. [4] give randomized truthful mechanisms that are competitive with the optimal constant price profit. In the "knapsack auction" setting, where each bidder may have a different demand and there is a fixed limited supply, Aggarwal and Hartline [1] give a randomized truthful mechanism that achieves a profit of $\alpha \mathbf{OPT}_{mono} - \gamma h \lg \lg \lg n$, where n is the number of players, \mathbf{OPT}_{mono} is the optimal monotone pricing profit, h is the maximum valuation of any bidder, and α and γ are constants. It is important to note that the mechanisms given in [6] and [1] use random sampling techniques and are not monotone. That is, they can quote customers different prices for identical orders. This can in some sense be justified by the result in [5] that shows that, for either the setting of digital goods or knapsack auctions, no truthful, fair pricing mechanism can be $o(\log n/\log \log n)$ competitive with the optimal constant price profit. That is, there is no mechanism that achieves all the properties of (1) truthfulness, (2) fairness, and (3) constant competitiveness with respect to profit.

Goldberg and Hartline [5] go on to show that if the fairness requirement is relaxed (and the auction is allowed to give non-envy-free outcomes with small probability), auctions competitive with the optimal constant price can be found. Guruswami et al. [7] show that in two auction settings closely related to ours, simply computing fair prices that maximize profit, without the requirement of truthfulness, is APX-hard. Very recently, Babaioff et al. considered an auction setting in which the supply arrives online, and generalized the definition of fairness to their setting, also showing an $\Omega(\log n/\log\log n)$ lower bound for even welfare maximization in this online setting [2].

Intuitively, the papers [6,1,4] (among others) study the profit competitiveness that can be achieved if one gives up on fairness. In contrast, the goal of this work is to understand the pricing techniques that are needed to maximize profit in a variety of settings, if monotonicity and truthfulness are objectives we are unwilling to sacrifice.

1.2 The Problem

In this work, we consider single-round, sealed-bid auctions with a set N = [n] of single-minded bidders. Each bidder i has a public size (or demand) x_i , and a private valuation v_i . We write $\mathbf{x} = (x_1, \dots, x_n)$ and $\mathbf{v} = (v_1, \dots, v_n)$. We will assume that smallest size x_i of any bidder is 1. Let $X = \sum_{i=1}^n x_i$. The demand of a player must be satisfied completely or not at all; we do not allow fractional allocations.

Definition 1 In a single-round, sealed-bid auction, each bidder i submits a bid b_i , which is the most she is willing to pay if she wins. We write $\mathbf{b} = (b_1, \ldots, b_n)$. Given \mathbf{b} and \mathbf{x} , the mechanism returns prices $\mathbf{p} = (p_1, \ldots, p_n)$ and an indicator vector $\mathbf{w} = (w_1, \ldots, w_n)$. If $w_i = 1$, we say player i wins; otherwise, we say she loses. Player i pays p_i if she is a winner and 0 if she is not. The mechanism is valid if and only if every winning bidder has $b_i \geq p_i$, and every losing bidder has $p_i \geq b_i$. The profit achieved by the mechanism depends on the capacity constraints under consideration (discussed below).

Note that the agents are indistinguishable to the auction mechanism, except for their size. We will only present truthful mechanisms, and so throughout the paper, we will assume $\mathbf{b} = \mathbf{v}$.

Definition 2 A deterministic auction is truthful if, for any \mathbf{x} , for all bidders $i \in N$, for any choice of \mathbf{v}_{-i} , bidder i's utility $(v_i - p_i)$ is maximized by bidding her true value v_i . We say that a randomized auction is truthful if it is a probability distribution over deterministic truthful auctions.¹

We require that our pricing schemes be fair, meaning monotone: we cannot charge some player more than another player with higher demand.

Definition 3 A deterministic auction's pricing is *monotone* if for any \mathbf{v} and \mathbf{x} , it assigns prices \mathbf{p} such that $p_i \geq p_j$ whenever $x_i \geq x_j$, for any i and j in N. A randomized auction's pricing is monotone if it is a distribution over deterministic monotone auctions.

One example of of a valid monotone pricing scheme is constant pricing, where each player is offered the same price $p = p_1 = p_2 = \ldots = p_n$, every player i with $v_i > p$ is a winner, and no player i with $v_i < p$ is a winner. Another monotone scheme, proportional pricing, fixes some value c and charges each player i a price $p_i = c \cdot x_i$; every player i with $v_i > p_i$ is a winner, and no player i with $v_i < p_i$ is a winner.

In all cases, we assume that the bidders are trying to maximize utility, that they know the mechanism being used, and that they don't collude. Our goal is to study truthful, monotone pricing mechanisms that maximize our profit, in a variety of capacity constraint settings:

 $^{^{1}}$ Such mechanisms are sometimes known as $\it ex-post$ or $\it universally~truthful.$ This guarantees that bidders do not regret truthful bidding even after the coin tosses of the mechanism are revealed, and is a stronger guarantee than truthfulness in expectation.

- In the *unlimited supply* setting, there is no limit on the total size of the bidders we can accept. A mechanism's profit here is just $\mathbf{p} \cdot \mathbf{w}$.
- In the *knapsack* setting, there is a hard limit S on the total size of the winning bidders. Here, the mechanism's profit is $\mathbf{p} \cdot \mathbf{w}$ if $\mathbf{x} \cdot \mathbf{w} \leq S$, and $-\infty$ otherwise.
- In the general cost setting, there is some non-decreasing convex cost function C on the size of the winning bidder set, and the profit of the auctioneer is the difference between the sum of the prices paid by the winning bidders and the cost of the size of the winning set. Here, we define the mechanism's profit to be $\mathbf{p} \cdot \mathbf{w} C(\mathbf{x} \cdot \mathbf{w})$.

Note that the unlimited supply problem is an instance of the knapsack problem with S>X. The knapsack problem is an instance of the general cost setting, with a cost function that takes value 0 for sizes less than S, and jumps to ∞ at S.

In all three cases, we compare our schemes with the optimal multiple-price omniscient allocation that is not constrained to be truthful nor envy-free. In the unlimited supply case, $\mathbf{OPT} = \sum_{i=1}^{n} v_i$. In the knapsack setting,

$$\mathbf{OPT} = \max_{\mathbf{w} \ | \ \mathbf{x} \cdot \mathbf{w} \le S} \mathbf{v} \cdot \mathbf{w}.$$

In the general cost setting,

$$\mathbf{OPT} = \max_{\mathbf{w}} \left(\mathbf{v} \cdot \mathbf{w} - C \left(\mathbf{x} \cdot \mathbf{w} \right) \right).$$

Setting \mathbf{w}^* to be the set of winners in an optimal general-cost solution, we will write $\mathbf{REV} = \mathbf{v} \cdot \mathbf{w}^*$ for the revenue of the optimal solution. As mentioned above, in general no truthful algorithm can achieve better than $\mathbf{OPT}-h$, where h is the value of the highest bid, so this will be our performance benchmark.

1.3 Unlimited Supply Auctions

In the unlimited supply setting, the auctioneer has an unlimited number of items to sell, at zero marginal cost (equivalently, if each item has some constant marginal cost, we may simply subtract this cost from the valuations of the bidders). Goldberg and Hartline gave a simple randomized mechanism that achieves a $\Theta(\log n)$ approximation to the profit obtained by the best constant price \mathbf{OPT}_c , and showed that this is almost optimal [5]. It is also known that \mathbf{OPT}_c can differ by a $\Theta(\log n)$ factor from \mathbf{OPT} . That is, if the profit obtained by the mechanism below is \mathbf{OPT}_c/α , and $\mathbf{OPT}_c = \mathbf{OPT}/\beta$, it is known that α and β can both take values as large as $\Theta(\log n)$, but no larger. This immediately shows that RANDOMPRICE (given below) is an $O(\log^2 n)$

² Actually, they show something slightly weaker, defining \mathbf{OPT}_c to be the optimal constant price when the mechanism is required to sell at least 2 items.

³ Consider n bidders with valuations v_1, \ldots, v_n with $v_i = 1/i$. **OPT** = H(n), but the best constant price obtains profit $i \cdot v_i = 1$ for all i. [1]

approximation to **OPT**. In fact, we see below that for any instance, $\alpha \cdot \beta = O(\log n)$. In other words, RANDOMPRICE gives a $\Theta(\log n)$ approximation to **OPT**.

RandomPrice(\mathbf{v}, \mathbf{x})

- 1 Choose $i \in \{1, 2, ..., \log n\}$ uniformly at random.
- 2 Let $g = 2^i$. Sell items to the g 1 highest bidders at price v_g (where $v_1 \geq v_2 \geq \ldots \geq v_n$).

Theorem 1 (implicit in [5]) For any set of bidder values, let P be the expected profit obtained by RANDOMPRICE. Let α be such that $P = \mathbf{OPT}_c/\alpha$, and let β be such that $\mathbf{OPT}_c = (\mathbf{OPT} - h)/\beta$. Then $\alpha \cdot \beta = O(\log n)$. Equivalently, $P \geq (\mathbf{OPT} - h)/O(\log n)$.

We note that although constant pricing is sufficient to obtain an $O(\log n)$ approximation, this mechanism is almost optimal over the set of all *monotone* pricing mechanisms. The following lower bound is implicit in the lower bound proved by Goldberg and Hartline in the context of digital goods auctions:

Theorem 2 (Goldberg and Hartline [5]) In the uncapacitated setting, no truthful mechanism using a monotone pricing scheme can achieve profit within a factor of $o(\log n / \log \log n)$ of $\mathbf{OPT} - c \cdot h$ for any constant c.

2 Knapsack Auctions

Knapsack auctions, first studied by Aggarwal and Hartline [1], model auctions for items for which there is a strict limit on supply: we are given a set of bidders with demands and valuations, and can only sell to a set of bidders whose total demand is smaller than our knapsack capacity S.

When supply is unlimited, we have seen that constant pricing is as powerful as monotone pricing in the sense that both can achieve within a $O(\log n)$ factor of $\mathbf{OPT} - h$, but no better. In this section, we show that in the knapsack case, when supply is limited, no valid constant pricing scheme can achieve profit within a factor of o(S) of $\mathbf{OPT} - o(n)h$. However, we show that proportional pricing is as powerful as monotone pricing in the sense that both can achieve $O(\log S)$ competitiveness with $\mathbf{OPT} - h$, but no better. Our result is also optimal over the set of proportional pricing schemes, even those that are not truthful: Aggarwal and Hartline [1] give an example in which the optimal proportional pricing is an $\tilde{\Omega}(\log S)$ factor off from $\mathbf{OPT} - h$.

⁴ Aggarwal and Hartline show that proportional pricing cannot in general approximate monotone pricing within a factor of o(n), and in their lower bound instance, use bidders with exponentially large demand (also showing that proportional pricing cannot in general approximate monotone pricing to within an $\tilde{\Omega}(\log X)$ factor, where X is the total demand of all players). Our result implies that there always exists a proportional pricing that approximates $\mathbf{OPT} - h$ (and not just the optimal monotone pricing) to within an $O(\log S)$ factor.

Theorem 3 In the knapsack setting, no mechanism which uses a valid constant pricing scheme can guarantee an approximation within a o(S) factor of $\mathbf{OPT} - o(n)h$.

Proof Consider an instance with knapsack capacity $S=1/\epsilon, \ 1/\epsilon$ bidders i with size $x_i=1$ and value $v_i=1$, and one bidder i^* of size $x_{i^*}=1/\epsilon$ and value $v_{i^*}=1+\epsilon$. Clearly, $\mathbf{OPT}=1/\epsilon$, which results from selling to each of the $1/\epsilon$ bidders $i\neq i^*$, and $\mathbf{OPT}-o(n)h=1/\epsilon-o(1/\epsilon)=\Theta(1/\epsilon)$. However, no constant price $p\leq 1$ results in a valid allocation, since $\sum_{i:v_i\geq p} x_i=2/\epsilon>S$. Therefore, $\mathbf{OPT}_c=1+\epsilon$, and results from selling only to bidder i^* . Therefore, $(\mathbf{OPT}-o(n)h)/\mathbf{OPT}_c=\Theta((1/(1+\epsilon)\epsilon)=\Theta(S))$.

We can also extend the lower bound for general monotone mechanisms from the unlimited supply setting, simply by choosing the total demand of the lower bound instance to be strictly less than the knapsack size.

Theorem 4 (Goldberg and Hartline [5]) In the knapsack setting, no truthful mechanism using any monotone pricing scheme can guarantee an $o(\log S/\log\log S)$ approximation to $\mathbf{OPT} - c \cdot h$ for any constant c.

Next, we give a truthful mechanism which achieves an $O(\log S)$ approximation to $\mathbf{OPT} - 2h$ by using a valid proportional pricing scheme. This pricing scheme is similar to the "RANDOM single price" algorithm proposed in [3] for an online, combinatorial auction setting. In [3] they show their algorithm is $O(\log(S) + \log(n))$ competitive with \mathbf{OPT} . Below we give a bound for our algorithm only in terms of S.

ProportionalKnapsack($\mathbf{v}, \mathbf{x}, S$)

- 1 Let π be an ordering of bidders in non-increasing density $d_i = v_i/x_i$. Denote by H the largest prefix of bidders in π that is satisfiable with supply S; note $\sum_{i \in H} x_i \leq S$.
 - Let $X_i := \sum_{j=1}^i x_{\pi(j)}; X_0 = 0.$
 - Let g be a function mapping points in the knapsack $x \in [0, S]$ to bidders, in order of π as follows: g(x) = i if $x \in [X_{i-1}, X_i)$
- 2 Choose $s \in \{0, ..., \lfloor \log(S) \rfloor \}$ uniformly at random. Consider the bidder $\pi(i^*) = g(2^s - 1)$ who corresponds to the point $2^s - 1$ in the knapsack; let $d^* = d_{\pi(i^*)}$ be its density.
- 3 Sell items to bidders $\pi(1), \ldots, \pi(i^*-1)$ at prices proportional to d^* .

The above mechanism is reminiscent of a multi-unit Vickrey auction in the sense that given a set amount of supply, the items are sold to the top bidders at a price equal to the top losing bidder's bid. In our mechanism, the amount of supply is randomized and the bids are sorted per unit demand.

Proposition 1 Proportional Knapsack is truthful and produces a valid proportional pricing.

Proof We observe that fixing any realization of the random coin flips, no winning bidder can become a losing bidder by raising his bid, since a bidder can only increase his rank in π by raising his bid. Our pricing scheme is truthful because the price we charge a player is independent of her bid, and all losing players have values below the price they would be offered if they raised their bid to a winning level. Validity follows since winning bidders are charged proportional to a rate that is at most their own density, and losing bidders are charged proportional to a rate that is at least their own density.

Theorem 5 Proportional Knapsack achieves expected profit at least $(\mathbf{OPT} - 2h)/O(\log S) - h$.

Proof Let $\mathbf{OPT}(H)$ refer to the value of the optimal solution if the set of bidders were comprised only of those in the set H. First observe that $\mathbf{OPT}(H) \geq \mathbf{OPT} - h$. To see this, note that \mathbf{OPT} can take value at most the value of the corresponding knapsack problem. By taking the largest density prefix that fits in our knapsack, we are preserving the value of the optimal solution to the fractional knapsack problem, minus at most the value of a single bidder (the first bidder according to π not included in H). Since we wish to be competitive with $\mathbf{OPT} - h$, for the rest of the argument we may restrict our attention to H and assume we are in the unlimited supply setting (since the available supply is larger than the total remaining demand).

We bound $\mathbf{OPT}(H)$ by considering the bidders in decreasing order of density $\pi(1), \ldots, \pi(|H|)$, and bounding the density of the optimal knapsack solution. Let f(x) denote the density of the bidder occupying position x in the knapsack. We have

$$\mathbf{OPT}(H) \le \int_0^S f(x) dx \le \sum_{i=0}^{\lfloor \log(S) \rfloor} f(2^i - 1) 2^i,$$

where the inequality follows since we have ordered the bidders such that their density is non-increasing. Similarly, we may bound the expected profit P obtained by our mechanism:

$$P = \frac{1}{\lfloor \log(S) \rfloor + 1} \left(\sum_{i=0}^{\lfloor \log(S) \rfloor} \left(f(2^i) 2^i - h \right) \right),$$

where we lose the h term since we cannot sell to the bidder from whom we've sampled the sale density d^* . Thus,

$$\mathbf{OPT}(H) \le h + 2((P+h)(|\log(S)| + 1)).$$

Recalling that $\mathbf{OPT}(H) \geq \mathbf{OPT} - h$, we get

$$P \ge \frac{\mathbf{OPT} - 2h}{2(|\log(S)| + 1)} - h.$$

3 General Convex Cost Auctions

In this section, we propose a general setting in which the mechanism incurs a cost, expressed as a function of the amount of supply sold. In the previous section, we showed that in the bounded supply setting, proportional pricing was sufficient to get essentially as good an approximation to $\mathbf{OPT} - h$ as was possible using any monotone pricing scheme. Here, we show that, in general, there is an unboundedly large gap between the profit attainable with proportional pricing and the profit obtainable by monotone pricing, even if we require the monotone pricing scheme to pay a higher cost. Note that this lower bound applies to all pricing schemes (which are not constrained by truthfulness), rather than mechanisms. We will show that surprisingly, if we wish to compete with $\mathbf{OPT} - h$, then proportional pricing is sufficient.

Theorem 6 In the general non-decreasing convex cost setting, for any value d, there exists a set of bidders and a convex cost function such that the optimal profit is obtained by selling supply S^* , yielding profit $\mathbf{OPT} = \mathbf{REV} - C(S^*)$, but no proportional pricing scheme is able to achieve any approximation to $\mathbf{REV} - d \cdot C(S^*)$.

Proof Consider an instance with a quadratic cost function: $C(x) = x^2$. There are two bidders: One bidder has size 1 and value d+2 (this bidder has density d+2). Let k refer to the size of the second bidder, and set his value to (d+2)k+1 (this bidder has higher density, d+2+1/k). The optimal monotone pricing sells only to the first bidder, and gets profit d+1. Note that in this case, even $\mathbf{REV} - d \cdot C(1) = 1 \geq 0$. However, for proportional pricing, it is impossible to sell to the first bidder without selling to the second, since the second bidder is denser. If we sell to the both, however, we get at most profit $(d(k+1)+1-(k)^2)$, which is negative for large enough k. If we sell to only the denser bidder, we get a most profit $(d+2)k+1-k^2$, again negative for large enough k. Thus, the best proportional pricing sells no items, and gets profit 0.

Fiat et al. [4] also consider a setting in the presence of a cost function and give a mechanism that is competitive only with the optimal revenue in some class, minus a multiplicative factor times the cost function, and conjecture that this is necessary. Below, we demonstrate a proportional pricing mechanism that is polylog competitive with $\mathbf{REV} - 3h - (1 + \epsilon)C(S^*)$, for any constant ϵ .

Given a non-decreasing convex cost function C, the profit P the algorithm obtains when assigning prices \mathbf{p} and allocation \mathbf{w} to players (\mathbf{v}, \mathbf{x}) is $P = \mathbf{p} \cdot \mathbf{w} - C(\mathbf{x} \cdot \mathbf{w})$. We assume C is continuous and that C(0) = 0.

We will write $\mathbf{OPT}_S(\mathbf{v}, \mathbf{x})$ for the maximum profit extraction obtainable from \mathbf{v} and \mathbf{x} with a knapsack restriction S in place of the cost function. Let S^* be the total size of the set of winning bidders under the optimal (non-truthful) allocation:

$$S^* = \mathbf{x} \cdot \mathbf{w}^*$$
, where $\mathbf{w}^* = \underset{\mathbf{w}}{\operatorname{argmax}} (\mathbf{v} \cdot \mathbf{w} - C(\mathbf{x} \cdot \mathbf{w}))$.

Our algorithm first attempts to guess S^* . We then seek to estimate a value for each bidder that reflects the amount of her valuation adjusted for the cost the auctioneer would incur for servicing her. We compute such "adjusted" valuations based on the bidders' stated valuations minus their contribution to the cost of the estimated solution size, and run our proportional-pricing mechanism on these estimated values. We then adjust the resulting prices to re-incorporate their share of the cost, and output this with the resulting allocation. In what follows, we assume that $h \geq 1$; if this is not the case, the algorithm can be easily adapted to begin guessing the value of c at h instead of 1, losing an additional $\log 1/h$ factor.

GENERALAUCTION($\mathbf{v}, \mathbf{x}, C$)

- Select cost c at random from among $\{1, 2, 2^2, \dots, 2^{\lceil \log C(X) \rceil}\}$.
- Set S to be the largest value such that C(S) = c, or ∞ if no such value exists.
- $v'_i := v_i \frac{C(S)x_i}{S}$ and $\mathbf{v}' := (v'_1, \dots, v'_n)$ $(\mathbf{p}', \mathbf{w}') := \text{ProportionalKnapsack}(\mathbf{v}', \mathbf{x}, S).$
- $p_i := p_i' + \frac{C(S)x_i}{S}$ Return $(\mathbf{p}, \mathbf{w}')$

Lemma 1 Suppose that $2C(S^*) \geq C(S) \geq C(S^*) \geq 1$. The optimal profit obtainable in an S-capacitated knapsack, given values \mathbf{v}' , is at least the optimal revenue minus twice the cost at that supply, on the original instance: $\mathbf{OPT}_S(\mathbf{v}', \mathbf{x}) \geq \mathbf{OPT}_{S^*}(\mathbf{v}, \mathbf{x}) - 2C(S^*).$

Proof We observe $\mathbf{OPT}_S(\mathbf{v}', \mathbf{x}) \geq \mathbf{OPT}_S(\mathbf{v}, \mathbf{x}) - C(S)$, since at worst, the optimal knapsack solution given \mathbf{v}' selects the exact same winners as the optimal knapsack solution given v. Now note that $S^* \leq S$ and $2C(S^*) \geq C(S)$, and so $\mathbf{OPT}_S(\mathbf{v}, \mathbf{x}) - C(S) \ge \mathbf{OPT}_{S^*}(\mathbf{v}, \mathbf{x}) - 2C(S^*)$.

Theorem 7 The General Auction algorithm obtains expected profit at least

$$\frac{\mathbf{REV} - 2C(S^*) - 3h}{O(\log(X)\log(C(X)))} - \frac{h}{\log(C(X))}.$$

Note that the denominator is $O(\log^2(X))$ when C is polynomially bounded.⁵

Proof Suppose that $2C(S^*) \ge C(S) \ge C(S^*) \ge 1$. Note that by the definitions of REV, S^* , and $\mathbf{OPT}_{S^*}(\mathbf{v}, \mathbf{x})$, we have $\mathbf{REV} \leq \mathbf{OPT}_{S^*}(\mathbf{v}, \mathbf{x})$. Hence, the optimal profit obtainable in an S-capacitated knapsack under $(\mathbf{v}', \mathbf{x})$, as shown above, is

$$\mathbf{OPT}_S(\mathbf{v}', \mathbf{x}) > \mathbf{REV} - 2C(S^*).$$

^a Note that if h > 1, we lose at most an additive h by assuming S^* is such that c > 1. Also note that it is possible to get a slightly stronger result by excluding values of S such that $S > \sum_{i=1}^{n} v_i$.

⁵ This can be improved to $(\mathbf{REV} - (1 + \epsilon)C(S^*) - 3h)/(O(\log(X)\log(C(X))))$ – $h/\log(C(X))$ for arbitrary constant ϵ simply by selecting c from among $\{1,(1+\epsilon),(1+\epsilon)\}$ ϵ)²,...}.

Then by the approximation ratio of the knapsack algorithm, Proportion-ALKNAPSACK returns a solution of value

$$\frac{\mathbf{OPT}_S(\mathbf{v}', \mathbf{x}) - 2h}{2\lfloor \log(S) \rfloor} - h \ge \frac{\mathbf{REV} - 2C(S^*) - 2h}{2\lfloor \log(S) \rfloor} - h.$$

If $C(S^*) < 1$, this becomes at worst

$$\frac{\mathbf{REV} - 2C(S^*) - 3h}{2|\log(S)|} - h$$

for $h \geq 1$. The additional profit obtained by prices \mathbf{p} over prices \mathbf{p}' in allocation \mathbf{w} is $\frac{C(S)S'}{S}$, where S' is the size of the solution selected by Proportional K-Napsack. The cost imposed by the cost function is C(S'). Thus, in this case, the profit General Auction obtains is at least

$$\frac{\mathbf{REV} - 2C(S^*) - 3h}{2|\log(S)|} - h + \left(C(S)\frac{S'}{S} - C(S')\right).$$

Since the cost function C is non-decreasing and convex, this second term is non-negative. Since $2C(S^*) \geq C(S) \geq C(S^*)$ holds with probability $O(1/\log(C(X)))$, and $X = \sum_{i=1}^n x_i$ is an upper bound on S, this completes the proof.

Proposition 2 General Auction is truthful.

Proof This is immediate, since it is a distribution over truthful mechanisms. (Specifically, it is a distribution over instances of Proportional Knapsack in which the prices have been modified by a bid-independent function, which preserves truthfulness).

Proposition 3 General Auction is a valid mechanism.

Proof Suppose player i is a winner. Then, by the validity of PROPORTION-ALKNAPSACK, $v'_i \geq p'_i$. Thus,

$$v_i = v'_i + \frac{C(S)x_i}{S} \ge p'_i + \frac{C(S)x_i}{S} = p_i.$$

Now, suppose player i loses. Then, by the validity of the knapsack mechanism, $p'_i \geq v'_i$. Thus,

$$p_i = p'_i + \frac{C(S)x_i}{S} \ge v'_i + \frac{C(S)x_i}{S} = v_i.$$

Proposition 4 GENERALAUCTION produces proportional pricing.

Proof Proportional Knapsack produces a proportional pricing scheme, and the prices returned by General Auction increase the proportional factor by $\frac{C(S)}{S}$.

Acknowlegments. We would like to thank several anonymous referees for thoughtful and helpful comments on earlier versions of this work.

References

- 1. Aggarwal G, Hartline J (2006) Knapsack Auctions. In: Proceedings of the Symposium on Discrete Algorithms
- 2. Babaioff M, Blumrosen L, Roth A (2009) Auctions with Online Supply
- 3. Balcan MF, Blum A, Mansour Y (2008) Item pricing for revenue maximization. In: EC '08: Proceedings of the 9th ACM conference on Electronic commerce
- Fiat A, Goldberg A, Hartline J, Karlin A (2002) Competitive generalized auctions. In: Proceedings of the thiry-fourth annual ACM symposium on Theory of Computing, ACM New York, NY, USA, pp 72–81
- Goldberg A, Hartline J (2003) Envy-free auctions for digital goods. In: Proceedings of the 4th ACM conference on Electronic commerce, ACM New York, NY, USA, pp 29–35
- Goldberg A, Hartline J, Karlin A, Saks M, Wright A (2006) Competitive auctions. Games and Economic Behavior 55(2):242–269
- Guruswami V, Hartline JD, Karlin AR, Kempe D, Kenyon C, McSherry F (2005) On profit-maximizing envy-free pricing. In: SODA '05: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms, Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, pp 1164–1173
- 8. Ramasastry A (2005) Web sites change prices based on customers' habits. Http://www.cnn.com/2005/LAW/06/24/ramasastry.website.prices/
- 9. Streitfeld D (2000) On the web, price tags blur: What you pay could depend on who you are. $\label{eq:http://www.washingtonpost.com/ac2/wp-dyn/A15159-2000Sep25}$
- 10. Wolverton T (2000) Amazon backs away from test prices. Http://news.cnet.com/2100-1017-245631.html