

Knapsack Auctions

Gagan Aggarwal*

Jason D. Hartline†

Abstract

We consider a game theoretic knapsack problem that has application to auctions for selling advertisements on Internet search engines. Consider n agents each wishing to place an object in the knapsack. Each agent has a *private valuation* for having their object in the knapsack and each object has a *publicly known* size. For this setting, we consider the design of auctions in which agents have an incentive to truthfully reveal their private valuations. Following the framework of Goldberg et al. [10], we look to design an auction that obtains a constant fraction of the profit obtainable by a natural optimal pricing algorithm that knows the agents' valuations and object sizes.

We give an auction that obtains a constant factor approximation in the non-trivial special case where the knapsack has unlimited capacity. We then reduce the limited capacity version of the problem to the unlimited capacity version via an approximately *efficient* auction (i.e., one that maximizes the social welfare). This reduction follows from generalizable principles.

1 Introduction

The ability to perform targeted advertising on Internet search engines has recently spurred a flurry of activity in studying pricing mechanisms for the problem. Given the dynamic nature of the problem, auctions are a natural solution that leading search engines such as Google, Yahoo Search, and MSN Search have adopted. The problem of designing a good auction for this setting is multi-faceted. Advertisers may have combinatorial preferences: it is normal to show multiple advertisements per search and advertisements may be shown for multiple combinations of keywords. Advertisers may have budgets limiting the amount they can spend on advertising [4]. In addition, the searches arrive over time which gives the problem an online matching flavor [13]. Finally, this is a game-theoretic problem and any solution must take into account the issue of *incentives*. Each of these aspects of the problem in themselves represents a significant challenge to auction design. In this paper we focus on a particular, yet

fundamental, aspect of this problem.

Consider the following abstract *private value* version of the *knapsack problem*. A profit-maximizing auctioneer is auctioning off space in a knapsack of fixed capacity C . Each agent would like to place exactly one object in the knapsack. Agent i values the placement of her object in the knapsack at v_i . A priori, the valuations are the *private data* of each respective agent. Each object takes up a certain amount of space in the knapsack, e.g., agent i 's object takes space c_i , and these sizes are publicly known. Thus, the c_i s are public values while the v_i s are private values.¹

The *knapsack auction problem* models several interesting applications. For example, consider running a single auction to sell advertising space on a web page over the course of a day. Suppose statistical information is available for each advertiser as to how many showings (a.k.a., *impressions*) are necessary for to result in a user *click-through* and as well how many times the web page itself will be viewed in a day. The number of impressions necessary to generate a click-through corresponds to the c_i s and the number of total views corresponds to the capacity of the knapsack, C . Assume that each advertiser wants exactly one click-through and we have an instance of the (fractional)² knapsack auction problem.

In this work we assume that agents are indistinguishable except for the fact that each agent's demand (i.e., the size of her object) is publicly known. This distinguishability allows an auction to charge agents with different sized objects different prices and raises the question of how prices should be related to the agent demands. We will consider three ways in this paper. The first, essentially ignoring the demands, is to offer a single price per agent. In this pricing model, we consciously decide not to discriminate against agents based on their demands. We refer to such pricing schemes as *constant pricing* as valid pricing functions (from sizes to prices) are only constant functions. Another natural candidate pricing strategy is to charge a single price per unit capacity. Thus, agents demanding more of the capacity pay proportionally more. We refer to such pricing schemes as *proportional pricing*. Finally, the most general pricing

*Google Inc., Mountain View, CA. and Computer Science Dept., Stanford University, Stanford, CA 94305. This work was done while author was visiting Microsoft Research. gagan@cs.stanford.edu

†Microsoft Research – Silicon Valley, Mountain View, CA 94043. hartline@microsoft.com

¹Notice that were the v_i s public as well, this problem would be the standard weighted knapsack problem – the auctioneer could charge each agent her full valuation if her object is selected.

²In this paper we restrict our attention to the integral knapsack problem. Naturally, all of our positive results apply to the fractional case as well, and some can be easily improved.

strategy we will consider follows the least restrictive natural assumption we could place on prices: that agents desiring more capacity not pay less than those desiring less. We refer to this as *monotone pricing* since the valid pricing functions are the class of monotone non-decreasing functions of object size. Clearly, both constant pricing and proportional pricing are special cases of monotone pricing. As such we will focus throughout the paper on monotone pricing noting that a constant approximation to it also gives a constant approximation to constant and proportional pricing.

Following Goldberg et al. [10], we consider analyzing knapsack auctions in the framework of *competitive analysis* by comparing the performance of the auction to the profit of an optimal pricing function. Accordingly, we define OPT to be the profit obtained by the best monotone pricing function for agents’ actual valuations when we assume that an agent pays the offered price if it is at most their valuation. Because it is not possible to obtain a constant fraction of OPT in the worst case, we look to design auctions that obtain at least a constant fraction of OPT less a small additive loss term, i.e., $\alpha \text{OPT} - \lambda h$ (where h is an upper bound on the highest agent’s valuation). Ideally both α and λ would be constants. In this paper, we achieve a constant α and $\lambda \in O(\log \log \log n)$, where n is the number of agents.

Through the study of the knapsack auction problem, we wish to develop a better understanding of how to design mechanisms for profit maximization when there are non-trivial constraints on the allocation. The non-trivial constraint we face is that objects selected for inclusion in the knapsack must all fit. A similar direction was attempted by Fiat et al. in [6] for the *multicast pricing* problem. They exploit inherent market segmentation in the problem definition to reduce the problem from a *private-value* optimization problem (mechanism design) to a *public-value* optimization problem (algorithm design). In the knapsack auction problem, however, there is no obvious market segmentation and indeed, figuring out how to segment the agents into markets in a truthful manner constitutes a key portion of our solution. To the best of our knowledge, this work represents the first solution to a non-trivial private-value optimization problem when market segments are not given in advance.

Along this vein, we outline a general approach for dealing with non-trivial optimization problems. The first step of this approach is to solve the *unlimited-supply* special case of the problem. For example, in the unlimited supply special case of the knapsack problem, the capacity of the knapsack, C , exceeds the total demand, $\sum_i c_i$. As this special case is not constrained in what objects it may select, it allows us to focus on the problem of how prices offered to agents relate to their demands. The second step of this approach is to reduce the *limited-supply* (a.k.a., general) version of the problem to the unlimited-supply version. This approach works in general for “monotone” optimization problems, where if an

allocation is feasible, then any subset of the allocation is also feasible. The reduction works by (a) selecting a set of agents that can all be allocated together, and (b) running the unlimited capacity solution on this selected set. This must be done carefully so as to preserve game theoretic properties and guarantee good performance.

The rest of the paper is organized as follows. In Section 2 we define the knapsack auction problem formally and discuss related work. In Section 3 we present a comparison of the different pricing rules that we consider in this paper. In Section 4, we discuss the algorithmic complexity of computing the optimal pricing function from a class when the agent valuations are public knowledge. We present an approximately optimal auction for the knapsack auction problem in Section 5. For this, we first show (in Section 5.1) how to reduce the limited-supply auction problem to the unlimited-supply auction problem with a small loss in approximation factor. Then in Section 5.2, we give an unlimited-supply auction that achieves a constant fraction of the benchmark revenue (with a small additive loss).

2 Preliminaries

In this section, we formally define the knapsack auction problem, discuss game theoretic constraints, and review prior results that we will be building on.

The Problem. Consider the following setting. There is a set of n agents $N = \{1, \dots, n\}$, each of whom has an object. Let c_i represent the publicly-known size of agent i ’s object. Each of these agents desires to have her object placed in a knapsack with total capacity C . Our goal is to design a single-round, sealed-bid auction for this setting. In this auction, winning agents have their objects placed into the knapsack and losing agents do not. Let v_i denote agent i ’s *valuation* for having her object placed in the knapsack. This valuation represents the benefit to the agent for winning. We assume that the agents attempt to maximize their *utility*, measured as the difference between their valuation and their payment (zero, if their item is not selected). We assume that all the agents’ valuations fall within a known range $[1, h]$.

We denote by $\mathbf{b} = (b_1, \dots, b_n)$ the vector of bids submitted by the agents and by $\mathbf{c} = (c_1, \dots, c_n)$ the vector of publicly known object sizes. We assume for convenience that the agents are indexed by size, i.e., $c_i \geq c_{i+1}$. Following [3], we sometimes refer to these object sizes as *attributes*.

Incentive Properties. We adopt the game-theoretic solution concept of *truthful mechanism design*. In a truthful mechanism it is a dominant strategy for every agent to report their true valuation as their bid, i.e., $b_i = v_i$. We will employ the following theorem to argue that our auctions are truthful. To obtain a unique payment rule, we make the standard assumptions of *no-positive-transfers*, that no agents are

“paid to play”; and *individual rationality*, that agents are not changed more than their bids.

THEOREM 2.1. [15] *A deterministic mechanism is truthful if and only if fixing the bids of all other agents, the selection rule is monotone in agent i 's bid, i.e., raising her winning bid also results in her winning; and i 's payment is her minimum winning bid value.*

We define a randomized mechanism to be truthful if it is a randomization over truthful deterministic mechanisms. In a randomized auction the auction's allocation, prices, and profit are all random variables.

Analysis Framework. Our goal is to design auctions that perform well in the worst case over all possible inputs. Following the competitive framework in [10], we analyze the performance of an auction by comparing it to a meaningful benchmark. For the case that the agents are indistinguishable, [10] uses the benchmark profit of the optimal constant price per agent. This seems reasonable when the agents are indistinguishable and [9] shows that indeed no truthful mechanism in a large natural class of auction mechanisms can outperform this benchmark. For knapsack auctions, the least restrictive natural pricing rule is to assume that agents with larger objects pay no less than those with smaller objects. This is monotone pricing and we let OPT denote the profit obtained by the optimal monotone pricing function. It is not possible to design a truthful auction that approximates OPT in all cases [9]. Instead, our goal is to design a knapsack auction that obtains a profit of $OPT / \beta - \gamma h$ for small β and γ (where h is an upper bound on the highest valuation).³

Related Work. The problem of designing profit-maximizing auctions for selling advertisements on Internet search engines has been a subject of recent interest. Mehta et al. [13] consider an online matching problem that ignores the game-theoretic issues of mechanism design and instead focus on the algorithmic problem of matching advertisers (with known valuations and budgets) to key word searches that arrive over time. Borgs et al. [4] study the (offline) mechanism design problem of selling multiple identical units when the agents are interested in obtaining as many units as possible while keeping their total payments within budget. In their work, both the valuation and the budget of an advertiser are considered private values. They give near optimal auctions under certain assumptions. Abrams has recently improved on both the severity of the assumptions and the approximation factor for this problem [1].

The private value knapsack problem was studied by Mu'alam and Nisan [14] with the objective of maximizing

social welfare, i.e., the sum of the valuations of selected objects, rather than the auctioneer's profit. They give a truthful auction that approximates their objective. Our result is based on a mechanism that is almost identical to theirs.

The limited supply digital good auction of [10, 9] solves the profit-maximizing knapsack auction problem for the special case where the object sizes are identical. We generalize their approach of solving general auction problems by reducing them to the unlimited supply special case.

A problem closely related to the unlimited supply version of the knapsack auction problem is the *attribute auction problem*. It was introduced by Blum and Hartline [3], who demonstrated that it is possible to get a higher profit than is possible with a single sale price when the auctioneer is able to distinguish between agents based on their differing, publicly observable, attribute values. They gave a solution for the unlimited-supply, single-dimensional-attribute auction problem, and analyzed its performance by comparing its profit to that of the optimal piecewise-constant (not necessarily monotone) pricing rule. We will henceforth refer to this attribute auction as the *general attribute auction*. We will make use of the following result.

THEOREM 2.2. [3] *The general attribute auction obtains a profit of at least $OPT_m / 16 - mh/2$ simultaneously for all m , where OPT_m defined as the total profit of the best piecewise-constant pricing function with m pieces.*

These results have been recently generalized to more general pricing functions over more general attribute spaces [2].

3 Pricing Rules

A pricing rule is a function $\pi(\cdot)$ which specifies the price that should be paid in order to have an object of a given size placed in the knapsack. For a given instance of the knapsack problem, we call a pricing rule and a selection of objects to be contained in the knapsack *valid* if and only if

1. the valuation of each object in the knapsack is at least the price set for objects of that size,
2. the valuation of each object not in the knapsack is at most the price set for objects of that size, and
3. the sum of the sizes of the selected objects is at most the knapsack capacity.

The *payoff* of a valid pricing rule, π , and selection, H , is simply the sum, over all objects in H , of the price π assigns the object. Given a class of pricing functions, the optimal pricing function from the class is the one that maximizes its total payoff. We note that the validity conditions can be viewed as a requirement that the pricing rule and selection be *envy-free* [12] in the sense that each agent prefers her outcome to the outcome of any other agent, or equivalently that none of the agents is envious of the outcome of another.

³By necessity [7], the auctions we obtain are pseudo-approximations. We do not require them to have monotone prices nor charge the same price for objects of the same size.

As mentioned in the introduction, we will be primarily interested in three classes of pricing functions: *constant pricing*, *proportional pricing*, and *monotone pricing*. The class of constant pricing functions are those that give a single sale price irrespective of an object’s size (i.e., constant functions). The class of proportional pricing functions contains those that charge a single price per unit size. Finally, our most general class of pricing functions is that of monotone pricing where the price is required to be a non-decreasing function of object size. One can view the restriction to monotone prices as an additional requirement for envy-freedom, since without monotone prices, a small object would be envious of a larger object being placed into the knapsack at a smaller price.

We now consider the worst case relationship between optimal constant pricing, optimal proportional pricing, and optimal monotone pricing. As constant and proportional pricing are a special case of monotone pricing, it is clear that the profit of the optimal monotone pricing is better than that of both the optimal constant and the optimal proportional pricing. We now get bounds on how much worse constant and proportional pricing can be.

LEMMA 3.1. *Constant pricing can be a $\log n$ factor worse than monotone and proportional pricing and this is tight.*

Proof. Consider n objects with $v_i = c_i = 1/i$ and $C = \infty$. Optimal monotone and proportional pricing use the pricing rule $\pi(c) = c$ for a total payoff of $\sum_i 1/i = \Theta(\log n)$. Constant pricing on the other hand must choose a single price $\pi(c) = p$. Since the number of objects with value at least p is at most $1/p$, the total payoff of constant pricing is at most 1. This provides the desired logarithmic factor separation. Tightness follows from the following observation: for any v_1, \dots, v_n reordered such that $v_i \leq v_{i+1}$, the payoff of the optimal constant price is given by $\max_i iv_i$. Since $v_i \leq \frac{\max_i iv_i}{i}$ for all i , the maximum possible payoff $\sum v_i \leq \log n \cdot \max_i iv_i$. \square

LEMMA 3.2. *Proportional pricing can be a linear factor worse than monotone and constant pricing and this is tight.*

Proof. Take $C = \infty$, $v_i = 1$, and $c_i = n^{-(i-1)}$ for $1 \leq i \leq n$. Optimal monotone and constant pricing set $\pi(c) = 1$ and obtain a payoff of n . The optimal proportional pricing function uses the pricing function $\pi(c) = c$. This gives a payoff of $\sum_i n^{-(i-1)} = O(1)$. To see tightness, note that proportional pricing can always obtain a payoff of at least $\max_i v_i$, which is at least $(1/n)^{th}$ of the optimal monotone and constant pricing payoffs. \square

4 Pricing Algorithms

In this section, we explore the non-game-theoretic problem of designing good *knapsack pricing algorithms*. For a particular instance of the knapsack problem, we would like

an algorithm for efficiently computing the optimal pricing function from a class. Note that this knapsack pricing problem differs from the conventional knapsack problem in that the payoff earned from placing an object in the knapsack is not the object’s value, but instead a price that is a function of the size of the object.

LEMMA 4.1. *Optimal constant pricing is in P.*

LEMMA 4.2. *Optimal proportional pricing is NP-hard.*

LEMMA 4.3. *Optimal monotone pricing is NP-hard.*

Lemma 4.1 follows from the following simple procedure for computing the profit for any constant price p . First add all objects with value strictly greater than p ; these must be in any knapsack when price p is offered. If this exceeds the capacity of the knapsack, then p is an infeasible offer price. Otherwise, add the objects with value equal to p to the knapsack from smallest to largest. This maximizes the number of objects in the knapsack given the offer price of p . Given this procedure, we can find the optimal constant offer price by searching through the n object values, v_1, \dots, v_n , as possible offer prices.

Lemmas 4.2 and 4.3 follow from the hardness of the subset-sum problem by the following simple reduction. Given an instance of the subset-sum problem with object i having size \hat{c}_i , create an instance of the knapsack pricing problem with the same number of objects, and set $v_i = c_i = \hat{c}_i$ for all i . Set the knapsack capacity, C , equal to the desired subset sum, S . The optimal pricing function is simply $\pi(x) = x$ (which is proportional and therefore monotone); however, the algorithm still has the discretion to “break ties” by choosing which subset of the objects to put in the knapsack (the validity conditions except for Condition 3 are satisfied for any subset of the objects). The reduction is complete when we observe that there exists a subset of objects with sum S if and only if the optimal profit for this knapsack pricing instance is S .

4.1 Pricing Algorithms for Unlimited Supply. An interesting special case of the knapsack pricing problem, referred to as the *unlimited-supply* problem, is the case where C is effectively infinite, i.e. $C \geq \sum_i c_i$. It turns out that the unlimited-supply cases of the knapsack pricing problem are relatively easy to solve in polynomial time. Contrast this with the envy-free pricing problems from [12], where even simple special cases of the unlimited-supply pricing problems considered are APX-hard, i.e., unless $P = NP$, there is no polynomial time approximation scheme (PTAS).

LEMMA 4.4. *Unlimited-supply proportional pricing is in P.*

Proof. To compute the optimal proportional pricing for objects $1, \dots, n$ with object i with value v_i and size c_i , we compute each object’s value per unit size, $d_i = v_i/c_i$. For each

price rate d_i , the payoff of the algorithm is $d_i \times \sum_{j: d_j \geq d_i} c_j$. Payoffs for all values of d_i can easily be computed in $O(n \log n)$ time by first sorting the objects by d_i . \square

LEMMA 4.5. *Unlimited-supply monotone pricing is in P.*

Proof. The proof of this lemma is from the correctness of the following dynamic programming algorithm. Intuitively, the table entry $T[i, p]$ corresponds to the optimal payoff from the smallest i objects using monotone prices less than and including price p . \square

DEFINITION 4.1. (USMP) *The Unlimited-Supply Monotone Pricing (USMP) algorithm works as follows.*

1. Order objects by increasing size, i.e., $c_1 \leq c_2 \leq \dots \leq c_n$.
2. Solve dynamic program for all $p \in \{v_1, \dots, v_n\}$.

$$T[0, p] = 0$$

$$T[i, p] = \text{profit}(v_i, p) + \max_{q \in \{v_1, \dots, v_n\}; q \leq p} T[i-1, q]$$

with $\text{profit}(v_i, p) = p$ if $v_i \geq p$ and 0 otherwise.

3. Output the pricing corresponding to the computation of $\max_{p \in \{v_1, \dots, v_n\}} T[n, p]$.

4.2 Limited-Supply Approximation via Reduction to Unlimited Supply. We now show how to approximate the optimal monotone knapsack pricing in the general case by using an optimal or approximate pricing algorithm for the unlimited-supply case. Similar results can be obtained for the proportional pricing variant of the problem. Consider the following technique for composing two pricing algorithms, \mathcal{A}_1 and \mathcal{A}_2 :

DEFINITION 4.2. (PRICING ALGORITHM COMPOSITION) *Given two pricing algorithms \mathcal{A}_1 and \mathcal{A}_2 , we define the composite algorithm $\mathcal{A}_1 \circ \mathcal{A}_2$ as:*

1. Run \mathcal{A}_1 to obtain pricing function $\pi_1(\cdot)$ and let H be the set of winners.
2. Run \mathcal{A}_2 on H to obtain pricing function $\pi_2(\cdot)$.
3. Output $\pi(x) = \max(\pi_1(x), \pi_2(x))$.

If algorithm \mathcal{A}_1 produces a set of winners H that is feasible (for our knapsack problem, feasibility means that all objects in H can fit in the knapsack simultaneously), then we can choose \mathcal{A}_2 as the optimal unlimited-supply monotone pricing algorithm. Since the feasible solutions to the knapsack problem are *closed under inclusion*, meaning that any subset of a feasible set is also feasible, \mathcal{A}_2 will always produce a feasible set. All we need to argue then is that the composite pricing algorithm yields a monotone pricing and that it performs well. The former is clear.

LEMMA 4.6. *If \mathcal{A}_1 and \mathcal{A}_2 are monotone pricing algorithms, then $\mathcal{A}_1 \circ \mathcal{A}_2$ also yields monotone pricing.*

DEFINITION 4.3. (PERFORMANCE PRESERVATION) *An algorithm \mathcal{A} approximately preserves a performance benchmark, OPT, if when given N , \mathcal{A} selects objects $H \subset N$ that satisfy $\text{OPT}(H) \geq \text{OPT}(N)/\beta - \gamma h$ for small constants β and γ , where h is the maximum valuation for any object.*

We next discuss monotone pricing algorithms that approximately preserve the performance of the optimal monotone pricing and produce a feasible selection of objects. Recall the standard weighted knapsack problem: given object values v_1, \dots, v_n , object sizes c_1, \dots, c_n , and knapsack capacity C , find the set of objects, H , with maximum total value that simultaneously fit in the knapsack. We present a pricing algorithm based on a natural greedy approximation algorithm for this standard knapsack problem.

DEFINITION 4.4. (APPROX. KNAPSACK ALGORITHM, AK)

1. Ignore large objects with $c_i > C/2$.
2. List the remaining objects in the order of decreasing value-per-unit-size, $d_i = v_i/c_i$.
3. Select the largest prefix of the object list that fits in the knapsack as the winner set H .
4. Let d^* be the largest value-per-unit-size of the losers. Output $\pi(x) = d^*x$ for $x \leq C/2$ and ∞ otherwise.

LEMMA 4.7. *AK approximately preserves the optimal monotone pricing. On input N , AK selects objects H satisfying $\text{OPT}(H) \geq \text{OPT}(N)/3 - h$.*

Proof. Let $N' \subset N$ be the objects with size at most $C/2$. At most one object from the set $N \setminus N'$ can fit in the knapsack. Thus, the algorithm can restrict its attention to the set N' without losing more than an additive term of h .

If all of N' fits into the knapsack then the theorem follows. Otherwise, the objects in the winner set H fill at least half of the knapsack. This is because there is some object in N' that could not fit into the remaining space of the knapsack, and the objects in N' have size at most $C/2$. Therefore, $\pi(x)$ is a monotone pricing rule that obtains a payoff of at least $d^*C/2$; this is because the value-per-unit-size of agents in H is at least d^* . Thus, $\text{OPT}(H) \geq d^*C/2$.

Let $L = N' \setminus H$ be the objects not included in the knapsack (which all have value-per-unit-size at most d^*). Clearly, $\text{OPT}(L) \leq d^*C \leq 2 \text{OPT}(H)$. Therefore, $3 \text{OPT}(H) \geq \text{OPT}(H) + \text{OPT}(L) \geq \text{OPT}(N')$. \square

THEOREM 4.1. *The algorithm composition of the approximate knapsack algorithm, AK, and the unlimited supply monotone pricing algorithm, USMP, achieves a payoff of at least $\text{OPT}/3 - h$.*

5 Approximately Optimal Knapsack Auctions

In this section, we extend the technique of composing pricing algorithms to mechanism design problems. These techniques suggest a general procedure for reducing limited-supply (or, constrained) problems to unlimited-supply (or, unconstrained) mechanism design problems.

5.1 Reduction via Composition. Consider any constrained profit maximization problem in a private-value setting, e.g., the single-parameter agent settings of [6, 8]. One can think of the unlimited-supply case as that where all outcomes are feasible; whereas the limited-supply case is constrained to produce some outcome in a restricted feasible set. In the case where the set of feasible outcomes (sets of agents) is *closed under inclusion*, meaning that all subsets of a feasible set are also feasible, the following general approach can be attempted: first find a good feasible set, then run an unlimited-supply auction on it. Below we formalize the game-theoretic issues that arise with this approach.

DEFINITION 5.1. (MECHANISM COMPOSITION) *Given two mechanisms \mathcal{M}_1 and \mathcal{M}_2 , we define the composite mechanism $\mathcal{M}_1 \circ \mathcal{M}_2$ as:*

1. Simulate \mathcal{M}_1 and let H be the set of winners.
2. Simulate \mathcal{M}_2 on the set H .
3. Offer a price to each winner of Step 2 that is the maximum of the price she is offered by \mathcal{M}_1 and \mathcal{M}_2 .

We will be looking to use this composition technique with a mechanism \mathcal{M}_1 that always outputs a set of winners for which all subsets are feasible, and a mechanism \mathcal{M}_2 which takes such a set of agents (i.e., a set with respect to which the mechanism effectively has unlimited supply) and computes offer prices with the goal of maximizing profit.

There are four potential issues when using this approach: correctness, truthfulness, and performance, and polynomial time computability.

Correctness. The technique is correct if it produces a feasible outcome. A mechanism for the unlimited-supply case, \mathcal{M}_2 , could output any subset of H as its final outcome; this immediately imposes the constraint that the set of feasible outcomes must be closed under inclusion. This condition, which is satisfied by the knapsack problem, is also sufficient as asserted by the following lemma.

LEMMA 5.1. *If the set of feasible outcomes are closed under inclusion and \mathcal{M}_1 produces a feasible outcome then $\mathcal{M}_1 \circ \mathcal{M}_2$ produces a feasible outcome.*

Truthfulness. We would also like the construction to yield a truthful mechanism. Unfortunately, the condition that \mathcal{M}_1

and \mathcal{M}_2 both be truthful is not enough to guarantee the truthfulness of the composite mechanism. In this discussion, we refer to the agents in H (Definition 5.1) as the *survivors* and the prices offered by Step 1 as the *survival prices*. Note that if \mathcal{M}_1 is truthful, the survival price of an agent does not depend on her bid. However, even for a truthful mechanism \mathcal{M}_1 , the winner set H could be a function of some survivor's bid value. In this case, such an agent could manipulate her bid to change the set H which could affect the price she is offered by \mathcal{M}_2 . Thus, we must require that \mathcal{M}_1 satisfy a stronger property than truthfulness.

DEFINITION 5.2. (COMPOSABILITY) *A mechanism is composable if it is truthful and the survivor set produced does not change as a winning agent's bid varies above her survival price.*

It turns out that the approximate knapsack algorithm, AK, which computes a monotone pricing and a selection of objects is not just a pricing algorithm. We can also consider running it as an auction. In fact, it is very closely related to an auction that Mu'alem and Nisan [14] show is both truthful and approximately optimal (for the objective of maximizing the sum of the valuations of selected objects).

LEMMA 5.2. *AK is composable.*

Proof. First we show truthfulness then we show composability. Given Theorem 2.1, we must only observe that the selection rule is monotone and agent's payment is the minimum bid necessary to be selected. It is easy to see that if an agent is selected with a particular bid and they raise their bid, their object will continue to be selected. For winning agent i , the minimum necessary bid value is precisely d^*c_i as set by the algorithm. Bidding above d^*c_i , agent i continues to win; while bidding below d^*c_i results in i losing (the bid with value-per-unit-size of d^* would then have priority over i and both these agents' objects do not simultaneously fit in remaining available space).

For composability, we need to show that when the bids of all the agents except one are fixed arbitrarily, the set of selected objects as a function of this one agent's bid is unchanged for all the winning bid values of this agent. Whenever the approximate knapsack algorithm, AK, selects agent i , the other agents selected are exactly those that would have been selected had we run the algorithm without agent i on a knapsack of size $C' = C - c_i$ (after ignoring agents with size greater than $C/2$). Since agent i cannot affect the outcome of this process, the algorithm is composable. \square

The rationale for the term "composable" comes from the following lemma.⁴

⁴Note that composability plays a role similar to *cancellability* in Fiat et al. [6]. In a cancellable auction, the auction's profit is not a function of

LEMMA 5.3. *If mechanism \mathcal{M}_1 is composable and mechanism \mathcal{M}_2 is truthful then the composite mechanism, $\mathcal{M}_1 \circ \mathcal{M}_2$, is truthful.*

Proof. Fix the values of all bids but that of agent i . By the composability of \mathcal{M}_1 , if agent i is selected by \mathcal{M}_1 then H is fixed. This fixes the monotone selection rule of “ \mathcal{M}_2 given H ”. Intersecting the monotone selection rule of \mathcal{M}_1 with that of “ \mathcal{M}_2 given H ” gives the selection rule used for i by the composite mechanism. It is monotone. \square

Performance. Given some benchmark for gauging performance, the feasible outcome produced by \mathcal{M}_1 should not output a solution that is substantially worse, in terms of the chosen benchmark, than the optimal solution on the full input. If this is indeed the case, then with an approximately-optimal unlimited-supply mechanism, \mathcal{M}_2 , the composite mechanism will approximate the chosen benchmark on the full input. Recall that this notion is made precise by the definition of *performance preservation* (Definition 4.3). Lemma 4.7 asserts that the approximate knapsack algorithm, AK, approximately preserves the performance of the optimal monotone pricing.

Polynomial Time Computability. Until we impose the constraint of polynomial time computability, another auction seems like an attractive candidate for \mathcal{M}_1 in our composite mechanism: the Vickrey-Clarke-Groves (VCG) [16, 5, 11] mechanism. The VCG mechanism always selects set of items with the maximum valuation sum (AK only approximates this solution). Further, like AK the VCG mechanism is composable and approximately preserves the performance of the optimal monotone pricing (proofs omitted). Unfortunately, given standard complexity assumptions, VCG is not polynomial time computable.

AK satisfies all the requirements for \mathcal{M}_1 in the construction of a (limited-supply) knapsack auction. The missing ingredient is an approximately-optimal unlimited-supply knapsack auction that can be used as \mathcal{M}_2 . We present such an auction in the next section which when composed with AK gives a constructive proof of following theorem.

THEOREM 5.1. *For constants α and γ , there exists a knapsack auction with expected profit at least*

$$\alpha \text{OPT} - \gamma h \lg \lg n.$$

5.2 Unlimited-Supply Knapsack Auction. In this section, we consider the knapsack auction problem when $C = \infty$. We first attempt to use the general attribute auction of Blum and Hartline [3] to solve this problem. Since the optimal monotone pricing rule might offer a different price to

the value of any winning bid. This allows the auction to be canceled as a function of its profit.

every agent, the number of piecewise-constant pieces needed to emulate this rule could be as high as the number of agents n . Thus, a direct application of the attribute auction result (Theorem 2.2) to the knapsack auction problem would only guarantee a minimum profit of $\text{OPT}/16 - nh/2 \leq 0$, where $\text{OPT} \leq nh$ is the payoff of the optimal monotone pricing rule. Still, the unlimited-supply knapsack auction problem remains closely related to the attribute auction problem, and we will be making use of Theorem 2.2 in this section.

Let n' be the number of winners for the optimal monotone pricing function. Our results come from observing Lemma 5.4 below, which implies that there is an approximately optimal monotone pricing function

- (a) that divides the size range into $\lg n'$ intervals and for each interval, offers the same price to all agents whose size lies in the interval, and
- (b) for which most (all but $O(\lg \lg n')$) of the intervals have many (at least $O(\lg \lg n')$) winners.

Simply using part (a) of this fact and applying the result of Blum and Hartline [3], we can obtain an auction that is $\text{OPT}/16 - h \lg n'/2$. The main result of this section will be to use part (b) of this fact to improve the additive loss term to $O(h \lg \lg n')$.

We obtain this improvement by analyzing two possible cases. In the first case, most of the payoff from our approximately optimal monotone pricing comes from intervals with at least $\Omega(\lg \lg n')$ winners. For these large intervals, we can apply random sampling techniques and the Chernoff bound to show that a generalization of the random sampling auction of [10] will obtain a constant fraction of the optimal monotone payoff.

In the second case, most of the payoff from our approximately optimal monotone pricing comes from the $\Theta(\lg \lg n')$ small intervals. Here, the result of Blum and Hartline can be applied to get an auction that obtains a constant fraction of OPT less an additive term that is linear in the number of relevant intervals. This gives an additive loss term of $O(h \lg \lg n')$.

A convex combination of these two approaches gives an auction that is good in both cases. We start with a definition and a lemma.

DEFINITION 5.3. *A monotone pricing rule with exponential intervals is a monotone pricing rule in which the winners can be partitioned into equal-priced intervals over the attributes such that the i^{th} interval (in decreasing order of attribute value) contains at least 2^{i-1} winners.*

LEMMA 5.4. *Given any monotone pricing rule, $\pi(\cdot)$, that obtains total payoff P on instance $(v_1, \dots, v_n; c_1, \dots, c_n; C = \infty)$, there is a monotone pricing rule with exponential intervals, $\pi'(\cdot)$ with payoff at least $P/2$.*

Proof. Order the winners of $\pi(\cdot)$ on the instance by decreasing size (breaking ties arbitrarily). Divide the attribute range into intervals such that the i^{th} interval has at least 2^{i-1} winners but strictly fewer than 2^{i-1} winners with size strictly bigger than the smallest winner in i . This can be done by considering the attributes in decreasing order and adding them to the current interval until the first time the number of winners in the interval becomes at least 2^{i-1} . At this point, we move on to the next interval. Let $c(i)$ be the size of this smallest object in interval i . Consider $\pi'(\cdot)$ defined such that all objects in interval i are offered price $\pi(c(i))$.

Now we show that the payoff of $\pi(\cdot)$ is no more than twice that of $\pi'(\cdot)$. The loss for interval i is the difference in payoff between $\pi'(\cdot)$ and $\pi(\cdot)$ over the attribute interval $[c(i), c(i-1))$. There is no loss from objects with size exactly $c(i)$ and the loss from other objects in interval i is bounded by $\pi(c(i-1)) - \pi(c(i))$. Since interval i contains fewer than 2^{i-1} objects with size strictly more than $c(i)$, the total loss is no more than $(2^{i-1} - 1) \times (\pi(c(i-1)) - \pi(c(i)))$. We charge this loss to the winners in all the previous intervals. There are at least $\sum_{j=1}^{i-1} 2^{j-1} = 2^{i-1} - 1$ such winners; so each winner is charged at most $\pi(c(i-1)) - \pi(c(i))$. Now consider the total amount charged to a winner in interval i by subsequent intervals. The charges made to any given winner in interval i sum to at most $\pi(c(i))$; thus the total loss can be accounted for by the total payoff of $\pi'(\cdot)$. Therefore the payoff of $\pi'(\cdot)$ is at least half that of $\pi(\cdot)$. \square

Now, we are ready to define the random-sampling part of the unlimited-supply auction.

DEFINITION 5.4. (RSK) *The random sampling knapsack auction, RSK, does the following:*

1. *Partition the agents into two sets A and B uniformly at random.*
2. *Compute the optimal monotone pricing rule with exponential intervals (restricting prices to powers of two) for each partition. Let the pricing rules for A and B be π_A and π_B respectively.*
3. *Use π_A as the pricing rule for B and vice versa.*

LEMMA 5.5. *RSK is truthful.* (proof omitted)

Let π_A on A have n_A winners. Let \bar{n}_A be the largest power of 2 that is no larger than n_A . Then, the winners are divided up into at most $\lg \bar{n}_A + 1$ equal-priced markets. A market is said to be *large* for A if it has at least $256 \lg \lg \bar{n}_A$ winners when π_A is applied to A . Note that all markets other than the first $\lg \lg \lg n_A + 8$ markets (by decreasing attribute value) are large. Markets that are not large are called *small*. We wish to analyze the performance of RSK on the large markets. Define π'_A to be the pricing rule that is the same as

π_A , except that it offers a price of ∞ to the small markets. Let $P(\pi, A)$ denote the total profit of pricing function π applied to set A . Let \mathcal{L} be an ordering of the agents in the decreasing order of attribute value (breaking ties arbitrarily). Let \mathcal{L}_p denote the ordering \mathcal{L} restricted to agents having bids p or higher.

DEFINITION 5.5. (BAD EVENT, BAD SET) *A Bad Event is said to have occurred in RSK if there exists an $\eta = 2^k$ for integer $k \geq 4$, a price $p = 2^r$ with $h/\eta^2 < p \leq h$ and r integer, and a subset X of agents, satisfying the following properties:*

- (i) *All the agents in X have bids p or higher, and appear consecutively in \mathcal{L}_p .*
- (ii) $|X| \geq \frac{3}{2} \max \left\{ \frac{m_p}{6 \lg \eta}, 256 \lg \lg \eta \right\}$, where m_p is the total number of agents with bid p or higher.
- (iii) *One of the two sets created by RSK has more than $2|X|/3$ of the agents in X .*

A set X that satisfies the first two properties is called a Potential Bad Set, while any set X that satisfies all the above properties is called a Bad Set.

In the subsequent lemma we will make use of the following specification of the Chernoff bound:

CLAIM 5.1. *Consider a set X of $3x$ agents. The probability that set A has more than $2x$ agents from X is no more than $e^{-x/12}$. (proof omitted)*

LEMMA 5.6. *The probability of a Bad Event occurring in RSK is no more than 0.01.*

Proof. We will prove that the probability of the existence of a Bad Set X for which set A gets more than $2|X|/3$ of the agents is no more than 0.005. By symmetry, the probability of the existence of a Bad Set X with respect to B is also no more than 0.005. Then, we can take the union bound to get the lemma.

Fix a number $\eta = 2^k$ for some integer $k \geq 4$ and a price $p = 2^r$ such that $h/\eta^2 \leq p \leq h$. Let m_p be the total number of agents with price p or higher. Arrange these agents by decreasing order of object size. Let $L_\eta = \max \left\{ \frac{m_p}{6 \lg \eta}, 256 \lg \lg \eta \right\}$. Consider a subset X of $3x$ consecutive agents where $2x \geq L_\eta$. By Claim 5.1, the probability that this subset splits such that set A has more than $2x$ of these agents is no more than $e^{-x/12}$. Taking the union bound, the probability of such a subset existing for these

fixed values of η and p is no more than

$$\begin{aligned}
& m_p \sum_{|X|=3L_\eta/2}^{m_p} e^{-|X|/36} \\
& \leq m_p e^{-L_\eta/24} / (1 - e^{-1/36}) \\
& \leq 36.6 * (6L_\eta \lg \eta) e^{-L_\eta/24} \\
& \leq 220L_\eta 2^{L_\eta/256} 2^{-1.44L_\eta/24} \\
& = (220L_\eta 2^{-L_\eta/30}) 2^{-1.44L_\eta/24 + L_\eta/256 + L_\eta/30} \\
& \leq e^{-L_\eta/44}
\end{aligned}$$

To get the last inequality, we used the fact that $L_\eta \geq 256 \lg \lg \eta \geq 512$ when $n \geq 2^4$. Taking the union bound over all possible values of p (there are at most $2 \lg \eta$ of them), we get that the probability of such a subset existing for a given value of η is no more than

$$\begin{aligned}
2(\lg \eta) 2^{-L_\eta/44} & \leq 2(2^{\frac{256 \lg \lg \eta}{44}} \lg \eta) \leq 2(\lg n)^{-\frac{256}{44} + 1} \\
& \leq 2(\lg \eta)^{-4.8}
\end{aligned}$$

Taking the union bound over all $\eta = 2^k$ for $k = 4, 5, \dots$, we get that the probability is no more than

$$\begin{aligned}
2 \sum_{k=4}^{\infty} k^{-4.8} & = 2(4^{-4.8} + 5^{-4.8} + 6^{-4.8} + \dots) \\
& \leq 0.005
\end{aligned}$$

The inequality is obtained by using an integral to approximate the summation. \square

We now prove the following lemma about the revenue of RSK. A similar lemma holds when the roles of A and B are interchanged.

LEMMA 5.7. *For RSK,*

$$\mathbf{E}[P(\pi'_A, B)] \geq \frac{0.99}{2} (P(\pi'_A, A) - \frac{1}{2}P(\pi_A, A) - \frac{h}{2}).$$

Proof. Assume that $\bar{n}_A \geq 16$ and $P(\pi_A, A) > h$ as otherwise the claim is trivially true. Recall that a market is *large* for A if it has at least $256 \lg \lg \bar{n}_A$ winners when π_A is applied to A . If π_A applied to a large market has a profit greater than $\frac{P(\pi_A, A)}{2 \lg \bar{n}_A}$, then that market is called *significant* for A . Since the number of large markets is at most $\lg \bar{n}_A$, the total profit on applying π'_A to the significant markets of A is at least $P(\pi'_A, A) - P(\pi_A, A)/2$. Thus, we can prove the lemma by showing that with constant probability, $P(\pi'_A, B)$ is at least a constant fraction of the profit from applying π'_A to the significant markets of A .

Let $n_i(\pi_A, A)$ denote the number of winners in the i^{th} market when π_A is applied to A . We will show that assuming no *Bad Event* (see Definition 5.5) has occurred, there is no significant market i of A , such that $n_i(\pi_A, A) > 2n_i(\pi_A, B)$. Since no *Bad Event* occurs with probability at least 0.99, it would immediately imply that

$$\mathbf{E}[P(\pi'_A, B)] \geq \frac{0.99}{2} (P(\pi'_A, A) - \frac{1}{2}P(\pi_A, A)).$$

Assume that no *Bad Event* has occurred. For a contradiction, suppose that there is a significant large market i of A that has $n_i(\pi_A, A) > 2n_i(\pi_A, B)$. Let $[a, b]$ be the attribute range corresponding to this market. Let the price offered to the i^{th} market by π_A be $p = 2^k$ for some integer k . Let m_p be the total number of agents with bid p or higher. We claim that $p > \frac{h}{n_A^2}$. Suppose to the contrary, the price $p \leq \frac{h}{n_A^2}$. Then the i^{th} market has a payoff of at most

$$\frac{h}{n_A} < \frac{h}{2 \lg \bar{n}_A} < \frac{P(\pi_A, A)}{2 \lg \bar{n}_A}$$

This would imply that market i is not a significant market, a contradiction to the supposition above.

Recall that \mathcal{L}_p is an ordering of all the agents with bids p or higher in decreasing order of attribute value. Consider a set X of $\frac{3}{2}n_i(\pi_A, A)$ agents with bids p or higher that appear consecutively in \mathcal{L}_p over the attribute range $[a, b]$. Then, by assumption, more than $\frac{2}{3}|X|$ agents from this set are in set A . We show that X is a *Potential Bad Set* with $\eta = \bar{n}_A$. We already know that $|X| \geq \frac{3}{2}(256 \lg \lg \bar{n}_A)$. Thus, all we need to show is that $|X| \geq \frac{m_p}{4 \lg \bar{n}_A}$, or alternatively, that $n_i(\pi_A, A) \geq \frac{m_p}{6 \lg \bar{n}_A}$. To see this, note that since

market i is significant for A , $pn_i(\pi_A, A) > \frac{P(\pi_A, A)}{2 \lg \bar{n}_A}$. In other words, $2pn_i(\pi_A, A) \lg \bar{n}_A > P(\pi_A, A)$. If more than $2n_i(\pi_A, A) \lg \bar{n}_A$ agents in set A had bids of p or higher, then offering a price of p to everybody would yield a profit of more than $P(\pi_A, A)$, contradicting the optimality of π_A for set A . Thus, the number of agents in set A with bid p or higher is no more than $2n_i(\pi_A, A) \lg \bar{n}_A$. Consider the set of all agents with bids p or higher. This is a *Potential Bad Set*. Since the *Bad Event* has not occurred, the third condition for a *Bad Event* is not satisfied. Thus, if set A has no more than $2n_i(\pi_A, A) \lg \bar{n}_A$ agents with bid p or higher, then the total number of agents with bid p or higher $m_p \leq 3(2n_i(\pi_A, A) \lg \bar{n}_A)$, or $n_i(\pi_A, A) \geq \frac{m_p}{6 \lg \bar{n}_A}$.

Thus, X is a *Potential Bad Set*. Since the *Bad Event* has not occurred, X does not satisfy the third condition of being a *Bad Set*, implying that the number of agents in $X \cap A$ is no more than $\frac{2}{3}|X|$, thus contradicting the supposition that $n_i(\pi_A, A) > 2n_i(\pi_A, B)$. \square

Consider the following combination of the general attribute auction with the random sampling knapsack auction.

DEFINITION 5.6. (USK) *The unlimited supply knapsack auction, USK, works as follows,*

1. Perform Step 1 of RSK.
2. With probability p , run the general attribute auction on the sets A and B separately. With the remaining probability, run the remaining steps of RSK.

LEMMA 5.8. *Auction USK is truthful.* (proof omitted)

THEOREM 5.2. *The revenue generated by USK is $\alpha \text{OPT} - \gamma h (\lg \lg \lg n_A + \lg \lg \lg n_B + 19)$, where α and γ are constants.*

Proof. Recall that OPT is the payoff of the optimal monotone pricing scheme. Using Lemma 5.4 and losing another factor of 2 due to restriction to prices that are powers of 2, $\text{OPT} \leq 4(\mathbf{E}[P(\pi_A, A)] + \mathbf{E}[P(\pi_B, B)])$.

Recall that any market with fewer than $256 \lg \lg \bar{n}_A$ winners is small for A . There are at most $\lg \lg \lg \bar{n}_A + 9$ small markets of A with respect to π_A . Similarly, there are at most $\lg \lg \lg \bar{n}_B + 9$ small markets of B with respect to π_B . Note that markets that are not small are *large* for their respective sets. Let $P(\pi_A, A_S)$ be the payoff of applying π_A to the small markets of A . Define $P(\pi_B, B_S)$ similarly. Then, $\mathbf{E}[P(\pi_A, A)] + \mathbf{E}[P(\pi_B, B)] = \mathbf{E}[P(\pi_A, A_S)] + \mathbf{E}[P(\pi_B, B_S)] + \mathbf{E}[P(\pi'_A, A)] + \mathbf{E}[P(\pi'_B, B)]$.

With probability p , we use the general attribute auction, in which case, by Theorem 2.2, we get an expected revenue of at least $\frac{1}{16}(\mathbf{E}[P(\pi_A, A_S)] + \mathbf{E}[P(\pi_B, B_S)]) - \frac{h}{2}(\lg \lg \lg n_A + \lg \lg \lg n_B + 18)$.

On the other hand, when we use RSK (which we do with probability $1 - p$), we can apply Lemma 5.7 and the same lemma with A and B interchanged, to show that we get an expected revenue of at least $\beta(\mathbf{E}[P(\pi'_A, A)] - \frac{1}{2}\mathbf{E}[P(\pi_A, A)] + \mathbf{E}[P(\pi'_B, B)] - \frac{1}{2}\mathbf{E}[P(\pi_B, B)] - h)$ for $\beta = 0.99/2$. Thus, the overall expected revenue is at least

$$\begin{aligned} & \frac{p}{16} (\mathbf{E}[P(\pi_A, A_S)] + \mathbf{E}[P(\pi_B, B_S)]) \\ & - \frac{ph}{2} (\lg \lg \lg n_A + \lg \lg \lg n_B + 18) \\ & + (1 - p)\beta(\mathbf{E}[P(\pi'_A, A)] + \mathbf{E}[P(\pi'_B, B)] \\ & - \mathbf{E}[P(\pi_A, A)]/2 - \mathbf{E}[P(\pi_B, B)]/2 - h) \end{aligned}$$

Setting $p = \frac{24\beta}{24\beta+1}$, we get an expected revenue of at least

$$\begin{aligned} & \frac{3\beta}{2(24\beta+1)} (\mathbf{E}[P(\pi_A, A_S)] + \mathbf{E}[P(\pi_B, B_S)]) \\ & - \frac{12\beta h}{24\beta+1} (\lg \lg \lg n_A + \lg \lg \lg n_B + 18) \\ & + \frac{\beta}{24\beta+1} (\mathbf{E}[P(\pi'_A, A)] + \mathbf{E}[P(\pi'_B, B)]) \\ & - \frac{\beta}{24\beta+1} \left(\frac{\mathbf{E}[P(\pi_A, A)]}{2} - \frac{\mathbf{E}[P(\pi_B, B)]}{2} - h \right) \\ & \geq \frac{\beta}{24\beta+1} (\mathbf{E}[P(\pi_A, A)] + \mathbf{E}[P(\pi_B, B)]) \\ & - \frac{12\beta h}{24\beta+1} \left(\lg \lg \lg n_A + \lg \lg \lg n_B + 18 + \frac{1}{12} \right) \end{aligned}$$

To get the inequality, we have used the fact that $\mathbf{E}[P(\pi_A, A)] + \mathbf{E}[P(\pi_B, B)] \leq \mathbf{E}[P(\pi_A, A_S)] + \mathbf{E}[P(\pi_B, B_S)] + \mathbf{E}[P(\pi'_A, A)] + \mathbf{E}[P(\pi'_B, B)]$. Putting in the value $\beta = 0.99/2$, and using the fact that $\text{OPT} \leq 4(\mathbf{E}[P(\pi_A, A)] + \mathbf{E}[P(\pi_B, B)])$, we get the theorem with $\alpha = 0.009$ and $\gamma = 0.47$. \square

Applying Theorem 5.2 with the observation that n_A and n_B are no more than n , we get Theorem 5.1.

References

- [1] Z. Abrams. Revenue maximization when bidders have budgets. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2006.
- [2] M.-F. Balcan, A. Blum, J. Hartline, and Y. Mansour. Mechanism Design via Machine Learning. In *Proc. of the 46th IEEE Symp. on Foundations of Computer Science*, 2005.
- [3] A. Blum and J. Hartline. Near-optimal online auctions. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1156–1163, 2005.
- [4] C. Borgs, J. Chayes, N. Immorlica, M. Mahdian, and A. Saberi. Multi-unit auctions with budget-constrained bidders. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, pages 44–51, 2005.
- [5] E. H. Clarke. Multipart Pricing of Public Goods. *Public Choice*, 11:17–33, 1971.
- [6] A. Fiat, A. Goldberg, J. Hartline, and A. Karlin. Competitive generalized auctions. In *Proceedings of the 34th Annual ACM Symposium on the Theory of Computing*, pages 72–81, 2002.
- [7] A. Goldberg and J. Hartline. Envy-free auctions for digital goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, pages 29–35, 2003.
- [8] A. Goldberg and J. Hartline. Collusion-resistant mechanisms for single-parameter agents. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 620–629, 2005.
- [9] A. Goldberg, J. Hartline, A. Karlin, M. Saks, and A. Wright. Competitive auctions. *Games and Economic Behavior*. Submitted for publication. An earlier version available as InterTrust Technical Report STAR-TR-99-01.
- [10] A. Goldberg, J. Hartline, and A. Wright. Competitive auctions and digital goods. In *Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 735–744, 2001.
- [11] T. Groves. Incentives in Teams. *Econometrica*, 41:617–631, 1973.
- [12] V. Guruswami, J. Hartline, A. Karlin, D. Kempe, C. Kenyon, and F. McSherry. On profit-maximizing envy-free pricing. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1164–1173, 2005.
- [13] A. Mehta, A. Saberi, U. Vazirani, and V. Vazirani. Adwords and generalized online matching. In *Proceedings of the 46th IEEE Symposium on Foundations of Computer Science*, 2005.
- [14] A. Mu'alem and N. Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. In *Proceedings of the 18th National Conference on Artificial Intelligence (AAAI)*, pages 379–384, 2002.
- [15] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.
- [16] W. Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *J. of Finance*, 16:8–37, 1961.