Revenue Maximization When Bidders Have Budgets *

Zoë Abrams†

Abstract

We study the problem of maximizing revenue for auctions with multiple units of a good where bidders have hard budget constraints, first considered in [2]. The revenue obtained by an auction is compared with the optimal omniscient auction had the auctioneer known the private information of all the bidders, as in *competitive analysis* [7]. We show that the revenue of the optimal omniscient auction that sells items at many different prices is within a factor of 2 of the optimal omniscient auction that sells all the items at a single price, implying that our results will carry over to multiple price auctions. We give the first auction for this problem, to the best of our knowledge, that is known to obtain a constant fraction of the optimal revenue when the *bidder dominance* (the ratio between the maximum contribution of a single bidder in the optimal solution and the revenue of that optimal solution) is large (as high as $\frac{1}{2}$). Our auction is also shown to remain truthful if canceled upon not meeting certain criteria. On the negative side, we show that no auction can achieve a guarantee of $\frac{1}{2-\epsilon}$ the revenue of the optimal omniscient multi-price auction. Finally, if the bidder dominance is known in advance and is less than $\frac{1}{5.828}$, we give an auction mechanism that raises a large constant fraction of the optimal revenue when the bidder dominance is large *and* is asymptotically close to the optimal omniscient auction as the bidder dominance decreases. We discuss the relevance of these results for related applications.

1 Introduction

This paper is concerned with maximizing revenue in auctions with multiple units of a good where bidders have hard constraints on the amount they are able to spend. These auctions arise often in practice. Examples include privatization in places such as Eastern Europe, China, and Russia. Budget constrained auctions have also occurred in the U.K. and U.S.A., where the Federal Communication Commission raised billions auctioning off licenses for telecommunication radio frequencies [13]. Another example is a company's auctioning of shares in an initial public offering. Recently, the advent of electronic commerce has given rise to auctions of internet advertisements for search engine queries at companies such as Yahoo! and Google. Yet, despite their prevalence and the potentially tremendous gains in revenue, there are still many aspects of budget constrained auctions that are not well understood.

We consider auctions where the goal is profit maximization (i.e. the auctioneer would like to raise as much revenue as possible). Traditional auction analysis focuses on profit maximization when there is some advance knowledge about the distribution of the bidders' valuations. Along these lines, there has been work in the economics community on designing budget constrained auctions in the Bayesian setting [3, 4, 9, 10]. In contrast, we are interested in maximizing profit in the worst case scenario, when the auctioneer does not have any information about the bidder valuations. We will employ *competitive analysis*, introduced in [7], where the metric used to gauge the performance of the auction is to compare its revenue with the revenue that could be raised by the *optimal omniscient auction* (the auction that raises the optimal revenue if the auctioneer knows the true valuations of all the bidders). The *competitive ratio* is the worst case ratio, over all possible inputs, between the auction's revenue and the revenue of the optimal omniscient auction.

The problem of focus in this paper is the multi-unit budget-constrained (MUBC) auction problem, proposed in [2], in which there are a limited number of goods, and bidders have hard budgets above which they are unwilling to spend. Each bidder has a private value for a unit of the good and a private overall budget. We assume that bidders want to maximize their own utility, and that they have unbounded negative utility if they exceed their budget. We make the following contributions to the study of the MUBC auction problem:

- **The revenue of the multi-price optimal omniscient auction is within 2 of the single-price optimal omniscient auction.** We show that the optimal revenue that can be achieved by selling the good at many different prices is at most twice the optimal revenue of selling the good at a single price.
- **A constant competitive auction when the bidder dominance is large.** The auction we will present has strong guarantees on performance even when a single bidder's budget is as much as half

*Research supported by NSF under ITR IIS-0205633
†Stanford University, email: zoea@stanford.edu

the maximum revenue.

This auction can be canceled. The concept of cancelability is defined in [5] as the property that an auction can be canceled if it does not raise enough revenue, and the resulting auction maintains truthfulness. Using the notion of composability from [1], we present an auction for the MUBC problem that can be canceled if some value at least half of the revenue raised is not more than some predetermined cutoff.

- **A lower bound on the competitive ratio.** We show that no MUBC auction can achieve a guarantee of more than $\frac{1}{2}$ the revenue of the optimal omniscient auction that sells the items at multiple prices.

- **An auction parameterized by the bidder dominance that is constant competitive in the worst case *and* asymptotically optimal.** As the bidder dominance decreases, this auction approaches the optimal solution extremely quickly. One drawback is that the auction is parameterized by the bidder dominance, and therefore the dominance must be known in advance. Also, the auction sells items at potentially many different prices and is infeasible if the bidder dominance is above $\frac{1}{5.828}$.

Related Work Our framework for studying the MUBC auction problem is inspired by work that has been done on the *digital goods auction problem* [7, 5, 8, 6], where the auctioneer has an unlimited supply of goods. Despite similarities in the techniques we use, particularly with regards to our auction that is constant competitive, there are considerable differences between the two problems. First, the utility functions of the bidders in the MUBC auction depend on two inputs and are not quasi-linear. Also, in the MUBC problem, a single bidder may purchase multiple items, whereas in the digital goods auction problem, each bidder only purchases one unit of the item. Because of these differences, the results in this paper require novel technical contributions and new theoretical insights to effectively adapt the techniques from [5] to the budget constrained scenario.

The focus of this paper is motivated by the work in [2], where the MUBC problem is proposed and an asymptotically optimal auction designed. The asymptotically optimal auction performs well when the bidder dominance is small, achieving performance that moves asymptotically closer to the optimal as the bidder dominance decreases. In particular, when the bidder dominance is below $\frac{1}{160}$, the asymptotic auction outperforms the profit of the auction presented in Section 4 of this paper. However, the auction is not known to be cancelable and there is no guarantee that it is constant competitive when the bidder dominance is large. In contrast with our work, there is no bound on the competitive ratio using the results from [2] when the bidder dominance is above $\frac{1}{24}$.

The paper is organized as follows. In Section 2 we present the problem and definitions. In Section 3 we prove the bound between the single and multi-price optimal omniscient auctions. Section 4 contains an auction that is constant competitive when the bidder dominance is large and proofs of its truthfulness, competitive ratio, and cancelability. A negative result on the impossibility of designing a randomized auction with competitive ratio below 2 follows in Section 5. We present an auction with high competitive ratio when the bidder dominance is known in advance, along with related proofs, in Section 6 and end with some concluding thoughts.

2 Problem and Definitions

In the *multi-unit budget constrained auction problem*, an auctioneer would like to sell m indivisible units of a good to n bidders. Each bidder i has a private valuation $t_i^v \in \mathbb{R}^+$, which is the price they are willing to pay per unit of the good, and a private budget constraint $t_i^b \in \mathbb{R}^+$. The budget constraint is hard: the bidder will not spend more than the budget under any circumstances. The total utility that a bidder receives from an allocation of x_i units of the good, for a total price of P_i, is $x_i t_i^v - P_i$ if $P_i \leq t_i^b$ and $-\infty$ otherwise. The goal of the bidder is to maximize their expected utility. We assume bidders do not collude, are rational, and have complete knowledge of the auction mechanism.

We consider a single round sealed bid auction that proceeds as follows. First, each bidder submits two values. One represents the maximum amount they are willing to pay per item (i.e. their value for one unit of the item) and the other the maximum they are willing to pay overall (i.e. their budget). The auctioneer receives vectors $v = (v_1, v_2, ...v_n)$ and $b = (b_1, b_2, ...b_n)$ of bids where v_i is the reported value and b_i the reported maximum budget from bidder i. Next, given vectors v and b, the auctioneer computes an output consisting of an allocation vector $x = (x_1, ..., x_n)$ and a price vector $P = (P_1, ..., P_n)$. Each element of the allocation vector indicates the number of items received by the corresponding bidder. Each element of the price vector indicates the payment made to the auctioneer.

We place some restrictions on the auctioneer. To ensure voluntary participation [11], prices must satisfy $0 \leq P_i \leq b_i$ and the expected utility of a truthful bidder must be non-negative. For a randomized auction, the auctioneer outputs a probability distribution over pairs of vectors x and P which satisfy $0 \leq P_i \leq b_i$ for all possible pairs of vectors that have non-zero probability. We also require that $|x| \leq m$.

The goal of the auctioneer is to raise as much revenue as possible. The revenue of the auction A selling m units of a good given bids v and b is $A(v,b,m) = \sum_i P_i$. Unless the submitted bid vectors bear some relationship to the true valuations, the auctioneer will be unable to make price and allocation decisions that provide a large amount of revenue. We will overcome this difficulty by constructing auctions that are truthful. An auction is *truthful* if, for every bidder i, the profit of bidder i is maximized by biding their true valuations, regardless of the values reported by bidders other than i.

We will determine the success of the auction by comparing the revenue raised with the optimal profit the auctioneer could have raised if the true valuations of the bidders were known in advance (i.e. the profit of the *optimal omniscient auction*). We consider value and budget vectors v and b (of values v_i and b_i respectively), sorted in decreasing order according to v. The optimal single price revenue of selling m items to bidders with values v and budgets b is $F = F(v,b,m) = min(\sum_{j=1}^{k^*} b_j, v_{k^*}m)$ where $k^* = k(v,b,m) = argmin_i(\sum_{j=1}^{i} b_j \geq v_{i+1}m)$. We refer to k^* as the stopping bidder. We refer to the amount of revenue collected from bidder i in the optimal solution as $\bar{b}_i$, so that $F = \sum_{i=1}^{k^*} \bar{b}_i$. If $v = t^v$ and $b = t^b$ than $F(v,b,m)$ is the optimal single-price omniscient revenue.

It will not be possible to compete with the optimal single-price omniscient profit for all vectors v and b. We can imagine an optimal solution that sells all the items to the bidder with the highest value. Due to the characterization in [2], we cannot hope to get a sufficient amount of revenue from this single bidder. Therefore, as in [2], the auction will provide guarantees for input vectors where the optimum sells to more than a single bidder. We define b_{max} as the largest budget from amongst the bidders receiving items in the optimal solution and define α as any value such that $\alpha \leq \frac{F}{b_{max}}$. Intuitively, b_{max} is a bound on the amount any single bidder contributes to the optimal solution and $\frac{1}{\alpha}$ is an upper bound on the fraction of the optimal revenue any single bidder contributes. We use terminology from [8] and say that auction A has competitive ratio β against z if for all bid vectors v and b such that $\alpha \geq z$, the expected profit of $A(v,b,m)$ satisfies $\beta \geq \frac{F(v,b,m)}{E[A(v,b,m)]}$.

We will assume hereafter that the goods are divisible. If goods are indivisible, then the auctioneer fixes the overall price and instead of allocating a fractional quantity c of a good, allocates a full unit with probability c as described in [2].

3 Multi-price Versus Single Price Auctions

We will be comparing our auction to the optimal single-price omniscient auction. This gives our benchmark a significant handicap since the problem formulation does not require that we sell all items at the same price. We would prefer a guarantee against the optimal auction that is allowed to sell items at many different prices. Surprisingly, we will show that auctions that perform well against a single-price optimal omniscient auction *also* perform well against a multi-price optimal omniscient auction. The *optimal omniscient multi-price auction* has stopping bidder $k^T = argmax_i(\sum_{j=1}^{i-1} \frac{b_j}{v_j} < m)$ and revenue $T(v,b,m) = \sum_{j=1}^{k^T-1} b_j + \left(m - \sum_{j=1}^{k^T-1} \frac{b_j}{v_j}\right) v_{k^T}$.

THEOREM 3.1. *For the budget constrained auction problem, $T(v,b,m) \leq 2F(v,b,m)$ and this bound is tight.*

Proof. The budget of bidder k^T is assumed to be exactly the amount of budget it contributes to $T(v,b,m)$ for notational convenience so that $T = T(v,b,m) = \sum_{i=1}^{k^T} b_i$, $\sum_{i=1}^{k^T} b_i/v_i = m$. There are two cases.
CASE I ($F = mv_{k^*}$): $\forall i, k^* \leq i \leq k^T$, bidder i has $m/F \leq 1/v_i$ because $mv_{k^*} = F$ and $v_i \leq v_{k^*}$ by definition. Combining equations, $\sum_{i=k^*}^{k^T} mb_i/F \leq \sum_{i=k^*}^{k^T} b_i/v_i \leq m \rightarrow \sum_{i=k^*}^{k^T} b_i \leq F$. Now, $T(v,b,m) = \sum_{i=1}^{k^*-1} b_i + \sum_{i=k^*}^{k^T} b_i \leq 2F$.
CASE II ($F = \sum_{i=1}^{k^*} b_i$): $\forall i, k^* < i \leq k^T$, bidder i has $m/F \leq 1/v_i$ because $mv_{k^*+1} \leq \sum_{j=1}^{k^*} b_j = F$ and $v_i \leq v_{k^*+1}$ by definition. Combining these equations, $\sum_{i=k^*+1}^{k^T} mb_i/F \leq \sum_{i=k^*+1}^{k^T} b_i/v_i \leq m \rightarrow \sum_{i=k^*+1}^{k^T} b_i \leq F$. Now, $T = \sum_{i=1}^{k^*} b_i + \sum_{i=k^*+1}^{k^T} b_i \leq 2F$.

To show that this bound is tight, consider an auction with two bidders and $n+1$ items. For the first bidder, $b_1 = 1, v_1 = 1$. The second bidder has $b_2 = 1, v_2 = \frac{1}{n}$. $T(v,b,n+1) = 2$ and $F(v,b,n+1) = 1 + \frac{1}{n}$, so as n gets large, $T(v,b,n+1)$ approaches twice the value of $F(v,b,n+1)$. ∎

4 Profit Extract Partition Auction: A Constant Competitive Auction For Large Bidder Dominance

The techniques used in this auction are inspired by [5]. However, our problem is different because each bidder desires many items (as opposed to only a single item) and contributes varying values to the optimal solution (as opposed to the price for only a single item). Due to this difference, we cannot use the same counting argument to show a bound on the expected min, over two partitions. Instead we must develop a new and technically involved analysis using the merging of budgets to handle variable bidder contributions. The competitive ratio of the auction we present is at most $\frac{4\alpha}{\alpha-1}$, which is always at most 8 for inputs where $\alpha \geq 2$. We first give a new version of $ProfitExtract_R$ from [8] that is designed for the budget constrained auction problem and will be used as a subroutine in the profit extract partition auction.

$ProfitExtract_{R,m}$ **Auction**
1. Set $p = \frac{R}{m}$.
2. Consider each bidder in random order. If bidder i has $v_i \geq p$, it receives either $\frac{b_i}{p}$ or all of the remaining goods, whichever is smaller, and is charged p per item received. Continue until all m goods allocated or all bidders considered.
3. If all items have not been sold, all bidders lose.

LEMMA 4.1. *$ProfitExtract_{(R,m)}$ is truthful.*

Proof. There are two cases.
Case 1 ($t_i^v < p$): If bidder i bids $v_i \geq p$, her utility is nonpositive if she receives the item and 0 if she does not receive the item. If she bids below p she will not receive items and her utility will be zero. Thus, bidding any value below p is preferable and, in particular, bidding her true value is a dominant strategy. If she bids $v_i = t_i^v$, the budget bid b_i is irrelevant. Therefore truthful reporting of the budget and value is a dominant strategy.
Case 2 ($t_i^v \geq p$): Assuming bidder i reports $v_i = t_i^v$, no other bid can increase bidder i's utility. If not all m are sold in Step 3, since the price is set independent of bidder i's bids, misreporting can only raise more revenue by over bidding budget and buying more items, leading to negative utility. If all m items are sold in Step 3, there is a non-zero probability that bidder i will receive items, and this probability is the same for all values of v_i above p, and is independent of b_i. If she is chosen to receive items, her expected utility in Step 2 is $min(\frac{t_i^b}{p}(t_i^v - p), x(t_i^v - p))$ (where x is the remaining number of items when i is considered), is independent of v_i. For the reporting of b_i, under reporting can only decrease total expected utility and over reporting results in an allocation that exceeds her budget (resulting in negative infinity utility).

LEMMA 4.2. *$ProfitExtract_{(R,m)}$ raises revenue R if it is possible (i.e. if $F(v, b, m) \geq R$) and zero otherwise.*

Proof. Take the optimal solution resulting in revenue $F = F(v, b, m)$. If $m' \leq m$ goods are sold, then the optimal solution sells the items at price $F/m' \geq R/m = p$. All winning bidders from the optimal solution, whose budgets sum to at least F, are willing to buy the goods at price p and therefore $ProfitExtract_{(R,m)}$ is guaranteed to sell all m items at price p and raise revenue exactly R. If $F < R$, then clearly no single price can raise R and not all items will be sold in Step 3. ∎

Profit Extract Partition Auction
1. Partition bids uniformly at random into two sets S' and S'' with corresponding value and budget vectors v', b', v'', and b''.
2. Compute $F' = F(v', b', m/2)$ and $F'' = F(v'', b'', m/2)$.
3. Run $ProfitExtract_{(F',m/2)}$ on the bidders in S'' and $ProfitExtract_{(F'',m/2)}$ on the bidders in S'.

THEOREM 4.1. *The Profit Extract Partition Auction is truthful.*

Proof. The revenue F'(or F'') to be extracted and the number of items $m/2$ are independent the reports of bidders in S'(or S'). Since $ProfitExtract$ with independent input parameters is truthful, the profit extract partition auction is truthful. ∎

We present Lemmas 4.3 and 4.4 that will be helpful later on in proving the competitive ratio of the Profit Extract Partition Auction. For ease of notation, the 1 subscript for the L1-norm is sometimes dropped. For a set of positive numbers S, we denote a vector v of these numbers with v^S. Given a vector b of n positive real values from the set $S = \{b_1, b_2, ..., b_n\}$, where each value has been assigned to either partition S' or S'' based on an independent toss of a fair coin, then $g(b) \triangleq E[min(|b^{S'}|_1, |b^{S''}|_1)]$. The function $merge(b, b_i, b_j)$, where $b_i \geq b_j$ and $b_i, b_j \in b$, creates a new bid vector b' of $n-1$ bids that includes all bids from b except b_j and replaces the value b_i with a new bid of value $b_i + b_j$.

LEMMA 4.3. *Given a vector b, for any i,j, and $b' = merge(b, b_i, b_j)$,$g(b) \geq g(b')$.*

Proof. We define Q to be $\{q | q \subset S \wedge [(|b^q| < \frac{|b|}{2}) \vee (|b^q| = \frac{|b|}{2} \wedge b_i \notin q)]\}$. In words, every collection of bidders whose sum of bids is less than the sum of bids of all other bidders, is a set in Q. To ensure that evenly split partitions are not counted twice, when there is an even split we include in Q only the side that does not include bid b_i. Let $Q_{b_i,b_j} \triangleq \{q \in Q | b_i, b_j \in q\}$, $Q_{\not b_i,\not b_j} \triangleq \{q \in Q | b_i, b_j \notin q\}$, $Q_{b_i,\not b_j} \triangleq \{q \in Q | b_i \in q, b_j \notin q\}$

, $Q_{\not{b}_i,b_j} \triangleq \{q \in Q | b_i \notin q, b_j \in q\}$, and let p_q be the probability that set q is the collection of bidders in the minimum sized partition.

$$g(b) = \sum_{q \in Q} |q| p_q =$$
$$\sum_{q \in Q_{b_i,b_j}} |q| p_q + \sum_{q \in Q_{\not{b}_i,\not{b}_j}} |q| p_q +$$
$$\sum_{q \in Q_{\not{b}_i,b_j}} |q| p_q + \sum_{q \in Q_{b_i,\not{b}_j}} |q| p_q$$

Now,

$$g(b') = \sum_{Q_{b_i,b_j}} |q| p_q +$$
$$\sum_{Q_{\not{b}_i,\not{b}_j}} |q| p_q + \sum_{Q_{\not{b}_i,b_j}} (|q| - b_j) p_q +$$
$$\sum_{\substack{q \in Q_{b_i,\not{b}_j}, \\ |q| \leq \frac{|b|}{2} - b_j}} (|q| + b_j) p_q +$$
$$\sum_{\substack{q \in Q_{b_i,\not{b}_j}, \\ |q| > \frac{|b|}{2} - b_j}} (|b| - |q| - b_j) p_q$$

For $q \in Q_{b_i,\not{b}_j}$ and $|q| > \frac{|b|}{2} - b_j$, we have that $|q| - (|b| - |q| - b_j) = 2|q| - |b| + b_j > -2b_j + |b| = -b_j$. Subtracting $g(b')$ and substituting the above inequality:

$$\begin{aligned} g(b) - g(b') \quad > \quad & - \sum_{q \in Q_{b_i,\not{b}_j}, |q| > \frac{|b|}{2} - b_j} b_j p_q \\ & - \sum_{q \in Q_{b_i,\not{b}_j}, |q| \leq \frac{|b|}{2} - b_j} b_j p_q \\ & + \sum_{Q_{\not{b}_i,b_j}} b_j p_q \\ = \quad & - \sum_{q \in Q_{b_i,\not{b}_j}} b_j p_q + \sum_{Q_{\not{b}_i,b_j}} b_j p_q \\ = \quad & b_j \Big(\sum_{Q_{\not{b}_i,b_j}} p_q - \sum_{q \in Q_{b_i,\not{b}_j}} p_q \Big) \end{aligned}$$

For every set $q \in Q_{b_i,\not{b}_j}$, $b_i \geq b_j$ and $q \neq \frac{|b|}{2}$ ensures the set $q' = q - b_i + b_j$ is in $Q_{\not{b}_i,b_j}$. Furthermore $p_{q'} = p_q$ because they are the same collection except one final element is swapped and element b_i has the same probability of being the final single element as does b_j. We define $Q_{map} = \{q' \in Q_{\not{b}_i,b_j} | \exists q \in Q_{b_i,\not{b}_j}\}$. In words, Q_{map} contains all sets from $Q_{\not{b}_i,b_j}$ that are a counterpart for some $q \in Q_{b_i,\not{b}_j}$. Because every $q \in Q_{b_i,\not{b}_j}$ has a unique counterpart $q' \in Q_{\not{b}_i,b_j}$, we can rewrite the above equation as: $g(b) - g(b') > b_j(\sum_{Q_{map}} 0 + \sum_{Q_{\not{b}_i,b_j} - Q_{map}} p_{q'})$, which is clearly positive and the claim is proved. ∎

LEMMA 4.4. *Given a vector b, $g(b)$ is at least $\frac{(\alpha-1)|b|_1}{4\alpha}$ if $\alpha b_{max} \leq |b|$.*

Proof. The Lemma 4.3 shows that for any merge operation, $g(b) \geq g(b')$. We partition the bids in b into two sets A and B by placing the next successive bid (sorted in decreasing order), in whichever partition has a smaller sum of bids. The resulting vectors are b^A and b^B respectively. Assuming without loss of generality that $|b^A| \leq |b^B|$ and that b_f is the last bid placed in partition B, $|b^A| \geq |b^B| - b_f$ because we chose to place b_f in b^B when A contained at most all of its items. Combining this inequality with $|b^A| + |b^B| = |b|$ we have $|b^A| \geq |b^B| - b_f \geq |b| - |b^A| - b_f$, which implies $|b^A| \geq \frac{|b| - b_f}{2}$. By assumption $\alpha b_f \leq |b|$ (or, $-b_f \geq \frac{|b|}{\alpha}$), giving $|b^A| \geq (\frac{\alpha-1}{2\alpha})|b|$.

Now, we continually apply the function merge to items in the same partition until we are left with two items, one with value $|b^A|$ and one with value $|b^B|$. The expected min over partitions of these two items is $\frac{|b^A|}{2}$ since the probability the items are split is $\frac{1}{2}$. By Lemma 4.3, $g(b)$ is larger than $g(\{|b^A|, |b^B|\})$, giving the lemma. ∎

THEOREM 4.2. *The Profit Extract Partition Auction has competitive ratio $\frac{4\alpha}{\alpha-1}$ against α.*

Proof. Consider the k bidders who contribute to the optimal solution F. By Lemma 4.4, the expected min sum of $\bar{b}_i$ from these k bidders once partitioned is at least $\frac{(\alpha-1)F}{4\alpha}$. But, Step 2 of the auction sells only half of the items to each partition, so we would like to show that even with fewer items, a large portion of the optimal solution's revenue is retrieved. Without loss of generality assume $F' \leq F''$. There are two possibilities:

- If $\sum_{\{i \in S', 1 \leq i \leq k\}} \bar{b}_i \geq F/2$, then $F(v', b', m/2) \geq F/2$ since at least $m/2$ items can be sold at price F/m.
- Otherwise, $\sum_{\{i \in S', 1 \leq i \leq k\}} \bar{b}_i < F/2$ and therefore $F(v', b', m/2) \geq \sum_{\{i \in S', 1 \leq i \leq k\}} \bar{b}_i$ since bidders from k can expend their entire budgets on less than m/2 items sold at price F/m.

Combining possibilities, $F' = min(F/2, \sum_{i=1}^{k} \bar{b}_i)$. Since $E[\sum_{i=1}^{k} \bar{b}_i] \geq \frac{(\alpha-1)F}{4\alpha}$, we have $E[F'] \geq \frac{(\alpha-1)F}{4\alpha}$.

When $ProfitExtract_{(F',m/2)}$ is run on S'' in Step 3, by lemma 4.2, the auction raises at least F' revenue. The main contribution of this section, namely, the competitive ratio of the profit extract partition auction, follows directly. ∎

4.1 Composition and Cancelability The process of combining auctions in sequence has been previously studied [12, 1]. We now generalize the notion of composability as presented in [1].

Definition. Generalized Composition Given two auctions $\mathcal{A}^1$ and $\mathcal{A}^2_Q$, we define the general composition of $\mathcal{A}^1 \circ \mathcal{A}^2_Q$ as:

1. Run $\mathcal{A}^1$ to compute allocation x' and prices P'. Compute $Q = q(v, b, x', P')$, possibly based on the execution of $\mathcal{A}^1$.
2. Run $\mathcal{A}^2_Q$ to compute new prices and allocations x'' and P''.
3. Output the allocation and price pair for each bidder, either (x', P') or (x'', P''), that minimizes the bidder's utility.

THEOREM 4.3. *If winners of $\mathcal{A}^1$ cannot affect Q without receiving non-positive utility and $\mathcal{A}^1$ and $\mathcal{A}^2_Q$ are truthful, then the composition is truthful.*

This formulation generalizes two concepts in the auction literature. First, it generalizes the cancelable auction of [5], where the auction is canceled if not enough revenue is raised. To instantiate the auction from [5], consider Q is the revenue of $\mathcal{A}^1$ and $\mathcal{A}^2_Q$ is the auction that sells the same number of items at the same price if $Q \geq R$ and rejects all bidders if $Q < R$. Second, composition as described in [1] is generalized by setting Q to be the winners of $\mathcal{A}^1$. Then, $\mathcal{A}^2_Q$ is computed over the previous winners.

In the current setting, we use generalized composition to create a type of cancelable auction for the MUBC problem. The auction $\mathcal{A}^1$ is defined to be the profit extract partition auction. Suppose the auction raises some amount of profit from each side of the partition, say T' and T'' from the respective partitions (we note that either one of the partitions has zero profit or $T' = T''$). We define $Q = max(T', T'')$. For some predefined amount of revenue C that must be raised, the auction $\mathcal{A}^2_Q$ either cancels the auction if $Q < C$ and outputs no winners or, if $Q \geq C$, then it outputs the same set of winners and the same price and allocation vectors as produced by $\mathcal{A}^1$.

THEOREM 4.4. *For $\mathcal{A}^1$ and $\mathcal{A}^2_Q$ as described above, their generalized composition is truthful.*

Proof. It suffices to show that a bidder receiving positive utility before $\mathcal{A}^2_Q$ is run cannot alter its bid to increase Q and still incur positive utility. For a bidder to increase Q and still have positive utility, it must increase the revenue of the partition to which it belongs. Since the price offered the bidder's partition is independent of the bidder's bid, it can only increase revenue by purchasing more items. But either there are no items left to purchase or doing so would exceed its budget, resulting in a bidder that no longer incurs positive utility. ∎

This procedure is simple and can easily be applied to other auctions such as RSPE [8]. Unlike the tie-breaking procedure in [5] [8], this approach does not necessarily require half the revenue is forfeited when the revenue from both partitions is equal.

5 An Impossibility Result: Competitive Ratio Lower Bound

The following theorem says that we cannot hope for any auction (even a randomized auction) to have competitive ratio strictly lower than 2 against $T(v, b, m)$ for $\alpha = 2$. The theorem is proved by showing that for any truthful randomized auction, there exists a pair of input vectors u and b with $\alpha \geq 2$ such that $\frac{T(v,b,m)}{E_A[A(v,b,m)]} \geq 2 - \epsilon$.

THEOREM 5.1. *For any truthful randomized auction A, the competitive ratio against the optimal omniscient multi-price auction when $\alpha \geq 2$ is at least $2 - \epsilon$.*

Proof. We first define a distribution over input vectors for which $\frac{E_{(v,b,m)}[T(v,b,m)]}{E_{(v,b,m)}[E_A[A(v,b,m)]]} \geq 2 - \epsilon$. As proved in [8], this implies that for every randomized auction A, there exists a particular triple (v, b, m) for which $\frac{T(v,b,m)}{E_A[A(v,b,m)]} \geq 2 - \epsilon$. The distribution is over (v, b, m) such that:

- The number of bidders is n and the auctioneer has exactly n items to distribute.
- The value of a bidder equals its budget.
- All values (and therefore budgets) are either 1 or c, where c is a predefined constant.
- There are two bidders with value c. So, the value and budget vectors are $(c, c, y_1, .., y_{n-2})$, where y_i is a random variable.
- $\forall i, 1 \leq i \leq n-2$, $pr(y_i = 1) = \frac{c-1}{c}$ and $pr(y_i = c) = \frac{1}{c}$. The outcome for each y_i is independent.

For this distribution over y, $E[y] = 2 - \frac{1}{c}$. We would like a upper bound on the expected profit A can obtain. The proof relies on the price offered to bidder i being independent of v_i and b_i for any truthful auction as shown in [2].

Suppose that for a given input our randomized auction A has chosen prices $P = (P_1, P_2, ...P_n)$ and item allocation $a = (a_1, a_2, ..., a_n)$. We perform several modifications of the vectors P and a that can only increase the revenue of auction A. We create a modified price vector for A of $(c, c, ..., P_n)$. The modified price vector can only increase revenue since whatever profit was obtained previously can be obtained with the new prices using less items. Furthermore, no other item can gain revenue more than c, so any revenue maximizing algorithm will sell exactly 1 item to bidder 1 and exactly 1 item to bidder 2, since this is the best value possible per item. We modify the allocation vector accordingly to have $a_1 = 1$ and $a_2 = 1$. For any price $P_i > 1$ we modify the price P_i to be equal to c and $a_i = 1$. Again, this only increases the auction's expected revenue since originally P_i will be rejected if the (v_i, b_i) input is $(1, 1)$ and will have revenue $min(P_i, c)$ if the input is (c, c). After modification, the price will again be rejected for input $(1, 1)$ but will have revenue c expending only 1 item (which is more profitable than $min(P_i, c)$ for a_i items). Also, if $P_i = 1$, then modifying the auction to offer i only 1 item cannot decrease the auction's revenue since i's budget could only afford 1 item anyway. We are left with all values in P either 1 or c and all values in a equal to 1 and have only increased auction A's revenue. Since v_i, b_i are each chosen independently, the expected revenue from i given the algorithm has chosen $P_i = 1$ is $\frac{c-1}{c}$ and the expected revenue given the algorithm has chosen $P_i = c$ is $c \times \frac{1}{c} = 1$. Therefore, $E_{(v,b,m)}[E_A[A(v, b, m)]] \leq 2c + n$.

For the optimal omniscient multi-price auction, the expected revenue is the expected sum of the utilities (or budgets). We have $E_{(v,b,m)}[T(v, b, m)] = 2c + (2 - \frac{1}{c})n$. Combining this with above inequality, $\frac{E_{(v,b,m)}[T(v,b,m)]}{E_{(v,b,m)}[E_A[A(v,b,m)]]} \geq \frac{2c+(2-\frac{1}{c})n}{2c+n} \geq 2 - \epsilon$, for large c and n. Applying the Lemma from [8], we have the theorem. ∎

This bound is less than 2, so it does not give a lower bound on the competitive ratio compared with the best single-price auction. If this bound could be improved above 2, this would answer the currently open question as to whether there exists an auction that raises as much revenue as the optimal omniscient single-price auction.

6 Masking Auction: Advance Knowledge Benefits

In this section we show that advance knowledge about the parameter α can lead to significant benefits in revenue. In previous work on MUBC auctions, although the parameter α is used in performance analysis, its value does not influence the execution of the algorithm. An auction based on α calculated from the input is not necessarily truthful. However, it may be possible to determine a useful bound for α based on market research or other external information.

We will first describe an auction that does well if the number of items is flexible and α is known in advance. Then, we will show how this auction, assuming $\alpha \geq 5.828$, can be used as a building block for an auction that never oversells the items, does well for small values of α, and approaches the optimal quickly as α increases.

Masking Auction

1. For every bidder i, compute $F_{-i} = F(v_{-i}, b_{-i}, m)$ and set $p_i = \frac{F_{-i}}{m}$.
2. For bidders i with $v_i > p_i$, each of $\frac{b_i}{p_i}$ items are allocated independently with probability $\frac{\alpha-1}{\alpha+1}$ at price p_i per item.
3. If $\sum_{i:v_i>p_i} \frac{b_i}{p_i} \geq m$, end the auction. Otherwise, allocate $m - \sum_{i:v_i>p_i} \frac{b_i}{p_i}$ items to bidders with $v_i = p_i$. Allocate arbitrarily, not exceeding budgets, at price p_i per item.

Theorem 6.1. *The masking auction is truthful.*

Proof. There are three cases.

Case 1 ($t_i^v > p_i$): Bidder i has positive utility if he receives items, and receives the same number of items at the same price with the same probability, regardless of what value above p_i is reported. Therefore, bidder i is indifferent between any $v_i > p_i$, and truthfulness is a dominant strategy for reporting v_i. If i receives items, $b_i < t_i^b$ decreases bidder i's expected utility. Reporting $b_i > t_i^b$ results in negative infinity utility since there is a non-zero probability of being charged b_i. Therefore, bidding t_i^b is a dominant strategy.

Case 2 ($t_i^v = p_i$): Bidder i has zero utility for any allocation that does not exceed the budget. Since bidding the budget truthfully ensures the bidder does not go over budget, truthful reporting of both budget and value is a dominant strategy.

Case 3 ($t_i^v < p_i$): Bidder i has negative utility if he receives items and zero otherwise. If $v_i = t_i^v$, bidder i will not receive items and maximize his utility. ∎

Theorem 6.2. *The masking auction has competitive ratio $\frac{(\alpha+1)(\alpha)}{(\alpha-1)^2}$ and never allocates more than $\frac{\alpha+1}{\alpha-1}m$ items.*

Proof. Here, $F = F(v, b, m)$ and $k = k(v, b, m)$, the revenue and stopping bidder, respectively, of the optimal single-price omniscient auction. It is first shown that $\sum_{i:v_i>p_i} b_i \leq F + b_{max}$ by proving if $v_i > p_i$, then $i \leq k$. Assume $v_i > p_i$ and $i > k$. Then masking i does not affect the optimal revenue and $mv_i > mp_i = F$. F equals either $\sum_{i=1}^k b_i$ or mv_k. If $F = \sum_{i=1}^k b_i$, then $mv_i > \sum_{i=1}^k b_i$, which contradicts the definition of k. If $F = mv_k$, then $v_i > v_k$ and $i > k$, which contradicts the definition of v (as sorted in decreasing order). Therefore, every bidder i with $v_i > p_i$ has $i \leq k$ and $\sum_{i:v_i>p_i} b_i \leq \sum_{i=1}^k b_i \leq F + b_{max}$. There are two cases in the auction, we use M to denote the random variable that is the number of items sold by the Masking Auction:

- **CASE I ($\sum_{i:v_i>p_i} \frac{b_i}{p_i} \geq m$):** Then $E[M] = \frac{\alpha-1}{\alpha+1}\sum_{i:v_i>p_i} \frac{b_i}{p_i} \geq \frac{\alpha-1}{\alpha+1}m$. We also have

$$\begin{aligned} E[M] &= \frac{\alpha-1}{\alpha+1}\sum_{i:v_i>p_i}\frac{b_i}{p_i} = \frac{(\alpha-1)m}{\alpha+1}\sum_{i:v_i>p_i}\frac{b_i}{F_{-i}} \\ &\leq \frac{(\alpha-1)m}{(\alpha+1)(F-b_{max})}\sum_{i:v_i>p_i} b_i \\ &\leq \frac{m(\alpha-1)(F+b_{max})}{(\alpha+1)(F-b_{max})} \\ &\leq \frac{m(\alpha-1)(\alpha+1)}{(\alpha+1)(\alpha-1)} \\ &= m \end{aligned}$$

  The last inequalities are due to the assumption that $b_{max} = \frac{F}{\alpha}$.

- **CASE II ($\sum_{i:v_i>p_i} \frac{b_i}{p_i} < m$):** It is first shown that all items in Step 3 of the auction are allocated. Because $v_i \geq p_i$ for all bidders that win in the optimal omniscient solution, $m \leq \sum_{i:v_i>p_i} \frac{m\bar{b}_i}{F} + \sum_{i:v_i=p_i} \frac{m\bar{b}_i}{F} \leq \sum_{i:v_i>p_i} \frac{b_i}{p_i} + \sum_{i:v_i=p_i} \frac{b_i}{p_i}$. This implies $m - \sum_{i:v_i>p_i} \frac{b_i}{p_i} \leq \sum_{i:v_i=p_i} \frac{b_i}{p_i}$. Therefore, the expected number of items allocated is $E[M] = \frac{\alpha-1}{\alpha+1}\sum_{i:v_i>p_i} \frac{b_i}{p_i} + m - \sum_{i:v_i>p_i} \frac{b_i}{p_i}$ and $\frac{\alpha-1}{\alpha+1}m \leq E[M] \leq m$.

For all i, $p_i = \frac{F_{-i}}{m} \geq \frac{F-b_{max}}{m} \geq \frac{F-\frac{F}{\alpha}}{m} = (\frac{\alpha-1}{\alpha})(\frac{F}{m})$. Therefore, the expected revenue in both cases is $E[R] \geq (\frac{\alpha-1}{\alpha+1})(m)(\frac{\alpha-1}{\alpha})(\frac{F}{m}) = \frac{(\alpha-1)^2}{(\alpha+1)(\alpha)}F$. In the worst case, the auction will never allocate more than $\frac{\alpha+1}{\alpha-1}m$ items since $\sum_{i:v_i>p_i} \frac{b_i}{p_i} \leq \frac{\alpha+1}{\alpha-1}m$ due to the arguments above. ∎

Finally, we use the Masking Auction to obtain the main theorem for this section.

THEOREM 6.3. *Assuming $\alpha \geq 5.828$, the masking auction using xm items, where $x = \frac{\alpha-1+\sqrt{(\alpha-1)^2-4\alpha}}{2\alpha}$, has competitive ratio $\frac{(x\alpha+1)(\alpha)}{(x\alpha-1)^2}$ and allocates at most m items.*

Proof. If xm items are sold, the bidder dominance is revised to $\hat{\alpha} \geq x\alpha$ and the revenue earned is $\hat{F} = F(v, b, xm) \geq xF$. Using Theorem 6.2, and plugging in above inequalities gives the theorem.

7 Conclusions

In Table 1 below, we list properties for three auctions designed for the MUBC auction problem: the asymptotically optimal auction from [2], the profit extract partition auction in Section 4, and the masking auction in Section 6. The performance of the asymptotically optimal auction in [2] is formulated as a bound on the tail of the distribution. Since we would like to compare with our performance, which is the expected competitive ratio, we consider the expected inverse of the competitive ratio for the auction from [2] to be $(1-\delta)(1-2exp(-\frac{\alpha\delta^2}{36}))$, where given a particular choice of α, δ is chosen to maximize the expectation.

We also graph the behavior of the three auctions as a function of the bidder dominance in Figure 1. The figure highlights the advantage of knowing α in advance, before observing bid vectors, since it would allow the auctioneer to choose the auction that gives it the most revenue for that α, and provides the option of using the highly competitive masking auction.

Acknowledgments I would like to thank Serge Plotkin, Tim Roughgarden, Bob McGrew, and Jason Hartline for helpful discussions. The notion of generalized composability in Section 4.1 is due to Jason Hartline.

References

[1] G. Aggarwal and J. Hartline. Knapsack Auctions. SODA 2006.

[2] C. Borgs, J. Chayes, N. Immorlica, M. Mahdian, and A. Saberi. Multi-unit Auctions with Budget-constrained Bidders. Electronic commerce 2005.

[3] Y. Che and J. Gale. Expected revenue of the all-pay auctions and first-price sealed-bid auctions with budget constraints. Economics Letters, 50:373-380, 1996.

[4] Y. Che and J. Gale. The optimal mechanism for selling to a budget-constrained buyer. Journal of Economic Theory, 92:198-233, 2000.

[5] A. Fiat, A. Goldberg, J. Hartline, and A. Karlin. Competitive Generalized Auctions. STOC, 2002.

[6] A. Goldberg, J. Hartline, A. Karlin, A. Wright, and M. Saks. Competitive Auctions. Games and Economic Behavior, 2002.

[7] A. Goldberg, J. Hartline, and A. Wright. Competitive Auctions and Digital Goods. SODA, 2001.

Auction	Competitive Ratio Inverse	Comments
Asymptotically Optimal	$max_\delta((1-\delta)(1-2exp(-\frac{\alpha\delta^2}{36})))$	· competitive ratio inverse asymptotically approaches 1 · for $\alpha \leq 24$, no known bound on the competitive ratio · sells exactly m items
Profit Extract Partition	$\frac{(\alpha-1)}{4\alpha}$	· asymptotically approaches competitive ratio 4 · $\forall\alpha \geq 10$, competitive ratio is less than 4.5 · sells exactly m items
Masking	$\frac{(x\alpha-1)^2}{\alpha(x\alpha+1)}$ $x = \frac{\alpha-1+\sqrt{(\alpha-1)^2-4\alpha}}{2\alpha}$	· competitive ratio inverse asymptotically approaches 1 · $\forall\alpha \geq 10$, competitive ratio is less than 2 · parameterized by α · may undersell items

Table 1: Properties of three auctions for the MUBC auction problem.

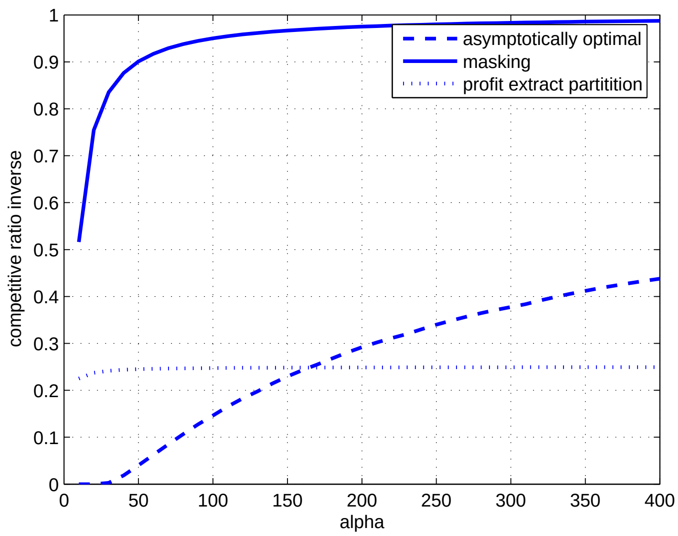

Figure 1: The inverse of the competitive ratio as a function of α.

[8] J. Hartline. Optimization in the Private Value Model: Competitive Analysis Applied to Auction Design. PhD Thesis, 2003.

[9] J.-J. Laffont and J. Robert. Optimal auctions with financially constrained buyers. Economics Letters, 52:181-186, 1996.

[10] E.S. Maskin. Auctions, development and privatization: Efficient auctions with liquidity-constrained buyers. European Economic Review, 44:667-681, 2000.

[11] H. Moulin, and S. Shenker. Strategyproof Sharing of Submodular Costs: Budget Balance vs. Efficiency. Economic Theory, 18:511-533, 2001.

[12] A. Mu'alem, and N. Nisan. Truthful Approximation Mechanisms for Restricted Combinatorial Auctions. AAAI 2002.

[13] Federal Communications Commission. Auctions. http://wireless.fcc.gov/auctions 2004.