Multi-unit auctions with budget-constrained bidders

Christian Borgs* Jennifer Chayes* Nicole Immorlica†
Mohammad Mahdian* Amin Saberi*

Abstract

We study a multi-unit auction with multiple agents, each of whom has a private valuation and budget. The truthful mechanisms of such an auction are characterized, in the sense that, under standard assumptions, we prove that it is impossible to design a non-trivial truthful auction which allocates all items, while we provide the design of an asymptotically revenue-maximizing truthful mechanism which may allocate only some of the items. Our asymptotic parameter is a *budget dominance parameter* which measures the size of the budget of a single agent relative to the maximum revenue. We discuss the relevance of these results for the design of online ad auctions.

1 Introduction

Budget constraints are a central feature of many real auctions. In the context of e-commerce, there is a great deal of interest in multi-item auctions of relatively low-value goods, such as the auction of online ads for search terms and content pages on MSN, Google, Yahoo, etc., to bidders with budget constraints. Indeed, it is widely believed that advertising will be the principal business model for online activity, and that budget-constrained auctions will be the primary means of realizing that revenue stream. Auctions with budget constraints have been considered previously in the context of privatization of high-value public goods, such as FCC auctions of telecommunications bands [1, 2, 7, 8]. However, the theoretical framework of budget-constrained auctions is currently substantially less well-developed than that of unconstrained auctions – which is unsatisfactory both from a theoretical viewpoint, and from a practical viewpoint, where the absence of an an appropriate framework leads to losses in revenue and efficiency. It is therefore of tremendous interest to design an incentive-compatible allocation (i.e., truthful) mechanism for budget-constrained auctions, and indeed, to determine the circumstances under which such a mechanism even exists.

In this paper, we consider the problem of a multi-unit auction with multiple bidders, each of whom has private valuation and budget. We prove both an impossibility result and a constructive, positive result. Throughout the paper, we assume the very natural conditions of enforcement of supply limits, individual rationality, and incentive compatibility (see Section 2 for definitions of these terms).

Existence of an incentive-compatible mechanism for budget-constrained bidders is a technically non-trivial problem. We assume that each bidder has a fixed private valuation and budget, such that if this budget is exceeded, then the bidder's total utility becomes unbounded below. We sometimes call this a "hard" budget constraint to distinguish it from "flexible" budget constraints considered by other authors [8], where the constraints can be exceeded under certain circumstances. Somewhat surprisingly, the well-known VCG mechanism is not truthful in the case with hard budgets – either with private hard budgets or in the *a priori* easier case with public hard budgets. This is an easy

*Microsoft Research; One Microsoft Way; Redmond, WA 98052. [borgs,jchayes,mahdian,saberi]@microsoft.com
†MIT CSAIL, 32G-596; 32 Vassar Street; Cambridge, MA 02139. nickle@theory.csail.mit.edu. Research conducted while author was visiting Microsoft Research in Redmond, WA.

consequence of the fact that the utilities are not quasi-linear, and will be demonstrated explicitly in Section 2.

Section 4 contains the proof of our impossibility result. We show that, under the assumption that the auction sells all items, then in the two-item, two-bidder case, the only incentive-compatible mechanism is the trivial bundling mechanism. In other words, there is no truthful mechanism that sells the items to distinct buyers. Our proof follows from a tedious, but straightforward analysis of the constraint equations. We note that, while the condition of selling all items seems quite restrictive, it turns out that in this case, it is implied by the often assumed condition of "independence of irrelevant alternatives" (see [**?**]).

Section 5 contains the proof of our principal result. There we relax the condition of selling all items, and construct an incentive-compatible mechanism which asymptotically achieves revenue maximization. We introduce the *budget dominance parameter*, defined to be the maximum budget of any single bidder divided by the optimal, omniscient revenue. As in the work of Goldberg *et. al.* [4, 5], we define the competitive ratio to be the ratio of the optimal revenue to the revenue of our mechanism. We then introduce an incentive-compatible mechanism for the general m-item, n-bidder auction, and prove that, as the budget dominance parameter tends to zero, the competitive ratio tends to one.

Our mechanism is inspired by lovely work of Goldberg *et. al.* [4, 5], but is different in several significant respects. First, our problem is a *two*-parameter problem, since each bidder specifies both a private valuation and a private budget. Second, our utility function is not quasi-linear. To our knowledge, there are no previously known incentive-compatible revenue-maximizing mechanisms in either of these cases. Finally, in the problem considered by [4, 5] at most one item was allocated to each agent, whereas our setting has no such restriction. Due to these differences, our proof is substantially more complicated than that of [4, 5], requiring delicate martingale arguments to achieve the necessary independence.

Let us also contrast our work with the some of previous work in the economics community which addressed Bayesian budget constraints [1, 2, 7, 8]. Che and Gale [1, 2] studied the single-item, single bidder case. Motivated by the goal of modelling efficient redistribution of public goods to the private sector, Maskin [8] studied the single-item, multiple bidder case with flexible budget constraints. The work closest to our context is that of Laffont and Robert [7], who treated a problem similar in some respects to the one studied here, but appropriate for a different set of applications. They too considered a two-parameter, non-quasi-linear, budget constrained problem. But they treated only single-item auctions with common public budgets. Moreover, whereas they examined the Bayesian problem, we consider *dominant strategy*. While their proposed mechanism, namely an all-pay auction, makes sense in the Bayesian context, it would not be appropriate as a dominant strategy, nor would it provide a reasonable mechanism in the case of online ad auctions.

2 Setting

We consider a setting in which an auctioneer has several indivisible units of a single good which he would like to auction off to n interested agents. Each agent i has a private utility $u_i \in \Re_+$ per unit of the good and a private budget constraint $b_i \in \Re_+$. We denote the vectors $(u_1, \dots, u_{i-1}, u_{i+1}, \dots, u_n)$ and $(b_1, \dots, b_{i-1}, b_{i+1}, \dots, b_n)$ by u_{-i} and b_{-i}, respectively. The budget constraint is a *hard* constraint, i.e., the agent cannot spend more than her budget under any circumstances. In other words, the total utility $u_i(j,p)$ that agent i derives from an allocation of j units at a total price of p is:

$$u_i(j,p) = \begin{cases} ju_i - p & \text{if } p \le b^i, \\ -\infty & \text{if } p > b^i. \end{cases}$$

The value $-\infty$ in the above definition means that this agent prefers receiving no item and paying nothing to any lottery with a non-zero risk of going over the budget.

An auction mechanism solicits a two-parameter bid from each agent. The first parameter is interpreted as that agent's announced utility per unit and the second parameter is that agent's announced budget. The mechanism then outputs an allocation and payment for each agent. We consider mechanisms that satisfy the following properties:

- *observe supply limits* – The mechanism never allocates more units than are available.
- *individual rationality* – An agent's utility from participating in the mechanism is non-negative.
- *incentive compatibility* or *truthfulness* – An agent's total utility is maximized by announcing his true utility and budget to the auction regardless of the strategies of the other agents.

These properties can be generalized for randomized mechanisms by replacing *utility* by *expected utility*. We call an auction mechanism satisfying the above properties a *truthful mechanism*. Notice that individual rationality and the definition of utility functions imply that a truthful mechanism never charges an agent an amount more than her budget.

One important distinction of the above setting compared to other models usually studied in the auction theory is that in our setting, the agents' utility functions are not *quasi-linear*. A quasi-linear utility function is a utility function of the form $u(x) - p$, where $u(x)$ is a function that only depends on the allocation and not on the payments, and p is the amount charged to the agent. The utility function $u_i(j, p)$ defined above cannot be written in this form. This makes much of the results in the auction theory inapplicable to our setting. In particular, the classical Vickrey-Clarke-Groves (VCG) mechanisms [9, 3, 6] are *not* incentive compatible in our setting. This fact is illustrated in the following example.

Example 1 A natural mechanism for auctioning m units of a good to budget-constrained buyers is to apply the VCG mechanism assuming that the utility of agent i for j units of the good is $\min(b_i, ju_i)$. A common mistake is to assume that since this mechanism is based on VCG, it is truthful. The following example shows that this is not the case: assume we have two units of the good to sell to two agents, and the bids of these agents are given by $(u_1, b_1) = (10, 10)$ and $(u_2, b_2) = (1, 10)$. The above mechanism assumes that the utility of the first agent for one unit of the good is 10, and therefore allocates one unit to each agent to maximize the total utility (which is $10 + 1$). The payment charged to the agents by this mechanism is 1 and 0. Therefore, the utility of the first agent is 9. However, if the first agent announces the bid $(5, 10)$, then the mechanism will allocate both items to this agent at a total price of 2. Thus, the first agent would achieve a utility of 18 by bidding untruthfully. This example shows that the above VCG-based mechanism is not truthful even if the agents are not allowed to lie about their budget.

It is easy to observe that in our setting, no truthful mechanism can always produce an *efficient* allocation, i.e., an allocation that maximizes the social welfare, even when there is only one good. The reason for this is that an efficient mechanism should always allocate the good to the bidder with the highest u_i, even if such a bidder has a zero budget and therefore cannot be charged any positive amount. Therefore, any agent can bid a high utility and zero budget to get the item for free. This simple impossibility result shows that we cannot require efficiency from a truthful mechanism.

3 Characterization

In this section, we give a simple price-based characterization of truthful auctions. It essentially claims that any truthful auction determines the allocation and price for agent i by comparing his bid to thresholds computed from the other agents' bids. The unconstrained budgets version of this proposition is a well-known folklore theorem.

Proposition 3.1 *For any truthful auction selling m units of a good to n agents, there exist mn functions $p_i^1, \ldots, p_i^m : \Re_+^{2(n-1)} \to \Re_+ \cup \{\infty\}$ such that agent i receives j units at price $p_i^j(u_{-i}, b_{-i})$ where j maximizes $ju_i - p_i^j(u_{-i}, b_{-i})$ subject to $p_i^j(u_{-i}, b_{-i}) \leq b_i$.*

Proof. For any $(u_{-i}, b_{-i}) \in \Re_+^{2(n-1)}$ and $j \in \{1, \ldots, m\}$, we define $p_i^j(u_{-i}, b_{-i})$ as the minimum, over the choice of (u_i, b_i) such that the auction allocates at least j items to i if agents bid (u, b), of the price that the mechanism charges to i at these bids. Let j^* be an index that maximizes $j^* u_i - p_i^{j^*}(u_{-i}, b_{-i})$ subject to $p_i^{j^*}(u_{-i}, b_{-i}) \leq b_i$. If when agents bid (u, b), the mechanism allocates j units to i at price p, then we must have $ju_i - p = j^* u_i - p_i^{j^*}(u_{-i}, b_{-i})$, since otherwise agent i would have an incentive to bid untruthfully to get j^* items at price $p_i^{j^*}(u_{-i}, b_{-i})$. □

By considering all cases for the relationship between the p_i^j's, the auction can be expressed as a concise set of inequalities. This is done for the case of two units of good and two buyers in the following corollary. We will use this corollary in the next section to prove that truthful mechanisms satisfying certain properties do not exist.

Corollary 3.1 *For any deterministic truthful auction selling 2 units of a good to 2 agents, there exist threshold functions $p_i^j : \Re_+^2 \to \Re_+ \cup \{\infty\}$, $1 \leq i, j \leq 2$, such that for $i = 1, 2$, the agent i receives*

- 2 *units at a total price of $p_i^2(u_{3-i}, b_{3-i})$ if $b_i \geq p_i^2(u_{3-i}, b_{3-i})$ and $u_i > p_i^2(u_{3-i}, b_{3-i}) - \min(p_i^1(u_{3-i}, b_{3-i}), p_i^2(u_{3-i}, b_{3-i})/2)$ (or if the latter inequality holds with equality, the mechanism can choose to allocate 2 units to i);*
- *else* 1 *unit at price $p_i^1(u_{3-i}, b_{3-i})$ if $b_i \geq p_i^1(u_{3-i}, b_{3-i})$ and $u_i > p_i^1(u_{3-i}, b_{3-i})$ (or if the latter inequality holds with equality, the mechanism can choose to allocate 1 units to i);*
- *else* 0 *units.*

Conversely, for any set of threshold function $p_i^j : \Re_+^2 \to \Re_+ \cup \{\infty\}$, $1 \leq i, j \leq 2$, the mechanism defined above satisfies incentive compatibility and individual rationality.

Proof. We prove the statement for $i = 1$ ($i = 2$ is analogous). Consider the threshold functions given by Proposition 3.1. Fix any bid (u_2, b_2) of the second agent. Suppose the true utility and budget of the first agent is u_1 and b_1, respectively. For simplicity, we use the notation $p_1^1 := p_1^1(u_2, b_2)$ and $p_1^2 := p_1^2(u_2, b_2)$. Notice that by the definition of p_1^1 and p_1^2 in the proof of Proposition 3.1, $p_1^1 \leq p_1^2$. The first agent's utility for an allocation of 0 units is 0, 1 unit is $u_1 - p_1^1$ assuming $b_1 \geq p_1^1$, and 2 units is $2u_1 - p_1^2$ assuming $b_1 \geq p_1^2$. The first agent receives two units if and only if she has enough budget to pay for it (i.e., $b_1 \geq p_1^2$), and her utility for receiving two items ($2u_1 - p_1^2$) is greater than or equal to her utility for receiving one item ($u_1 - p_1^1$) and zero item (zero). This can be written as $u_1 \geq p_1^2 - p_1^1$ and $u_1 \geq p_1^2/2$, or equivalently, $u_1 \geq p_1^2 - \min(p_1^1, p_1^2/2)$. Otherwise, if the first agent does not receive two items, then she receives one item if and only if she has the budget (i.e., $b_1 \geq p_1^1$), and her utility for one item ($u_1 - p_1^1$) is greater than or equal to her utility for zero items, or equivalently, $u_1 \geq p_1^1$. If these conditions do not hold, then the agent receives zero units. The converse follows easily from the definition of the mechanism. □

4 An Impossibility Result

In this section, we show that there is no truthful mechanism satisfying three properties which we define in this section, even if there are only two buyers and two units of the good. This result automatically generalizes to auctions with more buyers, by considering the situation where all but two of the buyers bid zero.

The first property is the following. This is similar to a property with the same name defined by Moulin [**?**] in the context of group-strategyproof mechanisms for cost sharing problems.

- *consumer sovereignty* – For any agent i and any vector of bids (u_{-i}, b_{-i}) for other agents, there is a bid (u_i, b_i) such that if agents bid according to (u, b), then agent i receives all units of the item.

Intuitively, consumer sovereignty requires that each agent must be able to win all items if she bids *high enough*. This precludes trivial mechanisms that for example sell at most one item to each bidder. In terms of the characterization in Proposition 3.1 and Corollary 3.1, this property is equivalent to saying that the threshold functions p_i^j are all finite.

The second property, which we call the *independence of irrelevant alternatives* (IIA), is a much weaker version of a property of the same name in Lavi et al. [**?**]. This property is defined as follows.

- *independence of irrelevant alternatives (IIA)* – For any agent i and a bid vector (u, b), if i receives no item at (u, b), then the allocation when every agent bids according to (u, b) is the same as the allocation when agent i bids $(0, 0)$ and others bid according to (u_{-i}, b_{-i}).

Intuitively, the above property states that if an agent who does not win the auction leaves, the allocation to other agents should not change (Their payment, however, might change). As we will see in the proof of Theorem 1, in the case of two buyers and two items, IIA is equivalent to the property that if bids of both agents are large enough (both the utility and the budget), then both units are allocated.

As we will see at the end of this section, there are truthful mechanisms not satisfying the IIA. In fact, the following example shows that even with IIA, there are mechanisms that are truthful.

Example 2 *Bundling mechanism:* Consider the mechanism that always *bundles* the two units, i.e., it allocates both items to the agent i such that $\min(2u_i, b_i) > \min(2u_{3-i}, b_{3-i})$, and charges her $\min(2u_{3-i}, b_{3-i})$. It is easy to see that this mechanism is truthful and satisfies the IIA.

However, we conjecture that the bundling mechanism is essentially the only truthful mechanism satisfying the above properties. In other words, we would like to show that there is no truthful mechanism satisfying the above properties and the following.

- *non-bundling* – there is a bid vector (u, b) such that the mechanism allocates one unit of the good to each buyer.

Unfortunately, we do not know how to prove this conjecture. However, we can prove this statement under the following stronger condition.

- *strong non-bundling* – for any bid (u_1, b_1) of the first agent, there is a bid (u_2, b_2) for the second agent such that if both agents bid according to (u, b), the mechanism allocates one unit of the good to each buyer.

The following theorem is the main result of this section.

Theorem 1 *There is no truthful auction for two buyers and two units of a good that satisfies consumer sovereignty, IIA, and strong non-bundling.*

Proof. The proof of Theorem 1 examines functional relations imposed by our assumptions on the threshold functions of any truthful auction. We obtain the impossibility result by showing that this set of functional relations has no solution.

The fact that our auction observes supply limits implies that whenever the threshold functions are such that the first (second) agent gets two items, then the second (first) agent must get zero items. The consumer sovereignty and IIA assumptions imply that these two situations are in fact equivalent in certain regions of the bid space, i.e. the mechanism always allocates all the units when the bids are large enough.

By consumer sovereignty, for each agent $i = 1, 2$, there is a bid (u_i^*, b_i^*) such that if i bids (u_i^*, b_i^*) and the other agent bids $(0, 0)$, then agent i wins both items. Furthermore, by Corollary 3.1, for every $u_i' \geq u_i^*$ and $b_i' \geq b_i^*$, if i bids (u_i', b_i') and the other agent bids $(0, 0)$, then i wins both units. Let $C = \max\{u_1^*, b_1^*, u_2^*, b_2^*\}$.

Claim 1 *For any set of bids (u_1, b_1) and (u_2, b_2) such that $u_1, b_1, u_2, b_2 \geq C$, the mechanism allocates both items when agents bid according to (u, b).*

Proof. Assume, for contradiction, that for one such bid vector the mechanism allocates at most one unit of the good to the first agent and zero units to the second agent. Now, by IIA, if the second agent bids $(0, 0)$, the first agent must still receive at most one item. This, however, contradicts the definition of C. □

Immediate from Corollary 3.1 is the fact that the allocations and payments given bid (α_i, β_i) holding bid $(\alpha_{3-i}, \beta_{3-i})$ fixed is constant for all $\alpha_i \geq 2\beta_i$ and for all $\alpha_i \geq \beta_i$. We will use this observation to make statements about the properties of the threshold functions as one of the inputs becomes irrelevant (i.e. sufficiently large). Let

$$\begin{aligned} r_i^j(x) &= p_i^j(x, 2x), \\ s_i^j(x) &= p_i^j(x, x) \end{aligned}$$

for $i, j = 1, 2$. By Corollary 3.1, all of the above functions are non-decreasing functions. Therefore, they can be discontinuous in at most a countable number of points. Let T denote the set of numbers more than C at which all of the above functions are continuous. Notice that this set is dense.

Claim 1 together with our characterization, Corollary 3.1, immediately imply the following functional relations:

Lemma 4.1 *For all $A, B \in T$,*

$$B < r_{3-i}^1(A) \Rightarrow A \geq (s_i^2 - \min(s_i^1, s_i^2/2))(B) \tag{1}$$

Proof. Suppose agent i bids $(A, 2A)$ and agent $(3-i)$ bids (B, B) and $B < r_{3-i}^1(A)$. Then agent $(3-i)$ receives zero units, so agent i must receive two units. As agent i's budget is essentially unconstrained, this implies that his utility is at least the utility threshold, or $A \geq (s_i^2 - \min(s_i^1, s_i^2/2))(B)$. □

Similarly, we can prove the following equations:

$$\forall\ A, B \in T,\ B < r_{3-i}^1(A) \Leftarrow A > (s_i^2 - \min(s_i^1, s_i^2/2))(B), \tag{2}$$

$$\forall\ A, B \in T,\ B \geq r_{3-i}^2(A) \Rightarrow A \leq \min(s_i^1, s_i^2/2)(B), \tag{3}$$

$$\forall\ A, B \in T,\ B \geq r_{3-i}^2(A) \Leftarrow A < \min(s_i^1, s_i^2/2)(B), \tag{4}$$

$$\forall\ A, B \in T, B > (r_{3-i}^2 - \min(r_{3-i}^1, r_{3-1}^2/2))(A) \Rightarrow A \leq \min(r_i^1, r_i^2/2)(B), \tag{5}$$

$$\forall\ A, B \in T, B \geq (r_{3-i}^2 - \min(r_{3-i}^1, r_{3-1}^2/2))(A) \Leftarrow A < \min(r_i^1, r_i^2/2)(B), \tag{6}$$

$$\forall\ A, B \in T, B \geq s_{3-i}^2(A) \iff A < s_i^1(B). \tag{7}$$

From these functional relations, we can derive the following inequalities.

Lemma 4.2 *For all $A \in T$,*

$$(r_2^2 - \min(r_2^1, r_2^2/2))(A) \geq (s_2^2 - \min(s_2^1, s_2^2/2))(A). \tag{8}$$

Proof. Choose $B \in T$, $B > (r_2^2 - \min(r_2^1, r_2^2/2))(A)$. Then relation 5 (with $i = 1$) implies $A \leq \min(r_1^1, r_1^2/2)(B) \leq r_1^1(B)$. Take $\epsilon > 0$ and note that relation 1 (with $i = 2$) implies $B > (s_2^2 - \min(s_2^1, s_2^2/2))(A - \epsilon)$. Taking the limit as ϵ goes to zero and using the continuity of s_2^1 and s_2^2 at A, we have that $B > (r_2^2 - \min(r_2^1, r_2^2/2))(A)$ implies $B \geq (s_2^2 - \min(s_2^1, s_2^2/2))(A)$. The lemma follows. □

Similarly, we can prove:

$$\forall\ A \in T, \min(r_2^1, r_2^2/2)(A) \geq \min(s_2^1, s_2^2/2)(A). \tag{9}$$

Our non-bundling assumption implies that for all $Z \in T$ the interval $(r_2^1(Z), r_2^2(Z))$ is non-empty. Select $t \in (r_2^1(Z), r_2^2(Z))$ and observe that the contrapositive of relations 2 (with $i = 1$) implies $Z \leq (s_1^2 - \min(s_1^1, s_1^2/2))(t)$. Let $\epsilon > 0$ and not that $t \in (r_2^1(Z - \epsilon), r_2^2(Z - \epsilon))$ for small enough ϵ by continuity. Thus the contrapositive of relation 4 with $i = 1$ implies $Z - \epsilon > \min(s_1^1, s_1^2/2)(t)$. Combining the two inequalities, we get $(s_1^2 - \min(s_1^1, s_1^2/2))(t) > \min(s_1^1, s_1^2/2)(t)$ and so

$$\forall\ Z \in T,\ \forall\ t \in (r_2^1(Z), r_2^2(Z)),\ s_1^1(t) \leq s_1^2(t)/2. \tag{10}$$

This fact together with relations 4 and 7 with $i = 1$ imply that

$$\forall\ Z \in T,\ \forall\ t \in (r_2^1(Z), r_2^2(Z)),\ t \geq r_2^2(Z) \iff t \geq s_2^2(Z). \tag{11}$$

Fix $Z \in T$ and consider $t = r_2^2(Z) - \epsilon$ for a small $\epsilon > 0$. If ϵ is small enough, $t \in (r_2^1(Z), r_2^2(Z))$ and so $s_2^2(Z) \geq t$. Taking the limit as $\epsilon \to 0$, this implies that $r_2^2(Z) \leq s_2^2(Z)$. On the other hand, summing Equations 8 and 9 implies that $r_2^2(Z) \geq s_2^2(Z)$. Therefore, $r_2^2(Z) = s_2^2(Z)$. Therefore, inequalities 8 and 9 must both attain equality at Z. Ranging over choice of $Z \in T$, we see that inequalities 8 and 9 must attain equality everywhere in T. Our contradiction arises from the observation that in fact for some $Z \in T$, inequality 9 is strict. By consumer sovereignty, prices are always nonzero, and so $C < (r_1^2 - \min(r_1^1, r_1^2/2))(A) < r_1^2(A)$ for some $A \in T$. Select such an A and $Z \in ((r_1^2 - \min(r_1^1, r_1^2/2))(A), r_1^2(A)) \cap T$. Notice that since T is a dense set, this intersection is nonempty. Note that relation 4 with $i = 2$ implies that $A \geq \min(s_2^1, s_2^2/2)(Z)$. Therefore, for any small $\epsilon > 0$, $A + \epsilon > \min(s_2^1, s_2^2/2)(Z)$. Similarly, note that relation 5 with $i = 2$ implies $A + \epsilon \leq \min(r_2^1, r_2^2/2)(Z)$. But this means that, for this particular Z, $\min(r_2^1, r_2^2/2)(Z) > \min(s_2^1, s_2^2/2)(Z)$, yielding our contradiction. □

5 An asymptotically optimal auction

As we saw in the previous sections, there appears to be no reasonable mechanism for allocating all the goods truthfully. In this section, we consider mechanisms that may allocate only some of the goods, and among them, seek the one that maximizes the expected revenue. We do this by designing a mechanism that has revenue comparable to the best omniscient posted-price auction.

The method that we use for design and analysis of our auction is inspired by the work of Goldberg et. al. [4, 5]. As in [4, 5], we take the *competitive ratio* to be the expected revenue of our mechanism over the revenue of the optimum posted-price auction, and attempt to design an auction which minimizes this ratio. However, our design analysis differs from that of [4, 5] in some important aspects. Unlike [4, 5] and subsequent results [], in our setting the mechanism may need to allocate more than one good to every agent. Moreover, the agents can lie about both their bid and their budget, which introduces significant complications.

An important parameter for the budget-constrained problem is what we call the *budget dominance parameter*. For each agent, we define her budget parameter as the ratio of her budget to the revenue of the optimum posted-price auction. The budget dominance parameter is the maximum of the budget parameter of all the agents. The algorithm we design will have the property that the competitive ratio tends to 1 as the budget dominance parameter tends to 0.

5.1 The Algorithm

As before, let n be the number of agents and m be the number of indivisible goods. Each agent i submits her utility value for one item u_i and her maximum budget b_i.

Note, although our goods are indivisible, we can assume that fractional allocations are possible by using the proper randomization: whenever the algorithm asks us to allocate a fraction c of a good to an agent, we instead allocate a full good to her with probability c.

Algorithm

- Partition the agents randomly into two sets A and B by independently putting each agent into either set uniformly at random with probability $\frac{1}{2}$.
- From the set of utility values u_i of agents $i \in A$, choose p_A to be the price which maximizes the revenue of selling at most $m/2$ items in A. In other words, if the u_i's are sorted in a decreasing order, for

$$i_0 = \min_i \sum_{j=1}^{i-1} b_j \geq \frac{u_i m}{2},$$

  define $p_A = u_{i_0 - 1}$. Compute p_B analogously.
- Consider the agents in A in a random order. In every step, if the utility of agent i satisfies $u_i \geq p_B$, allocate $\frac{b_i}{p_B}$ goods to i. Continue this process until all $\frac{m}{2}$ goods have been allocated, or until all the agents in A have been processed. Apply the same procedure to the set B using the threshold value p_A.

In the rest of this section, we will analyze the above algorithm, and prove that it is truthful and that its competitive ratio tends to 1 as the budget dominance parameter tends to zero.

5.2 Analysis of the algorithm

In this section, we assume without loss of generality that our goods are divisible. Indeed, as explained above, this assumption can be easily removed by a few additional coin flips.

First, we give a simple proof of the truthfulness of the algorithm.

Lemma 5.1 *The above algorithm is truthful, i.e., for every agent reporting the correct utility and budget values, the algorithm is a dominant strategy.*

Proof. Consider an agent i in A. First we argue that agent i does not have any incentive to misreport her utility value. We know that agent i receives a good only if $u_i \geq p_B$, and that she pays p_B for a good if she receives it. The two key observations are that (1) the threshold p_B is determined independently of all u_j and b_j, including $j = i$; and (2) when the supply of goods in A is inadequate to meet the demands of all agents in A whose utilities exceed p_B, then the allocation of goods to those agents is done in an arbitrary order, again independently of all u_j and b_j.

Finally, by reporting a budget below b_i, agent i would potentially decrease her allocation and hence her total utility, whereas by reporting a budget above b_i, she has non-zero probability of decreasing her total utility to negative infinity. In either case, she has no incentive to misreport. □

The rest of this section is an analysis of the revenue achieved by our algorithm. For notational convenience, without loss of generality, we assume that $u_1 > u_2 > \cdots > u_n$.

For any price p, we denote by $r_S(p,k)$ the revenue of allocating at most k goods to a set S of agents at price p:

$$r_S(p,k) = \min(kp, \sum_{\substack{j \in S \\ u_j \geq p}} b_j).$$

We will use the notation $r(p,k)$ in the case where we are allocating the goods to the whole set $A \cup B$. Finally, we also define $r(p) = r(p,\infty) = \sum_{u_i \geq p} b_i$.

Given the utility and budget values of the agents, one can find the optimum price $p*$ at which $r(p,m)$ is maximized, and allocate the goods at this price. We call this mechanism the optimum posted-price auction $OPT = r(p^*, m)$. In our argument, we will use the following properties of an optimum posted-price auction for allocating at most k goods, for any k.

1. There exists an agent i such that selling the goods at price $p = u_i$ results in the optimum revenue.
2. For any k, if p is the optimum price for allocating at most k goods, then $r(p,k) \leq r(p) = r(p) + b_{max}$ where $b_{max} = \max_i b_i$. In particular,

$$OPT \leq r(p^*) \leq OPT + b_{max}.$$

Let ϵ denote the ratio of the maximum budget of all agents, b_{max}, to the value of the optimum solution OPT. In some sense, ϵ is a *budget dominance parameter* since it captures the extent to which the budget of a single agent dominates the market in the optimum solution. As we will show, the probability of success of our algorithm is asymptotically controlled by ϵ.

Lemma 5.2 *Let $\delta > 0$. Then the probability that*

$$\left| r_A(u_l) - r_B(u_l) \right| < \delta OPT \quad \text{for all} \quad l \quad \text{with} \quad u_l \leq p^*$$

is at least $1 - 2e^{-\delta^2/(4\epsilon)}$.

Proof. Let k be such that $u_k = p^*$. Define α_i to be a random variable indicating whether agent i is in A, with $\alpha_i = 1$ when $i \in A$ and $\alpha_i = -1$ when $i \in B$. Let $S_i = \sum_{j \leq i} \alpha_j b_j$. Then $|r_A(u_l) - r_B(u_l)| = |S_l|$. Thus we need to bound the probability that the random variable S_i deviates by more than δOPT from its expectation 0.

Let $\tau(\delta) = \min_i \{|S_i| \geq \delta OPT\}$. We define the following martingale:

$$\tilde{S}_i = \begin{cases} S_i & \text{if } i \leq \tau(\delta) \\ S_{\tau(\delta)} & \text{otherwise} \end{cases}$$

Then we have

$$\begin{aligned} 1 - \Pr\left(|r_A(u_i) - r_B(u_i)| < \delta OPT, \ \forall i \leq k\right) &= 1 - \Pr\left(|S_i| < \delta OPT, \ \forall i \leq k\right) \\ &= \Pr\left(\exists\, i \leq k \ : \ |S_i| \geq \delta OPT\right) \\ &= \Pr\left(\tau(\delta) \leq k\right) \\ &= \Pr\left(|\tilde{S}_i| \geq \delta OPT\right) \end{aligned}$$

Now since $\tilde{S}_i$ is a martingale, by the Azuma-Hoeffding inequality we have:

$$\Pr(|\tilde{S}_k| \geq \delta OPT) \leq 2\exp(\frac{-\delta^2 OPT^2}{2\sum_{i \leq k} b_i^2}).$$

Bounding the sum $\sum_{i\leq k} b_i^2$ by $b_{max}r(p^*) \leq \epsilon OPT r(p^*)$ and using that $r(p^*) \leq OPT(1+\epsilon) \leq 2OPT$, we obtain the lemma.

□

From now on, we will say that an event happens with high probability if its probability is at least $1 - 2e^{-\delta^2/(4\epsilon)}$.

Corollary 5.1 *With high probability,* $r_A(p^*, \frac{m}{2}) \geq \frac{1-\delta}{2} OPT$.

Proof. Let us first note that $OPT = r(p^*, m) = \min\{p^* m, r(p^*)\}$, implying that $mp^* \geq OPT$ and $r(p^*) \geq OPT$. With the help of Lemma 5.2, we conclude that with high probability, $r_A(p^*) \geq \frac{1-\delta}{2} r(p^*) \geq \frac{1-\delta}{2} OPT$. Inserting the definition of $r_A(p^*, \frac{m}{2})$ and using $\frac{m}{2} p^* \geq OPT$ we get the lemma. □

Corollary 5.2 *With high probability, we have that*

$$r_B(u_k) \geq \min\{r_A(u_k) - \delta OPT, \frac{1-\delta}{2} OPT\} \qquad \text{for all} \quad k.$$

Proof.

By Lemma 5.2, we have that with high probability,

$$r_B(u_k) \geq r_A(u_k) - \delta OPT \quad \text{for all } k \text{ with } u_k \leq p^*. \tag{12}$$

Recalling that $r_A(p^*) + r_B(p^*) = r(p^*) \geq OPT$, we conclude that with high probability, both (12) and

$$r_B(p^*) \geq \frac{1-\delta}{2} OPT$$

hold simultaneously. By monotonicity, this implies the statement of the lemma. Indeed, either $u_k \leq p^*$ so that $r_B(u_k) \geq r_A(u_k) - \delta OPT$, or $u_k \geq p^*$ and $r_B(u_k) \geq r_B(p^*) \geq \frac{1-\delta}{2} OPT$, which gives the lemma. □

Theorem 2 *The mechanism described in the previous section is truthful. Furthermore, for all* $0 < \delta < 1$*, the algorithm has revenue at least* $(1-\delta)OPT$ *with probability* $1 - O(e^{-c\delta^2/\epsilon})$ *with* $c = 1/36$ *and* $\epsilon = b_{max}/OPT$.

Proof.

For all p, we have $r_A(p, \frac{m}{2}) \leq r_A(p_A, \frac{m}{2})$, so in particular $r_A(p_A, \frac{m}{2}) \geq r_A(p^*, \frac{m}{2})$. Combined with the last corollary, we conclude that with high probability, $r_A(p_A, \frac{m}{2}) \geq \frac{1-\delta}{2} OPT$, which in turn implies that

$$p_A \frac{m}{2} \geq \frac{1-\delta}{2} OPT$$

and $r_A(p_A) \geq \frac{1-\delta}{2} OPT$. Combined with Corollary 5.2, the last inequality gives

$$r_B(p_A) \geq \frac{1-3\delta}{2} OPT,$$

again with high probability. Inserting the definition of $r_B(p_A, \frac{m}{2})$, we therefore have that with high probability,

$$r_B(p_A, \frac{m}{2}) \geq \frac{1-3\delta}{2} OPT.$$

Exchanging the roles of A and B, we get the same result for $r_A(p_B, \frac{m}{2})$. Since $r_B(p_A, \frac{m}{2}) + r_A(p_B, \frac{m}{2})$ is the revenue of the algorithm, this establishes the theorem.

□

Acknowledgments. We thank Rakesh Vohra for bringing several references to our attention, and for a delightful discussion in which he helped us to hone our economic sensibilities.

References

[1] Y. Chae 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.

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

[3] E. Clarke. Multipart pricing of public goods. *Public Choice*, 8:19–33, 1971.

[4] A. Goldberg, J. Hartline, A. Karlin, A. Wright, and M. Saks. Competitive auctions. *preprint*, 2003.

[5] A. Goldberg, J. Hartline, and A. Wright. Competitive auctions and digital goods. *Proc. 12th Symp. on Discrete Alg.*, pages 735–744, 2001.

[6] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.

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

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

[9] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.