Sort-Cut: A Pareto Optimal and Semi-Truthful Mechanism for Multi-Unit Auctions with Budget-Constrained Bidders

I. Hafalir * R. Ravi * A. Sayedi *

July 16, 2009

Abstract

Motivated by sponsored search auctions with hard budget constraints given by the advertisers, we study multi-unit auctions of a single item. An important example is a sponsored result slot for a keyword, with many units representing its inventory in a month, say. In this single-item multi-unit auction, each bidder has a private value for each unit, and a private budget which is the total amount of money she can spend in the auction. A recent impossibility result [Dobzinski et al., FOCS'08] precludes the existence of a truthful mechanism with Pareto-optimal allocations in this important setting. We propose Sort-Cut, a mechanism which does the next best thing from the auctioneer's point of view, that we term *semi-truthful*. While we are unable to give a complete characterization of equilibria for our mechanism, we prove that *some* equilibrium of the proposed mechanism optimizes the revenue over all Pareto-optimal mechanisms, and that this equilibrium is the unique one resulting from a natural rational bidding strategy (where every losing bidder bids at least her true value). Perhaps even more significantly, we show that the revenue of **every** equilibrium of our mechanism differs by at most the budget of one bidder from the optimum revenue (under some mild assumptions).

1 Introduction

While billions of dollars are spent on keyword-based advertising in the web, a majority of it is cleared through advertisement auctions for sponsored search results in search pages [BCI$^+$05]. When a typical advertiser goes to a typical search engine company to sign up to bid in such auctions, they specify the set of keywords whose search result pages they are interested in bidding for, with a bid value per click; they also specify a total (monthly) budget for the total amount they are willing to spend across all these keywords in this search company's site. We focus our attention on the budget constraint which we model as being a hard constraint (cannot be exceeded) and attempt to design an auction mechanism for clearing this market. However, to better understand the difficulty of the hard budget constraint, we simplify the other aspects by restricting ourself to the problem for a single keyword, and within the keyword search results, restrict our attention to a single sponsored search result slot (in contrast with the whole ladder of slots on the right of the results page). Thus, one may view this, e.g. as the auction for the single shaded sponsored search slot right under the search query window and above the organic search results in Google.com.

Our model abstracts the auction for this single item (one slot for a keyword) as a multi-unit auction where the number of units is the inventory of such search results in the period for which

*Tepper School of Business, Carnegie Mellon University {isaemin, ravi, ssayedir}@cmu.edu

the budget is specified. While this is not the way that current sponsored results are allocated (rather an auction is run every time a query is made), it is quite a plausible scenario that the industry may move to, especially as the market for such results mature, and a few major advertisers wish to plan for their internet advertising campaigns in much the same way as for other media advertising (such as the annual Fall market clearing event for prime-time TV advertising in New York). Our study represents the first step towards the design of such a market.

Multi-unit auctions have been studied comprehensively in microeconomics, especially in auction theory [AC95, Wol98, Nau95]. However, the problem of budget-constrained bidders has been paid surprisingly little attention, despite the fact that in practice, bidders face natural budget constraints. One reason the budget constraints have not received much attention is the traditional economist's view that such constraints are unnatural and that if an advertiser can make positive payoff at her current true valuation by participating in the auction, then every extra dollar spent by her in the auction should yield even more returns resulting in virtually no budget constraints. However, reality is different and the valuations announced do not scale forever. Furthermore, practical considerations (business planning) also put a cap on how much can be spent by advertisers in each period. Current keyword auctions get around this by treating each search page in the ad inventory as an instance of an online matching of current advertisers to slots, where the bids are assumed to be small compared to the budgets [MSVV07], and adjusting participation as the budget gets close to being spent.

Another reason for this lack of attention to the problem may be the technical difficulty that the utility of obtaining the items is compared to the total price at which the items are procured to give a net payoff. The total price is now discontinuously influenced by a hard budget constraint. Perhaps because of this, the theoretical framework of budget-constrained auctions is currently substantially less well-developed than that of unconstrained auctions. This is unsatisfactory both from a theoretical viewpoint, and from the practical viewpoint, where the absence of appropriate framework might potentially result in losses in revenue and efficiency.

1.1 Model

In this paper we study multi-unit auctions with budget-constrained bidders. We suppose there are m identical divisible units of a single item for sale. Each bidder i has a private value ν_i for each unit, and a private budget limit b_i on the total amount she may pay. We assume that bidder i's utility from acquiring x_i units and paying price p_i is $u_i = x_i\nu_i - p_i$ as long as the price is within budget: $p_i \leq b_i$, and is negative infinity if $p_i > b_i$. (i.e. the budget constraint is *hard*.)

Throughout the paper, by v_i and b_i we mean the *submitted* value and the *submitted* budget of the i-th bidder; private value of the i-th bidder however, is shown by ν_i.

Our Contributions:

1. We propose a new mechanism, called Sort-Cut, (Section 2) for selling *all* the units. We prove that Sort-Cut mechanism is *semi-truthful* (Section 3), i.e. no agent can benefit from lying about her budget or understating her value, but may overstate her value at equilibrium (thus address all but one of the four ways in which a bidder may lie about her budget and value) . We show that the allocation of the Sort-Cut mechanism is Pareto-optimal (defined formally later - Section 4); hence, it is nearly the best possible result that can be obtained for this problem since the recent result of Dobzinski, Lavi and Nisan [DLN08] shows that there is no

truthful Pareto-optimal deterministic [1] mechanism for this problem.

2. We obtain an upper-bound R^* (which coincides with the revenue of ascending price auction with truthful bids) for the revenue at equilibrium for *any* Pareto-optimal mechanism (Section 5) and then show that if the sum of budgets of all bidders is at least twice of R^*, this upper-bound is achievable at an equilibrium of the Sort-Cut mechanism.

3. Assuming reasonable behavior of the bidders where every losing bidder bids at least her true value (defined as rational bidding - see definition 4), we show (also in Section 5) that any equilibrium of Sort-Cut has a revenue of at least $R^* - b_{max}$ where b_{max} is the maximum budget among the winners, and prove this bound is tight for equilibria of all Pareto-optimal mechanisms which are budget-truthful (bidders can not benefit from lying about their budgets).

4. We also study the properties of this auction under greedy bidding behavior (Section 6) and show that under some natural assumptions, if the behavior leads to an equilibrium, the unique one it leads to is the revenue-maximizing one (attains revenue R^*).

Previous Work:

The problem of multi-unit auctions with budget-constrained bidders was initiated by Borgs et al. in [BCI+05]. Our model is identical to theirs. They introduce a truthful mechanism that is asymptotically revenue-maximizing; however, it may leave some units unsold. The idea is to group the bidders randomly into two groups, and use the market clearing price of each group as an offering price to the other group, following [GHK+06]. Another paper that uses the same model is by Abrams [Abr06] - it uses techniques similar to [BCI+05] but improves upon it; however, it may still leave some units unsold.

A recent paper that analyzes this problem by Dobzinski et al. [DLN08], mainly proves an impossibility result. They assume that budgets of all players are publicly known, and give a truthful mechanism which solves the problem under this assumption. Their mechanism is a direct application of Ausubel's auction [Aus04]. Then they show that this mechanism is the unique mechanism which is both truthful and Pareto-optimal under the assumption of publicly known budgets. Finally by showing that their mechanism is not truthful if the budgets are private knowledge, they conclude that no mechanism for this problem can be both truthful and Pareto-optimal.

Both [BCI+05] and [DLN08] argue that lack of quasi-linearity (because of hard budget constraints) is the most important difficulty of the problem. Some papers, nonetheless, have tried to solve the problem by relaxing hard budget constraints [Mas00], or modeling the budget constraint as an upper bound on the value obtained by the bidder rather than her payment [MSVV07]. It has also been shown [BCI+05] that modeling budget constraints with quasi-linear functions can lead to arbitrarily bad revenue.

Another paper that has studied budget constraints, mainly for advertisement auctions, is the work of Feldman et al.[FMNP08]. They give a truthful mechanism for ad auctions with budget-constrained advertisers where there are multiple slots available for each query, and an advertiser cannot appear in more than one slot per query. Their work is related to our work because they also consider the game-theoretic aspects of the problem. However, the utility function that they use is very different from ours. In [FMNP08] they define advertisers to be click-maximizers, while in our

[1] We emphasize that our mechanism is not deterministic.

model, advertisers are profit-maximizers, which we believe to be more realistic in the case of ad auctions.

Other papers that have considered budgets in auctions include [AM04],[BK98],[CG96]. However, [AM04] only considers the offline optimization problem and does not study the game theoretic aspects of the problem. They also model budget constraints by value functions of the bidders, which means bidders are not willing to get value more than their budget. In [BK98], they study an auction for selling two single items to budget-constrained bidders. They mainly focus on the effect of bidding aggressively on an unwanted item with the purpose of depleting other bidders budget. A similar effect arises in our model as well, but the focus of our work is generally very different from theirs. Another paper [CG96] compares first-price and all-pay auctions in a budget-constrained setting and show that the expected payoff of all-pay auctions is better under some assumptions. However, they do not consider multi-unit items.

The rest of the paper is organized as follows. In Section 2 we describe our proposed mechanism and in Section 3 we show that our mechanism is semi-truthful. Then in Section 4 we show the Pareto-optimality of Sort-Cut mechanismusing a natural observation from [DLN08]. We step into revenue analysis in section 5 by proving some theorems which state Sort-Cut mechanism is almost revenue optimal. In Section 6 we show that if bidders use a natural greedy algorithm for bidding, the only equilibrium of the infinitely repeated game is revenue optimal.

2 Sort-Cut mechanism Description

In this section we describe how *Sort-Cut* mechanism allocates the units and charge the bidders. Throughout the paper, for simplicity of description we always assume there exists a bidder with value ε and budget $m\varepsilon$ (she has enough money to buy all the items with her value). As ε tends to zero, the revenue of this modified instance approaches that of the original, and hence this assumption is without loss of generality.

The idea of Sort-Cut is very similar to the idea of a second-price auction. In second-price auctions without budget constraints, the highest bidder (highest value person in a symmetric equilibrium) is allocated the object and what she pays to the auctioneer is the bid of the highest loser's bid. Uniform-price auction generalizes this idea to multi-unit auctions. The idea is to charge the winners by the opportunity cost: the losers' bids (or values in equilibrium). When the bidders have budget constraints, however, losers might not be able to buy all the items if they were offered, they might simply not afford it. Taken this into account, we modify the algorithm to charge the winners, per item, for the value of the highest value loser, but only up to highest loser's budget. After the highest value loser's budget is exhausted, she would not be able to afford any more items, so we start charging the winners the value of the second highest value loser, up to her budget and so on. Given this pricing idea , the winners and losers are determined via a cut-point to clear the market, i.e. to be able to sell all the available items.

There is a caveat here, which is that the lowest value winner might not be able to exhaust all her budget. Then all higher value bidders are charged first at the lowest value winner's value up to her unused budget. This makes sense as the lowest value winner is still a competitor to other winners to buy further items. The pricing for the lowest value winner, for the same reason, starts from the highest value loser. She cannot be a competitor to herself!

In our model, there are n bidders and m units available of the same item. The bidders have flat demands, i.e. bidder i's value per item is ν_i. Bidders submit their budget b_i and their value v_i to the mechanism and Sort-Cut decides about the allocation and pricing. Note that submitted budget and submitted value might be different from the bidder's actual budget and value. Throughout this section, by value and budget we mean submitted value and submitted budget of the bidders.

The Sort-Cut mechanism has two main features: sorting and cutting. First, we sort the bidders in non-increasing order of their values, and assume by relabeling that $v_1 \geq v_2 \geq \ldots v_n$. The second part is the cutting by which the algorithm assigns the available m units to the bidders 1 through k (k is not determined yet), and assigns nothing to the bidders $k+1$ through n. The bidders $1, \ldots, k-1$ must have exhausted all their budget, and bidder k may be left with part of her budget. Bidders $k+1, \ldots, n$ do not pay anything. We describe later how we determine the exact cut-point which is the index k and the portion of the budget of k that is used up in payments.

Suppose that the money left for bidder k is b'_k, and define $c_k = \frac{b'_k}{v_k}$. For $i > k$ define $c_i = \frac{b_i}{v_i}$. Here, c_i denotes the number of (fractional) units that bidder i can buy according to *her* value for one unit. We now describe the pricing rule for Sort-Cut: For all bidders whose values are greater than the cut point bidder (bidder i with $i < k$), we charge bidder i for the first c_k items that she wins a price of v_k, for the next c_{k+1} items that she wins a price of v_{k+1}, for the next c_{k+2} items that she wins a price v_{k+2} and so on until her budget is exhausted. For the cut-point bidder k, the pricing is slightly different: we start charging her a price of v_{k+1} for the first c_{k+1} items that she gets, v_{k+2} for the next c_{k+2} items that she gets and so on until she has exhausted all her allocated spending budget of $b_k - b'_k$. We give each bidder i ($i < k$) as many units as she can afford according to the pricing policy and her budget.

To finish the description, we need to specify how we determine k, and b'_k (the money which is left for the last winner), because they play an important role in our pricing mechanism. If the index k is very large (close to n), then the prices for the units will be very low and the number of units that each winner can afford increases, also the number of winners is large resulting in a shortage of supply. On the other hand, if k is chosen very small (close to 1), we have few winners, and relatively high prices, so a number of units will be left unsold. We seek to find the right point which determines bidder k and the amount of money b'_k left for her, such that the market clears at this point.

To be more precise (see algorithm 1), define $B = \sum_{i=1}^{n} b_i$. We are looking for a cut-point $R = \sum_{i=1}^{k} b_i - b'_k$ in the interval $[0, B]$ that determines for us both k and b'_k which will make the number of allocated objects exactly equal to m. We want to sell all items. We must also guarantee that the bidders $1, ..., k-1$ (which are determined by R) have exhausted all their budget. As we increase R, the prices decrease, the number of winners and units demanded to be allocated increase, and consequently the demand increases. We can find the solution by slowly increasing R from 0 until the demand becomes equal to supply (In algorithm 1 we use binary search). In other words, we increase R until the total number of units that bidders 1 through k want (assuming that bidder k can use only $b_k - b'_k$ of her budget) becomes equal to the units that we have to sell. (A formal proof of the fact as well as examples of such cut-points appear in Appendix A.)

Finally, to keep the bidders from overstating their budget, we add the following to Sort-Cut. Suppose that the Sort-Cut mechanism wants to charge a bidder i an amount equal to p_i, and her announced budget is b_i. Instead of charging her p_i, we charge her b_i with probability p_i/b_i and 0 otherwise. In this way, if somebody overstates her budget, she is accepting the risk of paying more than her budget which makes her expected utility equal to minus infinity (this is because we have hard budget constraints). To make this more practical, even if we perform this alternate pricing process with probability ε, and simply charge the bidder i a price of p_i in the rest of the cases, still bidders can not take the risk of overstating their budget[2].

3 Semi-Truthfulness

Although Dobzinski et al. [DLN08] shows that no truthful Pareto-optimal mechanism exists for this problem, it is still interesting to know *how truthful* a Pareto-optimal mechanism can be. In other words, we want to know how the different ways of lying can benefit them in a mechanism.

There are four ways that a bidder can lie: overstating budget, understating budget, overstating value and understating value. In different mechanisms, bidders may take different strategies and use either of these ways to increase their utility. We show that in Sort-Cut mechanism, the only way out of these four that the bidder can use to benefit from lying is by overstating value. This result is interesting because first, we know that some kind of lying must be profitable for the bidders if the mechanism is Pareto-optimal, and second, among four different ways of lying, this is the most desirable one for the auctioneer – giving good revenue properties (This is formalized in Section 5). It is easy to see that the revenue of Sort-Cut mechanism is monotone with respect to the vector of bids. (i.e. if bidders increase their stated values, the revenue of the mechanism does not decrease.)

Definition 1 *We say a mechanism A is semi-truthful if every strategy for the bidders that involves mis-stating the budget or understating the value is dominated by a strategy that does not misstate the budget and does not understate the value.*

Semi-truthfulness limits the set of undominated strategies to those that budget is stated truthfully and value is not understated. In other words, definition 1 states that in a semi-truthful mechanism one should not expect from the bidders to lie about their budget or understate their values. Babioff et al. [BLP09] and Jackson [Jac92] describe in more details how undominated strategy mechanisms generalize the more usual framework of dominant strategy mechanisms. In this paper, as in [BLP09] we rely on the fact that a strategy is *reasonable* if it is not dominated, and hence, we assume the bidders do not understate their values and do not misstate their budgets.

Theorem 1 *Sort-Cut mechanism is semi-truthful.*

Proof:

First it is easy to see that no bidder can benefit by overstating her budget, since by design, if she wins any allocation at all, there is positive probability that she will be charged more than her budget which makes her utility minus infinity.

Next, we argue that no bidder can benefit by understating her budget. Consider the winners; they (including the cut-point bidder) can pull the prices down by understating the budget. Consider $j < k$ (for current prices) who understates her budget from b_j to b'_j, and let the new allocation be $x'_1, \ldots, x'_n$. This deviation could be beneficial to the bidder only if she gets the x'_j units at lower average price than the average price of cheapest x'_j units she was allocated before the deviation. For this to be the case, we should have $R' \geq R$ (where R and R' represent winners' total payments before and after the deviation respectively.) because otherwise, she cannot get the extra units at lower prices after deviation. But it is easy to see that this cannot happen. Consider $R' \geq R$, then

[2] It is not hard to argue that it suffices to use this kind of randomization for the cut-bidder, bidder k only. We can also construct examples where the cut-bidder can benefit from over stating the budget without this modification.

$x'_i \geq x_i$ for all winners except j, because they have the same budgets but face smaller prices. On the other hand, $x'_j \geq x_j - (b_j - b'_j)/v_k$, and also $x'_k + \ldots + x'_n \geq x_k + (b_j - b'_j)/v_k$. Hence there would be excess demand if $R' \geq R$. Similar argument shows that deviation by understating the budget cannot be helpful for the cut-point bidder.

For the third kind of deviation (under-stating value), we argue in three cases:

1. Consider a bidder j where $j \leq k-1$. If bidder j understates her value, she may still remain among the first $k-1$ bidders which does not change anything for her, or, she may go to the boundary which makes the situation trickier, or she may go below the boundary which decreases her utility. So the only case that we must handle is when she goes to the boundary.

   Suppose bidder j moves to the boundary by announcing a value $v'_j \leq v_k$ and she spends b of her money while b' of her money will be remaining. (so we have $b + b' = b_j$) We know that before going to the boundary, with b'_k of her budget she bought the units for price v_k per unit, and after that she had unit prices $v_{k+1}, \ldots$. Now when she goes to the boundary, she buys for prices $v_{k+1}, v_{k+2}, \ldots$ which are lower and seem to be better for her, but as we will see, that is not the case because she is not using all her money when she is on the boundary. First note that if $b' \geq b'_k$, she can not benefit from going to the boundary (because previously she was using b'_k of her budget for getting the units for price v_k per unit, but now she has b' of her budget left unused). So we may assume $b' < b'_k$. Now, consider bidder k to see how many units she wins after leaving the boundary. Now the bidder with value v_k gets at least $\frac{b'_k}{v'_j} \geq \frac{b'_k}{v_k} \geq c_k$ units (for a price of v'_j) in addition to all the units that she had before (when she was on the boundary). The prices for all other winners is less than or equal to what it was before. Therefore, the number of units that bidder j wins after going to the boundary must be reduced by at least these c_k units. Bidder j's costliest c_k units were priced v_k units each for a total price of $b'_k = c_k v_k$. After understating her value, she is paying $b_j - b' > b_j - b'_k$ and getting at least c_k units less. Thus her average price per unit has increased so this is not an improvement.

2. Now consider a bidder j where $j > k$. It is obvious that bidder j can not benefit from understating her value, because it keeps her among the losers. However, overstating the value may be beneficial for her in some cases.

3. For the person on the boundary ($j = k$), it is clear that she can not benefit from understating her value, because it can not influence her price and she may even lose the units that she already wins (by reducing the price for earlier winners). Again, overstating the value may increase her utility in some cases.

□

4 Pareto Optimality

Definition 2 *An allocation* $\{(x_i, p_i)\}$ *is* Pareto Optimal *if for no other allocation* $\{(x'_i, p'_i)\}$ *are all players better off:* $u_i(x_i, p_i) > u_i(x'_i, p'_i)$ *(Recall that* $u_i(x_i, p_i) = x_i \nu_i - p_i$ *if* $p_i < b_i$ *and* $-\infty$ *otherwise), as well as the auctioneer:* $\sum_i p'_i \geq \sum_i p_i$*, with at least one inequality strict.*

Pareto optimality is simply implied by a proposition from [DLN08].

Proposition 2 *An allocation $\{x_i, p_i\}$ is Pareto-optimal in the infinitely divisible case if and only if (a) all units are completely sold, and (b) for all i such that $x_i > 0$ we have that for all j with $\nu_j > \nu_i$, $p_j = b_j$. I.e. a player may get a non-zero allocation only if all higher value players have exhausted their budget.*

Since Sort-Cut mechanism always allocates the units in decreasing order of the submitted values, a bidder with submitted value v_i may be allocated some units only if the bidders with higher submitted values $v_j > v_i$ have exhausted their budget. Therefore, the allocation of the Sort-Cut mechanism is Pareto-optimal with respect to submitted values by construction.

Notice that we are showing Pareto-Optimality with respect to the submitted vector of values v instead of actual vector of values ν. Obviously, the definitions are not equivalent because Sort-Cut is not truthful. However, in the following theorem we show that any ex post Nash equilibrium of Sort-Cut in which losers bid at least their true values is Pareto-optimal with respect to private vector of values.

Theorem 3 *Any ex post Nash equilibrium of Sort-Cut in which all losers and the cut-point bidder bid at least their true values is Pareto-optimal with respect to actual values.* [3]

Proof: Suppose the private values of the bidders $1, \ldots, n$ are $\nu_1, \ldots, \nu_n$ respectively. Note that since v_i's are sorted decreasingly, ν_i's are not necessarily in decreasing order. Take an arbitrary ex post Nash equilibrium of Sort-Cut and assume the submitted vectors of values and bids are v and b respectively. Suppose bidder with index i is a winner, bidder with index j is on the boundary and bidder with index k is a loser. We show that $\nu_i \geq \nu_j$ and $\nu_j \geq \nu_k$ which together imply that any ex post Nash equilibrium of Sort-Cut is Pareto-Optimal with respect to private values ν.

Assume for sake of contradiction that $\nu_i < \nu_j$. First consider the case where $b_i \leq b'_j$ (i.e. budget of bidder i is less than or equal to the left-over budget of bidder j). Bidder i has to pay at least $v_j \geq \nu_j$ (by theorem hypothesis) per unit for all units that she is getting; therefore, bidder i has negative utility which contradicts the assumption of Nash equilibrium, since a bidder can always get zero utility by bidding 0. Now assume $b_i \geq b'_j$. We show that bidding $v_j - \varepsilon$ is a profitable deviation for bidder i. By bidding $v_j - \varepsilon$, the allocation of bidder i changes by $\Delta x = -b'_j/v_j$ and her price changes by $\Delta p = -v_j b'_j/v_j$. Therefore, her utility changes by $\Delta u = -\nu_i b'_j/v_j + b'_j$. Contradiction assumption $\nu_i < \nu_j$ and theorem hypothesis $v_j \geq \nu_j$ imply that Δu is positive, and hence, contradicts the assumption of Nash equilibrium.

The argument for bidders j and k is very similar. Assume for sake of contradiction that $\nu_j < \nu_k$. Suppose that β is the sum of the budget of all bidders who are ranked between j and k, including bidder k and excluding bidder j. (i.e. $\beta = \sum_{l=j+1}^{k} b_l$.) If $b_j \leq \beta$, bidder j has to pay at least $v_k \geq \nu_k$ per unit for all units she is getting which yields to negative utility for her and consequently contradicts the assumption of Nash equilibrium. If $b_j \geq \beta$, bidder j can benefit by deviating and bidding $v_k - \varepsilon$. The deviation changes her allocation by $\Delta x = -\beta/v_k$ and her price by $\Delta p \leq -v_k\beta/v_k$. Therefore, her utility changes by $\Delta u \geq -\nu_j\beta/v_k + \beta$. Contradiction assumption $\nu_j < \nu_k$ and theorem hypothesis $v_k \geq \nu_k$ imply that the deviation strictly increases the utility of bidder j, and hence contradicts the Nash equilibrium assumption.

□

[3] In an ex post Nash equilibrium, players have no incentive to change their bids even after the bids of other players are announced.

5 Revenue Analysis

Definition 3 *The ascending price auction mechanism is defined as follows. The price starts at $p = 0$ and increases infinitesimally and continuously; At any time, the demand of each bidder i is $d_i = b_i/p$ if $p \leq v_i$ and is $d_i = 0$ if $p > v_i$. The price continues increasing as long as there is over-demand $\sum d_i > m$. The price $v^* = p$, the first point when the demand equals the supply m, is defined to be the market clearing price. All bidders i with value $v_i \geq v^*$ are allocated b_i/v^* units for price v^*. (The lowest valued one may be partially allocated, but still for price v^* per unit.) We define R^*[4] to be the revenue of ascending price auction in which bidders are bidding truthfully.*

Note that ascending price auction is Pareto-optimal, but is not truthful; Specifically, bidders can benefit by understating their budgets.

In the next lemma, we show an upper bound on the revenue of any mechanism which guarantees Pareto-optimality.

Lemma 4 *No individually rational, Pareto-optimal mechanism, in equilibrium, can guarantee revenue more than R^*.*

Proof: A mechanism will take announced values and bids and allocate the goods to some bidders at some prices. A mechanism, per-item, should not charge any bidder more than her value.

Suppose that v^* is the market clearing price and let l be the greatest index such that $v_l \geq v^*$. If a mechanism A generates a revenue more that R^*, it must charge some bidder i ($1 \leq i \leq l$) more than v^* per unit. But if bidder i decreases her bid down to $v^* + \varepsilon$, the mechanism still has to exhaust all her budget (otherwise, because of Pareto-optimality, it can not charge the bidders who have value v^* or less, and consequently can not even make revenue R^*) but now with price of at most $v^* + \varepsilon$. That means that at an equilibrium of the mechanism, no bidder can be charged more than v^* per unit. $\square$

Proposition 5 *If we assume that revenue R^* of ascending price auction for truthful bids is less than half of the total budget of all participants, i.e. $\sum_{i=1}^{n} b_i \geq 2R^*$, then there exists an equilibrium of the Sort-Cut mechanism in which the payments and utilities are like those of the ascending price auction with truthful bids.*

Proof: The following vector of bids will be a Nash equilibrium in the game of complete information. All those who have value greater than v^* bid truthfully, those who have value less than or equal to v^* bid v^*. Therefore, all those who are bidding v^* are the losers and will not be assigned anything, and the winners have to pay v^* per unit. (Note that if the last winner is partially using her budget, then v^* is equal to her value, and she has utility 0. Therefore she has no incentive to increase her bid for depleting the budget of other winners.)

The condition on the revenue is required so that there is enough budget of unallocated bidders to set the corresponding market clearing price for the ones that are allocated in the ascending price auction equilibrium. $\square$

The rest of this section obtains a lower bound for the revenue of Sort-Cut mechanism[5]. Before that, we need to introduce the concept of *Rational Bidding*. Since someone who is not winning

[4] R^* is defined in [DLN08] as revenue of a non-discriminatory monopoly that knows the true budgets and true values of the bidders, and has to determine a single unit-price at which the units will be sold.

[5] We do the revenue analysis for the announced values and budgets.

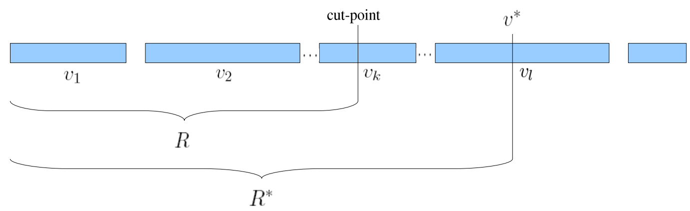

Figure 1: Revenue Comparison

anything in the Sort-Cut mechanism can never benefit from understating her value, we have the following definition.

Definition 4 *We say that bidders are bidding rationally*[6]*, if those who do not win anything bid at least their true value.*

Since Sort-Cut is a semi-truthful mechanism, as stated and proved in section 3, every strategy that understates value is dominated. Therefore, assuming that bidders bid rationally is weaker than assuming they play an undominated strategy.

Theorem 6 *Assuming rational bidding, the revenue of the Sort-Cut mechanism at any equilibrium is at least $R^* - b_{max}$ where b_{max} is the maximum budget among the winners.*

Proof: Suppose that Sort-Cut mechanism has used all the budget of bidders $1, ..., k-1$ and a part of the budget of k-th bidder. Also suppose that market clearing price for truthful bids is v^* where $v_l \geq v^* > v_{l+1}$. As we defined, $b_{max} = max_{1 \leq i \leq k} b_i$, where b_i is the budget of i-th bidder. Note that the revenue of Sort-Cut mechanism is $R = \sum_{i=1}^{k} b_i - b'_k$ where b'_k is the amount of money left unused by the k-th bidder. Also note that $R^* \leq \sum_{i=1}^{l} b_i$ (by the definition of ascending price auction).

Now, we are ready to prove the claim. Consider an output of the Sort-Cut mechanism. Since the revenue of Sort-Cut mechanism is less than R^* and both mechanisms sell all m units, there must be some winner i who is getting the item cheaper than v^* per unit. Also we can assume that $k > 1$ since otherwise the claim holds trivially. We can either have $i < k$ or $i = k$. First suppose that $i < k$. According to our pricing scheme, bidder i has to pay at least v^* per unit up to $R^* - R$ of her budget (See Figure 1). That means if i is paying less than v^* on average per unit, her budget must be more than $R^* - R$. Therefore, $b_i > R^* - R$ which implies $R > R^* - b_{max}$. Now consider the case where $i = k$. Here, bidder i is using $b_i - b'_i$ of her budget, and she has to pay at least v^* per unit up to $R^* - R - b'_i$ of her budget. Therefore, if she pays less than v^* on average per unit, the amount of her budget that she is using, $b_i - b'_i$, must be more than $R^* - R - b'_i$. That is $b_i - b'_i > R^* - R - b'_i$, equivalently, $R > R^* - b_i$. $\square$

Surprisingly, an independent analysis in [DLN08] shows exactly the same revenue result for a different, still non-truthful, mechanism for this problem. The expression they prove for the revenue

[6] Rational bidding is a special case of undominated strategy discussed in [BLP09] and mentioned in section 3.

of their mechanism reduces to

$$R \geq (1-\alpha)R^*$$

for the divisible case, where

$$\alpha = max_{i=1,\dots,n} \frac{x_i^*}{\sum_{j=1}^n x_j^*}$$

and x^* represents the allocation in ascending price auction. Since ascending price auction is single-price, α is equal to the highest budget among the winners divided by the total revenue R^*. Therefore, αR^* is simply the same as b_{max} and consequently,

$$(1-\alpha)R^* = R^* - b_{max}.$$

Our analysis in the above theorem is tight: a simple example with one bidder demonstrates that any Pareto-optimal mechanism which is truthful for the budgets (like Sort-Cut mechanism) can not guarantee a revenue higher than $R^* - b_{max}$. No budget-truthful Pareto-optimal mechanism can charge this bidder more than 0. Therefore, our bound of $R^* - b_{max}$ is tight, and the best achievable in budget-truthful Pareto-optimal mechanisms.

6 Revenue Optimal Equilibrium by Repeated Bidding

In this section, we are using the same approach that Edelman et al [EOS05] used to model the generalized second price (GSP) auctions. Like in their paper, we assume that the bidders are playing an infinitely repeated game, and use this to obtain some equilibria properties for Sort-Cut mechanism. We then take the approach of a consequent paper [CDE$^+$07] which shows that a natural bidding strategy played by all bidders leads to a unique Nash equilibrium of GSP, and that the Nash equilibrium coincides with the outcome of a VCG auction. Here, we show that a natural bidding strategy, called *Greedy Bidding*, when it converges to an equilibrium, leads to one that coincides with the outcome of ascending price auction with optimal revenue.

We focus on simple strategies and impose some assumptions and restrictions. First, we assume that all budgets and values are common knowledge: over time, advertisers are likely to learn all relevant information about each other. Second, since bids can be changed at any time, stable bids must be best responses to one another. We define *Greedy Bidding*, a simple and natural response algorithm for the bidders who are playing the infinitely repeated game without knowing anything about bids and budgets of other bidders. Then we show that if running this algorithm converges, it does so at an unique equilibrium with prices and allocations identical to those of the ascending price auction with truthful bids.

Definition 5 *(Greedy Bidding) Assume that each bidder always bids her true budget. Moreover, she revises her bid at each round of the infinitely repeated game by executing the following rules in order.*

1. *If what the bidder is paying per unit on average is higher than her value per unit (her bid is too much above her value), she decreases her bid continuously over time.*
2. *If all or part of her budget is left unspent (trying to deplete the budget of those who are above her, and winning more units), she increases her bid continuously over time.*

3. *If she is using all her budget (she can not influence her own price according to the pricing scheme), she does not change her bid.*

Note that we are not specifying any order for the bidders to change their bids. As it will be clear from the proof below, it does not matter how the bidders take turns to modify their bids as long as they converge to some equilibrium.

Lemma 7 *If all bidders use Greedy Bidding and converge to an equilibrium, this equilibrium has prices and allocations identical to those of the ascending price auction with truthful bids, which provides revenue R^*.*

Proof: First we claim that all losers have the same bid in an equilibrium. This is because if any of them bids slightly higher, she must be assigned something with a price higher than her value to send her back to her previous bid using rule 1. Because of Pareto-optimality, this can happen only if they all have the same bid. Moreover, this common bid must be the highest possible, otherwise they all can increase their bid. Those who have higher value than the common bid of the losers must bid higher than the losers. Therefore, the unique solution to this greedy bidding system is when all losers are bidding exactly (or slightly lower than) the market clearing price and this completes the proof. □

Acknowledgments: We are grateful to Arpita Ghosh, David Pennock, Mohammad Mahdian from Yahoo! Research, Moshe Babaioff from Microsoft Research and anonymous reviewers of earlier versions of this paper for their thoughtful suggestions and valuable comments.

References

[Abr06] Zoë Abrams. Revenue maximization when bidders have budgets. In *SODA '06: Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 1074–1082, New York, NY, USA, 2006. ACM.

[AC95] Lawrence M. Ausubel and Peter Cramton. Demand reduction and inefficiency in multi-unit auctions. Technical report, University of Maryland, Department of Economics - Peter Cramton, November 1995.

[AM04] Nir Andelman and Yishay Mansour. Auctions with budget constraints. *SWAT 04, Lecture Notes in Computer Science*, 3111:26–38, 2004.

[Aus04] Lawrence M. Ausubel. An efficient ascending-bid auction for multiple objects. *American Economic Review*, 94(5):1452–1475, December 2004.

[BCI+05] Christian Borgs, Jennifer Chayes, Nicole Immorlica, Mohammad Mahdian, and Amin Saberi. Multi-unit auctions with budget-constrained bidders. In *EC '05: Proceedings of the 6th ACM conference on Electronic commerce*, pages 44–51, New York, NY, USA, 2005. ACM.

[BK98] Jean-Pierre Benoit and Vijay Krishna. Multiple-object auctions with budget constrained bidders. Game Theory and Information 9805001, EconWPA, May 1998.

[BLP09] M. Babaioff, R. Lavi, and E. Pavlov. Single-value combinatorial auctions and algorithmic implementation in undominated strategies. 2009.

[CDE+07] Matthew Cary, Aparna Das, Ben Edelman, Ioannis Giotis, Kurtis Heimerl, Anna R. Karlin, Claire Mathieu, and Michael Schwarz. Greedy bidding strategies for keyword auctions. In *EC '07: Proceedings of the 8th ACM conference on Electronic commerce*, pages 262–271, New York, NY, USA, 2007. ACM.

[CG96] Yeon-Koo Che and Ian Gale. Expected revenue of all-pay auctions and first-price sealed-bid auctions with budget constraints. *Economics Letters*, 50(3):373–379, March 1996.

[DLN08] Shahar Dobzinsky, Ron Lavi, and Noam Nisan. Multi-unit auctions with budget limits. In *to appear: Foundation of Computer Science*, Philadelphia, PA, USA, October 2008.

[EOS05] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. Working Paper 11765, National Bureau of Economic Research, November 2005.

[FMNP08] Jon Feldman, S. Muthukrishnan, Evdokia Nikolova, and Martin Pal. A truthful mechanism for offline ad slot scheduling, 2008.

[GHK+06] Andrew V. Goldberg, Jason D. Hartline, Anna R. Karlin, Michael Saks, and Andrew Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242–269, May 2006.

[Jac92] M.O. Jackson. Implementation in undominated strategies: A look at bounded mechanisms. *The Review of Economic Studies*, pages 757–775, 1992.

[Mas00] Eric S. Maskin. Auctions, development, and privatization: Efficient auctions with liquidity-constrained buyers. *European Economic Review*, 44(4-6):667–681, May 2000.

[MSVV07] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. Adwords and generalized online matching. *J. ACM*, 54(5):22, 2007.

[Nau95] D. Nautz. Optimal bidding in multi-unit auctions with many bidders. *Economics Letters*, 48(3-4):301–306, June 1995.

[Wol98] Catherine D. Wolfram. Strategic bidding in a multiunit auction: An empirical analysis of bids to supply electricity in england and wales. *RAND Journal of Economics*, 29(4):703–725, Winter 1998.

A Sort-Cut Details

We first sketch here the proof that there exists k, b'_k such that the number of allocated items allocated is exactly m in the description of Short-Cut. First of all, x can span the whole interval $[0, B]$, that is there is continuity in (k, b'_k). The only discontinuity can be when we change k to $k+1$, but b'_k gives us enough continuity to span the whole interval, therefore the number of allocated objects is also going to be continuous in (k, b'_k). Secondly, when x is low (close to zero by setting k equal to 1 and b'_1 close to b_1), the number of allocated objects can be at most $\frac{b_1}{v_2}$ (which we assume to be smaller than m) and is very low. On the other hand, when x is high (close to B, by setting k equal to n) the number of allocated objects is very high (goes to infinity with the assumption that there exists a fictitious last bidder with very low value and high budget). Hence, from intermediate value theorem we can conclude that such k, b'_k exists. Algorithmically, we can use a simple binary search for finding the right value for x.

Example 1 *We give an example to show how our mechanism works for selling* 19 *units of a divisible item to a set of* 4 *bidders with the following (private) values per item and budgets.*

i	v_i	b_i
1	10	55
2	9	60
3	7	40
4	6	30

We start with $x = 128$*, so* $k = 3$ *and* $b'_3 = 27$*. This means that the price of each unit (for the first two bidders) is* 7 *for the first* 27 *that they spend, and after that for the next* 30 *that they spend, the price for each unit is reduced to* 6*, and finally, after that the price is* ε *for each unit. Therefore, the first bidder can afford* $x_1 = 27/7 + 28/6$ *units with a total price of* 55*. But the second bidder can afford all the remaining units now (which means nothing will be left for the third bidder, who must be assigned something according to our cut-point).* $x_2 = 27/7 + 30/6 + 3/\varepsilon$*. This means that* x *is too large for the cut-point.*

Our next guess is $x = 122$*, so* $k = 3$ *and* $b'_3 = 33$*. Here, the price (for the first two bidders) is* 7 *per unit for the first* 33 *that they spend, and* 6 *per unit for the next* 30 *that they spend, and* ε *per unit after that. Therefore, the first bidder can afford* $x_1 = 33/7 + 22/6$ *and the second bidder can afford* $x_2 = 33/7 + 27/6$ *units. The third bidder can use* 7 *of her money and she has to pay* 6 *per unit for the first* 30 *that she spends. Therefore, she can afford* $x_3 = 7/6$ *units. We can see that* $x_1 + x_2 + x_3 < m$*, therefore,* x *is too small for the cut-point.*

By continuing the same procedure, we will see that $x \simeq 123.11$ *is the right value for* x*. Therefore,* $x_1 \simeq 8.4$*,* $x_2 \simeq 9.25$ *and* $x_3 \simeq 1.35$*, and the prices they pay are* $p_1 = 55$*,* $p_2 = 60$ *and* $p_3 \simeq 8.11$*.*

Algorithm 1 The Sort-Cut Mechanism for divisible units to determine the cut-point x^*

$\{$Initialization$\}$ **Let** $B = \sum_i b_i$; **Initialize** $x = B/2$ and allocations $y_j = 0$ for all j.
repeat
  $B \leftarrow B/2$
  $\{$Cut-point determination$\}$ **Let** k be the largest index such that $\sum_{i=1}^{k-1} b_i \leq x$.
  $\{$Pricing$\}$ **Let** $b'_k = \sum_{i=1}^{k} b_i - x$, and define $c_k = b'_k / v_k$. For $i = k+1$ to n define $c_i = b_i / v_i$.
  $\{$Payments and Allocations until cut-point$\}$
  for i = 1 to $k - 1$ **do**
    Set the payment $p_i = b_i$; Initialize allocation $y_i = 0$ and $j = k$
    while $b_i > 0$ **do**
      $y_i = y_i + \frac{min(b_i, c_j.v_j)}{v_j}$
      $b_i = b_i - min(b_i, c_j.v_j)$
      $j = j + 1$
    end while
  end for
  $\{$Payments and Allocations for cut-point$\}$
  for bidder k **do**
    Let $b = b_k - b'_k$ and set the payment $p_k = b$; Initialize allocation $y_k = 0$ and $j = k + 1$
    while $b > 0$ **do**
      $y_k = y_k + \frac{min(b, c_j.v_j)}{v_j}$
      $b = b - min(b, c_j.v_j)$
      $j = j + 1$
    end while
  end for
  $\{$Binary Search Update$\}$
  if $\sum_{i=1}^{k} y_i > m$ **then**
    $x \leftarrow x - B/2$
  end if
  if $\sum_{i=1}^{k} y_i < m$ **then**
    $x \leftarrow x + B/2$
  end if
until the sum of allocations y_j is the supply m