

Maximizing Revenue in Sequential Auctions

Shaheen Fatima¹ and Edith Elkind²

¹ Department of Computer Science, University of Liverpool, U.K.

² School of Electronics and Computer Science, University of Southampton, U.K.

Abstract. We study sequential auctions for private value objects and unit-demand bidders using second-price sealed-bid rules. We analyze this scenario from the seller's perspective and consider several approaches to increasing the total revenue. We derive the equilibrium bidding strategies for each individual auction. We then study the problem of selecting an optimal agenda, i.e., a revenue-maximizing ordering of the auctions. We describe an efficient algorithm that finds an optimal agenda in the important special case when the revenue of each auction is guaranteed to be strictly positive. We also show that the seller can increase his revenue by canceling one or more auctions, even if the number of bidders exceeds the number of objects for sale, and analyze the bidders' behavior and the seller's profit for different cancellation rules.

1 Introduction

Market-based mechanisms such as auctions are now being widely studied as a means of allocating resources in multiagent systems. There are several reasons for their popularity: auctions are simple to implement and can also have desirable economic properties, probably the most important of which are their ability to generate high revenues to the seller and also allocate resources efficiently [1, 5, 6]. In many practical applications, the number of objects for sale is large, and the seller has to choose a suitable auction format, such as a *combinatorial* [7], *simultaneous*, or *sequential* [8, 4] auction. Each of these formats has unique advantages, and selecting the best mechanism for a given setting can be a challenging task. For example, while a combinatorial auction is a good choice when it is imperative to allocate the objects efficiently, simultaneous or sequential auctions are easier to implement, as one can use off-the-shelf systems for single-item auctions.

In this paper, we study sequential auctions, i.e., mechanisms in which objects are sold consecutively one at a time. Even though at any given moment there is only one item being auctioned, the bidders' behaviour in any individual auction strongly depends on the auctions that are yet to be conducted [8, 4]. In particular, even if the auctions are run using second-price rules, the bidders are not likely to bid their true value for the item if they expect to profit from subsequent auctions. Moreover, the bidding strategies for an auction strongly depend on the agenda (i.e., the order in which the objects are auctioned); if we change the agenda, then the bidding strategies and consequently the equilibrium outcome changes [8].

The model considered in this paper is motivated by the following scenario. Suppose that we are selling advertising space in a recurring event, such as several episodes of a TV show or all football games in a given season. Assume for simplicity that each event is associated with a single advertising slot, and all bidders are *ex ante* symmetric and

have unit demand, i.e., each of them only needs in one slot. I.e., if a bidder wins an auction, he does not participate in the future auctions. Clearly, some games or episodes of the show will be more attractive to the advertisers than others: a game between top teams will have a larger audience than a game between less successful teams, some episodes of the show may include appearances by celebrities, etc. In the beginning of the season, the seller and the advertisers have some estimates of the value of advertising in each slot, which, as argued above, can be different for different slots. One can associate such an estimate with a probability distribution over possible values. In the symmetric setting, it is natural to assume that these distributions are publicly known.

If the slots are to be auctioned off in the beginning of the season, the advertisers will have to bid based on these imprecise value estimates. However, it is possible to postpone selling each slot till the respective episode or game is about to be shown. At this point the bidders are likely to have a better understanding of how much this slot is worth to them, based on their current stock, existing orders, etc. Therefore, we can treat their value for a slot just before the screening as their actual value for this slot, i.e., a random variable drawn from the initial distribution associated with this slot.

In this situation, the seller may have (partial) control over the order of individual auctions. Indeed, in many TV shows, the order in which some (or all) of the episodes are shown is flexible. It is also possible to take into account revenue considerations when scheduling football games. Therefore, it is natural to ask whether the seller can order the individual auctions to optimize his revenue. An alternative approach to maximizing the total revenue, which is also feasible in the above-described scenario, is canceling some of the auctions. In this paper, we study both of these revenue maximization techniques.

The formal model that we use is described in Section 2. Note that we assume that the bidders can commit to participating in the entire series of auctions, i.e., they take part in all auctions until they win an object. We believe that in the above-described setting, this assumption can be justified. Also, in most of the paper we restrict ourselves to the case where all bidders are symmetric and the seller has full control over the ordering of individual auctions. In the end of the paper, we briefly discuss relaxing some of these assumptions.

We consider several approaches to increasing the total revenue. We assume that each auction is run using second-price rules, and compute the participants' equilibrium bids. Using this result as a starting point, we study the problem of selecting the optimal agenda, i.e., a revenue-maximizing ordering of the auctions. We describe an efficient algorithm that finds the optimal agenda whenever the revenue of each auction is guaranteed to be strictly positive. While the latter condition does not always hold, it is likely to be satisfied when the number of bidders is large; also, our approach may provide a reasonable heuristic in the general case.

We then study the complementary case when in some of the auctions there is at most one bidder who submits a non-zero bid. The second-price auction rules allocates the object to one of the highest bidders and charges him the second-highest bid. This means that when all bids are zero, we will give the object to an arbitrary bidder (e.g., one chosen at random) for free. Similarly, if there is exactly one bidder with a non-zero bid, he will get the item and pay nothing. Intuitively, in this situation we are better off canceling the auction altogether: it brings in no revenue anyway, and by allocating the object we eliminate one of the bidders who could otherwise submit a high bid in

a future auction. However, to formalize this intuition we have to take into account the ripple effect of this decision, i.e., its influence on bidders' behavior in other auctions. We study the changes in the bidders' strategies under the new rules. We also sketch an efficient algorithm that chooses in advance which auctions to cancel based on the respective value distributions rather than actual bids.

In Section 6, we consider a more general case of our problem, where bidders are not necessarily symmetric, and the seller's choice of agendas is restricted. We prove that in this case the problem becomes NP-complete.

2 The Auction Setting

There are m private value objects for sale. Each object is sold in a separate auction using the second-price sealed-bid rules, and the auctions are held sequentially. Initially, there are n risk-neutral bidders. For each object j , $j = 1, \dots, m$, the bidders' valuations are drawn independently from a distribution with a cumulative density function $V_j : R_+ \rightarrow [0, 1]$. Consequently, any bidder's valuations for different objects are independent random variables. However, each bidder only needs one object. Therefore, if he wins an auction, he does not participate in subsequent auctions.

The number of objects m , the initial number of bidders n , and the functions $V_j : R_+ \rightarrow [0, 1]$ are common knowledge to all bidders. However, each bidder draws his private value signal for auction j after the end of auction $j - 1$. This model was introduced in [4] in the context of sequential auctions for two private value objects. Here, we generalize the approach of [4] to $m > 2$ objects.

The sequential auctions are conducted as follows. The first object is sold in a second-price sealed bid auction. There are n bidders for this auction. The winner is announced at the end of the auction. He receives the object and leaves. All other bidders then draw their values for the second object and take part in the second auction. The process repeats until all objects are sold (if $n > m$) or until there are no more bidders (if $n \leq m$); without loss of generality we can assume $n \geq m$. Note that there are $n - j + 1$ bidders for auction j .

Note that in our model the private values are not correlated across the m objects. While in general this may not be the case, the special case of uncorrelated values captures the real-life scenarios in which the objects are sufficiently dissimilar, and also provides an interesting technical challenge.

3 Equilibrium Bids

In this section, we find the equilibrium bids for our setting. Since there is more than one auction, a bidder's behavior in an auction depends not only on that auction but also on the profit he expects to get from the future auctions. We first determine this profit and then find the equilibrium bids.

For a series of m auctions with n bidders in the first one, let $EP(j, m, n)$ denote the winner's expected profit for the j th auction, let $ER(j, m, n)$ denote the expected revenue of the j th auction, and let $ES(j, m, n) = ER(j, m, n) + EP(j, m, n)$ be the expected surplus from the j th auction, i.e., the total increase in social welfare that results from allocating the j th object. In what follows, we express these quantities as certain functions of the parameters of the problem and use them to derive an explicit expression for each bidder's bidding strategy.

Let $E(f_i^n)$ and $E(s_i^n)$ denote the expected first and second order statistics for n draws from the distribution V_i , i.e., set $X_i^n = \{x^1, \dots, x^n \mid x^j \sim V_i\}$ and let $f_i^n = \max X_i^n$, $s_i^n = \max X \setminus \{f_i^n\}$.

For any $1 \leq j \leq y \leq m$, let $\beta(y, j, m, n)$ denote a bidder's ex-ante probability of winning the y th auction in the series from the j th to the m th one before the j th auction begins. When the number of bidders and objects is fixed, we sometimes write $\beta(y, j)$ instead of $\beta(y, j, m, n)$.

For instance, $\beta(1, 1, m, n)$ is the probability of winning the first auction in the series of auctions from the first to the m th one. Since $\beta(1, 1, m, n)$ is the ex-ante probability (i.e., before the bidders draw their values for the first auction), each bidder has equal chances of winning the first auction, i.e., $\beta(1, 1, m, n) = 1/n$. Similarly, $\beta(2, 1, m, n)$ is the probability that a bidder wins the second auction in the series of auctions from the first to the m th one. We have $\beta(2, 1, m, n) = (1 - 1/n)(1/(n - 1)) = 1/n$: to win the second auction, a bidder has to lose the first one, which has probability $1 - 1/n$, and win the second one, which happens with probability $1/(n - 1)$.

In the same vein, we have

$$\beta(y, 1) = \frac{1}{n - y + 1} \prod_{k=1}^{y-1} \left(1 - \frac{1}{n - k + 1}\right) = \frac{1}{n}.$$

In general, for $j \leq y \leq m$, $\beta(y, j)$ is given by

$$\beta(y, j) = \frac{1}{n - y + 1} \prod_{k=j}^{y-1} \left(1 - \frac{1}{n - k + 1}\right) = \frac{1}{n - j + 1}.$$

Note that $\beta(y, j, m, n)$ does not depend on y . Intuitively, before the beginning of the j th auction, all bidders are symmetric with respect to winning the y th auction, and there are $n - j + 1$ bidders left at that point. Hence, each bidder's probability of winning the y th auction is $1/(n - j + 1)$.

Let $\alpha(j, m, n)$ denote a bidder's ex-ante expected profit from winning any one auction in the series of auctions from the j th (for $1 \leq j \leq m$) to the m th one. We have

$$\alpha(j, m, n) = \sum_{y=j}^m \beta(y, j) EP(y, m, n) = \frac{1}{n - j + 1} \sum_{y=j}^m EP(y, m, n). \quad (1)$$

Note that by definition, $\alpha(m + 1, m, n) = 0$ for any $n > 0$.

Theorem 1. *If each auction in a series is conducted using the second price rules, then the equilibrium bidding strategy for a bidder whose value in auction j is v is given by*

$$B_j(v) = \max\{0, v - \alpha(j + 1, m, n)\} \quad (2)$$

Proof. In order to find the equilibrium strategies, we begin with the last auction and then reason backwards. Recall that a bidder comes to know his valuation v just before auction j begins (i.e., after the previous $j - 1$ auctions are over).

Consider auction m . The number of bidders for this auction is $n - m + 1$. Since this is the last auction, the bidding strategies for it are the same as those for a single

object auction [1], i.e., each bidder's equilibrium strategy is to bid his true value $v = v - \alpha(m+1, m, n)$.

Now consider auction j ($1 \leq j < m$). Consider bidder 1 whose value is v and set $x = v - \alpha(j+1, m, n)$. Let $b^* = \max_{j \neq 1} b_j$ be the highest competing bid. Assume for simplicity that the draw resolution rule is lexicographic, i.e., bidder 1 wins as long as $x \geq b^*$; the analysis for other draw resolution rules is similar.

If $x < 0$, it means that the bidder's expected profit from the future auctions exceeds his valuation for the object that is being auctioned, so he prefers not to win the object. Therefore, bidder 1's equilibrium strategy is to make the lowest possible bid, i.e., 0.

Otherwise, if $x \geq b^*$, then by bidding $z = x$ bidder 1 wins and his profit from the current auction is $v - b^*$. As $x = v - \alpha(j+1, m, n) \geq b^*$, we have $v - b^* \geq \alpha(j+1, m, n)$, i.e., the bidder weakly prefers winning this auction to participating in the future auctions. Now, if bidder 1 bids any other amount $z > b^*$, he still wins and his profit does not change, whereas if he bids $z < b^*$, he loses the auction, so his expected profit is $\alpha(j+1, m, n) \leq v - b^*$.

If $x < b^*$, then by bidding $z = x$ or, in fact, any $z < b^*$, bidder 1 loses the auction, so his total profit from the entire series is $\alpha(j+1, m, n)$. If he bids $z \geq b^*$, he wins, but his total profit is $v - b^* < \alpha(j+1, m, n)$, i.e., this outcome is less preferable than losing the current auction. In all cases, bidding $z \neq x$ may decrease the bidder's profit, but cannot increase it, i.e., bidding $\max\{x, 0\}$ is an equilibrium strategy.

To characterize the bidding strategies, it remains to show how to compute $\alpha(j, m, n)$. For $j = m$ we have

$$\begin{aligned} ES(m, m, n) &= E(f_m^{n-m+1}) \\ ER(m, m, n) &= E(s_m^{n-m+1}) \\ EP(m, m, n) &= E(f_m^{n-m+1}) - E(s_m^{n-m+1}). \end{aligned}$$

For $j < m$, the value of $\alpha(j, m, n)$ can be computed inductively: in what follows, we describe how to compute $\alpha(j, m, n)$ given $\alpha(y, m, n)$ for $y = j+1, \dots, m$.

Fix $j < m$ and set $f = f_j^{n-j+1}$, $s = s_j^{n-j+1}$. We will consider three cases.

- All bidders bid 0.

This happens with probability $P_0 = (V_j(\alpha(j+1, m, n)))^{n-j+1}$. In this case, the item gets allocated to a random bidder who pays nothing. Set

$$E_0 = E(V_j | f < \alpha(j+1, m, n)).$$

We have

$$ES(j, m, n) = E_0, \quad ER(j, m, n) = 0, \quad EP(j, m, n) = E_0.$$

- Exactly one bidder makes a positive bid.

The probability of this event is $P_1 = (n-j+1)(V_j(\alpha(j+1, m, n)))^{n-j}(1 - V_j(\alpha(j+1, m, n)))$. In this case, the object is allocated to the bidder with a positive bid and the winner pays nothing. Set

$$E_1 = E(f_j^{n-j+1} | s < \alpha(j+1, m, n) < f).$$

We have

$$ES(j, m, n) = E_1, ER(j, m, n) = 0, EP(j, m, n) = E_1.$$

- Two or more bidders make a positive bid.

The probability of this event is $P_{>1} = 1 - P_0 - P_1$. Set

$$\begin{aligned} E_{>1,s} &= E(f|\alpha(j+1, m, n) < s) \\ E_{>1,r} &= E(s|\alpha(j+1, m, n) < s) - \alpha(j+1, m, n). \end{aligned}$$

We have

$$ES(j, m, n) = E_{>1,s}, ER(j, m, n) = E_{>1,r}, EP(j, m, n) = E_{>1} = E_{>1,s} - E_{>1,r}.$$

By combining these three cases, we obtain

$$EP(j, m, n) = P_0 E_0 + P_1 E_1 + P_{>1} E_{>1}, \quad (3)$$

i.e., given $\alpha(j+1, m, n)$ we can compute $EP(j, m, n)$. Hence, given $\alpha(y, m, n)$ for $y = j+1, \dots, m$, we can compute $\alpha(j, m, n)$ using formula (1).

Assuming that bidders' valuations are such that in each auction at least two bidders submit a strictly positive bid, i.e., the revenue of each auction is non-zero, this formula can be simplified considerably. Namely, in this case we have

$$\begin{aligned} ES(j, m, n) &= E(f) \\ ER(j, m, n) &= E(s) - \alpha(j+1, m, n) \\ EP(j, m, n) &= E(f) - E(s) + \alpha(j+1, m, n), \end{aligned}$$

where $f = f_j^{n-j+1}$, $s = s_j^{n-j+1}$. The advantage of this expression is that it does not use conditional expectations. In fact, the only information about the j th distribution that is required is the expectations of the first and second order statistics for $n-j+1$ bidders. Moreover, for large values of bidders, it is quite likely that each auction has non-zero revenue. Indeed, as the number of bidders increases, the profits from future auctions have to be divided among a higher number of potential winners, reducing α .

4 Choosing the optimal agenda

In this section, we consider the problem of choosing the agenda so as to maximize the seller's profit. Our focus here is on those cases where, for each auction, the two highest bids are greater than zero – i.e., all auctions have non-zero profit. Recall that we have seen that for a large value of n this scenario is quite likely.

To see how agenda can affect the revenue, consider first a simple example.

Example 1. Suppose that there are 2 items A and B and 3 bidders. The bidders valuation for A are drawn from $U[8, 20]$ and the bidders valuations for B are drawn from $U[3, 39]$. Suppose we sell these items in the order A, B . Our expected revenue from the second auction is $s_B^2 = 15$ and we have $\alpha(2, 2, 3) = (f_B^2 - s_B^2)/2 = 6$. Therefore, our expected revenue from the first auction is $s_A^3 - 6 = 8$, and the total revenue is 23. If, on the other hand, we sell these items in the order B, A , our expected revenue from the second auction is $s_A^2 = 12$ and we have $\alpha(2, 2, 3) = (f_A^2 - s_A^2)/2 = 2$. Therefore, our expected revenue from the first auction is $s_B^3 - 2 = 19$ and the total revenue is 31.

One can see that selecting the wrong agenda can substantially decrease the revenues. This motivates the following question: Given bidder's value distributions for all items, is it possible to select the optimal agenda in time polynomial in n and m ? Clearly, if we simply consider all possible agendas and compute the expected revenue for each one using the formulas derived in the previous section, we will identify the optimal agenda, but the running time of this procedure is exponential in m . In what follows, we show how to solve this problem more efficiently.

We start by introducing additional notation. Assume that the objects are numbered from 1 to m ; the optimal agenda is then a permutation of $1, \dots, m$. Let $\delta_1(i, n) = E(s_i^n) - E(s_i^{n-1})$ and $\delta_2(i, n) = E(f_i^n) - E(s_i^n)$. Also, as in this section the number of bidders and objects is fixed, we will write $\beta(y, j)$, $EP(k)$, and $ER(k)$ instead of $\beta(y, j, m, n)$, $EP(k, m, n)$, and $ER(k, m, n)$, respectively.

Proposition 1. *Consider two agendas $A^{(1)}$ and $A^{(2)}$ for the same set of m objects such that $A^{(2)}$ can be obtained from $A^{(1)}$ by changing the order of the first and second auction. Let i and j be the objects sold at the first and second auction in $A^{(1)}$, respectively. Let $R^{(k)}$, $k = 1, 2$, be the seller's expected total revenue from $A^{(k)}$. Then $R^{(1)} > R^{(2)}$ as long as*

$$[\delta_1(i, n) - \delta_1(j, n)] - \beta(2, 2)[\delta_2(i, n - 1) - \delta_2(j, n - 1)] > 0.$$

Proof. Let $\alpha = \alpha(3, m, n)$ be a bidder's expected ex ante profit from the last $m - 2$ auctions; obviously, this number is the same for $A^{(1)}$ and $A^{(2)}$. Also, let $\beta = \beta(2, 2) = 1/(n - 1)$; note that β does not depend on the agenda.

Under $A^{(1)}$, the seller's expected revenue from the second auction is $R_2^{(1)} = E(s_j^{n-1}) - \alpha$, and the winner's expected profit from the second auction is $E(f_j^{n-1}) - E(s_j^{n-1}) + \alpha = \delta_2(j, n - 1) + \alpha$. Therefore, in the first auction a bidder whose value is v is going to bid

$$v - (\delta_2(j, n - 1) + \alpha)\beta - \sum_{y=3}^m \beta(y, 2)EP(y);$$

note that the last term in this expression is independent of the agenda. Finally, the seller's expected revenue from the first auction is

$$R_1^{(1)} = E(s_i^n) - (\delta_2(j, n - 1) + \alpha)\beta - \sum_{y=3}^m \beta(y, 2)EP(y).$$

Similarly, under $A^{(2)}$, the seller's expected revenue from the second auction is $R_2^{(2)} = E(s_i^{n-1}) - \alpha$, and the winner's expected profit from the second auction is $E(f_i^{n-1}) - E(s_i^{n-1}) + \alpha = \delta_2(i, n - 1) + \alpha$. Therefore, in the first auction a bidder whose value is v is going to bid

$$v - (\delta_2(i, n - 1) + \alpha)\beta - \sum_{y=3}^m \beta(y, 2)EP(y).$$

Hence, the seller's expected revenue from the first auction is

$$R_1^{(2)} = E(s_j^n) - (\delta_2(i, n-1) + \alpha)\beta - \sum_{y=3}^m \beta(y, 2)EP(y).$$

Under both agendas, the seller's expected profit from the last $m-2$ auctions is the same. Hence, we have $R^{(1)} - R^{(2)} = R_1^{(1)} - R_1^{(2)} + R_2^{(1)} - R_2^{(2)}$. It is easy to see that $R_2^{(1)} - R_2^{(2)} = E(s_j^{n-1}) - E(s_i^{n-1})$. Furthermore, we have

$$R_1^{(1)} - R_1^{(2)} = E(s_i^n) - \beta\delta_2(j, n-1) - E(s_j^n) + \beta\delta_2(i, n-1).$$

Hence, we conclude that $R^{(1)} > R^{(2)}$ if and only if

$$\delta_1(i, n) - \delta_1(j, n) + \beta(\delta_2(i, n-1) - \delta_2(j, n-1)) > 0.$$

Proposition 1 describes the change in the revenue that corresponds to switching the order of the first two auctions on the agenda. If we change the relative order of the $(k-1)$ st and k th auction, we can still use Proposition 1 to compare the revenue from the last $m-k+2$ auctions under the two agendas. However, to compare the total revenues, we need additional techniques, as the choice of the agenda will influence the bidders' behavior in the first $k-2$ auctions. Fortunately, it turns out that as long as changing the order of the $(k-1)$ st and k th auction increases the revenue from the last $m-k+2$ auctions, it increases the total revenue as well.

Proposition 2. Consider two agendas $A^{(1)}$ and $A^{(2)}$ for the same set of m objects such that $A^{(2)}$ can be obtained from $A^{(1)}$ by changing the order of the $(k-1)$ st and k th auction. Let i and j be the objects sold at the $(k-1)$ st and k th auction in $A^{(1)}$, respectively. Let $R^{(i)}$, $i = 1, 2$, be the seller's expected total revenue from $A^{(i)}$, and let $R_{k-1,k}^{(i)}$, $i = 1, 2$, be the seller's expected revenue from the $(k-1)$ st and k th auction in $A^{(i)}$. Then $R^{(1)} > R^{(2)}$ whenever $R_{k-1,k}^{(1)} > R_{k-1,k}^{(2)}$.

Proof. For $i = 1, 2$, we have $R^{(i)} = \sum_{j=1}^m ER^{(i)}(j)$, where $ER^{(i)}(j)$ is the expected seller's revenue in the j th auction under agenda $A^{(i)}$. Clearly, for $j > k$ we have $ER^{(1)}(j) = ER^{(2)}(j)$, and we assume

$$ER^{(1)}(k-1) + ER^{(1)}(k) > ER^{(2)}(k-1) + ER^{(2)}(k).$$

We will now prove that for any $0 < j < k-1$, if

$$\sum_{t=j+1}^m ER^{(1)}(t) > \sum_{t=j+1}^m ER^{(2)}(t),$$

then

$$\sum_{t=j}^m ER^{(1)}(t) > \sum_{t=j}^m ER^{(2)}(t).$$

Applying this result inductively to $j = k-2, k-3, \dots, 1$ completes the proof.

Fix some $j < k - 1$. Let $\alpha^{(1)}(j)$ and $\alpha^{(2)}(j)$ be a bidder's expected ex ante profit from the auctions j, \dots, m under $A^{(1)}$ and $A^{(2)}$, respectively. For $i = 1, 2$, we have

$$\alpha^{(i)}(j) = \frac{1}{n - j + 1} \sum_{y=j}^m EP^{(1)}(y).$$

Under both agendas, the expected total surplus from the last $n - j$ auctions is the same, namely, $\sum_{t=j+1}^m E(f_t^{n-t+1})$. Hence, we have

$$\sum_{t=j+1}^m [ER^{(1)}(t) + EP^{(1)}(t)] = \sum_{t=j+1}^m [ER^{(2)}(t) + EP^{(2)}(t)].$$

Therefore,

$$\sum_{t=j+1}^m ER^{(1)}(t) > \sum_{t=j+1}^m ER^{(2)}(t)$$

implies

$$\sum_{t=j+1}^m EP^{(1)}(t) < \sum_{t=j+1}^m EP^{(2)}(t),$$

and consequently $\alpha^{(1)}(j + 1) < \alpha^{(2)}(j + 1)$. As $ER^{(i)}(j, m, n) = E(s_j^{n-j+1}) - \alpha^{(i)}(j + 1)$, it follows that in this case $ER^{(1)}(j) > ER^{(2)}(j)$. Hence,

$$\sum_{t=j}^m ER^{(1)}(t) > \sum_{t=j}^m ER^{(2)}(t),$$

as required.

We summarize our results in the following theorem.

Theorem 2. *For each object $j = 1, \dots, m$, define $r_j = \delta_1(j, n) - \delta_2(j, n - 1)/(n - 1)$. Then an optimal agenda can be obtained by ordering the auctions in order of non-increasing r_j .*

Clearly, this means that one can find an optimal agenda in polynomial time. Moreover, the only information about the distributions that is required is their first and second order statistics.

To illustrate Theorem 2, consider the case when the value distributions for two consecutive auctions i and j are $U[a, a + x]$ and $U[b, b + y]$, respectively, and $x > y$. We have $\delta_1(i, n) = x \frac{n-1}{n+1} - x \frac{n-2}{n} = \frac{2x}{n(n+1)}$, $\delta_1(j, n) = y \frac{n-1}{n+1} - y \frac{n-2}{n} = \frac{2y}{n(n+1)}$, $\delta_2(i, n) = \frac{x}{n+1}$, $\delta_2(j, n) = \frac{y}{n+1}$. Consequently, $\delta_1(i, n) - \delta_1(j, n) > 0$ if and only if $x > y$, and also $\delta_2(i, n) - \delta_2(j, n) > 0$ if and only if $x > y$. We conclude that if all values are drawn from uniform distributions the auctions should be run in the order of non-increasing distribution support size (or, equivalently, non-increasing variance), whereas the expected value of each object has no effect on the optimal ordering. This explains why in Example 1 the ordering B, A produced a higher revenue than A, B .

5 Selling a subset of items

We have seen that changing the agenda may considerably increase the revenue in a sequential auction. However, in some cases this approach is not feasible, since the objects have to be sold in a fixed order (e.g., they become available one by one and expire shortly after becoming available). In this case, we can try to increase the revenue by selling a subset of all available items. This approach is based on the idea that reducing supply may motivate the bidders to bid more aggressively. In this section, we assume that the agenda is fixed and consider two ways of deciding which auctions to cancel.

5.1 Dynamic cancellations

As the bidders shade their bids since they expect to profit from the future auctions, it may happen that in some auction in the series there is at most one strictly positive bid. In this case, the seller may be tempted to cancel the auction: he receives no revenue in the current auction anyway, and moreover, by doing so he increases the number of bidders in subsequent auctions (and hence the expected revenue from these auctions). In some cases, doing so may even increase the social welfare: when all bidders bid 0, the object is assigned to a random bidder who may have very little value for it, and there is a chance that if he is not eliminated now, in the future he will have a very high value for another object. However, one has to take into account that changing the auction rules will affect the bidders' behavior. In particular, if canceling auction j increases the bidders' expected profit from the last $m - j + 1$ auctions, the bidders will shade their bids more heavily in the first $j - 1$ auctions, and the net impact on the auction revenue is unclear. We illustrate the changes in the bidders' strategies with the following example.

Example 2. Consider a sequential auction with 2 items and 3 bidders, where the bidders' values for each object are distributed as $U[0, 1]$. It is easy to see that if the second auction has 2 participants, each bidder's ex ante expected profit from this auction is $1/6$, whereas if it has 3 participants, the ex ante expected profit is $1/12$. Hence, under the original rules, in the first auction a bidder with value v will bid $\max\{v - \alpha_1, 0\}$, where $\alpha_1 = 1/6$ is his expected profit from the second auction.

If the seller is allowed to cancel the first auction as long as he gets no profit from it, the bidder will bid $\max\{v - \alpha_2, 0\}$, where α_2 is his expected profit from the second auction under the new rules. We have

$$\alpha_2 = \frac{P_c}{12} + \frac{1 - P_c}{6}, \quad (4)$$

where P_c is the probability that the first auction is canceled. The quantity P_c can be expressed as a function of α_2 : a bidder bids 0 if his value is at most α_2 , so we have

$$P_c = \alpha_2^3 + 3(1 - \alpha_2)\alpha_2^2. \quad (5)$$

Combining (4) and (5), we obtain $2\alpha_2^3 - 3\alpha_2^2 - 12\alpha_2 + 2 = 0$. Solving this numerically and taking into account that $P_c = 2 - 12\alpha_2 \in [0, 1]$, we obtain $\alpha_2 \approx 0.161 < 1/6$. Hence, in this case, by allowing the seller to cancel the first auction, we increase his expected revenue both in the second auction (since there is some chance that it will have three bidders instead of two) and in the first auction (since the bidders expect less profit from the second auction, so they shade less).

The method for computing the bidding strategies described in Example 2 can be generalized to more than two items and general distributions. However, the bidders' computational problem becomes quite complex. Moreover, it is not clear if the inequality $\alpha_1 > \alpha_2$ holds in general. Therefore, it is not necessarily the case that the new format increases the seller's total revenue. Therefore, in practice, the seller may want to pre-compute the expected revenue from both formats and pick the better one rather than assume that canceling non-profitable auctions is always beneficial.

5.2 Static cancellations

A related, but easier-to-analyze approach is to cancel some auctions in advance based on their value distributions. For example, if the number of items m is equal to the number of bidders n , then the last auction will only have one participant who will therefore get the item for free. Hence, it is clear that in this case the auctioneer should sell at most $n - 1$ items. Moreover, even when $n > m$, it may be profitable for the auctioneer to sell less than m items.

Example 3. Consider a sequential auction with 3 bidders and 2 objects, where the values for auction i are drawn from a distribution F_i . Assume for now that in the first auction all bids are positive (the specific distributions we construct later satisfy this property). If both auctions are executed, the seller's expected revenue is $E(s_2^2) + E(s_1^3) - [E(f_2^2) - E(s_2^2)]/2$. If the seller only sells one object, his expected revenue is $E(s_1^3)$. Hence, it is more profitable to sell one object if $3E(s_2^2) < E(f_2^2)$.

Consider the probability distribution on $[0, 1]$ whose cumulative density function is given by $V_a(x) = x^a$. It is easy to check that

$$E(f^2) = \int_0^1 (2ax^{2a})dx = \frac{2a}{2a+1}$$

$$E(s^2) = \int_0^1 (-2ax^{2a} + 2ax^a)dx = \frac{2a^2}{(a+1)(2a+1)}.$$

We have $E(f^2)/E(s^2) = (a+1)/a$, so for $a < 1/2$ we have $3E(s_2^2) < E(f_2^2)$. Fix $a = 1/3$, and shift this probability distribution by 1 i.e., consider the distribution on $[1, 2]$ with probability density function $W_{1/3}(x) = V_{1/3}(x-1)$. If the probability distribution for the first object is $W_{1/3}(x)$ and the probability distribution for the second object is $V_{1/3}(x)$, then in the first auction all bids are positive as required, and canceling the second auction increases the total revenue. Note also that in this setting canceling the second auction is more profitable than canceling the first one.

Intuitively, under this distribution, the bidders expect to profit considerably in the second auction, and therefore shade heavily in the first auction, while our own profit from the second auction is relatively small. Therefore, one cannot assume that it is always profitable to sell at least $n - 1$ items, where n is the number of bidders.

Consequently, the seller needs to identify the optimal subset of items to sell. Computing the expected revenue for all possible subsets is not feasible, as there are exponentially many of them. While we do not have an exact solution for this problem, in what follows, we sketch an efficient algorithm that finds an approximately optimal subset

(i.e., one with almost optimal total revenue). Our algorithm is based on dynamic programming. Set $C = m \max_{i=1, \dots, m} E(f_i^n)$. Clearly, C is an upper bound on all participants' profit from the entire series. Fix a parameter $\epsilon = 1/N$, which corresponds to the approximation error. For $i = 1, \dots, n, j = 1, \dots, m, k = 0, \dots, NC$, we fill in the array $TR(K, i, j)$. Intuitively, the values of $TR(K, i, j)$ are approximations to the seller's maximal expected revenue from the last $n - i + 1$ auctions, assuming that there are j bidders before the start of the i th auction, and $K\epsilon \leq \sum_{y=i}^m EP(y, m, n) < (K + 1)\epsilon$. The values of $TR(K, i, j)$ can be computed inductively starting with $TR(K, n, j)$. In the end, we pick the entry $TR(K, 1, j)$ that corresponds to the maximal feasible revenue. For reasonable probability distributions, we can bound the error caused by rounding the winners' profit to the nearest multiple of ϵ . We omit the details due to lack of space.

6 Asymmetric bidders and restricted orderings

In many real-life scenarios, some of the assumptions made in this paper may fail to hold. In particular, in the advertisement scheduling problem described in the introduction, the bidders are not necessarily symmetric, i.e., different bidders's valuations for the same slot can differ even ex ante. Also, the auctioneer may not have full control over the ordering of the auctions. For example, in a typical TV show, some episodes have to be screened in a certain order. Unfortunately, in this more general setting, the problem of selecting the optimal agenda becomes NP-complete, even if we assume that we know each bidder's value for each item, and all bidders are myopic, i.e., they bid truthfully in each auction.

An instance of the problem OPTIMAL AGENDA WITH EXPIRATION DATES is given by n bidders, m items, mn integers $v(i, j)$, $i = 1, \dots, n, j = 1, \dots, m$, m integers $e(1), \dots, e(m)$, where $e(i) \in \{1, \dots, m\}$ for $i = 1, \dots, m$, and a rational number R . We interpret these numbers as follows: the i th bidder's valuation for the j th item is $v(i, j)$ if the j th item is sold at one of the first $e(j)$ auctions and 0 otherwise. It is a "yes"-instance if there exists an agenda that results in a total revenue of at least R and a "no"-instance otherwise.

Theorem 3. OPTIMAL AGENDA WITH EXPIRATION DATES is NP-complete.

The proof of this theorem can be found in the appendix. This hardness result explains why in most of the paper we chose to focus on the simplified case.

7 Conclusions and open problems

We have studied two methods for increasing the revenue of a sequential auction: choosing an optimal agenda, and, for a fixed agenda, identifying an optimal subset of objects to sell. It would be interesting to see if one can achieve further improvements in total revenue by combining these methods, i.e., choosing the subset of objects and then selecting the optimal ordering in which to sell these objects. Another tool for increasing revenue that should be studied in this context is reserve prices. Also, it would be useful to identify a large class of distributions for which canceling zero revenue auctions is guaranteed to increase the total revenue. An interesting special case of our problem is when all objects have the same value distribution. In this case, the seller simply has to decide how many objects to sell. Currently, we do not know if in this case it is always profitable to sell exactly $\min\{m, n - 1\}$ items, or there are value distributions for which the seller may want to sell fewer items.

References

1. W. Vickrey, Counterspeculation, auctions and competitive sealed tenders, *Journal of Finance*, 16:8–37, 1961.
2. J. Galambos, *The asymptotic theory of extreme order statistics*, John Wiley and Sons, 1978.
3. J. Edmonds, Optimum branchings. *Journal of Research of the National Bureau of Standards*, 1967
4. D. Bernhardt and D. Scoones, A note on sequential auctions, *American Economic Review*, 84(3): 653–657, 1994. number = 3,
5. P. Dasgupta and E. Maskin, Efficient auctions, *Quarterly Journal of Economics*, 115: 341–388, 2000.
6. M. P. Wellman and W. E. Walsh and P. R. Wurman and J. K. McKie-Mason, Auction protocols for decentralised scheduling, *Games and Economic Behavior*, 35: 271–303, 2001.
7. T. Sandholm and S. Suri, BOB: Improved winner determination in combinatorial auctions and generalizations, *Artificial Intelligence*, 145: 33–58, 2003.
8. W. Elmaghraby, The importance of ordering in sequential auctions, *Management Science*, 49(5):673–682, 2003.

A Hardness result for asymmetric case

Theorem 3. OPTIMAL AGENDA WITH EXPIRATION DATES is NP-complete.

Proof. It is easy to check that this problem is in NP: given a candidate agenda, one can efficiently compute the respective total revenue.

To prove NP-hardness, we reduce 3-SAT to OPTIMAL AGENDA WITH EXPIRATION DATES. Recall that an instance of 3-SAT is given by N variables x_1, \dots, x_N and M clauses c_1, \dots, c_M , where each clause consists of at most three literals. It is a “yes”-instance if there is a way to set the variables to 1 (“true”) or 0 (“false”) so that all clauses are satisfied.

Given an instance of 3-SAT, we construct an instance of OPTIMAL AGENDA WITH EXPIRATION DATES as follows. There are two groups of items that correspond to variables and clauses, respectively. Namely, for each variable x_i , there are two items X_i and \bar{X}_i , and for each clause c_j , there is an item C_j . The expiration date for all items of the form $X_i, \bar{X}_i, i = 1, \dots, N$, is N , the expiration date for all other items is $2N + M$. Similarly, the bidders can also be subdivided into two groups. namely, for each variable x_i there are two bidders b_i and \bar{b}_i (we will refer to them as the literal bidders), and for each clause c_j there is a bidder a_j (the clause bidders).

We will now describe the valuations of all bidders. To simplify notation, all valuations that are not described explicitly are equal to 0. For $X_i, i = 1, \dots, N$, we have $v(b_i, X_i) = 2M + 1, v(\bar{b}_i, X_i) = 2M$. For $\bar{X}_i, i = 1, \dots, N$, we have $v(b_i, \bar{X}_i) = 2M, v(\bar{b}_i, \bar{X}_i) = 2M + 1$. For any C_j we have $v(a_j, C_j) = 2$. Moreover, if $c_j = l_{i_1} \vee l_{i_2} \vee l_{i_3}$, where $l_{i_j} \in \{x_{i_j}, \bar{x}_{i_j}\}$, we set $v(b, C_j) = 1$ if and only if $b = b_i, i \in \{i_1, i_2, i_3\}$ and $l_i = x_i$ or $b = \bar{b}_i, i \in \{i_1, i_2, i_3\}$ and $l_i = \bar{x}_i$. Set $R = 2MN + M$.

Intuitively, under the optimal agenda we have to use the first N slots to sell N objects from the set $\{X_i, \bar{X}_i \mid i = 1, \dots, N\}$. Moreover, it is clear that after selling X_i , we can make no profit from selling \bar{X}_i , as there is only one bidder interested in this object, and vice versa. Therefore, we will not schedule both of these auctions during the first N slots. Hence, after the first N auctions we eliminate exactly one of the bidders b_i, \bar{b}_i for each $i = 1, \dots, N$, which is equivalent to choosing the truth values for x_1, \dots, x_N .

Moreover, the auction for C_j is profitable if and only if at least one of the three bidders that correspond to literals in c_j is still present, i.e., under the truth assignment induced by the first N auctions the clause c_j is satisfied.

Formally, suppose that the given instance of 3-SAT is satisfiable and let $\rho, \rho(x_i) = 0, 1$, be the corresponding truth assignment. Consider the following agenda: for $i = 1, \dots, N$, if $\rho(x_i) = 0$, use the i th slot to sell X_i , and if $\rho(x_i) = 1$, use the i th slot to sell \bar{X}_i ; use the next M slots to sell C_1, \dots, C_M (in arbitrary order); use the last N slots to sell the remaining items. Clearly, the total revenue from the first N auctions is $2NM$. Now, consider the auction for C_j . The clause c_j evaluates to 1 under the truth assignment ρ , so there must exist a literal l_k that appears in c_j such that $\rho(x_k) = 1$ if $l_k = x_k$ and $\rho(x_k) = 0$ if $l_k = \bar{x}_k$. Suppose for simplicity that $l_k = x_k$; the argument for the case $l_k = \bar{x}_k$ is similar. The bidder b_k participates in the auction for C_j : indeed, in auctions $1, \dots, k-1, k+1, \dots, N, \dots, N+j-1$ his bid was 0, so he did not win any of them, and since $\rho(x_k) = 1$, the k th auction was for \bar{X}_k , so it was won by \bar{b}_k . Hence, the auction for C_j has at least two participants, namely, a_j and b_k , so its revenue is at least 1.

Conversely, suppose that there is an agenda that results in a revenue of at least $2MN + M$. It is easy to see that this means that the first N slots are used to sell exactly one element of each pair (X_i, \bar{X}_i) , $i = 1, \dots, N$. Indeed, any other agenda results in a revenue of at most $2M(N-1) + 1$ from the first N auctions, and the only items that can be profitably sold during the last $M + N$ auctions are C_1, \dots, C_M , each of which contributes at most 1 to the total revenue. Therefore, the total revenue for any other agenda is at most $2MN - 2M + 1 + M < 2MN + M$. Moreover, the total revenue from the last $N + M$ auctions is at least M . As selling any item of the form X_i, \bar{X}_i , $i = 1, \dots, N$ during these auctions results in zero revenue, it follows that the revenue from the sale of each C_1, \dots, C_M is 1.

It follows that we can define a valid truth assignment ρ by setting $\rho(x_i) = 1$ if \bar{X}_i is sold in one of the first N auctions, and $\rho(x_i) = 0$ if X_i is sold in one of the first N auctions. Now consider a clause $c_j = l_{i_1} \vee l_{i_2} \vee l_{i_3}$. The revenue from the corresponding item C_j is 1, so at least one of the bidders that correspond to a literal in c_j was present during the auction for C_j . Therefore, the value of this literal under the truth assignment ρ is 1 and the clause c_j is satisfied.