Efficiency Levels in Sequential Auctions with Dynamic Arrivals

Ron Lavi
Faculty of Industrial Engineering and Management
The Technion – Israel Institute of Technology
ronlavi@ie.technion.ac.il

Ella Segev
Faculty of Industrial Engineering and Management
The Technion – Israel Institute of Technology
esegev@ie.technion.ac.il

August 11, 2008

Abstract

In an environment with dynamic arrivals of players who wish to purchase only one of multiple identical objects for which they have a private value, we analyze a sequential auction mechanism with an activity rule. If the players play undominated strategies then we are able to bound the efficiency loss compared to an optimal mechanism that maximizes the total welfare. We have no assumptions on the underlying distribution from which the players' arrival times and valuations for the object are drawn. Moreover we have no assumption of common prior on this distribution.

JEL Classification Numbers: C70, D44, D82

Keywords: Sequential Ascending Auctions, Undominated Strategies, Dynamic Arrivals.

1 Introduction

Models of dynamic mechanism design are attracting much attention in recent years, partly because of the new electronic markets over the Internet, that are much more dynamic than classic marketplaces. For example, Athey and Segal (2007) and Bergemann and Välimäki (2007) study a general multi-period allocation model, in which a designer needs to perform allocation decisions in each period $t = 1, 2, ...$, and players have private values for the different allocations, that may be stochastic and time-dependent. This model includes the possibility that players may arrive dynamically over time, by setting their valuations to zero until their "arrival time". This property of dynamic arrivals seems to be particularly appealing when one considers the model of sequential auctions, which is a special case of the general model. Indeed, the assumption of a static set of players throughout the sequence of auctions, which is the prevailing assumption in the classic literature on sequential auctions (a long sequence of works, initiated by Milgrom and Weber (2000)), seems to be problematic for the Internet reality. A recent paper by Said (2008) shows explicitly the connection between the general results of Bergemann and Välimäki (2007) and the model of sequential ascending auctions. These results show that full efficiency can be achieved under the assumption that the mechanism has correct information about future arrivals[1].

Another problematic aspect of the classic works on sequential auctions is the common-belief assumption, used to construct a Bayesian-Nash equilibrium. It has long been argued that assuming common knowledge of a prior (let alone a common belief) is problematic, see for example Wilson (1987) and Bergemann and Morris (2005). Requiring the players to be coordinated with respect to their beliefs about the types (and number) of players that will arrive in the future seems like an unacceptably strong and unrealistic assumption. Moreover, such an assumption critically affects the nature of the results regarding the efficiency of the mechanism. With common beliefs, full efficiency can be achieved. In reality, conflicting beliefs, and perhaps even contradicting beliefs (i.e. that cannot be derived from a common prior) do exist in many dynamic settings. This obviously leads to efficiency loss. An important question in such settings is therefore *how much* efficiency may be lost, and what is the effect of the specifics of the mechanism on reducing this loss.

This work describes an analysis of sequential auctions with dynamic arrivals, allowing players to have different, even contradicting, beliefs. As an example, let us consider the following scenario. A seller conducts two ascending ("English") auctions, one after the other. There are two players that participate in the first auction, each player desires one of the two items. The players do not discount time, and they are indifferent between the two items. There is a certain probability that a high-value bidder will join the second auction, and in this case the loser in the first auction will also lose the second auction. Clearly, if a player assigns a high probability to this event, she will be willing to compete (almost) up to her value in the first auction while if she assigns a low probability

[1] In Bergemann and Välimäki (2007) this is achieved by assuming that players report, before they arrive, the probability of their future arrival, while Said (2008) abstracts this from the mechanism by assuming a common-prior.

to this event, she will stop competing in the first auction at a low price. If the two players have significantly different beliefs regarding this event, one will retire early and the other one will win. This has a negative effect on the social welfare when the player with the higher value incorrectly under estimates the probability of the new second-period arrival. This loss of efficiency is *inherent* to the dynamic setting, a real phenomena that we do not wish to ignore, but rather to highlight and analyze. Thus, an important conceptual motivation of our work is the exploration of a way to correctly quantify the efficiency loss experienced in dynamic settings with incomplete information and no common priors.

To describe our exact results, let us state our model slightly more formally: an auctioneer conducts a sequence of K ascending auctions. Players have quasi-linear utilities and private values, with unit-demand, and are indifferent between the K items that are being sold. Each player has an "arrival time", and cannot participate in the auction (or communicate with the auctioneer) before that time. The efficient outcome maximizes the sum of values of players that receive an item (under the requirement that a player cannot be allocated an item before her arrival time). We ask how low can the sum of values be (w.r.t. the optimum) when the allocation is a result of the sequential auction game of incomplete information. As we do not make any assumptions on the priors of the players, we cannot hope to obtain equilibrium strategies in this dynamic setting (as the example above demonstrates). Instead, our analysis reasons about the efficiency loss in two steps: first, we characterize the set of all undominated strategies and, second, we bound the efficiency loss over all tuples of undominated strategies. To obtain a "nice" characterization of all undominated strategies, we add a simple activity rule to the ascending auction – players that drop out "too early" in the auction cannot continue to participate in future auctions (clearly, one has to be careful about what is "too early", as disqualifying high-value players will intensify the efficiency loss).

The efficiency loss over this set of undominated strategies is bounded using a worst-case analysis: imagine an adversary that is allowed to determine the number of players, their arrival times, their values, and their (undominated) strategies, in order to "fail" the auction. We ask what is the lowest possible ratio between the actual sum of values and the optimal sum of values that such an adversary can achieve. We show that no matter how the parameters are set, as long as the chosen strategies are undominated, the sequential auction will obtain at least 50% of the optimal efficiency, i.e. the above mentioned ratio is at least one-half.

What to make of this number - is it low or high? One possible intuitive answer is that, since the worst-possible case is a 50% loss, then on average the loss should be much smaller. To make this intuition more rigorous, we perform the following analysis, for the case of two items. Assume an adversary with powers almost as before, except that, instead of assigning arbitrary worst-case values to the players, the adversary now has to choose (any) probability distribution, and draw the values independently from this distribution. She can set all other parameters freely, as before. We show that, now, the ratio between the expected welfare of the sequential auction and the expected welfare of the optimal allocation is at least 70%. Of-course, this bound is achieved only for the *worst possible distribution*. For example, if we take the uniform distribution over some interval then the ratio will increase to 80%. One can continue further, and obtain the other parameters in a distributional way, and this will most probably decrease the efficiency loss even further.

This average-case analysis seems to strengthen the basic intuition, that a worst-case bound of 50% is a "good" bound, but we do not argue that the sequential ascending auction obtained is the "best" auction; this is a matter of further study. Our main argument is that the issue of different and conflicting beliefs and its influence on the nature of results in the dynamic setting, should be given more weight, and that our analysis is a necessary first step in this direction.

The results are based on our ability to identify a simple property that holds for every undominated strategy of a player: if a player plays an undominated strategy she will never drop out "too" early. In order to achieve this property we must introduce an activity rule. We leave for future research the question whether the activity rule defined here is an "optimal" activity rule in this setting. We wish to highlight the fact that it allows us to analyze the players' behavior and to bound the efficiency loss. Many researchers have argued in favor of using activity rules (e.g. Ausubel and Milgrom (2002)), for different reasons, and we show how this helps to bound the efficiency loss.

Our set of results is different than most results in the auction theory literature in two aspects: (i) it does not obtain equilibrium behavior, and (ii) it does not produce a dichotomic judgement, whether the outcome is efficient or not, but rather it produces a *quantitative* judgement, of how much efficiency is lost. In the context of first-price single-item auctions, a recent line of research by Battigalli and Siniscalchi (2003), Dekel and Wolinsky (2003), and Cho (2005) advocates the shift from equilibrium analysis to an analysis that is based on more fundamental assumptions. They assume that bidders are rational and strategically sophisticated, but they avoid the assumption that bidders share common beliefs. The first difference mentioned above follows their footsteps. However, these works still produce a dichotomic distinction, whether the outcome is efficient or not. We argue that, at least in the context of dynamic settings, such a distinction is too coarse. To differentiate between the many popular auction methods that we see in the real world, we advocate a quantitative assessment of the efficiency loss, as exemplified here.

While we analyze an existing (popular and natural) method to auction a set of items, a different approach would be to design a new mechanism that exhibits dominant strategies (or ex-post equilibrium) and a small efficiency loss in that detail-free equilibrium. This approach was largely adopted in the computer science literature, starting with Lavi and Nisan (2004). For example, Hajiaghayi, Kleinberg, Mahdian and Parkes (2005) show that if prices are charged only after all auctions end (and depend on all the sequence of the auctions) then dominant-strategies can be recovered. Cole, Dobzinski and Fleischer (2008) show that if we can restrict each player to participate in only one particular auction, then a certain choice rule of the "right" auction in which to participate will guarantee a total efficiency loss of at most 50%. Cavallo, Parkes and Singh (2007) suggest an improvement to the marginal contribution mechanism of Bergemann and Välimäki (2007), enabling players to be completely inaccessible before their arrival time. This current paper is mostly related to the work of Lavi and Nisan (2005), who perform a worst-case analysis over a large class of strategies (in contrast to all the previous papers in the CS literature), in a job-scheduling model that is a generalization of our sequential auctions model. Our focus here on the special case of sequential auctions enables us to give a significantly tighter game-theoretic justification to the set of strategies we analyze, as well as to significantly improve the bounds on the efficiency loss that we obtain.

Another modelling approach to dynamic mechanism design is taken by Gershkov and Moldovanu (2007). In their setting, players appear according to some fixed and known stochastic process (e.g. a Poisson arrival rate), and are impatient, i.e. must either be served upon their arrival, or not be served at all. This structure enables Gershkov and Moldovanu (2007) to characterize optimal dynamic mechanisms, with respect to both the social welfare and the seller's revenue.

The remainder of this paper is organized as follows. In section 2 we define our sequential auctions model, and analyze the resulting set of undominated strategies. Section 3 describes an analysis of the worst possible ratio between the resulting welfare in the sequential auction vs. the optimal welfare, assuming players play undominated strategies. Section 4 shows, for the case of two items, that this bound is significantly higher if players' values are drawn independently from some fixed distribution.

2 Strategic Analysis

2.1 Preliminaries

We study a setting where a seller sells K identical and indivisible items using a sequence of K single-item ascending auctions. The set of bidders is $I = \{1, .., N\}$.[2] Each bidder has unit demand, i.e. desires exactly one item out of the K items. We assume the standard model of private values and quasi-linear utilities: a bidder has value v_i for receiving an item; her utility is $v_i - p_i$ if she wins an item and pays p_i, and 0 if she loses. Bidders do not discount time – they are indifferent between winning the object at auction t or at any other auction in which they participate. The dynamic character of our setting allows bidders to arrive over time. Thus bidders (physically) cannot all participate in all auctions. Formally, bidder i's type includes, besides her value v_i, an arrival time r_i which is an integer between 1 and K, indicating that bidder i participates only in the auctions for items $r_i, ..., K$. The interpretation is that bidder i becomes aware of her need for an item just before auction r_i starts, or that some physical limitations prevent her from participating prior to that.

Thus, a bidder's type is a pair $\theta_i = (r_i, v_i)$, and the set of possible types for bidder i is denoted

[2] It will be convenient, and more interesting, to assume that $N >> K$, though our results hold for any N.

by Θ_i. In this section and in section 3 we make no assumptions as to how the bidders' types are drawn, i.e. they can be drawn from any joint distribution over all tuples of types, or even be set by an adversary so as to create the worst possible scenario. Moreover we do not assume a common prior belief for the bidders – they can have different beliefs regarding the underlying distribution of types. In section 4 we will add the assumption that bidders' *values* are drawn independently from some fixed distribution, as will be detailed there. We will evaluate the auction mechanism by comparing the resulting social welfare to the maximal possible social welfare, where the social welfare (also called social efficiency) of a specific allocation is the sum of the winners' values.

Each ascending auction is formally assumed to be a "Japanese" auction, where a "price clock" increments continuously and each bidder is free to drop at any price. The last bidder that remains is the winner, and she pays the price at which the second to last bidder dropped out. If some bidder drops at a certain price, then other bidders may respond by dropping as well. It will be convenient to structure the drops so that the bidders drop one at a time, in a well defined order, as formally described below. We use a "Stopping the Clock" assumption, as in Ausubel (2004), to describe this cascade of dropping decisions: whenever a bidder drops, the price clock pauses to enable other bidders' dropping. The clock then resumes its ascent at the same price where it stopped. The order in which bidders drop at a given price is important in the description of the history, and may determine the winner in case all bidders drop one after the other when the price clock is stopped. During the price ascent, bidders are able to observe how many bidders remain in the auction, at any price. This property will have a particular importance, as will become clear in the sequel. Thus, our auction model is of an extensive form game, and a strategy of a bidder is a function from her type, the history of all previous auctions, and the history of the current auction up to the current price, that outputs a binary decision whether to drop or to remain.

As a concrete example to this process, consider a single auction with four bidders, that have values $v_1 = 5, v_2 = 10, v_3 = v_4 = 8$. Assume the bidders are using the following strategies. Bidder 1 remains until her value, bidder 2 remains until her value, or until there remain at most two other bidders (the earliest of the two events), and bidders 3 and 4 remain until their value, or until there remain at most one other bidder. Given these strategies, the auction will proceed as follows. The price clock will ascend until it reaches a price of 5. Bidder 1 will then drop. The clock will stop to allow other bidders to drop, and indeed, bidder 2 will consequently drop (as after bidder 1 dropped only two other bidders, 3 and 4, still remain). Immediately after bidder 2 drops, both bidder 3 and bidder 4 will announce that they wish to drop. One of them will be chosen to actually drop, depending on the tie-breaking rule that is being used. The winner will be the remaining player. Her price will be 5. Note that, although all bidders dropped at 5, the tie-breaking rule affects the outcome only with respect to the choice between bidders 3 and 4.

The next subsection gives a more formal description of this process. Section 2.3 discusses the possible strategic choices of the bidders, showing that early drops (before one's value) may be a reasonable strategic choice. We introduce an activity rule to enable a tighter set of undominated strategies in section 2.4, and give the analysis itself in section 2.5.

2.2 Formalities

Recall that we conduct K sequential auctions, one for each item. The t'th auction (for any $t = 1, ..., K$) can be described by the sequence of bidder drops, as follows. The price clock ascends continuously until a bidder (or several bidders) announce that they wish to drop, at a price clock p. In this case the price clock is stopped, and one of these players is chosen to be the first "dropper" at price p. More bidders may wish to drop, either as a consequence of this drop, or independently of this event, because they already declared they wish to drop, but were not chosen by the tie-breaking rule. To capture this, the auctioneer asks all bidders simultaneously if they now wish to drop. All bidders reply simultaneously with a "yes" or "no", and again if more than one bidder wishes to drop the tie-breaking rule chooses one of them. The second bidder who dropped out is termed the second "dropper" at price p. This continues until no more players wish to drop, and the price clock then resumes its ascent. We assume that the number of droppers and their identities are public information, but the announcements themselves are not public (and so if a player announced that she wants to drop, but was not chosen to do so, the other players do not observe this). The tie-breaking rule is a total preference order on the set of bidders (i.e. a binary relation on the set of bidders which is reflexive, antisymmetric, transitive and complete) and is assumed to be public information as well.[3]

We denote by $D_t(p, k)$ the k'th dropper at price p in the t'th auction (this is a singleton set). Consider for example the scenario from above. Then $D_1(5, 1) = \{1\}, D_1(5, 2) = \{2\}, D_1(5, 3) = \{3\}$. The total number of droppers at price p is denoted by s_p^t (note that this is always smaller than the number of players). We have that $s_p^t = 0$ if and only if no bidder dropped at p, and then $D_t(p, 0) = \emptyset$. If $s_p^t > 0$ then $D_t(p, k) \neq \emptyset$ for every $k = 1, ..., s_p^t$.

Let X_t denote the set of bidders that participate in auction t, and $x_t = |X_t|$.[4] The entire information on a single auction can be described by the prices at which $D_t(\cdot, \cdot)$ is nonempty (a finite number of prices), the values of $D_t(\cdot, \cdot)$ in these prices and the preference order that was used to determine the identity of $D_t(\cdot, \cdot)$. The process of the t'th auction up to price p for which the

[3] A deterministic tie-breaking rule is using the same preference order throughout the auction while a random tie-breaking rule is also allowed and is using a (possibly randomly chosen) different preference order each block. However we assume that the auctioneer makes public the order that was chosen at every step (after the dropper was chosen).

[4] Later on we impose an activity rule which will cause this set to be possibly different than the set of all non-winners with $r_i \leq t$.

price clock stopped, and block k within p is fully described by the tuple (history) [5],[6]

$$h_t(p,k) = (t, X_t, (p', (D_t(p',k'), \succ_{(p',k')})_{k'=1,\dots,s^t_{p'}})_{p'\in\Re \text{ s.t. } s^t_{p'}>0 \text{ and } p'<p}, \big(D_t(p,k'), \succ_{(p,k')}\big)_{k'=1\dots k})$$

By slightly abusing notations we write $(p',k') \in h_t(p,k)$ to denote the fact that the history $h_t(p,k)$ contains a k' dropper at price p'. The set of bidders that are active in the t'th auction, when the price is p and after k players dropped at p is defined as $I_t(p,k) = X_t \setminus \cup_{(p',k')\in h_t(p,k)} D_t(p',k')$. Auction t ends when exactly one bidder remains (even if she is just about to drop next). Thus, if we let p^*_t denote the end price of auction t, then p^*_t is the price p for which there exists an index $k \geq 1$ such that $|I_t(p,k)| = 1$. The winner at auction t is the last player to remain, and we denote this player by i^*_t; she pays p^*_t for the object.

A history of the entire game up to a point (t,p,k) fully summarizes the game and is the tuple $h(t,p,k) = ((h_1(p^*_1, s^1_{p^*_1}), \dots, h_{t-1}(p^*_{t-1}, s^{t-1}_{p^*_{t-1}}), h_t(p,k)))$. Let $\mathcal{H}$ denote the set of all (non-terminal) histories. A pure strategy for bidder i is a function $b_i(\theta_i, h) : \Theta_i \times \mathcal{H} \to \{D, R\}$ that determines for any history h whether bidder i drops or remains. Note that even if the bidder's strategy tells her to drop at a certain point she might not be chosen to be the dropper at that point.

2.3 Early drops

At this point it may be useful to illustrate the fact that bidders may indeed find it beneficial to drop, in a certain auction, well before the price reaches their value. Consider an instance with two items ($K = 2$) and three bidders that arrive at time 1 (i.e. $r_1 = r_2 = r_3 = 1$). One bidder has a low value, $v_1 = L$, and the other two bidders have much higher values, say both have a value of $v_2 = v_3 = H >> L$. Assume that bidder 1 plays the strategy "in both auctions remain until your value" i.e. for every p,t,k, $b_1(\theta_1, h(t,p,k)) = D$ if and only if $p \geq v_1$. If the two high bidders [7] will remain in the first auction until their value, the utility of the bidder who wins the first item will be zero: they will both want to drop at price H, and one of them will be determined the winner according to the tie-breaking rule. She will pay H. If that one will drop out earlier, while the other one continues to play the same strategy, she will lose the first auction, but will win the second auction for a price of L, if in the second auction she remains until her value (bidder 1 will lose the first auction and will participate in the second). In fact, under the simplifying assumption that the high bidders have almost complete information about the above situation, and the only unknown is the value L, the symmetric equilibrium strategy (there are other nonsymmetric equilibria) of each

[5] If the tie-breaking rule is deterministic, then it is part of the description of the mechanism and can be omitted from the detailed description of the histories.

[6] If p is a price for which no player announced her will to drop out, then the history at p is given by

$$h_t(p,0) = (t, X_t, (p', (D_t(p',k'), \succeq_{(p',k')})_{k'=1,\dots,s^t_{p'}})_{p'\in\Re \text{ s.t. } s^t_{p'}>0 \text{ and } p'<p}, p)$$

[7] Throughout, we use "high bidders" instead of high-value bidders to shorten notation.

of the high bidders would be "drop out in the first auction when exactly one other bidder remains, then in the second auction, remain until your value". Playing this will enable both high bidders to win, each for a price of L, as explained above.

On the other hand, since we assume incomplete information, a bidder that drops early in the first auction faces the risk of losing in the second auction because additional high-value bidders may join in. For example, suppose there are four bidders in the above example, and that the fourth bidder arrives at time 2 ($r_4 = 2$). Suppose that there is some uncertainty about v_4: with some probability p, $v_4 = 2 \cdot H$, and with probability $1 - p$ the value is $v_4 = L$. Then it is no longer clear when bidders 2 and 3 should choose to drop in the first auction, since if bidder 4 indeed has a high value, bidders 2 and 3 will not win the second auction. This point can be made even sharper since we do not assume common priors. If bidder 2's belief assumes a very small p, then it seems likely that she will drop out early in the first auction. If bidder 3 assumes a very large p, then it seems likely that she will remain almost until her value in the first auction. One can complicate things much further by increasing the level of uncertainty.

This example shows that a bidder can rationalize any dropping point, in the first auction, lower than her value and higher than the point where exactly one other bidder remains active. We intentionally emphasize these two strategies since later on we show that, if bidders indeed restrict themselves to such strategies, then the loss of efficiency (as a result of the incomplete information) can be bounded. However, clearly, some strategic choices of the bidders *will* cause an unbounded efficiency loss. In particular, this will happen if bidders drop *very early* in the auction. In the extreme, if we have K items and all bidders drop at price 0 in all but the last auction, clearly the obtained welfare cannot be satisfactory. Unfortunately (and quite surprisingly) such extreme strategies can be weakly undominated in the most basic format of sequential auctions as described above. To rule them out, we introduce a certain "activity rule", as described next.

2.4 Adding an activity rule

The above discussion shows that there are at least two undominated strategies in the sequential auction: at each auction t, either: (1) drop when price reaches value, or (2) drop when the number of remaining bidders is equal to the number of remaining auctions, or when price reaches value, the earliest of the two. Unfortunately, there may be other undominated strategies, that result in an even earlier drop, even when the price is 0, and many players are still present (see Appendix A for a concrete example). These strategies will clearly damage the obtained social efficiency, as if many players drop at a price of 0, there is no way to guarantee that the high players win. A simple activity rule can fix this problem of the original formulation. We want to force the bidders to compete at least until the point where the number of remaining players is equal to the number of remaining items. We achieve this by basically saying that when the t'th auction ends, only the $K - t + 1$ highest bidders are qualified to participate in the next auction (note that after auction t

there remain $K - t$ additional auctions):

A sequential auction with an activity rule. Let (p_t^{ar}, k_t^{ar}) be the first point in auction t where there remain at most $K-t+1$ active bidders. In other words, $|I_t(p_t^{ar}, k_t^{ar})| \leq K-t+1$, and for any other point (p, k), $|I_t(p, k)| > K - t + 1$ if and only if $p < p_t^{ar}$ or $p = p_t^{ar}$ and $k < k_t^{ar}$. We modify the starting condition for the next auction (auction $t + 1$), as follows:

- The price clock at auction $t + 1$ starts from price p_t^{ar}.
- Bidders not in $I_t(p_t^{ar}, k_t^{ar})$ are disqualified from participating in auction $t + 1$. Equivalently,

$$X_{t+1} = I_t(p_t^{ar}, k_t^{ar}) \setminus \{i_t^*\} \cup \{j \mid r_j = t + 1\}.$$

Note that if $x_t < K - t + 1$, then the activity rule does not have any effect, i.e. $X_{t+1} = X_t \setminus \{i_t^*\} \cup \{j \mid r_j = t + 1\}$. However, if $x_t > K - t + 1$, then some players will be disqualified from continuing to the next auction, and the cutoff point (p_t^{ar}, k_t^{ar}) will have $|I_t(p_t^{ar}, k_t^{ar})| = K - t + 1$.

We denote by $Q_t = I_t(p_t^{ar}, k_t^{ar})$ the set of players at auction t that are qualified to continue to auction $t + 1$. This set contains also the winner of auction t that will eventually not participate in auction $t + 1$.

We will show that this activity rule enables a good characterization of the set of undominated strategies: in every auction $t = r_i, ..., K$, bidder i does not drop before there remain at most $K-t+1$ active bidders, unless the price reaches her value before that. As a consequence, we get that the bidders with the $K - t + 1$ highest values *among all the bidders that arrived up to time t and have not won yet* are qualified to continue on to auction $t + 1$ (besides of course the winner, that also belongs to this set). In the next sections we show that this property is sufficient to bound the efficiency loss of the auction.

The purpose of setting the start price of auction $t + 1$ to p_t^{ar} is to preserve the results of the past competition in future rounds. By disqualifying all but the $K - t + 1$ highest bidders, we reduce future competition in the following auctions. The price p_t^{ar} is the exact price at which the disqualified bidders stopped competing, while the qualified bidders were still willing to compete. Thus, the start price of auction $t + 1$ reflects the competition at auction t, and we do not lose by disqualifying some bidders. Moreover, as our analysis shows, the activity rule actually increases the competition at time t, since players that do not compete at this auction will not be allowed to enter the next auction. This enhanced competition will prevent the social welfare from deteriorating towards an infinitesimal fraction of the optimal one, *no matter what undominated strategies the players choose to play.*

Given this, one may be tempted to claim that a more demanding activity rule will yield an even better result. For example, one can disqualify **all** participants of auction t from participating in auction $t + 1$. This will clearly yield a higher price in auction t. However, by this, we may end up with many unsold items. For example, if all bidders arrive at time 1, this alternative rule will result in $K-1$ items being unsold, while with our activity rule all items will always be sold[8]. Therefore, if we care about social efficiency, our cutoff point seems to be the correct one.[9]

Another justification for this exact cutoff point comes from looking at the case where all bidders are known to arrive at time 1. Suppose that all bidders play the strategy that in each auction t, the bidder drops when there remain exactly $K-t+1$ bidders, or when the price reaches the bidder's value (the earlier event of the two). Notice that when all bidders arrive at time 1, the end result of this strategy will be equivalent to Vickrey's result: the K highest bidders win, each pays the $(K+1)$'th highest value. Thus, if it is common knowledge that all bidders arrive at time 1, then the strategy of dropping out when exactly $K-t+1$ bidders remain or when the price reaches the value (the earlier) is an ex-post equilibrium. Putting this differently, our cutoff point is the highest cutoff that does not rule out the possibility of Vickrey prices. As explained earlier, when bidders' arrivals are scattered throughout the different auctions, and their arrival times are unknown to the other bidders, it may be rational for a bidder to remain in auction 1 even when less than $K-1$ other bidders remain active. Our analysis quantifies the possible efficiency loss in this case.

One immediate property of the activity rule, that will become important later, is:

Proposition 1 *If $|X_t| < K-t+1$ then no player was disqualified at any auction $s \leq t$.*

Proof. We will prove that, if a player is disqualified at auction s then, for every $t \geq s$, $|X_t| \geq K-t+1$. We prove this by induction on t. For $t = s$, since some player was disqualified, then by definition $|X_s| \geq K-s+1$. Assume the claim is true for t, and let us verify it for $t+1$. Since $|X_t| \geq K-t+1$ then by definition $|Q_t| = K-t+1$; hence $|X_{t+1}| \geq |Q_t|-1 = K-t = K-(t+1)+1$, and the claim follows. ■

2.5 Analysis of undominated strategies

We start by showing that the sequential auction with our activity rule reduces the possible strategic choices of the bidders, so that the weakly undominated strategies are of the following form: in auction $t = r_i, ..., K$, bidder i does not drop before there remain at most $K-t+1$ active bidders, unless the price reaches her value before that.

Proposition 2 *Suppose bidder i plays some weakly undominated strategy $b_i(\cdot)$. Then, in any auction $t = r_i, ..., K$, bidder i does not drop before there remain at most $K-t+1$ active bidders, unless the price reaches her value before that. More formally, if $b_i(\cdot)$ is undominated, then it satisfies:*

$$b_i((r_i, v_i), h(t,p,k)) = D \text{ implies } |I_t(p,k)| \leq K-t+1 \text{ or } p \geq v_i.$$

[8] Unless no player shows up for some auction and all players that arrived previously are already winners.

[9] An interesting question, that we leave for a later study, is whether a higher cutoff point can increase the seller's revenue.

Proof. Suppose by contradiction that for some history h', at some point (t', p', k'), we have $b_i((r_i, v_i), h'(t', p', k')) = D$, but $I_{t'}(p', k') > K - t' + 1$ and $p' < v_i$. We will argue that, if player i indeed dropped as a result of this strategy, then clearly she could have stayed until her value, and if she did not drop out (because she was not chosen by the tie-breaking rule), then her action was ignored; declaring R will make no difference.

Formally, let $D_t^h(p, k)$ denote the player who dropped at the point (t, p, k) in history h, and let $i \succ_{(p,k)} D_t^h(p, k)$ denote the fact that the tie-breaking rule at (p, k) prefers i to $D_t^h(p, k)$. We will argue that the following strategy $\bar{b}_i(\cdot)$ weakly dominates $b_i(\cdot)$:

$$\bar{b}_i((r_i, v_i), h(t, p, k)) = \begin{cases} R & h'(t', p', k') = h(t, p, k) \\ R & h'(t', p', k') \text{ is a prefix of } h(t, p, k) \\ & \text{and } i \succ_{(p,k)} D_t^h(p, k) \text{ and } p < v_i \\ D & h'(t', p', k') \text{ is a prefix of } h(t, p, k) \\ & \text{and } i \succ_{(p,k)} D_t^h(p, k) \text{ and } p \geq v_i \\ b_i((r_i, v_i), h(t, p, k)) & \text{otherwise} \end{cases}$$

In words, on histories that do not start with $h'(t', p', k')$, $\bar{b}$ is identical to b. For the history $h'(t', p', k')$, $\bar{b}$ announces "remain" instead of "drop", and from then on it follows one of the two options: (1) if player i is preferred over the player that actually dropped at t', p', k' (i.e. player i would have dropped if she would have announced "drop") then $\bar{b}$ remains if and only if the price is smaller than the value, in auction t and in all following auctions, or (2) if the player that actually dropped at t', p', k' is preferred over player i (i.e. player i would not have dropped even if she would have announced "drop"), then $\bar{b}$ is identical to b. (In addition, $\bar{b}$ is identical to b for all types different than (r_i, v_i)).

Let us verify that $\bar{b}_i(\cdot)$ weakly dominates $b_i(\cdot)$. For any history such that $h'(t', p', k')$ is not a prefix of $h(t, p, k)$, the strategies are identical. For a history with a prefix $h'(t', p', k')$, we have two cases: (1) if $D_{t'}^h(p', k')$ is preferred over i, then $\bar{b}$ will yield an identical result to b's result, as the only difference is at the point (t', p', k'), and i's announcement at this point is completely ignored by the auction and by the other players; (2) if i is preferred over $D_{t'}^h(p', k')$, then in $b_i(\cdot)$ player i will be disqualified, and will thus obtain a zero utility, while $\bar{b}_i(\cdot)$ yields a nonnegative utility, since it never remains above v_i. A strictly positive utility is obtained, e.g. in the situation where v_i is the maximal value among all remaining players, and they all remain exactly until their price reaches their value. ■

Proposition 3 *Any undominated strategy $b_i(\cdot)$ satisfies:*

$$|I_t(p, k)| > K - t + 1 \textit{ and } p \geq v_i \textit{ implies } b_i((r_i, v_i), h(t, p, k)) = D$$

for any history h.

Proof. Suppose by contradiction that the above condition is violated, and consider the following strategy:

$$\bar{b}_i((r_i, v_i), h(t,p,k)) = \begin{cases} D & |I_t(p,k)| > K - t + 1 \text{ and } p \geq v_i \\ b_i((r_i, v_i), h(t,p,k)) & \text{otherwise} \end{cases}$$

We argue that $\bar{b}_i$ weakly dominates b: on any history such that the condition of the claim is not violated, the two strategies are identical. On any history in which the condition of the claim is violated, say at auction t, note that by the definition of the activity rule, the price will be larger than v_i throughout all remaining auctions. Thus the best possible utility for player i is zero, and the strategy $\bar{b}$ will indeed yield a zero utility, since it will announce "drop" when the price at auction t will reach v_i.[10] ∎

These two propositions imply that, as long as there exist more than $K-t+1$ remaining players, a player will drop if and only if the price reaches her value (assuming players play some tuple of undominated strategies). As a result, the winner's value is one of the $(K-t+1)$-highest values of the bidders in X_t, and, moreover, the other $(K-t)$-highest values belong to the bidders that are qualified to participate in the next auction. We next wish to claim something slightly stronger – that Q_t contains the $(K-t+1)$-highest values among all bidders that have arrived up to time t and have not won yet. This is (possibly) a larger set than X_t, since it also contains all the disqualified bidders from previous auctions. Formally, for the following, we assume any fixed profile of types and fixed tuple of undominated strategies. Denote by Λ_t the set of all bidders that arrived at or before auction t and did not win any of the items $1,..,t-1$, i.e. $\Lambda_t = \{j \in I \mid r_j \leq t \text{ and } j \neq i^*_k \text{ for } k = 1,...,t-1\ \}$.

Proposition 4 *Suppose all bidders play an undominated strategy. Then, for any two players $i \in \Lambda_t \setminus Q_t$ and $j \in Q_t$, we have $v_j \geq v_i$.*

Proof. We first show that for $i \in X_t \setminus Q_t$ and $j \in Q_t$ the claim holds: by the above propositions, in any undominated strategy, while more than $K-t+1$ players remain, a player drops if and only if her value is equal to the price. Thus all players not in Q_t have values smaller or equal to p^{ar}_t, and all players in Q_t have values greater or equal to p^{ar}_t, hence $v_j \geq v_i$. Note that this also implies that the $K-t+1$ highest-value bidders in X_t belong to Q_t.

We now prove the claim by induction on t. For $t=1$, $\Lambda_t = X_t$ and the claim follows from the above argument. Assume the claim is correct for any $t' < t$, and let us prove it for t. If $i \in X_t$ then again the above argument holds. Otherwise $i \in \Lambda_t \setminus X_t$, which implies that i arrived strictly before time t and was disqualified at or before time $t-1$. Note that by proposition 1 we have that

[10] With a random tie-breaking rule, b may sometimes yield a negative utility, as it can be that all players will drop before i, forcing her to pay more than her value. A deterministic rule is less effective in this respect.

$|Q_{t-1}| = K - (t-1) + 1$. Since player i was disqualified, we have $i \in \Lambda_{t-1} \setminus Q_{t-1}$. Let j' be the player with minimal value in Q_{t-1}. By the induction assumption we have that $v_{j'} \geq v_i$. There are $K - t + 2$ players in Q_{t-1}, out of them $K - t + 1$ continue to auction t (i.e. belong to X_t), all with values larger or equal to $v_{j'}$. Thus, again by the argument in the first paragraph, any player $j \in Q_t$ has $v_j \geq v_{j'} \geq v_i$, and the claim follows. ■

Corollary 5 *If all bidders play a weakly undominated strategy then at each auction t the $K-t+1$ highest-value bidders in Λ_t have the same set of values as the bidders in Q_t.*

3 Worst-Case Analysis

We now turn to analyze the social efficiency of the sequential auction mechanism, under the assumption that players may choose to play any tuple of undominated strategies. To rephrase this more formally, consider the following definitions. Fix any realization of players' types $\theta = (\theta_1, ..., \theta_n)$, where $\theta_i = (r_i, v_i)$. A subset of players W is a valid set of winners, for a given θ, if there exists an assignment of items to all players in W such that no item is assigned to more than one player, and if item t is assigned to player i then $t \geq r_i$. Let $\mathcal{W}(\theta)$ denote the set of all valid sets of winners. The value of some $W \in \mathcal{W}(\theta)$ is $v(W, \theta) = \sum_{i \in W} v_i$. We say that $OPT \in \mathcal{W}(\theta)$ is socially efficient if $v(OPT, \theta) = \max_{W \in \mathcal{W}(\theta)} v(W, \theta)$. Note that a socially efficient assignment is independent of the profile of strategies of the bidders.

For example, suppose three players and two items, and $r_1 = r_2 = 1, r_3 = 2$, $v_1 = \epsilon, v_2 = 1, v_3 = 1 + \epsilon$ where $0 < \epsilon < 1$. In other words, two players arrive at time 1, with values ϵ and 1. At time 2 another player arrives with value $1 + \epsilon$. Consider the following strategies, in the first auction: player 1 continues up to her value, while player 2 drops immediately. Recall that for each of these strategies the property in proposition 2 holds.[11] By the auction rules, if these strategies are played, then player 1 wins the first item and pays zero, while player 2 continues on to the next auction. In the next auction, both players 2 and 3 remain until their value (assuming they play an undominated strategy). Thus player 3 wins and pays 1. The winners of the auction are thus $A = \{1, 3\}$, and $v(A) = 1 + 2 \cdot \epsilon$. The efficient set of winners is $OPT = \{2, 3\}$, and $v(OPT) = 2 + \epsilon$. Consequently, when ϵ approaches zero, the auction results in a loss of half of the optimal social efficiency or, equivalently, the auction obtains half of the optimal social efficiency.

Is there any other combination of types and undominated strategies that can lead to a lower ratio (for two items)? In particular, can we decrease this ratio to zero, by appropriately setting types and strategies? It turns out that this is not possible. To see this, we use the following fact: for any realization of players' types, and for any tuple of undominated strategies, the player with the highest value wins the auction. If this player arrived for the first auction but was not a winner,

[11] As explained above, player 2 may rationalize such a strategy by the belief that player 1 has a high value, while player 3 has a low value.

then by corollary 5 she was qualified for the second auction. In the second auction, since all players remain up to their values, the highest player wins. In other words, the highest player always wins one of the items in the auction. Now, denote her value as x. Let A be the set of winners in the two auctions, and let OPT be a set of winners with maximal social efficiency. By the above argument, for any profile of undominated strategies, $v(A) \geq x$. Since we have two items, $|OPT| \leq 2$, and since for all i, $v_i \leq x$, we get $v(OPT) \leq 2 \cdot x$. Therefore $v(A)/v(OPT) \geq 1/2$, as claimed.

What happens if we consider more than two items? Can we then find a combination of types and undominated strategies that leads to a lower ratio? The answer is again no, and the analysis is slightly more subtle. Recall that the key observation in the two items case is that the highest player of OPT must win, if all players play undominated strategies. For three items, this is still true. Unfortunately, the *second highest* player of OPT need not necessarily win. Nevertheless, we show that if she does not win, then the third highest must win instead. This structure can be generalized to any number of items, as follows.

Fix any tuple of types θ. For simplicity of notation we omit repeating θ throughout. Let OPT be a valid assignment with maximal efficiency, and let A be an assignment that results from the sequential auction with the activity rule, when all players play some tuple of undominated strategies. Let $v_1^{OPT}, ..., v_K^{OPT}$ be the values of the winners of OPT, ordered in a non-increasing order (i.e. $v_1^{OPT} \geq v_2^{OPT} \geq \cdots \geq v_K^{OPT}$). (we also set $v_{K+1}^{OPT} = 0$ for notational purposes). Similarly, let $v_1^A, ..., v_K^A$ be the values of the winners of A, again in a non-increasing order.

Lemma 6 *Fix any index* $0 \leq l \leq \lfloor \frac{K}{2} \rfloor$. *Then* $v_{l+1}^A \geq v_{2 \cdot l+1}^{OPT}$.

Proof. Assume by contradiction that there are at most l winners in A with values that are larger or equal to $v_{2 \cdot l+1}^{OPT}$. Let $K - t$ be the last auction at which the winner in A has value strictly smaller than $v_{2 \cdot l+1}^{OPT}$. After this auction there remain exactly t more auctions, hence there are at least t players in A with value at least $v_{2 \cdot l+1}^{OPT}$. Thus, by the contradiction assumption, we have that $t \leq l$.

Let X be the set of players in OPT with the $2 \cdot l + 1$ highest values. Let $Y = \{i \in X \mid r_i \leq K - t\}$. Note that $|Y| \geq (2 \cdot l + 1) - t$: there are only t auctions after time $K - t$, so there are at least $(2 \cdot l + 1) - t$ players in X that receive an item in OPT at or before time $K - t$, and these must have an arrival time smaller or equal to $K - t$. Let Z be the set of players in Y that win in A before auction $K - t$. Thus $Y \setminus Z \subseteq \Lambda_{K-t}$.

Now, note that $|Z| \leq l - t$: from the definition of t, after auction $K - t$, all winners in A have values at least $v_{2 \cdot l+1}^{OPT}$, all players in Z are winners in A and they also have values at least $v_{2 \cdot l+1}^{OPT}$; by the contradiction assumption there are at most l such winners in A. Thus $|Y \setminus Z| \geq (2 \cdot l + 1 - t) - (l - t) = l + 1 \geq t + 1$.

Since $Y \setminus Z \subseteq \Lambda_{K-t}$ and $|Y \setminus Z| \geq t + 1$, then the $(t + 1)$-highest-value in Λ_{K-t} is larger or equal than the minimal value in $Y \setminus Z$. By corollary 5, the winner in A at auction $K - t$ (which belongs to Q_{K-t}) must have value at least as large as the $(t + 1)$-highest-value in Λ_{K-t} (note that

$K-(K-t)+1=t+1$). Thus the winner in A at auction $K-t$ has value at least as large as the minimal value in $Y \setminus Z$. But all players in $Y \setminus Z$ have values at least $v_{2\cdot l+1}^{OPT}$, hence this is a contradiction to our assumption that the winner in A at auction $K-t$ has value strictly smaller than $v_{2\cdot l+1}^{OPT}$. ■

This lemma immediately implies that $v(OPT)$ is at most twice $v(A)$: it shows that $v_1^{OPT}, v_2^{OPT} \leq v_1^A$, and that $v_3^{OPT}, v_4^{OPT} \leq v_2^A$, and so on and so forth, and thus $v(OPT) \leq 2 \cdot \sum_{k=1}^{\lfloor \frac{K}{2} \rfloor + 1} v_k^A \leq 2 \cdot v(A)$. We get:

Theorem 7 *Fix any tuple of types θ, and any tuple of undominated strategies. Let A be the resulting set of winners in the sequential auction, and let $v(OPT, \theta)$ be the optimal social efficiency with respect to θ. Then it must be the case that $v(A, \theta) \geq \frac{1}{2} v(OPT, \theta)$.*

We point out that the factor one-half lower bounds the *actual* ratio between the resulting efficiency of the auction, and the optimal efficiency for any realization of types. While, as the example in the beginning of the section shows, our analysis is tight in the sense that it is not possible to replace the factor one-half with a larger factor, it is not clear if such a low ratio will indeed be achieved for *most* realizations of types. Indeed, the example involved a specific tuple of types, and the adversarial choice of the parameter ϵ and the strategies of the bidders, greatly influenced the ratio of the resulting efficiency of the auction, to the optimal efficiency. It seems that, as the *worst-case* ratio is one-half, the average-case ratio should be much larger. To examine this conjecture, in the next section we engage in a distributional analysis that verifies this rough intuition, at least for the case of two items.

4 Average-Case analysis for $K = 2$

The analysis of the previous section is worst-case in the sense that even if we have an adversary that chooses the number of players, their arrival times, their values, and their strategies (restricted to the set of undominated strategies), the bounds on the efficiency loss still hold. Clearly, this is a very pessimistic viewpoint, and it would be more reasonable to assume that some of these variables are determined according to some underlying probability distribution. In this section we will concentrate on the special case where there are only two items for sale, and demonstrate that even a minor shift from the worst-case setting towards the average-case setting will improve the efficiency guarantee quite significantly.

Formally, we assume an adversary that is allowed to choose the number of players, n, and their arrival times. Thus, the adversary determines a number $r \leq n$, such that the first r players arrive for the first auction, and the remaining $n-r$ players arrive for the second auction. The adversary then chooses a cumulative probability distribution F with some support in $[0, \infty]$, and draws the values of the players from this distribution, i.e. the values are i.i.d. The adversary then determines the strategy of each player (as before, the choice of the strategy may depend on the random result of the players' values, as to "fail" the auction). Comparing this setup to the setup of the previous section, we can see that the only change is that now the adversary must draw the players' values from some fixed distribution (but the adversary can choose what distribution to use). We will show that this small modification towards an average-case setup implies a significant increase in the efficiency of the sequential auction: the auction will obtain at least 70% of the optimal efficiency, no matter the number of players, their arrival times, the chosen distribution of the players' values, and any choice of undominated strategies. Moreover, we will show that this bound is tight, i.e. that there exists a sequence of distributions that approach this efficiency guarantee in the limit.

The analysis is carried out in the following way. Fix any number of players, n. If $n = 1, 2$ then the auction must choose the optimal outcome, and so we assume that $n \geq 3$. Fix any number $r \leq n$ of players that arrive for the first auction. Again, if $r = 1$ then the auction must choose the optimal outcome, and so we assume that $r \geq 2$. Now fix any cumulative distribution F. Given these, we define two random variables: $OPT_{r,n}$ is equal to the highest value among all players that arrive at time 1 plus the highest value among all the remaining players (including those that arrive at time 2). Note that $OPT_{r,n}$ is indeed equal to the optimal efficiency, given a specific realization of the values. The second random variable, $\tilde{A}_{r,n}$, is equal to the second highest value among all players that arrive at time 1 plus the highest value among all the remaining players (including those that arrive at time 2). By corollary 5 the winner in the first auction has a value larger or equal to the second highest value among all players present in the first auction, and the winner in the second auction has the largest value among all remaining players (assuming all players play some tuple of undominated strategies). Thus, A's value (the sum of the values of the winners) is always larger or equal than the value of $\tilde{A}$. Using these settings, the main result of this section is that the expected efficiency of the sequential auction is at least 70% of the expected optimal efficiency:

Theorem 8 *For any choice of the parameters* n, r, F, $\frac{E_F[\tilde{A}_{n,r}]}{E_F[OPT_{n,r}]} \geq 0.707$.

We prove this in two parts. We first concentrate on the case of a Bernoulli distribution over the values $\{0, 1\}$, and bound the ratio of expectations over all such possible distributions. The second step is to show that the case of a Bernoulli distribution is, in some sense, the worst possible case. We show, in a formal way, how to use the obtained bound for the Bernoulli distribution to bound any other distribution.

4.1 A bound on any Bernoulli distribution

We have n players with i.i.d. values drawn from a Bernoulli distribution such that $\Pr(v_i = 0) = p$ and $\Pr(v_i = 1) = 1 - p$ for some $0 \leq p < 1$. Players $1, ..., r$ arrive for the first auction (at time 1), and players $r + 1, ..., n$ arrive for the second auction, at time 2, where p, n, r are parameters.[12]

[12] Since players are a priori symmetric it does not matter which players arrive at time 1 and which arrive at time 2; the only important parameter is the number of arrivals for each auction.

We ask what p, n, r will minimize the ratio $\frac{E_{F_p}[\tilde{A}_{n,r}]}{E_{F_p}[OPT_{n,r}]}$, where F_p denotes the above-mentioned Bernoulli distribution.

Observe that, since a player's value is either zero or one, the random variables $OPT_{n,r}$ and $\tilde{A}_{n,r}$ can take only the values $0, 1, 2$. We calculate:

$$\Pr(\tilde{A}_{n,r} = 0) = p^n,$$
$$\Pr(\tilde{A}_{n,r} = 1) = p^r(1 - p^{n-r}) + r(1-p)p^{r-1},$$
$$\Pr(\tilde{A}_{n,r} = 2) = 1 - p^r - r(1-p)p^{r-1}.$$

For example, $\tilde{A} = 1$ if all values at auction 1 are 0 and at least one value at auction 2 is 1 (this happens with probability $p^r(1 - p^{n-r})$), or if there exists exactly one value that is equal to 1 at auction 1, and then it does not matter what the values are at the second auction (this happens with probability $r(1-p)p^{r-1}$). Similarly, we also have:

$$\Pr(OPT_{n,r} = 0) = p^n,$$
$$\Pr(OPT_{n,r} = 1) = p^r(1 - p^{n-r}) + r(1-p)p^{n-1},$$
$$\Pr(OPT_{n,r} = 2) = 1 - p^r - r(1-p)p^{n-1}.$$

Using this, a lengthy calculation, detailed in appendix B, gives us:

Proposition 9 *For any n, r, and $0 \le p < 1$,*

$$\frac{E_{F_p}[\tilde{A}_{r,n}]}{E_{F_p}[OPT_{r,n}]} = \frac{2 - p^r - p^n - r(1-p)p^{r-1}}{2 - p^r - p^n - r(1-p)p^{n-1}} \ge 0.70711.$$

The calculations first show that this ratio decreases with n (for any r, p), so it suffices to compute a lower bound on the limit of the ratio of expectations when $n \to \infty$. In that case, a minimum is achieved for $r = 2$ and $p = 2 - \sqrt{2}$. Note that for $p = 1$, the two expectations become zero and the ratio is undefined.

While the worst-case scenario of section 3 requires only three players, here, to approach the minimal ratio of 0.7 we need the number of players to approach infinity, and these additional players should arrive only for the second auction. This is quite puzzling, at first, since we know that in the second auction the player with the highest value wins the sequential auction. So what is the effect of adding more players to the second auction? Looking at the probability distributions of $\tilde{A}_{n,r}$ and of $OPT_{n,r}$, one can see that as n increases, the probability of OPT and $\tilde{A}$ to be equal to 0 or 1 decreases, and the probability to be equal to 2 increases. However, the probability to have a value of 1 decreases faster for OPT. The events that explain this are those in which, at the first auction, exactly one player gets value 1 and the other players get value 0, and at the second auction there exists at least one additional player with value 1. This is a "good" scenario for OPT and a “bad” scenario for $\tilde{A}$. In fact, these are the *only* events that differentiate OPT from $\tilde{A}$. As the number of players, n increases (while keeping r constant), these events get more probability, hence the above-mentioned effect. This is not the only difference between the worst-case and the average-case settings, e.g. $r = 2$ is not necessarily the choice that minimizes the expectation ratio, given n and p. For some distributions, a larger r may actually decrease the ratio between the two expectations.[13]

4.2 Generalizing to any other distribution

To explore the case of a general distribution F with a support in $[0, \infty)$, we must take a closer look at the expression for the expectation of OPT and $\tilde{A}$. We denote by $X_{n-j:n}$ the j'th order statistic of the random variables $v_1, ..., v_n$ (the players' values), which denotes the $(j+1)$'th highest value of the players, i.e., $X_{n:n}$ is a random variable that takes the maximal value among $v_1, ..., v_n$; $X_{n-1:n}$ is a random variable that takes the second largest value among $v_1, ..., v_n$, and so on. If the highest value at time 1 is the $j+1$'th highest value among all the n players, then $OPT_{n,r} = X_{n-j:n} + X_{n:n}$. Hence

$$E[OPT_{n,r}| \text{ highest at time 1 is (j+1)-highest overall }] = E[X_{n-j:n} + X_{n:n}].$$

Denote by $q_j^{n,r}$ the probability that the highest value at time 1 is the $j+1$'th highest value among all the n players. It follows that:

$$E_F[OPT_{n,r}] = q_0^{n,r}(E_F[X_{n-1:n}] + E_F[X_{n:n}]) + \sum_{j=1}^{n} q_j^{n,r}(E_F[X_{n-j:n}] + E_F[X_{n:n}]). \tag{1}$$

We remark that the highest player among the players that arrive at time 1 is at least the $n-r+1$ highest player among all players; therefore $q_j^{n,r} = 0$ when $j > n - r$. It will be important for the sequel to verify that the probability $q_j^{n,r}$ does not depend on the distribution F. First, note that since the values are drawn i.i.d. then each value-ordering of the players has equal probability. Thus, the probability of any specific order of all the players is $1/n!$, and the probability that the order of values will satisfy any specific property is simply the number of orderings that satisfy this property, divided by $n!$. To find $q_j^{n,r}$, we thus ask in how many orderings, the highest player among the first r players is exactly the $j + 1$ highest among all players. To get one such ordering, one needs to choose one player (say i) out of the r players of time 1 (this is the highest player at time 1), to choose j players out of the $n - r$ players of time 2 (these are the players that are higher than i), to order them in one of the $j!$ orderings, then to place i, and then to order the remaining $n - j - 1$ players. Thus, for any $0 \leq j \leq n - r$,

$$q_j^{n,r} = \frac{1}{n!} \cdot r \cdot \binom{n-r}{j} \cdot j! \cdot (n-j-1)!$$

[13]E.g. for a Bernoulli distribution with $n = 7$ and $p = 0.9$, $r = 2$ yields a lower ratio than $r = 3$.

(and we set $q_j^{n,r} = 0$ for any $n - r + 1 \leq j \leq n$), which does not depend on F.

Similarly, given that the second-highest player at time 1 is the $j+1$'th highest player among all the n players, the expected welfare of $\tilde{A}$ is $E[X_{n-j:n} + X_{n:n}]$. Denoting by $p_j^{n,r}$ the probability that the second-highest player at time 1 is the $j+1$'th highest player among all the n players (where again this probability does not depend on F), it follows that:

$$E_F[\tilde{A}_{n,r}] = \sum_{j=1}^{n} p_j^{n,r}(E_F[X_{n-j:n}] + E_F[X_{n:n}]) \tag{2}$$

(and we set $q_j^{n,r} = 0$ for any $n - r + 2 \leq j \leq n$).

We now consider the terms $E[X_{n-j:n}]$. Let $F_{n-j:n}(x)$ be the probability distribution of $X_{n-j:n}$. The probability that $X_{n-j:n} \leq x$ is the probability that at most j values will be higher than x, and the remaining $n - j$ values will be smaller than x, or, in other words,

$$F_{n-j:n}(x) = \Pr(X_{n-j:n} \leq x) = \sum_{k=0}^{j} \binom{n}{k} (1 - F(x))^k (F(x))^{n-k}.$$

Therefore, $F_{n-j:n}(x)$ is a polynomial in $F(x)$, where the constants of the polynomial do not depend on the distribution F. A well-known formula for the expectation of an arbitrary nonnegative random variable Y with cumulative distribution G is $E[Y] = \int_0^\infty (1 - G(y))dy$. In particular, $E[X_{n-j:n}] = \int_0^\infty (1 - F_{n-j:n}(x))dx$. In other words, the expectation of the j'th order statistic is an integration over a polynomial in $F(x)$, i.e. there exist constants $w_l^{(j)}$ for $l = 0, ..., n$ and $j = 1, ..., n$ (that does not depend on the distribution F) such that

$$E_F[X_{n-j:n}] = \int_0^1 [\sum_{l=0}^{n} w_l^{(j)} (F(x))^l] dx.$$

Combining this equation with equations (2) and (1), we get that both $E_F[OPT_{n,r}]$ and $E_F[\tilde{A}_{n,r}]$ are an integration over a polynomial in $F(x)$, i.e. there exist constants $\beta_0^{(n,r)}, ..., \beta_n^{(n,r)}$ and $\gamma_0^{(n,r)}, ..., \gamma_n^{(n,r)}$, that do not depend on the distribution F, such that

$$E_F[OPT_{n,r}] = \int_0^\infty [\sum_{l=0}^{n} \beta_l^{(n,r)} (F(x))^l] dx$$

and

$$E_F[\tilde{A}_{n,r}] = \int_0^\infty [\sum_{l=0}^{n} \gamma_l^{(n,r)} (F(x))^l] dx.$$

One additional important observation is that $\sum_{l=0}^{n} \beta_l^{(n,r)} = \sum_{l=0}^{n} \gamma_l^{(n,r)} = 0$. To see this, take some distribution F with a bounded support, say $[0, 1]$. The above equality implies that

$E_F[OPT_{n,r}] > \int_1^\infty [\sum_{l=0}^n \beta_l^{(n,r)}]dx$, which is unbounded if $\sum_{l=0}^n \beta_l^{(n,r)} \neq 0$. But clearly $E_F[OPT_{n,r}]$ is a finite number since the support is bounded; hence it must be that $\sum_{l=0}^n \beta_l^{(n,r)} = 0$. The same argument implies that $\sum_{l=0}^n \gamma_l^{(n,r)} = 0$.

The Bernoulli distribution F_p $(0 \leq p < 1)$ gives a fixed function over the interval $[0,1)$, specifically $F_p(x) = p$ for any $0 \leq x < 1$, and $F_p(x) = 1$ for $x \geq 1$. Thus for this distribution the integration cancels out, and we get:

$$E_{F_p}[OPT_{n,r}] = \sum_{l=0}^{n} \beta_l^{(n,r)} p^l \quad , \quad E_{F_p}[\tilde{A}_{n,r}] = \sum_{l=0}^{n} \gamma_l^{(n,r)} p^l. \tag{3}$$

As an aside, we remark that when plugging the exact terms for all these constants, most terms cancel out, and it is possible to get a simple exact formula for the two expectations:

$$E_F[OPT_{n,r}] = \int_0^\infty [2 - r(F(x))^{n-1} + (r-1)(F(x))^n - (F(x))^r]dx$$

and

$$E_F[\tilde{A}_{n,r}] = \int_0^\infty [2 - (F(x))^n + (r-1)(F(x))^r - r(F(x))^{r-1}]dx,$$

where one may compare this with the explicit formula for the expectations in the case of the Bernoulli distribution, detailed in the previous subsection.

We next show how all the above implies:

Proposition 10 *Fix any α such that $\frac{E_{F_p}[\tilde{A}_{n,r}]}{E_{F_p}[OPT_{n,r}]} \geq \alpha$, for any n, r and $0 \leq p < 1$. Then, for any other cumulative distribution F with $E_F[OPT_{n,r}] > 0$, it must be that $\frac{E_F[\tilde{A}_{n,r}]}{E_F[OPT_{n,r}]} \geq \alpha$.*

Proof. We need to show that $\frac{E_F[\tilde{A}_{n,r}]}{E_F[OPT_{n,r}]} \geq \alpha$, or, equivalently, that $E_F[\tilde{A}_{n,r}] - \alpha E_F[OPT_{n,r}] \geq 0$. Using the above equations, this term becomes

$$\int_0^\infty [\sum_{l=0}^{n} \gamma_l^{(n,r)} (F(x))^l - \alpha \sum_{l=0}^{n} \beta_l^{(n,r)} (F(x))^l]dx.$$

We will show that, for every $x \geq 0$, $\sum_{l=0}^n \gamma_l^{(n,r)}(F(x))^l - \alpha \sum_{l=0}^n \beta_l^{(n,r)}(F(x))^l \geq 0$, which implies the above inequality. Fix some $x \geq 0$, if $F(x) = 1$ then indeed $\sum_{l=0}^n \gamma_l^{(n,r)}(F(x))^l - \alpha \sum_{l=0}^n \beta_l^{(n,r)}(F(x))^l = 0 - \alpha \cdot 0 = 0$. Otherwise, denote $p = F(x) < 1$. Thus

$$\sum_{l=0}^{n} \gamma_l^{(n,r)} (F(x))^l - \alpha \sum_{l=0}^{n} \beta_l^{(n,r)} (F(x))^l = E_{F_p}[\tilde{A}_{n,r}] - \alpha E_{F_p}[OPT_{n,r}] \geq 0,$$

where the equality follows Eq. 3, and the inequality follows from the assumption in the claim. ■

Corollary 11 *For any cumulative probability distribution F, and any n, r, $\frac{E_F[\tilde{A}_{n,r}]}{E_F[OPT_{n,r}]} \geq 0.707$.*

This completes the proof of theorem 8.

We wish to remark that this bound is achieved only for the worst distribution, and for most distributions the bound would be higher. For example, we have obtained that, for the uniform distribution over $[0, 1]$, the ratio of expectations is minimized for $r = 2$ and $n = 16$. For these parameters the ratio is slightly higher than 80%. For a uniform distribution on any other interval $[a, b]$ the ratio will be higher.

5 Conclusions

We have analyzed a common sequential auction structure. The results bound away from zero the efficiency loss in such a setting, without making any distributional assumptions. By imposing a simple to understand and to implement activity rule, we were able to characterize the undominated strategies of the players. The activity rule states that at a given auction t (out of K total auctions), only the $K - t$ highest bidders who did not win at auction t are qualified to continue to the next auction. Regardless of the underlying distribution from which the players' arrival times and valuations are drawn, and of their beliefs about this distribution, we show that a player does not drop out before there remain at most $K - t + 1$ other active bidders in the auction, unless the price reaches her value (assuming that players play any undominated strategy). This provides a sufficient level of competition, among the players, regardless of their beliefs about future auction, and allows us to give bounds on the efficiency loss.

The bounds that we provide hold both for a "worst case" scenario and for an "average case" scenario for two items. For the "worst-case" analysis, an adversary is allowed to determine the number of players, their arrival times, their values, and their (undominated) strategies, in order to "fail" the auction. When this is the case, we show that the sequential auction mechanism achieves at least 50% of the optimal efficiency. The efficiency (social welfare) of an allocation is the sum of values of players that won an item after their arrival. For the "average-case" analysis, an adversary is again allowed to determine the number of players, their arrival times, and their (undominated) strategies. However, here the adversary is forced to determine a distribution on $[0, \infty)$ from which she will independently choose the values of the players. In this case, for $K = 2$, we show that the expected efficiency of the mechanism is at least 70% of the expected optimal efficiency (when choosing the worst such possible distribution) and, for example, at least 80% of the total efficiency when the chosen distribution is the uniform distribution on some interval.

Our goal was to analyze a "real world" mechanism using as little as possible assumptions on the players' beliefs and behavior, and to give a quantitative assessment of its efficiency loss. Rather than constructing a mechanism and finding its equilibrium strategies, we characterize the set of undominated strategies of a real, common mechanism, and obtain bounds on the efficiency loss incurred when players choose arbitrary undominated strategies. We did not perform any revenue considerations in this work, and possible future research may wish to address this point. For example, one may compare the seller's revenue in our setting versus the revenue of the optimal (revenue-maximizing) allocation scheme described by Gershkov and Moldovanu (2007).

A The need for an activity rule via an example

To exemplify the need to add an activity rule to the auction, consider a setting of two items and three bidders that arrive in time 1. In the second auction, it is rather immediate that the strategy of "remaining until price reaches value" weakly dominates all other strategies[14]. As the example in section 2.3 shows, there are at least two weakly undominated strategies:

1. In both auctions, remain until price reaches value, i.e. for $t = 1, 2$

$$b_i^{value}((r_i, v_i), h(t, p, k)) = D \text{ iff } p \geq v_i.$$

2. In the first auction, remain until exactly one other bidder remains, and in the second auction, remain until your value. Formally,

$$b_i^{EPD}((r_i, v_i), h(1, p, k)) = D \text{ iff } \{|I_1(p, k)| = 1 \text{ or } p \geq v_i\},$$

and

$$b_i^{EPD}((r_i, v_i), h(2, p, k)) = D \text{ iff } p \geq v_i$$

(where EPD stands for Earliest Possible Dropping).

We show that these two strategies do not weakly dominate all other strategies, in the sequential auction without the activity rule. Suppose three bidders $1, 2, 3$ arrive at time 1, with $v_1 > v_2 > v_3$, and let us consider the strategy b_1^0 for bidder 1, in which she drops at price 0 in the first auction, and remains until her value in the second auction[15]. None of the above two strategies dominates b_1^0, due to the following reasoning.

Consider first the strategy b_1^{value}. This strategy performs strictly worse than b_1^0 in case both 2 and 3 choose to do the same (remain until their value in both auctions, i.e. play b_2^{value}, b_3^{value}), due to the following. By playing b_1^{value}, bidder 1 will win the first auction and will pay v_2. By playing b_1^0, bidder 1 will lose the first auction and will win the second auction for a lower price of v_3.

Consider next the strategy b_1^{EPD}. This strategy performs strictly worse than b_1^0 when bidder 3 plays the strategy b_3^{value} and 2 uses the following strategy: In the first auction, if bidder 1 drops at

[14] Formally, any strategy b_i^0 is weakly dominated by the strategy $b_i^*((r_i, v_i), h_1(p, k)) = b_i^0((r_i, v_i), h_1(p, k))$, and $b_i^*((r_i, v_i), h_1(p_1^*, s_{p_1^*}^1), h_2(p, k)) = D$ if and only if $p \geq v_i$.

[15] Formally, $b_1^0((1, v_1), h(t = 1, p, k)) = D$ for all p, k, and $b_1^0((1, v_1), h_1(p_1^*, s_{p_1^*}^1), h_2(p, k)) = D$ iff $p \geq v_1$.

price 0 then bidder 2 continues until her value, and if bidder 1 does not drop at price 0 then bidder 2 drops immediately after that. In the second auction, bidder 2 remains until her value[16]. In this case, if bidder 1 follows b_1^0 and drops at 0 then bidder 2 will win the first auction, bidder 1 will win the second auction, and will pay v_3. If bidder 1 follows b_1^{EPD} then bidder 2 drops, and bidder 1 drops immediately after that (since now only bidder 3 remains besides 1). Thus, bidder 3 wins the first auction, bidder 1 again wins the second auction, but this time pays v_2 which is larger than v_3.

B Proof of Proposition 9

We need to show that, for any n, r, and $0 \leq p < 1$,

$$\frac{E_{F_p}[\tilde{A}_{r,n}]}{E_{F_p}[OPT_{r,n}]} = \frac{2 - p^r - p^n - r(1-p)p^{r-1}}{2 - p^r - p^n - r(1-p)p^{n-1}} \geq 0.70711. \qquad (4)$$

We first derive this expression by n, to show that it decreases as n increases.

$$\frac{d}{dn}\left(\frac{\left(2 - p^r - p^n - r\left(1-p\right)p^{r-1}\right)}{\left(2 - p^r - p^n - r\left(1-p\right)p^{n-1}\right)}\right) = rp^{n-2}\left(\ln p\right)\left(1-p\right)\frac{\left(2p - p^r r + p^{r+1}\left(r-2\right)\right)}{\left(-rp^{n-1} - p^r + \left(r-1\right)p^n + 2\right)^2}.$$

We concentrate on the term

$$G\left(p, r\right) = 2p - p^r r + p^{r+1}(r-2)$$

and claim that it is nonnegative for every $2 \leq r \leq n$ and $p \in [0, 1]$. For $p = 0$ we have $G(0, r) = 0$ and for $p = 1$ we have $G(1, r) = 0$. Moreover

$$\frac{d^2}{dp^2}G\left(p, r\right) = rp^{r-2}\left(\left(r+1\right)\left(r-2\right)p - r\left(r-1\right)\right)$$

and since

$$\frac{r\left(r-1\right)}{\left(r+1\right)\left(r-2\right)} > 1$$

for every $r \geq 2$ we know that $\frac{d^2}{dp^2}G\left(p, r\right) \leq 0$ and $G\left(p, r\right)$ is concave in p. We thus conclude that $G\left(p, r\right) \geq 0$ for every $2 \leq r \leq n$ and $p \in [0, 1]$ and consequently that $\frac{d}{dn}\left(\frac{\left(2-p^r-p^n-r(1-p)p^{r-1}\right)}{\left(2-p^r-p^n-r(1-p)p^{n-1}\right)}\right) \leq 0$. We take n to infinity and get that

$$\lim_{n\to\infty}\left(\frac{\left(2 - p^r - p^n - r\left(1-p\right)p^{r-1}\right)}{\left(2 - p^r - p^n - r\left(1-p\right)p^{n-1}\right)}\right) = 1 - \frac{r\left(1-p\right)p^{r-1}}{\left(2 - p^r\right)}$$

[16] Formally, $b_2((1, v_2), h(t = 1, p, k)) = R$ iff $[p = 0$ or $(0 < p < v_3$ and $1 \notin I_1(0, s_0^1))]$.

We wish to find the minimum of

$$1-\frac{r\left(1-p\right)p^{r-1}}{\left(2-p^{r}\right)}$$

which will give us a lower bound for (4), for every n, since we obtained that (4) decreases towards the limit as n increases.

Equivalently, we look for the maximum of

$$H\left(p,r\right)=\frac{r\left(1-p\right)p^{r-1}}{\left(2-p^{r}\right)}.$$

Now

$$\frac{d}{dr}H\left(p,r\right)=-\frac{p^{r-1}\left(1-p\right)\left(-2r\ln p+p^{r}-2\right)}{\left(2-p^{r}\right)^{2}}$$

$$\frac{d}{dp}H\left(p,r\right)=rp^{r-2}\frac{\left(2r\left(1-p\right)+p^{r}-2\right)}{\left(2-p^{r}\right)^{2}}.$$

Therefore if there exists a global maximum at $0<p<1$ and $2<r$ then we must have

$$-2r\ln p=2-p^{r}$$

and

$$2r\left(1-p\right)=2-p^{r}$$

but this is not possible since for every $0<p<1$ we have $-\ln p>1-p$. We thus conclude that the maximum is achieved on the boundary. Now for $p=0$, we have $H\left(0,r\right)=0$ and for $p=1$, we have $H\left(1,r\right)=0$; therefore we conclude that the maximum is achieved on the boundary where $r=2$. We find p that solves

$$\max_{p\in(0,1)}H\left(p,2\right)=\max_{p}\frac{2\left(1-p\right)p}{\left(2-p^{2}\right)}$$

and the solution is

$$p^{*}=2-\sqrt{2}.$$

Finally, for $r=2,p^{*}=2-\sqrt{2}$ and $n\rightarrow\infty$ we have

$$\frac{E_{F_p}[\tilde{A}_{r,n}]}{E_{F_p}[OPT_{r,n}]}=\frac{1}{2}\sqrt{2}\simeq0.707\,11.$$

References

Athey, S. and I. Segal (2007). "An efficient dynamic mechanism." Working paper.

Ausubel, L. and P. Milgrom (2002). "Ascending auctions with package bidding." *Frontiers of Theoretical Economics*, **1(1)**, Article 1.

Ausubel, L. M. (2004). "An efficient ascending-bid auction for multiple objects." *American Economic Review*, **94**(5), 1452 – 1475.

Battigalli, P. and M. Siniscalchi (2003). "Rationalizable bidding in first price auctions." *Games and Economic Behavior*, **45**, 38–72.

Bergemann, D. and S. Morris (2005). "Robust mechanism design." *Econometrica*, **73**, 1771 – 1813.

Bergemann, D. and J. Välimäki (2007). "Dynamic marginal contribution mechanism." Working paper.

Cavallo, R., D. Parkes and S. Singh (2007). "Efficient online mechanisms for persistent, periodically inaccessible self-interested agents." Working paper.

Cho, I.-K. (2005). "Monotonicity and rationalizability in a large first price auction." *Review of Economic Studies*, **72(4)**, 1031 – 1055.

Cole, R., S. Dobzinski and L. Fleischer (2008). "Prompt mechanisms for online auctions." In *Proc. of the 1st International Symposium on Algorithmic Game Theory (SAGT'08).*

Dekel, E. and A. Wolinsky (2003). "Rationalizable outcomes of large private-value first-price discrete auctions." *Games and Economic Behavior*, **43(2)**, 175–188.

Gershkov, A. and B. Moldovanu (2007). "The dynamic assignment of heterogenous objects: A mechanism design approach." Working paper.

Hajiaghayi, M., R. Kleinberg, M. Mahdian and D. Parkes (2005). "Online auctions with re-usable goods." In *Proc. of the 6th ACM Conf. on Electronic Commerce (ACM-EC'05).*

Lavi, R. and N. Nisan (2004). "Competitive analysis of incentive compatible on-line auctions." *Theoretical Computer Science*, **310**, 159–180.

Lavi, R. and N. Nisan (2005). "Online ascending auctions for gradually expiring items." In *Proc. of the 16th Symposium on Discrete Algorithms (SODA).*

Milgrom, P. and R. Weber (2000). "A theory of auctions and competitive bidding, ii." In P. Klemperer, ed., *The Economic Theory of Auctions*, pp. 179 – 194. Edward Elgar Publishing, Cheltnam, UK.

Said, M. (2008). "Information revelation and random entry in sequential ascending auctions." Working paper.

Wilson, R. (1987). "Game-theoretic analyses of trading processes." In T. Bewley, ed., *Advances in Economic Theory: Fifth World Congress*, pp. 33 – 70. Cambridge University Press.