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
Department of Industrial Engineering and Management
Ben-Gurion University of the Negev
ellasgv@bgu.ac.il

February 4, 2011

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 a common prior on this distribution.

JEL Classification Numbers: C70, C72, D44, D82

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

*We thank Olivier Compte, Dan Levin, and Benny Moldovanu, for many helpful comments.
†This author is supported by grants from the Israeli Science Foundation and the Bi-national Science Foundation.

1 Introduction

The classic format of sequential auctions attracts much practical attention in recent years, being often used over the Internet as well as in many other more traditional settings. However the classic theoretical literature on sequential auctions (a long sequence of works, initiated by Milgrom and Weber (2000)) contains few assumptions that do not fit today's dynamic and changing environment. In particular, in this paper we focus on two such assumptions: (1) the assumption of a static set of players that are present throughout the auction, and (2) the common-prior assumption used to construct a Bayesian-Nash equilibrium. Over the Internet, and in other dynamic marketplaces that now become wide-spread, players usually arrive dynamically over time, and because of geographical and cultural spread may not share a common-prior.[1]

In an effort to incorporate these two new characteristics in a sequential auctions model, this paper describes an analysis of sequential auctions with dynamic arrivals and uses the solution concept of undominated strategies and therefore assumes nothing about common prior beliefs of the players. To get some intuition how these two new characteristics alter the nature of classic results, consider the following very simple example: 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, and is 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 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 (or efficiency) when the player with the higher value incorrectly underestimates the probability of the new second-period arrival. This loss of efficiency is *inherent* to the dynamic setting, a real phenomenon that we do not wish to ignore, but rather to highlight and analyze for this common mechanism. Thus, the main conceptual focus of this work is: *what is the "social damage" due to the use of a sequential ascending auction in an incomplete-information setting with dynamic arrivals and no-common-priors.*

Our answer relies on two relatively non-standard conceptual components, that we view as the main contributions of the paper. The first conceptual component is that instead of producing a dichotomic judgement, whether the outcome is efficient or not, we produce a *quantitative* judgement, of how much social welfare is lost. For this purpose we use the robust solution concept of undominated strategies, and analyze the sequential ascending auction, with an arbitrary number of items and any structure of players' arrivals. We assume that the bidder's valuation is private

[1]We are not unique in the attempt to remove the no common-priors assumption, following Wilson's Doctrine (Wilson, 1987); such attempts were recently being conducted in various settings under the name of "robust mechanism design" e.g. Bergemann and Morris (2005), and in the computer science literature. More details are given below.

information known only to her, but that she cannot conceal her true arrival time. We then bound the efficiency loss over the set of all undominated strategies using a worst-case analysis: An adversary is allowed to determine the number of players, their arrival times, their values, and their (undominated) strategies. We show that the lowest possible ratio between the actual sum of values of winners (actual welfare) and the optimal sum of values (optimal welfare) is at least one half, regardless of how this adversary sets the different parameters.[2] To demonstrate that the average bound on the loss of efficiency will be much higher, we also analyze, for two items, a small modification of the above setting, where the adversary must draw the players' values i.i.d. from some fixed distribution (to her choice). We show that this small step towards an average-case analysis already significantly increases the lower bound on the ratio of achieved efficiency to $\sqrt{2}/2 \simeq 70\%$, no matter what the chosen distribution of values is.[3] 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 result strengthens the basic intuition that a worst-case bound of 50% is a "good" bound.

The second conceptual component is a novel approach for using an "activity rule", as a way to guarantee some efficiency by forcing a minimal level of competition. Activity rules are gradually being recognized as a useful tool for auction design, and are by now being used both in theory and in practice (see for example the detailed discussion in Ausubel and Milgrom (2002)). Our specific activity rule eliminates too early drop out decisions, in the following exact sense: if there exist t remaining items for sale (including the current auction), a player that drops when there are more than t active bidders in the current auction is disqualified from participating in any future auctions. Indeed, since at this price-point, say p, there remain more players than items (demand is higher than supply), the competition should still exist, and no player that values the item by more than p should be allowed to drop and to expect positive gains from this action. By introducing this disqualification sanction, we are able to tightly characterize the set of undominated strategies, and as a result to obtain the bound on the worst-possible relative inefficiency.

To the best of our knowledge, activity rules that restrict the behavior of players between consecutive sequential auctions were never studied before in a formal settings like ours. However, the specific activity rule that we use resembles "indicative bidding", a rule that was recently used e.g. to auction assets in the electricity market in the US (Ye, 2007).[4] With indicative bidding, players place initial bids, and only the bidders with the few highest bids are allowed to continue to the final round of the auction, that determines the winner. In our model, the highest bidder for the

[2]We also show that it is not possible to design a dominant-strategy (or ex-post) mechanism that is guaranteed to always obtain a strictly higher fraction of the optimal welfare. It is possible to design dominant-strategy mechanisms that achieve this exact one-half bound by making additional assumptions that seem less realistic to us, see a more detailed discussion in the sequel.

[3]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%.

[4]Central Maine Power, Pacific Gas and Electric, and Portland General Electric have recently used this format.

current item wins, and only the few bidders with the next-highest bids are allowed to continue to the next auction, while the bidders with the low bids cannot continue.

The remainder of this paper is organized as follows. The rest of this section is devoted to a short summary of related literature. Section 2 defines our sequential auctions model in general terms, where all formal definitions and proofs are given in Appendix A. Section 3 proves that the worst possible ratio between the resulting welfare in the sequential auction vs. the optimal welfare is one half, assuming players play undominated strategies. Appendix B shows that no dominant-strategy mechanism can obtain a better bound. 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.

1.1 Related literature

Dynamic mechanism design is recently attracting attention. Athey and Segal (2007) and Bergemann and Välimäki (2007) in the economics literature, and Parkes and Singh (2003) and Cavallo, Parkes and Singh (2009) in the computer science literature, study a general multi-period allocation model, in which a designer needs to perform allocation decisions in each period and players have private values for the different allocations, that may be stochastic and time-dependent. 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. 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[5].

Another modeling approach to dynamic mechanism design is taken by Gershkov and Moldovanu (2008). 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 (2008) to characterize optimal dynamic mechanisms, with respect to both the social welfare and the seller's revenue.

Several papers analyze mechanisms that maximize the seller's revenue when buyers are patient. Pai and Vohra (2008) study a setting where a seller wishes to sell C identical units over T time periods, where buyers arrive and depart over time, and obtain the revenue-optimal Bayesian incentive-compatible mechanism. Board (2008) studies posted-price mechanisms and obtains the optimal price path. Board and Skrzypacz (2010) show that, if items expire at some common cut-off point, then the optimal mechanism is a sequence of posted-prices with a final auction at the cut-off point. Two earlier works that study similar (slightly simpler) models are Vulcano, van Ryzin and Maglaras (2002) and Gallien (2006). Many earlier works on revenue-management are surveyed in

[5]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.

Board (2008).

Dynamic auctions have been studied in the computer science literature starting with Lavi and Nisan (2004). While all above papers assume a Bayesian setting, the computer science literature usually assumes a worst-case scenario, and aims to design dominant-strategy mechanisms that exhibit a small efficiency loss in a detail-free and robust way. Hajiaghayi, Kleinberg, Mahdian and Parkes (2005) study a setting similar to ours, and show that if prices are charged only after all auctions end (and depend on all the sequence of the auctions) then dominant-strategies can indeed be obtained while guaranteeing an efficiency level of at least 50%. Cole, Dobzinski and Fleischer (2008) require prices to be charged at purchase time, and 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 again guarantee a total efficiency loss of at most 50%.

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 a job-scheduling model that generalizes the 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.

Most of the auction theory literature uses solution concepts that are based on a notion of an equilibrium. We are non-standard in this aspect too, and several previous papers suggested a conceptually-similar approach. 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) advocate 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. However, these works still produce a dichotomic distinction, whether the outcome is efficient or not. There are by now several examples of the quantitative approach (but with a solution concept that is based on an equilibrium notion), e.g. Neeman (2003) and Chu and Sappington (2007). A recent paper by Babaioff, Lavi and Pavlov (2009) combines this two aspects, as we do here, and quantifies the efficiency loss over the set of all undominated strategies, in the different context of combinatorial auctions.

2 Strategic Analysis

A seller sells K identical items using a sequence of K single-item ascending auctions. There are n unit-demand bidders with 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 does not win an item. We study a dynamic setting where bidders arrive over time. 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 may participate only in the auctions for items $r_i, ..., K$. Thus, a bidder's type is a pair $\theta_i = (r_i, v_i)$, and the set of possible types for bidder i is Θ_i. We denote $\Theta = \Theta_1 \times \cdots \times \Theta_n$. We assume that the auctioneer knows the true arrival time of the player when the player truly arrives. We wish to maximize the social welfare – sum of winners' values – and evaluate a given mechanism according to the worst-case ratio (over all tuples of types) between the optimal social welfare and the social welfare that the mechanism obtains. The formal description of the mechanism and the bidders' strategies appears in appendix A.

As was demonstrated in the Introduction, when using the original English auction format (using a price clock that ascends continuously), players may find it beneficial to quit the current auction at a price well below their value, expecting to win one of the subsequent items for a lower price. To increase competition among bidders we suggest the following "activity rule" that will force players to compete at least until a carefully-chosen cut-off point, where the number of remaining players is equal to the number of remaining items:

An English auction with an activity rule: At auction t (for $t = 1, ..., K$), let p_t be the price point at which there remain exactly $K - t + 1$ bidders who have not dropped yet from the current auction (this is the point where the number of remaining bidders is equal to the number of unsold items). If several bidders drop together at p_t, so that more than $K - t + 1$ bidders are active before p_t and less than $K - t + 1$ bidders are active after p_t, the auction orders them so that it has a set of players of size exactly $K - t + 1$ that dropped last. We refer to p_t as the cutoff price at auction t. Then,

1. Bidders that do not belong to the set of $K - t + 1$ bidders that dropped last are not allowed to participate in subsequent auctions. This implies that $K - t$ of the bidders that participate in this auction are qualified to participate in the next auction, and one additional bidder wins this auction.

2. The next auction $t + 1$ starts from price p_t.

Note that the start price of auction $t + 1$ can be strictly below the end price of auction t. Since player i obtains zero utility if she drops before the cutoff price, it is intuitively a dominated strategy to do so, and we provide a complete and formal analysis for this in appendix A. Thus, we assume throughout that in every auction $t = r_i, ..., K$, bidder i does not drop before the cutoff price p_t, unless her value is lower than p_t. If indeed the price reaches the player's value, and this point is lower than p_t, the player drops at this point since subsequent prices will not be lower than p_t, and thus the player cannot obtain positive utility regardless of her actions.

As a direct result of this activity rule, we show that the bidders with the $K - t + 1$ highest values *among all bidders that arrive up to time t and have not won yet* are qualified for auction $t+1$. Formally, let Λ_t denote the set of bidders that arrive up to time t and have not won any item $1, ..., t - 1$. Let $X_t \subseteq \Lambda_t$ denote the set of bidders that participate in auction t i.e. all bidders that arrived up to time $t - 1$ and in each auction they participated were qualified to participate in the following auction and all bidders that arrive at time t. Finally, let $Q_t \subseteq X_t$ be the set of players at auction t that are qualified for auction $t+1$. Note that by definition $|Q_t| = \min\{|X_t|, K-t+1\}$. Then:

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

Proof. We prove that if a player is disqualified at auction s then for every $t \geq s$, $|X_t| \geq K - t + 1$, 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. ■

Proposition 2 *If* $|X_t| \geq K - t + 1$ *then the* $K - t + 1$ *highest-value bidders in* Λ_t *have the same set of values as the bidders in* Q_t.

Proof. We show that for every player $i \in \Lambda_t \setminus Q_t$ and player $j \in Q_t$ we have $v_j \geq v_i$, which implies the claim.

First observe that this holds for $i \in X_t \setminus Q_t$ and $j \in Q_t$: all players in X_t but not in Q_t have values smaller or equal to the cutoff price p_t, and all players in Q_t have values greater or equal to p_t, hence $v_j \geq v_i$. 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 $|Q_{t-1}| = K - (t-1) + 1$. 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 previous paragraph, any player $j \in Q_t$ has $v_j \geq v_{j'} \geq v_i$, and the claim follows. ■

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 indeed the following characteristic holds: the bidder never drops at auction t if more than $K - t + 1$ bidders are active and the price did not reach her value. [6] 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 $v(A) = 1 + 2\epsilon$. 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, then by proposition 2 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 $v(OPT) \leq 2x$. 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

[6] 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.

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

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

Let X be the set of players in OPT with the $2l+1$ highest values. Let $Y = \{i \in X \mid r_i \leq K-t\}$. Note that $|Y| \geq (2l+1)-t$: 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_{2l+1}^{OPT}, all players in Z are winners in A and they also have values at least v_{2l+1}^{OPT}; by the contradiction assumption there are at most l such winners in A. Thus $|Y \setminus Z| \geq (2l+1-t)-(l-t) = l+1 \geq t+1$.

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 proposition 2, 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_{2l+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_{2l+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 $v(OPT) \leq 2\sum_{k=1}^{\lfloor \frac{K}{2} \rfloor+1} v_k^A \leq 2v(A)$. We get:

Theorem 4 *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* $v(A, \theta) \geq \frac{1}{2} v(OPT, \theta)$.

To complete the picture, in Appendix B we show that no incentive-compatible mechanism can guarantee a bound higher than one-half in any ex-post equilibrium. We leave it for further study to determine whether different mechanisms can obtain a better bound for any tuple of undominated strategies.

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,[7] 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. 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 section 4 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 section 3 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 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 undominated 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). Note that we do not assume that players share a common prior on the distribution of the valuations. We only assume that such a distribution exists and the adversary randomly chooses the values according to it. 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 $\sqrt{2}/2 \simeq 70\%$ of the optimal efficiency, no matter what are 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

[7] The example uses only two objects but it can be easily adopted to any number K of objects simply by setting all values of players that arrive before the last two periods to be sufficiently small, and in the last two periods replicate the same example.

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 proposition 2 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 since only the two highest bidders are qualified to participate in the second auction and one of them is determined to be the winner of the first auction. 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 to the value of $\tilde{A}$.[8] Using these settings, the main result of this section is that the expected efficiency of the sequential auction is at least $\sqrt{2}/2 \simeq 70\%$ of the expected optimal efficiency:

Theorem 5 *For any choice of the parameters n, r, F,* $\frac{E_F[\tilde{A}_{n,r}]}{E_F[OPT_{n,r}]} \geq \frac{\sqrt{2}}{2} \simeq 0.707$.

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.[9]

[8] There always exists a tuple of undominated strategies such that A's value exactly equals $\tilde{A}$'s value. Recall that the adversary here may choose the strategy after she knows the values of the bidders. Therefore she can choose the following tuple: each player, other than the player with the $K - t + 1$'s highest value at auction t, drops when the number of remaining bidders is equal to the number of remaining auctions, or when the price reaches the player's value, the earliest of the two events. The player with the $K - t + 1$'s highest value at auction t drops if and only if the price reaches his value. This ensures that the value of the winner at auction t is the $K - t + 1$ highest among all bidders who are present at auction t.

[9] Since players are ex-ante 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:

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

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:

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

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

Proposition 6 *For any n, r, and $0 \le p < 1$,*

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

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. Note that the minimum is achieved when only two players arrive for the first auction and in this case the activity rule has no effect. However, as explained above, without the activity rule we cannot bound the value of A from below by the value of $\tilde{A}$. Without the activity rule this ratio can even decrease to zero (see the discussion in appendix A)

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.[10]

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 player with the highest value at time 1 has the $j+1$'th highest value among all the n players[11], 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

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

[11] If the distribution is discrete we use an arbitrary deterministic tie-breaking rule to ensure that the events (indexed by j) "highest player at time 1 has the (j+1)-highest value overall" are mutually exclusive.

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 $0 \leq j \leq n-r$,

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

(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 $p_j^{n,r} = 0$ 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 at least $n-j$ 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 coefficients 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 coefficients $w_l^{(j)}$ 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 coefficients $\beta_0^{(n,r)}, ..., \beta_n^{(n,r)}$ 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 coefficients, 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 7 *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 8 *For any cumulative probability distribution F, and any n, r,* $\frac{E_F[\tilde{A}_{n,r}]}{E_F[OPT_{n,r}]} \geq \frac{\sqrt{2}}{2} \simeq 0.707$.

This completes the proof of theorem 5.

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 79.6%. For a uniform distribution on any other interval $[a, b]$ the ratio is higher.

5 Conclusions

We have analyzed a common sequential auction structure with an activity rule. The results bound 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 regardless 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 auctions, 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 $\sqrt{2}/2 \simeq 70\%$ 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 on which we impose an activity rule, and obtain bounds on the efficiency loss incurred when players choose arbitrary undominated strategies. A still open question is whether revenue loss can also be bounded in such a setting, comparing the seller's revenue in our setting versus the revenue of the optimal (revenue-maximizing) allocation scheme. Towards this end, one may rely and extend the work of Neeman (2003), that analyzes the revenue loss in an English auction for a single item. This work still relies on the classic Bayesian-Nash equilibrium concept, but identifies the distribution of values that achieves the worst revenue loss.

Our analysis compares the welfare of the sequential auction mechanism to the optimal welfare. Another possibility is to compare the welfare of the sequential auction mechanism to the welfare of some other baseline mechanism. One such baseline mechanism might be a random mechanism, i.e. a mechanism that randomly picks a bidder among those who are present at the current auction. Such an analysis is performed by McAfee (2002), in the context of matching mechanisms. He bounds the term $(E[A]-E[R])/(E[OPT]-E[R])$ from below, where R is the value of a random matching. This measures how much of the distance between the random mechanism and the optimal mechanism is "covered" by the mechanism being evaluated. In future research we plan to adopt this way of analysis to our setting. This involves several conceptual and technical problems. First, we need to carefully define the random mechanism and consider whether it is indeed reasonable to use it as a baseline for comparison. It is not a mechanism that can be easily implemented, as one needs to determine prices in such a way that bidders will want to participate. Another possibility is to look for a different baseline mechanism that is actually being used in reality. Second, even if the random mechanism is chosen to be the baseline mechanism, for some distributions it will perform very well, e.g. distributions with a very small variance. The result of the comparison will also be a function of the number of players and their arrival times, and a comparison between the two mechanisms should identify the settings in which the sequential auction performs significantly better than the random mechanism.

Appendices

A A Formal Strategic Analysis

A.1 The Formal Setting

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. A 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.[12]

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 following scenario. 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 i.e. until the price clock reaches 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

[12] 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 time the price clock stops. However we assume that the auctioneer makes public the order that was chosen at every step. In any case the tie-breaking rule cannot depend on players' announcements.

between bidders 3 and 4. Then $D_1(5,1) = \{1\}, D_1(5,2) = \{2\}, D_1(5,3) = \{3\}$ or $D_1(5,3) = \{4\}$.

The total number of droppers at price p is denoted by s^t_p (note that this is always smaller than the number of players). We have that $s^t_p = 0$ if and only if no bidder dropped at p, and then $D_t(p,1) = \emptyset$. If $s^t_p > 0$ then $D_t(p,k) \neq \emptyset$ for every $k = 1, ..., s^t_p$.

Let X_t denote the set of bidders that participate in auction t. Specifically, X_t contains all bidders that arrive at time t, plus all bidders that arrive prior to time t, are not winners in any prior auction, and are qualified to continue to auction t. Define $x_t = |X_t|$. The entire information on a single auction can be described by the prices at which $D_t(\cdot, 1)$ is non-empty (a finite number of prices), the values of $D_t(\cdot,\cdot)$ in these prices and the preference order $\succ_{(t,p',k')}$ 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 price clock stopped, and dropper k within p is fully described by the tuple (history) [13],[14]

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

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 $h(t,p,k) = ((h_1(p^*_1, s^1_{p^*_1}), ..., h_{t-1}(p^*_{t-1}, s^{t-1}_{p^*_{t-1}}), h_t(p,k)))$. 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.

In principle we do not wish to place any further restrictions on valid strategies $b_i(\cdot,\cdot)$. For example, a player can decide to drop at a certain price but then at a later price decide to stay if the previous drop announcement was discarded due to tie-breaking. (this gives freedom to various "strange" strategies and hence strengthens the undominated-strategies analysis). The only limitation that we do have to require follows since the current definition gives rise to a situation that cannot be interpreted in the real world, for example by setting $b_i(\theta_i, h(t,p,k) = D$ if and only if $p > 5$. This does not make sense since it is not clear at what price the player actually drops. To

[13] 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.

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

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

avoid this we require a valid strategy to satisfy the following condition: For any t, p, $p' > p$, and any history $h(t, p', k)$ that satisfies, for every $p < p'' \leq p'$:

1. $D_t(p'', 1) = \emptyset$, and,

2. $b_i(\theta_i, h(t, p'', 0)) = D$,

we must have $b_i(\theta_i, h(t, p, s_p^t) = D$. In words, if a player announces "drop" in $p + \epsilon$ for every sufficiently small $\epsilon > 0$, and no other player has dropped after p, then the player must announce "drop" already at price p after all droppers at p were revealed.

A.2 Early drops

We wish to illustrate the fact that bidders may indeed find it beneficial to drop, in a certain auction, well before the price reaches their value, by repeating the example from the Introduction in a formal and complete way. 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 [15] 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 (the low bidder 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 non-symmetric equilibria) of each 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 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, $v_4 = 2H$, 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.

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

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. Moreover one can construct examples in which it is an undominated strategy for a player to drop as early as price zero. For this reason we introduce the following activity rule.

A.3 The activity rule and analysis of undominated strategies

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 (note that p_t^{ar} may be strictly lower than the end price of auction t). 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 $x_t \leq K-t+1$ 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 $Q_t = I_t(p_t^{ar}, k_t^{ar}) \subseteq X_t$ 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 $t+1$. Then $|Q_t| = \min\{|X_t|, K-t+1\}$

The main claim of this section is that the sequential English auction with our activity rule restricts 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. Recall that a strategy b_i of player i is a mapping from the player type (r_i, v_i) and the history $h(t,p,k)$ to a decision, whether to drop (D) or to remain (R).

Proposition 9 *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:*

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

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.

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

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

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 from (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 10 *Any undominated strategy $b_i(\cdot)$ satisfies:*

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

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 or equal to 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.[16] ■

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).

B Ex-post Mechanisms Cannot Do Better

In a "detail-free"/"robust" setting, the literature commonly uses the solution concepts of dominant-strategies (for direct mechanisms) and ex-post equilibria (for indirect mechanisms). A natural question is whether such mechanisms can obtain higher worst-case efficiency than our sequential auction with an activity rule. In this appendix we give a negative answer to this question, and show that every dominant-strategy mechanism for our setting can obtain, in the worst-case, at most half of the optimal welfare. Lavi and Nisan (2005) and subsequently Hajiaghayi et al. (2005) and Cole et al. (2008) prove similar results for slightly different settings. In particular they all rely on the fact that players have a departure time to prove the impossibility. The argument that we devise here does not rely on this assumption and is therefore suitable for our sequential auction setting.

By the direct-revelation principle, we can focus on direct mechanisms in which truthful reporting of the type is a dominant-strategy. We term these "truthful mechanisms". We additionally assume ex-post Individual Rationality, i.e. that a winner is never required to pay more than her declared value. We show the impossibility for the very restrictive setting of two items and three players, where it is common knowledge that players 1 and 2 arrive for the first auction and player 3 arrives for the second auction. Clearly, this only strengthens the impossibility, since if one can freely determine the number of items and players and their arrival times then one can replicate this limited setting.[17]

[16] 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.

[17] Even if one is not free to determine the number of players and their arrival times, one can set the values of all

Definition 11 (A direct-mechanism for a limited dynamic setting) *A direct mechanism is a set of four functions:* $w_1(v_1, v_2)$ *determines the winner (either* 1 *or* 2*) of the first item, and she pays a price* $p_1(v_1, v_2)$*, where* $p_1(v_1, v_2) \leq v_{w_1(v_1,v_2)}$*.* $w_2(v_1, v_2, v_3)$ *determines the winner of the second item (either* 1*,* 2*, or* 3*, but not* $w_1(v_1, v_2)$*), and she pays a price* $p_2(v_1, v_2, v_3)$*, where* $p_2(v_1, v_2, v_3) \leq v_{w_2(v_1,v_2,v_3)}$*. Such a mechanism is truthful if it is a dominant-strategy of every player to report her true type.*

Theorem 12 *Every truthful mechanism for the limited dynamic setting obtains in the worst-case at most half of the optimal social welfare.*

Proof. Fix any $\frac{1}{2} \geq \epsilon > 0$. Suppose by contradiction that their exists a truthful mechanism $M = (w_1(\cdot, \cdot), p_1(\cdot, \cdot), w_2(\cdot, \cdot, \cdot), p_2(\cdot, \cdot, \cdot))$ that always obtains at least $\frac{1}{2} + \epsilon$ of the optimal social welfare. We reach a contradiction via a series of three claims.

Claim 13 *If* $v_2 > \frac{v_1}{2\epsilon}$ *then* $w_1(v_1, v_2) = 2$*, i.e. player* 2 *must be the winner of the first auction.*

Proof. Suppose by contradiction that there exists an instance (v_1, v_2, v_3) such that $v_2 > \frac{v_1}{2\epsilon}$ and $w_1(v_1, v_2) = 1$. Consider another instance $(\tilde{v}_1, \tilde{v}_2, \tilde{v}_3)$, where $\tilde{v}_1 = v_1, \tilde{v}_2 = v_2$, and $\tilde{v}_3 = v_2$. The optimal social welfare in this instance is $2v_2$. We have $w_1(\tilde{v}_1, \tilde{v}_2) = w_1(v_1, v_2) = 1$, and therefore the social welfare that the mechanism obtains is $v_1 + v_2$. But

$$\frac{v_1 + v_2}{2v_2} < \frac{1}{2} + \epsilon$$

which contradicts the fact that the mechanism always obtains at least $\frac{1}{2} + \epsilon$ of the optimal social welfare. ■

Claim 14 *In the instance* $(v_1 = 1, v_2 > \frac{1-2\epsilon}{1+2\epsilon}, v_3 = 0)$*, the winners must be players* 1 *and* 2*.*

Proof. Any other set of winners has welfare strictly less than a fraction of $\frac{1}{2} + \epsilon$ of the optimal social welfare of this instance. ■

Claim 15 *If* $v_1 = 1, v_2 > \frac{1-2\epsilon}{1+2\epsilon}$*, and* $w_1(v_1, v_2) = 2$*, then* $p_1(v_1, v_2) \leq \frac{1-2\epsilon}{1+2\epsilon}$ *(note that* $\frac{1-2\epsilon}{1+2\epsilon} < 1$*).*

Proof. Suppose a contradicting instance (v_1, v_2, v_3) where $p_1(v_1, v_2) > \frac{1-2\epsilon}{1+2\epsilon} + \delta$ for some $\delta > 0$. Note that player 2 wins item 1 and pays the same price in the instance $(v_1, v_2, 0)$ (call this "instance 2"). Consider a third instance $(\tilde{v}_1, \tilde{v}_2, \tilde{v}_3)$, where $\tilde{v}_1 = 1, \tilde{v}_2 = \frac{1-2\epsilon}{1+2\epsilon} + \delta$, and $\tilde{v}_3 = 0$. By claim 14 player 2 must be a winner in the third instance, and by individual rationality she pays at most $\frac{1-2\epsilon}{1+2\epsilon} + \delta$. Therefore, in instance 2, player 2 has a false announcement ($\frac{1-2\epsilon}{1+2\epsilon} + \delta$ instead of v_2) that strictly increases her utility, a contradiction to truthfulness. ■

players besides the last three to be zero, by this returning to the limited setting and yielding the impossibility.

We can now reach a contradiction and conclude the proof of the theorem. Consider the instance $(1, 1, 5)$. Suppose without loss of generality that $w_1(1, 1) = 1$. To obtain at least half of the optimal social welfare we must have $w_2(1, 1, 5) = 3$. Thus player 2 loses and has zero utility. However if player 2 will declare some $\tilde{v}_2 > \frac{1}{2\epsilon}$ instead of her true type $v_2 = 1$ then by claim 13 she will win the first item and by claim 15 she will pay a price of at most $\frac{1-2\epsilon}{1+2\epsilon} < 1$. Thus player 2 is able to strictly increase her utility by some false declaration, contradicting truthfulness. ■

C Proof of Proposition 6

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

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

We first differentiate this expression with respect to n, to show that it decreases as n increases.

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

We concentrate on the term

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

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

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

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 $n \to \infty$ we have

$$\frac{E_{F_p}[\tilde{A}_{r,n}]}{E_{F_p}[OPT_{r,n}]} = \frac{\sqrt{2}}{2} \simeq 0.707\,11.$$

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.

Babaioff, M., R. Lavi and E. Pavlov (2009). "Single-value combinatorial auctions and algorithmic implementation in undominated strategies." *Journal of the ACM*. To appear.

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.

Board, S. (2008). "Durable-goods monopoly with varying demand." *Review of Economic Studies*, **75**(2), 391.

Board, S. and A. Skrzypacz (2010). "Optimal dynamic auctions for durable goods: Posted prices and fire-sales." Working paper.

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

Cavallo, R., D. C. Parkes and S. Singh (2009). "Efficient mechanisms with dynamic populations and dynamic types." Tech. rep., Harvard University.

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

Chu, L. and D. Sappington (2007). "Simple cost-sharing contracts." *American Economic Review*, **97**(1), 419–428.

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.

Gallien, J. (2006). "Dynamic mechanism design for online commerce." *Operations Research*, **54**(2), 291.

Gershkov, A. and B. Moldovanu (2008). "Efficient sequential assignment with incomplete information." 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)*.

McAfee, R. P. (2002). "Coarse matching." *Econometrica*, **70(5)**, 2025 – 2034.

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.

Neeman, Z. (2003). "The effectiveness of english auctions." *Games and Economic Behavior*, **43**, 214 – 238.

Pai, M. and R. Vohra (2008). "Optimal dynamic auctions." Working paper.

Parkes, D. C. and S. Singh (2003). "An MDP-Based approach to Online Mechanism Design." In *Proc. 17th Annual Conf. on Neural Information Processing Systems (NIPS'03).*

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

Vulcano, G., G. van Ryzin and C. Maglaras (2002). "Optimal dynamic auctions for revenue management." *Management Science*, **48**(11), 1388–1407.

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.

Ye, L. (2007). "Indicative bidding and a theory of two-stage auctions." *Games and Economic Behavior*, **58(1)**, 181 – 207.