Speeding Up Ascending-Bid Auctions*

Yuzo Fujishima†, David McAdams‡, and Yoav Shoham†
Computer Science Department (†) and Graduate School of Business (‡)
Stanford University
Stanford, CA 94305

Abstract

In recent years auctions have grown in interest within the AI community as innovative mechanisms for resource allocation. The primary contribution of this paper is to identify a family of hybrid auctions, called *survival auctions*, which combine the benefits of both sealed-bid auctions (namely, quick and predictable termination time) and ascending-bid auctions (namely, more information revelation often leading, among other things, to better allocations and greater expected revenue). Survival auctions are multi-round sealed-bid auctions with an information-revelation component, in which some bidders are eliminated from the auction from one round to the next. These auctions are intuitive, easy to implement, and – most importantly – provably optimal. More precisely, we show that (a) the survival auction in which all but the lowest bidder make it into the next round (the auction lasts for $(n-1)$ rounds when there are n bidders) is strategically equivalent to the Japanese ascending-bid auction, which itself has been proven to be optimal in many settings, and that (b) under certain symmetry conditions, even a survival auction in which only the two highest bidders make it into the next round (the auction lasts only two rounds) is Nash outcome equivalent to the Japanese auction.

1 Introduction

Auction theory has recently captured the attention of computer scientists and especially of AI researchers. One reason, of course, is the explosion of auctions on the internet and to facilitate business-to-business trade [Hansell, 1998; Cortese and Stepanek, 1998]. A more specific reason for interest within AI, however, is the use of auctions as a distributed protocol to solve resource allocation problems. For example, auctions as well as other market mechanisms are used in parts configuration design, factory scheduling, operating system memory allocation, ATM network bandwidth allocation, and distributed QoS allocation. [Clearwater, 1996; Yamaki *et al.*, 1996]. Within AI, market-oriented programming (MOP) [Wellman, 1993; Mullen and Wellman, 1996] has excited many researchers with the prospect of market-based control.

MOP and related approaches leverage ideas from economics and game theory, but these ideas provide little help with computational issues. This paper is concerned with the speed with which an auction terminates. (In the following we assume familiarity with some auction theory. Unfamiliar readers should consult the brief primer provided in Section 2.)

Sealed-bid auctions are attractive since they last exactly one round, but in some situations they have substantial disadvantages. Bidders in an auction often possess information that would be useful to other bidders to assess the value of the good for sale. Sealed bidding denies bidders the opportunity to learn about others' information during the course of the auction, which (as we discuss later) can lead to bad outcomes. In particular, the wrong bidder will sometimes win, the auction will tend to yield lower revenues, and winners will be subjected to more uncertain payoffs. In such situations, an ascending-bid auction is generally preferable on information-revelation grounds. In an ascending-bid auction, bidders with high estimates of a good's value can see the drop-out points of other bidders who have lower estimates. This information – a sense of just how bad is the worst information of others – helps bidders to assess the good's worth to them. For example, in an auction of a purported Rembrandt painting, very aggressive bidding by several experts should convince a non-expert that the painting is unlikely to be a fake. In a sealed-bid auction, bidders can not similarly condition their bidding behavior on the behavior (and hence information) of others.

Yet ascending-bid auctions have a distinct disadvantage. They can take a long and unpredictable amount of time to terminate. This shortcoming can be devastating, for example, when using auctions to allocate time-shared resources in real-time environments.

*This work was supported in part by the DARPA CoABS program, contract number F30602-98-C-0214

The primary contribution of this paper is to present and analyze a new sort of auction, the "survival auction", which shares the information-revelation property of ascending-bid auctions but which looks essentially like a multi-round sealed-bid auction. Here is how a survival auction works. In the first round, all bidders submit a sealed bid. As a function of these bids, the auctioneer announces which bidders "survive" to the second round, a minimum bid for the second round, and, most importantly, all losing bids.

In the final round of a survival auction, the remaining bidders have seen bids submitted by all of the others. Before making their final bids, therefore, they have some sense of the information possessed by the others. In the final stage of an ascending-bid auction, similarly, the last two bidders have seen drop-out points of all of the others, from which they can and should infer similar information.

The usual challenge of auction design is not so much to propose a reasonable-looking scheme as to analyze its equilibrium properties – in particular, its success in allocating the good optimally and in raising maximal revenue. One advantage of survival auctions over other sorts of "accelerated auctions" is that we are able to analyze them with the tools of game theory.

We show that one survival auction is strategically equivalent to the Japanese ascending-bid auction. In this survival auction, all bidders except the lowest survive each round and the minimum bid for each round is the losing bid of the previous round. We call this "the $(n-1)$-round survival auction" since it requires $(n-1)$ rounds to complete when there are n bidders. Strategic equivalence of two auctions implies that, as long as bidders behave rationally and do not care about the superficial details of the auction in which they are participating (which would not be true, for example, if bidders derive "entertainment value" only out of ascending-bid mechanisms), the same bidder will always win and pay the same amount in both auctions. Other survival auctions have the advantage of ending even more quickly than the $(n-1)$-survival auction, but are not strategically equivalent to the Japanese auction. If we impose certain symmetry assumptions, however, these survival auctions are still "Nash outcome equivalent" to the Japanese auction. This includes an auction which lasts only two rounds.

The rest of the paper is organized as follows. Section 2 is an auction primer, which includes a description of the most common types of single-good auctions, a brief literature review pertaining to the comparison of ascending-bid and sealed-bid auctions, and an introduction to the two principles of rational bidding most essential to our equilibrium analysis in Section 3.2. Section 3 introduces survival auctions and presents our results. The first result, in subsection 3.1, is that the $(n-1)$-round survival auction is strategically equivalent to the Japanese ascending-bid auction. The second result, in subsection 3.2, is that all of our survival auctions are Nash outcome-equivalent to the Japanese auction in the context of a model with several symmetry assumptions. Section 3.3 outlines how our results can be extended to the setting of multiple-good auctions. We conclude with a few comments about the relevance of our results.

2 Auction Primer

The most common auctions fall into three categories: ascending-bid auctions, descending-bid auctions, and sealed-bid auctions. Ascending-bid and descending-bid auctions are known together as open outcry auctions.

Ascending-bid auctions are the most prevalent in practice, accounting for an estimated 75% of all auctions worldwide [Cassady, 1967]. In the English auction, participants make successively higher bids. The winner is the last bidder and pays the last bid. Another common, ascending-bid auction is the Japanese auction. Here, the auctioneer continuously raises the price. Each bidder decides when to drop out, and once a bidder drops out he can not reenter. The last bidder to remain is the winner and pays the final price.

In a sealed-bid auction, all bidders submit a single bid to the auctioneer in ignorance of others' bids. The winner is the one who submitted the highest bid. Payment is made as a function of the bids. The most common sealed-bid auctions are the first- and the second-price. In the first-price auction, the winner pays his own bid. In the second-price auction, the winner pays the second-highest bid.

Descending-bid auctions are the least common in practice. The most familiar descending-bid auction is the Dutch auction. The auctioneer continuously lowers the price until a bidder expresses willingness to buy. That bidder wins and pays the final price.

Auction theorists have shown that under certain symmetry assumptions, when bidders are risk-neutral, do not care about the information possessed by others, and receive statistically independent information, all of the common auctions described above are optimal in the sense that they all allocate the good optimally (the bidder with the highest value for the good always gets the good) and all raise the highest possible expected revenue (see for example [McAfee and McMillan, 1987]). Relaxing various of these assumptions induces rank orderings of the common auctions, both in allocation and revenue terms. In particular, when bidders care about the information possessed by others, ascending-bid auctions outrank sealed-bid auctions in three important settings.

(1) Let $v_i(s_1, ..., s_n)$ represent bidder i's valuation for the good, which is allowed to be a function of all bidders' information. When $\frac{\delta v_i}{\delta s_i}(\vec{s}) \geq \frac{\delta v_j}{\delta s_i}(\vec{s})$ for all $j \neq i$ and for all $\vec{s} = (s_1, ..., s_n)$, ascending-bid auctions are always efficient whereas sealed-bid auctions are not [Dasgupta and Maskin, 1998].

(2) In the model of Section 3.2, when bidders are risk averse the ascending-bid auction has higher expected revenue than the second-price auction. In an ascending-bid auction, the drop-out points of losing bidders reveals information that helps remaining bidders make a tighter estimate of the good's value to them. Since winning

bidders are thus exposed to less risk, all bidders will be willing to bid more aggressively on average for the chance to win. (For more on this, see the discussion of extensions of the Revenue Equivalence Theorem to risk averse bidders in any auction or game theory text, such as [Fudenberg and Tirole, 1991].)

(3) In the model of Section 3.2, when bidders are risk neutral with affiliated private information the ascending-bid auction has higher expected revenue than any sealed-bid auction [Milgrom and Weber, 1982]. Roughly speaking, affiliation means that the more optimistic a given bidder's information, the more likely other bidders are to have more optimistic information. See Milgrom and Weber for a rigorous exposition.

Finally, our analysis presumes that bidders are aware of and act in accordance with the principles of rational bidding. For example, each bidder should reflect: "I make a payment only if I win the auction. Thus, when I choose how much to bid, I ought to presume that I will be the winner". We will call this *the winning bidder principle.* When bidders fail to abide by this principle, they will suffer from "the winner's curse". Typically, the bidder to win an auction is the one with the most optimistic information. If a bidder works off of the presumption that his information is average, he will systematically overvalue the good those times when he does win.

More subtlely, each bidder should reflect upon the marginal relevance of his bid. In the Japanese auction, a bidder considering dropping out at price p when there are k other bidders left in the auction should reflect: "Why should I drop out exactly now? If I stay in the auction and drop out instead at $p + \triangle$, I will still not win unless all other bidders drop out in between p and $p + \triangle$." When deciding to drop out at p, therefore, a savvy bidder will presume that all other remaining bidders will drop out immediately. We will call this *the marginal bid relevance principle.*

3 Survival Auctions

First, we describe the class of multi-round single-good auctions that we have named "survival auctions". Each survival auction is characterized by a survival rule, a minimum bid rule, and a final price rule. The survival rule specifies which bidders are allowed to continue in later rounds of the auction. Once excluded from participating in a given round, a bidder is not allowed to return in later rounds. An initial bid minimum is in effect in the first round; in all later rounds, the minimum bid is set as a function of the bids in earlier rounds. The final price rule specifies the winner's payment.

1. All bidders are active. The auctioneer announces an initial minimum bid.
2. Each active bidder submits a sealed bid. The bid must be no less than the minimum bid.
3. The auctioneer announces which of the active bidders remain active.
4. If only one bidder remains active, this is the winner. The auctioneer announces how much he must pay.
5. If more than one bidder remains active, then the auctioneer announces a minimum bid for the next round and all bids of those who have become inactive. Repeat from 2.

Although some of our results extend to a broader family of survival auctions, we will focus in this paper on a special subclass of survival auctions, called "our survival auctions", in which: (1) the number of survivers in the second round, the third round, and so on is commonly known in advance; (2) the bidders who survive are those who submit the highest bids; (3) the minimum bid is always equal to the *lowest* losing bid in the previous round (and the initial minimum bid is $-\infty$); and (4) the final round has two bidders.

3.1 $(n-1)$-round Survival Auction

Consider the specific survival auction in which only one bidder is eliminated at a time and the minimum bid in each round is the losing bid of the previous round. We call this the "$(n-1)$-round survival auction" since it takes $(n-1)$ rounds to complete when there are n bidders. In this section we prove that the $(n-1)$-round survival auction is "strategically equivalent" to the Japanese auction. Strategic equivalence is the strongest possible formal relationship to establish between two auction mechanisms.

Two auctions are strategically equivalent if there exists an isomorphism between their strategy spaces which preserves payoffs. A strategy z_i of bidder i maps each of bidder i's decision points (also known as action nodes or information sets) into a feasible action. Each of these decision points corresponds to the information that bidder i will possess if he reaches a certain potential moment of decision. For example, suppose that in the first round of the $(n-1)$-round survival auction bidder 1 bids b_1 and is eliminated. In this scenario, bidders $2, ..., n$ each have a second-round decision point at which they have the information "bidder 1 bid exactly b_1 in the first round while all other bidders bid at least b_1 in the first round". The set of feasible actions of each of these bidders at these decision points is simply the set of bids no lower than b_1. The strategy space Z_i of bidder i is the set of all of his strategies. For formal convenience we will include in a bidder's set of decision points the set of his "terminal points". Each of these corresponds to information the bidder can have when the auction terminates. (No decision is made at a terminal point.) A bidder's payoff from an auction is a function both of the auction outcome and of the information available to that bidder at the end of the auction. An outcome consists of an allocation of the good and payments by the bidders.

To show that there exists an isomorphism between the strategy spaces of two auctions preserving payoffs, one must show that: (1) An isomorphism exists between each bidder's sets of decision points in the two auctions. The structure which must be preserved by the isomorphism is that of decision point precedence. One decision point precedes another iff the follower decision point can only be reached if the precedent decision point is reached first. We will say that decision points in the two auctions are "the same" if they are related by this isomorphism. Note that (1) implies that there must be a bijection between the sets of information that each bidder can acquire immediately after every pair of same decision points; (2) There exists an isomorphism between feasible action sets at same decision points which is consistent with the precedence structure. That is to say, there is a set of (vectors of) actions which are consistent with each follower decision point being reached after its immediate precedent. The isomorphism must induce a bijection between all such sets for all same follower-precedent pairs. We will say that actions in the two auctions are "the same" if they are related by this isomorphism; and (3) Bidder payoffs are always the same at same terminal points. That is to say, if $g(\cdot), h(\cdot)$ are the isomorphisms of feasible actions and decision points, then each bidder's payoff at a terminal point t following actions $\vec{a}$ in the first auction must be the same as his payoff at the same terminal point $h(t)$ following the same actions $g(\vec{a})$ in the second auction.

Theorem 1 *The $(n-1)$-round survival auction is strategically equivalent to the Japanese ascending-bid auction.*

[Proof] We follow the proof outline established above.

(1) In the survival auction, the new information that bidder i receives in between rounds k and $k+1$, if he survives round k, is the identity of the loser in round k and the losing bid in that round. Each of his decision points in the $(k+1)$st round can thus be described by the $2k$-dimensional vector of all losing bidders and losing bids in the first k rounds. In the Japanese auction, the new information that bidder i receives in between the $(k-1)$st and kth drop-out, if he is not himself the one to be the kth drop-out, is the identity of the kth bidder to drop out and his drop-out point. His decision after the kth drop-out can thus be described by the $2k$-dimensional vector of the first k bidders to drop out and their drop-out points in the first k rounds.[1]

The isomorphism of decision point sets is the identity mapping. Thus, for example, we will say that a bidder's decision of what to bid in the second round of the survival auction after observing that bidder 1 lost with a bid of b_1 in the first round is the same as his decision of how long to wait before being the next to drop out in the Japanese auction after observing that bidder 1 is the very first to drop out at the price b_1. (A decision in the Japanese auction is not how long to wait until dropping out but how long to wait before being the next one to drop out. If someone else drops out first, a bidder should revise his willingness to remain in the bidding.)

(2) In the survival auction, each bidder in the $(k+1)$st round can make any bid higher than the minimum allowable bid in that round, which equals the losing bid in the kth round. In the Japanese auction, each bidder after the kth drop-out can decide to wait until any price higher than the last drop-out before being the next to drop out. Thus, the feasible action sets are identical and consistency with the decision point isomorphism is obvious.

(3) If play in the two auctions reaches the same terminal point, then all actions will have been same actions at same decision points. Thus, (a) the same bidder will win, (b) this bidder will pay the same amount, and (c) the information available to all bidders at the end of the auction will be the same. ♣

Since strategic equivalence implies outcome equivalence, all of the allocation and revenue advantages of the Japanese auction over sealed-bid auctions carry over to the $(n-1)$-round survival auction. Since the $(n-1)$-round survival auction predictably takes $(n-1)$ rounds whereas an ascending-bid auction can conceivably require hundreds or even thousands of rounds of communication, this marks a significant improvement in speed and reliability.

3.2 2-round Survival Auction

Can we do better than the $(n-1)$-round survival auction? For example, can we construct a survival auction which requires $O(1)$ rounds to terminate and which is strategically equivalent to the Japanese auction? Unfortunately, No. In any survival auction which takes less than $n-1$ rounds, the winner makes fewer than $n-1$ decisions. Thus, requirement (1) of strategic equivalence can not be met. Nonetheless, under certain symmetry assumptions, we can prove that all of our survival auctions are "Nash outcome equivalent" to the Japanese.

Two auctions are Nash outcome equivalent if each auction possesses a Nash equilibrium such that, when those Nash equilibria are played out, the same bidder wins and pays the same price after every realization of uncertainty. Nash outcome equivalence is weaker than strategic equivalence since it is model-specific. Furthermore, bidders may not coordinate on Nash equilibrium play, especially if there are multiple equilibria.

Consider the following model. Each bidder i possesses one-dimensional private information s_i. All signals $\{s_i\}$ are drawn from a common distribution and possess a symmetric correlation structure. (The conditional distribution of all signals $\{s_j\}_{j\neq i}$ given $s_i = x$ is symmetric in all of the other signals and the same as the conditional distribution of all signals $\{s_j\}_{j\neq k}$ given $s_k = x$ for all $k \neq i$. Similarly, all conditional distributions given

[1] The possibility that bidders can drop out simultaneously in the Japanese auction presents a wrinkle. If we presume that bidders can react instantaneously to others' drop-outs, however, then this does not change the analysis. Sometimes, based on one bidder's drop-out, another bidder wants to drop out immediately. We allow such immediate reaction, with the proviso that the bidder who dropped out first can not reverse his decision to drop out after observing the other bidder's drop out. This way, every bidder is faced with a new decision (and new information) after every drop-out.

two signals, three signals, and so on are symmetric and equal.)

Each bidder i's willingness to pay for the good is a function of all private information, $v(s_i; \{s_j\}_{j \neq i})$. For each bidder, this function is symmetric in all other bidders' information, increasing in all signals, and satisfying

$$\frac{\partial v_i}{\partial s_i}(\vec{s}) \geq \frac{\partial v_j}{\partial s_i}(\vec{s})\ \forall j \neq i\ \forall \vec{s}$$

This inequality states (loosely) that "Each bidder always cares about his own information more than others do." Thus, given any set of information available to both bidders i and j, bidder i has a higher valuation of the good iff his signal is higher. Finally, bidders have identical expected utility functions defined over the difference between valuation and payment. That is to say, there exists a function $u : R \rightarrow R$ such that bidder i's utility from getting the good at price p when his valuation is v_i equals $u(v_i - p)$. This allows bidders to be risk averse or risk loving, as $u''(\cdot) < 0$ or $u''(\cdot) > 0$.

Theorem 2 *In the model of this section, all of our survival auctions are Nash outcome equivalent to the Japanese auction.*

[Proof] Recall that in "our survival auctions" the minimum bid equals the lowest losing bid of the previous round and the final round always has two bidders. To conserve space, some details have been left to the reader.

First, we construct a symmetric equilibrium of the Japanese auction. By the principles of rational bidding (see the primer) and model symmetry, at every decision point each bidder chooses the price at which he wants to be the next to drop out by working off of the presumption that all remaining bidders have the exact same signal as he does. Thus, in any symmetric equilibrium, bidder n must choose $v(s_n; s_n ..., s_n)$ as the price at which he will be the first to drop out. (The $n-1$ signals to the right of the semi-colon are those of bidders 1 through $n-1$, since this is bidder n's valuation function.) Now, if bidder 1 drops out first at b_1, all remaining bidders should infer that his signal s_1^* satisfies $v(s_1^*; s_1^*, ..., s_1^*) = b_1$. Bidder n must then wait until the price $v(s_n; s_1^*, s_n, ..., s_n)$ before being the next to drop out. And so on after every subsequent drop-out.

Consider now the one of our survival auctions with R rounds and k_r survivors of the rth round ($n > k_1 > ... > k_{R-1} = 2$). Again, by the principles of rational bidding and model symmetry, each bidder in round r must work off of the presumption that his signal is highest and equal to k_r others in any symmetric equilibrium. To minimize notational complexity, we will continue the proof under an additional assumption of risk neutral bidders.[2]

In the first round, then, bidder n wants to bid

$$E[v(s_n; s_1, ..., s_{n-1} | s_n \text{ high},\ k_1 \text{ equal})]$$

[2] If a bidder is not risk neutral, in round r he ought to bid so as to maximize his expected utility from paying his bid and winning the good given that his signal is the highest and equal to k_r other signals.

if all others follow a similar strategy, where "s_1 high, k_1 equal" means "$s_n \geq s_j\ \forall j$ and $s_n = s_{j_1} = ... = s_{j_{k_1}}$ for some set of k_1 other bidders, $\{j_1, ..., j_{k_1}\}$". Suppose that bidders $1, ..., n-k_1$ lose in the first round with bids $b_1, ..., b_{n-k_1}$. Given the specified first round bid strategy, all remaining bidders should infer each of their signals, $\{s_i, i = 1, ..., n-k_1\}$, by the relation

$$E[v(s_i; \{s_j\}_{j \neq i} | s_i \text{ high},\ k_1 \text{ equal})] = b_i.$$

Similarly, in the second round, bidder n wants to bid

$$E[v(s_n; s_1^*, ..., s_{n-k_1}^*, s_{n-k_1+1}, ..., s_{n-1} | s_n \text{ high},\ k_2 \text{ equal})]$$

if all others do the same, and so on in all subsequent rounds. These second-round bids will always be feasible since they will never be less than the lowest bid in the first round. ($s_n \geq s_1^*, ..., s_{n-k_1}^*$ and bidder n knows in the second round that $s_{n-k_1+1}, .., s_n \geq s_1^*, ..., s_{n-k_1}^*$ whereas the lowest losing bidder based his first round bid on less optimistic presumptions.) Thus, bidders have no incentive to bid less aggressively in the first round to avoid a possibly disadvantageous position in the second round, nor similarly in later rounds, and these strategies do indeed form a Nash equilibrium.

The winner in this equilibrium is always the same as in the equilibrium of the Japanese auction, since it is the bidder with the highest signal. Finally, in both equilibria the price when n wins and $n-1$ is second equals the final drop-out point / bid of $n-1$,

$$p = v(s_{n-1}; s_1^*, ..., s_{n-2}^*, s_n | s_n = s_{n-1}),$$

since the final round of the survival auction has two bidders. ♣.

Define the "2-round survival auction" to be the one in which only two bidders survive to the second round and in which the minimum bid is the lowest losing bid. This auction always takes two rounds to complete.

Corollary 1 *The 2-round survival auction is Nash outcome equivalent to the Japanese auction.*

3.3 Extension to Multiple Goods

All of our arguments in the context of single-good auctions easily extend to prove analogous claims about multiple-good auctions. In particular, consider the "k-good Japanese auction" in which each bidder may receive any number of units of the good at auction. Each bidder begins by choosing a number of "active bids". As the price rises, each bidder can lower his number of active bids. The auction ends when exactly k active bids remain. Similarly, define "k-good survival auctions" to be those in which each bidder can have multiple active bids and in which the last k active bids win. In particular, define the one-at-a-time k-good survival auction to be the one which eliminates one active bid at a time. (The number of rounds that this auction takes depends on the initial number of active bids.) Also, let "our k-good survival auctions" be those in which: (1) the number of surviving active bids in each round is known in advance as a function of the number of initial active bids; (2) the highest bidders survive; (3) the minimum bid equals the lowest losing bid in the last round; and (4) the final round has $k + 1$ bidders.

Theorem 3 *The one-at-a-time k-good survival auction is strategically equivalent to the k-good Japanese ascending-bid auction.*

[Proof] The proof is analogous to that of Theorem 1, but accounting for the richer strategy spaces of bidders since each must first choose his number of active bids. ♣

Theorem 4 *In the model of subsection 3.2, all of our k-good survival auctions are Nash outcome equivalent to the k-good Japanese auction.*

[Proof] The proof is analogous to that of Theorem 2, also accounting for the richer strategy space. ♣

It is important to note, however, that the k-good Japanese auction in the k-good setting does not share all of the desirable theoretical properties of the Japanese auction in the 1-good setting. In particular, the k-good Japanese auction will not always get the k goods to the bidders who value them most. Which multiple-unit auctions always allocate k-goods optimally is, as far as we are aware, still an open question.

4 Conclusion

There is a reason why the vast majority of auctions are ascending-bid auctions; they reveal information. Given rational bidding, this information revelation leads to better allocations and higher revenue, or, given irrational bidding, can at least mitigate the severity of the winner's curse. Their main practical disadvantage is a long and uncertain termination time. We have proposed a new class of auctions, *survival auctions,* which combine the speed and predictability of sealed-bid auctions with the desirable properties of ascending-bid auctions.

Other sorts of auctions appear to offer a similar blend. For example, in what we call "the binary price-search auction", the auctioneer queries all bidders whether they are willing to pay a given price. If only one bidder answers Yes, he gets the good at that price. If zero or more than one answer Yes, another query follows at a lower or higher price, in such a way to converge at logarithmic speed to a price that exactly one bidder will be willing to pay. Unfortunately, this and other "accelerated auctions" are difficult to analyze with the tools of game theory.

In any event, one should be cautious about applying our pure analysis in real-world situations, especially in consumer-based auctions. For example, it is not realistic to expect all bidders to be present at the very beginning of an auction, and it seems likely that psychological factors may lead human bidders to behave differently in auctions which are game-theoretically equivalent. We do believe, however, that survival auctions should provide a more desirable protocol than ascending-bid auctions in resource allocation problems with computerized agents.

5 Acknowledgments

We are grateful to participants of the Stanford CoABS and Stanford GSB Game Theory seminars for their helpful feedback, in particular Craig Boutilier, Moises Goldszmidt, Leonardo Rezende, Michael Schwarz, Robert Wilson, and Muhamet Yildiz. We especially thank Jeremy Bulow for his insightful comments.

References

[Cassady, 1967] R. Cassady. *Auctions and Auctioneering.* U. California Press, Berkeley, 1967.

[Clearwater, 1996] Scott H. Clearwater, editor. *Market-Based Control: A Paradigm for Distributed Resource Allocation.* World Scientific, 1996.

[Cortese and Stepanek, 1998] Amy E. Cortese and Marcia Stepanek. Good-bye to fixed pricing?: How electronic commerce could create the most efficient market of them all. *Business Week*, pages 71–84, May 4 1998.

[Dasgupta and Maskin, 1998] Partha Dasgupta and Eric Maskin. Efficient auctions. Technical report, Harvard University, 1998.

[Fudenberg and Tirole, 1991] Drew Fudenberg and Jean Tirole. *Game Theory.* MIT Press, Cambridge, Mass., 1991.

[Hansell, 1998] Saul Hansell. Auctions of goods and services thrive on internet. *New York Times*, April 2 1998.

[McAfee and McMillan, 1987] R. Preston McAfee and John McMillan. Auctions and bidding. *Journal of Economic Literature*, 25:699–738, June 1987.

[Milgrom and Weber, 1982] P. Milgrom and R. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50(5):1089–1122, September 1982.

[Mullen and Wellman, 1996] Tracey Mullen and Michael P. Wellman. Some issues in the design of market-oriented agents. In M.J. Wooldridge, J.P. Muller, and M. Tambe, editors, *Intelligent Agents II: Agent Theories, Architectures, and Languages*. Springer-Verlag, 1996. Also available at http://ai.eecs.umich.edu/people/mullen/atal-revised.ps.Z.

[Wellman, 1993] Michael P. Wellman. A market-oriented programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research*, 1:1–23, 1993. Avaliable at http://www.jair.org/.

[Yamaki *et al.*, 1996] Hirofumi Yamaki, Michael P. Wellman, and Toru Ishida. A market-based approach to allocating qos for multimedia applications. In *The Second International Conference on Multi-Agent Systems (ICMAS-96)*, pages 385–392, 1996.