

AGENT-BASED SIMULATION OF DYNAMIC ONLINE AUCTIONS

Hideyuki Mizuta

Tokyo Research Laboratory
IBM Research
1623-14, Shimotsuruma, Yamato-shi,
Kanagawa-ken 242-8502, JAPAN

Ken Steiglitz

Department of Computer Science
Princeton University
35 Olden Street
Princeton, NJ 08544-2087, U.S.A.

ABSTRACT

The need to understand dynamic behavior in auctions is increasing with the popularization of online auctions. Applications include designing auction mechanisms, bidding strategies, and server systems. We describe simulations of a typical online auction, where the duration is fixed, and the second-highest price is continuously posted and determines the winner's payment. We modeled agents of exactly two types, idealizations and simplifications of those observed in practice: *early bidders*, who can bid any time during the auction period, and *snipers*, who wait till the last moments to bid. This allows us to study the interactions of the two types of bidders during the course of auctions, and the effects of the two strategies on the probability of winning, the final price, and the formation of price consensus in iterated auctions. Results show that 1) early bidders can win with a lower price on average than snipers, but much less often; 2) the late bidding strategy of snipers is effective; and 3) in iterated auctions, adjustment feedback of motivational parameters can lead to effective price consensus with small fluctuations.

1 INTRODUCTION

The use of online auctions is rising at a dramatic rate, and in general many segments of the economy are becoming granulated at a finer and finer scale. There is a large and well developed literature on auctions, comprising a long sequence of remarkable theoretical papers (such as Vickrey 1961, Milgrom and Weber 1982, Riley and Samuelson 1981, Myerson and Satterthwaite 1983), more recent experimental work to test theory with human subjects in the laboratory (for example, Kagel and Dyer 1988), and more recently observing (presumably) human behavior in internet auctions (for example, Lucking-Reiley et al. 2000). At the same time an intense interest is developing in software agents (or "bots"; see for example, Kephart et al. 2000).

Thus, understanding behavior in auctions, and especially the interaction between the design of auctions, agent behavior, and the resulting allocations of goods and money has become increasingly important—first because we may want to design auctions that are as profitable as possible from the sellers' point of view, but also because we may want to bid in auctions, or design computer systems that respond well to the loads that auctions generate.

We describe in this paper some preliminary simulation work that attempts to bridge the gap between experimental work with human subjects and theory. In particular we describe a simulation model for bidder behavior in what is the most popular style online auction: a second-price auction that is *dynamic* in the sense that the current second price is posted, and there is a hard deadline at which time the auction ends. It is obvious—but important to remember—that this auction design thus differs from the Vickrey (sealed-bid second-price one-shot) auction (Vickrey 1961) in very significant ways, and in practice this has a strong effect on bidder behavior. For example, in a Vickrey auction with independent private values, there is no incentive to do other than bid one's own value. In practice we find some bidders in the online auction bidding throughout the bidding period (typically a week), sometimes many times, while other bidders wait until the final seconds. Participants in these auctions appear to be revising their values based on observations of who is bidding what, and sometimes appear to be behaving irrationally. Such dynamic interaction is not easily captured in any of the usual theoretical models, and provides a strong motivation for simulation.

2 SYSTEM DESCRIPTION

We consider a single auction with the sale of one item by one seller to one of n bidders, who submit their bids over time in the interval $[0, T)$ to an *auctioneer*, who awards the item to the highest bidder at closing time (see Figure 1). A bidder can send more than one bid during the auction.




Figure 1: Auctioneer and Bidders




Figure 2: Auction Procedure

Throughout this paper, the starting bid price is fixed at 1, and the duration of the auction is 500 time units. As with the major online auctions, we assume “proxy bidding,” with the bidder with the highest bid winning the item and paying the second highest price. We ignore the usual minimum increments imposed in practice to prevent the price from moving too slowly. At each time period, the auctioneer posts to all bidders the current high bidder’s identity and the current second-highest bid s . A bid which exceeds the indicated second-largest bid will be accepted, and the auction status updated.

Figure 2 shows the execution of the auction. First, each bidder determines his first impression of the item. At each time period $0 \leq t < T$, each bidder receives the status of the auction, can update his estimation at a fixed schedule or probabilistically, and can submit bids if the conditions of his strategy are satisfied. It’s worth noting that the bidder can modify his valuation for the item according to the information he receives about the bids of others, and so the frequency and the timing of bidding are very important. This is in contrast to the vast majority of the models in the theoretical auction literature, which assume that bidder values are fixed throughout the auction (although they may be initially independent or affiliated). Examples of bidding strategies will be given in the next section.

3 EARLY BIDDERS AND SNIPERS

Observation of online bidder behavior quickly suggests that there are—to a first approximation—two markedly different types of bidders. The first type acts as if his actions are independent of others’ actions; he bids once, presumably with something like his value estimate, and lets the auction run its course. If he bids his value and the mechanism were strictly sealed-bid and second-price, this is well known to be a dominant strategy (Vickrey 1961). We call such an agent an *early bidder*. Further observation, however, reveals the fact that it is not uncommon for early bidders to bid more than once, perhaps because they revise their values, or

perhaps because they misunderstand the rules and think of the process as similar to an English (outcry) action, where bidding stops when no further bids are received, and the item is awarded at the highest, not the second-highest price.

In contrast, some bidders wait until the last possible practical moment (allowing for possible network delay) and then submit one bid. The motivation seems to be to reveal as little information as possible to other bidders until it is too late for them to act on it. A very knowledgeable dealer with private information, for example, would have good reason to remain on the sidelines until the last possible moment so as not to draw attention to an item that is known to be valuable only by specialists. We call such bidders *snipers*. In the double auction study of Rust, Miller and Palmer (1993), for example, analogous automated bidding programs that wait until offers to buy and sell are close and then move in to “steal the deal” at the last moment beat out other bidding strategies.

We can briefly characterize the strategy of early bidders as watch/modify/bid, and that of snipers as wait/bid. We will next describe the results of simulations with a mixture of early bidders and snipers. As we will see the heterogeneity of the bidder population produces interesting and plausible behavior. We next fill in some details of the particular bidding strategies used in this paper.

Each early bidder and sniper is initially assigned a highest possible bid, or limit price, L , and a private value, v , both randomly drawn from identical distributions. The private value v of an early bidder is initially low, because of his lack of confidence in his valuation, but he may increase his valuation up to L as the auction proceeds if he receives the right signals from the posted second price.

Each early bidder watches the auction status and at each period, if he is outbid, he modifies his private valuation v with fixed *watch probability* w according to the following rule

$$v' = \min\{cv, L\},$$

where the coefficient c is a random number uniformly distributed in $c \in [1.2, 2.0]$. After the modification, if the modified valuation v' is greater than the posted second price s , the early bidder sends bids v' with fixed *bidding probability* b .

On the other hand, a sniper waits until the end of auction, $t \in [T - 20, T)$. The uncertainty in network transmission time makes it unrealistic to assume that he can bid at precisely the last time period. At that moment, the sniper sends a bid in the amount $\min\{2s, L\}$ if $s < L$. The reason why he does not submit a bid in the amount $s + 1$ is that s is the second largest price and the winner's bid price may exceed it widely. The sniper therefore allows for some margin above twice the posted second price but no greater than his limit.

4 SIMULATION RESULTS

An example with 7 early bidders and 3 snipers is shown in Figures 3 and 4. There is a price jump at the last moment in Figure 3, but not in Figure 4.

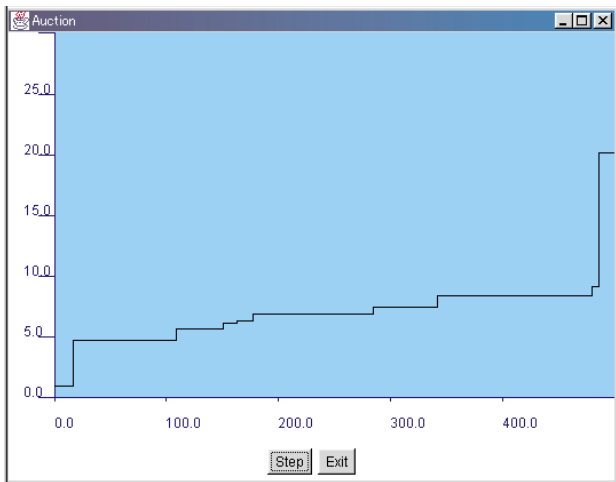


Figure 3: Simulation Result with a Mixture of Early Bidders and Snipers, Showing a Jump at the Last Moment

Figure 5 shows the histogram of winning price for 10000 simulation runs. We can see two main peaks here, one at low prices and one at high. Figure 6 aggregates the results when an early bidder wins, and Figure 7 aggregates the results when a sniper wins. This makes it clear that early bidders can actually get some items at low prices. But there is a crucial compensating factor: the vertical scale shows that the sales at low prices enjoyed by early bidders occur much less frequently. That is, snipers win most of the auctions, and early bidders get bargains, but relatively infrequently. It also happens fairly often that an early bidder wins at a price driven up by others.

Now we investigate the effects of bidding time on the winning price and the probability of winning. We change

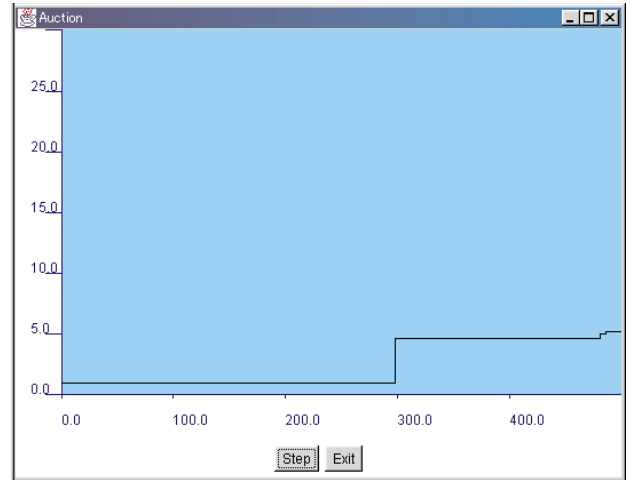


Figure 4: Simulation Result with a Mixture of Early Bidders and Snipers, Showing No Jump at the Last Moment




Figure 5: Histogram of Auction Price

one of the 3 snipers to one who bids once with his limit value L , and study the effects of his bidding at various predetermined times. The bidding times we tried were $t = 0, 250, 450,$ and 495 . For each case, we executed 10000 simulations with 7 early bidders and 3 snipers, one of which is special in the way described. The results are given in Table. 1. We can easily see that as the special sniper delays his bid, his winning price decreases and winning probability rises.

5 ITERATED AUCTIONS

In this section, we consider an auction repeated many times for identical items. We modify the bidders using a “motivation” parameter that affects the bidders’ other parameters and changes in ways dependent on the previous auction. Our aim here is to model the process of price formation over successive auctions through feedback from auction to auction.




Figure 6: Histogram of Auction Price when the Early Bidder Wins




Figure 7: Histogram of Auction Price when the Sniper Wins

Table 1: Effects of Bidding Time on the Winning Price and Probability of a Second Sniper

Bid Time	Win Price	Win Probability
0	17.86	0.380
250	16.97	0.454
450	15.82	0.537
495	14.21	0.683

Each auction is executed as before, that is, with fixed duration and second price winning. In this iterated model, both kinds of bidders have a motivation parameter M . The modified early bidders are called $M1$ bidders, and the modified snipers are called $M2$ bidders. $M1/M2$ bidders estimate some hypothetical fundamental price P with an exponential smoothing algorithm, $P = (1 - \alpha)P + \alpha p$, where p denotes the previous auction price and α is a small coefficient (typically $\alpha = 0.01$).

Parameters in early bidders are modified in $M1$ bidders as follows; The watch probability becomes $w = 0.02(1 + M)$, the limit value $L = (0.9 + 0.3M)P$, and the initial

internal valuation $v = 0.3P$, where P is the estimate of fundamental value, as mentioned above. Thus, a high motivation M causes the bidder to check the auction status more often on the average and increase his limit. His estimate of fundamental value increases his limit and private valuation.

The $M1$ bidders look at the auction status with probability w . If the current posted second price is $s > v$, v is updated to $\min\{L, cs\}$. The coefficient c is determined randomly at each time, $c = N(1.5, 0.1)$, where $N(\mu, \sigma)$ is normal with mean μ and standard deviation σ . Then, if $v > s$, the bidder bids v with probability $0.1L/s$.

After an auction, the motivation parameter M is updated to

$$M' = M + N(d, 0.02) + 0.05(1 - p/P),$$

where the mean drift d is given by $0.1/p$. If this bidder wins the item, his motivation M decreases according to $M' = M - 0.05$. The drift in the motivation is meant to reflect uncertainties that may arise in the interval between two successive auctions. The motivation of a buyer will increase if he waits for an item a long time. If the auction price is high, it's likely that the next seller will appear soon and the auction will be executed before the motivation of a buyer grows, and vice versa. The stationary condition of M and P determines the fundamental price.

We also created $M2$ bidders who wait until the last moment and send a bid with $\min\{P, 2s\}$. $M1$ and $M2$ bidders can compete and this causes interesting dynamics in the auction price.

In experiments with these bidders, starting with random values of M and P for each bidder, the variance of the clearing price decreases with time and the price approaches the fundamental price (see Figure 8).




Figure 8: Simulation Result of Iterated Auctions: Price and Motivation

6 FUTURE WORK

The work reported here is preliminary and attempts to capture behavior that has been observed anecdotally by the authors in some online auctions. The rise of the online venue for auctions has made available a tremendous amount of real data for studies of auction behavior, and has opened the way for empirical study and simulation models such as the one described here. The rigorous verification of simulation predictions by empirical data, however, presents some interesting problems, and is a challenging logical next step in understanding auction behavior through modeling.

The most important problem that arises is the calibration of the results across many auctions in situations when the items are not identical. For example, when collectibles are being sold, it may be impossible to find sets of objects that are effectively interchangeable. Even if the items have identical descriptions, or even appear superficially to be identical, they may differ in important ways, especially in condition. Sellers differ greatly in their ability and propensity to report condition, and uncertainties are introduced by variations in photograph quality. There is also the uncertainty in the likelihood that a given seller deliver the item fast, or even deliver it at all. Lucking-Reiley et al. 2000, in their study of the effect of seller reputation on fetched price, look at the prices realized for what is ostensibly a standardized item, an Indian-Head U.S. cent of a given year and condition. But any coin collector knows that the market value of a coin is a very critical function of condition, and that grading is very subjective. On the other hand, it seems probable that dynamic effects are less interesting and dramatic in auctions of new, commodity items (like computers), where posted prices in other venues give buyers a much better idea of likely market value, and where condition is much less of an issue.

In future work we plan careful comparison of simulation results with real-world data, but that could be a much slower and labor intensive process than first suggested by the volume of data that presents itself. Finding the right hypotheses to test and then testing them in a rigorous way is not an easy problem, and one that must be approached with circumspection.

7 CONCLUDING REMARKS

We have described simulations of a popular style online auction, where there is a fixed duration, a continuously posted second price, and bidders can bid more than once, being in general influenced by the bids of others. We modeled agents of exactly two types, idealizations and simplifications of those observed in practice: *early bidders*, who can bid any time during the auction period, and *snipers*, who wait till the last moments. This allows us to study the interactions of the two types of bidders during the course

of auctions, and the effects of the two strategies on the probability of winning, the final price, and the formation of price consensus in iterated auctions.

To summarize our main findings:

1. Early bidders can win with a lower price than snipers on average, but with a low probability;
2. The winning prices and win probabilities of snipers using several pre-determined bid times clearly show that the late bidding strategy of snipers is effective;
3. In iterated auctions, adjustment feedback of motivational parameters can lead to effective price consensus with small fluctuations.

The need to understand dynamic behavior in auctions is increasing with the popularization of online auctions. We believe that agent-based simulations such as those described here can be an important tool for understanding the complex economic phenomena in such auctions, where bidders interact over time via shared posted price signals. In future work we plan both mathematical analysis of knowledge acquisition by heterogeneous agents, and further study of empirical data from experimental and real auctions to refine and verify the model.

REFERENCES

- Kagel, J. H., and D. Dyer. 1988. Learning in common value auctions. In *Bounded Rational Behavior in Experimental Games and Markets*, ed. R. Tietz, W. Albers, and R. Selten, 184–197. Berlin: Springer-Verlag.
- Kephart, J. O., J. E. Hanson, and A. R. Greenwald. 2000. Dynamic pricing by software agents. To appear in *Computer Networks*.
- Lucking-Reiley, D., D. Bryan, N. Prasad, and D. Reeves. 2000. Pennies from eBay: the determinants of price in online auctions. Working paper, Vanderbilt University, Nashville, Tennessee.
- Milgrom, P. R., and R. J. Weber. 1982. Theory of auctions and competitive bidding. *Econometrica* 50 (5): 1089–1122.
- Myerson, R. B., and M. A. Satterthwaite. 1983. Efficient mechanisms for bilateral trading. *J. Economic Theory*, 29: 265–281.
- Riley, J. G., and W. F. Samuelson. 1981. Optimal auctions. *American Economic Review* 71 (3): 381–392.
- Rust, J., J. H. Miller, and R. Palmer. 1993. Behavior of trading automata in a computerized double auction market. In *The double auction market: Institutions, theories, and evidence*, ed. D. Friedman and J. Rust, Chapter 6. Massachusetts: Addison-Wesley.
- Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16 (1): 8–37.

AUTHOR BIOGRAPHIES

HIDEYUKI MIZUTA is a researcher of Tokyo Research Laboratory at IBM Japan. He received B.S., M.S. and Ph.D. degrees in physics from the University of Tokyo. He is a member of IPSJ and JPS. His research interests include dynamic economics with heterogeneous agents and e-Marketplace. His email and web addresses are <e28193@jp.ibm.com> and <www.tr1.ibm.co.jp/>.

KEN STEIGLITZ is a Professor of Computer Science at Princeton University. He received his B.E.E., M.E.E., and Eng.Sc.D degrees from New York University. A fellow of the IEEE and the ACM, he has received the IEEE Centennial medal, the IEEE Third Millenium medal, and the IEEE Signal Processing Society's Technical Achievement and Society awards. His research interests include e-commerce and computational physics. His email and web addresses are <ken@cs.princeton.edu> and <www.cs.princeton.edu/~ken/>.