Strategic bidding for multiple units in simultaneous and sequential auctions

Stéphane Airiau & Sandip Sen
Mathematical & Computer Sciences Department
University of Tulsa
600 S. College Avenue, Tulsa, OK 74104.
Phone: 918-631-2985
e-mail:stephane@utulsa.edu, sandip-sen@utulsa.edu

1. Introduction

Auctions provide an efficient mechanism to reach an economically efficient allocation of goods, services, resources, etc. between agents with local interests (Klemperer, 1999; Milgrom, 1989; Monderer and Tennenholtz, 2000). An amazing variety and quantity of goods, services are traded everyday in online auctions (Lucking-Reiley, 2000). These trades occur between and among businesses and consumers. In this paper we use the generic term *buyer* to represent an agent that bids in an auction. In practice, the agents can represent either business entities or individual consumers who are interested in purchasing a bundle of items from multiple online auctions.

The growth of online auction market, in size and variety, provides new challenges for buyers. A buyer's goal is to obtain the best deal possible, and to achieve this, he must keep track of multiple auctions at many different sites. Even if the buyer follows only one item in multiple, simultaneous auctions, it might lead to information overload and associated sub-optimal purchase decisions. To confound the problem, a buyer may want to buy not just one, but a bundle of items. In this case, the buyer is interested only in obtaining all the items in the bundle and not any proper subset. On a larger scale, this situation also describes an industrial producer who needs raw material and services for production.

Different aspect of the problem of bidding in sequential and/or simultaneous auctions has been studied: increasing the performance of simultaneous auctions (Matsumoto and Fujita, 2001; Preist et al., 2001; Shehory, 2001), defining strategy when auctions do not use the same mechanism (Byde et al., 2002), defining strategy to bid in sequential auctions using past history (C.Boutilier et al., 1999; Tesauro and Bredin, 2002), etc. Combinatorial auctions, i.e., auctions that offer bundles of goods have received particular attention from researchers, and facilitate the purchase of bundles by users. The allocation of bundles to bidders so as to maximize total revenue for the sellers, however, is known to be an NP-complete problem. Various approximation schemes, as well as exact schemes for limited bid types have been investigated (Hoos and Boutilier, 2000; K.Fujishama et al., 1999; Parkes, 1999; H.M.Rothkopf et al., 1998; Sandholm, 1999; Yokoo et al., 2001). In general, however, online auction marketplaces, such as eBay (eBay,), host simultaneous auctions that sell only one item at a time. Therefore, the buyers at such auctions have to put together bundles for themselves through buying the bundle constituents at different auctions.

The general *bundle bidding problem* involves bidding for multiple units of different items that constitute a bundle. An important problem may be to weight the value of one element of the bundle. Since only the valuation of the association of each element forming the bundle is known, each element does not have an intrinsic valuation. Obtaining a proper subset may not represent any value for the buyer. Therefore, the bidding problem for such individual buyers or producers become even more complex and involves the following decisions:

– select the auctions to bid in, from the numerous auctions being held at various auction sites,

– decide how much to bid and for what quantity in each of the selected auctions,

– factor in considerations of future auctions.

Related considerations include possibilities of obtaining too many or too few items. Because of the dynamics of online auctions, time constraints may not allow the buyer to compute an optimal decision, and the buyer may have to accept sub-optimal results.

We believe the bundle bidding problem provides a key opportunity for applying intelligent agent technology. An agent can automate the task of bidding for bundles on behalf of the associate user, i.e., the agent can take preferences for bundles from the user and try to put together the bundle by bidding at multiple simultaneous or sequential auctions held at different online auction sites. The research goal is to develop a bundle bidding strategy that takes into account the user preferences or bundle valuations, budget constraints, etc. We believe the use of such an agent has great potential to enhance user's profit and satisfaction.

We assume that our agent has some expectation of the number of auctions selling a given item type in the near future. In addition, the agent has expectations of closing prices, or valuations of other bidders, of the items in those auctions. Based on previous auctions, one may design a probabilistic model for future auctions. In this paper, the agents represent buyers who want to obtain a certain number of a specific item type over a given time period. Therefore, we are not addressing the completely general bundle bidding problem, but focusing our efforts on obtaining multiple units of a single item. The buyers may not prescribe an exact number of items to buy. Rather, they present a utility function mapping number of units of a given item to a numeric utility.

We consider different multiple unit English auctions that sell multiple units of this item type and that run simultaneously or sequentially over a period of time. Because of the number of auctions, the possible combinations to obtain a number of items may be large, and the bidding problem quickly turns out to be intractable for a buyer. The use of agent technology offers a solution for the buyer. Expectations on the future auctions and their closing price allow an automated agent to bid strategically on behalf of a user. The agent chooses the auctions to place a bid, decides how many items to bid for in an auction and the corresponding bid price.

This paper has two primary focus:

Effect of lookahead: We experiment with agents that have different time horizons of lookahead. Agents with longer lookahead are expected to perform better as they can be less myopic in their decision making.

Effect of risk-attitude: We want to evaluate the relative performance of risk-neutral, risk-seeking and risk-averse agents. Given a probability distribution of valuations of other bidders in the marketplace, different risk attitudes influence the highest bid an agent is willing to place in current auctions. This in turn can influence the relative profitability of the corresponding agents.

We first evaluate individual strategic agents interacting with non-strategic bidders who bid up to their fixed valuation, drawn from a distribution, for unit items. Next, we evaluate the performance of a pair of strategic agents competing against other non-strategic agents. Such head-to-head competition helps us evaluate relative merits of different lookahead or risk-attitudes in more demanding or competitive environments.

2. Auction framework

For our simulation, we designed the following scenario: auctions take place over $\mathcal{D} = \{D_1, \ldots, D_k\}$ days, on each day, l auctions take place, where each auction sells n identical units of a particular item. Each auction is then a multiple unit English auction. In contrast to *strictly ascending* auctions, where each bid must be strictly higher than all previous bids, we experiment with *sufficiently ascending auctions*, where any bid that will win an item is allowed (this means a valid bid is one that is higher than at least one active bid). We also assume that the time at which other agents have placed bids are not known to a user. This means that if an agent's bid is tied with another bidder's bid, and only one of the two will be cleared, the agent is uncertain about winning the item.

Each agent, i represents a buyer, B_i, with value function, v_i that maps number of units of the item to a real value (we will drop the subscript when the buyer reference is obvious from context). The goal of a buyer is to maximize the difference between the value of the number of units bought according to its valuation function and the cost of obtaining those units. To achieve this goal, each buyer may use the expectation of the closing price of auctions in the future.

We assume that the current ask price of every unit in each auction is known to all bidders, but the name of the buyer who has each bid is unknown. Each bidder has access to his current number of active bids. The auctions taking place in the same day are simultaneous: they have the same opening time, and they terminate when no buyer place a bid in any of the auctions. To simulate this, the auction house picks randomly one buyer from the set of buyers. This buyer has the right to place one bid in each auction. To place a new bid, a buyer announces how many units it wants to obtain, and how much it is bidding. The bid is valid when for each unit, the bid is a minimum increment above the the ask price for the auction. Invalid bids are rejected. When all of a buyer's bids are placed or rejected, the auction house picks randomly another buyer and gives it the hand. In real life of course, agent would react to events such as being outbid or having too few active bids. In this paper, we do not address the problem of making the decision of when to make new bids. Hence this random selection of bidders seems to be appropriate enough for our purpose. The sequence of selection of buyers will of course have an influence on the outcome of each auctions.

Let us emphasize two details here. Each buyer can know how many bids are active in a given auction, and it has access to the current price. But, if there are more bids than it has at the same price, it does not know the rank of its bids. If it wants to place a bid to obtain more units in the same auction, it can place a new bid with only one unit and wait to see whether it has out-bid one of his bid. As the auction house only allow one bid per auction, this is not a good strategy. As a consequence, if a buyer wants to obtain more units from one auction, it may have to out-bid its own bids to ensure having the desired number of active bids. Also, as the auction house can keep track of who got the hand in the past, it can detect when all the participants have successively refused to place new bids. Then, the auction house conclude that everybody is satisfied with the current state and closes all the auctions.

In the simulations, we are primarily interested in the buyers who want to buy a certain number of units over a given number of days. We also use non-strategic, dummy buyers with specific valuation functions in each auction. Each dummy buyer bids on the auctions of a particular day. These dummy buyer's bid up to their valuations, where these valuations are drawn from a probability distribution, specific to a particular day. The same probability distributions are used by strategic buyer's to form expectations of closing prices of auctions on a particular day.

3. Bidding Strategy

Let assume a buyer has expectations over the next d days. This means that it has some expectations on the number of units for sale in the next d days, and for each unit, it has an estimated closing price. More precisely, a buyer who looks forward to the k^{th} day has an expectation of the closing prices of all the auctions held between now and day k.

A buyer also has access to the asking prices of all current auctions. To bid strategically, we break the problem into two decisions: how much does it cost to buy k items now and how much does it cost to buy l items in the different auctions in the future? Having answered these two problems, it is easy to find k and l such that the buyer optimizes its benefit over the d days, i.e., maximize the difference between the valuation of buying $k+l$ items and the cost to obtain these $k+l$ items.

3.1. Obtaining k more items in current auctions

We breakdown the steps to calculate the *additional* amount to be bid to purchase k additional items in the simultaneous auctions being held on the current day:

– **Bidding for one more item in an auction:** Assume that there are n units for sell in a given auction. A buyer knows the current ask price for each unit, and knows how many active bids it has. δ denotes the increment to add to place a new bid, and $AP(i)$ denotes the i^{th} cheapest ask price of the auction. There can be three distinct cases when an agent tries to bid for one additional item;

  1. If the buyer does not have any active bid, then the additional cost to get one more unit is the lowest ask price, $AP(1)$, plus the increment δ.

  2. The buyer already has m active bids. By design, as will be clear from the following discussion, these are also the m cheapest bids in the auction. Hence the agent must out-bid its own bids to make sure it will out-bid the $(m+1)^{th}$ cheapest bid. The goal is to displace all of its previous m bids and one active bid by some other bidder. The resultant effect is that the agent will

place $m+1$ new bids, all at the value of $AP(m+1)+\delta$. The additional cost to obtain one additional item is then composed of two parts:

a) the amount of the new bid, $AP(m+1)+\delta$,

b) the additional value it must add to the bid values of each of its previous m active bids: $\sum_{i=1}^{m}\left(AP\left(m+1\right)+\delta-AP\left(i\right)\right)$.

3. The buyer owns all the active bids, there is no more items it can get from that auction.

– **Bidding for one more item in the auctions of the day:** Since we can compute for each auction of the day the price of getting an additional item, we can find the auction to bid on that will minimize the cost to obtain the additional item. If the buyer already have all the active bids in all auctions on that day, he can not buy an extra unit that day.

– **Bidding for k more items at the same time during a day:** To find the cost of obtaining k additional items, one needs to repeat k times the process of calculating the cost of obtaining one additional item, where each iteration is performed with the updated auction states after simulating the placement of the last bids.

3.2. Obtaining l more items in the future

Let us assume the buyer has maximum valuation of other buyer's for items in auctions to be held over the next few days. This translates into closing price expectations when the agent is not bidding in that auction. We assume that adding l bids above the expected closing price will enable the agent to win those l items in that auction, i.e., the adding of bid by this agent does not prompt other agent's to raise their bids in those auctions.

If one applies the same analysis for the current day's auction for the future auctions, then a buyer can form an estimate of the price it is likely to pay to obtain l more items in the future. But this cost is associated with two uncertainties: the buyer's expectations about other's valuations may be wrong and other agents may raise their bids in response to the buyer's bids above the predicted closing prices.

3.3. Different lookahead

We designed three different strategies corresponding to how far an agent looks ahead to the future:

- *oneday* is a strategy where no expectation on the future is used. It chooses bidding decisions to maximize utility only over the current day.

- *twodays* is a strategy where the buyer looks ahead one day. The buyer using *twodays*, unlike one using *oneday*, optimizes its utility over two days.

- *threedays* strategy looks two days ahead, and chooses bids to optimize utility over three days.

3.4. Different risk attitudes

The strategic agent has full information about the probability distribution from which the valuations of the non-strategic (dummy) buyers are drawn. We have used Gaussian distributions characterized by means, μ, and standard deviations, σ, to choose buyer valuations. *Risk neutral* (RN) agents make bidding decisions with the expectation of auction closing prices to be μ. *Risk seeking* (*risk averse*) agents expect auction closing prices to be less (more) than μ. In our experiment we study different degrees (levels) of risk seeking and risk neutral agents. In particular, severe risk seeking (SRS) and risk seeker (RS) bidders expect auction closing prices to be $\mu - 2\sigma$ and $\mu - \sigma$ respectively. Similarly, severe risk averse (SRA) and risk averse (RA) bidders expect auction closing prices to be $\mu + 2\sigma$ and $\mu + \sigma$ respectively. Compared to RN agents we expect RA agents to buy more items up front and RS agents to wait till later days to buy items. There are advantages and disadvantages to following either non-risk-neutral behaviors. We wanted to evaluate experimentally how these risk attitudes affected their relative performance, both when agents with different risk attitudes compete and when they individually compete with only non-strategic agents.

3.5. Agent valuation

The question we want to answer in the experiments is the following: if one believes the market will evolve in a certain way, what is the best strategy to bid now. We focus on cases where one expects the prices to vary significantly over days.

Let us motivate our experiments with a realistic scenario. One illustration of the problem we are considering may be to bid for supplies on behalf of a quality restaurant. Assume that the restaurant specializes in providing fresh fruits (and for our current simulation, we assume that the restaurant is specialized in only one fruit). Based on past weather, supply disruption news, etc. the owner may expect the quality or the quantity of the produces in the market to be high or low, the expected price varying consequently. We also consider that the restaurateur cannot look ahead too far: he may guess the state of the market for the next day, or for the two next days. His aim is to buy a certain quantity of fruits where a very small or a very large quantity is of less interest to him. This information is represented by the valuation function v of the restaurant. We use the following valuation function:

$$v(n) = \frac{c}{1 + e^{\frac{n_1 - n}{n_2}}},$$

where c is the maximum amount of money that a buyer is ready to pay for obtaining goods, n_1 and n_2 controls how many units are wanted, and how tight this number is. For the following experiments, we have used $c = 1500$, $n_1 = 20$ and $n_2 = 15$.

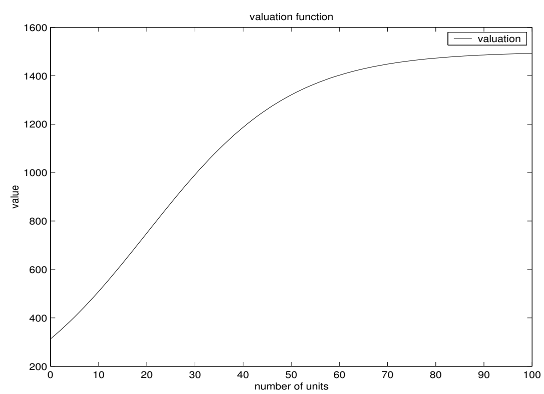

Figure 1. Valuation function used by "smart buyers".

The owner has contracted a supplier (a role which will be played by our agent) to buy the produces, etc. over the next D days (we use

$D = 5$). He may attend a market where producers sell their produces in multiple unit English ascending auctions. The auctions for a given day are held simultaneously. At the end of the D days, the supplier obtains n units, each unit i for a price of $c(i)$. The goal of the supplier is to maximize the restaurant owner's utility (we will also refer the utility as his gain) i.e. $Gain = v(n) - \sum_{i=1}^{n} c(i)$.

4. Experiments

In the setting of our experiments, we consider $l = 5$ auctions each day, each auction selling n units. We choose the valuations for the dummy buyers from Gaussian distributions that have different means on the different days. The means are chosen such that prices can go down significantly in the future. This variation allows strategic buyers to benefit by looking ahead of the future. The means of distributions representing the valuations of the dummy buyers over the D different days are described in Figure 2. The standard deviation σ is set to 1.5 for the next experiments.

In our experiments, the valuation functions and the price expectations are identical for all strategic bidders. The results are averaged over 50 runs with random non-strategic buyer valuations, drawn from given probability distributions. The order of bidder selection in each round of bidding is also randomly generated.

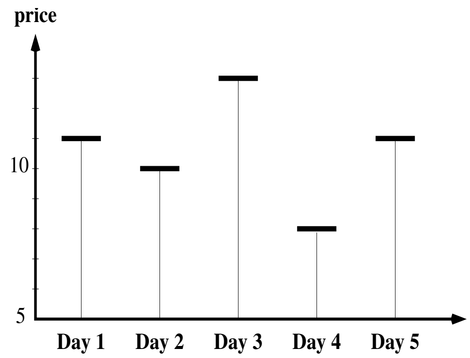

Figure 2. Setting for the dummy buyers.

The goal of this setting is to encourage the spread of the purchase of units over several days. A buyer who is considering only the current or first day will try to buy as much as he can during this day. An agent capable of looking at one or two days ahead may find better to wait for future opportunities, or to buy some units now and buy more later. From Figure 2 we see that at the start of the second day, an agent who is looking one day ahead knows the prices will be high on the next day, and hence has the incentive to buy as many items as possible on the second day itself. The situation, however, is quite different for a buyer who looks two days ahead, as it can predict that the market will be comparatively cheaper on day 4, and hence it does not rush to complete purchases on day 2. From this example, we see having knowledge of the future can be of advantage to strategic bidders. Of course, the advantage of such strategic bidding is lost if the information is incorrect or if a significant percentage of bidders in the marketplace bid strategically.

5. Results

In the first set of experiments, we compare the performance of the strategic bidders with different lookahead capacities. The results confirm that considering two days in the future yields better results than considering only the current day and/or the next day.

5.1. Study of different lookahead capacities

In the first set of experiments, we study the performance of one strategic risk-neutral bidder competing against dummy buyers. In this scenario, if the strategic agent can win by bidding above the valuation of dummy bidders. From Table I we see that strategic bidders with longer lookahead are successful in obtaining more items and higher utility.

Next, we run experiments where a *twodays* and a *threedays* lookahead strategic bidder compete with other non-strategic bidders. In this set of experiments we allowed different auctions to sell different number of

Table I. One strategic bidder against dummy bidders: utility and number of units obtained.

	utility	# units
oneday vs dummies	642	32
twodays vs dummies	736.7	37.4
threedays vs dummies	803.5	39.3

items. Table II presents the number of items bought by each agent in each auction and the average closing prices for auctions in each day. The entries of the form p/q represent the p units bought by the agent over the q units for sell in that auction. The column $\bar{p}$ contains the mean of the price of one unit during the day. The *threeday* bidder outperform the *twodays* bidder by a measure of 707.46 to 674.62. The *threeday* bidder, who looks two days ahead, performs better since during day 2, it foresees that day 4 will be cheap, so it waits for that day to get better deals. The *twoday* bidder, who considers only the next day, has to compete to obtain as many units as it can during day 2 since it only sees that day 3 will be much more expensive than day 2. When day 3 arrives, it has already bought enough units.

Table II. Competition *twodays* vs *threedays*: Breakdown of the units obtained over different days and average closing prices $\bar{p}$.

units obtained by *twodays*

day	a1	a2	a3	a4	a5	$\bar{p}$
1	0/1	1/1	7/7	2/7	7/7	12.20
2	2/7	10/10	1/10	1/1	1/1	10.18
3	0/3	0/2	0/5	0/1	0/1	13.03
4	0/5	0/7	0/5	0/5	0/15	8.70
5	0/7	0/7	0/7	0/1	0/1	11.10

units obtained by *threedays*

day	a1	a2	a3	a4	a5	$\bar{p}$
1	1/1	0/1	0/7	5/7	0/7	12.20
2	1/7	0/10	0/10	0/1	0/1	10.18
3	0/3	0/2	0/5	0/1	0/1	13.03
4	0/5	3/7	5/5	0/5	15/15	8.70
5	0/7	0/7	0/7	0/1	0/1	11.10

A second set of experiments uses a small number of strategic agents in the auctions. The presence of other strategic agents violate the assumptions of other bidders' valuations and that the others are not going to respond to higher bid by one agent. In terms of strategy, the expected prices of the next days should be thought of as the minimum price of the auctions for the next days. The uncertainty lies on how much these prices will increase. A buyer looking two days ahead performs at least as well as a buyer who is looking only one day ahead. But the competition is likely to adversely affect the utility of all the buyers. However, since the expected closing price may be significantly smaller than the actual closing price, one agent may be forced to give up bidding and wait another day, which may allow him to take advantage of new opportunities in the future.

We experimented with all possible combinations of two strategic buyers with non-strategic buyers and also the scenario where there were three distinct strategic buyers competing with non-strategic buyers. The table III presents the average gain of strategic bidders in different competitive environments. The percentage of loss for an agent using strategy *kdays* is the percentage decrease in utility in the current scenario with that obtained by the same agent type when competing only against dummy buyers (using the results contained in table I). As expected, a buyer using the *threedays* strategy performs better than

Table III. Experiment set 2: utility and number of units obtained.

	gain	% of loss
1*day* vs 1*day*	473.5	26.2
2*days* vs 2*days*	640.2	13.1
3*days* vs 3*days*	647.8	19.4
1*day* vs 2*days*	avg 666.9 1*day* 618.5 2*days* 715.1	3.7 3
1*day* vs 3*days*	avg 698.3 1*day* 606.4 3*days* 790.2	5.4 1.7
2*days* vs 3*days*	avg 765 2*days* 731 3*days* 799.5	0.7 0.5
1*day* vs 2*days* vs 3*days*	avg 680.5 1*day* 609.3 2*days* 647.1 3*days* 785	5.1 12.2 2.3

an agent using the *twodays* or *oneday* strategy. The utility is the lowest when two buyers using the same strategy compete against each other. This is because their identical reasoning processes lead them to bid in the same auctions, thus increasing competition and hence winning prices. However, the loss of gain encountered when two agents using different strategy compete against each other may not be as severe. When two agents using different strategies are competing, if one decides to give up bidding, the other may not take the same decision, and thus it obtains the units. There are two main effects to this, since one agent gets some units, he may not compete intensively to obtain more units in the future. Also, this allows the other agent to take advantage of new opportunities later. Consequently, the increase of competition does not necessarily implies an important loss of gain for the buyers. For example, the gain of a buyer using the *threedays* strategy does not decrease significantly (no more than 2.3% of the gain obtained in the first set of experiments).

5.2. Evaluating Risk Attitudes

In this section, we only consider agents using the *threedays* strategy. However, instead of using only the mean of the normal distribution to predict the closing prices of auctions in the future, they can consider also the standard deviation. We assume that each strategic agent has the same estimate of the standard deviation σ and the mean μ. Each strategic agent may not completely trust the estimate he has about the closing price of future options. He may revised his estimates. Since we know the estimate are correct, decreasing the estimate of the closing price may be considered as a risk seeking strategy, since the agent believes that the prices in the future will be cheaper than they may be, the future looks more attractive than it may be. On the other hand, increasing the estimate can be considered as a risk averse strategy, since the agent believes the price will be more expensive in the future. We defined the following risk levels:

severe risk seeking (SRS) Agent may believe his estimation is an overestimate, accordingly, he will believe that the actual closing price is $\mu-2\sigma$, auctions in the future will look much more attractive than they may be.

risk seeking (RS) The agent will believe that the actual closing price will be $\mu - \sigma$.

risk neutral (RN) The strategic agent will believe that the closing price is the mean μ, which is the most reasonable assumption.

risk averse (RA) The agent may not trust must his estimate for future auctions, believing it is underestimated and too risky. Thus he may overestimate the closing price by believing the closing price is $\mu + \sigma$. Auctions in the future will appear less attractive than they may be.

severe risk averse (SRA) The agent really believes his estimate is too low to be possible, and he increases it more severely, believing the closing price is $\mu + 2\sigma$.

Notice that we can also read these results as a study of the error caused by an inaccurate estimation of the mean of the closing price of the future auctions. Risk seeking corresponds to the case where the agent underestimates the mean, risk averse corresponds to the overestimate.

5.2.1. *A single strategic buyer competing with dummy buyers*

We first run experiment where a single strategic bidder competes against dummy bidders. From Table IV we see that risk neutral is the maximally profitable risk attitude. However, It is interesting to notice that it is better to be risk seeking than risk averse. This result stems from the fact that by underestimating the price, risk seeking can lose out on auctions in the first few days and still be able to make profits by winning auctions at a later time. In its eagerness to buy up items, risk averse strategies may end up paying more than if it had waited for an auction at a later day.

5.2.2. *Two strategic buyers competing with dummy buyers*

In this section, we present results with two strategic bidders with identical or different risk attitudes[1]. As we saw in Section 5.1, introduction

[1] This is equivalent to studying agents who have inaccurate expected price estimates for future auctions.

Table IV. Utility obtained by strategic buyers with different attitudes competing against dummy buyers.

threedays Vs *dummies*	utility
Severe Risk Seeking (SRS)	740.29
Risk Seeking (RS)	768.36
Risk Neutral (RN)	803.5
Risk Averse (RA)	677.89
Severe Risk Averse (SRA)	653.96

of multiple strategic buyers result in their losing utility compared to when they were competing with dummy buyers only. Table V present the results of the simulations of the auctions. Results are averaged over 25 runs only. We present the average utility achieved by each agent, and, in parenthesis, the average number of items that each agent obtained.

The primary observations from the matrix can be summarized as follows:

- The risk neutral strategic bidders perform better when pitted against players with all other risk attitudes.

Table V. Utility and number of items obtained by two strategic buyers with risk attitudes competing against dummy buyers.

	SRS	RS	RN	RA	SRA
SRS	659.5				
RS	697.2 (32.5) 749.8 (35.1)	686.3 (31.8)			
RN	693.3 (32.2) 753.3 (35.6)	679.6 (31.7) 736.2 (36.8)	647.8 (30.4)		
RA	686.4 (29.6) 683.9 (31.4)	697.7 (30.9) 676.9 (31.9)	722.1 (33.9) 678.3 (32.2)	685.58 (31.9)	
SRA	730.9 (31) 643.1 (28.9)	765.6 (34.8) 629.6 (29.8)	758.0 (35.5) 603.0 (28.8)	734.8 (34.5) 617.2 (29.3)	654.77 (31.2)

– Agents who are further apart in risk attitudes tend to gain more by competing than agents who are closer in their risk attitudes.

– Risk seeking attitudes perform better than risk averse attitudes.

– Utility gains come primarily from buying more items.

We believe the above observations will also hold with *twoday* lookahead agents, though the difference in performance may not be as pronounced.

5.2.3. *Study of difference in predictability*

In all the results presented so far, we used the means pictured in Figure 2 and the same standard deviation $\sigma = 1.5$. This information is used to draw the valuation of the dummy buyers, and to define the valuation of the strategic buyers with respect with their risk attitude. In this section, we keep the same means, but we study the variation of σ. For small values of σ, there is less uncertainty, it is easier to predict the closing prices.

The results are consistent with the observation made in the previous sections. In addition, we can note that the difference in performance between two agents, and the overall performance decrease. When we study the individual variation, as expected, the global trend is a decrease of utility with the increase of uncertainty (see Figure 3).

However, in few cases, some strategies increase their performance. Because of a large uncertainty, the performance of agents becomes closer(in some cases, the difference is not statistically different), but in this process, some agent can gain. It is the case when risk adverse agent competes versus a (severe) risk seeking agents. In both configurations, the risk averse agent increases its utility. With more uncertainty, the risk taken by a risk seeking agent is bigger, which naturally yields the decrease in performance. Meanwhile, this gives more opportunity to an risk averse agent, and this one takes advantage of it. We also observes this phenomenon when a risk seeking agent competes with a risk neutral, the risk seeking agent increases its performance (see Figure 3).

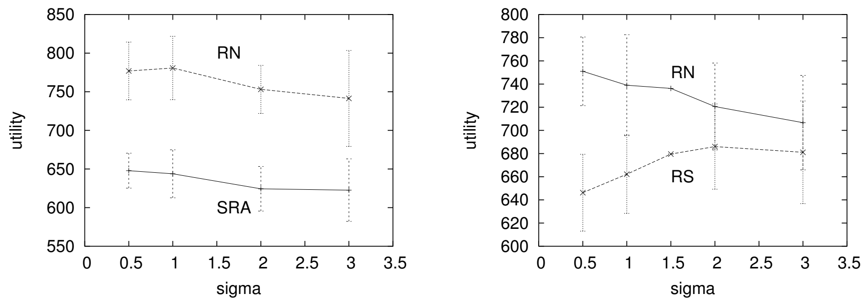

SRA vs RN RS vs RN

Figure 3. variation on sigma

6. Relation with other work

Ebay (EBay.Buyer Guide,), perhaps the most prolific online auction site, provides bidders with a "proxy-bidding" agent. The bidder sets a maximum amount he or she is willing to pay for the item and the agent automatically raises the bid until the bid has exceeded all the other bids or reached the bidder's maximum amount. The Michigan Auction-Bot (P.Wurman et al., 1998) provides an automated auction house for experimentation with bidding algorithms. Both these systems reduce the effort required of the bidder in monitoring the prices of a single item that is up for auction. These systems, however, are not appropriate for multi-unit auctions. The Spanish FishMarket (Rodriquez-Aguilar et al., 1997) is designed for comparison among different bidding strategies in Dutch auction, where a variety of lots are offered sequentially. Park *et al.* (S.Park et al., 1998; S.Park et al., 1999) presents a stochastic-based algorithm for use in the University of Michigan Digital Library, which is a many-to-many market. Gjerstad *et al.* (S.Gjerstad and J.Dickhaut, 1998) uses a belief based modeling approach to generating appropriate bids in a double auction. However, it is applied to a single double auction market place and does not allow agents to bid in multiple auctions. Vulkan *et al.* (N.Vulkan and C.Priest, 1999) uses a more sophisticated learning mechanism that combines belief based learning with reinforcement learning. Unlike Gjerstad's work, this work mainly focuses on learning the distribution of the equilibrium prices. Garcia *et al.* (Garcia et al., 1998) develop bidding strategies in the context of Spanish Fishmarket tournament. Agents compete in a sequence of Dutch auctions, and use a combination of utility modeling and fuzzy heuristics to generate their bidding strategy. Their work focuses on Dutch rather than English auctions, and on a sequence of auctions run by a single auction house rather than parallel auctions run by multiple auction houses. Preist *et al.* (Preist et al., 2001) evaluate agent strategies operating in multiple concurrently running auctions to create a more efficient market. However, they do not consider future auctions and hence the decision procedure is myopic. Also, they assume that a fixed number of items need to be bought. Byde *et al.* present a theoretical treatment of bidding strategies for bidding in multiple, heterogeneous single-unit auctions mechanisms (Byde et al., 2002).

Fujishama (K.Fujishama et al., 1999), Rothkopf (H.M.Rothkopf et al., 1998), Sandholm (Sandholm, 1999) and Varian (Varian, 1995) have researched auction mechanisms that assist bidders in bidding for combinations of items. These mechanisms are very useful, because bidders need not care about winning one item but losing the other when they want both the items but either item is worthless on its own. As a result, the bidders might get more profit at the auctions than if they did not use such combinatorial mechanisms. However, the determination of an optimal winner is an NP-complete problem, although approximations have been proposed by Fujishama and Sandholm. Moreover, actual auctions are held at various sites by various authorities, and it is hard to integrate bidding for auctions at such a variety of sites. Boutilier et al. (C.Boutilier et al., 1999) and Matsumoto et al. (Matsumoto and Fujita, 2001) has developed a model where agents bid for required resources sequentially, i.e., auctions are totally ordered in time. Boutilier et al. provides a dynamic programming model for agents to compute *bidding policies* based on the estimated distributions over prices. However, their work is built around sealed-price auctions, and so their algorithm determines bid price by which the expected profit will be maximized. Matsumoto et al., on the other hand, has examined open-price ascending auctions, such as English auctions. Their work, however, does not consider the more general setting of overlapping or asynchronous auctions. Shehory (Shehory, 2001) incorporates the degree of risk (e.g., how risk averse) an agent is willing to take in bidding for the item. Although he deals with asynchronous auctions, only single item but not bundles are covered.

7. Future Work

The goal of the current work was to develop a bidding strategy that can utilize knowledge of valuations of other bidders in future auctions to enhance buyer's utility when purchasing multiple units of an item from several auctions. We have developed a strategic bidding agent that utilizes a user valuation function for different quantities of an item, and the knowledge of valuations of other buyers in auctions to be held in the near future. In particular, we have studied strategic agents with knowledge of different time horizons purchasing units of a given item from several multiple-unit English auctions held over successive days.

We demonstrate that strategic agents with longer lookahead perform better not only when competing against non-strategic agents, but also when their expectations of closing prices are violated because of the presence of another strategic agent in the marketplace. Moreover, an increase of the competition does not necessarily implies a significant decrease buyers'gain using different strategies.

We also evaluated various risk attitudes in the above context and found that risk neutral agents perform better than risk seeking or risk averse agents when there is a single strategic agent in the marketplace. When two strategic agents compete in the presence of non-strategic agents, the risk neutral agents dominate. Also, risk seeking agents gain from competing with risk averse agents.

However, all these results are based upon one experimental setting. We believe that the observations hold in more general settings. We need to test our results by performing the same experiments in a wide range of settings. Also, an extension of the experimentation will be to test the effect of more strategic agent on each others utility. Though it is likely that their performances will suffer, it will be interesting to note the rate and nature of degradation.

In the current work, an agent uses the expected valuation of other bidders as the closing price in a future auction. This expected valuation is different for agents with different risk seeking attitudes, but we believe that a more effective decision mechanism needs to use the entire probability distribution over possible closing prices or buyer valuations. We plan to investigate this modification to our work in the future.

An interesting variation is the use of confidence estimate an agent has for expectations of auction closing prices in the near future, e.g., next day, compared to farther in the future.

Acknowledgments: We would like to thank Grégoire Richard for his help in the design and implementation in the early stage of this work.

This work has been supported in part by an NSF CAREER award IIS-9702672 and IIS-0209208.

References

Byde, A., C. Preist, and N. R. Jennings: 2002, 'Decision Procedures for Multiple Auctions'. In: *First International Joint Conference on Autonomous Agents and Multi-Agent Systems, Bolgna, Italy.* pp. 613–620.

C.Boutilier, M.Goldszmidt, and B.Sabata: 1999, 'Sequential auctions for the allocation of resources with complementarities'. In: *Proceedings of the Sixteenth International Conference on Artificial Intelligence.* pp. 527–534.

eBay, 'eBay'. URL: http://www.ebay.com/.

EBay.Buyer Guide, 'E Bay.Buyer Guide'. http://pages.ebay.com/help/buyerguide/index.html.

Garcia, P.Giminez, E.Godo, and J. Rodrigues-Aguilar: 1998, 'Possibilistic-based design of bidding strategies in Electronic Auctions'. In: *Proceedings of the 13th Biennial European Conference on Artificial Intelligence*.

H.M.Rothkopf, A.Pekeficc, and M.R.Harstad: 1998, 'Computationally manageable combinatorial auctions'. *Management Science* **44**(8), 1131–1147.

Hoos, H. H. and C. Boutilier: 2000, 'Solving Combinatorial Auctions Using Stochastic Local Search'. In: *AAAI/IAAI*. pp. 22–29.

K.Fujishama, Y.Leyton-Brown, and Y.Shoham: 1999, 'Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches'. In: *Proceedings of the Sixteenth International Conference on Artificial Intelligence*. pp. 548–553.

Klemperer, P.: 1999, 'Auction Theory: A guide to the Literature'. *Journal of Economic Surveys*.

Lucking-Reiley, D.: 2000, 'Auctions on the Internet: What's Being Auctioned, and How?'. *The Journal of Industrial Economics* **48**(3), 227–252.

Matsumoto, Y. and S. Fujita: 2001, 'An Auction Agent for Bidding on Combination of Items'. In: *Proceedings of the 5th International Conference on autonomous Agents*. pp. 551–559.

Milgrom, P.: 1989, 'Auctions and Bidding: A primer'. *Journal of Economic Perspectives* **3**(3), 3–22.

Monderer, D. and M. Tennenholtz: 2000, 'Optimal auctions revisited'. *Artificial Intelligence* **120**(1), 29–42.

N.Vulkan and C.Priest: 1999, 'Automated Trading in Agents-based Markets for Communication Bandwidth'. In: *Proceedings of UKMAS*.

Parkes, D. C.: 1999, 'iBundle: An Efficient Ascending Price Bundle Auction'. In: *Proceedings of the First ACM Conference on Electronic Commerce*. pp. 148–157.

Preist, C., A. Byde, and claudio Bartolini: 2001, 'Economic Dynamics of Agents in Multiple Auctions'. In: *Proceedings of the 5th International Conference on autonomous Agents*. pp. 545–551.

P.Wurman, M. Wellman, and W. Walsh: 1998, 'The Michigan Internet AuctionBot: A Configurable Auction Server for Human and Software Agents'. In: *Proceedings of the 2nd International Conference on Autonomous Agents*.

Rodriquez-Aguilar, J.Noriega, C.Sierra, and J.Padget: 1997, 'Fm96.5: A Java-based electronic auction house.'. In: *Proceedings of the Second International Conference on the Practical Application of Intelligent Agents And Multi-Agent Systems*. pp. 207–224.

Sandholm, T.: 1999, 'An algorithm for optimal winner determination in combinatorial auctions.'. In: *Proceedings of the Sixteenth International Conference on Artificial Intelligence*. pp. 542–547.

S.Gjerstad and J.Dickhaut: 1998, 'Price formation in double auctions'. *Games and economic behaviour* **22**, 1–29.

Shehory, O.: 2001, 'Optimality and risk in purchase from multiple auctions'. In: *CIA (Cooperative Information Agents) Workshop*.

S.Park, E.Durfee, and W.Birmingham: 1998, 'Emergent Properties of a Market-Based Digital Library with Strategic agents.'. In: *Proceedings of the International Conference on Multi Agent Systems*.

S.Park, E.Durfee, and W.Birmingham: 1999, 'An Adaptive Bidding Strategy based on Stochastic Modelling'. In: *Proceedings of the 3rd International Conference on Autonomous Agents*.

Tesauro, G. J. and J. Bredin: 2002, 'Strategic Sequential Bidding in Auctions using Dynamic Programming'. In: *First International Joint Conference on Autonomous Agents and Multi-Agent Systems, Bologna, Italy*. pp. 591–598.

Varian, H.: 1995, 'Mechanism Design for Computerized Agents'. In: *USENIX Workshop on Electronic Commerce*.

Yokoo, M., Y. Sakurai, and S. Matsubara: 2001, 'Robust Combinatorial Auction Protocol against False-name Bids'. *Artificial Intelligence Journal (to appear).*