

Genetic Algorithms versus Human Bidding Strategies for Auctions



Asunción Mochón¹, David Quintana², Pedro Isasi² and Yago Sáez²

amochon@cee.uned.es, dquintan@inf.uc3m.es, {yago.saez|pedro.isasi}@uc3m.es

¹Department of Applied Economics, UNED, ²Department of Artificial Intelligence, University Carlos III of Madrid,

Abstract:

The increasing use of auctions has led to a growing interest in the subject. The objective of this paper is to develop a genetic algorithm to find optimum bidding strategies for Ausubel auctions. The algorithm provides the strategy (defined as the action to be taken under different market conditions) that maximizes pay off to the bidder. The algorithm is tested under a series of scenarios that differ in their composition of utility functions, number of bidders and quantity of lots auctioned. The results suggest that the approach leads to strategies that outperform sincere bidding systems¹.

Keywords: Genetic Algorithms, Ascending multi-unit auctions, bidding strategy

1. Introduction

Nowadays auctions have become an important method of buying and selling a great variety of assets in different markets around the world. These assets range from the wide variety of things being auctioned in eBay, the selling of companies by governments, the selling of mobile telephone licenses [1] and the selling of emissions permits [2], etc. As a result auction theory and experimental examinations of these theories are of growing interest. A recent method used for carrying out these examinations has been the performing of “computational experiments” on systems of Artificial Adaptive Agents (AAA). The analysis of these systems gives us new approaches to the understanding of the economic and social behavior of auctions.

One way to studying adaptive learning model behavior with regard to AAA’s is through the use of Genetic Algorithms (GAs). GAs were developed in Holland [3] as a robust method of adaptive searching, learning and optimization in complex problem domains. This method has already been implemented in research that in-

¹ This article has been financed by the Spanish founded research MCyT project TRACER, Ref: TIC2002-04498-C05-04

cludes auctions. For instance, Andreoni and Miller [4] created and analyzed a GA model to capture the bidding patterns evident among human subjects in experimental auctions and also compared the different results for first and second price auctions formats. Moreover, Anthony and Jennings [5, 6] employed a GA model to search for effective strategies with multiple auctions (English, Dutch and Vickrey). For further details about different types of auctions and auction theory see [7,8,9,10].

In this paper we are going to use GAs to find whether a canonical behaviour (Sincere Biding) in Ausubel auctions can be outperformed, from the bidder point of view, in some specific conditions. The work also tries to find out whether better strategies can be achieved automatically, by means of GAs.

The remainder of the article is structured in the following manner: A definition of the auction model selected and the equilibrium bidding strategies are given in section 2. In section 3 we described the fundamentals of the bidding algorithm and various market settings that affect its performance. Section 4 describes and evaluates the experimental results of the GA in several environments and compares its' pay off with that of strategies used by other bidders. In section 5 we present our conclusions.

2. Bidding strategy framework

In multiple-object environments where individual bidders may demand more than one object there are several types of auctions. The most common multi-unit auctions are the standard ascending-bid auctions, the pay-your-bid auctions, the uniform-price auctions and the Vickrey auctions.

This paper focuses on a particular specification of rules for what might be referred to generically as an “ascending-clock auction”. In the standard ascending-clock auction (where demands are required to be non incremental in price) the auctioneer operates a continuously ascending clock and the price goes up. For each price, p , each bidder i simultaneously indicates the quantity $q_i(p)$ she desires. When the price p^* is reached, such that aggregate demand no longer exceeds supply, the auction is over. Each bidder i is then assigned the quantity $q_i(p^*)$ and is charged a unit price of p^* . The description of the alternative ascending-bid auction developed by Ausubel [11] coincides with the standard specifications of this type of auction but what is not specified however are the payments owed to each of the winners. As we have mentioned, when the auction is over each bidder i is assigned the quantity $q_i(p^*)$ but is charged the standing prices at which the respective objects were “clinched”. The bidder i clinches an object when the total demand of all bidders excluding i is lower than the total supply (but the total supply is covered by the total demand). In this way the bidder is guaranteed the object irrespective of what the other bidders do.

For a better understanding of this mechanism, table 1 includes an example of the alternative ascending-bid auction where we suppose that there are five identical licenses to sell. Each bidder has taste for more than one license, but bidders are each limited to winning at most three licenses. There are six bidders and their marginal values are given as follows (where numbers are expressed in millions of dollars): Bidder A: $v_{A,1} = 123$, $v_{A,2} = 113$, $v_{A,3} = 103$; Bidder B: $v_{B,1} = 75$, $v_{B,2} = 5$, $v_{B,3} = 3$; Bidder C: $v_{C,1} = 125$, $v_{C,2} = 125$, $v_{C,3} = 49$; Bidder D: $v_{D,1} = 85$, $v_{D,2} = 65$, $v_{D,3} = 7$; Bidder E: $v_{E,1} = 45$, $v_{E,2} = 25$, $v_{E,3} = 5$; Bidder F: $v_{F,1} = 49$, $v_{F,2} = 9$, $v_{F,3} = 3$. These are marginal values for a first, second and third license, respectively.

Price exceeded (millions)	Quantities demanded by bidders						Aggregate demand	Clinching of licenses
	A	B	C	D	E	F		
10	3	1	3	2	2	1	12	
25	3	1	3	2	1	1	11	
45	3	1	3	2	0	1	10	
49	3	1	2	2	0	0	8	
65	3	1	2	1	0	0	7	A "clinches" a license
75	3	0	2	1	0	0	6	A & C "clinch" licenses
85	3	0	2	0	0	0	5	A & C "clinch" licenses
Source: Ausubel [11]								

Table 1. Progression of bidding in alternative ascending-bid auction for example one.

In this auction each bidder is going to bid according to their real valuation of the licenses. This means that bidders will demand licenses as long as their marginal value is lower than the current price. Therefore, for the initial price, 10, the aggregate demand is 12, which is bigger than the available supply, 5, so the auction must proceed further. The auctioneer begins to continuously increment the clock and price and the first important change occurs when the price reaches 65 millions. The aggregate demand at this price is 7 and the supply is 5, so bidder A has now mathematically guaranteed herself at least one license as the aggregate demand of all competitors other than bidder A has dropped to 4. Bidder A has clinched winning a license. The auction continues and as the prices goes up bidders A and C clinched one more license. The final outcome is that bidder A wins three licenses for 65, 75 and 85 millions of dollars each and bidder C wins two licenses for 75 and 85 millions of dollars each.

In each auction the seller offers M number of indivisible units of a homogeneous good to n number of bidders. Each bidder i obtains a marginal utility of $V_{i,k}$ for the k th unit of the good, for $k=1, \dots, M$. Thus if bidder i gets q_i units of the good for a total payment of P_i , he obtains a payoff of:

$$\sum_{k=1}^{q_i} V_{i,k} \leq P_i, \text{ for } i=1, \dots, n \text{ and } q_i=1, \dots, M. \quad (1)$$

All bidders are assumed to exhibit a (slightly) diminishing marginal utility, i.e., $V_{i,k} \geq V_{i,k+1} \geq 0$, for all $i=1, \dots, n$ and $k=1, \dots, M-1$. In this research each bidder has a different utility function which is generated randomly. We assume that all bidders have pure private values in the sense that the estimation of her own value made by bidder i is not affected by bidder j 's estimation.

In our model we have included one possible specification of bidders' information which is that all bidders have full bid information. This means that they have the complete history of all bids made by all bidders in the auction. According to Ausubel [11], in an alternative ascending-bid auction with pure private values, weakly diminishing marginal utilities and full or partial bid information sincere bidding is an efficient equilibrium.

In our model all but one of the bidders fulfill the Sincere Bidding (SB) hypothesis as it is an efficient equilibrium strategy, their fixed bidding strategy would be to demand a license until the price reaches their marginal value. There is only one bidder whose strategy will depend on the evolution of the auction. The bidder using our GA that learns the best strategy to implement from one auction to another and that indicates how the bidder should bid according to its results.

It is possible that for a certain increase of price the decrease of the demand implies that supply is not covered. In these circumstances a rationing rule is introduced. We turn to the previous price and allocate all the units to the bidders. We then calculate the excess demand and, beginning with the bidder who has the highest quantity, start to remove one unit from each until the total demand equals supply. We will not be able to remove units from bidders that clinched their units in previous rounds.

3. Designing the GA for the bidding strategy

The aim of the GA is to find a set of bidding rules that maximize the profit of a bidder according to equation (1) for a specific scenario. Each scenario is defined by the number of items to be auctioned, the initial price and the set of bidders. Each bidder has a different utility curve and a specific strategy. It is under these circumstances that the resulting strategy would be valid.

The sequence of events in our experiments will be to:

- Specify a scenario.
- Pick one of the bidders and let their strategy evolve from the initial SB.

- Compare the outcome of applying the GA strategy with the result that would have been obtained using SB.

We will therefore start providing details regarding the scenario settings and will then introduce the suggested strategy specification approach before finally discussing some issues related to the GA.

3.1. Scenario settings

All the scenario settings are defined as external variables except for the utility functions which are generated by the algorithms described in this section.

All the participants have a fixed strategy: SB. Therefore, they only compare the actual price with their utility function. The bids of the bidder who follows the GA strategy will depend on several actions that we will define as deviations from the SB strategy. Hence, our first step is to generate a utility function for every participant.

Each utility function consists of a set of values organized from high to low that specify the utility that is expected from the consumption of each unit. We force each bidder to define the utility for at least as many items as the total supply quantity. As we mentioned previously, the utility of each unit must always be equal to or lower than the previous one. Since the results might be affected by the elasticity of the utility curves we tested two elasticity function building algorithms. One of them results in very elastic demand curves and the other generates inelastic curves.

The first algorithm starts with a random number between 0 and 200 which will be assigned to the utility of the first unit. For each additional unit we reset the upper band by lowering it at a fixed rate equal to 200 divided by the total number of units auctioned multiplied by two. The utility for each unit is the minimum of a random number generated under the new limit and the utility of the preceding unit.

The second method starts just like the previous one. Once we have the utility of the first unit we generate a random number between 0 and that utility. The value assigned to the next unit would be the maximum of that random number and the utility of the previous unit divided by two.

3.2. Bidding strategy

The definition of the bidding strategy that we suggest requires the identification of actions linked to specific auction conditions. Each action is defined in terms of deviations from the SB strategy which demanded quantity is represented by q . To this end we consider the following four actions: bid half of the SB; bid the SB; bid 50% more of the SB or bid twice as much as the SB. All these strategies have an

upper band that is the lowest of either the number of units being auctioned or, alternatively, the units demanded in the previous round. The lower band is the number of units that the participant has already clinched. Bidders face 61 possible states of the auction. One of them represents the initial market condition where no relevant information is available. The rest are made up from the combination of the potential values of three different indicators which can be described as follows (see table 2).

$\varepsilon_{p,d}$		AB(%)		OM			
$\geq 1,5$	+	$< 100\%$	+	≥ 2	+	Initial	61
$< 1,5$		$< 50\%$		< 2			
1		$< 25\%$		$< 1,25$			
< 1				< 1			
$< 0,5$							

Table 2. Bidders face 61 possible states of the auction made up from the combination of these three different indicators.

The first one is the trend of price elasticity of demand from the competitors ($\varepsilon_{p,d}$). This is the evolution of the competing bidders sensitivity of demand according to price increases. Formally, the definition of price elasticity responds to formula (2).

$$\varepsilon_{p,q} = \frac{\frac{q^{l-1}}{q^{l-2}}}{\frac{p^{l-1}}{p^{l-2}}} \quad (2)$$

Where p^{l-1} denotes the price at the $(l-1)$ occasion on which one or more of the bidders strictly decreased the quantity demanded, and q^{l-1} denotes the sum of the quantities demanded by competing bidders $1, \dots, n$ beginning at the time. In the same way, the indexes $(l-2)$ refer to the previous round.

Since we are concerned with the evolution of this figure the magnitude that we consider is the ratio of the mentioned elasticity divided by the same elasticity in which we consider $(l-2)$ and $(l-3)$ data. These ratios are classified into five categories: ≥ 1.5 ; < 1.5 ; $= 1$; < 1 and < 0.5 .

The second indicator is the percentage of active bidders (AB%). This is the number of bidders that ask for at least a unit over the number of bidders at the beginning of the auction. We have classified this indicator in three categories: $< 100\%$; $< 50\%$ and $< 25\%$. The third indicator is the operating margin of the strategic bidder (OM). This variable is defined as the ratio of the total number of units that no one clinched in the previous round over the number of units demanded by the bidder using the GA strategy in that previous round. This indicator is also classified into four categories: ≥ 2 ; < 2 ; < 1.25 and < 1 .

3.3. Genetic Algorithm

The translation of strategies to individuals is direct. As we mentioned before, all strategies consist of a set of 61 integers valued from 0 to 3, depending on the action to be taken which in turn depends on the state of the market. Since the actions can be encoded in 2 bits each strategy could be represented as an array of 122 bits. This includes the SB strategy. This strategy means that under all circumstances the appropriate action is to bid for the amount that doesn't deviate from the SB, which is derived straight from the utility function. Since we define the action of following the SB strategy as 1 (01), the strategy of a SB bidder would be represented as a string of 122 bits (where the value of every odd bit is 0 and 1 otherwise).

The assessment of a strategy is made by running an auction against competing bidders and evaluating the payoff as defined in (1). The bigger the sum of the difference between the maximum price the bidder would potentially pay (utility) and the actual price paid for each item, the better it is. In this process it is assumed that the rest of the bidders do not evolve and have a fixed strategy that makes them use SB at all times.

The search for the best strategy for each scenario was performed using populations of 30 individuals that evolved for 1,000 generations. The GA used elitism, a roulette wheel selection mechanism in addition to mutation and crossover probabilities of 0.9 and 0.01 respectively. In general, the convergence of the population was fast and the best individual was found within the first 10 generations, something that was favoured by the fact that most market situations are not applicable to a specific scenario. Since only the value of a few bits needs to be specified the search space is not as big as it would be expected.

In addition to this there is another factor that contributes to these results, the fact that under some scenarios several actions might be equivalent. To illustrate this point we may consider a GA bidder whose utility function specifies utilities that are higher than the initial price for every unit. In this case, three out of four actions in the first price model, scenario 61 (no information), would lead to the same results. Given the rule that establishes that demand is limited by the number of units and the fact that SB is a strategy that would suggest bidding for all the units, SB, SB+50% or 2*SB actions would result in a bid for the same amount.

4. Evaluating the bidding strategy

In this section we evaluate the outcomes obtained by the bidder who acts according to the GA described in the previous section. We have run the experiments with two different types of utility functions, the first being referred to as elastic and the second as non elastic. In the first case, the quantities demanded by the bidders decrease quite quickly as the price goes up. In the second case, the quantities de-

manded by the bidders don't decrease significantly as price goes up. For each type of utility function we have calculated the outcome for the bidder who acts according to the GA strategy. Moreover, we compare this outcome with the pay off they would have obtained with the equilibrium strategy (the sincere bidding method). Experiments were conducted separately for different numbers of bidders ($n=4, 6, 8$); different numbers of objects auctioned ($k=10, 15, 20$, and 25); and different starting prices ($p^0=1, 10$). For each environment we ran the algorithm 25 times thus letting a population of 30 individuals evolve for 1000 generations.

The results of the experiments show that the GA strategy performs at least as well as the SB strategy and, in some cases, even better. We have started to analyze the results by doing a descriptive statistic analysis of the pay off obtained with the two different bidding strategies. The average for the SB method is 138.79 and for the GA method is 153.88. An analysis of variance (ANOVA) was done to test the hypothesis regarding the differences between this average and shows that the differences are statistically significant ($F = 8.806, p = 0.03$). As a result, it is demonstrated that at least in some cases the pay off to the bidders can be improved by using our GA strategy rather the SB one.

The next step in our evaluation is to find out in which particular environments our GA can mean an improvement on the SB strategy. To this end a new variable has been created: the difference between the pay off generated by the GA strategy and the SB strategy. Analyzing this variable gives us several results.

Firstly, the analysis of this variable shows that differences between the average for the elastic and non elastic utility functions are statistically significant ($F = 4.191, p = 0.041$). As they have a positive significant correlation ($r = 0.059, p = 0.041$) we can conclude that the GA solution is better when the bidders have elastic utility functions in which the private valuations of the bidders are more sensitive to the evolution of the price. This means that there is a greater probability of finding an alternative bidding strategy than in the SB method that maximized the income of the bidder.

Secondly, this variable also has a positive significant correlation with the minimum starting price ($r = 0.124, p = 0.00$) which implies that the bigger the minimum price is the better the GA model works (with respect to the SB model). This result is due to the fact that as the minimum price goes up the probability that the total demand won't cover the supply increases. Under these circumstances, in which the available quantity was not fully subscribed at the initial price, the auction has failed and therefore the probabilities of getting a higher pay off with the GA strategy increases.

The third remark concerns the relationship of this variable with respect to the number of bidders and the number of objects auctioned. The difference between the pay off generated by the GA strategy and the SB strategy has an inverse significant correlation with the number of bidders ($r = -0.217, p = 0.00$) and a posi-

tive significant correlation with the number of objects auctioned ($r = 0.207$, $p = 0.00$). Figure 1.A shows the value of the difference between the pay off generated by both strategies and the number of bidders included in each experiment. As can be seen our algorithm performs better when there are less bidders.

Nevertheless, when we analyze the relationship between this variable and the number of objects auctioned (Figure 1.B) we find that our GA strategy works better as the number of objects auctioned increases. Consequently, when there are few objects to auction the probability of finding a better bidding strategy than the SB is reduced. In these environments the allocations of the objects depends more on the private valuations of the bidders than on their possible alternative strategies.

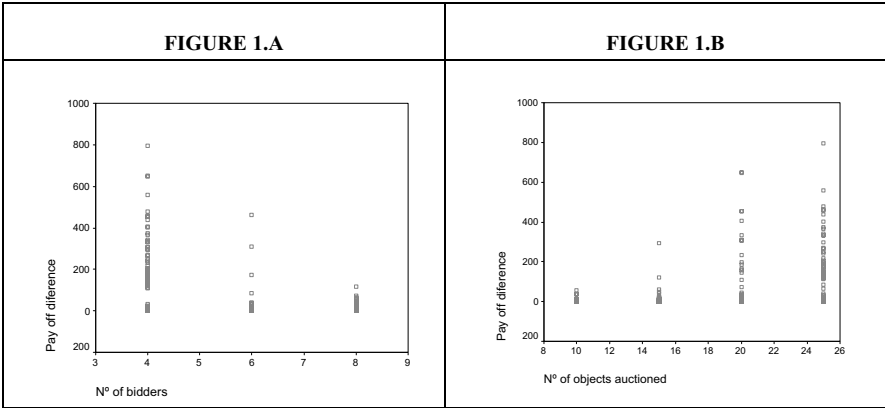


Fig. 1. Figures 1.A and 1.B shows the differences between the payoff generated by the GA strategy and the SB strategy by the number of bidders (1.A) and the number of objects auctioned (1.B)

5. Conclusions

This research presents a GA strategy that can be successfully employed to evolve bidding strategies in alternative ascending-bid auctions that use pure private valuations. The algorithm generates different strategies according to the market setting which are defined by three indicators: tendency towards elasticity in price demand, the number of bidders active in the auction with respect to the total number of bidders at the beginning of the auction and the number of objects that have not been clinched with respect to the number of objects demanded by the bidder.

By relating the GA strategy to a bidder a successful participation will be guaranteed in particular environments (in all cases the performance of the GA model is

equal to or better than the SB one). We have run the experiments with two different types of utility functions: elastic and non elastic. The experiments were conducted separately for different numbers of bidders ($n=4, 6, 8$); different numbers of objects auctioned ($k=10, 15, 20$, and 25); and different starting prices ($p^0=1, 10$). For each environment we ran the algorithm 25 times which allowed a population of 30 individuals to evolve for 1000 generations.

In the evaluation of the GA strategy it was discovered that it works better when the utility functions are more elastic, when the starting price increases, when the number of bidders decreases and when the quantity of objects auctions is bigger. Once we have analyzed these preliminary results it seems to happen that with the GA strategy the outcome of the bidder is better because it push the auction to a situation of rationing. If this situation is reached, the bidder with the GA strategy has more options to improve the SB outcome. Given these findings an alternative line of future enquiry could be to search for common patterns in different environments and try to work out general behaviours.

References

- [1] Klemperer, P. (1999) "Auction Theory: a Guide to Literature" *Journal of Economics Survey*, vol. 13, 3, pp 227-286.
- [2] Sunnevag, K. J. (2001) "Auction design for the allocation of emission permits", *Working Paper*, UCSB.
- [3] Holland, J. (1975) "Adaptation in Natural and Artificial Systems," *Ann Arbor*., University of Michigan Press.
- [4] Andreoni J., Miller J.H. (1995) "Auctions with Artificial Adaptive Agents," *Games and Economic Behavior*, Vol.10, 39-64.
- [5] Anthony, P., Jennings, N.R. (2002) "Evolving Bidding Strategies for Multiple Auctions," *Proceedings of the 15th European Conference on Artificial Intelligence*. 178-182.
- [6] Anthony, P., Jennings, N.R. (2003) "Developing a Bidding Agent for Multiple Heterogeneous Auctions," *ACM Transactions on Internet Technology*. Vol. 3, No 3.185-217.
- [7] Klemperer, P. (2000): "The Economic Theory of Auctions. Volume I". Cheltenham, UK: Edward Elgar.
- [8] Klemperer, P. (2000): "The Economic Theory of Auctions. Volume II". Cheltenham, UK: Edward Elgar.
- [9] Klemperer, P. (2003) "Auctions: Theory and Practice," Princeton University Press.
- [10] Milgrom, P. (2004) "Putting Auction Thoery to Work," Cambridge University Press.
- [11] Ausubel, L.M. (1997) "An Efficient Ascending-Bid Auction for Multiple Objects," University of Maryland. Working Paper No. 97-06.