Comparing the Behavior of Agents to Human Subjects in a Uniform Price Auction

HyungSeon Oh[1]   Robert J. Thomas[1]   Bernie Leiseutre[1]   Timothy D. Mount[2]
hso1@cornell.edu   rjt1@cornell.edu   bcl25@cornell.edu   tdm2@cornell.edu

1. School of Electrical and Computer Engineering, Cornell University, Ithaca, NY 14853
2. Applied Economics and Management, Cornell University, Ithaca, NY 14853

Abstract

The idea that large-scale generating units will operate at marginal cost when given the ability to offer their power for sale in a uniform price auction is at best wishful thinking. In fact, both real and experimental data show that the more uncertainty a supplier faces (e.g., load uncertainty, uncertainty of other suppliers, etc.) the more they will try to increase their profits by submitting offers to sell higher than marginal cost and by withholding units if permitted. This makes predicting unit commitment and dispatch ahead of time difficult. This paper explores characteristics of software agents that were designed based on the outcome of tests with human subjects using a uniform price auction with stochastic load. The agent behavior is compared to the behavior of the subjects. Both subject and agent behavior is classified based on the data. Differences and similarities are noted and explained.

I. Introduction

For the last decade, deregulation in the electric power market has been taking place in many countries. In many deregulated markets, an auction plays a major role in determining the price for electricity, using an online auction over the internet. The auction-based market is thought to be more economically efficient than a traditional regulated market due to the interaction and easy access of different suppliers to the market. For developing a tool for power system planning, it is necessary to simulate a web-based auction in which human agents participate. However, participation of human agents in the market needs a lot of time and cost. A well-designed software agent can be a substitute to emulate the offer behavior of human agents. For a simulation, a limited number of different software agents participate, and each of the agents represent a firm that owns several generators. Thus, it is necessary to select agents for a simulation from all the software agents designed. Since it takes too much time to test all the possible permutations of all the possible software agents, a way is needed to classify the agents into a small number of groups based on the effect of the agents on the market outcomes. Several different types of strategy used by human agents have been observed in tests of markets using a uniform price auction. The most extreme strategies are to offer marginal cost and to speculate. For the sake of simplicity, earnings can be determined by the market clearing price and the quantity dispatched. A marginal cost offer agent wants to maximize the quantity dispatched by offering low and hopes that someone else sets a high price. A speculator wants to increase the market clearing price, and takes the risk of not getting as much capacity dispatched. It is fairly easy to model a marginal cost offer agent since it offers all the blocks at the marginal cost. On the other hand, there are many different types of speculators depending on the degree that they speculate.

In this study, five standardized agents were designed for simulation and classification - four different types of speculators and a marginal cost offer agent. A human subject and a software agent competed against combinations of the standardized agents. Based on their performance (earnings in each period), the subjects and the agents were classified into the five different groups

II. Electricity market

Agents develop auction rules for themselves based on the rules of the auction they are participating in and, in repeated auctions, based on the actions of their competitors. In the design considered here the electricity

market was assumed to be a uniform price auction with an inelastic but time varying load demand. In this market an independent system operator (ISO) provides a load forecast and collects offers submitted by six participating agents. The ISO then clears the market and checks the security of the system.

In every period, each agent is asked to submit a price and quantity. No price can exceed a reservation price meant to represent the price above which no load would be willing to pay for power. The offers submitted by all the agents are then ranked according to the offer price from lowest to highest. Then, the ISO dispatched blocks beginning with the lowest offer until actual demand (which is different than forecasted demand) is met. If two or more blocks were offered at the same price, the ISO randomly selected which block(s) to be dispatched. All the winning agents were paid according to a second price auction, meaning that winners were paid at the same price (uniform price auction). If the actual demand were larger than the capacity offered, ISO would recall short of capacity from the blocks withheld at the price of the last accepted offer. The agent whose block was recalled would be charged a recall cost. After clearing the market, ISO published the market clearing price and quantity dispatched to corresponding agents. Each agent received information only related to its own generator such as the dispatch quantity and price. One scenario was comprised of 200 periods.

Six agents each had the same capacity with 5 blocks. Their generators had identical operating costs including fuel cost and standby cost as well as interest charges. For the sake of simplicity, startup costs were not taken into account. Based on its maximization algorithms, available history data and load forecast, each agent decided how many blocks to offer and the offer price of a block if offered. Exchange of information among agents was not allowed.

III. Standardized agents

Five standardized agents consisting of one marginal cost offer agent and four speculators were designed to be used in a test bed whose purpose is to a classify other software or human agents. That is, the thesis is that an agent with unknown behavior can be classified based on its play with known agent types. The marginal cost offer agent (MC) is an agent that offers all five blocks at marginal cost without any withholding. The four speculators had different degrees of speculation. In order to be a speculator, at least one block must be offered at a high price.

It is crucial to an ability to implement a speculator to be able to determine which block or blocks are to be offered at a high price. For simplicity, any offer submitted at a high price was made at the same price regardless of the

type of speculation. A fair share of the market was calculated based on the load forecast. The block in which the fair share quantity falls is termed the "fairshare block". If this were the last block chosen for the unit by the auction, then it would be the units' marginal block. Thus, the fair share calculation is just a means for trying to predict a unit's marginal block a period ahead and any calculation that accomplishes that prediction is suitable for the purpose we have in mind. Since all the competitors in the market considered here have the same capacity, fairshare was calculated simply by dividing the load forecast by the number of market participants. If there were differences in the generating capacity being represented by an agent, the formula for a fair share is more complicate. Also, if some agents have a locational benefit over others, their fairshare should not be a simple dividend of a forecast. In such a case, fairshare could be calculated in a following way. Suppose all the agents that have the same locational benefit submit offers in order for them to get dispatched in the same fraction, which is the ratio of quantity dispatched and the total capacity. For three speculators, only one block was offered at high price, and the blocks with a lower operating cost than the fairshare block were offered at the marginal cost. The blocks with a higher operating cost than the fairshare block were withheld from the market. There are several reasons why a speculator withholds its capacity from the market. First, a speculator may suspect that the withheld block will not be dispatched if offered. In such a case, the speculator may only pay a standby cost which results in decreasing profit. Another reason is that physically withholding capacity increases the chance that a high offer will need to be dispatched since load must be met. If standby costs are ignored, the effect of withholding is essentially that of submitting an offer higher than the reservation price.

The strategies for offers of the standardized agents is shown in Table 1.The standardized agent with the weakest degree of speculation, called a weak speculator (WS), was designed to speculate with the block that is adjacent to and more expensive than its fairshare block. If the load forecast had no significant error (i.e., if the forecast was similar to the actual demand), the behavior of WS was found to be similar to that of MC with some withholding capacity. Since no speculator could speculate less than WS, the agent was called weak speculator. The agent with stronger degree of speculation, strong speculator (SS), offered a high price for its fairshare block. This agent took the risk not being dispatched for a higher market clearing price. Two stronger speculators (SS2, SS3) were also implemented. One of them (SS2) offered at high price for the block before the fairshare block while the other did from the first to the fairshare block. Figure 1 shows the result of Table for units with 5 blocks.

Table 1. Offer strategies of the standardized agent when the fairshare block is the j^{th} block. Here MCO, S and W stand for the marginal cost offer, speculate and withhold, respectively

	base unit	$(j-1)^{th}$ block	j^{th} block	$(j+1)^{th}$ block	higher block
MC	MCO	MCO	MCO	MCO	MCO
WS	MCO	MCO	MCO	S	W
SS	MCO	MCO	S	W	W
SS2	MCO	S	W	W	W
SS3	S	S	S	W	W

Fig. 1. Schematic diagram showing the offer functions of the different types of standardized agents

IV. Classification of an agent

In a simulation five agents composed of some mix of the standardized agents and one agent of interest that was either a software design or a human agent were used. A specific combination of the five standardized agents composed one scenario. It turned out that only six different scenarios were needed for a classification. Each software and human agent participated in the chosen six scenarios at a time. For each scenario, the six agents participated in 200 periods, and their earnings were collected and plotted as a function the earning of the agent of interest at each period. Figure 2 shows one simplified plot of the earnings of all participating agents. The six lines show how the corresponding agents performed in each period. All the lines have different slopes, which characterizes the type of agent. Among the lines, the line showing $y = x$ represents the earning of the agent of interest. If the $y = x$ line is "close" to one of lines showing the earnings of a standardized agent, the agent of interest is classified as an agent whose behavior is similar to that of the standardized agent that produced the close line. For example, the agent shown in Figure 2 is classified as a strong speculator (SS). In the scenario that produced this plot, the MC (no speculating) agent earned the most while SS3 (the speculator with the strongest degree of speculation) earned the least. In simulations with software agents, this feature was found to be true in general. However, an agent with a less degree of speculation made the market more competitive and consequently made everyone including the agent itself earn less. This might encourage an agent to speculate if it wants to maximize its own profit without concern for the profits of others.
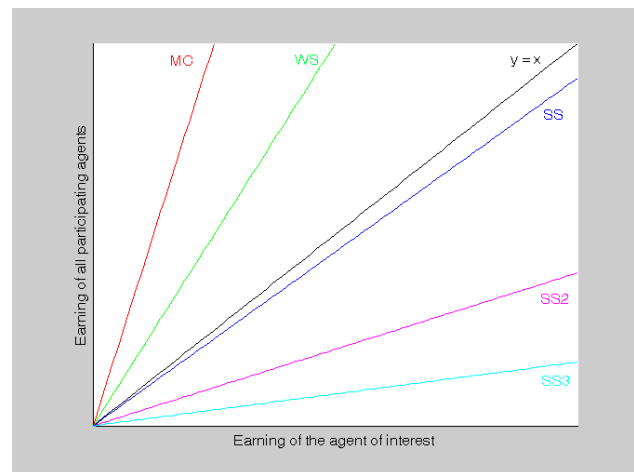

Fig. 2. Earnings of the standardized agents and the agent of interest

V. Expected earnings

It was assumed that the earning of a standardized agent was highly correlated to that of an actual agent of the same type. To calculate the earnings of six different types of agents, an electricity market was simulated with the standardized agents. Note that there was no individual software or human agent in this simulation. From the simulation, the earnings of participating agents were obtained for different types of agents.

Expected earnings of the software agents were calculated based on the actual distribution of the software agents once they were classified. After classification, one could calculate the earning of each agent from each scenario, and then multiply the earning by a weight factor. The weight factor could be calculated based on the probability that the agent might be in the same group in

agent competition as the competition where it earned the profit considered now. For example, suppose that there were 24 agents. Suppose we had classified them as 5 speculators and 19 marginal cost offer agents. Now, suppose we were interested in one of the speculators competing with five other agents from the group of 24. The following enumerate the choices: Number of possible choices when selecting 5 agents without regard to type from the 23 agents left in the pool (# of different choices) is:

$$_{23}C_5 \times {}_1C_1 = \frac{23!}{(23-5)!5!} \times \frac{1!}{(1-1)!1!} = 33,649 \,.$$

The number of choices that have no speculator in a group is 11,628 ($= {}_4C_0 \times {}_{19}C_5 \times {}_1C_1$). From similar calculations the possible number of choices can be calculated for other mixes of agents. The corresponding probabilities can also be calculated. For example, the probability that the agent of interest participates in a market with no speculator is 0.3456 (=11,628/33,649). The probabilities that the market has one, two, three and four speculators are 0.4608, 0.1728, 0.0203 and 5.65×10^{-4}, respectively. That is, the probability that all marginal agents are competing with the chosen speculative agent (i.e., there are no speculators in the competition other than the chosen speculative agent) is 0.3456. If, for example, the agent of interest earns $100, $300, $700, $1,800 and $2,500 in each of 5 competitions where each has a different mix of competing agents as listed above, then the weighted earning of agent k, E^k, is about

$$E^k = \sum_{\substack{i \in possible \\ group}} p_i^k \times e_i^k \approx \$332$$

where p_i^k and e_i^k stand for probability that agent k is in group i and the earnings for agent k is in the group i, respectively. The expected earnings obtained in this way were used for a further comparison of the actual earnings.

VI. Simulation results

In the fall 2002, fourteen different software agents were submitted by the students taking the class ECE 551/AEM655 at Cornell University. These agents were competed in a class competition and subsequently used as early tests of the classification ideas presented here. From experiments performed in the same class with the students, it was believed that MC, WS and SS were the most competitive types of agents. Therefore, only those types of standardized agents were used. After performing simulations in which all possible combinations of the three standardized agents were used, the classifications of each agent of interest by certain of those simulations were found to be redundant, i.e., classifications using one scenario and that by using another different one was identical. It was found that of the all the combinations of 3

agents choose 5 that are possible, only six were enough to produce distinctive classifications. The following scenarios were selected since they are a complete set for the classification in a consistent way:

4 WS + 1 MC, 3 WS + 2 MC, 4 SS + 1 MC, 3 SS + 2 MC, 1 SS + 2 WS + 2 MC and 1 SS + 1 WS + 3 MC.

One randomly selected set of the forecasted and actual load was assigned for one scenario. Average load was 470 MW, and the maximum error between forecast and actual load was 20 MW.

Each of the fourteen software agents and five standardized agents formed a group for the simulation, and corresponding plots were generated based on the results of the simulations. According to the plots, the fourteen agents were classified into three groups – five MC, 4 WS and 5 SS. It seemed that most agents speculated to some extent with the degree of speculation somewhere between WS and SS. It is worthwhile to note that the earnings from one scenario of the agent was close to that of the standardized agent of the same type. Figure 3 shows one example of the plots of the earning of a randomly selected software agent classified as SS. The classification of the software agents was fairly easy since a strategy used seemed consistent in a given scenario. For most of the agents, strategies seemed not to change for different scenarios, i.e., type of competitors. It was also found that no agents developed by the students used learning algorithms which would alter the results significantly.
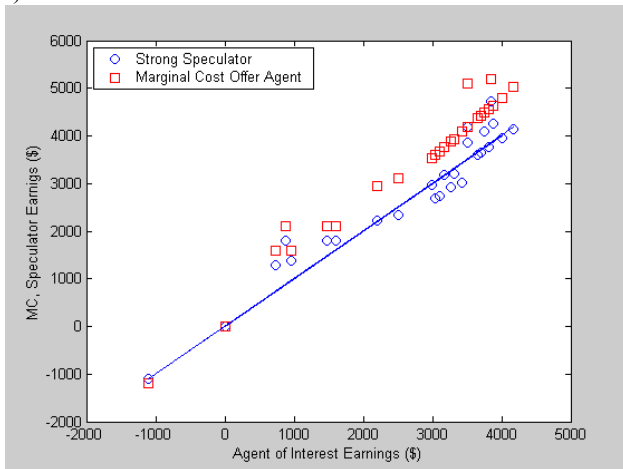
a) 4WS + 1 MC

b) 3 WS + 2 MC

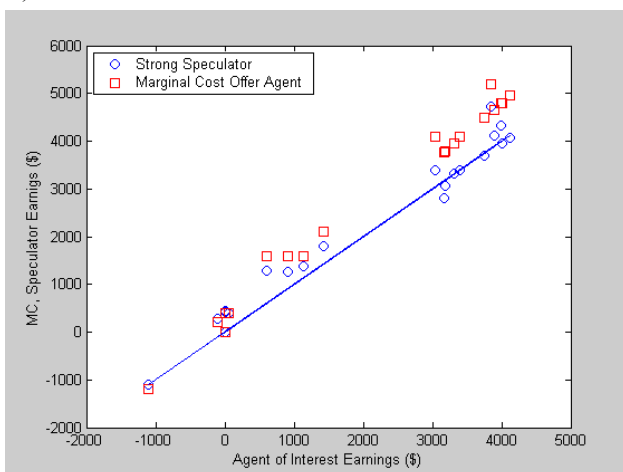c) 4 SS + 1 MC

d) 3 SS + 2 MC

Fig. 3. Example of a performance of the software agents: in the plot, red square, green and blue circle stand for the earning in a period of MC, WS and SS, respectively

For a simulation with a human agent, twenty students were recruited from ECE 451, electric power systems, in Cornell University. Each of twenty students participated in the simulation with five standardized agents just like the software agents. The purpose of this experiment was to find out if the same technique that was successful for classifying software agents could be used to determine human strategies. The same sets of the forecast and actual load were used for the simulation. They learned from experience, and were consistent only in some scenarios. Therefore, the data obtained only after a learning period were useful for the classification for the scenarios. After examining earning data, ten periods were assigned to the learning period. It was also found that one behaved SS in some scenarios while the same person did WS in other scenarios, i.e. different strategies were used for different types of competitors. Strategies other than ones used by the standardized agents were also observed. The conclusion was that the set of standardized agents was not rich enough and that it was possible to classify some of the different strategies by adding by the speculating agents SS2 and SS3 to the mix. A typical result of the simulation is shown in Figure 4.

In the case of a) and b), one was classified as SS2 and SS3 while the same one was classified WS and SS in the case of c) and d), respectively. When SS3, a standardized agent with the strongest degree of speculation, participated in a scenario described above, the plot b) was a common feature. What SS3 did in the market was effectively withholding its whole capacity from the market unless the market clearing price was high. Therefore the market clearing price was high even in low demand period, which lead the earning of all the competitors to increase a way high. Even though this type of strategy seemed not reasonable, it was often observed especially when the market was very competitive, i.e., a market with agent of less degree of speculation – the simulation with 3 WS + 2 MC in this study. For a little less competitive market such as 4 WS + 1 MC or 3 SS + 2 MC, the strategy was rarely used.

a) 4 WS + 1 MC

b) 3 WS + 2 MC

c) 4 SS + 1 MC

d) 3 SS + 2 MC

Fig. 4. Example of a performance of the human agents: red square, green and blue circle stand for the earning in a period of MC, WS and SS, respectively

For the case in which it was possible to classify a human agent, the total earning of a human agent from the scenario was compared to that of the standardized agent of the same type from the same scenario. The comparison between the two earnings was shown in Fig. 5. The red line corresponds the perfect correlation, which is y = x. In the Section V, it was assumed that standardized agent earnings were highly correlated to the earnings of actual agent of the same type. Fig. 5 shows the assumption was satisfied in the experiments performed in this study. The correlation between two earnings was checked for the both with a software agent and with a human agent as long as it was possible to classify the agent of interest.


Fig. 5. Actual earning vs. Expected earnings calculated from simulation

There was an interesting software agent worthy of special note. It offered some capacity into the market at marginal cost, but started to withhold some from the fairshare block. Therefore, its offer function was similar to that of SS except for withholding capacity from the fairshare block instead of offering at a high price. This offer behavior is known as Cournot type [1-2]. This agent was classified into SS as long as other speculator(s) exist(s) in the market regardless of type such as WS or SS. For a further investigation, other types (agents offering marginal cost with some withholding) were implemented. For example, an agent offered just like what WS did with withholding instead of submitting high offer. The agent was classified as WS in the same condition described before. It was concluded that the degree of speculation was closely related to which capacity an agent starts to deviate from the marginal cost (or low offer).

In an agent simulation, it was found in general that a higher earning for everyone was made, as the speculation got stronger. However, in a given scenario, the agent who earns most was the least speculating agent – MC, WS, SS, SS2 and SS3 in decreasing order. For the agent simulation,

the objective of each agent was the profit maximization. The best strategy of an agent to serve the objective depended on the scenario in which the agent was participated. Therefore, it is important to figure out the type of the competitors in the market.

In the class of ECE551/AEM655, a round robin type tournament was designed to determine a winner among submitted agents based on the earning. In the tournament, submitted agents were randomly divided into three groups of six agents. A group of six agents participated in a simulation. Based on the earning from the simulation, two agents from each group were selected for a final simulation. The winner was nominated from the final competition composed of two winning agents from each of three groups. The winner was classified by using the classification method, and turned out to be a type of MC. In such a competition, not many combinations were given to agent even though the group selection was random. Therefore, it is reasonable why the winner was the type of MC when one considers that the least speculating agent (MC) in a given scenario is the most rewarding agent.

When all possible combinations were given (complete search), the winner was a type of an SS. The method seems fair to all the agents, but it takes too much time because it needs to perform large number of simulations. For an alternative method, it was suggested that one should select only small number of agents, and then give all the combination for the selected agents. It is important how effectively and fairly one can select the small number of agents out of all the agents. Based on the expected earning, E^k, obtained by using the method described in the Section V, one can rank all the agents by assuming that the actual earning of agent k has a good correspondence with the expected earning, E^k. The rank is to be used for a selection of small number of agents. By using this method, ten agents were selected for the final competition. Eight out of ten selected agents were ranked in top ten from complete search method. The winner determined by this method was also turned out to be the winner from the complete search.

VII. Conclusions

In this paper, several simplifications have been made for the system for both a market and agent used here such as all equal marginal cost, equal capacity, no startup cost and no line constraints. From a market simulation with those simplifications, offer strategies under a uniform price auction are classified. Under the auction rule, the last accepted block determines the market clearing price, second price auction. The earning is approximately determined by the quantity dispatched as well as the market clearing price. To maximize earning, a software and a human agent choose several different strategies. Each strategy produces different offer function. The main

results of this paper describe how to classify the strategy not by inspecting individual offer function but by comparing the result of simulation with its competitors. Different types of agent can be characterized by their degrees of speculation. The degree of speculation is closely related to where its offer function deviates from the low offer or marginal cost offer. This paper also shows that only a small number of standardized agents can be used for the classification, and their earnings have a good correspondence with the earning of an actual agent.

VIII. Future works

In this study, the market setup was simplified for convenience and to handle the problems discussed here. In a real market, there are many constraints that one must satisfy. Some of the constraints are closely related to the locational benefit. One possible way to implement those constraints is to formulate a proper equation for a fairshare. Another big simplification was a discrete offer submitted by an agent, i.e., an agent was not allowed to change a quantity of each block. The discreteness may restrict the behavior of an agent. It is desirable to allow an agent to decide quantity as well as offer price to optimize its profit. Both studies are undergoing.

Acknowledgment

References

[1]. A. Cournot, Researches Into the Mathematical Principles of the Theory of Wealth, New York: Macmillan (1897).

[2]. A. Daughety, Cournot oligopoly: characterization and applications, Cambridge University Press, New York (1988).