Software Agents for Market Design and Analysis

Final Project Report

Power Systems Engineering Research Center

A National Science Foundation
Industry/University Cooperative Research Center
since 1996

Power Systems Engineering Research Center

Software Agents for Market Design and Analysis

Project Team

Sarosh Talukdar, Project Leader
Carnegie Mellon University

Tim Mount, Robert J. Thomas
Cornell University

Shmuel Oren
University of California at Berkeley

PSERC Publication 05-37

June 2005

Information about this project

For information about this project contact:

Sarosh N. Talukdar
Professor
Electrical and Computer Engineering
Carnegie Mellon University
Pittsburgh, PA 15213-3891
Phone: 412-268-8778
Email: talukdar@ece.cmu.edu

Power Systems Engineering Research Center

This is a project report from the Power Systems Engineering Research Center (PSERC). PSERC is a multi-university Center conducting research on challenges facing a restructuring electric power industry and educating the next generation of power engineers. More information about PSERC can be found at the Center's website: http://www.pserc.org.

For additional information, contact:

Power Systems Engineering Research Center
Cornell University
428 Phillips Hall
Ithaca, New York 14853
Phone: 607-255-5601
Fax: 607-255-8871

Notice Concerning Copyright Material

Acknowledgements

The work described in this report was sponsored by the Power Systems Engineering Research Center (PSERC). It is the final report for the PSERC project "Bidding Agents". We express our appreciation for the support provided by PSERC's industrial members and by the National Science Foundation under grant NSF EEC-0109543 at Carnegie Mellon University, and grant NSF EEC-0118300 at Cornell University received through the Industry/University Cooperative Research Center program.

Executive Summary

Software agents can be of help in the design, analysis and verification of markets. The agents described in the four chapters of this report do two things. First, they identify some of the uses of software agents. Second, they show that competitive behavior is not guaranteed for existing electricity auctions. Rather, suppliers can drive prices to well above competitive levels and make excessive profits.

Chapter-1 covers software agents designed to mimic test-results obtained from human subjects in simulated electricity markets. The agent behavior is compared to the behavior of the subjects. Both human and software-agent behavior are classified based on the data. Differences and similarities are noted and explained. The agents are used to demonstrate that the hope that large-scale generating units will operate at marginal cost 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 the supplier will try to increase its profits by submitting offers to sell at higher than marginal cost and by withholding units, if permitted. This makes predicting unit commitment and dispatch ahead of time difficult.

Chapter-2 covers the use of a multi-agent system (MAS) to simulate a spot market with several supply firms. The firms submit offers to maximize their own expected profits, and an Independent System Operator (ISO) clears the market for a predetermined load in a uniform price auction. The firms learn about the structure of the market and the behavior of their competitors by comparing actual market outcomes with predicted outcomes based on an estimate of their own residual demand curve. This estimated demand curve is updated each period using a Kalman filter. The simulations were used: 1) to determine which characteristics of a deregulated wholesale market for electricity make price spikes likely to occur, and 2) to determine a structure of firms, in terms of size and type, that replicates the infrequent price spikes observed in the PJM market during the summer of 1999. The two most important characteristics of the market for creating price spikes are 1) uncertainty about the system load is the primary determinant of the observed speculative behavior (i.e. submitting offer curves shaped like a hockey stick), and 2) all firms eventually become speculators, and it is unrealistic to expect firms to behave like price takers in a market with only a few suppliers. In other words, given the intrinsic uncertainty about the load (and the possible outages of generators), speculative behavior by suppliers is entirely rational. Overall, the results show that the analytical framework of a MAS can replicate observed price behavior in a deregulated electricity market, can provide new insights into how suppliers behave, and has the potential for evaluating a wide range of policy options for mitigating high prices.

Chapter-3 investigates how spot prices are affected by forward contracts. Experimental economics and agent-based simulations are used. The results show that holding a forward contract is an effective way to mitigate high prices if the same contract is held for all trading periods and the price of this contract is independent of the spot prices. However, when a forward contract is renewed periodically and the spot prices influence the forward

price, there is more speculation and the spot prices may be higher than the base case with no forward contracts. The simulation results also show that software agents were able to replicate the behavior of students effectively in the experiments.

Chapter-4 uses software agents to examine repeated auctions. The purpose of the agents is not to mimic the behavior of humans. Rather, it is to reveal those fundamental properties of repeated auctions that are independent of human idiosyncrasies; just as wind tunnels and finite element programs are used to reveal the fundamental properties of new airplanes before test pilots are made to fly them. Simple experiments show that existing designs of repeated auctions have at least one major flaw: The sellers, even though they work without collusion, can learn strategies that raise prices and profits far above competitive values. This happens even when the demand is price-responsive. A remedy is to make the auctions symmetric, that is, to provide the buyers with as much autonomy as the sellers have, so the buyers, or their surrogates, can learn, and otherwise make decisions as quickly and cleverly as the sellers. However, this remedy is unlikely to be the only one needed. Complex artifacts, unless their designs are thoroughly verified, invariably suffer from major flaws. Thorough verification requires a set of tests that spans both the operating conditions and the desired behaviors of the artifact. What passes for verification in market design—opinion surveys and a few random experiments with human subjects—is far from adequate.

Table of Contents

Acknowledgments.. i
Executive Summary ii
Table of Contents.. iii
Table of Tables v
Table of Figures vii

1. Comparing the Behavior of Agents to Human Subjects in a Uniform Price Auction.. 1-1

1.1 Introduction 1-1
1.2 Electricity Market 1-2
1.3 Standardized Agents 1-2
1.4 Classification of an Agent 1-4
1.5 Expected Earnings 1-5
1.6 Simulation Results 1-6
1.7 Conclusions 1-13
1.8 Future Work 1-13
1.9 Acknowledgment 1-13
1.10 References 1-14

2. Using Software Agents to Replicate Observed Price Behavior in the PJM Electricity Market 2-1

2.l Introduction 2-2
2.2 A Multi-Agent System for Testing Electricity Markets 2-3
2.2.1 The Model of a Deregulated Electricity Market 2-3
2.2.2 Using Adaptive Autonomous Agents for Firms 2-4
2.2.3 Limitations of Deterministic Profit Maximization 2-5
2.3 Creating Price Spikes in a MAS 2-7
2.3.1 Salient Features of the Market 2-7
2.3.2 Speculative Behavior is Almost Inevitable 2-8
2.3.3 Speculation Increases when Load is Uncertain 2-11
2.4 Replicating Observed Price Behavior in the PJM Market 2-12
2.4.1 Using the Observed Pattern of Load in PJM 2-12
2.4.2 Too Many Firms are Required to Make Textbook Competition Practical
.... 2-13
2.4.3 The Effect of Including Vertically Integrated Firms is Relatively Small ....
.... 2-14
2.4.4 Replicating the Structure of firms in PJM 2-15
2.5 Conclusions 2-17
2.6 References 2-19

Table of Contents
(continued)

3. Using Software Agents to Supplement Tests Conducted by Human Subjects: How Do Forward Contracts Affect Behavior of Suppliers in a Wholesale Electricity Market? 3-1

3.1 Introduction 3-1
3.2 Optimum Offer Behavior in a Spot Market 3-3
3.3 A Supplier's Expected Profit 3-3
3.4 The Optimum Offer 3-4
3.5 The Experimental Framework 3-6
3.6 Analysis of the Test Results 3-8
3.7 The Performance of Computer Agents 3-12
3.8 Summary and Conclusions 3-15
3.9 References 3-17

4. Symmetry and Verification : Critical Parts of Market Design 4-1

4.1 Introduction 4-1
4.2 The Dynamics of Repeated Electricity Auctions 4-2
4.3 Overview of the Experiments 4-4
4.4 Experiment-1: Price-Invariant Demand 4-6
4.5 Experiment-2: Price-Responsive Demand 4-9
4.6 Experiment-3: Cooperation 4-10
4.7 Experiment-4: Active Demand 4-11
4.8 Conclusions 4-15
4.9 Acknowledgement 4-17
4.10 References 4-17

Table of Tables

Table 1.1 Offer strategies of the standardized agent when the fairshare block is the jth block. Here MCO, S and W stand for the marginal cost offer, speculate and withhold, respectively. 1-4

Table 2.1 Optimum Solutions for Profit Maximization by a Firm 2-5

Table 2.2 The Number of Firms at the Beginning and End of a Simulation by their Belief about High Prices 2-10

Table 2.3 Average Price for Different Numbers of Firms (Competitive price is \$39/MWh) 2-13

Table 2.4 Average Price for Different Numbers of VIF (Competitive price is \$39/MWh) 2.15

Table 3.1 The Optimum Offers in the Spot Market 3-5

Table 3.2 Impact of Contract Quantity on Optimum Offers 3-6

Table 3.3 Regression Results 3-10

Table 3.4 Average Earnings (\$ in periods 11 to 20) 3-11

Table 3.5 Average Market Prices (\$/MWh) 3-12

Table 3.6 Average Earnings (\$MWh)* for All-Agent Tests 3-15

Table 3.7 Average Market Prices (\$MWh)* for All-Agent Tests 3-15

Table of Figures

Figure 1.1 Schematic diagram showing the offer functions of the different types of standardized agents. 1-4

Figure 1.2 Earnings of the standardized agents and the agent of interest........ 1-5

Figure 1.3 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. 1-7

Figure 1.4 Example of the performance of human agents: the red square, green and blue circles represent a single period earnings of MC, WS and SS agents, respectively. 1-10

Figure 1.5 Actual earnings vs. expected earnings calculated from a simulation. 1-12

Figure 2.1 A Multi-Agent System for an Electricity Spot Market 2.4

Figure 2.2 Simulated Market Price 2-8

Figure 2.3 Load and Price for Case E3/S5........ 2-9

Figure 2.4 The Effect of Different Load Forecasting Errors (0% and 5%) on the Optimum Offers (OPT) of a Firm (MC is the Marginal Cost) 2-11

Figure 2.5 The Probability of Rejecting an Offer in the Auction (5% Forecasting Error)........ 2-12

Figure 2.6 Actual Load and Peak Price in the PJM Market (April to September, 1999) 2-13

Figure 2.7 Actual Market Prices and Simulated Market Prices Using Naïve Agents for Firms........ 2-16

Figure 2.8 Simulated Market Prices Using Naïve and Trained Agents for Firms 2-16

Figure 3.1 Optimum Offer Price with Contracts 3-5

Figure 3.3 Comparison of Average Earnings 3-13

Table of Figures
(continued)

Figure 4.1 The rows in the matrix, A_{kM}, represent offers from the sellers for the k-th auction; a "1" denotes a generator whose energy is offered for sale, a "0" denotes a generator that is withheld. The pattern of 1's and 0's in A_{kM} produces the maximum total profit for the sellers, $3860, just as they would obtain if their generators were controlled by a monopolist with perfect information, including perfect predictions of demand. 4-3

Figure 4-2 Twelve demand curves of which curves 1 - 3 are price-invariant, 4-12 are price-responsive. 4-7

Figure 4.3 Aggregate supply curves from Experiment-1 with the demand fixed at 50 MW. Initially, all the suppliers offer all their generators at cost. The resulting aggregate supply curve is represented by the broken line and the clearing price in the first auction is the competitive price: $3.13/MWh. As the auctions proceed, the suppliers learn to withhold generators. At repetition 1000, the withholding is so effective that the aggregate supply curve (solid line) intersects the demand at the maximum possible price, $100/MWh. 4-8

Figure 4.4 A trajectory of clearing prices from Experiment-1, with the demand fixed at 50 MW. The fluctuations in the trajectory decrease as the trajectory develops, but never completely disappear because the learning algorithm continues to makes random changes in the suppliers' offers. By repetition 1000, the average clearing price (averaged over 50 consecutive repetitions) has become $98/MWh, and it remains at this value thereafter. 4-9

Figure 4.5 Profits and learning yields from Experiments-1 and –2. Whether the demand is price-invariant (demand curves 1-3) or price-responsive (demand curves 4–12), the sellers are able to obtain between 72% and 87% of the monopolistic profit. 4-10

Figure 4.6 Profits and learning yields from Experiment-3, in which sellers cooperate by a) exchanging data on their individual profits, and b) seeking to maximize total, rather than individual, profit. 4-11

Figure 4.7 Some results from Experiment-4. Both the suppliers and the buyers are allowed to learn. Consequently, both the offers and the bids change from one repetition to the next. The broken curves represent the competitive supply and the actual demand. The solid curves show the supply that is offered and the demand that is bid at repetition-1000. 4-12

Table of Figures
(continued)

Figure 4.8 More results from Experiment-4. The broken lines represent the supply and demand curves at repetition-1. (The demand corresponds to Curve-10 in Fig 2). The solid lines show these curves in repetition-1000. The changes are learned. 4-13

Figure 4.9 A trajectory of clearing quantities from Experiment-4. with the demand corresponding to of Curve-4 in Fig. 2. 4-14

Figure 4.10 A clearing price trajectory from Experiment-4 with the demand corresponding to Curve-4 of Fig.2. 4-15

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

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. In this Chapter we explore 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 are classified based on the data. Differences and similarities are noted and explained.

1.1. Introduction

For the past 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

1.2. 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 five 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.

1.3. 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 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 1.1. 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 1.1 shows the result of Table for units with 5 blocks.

Table 1.1 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

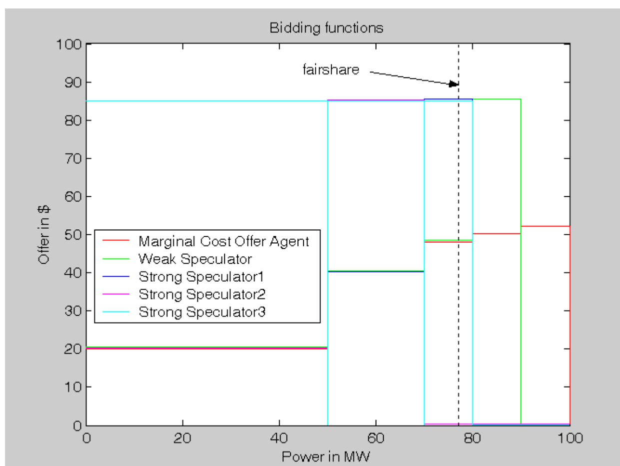

Fig. 1.1 Schematic diagram showing the offer functions of the different types of standardized agents

1.4. 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 1.2 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 1.2 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.

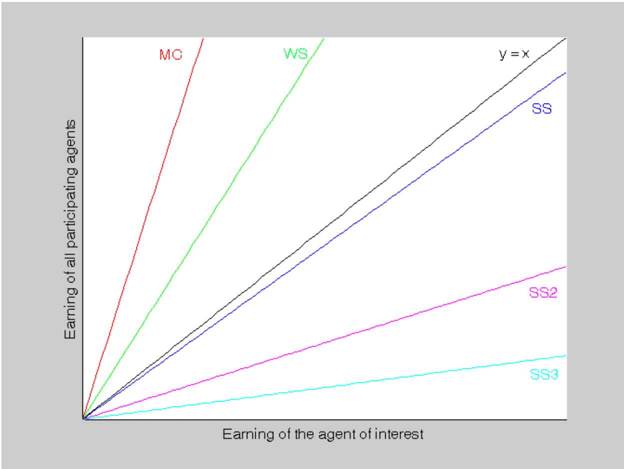

Fig. 1.2 Earnings of the standardized agents and the agent of interest

1.5. 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.

1.6. 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 1.3 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

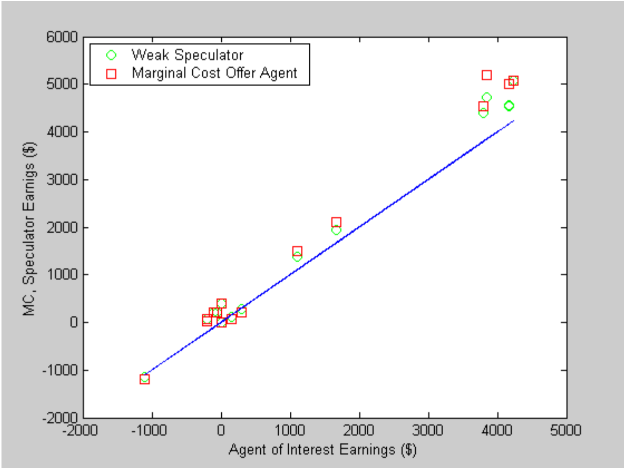

Fig. 1.3 Example of a performance of the software agents: in the plot, red square, green and blue circle stand for the earning in a period of MC, WS and SS, respectively

b) 3 WS + 2 MC

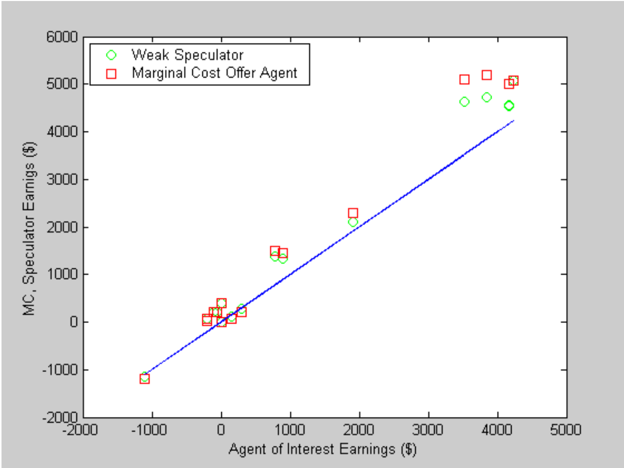

c) 4 SS + 1 MC

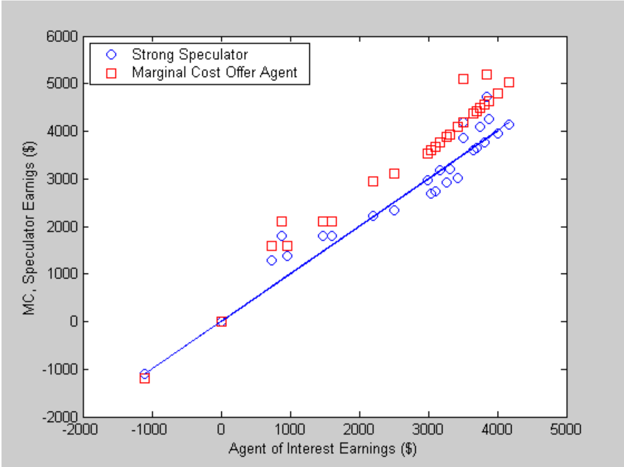

d) 3 SS + 2 MC

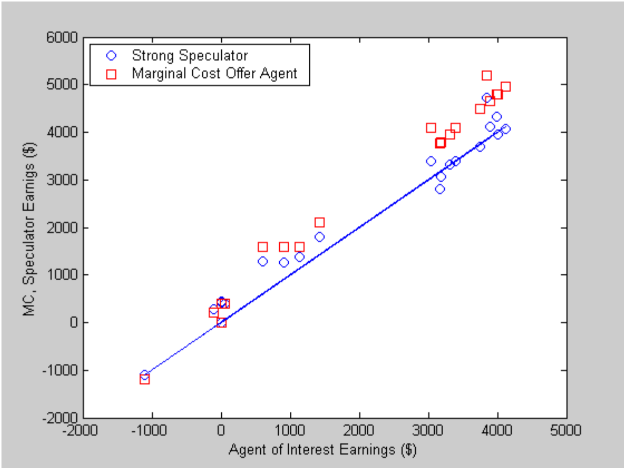

Fig. 1.3 Example of a performance of the software agents: in the plot, red square, green and blue circle stand for the earning in a period of MC, WS and SS, respectively (continued)

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 1.4.

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

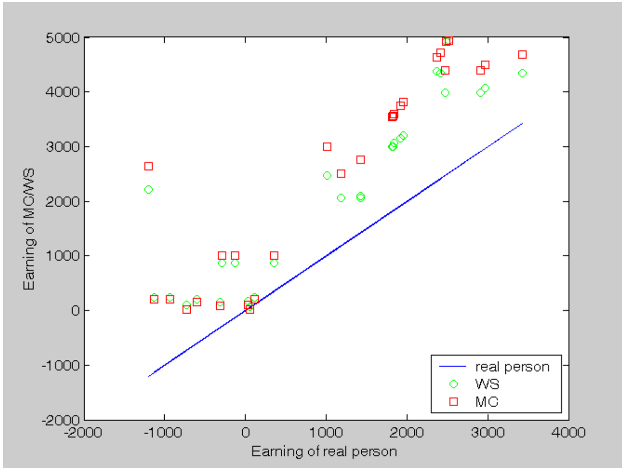

b) 3 WS + 2 MC

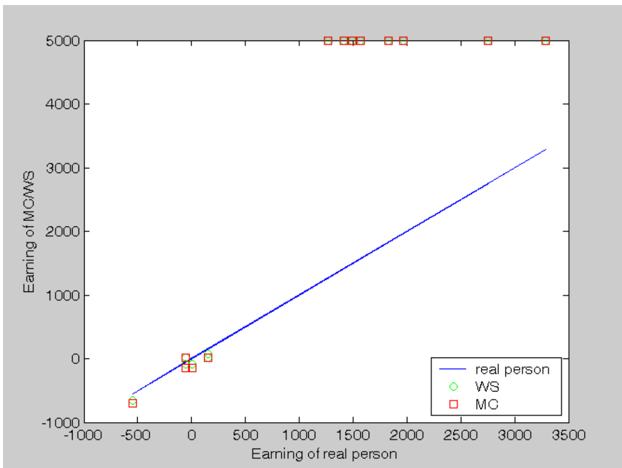

Fig. 1.4 Example of the performance of human agents: the red square, green and blue circles represent a single period earnings of MC, WS and SS agents, respectively

c) 4 SS + 1 MC

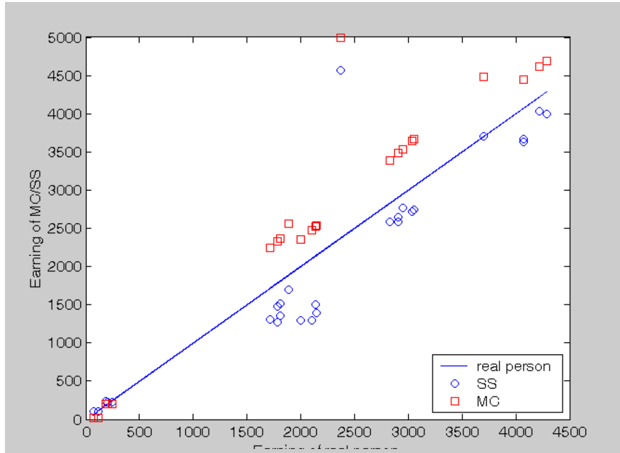

d) 3 SS + 2 MC

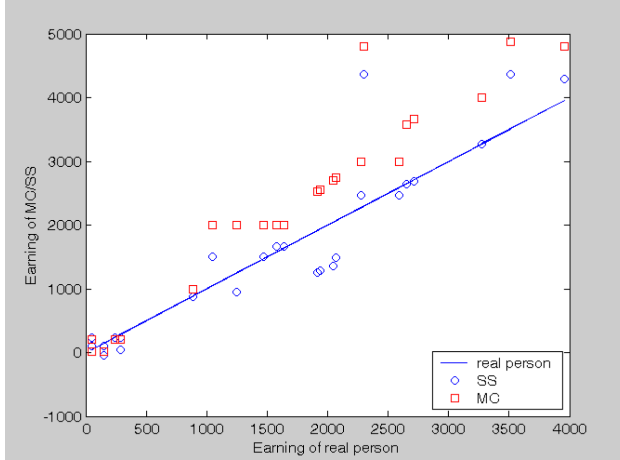

Fig. 1.4 Example of the performance of human agents: the red square, green and blue circles represent a single period earnings of MC, WS and SS agents, respectively (continued)

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. 1.5. 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. 1.5 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.

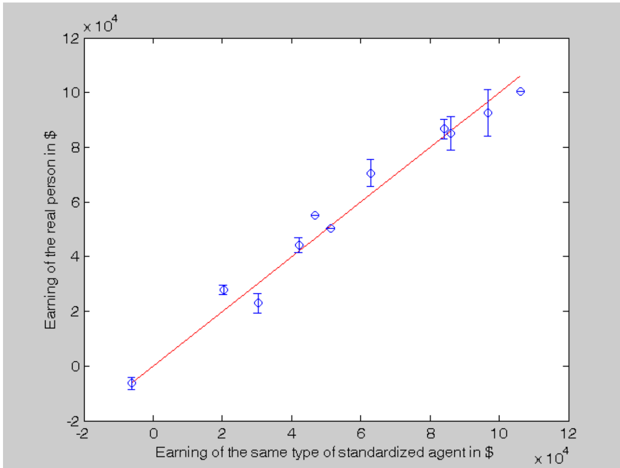

Fig. 1.5 Actual earning vs. expected earnings calculated from a simulation

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.

1.7. Conclusions

In this chapter, 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 agent and a human agent choose several different strategies. Each strategy produces different offer functions. The main results of this chapter describe how to classify the strategy not by inspecting individual offer functions 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. The results of this chapter also shows that only a small number of standardized agents need be used for the classification, and their earnings are highly correlated with the earning of an actual agent.

1.8. Future work

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 on-going.

1.9. Acknowledgment

The authors would like to thank Peter Wong for his valuable help on the computer simulation for this work. They also thank all the students from ECE 551/AEM 655 and ECE 451 for participating the experiments.

1.10. 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).

2. Using Software Agents to Replicate Observed Price Behavior in the PJM Electricity Market

Abstract

The main objectives of this chapter are 1) to determine which characteristics of a deregulated wholesale market for electricity make price spikes likely to occur, and 2) to determine a structure of firms, in terms of size and type, that replicates the infrequent price spikes observed in the PJM market during the summer of 1999. The analysis employs a Multi-Agent Simulation (MAS) to replicate a spot market with six supply firms, represented by adaptive autonomous agents. These firms submit offers to maximize their own expected profits, and an Independent System Operator (ISO) clears the market for a predetermined load in a uniform price auction. The firms learn about the structure of the market and the behavior of their competitors by comparing actual market outcomes with predicted outcomes based on an estimate of their own residual demand curve. This estimated demand curve is updated each period using a Kalman filter. An additional objective of the paper is to show that a MAS provides a potentially useful analytical framework for evaluating the effects of specified modifications to a deregulated electricity market.

The two most important characteristics of the market for creating price spikes are 1) uncertainty about the system load is the primary determinant of the observed speculative behavior (i.e., submitting offer curves shaped like a hockey stick), and 2) all firms eventually become speculators, and it is unrealistic to expect firms to behave like price takers in a market with only six suppliers. In other words, given the intrinsic uncertainty about the load (and the possible outages of generators), speculative behavior by suppliers is entirely rational.

With six suppliers controlling the same amount of generating capacity in a MAS, high prices can eventually be sustained in the market when the load is high. When all six firms speculate, the resulting price behavior does not replicate the relatively infrequent price spikes seen in a market like PJM. Using the actual pattern of daily peak load in the summer of 1999 for PJM, the results also show that it is impractical to rely solely on 1) increasing the number of competitive suppliers, or 2) increasing the number of vertically integrated firms to bring average spot prices close to competitive levels. In contrast, the effect of small firms in the market is important for mitigating high prices because these firms increase supplies when prices are high. Small firms do not own enough capacity to make speculation profitable. With a modified structure of firms, realistic price behavior can be replicated if 1) two of the six firms are vertically integrated, and 2) two of the firms are each replaced by four small firms that behave more like price takers. Since this modified structure of firms is roughly the same as the actual structure of firms in the PJM market, and the simulated price volatility is very similar to the actual volatility, this modified structure of firms could be used as the basis for evaluating different policies for mitigating high prices in a deregulated market. Overall, the results show that the analytical framework of a MAS can replicate observed price behavior in a deregulated electricity market, can provide new insights into how suppliers behave, and has the potential for evaluating a wide range of policy options for mitigating high prices.

2.1. Introduction

Following the California crisis in 2000-01, regulators in the USA have become more willing to intervene directly in electricity markets and modify the behavior of suppliers (Bushnell, 2003). For example, Automatic Mitigation Procedures (AMP) are used in the northeastern markets. However, requiring suppliers to justify high offers to sell capacity inevitably leads to suppliers exaggerating their true costs (Wolak, 2003). The overall result is quasi-regulation and endless litigation about the true costs. Markets should be more robust to speculation, and the best way to make a market competitive is to ensure that any firms submitting high, speculative offers lose market share. While there are special circumstances in electricity markets that give some suppliers market power (e.g. being assigned to must-run-for-reliability), it is still appropriate to consider how to make electricity markets more self-regulating under normal circumstances. This goal puts greater emphasis on getting the right structure of a market (e.g. how many firms are needed?) and less on restricting the behavior of suppliers. In this analysis, suppliers try to maximize profits with no fear of regulatory intervention. In all cases, the firms can withhold capacity and speculate with high offers if it is in their own interest to do so. The results of the analysis show that it is impractical to rely solely on 1) increasing the number of suppliers, or 2) increasing the number of vertically integrated firms in the market to eliminate speculative behavior. However, having a number of small firms in the market is an effective way to reduce the number of price spikes.

The analysis is based on simulations of a wholesale market for electricity run by an Independent System Operator (ISO). Suppliers submit offers into a central auction, and the ISO determines the optimum pattern of dispatch to minimize the cost of meeting load. A uniform price auction is used to determine the market price. The role of suppliers in the auction is taken by software agents. These agents learn about the market and adapt their behavior in response. The adaptation involves updating an estimate of the residual demand curve faced by each firm, and this curve is used by the firm to determine the optimum set of offers to maximize expected profits in the next round of the auction.

One advantage of using software agents to test markets, rather than people, is that it is practical to run a much more extensive range of tests. However, there are some important restrictions on the design of the computer agents for our analysis. An underlying objective is to design the agents to match the observed behavior of suppliers in electricity markets and in laboratory tests of markets. The main restrictions on the design of an experiment are that viable results must be obtained using a relatively small number of suppliers and a relatively small number of trading periods. For example, the standard market test using PowerWeb at Cornell involves only six firms, and tests with more than 50 trading periods are rare. These restrictions on the design of an agent can be met by specifying a specific functional form for the residual demand curve faced by a firm. The form chosen in the analysis allows firms to behave like price takers or like speculators, depending on the estimated parameter values. Updating the estimates of the demand parameters in response to market outcomes allows firms to modify their optimum strategy during a simulation. An interesting result from the analysis is showing that the uncertainty about actual levels of load purchased (due to forecasting errors or contingencies) results in supply curves that are shaped like a hockey stick. This is exactly the type of behavior observed in deregulated electricity markets and it accounts for the existence of price spikes as a typical feature of these markets.

Section 2 of the paper discusses the modeling framework, followed by an explanation of why price spikes occur in Section 3. In most situations, firms learn to speculate, and high market prices are very persistent compared to real markets like PJM. In Section 4, the model is modified to replicate the type of volatile price behavior with infrequent price spikes seen in the PJM market. The most important modification is to change the structure of six identical firms by including vertically integrated firms and small firms. Simply increasing the number of identical firms to get prices close to competitive levels requires an increase from 6 to over 20 firms. Similarly, competitive prices can be obtained if 5 of the 6 firms are vertically integrated. Neither of these specifications is practical as a standard for designing a competitive market. In contrast, having a few small firms in the market is an effective way to mitigate price spikes. Converting two of the six firms to vertically integrated firms and splitting another two firms into eight small firms results in a spot market with infrequent price spikes that is very similar to the price behavior observed in the PJM market during the summer of 1999.

2.2. A Multi-Agent System for Testing Electricity Markets

2.2.1 The Model of a Deregulated Electricity Market

A Multi-Agent System (MAS) consists of three main elements: an environment, a set of agents, and a set of tasks. Figure 2.1 shows how these elements are combined together to simulate an electricity market. The *environment* represents the domain in which the agents (i.e., decision-makers) interact, and in this case, it is a wholesale spot market for electricity. Assuming the simplest form of market, the characteristics of the environment are 1) all generation capacity is dispatched in a central auction, 2) the last accepted offer (i.e., the highest offer) sets the market price which is paid to all accepted offers (i.e., a uniform price auction), 3) the market is a one-settlement system and there are no bilateral contracts, 4) any offer to sell is allowed if it does not exceed a specified price cap, 5) demand (system load) is determined exogenously, 6) in cases of a supply shortage, withheld generators are randomly recalled to meet load, and 7) there are no imports and exports.

An *agent* is a computer program that is capable of autonomous action to meet specified objectives in the environment (Weiss, 1999). Figure 2.1 shows a group of *firm* agents representing suppliers (Agent I - Agent VI) and an *Independent System Operator* agent (ISO). The firms are *Adaptive Autonomous Agents* (AAA) since they learn from previous experiences in the market. In contrast, the ISO agent is not an AAA agent but simply applies a fixed set of rules to determine the market price and forecast load.

The *task* of the ISO is to operate the electricity market using the rules of a uniform price auction. Market operations include the following tasks: 1) aggregate offers submitted by the firms, 2) calculate the total supply curve (S in Figure 2.1), 3) determine the optimum dispatch to minimize the cost of meeting the current load (D in Figure 2.1), 4) recall generators withheld from the auction if the total offered capacity is less than the load, 5) tell each firm how much capacity it sells and the market price, and 6) post a new load forecast for the next trading period. Note that the task of the ISO does not include intervention in the market to mitigate high prices, for example. This implies that the firms do not face any regulatory threat, and they can exploit market power fully with no adverse consequences.

Environment: Spot Market

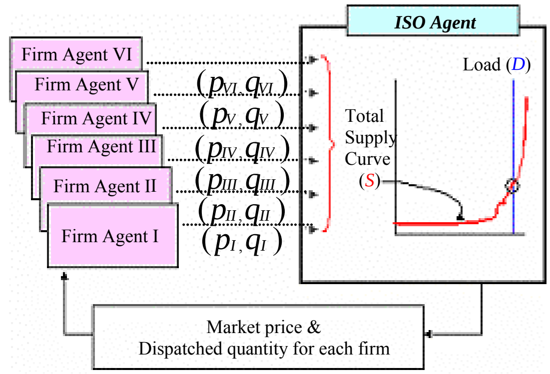

Fig. 2.1 A Multi-Agent System for an Electricity Spot Market

2.2.2 Using Adaptive Autonomous Agents for Firms

The main characteristics of each firm are 1) a set of learning and decision algorithms, 2) a specified amount of generating capacity owned, 3) the corresponding operating costs, and 4) an initial market perspective about the behavior of other suppliers. Throughout the entire set of simulations, the first two characteristics are fixed and identical for the six firms, but the operating costs and initial market perspectives can be quantitatively different. The goal of each firm is to submit price and quantity offers (p and q in Figure 2.1) to maximize expected profits using an estimate of the "residual demand" curve faced by the firm. Conditionally on the residual demand curve, the main source of uncertainty is the error in the load forecast provided by the ISO.

An important role of the firms in our analysis is to represent the behavior of suppliers in a framework that is consistent with economic theory (e.g. the approach used by Mayer and Klemperer (1989)). Using an explicit functional form for the residual demand curve represents a departure from the typical specifications of an agent used in a MAS, and in this respect, our firm agents are really *Adaptive Autonomous Optimizing Economic Agents* (AAOEA). The main advantage of our approach is that the behavior of each firm agent can be evaluated directly using conventional economic criteria.

In the simulations, each firm owns five generating blocks and submits offers to maximize the expected profit for the next auction. The capacity and operating costs are known for each block of capacity and are fixed over time. For each time period, each firm observes the market price, the system load and the quantity of its own capacity dispatched in the last auction. In addition, the ISO posts a load forecast for the next auction. For each firm, a residual demand curve is specified as an inverse function of the "excess" capacity offered into the auction (i.e., the available capacity that is offered but is not dispatched). This form of residual demand allows for a wide range of market behavior from competitive to the type of speculation implied by "hockey stick" supply curves (see Oh (2003)). For each firm, the residual demand curve for the next auction can be written:

$$\begin{aligned} \mathrm{P} &= 1/(a_t + b_t\,(OC_t + q - \hat{Q}_t)) \\ &= 1/(a_t + b_t OC_t - b_t\,(\hat{Q}_t - q)) \\ &= 1/(\alpha_t - \beta_t\,(\hat{Q}_t - q)/IC) \end{aligned}$$

where:

P is the market price,
$\hat{Q}_t$ is the forecasted system load,
OC_t is the offered capacity from other firms,
IC is the installed capacity of other firms,
q < $q_{\max}$ is the own capacity dispatched, and
$a_t > 0$ and $b_t > 0$ are the subjective parameter values of the firm.

The re-parameterization to $\alpha_t = a_t + b_t\, OC_t$ is convenient because OC_t is unobserved and it avoids the computational problems of getting $a_t < 0$ when updating (b_t is also used in the updating process, but β_t is specified here because the values are easier to interpret). $P_L = 1/\alpha_t$ corresponds to the low market price if the firm could undercut the offers of all other firms and cover all of the load (i.e., $q = \hat{Q}_t$). Clearly, the firm's own installed capacity, $q_{\max}$, is the maximum that can actually be offered into the auction by a firm. For the other parameter ($\beta_t = b_t\, IC$), $P_H = 1/(\alpha_t - \beta_t)$ corresponds to the highest possible price in the market when $q = \hat{Q}_t - IC$ (i.e., the price for the first unit of capacity dispatched in the market). In a truly competitive market, $P_L = P_H$ and $\beta_t = 0$. When $0 < \beta_t < \alpha_t$, $P_H > P_L$ and the firm believes that it has some market power. As $\beta_t \to \alpha_t$, P_H increases and values greater than the price cap in the market can be interpreted as other firms withholding capacity from the auction. This type of withholding can be sufficiently large to make the firm "pivotal" (i.e., essential for meeting the load when $OC_t < \hat{Q}_t$). The restriction $0 < \beta_t\, \hat{Q}_t / IC < \alpha_t$ ensures that prices are positive and finite for $0 \le q \le q_{\max}$, which is the relevant range of quantity offers for the firm ($\beta\, \hat{Q}_t / IC > \alpha_t$ makes the firm pivotal).

Unlike human subjects, the firm agents are incapable of learning about the structure of the market by employing complex counterfactual scenarios or deep introspection, and they rely on simple adaptive learning using a Kalman filter to update the residual demand curve.

By specifying a model with time-varying parameters for the residual demand curve, current estimates of the parameters can be revised after comparing the difference between the last actual market outcome and the predicted outcome. Changes in the parameters are proportional to the size of this prediction error. Using the updated parameters and the load forecast for the next auction, the firm agent then determines the optimum offer for each block of capacity to maximize expected profits. A numerical search is used to determine the set of optimum offers, and some blocks may be withheld.

2.2.3 Limitations of Deterministic Profit Maximization

If profit maximization is deterministic (i.e., α_t, β_t and $\hat{Q}_t$ are fixed) and the marginal cost of generation for the firm is constant, the optimum q and the market price P can be determined as the solution to a quadratic equation. However, the implied optimum behavior of the firm is quite different from actual behavior in a typical market like PJM. There is more withholding of capacity and less speculation, and the offered supply curve does not look like a hockey stick. Hence, an important objective in the next section is to show that optimum behavior in a MAS is consistent with observed behavior when load is stochastic and capacity is offered in discrete blocks.

A simple example of deterministic profit maximization is provided here as a benchmark for the MAS simulations in the next section. Assume that other suppliers have 50GW of capacity (OC), the load is 40GW (Q), and the firm owns 10GW of capacity ($q_{\max}$) with a fixed marginal cost of $50/MWh. In a competitive market, the residual demand curve would have $\beta=0$, and $q_{\max}$ would be the optimum dispatch if the market price P > $50/MWh ($\alpha$ < 0.02). Assuming that the residual demand is always $q_{\max}$ when the price is $50/MWh, values of $0 < \beta < 0.1$ (and $\alpha = 0.02 + 0.6\beta$) correspond to different degrees of market power for the firm. Assuming that the ISO imposes a price cap of $1,000/MWh, the optimum solutions are summarized in Table 2.1.

Table 2.1 Optimum Solutions for Profit Maximization by a Firm (Load = Q = 40GW, Installed Other Capacity = IC = 50GW, Own Installed Capacity = $q_{\max}$ = 10GW)

β	α	Implied Other Capacity (MW)*	Optimum Dispatch (MW)	Optimum Price ($/MWh)
.00	<.02	50,000**	10,000**	50
.01	.026	50,000**	4,868	53
.05	.05	49,000	4,142	71
.09	.074	40,556	2,403	158
.095	.077	40,000	1,827	224
.0995	.0797	39,548	610	707
.0999	.07994	39,510	490	1,000**
≥.1	≥.08	≤39,500	Pivotal	Infinite

*Implied capacity offered by other firms at the price cap of $1000/MWh.
**Maximum value allowed under the market specifications.

Optimum prices at the price cap of $1,000/MWh do not occur as β increases from zero until the firm is very close to being pivotal. With regard to withholding, all values of $\beta > 0$ imply that at least half of the firm's capacity is withheld. When the optimum price is at the price cap (β = 0.0999), less than 5% of own capacity is dispatched. For the same set of residual demand curves, setting the marginal cost at $20/MWh instead of $50/MWh results in less withholding and less speculation. For example, when β = 0.06 or less, all 10GW of capacity is dispatched and the optimum price is $50/MWh. However, when β = 0.0999, only 490 MW are dispatched and the optimum price is $1,000/MWh. Hence, there are still substantial differences between these deterministic optimum solutions and observed offer behavior. In the simulations discussed in the next section, high offers and supply curves shaped like a hockey stick are a standard feature of our MAS results under a wide range of market conditions.

2.3. Creating Price Spikes in a MAS

2.3.1 Salient Features of the Market

An important objective of this section is to gain insight into why price spikes are a common feature of wholesale markets for electricity. Under the initial specifications used for our MAS, it turns out that high prices are much more persistent than they are in a market like PJM. Section 4 provides an additional analysis of factors that can mitigate high prices and lead to market outcomes that are more consistent with observed behavior. The basic question posed in this section is what characteristics of the market are responsible for the type of speculative offers seen in real markets? The simple answer is that the uncertainty about actual system load in the next auction implies that there is always a wide range of possible dispatch levels for a given firm. In addition, there are substantial differences in the true marginal costs of units in low-load periods and high-load periods. Consequently, any firm knows that supply curves are not infinitely elastic even if all other firms submit "honest" offers equal to the true marginal costs. Consistent with the results of a repeated game, all firms will eventually learn how to exploit market power. The main difference among the simulation results presented in this section is in the number of auction periods needed to get the first price spike.

All scenarios assume that there are six suppliers (firms), and each one controls 10GW of generating capacity, split into five blocks (4, 2.5, 2, 1, 0.5) with similar cost structures for all firms ranging from base load to peaking capacity. The share of total capacity controlled by each firm (17%) corresponds roughly to the largest firms in a market like PJM. (It is also consistent with the number of firms in the standard market tests conducted at Cornell using PowerWeb.) The main difference among the firms in this section is in their perception of the market (i.e., the shape of the residual demand curve faced by a firm). Actual market outcomes are evaluated by each firm and used to update the shape of that firm's residual demand curve through a Kalman filter. After updating, each firm submits offers to maximize their own expected profits.

Each simulation involves a series of 135 auction periods, corresponding roughly to daily auctions from mid April to August. There are three different sets of initial conditions for a firm (market environment) and each one is run with five different sets of market conditions (scenarios) to give a total of 15 combinations. The three different environments for firms are:

E1) All firms are price takers, and the highest marginal cost is $55/MWh (**Initial price takers**, with low marginal costs)

E2) All firms are price takers, and the highest marginal cost is $100/MWh (Initial price takers, with **high marginal costs**)

E3) Four price takers and two latent speculators with low marginal costs (**two latent speculators,** with low marginal costs)

The marginal costs for the five blocks are (10, 20, 30, 50, 55) under the low cost conditions (E1 and E3), and under the high cost conditions (E2), the cost of the most expensive unit is 55, 65, 75, 85, 95 and 100 for the six different firms. The initial values of the two parameters defining the residual demand curve (α and β in the previous section) vary stochastically among firms. Initial price takers believe that the residual demand curve is relatively flat ($\beta > 0$ is relatively small). In contrast, a latent speculator believes that high prices are possible even though they may be highly unlikely ($\beta > 0$ is relatively large).

Changes in the market conditions in the five different scenarios are cumulative, and each change increases the likelihood that price spikes will occur. In the base scenario, the system load is fixed at 40GW (corresponding to two-thirds of the installed capacity). Load is completely inelastic, and there are no forecasting errors. The only costs of generation are the known operating costs when a unit is dispatched. The five different scenarios are:

S1) Base (Load fixed at 40 GW)
S2) Standby costs of $5/MW charged for all capacity submitted to the auction
S3) Actual load varies stochastically around the forecasted load of 40GW (5% error)
S4) Forecasted load follows a smooth seasonal cycle (average load is still 40GW)
S5) Forecasted load follows a stochastic seasonal cycle (average load is still 40GW)

Since the changes in the five scenarios are cumulative, standby costs are charged in all scenarios except S1, and load is forecasted with a 5% error in S3, S4, and S5. Standby costs increase the incentive for firms to withhold capacity from the auction. Forecasting error and load variability provide more information about the structure of the market (i.e., better estimates of the shape of the residual demand curve). Both of these effects make speculative behavior and price spikes more likely. A price cap of $1,000/MWh is imposed on the market.

2.3.2 Speculative Behavior is Almost Inevitable

For comparison purposes, efficient market prices are computed for each of the five scenarios (E0). These prices correspond to all firms submitting offers equal to the true marginal costs, with no capacity withheld. The results are summarized in Figure 2.2 in terms of the average market price for the different simulations. It is clear that there are two general outcomes. Either the average price is close to the efficient price, or it is close to the price cap. Once a price spike occurs, high prices become relatively persistent. Figure 2.3 shows the load pattern for S5 and the corresponding prices for E3/S5 to illustrate the relatively rapid change from low prices to high prices after the load increases above 40 GW. Note that the high prices persist after the load falls below 40GW because the firms learn how to exploit market power more effectively during the simulation.

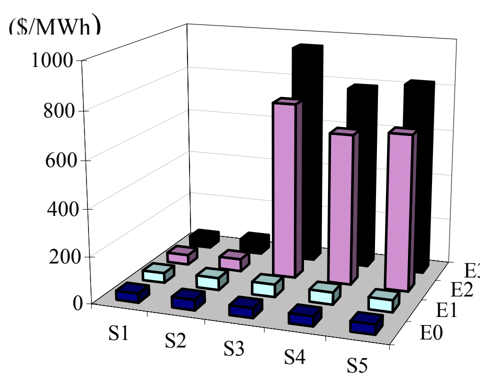

Fig. 2.2 Simulated Market Price

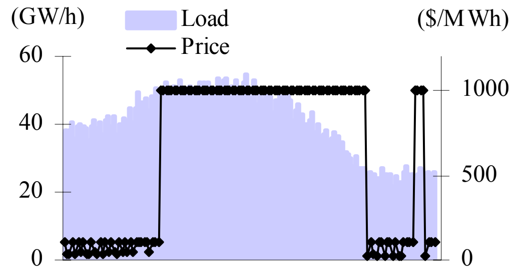

Fig. 2.3 Load and Price for Case E3/S5

When all firms start as price takers (E1), the average prices are low in all five scenarios. When marginal costs are high (E2) or there are latent speculators (E3), the average prices are high in S3, S4 and S5 (i.e., when firms experience a wider range of load levels), but not in S1 and S2. A fixed level of load does not provide firms with enough information to learn about the structure of their residual demand curves and exploit market power. This result is consistent with earlier tests of auctions using PowerWeb. Setting load at a fixed level, Bernard et al. (2003) found that market prices were close to competitive levels with six identical firms in the market. However, having fewer than six firms led to higher prices. An additional result in our simulations shows that charging standby costs in S2 does not change the outcome significantly from S1 even though there is a greater incentive to withhold capacity in S2.

The overall conclusion from the results in Figure 2.2 is that seeing higher prices in the market provides the information needed to exploit market power. This information may come from the physical characteristics of the market (i.e., high marginal costs in E2) or from the belief by some firms that high prices are possible (i.e., two latent speculators in E3). It is unrealistic to assume that firms can be shielded from this information.

Submitting high offers for some capacity is an inherent characteristic of a uniform price auction when there is uncertainty about the actual level of load purchased. This conclusion is illustrated by the changes in the shapes of the residual demand curves summarized in Table 2.2. Evaluating the residual demand curve at q = (Forecasted Load – Installed Capacity of Other Firms) gives the highest price payable in the market (i.e., the price when quantity is zero for a standard demand curve). The three columns in Table 2.2 correspond to different beliefs about how high prices can go. A price taker corresponds to beliefs that the highest price is <$100/MWh, and a speculator corresponds to beliefs that the highest price is >$100/MWh. High prices >$1000/MWh, the price cap, imply that a firm believes other firms are withholding capacity from the auction. Table 2.2 shows the number of firms for each type of belief (<100, 100-1000, >1000) at the beginning of a simulation and at the end of the simulation.

When all firms start as price takers (E1), they do not necessarily remain as price takers. When higher load levels occur in S3, S4 and S5, beliefs change and firms become latent speculators by the end of the simulation. This is true even though the market prices are low, and it is only in S1 and S2, when load is fixed at 40GW, that the firms remain as price takers. The implication from the low average prices in S3, S4, and S5 is that the firms became speculators too late to exploit market power when load was high. However, if the same pattern of load was repeated, the firms would begin as latent speculators (like E3), and high prices would occur. Seeing the first price spike provides a major source of new information for price takers, but price spikes are not necessary for price takers to adapt and become speculators.

The typical learning process for firms occurs in one direction from price taker to speculator. Speculators do not get discouraged when there are no price spikes because there are virtually no costs to holding speculative beliefs (in terms of lower profits). The belief that high prices are possible does not necessarily lead to speculative offers, but if market conditions are right, speculation will occur. In E3, the two latent speculators (and the four initial price takers) do not change their beliefs in S1 and S2 even though average prices are low. However, all firms in E3 are speculators by the end of S3, S4 and S5 when average prices are high. The overall conclusion is that it is unrealistic, under our specified set of market characteristics, to assume that a firm should behave like a price taker unless additional restrictions are imposed on the firm. For example, a firm controlling only one small generator or only nuclear capacity would be unwilling to speculate. Price speculation carries with it a larger probability that a high offer will be rejected, but it is still rational behavior for profit maximization when there is uncertainty about actual system load.

Table 2.2 The Number of Firms at the Beginning and End of a Simulation by their Belief about High prices

		<\$100	\$100 - \$1,000	>\$1,000
E1	S1	0 → 6	0 → 0	0 → 0
	S2	0 → 6	0 → 0	0 → 0
	S3	6 → 0	0 → 6	0 → 0
	S4	6 → 0	0 → 0	0 → 6
	S5	6 → 0	0 → 0	0 → 6
E2	S1	6 → 6	0 → 0	0 → 0
	S2	6 → 6	0 → 0	0 → 0
	S3	6 → 0	0 → 0	0 → 6
	S4	6 → 0	0 → 0	0 → 6
	S5	6 → 0	0 → 0	0 → 6
E3	S1	4 → 4	0 → 0	2 → 2
	S2	4 → 4	0 → 0	2 → 2
	S3	4 → 0	0 → 0	2 → 6
	S4	4 → 0	0 → 0	2 → 6
	S5	4 → 0	0 → 0	2 → 6

2.3.3 Speculation Increases when Load is Uncertain

In Section 2, for a firm facing no uncertainty, optimum offers exhibit substantial withholding and relatively low offer prices, unless the firm is very close to being pivotal. In contrast, the firms in many of the simulations presented in this section submit much higher offers and withhold less. A partial explanation for the differences in behavior is that the structure of costs is more complex in the simulations, but this is not the most important reason. Speculative behavior occurs when the load varies stochastically (S3, S4 and S5). Even if the typical forecasting error of total load is small in percentage terms, it will be relatively large for a firm. For example, a 5% error when load is 40GW corresponds roughly to a range of 8GW in the actual load. This range is substantial for a firm controlling 10GW of capacity. It increases the probability that a high offer will be accepted and makes speculation pay off.

The true marginal costs and the optimum offers are shown in Figure 2.4 for two cases with 1) no forecasting error, and 2) a 5% forecasting error. It is clear that there is much more speculation in the second case, and the offer curve is shaped like a hockey stick. In both cases, the forecasted load is 40GW and the residual demand curves are the same. The only difference is in the forecasting error. When there is no uncertainty about the load, the firm knows for sure that offers for the fourth block above $95/MWh will not be accepted. When load is uncertain, there is a small probability that offers as high as $1000/MWh will be accepted. For example, Figure 2.5 shows the probability of rejecting different offer prices in terms of the likely range of actual loads when the forecast of load is 40GW. Since loads above 40GW are possible, it is rational for a firm to speculate, and a hockey stick supply curve is an inevitable consequence. Furthermore, in many of our simulations, firms are able to sustain high prices (e.g. see Figure 2.3). The next section addresses the issue of how to modify the model so that it replicates a market which has only a few price spikes, like the PJM market.

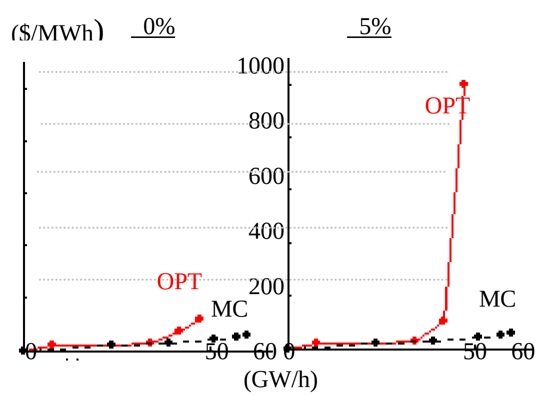

Fig. 2.4 The Effect of Different Load Forecasting Errors (0% and 5%) on the Optimum Offers (OPT) of a Firm (MC is the Marginal Cost)

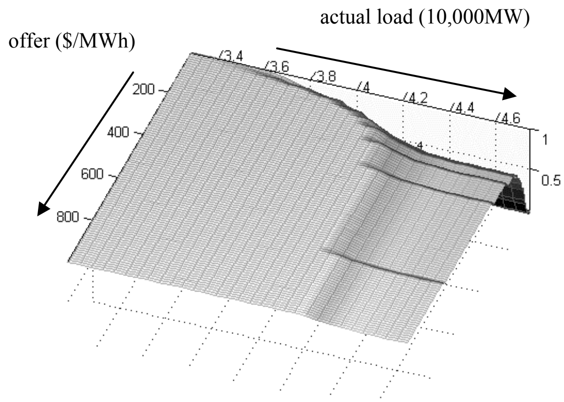

Fig. 2.5 The Probability of Rejecting an Offer in the Auction (5% Forecasting Error)

2.4. Replicating Observed Price Behavior in the PJM Market

Since the simulated patterns of market prices in the previous section (e.g. Figure 2.3) are very different from the observed price behavior in the PJM market (Figure 2.6), the objective of this section is to modify the environment of the market in the MAS to correspond more closely to the actual conditions in PJM during the summer of 1999 (this summer was chosen because the market had a relatively simple single-settlement structure at that time and there were a relatively large number of price spikes). This modified environment for the MAS will replicate the type of volatile price behavior seen in the summer of 1999. In addition, it will also be the base case for evaluating a number of alternative policies for mitigating high prices that will be discussed in Section 5.

2.4.1 Using the Observed Pattern of Load in PJM

The patterns of load used in the previous section are highly stylized because they were designed in a simple series of simulations from a fixed load of 40GW to a stochastic load following a seasonal cycle. In particular, the final levels of load in Scenarios 4 and 5 are very low to keep the average load for all periods equal to 40GW (see Figure 2.3). In a typical market like PJM, the levels of load before and after the summer months are similar, and the levels of load during the summer are higher. The first modification of the MAS environment is to use the actual levels of load in PJM for the summer of 1999 shown in Figure 2.6.

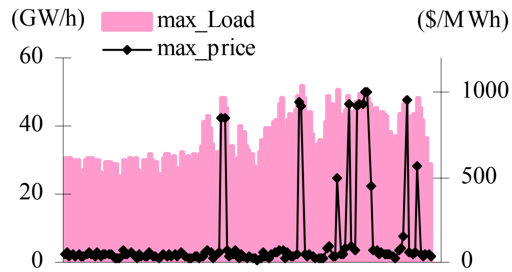

Fig. 2.6 Actual Load and Peak Price in the PJM Market (April to September 1999)

Using the PJM pattern of load in the MAS and the same six identical firms (latent speculators) from Section 3, the average market price is \$588/MWh and price spikes occur 57% of the time (see Table 3). These values are much higher than the observed average price of \$129/MWh in PJM with price spikes occurring only 10% of the time. Hence, using the stylized load patterns in Section 3 is not a sufficient explanation of why high prices are so persistent in the simulations. Furthermore, the size of each firm in the MAS is not unrealistically large compared to the largest firms in PJM. Hence, other features of the market must account for the difference in price behavior, and the next modifications to the market environment change the structure of the firms in terms of both their sizes and objectives.

2.4.2 Too Many Firms are Required to Make Textbook Competition Practical

The standard remedy for dealing with market power in an economics textbook is to increase the number of firms in the market. This is relatively easy to do in the MAS, and the amount of generating capacity owned by a firm can be scaled to keep the total installed capacity constant at 60GW. The basic rationale is that smaller firms have less incentive to speculate because setting a high price with a marginal block of capacity affects a relatively small number of MW of their own capacity. The higher probability of having the marginal block rejected when speculating, and the resulting foregone profits on that block, are harder to justify for a small firm. The basic question is this: how many firms are needed to make an electricity market competitive?

Table 2.3. Average Price for Different Numbers of Firms (Competitive price is \$39/MWh)

	Number of Firms				
	6	10	12	15	20
Ave	\$588	\$420	\$398	\$178	\$51
%*	57%	39%	37%	14%	<1%

* percentage of price spikes (>=\$450)

In the MAS simulations, the average price is only \$39/MWh in a competitive market (when the offers are equal to the low marginal costs used in E1 and E3 in the previous section). The corresponding average prices, and the percentage of prices above \$450/MWh, are summarized in

Table 2.3 for different numbers of firms. It is clear from these results that over 20 firms are needed to make the market truly competitive and eliminate price spikes. With 20 firms, the average price of $51/MWh is still 30% above the competitive price. It should be noted that the firms are specified as latent speculators, and lower prices would occur if they were specified as initial price takers. However, the results in Section 3 show that firms will eventually learn to speculate, and in reality, traders know that high prices can occur due to events like the California crisis in 2000-01. As a result, it is misleading to evaluate a market using results that are dependent on firms believing that high prices cannot occur.

The overall conclusion is that a large number of firms (>20 in this example) is needed to make an electricity market competitive when the load is stochastic. This result is consistent with the conclusions reached by Rudkevich et al. (1998). It is unlikely that regulators would be willing to impose a rule that no firm can own more than 5% of total capacity in a region. While adding new firms does help to mitigate high prices, it does not provide a practical solution on its own.

2.4.3 The Effect of Including Vertically Integrated Firms is Relatively Small

One explanation of why there are relatively few price spikes in PJM is that some of the largest firms are Vertically Integrated Firms (VIF). A VIF has a regulatory mandate to meet some load and to receive a fixed rate for these sales. Consequently, high spot prices do not affect the revenues from these sales, and there is less incentive for a VIF to speculate. However, under our specifications, a VIF is still a latent speculator (i.e., believes that high prices are possible), and can speculate with any generating capacity that is not used to meet the regulated load. The regulated load of a VIF is proportional to the system load, and as a result, it varies from period to period and is quite different from a fixed quantity contract. Two different levels of regulated load are specified in terms of the "fair" share load with six firms (i.e., 1/6 of the system load). The two levels are 100% and 80% of the fair-share load.

Table 2.4 summarizes the average prices for different numbers of VIF in a market with six firms of the same size. High prices are mitigated effectively when five firms are VIF at the 100% level. When all six firms are VIF at the 80% level, the average price is still over twice as high as the competitive price. These results are surprising and show that speculative behavior is hard to mitigate. There is, however, an important change in price behavior that occurs when the number of VIF increases. Most price spikes occur when the load is high, and as a result, price behavior corresponds to frequent price spikes rather than to persistent high prices. This is a more realistic market environment in which low prices can occur after price spikes.

Table 2.4 Average Price for Different Numbers of VIF (Competitive price is $39/MWh)

Fair share %		Number of VIFs				
		2	3	4	5	6
100%	Ave	$445	$430	$305	$31	$26
	%*	41%	41%	30%	4%	0%
80%	Ave	$473	$452	$341	$91	$84
	%*	45%	44%	32%	6%	5%

* percentage of price spikes (>=$450)

It is possible that the VIF in PJM behave like price takers and do not speculate. This type of behavior would help to further mitigate high prices. However, markets should be designed to withstand speculative behavior, and dependence on "good" behavior by some firms is not a satisfactory way to define a robust market. Nevertheless, in some markets like New York, good behavior by the power authorities, for example, is probably an important feature that does mitigate high prices. The overall conclusion is that having a few VIF in a market does help to mitigate high prices, but it is unlikely that it is the primary reason why the PJM market has relatively few price spikes compared to the MAS results.

2.4.4 Replicating the Structure of Firms in PJM

It is possible to maintain the basic structure of six firms in the MAS and to modify it to correspond more closely to the structure of firms in PJM. The objective is not to replicate the PJM structure exactly by, for example, using estimates of the actual costs for each firm (e.g. the approach used by Mansur (2003)). Instead, the objective is to make only a few changes to the structure of firms to keep the results more generic. The basic criterion is to generate realistic price behavior with the number of price spikes roughly the same as the number observed in PJM. This structure will then be used in the next section as the base case for the policy analysis.

Compared to the simulations presented in Section 3, the changes in the market environment are 1) using the actual pattern of load in PJM during the summer of 1999 (discussed in Section 4.1), 2) redefining the small capacity blocks for each firm to have equal sizes (the 10GW of capacity are now divided into six blocks: 5,1,1,1,1,1), and 3) setting the market price to the price cap when insufficient capacity is offered into the auction to meet load (instead of the highest offer). This latter change corresponds to the rule used in PJM, but it's effect on the results is relatively small.

The simulations in Section 3 use six firms, and each firm has 10GW of capacity and the same objective of maximizing expected profits. The main changes made to this structure of firms are 1) two of the six firms are now VIF (100% of the fair-share load), and 2) two of the firms are each replaced by four identical small firms (with only 2.5GW of capacity), and therefore, behave more like price takers. This modified structure of firms is similar to the actual structure of firms in PJM in 1999. No changes are made to the structure of costs of the firms (other than changing the sizes of capacity blocks).

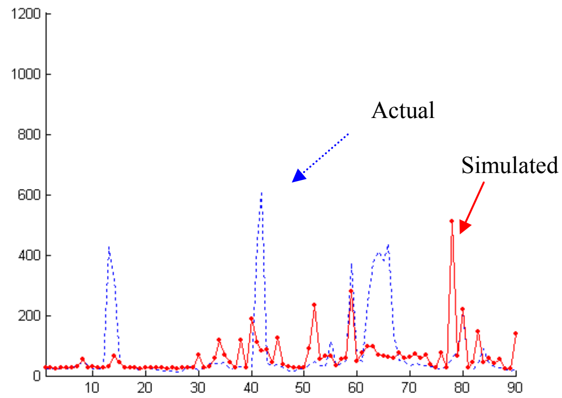

Fig. 2.7 Actual Market Prices and Simulated Market Prices Using Naïve Agents for Firms

The market prices from the MAS and the actual maximum daily prices in PJM in the summer of 1999 are shown in Figure 2.7. The average MAS price is \$62/MWh, which is only half the observed price of \$129/MWh in PJM. Price spikes above \$150/MWh occur 6% of the time in the MAS compared to 15% in PJM. (From this point on, the definition of a price spike is reduced from \$450/MWh to \$150/MWh because there fewer price spikes in the comparisons and almost none over \$450/MWh.) However, at the beginning of the simulation, the firms are relatively naïve about the possibility of price spikes (i.e., they correspond to "Latent Speculators", defined in Section 3). The firms learn to speculate more effectively during the simulation.

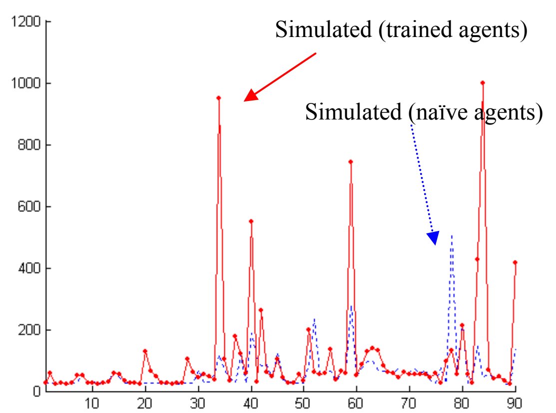

Fig. 2-8 Simulated Market Prices Using Naïve and Trained Agents for Firms

To illustrate the effects of learning, the estimated parameters of each firm's residual demand curve after the final period of the simulation in Figure 2.7 are used as the initial values for a new simulation with the same pattern of load. The new initial values imply that the firms have been trained in the market. Figure 2.8 compares the two sets of simulated market prices using the naïve and trained agents. With trained agents, the average price is now $103/MWh and price spikes occur 11% of the time. These values are more like the actual values observed in PJM because the first price spike occurs much earlier in the summer with trained agents. Since we have argued above that it is unrealistic to assume that profit-maximizing firms will not speculate in electricity markets, it is reassuring to find that the MAS, using the modified structure of firms with trained agents, closely replicates the type of price volatility seen in the PJM market. Therefore, we are now confident that this final structure of the MAS could be used effectively to evaluate different ways of mitigating high prices.

2.5. Conclusions

The results presented in Sections 3 and 4 demonstrate that a MAS can be designed to simulate realistic price behavior in a deregulated wholesale market for electricity, and this form of MAS has the potential for evaluating the effects of modifying the structure of a market on price behavior. The analysis also provides new insights into why "hockey stick" supply curves and price spikes are observed in many deregulated electricity markets. The use of a structured form of Adaptive Autonomous Agent to represent suppliers made it possible to keep the number of firms in the market and the number of trading periods relatively small. Consequently, the MAS results were comparable with the results obtained from the standard market experiments conducted at Cornell using PowerWeb. In other words, the results using computer agents and human subjects can be compared, and an important objective for future research is to use the MAS to supplement the results of tests using human subjects. In addition to replicating observed behavior, it is also possible to use a MAS to test the sensitivity of market outcomes to many different modifications to the structure of a market. With human subjects, there are severe practical limits on the number of replications for one market structure and on the number of different modifications that can be tested.

The initial structure of the MAS uses six identical software agents to represent supply firms in a uniform price auction run by an ISO. Patterns of load are exogenous, and there are 135 trading periods in each simulation. The firms use observed market outcomes to estimate the residual demand curve for their generating capacity in the next trading period. These estimates are updated after each trading period using a Kalman filter, and as a result, the behavior of a firm can change substantially during a simulation. The objective of a firm is to maximize the expected profits in the next trading period. The analysis in Section 3 focuses on creating price spikes and the analysis in Section 4 shows why the structure of firm sizes and firm types is important for replicating observed price behavior with infrequent price spikes.

With a relatively small number of firms, it is very difficult to avoid getting persistent high prices when load is high and stochastic. In a deterministic optimization with a fixed load, six suppliers are relatively competitive for all 135 trading periods. These results from the MAS are consistent with previous experimental results with human subjects using PowerWeb. Varying load provides firms with more information about the market. Once a firm realizes that the residual demand curve is not flat (i.e., not perfectly price elastic), offers become more speculative with very high offers on marginal blocks of capacity. Uncertainty about the load in the next period opens up opportunities for speculation, and firms submit offer curves shaped like a hockey stick. This is exactly the type of behavior observed in many deregulated markets. Once firms learn to speculate, they remain as latent speculators even if loads and market prices are consistently low. The learning process is essentially in one direction from price taker to speculator.

Given the characteristics of a typical deregulated market, it is rational for some firms to submit high offers. However, regulators have tended to treat these high offers as unacceptable behavior, particularly in the USA. It would be more realistic to accept the inevitability of hockey stick supply curves, and to look for other ways to make a market more robust to speculation. From the first principles of economic theory, the structure of the ownership of firms in a market is a key determinant of how well a market works.

Using the daily pattern of loads in the PJM market for the summer of 1999, the simulations with six identical firms result in persistent high prices and not to a pattern of infrequent price spikes observed in the PJM market. The high prices are substantially above competitive levels. In order to bring prices down to competitive levels, the results in Section 4 show that it is impractical to rely solely on 1) increasing the number of firms, or 2) changing profit maximizing firms to vertically integrated firms (i.e., sell a substantial portion of their load to customers at a regulated rate). Competitive prices do occur if there are more than 20 identical firms (profit maximizers) or if five of the six firms are vertically integrated. Both of these criteria are far too extreme to be acceptable for regulatory purposes.

It is possible to replicate observed price behavior in the PJM market by modifying the initial structure of six firms. This modified structure of firms 1) converts two of the six firms to vertically integrated firms, and 2) divides another two firms into eight small firms. Even though these small firms maximize their expected profits, they behave like price takers. This modified firm structure is similar to the structure in the PJM market, and the simulation replicates the observed price behavior with infrequent price spikes. The presence of small firms makes the market more robust to speculation by the large firms because the small firms tend to submit all of their capacity into the auction when the forecasted load is high. In contrast, it is rational for large firms to withhold some capacity to make it more likely that a price spike will occur.

The current approach favored by many regulators in the USA for mitigating high prices is to make it more difficult for firms to submit high price offers. Eliminating high price offers is an effective way of eliminating price spikes, but the results may be counterproductive. The primary regulatory mechanism is to require that high price offers must be cost justified. For example, the Federal Energy Regulatory Commission (FERC) proposed this approach when FERC intervened in California in the fall of 2000. The problem, however, is that this policy creates an obvious incentive for firms to exaggerate operating costs, and this opens the door to endless litigation. A self-regulating market disciplines firms that submit high price offers by reducing the amount of their capacity that is dispatched. For example, making load more price responsive is a highly effective way to discipline speculative behavior in a uniform price auction. In some circumstances, regulatory intervention may be needed when, for example, a specific power plant is essential for providing ancillary services. In other situations, regulating how firms behave in the auction is an indication that the structure of the market is flawed in some way. As recent experience in the Australian market has shown, high price spikes can still be consistent with low average prices. Furthermore, a few high price spikes may be necessary to attract investment in new peaking capacity.

2.6. References

[1] Bernard, J. C., Schultz W. and Mount, T.D. (2003), Comparing the Multiple-Unit Vickrey and Uniform Price Auctions in a Repeated Procurement Setting, Working paper, Cornell University.

[2] Borenstein, S., J. Bushnell and F. Wolak (2002). “Measuring Market Inefficiencies in California’s Restructured Wholesale Electricity Market,” *American Economic Review*, December 92(5), pp. 1376-1405.

[3] Bushnell, Jim (2003). “Looking for Trouble: Competition Policy in the U.S. Electricity Industry,” Center for the Study of Energy Markets (CSEM) Working Paper 109, University of California Energy Institute, Berkeley, California, April. http://www.ucei.org.

[4] Bushnell, Jim (2004). “California’s Electricity Crisis:A Market Apart?,” Center for the Study of Energy Markets (CSEM) Working Paper 119, University of California Energy Institute, Berkeley, California, November. http://www.ucei.org.

[5] Federal Energy Regulatory Commission (FERC), (2003). “Final Report on Price Manipulation in Western Markets,” Staff Report on Western Markets Investigation, Docket No. PA 02-2-2000, March 26.

[6] Federal Energy Regulatory Commission (2002). “Initial Report on Company-Specific Separate Proceedings and Generic Reevaluations; Published Natural Gas Price Data; and Enron Trading Strategies: Fact-Finding Investigation of Potential Manipulation of Electric and Natural Gas Prices,” Docket No. PA02-2000.

[7] Federal Energy Regulatory Commission (FERC), (2000). “Commission Proposes to Reshape California’s ‘Seriously Flawed’ Electricity Markets with Sweeping Changes,” News Release, November 1.

[8] Klemperer, Paul and Meyer, Milgrom (1989). “Supply Function Equilibria in Oligopoly Under Uncertainty.” *Econometrica,* Vol. 57, No. 6, pp.1243-1277.

[9] Kolstad, J.T. and F.A. Wolak. (2003) “Using Environmental Emissions Permit Prices to Raise Electricity Prices: Evidence from the California Electricity Market, Stanford University, Stanford, California. March 10, [online]: http://www.stanford.edu/~wolak.

[10] Lillis, Kevin (1997), “Electricity Reform Abroad and U.S. Investment: The England and Wales Power Pool”, [online]:
http://www.eia.doe.gov/emeu/pgem/electric/ch212.html.

[11] Mount, Timothy (2000). “Strategic Behavior in Spot Markets for Electricity when Load is Stochastic.” Presented as a selected paper at the Hawaii International Conference on System Science, Maui.

[12] Mount, T. D., C.A Vossler, R.J. Thomas, and R.D. Zimmerman (2003). “The California Electricity Crisis: An Experimental Investigation of the Effects of Changing the Market Rules,” Working Paper, Presented at Eighth Annual POWER Research Conference, University of California, Berkeley, California. March 14.

[13] Neenan, B.F., D.K. Pratt, P.A. Cappers, R.N. Boisvert and K.R. Deal (2002). “NYISO Price-Responsive Load Program Evaluation Final Report,” NISCO Report prepared by Neenan Associates, LLC, January.

[14] Oh, Hyunga (2003). “Simulation Methods for Modeling Offer Behavior and Spot Prices in Restructured Markets for Electricity”. Doctoral thesis, Cornell University.

[15] Oren, Shmuel (2001). “Market Based Risk Mitigation; Risk management vs. Risk avoidance”, Presented at a White House OSTP/NSF Workshop on Critical Infrastructure Interdependencies.

[16] Rudkevich A., Duckworth M., Rosen R (1998). “Modeling Electricity Pricing in a Deregulated Generation Industry: The Potential for Oligopoly Pricing in a Poolco.” *The Energy Journal* 19(3): 19-48.

[17] Singh, Harry and Jacobs, Johnathan (2000), “ISO Markets and Capacity Products”, Presented in the fifth annual Pserc conference.

[18] Weiss, Gerhard (1999). Multiagent systems: A Modern Approach to Distributed Artificial intelligent. The MIT Press.

[19] Wolak, F. A. and R. Nordhaus (2000). “An Analysis of the June 2000 Price Spikes in the California ISO’s Energy and Ancillary Services Markets,” Market Surveillance Committee (MSC) of the California Independent System Operator (ISO), September 6.

[20] Wolak, F.A. (2003a). “Lessons from the California Electricity Crisis,” Center for the Study of Energy Markets (CSEM) Working Paper 110, University of California Energy Institute, Berkeley, California, April. http://www.ucei.org.

[21] Wolak, F.A. (2003b). “Measuring Unilateral Market Power in Wholesale Electricity Markets: The California Market 1998 to 2000,” American Economic Review, May, forthcoming, [online]: http://www.stanford.edu/~wolak.

3. Using Software Agents to Supplement Tests Conducted by Human Subjects: How Do Forward Contracts Affect the Behavior of Suppliers in a Wholesale Electricity Market?

Abstract

The objective of this chapter is to investigate how spot prices are affected by forward contracts using experimental economics and agent-based simulations. The results show that holding a forward contract is an effective way to mitigate high prices if the same contract is held for all trading periods and the price of this contract is independent of the spot prices. However, when a forward contract is renewed periodically and the spot prices influence the forward price, there is more speculation and the spot prices may be higher than the base case with no forward contracts. The simulation results also show that software agents were able to replicate the behavior of students effectively in the experiments.

3.1 Introduction

High spot prices during the energy crisis in California (2000/01) provided an incentive for regulators to introduce a number of different ways to mitigate high prices in other restructured markets for electricity. One way, that was identified explicitly when the Federal Energy Regulatory Commission (FERC) intervened in the Californian market during the crisis, was to encourage suppliers and the incumbent utilities to hold forward contracts. A forward contract guarantees a certain level of earnings for a supplier and reduces the market risk of relying on the spot market for a buyer (primarily the incumbent electric utilities in California). As Carlton (1979) argued, a buyer must rely on the spot market unless the level of demand is known in advance and covered by forward contracts. During the energy crisis in California, the spot market was extremely volatile and high prices persisted. Since most of the incumbent utilities were paid relatively high regulated rates by customers, they did not think it was necessary to hold forward contracts with suppliers as well. However, when the spot prices were persistently higher than the regulated rate, one major utility (Pacific Gas and Electric) became bankrupt.

In the middle of the crisis period, the Market Surveillance Committee (MSC) of the California Independent System Operator reached the following conclusion (September 2000):

"The June 2000 price spikes can be attributed to several factors. The primary cause was the lack of sufficient forward energy and ancillary service purchases for the month of June 2000 by utility distribution companies…. There are significant benefits to load-serving entities from purchasing forward financial contracts. First, forward market purchases limit the spot price exposure faced by a load-serving entity. It is subject to spot price risk on its real-time energy requirements only to the extent that they differ from its forward market purchases. The second source of benefits is that forward market purchases from a generation unit owner limit the incentives this supplier has to exercise market power in the spot market. …The Independent System Operator argued that unrestricted forward contracting by the utility distribution companies would have substantially mitigated the market power in the past and would do so in the future."

In other words, these regulators believed that forward contracts would reduce the incentives for suppliers to speculate and exploit market power in an electricity auction (Wolak and Nordhaus 2000a; Wolak 2000b; Wolak 2003a; Wolak 2003b; Allaz and Vila, 1993). The implication is that spot prices should be lower when suppliers hold forward contracts, and that unrestricted trading in the forward market should eliminate speculation.

A countervailing argument is that market power in the spot market results in higher forward prices. For example, Harvey and Hogan (2000) argued that a forward contract would be an effective tool to mitigate market power if demand is price responsive and the forward price is not influenced by the spot price. Otherwise, the attractions to a supplier of holding a long-term forward contract substantially decrease. Further, they argued that the simple fact that a hedging contract reduces exposure to spot market prices is not sufficient to conclude that hedging contracts would reduce the average price paid for electricity. Using the actual offer data in the UK market, Wolfram (1999) also argued that the effects of forward contracts on restraining spot prices were quite limited. Mount and Lee (2003) explained why high forward prices were paid in California during the crisis. The risks faced by buyers and sellers (generators) in a deregulated electricity market are very asymmetric. Suppliers hold call options because they are not obligated to generate if the spot price is lower than their operating cost but, in contrast, buyers, with an obligation to serve load, have no such protection. Under these circumstances, buyers were willing to pay high risk premiums in forward contracts to stabilize the cost of their future purchases. Risk premiums are high since suppliers have less incentive to commit production to forward contracts when high prices persist in the spot market (Carlton 1979). In other words, the long-term forward contracts may not mitigate high spot prices when spot prices are abnormally high.

The objective of this chapter is to use experimental economics to test how suppliers change their strategies when they hold a forward contract and how the spot price is affected. One set of tests uses students to represent suppliers in an electricity auction with 1) no forward contracts, 2) permanent forward contracts (i.e. the same contract is held for all trading periods and the price of this contract is independent of the spot prices), and 3) renewable forward contracts (i.e. a forward contract is renewed periodically and spot prices influence the forward price). In the latter test, the price of a new forward contract is affected by conditions in the spot market. An identical set of tests was also conducted using software agents (i.e. artificial intelligence) to represent all of the suppliers. The objective is to demonstrate that software agents can be used effectively to test electricity auctions, do additional sensitivity tests and supplement the results obtained using humans.

The results using software agents were encouraging. In the first set of tests, two students competed in each market with four software agents. In almost all cases, the average earnings of the software agents were higher than the average earnings of the students. In the tests with all software agents, two software agents replaced the students, and these agents (latent speculators) were more likely to speculate than the other software agents. The average spot prices and average earnings with all software agents corresponded closely to the highest values obtained by the students. In both tests, holding a forward contract is an effective way to mitigate high prices if the forward contract is permanent. However, there is more speculation and the spot prices are higher when a forward contract is renewable. The tests were run on a computer platform, *PowerWeb*, developed by researchers in the Power System Engineering Research Center (PSERC), that replicates the physical constraints of meeting loads on an electrical grid and determines the nodal spot prices in a typical uniform price auction for electricity.

3.2 Optimum Offer Behavior in a Spot Market

3.3 A Supplier's Expected Profit

This section shows how a supplier's optimum offers for different generators are affected by holding a forward contract. In a wholesale spot market, the framework assumes that suppliers possess some market power, and therefore, that offers influence the spot price. In addition, since forward prices reflect expectations about future spot prices, the current spot price influences forward prices. Let P_t^S denote the wholesale spot price at time t, Q_t^S denote the quantity sold by supplier i in the spot market, N denote the duration of the forward contract, Q_t^F denote the quantity that supplier i has previously agreed to deliver in a forward contract at a forward price $\overline{P}^F$ and P_t^F denote the forward price for a new contract with delivery in periods *t+1, t+2,…,t+N*. There are several simplifying assumptions: $Q_t^F = \overline{Q}^F$ is predetermined; the cost function is linear, $C(Q) = c \cdot Q$; all suppliers are risk neutral; the discount rate is zero.

Using the assumptions made above, the expected profit of a supplier with a forward contract, for a predetermined quantity $\overline{Q}^F$ and price $\overline{P}^F$, is the sum of 1) the expected profit from the spot market at time *t*, 2) the fixed profit from the contract sale at time *t*, and 3) the expected profit from renewing the contract. The third term is included because spot prices may influence forward prices. The contract terminates at time *T* and is automatically renewed for N periods. The expected profit at time $t < T$ is:

$$(1)\quad \pi_t = Q_t^S (P_t^S - c) + \overline{Q}^F (\overline{P}^F - c) + N\overline{Q}^F (P_T^F - c)$$

If the supplier does not hold any forward contracts (i.e. $\overline{Q}^F$ =0), the expected profit simplifies to $P_t^S (Q_t^S - c)$.

P_t^F is the forward price at time *t* for buying a new N-period contract starting in period *t+1*. It is specified to be an exponentially weighted moving average of the current and past spot prices:

$$(2)\quad P_t^F = \lambda \cdot \sum_{j=0}^{\infty} (1-\lambda)^j \cdot P_{t-j}^S$$
$$= (1-\lambda) \cdot P_{t-1}^F + \lambda \cdot P_t^S$$

where $0 \le \lambda \le 1$ is a smoothing constant. The suppliers observe P_t^F after the spot market at time t has cleared. For T > t, the forward price at time t for delivery in periods T+1 to T+N is:

$$(3)\quad P_{T|t}^F = (1-\lambda)^{T-t+1} P_{t-1}^F + \lambda \sum_{j=t}^{T} (1-\lambda)^{T-j} P_j^S$$

If the value of λ is zero, the forward price ($P_{T|t}^F$) is not influenced by the current spot price. If $0 < \lambda \le 1$ and a forward contract is renewed at period T, the forward price is influenced by the current spot price, P_t^S, and the coefficient for P_t^S is $\lambda \cdot (1-\lambda)^{T-t}$.

Earlier research (Mount and Oh, 2004) has shown that the observed behavior of spot prices can be predicted by a supplier using an inverse form of residual demand curve:

(4) $P_t^S = \dfrac{1}{\beta_{1,t} + \beta_{2,t} \cdot (Q_t - D_t)}$

where $Q_t = Q_t^S + \overline{Q}^F$, D_t is the load and the two parameters ($\beta_{1,t} > 0, \beta_{2,t} > 0$) are updated by a Kalman filter, using realized sales and observed spot prices. From (4), Q_t can be written as follows:

(5) $Q_t = D_t + ({}^{1}\!/\!{}_{P_t^S} - \beta_{1,t}) / \beta_{2,t}$

3.4 The Optimum Offer

The optimum offer P_t^o can be derived by differentiating (1) with respect to P_t^S . The first-order conditions for maximizing the expected profit imply that the optimum offer prices are:

- No forward contract ($\overline{Q}^F$ = 0):

(6) $P_t^o = \sqrt{\dfrac{c}{\beta_1 - \beta_2 \cdot D}}$

- Permanent forward contract ($\overline{Q}^F$ >0, $T = \infty$):

(7) $P_t^o = \sqrt{\dfrac{c}{\beta_1 - \beta_2 \cdot (D - \overline{Q}^F)}}$

- Renewable forward contract ($\overline{Q}^F$ >0, $N < \infty$):

(8) $P_t^o = \sqrt{\dfrac{c}{\beta_1 - \beta_2 \cdot (D - \overline{Q}^F + \lambda(1-\lambda)^{T-t} N\overline{Q}^F)}}$

Comparing (6) and (7), the optimum offer price is always lower with a permanent contract than it is with no contract since $\beta_2 \cdot \overline{Q}^F > 0$. This implies that a supplier with a permanent forward contract has less incentive to speculate in the spot market. These results are consistent with Wolak (2000).

The optimum offer price with a renewable contract is more complicated. Comparing (7) and (8), the optimum offer price is always higher with a renewable forward contract than it is with a permanent contract since $\lambda(1-\lambda)^{T-t} \beta_2 N\overline{Q}^F > 0$. However, comparing (6) and (8) shows that the optimum offer price with a renewable contract is higher than it is with no contract if N is large, T-t is small and the following condition holds.

(9) $\lambda \cdot (1-\lambda)^{T-t} > \dfrac{1}{N}$

Since the maximum of the LHS of (9) is 0.25 (when λ =0.5 and $t = T - 1$), contracts with N<5 always reduce speculation compared to (6).

The following example demonstrates how a supplier changes her behavior by holding a forward contract: β_1 =0.0413, β_2 =0.000067, *D*=526MW, c=$45/MWh, Q^F =50MW and the installed capacity of a supplier is 150MW/h. A price cap of $100/MWh is imposed in the market. The optimum offers (P_t^o, Q_t^o) with both no contract and a permanent contract are unchanged by the number of periods to renewal, *T-t*. Assume T=10 and t=7. With a permanent contract of 50MW, the spot price is 20% lower (P^o =$86/MWh→$69/MWh) and the quantity dispatched is 33% higher (Q^o = 83MW/h → 125 MW/h) than it is with no contract.

Table 3.1 The Optimum Offers in the Spot Market

	No Contract		Permanent Contract		Renewable Contract	
T-j	P*	Q*	P*	Q*	P*	Q*
9	86	83	69	125	72	118
7	86	83	69	125	74	112
5	86	83	69	125	78	101
3	86	83	69	125	88	80
1	86	83	69	125	100	59

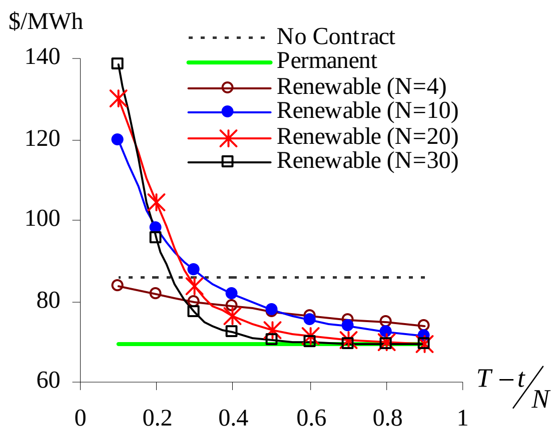

Fig. 3.1 Optimum Offer Price with Contracts

The optimum offer price with a renewable contract in Table 3.1 increases as the current period approaches the renewal period (t→T). When T-t < 4, the optimum offer price is higher than it is with no forward contract and a supplier will exploit market power and speculate more aggressively. This type of behavior is presented in Figure 3.1 for different values of N. The RHS of Figure 3.1 (T-t→N) supports Allaz and Vila (1993) who argued that spot prices tend towards a competitive equilibrium when N is large. However, when t approaches the renewal period T (the LHS of Figure 3.1), the optimum offer price with a large N is higher than it is with a small N. In other words, the situation is reversed and there is more speculation when N is large.

According to the MSC September Report (2000), a larger forward commitment will make a supplier less aggressive in the spot market. The MSC, therefore, expects lower spot prices when the quantity of forward contracts increases. Harvey and Hogan (2000) challenged this conclusion by asking the question: if suppliers have market power, why would they voluntarily surrender that market power in a forward contract? In contrast to the MSC, they expected suppliers to exploit market power and raise the forward price.

Table 3.2 shows that the MSC is correct if a supplier holds a permanent contract (and when the number of periods to the renewal of a contract is large). However, if the renewal period is close at hand, a supplier speculates more aggressively with a large contract than with a smaller contract. This happens because the effect of the current spot price on the future profits from the forward contract in (1) is large enough to compensate for lower profits in the spot market.

Table 3.2 Impact of Contract Quantity on Optimum Offers

T-t	Q	No Contract		Permanent Contract		Renewable Contract	
		P_t^o	Q_t^o	P_t^o	Q_t^o	P_t^o	Q_t^o
	30			75	109	78	101
7	50	86	83	69	125	74	112
	80			63	147	68	128
	30			75	109	93	71
2	50	86	83	69	125	98	62
	80			63	147	100	59

3.5 The Experimental Framework

All of the experiments were conducted using *PowerWeb* (an interactive, distributed, Internet-based simulation platform developed by PSERC researchers at Cornell University) to test different electricity markets using human decision makers and/or computer agents. An Independent System Operator (ISO) determines the optimum dispatch of generators and the spot (nodal) prices paid to suppliers. The *PowerWeb* environment is designed to run unit commitment and optimal power flow routines to minimize the cost of meeting load subject to the physical constraints of an AC network. However, for our experiments, the network constraints are not binding, and in each trading period, the same spot price is paid to all suppliers using a uniform price auction (last accepted offer).

After each trading period, the ISO announces a forecast of the load in the next trading period. Load is completely price inelastic but it does vary from period to period. The forecasted load is generated randomly using a uniform distribution from 430 MW to 550 MW. The actual load is also generated randomly using a uniform distribution (Forecast ± 20 MW). The average load is 82% of the total installed capacity, which corresponds to realistic conditions in the summer when the load is relatively high.

For each trading period, each supplier submits offers to sell (or withhold) five blocks of capacity into the auction. A price cap (maximum offer allowed) of $100/MWh is enforced by the ISO (to keep the payments to participants in the test reasonably low). If the total capacity offered into the auction is less than the actual load, the ISO covers the shortfall by purchasing expensive imports from another market. However, when a shortfall occurs, the spot price is set by the highest offer and not by the price of imports.

Each supplier owns five blocks of generating capacity with capacities 50, 20, 10, 10, 10 (MW) and production costs 20, 40, 48, 50, 52 ($/MWh generated), respectively. In addition, there is a fixed standby cost of $5/MW to cover the opportunity cost of being available when a block is offered into the market. Withholding a block from the auction is the only way to avoid the standby cost for that block. There is also a fixed cost charged each period to cover capital costs ($1200/period, to make earnings roughly equal to profits in excess of competitive levels). These capacity and cost structures are the same for all six suppliers and they remain the same in all of the markets tested.

Three different markets were tested during the fall semester, 2003 using 20 students majoring in applied economics or electrical engineering (the tests were part of a course on electricity markets). Each student represented a supplier in the market. In addition, some of the suppliers were represented by computer agents. The three tests were 1) no forward contracts, 2) a permanent forward contract, and 3) a renewable forward contract.

In a recent paper, Mount and Oh (2004) have shown that computer agents can replicate the type of volatile price behavior seen in the PJM market during the summer of 1999. Each agent represented a supplier and estimated the residual demand curve for their capacity in the market. This estimate was updated each period using a Kalman filter. An important requirement for generating realistic prices was having an appropriate structure of firms in the market because six identical firms resulted in too many price spikes.

In our experiments, two students compete with four agents. Two of these agents represent Vertically Integrated Firms (VIF) that have to meet a fixed proportion of load and are paid a regulated price (=$60/MWh) for this load. These firms have less incentive to speculate than the other firms. The other two agents are "Initial Price Takers (IPT)" (believe that price spikes can not occur). However, initial price takers can learn to speculate if high prices do occur, and in this sense, these two agents reinforce the behavior of the students. If the students do not speculate, none of the computer agents speculate, but if the students speculate, the agents learn to speculate and make the market easier to exploit.

An identical set of tests was also conducted using computer agents to represent all six suppliers by replacing the students by two "latent speculators" (believe price spikes can occur). These agents are more likely to speculate than initial price takers, and initial price takers evolve into latent speculators if there are high prices. The primary objective of the tests with all agents was to determine how well the computer agents can replicate the typical offer behavior of the students.

The students in the tests represented experienced traders, and they received an initial briefing about how suppliers behave in the PJM market. Hence, the students understood the rationale for speculating, and why hockey-stick offers cause price spikes. The first test (no contract) consisted of 25 trading periods, and the next two tests, conducted a week later, consisted of 20 trading periods each. Each student was paid in direct proportion to her cumulative earnings and told that the objective of the tests was to earn as much money as possible. The results from the initial trading periods were treated as a learning period for developing an offer strategy, and the average results from the last 10 periods in each test were used in the analysis.

3.6 Analysis of The Test Results

The Experimental Design

In Test 1, none of the six suppliers holds a forward contract, but the two vertically integrated firms (agents) must meet one sixth of the load at a predetermined price. In Test 2, regulations require that each student must hold a forward contract for half of her capacity, and has already signed a contract for 50MW (the first block of capacity) at a fixed price of \$60/MWh. These contracts are in place for all periods. Hence, the objective is to maximize the profits from selling the remaining four blocks of capacity (the first block is submitted automatically). In all other respects, the conditions are the same as Test 1.

In Test 3, each student has to renew a 10-period forward contract of 50MW every 10th period during the test. The forward price is given in (2) with λ=0.25 and a random residual added. The value of λ is not a priori information for the suppliers. The computer agents are designed to estimate λ based on the previous spot and forward prices. Simulation results show that the agents' estimates of λ are accurate after 3 periods. In Test 2 and Test 3, the students were paid for the forward contracts as well as for the earnings in the spot market. In Test 3, the first contract price was set at \$60/MWh for periods 1 to 10, and this contract was renewed in periods 10 and 20. The students' earnings were computed to reflect the forward prices in the two new contracts and not the initial contract (the actual revenue from the first two contracts was augmented by $50 \cdot 10 \cdot (P_{20}^{F} - 60)$. The reason for doing this was to provide the students with the same incentive to increase the forward price of a new contract throughout the test.

The Regression Model

The results from the three tests were summarized using the cumulative earnings for each supplier during the final 10 periods of each test. For Test 3, the revenue from the forward contract in periods 11 to 20 was revised to reflect the average revenue from the last two contracts ($50 \times 10 \times P_{10}^{F}$ is replaced by $50 \times 10 \times (P_{10}^{F} + P_{20}^{F})/2$). In most cases, $P_{20}^{F} > P_{10}^{F}$ because the students' ability to speculate successfully improved during the test. For each test, there are 10 sessions with six suppliers in each session (two students, two VIF agents and two IPT agents). The lowest profit session in each test is excluded so that the results correspond more closely to the behavior of the professional traders in real markets.

Since the primary objective of the analysis is to determine how forward contracts affect the behavior of the students, the regression model is specified to make it easy to test hypotheses about the earnings of the students as follows:

(10)

$$y_{ijk} = \mu + \sum_{i=2}^{3} \alpha_i M_i + \sum_{i=1}^{3} \sum_{j=2}^{3} \beta_{ij} F_{ij} + \sum_{i=1}^{3} \sum_{k=1}^{K-1} \gamma_{ik} S_{ik} + e_{ijk}$$

where y_{ijk} = log earnings for periods 11 to 20 for firm type *j* in session *k* of market *i.*

M_i = 1 for Test *i,* 0 otherwise

F_{ij} = 1 for Test *i* and Firm *j* , 0 otherwise

j = 1 for a student, 2 for a VIF agent and 3 for IPT agent

s_{ik} = 1 for Test *i* and Session *k*, 0 for Test *i* and Session not *k* or *K* and -1 for Test *i* and Session *K*

The dummy variables for Test 1 (α_1) and for students $\times$ markets (β_{11}, β_{21} and β_{31}) are set implicitly to zero, and the coefficients for sessions $\times$ markets sum to zero ($\sum_{k=1}^{K} \gamma_{ik} = 0$ for *i*=1,2,3). Consequently, the earnings of a student are represented by μ in Test 1, $\mu + \alpha_2$ in Test 2 and $\mu + \alpha_3$ in Test 3. β_{i2} measures the additional earnings of a VIF agent in market *i* compared to a student, and β_{i3} measures the same value for an IPT agent.

The five main hypotheses of interest in the analysis are as follows:

1. Permanent contracts reduce the earnings of students ($\alpha_2 < 0$).
2. Renewable contracts increase the earnings of students ($\alpha_3 > 0$).
3. For Test 1 (no contract), VIF agents earn less than students ($\beta_{12} < 0$) and IPT agents earn the same ($\beta_{13} = 0$).
4. For Test 2 (permanent contract), VIF agents earn less than students ($\beta_{22} < 0$) and IPT agents earn more ($\beta_{23} > 0$).
5. For Test 3 (renewable contract), VIF agents earn less than students ($\beta_{32} < 0$) and IPT agents may earn more or less than the students ($\beta_{33} = 0$).

Estimation Results

Table 3.3 summarizes the parameter estimates and the corresponding t-statistics for the model (10). The critical values of the t-statistic at the 5 percent level of significance (for 162 observations and 32 regressors) are 1.98 for a two-tailed test and 1.66 for a one-tailed test.

The estimated value of the intercept, $\hat{\mu}$, is 9.44, which measures the average earnings (in logarithms) of students in Test 1. Hypothesis 1 is supported by the data because α_2 is significantly less than zero. In other words, the average earnings of students are lower when they hold a permanent contract in Test 2. Hypothesis 2 is also supported by the data because α_3 is significantly greater than zero. The average earnings of students are higher in Test 3 when they hold a renewable contract and the level of past spot prices affects the forward price in a new contract.

For Hypothesis 3, the significant negative value of $\hat{\beta}_{12}$ is consistent with the hypothesis. The earnings of the VIF agents in Test 1 are less than the earnings of the students when they do not hold a contract. The second part of Hypothesis 3 is also supported by the data because $\hat{\beta}_{13}$ is not significantly different from zero. The earnings of students and the IPT agents are statistically the same in Test 1.

When the students hold permanent contracts in Test 2, both the VIF and IPT agents earn significantly more than the students ($\hat{\beta}_{22} > 0$ and $\hat{\beta}_{23} > 0$). This result for IPT agents supports Hypothesis 4 but the result for the VIF agents rejects it. Note that the objectives of the VIF and IPT agents in Test 2 are the same as in Test 1 because only the students hold the permanent contracts in Test 2. Hence, the results of Hypothesis 4 reinforce the results of Hypothesis 1. Holding a permanent forward contract reduces earnings in absolute terms and also in relative terms to the earnings of the VIF and IPT firms.

Table 3.3 Regression Results

Variable (Parameter)	Parameter Estimate	t-value	(Pr>|t|)
Intercept (μ)	9.441	252.33	<.0001
Market2 (α_2)	-0.242	-4.57	<.0001
Market3 (α_3)	0.374	7.07	<.0001
Market1VIF (β_{12})	-0.151	-2.85	0.0051
Market1IPT(β_{13})	0.024	0.46	0.6445
Market2VIF (β_{22})	0.130	2.46	0.0151
Market2IPT (β_{23})	0.510	9.64	<.0001
Market3VIF (β_{32})	-0.407	-7.69	<.0001
Market3IPT (β_{33})	0.141	2.66	0.0088
Session1_2 ($\gamma_{1,2}$)	0.013	0.22	0.8289
Session1_3 ($\gamma_{1,3}$)	0.086	1.41	0.1597
Session1_4 ($\gamma_{1,4}$)	0.094	1.54	0.1253
Session1_5 ($\gamma_{1,5}$)	-0.096	-1.57	0.1183
Session1_6 ($\gamma_{1,6}$)	0.206	3.36	0.0010
Session1_7 ($\gamma_{1,7}$)	-0.050	-0.82	0.4149
Session1_8 ($\gamma_{1,8}$)	0.138	2.26	0.0256
Session1_9 ($\gamma_{1,9}$)	-0.130	-2.12	0.0358
Session2_2 ($\gamma_{2,2}$)	0.099	1.62	0.1068
Session2_3 ($\gamma_{2,3}$)	-0.232	-3.80	0.0002
Session2_4 ($\gamma_{2,4}$)	0.056	0.92	0.3596
Session2_5 ($\gamma_{2,5}$)	-0.034	-0.56	0.5798
Session2_6 ($\gamma_{2,6}$)	0.032	0.52	0.6069
Session2_7 ($\gamma_{2,7}$)	-0.064	-1.05	0.2964
Session2_8 ($\gamma_{2,8}$)	-0.077	-1.26	0.2086
Session2_9 ($\gamma_{2,9}$)	0.221	3.62	0.0004
Session3_2 ($\gamma_{3,2}$)	0.208	3.40	0.0009
Session3_3 ($\gamma_{3,3}$)	0.055	0.91	0.3659
Session3_4($\gamma_{3,4}$)	-0.342	-5.59	<.0001
Session3_5 ($\gamma_{3,5}$)	-0.023	-0.37	0.7091
Session3_6($\gamma_{3,6}$)	-0.105	-1.71	0.0891
Session3_7($\gamma_{3,7}$)	0.098	1.61	0.1098
Session3_8($\gamma_{3,8}$)	-0.012	-0.19	0.8492
Session3_9($\gamma_{3,9}$)	0.119	1.95	0.0537

The results are also mixed for Hypothesis 5 when the students hold renewable contracts in Test 3. The earnings of the VIF agents are consistent with the hypothesis and significantly less than the earnings of the students ($\beta_{32} < 0$) but the earnings of the IPT agents are significantly higher ($\hat{\beta}_{33} > 0$).

Some of the estimated session effects in Table 3.3 are relatively large (in absolute terms). For Test 1, one of the eight session coefficients is bigger than 0.2 (i.e. over 20%), and for Test 2 and Test 3, two of the coefficients are large. For Test 3, one corresponds to earnings that are less than 30% of the average earnings. The large estimates of some session effects imply that there were substantial differences in the ability of students to exploit the market even though they faced the same market conditions (i.e. system load and type of contract) as the other students in different sessions. In all three tests, the earnings of the IPT agents were higher than the earnings of the students. In contrast, the earnings of the VIF agents were lower than the students in Tests 1 and 3 but not in Test 2 when the students had permanent contracts. Table 3.4 summarizes the main results in a compact form and shows the average earnings for each type of firm in the last ten periods of each test (the same data used to estimate the model in Table 3.3).

The average earnings of the VIF agents were relatively stable for different sessions in the same test and also across tests (the range is roughly $1400). In contrast, the earnings of the IPT agents differed by over $8200 in the three tests, and by over $5600 for the students. In addition, there was substantial variation in earnings among the IPT agents and among the students in each test. Considering that all firms had the same cost structure for their generators, and the students and the IPT agents faced identical market conditions in Test 1, it is surprising how well the IPT agents did compared to the students. This provides some preliminary evidence that computer agents can be used effectively to test the performance of simple electricity markets. Additional discussion of this issue is presented in the following section.

The results in Table 3.4 show that the earnings increased for both the VIF agents and the IPT agents from Test 1 to Test 2 and also from Test 2 to Test 3. In contrast, the earnings of the students in Test 2 were lower than in Test 1, and in Test 3, they were much higher. The overall effects of the tests on the average spot prices in the market are summarized in Table 3.5 (note that all six firms in a session were paid the same spot prices). In Test 1, the average spot price was $67.0/MWh, which was 25% higher than the competitive price ($53.8/MWh). Although requiring the students to hold permanent contracts in Test 2 lowered their earnings, it did not lower the spot price. The spot price was 5% higher in Test 2 ($70.3/MWh). One explanation for this result is that the students gained experience in Test 1, and they were largely responsible for getting higher prices in Test 2. However, the main

Table 3.4 Average Earnings ($ in periods 11 to 20)

Variable	Students	VIF-Agents	IPT-Agents
Test 1	13,019	10,845	13,170
	(3,234)	(633)	(2,701)
Test 2	10,230	11,295	17,031
	(2,703)	(894)	(4,713)
Test 3	18,625	12,210	21,494
	(3,532)	(866)	(4,130)

* The standard deviation is given in parentheses

Table 3.5 Average Market Prices ($/MWh)

Experiment	Average Spot Price	Average Forward Price
Cost-based	53.8	-
Test 1: No Contracts	67.0	-
Test 2: Permanent Contracts	70.3	60.0
Test 3: Renewable Contracts	76.4	72.4

benefits of the high prices in Test 2 went to the IPT agents because the students were paid only $60/MWh for 50MW in the permanent contract instead of the spot price ($70.3/MWh). Overall, these results show clearly that requiring suppliers to hold forward contracts does not guarantee that spot prices will be reduced. When the students held renewable contracts in Test 3, the average spot price ($76.4/MWh) was 14% higher than in Test 1. The students were able to increase the contract price from $60/MWh to $72.4/MWh in Test 3. This result supports our initial hypothesis about forward contracts (presented in Section 2) and the basic rationale for writing this paper. If there is market power in the spot market and the spot prices influence the forward prices, renewing contracts will ensure that the market power persists in both the forward and spot markets.

3.7. The Performance of Computer Agents

If computer agents can replicate human behavior, they can be used effectively to test electricity auctions and to supplement the results obtained using humans by doing additional sensitivity tests. In the first part of this section, the performance of the computer agents in the experiments is evaluated by comparing the average earnings of the two IPT agents and the two students in the three tests. In the second part, the results of experiments using all agents are presented.

We have already shown in the previous section that the average earnings of the IPT agents in different sessions were higher than the earnings of the students in all three tests, and significantly higher in Tests 2 and 3. The results for individual sessions are summarized in Figure 3.3 for the three tests. A circle above the 45 degree line implies that the two IPT agents in a specific session earned more than the two students. This was true for most sessions in all three tests. For the nine sessions in a test, the earnings of the IPT agents were lower in 3 sessions for Test 1, but higher in every session for Test 2 and Test 3. Although the earnings of the students and the IPT agents were similar in Test 1, this performance measures the effect after a learning period of 15 periods. The average earnings during the first 15 periods of Test 1 show that the IPT agents earned more than the students in 8 out of the 9 sessions.

These results are particularly important for Test 1 because the optimization problem was exactly the same for the IPT agents and the students. The conclusion is that their performances were similar in Test 1. In Test 2, the IPT agents had a distinct advantage because the students held a contract for 50MW at a predetermined price. The objectives for the IPT agents and the students in Test 3 were different, but there was no clear advantage for the IPT agents. The results demonstrate that the IPT agents are i) just as good as the students in a simple market (Test 1) and ii) are able to exploit market conditions effectively (Test 2 and Test 3) when given the opportunity.

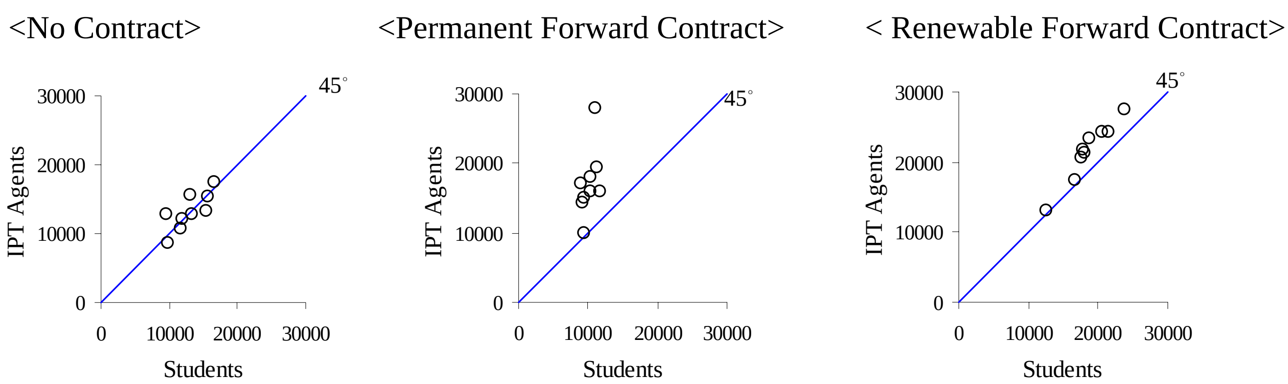

Fig. 3.3 Comparison of Average Earnings

It is also clear from Figure 3.3 that the earnings of students and IPT agents are positively correlated. This explains why the estimated session effects in Table 3.3 are relatively large. When one pair of students is successful in raising the market price, the IPT agents learn to speculate and reinforce the students' behavior. As a result, all firms get higher earnings.

The IPT agents can also exploit unusual situations effectively. For example, in Session 2_9 of Test 2, one student sold only the contracted 50MW and withheld everything else from the auction, and the other student submitted three of the four non-contracted blocks (40MW in addition to the 50MW contracted) at very high offer prices. As a result, the spot prices were high. These high prices persisted because the two students did not change their behavior. (Session 2_9 had the highest positive session effect in Table 3.3 among all sessions in all tests) Under these circumstances, the IPT agents earned more than 2.5 times as much as the students. The IPT agents withheld less capacity than the first student and submitted lower price offers than the second student. It was not necessary for the IPT agents to speculate because the students were speculating so aggressively. This is exactly the type of strategy that was followed by Eastern in the UK market during the 1990's when high market prices were set on a predictable basis by two other firms.

When the three tests were repeated using computer agents to replace the two students, the average earnings of firms and the average market prices were higher than the corresponding values in Tables 3.4 and 3.5 in all cases. For these tests, the students' firms were represented by Latent Speculators (LS) in the all-agent tests. LS agents are more likely to speculate than IPT agents, but when high prices occur, the IPT agents adapt to the new market conditions and evolve into LS agents.

The average earnings by the type of firm are summarized for the all-agent tests in Table 3.6. The percentage changes from the corresponding values in Table 3.4 are also shown, and in eight out of the nine cases, these changes are positive. The small negative change for the VIF agents in Test 2 is the only exception. The positive changes for Test 1 (no contract) and Test 3 (renewable contract) are very large (ranging from 13% to 70%). It is only when the LS agents have a permanent contract in Test 2 that the changes are relatively small (ranging from -1% to 21%). A comparison of the average earnings of the LS agents in Table 3.6 to the corresponding session values for the students in Figure 3.3 shows that the values for the LS agents fall in the ranges observed for the students. For Tests 1 and 3, the earnings of the LS agents are similar to the highest earnings of the students, but for Test 2, they are only slightly above the median value. The general conclusion is that the LS agents were able to exploit market power effectively when the opportunity arose in Tests 1 and 3. However, more students were able to get higher earnings than the LS agents in Test 2 when it was relatively difficult to exploit market power.

Using a Chow Type II test, it is possible to test whether or not the 18 new observations obtained from the all-agent tests deviated from the sample of 162 observations using students. The first null hypothesis assumed that the earnings of the all-agent firms were equal to the **average** earnings of the students (i.e. by setting the session effects for the all-agent tests to zero). The parameters in model (10) were estimated using the pooled data set of 180 observations. The computed F statistic (4.39) is large and the null hypothesis is rejected (the critical value for an $F_{(18,130)}$ is 1.70 at the 5% level of significance). This implies that the earnings of the all-agent firms were statistically different from the average earnings in the tests using students.

The second null hypothesis assumed that the earnings of the all-agent firms were equal to the sessions with the **highest** earnings obtained by the students (i.e. by selecting the sessions with the largest positive session-coefficients for each market in Table 3.3 (Sessions 6, 9 and 2 for Markets 1, 2 and 3, respectively)). In this case, the computed F statistic (0.95) is small and it supports the null hypothesis (note that the critical value of 1.70 is still valid). In other words, the earnings in the all-agent firms were statistically equivalent to the sessions with the highest earnings obtained by the students.

The average prices in the all-agent tests are summarized in Table 3.7, and the percentage changes from the corresponding average prices in Table 3.5 are also reported. The price changes in the three tests are all positive and equal to 14%, 6% and 28%, respectively. Unlike the tests with the students, the average price in Test 2 is lower than the price in Test 1. This is consistent with the arguments of the California MSC (discussed in the Introduction) that forward contracts will mitigate market power and lower spot prices. However, the reduction of (73.5– 75.9) = -$2.4/MWh is only 3% lower than the price in Test 1. This is roughly one tenth of the reduction needed to get competitive prices ((53.8 – 75.9) = -$22.1/MWh (-29%)). With a renewable contract in Test 3, the average price in the all-agent test ($95.7/MWh) is substantially higher than the average price ($74.9/MWh) obtained by the students. Once again, this evidence supports our argument in Section 2 that market power will be transferred from the spot market to the forward market when contracts are renewed.

Overall, the results from the all-agent tests are encouraging and show that computer agents do provide a valid way to evaluate the performance of electricity markets. The computer agents were able to match the earnings of the best students. It will be interesting to find out in the future whether this is also true for more complicated market structures, such as joint markets for energy and ancillary services.

Table 3.6 Average Earnings ($/MWh)* for All-Agent Tests

Variable	LS-Agents	VIF-Agents	IPT-Agents
Test 1	18,154 (+39%)	12,297 (+13%)	20,892 (+59%)
Test 2	12,375 (+21%)	11,200 (-1%)	18,235 (+7%)
Test 3	26,257 (+41%)	16,080 (+32%)	36,512 (+70%)

* The percentage change from the corresponding value in Table 3.4 is given in parentheses

Table 3.7 Average Market Prices ($/MWh)* for All-Agent Tests

Experiment	Average Spot Price	Average Forward Price
Cost-based	53.8 (0%)	-
Test 1: No Contracts	75.9 (+14%)	-
Test 2: Permanent Contracts	73.5 (+6%)	60.0 (0%)
Test 3: Renewable Contracts	95.7 (+28%)	92.9 (+30%)

* The percentage change from the corresponding value in Table 3.5 is given in parentheses

3.8. Summary and Conclusions

The primary objective of this paper is to investigate whether or not holding a forward contract is an effective way to mitigate market power and reduce spot prices in an electricity market. A second objective is to determine how well computer agents can replicate the behavior of human subjects in tests of electricity auctions. Using PowerWeb to simulate the operation of a uniform price auction run by an ISO, four computer agents and two human subjects (graduate students) represent six supply firms in three different market situations. In each case, the patterns of load are exogenous and there are 20 trading periods (25 for Test 1). The three tests are 1) no forward contract (all dispatched capacity is paid the spot price), 2) the two students hold a permanent forward contract (the contract price is fixed), and 3) the two students hold a renewable forward contract (the current spot price influences the forward price used to renew the contract). The specified structure of this forward market is simpler than an actual market because the quantities of the forward contracts and the renewal periods are preset. In a second experiment, the three tests are repeated with two additional computer agents replacing the students.

With a permanent forward contract in Test 2, it is expected that the high spot prices in Test 1 will be mitigated. Since the largest block of capacity (50% of each firm's installed capacity) has already been committed at a fixed price in the contract, the supplier has less incentive to speculate because getting a high spot price only affects the remaining capacity. This would also be true for a renewable forward contract if the spot price does not affect the forward price that determines the contract price when it is renewed. This is effectively the case when the number of periods to the renewal date is large. In general, however, high spot prices affect expectations about future spot prices, and therefore, they also affect the forward price for a specified delivery period. In Test 3, ten-period contracts are renewed in periods 10 and 20.

When high spot prices increase the expected future price for renewing a forward contract, market power in the spot market is transferred to the forward market because firms holding renewable contracts have greater incentives to speculate. High spot prices increase both current earnings from non-contracted sales in the spot market and future earnings from the next contract. Based on the theory developed in Section 2, our prior expectations were 1) the spot price in Test 2 would be lower than the spot price in Test 1, and 2) the spot price in Test 3 would be higher than the spot price in Test1. In other words, high spot prices would be mitigated by a permanent contract, but not by a renewable contract because persistently high spot prices increase expected future spot prices and forward prices.

During the fall 2003, a class of 20 graduate students in engineering and economics were divided into pairs, and each pair tested the three different markets. The average results from the final 10 periods from the test of each market were evaluated for each pair of students. The average prices ($/MWh from Table 3.5) for Tests 1-3 were 65, 69 and 75, respectively. Hence, the prices were higher with both a permanent and a renewable contract in Tests 2 and 3 compared to the price in Test 1. The higher price in Test 2 was the opposite of our prior expectations and inconsistent with the widely held rationale for requiring suppliers to hold forward contracts. However, the price difference between Tests 1 and 2 was relatively small. We attribute this unexpected result to the students learning how to exploit the market more effectively during the series of tests. The results from Test 3 showed clearly that renewing a forward contract provided additional incentives to speculate and raise both the spot price and the forward price.

The results for the three all-agent tests were more consistent with our prior expectations, and the prices for the three markets ($/MWh from Table 3.7) were 76, 74 and 96, respectively. However, the price difference between Test 1 and Test 2 was small. None of the results, with students or with all agents, support the conventional wisdom that making suppliers hold forward contracts will automatically mitigate market power in the spot market. One could specify conditions under which lower spot prices could be sustained in the presence of market power (e.g. by committing most generating capacity to long-term bilateral contracts), but in general, market power will lead to high spot prices and high forward prices. This is exactly what happened in the California energy crisis when new forward contracts were executed in the winter of 2001 (Mount and Lee, 2002). The high contract prices reflected the expectations of traders at that time that federal regulators were not prepared to enforce a price cap to deal with market power in the Western Interconnection. This result is consistent with our main conclusion about the market power in this paper. The market power exhibited by persistently high prices in the spot market resulted in higher expectations about future spot prices, and as a result, higher prices in the forward contracts.

The results for the computer agents were reassuring. Using the maximization of expected profits as the objective criterion for submitting offers by an agent, it is possible to modify the general form of a computer agent to represent different types of firm, such as a vertically integrated firm, and to deal with the different tests of forward contracts. In Test 1 with two students and four computer agents, the IPT agents and the students faced identical cost conditions. In 6 out of the 9 sessions in Test 1, the IPT agents had higher average earnings than the students. In the all-agent tests, the earnings of the LS agents that replaced the students were higher than the corresponding average earnings of the students in all three tests. In Tests 1 and 3, the earnings of the LS agents were similar to the highest earnings obtained by the students.

Our conclusion about the second objective of the paper is that computer agents can replicate the behavior of students in an electricity auction effectively. In fact, the agents were also able to exploit unusual situations by, for example, behaving as free riders when the students in a session speculated aggressively. These results suggest that it is appropriate to do additional sensitivity tests using all computer agents. This is a promising line of research that is essential for developing realistic simulation models of deregulated electricity markets. Relying exclusively on human subjects to test different market structures will, due to the practical difficulty of recruiting enough people, limit the scope of the tests. Computer agents that can replicate realistic behavior can be used to extend the range and number of tests conducted by human subjects. This capability has tremendous potential for identifying potential flaws in market designs and finding effective ways to improve the performance of electricity markets before a specific market design is imposed on the public.

3.9. References

Allaz Blaise and Jean-Luc Vila (1993) "Cournot Competition, Forward Markets and Efficiency," *Journal of Economic Theory*, Vol. 59, pp.1-6.

Dennis W. Carlton (1979), "Contracts, Price Rigidity and Market Equilibrium", *Journal of Political Economy*, Vol.87, no. 5, pp 1034-1062.

Scott M. Harvey and William W. Hogan (2000), "California Electricity Prices and Forward Market Hedging", mimeo, Harvard University.

Timothy Mount and Yoo-Soo Lee (2003), "The Effects of the Dysfunctional Spot Market for Electricity in California on the Cost of Forward Contracts", mimeo, Cornell University.

Timothy Mount and Hyungna Oh (2004), "On the First Price Spike in Summer", Presented at the Hawaii International Conference on System Sciences, Big Island.

Hany Shawky, Achla Marathe and Christopher Barret (2003), " A First Look at the Empirical Relation Between Spot and Futures Electricity Prices in the United States", *the Journal of Futures Markets,* Vol.23, No.1, pp.931-955.

Afzal Siddiqui (2002), "Managing Electricity Reliability Risk Through the Forward Markets", Lawrence Bekerly National Laboratory.

David Watts and Fernando Alvarado (2004), “The Influence of Futures Markets on Real Time Price Stabilization in Electricity Markets”, Presented a at the Hawaii International Conference on System Science, Maui.

Frank Wolak and Robert Nordhaus (2000a), “An Analysis of the June 2000 Price Spikes in the California ISO’s Energy And Ancillary Services Markets”, Market Surveillance Committee (MSC) of the California Independent System Operator.

Frank Wolak (2000b), “An Empirical Analysis of the Impact of Hedge Contracts on Bidding Behavior in a Competitive Electricity Market”, *International Economic Journal*, Vol. 14, No.2, pp. 1-39.

Frank Wolak (2003a). “Lessons from the California Electricity Crisis,” Center for the Study of Energy Markets (CSEM) Working Paper 110, University of California Energy Institute, Berkeley, California. http://www.ucei.org.

Frank Wolak (2003b). “Measuring Unilateral Market Power in Wholesale Electricity Markets: The California Market 1998 to 2000,” *AEA Papers and Proceedings*, May 2003, [online]: http://www.stanford.edu/~wolak.

Catherine D. Wolfram (1998), “Strategic Bidding in a Multiunit Auction: An Empirical Analysis of Bids to Supply Electricity in England and Wales”, *The Rand Journal of Economics*, Vol.29, No.4.

Catherine D. Wolfram (1999), “Measuring Duopoly Power in the British Electricity Spot Market”, *American Economic Review*.

4. Symmetry and Verification: Critical Parts of Market Design

Abstract

Most electricity markets contain auctions that are frequently repeated. We examine such auctions with the aid of software agents. The purpose of these agents is not to mimic the behavior of humans. Rather, it is to reveal those fundamental properties of repeated auctions that are independent of human idiosyncrasies; just as wind tunnels and finite element programs are used to reveal the fundamental properties of new airplanes before test pilots are made to fly them. We show, by means of simple experiments, that existing designs of repeated auctions have at least one major flaw: The sellers, even though they work without collusion, can learn strategies that raise prices and profits far above competitive values. This happens even when the demand is price-responsive. A remedy is to make the auctions symmetric, that is, to provide the buyers with as much autonomy as the sellers have, so the buyers, or their surrogates, can learn, and otherwise make decisions as quickly and cleverly as the sellers. However, this remedy is unlikely to be the only one needed by current designs. Complex artifacts, unless their designs are thoroughly verified, invariably suffer from major flaws. Thorough verification requires a set of tests that spans both the operating conditions and the desired behaviors of the artifact in question. What passes for verification in market design—opinion surveys and a few random experiments with human subjects—is far less than adequate.

4.1 Introduction

Both theory and practice suggest that competitive markets can benefit consumers by lowering prices and costs. Both theory and practice also suggest that if one or more firms has the power to set prices, then the result will be high prices for consumers. When FERC and many states began deregulating the electricity industry, they sought to gain the benefits of market competition by transferring authority from regulators to markets. Many legislators, sellers and customers hoped that competition would relieve bureaucracy and do away with misplaced incentives, providing higher profits for sellers and lower prices for buyers.

The California debacle of 2000 has chilled these hopes, leading other states to either renounce deregulation or pull back. Everyone has been searching for explanations for what went wrong. A series of articles and books has blamed the high prices on (1) the drought and the consequent shortage of hydroelectric power, (2) the refusal to allow utilities to enter into long term contracts for power, (3) the pricing structure in which customers' prices did not rise with costs, (4) the shortage of generation capacity, and (5) high natural gas prices. We add two other culprits to the list:

- asymmetric decision-making (allowing the sellers more autonomy, that is, more decision-making freedom, than the buyers), and
- unverified designs (designs that are implemented before being thoroughly tested to confirm that they meet their specifications and are free of major flaws).

4.2 The Dynamics of Repeated Electricity Auctions

The typical electricity auction is uniform and one-sided. The auctioneer (generally the ISO, or Independent System Operator), estimates the demand for electricity over the period of concern (each of the 24 hours of the next day, for instance), collects offers from the suppliers for this period, assembles these offers into an aggregate supply curve, and determines the clearing price (the price at which supply just meets demand). This clearing or wholesale price is paid for all the electricity that is bought from the sellers. The auction is repeated many times.

The sellers in electricity auctions are allowed to act as autonomous agents. The dynamics of systems containing such agents are emergent--the dynamics evolve, bottom-up, from the interactions of the agents, each deciding for itself what it is going to do. The dynamics can be represented by trajectories through the decision space of the autonomous agents. These trajectories can be quite complex. To illustrate, consider a scenario in which:

- An hourly, uniform auction is repeated many times.
- In every repetition, the demand is price-invariant with a value of 60 MW.
- There are 10 sellers, each with 10 generators. Every generator has the same capacity: 1 MW. The n-th generator of the m-th seller costs 10n dollars per MWh.
- At each repetition, each seller may choose to offer the output of some of its generators for sale, while withholding the others. But those generators that are offered, must be offered at cost.
- Let A_k be a binary matrix representing the offers in the k-th repetition. Let A_{kM} be one of the several values of A_k that maximizes the total profits of the sellers (see Fig. 4.1).
- If the sellers offer insufficient generation to meet the demand, the ISO can recall any generators it wishes. When this happens, the total profit for the sellers is the same as would be produced by the offer A_{kC} in which all the entries are "1". This value, called a competitive offer, produces the minimum total profit: $1500.

When the demand curve remains the same from one repetition to the next, the total profit produced by any offer, $\Pi(A_k)$, lies in the interval between the competitive profit, $\Pi(A_{kC})$, and the monopolistic profit, $\Pi(A_{kM})$.

Of course, high prices are necessary but not sufficient for high profits. For instance, the offer, A_{kR}, whose first four columns contain 0's and whose other columns contain 1's, produces the highest possible clearing price, $100/MWH, and the lowest possible total profit, $1500.

A_{kM} is a Nash equilibrium with respect to the absolute values of both total and individual profits. (A Nash equilibrium is a value of A_k in which no seller can improve its position by making a small and unilateral change in its offer, that is, by switching one entry in its offer.) However, A_{kM} is not a Nash equilibrium with respect to relative profits or profit-ranking. The clearing price of A_{kM} is set by seller-1, the same seller that makes less profit than anyone else. If seller-1 wishes to improve its profit-ranking, it could do so by

changing the "0" in the first row and sixth column of A_{kM} to a "1", and thereby, moving from being behind everyone else to a position of equality.

An aggregate offer: A_{kM} Profits: $U_k(A_{kM})$

1	1	1	1	1	0	0	0	0	1
1	1	1	1	1	1	0	0	0	0
1	1	1	1	1	1	0	0	0	0
1	1	1	1	1	1	0	0	0	0
1	1	1	1	1	1	0	0	0	0
1	1	1	1	1	1	0	0	0	0
1	1	1	1	1	1	0	0	0	0
1	1	1	1	1	1	0	0	0	0
1	1	1	1	1	1	0	0	0	0
1	1	1	1	1	1	0	0	0	0

350
390
390
390
390
390
390
390
390
390

Total profit: $\Pi(A_{kM})$ = 3860

Figure 4.1 The rows in the matrix, A_{kM}, represent offers from the sellers for the k-th auction; a "1" denotes a generator whose energy is offered for sale, a "0" denotes a generator that is withheld. The pattern of 1's and 0's in A_{kM} produces the maximum total profit for the sellers, \$3860, just as they would obtain if their generators were controlled by a monopolist with perfect information, including perfect predictions of demand.

In contrast to A_{kM}, the offer, A_{kC} is a Nash equilibrium with respect to profit-rank, but not with respect to absolute profits.

Let α be the set of all the 2^{100} distinct values of A_k; let β be the subset that contains the 10 values of A_k that are Nash equilibria w.r.t. profits and produce monopolistic prices; let γ be the subset that contains the 2^{40} values of A_k that are Nash equilibria w.r.t. rank, that yield competitive profits, and meet demand without any recall.

The set α is the decision space for the repeated auctions of the scenario being considered. The repetitions can be visualized as tracing a trajectory through this space. Equilibria serve as attractors for these trajectories. The stronger their attraction, the more trajectories that will converge to them.

Trajectories can converge to more complex features than equilibria. Limit cycles are an example. Consider the 10 elements of β. Each requires a different seller to sacrifice some profit in order to set the clearing price to its highest possible value. If the sellers are given to reciprocal altruism, they can take turns setting this clearing price. The resulting trajectory would cycle endlessly through the elements of β.

The question is: what form will the market's behavior take? Are there strong attractors in the decision-space of the sellers? Will the structure of the market draw the sellers' offers to the elements of β and high prices, or to the elements of γ and competitive prices, or to some quite different attractors?

We will seek experimental answers to these questions. Emergent dynamics are difficult, if not impossible, to analyze by a priori reasoning. Nor is any practicable list of design-rules likely to be sufficient for acceptable dynamics. Consider, for instance, the list that many believe to be sufficient for competitive behavior:

a) the market must include multiple, profit seeking sellers,
b) there must be no collusion among the sellers, nor any seller with market power, and
c) the aggregate demand must be elastic.

We will show by means of simple experiments, that this list is insufficient for holding prices at, or even near, competitive levels.

In addition, we will demonstrate some of the effects of cooperation and learning market performance. Cooperation among sellers is, of course, illegal. But learning is not, and variations in what and how the sellers learn can profoundly change the shapes of the trajectories of their offers, and the attractors to which they are drawn.

4.3 An Overview of the Experiments

The demand for electricity is strongly influenced by quasi-repetitive factors, such as the weather and peoples' work schedules. Consequently, the demand is often quasi-repetitive. (A sequence of demands is quasi-repetitive if it is composed of subsequences that are repetitive or very nearly so.) A seller, faced with a quasi-repetitive demand, can increase its profits by recognizing the repetitive subsequences, and learning a strategy for each subsequence. For instance, a seller might notice that the demand every Tuesday morning from ten to eleven is much the same, and learn a good strategy for this demand. (Learning is most effective in repetitive circumstances.)

We report, briefly, on four experiments designed to examine the effects of learning on repeated auctions. Buyers, sellers and auctioneers are represented by software agents. The purpose of these agents is not to mimic humans. Rather, it is to demonstrate that the results of repeated auctions are extremely sensitive to learning.

In each, there are 10 identical sellers, with total capacity well in excess of total demand. Therefore, no seller has market power. In the first experiment, the demand is price-invariant; in the second, demand is price-responsive; in the third, sellers are allowed to cooperate to very small; and in the fourth, both sellers and buyers actively bid into the market.

Work by Smith (1966), Rassenti et al (2003) and Mount et al (2001) suggest a flaw in current market designs. With repeated auctions and price-invariant demand, undergraduates acting as sellers, quickly learn to raise prices and profits. Our experiments extend this finding to price-responsive demands; they can be performed in a fraction of the time required for human subjects, and can easily be used to test many more combinations of market parameters than is possible with human subjects.

Terminology

Consider the k-th repetition of a two-sided auction. The decision problems for the m-th seller and the n-th buyer are to come-up with an offer a_k, and a bid g_k, so as to:

$$\underset{a_k}{\text{Maximize}}\ u_k(A_k, G_k) \tag{1}$$

$$\underset{g_k}{\text{Maximize}}\ h_k(A_k, G_k) \tag{2}$$

where:

u_k is the profit of the m-th seller in the k-th auction

a_k is a 10-dimensional, binary vector, representing the m-th seller's offer in the k-th auction. As before, a "1" represents a generator that is offered at cost, a "0"represents a generator that is withheld.

A_k is the matrix of all the sellers offers for the k-th auction; the m-th row of A_k is a_k.

g_k is a 10-dimensional vector, representing the n-th buyer's bid in the k-th auction. The j-th element of this vector is an integer between 0 and 100, and represents the price the buyer is prepared to pay for the j-th MWh of energy.

G_k is the matrix of all the buyers' bids for the k-th auction; the n-th row of G_k is g_k.

h_k is the consumer-surplus for the n-th buyer in the k-th auction.

In the experiments that follow, simple learning algorithms embedded in software agents (representing the buyers and sellers) will be used to calculate:

P_M, Π_M: the monopolistic clearing price and total profit, in \$/MWh and \$, that is, the price and profit an all knowing monopolist would obtain when the demand curve remains the same from one repetition to the next ($G_{k+1} = G_k\ \forall$ k).

P_C, Π_C: the competitive clearing price and total profit, in \$/MWh and \$, that is, the price and total profit when all the sellers offer all their energy at cost and the demand curve remains the same from one repetition to the next ($G_{k+1} = G_k\ \forall$ k).

P_L, Π_L: the learned clearing price and total profit, in \$/ MWh and \$, averaged over the last 50 repetitions of an extended sequence of repetitions obtained when the sellers and/or the buyers use automatic learning algorithms

An Evolutionary Learning Algorithm

The payoff for each agent depends, not only on what it does, but also, on what all the other agents choose to do. How is an agent to know what the others might do? An intelligent agent will learn to predict their future actions from whatever it knows about their past actions.

We will use a very simple learning algorithm that develops decision-vectors (offers for sellers, bids for buyers) by combining three ideas from Genetic Algorithms with one idea from Simulated Annealing. The ideas from genetic algorithms are: "survival-of-the-fittest" (meaning that only the most successful decision-vectors from previous repetitions are retained, the rest are discarded), "crossover" (meaning that a new decision-vector is formed by randomly breaking two old decision-vectors, into left- and right-parts, and combining the left-part of one with the right-part of the other), and "mutation" (which means choosing a random entry of a decision-vector and changing it). The idea from simulated annealing is "cooling" (which means that p_k, the probability of mutation, is slowly decreased as the repetition count, k, increases, much as the probability of change in the crystalline structure of an alloy decreases as it is annealed).

In the first auction, the sellers offer all their generators and the buyers bid their true valuations. The next 49 auctions are run with random changes to these initial offers and bids. For subsequent auctions, the buyers and sellers use the same learning algorithm.

Consider any agent preparing to make an offer or a bid in auction-k, with $k > 50$. Let:

$f_i(d)$ be a fitness function for the decision vector, $d = a_i$ or g_i. The value of f_i increases with profit for sellers, and with consumer surplus for buyers.

X, Y be sets of previous decision-vectors for the agent, such that X contains the 10 best vectors (w.r.t. f_i), and Y contains the 40 next best vectors.

The steps in calculating a_k and g_k are:

1. Randomly select two decision-vectors, x and y, such that $x \in X$ and $y \in Y$. Calculate a new decision-vector, c = Crossover[x, y].
2. Calculate another new decision-vector, e = Mutation[c] with probability p_k. (Randomly choose an element of c. Flip a coin, biased so the probability of heads is p_k; where $p_k = 0.2$ or $5^{-k/1000}$, whichever is bigger. If the coin comes-up heads and the decision-maker is a seller, change the value of the selected element to its binary complement; if the coin comes-up heads and the decision-maker is a buyer, change the value of the selected element to any randomly chosen integer between 0 and 100; otherwise, leave the selected element unchanged.)
3. Set E to e with probability p_k, and to x with probability $(1-p_k)$. (In other words, flip the biased coin again. If it comes-up heads, set E = e; else, set E = x.)
4. If the decision-maker is a seller, set $a_k = E$, else set $g_k = E$

4.4 Experiment-1: Price-Invariant Demand

Three fixed demands are considered (curves 1-3 of Fig. 4.2). Since these demands correspond to 50, 65, and 80% of total generating capacity, and since this capacity is

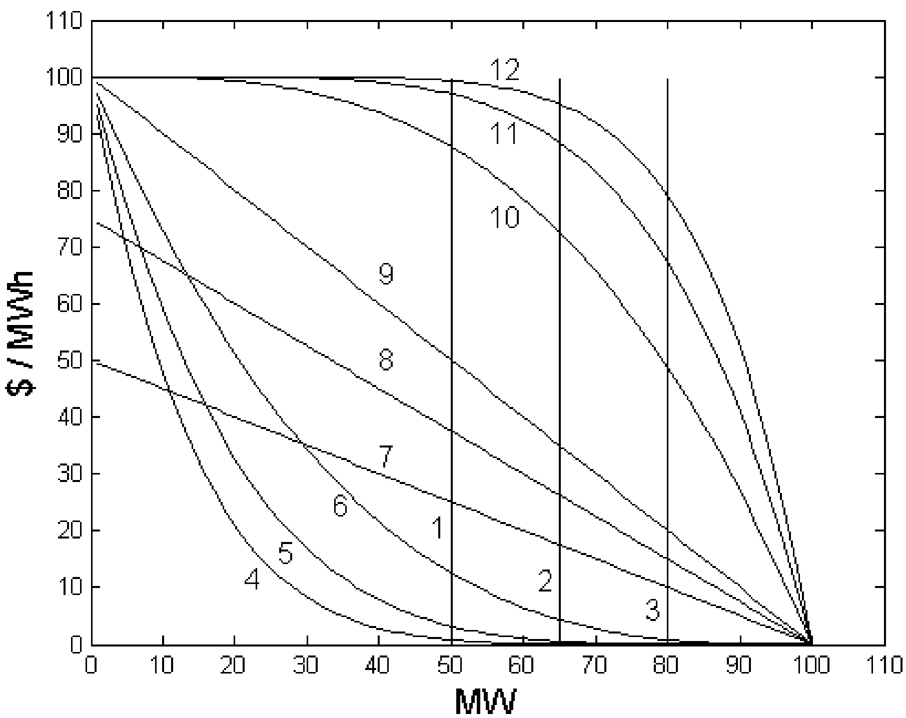

Figure 4.2 Twelve demand curves of which curves 1 - 3 are price-invariant, 4-12 are price-responsive.

equally distributed among 10 identical sellers, no seller has market power, at least not in the sense of the HHI index.

Each seller has ten 1-MW generators. The aggregate supply curve is represented by the broken line in Fig. 4.3. At the end of each repetition, each seller is told only two things: which of its generators were accepted by the ISO, and the market-clearing price. It then calculates the profit and uses a weighted average of price and profit for the fitness functions, that is, for deciding whether to put that choice in the set X, or in the set Y, or to discard it (the weights make price and profit equally important in this decision). Each seller then invokes the learning algorithm to calculate its offer for the next repetition.

Typical results are shown in Figures 4.3 and 4.4. Since the learning algorithm makes random changes at every repetition, prices will not converge to a single value. But the price-trajectories do seem to be drawn to complex attractors--bounded patterns--as suggested by Figure 4.4.

These patterns are sensitive to the form of the fitness function. If this function combines price and profit, as in this experiment, then the prices, profits, and learning yields evolve to high values (Figures 3.3 and 4.4). When the fitness function includes only profits, prices and profits evolve to competitive values (Lye, 2004). In other words, when the sellers are

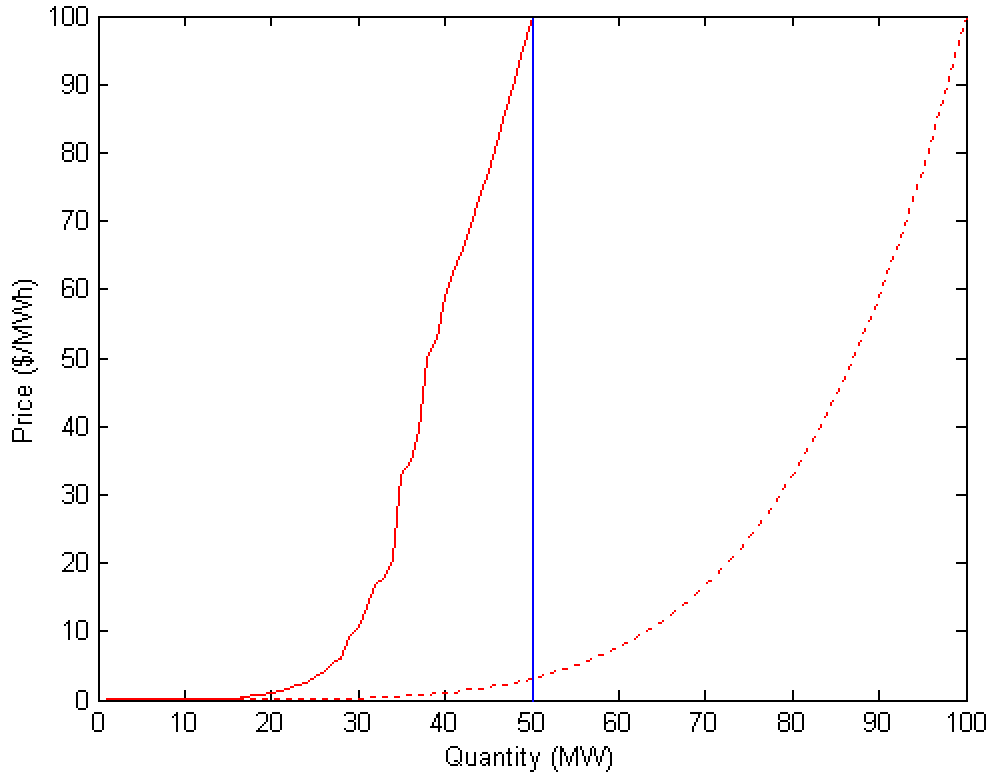

Figure 4.3 Aggregate supply curves from Experiment-1 with the demand fixed at 50 MW. Initially, all the suppliers offer all their generators at cost. The resulting aggregate supply curve is represented by the broken line and the clearing price in the first auction is the competitive price: $3.13/MWh. As the auctions proceed, the suppliers learn to withhold generators. At repetition 1000, the withholding is so effective that the aggregate supply curve (solid line) intersects the demand at the maximum possible price, $100/MWh.

concerned only with their own profits, the clearing prices are drawn to the competitive price, as might be expected in a market with ten sellers of equal size and no collusion. (Experiments by Dolbear, et al, 1968, suggest that the same sort of thing happens in an N-person prisoner's dilemma). However, when the sellers take both price and profit into account, they can obtain very high profits, even though they work independently.

4.5 Experiment-2: Price-Responsive Demand

This scenario is the same as Experiment-1, except that the price-invariant demands are replaced by a variety of downward sloping demand schedules. Figure 4.5 indicates that price-responsive demand does not reduce the ability of the sellers to extract high profits. Surprisingly, the yields are somewhat greater—80% to 90% for price-responsive demands, as compared to 70% to 80% for price-invariant demands.

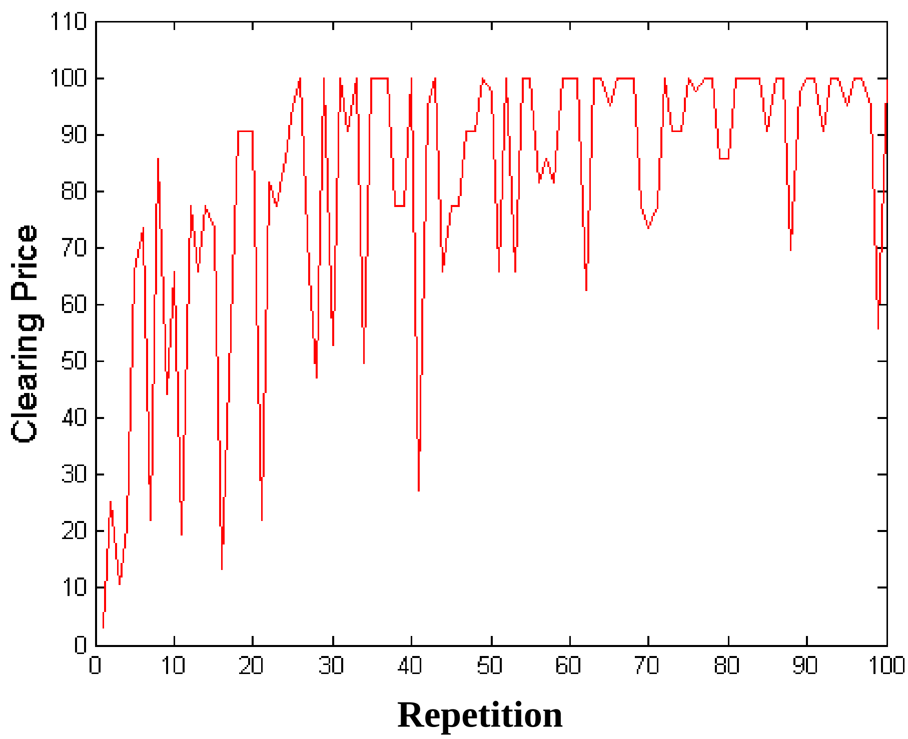

Figure 4.4 A trajectory of clearing prices from Experiment-1, with the demand fixed at 50 MW. The fluctuations in the trajectory decrease as the trajectory develops, but never completely disappear because the learning algorithm continues to makes random changes in the suppliers' offers. By repetition 1000, the average clearing price (averaged over 50 consecutive repetitions) has become $98/MWh, and it remains at this value thereafter.

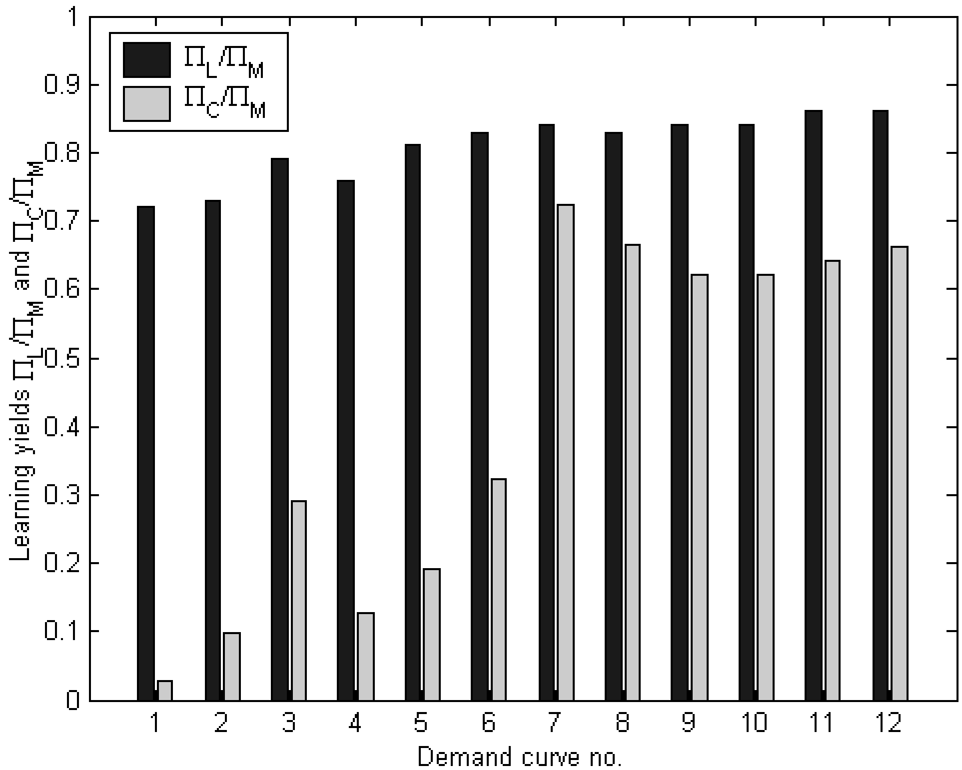

Figure 4.5 Profits and learning yields from Experiments-1 and –2. Whether the demand is price-invariant (demand curves 1-3) or price-responsive (demand curves 4–12), the sellers are able to obtain between 72% and 87% of the monopolistic profit.

4.6 Experiment-3: Cooperation

This scenario is the same as Experiment-1, except that the sellers cooperate by exchanging information on their individual profits, and by adopting a common goal: to maximize the total profit in each auction without regard to who gets the profit. (In California in 2000, each seller knew the heat rate of every natural gas generator as well as the price of natural gas. Thus, knowing the clearing price of auction-i, each seller could estimate the profits made by every other seller in that auction, and could, therefore, replicate the conditions of Scenario-3 without actually exchanging any information).

The sellers achieve higher profits: 80% to 98% of the monopolistic profits, when each seller seeks to maximize total profit, as compared to 70% to 90%, when each seller seeks to maximize its own profit (Figures 4.5 and 4.6). To divide the profits more evenly, the sellers seem to learn a simple form of reciprocal altruism, taking turns to set the clearing price.

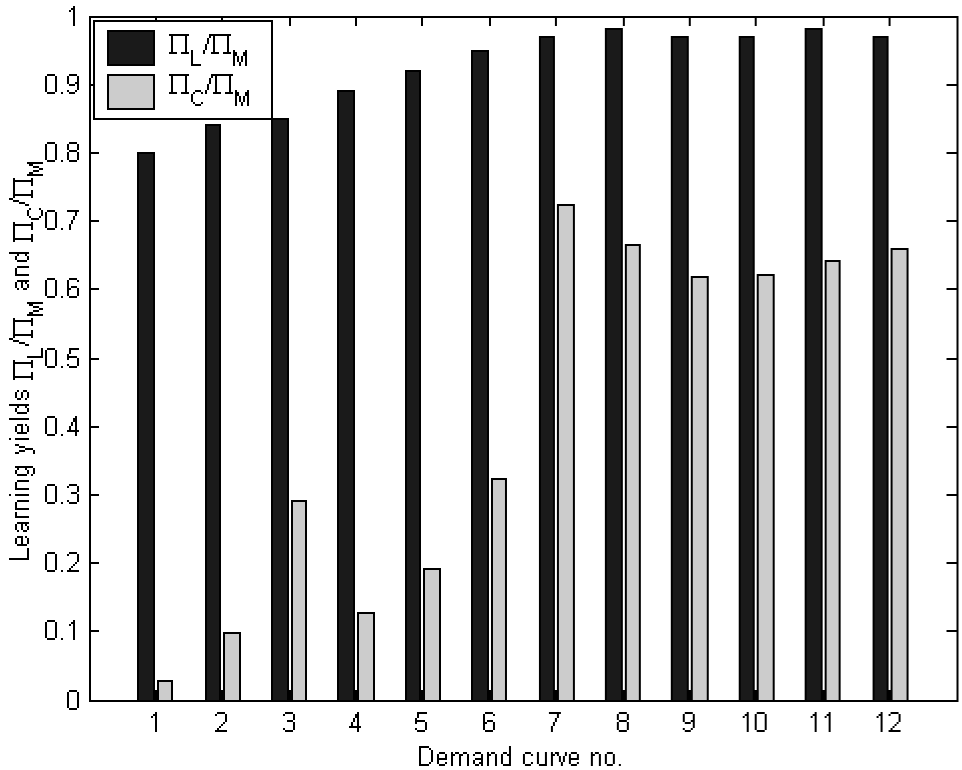

Figure 4.6 Profits and learning yields from Experiment-3, in which sellers co-operate by a) exchanging data on their individual profits, and b) seeking to maximize total, rather than individual, profit.

4.7 Experiment-4: Active Demand

The 10 sellers now face 10 identical buyers, all of whom are allowed to learn. Each buyer is assigned a demand of up to 10 MWh at prices such that the aggregate demand coincides with one of the previous demand schedules (the experiment is repeated so all the price-responsive demand curves are selected, one after the other).

Each seller works independently, as does each buyer. Both use the evolutionary learning algorithm to produce an offer or bid for the next auction. But the arrangement is not perfectly symmetric. The buyers have larger decision-spaces, that is, more autonomy. While each seller can choose only to offer or withhold each of its MWh of supply, each buyer can choose from among 101 different prices (ranging from 0 to $100) to bid for each of its MWh of demand.

The objective (fitness function) of each seller is to maximize a weighted average of its profit and the price in each auction. The objective (fitness function) of each buyer is maximize the weighted difference of its consumer surplus and the price in each auction. The weights are chosen to give approximately equal importance to profit and price for sellers, and to consumer surplus and price for buyers.

The sellers learn to drive prices and profits up; the buyers learn to drive prices down and consumer-surpluses up. But the buyers use their greater autonomy to get a better than equal share of the gains (Figures 4.7-4.10).

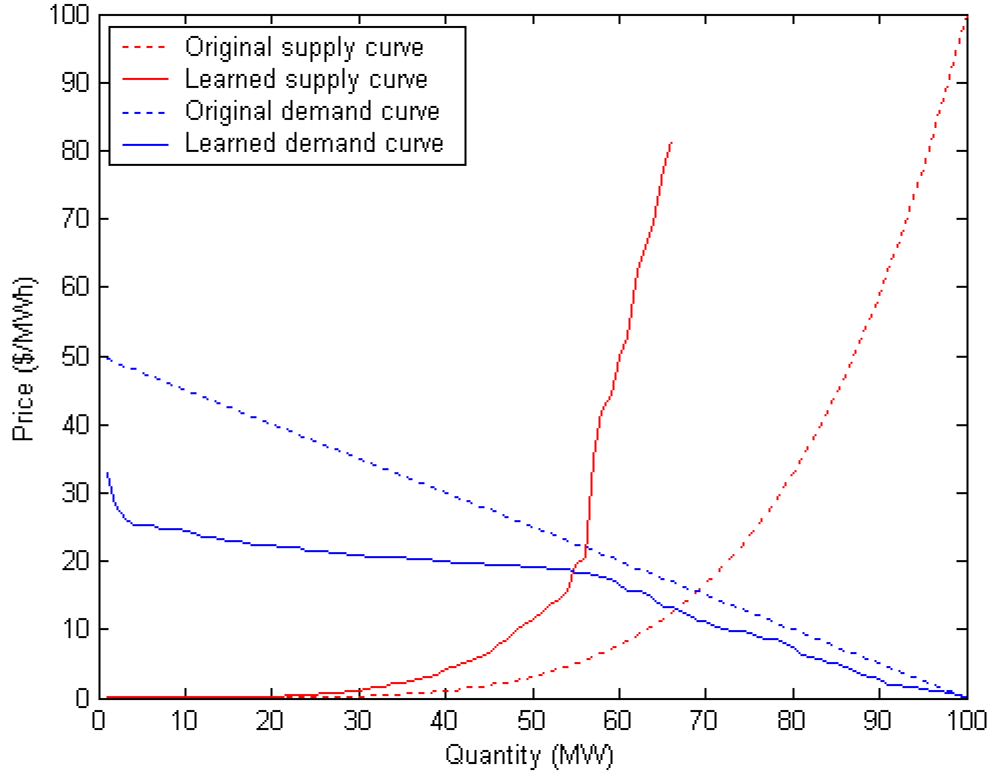

Figure 4.7 Some results from Experiment-4. Both the suppliers and the buyers are allowed to learn. Consequently, both the offers and the bids change from one repetition to the next. The broken curves represent the competitive supply and the actual demand. The solid curves show the supply that is offered and the demand that is bid at repetition-1000.

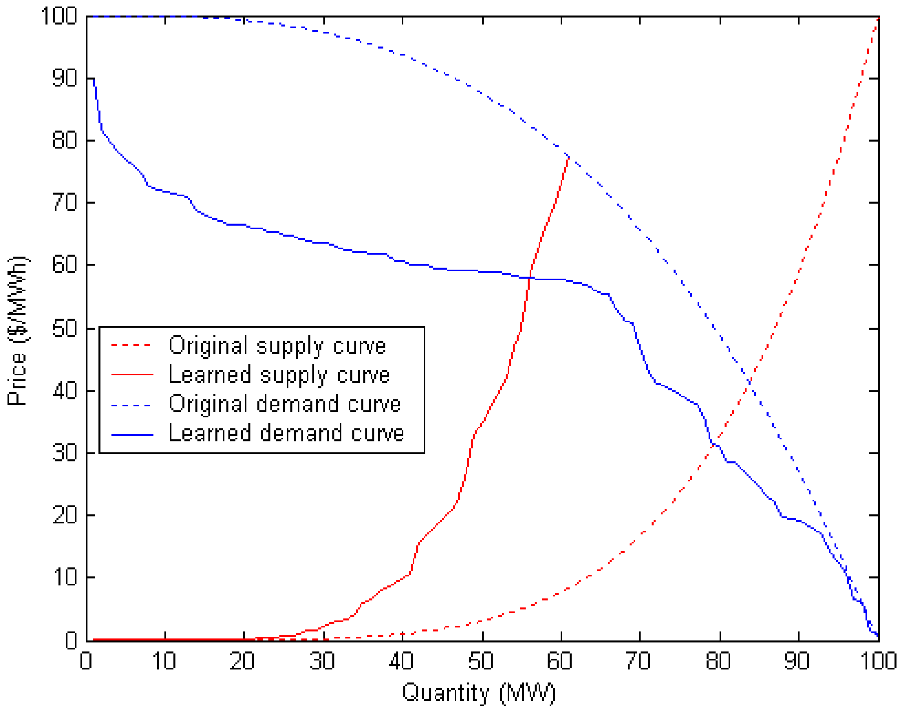

Figure 4.8 More results from Experiment-4. The broken lines represent the supply and demand curves at repetition-1. (The demand corresponds to Curve-10 in Fig 4.2). The solid lines show these curves in repetition-1000. The changes are learned.

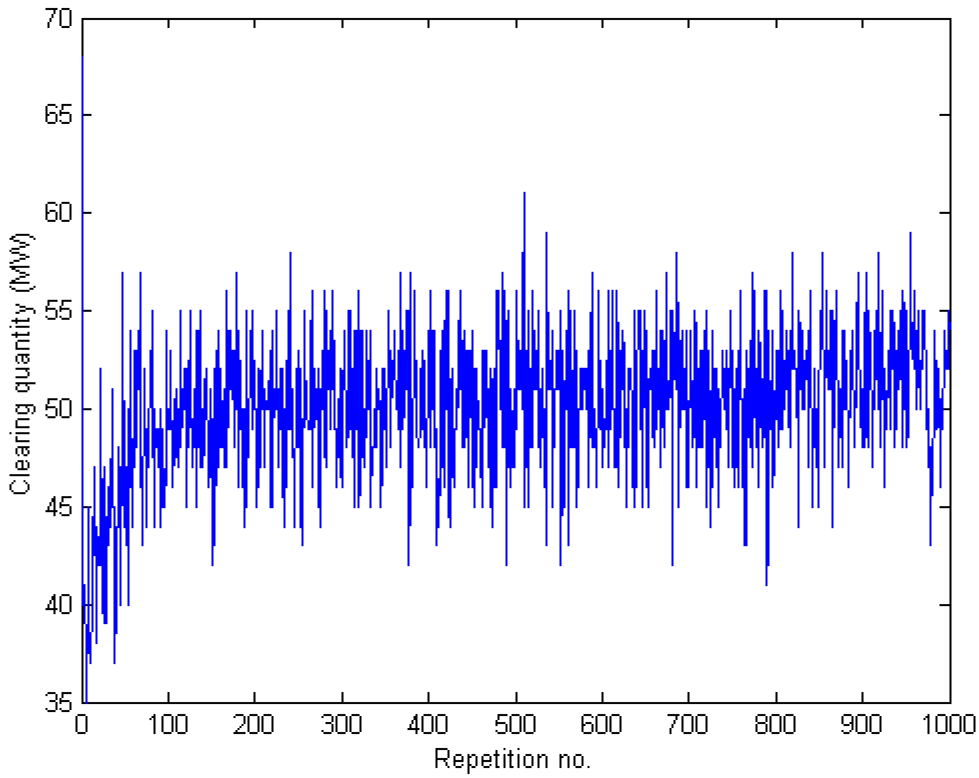

Figure 4.9 A trajectory of clearing quantities from Experiment-4 with the demand corresponding to of Curve-4 in Fig. 4.2.

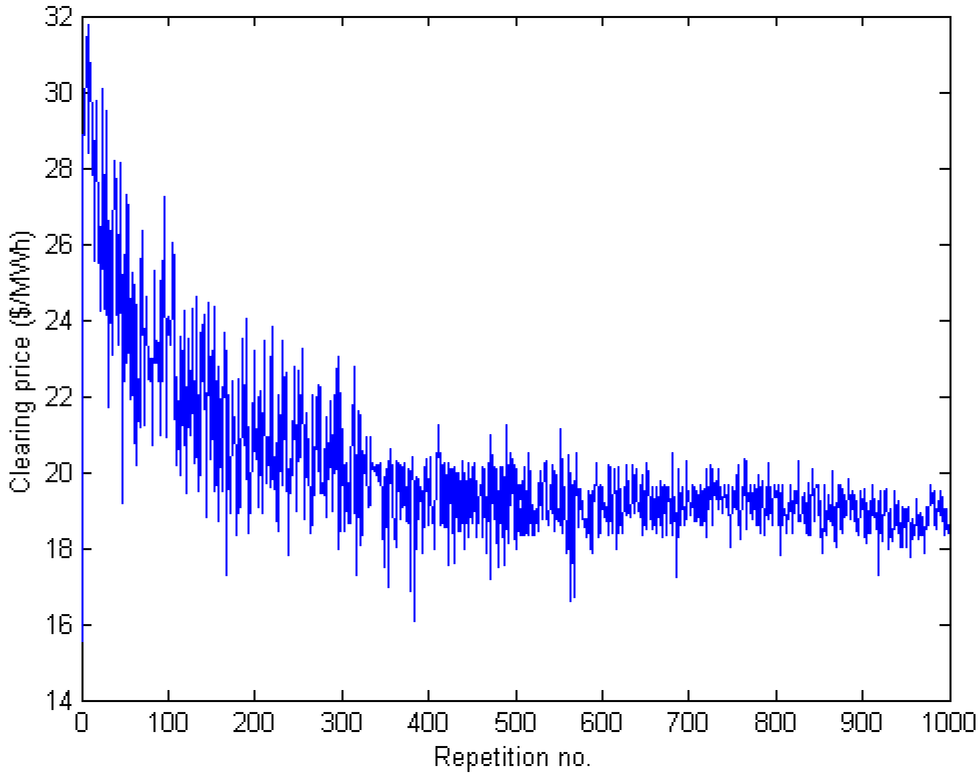

Figure 4.10: A clearing price trajectory from Experiment-4 with the demand corresponding to Curve-4 of Fig. 4.2.

4.8 Conclusions

Lessons from Experiments 1-4

The performance of repeated auctions is extremely sensitive to the effects of learning by the participants. If an auction is repeated many times, and if only the suppliers are allowed to learn, then, even though the suppliers work independently, they can learn strategies that produce high prices and profits.

The benefits of learning are increased, if each supplier is able to estimate the profits of all the others, and if each supplier is given to reciprocal altruism, at least to the extent of seeking to increase total rather than individual profits, and taking its turn to set the clearing price.

A price-responsive demand, even an elastic demand, is not a reliable counter to the advantages that suppliers can gain from learning. While a price-responsive demand can reduce the absolute value of the profits, it cannot keep suppliers from extracting almost as much profit as a monopolist would. (This result holds over a far wider range of experiments than those reported here, including the injection of random noise into the demand, and allowing only some of the suppliers to learn, while others make competitive offers [Kong-Wei Lye, Ph.D thesis, 2004]).

One way to reduce the advantage that sellers gain from learning and reciprocal altruism, is to make the market symmetric, that is, to allow the buyers as much autonomy as the sellers, so the buyers can adjust their demands, place bids for energy in real-time, and learn strategies, just as the sellers do. Active (autonomous) buyers facing passive sellers can push the price down below the competitive equilibrium. These active buyers are able to collect monopsony profit. A market that puts buyers and sellers on the same footing should result in a market clearing quantity that is close to the competitive equilibrium, and a price that could be anywhere between the monopsony and monopoly prices, depending on the bargaining power of each side.

Open vs. closed-loop design

Markets, like bridges, airplanes and computers, are complex artifacts. But unlike these other artifacts, new market-designs are not carefully verified before being implemented.

Good designs of complex artifacts cannot be obtained by deductive (open loop) reasoning alone. The specifications of complex artifacts cannot be reduced to manageable sets of constraints that can be directly translated into designs. Take the specification of competitive behavior, for instance. Many people seem to believe that it can be reduced to three constraints:

a) multiple, profit-seeking suppliers,
b) no collusion among the suppliers, nor any supplier with market power, and
c) elastic demand.

But, as our experiments show, these constraints are far from sufficient for competitive behavior in repeated auctions.

Abduction (iterative, closed-loop reasoning) is vastly more effective in the design of complex artifacts than open-loop deduction. The abductive process contains the following steps: develop specifications, develop a design, develop a set of tests, verify the design (subject it to the tests), and repeat till the design passes all the tests.

A design is flawed to the extent by which it can violate its specifications. For the abductive process to identify and correct flaws, two things are essential: precise and complete specifications, as well as comprehensive and reliable tests. Neither has been available for electricity markets. Therefore, it is to be expected that existing market designs contain more flaws than have yet been detected. It is a matter of chance, which of these flaws will remain hidden, and which will be revealed through future mishaps.

The evolutionary learning algorithm used in our experiments was designed for simplicity. It is not good enough to serve as a verification tool. Much more work will be required to develop a suite of computer-based verification tools for electricity auctions. But until such a suite is available, electricity auctions cannot expect to reach the levels of performance and reliability of other engineered artifacts--bridges, airplanes and computers, for instance, unless market designers forego novel designs in favor of incremental changes to existing and proven designs.

4.9 Acknowledgement

The work reported in this chapter was performed by Kong-Wei Lye, Lester Lave, Eswaran Subrahmanian and Sarosh Talukdar

4.10 References

[1] Apt, J. "Electricity Prices After Deregulation: A Comparison of Regulated and Deregulated States," CEIC Working Paper, Carnegie Mellon, 2004

[2] Bacon, R.W. and J. Besant-Jones, *Global Electric Power Reform, Privatization, and Liberalization of the Electric Power Industry in Developing Countries.* Annual Revue of energy and Environment, 2001. 26: p. 331-359.

[3] Bjornsson, H., R. Crow, and H. Huntington, *International Comparison of Electricity Restructuring Considerations for* Japan. 2004, Center for Integrated Facility Engineering: Stanford.

[4] Dolbear, F. T.; L. B. Lave; G. Bowman; A. Lieberman; E. Prescott; F. Rueter; R. Sherman "Collusion in Oligopoly: An Experiment on the Effect of Numbers and Information," *The Quarterly Journal of Economics*, Vol. 82, No. 2. (May, 1968), pp. 240-259

[5] William W. Hogan, "Successful Market Design (SMD) and Failure Diagnosis: Blackouts and Lampposts in Regulating Electricity markets, Kennedy School, Harvard University, 2003.

[6] Joskow, Paul L., "California's Electricity Crisis," working paper, Department of Economics, MIT, 2001.

[7] Joskow, Paul L., "Regulation and Deregulation After 25 Years: Lessons Learned for Research in Industrial Organization," Working paper, Department of Economics, MIT, 2004.

[8] Lye, K. W., "Learning Agents in Power Markets," Ph.D thesis, Carnegie Mellon University, 2004

[9] Marsall, K. C., S. Talukdar, "Evolutionary Learning in Uniform Electricity Auctions," working paper, Carnegie Mellon University, 2003.

[10] Mount, T. W.D. Schulze, R.J.Thomas, R.D. Zimmerman, "Testing the Performance of Uniform Price and Discriminatory Auctions," Working paper, Department of Applied Economics and Management, Cornell University, 7-16-01

[11] Rassenti, Stephen J., *Vernon Smith*, & B. J. Wilson, "Discriminatory Price Auctions in Electricity Markets: Low Volatility at the Expense of High Price Levels," Journal of Regulatory Economics, 2003, vol. 23, issue 2, pages 109-23.