On the First Price Spike in Summer

Tim Mount tdm2@cornell.edu
Hyungna Oh ho26@cornell.edu

Department of Applied Economics and Management
Cornell University

Abstract

The objective of this paper is to determine why price spikes occur in deregulated wholesale markes for electricity, and how effectively they can be mitigated by different modifications to the market.. The analysis employs a multi-agent system (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 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 main results for creating price spikes are 1) uncertainty about the system load is an important determinant of observed behavior (i.e. offer curves shaped like a hockey stick) that is replicated in the MAS, and 2) all firms eventually become speculators, and it is unrealistic to expect firms to behave like price takers in a market with six suppliers. The main results for mitigating high prices are 1) it is impractical to rely on more suppliers, vertically integrated firms, capacity payments or fixed contracts for power to eliminate speculative behavior, and 2) price responsive load is an effective way to mitigate high prices, particularly if some load responds at relatively low prices. Overall, the results show that a MAS can be used to evaluate a wide range of policy options and supplement the results of market tests using human subjects.

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 inevitably leads to suppliers exaggerating their true costs (Wolack, 2003). The overall result is quasi-regulation and endless litigation about costs. Markets should be more robust to speculation, and more specifically, firms submitting high offers should 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 our analysis, suppliers try to maximize profits with no fear of regulatory intervention.

Different modifications to the market are evaluated, and suppliers can withhold capacity and speculate with high offers if it is in their own interest to do so. The results of our analysis show that it is impractical to rely on 1) adding more suppliers, 2) having vertically integrated firms, 3) making capacity payments, and 4) contracting for additional power to eliminate speculative behavior. These modifications to the market do reduce average prices, but the effects are relatively small. The best way to mitigate high prices is to introduce price responsive load, and the effectiveness of this modification depends on how load responds. In particular, small responses at low prices are more effective than a large response at a high price.

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 computer 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

0-7695-2056-1/04 $17.00 (C) 2004 IEEE

firm to determine the optimum set of offers to maximize expected profits in the auction.

One advantage of using computer 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 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 can be met by specifying a specific functional form for the residual demand curve. 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 in response to market outcomes allows firms to modify their optimum strategy during a simulation. An interesting result from the analysis is showing why supply curves are shaped like a hockey stick when load is stochastic. This is exactly the type of behavior observed in electricity markets and it accounts for the existence of price spikes as a typical feature of these markets.

The next section of the paper discusses the modeling framework, followed by an explanation of why price spikes occur in Section 3. The effects of different modifications are summarized in Section 4, and the conclusions are given in the final section.

2. A Multi-Agent System for Testing Electricity Markets

A Multi-Agent System (MAS) consists of three main elements: an environment, a set of agents, and a set of tasks. Figure 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.

Environment: Spot Market

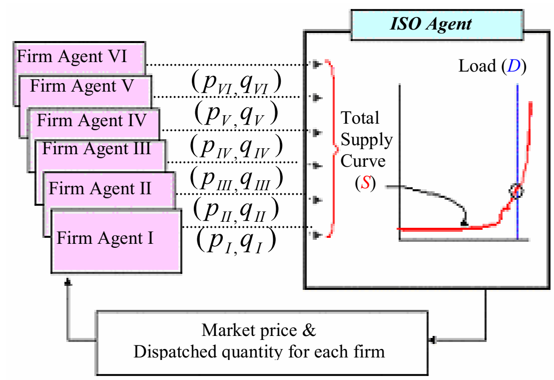

Figure 1. A Multi-Agent System for an Electricity Spot Market

An *agent* is a computer program that is capable of autonomous action to meet specified objectives in the environment (Weiss, 1999). Figure 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 1), 3) determine the optimum dispatch to minimize the cost of meeting the current load (D in Figure 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.

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 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. First, 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 reparameterization to $\alpha_t = a_t + b_t\,OC_t$ is convenient because OC_t is unobserved and itavoids 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. 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 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 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.

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 1.

Table 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 β > 0 imply that at least half of the firm's capacity is withheld. When the 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 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.

3. Creating Price Spikes in a MAS

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 (40, 25, 20, 10, 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.

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 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 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 in this case, because the firms learn how to exploit market power more effectively during the simulation.

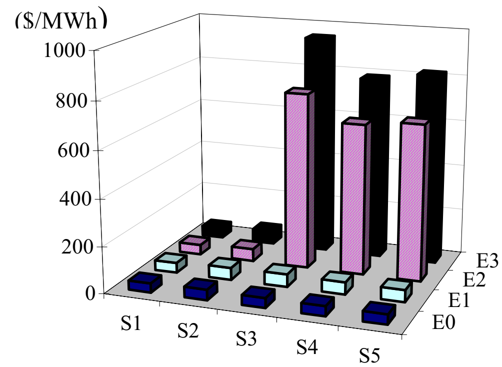

Figure 2. Average Energy Price (Simulated)

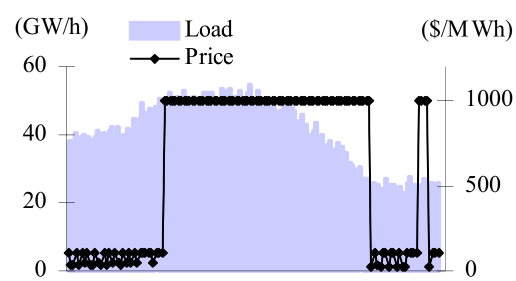

Figure 3. Load and Energy Price for E3/S5

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 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. 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 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 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 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 spice provides a major source of new information for price takers, but price spikes are not necessary for price takers to adapt and become speculators.

Table 2. The Number of Firms by Belief

		<$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

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.

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 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 5 shows the probability of rejecting different offers in terms of the likely range of actual loads when the forecast 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. The next section addresses the issue of how to mitigate high prices and replicate a market with only a few price spikes, like PJM.

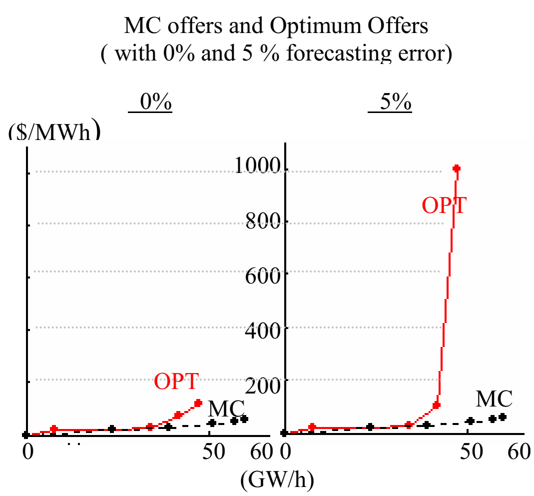

Figure 4. Load Forecasting Error and the Optimum Offer

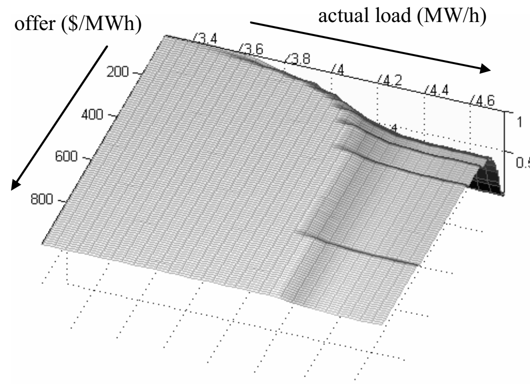

Figure 5. Probability of Rejection

4. Mitigating High Prices in Electricity Markets

4.1 Modify the MAS Environment to Represent the Pattern of Load in PJM

4.2 Too Many Firms are Required to Make Textbook Competition Practical

4.3 The Effect of Vertically Integrated Firms is Relatively Small

4.4 Capacity Payments for Availability Reduce Withholding but may Increase Prices

4.5 The Effect of Purchasing a Fixed Quantity of Imported Power is Relatively Small

4.6 The Effect of Price Responsive Load is Much Larger

5. Conclusions

The results presented in Sections 3 and 4 demonstrate that a MAS for simulating a deregulated wholesale market for electricity is a viable tool for evaluating how well different markets perform. New insights into why price spikes are observed in deregulated markets and how to mitigate them were uncovered during the analysis. The use of a structured form of Adaptive Autonomous Agent to represent suppliers replicated the type of behavior seen in deregulated markets. Furthermore, results were obtained within the desired limits on the number of suppliers in the market and the number of trading periods. Keeping these numbers small makes it practical to compare the results from the MAS with results from the standard market tests at Cornell using PowerWeb. In other words, computer agents and human subjects can be compared, and our 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 practical to use a MAS to test the sensitivity of results to changes in the market environment. With human subjects, there are severe limits on the number of replications for one market environment and on the number of different modifications that can be tested.

The structure of the MAS uses six computer 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. These estimates are updated using a Kalman filter, and as a result, the behavior of the firms can change during a simulation. The objective of each firm is to maximize expected profits. The analysis in Section 3 focuses on creating price spikes and the analysis in Section 4 deals with their mitigation.

With a relatively small number of firms, it is very difficult to avoid getting price spikes when load is stochastic. In a deterministic optimization with a fixed load, six suppliers are relatively competitive, and these results from the MAS are consistent with previous results 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 load opens up opportunities for speculation, and firms submit offer curves shaped like a hockey stick. This is exactly the type of behavior observed in deregulated markets. Once firms learn to speculate, they remain as latent speculators even if 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 firms to submit high offers. However, regulators have tended to treat these high offers as abhorrent behavior, particularly in the USA. It would be more realistic to accept the inevitability of hockey stock supply curves, and to look for other ways to mitigate high prices. This is the subject of Section 4.

Different market modifications were tested by the MAS, and each one was expected to reduce the number of high prices. For these simulations, the pattern of load corresponded to the daily pattern for PJM over the summer of 1999. In the base case with six suppliers, high prices occur most of the time when the load is high, and the average price is much higher than the observed average price in PJM, and over ten times the competitive level. This is a highly dysfunctional market. Nevertheless, it was still surprising how difficult it was to lower prices. One reason is that a single price spike of \$1000/MWh corresponds roughly to an increase of \$7/MWh in the average price, and having price spikes only 5% of the time would make the price twice the competitive level.

The results in Section 4 show that it is impractical to rely on some of the standard remedies for reducing high prices. Increasing the number of firms, changing suppliers to vertically integrated firms, making capacity payments for availability, and shifting the supply curve by importing power have relatively small effects over realistic ranges. One major surprise was that the average price increased with capacity payments because firms found it easier to speculate (there was no threat of regulatory intervention in our simulations).

The best way to mitigate high prices is to introduce price responsive load. This type of response is effective over the inelastic portion of a hockey stick supply curve, and this is exactly where price mitigation is needed. Our results show that the structure of load response matters, and small reductions at relatively low prices are more effective than a large reduction at a high price. The latter type of response is similar in form to a Demand Reduction Program that pays customers to reduce their load. This type of program has been very effective in New York (see Neenan et al., 2002), for example, and our results suggest that other types of program may be even more effective. The reluctance of incumbent utilities to develop innovative programs for reducing load when load is high remains a problem for customers in deregulated markets. Regulators should put more emphasis on ways to encourage load response and less on trying to modify the behavior of suppliers in the market.

The current approach used by regulators to mitigate high prices is to make it difficult for firms to submit high price offers. Keeping price offers low is an effective way to keep market prices low, but the implications are ominous. The primary regulatory mechanism is to require that high price offers must be cost justified. For example, this approach was used by the Federal Energy Regulatory Commission (FERC) when FERC intervened in California in the fall of 2000. The problem with this approach is that firms have an obvious incentive to exaggerate costs. Trying to regulate how firms behave in an auction in this way opens the door for endless litigation. A self-regulating market should discipline firms that submit high price offers by reducing the amount of their capacity that is dispatched. Making load more price responsive is a highly effective way to do this. In some circumstances, regulatory intervention may be needed when, for example, a specific power plant is essential for providing ancillary services. In most other situations, interventions by regulators in the market should not be necessary.

6. Acknowledgements

This work was supported in part by the US Department of Energy through the Consortium for Electric Reliability Technology Solutions (CERTS) and in part by the National Science Foundation Power Systems Engineering Research Center (PSERC).

7. 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. <www.ucei.org>.

[4] 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.

[5] 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.

[6] Federal Energy Regulatory Commission (FERC), (2000). "Commission Proposes to Reshape California's 'Seriously Flawed' Electricity Markets with Sweeping Changes," News Release, November 1.

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

[8] 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.

[9] 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.

[10] 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.

[11] 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.

[12] 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.

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

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

[15] 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.

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

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

[18] 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.

[19] 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. <www.ucei.org>.

[20] 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.