An Analysis of Price Volatility in Different Spot Markets for Electricity in the U.S.A.

by

Tim Mount*
Yumei Ning**
Hyungna Oh**

Department of Agricultural, Resource, and Managerial Economics
Cornell University
e-mail: tdm2@cornell.edu

Abstract

Earlier research has shown that the behavior of spot prices in the new auction markets for electricity can be described by a stochastic regime-switching model. This model captures the observed price spikes that occur in these markets, particularly during the summer months when levels of load are high. The first part of the paper shows how the exploitation of market power can lead to offers to sell power that are consistent with price spikes. An important feature of the model is that some suppliers are indifferent to having marginal units dispatched when they have sufficient market power. Given this analytical framework, the second part of the paper extends the regime switching model of prices by making key parameters functions of forecasted load. The first application shows how the structure of the PJM market changed when market-based offers were allowed, resulting in higher price spikes. The second application compares price behavior in PJM, New England and California. The transition probabilities in the three markets have similar relationships to load. The main differences among markets are the levels of the means in the high-price regime, and in this respect, PJM is quite different from the other two markets. Efforts to associate price spikes with errors in the forecasts of load or changes of actual load were not successful. The conclusion is that more research is needed to understand the motivation of suppliers submitting offers into an auction market.

* Professor of Resource Economics
** Graduate Research Assistants

1. Introduction

The institutional designs of restructured markets for electricity in the USA vary in fundamental ways. In particular, markets in the East, such as New England, New York and Pennsylvania, New Jersey and Maryland (PJM), place the responsibility for unit commitment and for dispatching generators on an Independent System Operator (ISO). In contrast, the system on the West Coast is much less centralized, and suppliers and regional coordinators can specify their own dispatching schedules based, in many cases, on bilateral trading. The ISO is responsible for making incremental changes to the planned dispatching pattern in order to, for example, avoid violating constraints on the capacity of the transmission network.

In spite of these differences in the structures of the markets in the USA, they share one important feature. Spot prices for electricity have been very volatile with dramatic price spikes occurring at certain times (see Figure 1, and Section 5 for a comparison of prices in three markets for summer 1999). While the degree of volatility has been more severe in some markets (e.g. PJM) than others (e.g. California), it is still true that prices reach levels that are much higher than the typical values of marginal costs (system lambda) under regulation. The primary objective of this paper is to provide new insight into the behavior of spot prices for electricity, and in particular, to model the type of volatility shown in Figure 1. For electricity, volatility is associated with price spikes, and higher volatility implies higher average prices as well. This is bad news for customers and good news for suppliers.

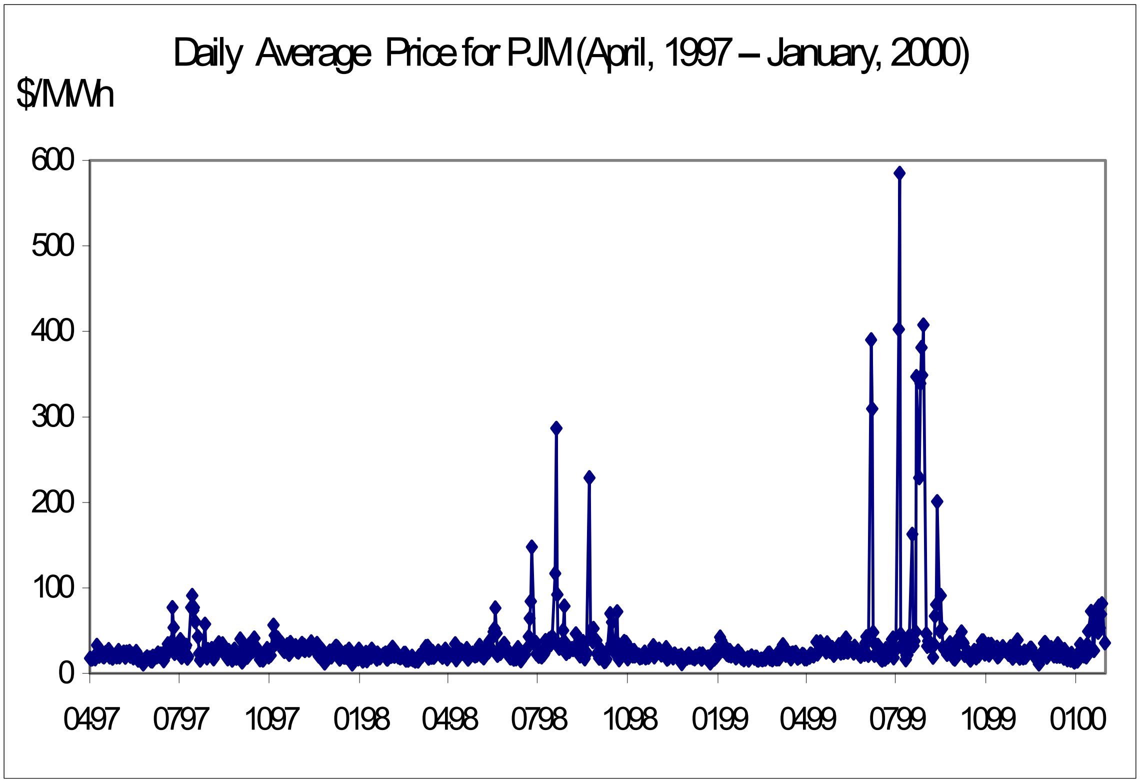

Figure 1

Simple regulatory "solutions" to price spikes may be counter productive. For example, putting a low cap on prices is one way to reduce those prices. However, this mechanism also eliminates the legitimate role of high prices by sending a signal that additional supply is needed, in response to a forced outage, for example. The real problem for customers occurs when suppliers are able to produce high prices artificially. The most obvious way to get high prices is by exploiting market power. This situation has proved to be a chronic problem in the UK market, which is still dominated by two suppliers. The blatant use of market power in the UK market has led to a radical restructuring from a centralized ISO to a design that is much more like the market in California. In contrast to the UK experience, the market in Australia was designed in such a way that market power is not a major problem (each power plant in Victoria was sold to a different company). The lesson is that it is naive to assume that reliable and efficient markets can be designed when the actual motivation of suppliers is poorly understood. In contrast to Australia, the restructuring process in the USA showed little concern about the consequences of market power when incumbent utilities sold power plants in groups to individual buyers.

The actual problems with restructured markets for electricity are even more serious than the volatility of the spot prices for energy. Energy represents just one of the markets for supplying electricity to customers. The analytical and public policy problem is that there is more scope for gaming in complicated markets. The UK market provides an excellent example of how interdependencies among markets can be exploited (see Newbery, and Wolak and Patrick). The overall performance of restructured markets in the USA has not been a great success to date. As a result, substantial changes will be made this year in the designs of the markets in PJM, New England, and California. Although the focus of this research is on energy only, it is fair to say that an economically efficient and reliable market for energy is necessary if the overall market is to perform in a satisfactory way. Furthermore, the type of behavior exhibited by suppliers in the energy market is likely to be repeated in the markets for ancillary services. The specific objectives of this paper are to present 1) a theoretical model of how market power affects optimum offers to sell energy, and 2) a statistical model of price behavior that is used a) to evaluate the increasing price volatility in the PJM market (see Figure 1), and b) to compare the performance of the markets in PJM, New England, and California.

2. Market Power and Optimum Offers to Sell Power

If it is assumed that suppliers follow the standard economic principle of maximizing profits, it is possible to derive optimum offers for selling power in a uniform price auction. The implications of this type of behavior are explained by Mount (2000) when suppliers face uncertainty about the actual load that will determine the market clearing price. The analysis compares the behavior of a "large" supplier controlling 20% of the expected load with a "small" supplier controlling only 4%. Both suppliers are assumed to have the same structure of costs for their own generating capacity and face the same competitive market as other suppliers (i.e. other suppliers submit cost-based offers into the auction).

The actual marginal costs and the optimum offers are shown in Figure 2 for the two suppliers. (To make comparisons easier, the horizontal scale is the percentage of owned capacity.) The large supplier owns 5000 MW and the small supplier owns 1000 MW). The two important conclusions from Figure 2 are 1) that the offers for the large supplier are substantially above the offers for the small supplier even though the structure of costs is identical, and 2) the difference between an offer and the true marginal cost is greater for the unit with the highest cost for both suppliers. The overall implication is that it is rational for a supplier with market power to speculate with a marginal unit by setting the corresponding offer at a high level. This is essentially what happened in the PJM market last summer. Nevertheless, the optimum offers shown in Figure 2 are still much lower than the actual offers that set prices in the PJM market during summer 1999.

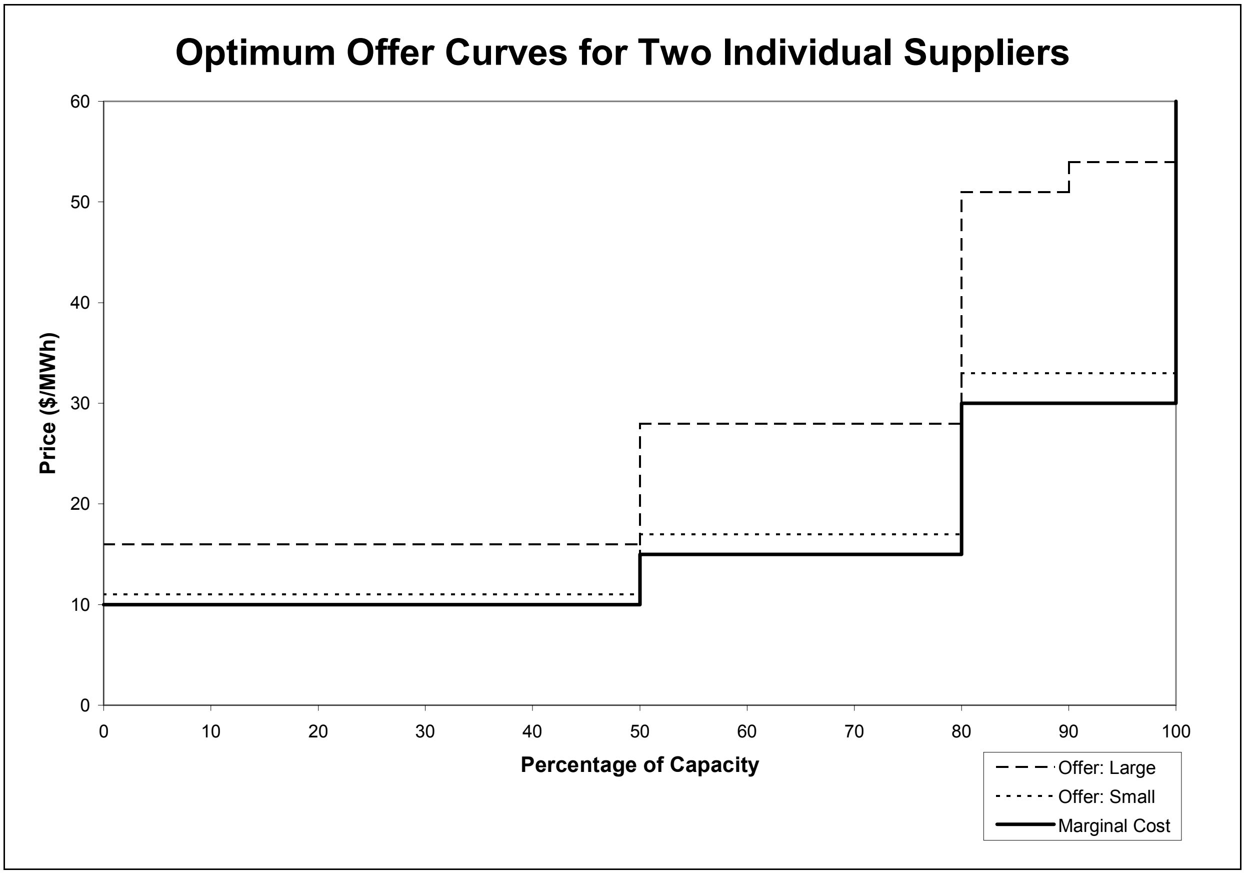

Figure 2

An explanation for why actual offers for marginal units were so high can be determined from Figure 3, which shows the expected excess profits per hour above the competitive level (i.e. the effect of submitting offers for the most expensive unit above the true marginal cost). The maximum expected profit for the large supplier ($3700/hour at Offer = $54/MWh) is over 100 times higher than the maximum for the small supplier ($30/hour at Offer = $33/MWh) even though the large supplier only controls five times the capacity of the small supplier. For the small supplier, the maximum at 33 is obvious in Figure 3, but for the large supplier, the maximum at 54 is almost identical to the levels of excess profit for any Offer > 54. The reason is that setting Offer = 54 is virtually equivalent to withdrawing the unit from the market because the probability of having the

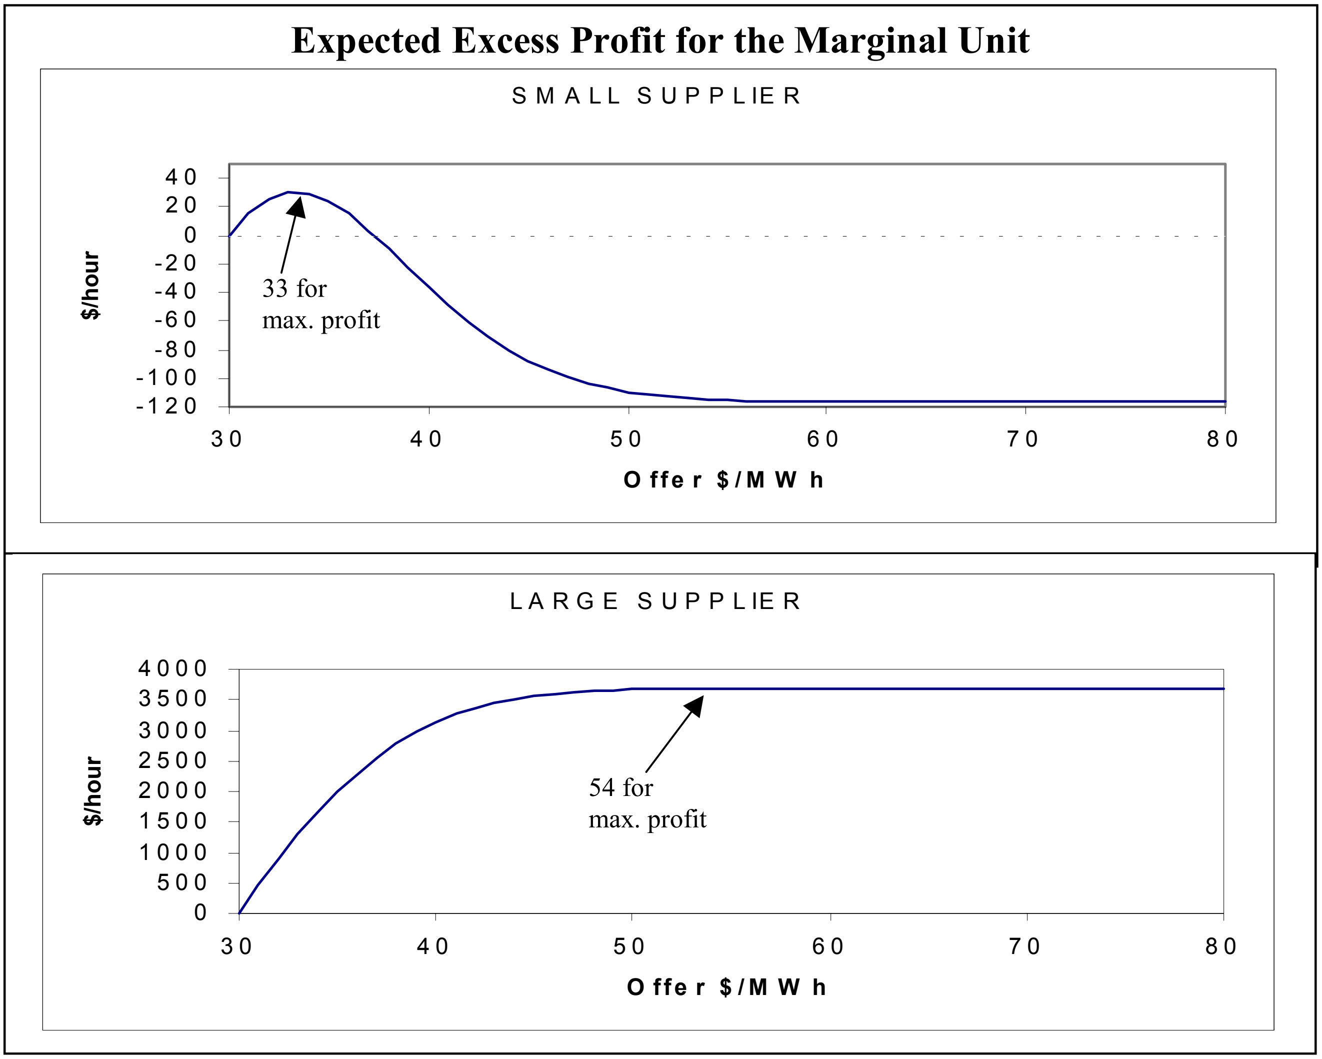

Figure 3

unit dispatched is close to zero. Consequently, the penalty from submitting high offers, in terms of lower expected profit, is trivially small for the large supplier. This is not the case for the small supplier. For Offer > 37, the expected excess profit is negative and the possible loss of up to \$116/hour is much larger than the potential gain of \$30/hour. For all practical purposes, the large supplier is completely indifferent between the optimum offer and any higher value of the offer. If there is a small probability that all available capacity will be needed to meet load, then it is rational for the large supplier to put in an offer for the marginal unit at the maximum allowed level. (In the model underlying Figure 3, it is assumed that other sources of generation are always available, and the demand curve faced by an individual suppler is never completely inelastic).

Table 1. Probabilities for a Marginal Unit

Dispatch Status	A	B	C
1. Small suppliers (MC = 30, Offer = 33)			
Competitive	0.50	0.02	0.48
Optimum Offer	0.33	0.02	0.65
2. Large supplier (MC = 30, Offer = 54)			
Competitive	0.50	0.10	0.40
Optimum Offer	<0.01	<0.01	>0.99
3. Large Supplier (MC = 10, Offer = 29)			
Competitive	>0.99	<0.01	<0.01
Optimum Offer	0.55	0.10	0.35

A Probability {Unit is fully dispatched} (i.e. Market Price > Offer)
B Probability {Unit sets price} (i.e. Market Price = Offer)
C Probability {Unit is not dispatched} (i.e. Market Price < Offer)

The best explanation for why price spikes occur is that some suppliers are indifferent to whether or not marginal units are dispatched. Once an offer is high enough to make it almost certain that the corresponding unit will not be dispatched, it does not make any difference to the expected profit if an even higher offer is submitted. These conclusions are illustrated by an example in Table 1 which summarizes the probabilities of a marginal unit 1) being fully dispatched, 2) setting the market price, and 3) not being dispatched for three different cases. In each case, the probabilities for the competitive solution and the optimum offer are computed, and the probability of not being dispatched (column C) is always higher using the optimum offer. For the small supplier (Case 1), there is a penalty for submitting an offer that is too high. This is also true for a large supplier if the marginal cost of the most expensive unit is well below the competitive price (Case 3). In both Case 1 and 3, expected profit can be much lower than the competitive solution if the offer is increased above the optimum level.

Case 2 in Table 1 is the exception, and there is almost no penalty on the expected profit if a higher offer is submitted because the unit will not be dispatched anyway (see Figure 3). The implication that Case 2 is potentially a problem supports the conclusion of Wolak and Patrick that a uniform-price auction is more vulnerable to market power if large suppliers control both base-load and peaking units. For suppliers owning units like Case 2, substantial excess profits can be earned even though there is less than a 1% chance of having the unit dispatched. It is equivalent to a reward for idleness. For all practical purposes, the offer could be $54/MWh or $5000/MWh without affecting expected profits. There is simply no incentive for the supplier to operate the unit. In the next section, a statistical model of price behavior is presented which reflects the structure of observed offer curves by using stochastic switching between a high-price and a low-price regime.

3. A Stochastic Model of Price Behavior

Recently, data for the actual offers submitted into the PJM market have been released for the spring and summer of 1999. The offer curves (actual supply) and the forecasted loads (demand) for three typical days (the first Tuesday of each month) are shown in Figure 4. The offer curves cover the transition from cost-based offers (prior to 4/1/99) to market-based offers. All three offer curves exhibit a distinct kink after which offers rise almost vertically. The actual load (demand) on July 6 resulted in a substantial price spike with some hourly prices reaching the maximum permitted level of $999/MWh (see Figure 1). The general shape of the offer curves provides direct evidence that market prices are set in one of two possible regimes. Whenever the forecasted load is close to the kink in the offer curve, price spikes are likely to occur if actual load is higher than the forecasted level. It is interesting, and somewhat ominous, that the kink for April, when load is low, occurs at a much higher level of capacity than the kink for July, when load is high.

The kinked shape of the offer curves in Figure 4 is extreme, and it probably reflects more than the simple exploitation of market power described in Section 2. Suppliers in PJM were permitted to sell capacity in the PJM capacity market and to sell non-firm energy from the same units in another region. (This is also the case in the current New York market, which started in November 1999.) Units accepted in the capacity market are subject to recall if capacity is needed to meet load in PJM. Under these circumstances, any contracts for the non-firm energy would be broken. However, recall could only occur when all units offered into the PJM energy market have been dispatched. Hence, offering a marginal unit into the PJM energy market at $999/MWh acts as an attractive hedge against having capacity recalled.

The behavior of prices in Figure 1 implies that the standard stochastic model of geometric brownian motion (i.e. the logarithm of price is a random walk) used in financial analyses is not appropriate. Mean reversion is a better choice for electricity spot prices as noted in Pilipovic, Tseng & Barz, Deng, Johnson, & Sogomonian, and Ethier. This model incorporates information about a long-run mean value and has a bounded long-run variance.

Ethier and Mount show that Hamilton's (1989) Markov regime switching model is well suited for capturing the behavior of electricity spot prices. Adding Markov regime switching to the model allows for stochastic price jumps. Each regime is a mean

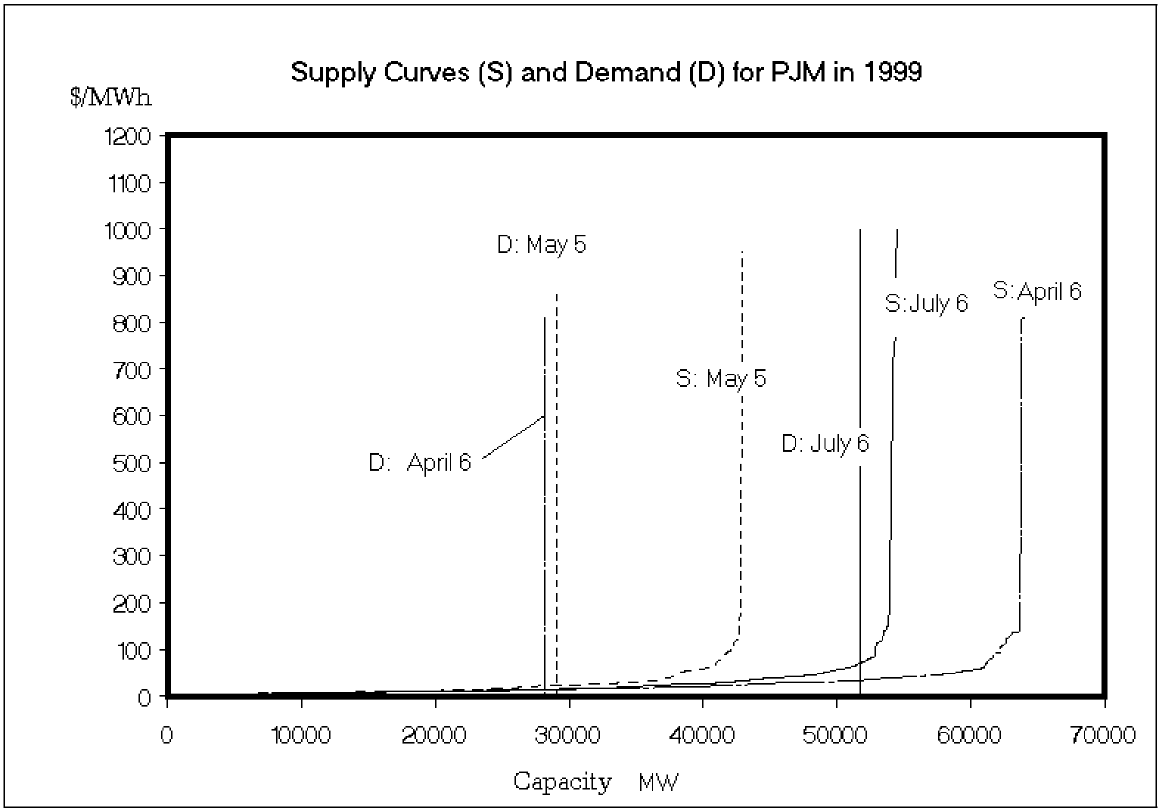

Figure 4

reverting AR(1) process. One regime represents the relatively flat part of the offer curve (the low-price regime), and the other regime represents the steeply sloped part (the high-price regime). In their analysis, a model with constant parameters is estimated separately for each season. The regimes show that the estimated parameters have distinct seasonal patterns and conclude that a model with varying coefficients would be more appropriate. Furthermore, Diebold, Lee and Weinbach have argued that treating transition probabilities as fixed parameters is restrictive. Consequently, a model with time varying coefficients was developed for this paper.

In Hamilton's original model, the current mean of y_t (the logarithm of price) is a function of both the current regime and the lagged regime. The regression relationship for y_t can be written as follows:

$$y_t - \mu_{s_t} = \phi\left(y_{t-1} - \mu_{s_{t-1}}\right) + \varepsilon_t \tag{1}$$

where S_t = 1 or 2 identifies the regime

μ_{s_t} is the mean for regime S_t

$0 \le \phi < 1$ is the AR coefficient

ε_t is an unobserved residual that is $N(0, \sigma^2_{s_t})$.

Specifying that the current price is a function of both the current and lagged mean creates computational complexity for estimation (particularly when more then two states are included in the model). Hamilton (1994) and Gray propose a computationally more tractable model in which the current price is a function of the lagged price, but not the lagged mean (i.e. it is not a function of the lagged state S_{t-1}). It turns out that this model has a higher likelihood value and a lower mean-squared error using data for PJM (summer 1999) compared with the original model.

The additional feature of the new model is that the key parameters are functions of time, and specifically of forecasted load. A general specification for the model can be written as follows:

Conditional Distributions

y_t is $N(\mu_{1t}, \sigma^2_1)$ if $S_t = 1$

y_t is $N(\mu_{2t}, \sigma_2^2)$ if $S_t = 2$

where y_t is the logarithm of price

μ_{it} is the conditional mean for regime $S_t = i$

σ_i^2 is the variance for regime $S_t = i$

In regression form, the model for y_t is mean reverting and a function of an additional regressor:

Regression Form

$$y_t = \mu_{it} + \varepsilon_{it} = \alpha_i + \phi_i y_{t-1} + \gamma_i x_t + \varepsilon_{it} \tag{2}$$

where α_i, ϕ_i and γ_i are parameters for regime $S_t = i$

x_t is the forecasted load

The implication is that the conditional mean of y_t ($E[y_t \mid x_t, y_{t-1}, s_t, = i] = \mu_{it}$) varies with time because forecasts of load change as well as varying from the dynamic properties of an AR(1) model. The use of forecasted load rather then the actual load was chosen to make the conditional means follow a relatively smooth time path.

The final component of the model specifies the transition probabilities for regime switching as logistic functions:

Transition Probabilities

$$\begin{aligned} \Pr[S_t = 1 \mid S_{t-1} = 1] &= P_{1t} \\ \Pr[S_t = 2 \mid S_{t-1} = 1] &= 1 - P_{1t} \\ \Pr[S_t = 2 \mid S_{t-1} = 2] &= P_{2t} \\ \Pr[S_t = 1 \mid S_{t-1} = 2] &= 1 - P_{2t} \end{aligned} \tag{3}$$

and $$P_{it} = \frac{\exp(c_i + d_i x_t)}{1 + \exp(c_i + d_i x_t)} \quad \text{for } i = 1.2$$

Hence, the transition probabilities are also affected by the forecasted load. Looking at the data for PJM in Figure 1, the expectation is that the probability of switching from the low to the high price regime will be close to zero in the spring and fall when the load is low, and much higher in summer when load is high. An important policy question is whether the high spikes observed in the summer 1999 are due to the high load or to changes in price behavior or both.

One restrictive feature of the specification in (2) is that the variances for each regime are constant, even though Hamilton's standard model allows for ARCH residuals. There are some underlying reasons for specifying constant variances. The first is that the empirical evidence for ARCH behavior in Figure 1, for example, is quite limited compared to price behavior for other forms of energy (see Duffie and Gray). A second reason is that there are distinct analytical advantages from having constant variances when evaluating financial derivatives, such as the price of an option to sell. It should also be noted that the variance of y_t does vary over time, because the probabilities of being in one or the other regime and the conditional means for each regime do vary with time. The variance of y_t can be written as follows:

$$\begin{aligned} \mathrm{Var}[y_t \mid \Phi_t] &= E[y_t^2 \mid \Phi_t] - E[y_t \mid \Phi_t]^2 \\ &= p_t(\mu_{1t}^2 + \sigma_1^2) + (1-p_t)(\mu_{2t}^2 + \sigma_2^2) - (p_t\mu_{1t} + (1-p_t)\mu_{2t})^2 \end{aligned} \tag{4}$$

where $\Phi_t = [y_1, y_2, \ldots, y_{t-1}, x_1, x_2, \ldots, x_t]$ represents the information needed to make a one-step ahead forecast of y_t in each state, and $p_t = \Pr[S_t = 1 \mid \Phi_t]$ is the conditional probability of y_t being in State 1. If State 1 represents the high-price regime with $\mu_{1t} > \mu_{2t}$ and $\sigma_1^2 > \sigma_2^2$, then both the mean and the variance of y_t will be larger when p_t is larger.

Smoothed values of p_t, as well as estimates of the parameters in (2) and (3), are standard outputs from the model, conditional on the full sample. The model is estimated by maximum likelihood using non-linear optimization routines in the computer package GAUSS. A full description of the algorithms used for estimation will be available in a future publication. The following two sections present results from estimating the regime switching model for the spot prices of electrical energy.

4. Structural Changes in the PJM Market

Prior to April, 1999 offers to sell electricity in the PJM auction were cost-based. After 4/1/99, suppliers were allowed to submit market-based offers. Hence, it is interesting to determine whether this change was responsible for the high price spikes shown in Figure 1. An alternative explanation is that the higher loads in the summer, 1999 were responsible.

The regime switching model described in Section 3 includes forecasted load as one of the explanatory variables. For the following analysis, a simple time-series model with two Sine/Cosine cycles was fitted to the daily peak load for PJM to capture the seasonal pattern of load (weekends were dropped from the data). The corresponding price data are the average daily on-peak prices for weekdays (6 a.m. to 10 p.m.).

Parameters were estimated using prices for 4/1/97 to 3/31/99 (Pre-market based offers) and for 4/1/99 to 1/31/2000 (Post-market based offers). The estimated coefficients are shown in Table 2. Since State 1 is chosen to be the high-regime, representing the price spikes, the corresponding mean is positively related to the load $(\gamma_1 > 0)$, particularly after market based offers are allowed in 1999. The equivalent effects for the low-regime are small and negative. At the same time, the probabilities of switching to the high-price regime ($1-P_{2t}$ in (3)) and staying in the high-price regime (P_t in (3)) are both positively related to load because $d_2 < 0$. The effects of load on the probabilities of staying the high regime are relatively small. The variances of the high-regime are larger than the variances of the low-regimes ($\sigma_1 > \sigma_2$), even though y_t is in logarithms. Finally, the adjustment coefficients (ϕ_i) have similar positive values ranging from 0.31 to 0.57.

The magnitudes of the parameters in Table 2 are relatively difficult to interpret. Consequently, the derived mean prices and ergodic probabilities (unconditional probability of being in the high-price regime) are shown in Table 3 and Figure 5. In Table 3, each characteristic is computed for the maximum, average, and minimum levels of load observed for the market. The main difference between the pre- and post-market based offers is that the mean in the high regime is much higher when market-based offers are allowed. In contrast, the probabilities ($1-P_{2t}$) and P_{1t} are similar between the two models for PJM. In Figure 5, the projected means and probabilities are shown for the post-market period using the estimated model for the pre-market period. The probabilities of switching to the high-regime are slightly lower in the post-market period than the projected values, but the mean price in the post-market period is much higher than the projected mean. Consequently, the weighted averages of the two conditional mean prices (using the ergodic probabilities) are substantially higher during the post-market period than during the projected values.

Table 2. Parameter Estimates for Regime Switching Models

Parameters	PJM Pre-Market	PJM Post-Market	New England	California
α_1	-9.4127	-31.1497	13.5714	-9.8112
	(9.2576)	(32.5712)	(37.6429)	(4.8505)
α_2	3.2386	3.1756	0.5036	0.4193
	(0.2985)	(2.1597)	(1.5062)	(1.9629)
ϕ_1	0.4182	0.5771	0.0000	0.7207
	(0.1477)	(0.1858)	(0.1619)	(0.0488)
ϕ_2	0.4687	0.3142	0.6419	0.6599
	(0.0465)	(0.0385)	(0.0547)	(0.0488)
γ_1	1.1068	3.1441	-0.9589	1.0346
	(0.8206)	(3.1047)	(3.8106)	(0.4711)
γ_2	-0.1540	-0.0926	0.0740	0.0695
	(0.1237)	(0.2090)	(0.1548)	(0.1883)
σ_1	0.5000	0.7431	0.7503	0.6123
	(0.0513)	(0.1048)	(0.1059)	(0.0373)
σ_2	0.2041	0.2400	0.1313	0.1825
	(0.0094)	(0.0134)	(0.0116)	(0.0147)
c_1	19.9710	-3.3616	87.1788	25.1198
	(61.1410	(28.0010)	(138.1942)	(71.3156)
c_2	154.3434	183.4200	183.8625	105.4641
	(47.2509)	(74.0406)	(112.2602)	(70.0729)
d_1	-1.8083	0.4104	-8.7953	-2.1612
	(5.7980)	(2.6359)	(14.0536)	(6.8510)
d_2	-14.4380	-17.1136	-18.4643	-9.8774
	(4.4901)	(6.9854)	(11.4020)	(6.7301)
Log likelihood	458.4724	137.6429	218.0255	317.7057

The overall conclusion from the comparisons in Table 3 and Figure 5 between the pre and post-market periods is that the change from cost-based to market based offers is the primary cause of the high price spikes in the PJM market during the summer 1999. The effect of the higher load was relatively small. The analysis in the next section compares the price behavior in three different markets for June to September 1999.

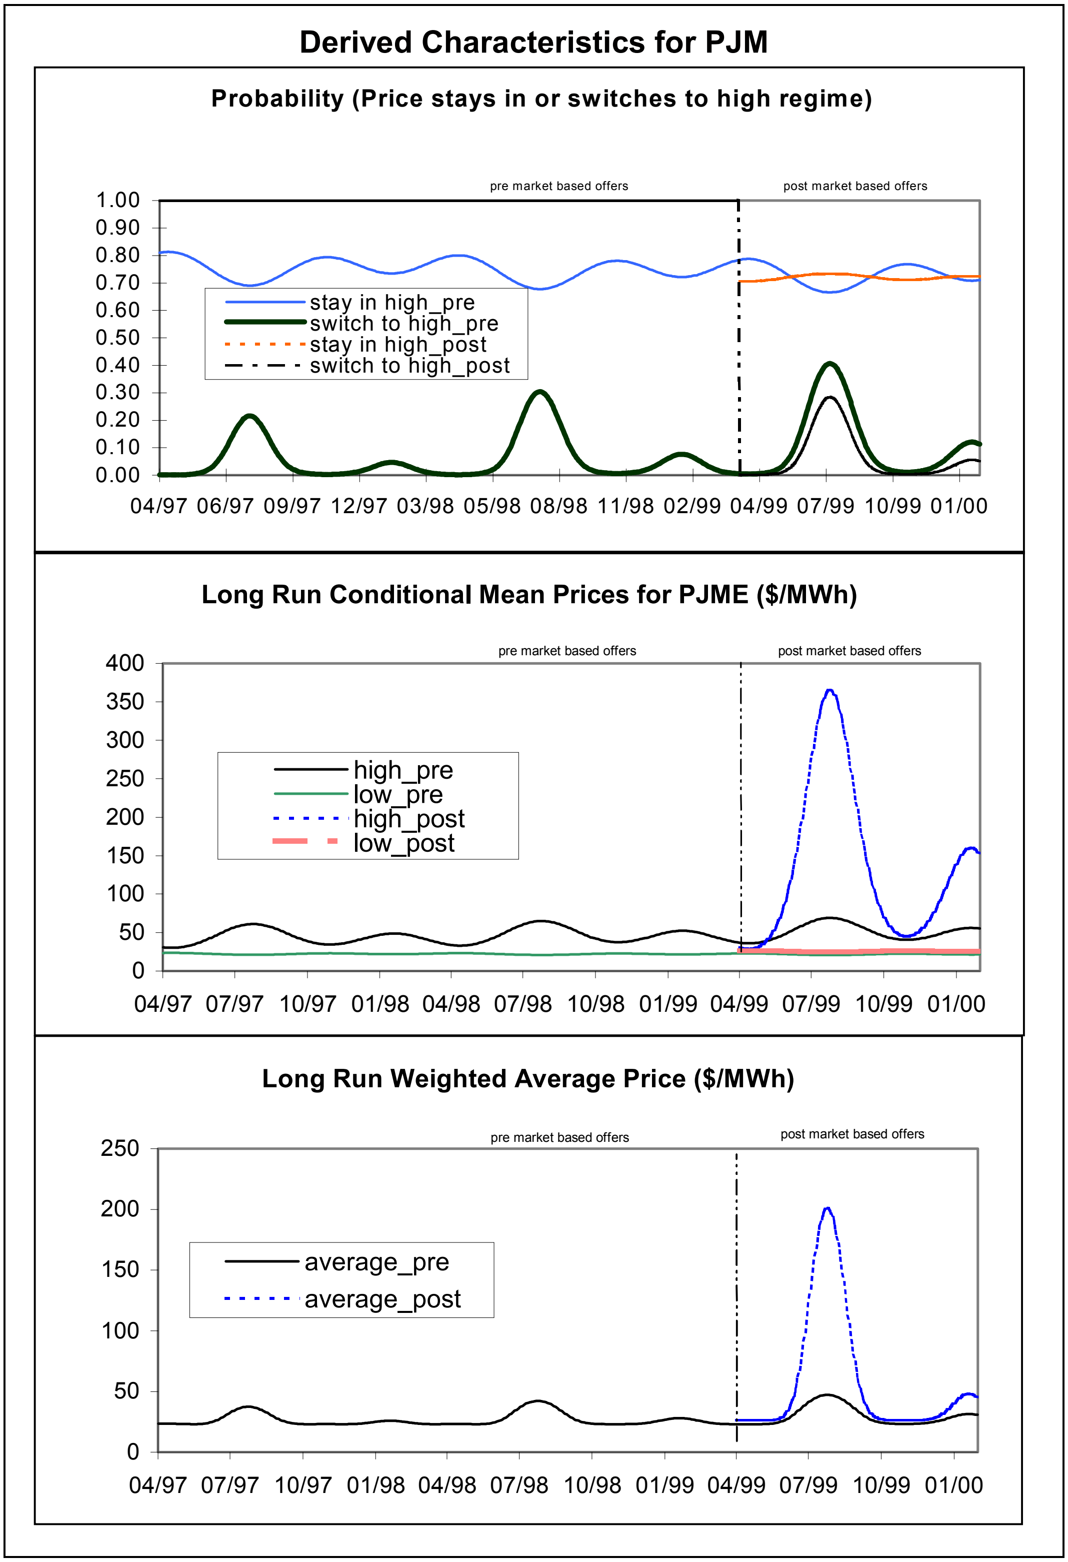

Derived Characteristics for PJM
Probability (Price stays in or switches to high regime)
pre market based offers
post market based offers
1.00
0.90
0.80
0.70
0.60
0.50
0.40
0.30
0.20
0.10
0.00
stay in high_pre
switch to high_pre
stay in high_post
switch to high_post
04/97 06/97 09/97 12/97 03/98 05/98 08/98 11/98 02/99 04/99 07/99 10/99 01/00
Long Run Conditional Mean Prices for PJME ($/MWh)
pre market based offers
post market based offers
400
350
300
250
200
150
100
50
0
high_pre
low_pre
high_post
low_post
04/97 07/97 10/97 01/98 04/98 07/98 10/98 01/99 04/99 07/99 10/99 01/00
Long Run Weighted Average Price ($/MWh)
pre market based offers
post market based offers
250
200
150
100
50
0
average_pre
average_post
04/97 07/97 10/97 01/98 04/98 07/98 10/98 01/99 04/99 07/99 10/99 01/00

Figure 5

Table 3. Derived Characteristics

	Load	PJM Pre-Market	PJM Post-Market	New England	California
High mean price $/MWh	Max	64.96	365.48	80.40	63.46
	Average	44.89	108.01	87.76	32.96
	Min	30.22	28.48	101.57	17.43
Low mean price $/MWh	Max	20.78	25.00	31.64	30.16
	Average	21.99	25.56	31.05	29.09
	Min	23.35	26.18	30.09	28.08
Weighted average price $/MWh	Max	42.24	200.91	45.81	53.98
	Average	24.12	32.10	38.52	30.77
	Min	23.40	26.19	32.15	26.25
Observed average price $/MWh		26.75	46.74	39.68	35.73
Pr[high regime]	Max	0.49	0.52	0.29	0.72
	Average	0.09	0.08	0.13	0.43
	Min	0.01	0.01	0.03	0.17

5. A Comparison of Three Markets

The regime switching model, estimated in the previous section for the PJM market, was also applied to price data in the New England and California markets. The prices for summer 1999 in the three markets are shown in Figure 6. Price spikes in PJM are the highest and the price spikes in California are the lowest among the markets. The estimated parameters and derived characteristics are summarized in Tables 2 and 3. The data used for estimation are for average on-peak weekday prices from 5/1/99 to 1/31/00 for New England and 4/1/98 to 3/31/00 for California.

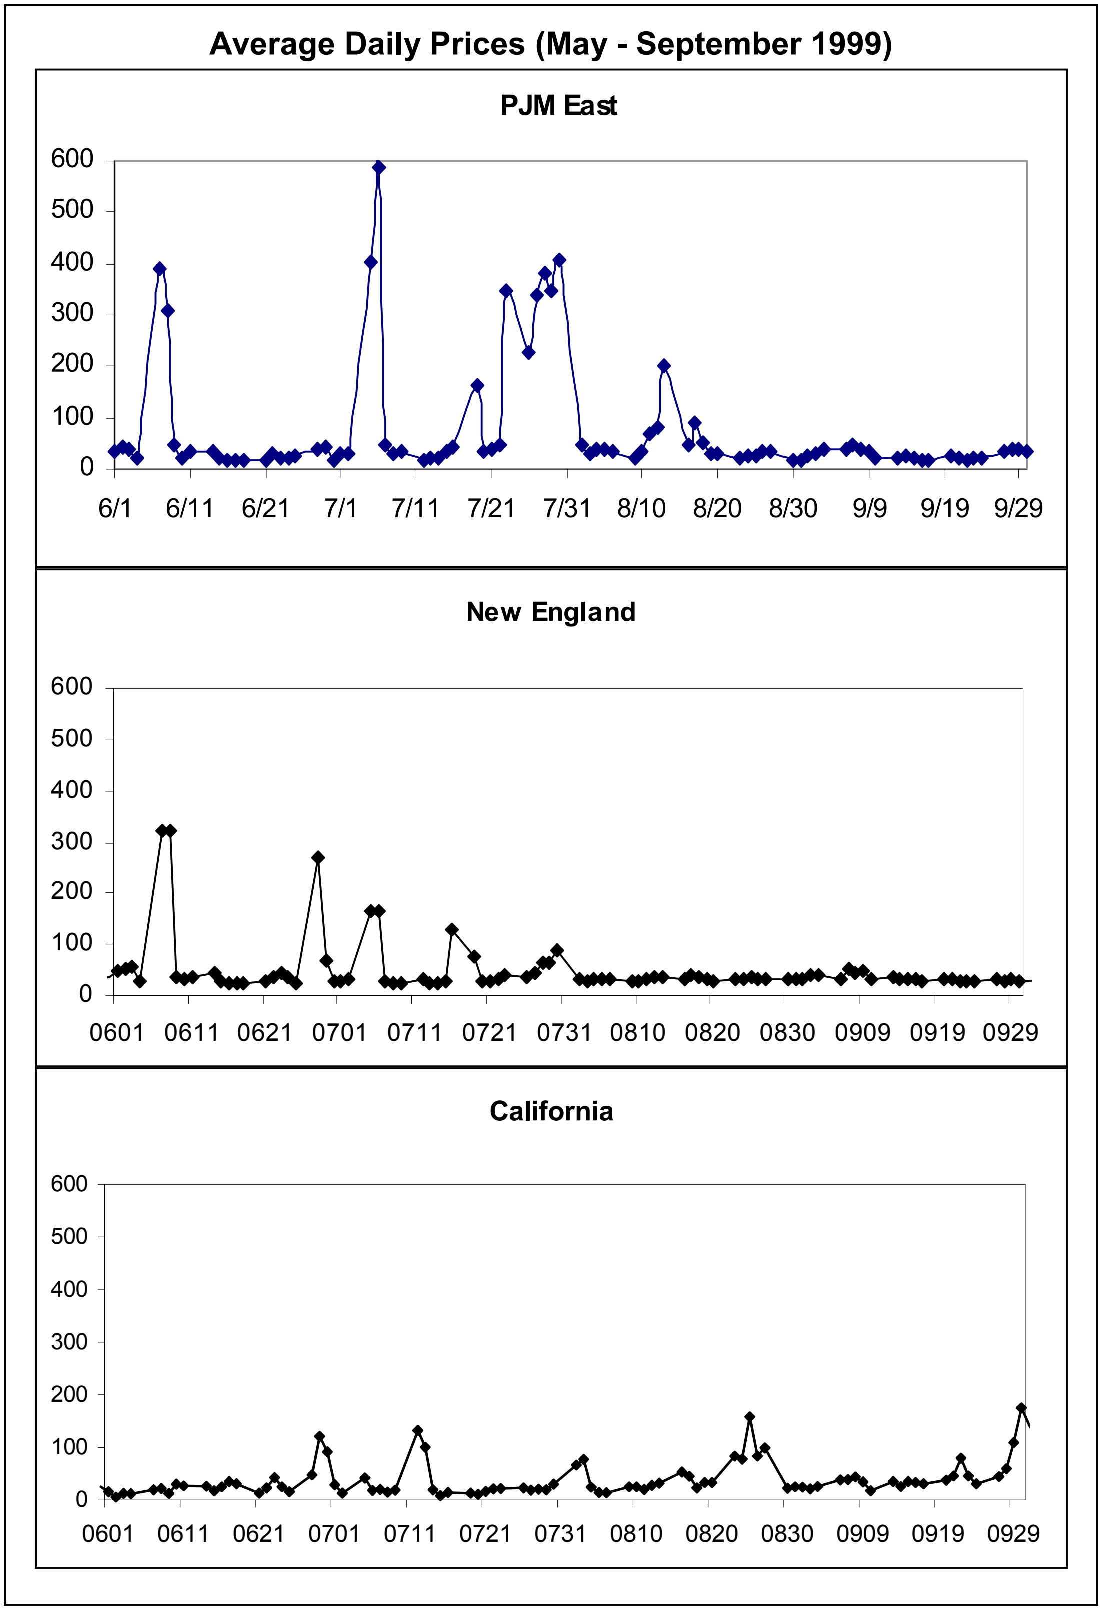
Average Daily Prices (May - September 1999)
PJM East
600
500
400
300
200
100
0
6/1
6/11
6/21
7/1
7/11
7/21
7/31
8/10
8/20
8/30
9/9
9/19
9/29
New England
600
500
400
300
200
100
0
0601
0611
0621
0701
0711
0721
0731
0810
0820
0830
0909
0919
0929
California
600
500
400
300
200
100
0
0601
0611
0621
0701
0711
0721
0731
0810
0820
0830
0909
0919
0929

Figure 6

The estimated parameters for the three markets (using the model of post-market based offers for PJM) differ in both signs and magnitudes. Nevertheless, there are some consistencies across markets. The variances in the high-price regimes are the largest ($\sigma_1 > \sigma_2$). More importantly, the probabilities of switching to the high-price regime ($1 - P_{2t}$ in (3)) have strong positive relationships to the forecasted load (the probability of staying in the high-price regime (P_{1t} in (3)) has a weak positive relationship in PJM and negative relationships in New England and California). When the logistic functions of load are plotted, the probabilities of switching to the high-price regime (see Figure 7) are remarkably similar relative to the minimum and maximum loads observed in the data sets. This suggests that the switching characteristics of the price are related to common features of the supply system and not to market power. One might have expected that market power would increase the number of price spikes, but this appears not to be the case. In fact, the probability of being in the high regime is largest for California, because price spikes are retained longer there, even though the magnitudes of the price spikes are relatively small. Note however, that the forecasted maximum loads used to estimate the regime switching parameters, are much lower than the observed maximum loads. Hence price spikes are not as predictable as the results in Figure 7 suggest.

The characteristics of the mean prices in the two regimes are very different across markets. The mean prices for the high and low regimes and the weighted average (using the unconditional probability of being in one regime or the other) are shown in Figure 8. The means for the low regime are all relatively flat (i.e. not sensitive to load). The main distinguishing feature of each market is the behavior of the mean in the high regime. The high mean increases dramatically with load in PJM, but the high means are relatively

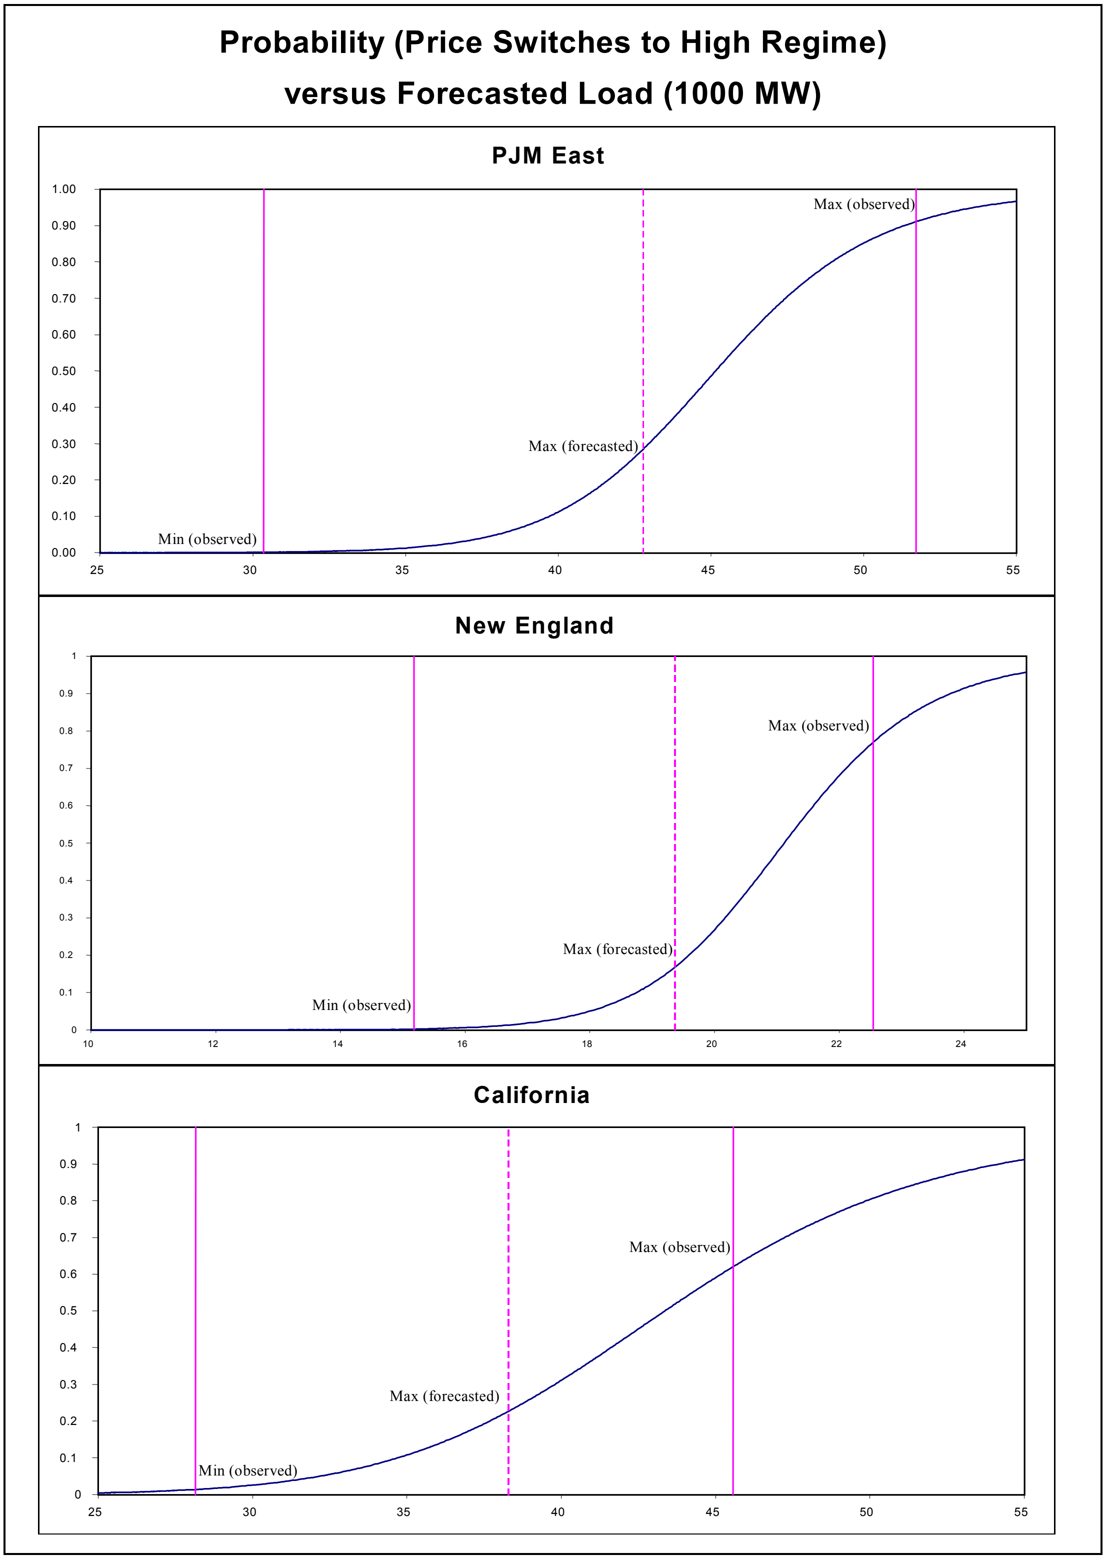
Probability (Price Switches to High Regime)
versus Forecasted Load (1000 MW)
PJM East
Min (observed)
Max (forecasted)
Max (observed)
1.00
0.90
0.80
0.70
0.60
0.50
0.40
0.30
0.20
0.10
0.00
25
30
35
40
45
50
55
New England
Min (observed)
Max (forecasted)
Max (observed)
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
10
12
14
16
18
20
22
24
California
Min (observed)
Max (forecasted)
Max (observed)
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
25
30
35
40
45
50
55

Figure 7

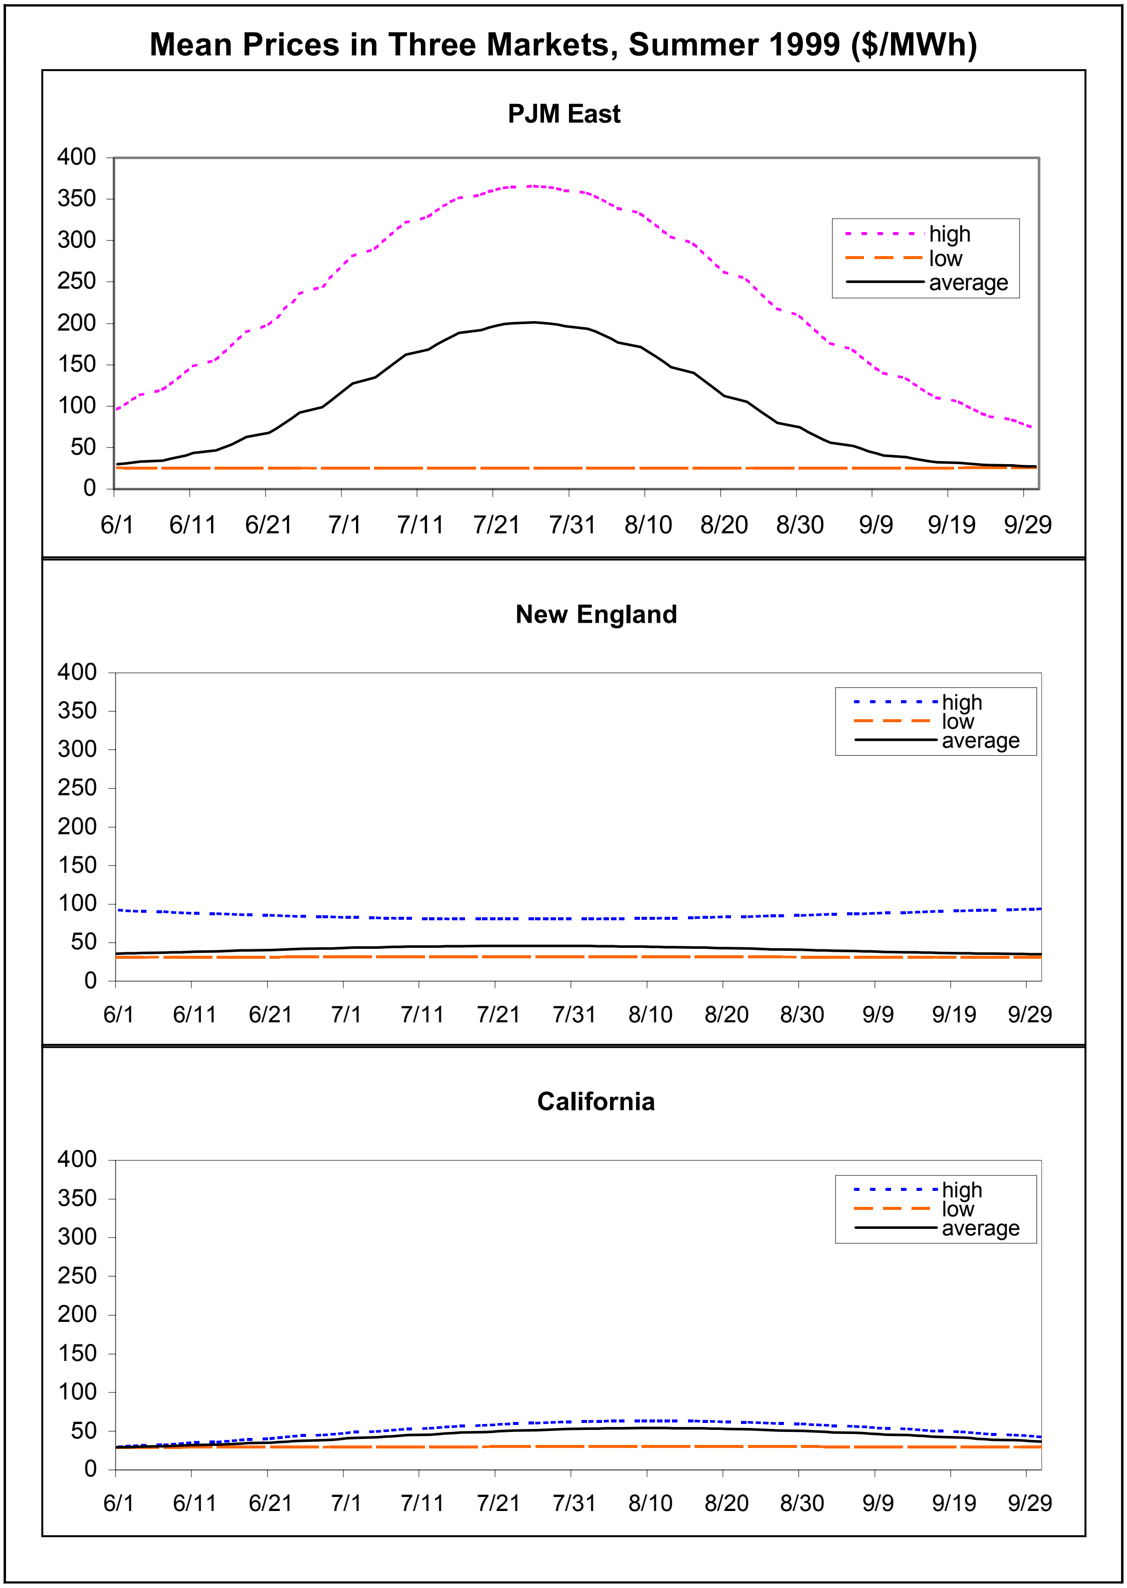
Mean Prices in Three Markets, Summer 1999 ($/MWh)
PJM East
New England
California
high
low
average
400
350
300
250
200
150
100
50
0
6/1
6/11
6/21
7/1
7/11
7/21
7/31
8/10
8/20
8/30
9/9
9/19
9/29

Figure 8

insensitive to load in the other two markets. The overall implication is that the characteristics of the PJM market in the high-price regime are unusual, and probably due to specific institutional characteristics of the market. Allowing suppliers to sell capacity in PJM and non-firm energy from the same unit in another market is one possible characteristic. Under these rules, a high offer submitted to the PJM market acts as compensation for the possibility of having capacity forcibly withdrawn from another market by the PJM ISO. The objective of the next section is to determine whether price spikes can be predicted or explained further.

6. Are Price Spikes Predictable?

The regime switching model provides a useful representation of the stochastic properties of prices that can be used to conduct financial analyses. However, the model does not provide much information about whether a price spike will occur tomorrow, for example, beyond the predicted probability of switching to the high regime (i.e. $(1-P_{2t})$ in (3)). Since this probability is positively related to the forecasted load, it increases at the beginning of a summer or winter season and then decreases. Figures 9-11 show that the probability of switching to the high regime reaches a maximum of 0.2 to 0.3 in the three markets. Since these probabilities are so low, they are not very helpful for predicting when price spikes will occur. A high expected load is a useful indicator of the possibility of a price spike, but most of the time price spikes will not occur even when the load is high. Other characteristics associated with price spikes are needed.

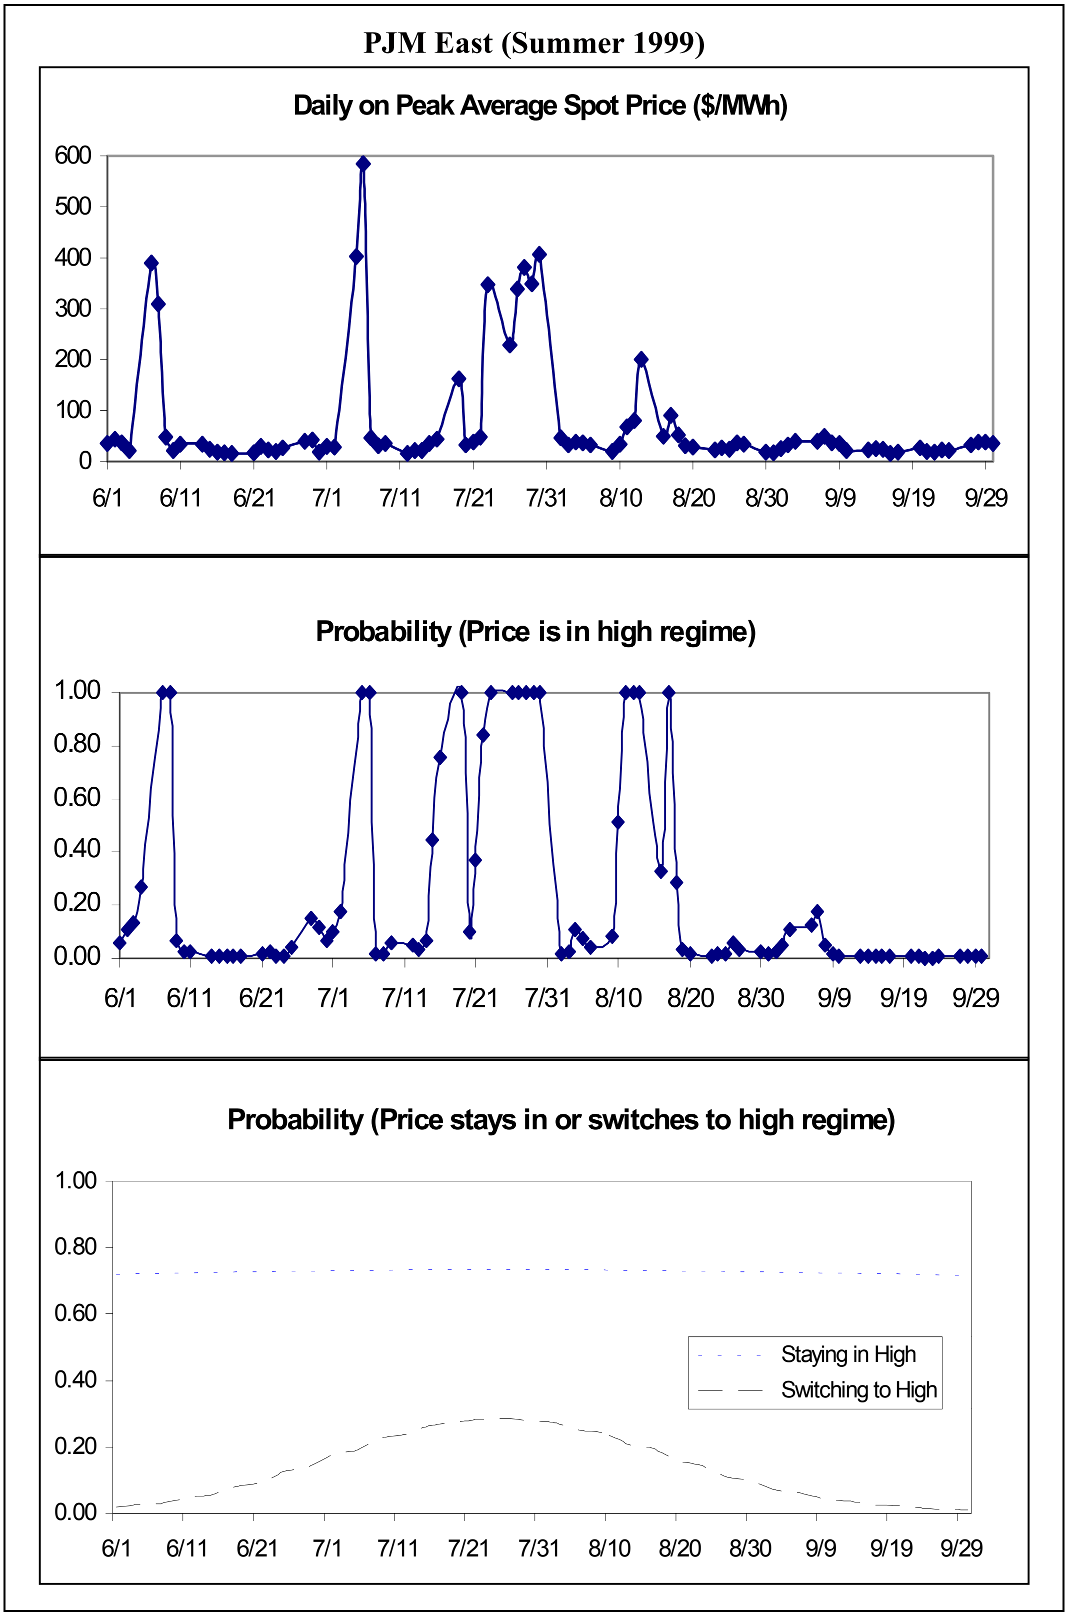
PJM East (Summer 1999)
Daily on Peak Average Spot Price ($/MWh)
600
500
400
300
200
100
0
6/1
6/11
6/21
7/1
7/11
7/21
7/31
8/10
8/20
8/30
9/9
9/19
9/29
Probability (Price is in high regime)
1.00
0.80
0.60
0.40
0.20
0.00
6/1
6/11
6/21
7/1
7/11
7/21
7/31
8/10
8/20
8/30
9/9
9/19
9/29
Probability (Price stays in or switches to high regime)
1.00
0.80
0.60
0.40
0.20
0.00
Staying in High
Switching to High
6/1
6/11
6/21
7/1
7/11
7/21
7/31
8/10
8/20
8/30
9/9
9/19
9/29

Figure 9

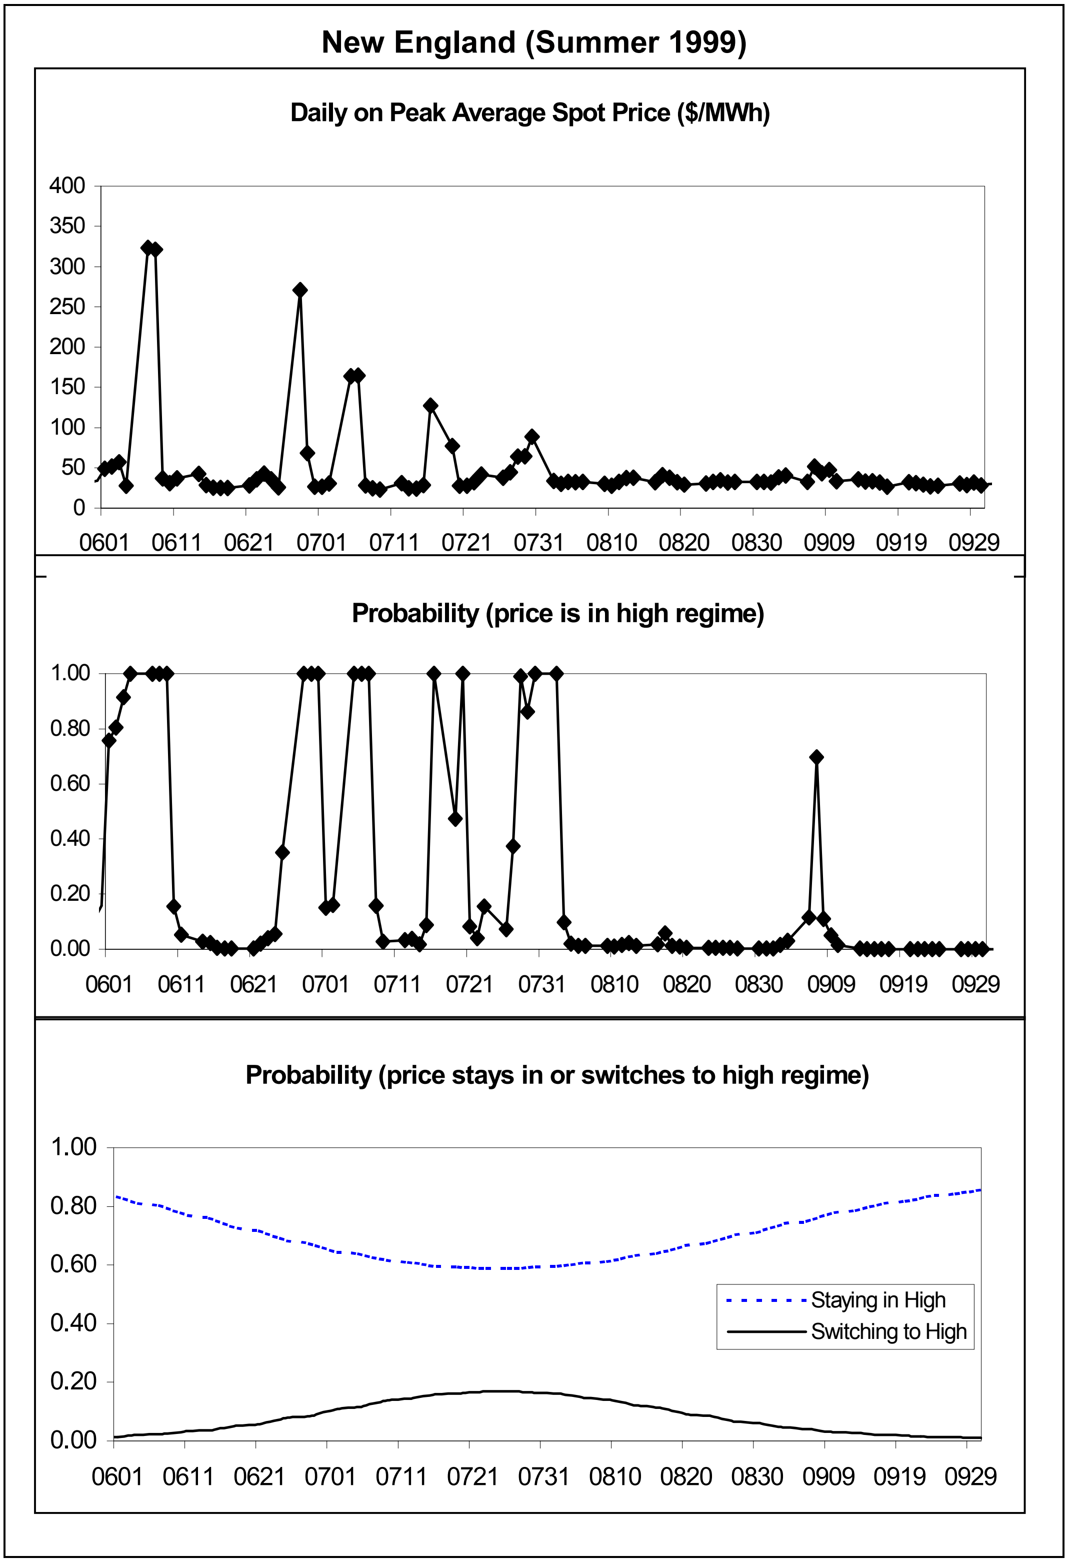

New England (Summer 1999)
Daily on Peak Average Spot Price ($/MWh)
400
350
300
250
200
150
100
50
0
0601 0611 0621 0701 0711 0721 0731 0810 0820 0830 0909 0919 0929
Probability (price is in high regime)
1.00
0.80
0.60
0.40
0.20
0.00
0601 0611 0621 0701 0711 0721 0731 0810 0820 0830 0909 0919 0929
Probability (price stays in or switches to high regime)
1.00
0.80
0.60
0.40
0.20
0.00
Staying in High
Switching to High
0601 0611 0621 0701 0711 0721 0731 0810 0820 0830 0909 0919 0929

Figure 10

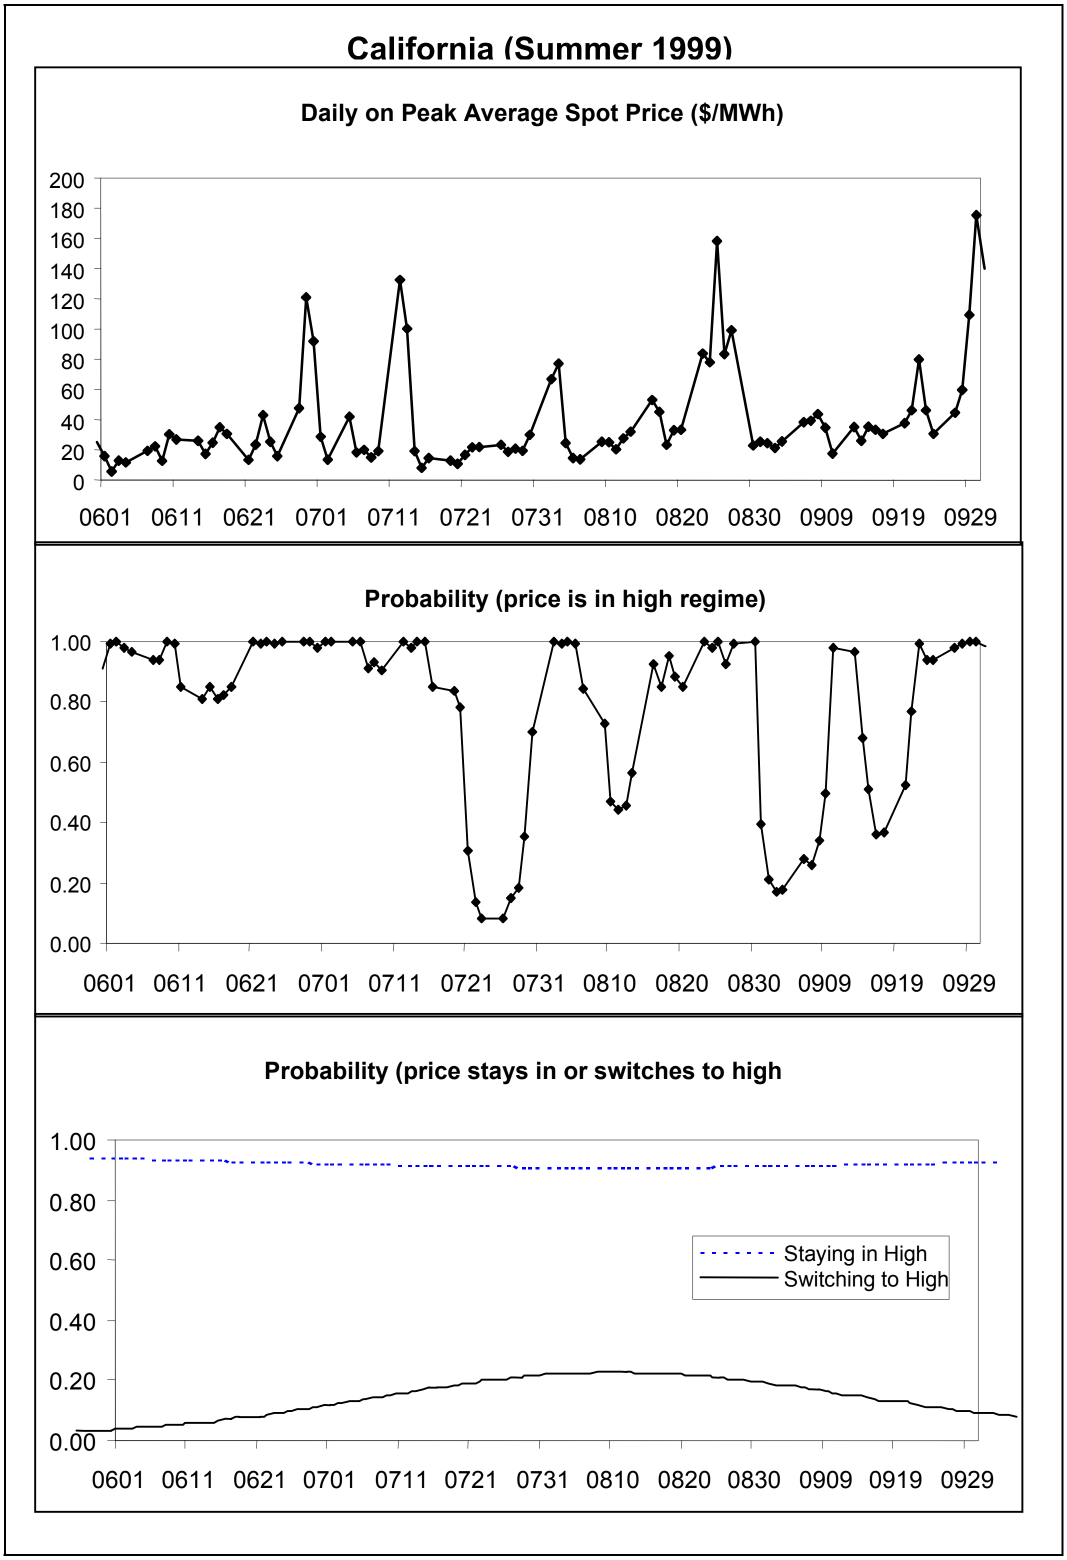

California (Summer 1999)
Daily on Peak Average Spot Price ($/MWh)
200
180
160
140
120
100
80
60
40
20
0
0601 0611 0621 0701 0711 0721 0731 0810 0820 0830 0909 0919 0929
Probability (price is in high regime)
1.00
0.80
0.60
0.40
0.20
0.00
0601 0611 0621 0701 0711 0721 0731 0810 0820 0830 0909 0919 0929
Probability (price stays in or switches to high
1.00
0.80
0.60
0.40
0.20
0.00
Staying in High
Switching to High
0601 0611 0621 0701 0711 0721 0731 0810 0820 0830 0909 0919 0929

Figure 11

Figures 9-11 show the observed average on-peak prices for June to September, 1999, and the estimated probabilities of being in the high regime, conditional on all observed prices (i.e. smoothing values). These probabilities represent ex-post predictions of how often the markets were in a high price regime. It is interesting to note how different California is compared to the other two markets. Since the probability of staying in the high regime is relatively high in California, the high regime is persistent, and jumps to the low regime are the unusual feature. For PJM and New England, the relationships between the probabilities and the observed price spikes are much more obvious. An additional feature of California is that the mean price for the high regime is slightly lower than the mean price for the low regime when the forecasted load is very low. As a result, the high probabilities of being in the high regime during periods of low load (i.e. the spring and fall) are not associated with price spikes. Since the price behavior in California is so different, most of the following analysis of price spikes focuses on prices in PJM and New England.

Given the characteristics of the offer curves submitted to the PJM auction (see Figure 4), it is sensible to consider whether high prices occur when actual load is higher than expected. In other words, do large positive forecasting errors (observed > predicted) result in the market price being set by the steeply sloped part of the offer curve? In this situation, better forecasts might help to reduce the number of price spikes. An alternative argument is that the price spikes occur when load increases above an already high level. In this situation, if load shifts to the right more than the offer curve shifts, price spikes could occur.

As a first step, the hourly prices for the summer, 1999 are plotted against the corresponding loads for the three markets in Figure 12. It is clear that high prices occur more often when the load is high, but the relationships between high prices and high loads are not particularly strong. Once again, the relationship for California is not as clear as it is for the other two markets. The results in Figure 12 are used to provide a rough division between low loads and high loads for the following analysis of prices in PJM and New England. The cutoff values are 45 and 20 thousand MW for PJM and New England, respectively.

Since daily forecasts of load in the PJM and New England markets are not saved in public archives, it is not possible to evaluate the performance of the actual forecasts used by the ISOs. Consequently, a time-series model was fitted to hourly load data to capture all seasonal, weekly and daily cycles. Sine and Cosine variables were used to capture seasonal and daily cycles, and arctangents were used for weekends and special holidays. The weather related departures of load from normal levels were captured by ARMA processes. A full discussion of these models will be presented in another publication, but it should be noted that both models fitted the data very well with $R^2 > 0.99$.

The good fit of the time-series models of load illustrates the ability of the ARMA processes to model deviations from normal weather patterns. In other words, it is relatively easy to predict load one hour ahead, conditionally on the current load. To approximate the forecasts of load made by an ISO, however, the models are used to forecast 30 hours ahead. This is equivalent to assuming that the ISO forecasts the pattern of load for the following day at 6 p.m. each day. The observed maximum load for each

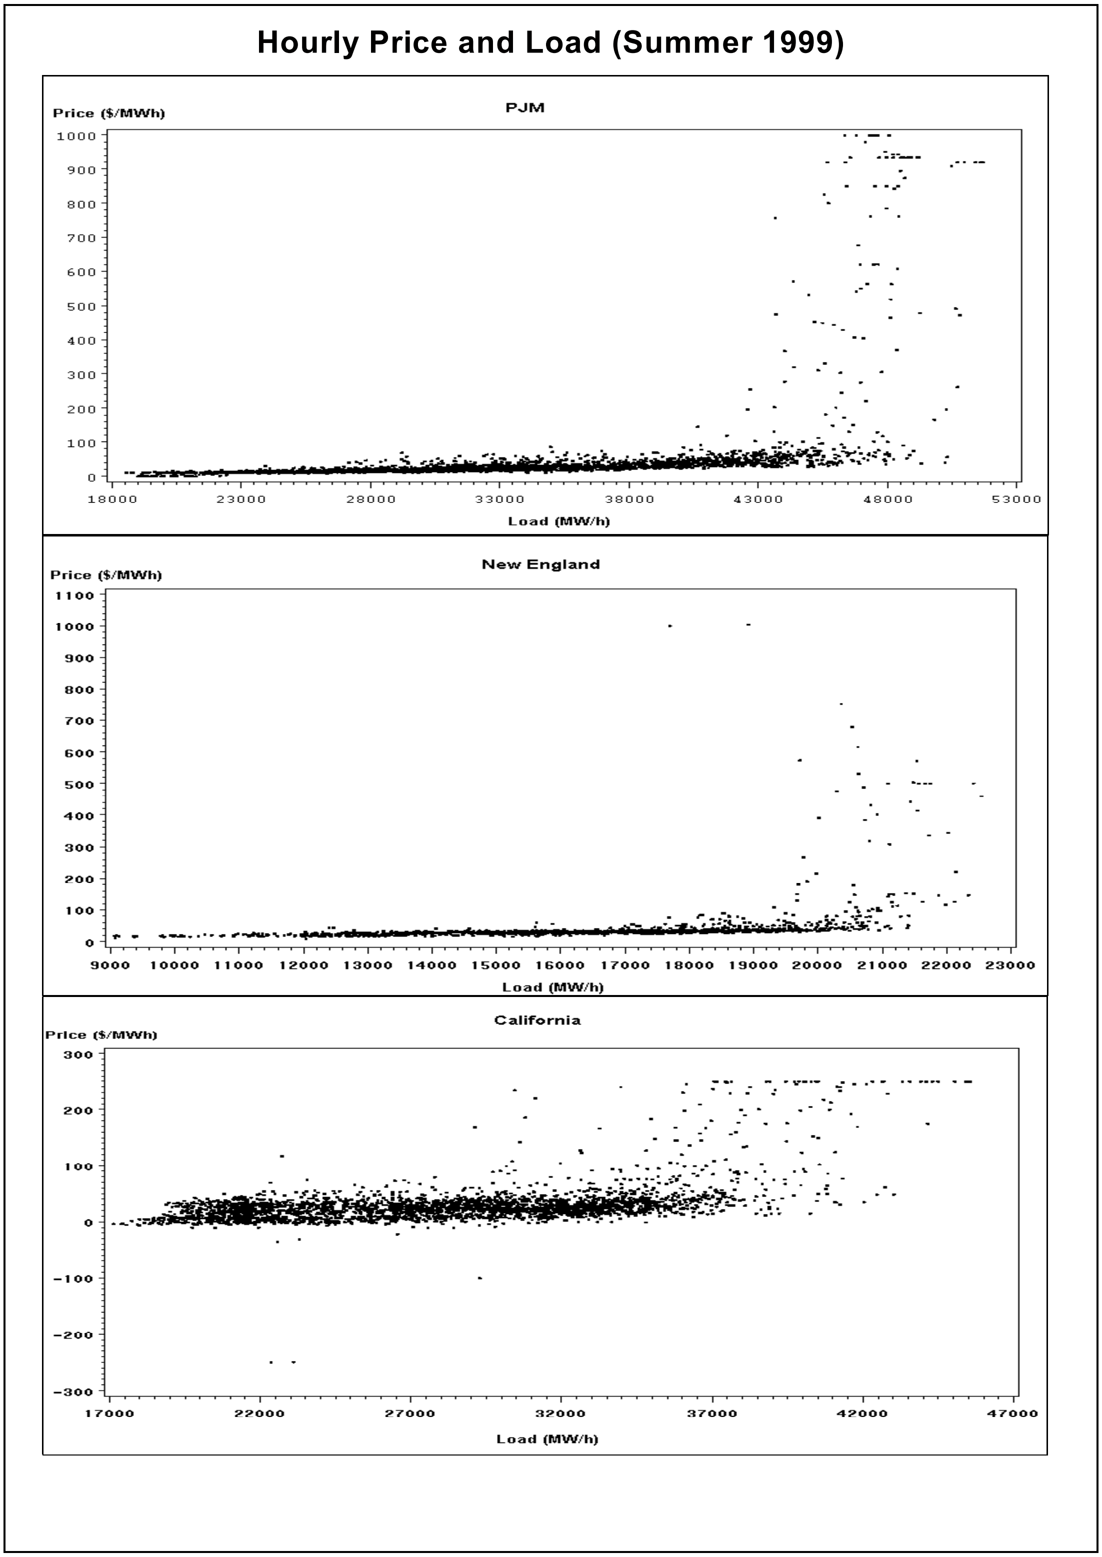
Hourly Price and Load (Summer 1999)
PJM
Price ($/MWh)
Load (MW/h)
New England
Price ($/MWh)
Load (MW/h)
California
Price ($/MWh)
Load (MW/h)

Figure 12

day is compared to the maximum forecasted load (the time periods need not coincide) to compute the forecasting error. Typical errors are about 5% compared to the hour-ahead errors of about 1%.

Figures 13 and 14 show 1) the actual peak loads and forecasting errors (the shaded segments correspond to Probability(in a high price regime) > 0.8), 2) prices plotted against forecast errors and 3) prices plotted against the changes in actual peak load. (The diamonds correspond to low load days and the circles to high load days.) Solid circles correspond to being in a high price regime, and S is the first day of a high price regime (NS is the first day of a high price regime when the load is low). Inspection of Figures 13 and 14 show that there are no strong positive relationships between high prices and forecasting errors or changes of load. For PJM, the first days of the price spikes are evenly split between positive and negative forecasting errors (four with positive errors and three with negative errors). In contrast, all but one of the first days are associated with positive changes of load. For New England, most first days are associated with both positive forecasting errors and with positive changes of load, and the latter shows a relatively strong positive relationship.

The overall conclusion is that there is no convincing evidence that price spikes are caused by the load being unexpectedly high (i.e. when forecasting errors are positive). If load is high, price spikes are more likely to occur when load increases further. In this sense, price spikes are predictable, and therefore it is reasonable to expect that the vulnerability of a market to price spikes can be reduced by modifying the rules that govern the market. However, the solutions are unlikely to be simple quick fixes due to the complicated designs chosen for electricity auctions in the USA, compared, for

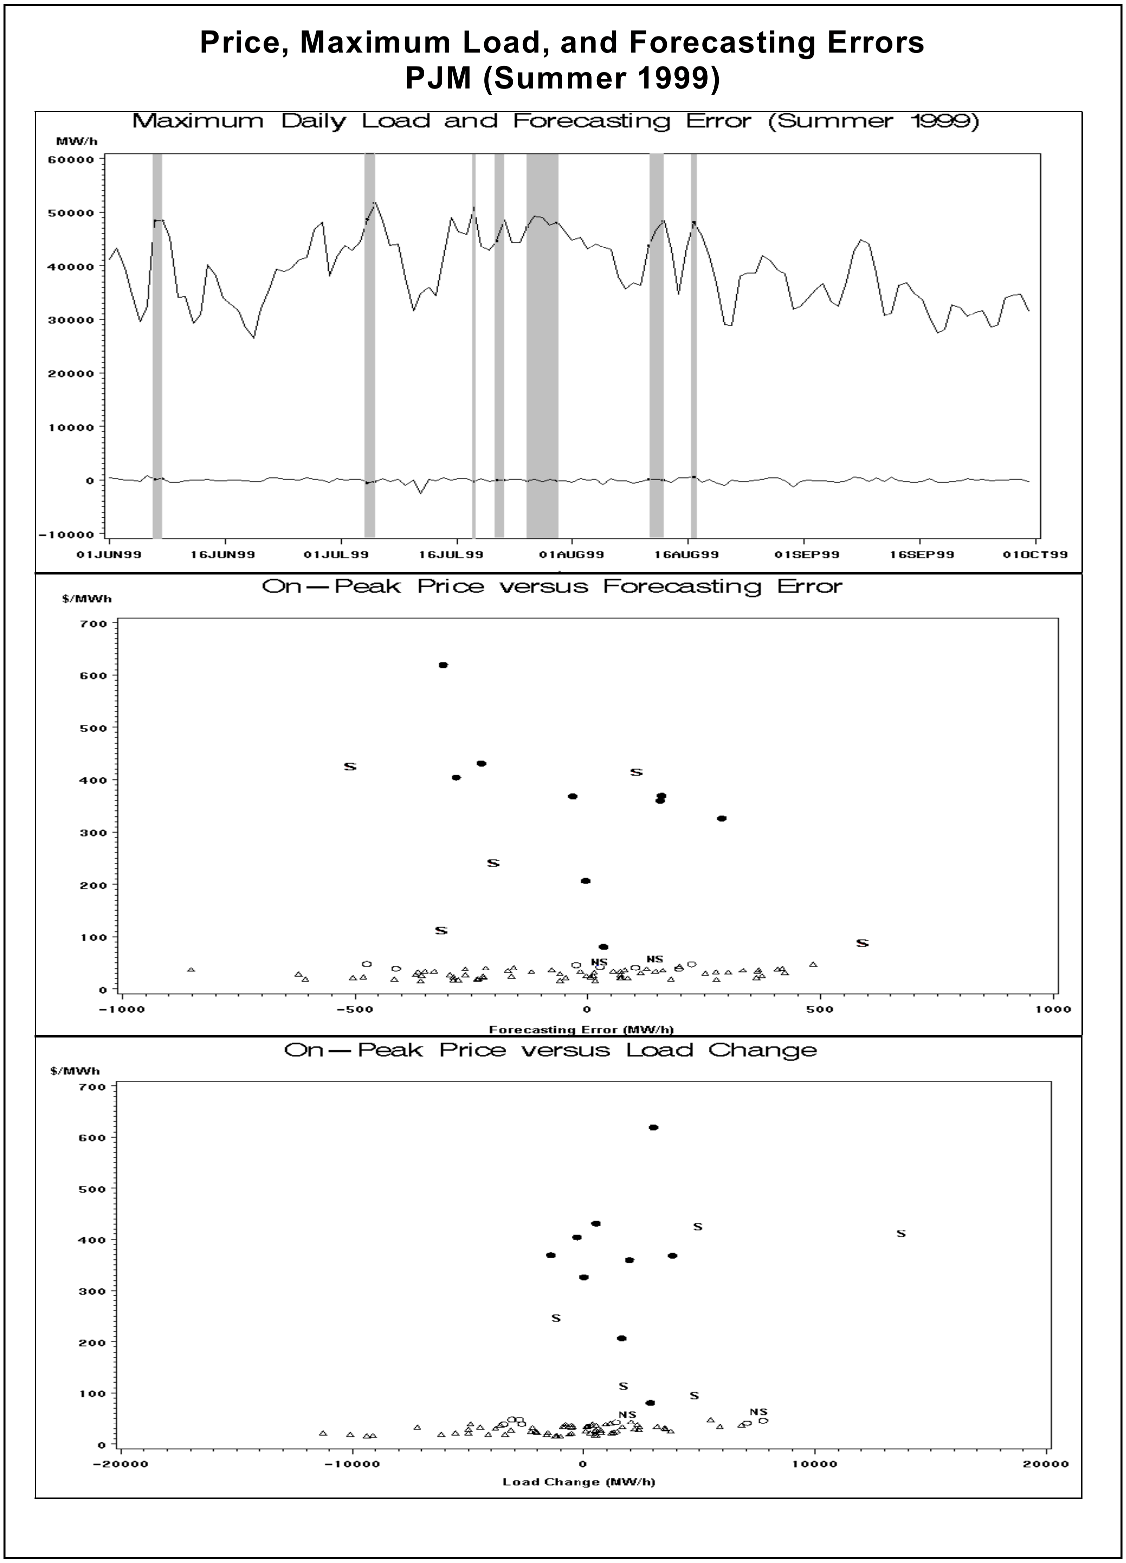
Price, Maximum Load, and Forecasting Errors
PJM (Summer 1999)
Maximum Daily Load and Forecasting Error (Summer 1999)
MW/h
60000
50000
40000
30000
20000
10000
0
-10000
01JUN99
16JUN99
01JUL99
16JUL99
01AUG99
16AUG99
01SEP99
16SEP99
01OCT99
On—Peak Price versus Forecasting Error
$/MWh
700
600
500
400
300
200
100
0
S
NS
-1000
-500
0
500
1000
Forecasting Error (MW/h)
On—Peak Price versus Load Change
$/MWh
700
600
500
400
300
200
100
0
S
NS
-20000
-10000
0
10000
20000
Load Change (MW/h)

Figure 13

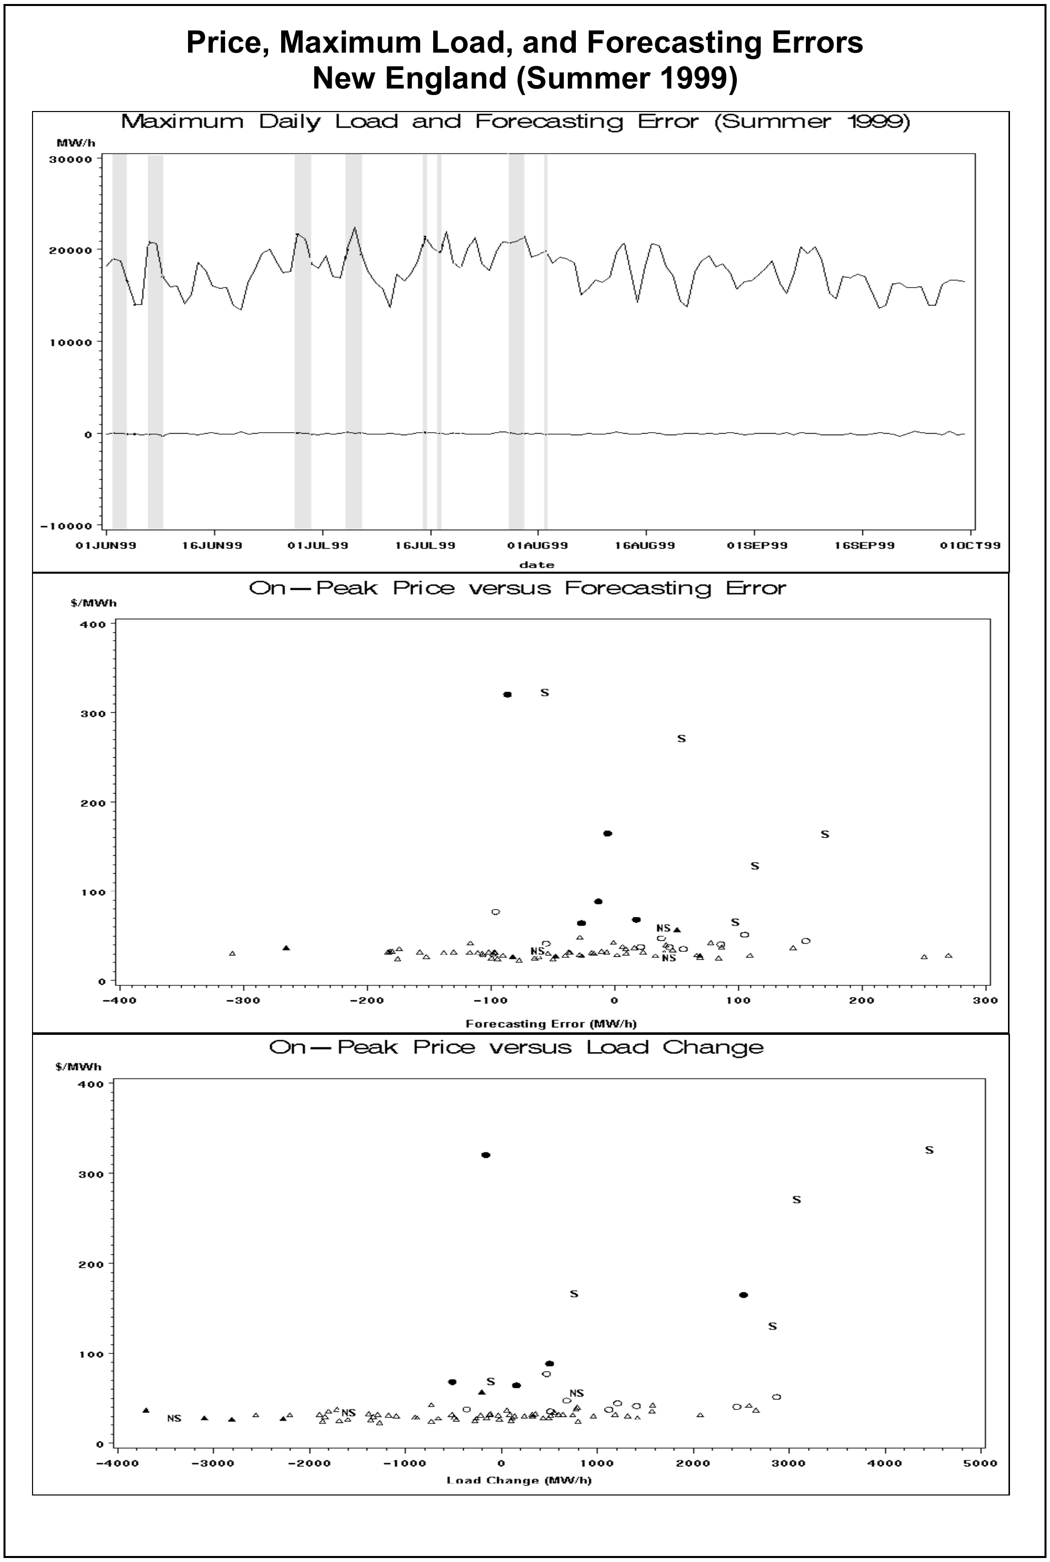
Price, Maximum Load, and Forecasting Errors
New England (Summer 1999)
Maximum Daily Load and Forecasting Error (Summer 1999)
MW/h
30000
20000
10000
0
-10000
01JUN99
16JUN99
01JUL99
16JUL99
01AUG99
16AUG99
01SEP99
16SEP99
01OCT99
date
On-Peak Price versus Forecasting Error
$/MWh
400
300
200
100
0
-400
-300
-200
-100
0
100
200
300
Forecasting Error (MW/h)
On-Peak Price versus Load Change
$/MWh
400
300
200
100
0
-4000
-3000
-2000
-1000
0
1000
2000
3000
4000
5000
Load Change (MW/h)

Figure 14

example, to Australia. It is safe to say that there is a lot of room for improvement over the current solution of suspending the operation of a market when the regulators or the incumbent utilities do not like the prices.

7. Conclusions

The main objective of this paper is to provide a better understanding of why spot markets for electricity are so volatile. The paper shows that a supplier with some market power may be completely indifferent about having marginal units dispatched (Section 2). Under these circumstances, there is no penalty to the supplier of setting offers for marginal units at very high levels. As a result, relatively minor institutional characteristics of a market may have large effects on suppliers' offers. For example, the reason that price spikes in PJM are so high may be because suppliers with contracts in the capacity market are creating hedges against being recalled from contracts to sell energy in other markets.

A regime switching model of price behavior is specified to allow for time-varying parameters. The mean prices and the transition probabilities are all functions of forecasted load in this model (Section 3). Applying this model to prices in PJM, using data pre and post the introduction of market-based offers on 4/1/99, shows that structural changes did occur (Section 4). In particular, the mean in the high regime is much higher and more sensitive to load after market-based offers were allowed. The probability of switching to the high regime, however, was not affected very much. Hence, it appears that market power enabled suppliers to set higher price spikes, but not to get price spikes to occur more often.

Comparing the price behavior in PJM (post market-based offer) with the New England and California markets shows that the sensitivity of the transition probabilities to load are similar in the three markets (Section 5). The high mean in the high regime is the main feature that distinguishes PJM from the other two markets. The low price spikes in California are not surprising given the importance of bilateral trading in this market. Bilateral trades have properties similar to discriminatory auctions, in which suppliers are paid actual offers rather than a uniform price set by the last accepted offer. Mount (1999) has argued that a discriminatory auction would be an effective way to reduce price volatility associated with market power. A key question for research is why are price spikes in PJM so much higher than they are in New England, even though the auction designs are quite similar?

The final analysis attempts to identify whether price spikes in the PJM and New England markets are associated with errors in forecasting load or with increases in load from one day to the next (Section 6). Given the kinked shape of the offer curves submitted into the PJM market, under-forecasting or unexpected increases of load could result in prices being set on the steeply sloped section of an offer curve. However, there appear to be no systematic relationships between high prices and forecasting errors. Positive changes of load when the current load is already high are related to price spikes. Hence, it would be an interesting question for further research to determine whether the regime switching model for prices can be modified to use day-ahead forecasts of load rather than seasonal forecasts. Our initial assumption was that the parameters of the model should not change erratically from day to day (initial efforts to estimate such a model were not successful). It is possible, however, that the probability of switching to a high regime is sensitive to current load, and in this situation, the probabilities would be much higher and more informative at certain times that the values shown in Figures 9-11.

While modifying the regime switching model would be a useful step forward, the most important information for an ISO is contained in the actual offer curves submitted to a market because they determine the effective reserve margin for a given level of expected load. In this respect, announcing a forecast of a future spot price, and allowing suppliers to submit new offers, is likely to be a more effective way to deal with price spikes than setting prices in a day-ahead market. If market power becomes a chronic problem, it may be preferable to abandon uniform price auctions in favor of discriminatinatory auctions. However, getting more players into the market is still the best way to deal with market power.

References

Baker, M., S. Mayfield and J. Parsons (1998). "Alternative Models of Uncertain Commodity Prices for Use with Modern Asset Pricing Methods." *The Energy Journal* 19 (1): 115-148.

Ball, C. and W. Torous (1983). "A Simplified Jump Process for Common Stock Returns." *Journal of Financial and Quantitative Analysis* 18 (1): 53-65.

Barz, G. and B.Johnson (1998). "Modeling the Price of Commodities that are Costly to Store: the Case of Electricity." Presented at the Chicago Risk Management Conference. Chicago, IL. May.

Cecchetti, S., L. Pok-Sang, and M. Nelson (1990). "Mean Reversion in Equilibrium Asset Prices." *The American Economic Review* 80 (3): 398-418.

Deng, S., B. Johnson and A. Sogomonian (1998). "Exotic Electricity Options and the Valuation of Electricity Generation and Transmission." Presented at the Chicago Risk Management Conference, Chicago, May.

Deng, S. (1998). "Stochastic Models of Energy Commodity Prices and Their Applications: Mean Reversion with Jumps and Spikes." PSERC Working Paper 98-28. Available at <http://www.pserc.wisc.edu/psercbin/test/get/publicatio/>.

Diebold, F.X., J.H. Lee, and G.D. Weinbach (1994). Regime-Switching with Time-Varying Transition Probabilities, in: C. Hargreaves (ed.), *Nonstationary Time Series Analysis and Cointegration*, (Oxford University Press) 283-302.

Dixit, A. and R. Pindyck (1994). *Investment Under Uncertainty*. Princeton, NJ: Princeton University Press.

Dorris, G. and R. Ethier (1998). "Modeling the Electricity Spot Price." Presented at The Energy Modeler's Forum. Houston, Texas. April.

Duffie, D. and S. Gray (1995). "Volatility in Energy Prices." in R. Jameson (ed.), *Managing Energy Price Rosk.* Risk Publications, London.

Engel, C. and J. Hamilton (1990). "Long Swings in the Dollar: Are they in the Data and Do Markets Know It?" *The American Economic Review* 80 (4): 689-713.

Ethier, R. (1997). "Competitive Electricity Markets and Long Term Investment." Departmental Seminar, Department of Agricultural, Resource, and Managerial Economics, Cornell University. Ithaca, NY. December 5.

Ethier, R. (1999). “Chapter Two: A Real Options Model Suitable for Competitive Electricity Markets.” Unpublished Doctoral Dissertation titled “Competitive Electricity Markets, Prices and Generator Entry and Exit.” Dept. of Agricultural, Resource, and Managerial Economics, Cornell University, May.

Ethier, R. and T. Mount (1998). “Winners and Losers in a Competitive Electricity Industry: An Empirical Analysis.” *The Energy Journal*, Special Issue on Distributed Resources, January: 161-186.

Either, R. and T. Mount (1999). "Estimating the Volatility of Spot Prices in Restructured Electricity Markets" Working Paper, Department of Agricultural, Resource, and Managerial Economics, Cornell University.

Gray, S.F.: Modeling the Conditional Distribution of Interest Rates as a Regime-Switching Process. *Journal of Financial Economics* 42, 27-62(1996).

Hamilton, J. (1989). “A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle.” *Econometrica* 57 (2): 357-384.

Hamilton, J. (1994). *Times Series Analysis*. Princeton, NJ: Princeton University Press.

Hamilton, J. (1996). “Specification Testing in Markov-switching Times Series Models.” *Journal of Econometrics* 70: 127-157.

Hamilton, J. and R. Susmel (1994), "Autoregressive Conditional Heteroskedasticity and Changes in Regime." *Journal of Econometrics* 64: 307-333.

Hansen, B.E. (1992). “The Likelihood Ratio Test Under Nonstandard Conditions: Testing the Markov Switching Model of GNP.” *Journal of Applied Econometrics* 7: s61-s82.

Laughton, D. (1998). “The Potential for Use of Modern Asset Pricing Methods for Upstream Petroleum Project Evaluation: Concluding Remarks.” *The Energy Journal* 19 (1): 149-153.

Lo, A. and J. Wang (1995). “Implementing Option Pricing Models When Asset Returns Are Predictable.” *The Journal of Finance* L (1): 87-129.

Mount, T. (1999). “Market Power and Price Volatility in Restructured Markets for Electricity.” Presented as a selected paper at the Hawaii International Conference on System Science. HI. January.

Mount, T. (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, HI. January.

Newbery, D. M. (1995). "Power Markets and Market Power," *The Energy Journal*, 16 (3): 39-66.

Pilipovic, D. (1997). *Energy Risk: Valuing and Managing Energy Derivatives*. New York, NY: McGraw-Hill.

Rudkevich, A., M. Duckworth, and R. Rosen (1998). "Modeling Electricity Pricing in a Deregulated Generation Industry: The Potential for Oligopoly Pricing in a PoolCo." *The Energy Journal* 19 (3): 19-48.

Schwartz, E. (1997). "The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging." *The Journal of Finance* LII (3): 923-973.

Tseng, C-L. and G. Barz (1998). "Short-Term Generation Asset Valuation." Working paper, Department of Civil Engineering, University of Maryland. College Park, MD.

Violette, D., M. King and R. Ethier (1998). "Electricity Price Forecasts and the Forward Price Curve for Electricity." Presented at the Electric Power Research Institute's 5th Biennial Conference on Innovative Approaches to Electricity Pricing: Pricing Energy in a Competitive Market. Washington, D.C. (June).

Wolak, F. (1996). "Market Design and Price Behavior in Restructured Electricity Markets: An International Comparison." Department of Economics, Stanford University. Palo Alto, CA.

Wolak, F. A. and R. H. Patrick (1997). "The Impact of Market Rules and Market Structure on the Price Determination Process in the England and Wales Electricity Market." February. Selected Paper presented at the POWER Conference, University of California, Berkeley. March.