An Empirical Examination of Deregulated Electricity Prices

Christopher R. Knittel and Michael R. Roberts*

Draft: October 30, 2001

Abstract

In this paper, we present an empirical analysis of deregulated electricity prices. We begin by examining the distributional and temporal properties of the price process in a non-parametric framework. This analysis is followed by comparing the forecasting ability of several different statistical models. The findings reveal several characteristics unique to electricity prices, including deterministic components of the series at different frequencies and a high degree of persistence in the price level. An "inverse leverage effect" is also found, where positive shocks to the price series result in larger increases in volatility than negative shocks. Results consistent with other asset prices, such as time-varying volatility are also uncovered. We find that existing financial models of asset prices fail to forecast the extremely erratic nature of electricity prices. Non-Markovian specifications, in conjunction with exogenous information (e.g. weather), are a necessary starting point for practical applications, such as security pricing.

*This paper has benefited from conversations with Ravi Bansal, John Graham, Pete Kyle and Tom Rothenberg, and comments from Mike Lemmon. Knittel is from the School of Management at Boston University. Roberts is from the Fuqua School of Business at Duke University. Correspondence may be sent to Christopher Knittel at knittel@bu.edu.

1 Introduction

Currently, 22 states and the District of Columbia have deregulated their electricity markets; another 19 are in the process of deregulating their markets. In addition, deregulation has spread internationally to countries such as Great Britain, Australia, Spain and Norway.

One of the many ramifications resulting from this change has been an increase in the importance of modeling and forecasting electricity prices. Under regulation, prices are set by state public utility commissions (PUC's) in order to curb market power and ensure the solvency of the firm. Price variation is minimal and under the strict control of regulators, who determine prices largely on the basis of average costs. This setting focuses the utility industry's attention on demand forecasting, as prices are held constant between PUC hearings. Market entry is barred and investment in new generation by incumbent firms is largely based on demand forecasts. In addition, there is little need for hedging electricity price risk because of the deterministic nature of prices.

Deregulation removes price controls and openly encourages market entry. Consequently, price variation has skyrocketed and the number of energy-based financial products has exploded, as purchasers of electricity hedge price risk and investors capitalize on new opportunities. Electricity and weather derivative markets have been established at various locations in the United States, as well as in Europe. The recent experience in California has illuminated the importance of such markets and implementing hedging strategies.[1] However, the transaction volume in these markets has been less than anticipated due, in part, to the difficulty in valuing electricity contracts.[2] A precursor for reliable valuation of electricity contracts is an accurate description of the underlying price

[1]Beginning in the summer of 2000, California electricity prices fluctuated wildly and capacity shortages existed during a number of hours. In an effort to curb demand, rolling blackouts were instituted. Because of regulatory constraints, the three investor owned utilities (IOUs), Pacific Gas & Electric, Southern California Edison and San Diego Gas & and Electric, felt the brunt of these experiences. When the market was initially designed, two rules were set in place that left the IOUs unable to hedge against volatility. First, the IOUs were not allowed to sign long term contracts for wholesale electricity. Secondly, the IOUs retail rates were largely fixed. Therefore, as prices began to fluctuate and rise, the IOUs were both unsheltered from the fluctuations and unable to pass any wholesale price increases onto consumers. As a consequence, California's largest utility, Pacific Gas & Electric, recently filed for bankruptcy and the other IOUs have compiled huge debts.

[2]We recognize that utilities in California were precluded from entering into forward contracts. However, this ban was recently repealed and, almost immediately, long term forward contracts were instituted at prices that many believe to be excessively high.

process. The analysis of this paper provides this description and suggests that valuation difficulties stem from the unique stochastic properties of electricity prices.

In this paper, we investigate the behavior of California's deregulated electricity prices. We begin by analyzing electricity prices and comparing their salient features with those of equities and other commodities. Despite some distributional similarities, electricity prices are dramatically different from equity prices, and even other commodity prices. Specifically, electricity prices display the following distinct characteristics:

1. a high degree of persistence in both the price level and squared prices,

2. pronounced intraday, day of week, and seasonal cycles, and

3. censoring from above.

The seasonal nature of the series is clearly reflected in the first two characteristics listed above. Interestingly, prices are censored from above but not from below. Price caps were instituted during several periods to avoid further financial strain on utilities. Also in the sample are negative prices, due to a combination of excess supply, the inability to store electricity, and high start-up costs faced by generators.

Electricity prices also share several characteristics in common with equity prices, including a high kurtosis and persistence in the square of prices. If not for the price caps, estimated kurtosis is likely to be far greater in electricity prices than equity prices. Electricity prices exhibit dramatic swings of over 10 standard deviations within only a few hours. Aside from the seasonal component, persistence in the square of prices is similar to that in squared equity prices.

After the preliminary data analysis, we analyze and compare several statistical models often used to represent asset price processes. These models serve to further contrast electricity prices with equity prices. They also illustrate the difficulty practitioners and academics face when modeling electricity prices. Our results demonstrate that common models of asset prices offer a poor representation of the electricity price process, in part, because of their failure to capture many of the characteristics discussed above. We also find that the forecasting performance of these models is unacceptable for any practical purpose, such as the pricing of energy-based financial products.

After altering the models to account for seasonal effects and persistence, little is changed. Forecasting performance, though markedly improved, is still poor. However, another characteristic unique to electricity prices is uncovered, which we term the "inverse leverage effect." This terminology is derived from Black's (1976) "leverage effect" and describes the asymmetric response of volatility to positive and negative shocks.[3] With electricity prices, volatility tends to rise more so with positive shocks than negative shocks; a result of convex marginal costs. Also of interest is the negligible effect of incorporating local temperatures into the specification, as it adds little to forecasting ability.

From the results of this study, several conclusions are clear. First, conditionally Gaussian models have little chance of accurately representing the data generating process because of their inability to capture very large changes in prices. Heavy-tailed random components, such as Student's-t and Levy processes, are likely candidates for future specifications. Second, univariate Markovian specifications are unable to accommodate the persistence found in prices. Depending on the frequency of one's data, higher order lags are needed to capture the autocorrelation in the price series. Finally, additional exogenous information, such as marginal costs, is likely to be another important addition to future modeling efforts.

While there has been a significant amount of research on commodity prices, because of the relative youth of deregulation, there have been few empirical studies focusing entirely on electricity prices.[4] Hoare (1996) presents a discussion of the UK electricity market. The collection of papers presented in *The US Power Market: Restructuring and Risk Management* provides a thorough introduction to the electricity industry and related markets. In that collection, the paper by Leong discussing the electricity forward curve touches on some of the models presented below. Bessembinder and Lemmon (2001) present a general equilibrium model of pricing and hedging in electricity forward markets in conjunction with empirical analysis. The paper most similar in spirit to ours is Bhanot (2000) who looks at the behavior of daily electricity prices across U.S. markets.

The remainder of the paper is organized as follows. Section 2 analyzes the data and discusses the characteristics of electricity markets. Section 3 presents several models of

[3]The leverage effect refers to the asymmetric behavior of equity prices – negative shocks tend to increase volatility more than positive shocks. The connection to leverage is that a lower stock price reduces the value of equity relative to debt, thereby increasing the leverage of the firm and consequently the risk of holding the stock.

[4]For a theoretical treatment of commodity prices see the papers by Chambers and Bailey (1996), Deaton and Laroque (1992, 1996), and Williams and Wright (1991). Duffie and Gray (1996) examine the volatility of energy prices. Kirk (1996) examines correlation in the energy markets.

electricity prices. Section 4 concludes the paper with a summary of the findings and directions for future research.

2 Data Analysis

The data used in this study consists of hourly electricity prices from California. Because of transmission capacity constraints, California is partitioned into 24 "zones" with a separate market price for each zone. When congestion does not exist, arbitrage opportunities restrict the prices in each zone to be equal.[5] Since prices across zones are likely to behave similarly, we focus on just one zone, denoted NP15, that corresponds to Northern California.[6] The sample begins on April 1st, 1998 (the opening of the market) and ends on August 30th, 2000 for a total of 21,216 observations.

The spot market for electricity is an auction. Generators and distributors submit hourly supply and demand curves. The equilibrium price is the resulting market price, and makes up the price series examined in this paper.

2.1 The Electricity Market

The nature of electricity and the behavior of electricity prices differs from that of other commodity markets. One reason for this difference is that electricity is a non-storable good, implying that inventories cannot be used to arbitrage prices across over time.[7] This inability to use arbitrage arguments for pricing securities creates a need for accurate forecasts of electricity prices and a greater understanding of the price process relative to its equity counterparts.

The behavior of prices is characterized by several distinguishing features, beginning with its regular intraday variation. This is seen quite clearly in the first two figures of

[5]See Borenstein, Bushnell, Knittel, and Wolfram (2001) for a further discussion of this.

[6]The size of the 24 zones are not equal. In fact, two zones, SP15 and NP15, corresponding to Southern and Northern California, constitute a large fraction of the state and its population. Further, roughly 90 percent of the state's electricity consumption occurs within these two zones. Data for Southern California, SP15, was examined in a similar manner and yielded results similar to Northern California. As such, these results are not presented here.

[7]Hydroelectric resources are arguably storable. Water is stored in a reservoir and then released to produce electricity. However, hydroelectric resources require large river systems and are thus infeasible in many regions. In California hydroelectric power represents a relatively small fraction of total electricity generation, compared to nuclear and fossil-fuelled generators.

the appendix. Figure 1 presents average hourly electricity prices measured in dollars per megawatt hour ($/MWh) for weekdays and weekends. As expected, prices are, on average, higher during the week when demand is greater. The price begins to increase at roughly 6:00 a.m., as the populace wakes and the workday begins. This price increase continues throughout the day as demand builds, peaking at 4:00 p.m. Prices begin to fall thereafter as the workday ends and demand shifts to primarily residential usage. Figure 2 presents a sample of hourly price data and demand data for the time period, July 1, 1998 to July 10, 1998. The left vertical axis units are $/MWh and the right vertical axis units are gigawatt hours (GWh). The horizontal axis categories correspond to midnight of that particular day. This figure illustrates more clearly the daily usage pattern and its persistence over time. In addition, it is apparent that prices are mimicking demand.

Electricity prices also contain a strong seasonal component, reflecting heating and cooling needs. This feature is illustrated in Figure 3, which plots hourly weekday averages for each of the four seasons. Northern California's primary electricity consumption occurs during the summer when air conditioning is needed. The swing months, which comprise spring and fall, show only a minor seasonal effect which are likely due to residual cooling needs. Since high temperatures often extend into fall months, we see that the electricity is less expensive in the spring than in the fall. Winter electricity prices are the lowest of the seasons since most heating needs are met by natural gas consumption.

Finally, electricity markets are characterized by distribution and transmission constraints. Once generated, electricity travels along a network of distribution and transmission lines designed to take the particles from the generation source to the demand source. Each line within this network has a "capacity" or a maximum amount of electricity that it can carry at a given moment. Once constrained, the marginal cost of transmission becomes infinite. This implies, and is often the case, that sections of an electricity market can become isolated from the rest of the market. Once this occurs, the generators in the isolated market enjoy a greater level of market power, or influence over prices. Not surprisingly, when a market becomes isolated, the price of electricity rises rapidly as generators exploit their position.[8] Indeed, California electricity prices are consistent with this theory. Within a time span of 24 hours, the spot price for electricity in California can move from a price of less than $5 per megawatt hour to $750 per megawatt hour.

Figure 4 plots the entire hourly price series from April 1, 1998 to July 31, 2000, and illustrates two points. First, prices make dramatic swings which tend to occur in clusters.

[8]See Borenstein, Bushnell and Knittel (1999) for a discussion of this idea.

This is a result of demand approaching – and in some cases exceeding – system capacity. Second, there are several negative prices, which are a consequence of the inability to freely dispose electricity coupled with non-trivial start up costs for generators.

2.2 Distributional Properties

In this subsection, we discuss the distributional properties of electricity prices. Figure 5 presents the empirical histogram for our price series, overlaid with a normal density curve. Figure 6 presents a QQ-plot of the data. The superimposed line joins the first and third quartile of the data and is a robust linear fit of the sample order statistics. Normally distributed data will appear linear in this plot. Both figures illustrate the deviation from normality. Figure 6 shows the heavy tail of the distribution and, to a lesser extent, the atom at 0.

Table 1 presents summary statistics for several subperiods of the sample:

1. the pre-crisis period from April 18, 1998 to April 20, 2000,

2. the crisis period from May 1, 2000 to August 31, 2000, and

3. the months of May through August for the pre-crisis period.

The last period allows for a more accurate comparison between the pre-crisis and crisis periods, by controlling for seasonal effects. The impact of the crisis is evident in each of the sample moments beginning with a fourfold increase in the mean price of electricity, even after controlling for the same months in the pre-crisis period. There is also a significant increase in the volatility of prices. The kurtosis falls significantly during the crisis period, but this is a direct result of price caps and the increase in the standard deviation. Relative to a normal distribution, the kurtosis is significantly larger in both time periods; large deviations from the mean are a relatively common occurrence. To a lesser extent, the distribution is skewed with a long right tail.

Several other distributional features are evident in Figure 4. The conditional mean price is varying over time in a systematic fashion. This fact is even more apparent in Figures 1 and 2. Figure 4 also shows that the variance of electricity prices is time varying and exhibits clustering.

2.3 Temporal Properties

While the price series has several distributional similarities with other tradeable assets, their temporal properties are quite different. Figure 7 plots the autocorrelation function for the level of prices. The autocorrelations are statistically significant even beyond 1000 lags. Also clear from the correlogram is the intraday usage pattern and to a lesser extent a weekend/weekday cycle. With a longer lag length, a seasonal pattern emerges as well, but the plot is truncated at 1000 periods for visual ease. This result is in stark contrast to the predictability of equity prices, which are commonly assumed to follow a random walk with drift.[9]

Though prices have an extremely long memory, visual inspection of the correlogram is consistent with the intuition that prices do not appear to be exploding. The decay in the autocorrelation function is fairly rapid, at least initially.[10] This hypothesis is tested more formally as follows. Consider the simple approximation to the price generating process:

$$p_t = \alpha + \beta p_{t-1} + \eta_t \quad (1)$$

$$\eta_t = \gamma \eta_{t-1} + \varepsilon_t \quad (2)$$

where p_t is the price at time t, α, β, and γ are unknown coefficients, and $\{\varepsilon_t\}$ is a Gaussian white noise process with variance σ_η^2.[11] In the presence of serially correlated errors, Phillips and Perron (1988) show that the parameters of (1) can be consistently estimated by ordinary least squares (OLS). The test statistic for a unit root, however, must be modified to take serial correlation into account. Using the Newey-West estimator, the corrected t-stat under the null hypothesis of a unit root in the presence of serial correlation is -1153.[12] With a 5% critical value of -2.89, the null of a unit root is rejected at all standard significance levels.

[9] We recognize that studies have found a predictable element to stock prices. However, the autocorrelations, though statistically significant, are very small in comparison to the autocorrelations present in the electricity series.

[10] In fact, a statistically significant but very small time trend is found in the series. Because of the magnitude of the coefficient and lack of economic support for the presence of such a trend, the statistical modeling is performed without incorporating a time trend.

[11] A Gaussian white noise process is a sequence of independent, normally distributed, mean zero random variables.

[12] See Hamilton (1994), Chapter 17 for details of this test and the Newey-West estimator. In fact, we ran another test allowing for fourth order serial correlation and the results are unchanged. The null of a unit root is rejected at all standard significance levels.

Figure 8 plots the autocorrelation function for the square of prices, which bears a striking resemblance to Figure 7. This plot confirms the observation made above regarding the volatility clustering. The second moment exhibits a high degree of persistence even after several hundred lags. As with the correlogram for the price level, the intraday usage pattern is immediately clear, as is the weekend/weekday cycle. This result is similar to equity markets where volatility persistence is a common phenomenon, although without the distinct seasonality.

From this preliminary analysis of the data, it is clear that any modeling effort should take into account the following characteristics of the prices series:

1. mean reversion,
2. time of day effects,
3. weekend/weekday effects,
4. seasonal effects,
5. time varying volatility and volatility clustering, and
6. extreme values.

We consider the right censoring and negative prices to be of less importance. The effect of censoring on the dynamics of prices is likely to have a secondary effect, and price caps have recently been abandoned. Accurately capturing the censoring would require a latent variable model, or a much more complicated statistical specification; both, exercises for future work. Negative prices are an increasingly rare occurrence, whose implications for pricing financial securities are of little consequence.

3 Models and Results

This section presents several different models of electricity prices. Each model is first motivated and discussed in the context of the preliminary data analysis. This is followed by an analysis of the estimation results and forecasting performance of the model.

The models are estimated using two subsamples: the pre-crisis period defined as April 1, 1998 to April 30, 2000, and the crisis period defined as May 1, 2000 to August 31,

2000. For each period, the last week of data are withheld in order to measure the out-of-sample forecasting ability. For example, Model 1 is first estimated using data from April 1, 1998 to April 23, 2000. Hourly forecasts over the period April 24, 2000 to April 30, 2000 are computed and then compared to actual prices. The same model is then estimated using data from May 1, 2000 to August 24, 2000. Forecasts are produced for the period August 25, 2000 to August 31, 2000 and again compared to actual prices. A one week forecast horizon is chosen or two reasons. First, the frequency of our data dictates that only short-term forecasts are feasible. Second, most electricity contracts are short-term, ranging from one day to several months. All models are estimated by conditional maximum likelihood.[13]

We discuss each model in the context of the data analysis performed above and the forecasting performance of the model. Each subsequent model builds on the previous model by introducing new aspects whose purpose is to capture another feature of the data.

3.1 Model 1: Mean-Reverting Processes

Traditional financial models typically begin with the Black-Scholes assumption of geometric Brownian motion or log normal prices. This assumption is inappropriate in the context of electricity prices for many reasons, primarily because of the predictability of electricity prices. An alternative model used in practice is the Ohrnstein-Uhlenbeck process. This continuous time model allows for autocorrelation in the series by specifying prices as:

$$dp(t) = \kappa [\mu - p(t)]\, dt + \sigma db(t), \qquad p(0) = p_0, \tag{3}$$

where $p(t)$ is the price of electricity at time t, κ, μ, and σ are unknown parameters, and $\{b(t)\}$ is a standard Wiener process. The intuition behind this specification is that deviations of the price from the equilibrium level, $[\mu - p(t)]$, are corrected at rate κ and subject to random perturbations, $\sigma db(t)$.[14]

Equation (3) is simply a first order autoregressive model in continuous time. This

[13]The conditioning is due to the presence of lagged dependent variables on the right hand side. Under the assumption of stationarity, the unconditional density for these initial observations could be specified. However, because of the large number of observations, the impact of the first few observations is likely to be negligible, even in the presence of long-term persistence.

[14]This assumes that $\kappa > 0$, a requirement for stationarity of the process.

may be seen by integrating equation (3) to obtain:

$$p(t) = e^{-\kappa t} p_0 + \mu \left(1 - e^{-t\kappa}\right) + \int_0^t e^{\kappa(s-t)} \sigma db(s) . \tag{4}$$

Algebra produces the "exact" discrete time version of equation (4):

$$p_t = \alpha_0 + \beta_1 p_{t-1} + \eta_t, \tag{5}$$

where $\alpha_0 = \mu\left(1 - e^{-\kappa}\right)$, $\beta_1 = e^{-\kappa}$, and $\eta_t = \int_{t-1}^t e^{\kappa(s-t)} \sigma db(s)$.[15] The error term, η_t, in equation (5) is Gaussian white noise with variance σ_η^2 equal to $\sigma^2\left(1 - e^{-2\kappa}\right)/2\kappa$, by Ito isometry. Thus, prices are Markovian with a Gaussian transition density. The conditional mean is $\alpha_0 + \beta_1 p_{t-1}$ and conditional variance is σ_η^2.

Because the modeling aspect of this study is focused on forecasting, as opposed to parameter interpretation, we estimate the discrete time parameters in equation (5): α_0, β_1, and σ_η^2.[16] The results are presented in Tables 2 and 5.

While this model captures some of the autocorrelation present in the price series, it suffers from several serious shortcomings. First, it ignores all cycles present in the series: intraday, weekend/weekday and seasonal. Second, it assumes the error structure is independent across time. Third, it assumes that the volatility is constant over time. Fourth, the normality assumption cannot reproduce the extreme swings found in the data. All of these shortcomings appear in the forecast results presented in Figures 9 and 10 for the pre-crisis and crisis periods, respectively. The figures present the actual and forecasted price series represented by the dashed and thin, solid line, respectively. The two thick lines represent two standard errors around the forecasted value. Quickly, the forecast becomes fixed at the unconditional mean of the process. In addition, the standard errors are enormous in the crisis period.

[15]The terminology "exact discrete time representation" is borrowed from Bergstrom (1984) and is intended to differentiate this manipulation from an approximation such as an Euler discretization.

[16]Estimates of the continuous time parameters (κ, μ, σ) may be obtained using several methods. One approach, given maximum likelihood estimates of the discrete time parameters, simple algebra using the mappings between the two sets of parameters will produce maximum likelihood estimates of the continuous time parameters by the invariance property of maximum likelihood. Asymptotic standard errors may be obtained using the delta method. Another approach is to directly maximize the likelihood function with respect to the continuous time parameters as opposed to the discrete time parameters. Because of the one-to-one mapping between the discrete time and continuous time parameters, both approaches should yield equivalent estimates.

3.2 Model 2: Time-Varying Mean

The second model addresses the systematic variation found in electricity prices. We consider the following extension to equation (3)

$$dp(t) = \kappa(\mu(t) - p(t))\,dt + \sigma db(t), \tag{6}$$

where

$$\begin{aligned} \mu(t) &= \alpha_1 1(t \in Peak) + \alpha_2 1(t \in Off\ Peak) \\ &\quad +\alpha_3 1(t \in Weekend) + \alpha_4 1(t \in Fall) + \alpha_5 1(t \in Winter) \\ &\quad +\alpha_6 1(t \in Spring) \end{aligned} \tag{7}$$

and $1(\cdot)$ denotes the indicator function. For example,

$$1(t \in Peak) = \begin{cases} 1 & \text{if the hour of the day falls between 6:00 AM \& 10 PM., and} \\ 0 & \text{otherwise.} \end{cases}$$

This specification implies that $\mu(t)$ is a step function, constant across any one hour. Integrating equation (6) yields

$$\begin{aligned} p(t) &= e^{-\kappa t}p_0 + \int_0^t e^{-\kappa(t-s)}\kappa\mu(s)\,ds + \int_0^t e^{-\kappa(t-s)}\sigma db(s) \\ &= e^{-\kappa}p(t-1) + \int_{t-1}^t e^{-\kappa(t-s)}\kappa\mu(s)\,ds + \int_{t-1}^t e^{-\kappa(t-s)}\sigma db(s). \end{aligned} \tag{8}$$

Considering one unit of time as an hour, $\mu(t)$ is constant over the interval $[t-1, t)$. Therefore, the exact discrete time version of (8) is

$$p_t = \alpha_t + \beta_1 p_{t-1} + \eta_t, \tag{9}$$

where $\alpha_t = \mu(t)(1 - e^{-\kappa})$, $\beta_1 = e^{-\kappa}$ and $\eta_t = \int_{t-1}^t e^{-\kappa(t-s)}\sigma db(s)$. The only difference between equations (9) and (5) is in the intercept. As such, prices are again Markovian with a Gaussian transition density.

Equation (9) may be viewed as an ARMAX(1,0) model with the exogenous variables consisting of six binary variables. Two of the binary variables indicate whether the observation falls in a peak or off-peak time period ($Peak_t$, $OffPeak_t$), one of the binary variables indicate whether the observation falls on a weekend or not ($Weekend_t$), and three of the binary variables indicate in which season the observation occurs ($Fall_t$,

$Winter_t$, $Spring_t$).[17] More explicitly, equation (9) may be written as

$$\begin{aligned} p_t &= \alpha_1 Peak_t + \alpha_2 OffPeak_t + \alpha_3 Weekend_t + \alpha_4 Fall_t \\ &\quad + \alpha_5 Winter_t + \alpha_6 Spring_t + \beta_1 p_{t-1} + \eta_t. \end{aligned} \tag{10}$$

The estimates for the pre-crisis and crisis periods are presented in Tables 2 and 5. The coefficients on the peak/off-peak indicators reflect intraday usage patterns. Interestingly, the sign on the fall variable in the pre-crisis period is positive suggesting that prices are, on average, higher during fall months than summer months. This result is due to a combination of a relatively cool June and unseasonably warm September in 1998 and 1999. An indicator for spring is the only seasonal binary variable included in the crisis regression since the period only encompasses spring and summer. From the estimates, the majority of the high prices occurred during June through August. Inspection of the Durbin-Watson t-statistic suggests that the residuals are correlated.[18] This is not a surprising result given the preliminary data analysis.

While all but one of the estimated coefficients is highly statistically significant, the improvement over the previous model in terms of forecasting performance is minimal. Out-of-sample forecasts for the time-varying mean model in both pre-crisis and crisis periods are found in Figures 11 and 12, respectively. The forecasted price series still fails to capture the erratic nature of the true price series.

3.3 Model 3a: Jump-Diffusion Process

As a first attempt to capture the leptokurtosis present in the price series, we turn to a popular extension of the standard diffusion process: the jump-diffusion process. Our price process is now specified by appending an additional term to Equation (6), yielding:

$$dp\,(t) = \kappa\,(\mu\,(t) - p\,(t))\,dt + \sigma_b db\,(t) + z dq\,(t) \tag{11}$$

where $q\,(t)$ is a Poisson process with intensity λ, z is a draw from a normal distribution with mean μ_z and standard deviation σ_z. We assume that the Wiener process, Poisson process and jump size are mutually independent.

[17] We also estimated the models using separate binary variables for each hour. However, there was no appreciable improvement in forecasting ability.

[18] Since the specification in (9) contains a lagged dependent variable, the Durbin Watson statistic is biased towards 2 and has reduced power. As such, we use Durbin's t-statistic which is asymptotically equivalent to the Durbin h-statistic. (See Durbin 1970)

As Ball and Torous (1983) note, empirical implementation of (11) is difficult. As such, we follow their approach and approximate this model with a mixture of normals. The intuition behind such a model can be explained in terms of a coin tossing experiment. At each time period, a λ-coin is flipped. That is, with probability λ, the coin shows a head and, with probability $(1-\lambda)$, the coin shows a tail. If the coin toss results in a tail, then no jump has occurred during the observation interval and the price process has behaved according to equation (6). This is equivalent to drawing the price at time t from a normal distribution with mean $\alpha_t + \beta_1 p_{t-1}$ and variance σ_η^2. If the coin toss results in a head, than a jump has occurred during the observation interval. Now the price is drawn from a normal distribution with mean $\alpha_t + \beta_1 p_{t-1} + \mu_2$ and variance $\sigma_\eta^2 + \sigma_z^2$. Note, that while the mean may rise of fall when a jump occurs, the variance always increases.[19]

The conditional likelihood function is thus:

$$\begin{aligned} L &= \prod_{t=2}^{T} (1-\lambda)\,\phi\left(\frac{P_t - (\alpha_t + \beta_1 P_{t-1})}{\sigma_b}\right)\sigma_b^{-1} \\ &\quad +\lambda\phi\left(\frac{P_t - (\alpha_t + \beta_1 P_{t-1} + \mu_z)}{\left(\sigma_b^2 + \sigma_z^2\right)^{1/2}}\right)\left(\sigma_b^2 + \sigma_z^2\right)^{-1/2}, \end{aligned}$$

where $\phi(\cdot)$ is the standard normal density.

To ease in interpreting the results, and without loss of generality, we can write the likelihood function as,

$$\begin{aligned} L &= \prod_{t=2}^{T} (1-\lambda)\,\phi\left(\frac{P_t - (\alpha_t + \beta_1 P_{t-1})}{\sigma_b}\right)\sigma_b^{-1} \\ &\quad +\lambda\phi\left(\frac{P_t - (\alpha_t + \beta_1 P_{t-1} + \mu_z)}{(\xi\sigma_b)}\right)(\xi\sigma_b)^{-1}, \end{aligned}$$

where ξ represents the proportional increase in volatility during hours that contain a jump.

The results are presented in Tables 2 and 5. Beginning with the pre-crisis period, the probability of a jump occurring during any hour, λ, is estimated as 0.03. Thus, a jump occurs, on average, every 33 hours. The average size of a jump, μ_z, is \$11.06. This suggests that prices jump up more often than they jump down. During hours that contain a jump, the price volatility is 62 times larger than when a jump does not occur:

[19] Theoretically, any number of jumps may occur in any time interval. For simplicity, we assume that at most one jump can occur during any hour.

an increase from \$6.45 to \$401.14. The parameter estimate of the slope coefficient, β_1, is consistent with the previous model's estimates, as are estimates of α_1 through α_6.

The crisis period estimate of λ reflects an increase in the probability of a jump to one every 20 hours. The estimated average jump size has fallen to \$0.66, reflecting an increase in negative jumps off-setting positive price spikes. The estimated price volatility is 45 times larger when a jump occurs. Again, parameter estimates of α_1 through α_6 are consistent with previous models. However, there is an increase in the estimated AR coefficient, β_1, to over 0.96. While significantly larger than previous model estimates, it is still statistically far from the boundary of non-stationarity.

Generating useful forecasts from this model is difficult. One must simulate a forecasted path because of the model's dependency on random jumps. Of course, there are a continuum of possible paths that may be simulated by, intuitively, flipping a λ-coin each time period, and drawing from the appropriate conditional distribution. Combining many simulated paths averages out the excess variation induced from the jumps, and leaves a very smooth forecast representing a number falling somewhere between the means of the two conditional distributions that make up the mixture. This result is illustrated in Figures 13 and 14.[20]

3.3.1 Model 3b: Time Dependent Jump Intensity

We now refine the model specification above by allowing the jump intensity parameter to vary over time. There are several reasons for doing this including the fact that jumps are more likely to occur when transmission lines become congested. This suggests that during high demand periods a jump in prices is more probable. Thus, we allow the jump intensity to vary by the time of day and season.

$$\lambda(t) = \lambda_0 + \lambda_{peak} Peak_t + \lambda_{weekend} Weekend_t + \lambda_{fall} Fall_t + \lambda_{win} Winter_t + \lambda_{spr} Spring_t$$

The results listed in Tables 2 and 5 confirm our observation. In both the pre-crisis and crisis samples, the probability of a jump increases during peak hours and decreases during the spring and winter months. In addition, we also find a significant weekend effect in the jump intensity.

[20] We ran 5000 Monte Carlo simulations of forecasted price paths. We then average the price paths over time to obtain a final forecasted path. Because of the high variance when a jump occurs, forecasted prices above \$250 and under –\$250 during the crisis period are set to \$250 and –\$250, respectively. This is done to respect the price cap and ignore impossibly low prices.

The same forecasting issues discussed above are present here. Figures 15 and 16 present the forecast results. The performance, as measured by RMS forecast error, is poor.

3.4 Model 4: ARMAX

The fourth model takes a more traditional time series approach to modeling electricity prices. We begin by relaxing the Markovian assumption on prices by introducing serial correlation in the error term. Working in a discrete time framework, price dynamics are now specified as

$$p_t = \alpha_t + \eta_t \tag{12}$$

$$\beta(L)\,\eta_t = \delta(L)\,\varepsilon_t \tag{13}$$

where $\beta(L)$ and $\delta(L)$ are the autoregressive and moving average polynomials in the lag operator L, respectively. These operators are defined as

$$\beta(L) = 1 - \beta_1 L - \beta_2 L^{24} - \beta_3 L^{25} \tag{14}$$

$$\delta(L) = 1 - \delta_1 L - \delta_2 L^{24} - \delta_3 L^{25} \tag{15}$$

The mean α_t is as specified in equation (9) and $\{\varepsilon_t\}$ is Gaussian white noise with variance parameter σ^2. The motivation for equations (12) through (15) follows from an examination of the correlogram, which shows high correlation between the current price and the previous day's prices.

Clear from the parameter estimates in Tables 2 and 5 is that electricity prices are not reasonably approximated by a univariate Markov process. All lag parameters are highly statistically significant.[21] The estimated deterministic component of prices is consistent with Model 2's results in both subsamples.

Improvement in forecast accuracy over pervious models is significant in the pre-crisis period as measured by root mean square forecast error. In the crisis period, the difference in RMS forecast error across models is relatively small. Forecasted and actual prices are plotted in Figures 17 and 18 along with asymptotic 95% confidence intervals. The improvement in forecasting accuracy is due primarily to the incorporation of higher order lags.

[21] We examined several other specifications incorporating higher order lags at 24 hour intervals. While the estimated coefficients are statistically significant, the contribution to forecasting accuracy is minimal.

3.5 Model 5: EGARCH Processes

The preliminary data analysis revealed that electricity prices exhibit volatility clustering. In addition, intuition tells us that it is also possible that innovations to the prices series have an asymmetric impact on the price volatility. *A priori*, we expect positive price shocks to increase volatility more than negative surprises. The intuition behind this is that a positive shock to prices is really an unexpected positive demand shock. Therefore, since marginal costs are convex, positive demand shocks have a larger impact on price changes relative to negative shocks. To test for this effect, we begin by specifying the price level as the sum of a deterministic component and a stochastic component

$$p_t = \alpha_t + \eta_t \tag{16}$$

where α_t is unchanged from above. The random term η_t is assumed to follow an autoregressive process

$$\beta(L)\,\eta_t = \nu_t \tag{17}$$

where $\beta(L)$ is the lag operator defined in equation (14) above. To capture the conditional heteroscedasticity, we adopt the EGARCH model of Nelson (1991), modeling ν_t as

$$\nu_t = \sqrt{h_t}\varepsilon_t \tag{18}$$

$$\ln(h_t) = \theta + \sum_{i\in\{1,24,25\}} \kappa_i g(z_{t-i}) + \gamma_1 \ln(h_{t-1}), \tag{19}$$

where

$$\begin{aligned} g(z_s) &= \{\psi z_s + |z_s| - E(|z_s|)\} \\ z_s &= \nu_t/\sqrt{h_t} \end{aligned}$$

and $\{\varepsilon_t\}$ is Gaussian white noise with unit variance. The coefficient ψ controls the degree of asymmetry. When $\psi = 0$, there is no asymmetric effect of past shocks on current variance. If $-1 < \psi < 0$, then a positive shock increases variance less than a negative shock. If $\psi < -1$ then positive shocks reduce variance while negative shocks increase variance. Our prediction is that $\psi > 0$, implying that the effect of positive shocks on the variance of prices is amplified over negative shocks.

Parameter estimates for equations (16) through (19) are found in Tables 3 and 6. In the pre-crisis period, estimated coefficients in the deterministic component of prices are consistent with previous specifications. The autoregressive coefficients in both periods are consistent in terms of signs, although the EGARCH model's estimates are a bit smaller in magnitude when compared with previous models. As anticipated, the asymmetry parameter is positive and significant, suggesting the presence of an "inverse leverage effect". Thus, positive shocks to prices amplify the conditional variance of the process more so than negative shocks.

The forecasting ability of this model in the pre-crisis period is actually the poorest of all models considered thus far. In the crisis period, the EGARCH specification has the best forecasting performance, but not dramatically so. Figures 19 and 20 plot the forecasted and actual price series along with the asymptotic 95% confidence interval.

3.6 Model 6: Incorporating Weather Data

The final model extends the ARMAX specification above (Model 4) by incorporating temperature data. Before specifying the model, however, several issues must be resolved. Northern California is both large and geographically diverse. There are inland valleys, mountainous regions, and coastal areas; each of which has a unique climate. As such, a single temperature from any one area is inappropriate.

Despite this diversity, the majority of electricity consumption is concentrated in a small number of areas. We gathered hourly temperature data from the National Oceanographic and Atmospheric Association (NOAA) corresponding to reading stations at San Francisco, Sacramento, and Fresno.[22] Figure 21 shows a scatter plot of price vs temperature, as measured in San Francisco, overlaid with a univariate regression line. Several characteristics are evident in the plot, including a nonlinear relationship. The price temperature relationship is negative for temperatures below 50 degrees, when heating becomes necessary. Since electric heating is rare in California, this relationship is subtle. After 55 degrees, the relationship turns positive as commercial cooling needs begin. The sensitivity is also much greater at higher temperatures than lower temperatures. Price caps are evident from the rows of data points at the $250, $500, and $750 marks. Scatter plots of price and temperature in other areas are not presented here because of their similarity to Figure 21.

To capture the diversity of regions, each of the three temperature series are initially included in the specification. To capture the nonlinearity, the square and cube of each temperature series are included as well. However, the additional explanatory power of

[22]Actually, San Francisco is not included in zone NP15. However, its proximity and similar climate to other areas in the zone motivate its inclusion.

additional temperature series is negligible – as is their impact on forecasting accuracy. Thus, the pre-crisis and crisis results for the following specification are presented in Tables 4 and 7.

$$p_t = \alpha_t + \eta_t \tag{20}$$

$$\beta(L)\eta_t = \delta(L)\varepsilon_t \tag{21}$$

where

$$\begin{aligned}\alpha_t &= \alpha_1 1(t \in Peak) + \alpha_2 1(t \in Off\ Peak) + \alpha_3 1(t \in Weekend) \\ &+ \alpha_4 1(t \in Fall) + \alpha_5 1(t \in Winter) + \alpha_6 1(t \in Spring) \\ &+ \alpha_7 Temp_t + \alpha_8 Temp_t + \alpha_9 Temp_t^3,\end{aligned} \tag{22}$$

the AR and MA polynomials are unchanged from above (i.e. 1, 24, and 25 period lags), and $\{\varepsilon_t\}$ is Gaussian white noise. The temperature variable used is an equally weighted average of the temperatures from each of the three cities.[23]

Referring to Tables 4 and 7, the temperature variables are all highly statistically significant during the pre-crisis period. The RMS forecast error is also the lowest of all models examined, though the difference from the previous ARMAX model is small, roughly 2.0. Forecasts are presented in Figure 22 and are similar to the forecasts of Model 4 (Figure 17). During the crisis period, the price-temperature association breaks down. All parameters are insignificant suggesting that other forces were at work during this period. This result highlights the limitations of simple statistical models relative to structural models. It is extremely difficult to predict this sort of breakdown or shift in the data generating process without information other than past prices and weather. Marginal cost data and the competitive nature of the market become crucial in understanding price response to demand shocks.

3.7 Non-normal Distributions

As an initial attempt to recognize the non-normality of the transition densities, the models are re-estimated and new forecasts are generated using the natural logarithm of prices. This transformation implies that the original price process is lognormal. In order

[23]We examined other weighting schemes, including population based weights and demand based weights. Each had a minimal impact on the results.

to perform this transformation, nonpositive prices are set to missing and dropped from the analysis.[24]

The estimation results offer no meaningful change in terms of estimated parameter significance or direction of association. And, the transformation had a negligible effect on forecasting performance. The reason for this lack of improvement is that the kurtosis is the dominant feature of the price series. The skewness, while clearly present, is not responsible for the forecasting performance of any model. The frequency of large price deviations from the conditional mean create big forecast errors, which translate into high RMS forecast errors. Even after using an alternative measure of forecasting performance (average absolute deviation), forecasting performance is minimally affected by the transformation.[25]

4 Conclusions and Directions for Further Research

The events of the past two years in California have made understanding the stochastic properties of deregulated electricity prices of the upmost importance. Retail electricity companies, large consumers, and entrants are increasing their use of electricity derivatives to hedge against price risk in this new era. However, the idiosyncracies of electricity prices make existing statistical models of asset prices of little practical use in modeling electricity prices.

In this paper, we have provided a detailed examination of deregulated prices. Unlike other commodity prices, electricity prices show a high degree persistence in both price levels and squared prices. In addition, because electricity prices closely track demand movements, we also find strong deterministic cycles including, intraday, day of week, and seasonal effects. Finally, the large values of higher order moments relative to a Gaussian distribution renders models based on normality and log-normality of limited use in representing electricity prices.

Forecasting performance, a crucial component of security valuation in the electricity industry, is relatively poor for most standard asset pricing models. Non-Markovian specifications are necessary to capture the high persistence in the level of prices and volatility. We also document an inverse leverage effect where positive price shocks increase price volatility more than negative shocks.

[24]This represented less than 1% of the data.

[25]For space considerations, the results are not presented here.

The next step in empirical research involves expanding the model set. Structural models of supply and demand should bring insight into the price generating process and potentially produce more accurate forecasts. Relaxing the normality assumption should also aid in better representing the stochastic properties of prices. Given the growing market for securities written on electricity and the empirical challenges modeling the price series brings, we expect future empirical research of electricity prices to be an active area.

References

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

[2] Bergstrom, A. R., 1984, "Continuous Time Stochastic Models and Issues of Aggregation over Time." in Griliches, Zvi, ed. and Intriligator, Michael D., ed. *Handbook of Econometrics, Volume II*, Amsterdam; New York and Oxford: North-Holland, Elsevier Science 1984, pp. 1146-1212.

[3] Bessembinder, H. and M. J. Lemmon, 2001, "Equilibrium Pricing and Optimal Hedging in Electricity Forward Markets," forthcoming, *The Journal of Finance.*

[4] Bhanot, K., 2000, "Behavior of Power Prices: Implications for the Valuation and Hedging of Financial Contracts," *The Journal of Risk*, **2**(3), 43-62.

[5] Black, F., 1976, "Studies in Stock Price Volatility Changes," *Proceedings of the Business and Economics Statistics Section*, Amercian Statistical Association, 177-181.

[6] Borenstein, S., Bushnell, J. and C. R. Knittel, 2000, "Market Power in Electricity Markets: Beyond Concentration Measures," *The Energy Journal*, **20**(4). 65-83.

[7] Borenstein, S., Bushnell, J., Knittel, C. R. and C. Wolfram, 2001, "Market Inefficiency in California's Electricity Markets," mimeo, Boston University.

[8] Chambers, M. and R. Bailey, 1996, "A Theory of Commodity Price Fluctuations," *Journal of Political Economy*, 104, 924-957.

[9] Deaton, A. and G. Laroque, 1992, "On the Behavior of Commodity Prices," *Review of Economic Studies*, 59, 1-23.

[10] Deaton, A. and G. Laroque, 1996, "Competitive Storage and Commodity Price Dynamics," *Journal of Political Economy*, 104, 826-923.

[11] Duffie, D. and S. Gray, 1996, "Volatility in Energy Prices," *Managing Energy Price Risk*, Risk Publications.

[12] Durbin, J., 1970, "Testing for Serial Correlation in Least-Squares Regression When Some of the Regressors are Lagged Dependent Variables," *Econometrica*, 38, 410-421.

[13] Hamilton, J., 1994, *Time Series Analysis*, Princeton University Press.

[14] Hoare, J., 1996, "The UK Electricity Market," *Managing Energy Price Risk*, Risk Publications.

[15] Kirk, E., 1996, "Correlation in Energy Markets," *Managing Energy Price Risk*, Risk Publications.

[16] Nelson, D. B., 1991, "Conditional Heterskedasticity in Asset Returns," *Econometrica*, 59, 347-370.

[17] Phillips, P. C. B. and P. Perron, 1988, "Testing for a Unit Root in Time Series Regressions," *Biometrika*, 75, 335-346.

[18] Priestley, M. B., 1981, *Spectral Analysis and Time Series*, Academic Press.

[19] Risk Publications, 1996, *The US Power Market.*

[20] Wright, B. and J. Williams, 1989, "A Theory of Negative Prices of Storage," *Journal of Futures Markets*, 9, 1-13.

A Figures & Tables

Figure 1: Average Hourly Electricity Prices Across the Entire Sample

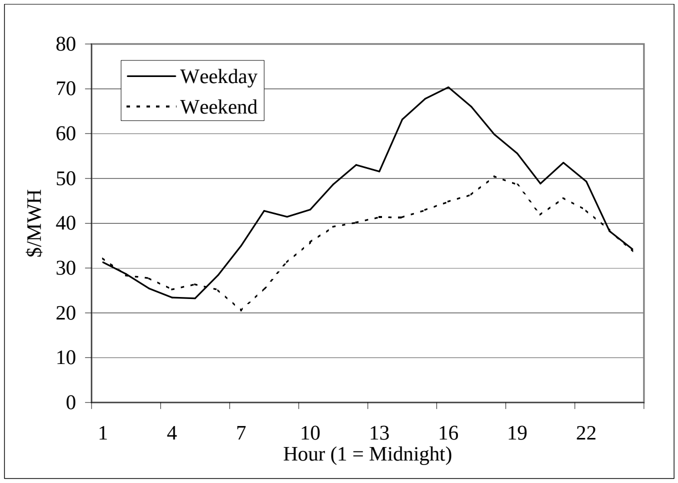

Figure 2: A Sample of Hourly Electricity Prices for the Period July 1, 1998 to July 10, 1998

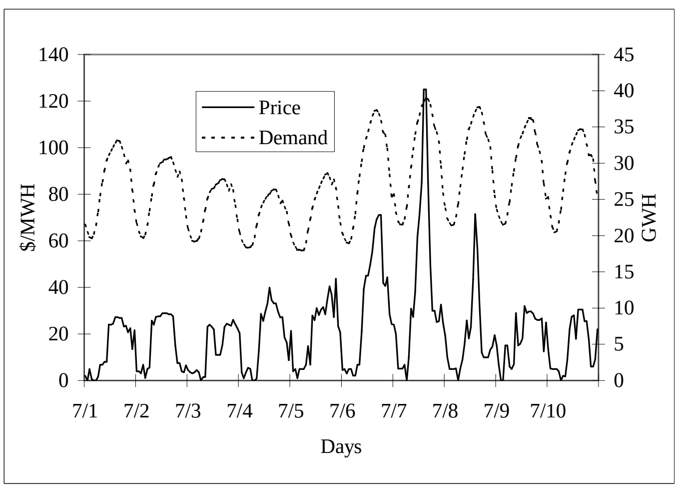

Figure 3: Average Weekday Hourly Electricity Prices by Season

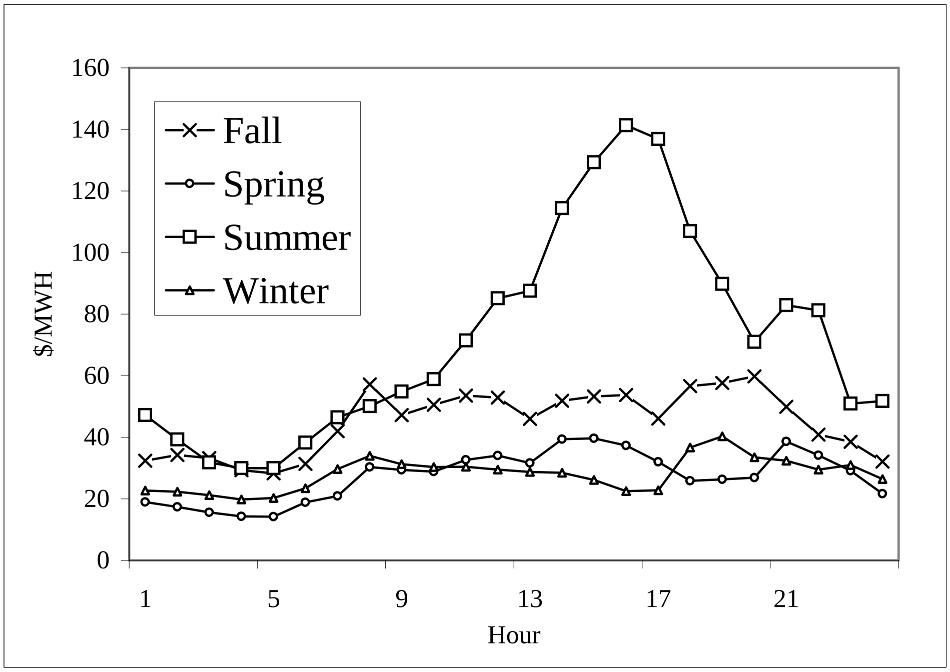

Figure 4: Hourly Electricity Prices for April 1, 1998 to July 31, 2000

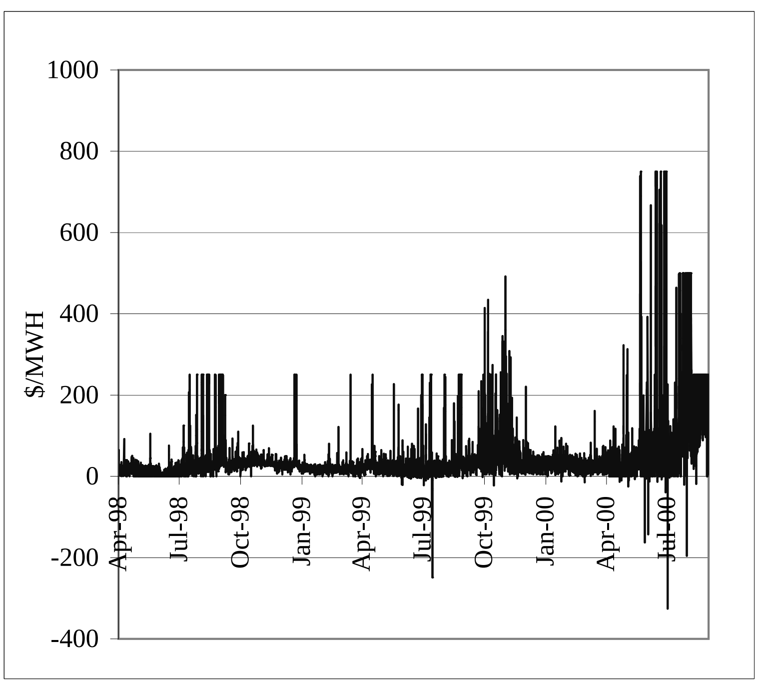

Figure 5: Empirical Histogram of Electricity Prices with Normal Density Superimposed

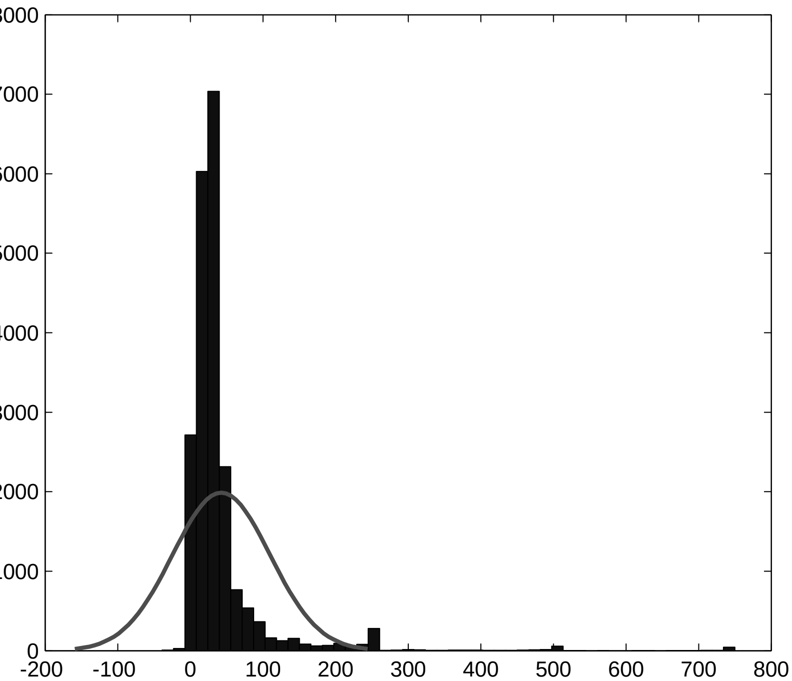

Figure 6: QQ-Plot of Electricity Prices

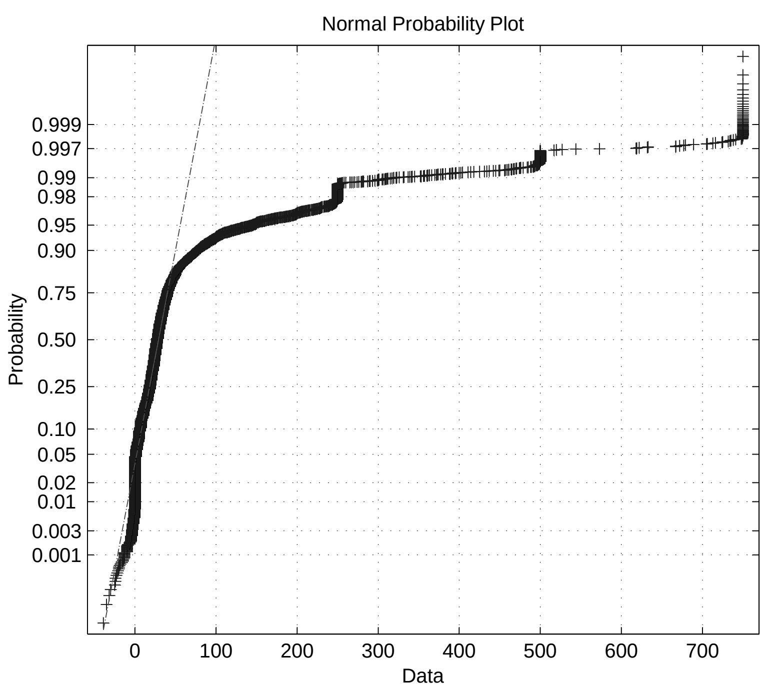

Figure 7: Autocorrelation Function for Electricity Prices

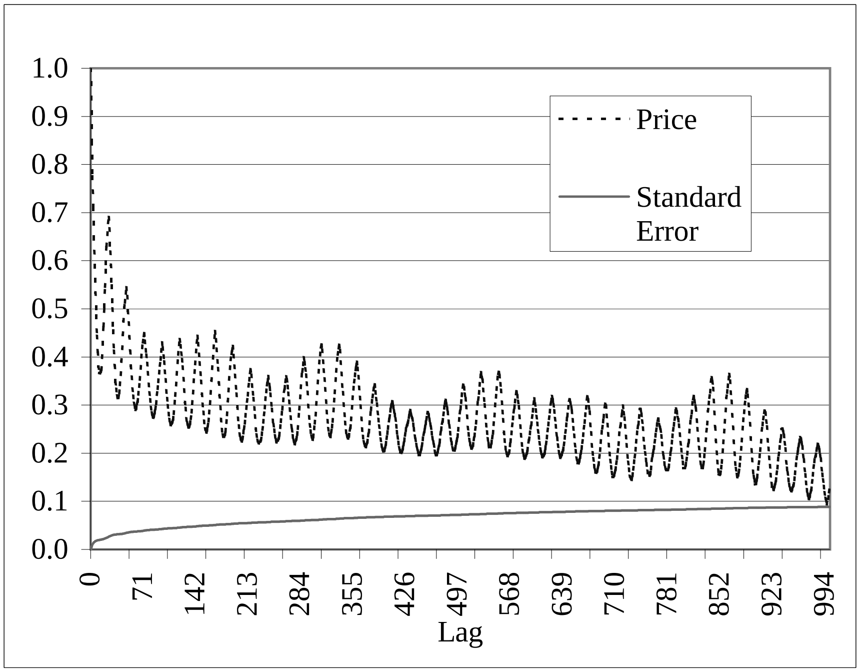

Figure 8: Autocorrelation Function for the Square of Electricity Prices

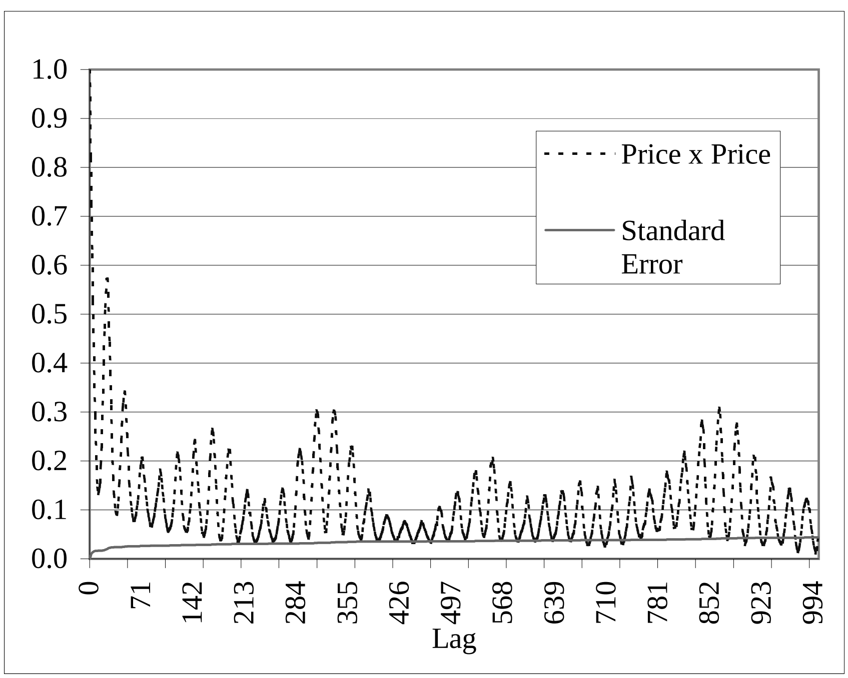

Figure 9: Week-Ahead Forecasts Over the Pre-Crisis Period for the Mean Reverting Model (Model 1)

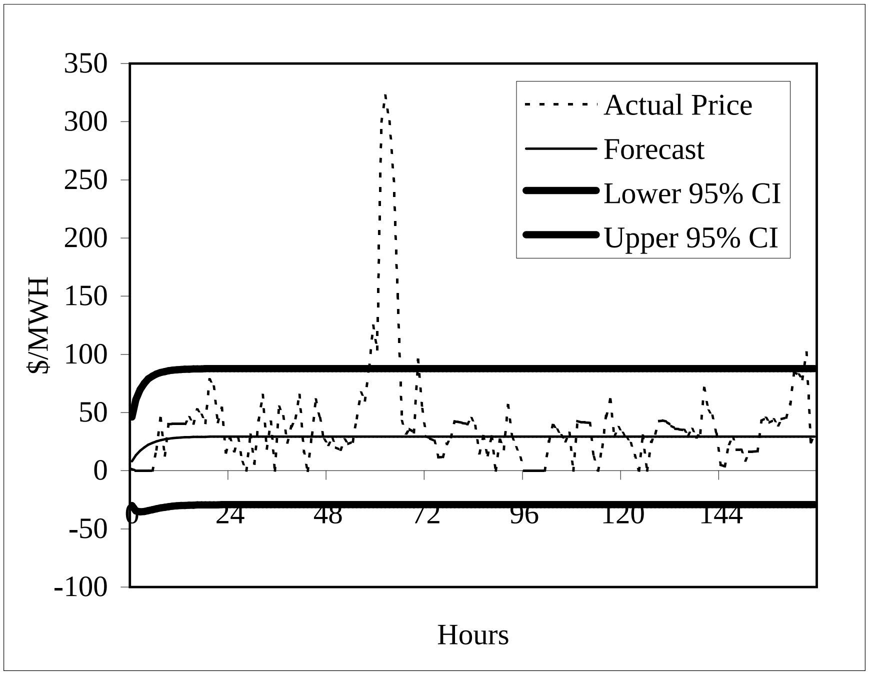

Figure 10: Week-Ahead Forecasts Over the Crisis Period for the Mean Reverting Model (Model 1)

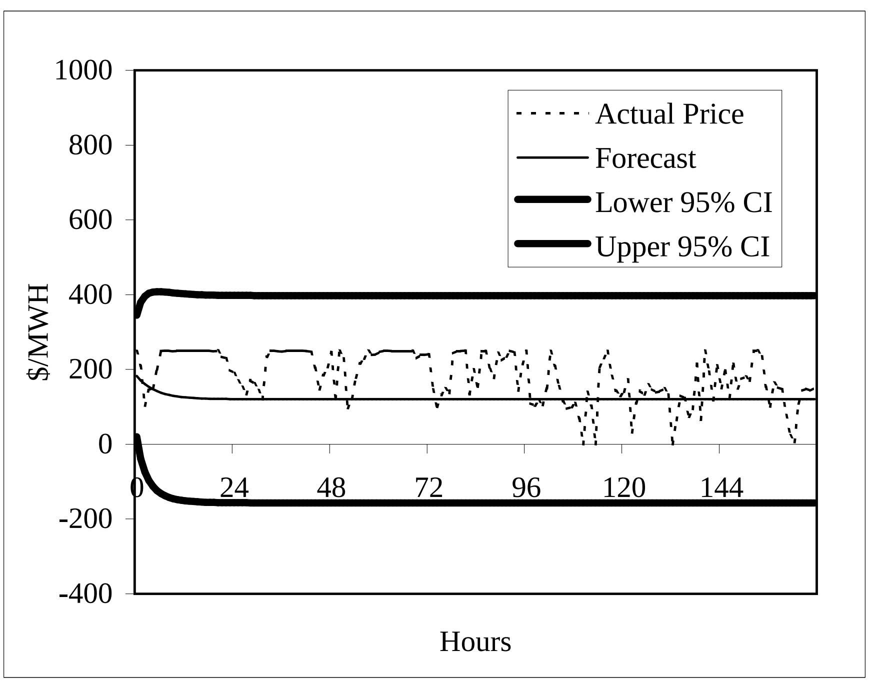

Figure 11: Week-Ahead Forecasts Over the Pre-Crisis Period for The Time-Varying Mean Model (Model 2)

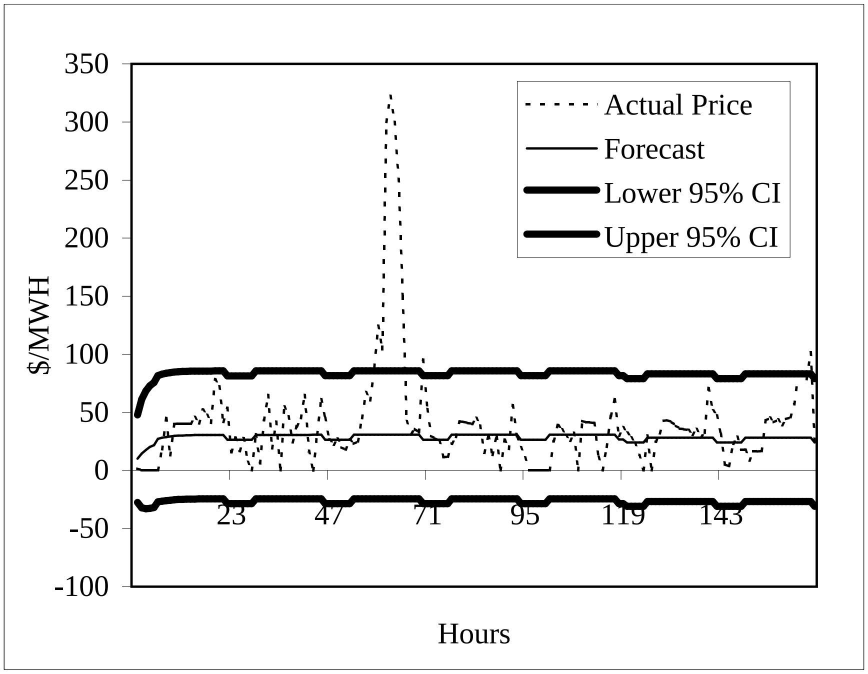

Figure 12: Week-Ahead Forecasts Over the Crisis Period for The Time-Varying Mean Model (Model 2)

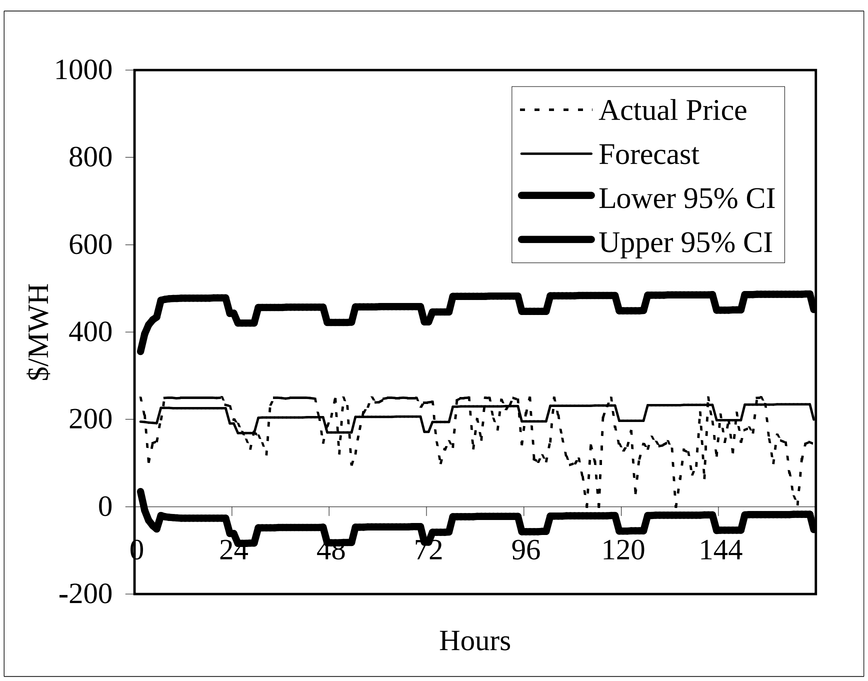

Figure 13: Week-Ahead Forecasts Over the Pre-Crisis Period for the Jump-Diffusion Model (Model 3a)

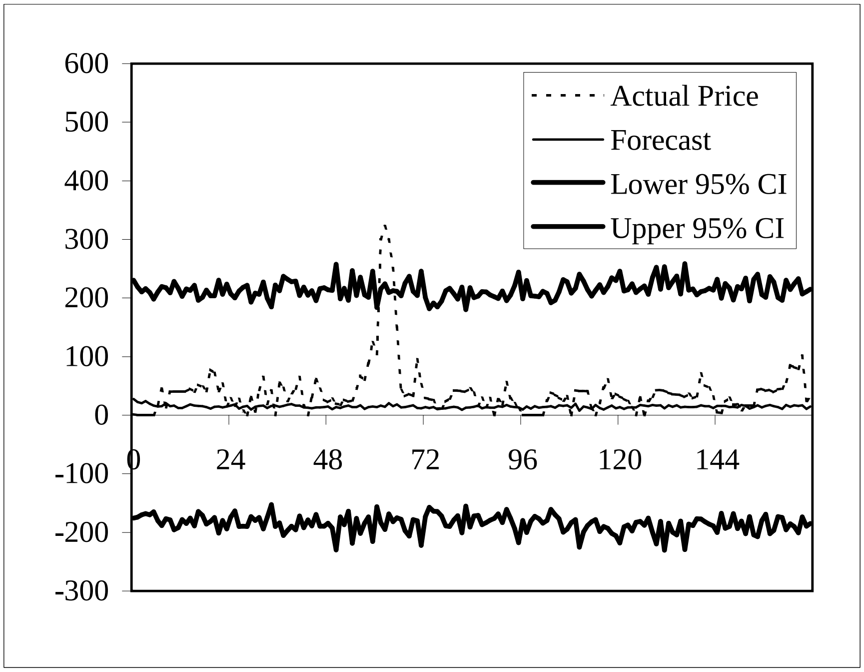

Figure 14: Week-Ahead Forecasts Over the Crisis Period for the Jump-Diffusion Model (Model 3a)

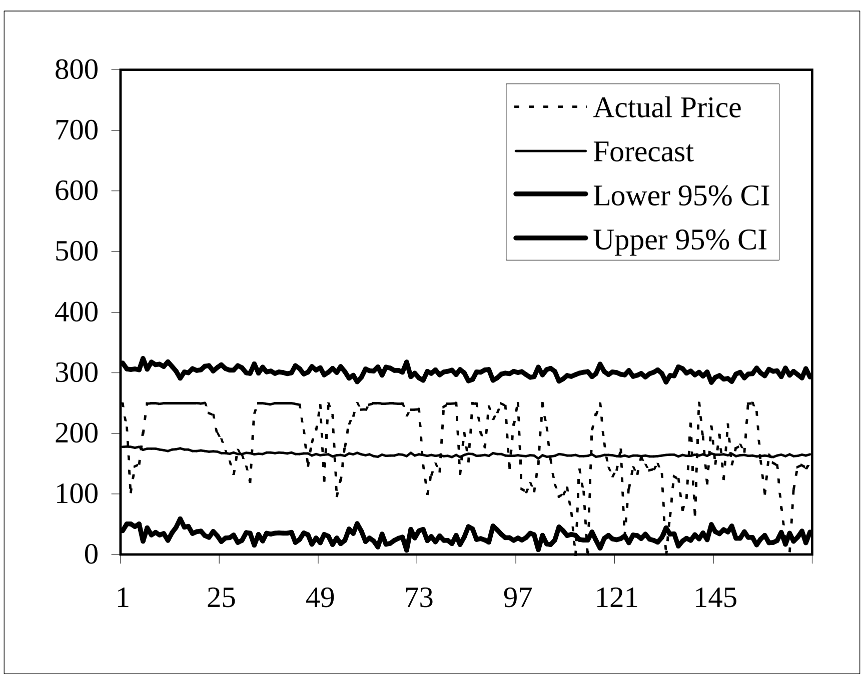

Figure 15: Week-Ahead Forecasts Over the Pre-Crisis Period for the Jump-Diffusion Model (Model 3b)

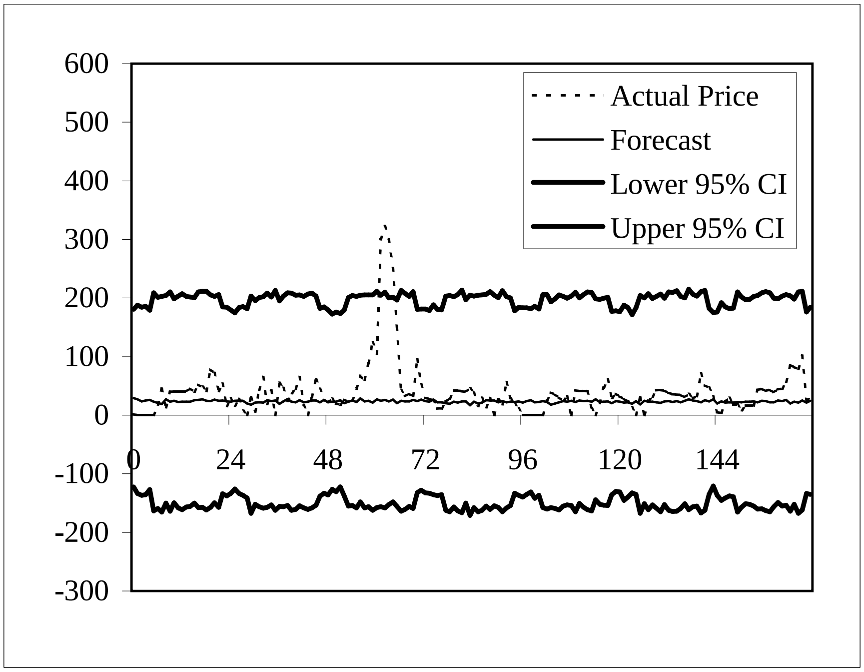

Figure 16: Week-Ahead Forecasts Over the Crisis Period for the Jump-Diffusion Model (Model 3b)

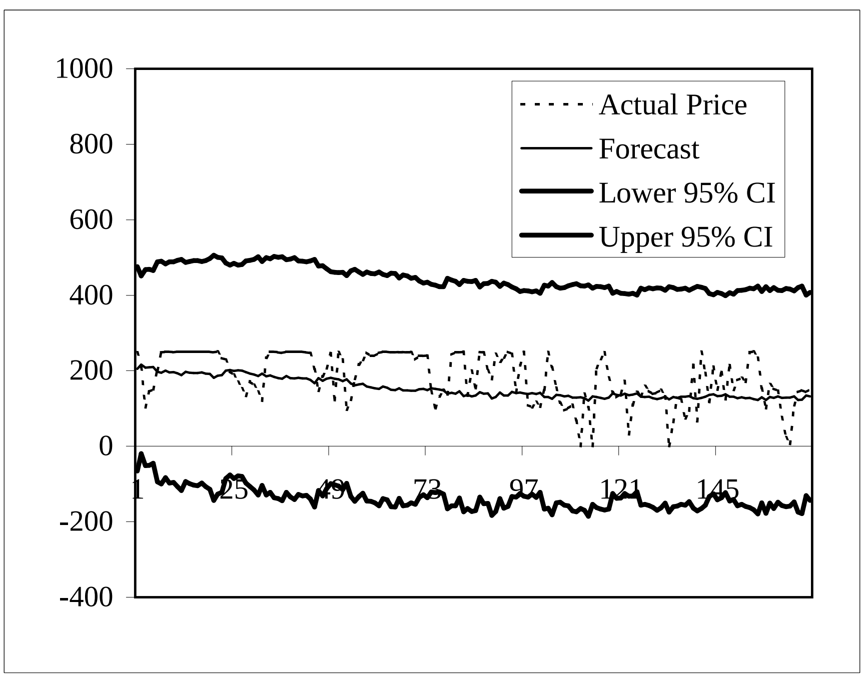

Figure 17: Week-Ahead Forecasts Over the Pre-Crisis Period for the ARMAX Model (Model 4)

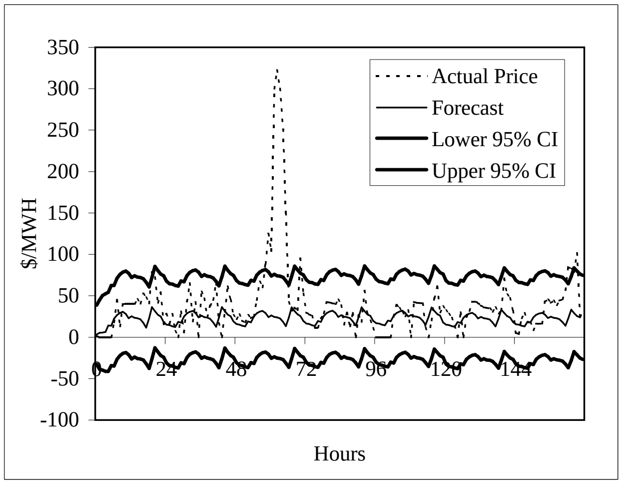

Figure 18: Week-Ahead Forecasts Over the Crisis Period for the ARMAX Model (Model 4)

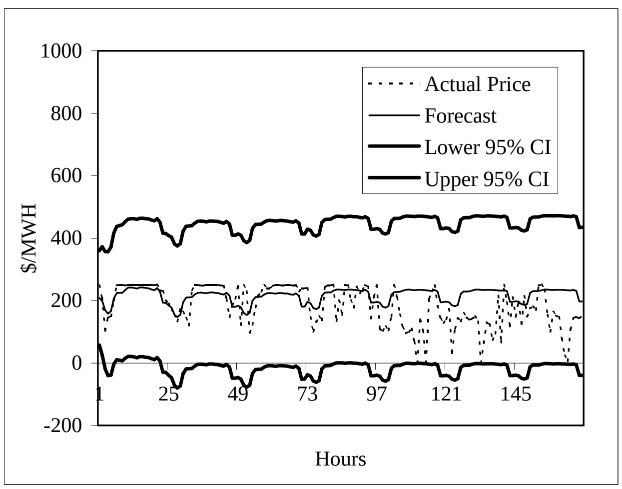

Figure 19: Week-Ahead Forecasts Over the Pre-Crisis Period for the EGARCH Model (Model 5)

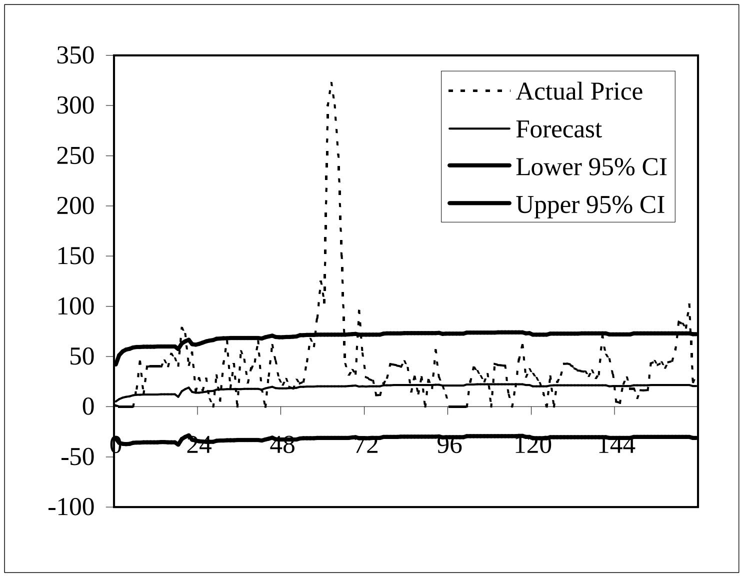

Figure 20: Week-Ahead Forecasts Over the Crisis Period for the EGARCH Model (Model 5)

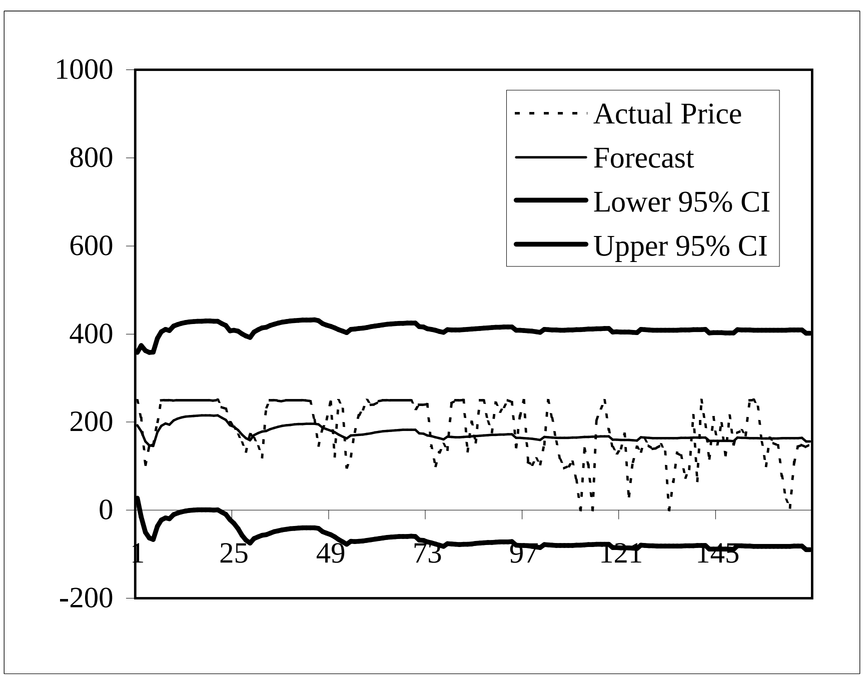

Figure 21: Scatter Plot of Price vs. Temperature (San Francisco)

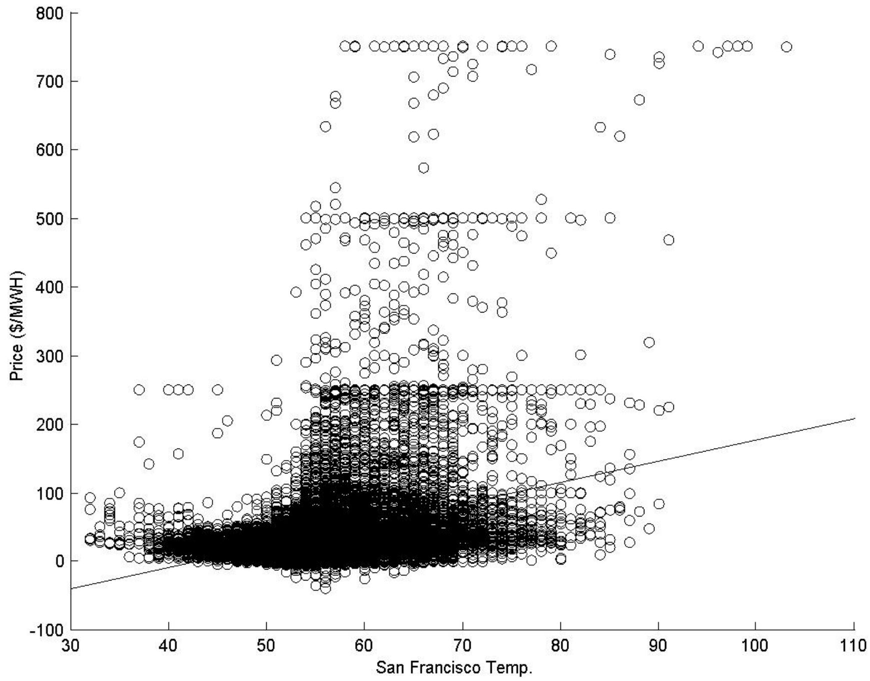

Figure 22: Week-Ahead Forecasts Over the Pre-Crisis Period for the ARMAX Model Incorporating Weather (Model 6)

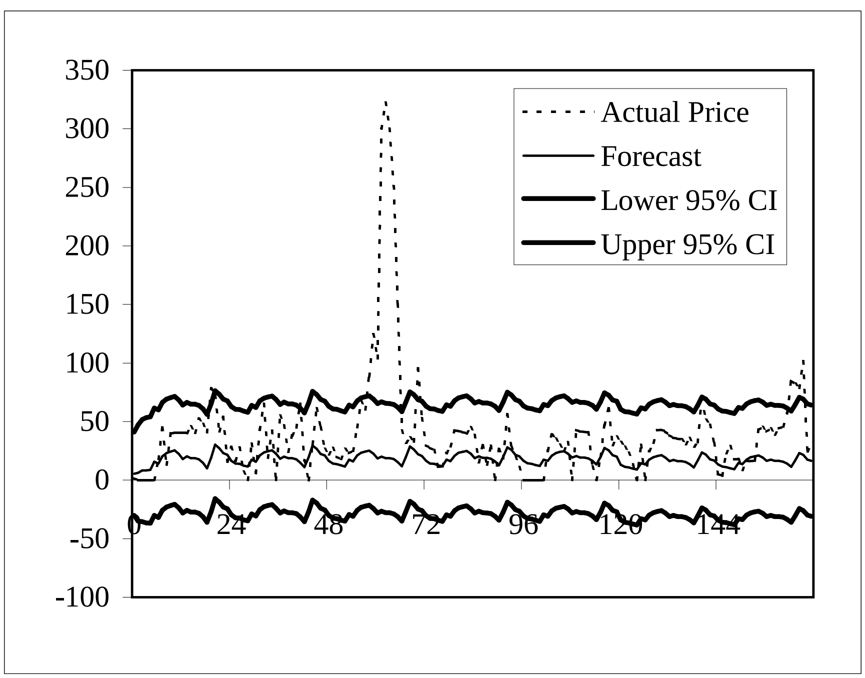

Figure 23: Week-Ahead Forecasts Over the Crisis Period for the ARMAX Model Incorporating Weather (Model 6)

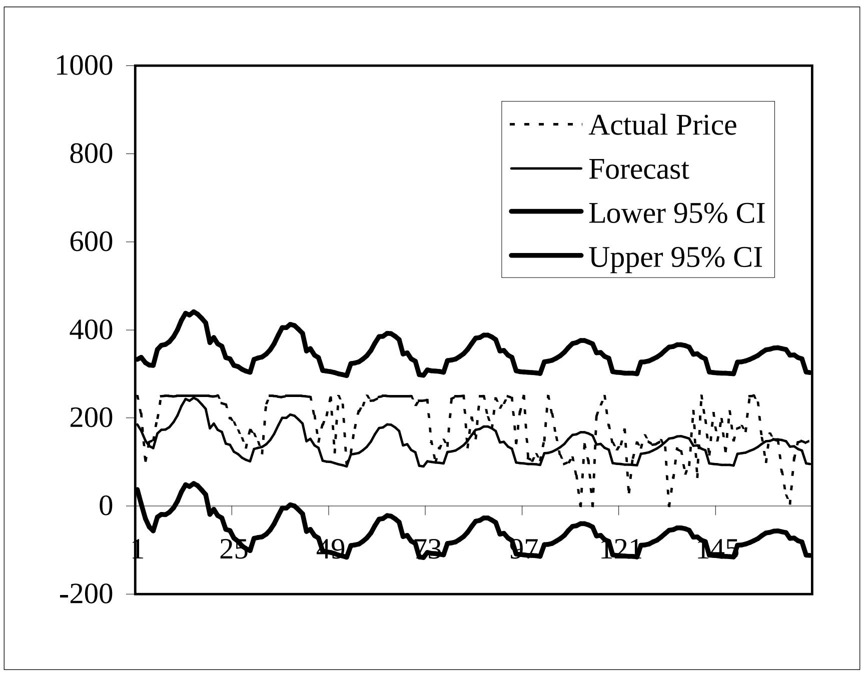

Table 1: Electricity Price Summary Statistics. Descriptive statistics for hourly NP15 (Northern California) electricity spot prices over the pre-crisis and crisis period, which are delineated by May 1, 2001.

Statistic	Pre-Crisis	Pre-Crisis May-Aug Months	Crisis
Mean	29.32	120.29	23.93
Minimum	−249.00	−325.60	−249.00
Maximum	492.20	750.00	250.00
Standard Deviation	29.85	141.6	34.19
Skewness	4.79	2.54	4.25
Kurtosis	35.70	7.197	25.72

Table 2: Pre-Crisis Period Parameter Estimates of Models 1 Through 4. The pre-crisis estimation period is 4/1/1998 - 4/23/2000. Standard errors are in parentheses.

Parameter	Model 1	Model 2	Model 3a	Model 3b	Model 4
α_0 (Mean)	29.32 (0.6010)				
α_1 (Peak Mean)		36.90 1.1140	23.86 0.5098	25.51 0.5524	28.18 3.201
α_2 (Off-Peak Mean)		32.7326 1.1604	21.7403 (0.5291)	23.32 0.5673	26.42 (3.314)
α_3 (Weekend Effect)		-2.1663 (0.9323)	0.1473 (0.4094)	0.1995 (0.3995)	-2.5071 (0.8114)
α_4 (Fall Effect)		4.6192 (1.5191)	5.0310 (0.7524)	5.1734 (0.7472)	7.4213 (03.720)
α_5 (Winter Effect)		-9.9530 (1.5537)	-1.7448 (0.7247)	-1.3876 (0.7048)	1.5572 (3.9731)
α_6 (Spring Effect)		-16.72 (1.4778)	-9.838 (0.7130)	-9.5840 (0.6962)	-0.9111 (3.444)
β_1 (AR 1)	0.7602 (0.0048)	0.7650 (0.0070)	0.7748 (0.0055)	0.7648 (0.0053)	0.7323 (0.0051)
β_2 (AR 24)					0.9592 (0.0039)
β_3 (AR 25)					-0.7323 (0.0079
δ_1 (MA 1)					0.1574 (0.0106)
δ_2 (MA 24)					0.8744 (0.0066)
δ_3 (MA 25)					0.1656 (0.0101
λ_0 (Jump Probability)			0.0327 (0.0092)	0.1396 (0.0091)	
λ_{fall} (Fall Effect)				0.0070 (0.0100)	
λ_{win} (Winter Effect)				-0.1133 (0.0085)	
λ_{spr} (Spring Effect)				-0.0954 (0.0085)	
$\lambda_{weekend}$ (Weekend Effect)				-0.0213 (0.0049)	
λ_{peak} (Peak Effect)				0.0489 (0.0053)	
μ_v (Jump Mean)			11.06 (1.1794)	12.45 (1.2450)	
ξ (Volatility Multiplier)			62.2042 (2.0934)	61.8939 (2.1492)	
σ (Volatility)	19.4037 (0.1019)	19.2475 (0.1011)	6.4533 (0.0793)	6.6365 (0.0732)	18.2501 (0.0535)
RMS Forecast Error	47.5088	47.1523	49.4037	44.3501	25.4839

Table 3: Pre-Crisis Period Parameter Estimates of Model 5. The pre-crisis estimation period is 4/1/1998 - 4/23/2000. Standard errors are in parentheses.

Parameter	**Model 5**
α_1 (Peak Mean)	29.2703
	(0.0047)
α_2 (Off-Peak Mean)	25.6694
	(0.0125)
α_3 (Weekend Effect)	-2.8401
	(0.0018)
α_4 (Fall Effect)	14.5210
	(0.0172)
α_5 (Winter Effect)	0.1660
	(0.0121)
α_6 (Spring Effect)	-6.2021
	(0.0398)
β_1 (AR 1)	0.6320
	(0.0068)
β_2 (AR 24)	0.2293
	(0.0070)
β_3 (AR25)	-0.0750
	(0.0065)
θ (ARCH Intercept)	4.9984
	(0.0020)
κ_1 (ARCH 1)	1.5042
	(0.0038)
γ_1 (GARCH 1)	0.2553
	(0.0034)
γ_{24} (GARCH 24)	-.0750
	(0.0027)
γ_{25} (GARCH 25)	0.0069
	(0.0040)
ψ (Asymmetry)	0.0034
	(0.0003)
RMS Forecast Error	52.1893

Table 4: Pre-Crisis Period Parameter Estimates of Model 6. The pre-crisis estimation period is 4/1/1998 - 4/23/2000. Standard errors are in parentheses.

Parameter	**Model 5**
α_1 (Peak Mean)	29.2703
	(0.0047)
α_2 (Off-Peak Mean)	25.6694
	(0.0125)
α_3 (Weekend Effect)	-2.8401
	(0.0018)
α_4 (Fall Effect)	14.5210
	(0.0172)
α_5 (Winter Effect)	0.1660
	(0.0121)
α_6 (Spring Effect)	-6.2021
	(0.0398)
α_7 (Temperature)	10.0729
	(1.6111)
α_8 (Temperature^2)	-0.2220
	(0.0266)
α_9 (Temperature^3)	0.0016
	(0.0001)
β_1 (AR 1)	0.71603
	(0.0084)
β_2 (AR 24)	0.9192
	(0.0071)
β_3 (AR 25)	-0.6530
	(0.0102)
δ_1 (MA 1)	0.1459
	(0.0118)
δ_1 (MA24)	0.8145
	(0.0104)
δ_1 (MA25)	-0.1302
	(0.0112)
σ (Error SD)	18.2067
	(0.0532)
RMS Forecast Error	23.4787

Table 5: Crisis Period Parameter Estimates of Models 1 Through 4. The crisis estimation period is 5/1/2000 - 8/24/2000. Standard errors are in parentheses.

Parameter	Model 1	Model 2	Model 3a	Model 3b	Model 4
α_0 (Mean)	120.5084				
	(8.2331)				
α_1 (Peak Mean)		159.9843	173.6223	173.4912	154.2709
		(9.2713)	(22.8593)	(22.8502)	(17.2948)
α_2 (Off-Peak Mean)		125.1324	164.9730	164.8100	121.3176
		(9.9646)	(22.6932)	(23.0237)	(18.2001)
α_3 (Weekend Effect)		-21.4966	-5.631	-5.9930	-11.8851
		(11.5998)	(6.5710)	(6.1973)	(10.0701)
α_6 (Spring Effect)		-86.8142	-98.3090	-98.8955	-71.9011
		(16.2252)	(22.9144)	(22.7052)	(28.7389)
β_1 (AR 1)	0.8091	0.7853	0.9617	0.9615	0.8562
	(0.0111)	(0.0118)	(0.0056)	(0.0052)	(0.0133)
β_2 (AR 24)					0.8028
					(0.0333)
β_3 (AR 25)					-0.6917
					(0.0336)
δ_1 (MA 1)					0.2899
					(0.0237)
δ_1 (MA24)					0.6687
					(0.0410)
δ_1 (MA25)					-0.2975
					(0.0238)
λ_0 (Jump Probability)			0.0530	0.2527	
			(0.0032)	(0.0252)	
λ_{spr} (Fall Effect)				-0.1301	
				(0.0197)	
$\lambda_{weekend}$ (Weekend Effect)				-0.0895	
				(0.0196)	
λ_{peak} (Peak Effect)				0.0564	
				(0.0194)	
μ_v (Jump Mean)			0.6587	0.7563	
			(8.2091)	(8.0784)	
ξ (Volatility Multiplier)			44.95	44.86	
				(3.3039)	(3.2811)
σ (Volatility)	83.2143	82.2609	27.6667	27.4639	77.1476
	(1.1204)	(1.1072)	(0.8602)	(0.8644)	(0.0657)
RMS Forecast Error	88.5649	76.1083	73.3400	83.3012	66.6337

Table 6: Crisis Period Parameter Estimates of Model 5. The crisis estimation period is 5/1/2000 - 8/24/2000. Standard errors are in parentheses.

Parameter	**Model 5**
α_1 (Peak Mean)	139.0722
	(0.1666)
α_2 (Off-Peak Mean)	32.6045
	(1.1498)
α_3 (Weekend)	4.4532
	(2.5943)
α_6 (Spring Effect)	-75.3778
	(0.2846)
β_1 (AR 1)	0.6935
	(0.0156)
β_2 (AR 24)	0.2362
	(0.0144)
β_3 (AR 25)	-0.0100
	(0.0223)
θ (ARCH Intercept)	0.7700
	(0.0743)
κ_1 (ARCH 1)	0.7434
	(0.0378)
γ_1 (GARCH 1)	0.7558
	(0.0172)
γ_2 (GARCH 24)	0.2726
	(0.0334)
γ_3 (GARCH 25)	-0.1113
	(0.0332)
ψ (Asymmetry)	0.3452
	(0.0311)
RMS Forecast Error	61.3005

Table 7: Crisis Period Parameter Estimates of Model 6. The crisis estimation period is 5/1/2000 - 8/24/2000. Standard errors are in parentheses.

Parameter	**Model 5**
α_1 (Peak Mean)	29.2703
	(0.0047)
α_2 (Off-Peak Mean)	25.6694
	(0.0125)
α_3 (Weekend Effect)	-2.8401
	(0.0018)
α_4 (Fall Effect)	14.5210
	(0.0172)
α_5 (Winter Effect)	0.1660
	(0.0121)
α_6 (Spring Effect)	-6.2021
	(0.0398)
α_7 (Temperature)	-30.7134
	(29.9954)
α_8 (Temperature^2)	0.2422
	(0.4084)
α_9 (Temperature^3)	0.0003
	(0.0019)
β_1 (AR 1)	0.8008
	(0.0180)
β_2 (AR 24)	0.3562
	(0.0885)
β_3 (AR 25)	-0.2531
	(0.0794)
δ_1 (MA 1)	0.2864
	(0.0280)
δ_1 (MA24)	0.2046
	(0.0924)
δ_1 (MA25)	-0.1437
	(0.0395)
σ (Error SD)	75.4821
	(0.0625)
RMS Forecast Error	75.3501