

A Dynamic Supply-Demand Model for Electricity Prices

Manuela Buzoianu*, Anthony E. Brockwell, and Duane J. Seppi

Abstract

We introduce a new model for electricity prices, based on the principle of supply and demand equilibrium. The model includes latent supply and demand curves, which may vary over time, and assumes that observed price/quantity pairs are obtained as the intersection of the two curves, for any particular point in time. Although the model is highly nonlinear, we explain how the particle filter can be used for model parameter estimation, and to carry out residual analysis. We apply the model in a study of Californian wholesale electricity prices over a three-year period including the crisis period during the year 2000. The residuals indicate that inflated prices do not appear to be attributable to natural random variation, temperature effects, natural gas supply effects, or plant stoppages. However, without ruling out other factors, we are unable to argue whether or not market manipulation by suppliers played a role during the crisis period.

1 Introduction

The energy industry provides electrical power to consumers from a variety of sources, including gas-based and hydroelectric plants, as well as nuclear and coal-based power plants. The price of electricity in any given area is governed by complex dynamics, which are driven by many factors, including, among other things, day-to-day and seasonal variation in demand, seasonal variation in temperature, availability of electricity from surrounding regions, and cascade effects when plants are shut down. In addition, major changes can be introduced for legal reasons, for instance, when deregulation takes place, or when emission permits are issued, etc.

A number of studies have analyzed electricity price time series and market dynamic. Knittel and Roberts (2001) and Deng (2000) attempt to build models that can capture the volatile behavior of the prices as well as their high jumps. In particular, Knittel and Roberts (2001) emphasize that market models can be improved when the prices are specified in relation with exogenous information. In addition, Joskow and Kahn give a nice analysis of the economics of the California market during the crisis of 2000. They explain market behavior and its prices by examining a number of factors involved in supply and demand, but do not explicitly pose time series models for the data.

*Address: Department of Statistics, Carnegie Mellon University, Pittsburgh, PA 15213. E-mail: manuela@stat.cmu.edu.

In this paper, we examine the use of a new model, which we refer to as a *dynamic supply-demand model*, to simultaneously capture electricity price and usage time series. This model is built on the basic economic principle that on each day, the price and quantity in a competitive market can be determined as the intersection of supply and demand curves. The model incorporates temperature and seasonality effects and gas-availability as factors by expressing the supply and demand curves as explicit functions of these factors. It also allows for intrinsic random variation of the curves over time. Since the model is nonlinear and non-Gaussian, and the supply and demand curves are not directly observable, traditional methods for parameter estimation and forecasting are not applicable. However, the particle filtering algorithm (see, e.g. Salmond and Gordon, 2001; Kitagawa, 1996; Doucet et al., 2001) can be used in this context.

The model is fit to the California electricity market data from April 1998 to December 2000. During this time the California market was restructured by divesting part of power plants to private firms and allowing competition among producers. But it experienced skyrocketing price values during the second half of 2000, the period often referred to as the “California power crisis”.

The particle filter yields one-step predictive distributions of power prices and quantities. Evaluating predictive cumulative distribution functions at observed data points gives a type of residual which can be used to evaluate goodness-of-fit. Computing these residuals, we find that the model fits reasonably well, except that it cannot explain skyrocketing prices during the crisis period. Moreover, some of the discrepancies between predicted and observed price values during the crisis are inconsistent with those from the pre-crisis period. One could argue that this might be due to the so called “perfect storm” conditions, manifested in the form of a dry climate in the Pacific Northwest that reduced hydroelectric power generation, leading to an increase in natural gas prices, and causing demand to exceed supply. However, since our model accounts for most of these factors, the discrepancies suggest that market structure may have changed during the second half of 2000. This is of particular interest to those who would argue that some production firms manipulated the market by withholding supply to drive profit margins up.

The paper is organized as follows. In Section 2, we provide a general description of supply-demand pricing principles and introduce our dynamic supply-demand model. In Section 3, we apply the model to analysis of the California energy market, developing specific functional forms for supply and demand curves, and using a particle-filtering approach to estimate the curves and obtain residuals. The paper concludes with further brief discussion in Section 4.

2 General Methodology

2.1 Principles of Supply-Demand Pricing

Supply and demand are fundamental concepts of economics. In a competitive market, basic principles of supply and demand determine the price and quantity sold for goods, securities and other tradeable assets. The market for a particular commodity can be regarded

as a collection of entities (individuals or companies) who are willing to buy or sell it. Under the assumption that the market is competitive, continuous interaction between suppliers and demanders establishes a unique price for the commodity (see, e.g. Mankiw, 1998).

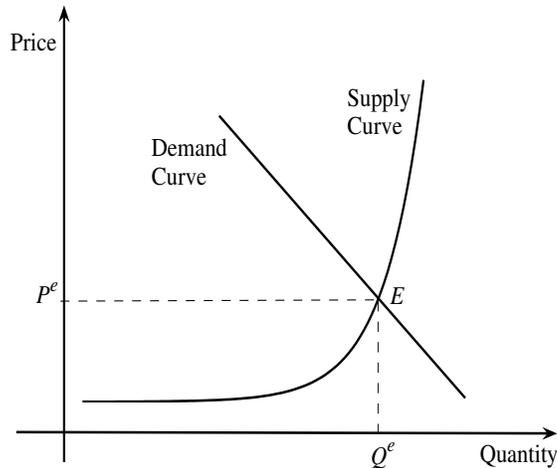


Figure 1: Typical supply and demand curves. $E = (P^e, Q^e)$ is the equilibrium point.

Market supply and demand are established by summing up individual supply and demand schedules. These schedules represent the quantity individuals are willing to trade at any unit price (see Figure 1-a). The supply curve describes the relationship between the unit price and the total quantity offered by producers, and is upward-sloping. In the case of electricity, this upward slope can be explained by the fact that small quantities will be supplied using the most efficient plant available, but that as quantity increases, producers will have to use less efficient plants, with higher production costs. The demand curve describes the relationship between the unit price and the total quantity desired by consumers. It is downward sloping since the higher the price, the less people will want to buy (see, e.g. Parkin, 2003). The intersection of these two curves is the point where supply equals demand, and is called the market equilibrium (Figure 1-a). This is the point where the price balances supply and demand schedules. Under the assumption that the market operates efficiently, observed prices will adjust rapidly to this equilibrium point.

When markets are not perfectly competitive, sellers may be able to exercise market power to maintain prices above equilibrium levels for a significant period of time. In the electricity market, this may be possible, due to an oligopolistic structure (e.g. centralized power pools) in which one or few companies have control over production, pricing, distribution etc. Market power can also be exercised by suppliers by moving to other markets that pay better, causing the current market to become less competitive. One way to model this is to factor the supply curve into two components: a market-free supply component and a “manipulation” component.

2.2 A Dynamic Supply-Demand Model

Supply and demand curves vary over time, due to various changing conditions (e.g. weather, population income) causing equilibrium prices and quantities to fluctuate. For instance, an increase in population income will generate a higher demand for goods. A reduction in production cost will increase the supply of goods. We therefore introduce supply and demand curve processes,

$$S_t(q; E_t, \theta), \quad D_t(q; E_t, \theta). \quad (1)$$

E_t represents the measurement of an exogenous process at time t , θ is a vector of parameters, $S_t(q; E_t, \theta)$ is the unit price producers charge for a quantity q at time t and $D_t(q; E_t, \theta)$ is the unit price consumers are willing to pay for a quantity q at time t .

In general, the supply and demand curves are not directly observable. However, prices and quantities traded are. We therefore regard $S_t(\cdot; \cdot, \cdot)$ and $D_t(\cdot; \cdot, \cdot)$ as latent random processes, and assume that we only directly observe the equilibrium prices and quantities, denoted by P_t and Q_t , respectively. As explained in the previous subsection, P_t and Q_t are assumed to represent the equilibrium point at the intersection of the the supply and demand curves. In other words,

$$\begin{aligned} Q_t &= \{q : S_t(q; E_t, \theta) = D_t(q; E_t, \theta)\} \\ P_t &= D_t(Q_t). \end{aligned} \quad (2)$$

(Note that P_t and Q_t are uniquely defined since the demand curve is monotonically decreasing and the supply curve is monotonically increasing.)

This idealized model assumes that supply and demand curves are deterministic functions of some unknown parameters θ , as well as exogenous explanatory variables E_t , and that observed prices and quantities are exactly the intersections of these curves. In reality we expect to see random fluctuations of the prices and quantities around these predicted values, as well as potential slow change in parameters, over time. We therefore replace θ in the equations above by θ_t , and assume that $\{\theta_t\}$ is a Markovian stochastic process, varying slowly over time.

2.3 Parameter Estimation and Forecasting

In the proposed model, supply and demand curves are not directly observable. We want to estimate them, or equivalently, estimate the sequence of parameter vectors $\{\theta_1, \dots, \theta_T\}$, using observed price and quantity time series, $Y_t = (P_t, Q_t), t = 1, \dots, T$, for some number of observations $T > 0$. Since the model (1,2) is nonlinear and non-Gaussian, standard methods for estimation and prediction are not applicable. However, a particle filtering approach can be used. We can regard $\{\theta_t, t = 1, \dots, T\}$ as a latent Markov process, and $\{Y_t, t = 1, 2, \dots, T\}$ as a sequence of observations, each of which has a (complicated) conditional distribution, given the corresponding underlying value of θ_t .

We are particularly interested in obtaining the so-called *filtering distributions* $p(\theta_t|y_1, \dots, y_t)$, and *predictive distributions* $p(\theta_t|y_1, \dots, y_{t-1})$. These distributions are approximated recursively, using the particle filter. The idea is to generate samples $\{\theta_t^{p,1}, \dots, \theta_t^{p,N}\}$ and $\{\theta_t^{f,1}, \dots, \theta_t^{f,N}\}$, called "particles", which have approximate distributions:

$$\begin{aligned} \{\theta_t^{p,1}, \dots, \theta_t^{p,N}\} &\sim p(\theta_t|y_s; s = 1, \dots, t-1) \\ \{\theta_t^{f,1}, \dots, \theta_t^{f,N}\} &\sim p(\theta_t|y_s; s = 1, \dots, t). \end{aligned}$$

In addition, we will draw (approximate) samples from the predictive distributions $p(y_t|y_s; s = 1, \dots, t-1)$.

$$\{y_t^1, \dots, y_t^N\} \sim p(y_t|y_s; s = 1, \dots, t-1).$$

N here is a "number of particles", which can be used to control approximation errors. Larger values of N give smaller approximation errors, but require more computation time. The samples are obtained using the following procedure.

1. Draw $\theta_0^{f,j} \sim p(\theta_0)$, for $j = 1, \dots, N$
2. Repeat the following steps for $t = 1, \dots, T$
 - (a) For $j = 1, \dots, N$, draw $\theta_t^{p,j}$ from $p(\theta_t|\theta_{t-1} = \theta_{t-1}^{f,j})$.
 - (b) Compute $w_t^j = p(y_t|\theta_t = \theta_t^{p,j})$, for $j=1, \dots, N$, using the observation equation (2). (This requires computation of the supply and demand curves corresponding to each value $\theta_t^{p,j}$, and the intersections of these curves.)
 - (c) Generate $\{\theta_t^{f,j}, j = 1, \dots, N\}$ by drawing a sample, with replacement, of size N , from the set $\{\theta_t^{p,1}, \dots, \theta_t^{p,N}\}$ with probabilities proportional to $\{w_t^j, j = 1, \dots, N\}$.

Step 2(b) requires exogenous information E_t to evaluate the required supply and demand curves. In some cases, when E_t is not available at time t , we replace E_t in this step with a predictor of E_t based on information that is available. We use temperature, for instance, as an exogenous variable, and usually, one day's temperature is a reasonably good predictor of the next day's temperature. Note also that to implement this procedure the prior $p(\theta_0)$ and the transition equation $p(\theta_t|\theta_{t-1})$ need to be specified.

At the end of this procedure, we have estimates of the vectors $\{\theta_t, t = 1, 2, \dots, T\}$, which directly determine our estimates of the supply and demand curves at each time point. As a by-product, we are also able to obtain one-step predictive distributions for the observations, which can be used to compute a "generalized residual", to be used for diagnostic purposes. Since the model is nonlinear and non-Gaussian, there is no standard definition of the residuals, but we can define them using the predictive distribution functions of price and quantity variables. Specifically, let $F_t^P(y)$ denote the one-step empirical predictive cumulative distribution function (cdf) for the price P_t , and let $F_t^Q(y)$ be the one-step empirical predictive cdf for the quantity Q_t . Then if the model is correct, $\{F_t^P(P_t), t = 1, \dots, T\}$ and $\{F_t^Q(Q_t), t = 1, \dots, T\}$ will form sequences of independent and identically distributed

Uniform(0,1) random variables (this follows from results which can be found, among other places, in Rosenblatt, 1952). If desired, these can be converted to Gaussian random variables by applying the inverse standard normal cdf. Thus standard tests can be applied to the residuals to check for goodness-of-fit.

3 Analysis of Californian Electricity Prices

In this section we apply the techniques described in the previous section in an analysis of Californian electricity prices during the period from April 1st, 1998, to December 31st, 2000.

The California energy market underwent a major structural change in April 1998, when the electricity generation component was opened to competition with the intent of reducing consumer prices. The wholesale market, called the “Power Exchange”, was organized as a daily auction. Each day generators and utility companies bid their price and quantity of energy, and the Power Exchange calculated a uniform price/quantity pair for all its participants. Utilities were required to satisfy the demand of end-consumers. They also had to sell power from the stations they owned or controlled through the Power Exchange. A major crisis occurred in late 2000, when the utilities were required to purchase power at extremely high prices and sell to end-users at a substantially lower price.

3.1 Data

During 1998-2000, the California Power Exchange operated a day-ahead hour-by-hour auction market. This market ran each day to determine hourly wholesale electricity prices for the next day. By 7 a.m. each day, generators and retailers submitted a separate schedule of price-quantity pairs for each hour of the following day. The Power Exchange assembled these schedules into demand and supply curves and established the equilibrium point by finding the intersections of these curves. We focus on this (day-ahead) market since it contains the majority of trades. Hourly day-ahead market data are available from the University of California Energy Institute (<http://www.ucei.berkeley.edu/>).

Plots of prices and quantities are given in Figures 2 and 3. They exhibit annual, weekly, and daily seasonal components. The biggest volume of trades of each day are registered during the high-demand period from 9 a.m. to 5 p.m. To avoid the need to deal with the daily cycle, we consider the daily average over the high-demand period. Figure 4 shows the resulting daily prices and quantities. (We will be particularly interested in the large jumps are registered near the end of the year 2000.)

In order to explain prices and quantities, we will use a number of exogenous variables, incorporating them into a particular functional form for the supply and demand curves $S_t(\cdot; \cdot, \cdot)$ and $D_t(\cdot; \cdot, \cdot)$. These include temperature, natural gas price, quantity of energy produced from natural gas, and total production of energy. Since daily temperatures are available for each California county, we aggregated them to state-wide temperature index by taking a weighted average based on population weights. Maximum daily temperatures are

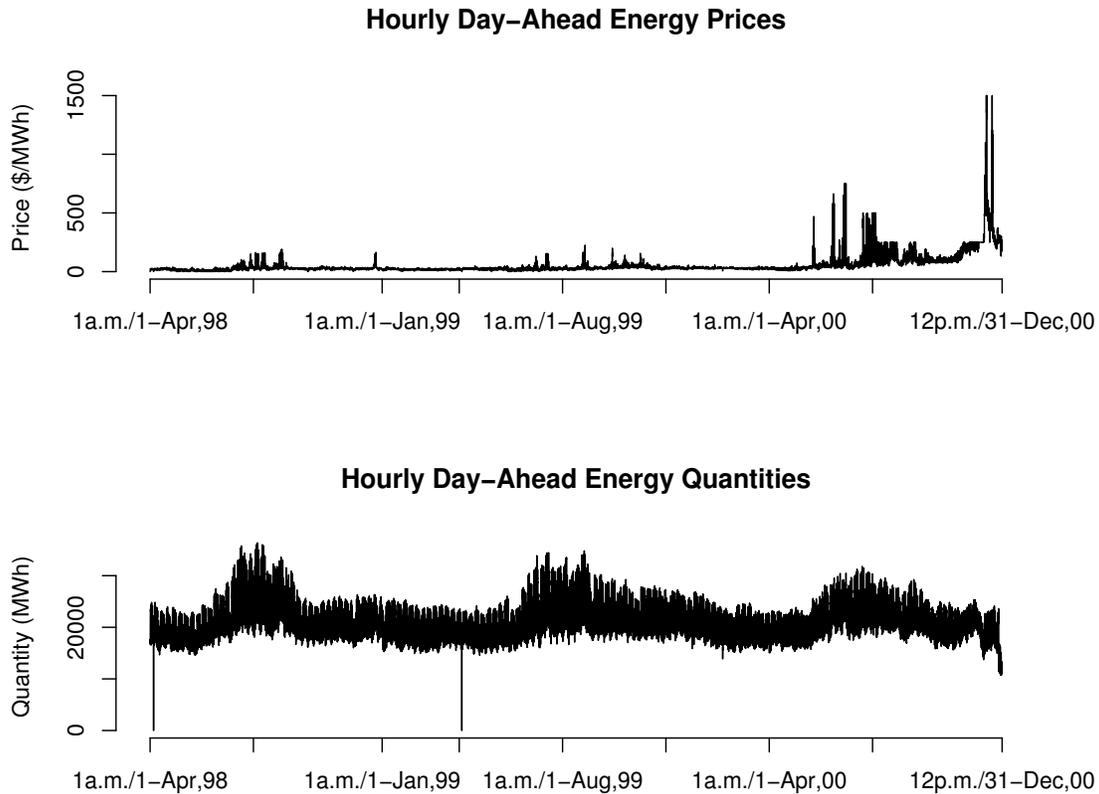


Figure 2: Hourly prices and quantities in the day-ahead market during April 1, 1998 - December 31, 2000

shown in Figure 5. Natural gas was also an important source of fuel for electricity generation plants in California in 1998-2000, and its availability (or lack thereof) had a significant effect on supply. Power generation based on natural gas accounted for about 75% of the power production private sector, and for about 35% of the total power production. Generators had to buy natural gas on the gas market, so gas prices had a direct impact on the power supply. Natural gas daily prices and the hourly gas based energy quantities are also publicly available, The latter (reduced to daily on-peak data) are shown in Figure 5. Note that natural gas prices increased in late 2000, while energy production from natural gas also increased (also Figure 5).

3.2 Specific Form of Supply-Demand Curves

In order to carry out our analysis, we need to determine functional forms for the supply and demand curves, as a function of parameter vectors θ_t and also of exogenous variables E_t .

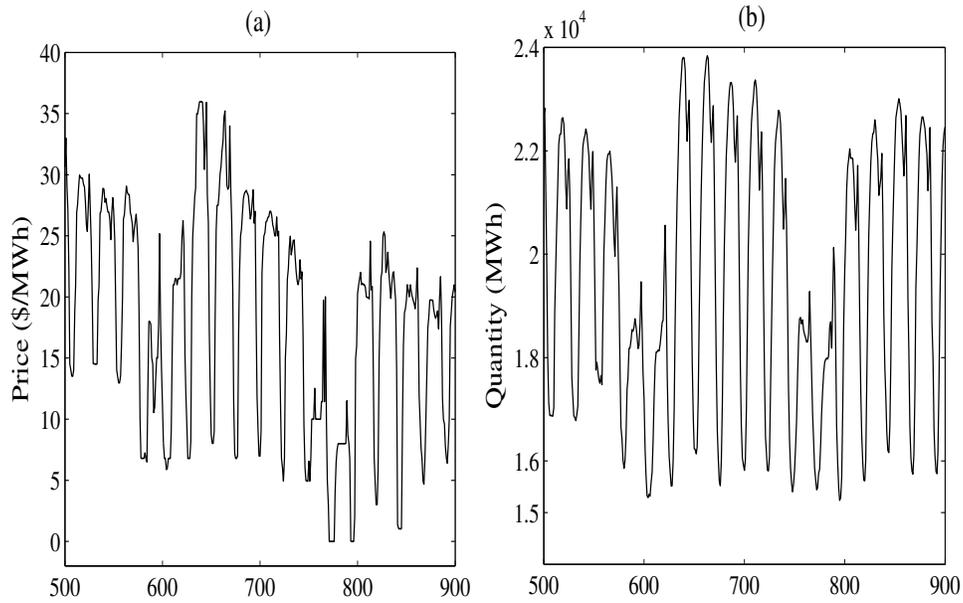


Figure 3: A 400-hour window of electricity price and quantity observations.

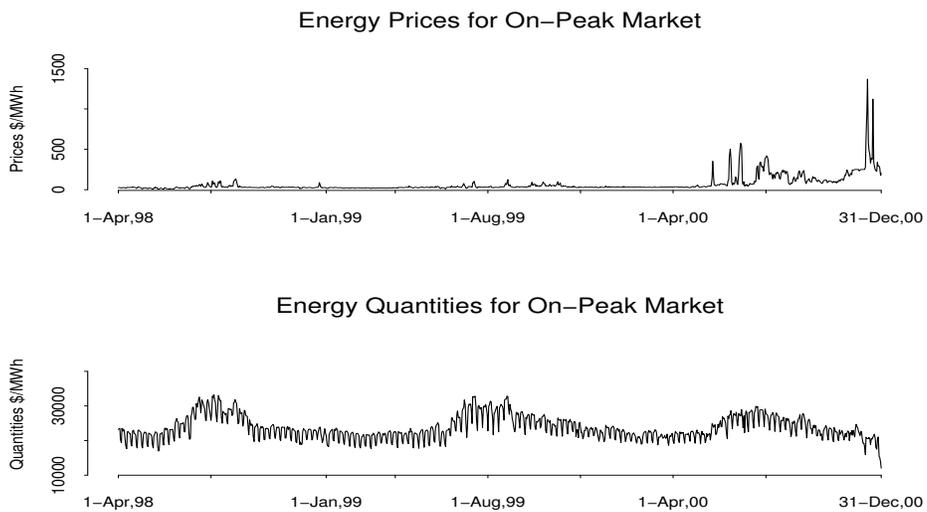


Figure 4: On-peak prices and energy quantities in the day-ahead market during 1 April,1998 - 31 December, 2000

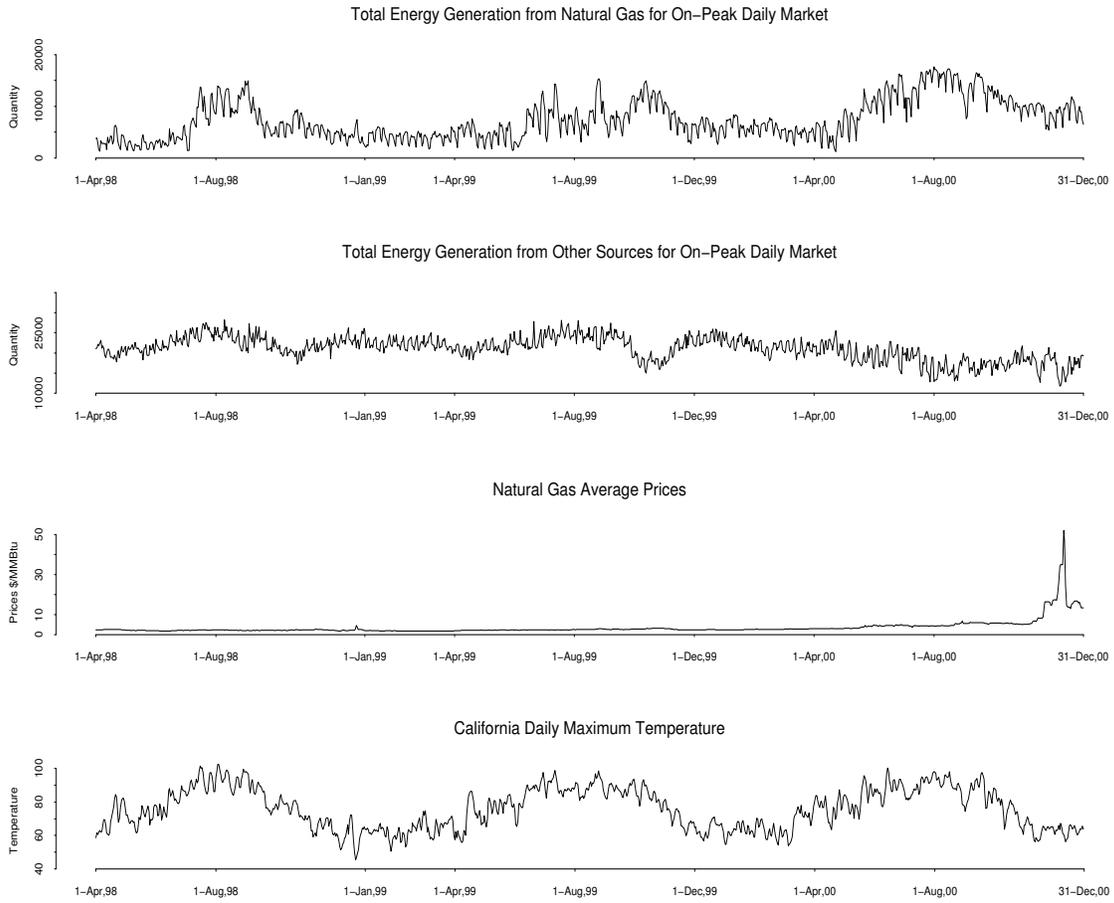


Figure 5: Total energy generated from natural gas; total energy generated from other sources; natural gas prices; California maximum temperature (April 1,1998 - 31 December,2000)

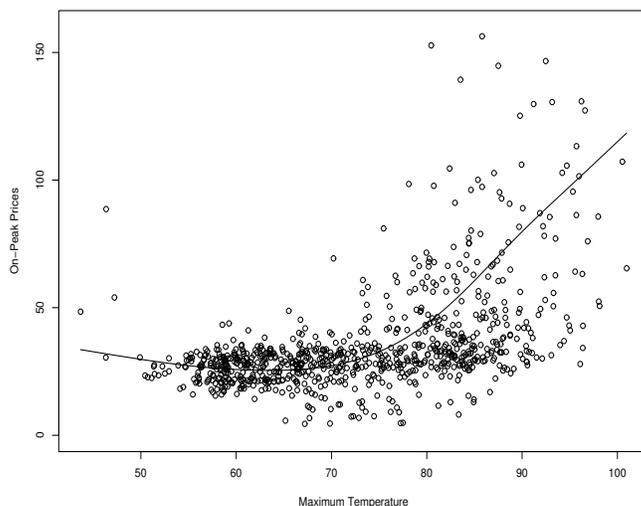


Figure 6: Daily electricity prices versus temperature.

As is apparent in Figure 6, temperature has a noticeable effect on electricity prices. A standard explanation for this is that demand increases in summer, as consumers turn on their air conditioners. Thus, higher temperatures will shift the demand curve to the right. The relationship between observed prices and temperatures appears to be roughly quadratic. The prices appear lowest when the temperature is around $65^{\circ}F$, and highest when the temperature rises to around $100^{\circ}F$. Therefore, we incorporate a term of the form $(T_t - 65)^2$, where T_t denotes temperature at time t , into the demand curve $D_t(\cdot; \cdot, \cdot)$.

Another driver of demand is the weekly seasonal component. The demand is high during week-days and low during week-ends. We deal with this is by defining a deterministic seasonal component SE_t , which is 1 for week-days and 0 for week-ends and holidays.

Natural gas prices also affect electricity generation. An increase in natural gas prices often leads to an increase in production cost of electricity, and consequently, a decrease in supply. This situation is depicted in Figure 7, where the supply curve shifts left and the equilibrium point moves from (P_1, Q_1) to (P_2, Q_2) , showing that the quantity decreases and the price increases. In this figure the cutoff point QH_t is the energy produced from all sources except natural gas. Above this threshold the total quantity of power is calculated by adding the production from gas based power plants. Let GP_t be the natural gas price at time t . In order to account for supply curve shifts due to gas availability and cost factors, we introduce a term $(q - QH_t)I_{q > QH_t}GP_t$ into the supply curve $S_t(\cdot; \cdot, \cdot)$.

Another change in production occurs when one or more gas power plants go off-line in a particular day due to maintenance requirements. Then the total generation from natural gas decreases. In this situation, the supply curve shifts to the left because producing firms increase the energy price in order to cover the cost for restarting the power plants the next day. The power-plants are also subject to maintenance requirements after a period of ceaseless activity. We handle this shift by including in the supply curve a term which

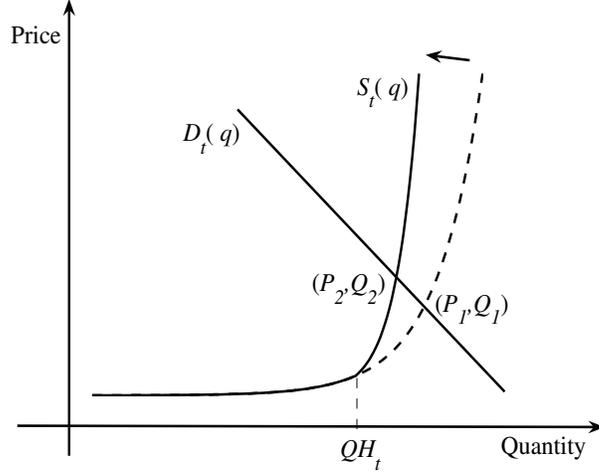


Figure 7: A shift in the supply curve due to changes in production.

represents the total production in one, two or three months previous to the current time t , that is, $\sum_{t-30(60,90) \leq \tau < t} GQ_\tau$.

We also assume that, apart from dependence on exogenous variables, the demand curve $D_t(q; \cdot, \cdot)$, is a linear function of q , and the supply curve for sellers, $S_t(q; \cdot, \cdot)$, is a piecewise exponential function of the supplied quantity q .

Our final expressions for supply and demand curves, including all the aforementioned components, are

$$S_t(q) = \alpha_{0,t} \exp \left(\alpha_{1,t}q + \alpha_{2,t}(q - QH_t)I_{q > QH_t}GP_t + \alpha_{3,t} \sum_{t-30(60,90) \leq \tau < t} GQ_\tau \right) \quad (3)$$

$$D_t(q) = \beta_{0,t} + \beta_{1,t}(T_t - 65)^2 + \beta_{2,t}q + \beta_{3,t}SE_t + R_t, \quad (4)$$

where α_j, t and β_j, t are coefficients to be estimated, QH_t represents the total amount of energy supplied from non-natural gas sources, GP_t denotes natural gas price, GQ_t represents the total quantity of electricity generated from natural gas, T_t is the daily maximum state temperature, and SE_t is the weekend/holiday indicator described above. The process $\{R_t, t = 1, 2, \dots, T\}$ is a first-order autoregressive process satisfying

$$R_{t+1} = \phi_t R_t + \epsilon_{t+1}^R,$$

where $\{\epsilon_t^R\}$ is an iid sequence of normal random variables with mean 0 and variance σ_R^2 . This process allows for serially-correlated random variation of the demand curve around its mean.

The parameters in the resulting model are lumped together into the vector

$$\theta_t = (\alpha_{0,t}, \alpha_{1,t}, \alpha_{2,t}, \alpha_{3,t}, \beta_{0,t}, \beta_{1,t}, \beta_{2,t}, \beta_{3,t}, \phi_t)^T.$$

Since the supply curve is always upward-sloping we impose constraints $\alpha_0 > 0$ and $\alpha_1 > 0$. Also the parameters of the shift terms in the supply curve model, α_2, α_3 , should be positive so that we obtain a shift to the left. Since the demand curve is downward-sloping, we constrain $\beta_2 < 0$. Furthermore, to ensure that high temperatures lead to an increase in demand, we require that $\beta_1 > 0$. Finally, we impose the constraint $|\phi| < 1$, so that the autoregressive process is stationary.

The state of our model at each time t includes the value of the autoregressive process R_t and the parameter vector θ_t . In order to account for the parameter constraints and time dependence, we use the state transition equation

$$\begin{bmatrix} R_t \\ \log(\alpha_{0,t}) \\ \log(\alpha_{1,t}) \\ \log(\alpha_{2,t}) \\ \log(\alpha_{3,t}) \\ \beta_{0,t} \\ \log(\beta_{1,t}) \\ \log(-\beta_{2,t}) \\ \beta_{3,t} \\ f(\phi_t) \end{bmatrix} = \begin{bmatrix} \phi_{t-1}R_{t-1} \\ \log(\alpha_{0,t-1}) \\ \log(\alpha_{1,t-1}) \\ \log(\alpha_{2,t-1}) \\ \log(\alpha_{3,t-1}) \\ \beta_{0,t-1} \\ \log(\beta_{1,t-1}) \\ \log(-\beta_{2,t-1}) \\ \beta_{3,t-1} \\ f(\phi_{t-1}) \end{bmatrix} + \begin{bmatrix} \epsilon_t^R \\ u_{0,t} \\ u_{1,t} \\ u_{2,t} \\ u_{3,t} \\ u_{4,t} \\ u_{5,t} \\ u_{6,t} \\ u_{7,t} \\ u_{8,t} \end{bmatrix}, \quad (5)$$

where $u_{i,t} \sim N(0, \xi_i^2)$, $i = 0, \dots, 8$ are independent Gaussian white noise processes, and $f(x) = \log(\frac{1-x}{1+x})$. The log transformations, and the function f , are used to ensure that parameters stay within their allowable ranges.

In order to use the particle filter to analyze our data, we need to make a minor adjustment to the observation model. We use equation (2), but we assume that instead of observing (P_t, Q_t) directly, we observe

$$\begin{aligned} P_t^* &= P_t + \epsilon_t^P \\ Q_t^* &= Q_t + \epsilon_t^Q, \end{aligned} \quad (6)$$

where $\{\epsilon_t^P\}$ and $\{\epsilon_t^Q\}$ are iid sequences of Gaussian random variables with mean zero and variances chosen to be very small. This ensures that the weights w_t^j computed in Step 2(b) of the particle filtering algorithm are always positive. Without this adjustment, the weights are all equal to zero with probability one, and the particle filter is rendered unusable (this is a well-known problem with particle filtering, discussed in more detail in several papers in Doucet et al., 2001).

3.3 Results

We employed the particle filtering algorithm described in Section 2.3, with $N = 30000$ particles at each point in time, along with the model (1,2,6), in order to find one-step predictive distributions for price-quantity pairs, along with estimates of the supply and demand curves at each point in time.

To initialize the filter, starting values were required for the state vector θ_0 . These were obtained using crude estimation procedures based on analysis of one month of data from April 1998, and are given in Table 1 below. Standard deviations of the white noise components in equation (5) were taken to be $\xi_0 = 0.02, \xi_1 = 0.012, \xi_2 = 0.035, \xi_3 = 0.013, \xi_4 = 1, \xi_5 = 0.04, \xi_6 = 0.08, \xi_7 = 0.5, \xi_8 = 0.08$. The other required parameters were $\sigma^R = 0.001, \sigma_P = 1.8, \sigma_Q = 0.9$.

α_0	α_1	α_2	α_3	β_0	β_1	β_2	β_3	ϕ
5.412	0.000043	0.00001388	0.00000395	424.277	0.266	-0.02	86.675	0.7824

Table 1: Initial values θ_0 .

The 0.05 and 0.95 quantiles of the one-step predictive distributions for prices and quantities for the data are shown in Figures 8 - 9, along with the observed data. The observed values lie mostly within the 0.05-0.95 quantile range, except during the 2000 power crisis. Furthermore, the difference between the real and predicted (mean of predictive distribution) prices is significantly higher during the second half of 2000 than in the pre-crisis period (Figure 10).

Three different quantities measuring the energy cost differences are used as additional measures of the difference between the crisis and pre-crisis periods. These are $\sum_t Q_t(P_t - \bar{P}_{pc})$, representing the difference between total real energy cost and the cost based on the pre-crisis price average, $\sum_t Q_t(\hat{P}_t - \bar{P}_{pc})$, the difference between total predicted cost and cost based on the pre-crisis price average, and $\sum_t Q_t(P_t - \hat{P}_t)$, the difference between the total real and predicted costs. Table 2 gives these quantities, as calculated using (a) the entire data set and (b) the data without the 22 highest isolated price spikes registered during the crisis. The values given in the table are suggestive of a change in structure of the time series during the crisis period.

Estimates of α_0 and α_1 indicate an increase during the crisis period. The increase in α_1 suggests that the supply curve becomes very steep during the critical period. Also, the estimates of the demand curve slope, β_2 , indicate that the demand during August-December 2000 was higher than the rest.

Figure 11 shows the generalized residuals we defined earlier, for the fitted model, based on

	Cost Difference	
	Whole Time Period	Period with Extreme Values Removed
$\sum_t Q_t(P_t - \bar{P}_{pc})$	$869.52 \cdot 10^6$	$628.19 \cdot 10^6$
$\sum_t Q_t(\hat{P}_t - \bar{P}_{pc})$	$678.60 \cdot 10^6$	$613.68 \cdot 10^6$
$\sum_t Q_t(P_t - \hat{P}_t)$	$190.92 \cdot 10^6$	$145.09 \cdot 10^5$

Table 2: Measures of difference between pre-crisis and overall behaviour of prices and quantities.

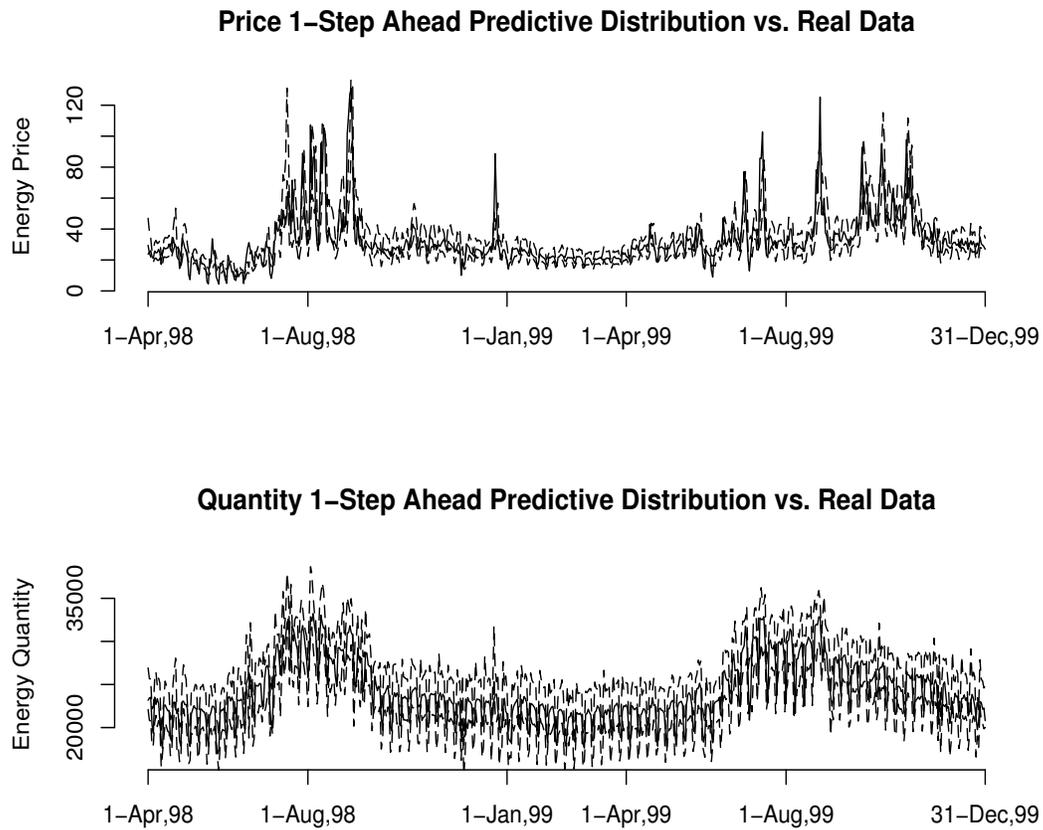


Figure 8: 0.05 and 0.95 quantiles of one-step predictive distributions (dashed lines) for price and quantity (dashed lines), along with the observed data (solid lines) from May 1, 1998 to December 31, 1999.

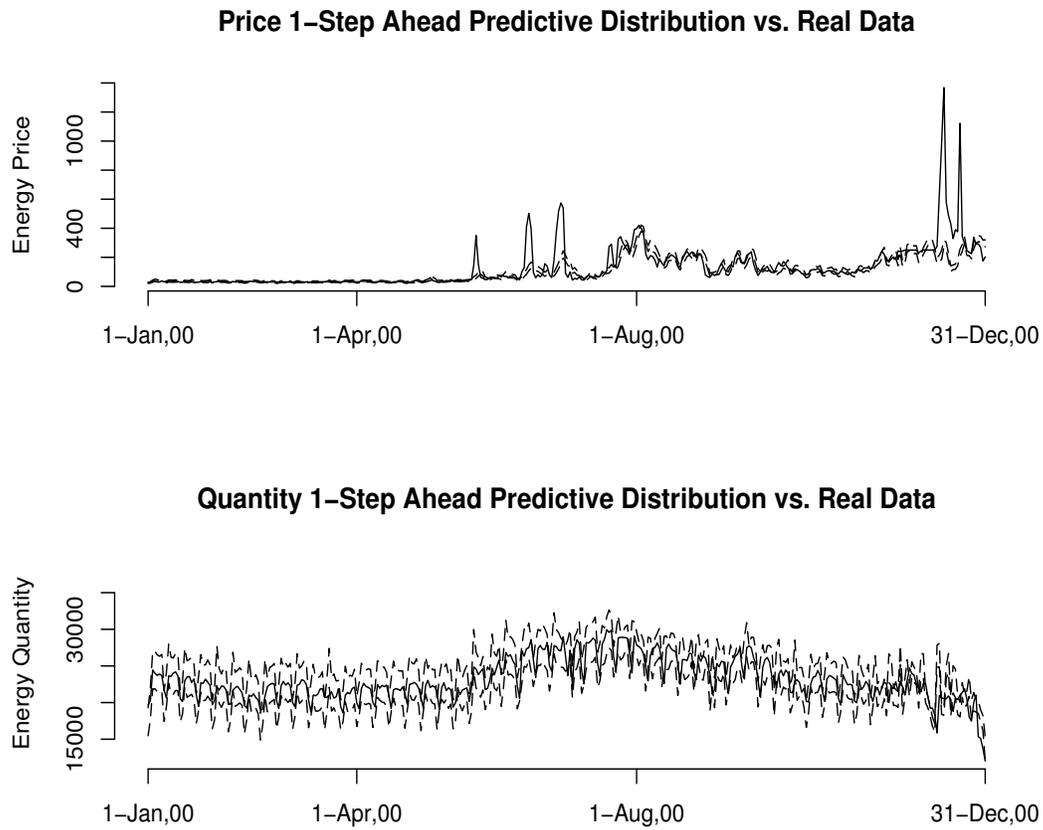
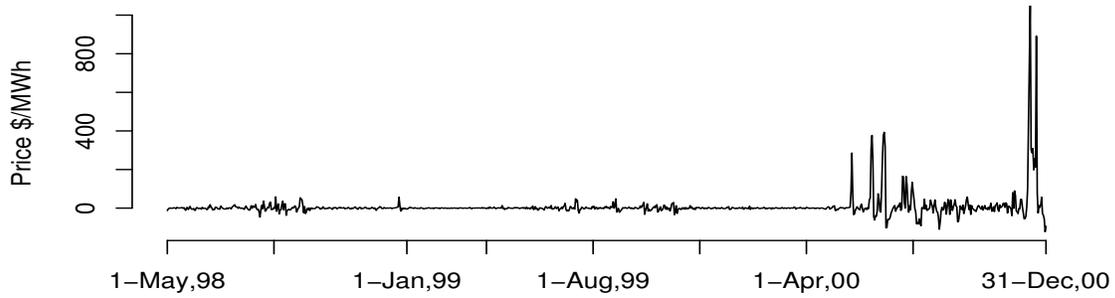
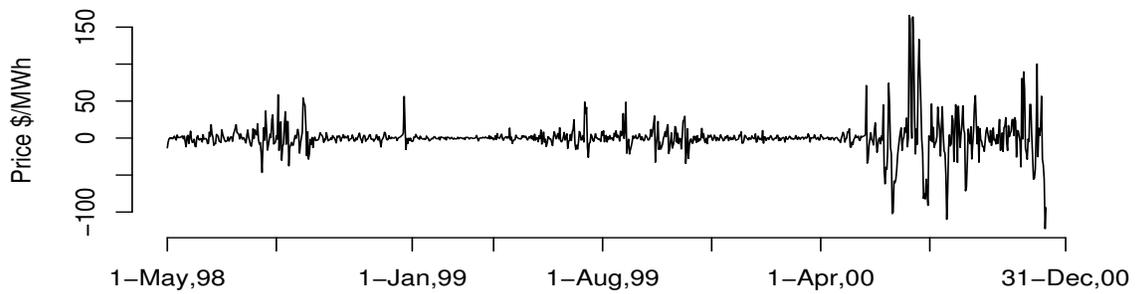


Figure 9: 0.05 and 0.95 quantiles of the one-step predictive distributions for price and quantity (dashed lines) along with the observed data (solid lines) for the year 2000



(a)



(b)

Figure 10: Differences between the real and one-step predictive prices for (a) all data and (b) for the data without the 22 highest price spikes.

the price and quantity one-step predictive cdfs. The scatter plots and quantile-quantile plots indicate reasonable goodness-of-fit, although the curve in the q-q plot for prices suggests that further improvement could be made.

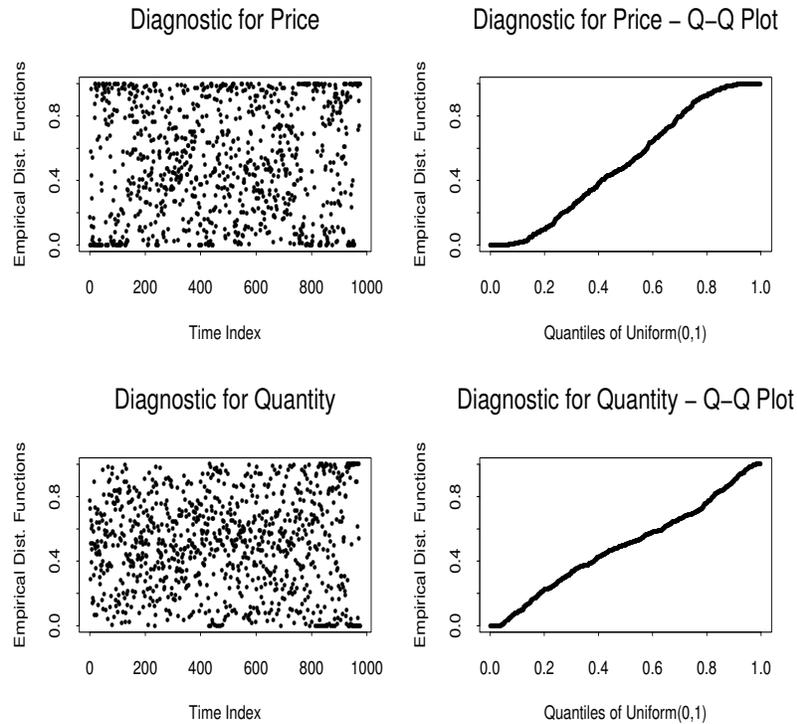


Figure 11: Generalized residuals (which should be iid Uniform(0,1) under the null hypothesis that the model is correct) for the fitted model.

4 Discussion

This paper has introduced a new dynamic supply and demand-based model for electricity prices. The model incorporates quantity data and a number of exogenous variables to describe the market dynamics. The model was fit to the California energy market data and appeared to fit the pre-crisis data well, but not the during-crisis data. Since we have accounted for temperature, natural gas supply, and plant stoppages, along with implicit random variation, the explanation for high prices during the crisis period would appear to be something else. One possible explanation is that producing firms may have withheld energy supply from the day-ahead market in order to drive the prices up. However, there could be other explanations as well, including the existence of factors which we have not accounted for. Without further study, and elimination of all other possible factors, it would seem impossible either to confirm or rule out the possibility of market manipulation by suppliers during the crisis.

References

- S. Deng. Pricing electricity derivatives under alternative stochastic spot price models. In *Proc. of the 33rd Hawaii International Conference on System Sciences*, 2000.
- A. Doucet, N. de Freitas, and N. Gordon, editors. *Sequential Monte Carlo Methods in Practice*. Springer, New York, 2001.
- P. Joskow and E. Kahn. A quantitative analysis of pricing behavior in california's wholesale electricity markey during summer 2000: The final word. Technical report, Massachusetts Institute of Technology.
- G. Kitagawa. Monte Carlo filter and smoother for non-Gaussian nonlinear state space models. *Journal of Computational and Graphical Statistics*, 5(1):1–25, 1996.
- C.R. Knittel and M.R. Roberts. An empirical examination of deregulated electricity prices. Technical Report PWR-087, Univ. of California Energy Institute, 2001.
- G. Mankiw. *Principles of Economics*. The Dryden Press, 1998.
- M. Parkin. *Economics*. Pearson Eduction, Inc., sixth edition, 2003.
- M. Rosenblatt. Remarks on a multivariate transformation. *Annals of Mathematical Statistics*, 23:470–472, 1952.
- D. Salmond and N. Gordon. Particles and mixtures for tracking and guidance. In *Sequential Monte Carlo Methods in Practice*, pages 517–528. Springer, 2001.