Available online at www.sciencedirect.com
SCIENCE DIRECT®
Energy Economics 27 (2005) 791–817
Energy Economics
www.elsevier.com/locate/eneco

An empirical examination of restructured electricity prices

Christopher R. Knittel[a,b,*], Michael R. Roberts[c]

[a]Department of Economics at the University of California, Davis, United States
[b]University of California Energy Institute and NBER, United States
[c]Wharton School at the University of Pennsylvania, United States

Available online 17 May 2005

Abstract

We present an empirical analysis of restructured electricity prices. We study the distributional and temporal properties of the price process in a non-parametric framework, after which we parametrically model the price process using several common asset price specifications from the asset-pricing literature, as well as several less conventional models motivated by the peculiarities of electricity prices. The findings reveal several characteristics unique to electricity prices including several deterministic components of the price series at different frequencies. An “inverse leverage effect” is also found, where positive shocks to the price series result in larger increases in volatility than negative shocks. We find that forecasting performance is dramatically improved when we incorporate features of electricity prices not commonly modelled in other asset prices. Our findings have implications for how empiricists model electricity prices, as well as how theorists specify models of energy pricing.

JEL classification: C51; G12

Keywords: Financial models; Restructured electricity prices; Price process

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

* Corresponding author. Department of Economics at the University of California, Davis, United States.
E-mail addresses: crknittel@ucdavis.edu (C.R. Knittel), mrrobert@wharton.upenn.edu (M.R. Roberts).

0140-9883/$ - see front matter
doi:10.1016/j.eneco.2004.11.005

One of the many consequences of restructuring is an increase in the importance of modeling and forecasting electricity prices.

Under regulation, prices were set by state public utility commissions (PUCs) in order to curb market power and ensure the solvency of the firm. Price variation was minimal and under the strict control of regulators, who determined prices largely on the basis of average costs. This environment focused the industry's attention on demand forecasting, as prices were held constant between PUC hearings. Market entry was barred and investment in new generation by incumbent firms was largely based on demand forecasts. In addition, there was little need for hedging electricity price risk because of the deterministic nature of prices.

Restructuring removes price controls and openly encourages market entry. Consequently, price variation has skyrocketed, giving birth to a new market for energy-based financial products as purchasers of electricity attempt to 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 large part, to the difficulty in understanding the behavior of electricity prices.[2] To this ends, we investigate the behavior of California's restructured electricity prices. We begin by analyzing electricity prices and comparing their salient features with those of equities and other commodities. Despite a few distributional similarities, electricity prices are dramatically different from equity prices, and even other commodity prices. Specifically, electricity prices display the following distinct characteristics:

1. stationarity in both the price level and squared prices,
2. pronounced intraday, day of week, and seasonal cycles, and
3. extreme price swings in a short period of time,
4. censoring from above,
5. negative prices.

Most of these characteristics represent departures from equity prices and interest rates, which form the foundation for much of empirical asset-pricing research. As such, we show

[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 and 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 and 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.

that statistical models developed for the purpose of modeling equity prices and interest rates fail to provide a reasonable description of the data generating process. Rather, richer models, in terms of the range of dynamics that they produce, are needed.

We are also able to test a number of empirical predictions from equilibrium electricity pricing models. For example, we confirm that power spot prices contain a positive skew that is larger (smaller) during periods of high (low) demand variability, as suggested by Bessembinder and Lemmon (2002). A similar pattern exists for the volatility of spot prices, higher (lower) during periods of high (low) demand, also predicted by Bessembinder and Lemmon (2002). We find mean reversion in spot prices, as implied by the equilibrium model of Routledge et al. (2001). However, we also note that electricity prices also contain what we refer to as, an "inverse leverage effect."[3] Electricity price volatility tends to rise more so with positive shocks than negative shocks; this is a result of convex marginal costs.

While there has been a significant amount of research on commodity prices, because of the relative youth of electricity market restructuring, 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. 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 1 discusses the market for electricity and the potential for an noncompetitive environment. Section 2 presents the data and a discussion of the distributional and temporal properties of electricity prices. Section 3 presents the statistical models of electricity prices. Section 4 concludes with a summary of the findings and directions for future research.

1. Electricity markets and market power

1.1. The electricity market

To date, a number of electricity markets around the world have restructured. The majority of restructured electricity markets set market clearing spot prices through an auction. Generators and demanders submit supply and demand curves for given time period; the equilibrium price is the resulting market price. For example, in the California electricity market, generators and demanders submit hourly supply and demand curves and the Independent System Operator (ISO) calculates the market clearing price. The auction is

[3] This terminology is derived from Black's (1976) "leverage effect" and describes the asymmetric response of volatility to positive and negative shocks.

[4] For a theoretical treatment of commodity prices see the papers by Chambers and Bailey (1996), Deaton and Laroque (1992, 1996), and Wright and Williams (1989). Duffie and Gray (1996) examine the volatility of energy prices. Kirk (1996) examines correlation in the energy markets.

a uniform price auction in which all suppliers and demanders receive and pay the same price, respectively. We use this market clearing price in this paper and discuss it in more detail below.

The nature of electricity and the behavior of electricity prices differ 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.[5] 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. Electricity markets also face extreme capacity constraints; this feature along with the non-storability of electricity implies the supply of electricity is often extremely inelastic.

The demand for electricity, at least in the short term, is also extremely in inelastic. Indeed, during our sample the California market retail customers faced a fixed price, independent of the wholesale price; therefore, their demand was completely inelastic. Combined with inelastic supply, small changes in either the supply or demand for electricity can have huge effects on market clearing prices, generating a tremendous degree of volatility.

1.2. Market power in electricity markets

Another distinguishing feature of electricity markets is the potential for suppliers to exercise market power. A number of studies have examined this issue by using one of three approaches. The first is to use data on the engineering marginal cost function and simulate a game-theoretic model of firm behavior, and compare the simulated strategic prices to simulated perfectly competitive prices. For example, Green and Newbery (1992) simulate supply function equilibria for the UK market and find that prices can substantially exceed marginal costs.[6] Borenstein and Bushnell (1999) simulate California's electricity market and find that during peak hours, firms have substantial market power. Borenstein et al. (1999) simulate the California and Pennsylvania, New Jersey and Maryland (PJM) markets and find similar results; furthermore, these authors find that traditional measures used to screen for market power do a poor job.

The second approach is more direct. A number of studies compute hourly engineering marginal cost and compare them to actual market prices. For example, Wolfram (2000) compares market prices to market marginal costs for the UK restructured electricity market and finds that the average Lerner index is 0.24.[7] However, much of the Lerner index can

[5] 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.

[6] The supply function equilibria literature began with Klemperer and Meyer (1989) and assumes firms submit supply functions rather than specific quantities or prices to an auctioneer under demand uncertainty. Klemperer and Meyer show that as the degree of uncertainty approaches zero, the set supply function equilibria converges to the Cournot equilibrium.

[7] The Lerner index is the typical measure of market power and is defined as $(P_t - mc_t)/P_t$.

be explained by the inelastic nature of electricity demand. After controlling for the elasticity of demand, the average elasticity adjusted Lerner index is 0.05.[8] More recently, Borenstein et al. (2002), and Joskow and Kahn (2002) compare California market prices to marginal costs. Both papers find that mark-ups over marginal cost increased during the crisis period.

Kim and Knittel (2003) use a third approach for estimating market power. They use "New Empirical Industrial Organization" techniques to estimate market power levels in the California market. Kim and Knittel (2003) find that mark-ups increased during the crisis period. New Empirical Industrial Organization techniques identify both marginal costs and market power levels, without data on marginal costs, by estimating the firms' first order conditions and treating marginal costs as functions of unknown parameters. They find that market power levels have been moderate.

While these studies are of great importance, our analysis does not speak to the level of market power in the industry. In this paper, we analyze the pricing behavior using traditional statistical and financial models of electricity prices. These pricing models are necessarily reduced form and independent of cost or market power levels. As mentioned above, our goal in this paper is to better understand electricity prices for the purpose of building empirical models, using these models in valuing energy securities, and testing equilibrium models of electricity pricing. However, we do note in the analysis below, a shift in the statistical properties of California prices in May of 2000, which is consistent with the increase in market power found by Borenstein et al. (2002), Joskow and Kahn (2002) and Kim and Knittel (2003).

2. Data and electricity price properties

The data used in this study consist of hourly electricity prices from California. Because of transmission capacity constraints, California is partitioned into 26 "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.[9] Since prices across zones are likely to behave similarly, we focus on just one zone, denoted NP15, that corresponds to Northern California.[10] 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.

[8] The elasticity adjusted Lerner index controls for cross-sectional or time series deviations between prices and marginal costs driven by differences in the elasticity of demand. It is defined as $(P_t - mc_t)\eta_t / P_t$, where η_t is the elasticity of demand at time t. A convenient benchmark is that the elasticity adjusted Lerner index for a symmetric Cournot game with N firms is $1/N$. Therefore, Wolfram (2000) results suggest the market is acting as though there are 20 symmetric firms, whereas there were three firms, one of which dominated the market with 52% of capacity.

[9] See Borenstein et al. (2001) for a further discussion of this.

[10] The sizes of the 26 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. During our sample, roughly 90% 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.

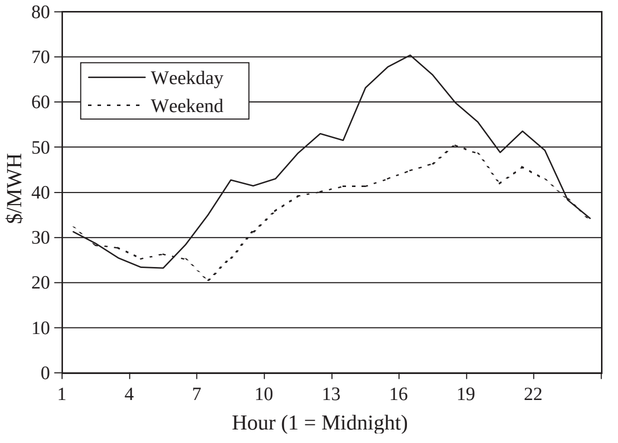

Fig. 1. Average hourly electricity prices across the entire sample.

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. Fig. 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. Fig. 2 presents a sample of hourly price and quantity 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.

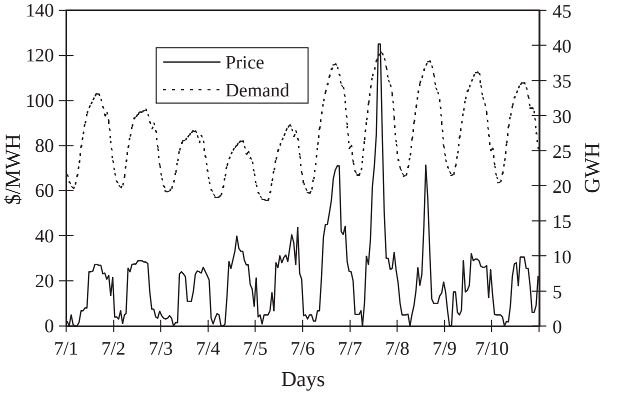

Fig. 2. A sample of hourly electricity prices for the period July 1, 1998 to July 10, 1998.

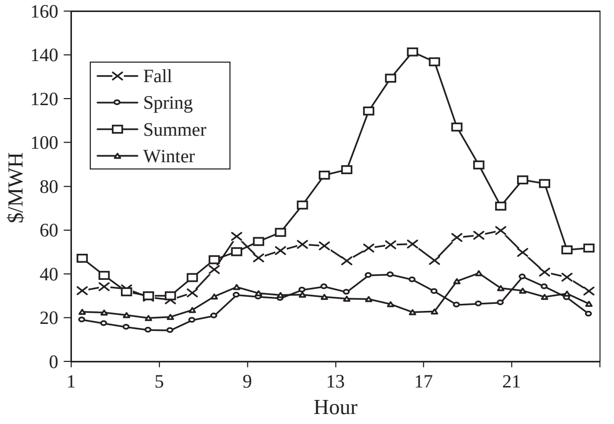

Fig. 3. Average weekday hourly electricity prices by season.

Electricity prices also contain a strong seasonal component, reflecting heating and cooling needs. This feature is illustrated in Fig. 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.[11] Indeed, California electricity prices are consistent with this theory. Within a time span of 24 h, the spot price for electricity in California can move from a price of less than \$5/MWh to \$750/MWh.

Fig. 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. 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.

[11] See Borenstein et al. (1999) for a discussion of this idea.

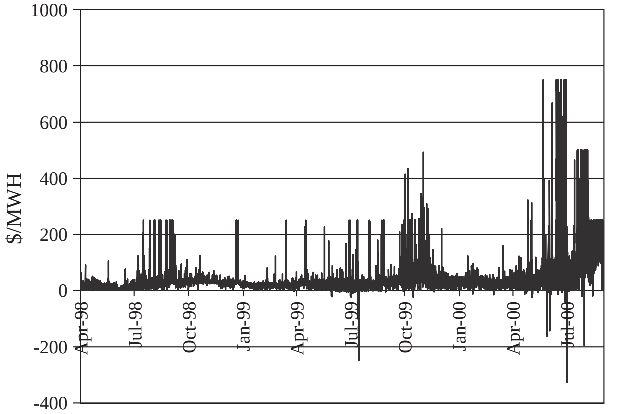

Fig. 4. Hourly electricity prices for April 1, 1998 to July 31, 2000.

1.3. Distributional properties

In this subsection, we discuss the distributional properties of electricity prices. Fig. 5 presents the empirical histogram for our price series, overlaid with a normal density curve. Fig. 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. Fig. 6

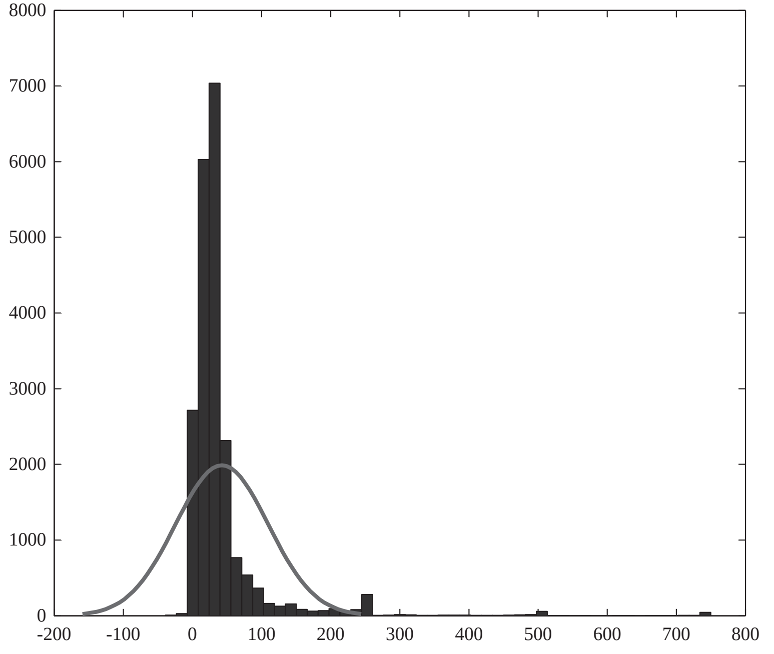

Fig. 5. Empirical histogram of electricity prices with normal density superimposed.

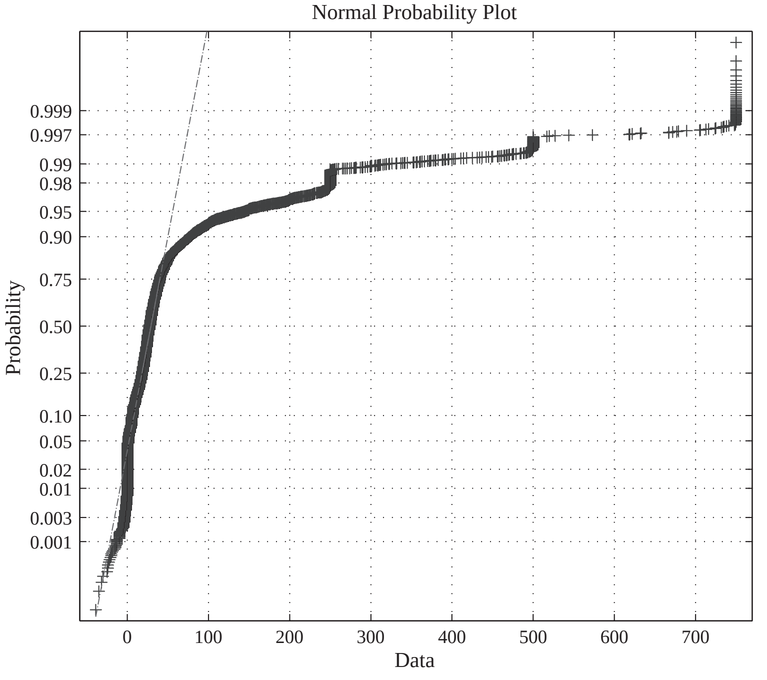

Fig. 6. QQ-plot of electricity prices.

shows the heavy tail of the distribution and mass of observations at zero, resulting from the features of electricity described in Section 1.1.

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

Table 1
Electricity price summary statistics

Statistic	Pre-crisis	Crisis	Pre-crisis May–August months
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

Descriptive statistics for hourly NP15 (Northern California) electricity spot prices over the pre-crisis and crisis period, which are delineated by May 1, 2000. To control for temporal changes in electricity demand, pre-crisis descriptive statistics for the months May through August are also included.

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 Fig. 4. The conditional mean price is varying over time in a systematic fashion. This fact is even more apparent in Figs. 1 and 2. Fig. 4 also shows that the variance of electricity prices is time varying and exhibits clustering.

1.4. Temporal properties

While the price series has several distributional similarities with other tradeable assets, their temporal properties are quite different. Fig. 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.[12]

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.[13] This hypothesis is tested more formally as follows. Consider the simple approximation to the price generating process:

$$p_t = \alpha + \beta p_{t-1} + \eta_t \tag{1}$$

$$\eta_t = \gamma \eta_{t-1} + \varepsilon_t \tag{2}$$

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_ε^2.[14] In the presence of serially correlated errors, Phillips and Perron (1988) show that the parameters of Eq. (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.[15] With a 5% critical value of -2.89, the null of a unit root is rejected at all standard significance levels.

[12] 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.

[13] 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.

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

[15] 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.

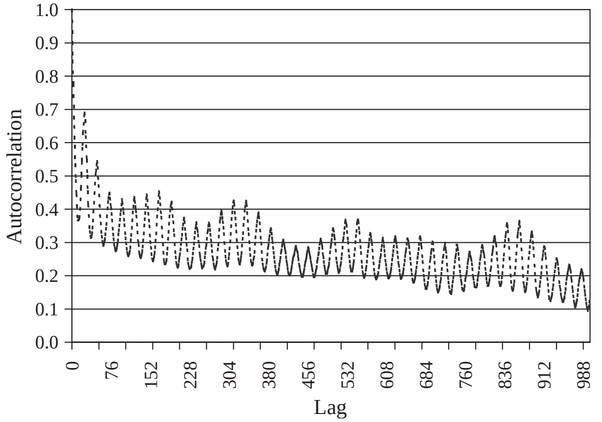

Fig. 7. Autocorrelation function for electricity prices.

Fig. 8 plots the autocorrelation function for the square of prices, which bears a striking resemblance to Fig. 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,

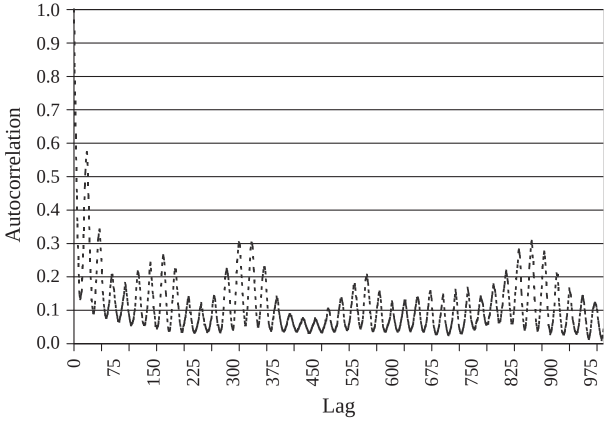

Fig. 8. Autocorrelation function for the square of electricity prices.

5. time-varying volatility and volatility clustering, and
6. extreme values.

We consider the right censoring and negative prices to be of less importance as the effect of censoring on the dynamics of prices is likely to have a secondary effect. 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 1-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 1 day to several months. All models are estimated by conditional maximum likelihood. 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.

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.

1.5. 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:

$$\mathrm{d}p(t) = \kappa[\mu - p(t)]\mathrm{d}t + \sigma \mathrm{d}b(t), \qquad p(0) = p_0, \tag{3}$$

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, $\sigma \mathrm{d}b(t)$.[16]

Eq. (3) is simply a first order autoregressive model in continuous time. This may be seen by integrating Eq. (3) to obtain:

$$p(t) = \mathrm{e}^{-\kappa t} p_0 + \mu(1 - \mathrm{e}^{-t\kappa}) + \int_0^t \mathrm{e}^{\kappa(s-t)} \sigma \mathrm{d}b(s). \tag{4}$$

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

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

where $\alpha_0 = \mu(1 - e^{-\kappa})$, $\beta_1 = e^{-\kappa}$, and $\eta_t = \int_{t-1}^{t} \mathrm{e}^{\kappa(s-t)} \sigma \mathrm{d}b(s)$.[17] The error term, η_t, in Eq. (5) is Gaussian white noise with variance σ_η^2 equal to $\sigma^2 \, (1 - e^{-2\kappa})/2$, 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.

We estimate the discrete time parameters in Eq. (5): α_0, β_1, and σ_η^2. The log likelihood function for the discrete time model is:

$$\log L(\alpha_0, \beta_1, \sigma_\eta) = -\frac{T-1}{2} \log(2\pi) - \frac{T-1}{2} \log \sigma_\eta^2 - \frac{1}{2\sigma_\eta^2} \sum_{t=2}^{T} (p_t - \beta_1 p_{t-1})^2. \tag{6}$$

To estimate the model, we condition on the first observation (i.e. treat is as known). 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 (e.g., $\alpha_0 = \mu(1 - \mathrm{e}^{-\kappa})$,). 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 for forecasting purposes.

The results are presented in Tables 2 and 3. 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 that 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. The root mean squared (RMS) error for the week ahead out-of-sample forecast during the non-crisis period is 47.51. As we noted above, given the stochastic nature of electricity prices, the forecast error will necessarily be high; we focus on the change in the forecast error across different models. During the crisis period, the RMS increases to 88.56. This increase is due to both an increase in the mean and volatility of electricity prices. We use Model 1 as a baseline to compare the remaining models.

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

[17] 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.

Table 2
Pre-crisis period parameter estimates of Models 1 through 4

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

1.6. Model 2: time-varying mean

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

$$\mathrm{d}p(t) = \kappa(\mu(t) - p(t))\mathrm{d}t + \sigma \mathrm{d}b(t), \tag{7}$$

where

$$\mu(t) = \alpha_1 I(t \in \text{Peak}) + \alpha_2 I(t \in \text{OffPeak}) + \alpha_3 I(t \in \text{Weekend}) + \alpha_4 I(t \in \text{Fall}) + \alpha_5 I(t \in \text{Winter}) + \alpha_6 I(t \in \text{Spring}) \tag{8}$$

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

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

Notes to Table 2:
The sample of data is comprised of hourly electricity prices for Northern California during the "pre-crisis" period: 4/1/1998–4/23/2000. The estimated models are:

1: $\mathrm{d}p(t) = \kappa(\mu - p(t))\mathrm{d}t + \sigma \mathrm{d}b(t)$

2: $\mathrm{d}p(t) = \kappa(\mu(t) - p(t))\mathrm{d}t + \sigma \mathrm{d}b(t)$, $\mu(t) = \alpha_1 I(t \in \text{Peak}) + \alpha_2 I(t \in \text{OffPeak}) + \alpha_3 I(t \in \text{Weekend}) + \alpha_4 I(t \in \text{Fall}) + \alpha_5 I(t \in \text{Winter}) + \alpha_6 I\ (t \in \text{Spring})$

3: $\mathrm{d}p(t) = \kappa(\mu(t) - p(t))\mathrm{d}t + \sigma_b \mathrm{d}b\ (t) + z\mathrm{d}q(t)$

3a: $\mathrm{d}p(t) = \kappa(\mu(t) - p(t))\mathrm{d}t + \sigma_b \mathrm{d}b(t) + z\mathrm{d}q(t)\lambda$
$(t) = \lambda_0 + \lambda_{\text{peak}} I(t \in \text{Peak}) + \lambda_{\text{weekend}} I(t \in \text{Weekend}) + \lambda_{\text{fall}} I(t \in \text{Fall}) + \lambda_{\text{win}} I(t \in \text{Winter}) + \lambda_{\text{spr}} I(t \in \text{Spring})$

4: $p_t = \mu_t + \eta_t$
$\beta(L)\eta_t = \delta(L)\varepsilon_t$
$\beta(L) = 1 - \beta_1 L - \beta_2 L^{24} - \beta_3 L^{25}$
$\delta(L) = 1 - \delta_1 L - \delta_2 L^{24} - \delta_3 L^{25}$,

where $p(t)$ is the electricity price at time t, $b(t)$ is Brownian motion, $I(\cdot)$ is the indicator function, $q(t)$ is a Poisson process with jump intensity λ, z, ε_t and η_t are standard normal random variables, L is the lag operator defined as $L^n x_t = x_{t-n}$, Peak hours are from 6 a.m. to 10 p.m. All parameters are estimated via maximum likelihood. Standard errors are in parentheses.

Notes to Table 3:
The sample of data is comprised of the "crisis" period: 5/1/2000–8/24/2000. The estimated models are:

1: $\mathrm{d}p(t) = \kappa(\mu - p(t))\mathrm{d}t + \sigma \mathrm{d}b(t)$

2: $\mathrm{d}p(t) = \kappa(\mu(t) - p(t))\mathrm{d}t + \sigma \mathrm{d}b(t)$, $\mu(t) = \alpha_1 I(t \in \text{Peak}) + \alpha_2 I(t \in \text{OffPeak}) + \alpha_3 I(t \in \text{Weekend}) + \alpha_4 I(t \in \text{Fall}) + \alpha_5 I(t \in \text{Winter}) + \alpha_6 I\ (t \in \text{Spring})$

3: $\mathrm{d}p(t) = \kappa(\mu(t) - p(t))\mathrm{d}t + \sigma_b \mathrm{d}b\ (t) + z\mathrm{d}q(t)$

3a: $\mathrm{d}p(t) = \kappa(\mu(t) - p(t))\mathrm{d}t + \sigma_b \mathrm{d}b(t) + z\mathrm{d}q(t)\lambda$
$(t) = \lambda_0 + \lambda_{\text{peak}} I(t \in \text{Peak}) + \lambda_{\text{weekend}} I(t \in \text{Weekend}) + \lambda_{\text{fall}} I(t \in \text{Fall}) + \lambda_{\text{win}} I(t \in \text{Winter}) + \lambda_{\text{spr}} I(t \in \text{Spring})$

4: $p_t = \mu_t + \eta_t$
$\beta(L)\eta_t = \delta(L)\varepsilon_t$
$\beta(L) = 1 - \beta_1 L - \beta_2 L^{24} - \beta_3 L^{25}$
$\delta(L) = 1 - \delta_1 L - \delta_2 L^{24} - \delta_3 L^{25}$,

where $p(t)$ is the electricity price at time t, $b(t)$ is Brownian motion, $I(\cdot)$ is the indicator function, $q(t)$ is a Poisson process with jump intensityλ, z, ε_t and η_t are standard normal random variables, L is the lag operator defined as $L^n x_t = x_{t-n}$, Peak hours are from 6 a.m. to 10 p.m. All parameters are estimated via maximum likelihood. Standard errors are in parentheses.

This specification implies that $\mu(t)$ is a step function, constant across any 1 h. Integrating Eq. (7) yields:

$$p(t) = \mathrm{e}^{-\kappa t} p_0 + \int_0^t \mathrm{e}^{-k(t-s)} \kappa \mu(s) \mathrm{d}s + \int_0^t \mathrm{e}^{-\kappa(t-s)} \sigma \mathrm{d}b(s)$$
$$= \mathrm{e}^{-\kappa} p(t-1) + \int_{t-1}^t \mathrm{e}^{-\kappa(t-s)} \kappa \mu(s) \mathrm{d}s + \int_{t-1}^t \mathrm{e}^{-\kappa(t-s)} \sigma \mathrm{d}b(s). \quad (9)$$

Table 3
Crisis period parameter estimates of Models 1 through 4

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)
α_1 (MA 1)					0.2899
					(0.0237)
α_1 (MA24)					0.6687
					(0.0410)
α_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	(3.2811)
				(3.3039)	
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	73.3012	66.6337

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 Eq. (9) is:

$$p_t = \alpha_t + \beta_1 p_{t-1} + \eta_t, \tag{10}$$

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 Eqs. (5) and (10) is in the intercept. As such, prices are again Markovian with a Gaussian transition density.

Eq. (10) 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).[18] More explicitly, Eq. (10) may be written as:

$$\begin{aligned} p_t = {} & \alpha_1 \text{Peak}_t + \alpha_2 \text{OffPeak}_t + \alpha_3 \text{Weekend}_t + \alpha_4 \text{Fall}_t + \alpha_5 \text{Winter}_t \\ & + \alpha_6 \text{Spring}_t + \beta_1 p_{t-1} + \eta_t. \end{aligned} \tag{11}$$

Model 2 is also estimated via conditional maximum likelihood. The likelihood function differs from Model 1 only through α_t, which includes the six binary variables, instead of being a constant as in Model 1. The estimates for the pre-crisis and crisis periods are presented in column 2 of Tables 2 and 3. 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.[19] 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. The RMS during the pre-crisis period is 47.15, an improvement of only 0.36 over Model 1; the forecasted price series still fails to capture the erratic nature of the true price series. During the crisis period, the improvement is greater as the RMS is 76.11 compared to 88.56 for Model 1. This is due to the larger disconnect between the peak and off-peak parameters during the crisis period, α_1 and α_2 (159.98 vs. 125.13 in the crisis period compared to 36.90 vs. 32.73 in the pre-crisis period).

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

[19] Since the specification in Eq. (10) 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).

1.7. 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 Eq. (7), yielding:

$$\mathrm{d}p(t) = \kappa(\mu(t) - p(t))\mathrm{d}t + \sigma_b \mathrm{d}b(t) + z\mathrm{d}q(t) \tag{12}$$

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.

As Ball and Torous (1983) note, empirical implementation of Eq. (12) 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 Eq. (7). 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, then 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 or fall when a jump occurs, the variance always increases.[20]

The conditional likelihood function is thus:

$$\log L(\theta) = \prod_{t=2}^{T} (1-\lambda)\phi\left(\frac{P_t - (\alpha_t + \beta_1 P_{t-1})}{\sigma_b}\right)\sigma_b^{-1} \tag{13}$$

$$+ \lambda\phi\left(\frac{P_t - (\alpha_t + \beta_1 P_{t-1} + \mu_z)}{\left(\sigma_b^2 + \sigma_z^2\right)^{1/2}}\right)\left(\sigma_b^2 + \sigma_z^2\right)^{-1/2}, \tag{14}$$

where $\phi(\cdot)$ is the standard normal density and μ are the unknown parameters.

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

$$\log L(\theta) = \prod_{t=2}^{T} (1-\lambda)\phi\left(\frac{P_t - (\alpha_t + \beta_1 P_{t-1})}{\sigma_b}\right)\sigma_b^{-1} \tag{15}$$

$$+ \lambda\phi\left(\frac{P_t - (\alpha_t + \beta_1 P_{t-1} + \mu_z)}{(\xi\sigma_b)}\right)(\xi\sigma_b)^{-1}, \tag{16}$$

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

The results are presented in Tables 2 and 3. 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 h. 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,

[20] 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.

the price volatility is 62 times larger than when a jump does not occur: 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 h. 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.

To generate useful forecasts from this model, 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.[21] The forecasting performance of the jump-diffusion model is poor in the pre-crisis period compared to Model 1. The RMS is 49.40, greater than the simple mean-reverting model. During the crisis period, the forecasting performance of the jump-diffusion model improves compared to Models 1 and 2; the RMS is 73.34, lower than the two preceding models. These results are consistent with our priors that during the crisis period jumps became more commonplace, increasing their importance in a model of the price process.

1.8. 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, by specifying:

$$\lambda(t) = \lambda_0 + \lambda_{\text{peak}}\text{Peak}_t + \lambda_{\text{weekend}}\text{Weekend}_t + \lambda_{\text{fall}}\text{Fall}_t + \lambda_{\text{win}}\text{Winter}_t + \lambda_{\text{spr}}\text{Spring}_t.$$

The results listed in Tables 2 and 3 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.

The same forecasting issues discussed above are present here. The performance of this model improves markedly during the pre-crisis period over the more simple jump intensity specification. The RMS decreases to 44.35, lower than Models 1 and 2. There is little

[21] 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.

improvement during the crisis period, however. The RMS is 73.30, virtually identical to the more simple jump-diffusion model.

1.9. 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:

$$p_t = \alpha_t + \eta_t \tag{17}$$

$$\beta(L)\eta_t = \delta(L)\varepsilon_t. \tag{18}$$

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

$$\beta(L) = 1 - \beta_1 L - \beta_2 L^{24} - \beta_3 L^{25} \tag{19}$$

$$\delta(L) = 1 - \delta_1 L - \delta_2 L^{24} - \delta_3 L^{25}. \tag{20}$$

The mean α_t is as specified in Eq. (10) and $\{\epsilon_t\}$ is Gaussian white noise with variance parameter α_2. The motivation for Eqs. (17)–(20) follows from an examination of the correlogram, which shows high correlation between the current price and the previous day's prices.

Model 4 is estimated via conditional maximum likelihood. Because there the model includes 25 lags, the first 25 observations are treated as fixed. The likelihood function can be derived as follows. First define the hourly error terms as:

$$\varepsilon_t = p_t - \alpha_t - \beta_1 p_{t-1} - \beta_2 p_{t-24} - \beta_3 p_{t-25} - \delta_1 \varepsilon_{t-1} - \delta_2 \varepsilon_{t-24} - \delta_3 \varepsilon_{t-25}. \tag{21}$$

The log likelihood function is given as:

$$\log L(\theta) = -\frac{T-25}{2}\log(2\pi) - \frac{T-25}{2}\log\sigma^2 - \frac{1}{2\sigma^2}\sum_{t=26}^{T}\varepsilon_t^2, \tag{22}$$

where θ is the vector of unknown parameters.

Clear from the parameter estimates in Tables 2 and 3 is that electricity prices are not reasonably approximated by a univariate Markov process. All lag parameters are highly statistically significant.[22] 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 pre-crisis period, the RMS is, by far, the lowest of the first five models -25.48. This improvement is also present in the

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

crisis period, where the RMS is 66.63, compared to 73.30 the lowest of the first four models. These results underline the importance of incorporating higher order autocorrelation in models of restructured electricity prices.

1.10. 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:

$$p_t = \alpha_t + \eta_t, \tag{23}$$

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

$$\beta(L)\eta_t = v_t, \tag{24}$$

where $\beta(L)$ is the lag operator defined in Eq. (19) above. To capture the conditional heteroscedasticity, we adopt the EGARCH model of Nelson (1991), modeling v_t as

$$v_t = \sqrt{h_t}\varepsilon_t, \tag{25}$$

$$\ln(h_t) = \theta + \sum_{i \in \{1,24,25\}} \kappa_i g(z_{t-i}) + \gamma_1 \ln(h_{t-1}), \tag{26}$$

where

$$g(z_s) = \{\psi z_s + |z_s| - E(|z_s|)\}$$

$$z_s = v_t / \sqrt{h_t},$$

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 Eqs. (23)–(26) are found in Tables 4 and 5.[23] 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

[23] See Hamilton (1994, pp. 668–669) for the derivation of the log likelihood function.

Table 4
Pre-crisis period parameter estimates of Model 5

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

The pre-crisis estimation period is 4/1/1998–4/23/2000. Standard errors are in parentheses. Model 5 is an EGARCH model with a ARCH component of one lag in the volatility and GARCH components of 1, 24 and 25 lags.

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, with an RMS of 52.19. In contrast, in the crisis period, the EGARCH specification has the best forecasting performance, but not dramatically so

Table 5
Crisis period parameter estimates of Model 5

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

The crisis estimation period is 5/1/2000–8/24/2000. Standard errors are in parentheses. Model 5 is an EGARCH model with a ARCH component of one lag in the volatility and GARCH components of 1, 24 and 25 lags.

(RMS of 61.30 compared to 66.63 for Model 4). These results suggest that when the electricity market is supply constrained–as it was during the crisis period–incorporating volatility will be most important, since persistent positive shocks to demand will result in high prices and therefore clustering.

1.11. 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.[24] Fig. 9 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°, when heating becomes necessary. Since electric heating is rare in California, this relationship is subtle. After 55°, 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 Fig. 9.

To capture the diversity of regions, each of the three temperature series is 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 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 6 and 7. Specifically, we specify:

$$p_t = \alpha_t + \eta_t \tag{27}$$

$$\beta(L)\eta_t = \delta(L)\varepsilon_t, \tag{28}$$

where

$$\begin{aligned}\alpha_t = {} & \alpha_1 1(t \in \mathrm{Peak}) + \alpha_2 1(t \in \mathrm{OffPeak}) + \alpha_3 1(t \in \mathrm{Weekend}) + \alpha_4 1(t \in \mathrm{Fall}) \\ & + \alpha_5 1(t \in \mathrm{Winter}) + \alpha_6 1(t \in \mathrm{Spring}) + \alpha_7 \mathrm{Temp}_t + \alpha_8 \mathrm{Temp}_t + \alpha_9 \mathrm{Temp}_t^3,\end{aligned} \tag{29}$$

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

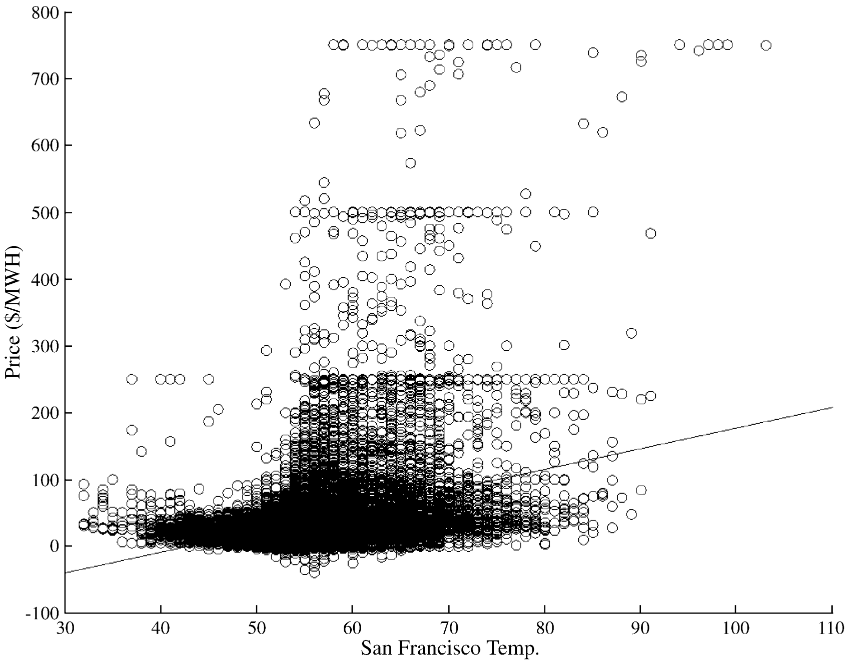

Fig. 9. Scatterplot of hourly spot prices vs. hourly San Francisco temperatures.

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.[25]

Referring to Tables 6 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. 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.

1.12. 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 to perform this transformation, nonpositive prices are set to missing and dropped from the analysis.[26]

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

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

Table 6
Pre-crisis period parameter estimates of model 6

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

The pre-crisis estimation period is 4/1/1998–4/23/2000. Standard errors are in parentheses. Model 6 is a ARMAX model with AR and MA terms of 1, 24, and 25 lags that incorporates a polynomial of the temperature of order 3.

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

Table 7
Crisis period parameter estimates of Model 6

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 (MA1)	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

The crisis estimation period is 5/1/2000–8/24/2000. Standard errors are in parentheses. Model 6 is a ARMAX model with AR and MA terms of 1, 24, and 25 lags that incorporates a polynomial of the temperature of order 3.

responsible for the forecasting performance of any model. The frequency of large price deviations from the conditional mean creates 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.[27]

4. Conclusions

The events of the past 2 years in California have made understanding the stochastic properties of restructured 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 restructured 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 render 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, can be greatly improved by incorporating the most salient features of electricity prices. Specifically, volatility clustering and higher order autocorrelation are two of the most important features. We also document an inverse leverage effect where positive price shocks increase price volatility more than negative shocks.

Acknowledgements

We thank an anonymous referee for insightful comments. This paper has benefited from conversations with Ravi Bansal, John Graham, Campbell Harvey, Pete Kyle and Tom Rothenberg, and comments from Mike Lemmon.

References

Ball, C., Torous, W.N., 1983. A simplified jump process for common stock returns. Journal of Financial and Quantitative Analysis 18 (1), 53–65.

Bergstrom, A.R., 1984. Continuous time stochastic models and issues of aggregation over time. In: Griliches, Zvi, Intriligator, Michael D. (Eds.), Handbook of Econometrics, vol. II. Elsevier, Amsterdam, pp. 1146–1212.

Bessembinder, H., Lemmon, M.J., 2002. Equilibrium pricing and optimal hedging in electricity forward markets. The Journal of Finance 57 (3), 1347–1382.

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

Bhanot, K., 2000. Behavior of power prices: implications for the valuation and hedging of financial contracts. The Journal of Risk 2 (3), 43–62.
Black, F., 1976. Studies in stock price volatility changes. Proceedings of the Business and Economics Statistics Section. Amercian Statistical Association, pp. 177–181.
Borenstein, S., Bushnell, J., 1999. An empirical analysis of market power in a deregulated California electricity market. Journal of Industrial Economics 47, 285–323.
Borenstein, S., Bushnell, J., Knittel, C.R., 1999. Market power in electricity markets: beyond concentration measures. The Energy Journal 20 (4), 65–83.
Borenstein, S., Bushnell, J., Knittel, C.R., Wolfram, C., 2001. Market Inefficiency in California's Electricity Markets, mimeo, UC Davis.
Borenstein, S., Bushnell, J., Wolak, F.A., 2002. Measuring market inefficiencies in California's wholesale electricity market. American Economic Review 92, 1376–1405.
Chambers, M., Bailey, R., 1996. A theory of commodity price fluctuations. Journal of Political Economy 104, 924–957.
Deaton, A., Laroque, G., 1992. On the behavior of commodity prices. Review of Economic Studies 59, 1–23.
Deaton, A., Laroque, G., 1996. Competitive storage and commodity price dynamics. Journal of Political Economy 104, 826–923.
Duffie, D., Gray, S., 1996. Volatility in energy prices. Managing Energy Price Risk. Risk Publications.
Durbin, J., 1970. Testing for serial correlation in least-squares regression when some of the regressors are lagged dependent variables. Econometrica 38, 410–421.
Green, R.J., Newbery, D.M., 1992. Competition in the British electricity spot market. Journal-of-Political-Economy 100, 929–953.
Hamilton, J., 1994. Time Series Analysis. Princeton University Press.
Hoare, J., 1996. The UK electricity market. Managing Energy Price Risk. Risk Publications.
Joskow, P.L., Kahn, E., 2002. A quantitative analysis of pricing behavior in California's wholesale electricity market during summer 2000. The Energy Journal 23, 1–35.
Kim, D.W., Knittel, C.R., 2003. Biases in Static Oligipoly Models? Empirical Evidence from the California Electricity Market, mimeo, UC Davis.
Kirk, E., 1996. Correlation in energy markets. Managing Energy Price Risk. Risk Publications.
Klemperer, P.D., Meyer, M.A., 1989. Supply function equilibria in oligopoly under uncertainty. Econometrica 57, 1243–1277.
Nelson, D.B., 1991. Conditional heterskedasticity in asset returns. Econometrica 59, 347–370.
Phillips, P.C.B., Perron, P., 1988. Testing for a unit root in time series regressions. Biometrika 75, 335–346.
Routledge, B.R., Seppi, D.J., Spatt, C.S., 2001. The 'Spark' spread: cross-commodity equilibrium restrictions and electricity, mimeo, Carnegi Mellon.
Wolfram, C.D., 2000. Measuring duopoly power in the British electricity spot market. American Economic Review 89, 805–826.
Wright, B., Williams, J., 1989. A theory of negative prices of storage. Journal of Futures Markets 9, 1–13.