Geman, H. and Roncoroni, A. (2006). Understanding the fine structure of electricity prices. *Journal of Business* **79** (3) 1225-1261.

Hélyette Geman
Birkbeck, University of London and ESSEC Business School

Andrea Roncoroni
ESSEC Business School

Understanding the Fine Structure of Electricity Prices*

This paper analyzes the special features of electricity spot prices derived from the physics of this commodity and from the economics of supply and demand in a market pool. Besides mean reversion, a property they share with other commodities, power prices exhibit the unique feature of spikes in trajectories. We introduce a class of discontinuous processes exhibiting a "jump-reversion" component to properly represent these sharp upward moves shortly followed by drops of similar magnitude. Our approach allows to capture—for the first time to our knowledge—both the *trajectorial* and the *statistical* properties of electricity pool prices. The quality of the fitting is illustrated on a database of major U.S. power markets.

I. Introduction

A decade ago, the electricity sector worldwide was a vertically integrated industry in which prices were set by regulators and reflected the costs of generation, transmission, and distribution. In this setting, power prices used to change rarely, and in an essentially deterministic manner. Over the last 10 years, major countries have been experiencing deregulation in generation and supply activities. One of the important consequences of this restructuring is that prices are now determined according to the fundamental economic rule of *supply and demand*: there is a "market pool" in which bids placed by generators to sell electricity for the next day are compared to purchase orders.

In a parallel way, deregulation of the energy industry has paved the way for a considerable amount of trading activity, in both the spot and derivative markets. Price risk in particular has forced the industry

* Both authors wish to thank Alexander Eydeland, Valentine Genon-Catalot, the editor (Albert Madansky), and an anonymous referee for helpful comments on earlier versions of the paper. Any remaining errors are our own. The work was financially supported by CERESSEC. Contact the corresponding author, Hélyette Geman, at geman@essec.fr and the coauthor, Andrea Roncoroni, at roncoroni@essec.fr.

[*Journal of Business*, 2006, vol. 79, no. 3]

0021-9398/2006/7903-0008$10.00

to identify, price, and hedge the options granted in energy contracts that have been written for decades.

Given the unique nonstorability (outside of hydro) of this commodity, electricity prices are much more likely to be driven by spot demand and supply considerations than any other good, with demand in the short-term market being fairly inelastic. As a result, sizable shocks in production or consumption may give rise to the price jumps that have been observed since 1998 in various parts of the United States. If one leaves aside the California 2000 events that were possibly driven by flaws in market design and wrongdoings on the part of some major players, spike prices have been motivated by disruption in transmission, generation outages, extreme weather, or a conjunction of these circumstances.

Today, an important fraction of the literature on electricity belongs to the economics arena and analyzes deregulated electricity markets from the regulatory viewpoint (see, e.g., Joskow and Kahn 2001). It is clear that a proper mathematical representation of spot prices is at the same time a necessary exercise and the cornerstone for the optimal scheduling of physical assets and the valuation of financial and real options in the electricity industry.

Some initial papers on the modeling of power price processes include Deng (1999), Bhanot (2000), Knittel and Roberts (2001), Pirrong (2001), Barlow (2002), Barone-Adesi and Gigli (2002), Escribano, Peña, and Villaplana (2002), and Lucia and Schwartz (2002). We extend this literature by proposing a family of stochastic processes meant to represent the *trajectorial* and *statistical* features displayed by electricity spot prices in deregulated power markets. We also introduce an effective method to identify spikes in historical raw data. In order to empirically investigate the information content of observed power price dynamics, we design a procedure for best fitting our model to market data in terms of both trajectories and moments. Since we focus on an analysis of the empirical properties of electricity prices, we shall work solely under the real probability measure. Yet, our concern is to preserve the Markov property in the view of future developments on the valuation of derivatives.

The paper is organized as follows. Section II discusses the main features of power prices and of the stack function. Section III introduces a class of processes that may encompass prices observed in a variety of regional markets. Section IV contains the description of data for three major U.S. power markets that exhibit different degrees of mean reversion and spike behavior. Section V analyzes the statistical methods allowing us to select a process matching observed spot prices. Section VI presents empirical results for all models and markets under investigation. Section VII concludes with a few comments and suggestions for future research.

II. The Key Features of Power Prices

Most of the important literature on commodities has focused on storable commodities (see, e.g., Fama and French 1987). The same property applies to the specific case of energy commodities since deregulated power markets were established fairly recently. Bessembinder and Lemmon (2002) build an equilibrium model for electricity forward markets derived from optimal hedging strategies conducted by utilities. In this setting, they compare forward prices to future spot prices and show that the forward premium increases when either expected demand or demand risk is high. Geman and Vasicek (2001) empirically confirm Bessembinder and Lemmon's findings and demonstrate, with a U.S. database, that short-term forward contracts are upward-biased estimators of future spot prices, in agreement with the high volatility and risk attached to U.S. spot power markets.

Our perspective in this paper is complementary and distinct at several levels. First, we are interested in the modeling of the spot price of electricity since we believe that in the wake of deregulation of power markets, a proper representation of the dynamics of spot prices becomes a necessary tool for trading purposes and optimal design of supply contracts. As discussed in Eydeland and Geman (1998), the nonstorability of electricity implies the breakdown of the spot-forward relationship and, in turn, the possibility of deriving the fundamental properties of spot prices from the analysis of forward curves. Moreover, as exhibited empirically in markets as different as the Nordpool, the United Kingdom, or the United States, electricity forward curve moves are much less dramatic than spot price changes.

If we turn to the wide literature dedicated to commodity prices in general, we observe that the convenience yield plays an important role in many cases. The interesting concept of convenience yield was introduced for agricultural commodities in the seminal work by Kaldor (1939) and Working (1949) to represent the benefit from holding the commodity as opposed to a forward contract. Our view is that a convenience yield does not really make sense in the context of electricity: since there is no available technique to store power (outside of hydro), there cannot be a benefit from holding the commodity, nor a storage cost. Hence, the spot price process should contain by itself most of the fundamental properties of power, as listed below.

A first characteristic of electricity (and other commodity) prices is mean reversion toward a level that represents marginal cost and may be constant, periodic, or periodic with a trend. Pindyck (1999) analyzes a 127-year period for crude oil and bituminous coal and a 75-year period for natural gas. He concludes that prices that deflated (and that are represented by their natural logarithms) exhibit mean reversion to a stochastically fluctuating trend line. In the case of power and with a few years' horizon in mind, we propose to represent the diffusion part of the price process as mean reverting to a deterministic periodical trend driven by seasonal effects. As we shall see, the mean reversion will be more or less pronounced across different markets.

A second feature of the price process, unsurprisingly, is the existence of small random moves around the average trend, which represent the temporary supply/demand imbalances in the network. This effect is locally unpredictable and may be represented by a white-noise term affecting daily price variations.

A third and intrinsic feature of power price processes is the presence of so-called spikes, namely, one (or several) upward jump shortly followed by a steep downward move, for instance, when the heat wave is over or the generation outage is resolved. Since shocks in power supply and demand cannot be smoothed away by inventories, our view is that these spikes will persist beyond the transition phase of power deregulation. The California situation has been widely discussed over the last two years, but many studies neglect to mention that the first event of this nature was totally unrelated to the possible exercise of market power by some key providers: in the East Center Area Reliability (ECAR) region (covering several midwestern states of the United States), prices in June 1998 went to several thousand dollars up from $25 per megawatt-hour. This spectacular rise was due to the conjunction of a long heat wave, congestion in transmission of hydroelectricity coming from Canada, and production outage of a nuclear plant. Within two days, prices fell back to a $50 range as the weather cooled down and transmission capacity was restored. In Europe, where weather events are usually less dramatic than in the United States (and capacity reserves probably higher), prices went from 25 to 500 on the Leipzig Exchange during a long cold spell in December 2001. From an economic standpoint, this phenomenon is illuminated by the graph of the marginal cost of electricity supply, called power stack function (see Eydeland and Geman 1998). Knowing the characteristics of the different plants in a given region, one can build the supply function by stacking the units in "merit order" from the lowest to the highest cost of production. The part corresponding to low-cost plants (coal-fired or hydro) is fairly flat or has a small upward slope; then the curve reaches a point at which there is an exponential price increase corresponding to the activation of very expensive units such as "peakers."

Figure 1 represents the merit order stack for the ECAR region and shows the electricity price as determined by the intersection point of the aggregate demand and supply functions. A forced outage of a major power plant or a sudden surge in demand due to extreme weather conditions would either shift the supply curve to the left or lift up the demand curve, in both cases causing a price jump. These spikes are a major subject of concern for practitioners and a key characteristic of electricity prices. Hence, they deserve to be the subject of a careful analysis.

Following the jump diffusion model proposed by Merton (1976) to account for discontinuities in stock price trajectories, a number of authors have introduced a Poisson component to represent the large upward moves of electricity prices; then, the question of bringing prices down is posed. Deng (1999) introduces a sequential regime-switching representation that may be a good way of addressing the dramatic changes in spot prices; the trajectories pro-

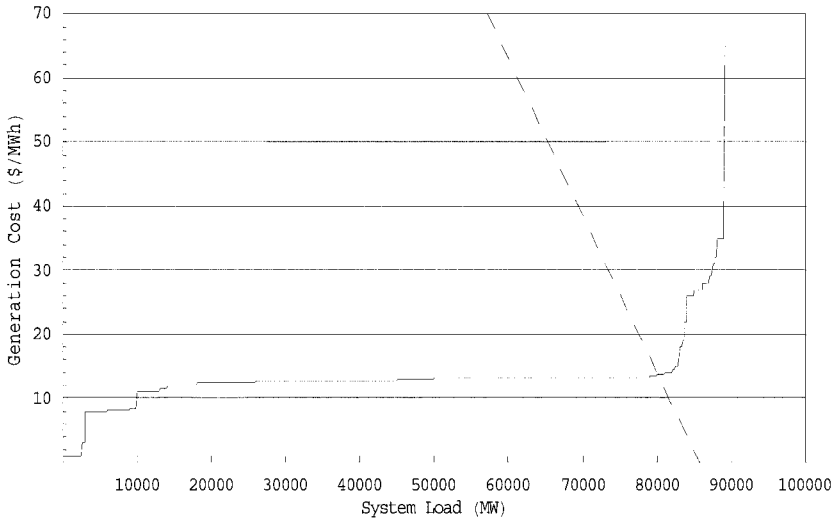

FIG. 1.—The power stack function for the ECAR market. The generation cost is mildly increasing until a load threshold is reached; then the supply curve exhibits strong convexity.

duced by the model, however, are fairly different from the ones observed in the market.

Lucia and Schwartz (2002) examine the Nordpool market and choose not to introduce any jump component in the price process. Data from this market show that despite the significant part of hydroelectricity in the northern part of Europe, power prices do not have continuous trajectories; for instance, there is a quasi-yearly violent downward jump in early April at the end of the snow season when uncertainty about reservoir levels is resolved. This tends to support our view that jump components are hard to avoid when modeling power prices, since they are structurally related to the physical features of this commodity. The class of models presented below is meant to translate the existence of several regimes for electricity prices, corresponding respectively to the quasi-flat and sharply convex parts of the merit order stack. Under the normal regime, the aggregate capacity of generation in the region under analysis is sufficient to meet consumers' demand. In the case of a plant outage, inelastic demand drives spot prices to very high values until the supply problem is resolved; hence the observed large spikes. We can note that in the case of storable commodities (such as oil or wheat), prices are determined not only by supply from existing production and demand for current consumption but also by the level of inventories. The buffering effect of these inventories does not exist in the case of electricity.

We argue in this paper that the classical setting of continuous-path diffusion processes does not deliver a viable solution to this problem for reasons linked

to trajectorial and statistical features of daily power prices. A jump component should account for the occurrence of spikes through an appropriate jump intensity function and also explain the significant deviations from normality in terms of high-order moments observed in logarithmic prices. As an example, figure 2 compares the empirical distribution in the ECAR market to a normal density with the same mean and variance.

III. The Model

We model the behavior of the price process of 1 megawatt-hour of electricity traded in a given pool market.[1] In order to ensure strict positivity of prices and enhance the robustness of the calibration procedure, we represent the electricity spot price in natural logarithmic scale.[2] Throughout the paper, except for the pictures representing trajectories, the term price will refer to "log price."

The spot price process is represented by the (unique) solution of a stochastic differential equation of the form

$$dE(t) = D\mu(t)dt + \theta_1[\mu(t) - E(t^-)]dt + \sigma dW(t) + h(t^-)dJ(t), \qquad (1)$$

where D denotes the standard first-order derivative and $f(t^-)$ stands for the left limit of f at time t.

The deterministic function $\mu(t)$ represents the predictable seasonal trend of the price dynamics around which spot prices fluctuate. The second term ensures that any shift away from the trend generates a smooth reversion to the average level $\mu(t)$. The positive parameter θ_1 represents the average variation of the price per unit of shift away from the trend $\mu(t)$ per unit of time. Note that the speed of mean reversion depends on the current electricity price level since the constant θ_1 is multiplied by $\mu(t) - E(t^-)$, a difference that may be quite large in electricity markets (in contrast to interest rates or stochastic volatility models for which mean reversion is classically present). The process W is a (possibly n-dimensional) standard Brownian motion representing unpredictable price fluctuations and is the first source of randomness in our model. The constant σ defines the volatility attached to the Brownian shocks. Note that the instantaneous squared volatility of prices is represented by the conditional second-order moment of absolute price variations over an infinitesimal period of time: in the present context, it is the sum of the squared Brownian volatility and a term generated by the jump component (see, e.g., Gihman and Skorohod 1972).

The discontinuous part of the process reproduces the effect of periodically recurrent spikes. A spike is a cluster of upward shocks of relatively large size with respect to normal fluctuations, followed by a sharp return to normal price

1. In most markets, this price for date t is defined the day before by the clearing of buy and sell orders placed in the pool.

2. Up to now, negative electricity prices have rarely been observed.

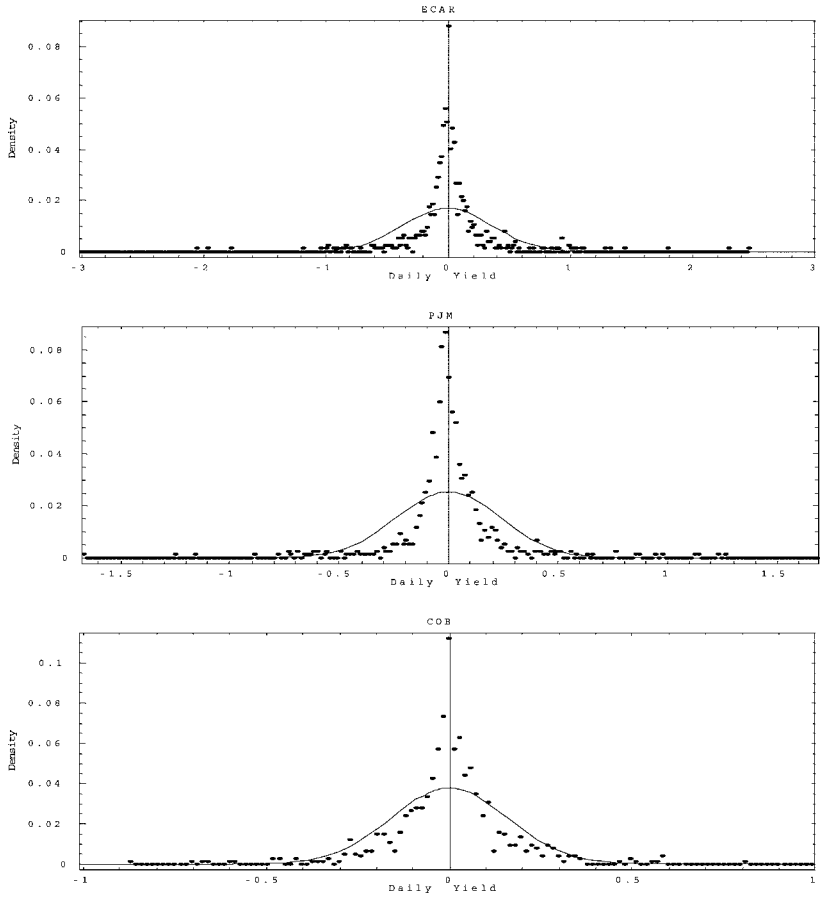

FIG. 2.—Empirical price returns distributions vs. normal distributions with equal means and variances. For each market, the empirical density of price returns is reported together with a normal density matching the first two moments. All markets display strong deviations from normality due to the presence of upward and downward jumps.

levels. We represent this behavior by assigning a level-dependent sign for the jump component. If the current price is below some threshold, prices are in the normal regime and any forthcoming jump is upward directed. If instead the current price is above the threshold, the market is experiencing a period of temporary imbalance between demand and supply reflected by abnormally high prices, and upcoming jumps are expected to be downward directed.

Jumps are characterized by their time of occurrence, size, and direction. The jump times are described by a counting process $N(t)$ specifying the number of jumps experienced up to time t. There exists a corresponding intensity

process ι defining the instantaneous average number of jumps per time unit. We choose for ι a deterministic function that we write as

$$\iota(t) = \theta_2 \times s(t), \tag{2}$$

where $s(t)$ represents the normalized (and possibly periodic) jump intensity shape, and the constant θ_2 can be interpreted as the maximum expected number of jumps per time unit.

The jump sizes are modeled as increments of a compound jump process $J(t) = \sum_{i=1}^{N(t)} J_i$. Here the J_i's are independent and identically distributed random variables with common density

$$p(x;\ \theta_3,\ \psi) = c(\theta_3) \times \exp[\theta_3 f(x)], \quad 0 \le x \le \psi, \tag{3}$$

where $c(\theta_3)$ is a constant ensuring that p is a probability distribution density, and ψ is the maximum jump size. The choice of a truncated density within the exponential family is meant to properly reproduce the observed high-order moments.

The jump direction determines the algebraic effect of a jump size J_i on the power price level. It is represented by a function h, taking values plus one and minus one according to whether the spot price $E(t)$ is smaller or greater than a threshold $\mathcal{T}$:

$$h(E(t)) = \begin{cases} +1 & \text{if } E(t) < \mathcal{T}(t) \\ -1 & \text{if } E(t) \ge \mathcal{T}(t). \end{cases} \tag{4}$$

This function plays an important role in our model for two sets of reasons related to the trajectorial properties of the process and the descriptive statistics of daily price returns. Some authors have proposed to model spikes by introducing large positive jumps together with a high speed of mean reversion, in particular Deng (1999), who was among the first ones to address the specific features of electricity prices. However, models with upward jumps only are deemed to display a highly positive skewness in the price return distribution, in contrast to the one observed in the markets. Other authors model spikes by allowing signed jumps (e.g., Escribano et al. 2002), but if these jumps randomly follow each other, the spike shape has obviously a very low probability of being generated. Finally, another type of solution proposed in particular by Huisman and Mahieu (2001) and Barone-Adesi and Gigli (2002) is the introduction of a regime-switching model. This representation does not allow the existence of successive upward jumps; moreover, a return to normal levels through a sharp downward jump would require in this case a non-Markovian specification. As a consequence, calibrating a regime-switching model is often quite problematic.

In our setting a proper choice of the barrier $\mathcal{T}$ coupled with a high jump intensity can generate a sequence of upward jumps leading to high price levels, after which a discontinuous downward move together with the smooth mean reversion brings prices down to a normal range. Moreover, our representation

has the merit of preserving the Markov property in a single state variable (see Roncoroni 2002).

Let us observe that general results about stochastic differential equations of the proposed type ensure that equation (1) admits a unique solution (see Gihman and Skorohod 1972). Hence, the level-dependent signed-jump model with time-varying intensity is fully described.

IV. Electricity Data Set

We calibrate the model on a data set consisting of a series of 750 daily average prices compiled from the publication *Megawatt Daily* for three major U.S. power markets: COB (California Oregon Border), PJM (Pennsylvania–New Jersey–Maryland), and ECAR (East Center Area Reliability coordination agreement). These markets may be viewed as representative of most U.S. power markets because of their various locations (California, East Coast, and Midwest), because of the different mix of generation (e.g., an important share of hydroelectricity in California), and, finally, because of the type of transmission network servicing the region. Moreover, the market design in ECAR and PJM has proved to have functioned properly so far; the choice of the period of analysis (ending in 1999) was meant to leave aside the California crisis and its effects on the COB pool. In terms of price behavior, the COB market is typical of "low-pressure" markets (such as Palo Verde, Mid Columbia, and Four Corners), with high prices ranging between \$90 and \$115 per megawatt-hour in the examined period. The PJM market represents a "medium-pressure" market (such as western New York, eastern New York, and Ercott), with highs between \$263 and \$412 per megawatt-hour during that period. Finally, the ECAR market portrays "high-pressure" markets (such as the Manchester Area Aggregation Program, Georgia-Florida Border, North and South Southwest Power Pool, and Middlewest Area Interregion Network), experiencing spikes between \$1,750 and \$2,950 per megawatt-hour.

Figures 3, 4, and 5 depict *absolute* historical price paths in these markets for the period between January 6, 1997, and December 30, 1999. As stated earlier, our goal is to adjust our class of processes to both trajectorial features (i.e., average trend, Brownian volatility, periodical component, and spikes) and statistical features (i.e., mean, variance, skewness, and kurtosis of daily price returns) of historical prices.

In order to start the calibration procedure, we need to detect jumps in the raw market data. The estimation of a mixed jump diffusion over a discrete sequence $\mathbf{E} = (E_1, \ldots, E_n)$ of observations may result in an ill-posed problem: standard methods in statistical inference require samples to represent whole paths over a time interval. In the case of discretely sampled observations, there are infinitely many ways a given price variation over a discrete time interval can be split into an element stemming from the continuous part of the process and another from the discontinuous one. Hence, the problem of

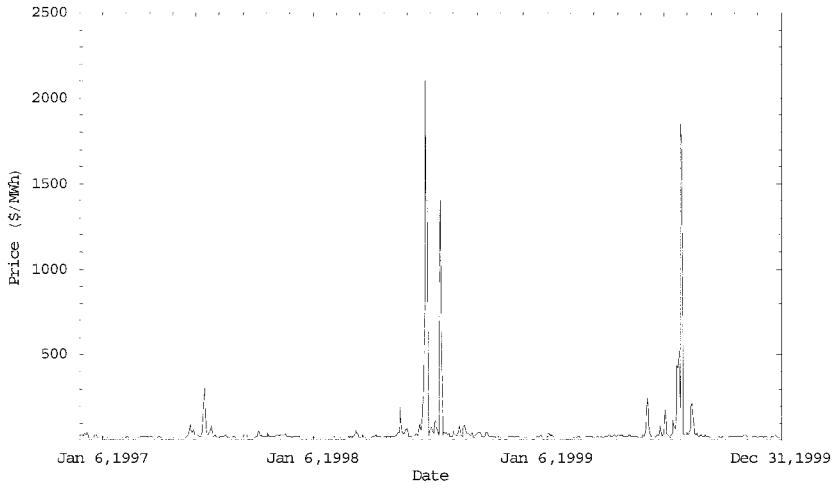

FIG. 3.—ECAR price path (January 6, 1997–December 30, 1999). Spikes concentrate in summer, when prices may rise as high as US$2,000 per kilowatt-hour.

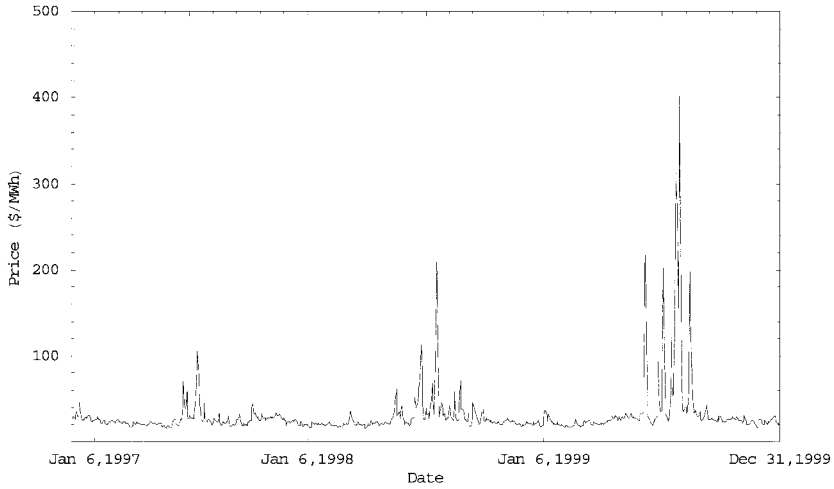

FIG. 4.—PJM price path (January 6, 1997–December 30, 1999). Spikes concentrate in summer, when prices move up to a level of US$400 per kilowatt-hour.

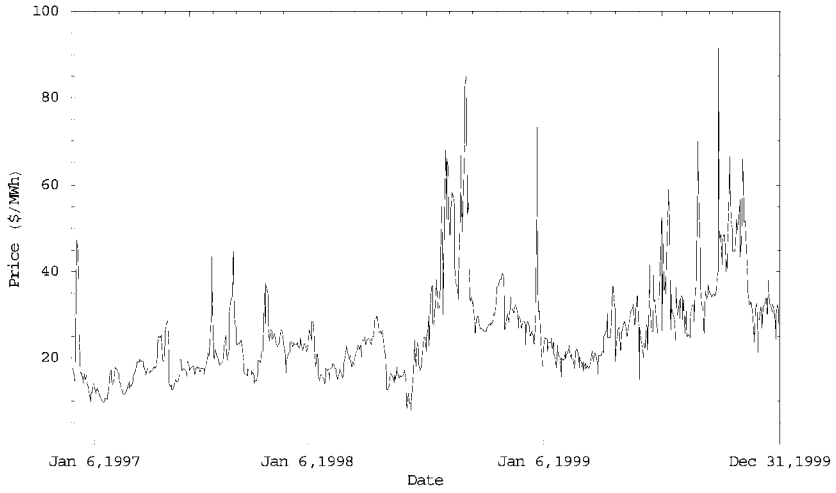

FIG. 5.—COB price path (January 6, 1997–December 30, 1999). Spikes concentrate in summer, when prices rise to values around US$100 per kilowatt-hour.

disentangling these two components on a discrete sample cannot be resolved in a theoretically conclusive way; yet, the situation is better in a continuous-time representation, which is our case. All examined data exhibit excess kurtosis in the empirical distribution of daily price variations. These changes tend to cluster close to either their average mean or the largest observed values (see fig. 2). In other words, data suggest that either there is a jump, in which case the variation due to the continuous part of the process is negligible, or there is no jump and the price variation is totally generated by the continuous part of the process. This observation leads us to identify a price change threshold Γ allowing one to discriminate between the two situations. In this order, we extract from the observed data set two important elements of the calibration procedure: the set $\Delta\mathbf{E}^d$ of sampled jumps and the "Γ-filtered" continuous sample path $\mathbf{E}^c$ obtained by juxtaposition of the continuous variations starting at the initial price.[3] A discussion of possible selection schemes in a general mathematical setting may be found in Yin (1988). We include Γ as a parameter to be estimated within the calibration procedure: for each market under investigation, we perform our calibration procedure over different Γ-filtered data sets for values of Γ chosen in the set of observed daily price variations. Then we select the value of Γ leading to the best calibrated model in view of its ability to match descriptive statistics of observed daily price

3. If t is a jump date, the continuous part of the path is assumed to be constant between t and the next sample date. Since spikes are rare and typical price variations are much smaller than those occurring during a spike, this simplification does not introduce any significant bias in the estimation procedure.

variations. From now on, we suppose that a value for Γ has been identified for the market under analysis and input data are described by the corresponding Γ-filtered pair $(\mathbf{E}^c, \Delta\mathbf{E}^d)$.

V. Calibration

We propose a two-step calibration procedure. The first step is the assignment of a specific form for the "structural" elements in the dynamics described in equation (1) and defined as (1) the mean trend $\mu(t)$, (2) the jump intensity shape $s(t)$, (3) the threshold $\mathcal{T}$ defining the sign of the jump, and (4) the jump size distribution $p(x)$. These quantities translate into path properties of the price process.

The second step consists in statistically estimating the four parameters of the selected model, namely, (1) the mean reversion force θ_1, (2) the jump intensity magnitude θ_2, (3) the jump size distribution parameter θ_3, and (4) the Brownian volatility σ. The resulting parametric model is fit to the filtered prices by a new statistical method described further on and based on likelihood estimation for continuous-time processes with discontinuous sample paths.

We now illustrate the implementation of this calibration procedure for the ECAR market, propose possible alternatives to the resulting model, and defer results and comments to the next section.

A. Selection of the Structural Elements

The mean trend $\mu(t)$ can be determined by fitting an appropriate parametric family of functions to the data set. As mentioned earlier, power prices exhibit a weak seasonality in the mean trend and a sharper periodicity in the occurrences of turbulence across the year. The latter periodicity may be an effective estimate for the one displayed by mean trend: for instance, ECAR market data show price pressure once a year, during the warm season. Some markets display price pressure twice a year, with winter average prices lower than summer average prices (which requires a lower local maximum in the former case). In general, we find that a combination of an affine function and two sine functions with, respectively, a 12-month and a six-month periodicity is appropriate for the U.S. historical data under investigation. We accordingly define the mean trend by a parametric function:

$$\mu(t;\ \alpha,\ \beta,\ \gamma,\ \delta,\ \varepsilon,\ \zeta) = \alpha + \beta t + \gamma \cos(\varepsilon + 2\pi t) + \delta \cos(\zeta + 4\pi t). \quad (5)$$

The first term may be viewed as a fixed cost linked to the production of power. The second one drives the long-run linear trend in the total production cost. The overall effect of the third and fourth terms is a periodic path displaying two maxima per year, of possibly different magnitudes. Observed prices over the three-year period are averaged into a one-year period and bounded from above by a suitable quantile ν of their empirical distribution. The trend function μ is fitted to the resulting average data by a sequential

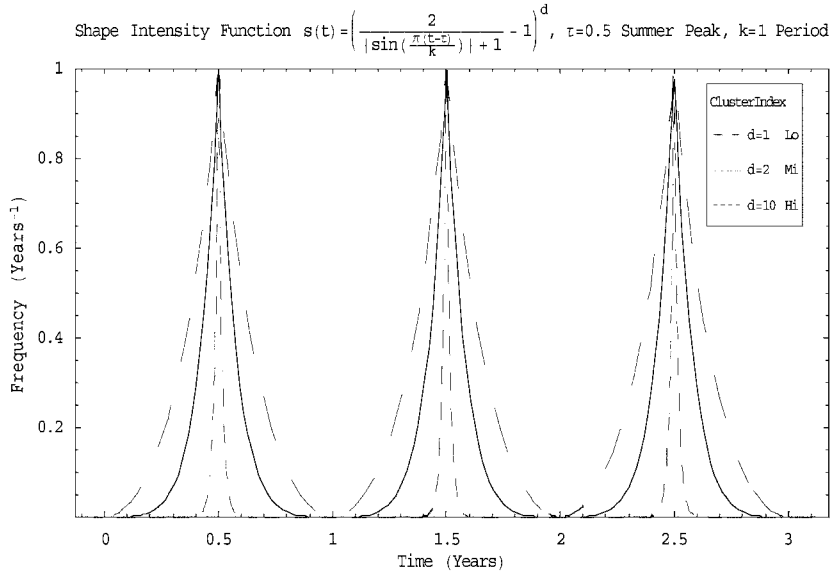

FIG. 6.—Time-dependent jump intensity function. The time-dependent jump intensity function is designed to concentrate jump occurrence during the warm season. Parameter d drives the degree of cluster.

ordinary least squares (OLS) method providing parameters α, β, γ, δ, ε, and ζ.

We now turn to the identification of the jump intensity shape s. Since spikes occur over short time periods, we select an intensity function exhibiting pronounced convex peaks with annual periodicity. This is meant to reflect the shape of the power stack function, which, as shown in figure 1, becomes very convex (and quasi-vertical) at some demand level. Sharp convexity also ensures that the occurrences of price jumps tend to cluster around the peak dates and rapidly fade away. In this order we choose

$$s(t) = \left\{\frac{2}{1 + |\sin[\pi(t-\tau)/k]|} - 1\right\}^{d}. \tag{6}$$

Here the jump occurrences exhibit peaking levels at multiples of k years, beginning at time τ.[4] For instance, price shocks concentrating twice a year at evenly spaced dates, with a maximum on August 1, are recovered by the choice $\tau = 7/12$ and $k = 1/2$. The exponent d allows us to adjust the dispersion of jumps around peaking times and is included among parameters to be estimated within the calibration procedure. Figure 6 shows intensity functions across different coefficients d, and figure 7 reports a sample of jump times.

4. Time τ is called "the phase" in the language of sinusoidal phenomena.

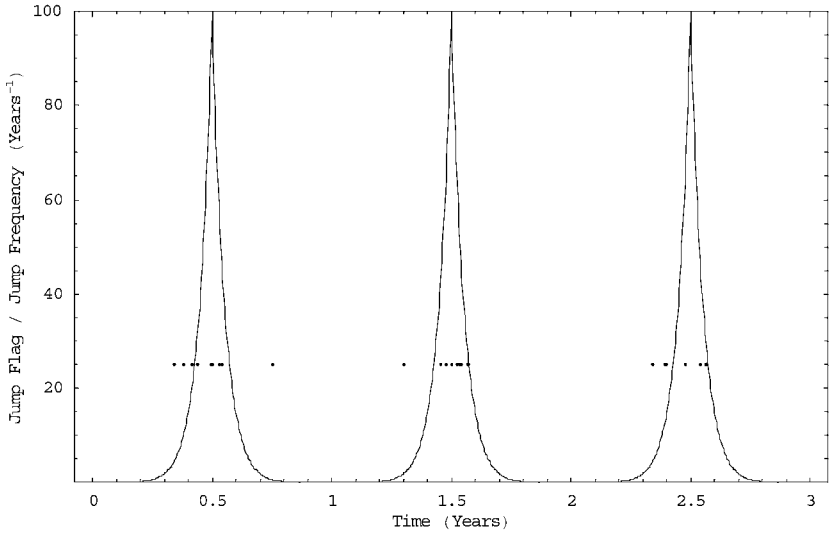

FIG. 7.—Sample jumps of a time-dependent jump intensity function. The time-dependent jump intensity function is designed to concentrate jump occurrence during the warm season. Dotted tags signal the sample jump times of a Poisson process corresponding to the displayed time-dependent intensity function.

We found that in all three examined markets the best value for d is two. We have discussed earlier the introduction of a barrier $\mathcal{T}$ above which all occurring jumps are downward directed. This threshold may reasonably be defined by a constant spread Δ over the selected average trend:

$$\mathcal{T}(t) = \mu(t) + \Delta. \tag{7}$$

The choice of Δ results from a balance between two competing effects: the greater the value of Δ, the higher the level power prices may reach during pressure periods; the smaller this value, the sooner the downward jump effect moves toward normal levels. Equally important, this number has an impact on the moments of daily price variations. We select Δ in such a way that the corresponding calibrated model generates paths whose average maximal values equal the maximal prices observed in the market under analysis.

The last structural element to be determined is a probability distribution for the jump size. We select a truncated version of an exponential density with parameter θ_3:

$$p(x;\ \theta_3,\ \psi) = \frac{\theta_3 \exp(-\theta_3 x)}{1 - \exp(-\theta_3 \psi)}, \quad 0 \le x \le \psi, \tag{8}$$

where ψ represents an upper bound for the absolute value of price changes. This distribution belongs to the family described in equation (3), where

$c(\theta_3) = \theta_3/[1 - \exp(-\theta_3\psi)]$ and $f(x) = -x$. The resulting price process is a "special semimartingale," a property required to obtain the statistical estimator proposed in the next section. This completes the first calibration step.

B. Model Parameter Estimation

The issue of estimating discontinuous processes has been the subject of particular attention in the financial econometric literature. The proposed methods mainly draw on the extension of statistical techniques well established in the case of continuous processes. Beckers (1981), Ball and Torous (1983), and Lo (1988) develop estimators based on moment matching; Johannes (1999) and Bandi (2000) propose nonparametric methods based on higher-order conditional moments of instantaneous returns. We choose to focus on maximum likelihood methods. The transition densities they typically require can rarely be computed in analytical terms; in our case, the mixed effect of continuous and jump terms makes the task even more arduous since one has to deal with mixtures of probability distributions. Several numerical devices have been recently proposed in order to overcome these difficulties. Broadly speaking, these methods start by discretizing the process and then computing approximated versions of the targeted transition densities. Pedersen (1995) explores simulation-based schemes, and Andersen, Benzoni, and Lund (2002) make use of auxiliary model approximations. Unfortunately, all these methods suffer from computational complexity because of the necessary double approximation of the process and of the transition densities.

We propose an estimator based on the exact likelihood of the unknown process with respect to a prior process chosen as a reference within the same class. Plugging a piecewise constant sample path agreeing with actual data at the sample dates into this likelihood delivers an approximated likelihood function process. The estimator is provided by the parameter vector maximizing this process over a suitable domain. This method has two major advantages: first, the analytical form of the exact likelihood function under continuous-time observations can be computed for nearly all semimartingales through a generalized version of the Girsanov theorem (see Roncoroni 2002). Second, the discrete sample estimator converges to the continuous sample one, and a well-established estimation theory exists in this latter case. We now explain the details of the procedure.

We compute the log likelihood function $\mathcal{L}$ for the law of the diffusion process corresponding to an arbitrary parameter vector $\boldsymbol{\theta}$ with respect to the law of the process under a prior reference parameter $\boldsymbol{\theta}^0$. Its exact analytical expression is derived in Appendix A. We decide to choose as starting parameter values $\theta_1 = 0$, $\theta_2 = 1$, and $\theta_3 = 1$, which correspond to an absence of drift, a normalized jump intensity, and a jump amplitude drawn from a truncated exponential distribution with parameter one.

The approximate logarithmic likelihood function reads as

$$\begin{aligned}\mathcal{L}_D(\boldsymbol{\theta}|\boldsymbol{\theta}^0, \mathbf{E}) = &\sum_{i=0}^{n-1} \frac{[\mu(t_i) - E_i]\theta_1}{\sigma^2}[\Delta E_i^c - D\mu(t_i)\Delta t] \\ &- \frac{\Delta t}{2}\sum_{i=0}^{n-1}\left\{\frac{[\mu(t_i) - E_i]\theta_1}{\sigma}\right\}^2 \\ &- (\theta_2 - 1)\sum_{i=0}^{n-1} s(t_i)\Delta t + \log\theta_2 N(t) \\ &+ \sum_{i=0}^{n-1}\left[-(\theta_3 - 1)\frac{\Delta E_i^d}{h(E_i)}\right] + N(t)\log\left[\frac{1 - e^{-\theta_3\psi}}{\theta_3(1 - e^{-\psi})}\right],\end{aligned}$$

where $D\mu(t_i)$ denotes the first-order derivative of μ at time t_i. The first part is a discretized version of the Doléan-Dade exponential for continuous processes. The remaining terms come from the jump part of the process. The log likelihood function explicitly depends on θ_1, θ_2, and θ_3 and on the filtered data set ($\Delta\mathbf{E}^d$, $\mathbf{E}^c$), which in turn is derived from the original market data set $\mathbf{E}$ and the choice of parameter $\boldsymbol{\Gamma}$. We maximize this function with respect to $\boldsymbol{\theta}$ over a bounded parameter set $\boldsymbol{\Theta}$ identified through economic interpretation of the model parameters. One may alternatively use Monte Carlo simulated samples to infer a reasonable parameter domain and starting values for the numerical optimization algorithm. We finally obtain a nonlinear maximization program of a continuous function over a compact set, and classical theorems ensure the existence of a local maximum, which will be our estimate for $\boldsymbol{\theta}^*$.

The constant Brownian volatility over observation dates $0 = t_0 < t_1 < \cdots < t_n = t$ can be obtained as

$$\sigma = \sqrt{\sum_{i=0}^{n-1} \Delta\overline{E}(t_i)^2}, \tag{9}$$

where each summand $\Delta\overline{E}(t_i)^2$ represents the square of the continuous part $\Delta E^c(t_i)$ of observed price variations (in a logarithmic scale) between consecutive days t_i and t_{i+1}, net of the mean reversion effect $|\theta_1 \times [\mu(t_i) - E(t_i)]|$ for all times t_i in which no jump occurs (i.e., $\Delta E^d(t_i) = 0$).[5] This estimator converges to the exact local covariance estimator for diffusion processes under continuous-time observations (Genon-Catalot and Jacod 1993). We note that numerical experiments not reported here suggest that a time-dependent volatility does not produce a significant improvement in the estimated process (given the other specifications of our model); moreover, in this case a joint

5. Note that in contrast to classical settings in which the mean reversion feature was introduced (e.g., interest rates, stochastic volatility), the difference $\mu(t) - E(t)$ may be quite large in the case of electricity prices. This observation was made in Sec. III of the paper.

estimation of volatility and mean reversion parameters would become necessary.

C. Alternative Model Specifications

We now consider two models displaying in their discontinuous component features either proposed in existing papers on electricity or that may be envisioned as improvements of some kind.

First, by setting to plus one the jump direction function h defined in formula (4), we obtain a restricted model in which upcoming jumps are all upward directed and reversion to normal levels is exclusively carried over by the smooth drift component. This upward-jump model represents the classical jump diffusion extension of the continuous diffusion models proposed over the years by Pilipovich (1997), Barlow (2002), and Lucia and Schwartz (2002). All the remaining model specifications are the same as those of our signed-jump model. As a consequence, calibration to market data follows the steps described above, with one major exception: price variations of negative size all enter the estimation of the continuous part of the process (i.e., $\Delta \mathbf{E}^d$ contains only positive jumps).

Alternatively, we may allow the jump intensity function ι defined in formula (2) to be stochastic. In order to account for the dependence of the likelihood of jump occurrence on the price level following upward shocks, we consider the following function of the spot price and time:

$$\iota(t,\, E(t^-)) = \theta_2 \times s(t) \times (1 + \max\{0,\, E(t^-) - \overline{E}(t)\}). \tag{10}$$

As in the case of the threshold $\mathcal{T}(t)$ defining the sign of the jump, we define $\overline{E}(t)$ as a constant c over the mean trend μ. If the spot price is below the mean trend μ plus this spread c, then intensity is purely time dependent. Each price unit beyond this boundary amplifies accordingly the time-dependent intensity. We determined that the best intensity function was provided by a choice of c equal to $\Delta/2$ (i.e., an increasing jump occurrence when prices are above the median line between the mean trend $\mu(t)$ and the threshold $\mathcal{T}(t)$). The "max" function ensures that the jump intensity never goes below the "standard level" $\theta_2 \times s(t)$ (which may be viewed as the effect of random outages that strike power plants). This effect is depicted in figure 8, where stochastic intensity is displayed as a function of time and log price. In this signed-jump model, jump occurrence is both time and level dependent. Because all the other model specifications are the same as those in the signed-jump model with deterministic intensity, calibration to market data follows the same steps as de-

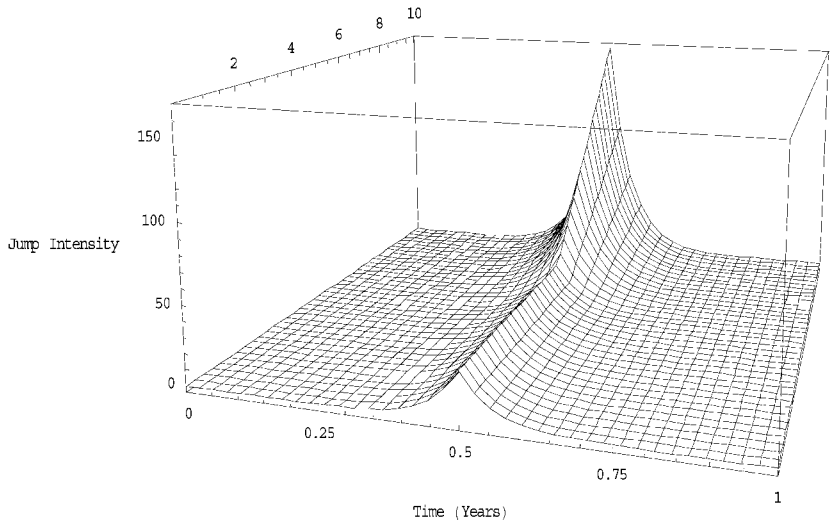

FIG. 8.—Stochastic jump intensity function. Jump intensity depends on time and electricity price level. If the spot price is below the mean trend μ plus the spread $\Delta/2$, then intensity is only time dependent. Each price unit beyond this boundary amplifies accordingly the time-dependent intensity.

scribed above. The log likelihood estimator, denoted in this case as $\mathcal{L}_R$, becomes slightly more complex:

$$\begin{aligned}\mathcal{L}_S(\boldsymbol{\theta}|\boldsymbol{\theta}^0, \mathbf{E}) = &\sum_{i=0}^{n-1} \frac{[\mu(t_i) - E_i]\theta_1}{\sigma^2}[\Delta E_i^c - D\mu(t_i)\Delta t] \\ &- \frac{\Delta t}{2}\sum_{i=0}^{n-1}\left\{\frac{[\mu(t_i) - E_i]\theta_1}{\sigma}\right\}^2 \\ &- (\theta_2 - 1)\sum_{i=0}^{n-1} s(t_i)\max\{1, E_i - \mu(t_i) - c\}\Delta t + \log\theta_2 N(t) \\ &+ \sum_{i=0}^{n-1}\left[-(\theta_3 - 1)\frac{\Delta E_i^d}{h(E_i)}\right] + N(t)\log\left[\frac{1 - e^{-\theta_3\psi}}{\theta_3(1 - e^{-\psi})}\right].\end{aligned}$$

This expression shows that parameters θ_1 and θ_3 are unaffected by a change in the jump intensity function since the corresponding term can be factored out of the likelihood estimator in absolute scale $\exp(\mathcal{L}_S)$.

VI. Empirical Results

The calibration procedure has been implemented on the U.S. data set described in Section IV. We first present results for the signed-jump model, then discuss the quality of our assessments on the data set under analysis, and finally

TABLE 1 Estimated "Structural" Elements

	Interpretation	ECAR	PJM	COB
α	Average log price level	3.0923	3.2002	2.8928
β	Average log price slope	.0049	.0036	.1382
γ	Yearly trend	−.1300	.0952	.1979
δ	6-month trend	.0292	.0217	.0618
ε	Yearly shift	.3325	2.4383	1.7303
ζ	6-month shift	.7417	.2907	1.7926
ν	0.7 average distribution quantile	3.2762	3.3232	3.3586
Δ	Jump regime level	2.5000	1.5000	1.0000
ψ	Maximum jump size	3.3835	1.6864	1.0169
k	Jump periodicity	1.0000	1.0000	1.0000
τ	Intensity phase	.5000	.5000	.5000

NOTE.—The electricity log price model, eq. (1), with average trend function (5) and jump component (4) (direction), $J(t) = \sum_{i=1}^{N(t)} J_i$, with $J_i \overset{i.i.d.}{\sim} p(x;\ \theta_2,\ \psi) \propto e^{\theta_3 s(x)}$, $0 \le x \le \psi$ (size), and eq. (10) (intensity) is calibrated to a data set including daily observations between January 6, 1997, and December 30, 1999. Observed log prices over the three-year period are averaged into a one-year period and bounded from above by the 0.7 quantile ν of their empirical distribution. The trend function μ is fitted to the average data by a sequential OLS providing parametersα, β, γ, δ, ε, and ζ. The regime-switching threshold $\mathcal{T}$ is set as a spread Δ over the average trend μ. The jump size distribution takes values in the interval $[0, \psi]$, where ψ is chosen as the observed maximal daily absolute variation in log prices. The shape of the jump intensity is described through the parameters k and τ.

conclude on a comparison with the alternative models introduced at the end of Section V.

As mentioned before, the first step is the functional estimation of the structural elements μ, s, $\mathcal{T}$, and p. The values α, β, γ, δ, ε, and ζ characterizing the average trend function $\mu(t)$ defined in formula (5) are reported in table 1. The jump intensity shape $s(t)$ has the form defined in equation (6), with $k = 1$, $\tau = 0.5$, and $d = 2$; this corresponds to a jump occurrence displaying an annual peak strongly clustered around the middle of the year, as observed in all examined markets. The threshold $\mathcal{T}(t)$ is defined by a spread Δ over the deterministic trend $\mu(t)$, where Δ is chosen in the order of 50% of the range spanned by the observed log prices. We observe that both ECAR and PJM reveal no significant linear trend over the three-year sample period, whereas COB shows a small positive linear trend expressed by the coefficient β.

In all cases, the annual periodicity expressed by the coefficient γ prevails over the semiannual component described by the coefficient δ. Figure 9 represents the average paths for the three regional markets in a joint graph; clearly, the annual component is predominant in the COB market, whereas an additional semiannual component is significant in the ECAR and PJM markets. A clear difference between the three markets is represented by the maximum size ψ of daily price variations: for instance, ECAR displays jumps that may be more than three times greater than the maximum value observed in the COB market. In this market, the high percentage of hydrogeneration and the reservoir capacity allow it to go through the year—the cold season in particular—with no or mild spikes. In contrast, the PJM and ECAR markets experience both very warm summers and cold winters; this leads to the semiannual periodicity of observed power prices in these regions. However, PJM

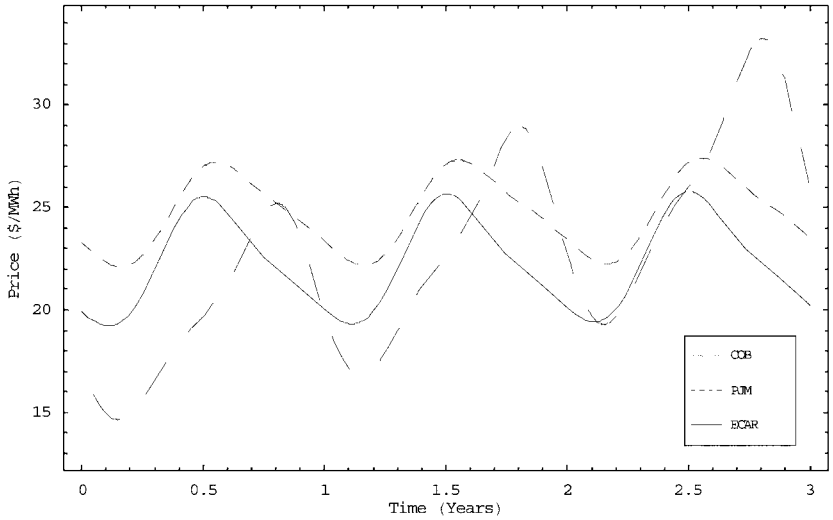

FIG. 9.—Estimated average trends in the observed log price paths (January 6, 1997–December 30, 1999). The PJM and ECAR markets exhibit overlapping periodicities with periods equal to six and 12 months. COB essentially displays an annual periodicity.

benefits from a fuel mix in power generation and also from a rich transmission network that has been very efficient since the start of deregulation; hence the less dramatic price spikes observed.

The second step of the calibration procedure is the statistical estimation of parameters θ_1, θ_2, θ_3, σ, Γ, and d. The approximated likelihood estimation detailed in the previous section has been implemented by the Levenberg-Marquardt nonlinear maximum search algorithm. Final results are reported in table 2.

All markets exhibit some amount of smooth mean reversion. Note that the value of the reversion force in PJM is significantly greater than that in ECAR. It is worth emphasizing again that the overall reversion displayed by our model is created by the joint effect of the classical mean reversion and an effect due to the downward jumps. Since ECAR displays more jumps than PJM, the overall reversion effect is higher than the one observed in the PJM market. This is statistically consistent with the fact that, in PJM, both skewness and kurtosis of daily price increments are lower since the smooth reversion suffices most of the time to ensure return to the average trend. We remark that the expected number of jumps per year is represented by the integral of the calibrated intensity function over one year.

We now turn to the assessment of the quality of the estimated processes. This is performed according to four criteria:

1. We analyze simulated sample paths together with empirically observed

TABLE 2 Estimated Model Parameters

	Interpretation	ECAR	PJM	COB
θ_1	Smooth mean reversion force	38.8938	42.8844	13.3815
θ_2	Maximum expected number of jumps	59.5210	63.9301	13.2269
θ_3	Reciprocal average jump size	.3129	.5016	1.0038
σ	Brownian local volatility	1.8355	1.4453	1.3631
Γ	Jump threshold	.9200	.6000	.6200
$N(1)$	Average number of jumps	9.0000	9.6667	2.0000
n_j	Number of filtered jumps	27	29	6

NOTE.—The model parameters θ_1 (smooth mean-reversion force), θ_2 (maximum expected number of jumps), and θ_3 (reciprocal expected jump size) are selected by an approximated maximum likelihood estimator. The Brownian volatility σ is calculated as a discrete-time observation approximation of the standard cumulated covariance estimator on the continuous path obtained by deleting observations of size larger than Γ. The jump threshold Γ is chosen in such a way that the resulting model matches the fourth moment of the daily log price return distribution. An estimate of the expected number of jumps over one year $N(1)$ is provided by the integral of the intensity function over a one-year period. The quantity n_j denotes the number of observed daily price variations attributed to the jump component of the process according to the selected jump size threshold Γ.

trajectories and make a judgment about the fitting quality of the trajectorial properties.

2. We compare simulated moments of the daily increments distribution with the empirical values displayed by each market under investigation.
3. We check for the robustness of the procedure by reestimating simulated sample paths generated by the calibrated model.
4. We test our model against the most popular representation of electricity spot prices so far, namely, a jump diffusion process with positive jumps only and smooth mean reversion.
5. We examine the effect of introducing a price-dependent jump intensity on both trajectorial and statistical properties displayed by the most irregular market in our data set (ECAR).

Figures 10, 11, and 12 show trajectories of the estimated model for the three markets. For the purpose of comparison, both historical and sample paths are reported at various scales. The dashed line represents the average mean trend $\mu(t)$. These pictures show that the proposed family of processes is capable of reproducing quite consistently the qualitative features exhibited by power paths in all three examined markets.

Table 3 reports the mean, standard deviation, skewness, and excess kurtosis of observed and simulated daily price variations. We see that all statistics of the simulated trajectories are quite satisfactory; there is, however, a small positive skewness that has no counterpart in the empirical data, suggesting that the reverting component ought to be more pronounced. The most important effect of the signed-jump model is the excellent fit of the leptokurtosicity of the distribution. The relevance of the incorporation of jumps in equity return modeling has been analyzed and exhibited in a number of recent papers of the financial economics literature (see, e.g., Carr et al. 2002). In

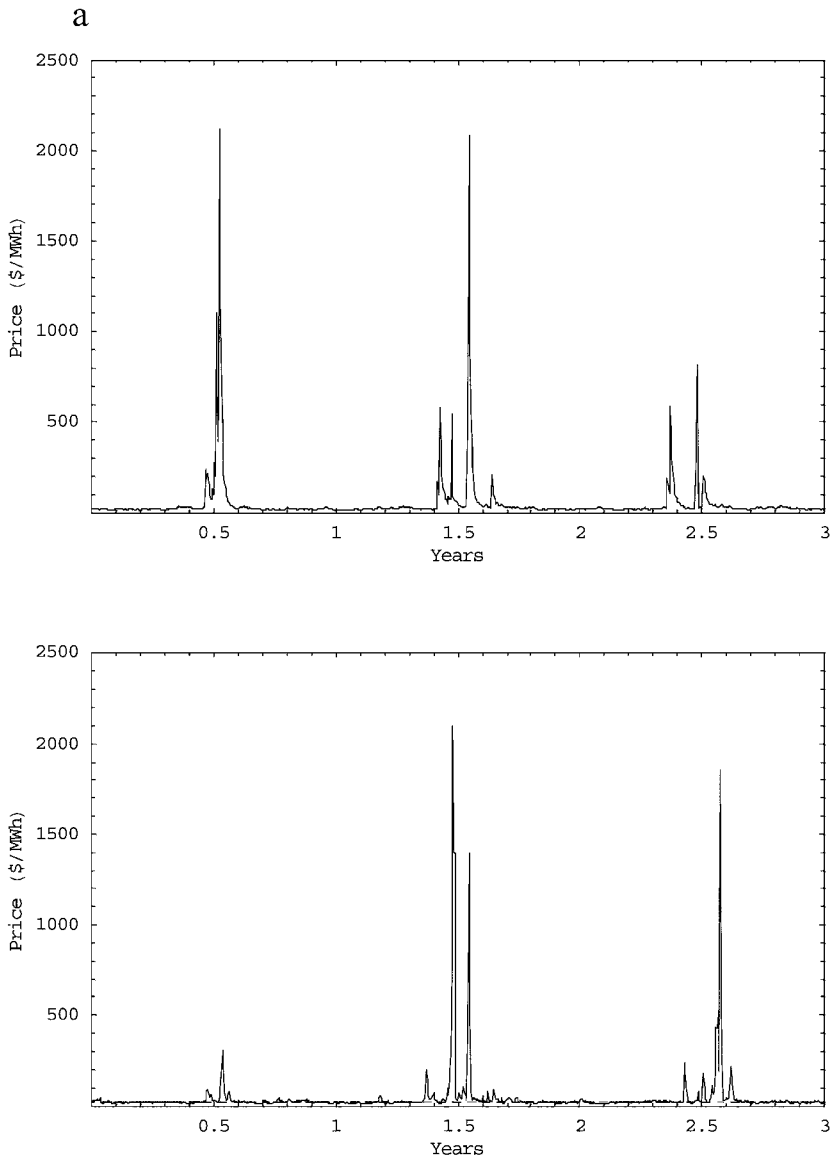

FIG. 10.—ECAR simulated price path vs. empirical path: *a*, absolute scale 0–2,500

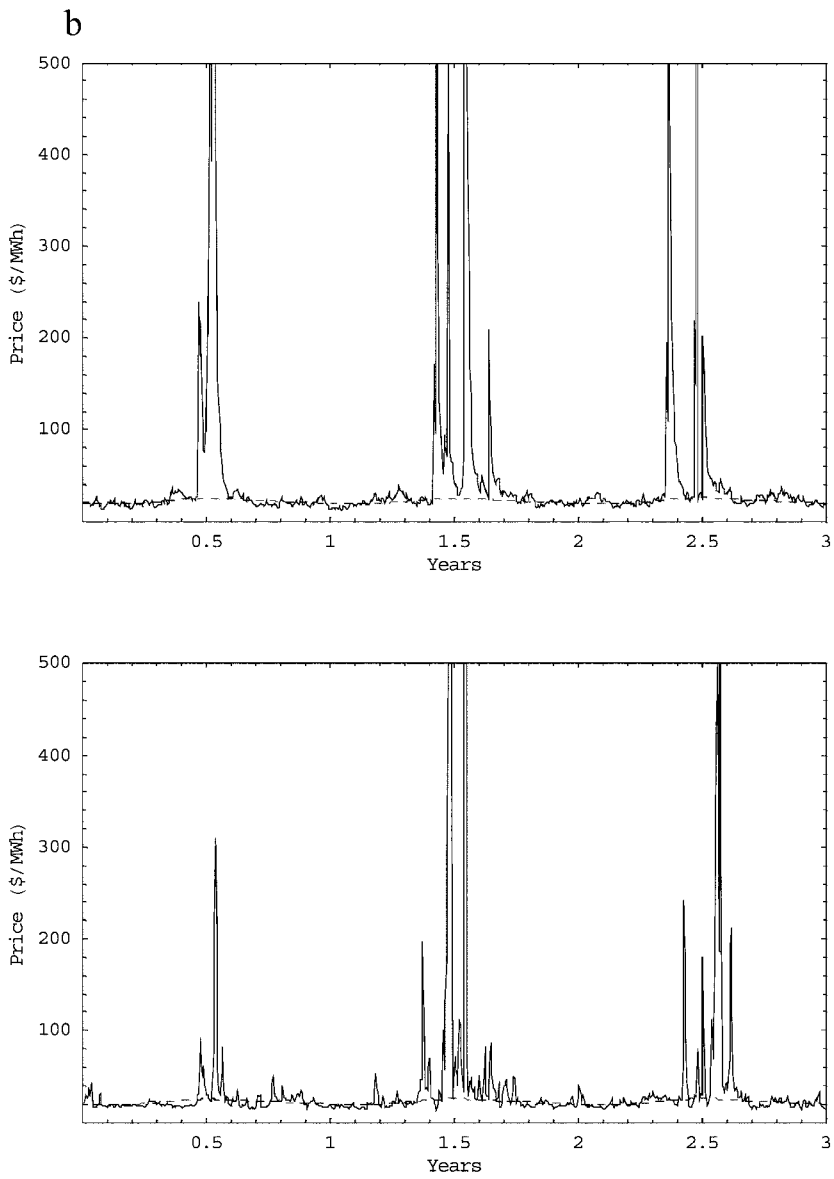

FIG. 10.—(*Continued*) ECAR simulated price path vs. empirical path: *b*, absolute scale 0–500.

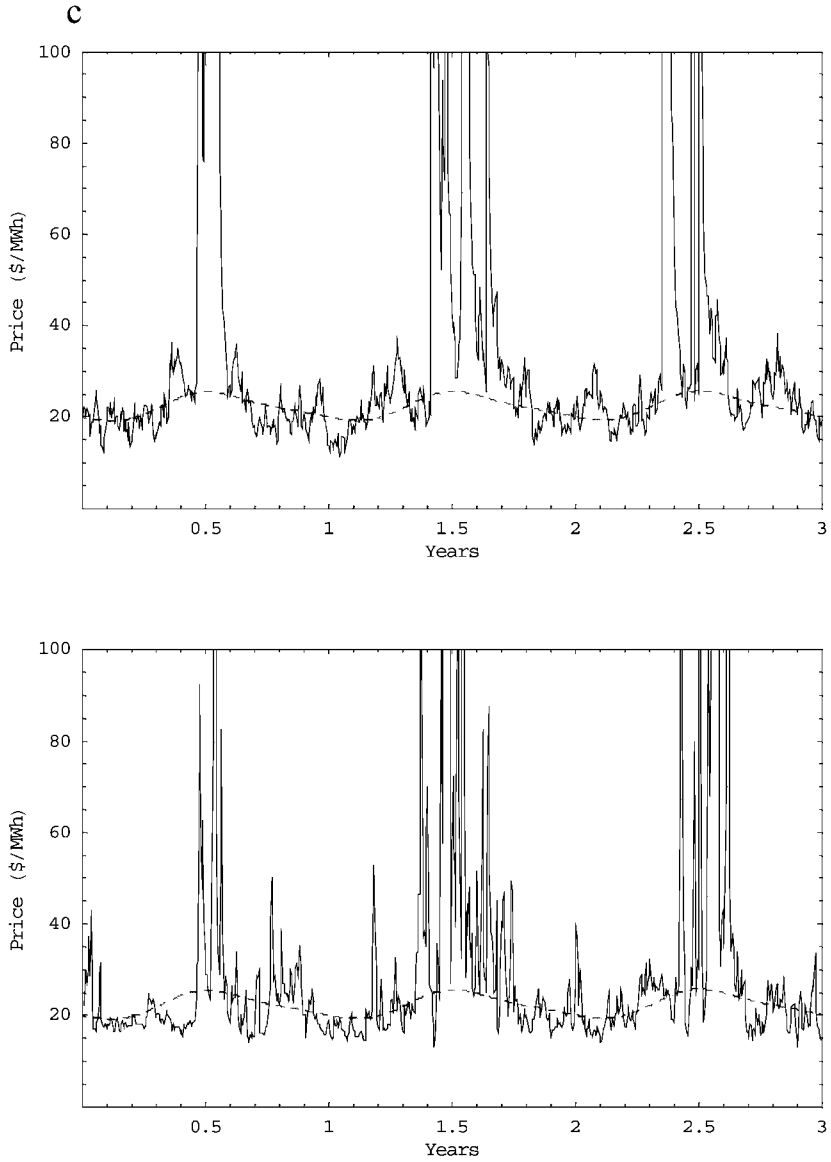

FIG. 10.—(*Continued*) ECAR simulated price path vs. empirical path: *c*, absolute scale 0–100.

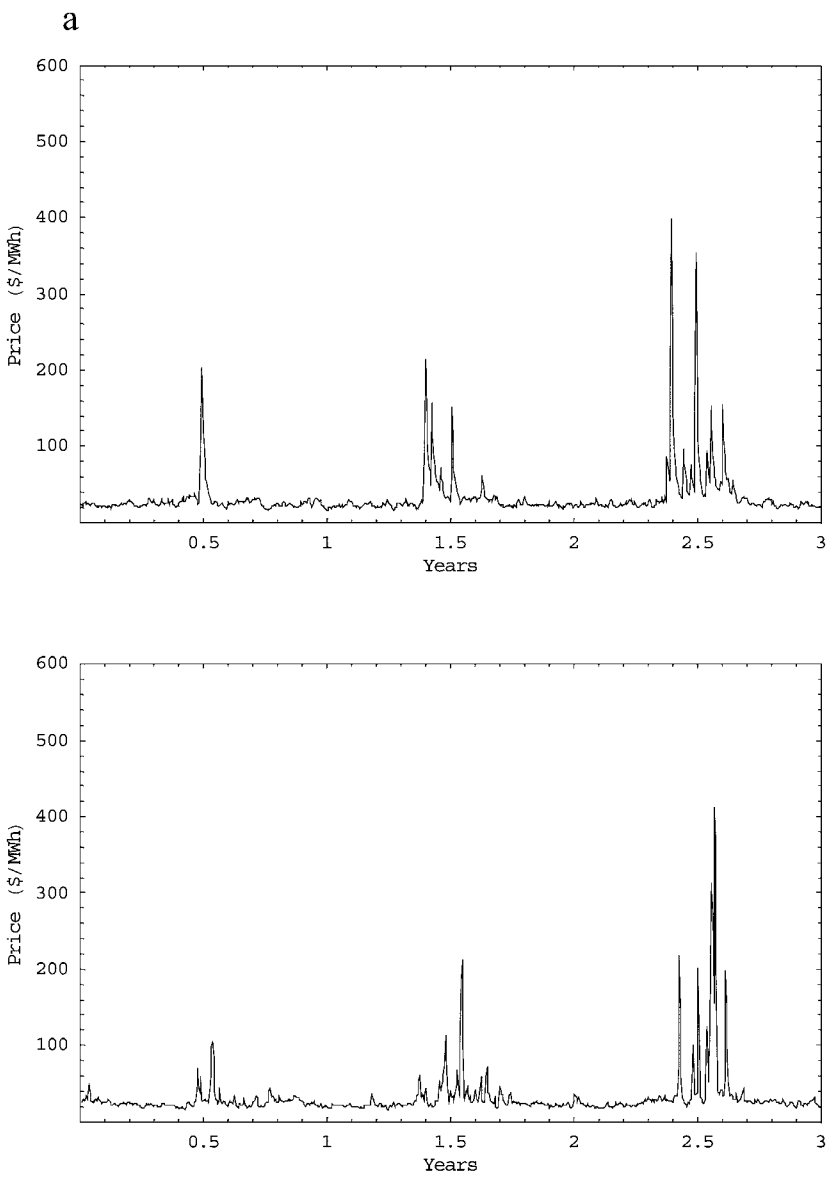

FIG. 11.—PJM simulated price path vs. empirical path: *a*, absolute scale 0–600

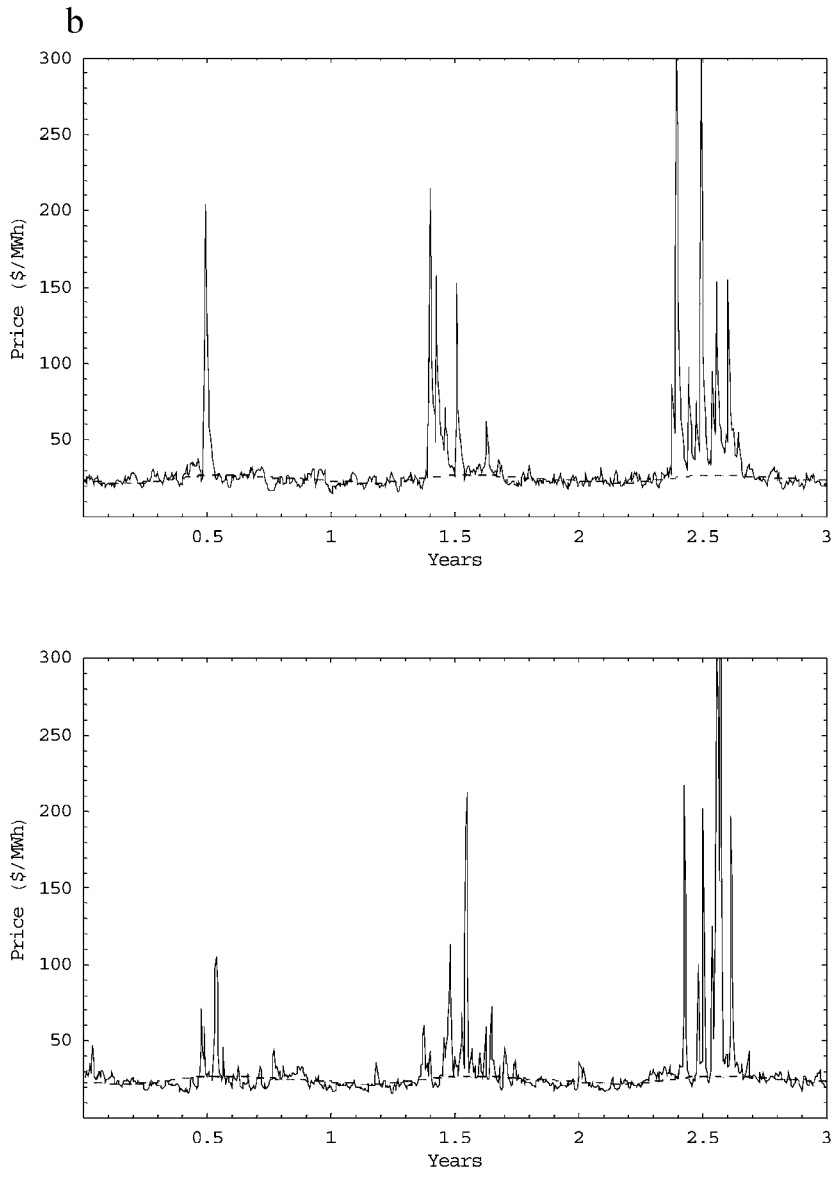

FIG. 11.—(*Continued*) PJM simulated price path vs. empirical path: *b*, absolute scale 0–300.

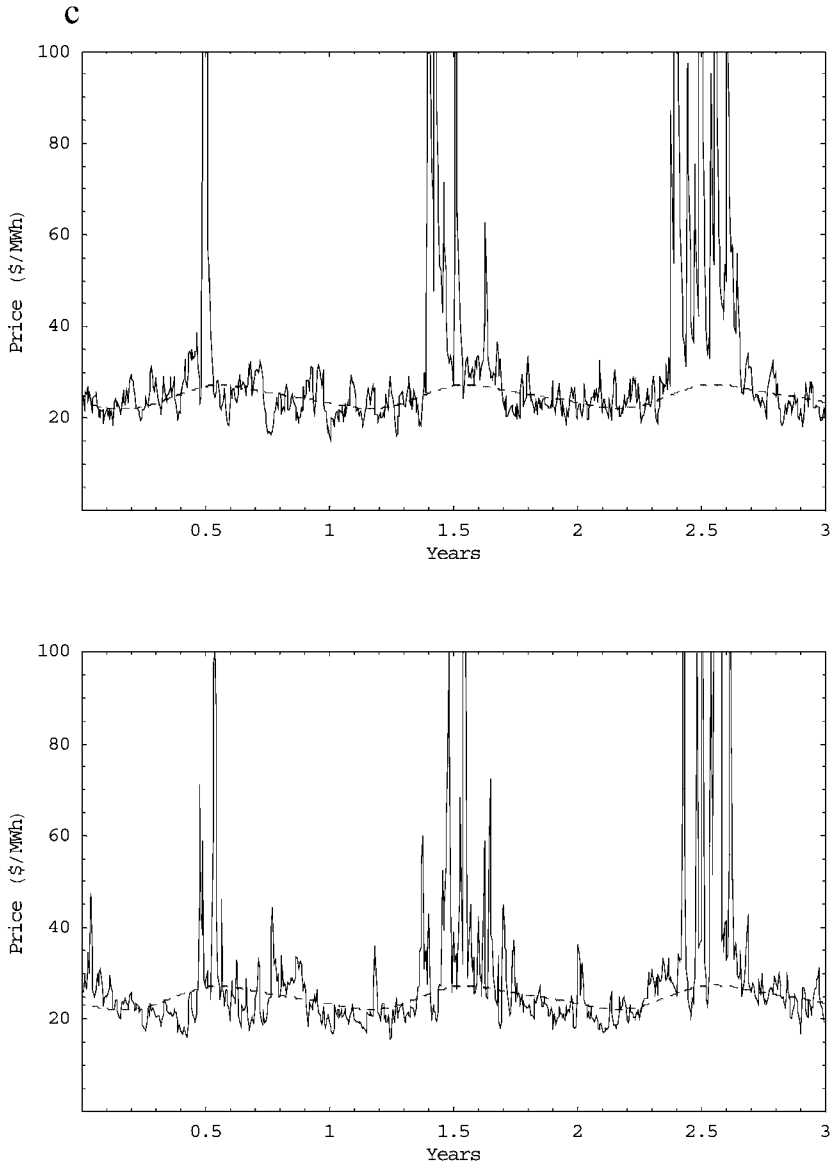

FIG. 11.—(*Continued*) PJM simulated price path vs. empirical path: *c*, absolute scale 0–100.

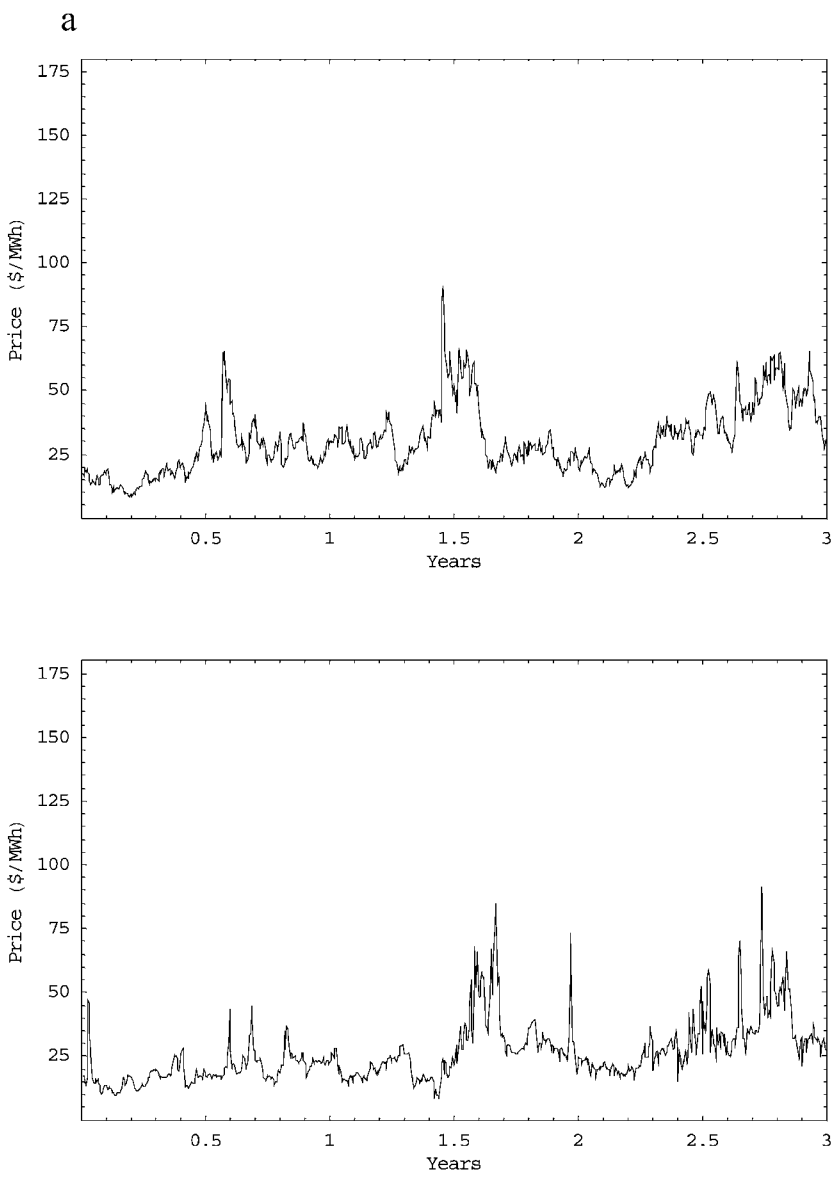

FIG. 12.—COB simulated price path vs. empirical path: *a*, absolute scale 0–175

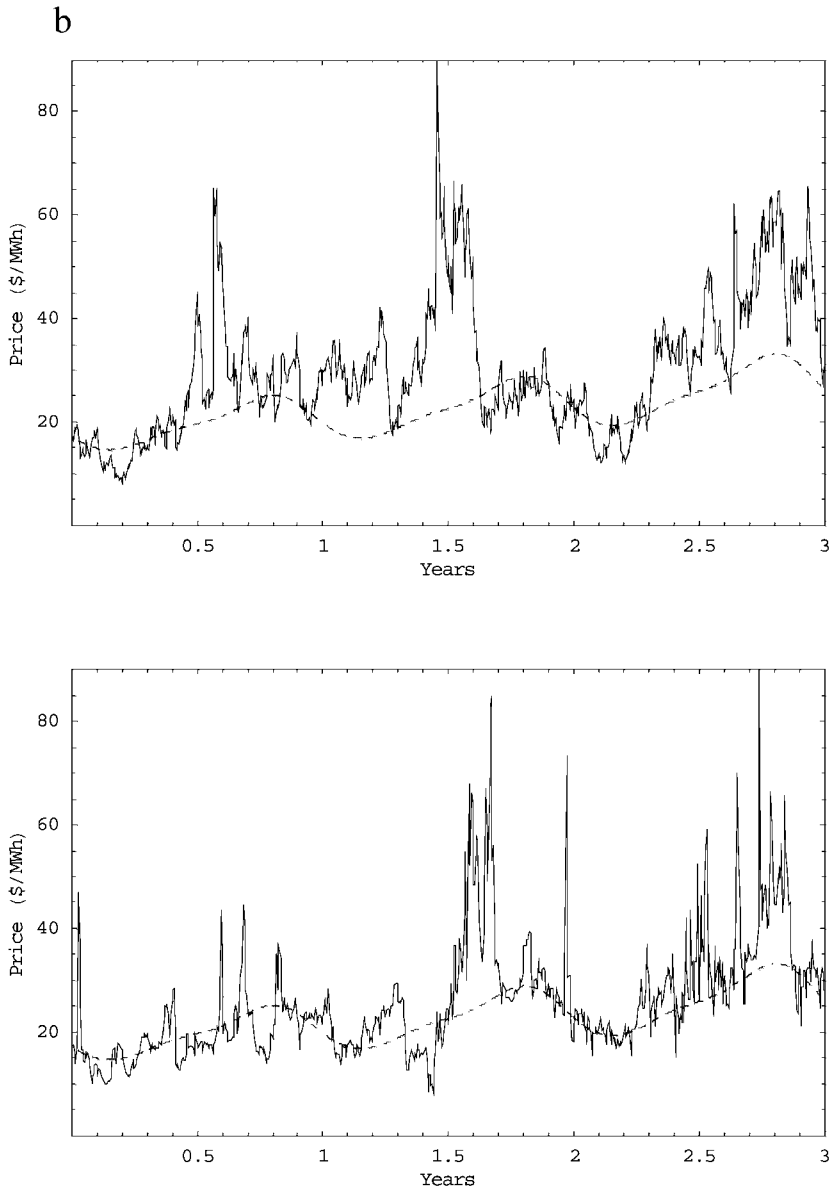

FIG. 12.—(*Continued*) COB simulated price path vs. empirical path: *b,* absolute scale 0–90.

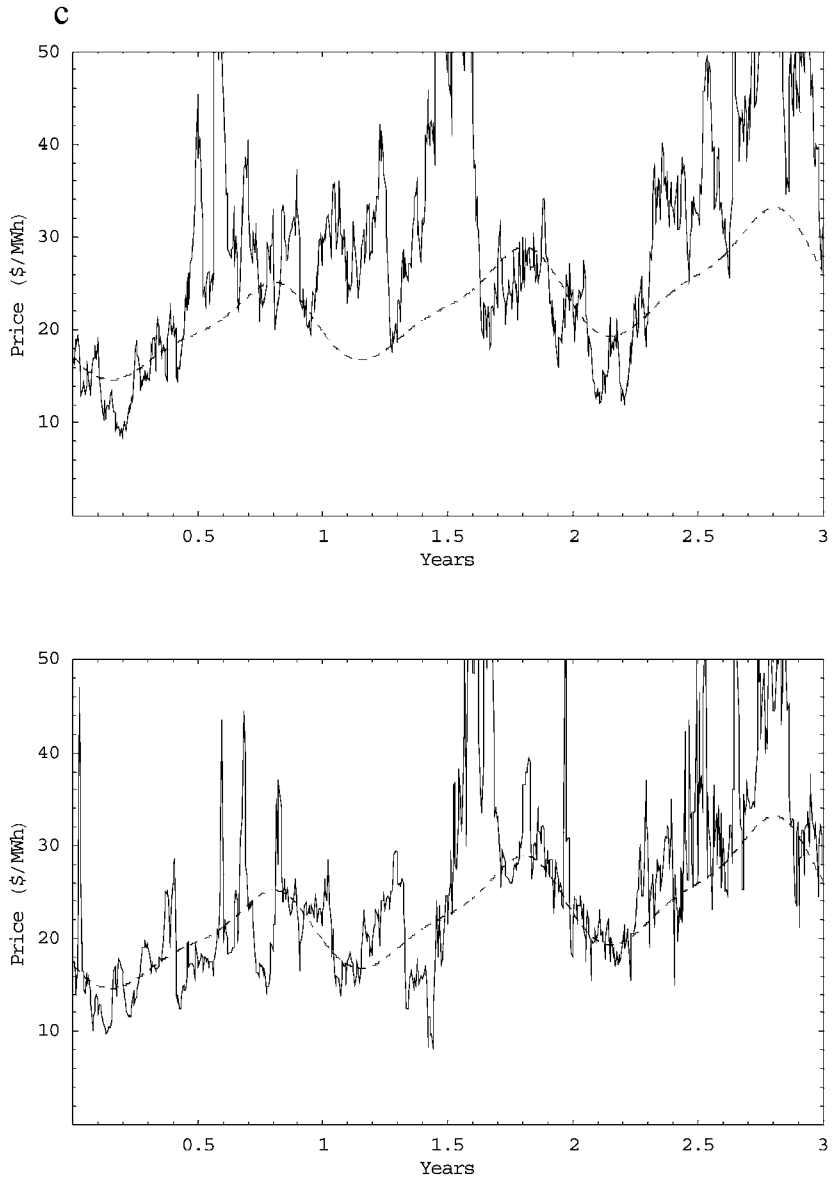

FIG. 12.—(*Continued*) COB simulated price path vs. empirical path: *c*, absolute scale 0–50.

TABLE 3 Moment Matching

	ECAR		PJM		COB	
	Empirical	Simulated	Empirical	Simulated	Empirical	Simulated
Average	−.0002	−.0001	−.0006	.0000	.0009	.0006
Standard deviation	.3531	.3382	.2364	.2305	.1586	.1121
Skewness	−.5575	2.1686	.3949	1.6536	.1587	.9610
Kurtosis	21.6833	22.5825	13.1507	14.8429	6.7706	6.5402

NOTE.—For each model estimated by maximum likelihood, descriptive statistics are computed for the empirical vs. simulated (after calibration) logarithmic price variations. Statistics include the mean, standard deviation, skewness, and excess of kurtosis. Simulations have been performed 1,000 times.

TABLE 4 Parameter Estimation Stability

	ECAR		PJM		COB	
	Original	Reestimation	Original	Reestimation	Original	Reestimation
θ_1	38.8938	37.7559	42.8844	40.0285	13.3815	11.7956
θ_2	59.5210	57.9367	4.1578	4.0188	2.5822	2.4001
θ_3	.3129	.2957	.5016	.4800	1.0038	1.1897
σ	1.8355	2.1355	1.4453	1.7822	1.3631	1.3882

NOTE.—The parameters θ_1, θ_2, θ_3, and σ have been reestimated by approximated maximum likelihood over 300 simulated paths. Results have been averaged over all samples.

the case of electricity prices, the nonnormality of distributions is widely recognized, and kurtosis naturally becomes a key parameter: in these markets in which extreme events provide the rationale for building small and flexible power plants called peakers, a proper representation of the spikes and their probability of occurrence (i.e., of the tail of the distribution) is the first requirement a model must satisfy.

We further test the robustness of the estimators by simulating 1,000 paths from the estimated process and then using the corresponding increments to reassess the values of the parameters θ_1, θ_2, θ_3, and σ. The simulation method is detailed in Appendix B, and the results are described in table 4.

For all estimated models the procedure is satisfactorily stable. We do not report the values for Γ because they are all identical to the original ones. The only slight mismatch occurs for the jump size parameter θ_3 in the case of the COB market; this may be due to the very low number of jumps, which makes the estimator sensitive to outliers in the simulated paths. This result is of minor importance, to the extent that the jump component is almost irrelevant for modeling COB prices. In general, we conclude that the procedure is not only statistically but also numerically robust.

Returning to the alternative specifications discussed in Section V.*C*, we also calibrated the *upward-jump* model with *deterministic* intensity and the *signed-jump* model with *stochastic* intensity to the ECAR market data. For the purpose of comparison, table 5 shows the quality assessment of these two models with respect to the benchmark defined by the signed-jump model with a deterministic intensity. It is clear that all three models account quite well for the first two moments of daily average prices, with an excess in volatility and

TABLE 5 Moment Matching of Alternative Models in the ECAR Market

	Market: ECAR (1)	Existing Literature: Upward-Jump Deteterministic Intensity (2)	Model I: Signed-Jump Deterministic Intensity (3)	Model II: Signed-Jump Stochastic Intensity (4)
Average	−.0002	.0000	−.0001	−.0000
Standard deviation	.3531	1.3238	.3382	.37821
Skewness	−.5575	3.5688	2.1686	−.0119
Kurtosis	21.6833	8.3542	22.5825	28.0288

NOTE.—A comparison between descriptive statistics of empirical data and corresponding statistics is produced for the three models. Col. 1 reports statistics of the ECAR market between January 6, 1997, and December 30, 1999. Col. 2 refers to the standard jump-diffusion model in the existing literature: a smooth mean-reverting diffusion with an upward jump component only. Col. 3 relates to our benchmark model (model I): a jump-diffusion model with a deterministic jump intensity, and col. 4, to a jump-reverting diffusion with a stochastic jump intensity (model II). Each model is estimated by maximum likelihood. Model-generated statistics are computed over a sample of 1,000 simulated paths. The simulation algorithm for both deterministic and stochastic jump intensities is detailed in App. B.

positive skewness, however, for the upward-jump model. The signed-jump model with stochastic intensity compared to the one with deterministic intensity slightly improves the value of the skewness, and our view is that this extra complexity does not bring any decisive improvement. As for the upward-jump model (with deterministic intensity), which is quite popular in the literature on electricity spot price modeling, it generates a kurtosis four times smaller than the real one; this misspecification may translate into a wrong estimation of value at risk numbers and have severe consequences in markets in which some inefficient plants continue to exist only because of these rare events. In all industries a wrong estimation of reserves leads to harmful consequences.

VII. Conclusion

We have proposed in this paper a family of discontinuous processes featuring upward and downward jumps to model electricity spot prices. Our approach is rooted in the physical properties of electricity, in particular its nonstorability, and their consequences on the short-term supply and demand equilibrium in the pool market.

Given the number of state variables that explain power prices in a pool (i.e., temperature, fuel mix, type of transmission network) and their distributional complexity (e.g., the occurrence of plant outages), we chose a reduced-form representation in order to get a tractable and efficient tool allowing us to handle the random evolution of spot prices and the related management decisions. The calibrated processes exhibit the expected mean reversion property, however, in an unevenly pronounced manner depending on the market. All analyzed trajectories show price spikes resulting from a momentary imbalance between offered generation and volume of demand.

The fitting performed on three major U.S. markets allows us to conclude positively the quality of the model, in terms of both its statistical and trajectorial properties.

Appendix A

Likelihood Estimator

The following proposition is an important result for the estimation of jump processes, from both a theoretical and an operational standpoint, and an original contribution of the paper (to our knowledge, at least).

PROPOSITION. Let μ, s, f, c, and σ be sufficiently regular functions for the stochastic differential equations (1)–(4) to admit a unique weak solution $E^{\boldsymbol{\theta}}$ for all $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3)$ in a compact subset of $\mathbb{R}_+^3$. Let $\mathbf{E} = \{E(t),\ t_0 \le t \le t\}$ be an observed path over the continuous-time interval $[t_0, t]$ and $\boldsymbol{\theta}^0 = (\theta_1^0, \theta_2^0, \theta_3^0)$ a starting parameter set. Then the log likelihood of observing a realization of the process $E^{\boldsymbol{\theta}}$ with respect to the process $E^{\boldsymbol{\theta}^0}$ is given by

$$\begin{aligned}\mathcal{L}(\boldsymbol{\theta}|\boldsymbol{\theta}^0, \mathbf{E}) = &\int_{t_0}^{t} \frac{[\mu(u) - E(u^-)](\theta_1 - \theta_1^0)}{\sigma(u)^2} d[E^c(u^-) - \mu(u)] \\ &- \frac{1}{2}\int_{t_0}^{t} \left\{\frac{[\mu(u) - E(u^-)]\theta_1}{\sigma(u)}\right\}^2 du \\ &- \left(\frac{\theta_2}{\theta_2^0} - 1\right)\int_{t_0}^{t} s(u)du + (\log\theta_2 - \log\theta_2^0)N(t) \\ &+ \sum_{u\le t, \Delta E \neq 0}\left[(\theta_3 - \theta_3^0) f\left(\frac{\Delta E(u)}{h(E(u^-))}\right) - \log c(\theta_3) + \log c(\theta_3^0)\right], \qquad \text{(A1)}\end{aligned}$$

where E^c is the path process devoid of its jump component:

$$E^c(u^-) = E(u^-) - E_0 - \sum_{s\le u, \Delta E(s)\neq 0} \Delta E(s), \qquad \text{(A2)}$$

E_0 is the starting point $E(t_0)$, $\Delta E(s)$ is the observed jump size at time s (if any), and $N(t)$ is the number of jumps that occurred up to time t.

Proof. For notational simplicity, we write equation (1) as

$$dE = (\alpha + \theta_1\beta)dt + \sigma dW + h dJ, \qquad \text{(A3)}$$

with $\alpha = D\mu(t)$, $\beta = \mu(t) - E(t^-)$, and $h = h(E(t^-))$, and we set $t_0 = 0$. We also denote $E(t^-)$ by E^-.

We divide the proof in two steps. First, we compute the semimartingale characteristic triplet $(B_{\boldsymbol{\theta}}, C, \nu_{\boldsymbol{\theta}})$ of the jump diffusion process E corresponding to a given choice of the parameter $\boldsymbol{\theta}$. Second, we calculate the likelihood by applying a general semimartingale version of the Girsanov theorem (see Jacod and Shiryaev 1987).

Step 1: Since N is independent of J_i for all i, $\mathbb{E}^{\boldsymbol{\theta}}(N(t)) = \iota(t)$, $J_i \overset{\text{i.i.d.}}{\sim} p(x;\ \theta_3)$, and the additive compensator of the purely discontinuous part of the semimartingale E is given

by

$$\begin{aligned}\nu_\theta(dt \times A,\ t) &= \nu_\theta(A,\ t)dt \\ &= \left(\mathbb{E}^\theta\left\{h(E^-)d_t\left[\sum_{i=1}^{N(t)} J_i\right]\right\}\right)dt \\ &= \left(\theta_2 s(t)\int_{[0,\psi]} dx\{1_{A\setminus\{0\}}[h(E^-)x]p(x;\ \theta_3)\}\right)dt \\ &= \left[\theta_2 s(t)\int_{\mathcal{X}} \frac{x}{h(E^-)}p\left(\frac{x}{h(E^-)};\ \theta_3\right)\right]dxdt.\end{aligned}$$

where $\mathcal{X} = ([0,\ \psi/h(E^-)] \cap A) \setminus \{0\}$. Since all coefficients are bounded functions, the process E is a special semimartingale. Consequently, the canonical representation of equation (A3) follows by adding and subtracting the compensator ν_θ to the jump measure $d\mu_\theta = h(t)dJ(t)$ and gathering the absolutely continuous terms:

$$dE = \left[\alpha + \beta\theta_1 + \int_{[0,\psi]} d\nu_\theta(x,\ t)\right]dt + \sigma dW + d\overline{\mu_\theta},$$

where $\overline{\mu_\theta}$ is a martingale measure under $\mathbb{P}_\theta$. From this expression we immediately identify the term of the semimartingale triplet corresponding to $\boldsymbol{\theta}$:

$$B_\theta(t) = \int_0^t \left[\alpha(u) + \beta(u)\theta_1 + \int_{[0,\psi]} h(E^-)x d\nu_\theta(x,\ u)\right]du. \tag{A4}$$

Step 2: The semimartingale process under the prior probability $\mathbb{P}^{\theta^0}$ is determined by the characteristic triplet $(B_{\theta^0},\ C,\ \nu_{\theta^0})$. Since

$$\begin{aligned}\nu_\theta(dt \times A) &= dt\int_{\mathcal{X}} dx\left\langle\frac{\theta_2}{\theta_2^0}\exp\left\{(\theta_3 - \theta_3^0)f\left(\frac{x}{h(E^-)}\right) - [\log c(\theta_3) - \log c(\theta_3^0)]\right\}\right. \\ &\quad \left.\times\ \theta_2^0 s(t)\frac{x}{h(E^-)}\exp\left[\theta_3^0 f\left(\frac{x}{h(E^-)}\right) - \log c(\theta_3^0)\right]\right\rangle \\ &= \int_{\mathcal{X}} \frac{\theta_2}{\theta_2^0}\exp\left[(\theta_3 - \theta_3^0)f\left(\frac{x}{h(E^-)}\right) - \log c(\theta_3) + \log c(\theta_3^0)\right]\nu_{\theta^0}(dt \times dx),\end{aligned}$$

the density of $d\nu_\theta$ with respect to $d\nu_{\theta^0}$ is given by

$$d_\theta(t,\ x) = \frac{\theta_2}{\theta_2^0}\exp\left[(\theta_3 - \theta_3^0)f\left(\frac{x}{h(E^-)}\right) - \log c(\theta_3) + \log c(\theta_3^0)\right].$$

By substituting this expression into (A4), we see that the drift term under $\mathbb{P}_\theta$ can be represented as the sum of the drift term under $\mathbb{P}_{\theta^0}$ and a term denoted as $c_\theta(t)\sigma(t)$, where

$$c_\theta(t) = \beta(t)(\theta_1 - \theta_1^0)\sigma(t)^{-1}.$$

Let $\mathbb{P}_\theta|_{\mathcal{F}_t}$ be the probability measure induced by E^θ over the path space and restricted to events up to time t. Given the set $\mathbf{E}$ of continuous-time observations, the corre-

sponding density of $\mathbb{P}_\theta|_{\mathcal{F}_t}$ with respect to the prior probability $\mathbb{P}_{\theta^0}|_{\mathcal{F}_t}$ is given by the Radon-Nikodym derivative:

$$\left.\frac{d\mathbb{P}_\theta}{d\mathbb{P}_{\theta^0}}\right|_{\mathcal{F}_t} = \exp\left(\int_0^t \left\{c_\theta dW - \frac{1}{2}c_\theta^2 du - \int_{\mathcal{X}} [(d_\theta - 1)d\nu_{\theta^0} + \log d_\theta d\mu]\right\}\right).$$

This is a consequence of the Girsanov theorem on measure changes for general semimartingales (see Jacod and Shiryaev 1987). The first two factors can be written as

$$\exp\left[\int_0^t \left(c_\theta dW - \frac{1}{2}c_\theta^2 du\right)\right] = \exp\left\{\int_0^t \frac{\beta(\theta_1 - \theta_1^0)}{\sigma^2} d\left[E(u) - \int_0^u \alpha(v)dv - \sum_{s\le u, \Delta E(s)\neq 0} \Delta E(s)\right] - \frac{1}{2}\int_0^t \frac{\beta^2(\theta_1 - \theta_1^0)^2}{\sigma^2} du\right\}.$$

The third factor is

$$\begin{aligned}\exp\left[-\int\int (d_\theta - 1)d\nu_{\theta^0}\right] &= \exp\left\{-\int_0^t s(u)\left[\frac{\theta_2}{\theta_2^0}\int_{\mathcal{X}} p\left(\frac{x}{h(E^-)};\ \theta_3\right)dx - \int_{\mathcal{X}} p\left(\frac{x}{h(E^-)};\ \theta_3^0\right)dx\right]du\right\} \\ &= \exp\left[-\left(\frac{\theta_2}{\theta_2^0} - 1\right)\int_0^t s(u)du\right],\end{aligned}$$

where we use the property $\int_{\mathcal{X}} p(x/h(E^-);\ \theta)dx = 1$.

The fourth factor is

$$\begin{aligned}\exp\left(\int_0^t\int_{\mathcal{X}} \log d_\theta d\mu\right) &= \exp\left\{\int_0^t\int_{\mathcal{X}}\left[(\theta_3 - \theta_3^0)f\left(\frac{x}{h(E^-)}\right) - \log c(\theta_3) + \log c(\theta_3^0)\right]d\mu \right. \\ &\quad \left. + (\log\theta_2 - \log\theta_2^0)\int_0^t\int_{\mathcal{X}} d\mu\right\} \\ &= \exp\sum_{u\le t, \Delta E(u)\neq 0}\left[(\theta_3 - \theta_3^0)f\left(\frac{\Delta E(u)}{h(E^-)}\right) - \log c(\theta_3) + \log c(\theta_3^0)\right] \\ &\quad + (\log\theta_2 - \log\theta_2^0)N(t),\end{aligned}$$

where the last equality stems from the relation between the process and the measure representation of any marked point process. Substituting the expressions for α and β leads to the log likelihood function (A1). QED

Appendix B

Simulation Algorithm

Monte Carlo simulations of trajectories described in equation (1) serve three purposes. First, they provide a starting value $\boldsymbol{\theta}^0$ for the maximum likelihood search algorithm. This is accomplished by sampling trajectories for several parameter sets until we find

one whose corresponding simulated paths show qualitative features comparable with those displayed in the empirical observations. Second, sample trajectories allow one to judge on the qualitative performance of the calibrated model and to compute simulated moments of various orders for the daily price variations. This is used for moment matching in the last step of the calibration procedure. Third, simulations provide a robustness analysis of the estimation procedure: parameters of a calibrated model can be reestimated over simulated paths. The closer to the original values the reestimated ones are, the more robust the likelihood estimation procedure is. We detail a simulation algorithm for sampling a path defined by equation (1). The Euler approximation of the stochastic differential equation (1) over a discrete set of evenly spaced sample times $t_1, \ldots, t_N$ is

$$E_{k+1} = E_k + D\mu(t_k) \times \Delta + \theta_1[\mu(t_k) - E_k] \times \Delta + \sigma\sqrt{\Delta}\mathcal{N} + h(t_k) \times \mathbf{1}_i \times J,$$

where $\mathcal{N}$ is a sample from a standard normal distribution and J is a sample from $p(\cdot, \theta_3)$. The function $\mathbf{1}_i$ is either one or zero according to whether t_i is, or is not, a jump time of the process. In order to sample jump times of a point process with nonconstant deterministic intensity, we may first simulate jump times of a constant intensity Poisson process and then use a variation of the "acceptance-rejection" method to make sure that these are statistically identical to the required sample set of times. More precisely, on a given horizon $[0, T]$, we generate interarrival times ε_i until their sum exceeds T. Each ε_i is a sample from an exponential distribution with parameter $\iota^* = \max_{t\in[0,T]}\iota(t)$. Candidate jump times τ_k' are defined by approximating each $\sum_{i=1}^{k}\varepsilon_i$ to the closest element in the set of sample times $\{t_1, \ldots, t_N\}$. For each k, we draw a uniform random variable U_k on $[0, \iota^*]$ and accept τ_k' if $U_k \le \iota(\tau_k')$; otherwise we reject it. The set of selected times is hence a sample sequence $(\tau_1, \ldots, \tau_n)$ of the jump times for a compound jump process with intensity function $\iota(t)$. Consequently, $\mathbf{1}_i = 1$ if $t_i = \tau_k$, for some $k = 1, \ldots, n$. This completes the description of the simulating algorithm for any calibrated solution of equation (1).

References

Andersen, Torben G., Luca Benzoni, and Jesper Lund. 2002. An empirical investigation of continuous-time equity return models. *Journal of Finance* 57:1239–84.

Ball, Clifford A., and Walter N. Torous. 1983. A simplified jump process for common stock returns. *Journal of Financial and Quantitative Analysis* 18:53–65.

Bandi, Federico. 2000. On the functional estimation of jump-diffusion models. Unpublished manuscript, University of Chicago.

Barlow, M. T. 2002. A diffusion model for electricity prices. *Mathematical Finance* 12, no. 4: 287–98.

Barone-Adesi, Giovanni, and Andrea Gigli. 2002. Electricity derivatives. Unpublished manuscript, Università della Svizzera Italiana, National Center of Competence in Research.

Beckers, Stan. 1981. A note on estimating the parameters of the jump-diffusion model of stock returns. *Journal of Financial and Quantitative Analysis* 16:127–40.

Bessembinder, Hendrik, and Michael L. Lemmon. 2002. Equilibrium pricing and optimal hedging in equilibrium electricity forward markets. *Journal of Finance* 57:1347–82.

Bhanot, Karan. 2000. Behavior of power prices: Implications for the valuation and hedging of financial contracts. *Journal of Risk* 2:43–62.

Carr, Peter, Hélyette Geman, Dilip Madan, and Marc Yor. 2002. The fine structure of asset returns: An empirical investigation. *Journal of Business* 75, no.2:305–33.

Deng, Shijie. 1999. Stochastic models of energy commodity prices and their applications: Mean reversion with jumps and spikes. Unpublished manuscript, Georgia Institute of Technology.

Escribano, Álvaro, Juan Ignacio Peña, and Pablo Villaplana. 2002. Modeling electricity prices: International evidence. Working Paper no. 02-27, Economic Series 08, Universidad Carlos III de Madrid, Departamento de Economia.

Eydeland, Alexander, and Hélyette Geman. 1998. Pricing power derivatives. *Risk* (October), pp. 53–60.

Fama, Eugene, and Kenneth French. 1987. Commodity futures prices: Some evidence on forecast power, premiums and the theory of storage. *Journal of Business* 60, no.1:55–73.

Geman, Hélyette, and Oldrich Vasicek. 2001. Forwards and futures contracts on non-storable commodities: The case of electricity. *Risk* (August), pp. 12–27.

Genon-Catalot, Valentine, and Jean Jacod. 1993. On the estimation of the diffusion coefficient for multidimensional diffusion processes. *Annales de Institut Henri Poincaré* 29, no.1:119–51.

Gihman, Iosif Il'Ich, and Anatoli V. Skorohod. 1972. *Stochastic differential equations*. NewYork: Springer-Verlag.

Huisman, Ronald, and Ronald Mahieu. 2001. Regime jumps in electricity prices. Working Paper no. ERS-2001-48-F&A, Erasmus University, Rotterdam, Rotterdam School of Management, Erasmus Research Institute of Management.

Jacod, Jean, and Albert Shiryaev. 1987. *Limit theorems for stochastic processes*. New York: Springer-Verlag.

Johannes, Michael. 1999. Jumps in interest rates: A nonparametric approach. Unpublished manuscript, University of Chicago.

Joskow, Paul L., and Edward Kahn. 2001. A quantitative analysis of pricing behavior in California's wholesale electricity market during summer 2000. Working paper, American Enterprise Institute–Brookings Joint Center for Regulatory Studies, Washington, DC.

Kaldor, N. 1939. Speculation and economic stability. *Review of Economic Studies* 7:1–27.

Knittel, Christopher R., and Michael R. Roberts. 2001. An empirical examination of deregulated electricity prices. Unpublished manuscript, Boston University.

Lo, Andrew. 1988. Maximum likelihood estimation of generalized Ito processes with discretely sampled data. *Econometric Theory* 4:231–47.

Lucia, Julio J., and Eduardo S. Schwartz. 2002. Electricity prices and power derivatives: Evidence from the Nordic Power Exchange. *Review of Derivatives Research* 5:5–50.

Merton, Robert. 1976. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* 3:125–44.

Pedersen, Asger Roer. 1995. A new approach to maximum likelihood estimation for stochastic differential equations based on discrete observations. *Scandinavian Journal of Statistics* 22: 55–71.

Pilipovich, Dragana. 1997. *Energy risk: Valuing and managing energy derivatives*. New York: McGraw-Hill.

Pindyck, Robert. 1999. The long-run evolution of energy commodity prices. *Energy Journal* (April), pp. 54–79.

Pirrong, Craig. 2001. The price of power: The valuation of power and weather derivatives. Unpublished manuscript, Oklahoma State University.

Roncoroni, Andrea. 2002. Essays in quantitative finance: Modelling and calibration in interest rate and electricity markets. PhD diss., University Paris Dauphine.

Working, Holbrook. 1949. The theory of the price of storage. *American Economic Review* 39: 1254–62.

Yin, Y. Q. 1988. Detection of the number, locations and magnitude of jumps. *Communications in Statistics: Stochastic Models* 4:445–55.