

Practical stochastic modelling of electricity prices

MICHEL CULOT

Electrabel SA, Belgium

VALÉRIE GOFFIN

FUNDP, University of Namur, Belgium

STEVE LAWFORD*

Department of Economics and Econometrics, ENAC, France

SÉBASTIEN DE MENTEN

Electrabel SA, Belgium

AND YVES SMEERS

Department of Mathematical Engineering, UCL and CORE, Belgium

December 19, 2011.

*Correspondence can be addressed to Steve Lawford, Department of Economics and Econometrics (LH/ECO), ENAC, 7 avenue Edouard Belin, BP 54005, 31055, Toulouse, Cedex 4, France; email: steve_lawford@yahoo.co.uk. Culot, Goffin, Lawford and de Menten were with the Strategy, Research and Development team, Electrabel SA, when the majority of this research was undertaken. Smeers is Tractebel Professor of Energy Economics at UCL/CORE, and advisor to Electrabel. We thank the Editor, Derek Bunn, and an anonymous referee, for comments that led to an improved version of the paper. For helpful discussions on various versions of this work, and for institutional knowledge, we are grateful to Paul Bekker, Kees Bouwman, Russell Davidson, Freddy Delbaen, Darrell Duffie, Robert Engle, Valérie Limpens, Michel Lubrano, Marius Ooms, and seminar participants at CORE, GREQAM, Groningen, the 2006 Energy Risk conference in Rotterdam, and the 2006 Econometric Society European Meeting in Vienna. Special thanks are due to Pavel Diko and Slimane Grine, for their support, and for carefully working through the proof of Proposition 1. We acknowledge support from Jacqueline Boucher and André Bihain at Electrabel SA, and Nathalie Lenoir at ENAC, where the final version was prepared. The views presented in this paper are those of the authors, and do not necessarily reflect the position of Electrabel SA. This paper is a heavily

Abstract

We develop a flexible multi-factor stochastic model with three diffusive and three spike regimes, for daily spot and forward electricity. The model captures various stylized features of power prices, including mean reversion and seasonal patterns, and short-lived spikes. We estimate parameters through a practical two-step procedure, that combines pre-calibration of deterministic elements and spikes, and state-space estimation of diffusive factors. We use several results on affine jump diffusions to combine the spike and diffusive components, and to provide convenient closed-form solutions for important power derivatives. We also propose a simple nonparametric model for hourly spot prices, based on hourly profile sampling from historical data. This model can reproduce complicated intraday patterns. We illustrate the performance of the daily and hourly models using data from the Amsterdam Power Exchange.

Keywords: Affine jump diffusions, Efficient option pricing, Electricity and energy markets, Regime-switching spikes, State-space (Kalman filter) estimation.

JEL classification: C10, C50, G12, G13, L94.

1 Introduction

We develop a practical multi-factor stochastic model for daily spot and forward electricity prices. The set-up captures some of the well known market-specific features of power price dynamics, including mean reversion and seasonal patterns, and short-lived price spikes. We estimate model parameters by a flexible two-step procedure, that combines pre-calibration of deterministic elements and spikes, and state-space estimation of a three-factor model for the short, medium and long-term driving diffusive factors. In a departure from the literature, we use a three-state regime-switching model for the *spikes*, which enables us to reproduce spike arrival frequencies, magnitudes and duration. We use several transform results on affine jump diffusions (AJDs), to derive convenient closed-form solutions for

revised version of an earlier manuscript that was circulated under the title “An affine jump diffusion model for electricity.” The paper was typed by the authors in MiKTeX and WinEdt. Numerical results were calculated using EViews, Mathematica, and Python.

important contingent claims, and show how the spike and diffusive components can be combined.

We also propose a simple nonparametric model for hourly spot prices, based on hourly profile sampling from historical data. This model can reproduce complicated intraday patterns. We illustrate the performance of the daily and hourly models, using data from the Amsterdam Power Exchange (APX), and perform a simulation-based assessment. Our modelling approach allows for various extensions, including more complicated models for the individual components, and time-varying dependence between power and fuel markets.

1.1 Features of electricity prices

The opening of continental European and North American electricity markets, and the continued increases in exchange based and over-the-counter volumes of trade, has exposed both energy producers and industrial end users to new forms of market and price risk (see Joskow (1997) and Mork (2001) for discussion of deregulation). Traders and risk managers rely upon accurate models of electricity spot and forward prices, and the construction of reliable forecasts and price scenarios, and tools for pricing energy derivatives. These models are important when evaluating hedging products and physical assets.

Despite apparent similarities with financial asset prices, such as heavy-tailed returns, electricity has very different stochastic properties to both standard securities and storable commodities.¹ It must be generated continuously for actual delivery and consumption, and meaningful quantities cannot often be stored at reasonable cost, or easily transported. This nonstorability, and lack of recourse to inventories, makes prices particularly sensitive to demand and supply shocks, that include unusually high temperatures, and technical problems such as power plant failure or transmission line overload. When major load and

¹For recent empirical research on power markets, see Karakatsani and Bunn (2005), Knittel and Roberts (2005), Atkins (2006), Koopman et al. (2007), and Escribano et al. (2011).

generation problems arise simultaneously, as during the European summer heatwave in 2003, these spikes can be extreme, e.g., between 10 and 11 August 2003, spot prices on the Dutch APX rose by over 3,000% (from 19.18 euros/MWh to 660.34 euros/MWh), only to return to original levels 5 days later. Spikes are generally short-lived, rather than leading to sustainable higher prices, but can nevertheless last for several days or more.

Price spikes typically lead to identification problems for single-factor jump diffusion spot models that were inspired by work on asset prices and interest rate dynamics. For instance, the Ornstein-Uhlenbeck process with additive Poisson jumps, a mean-reverting spot price, and volatility driven by Brownian motion, i.e., $d \ln S_t = \kappa(\theta - \ln S_t)dt + \sigma(S, t)dW_t + J_t dq_t$, usually requires a high speed of mean reversion in order to reduce the spot price following a large positive jump, which removes too much variability in the non-jump periods of the sample. See Huisman and Mahieu (2003), Weron et al. (2004), Borovkova and Permana (2006), and de Jong (2006) for background. Our regime-switching spike framework has three levels, that we identify from the spot prices. The levels are interpreted following calibration as small, medium and large spikes. A one-day transition matrix gives the probability of moving between the non-spike and the spike levels.

Electricity is also subject to natural and human phenomena, that result in periodic dependencies in the data. Examples include demand-driven annual patterns due to changing daylight hours, and intraweek and intraday periodicity influenced by industrial activity. In markets that are heavily dependent on hydroelectric generation, supply-side patterns are important: spot prices on the Scandinavian Nord Pool exchange are affected by precipitation and snowmelt. Practitioners and researchers commonly model annual and intraweek patterns with deterministic functions, such as truncated Fourier series (Pilipović (1998)) and sinusoids (Geman and Roncoroni (2006), Weron (2008), and Erlwein et al. (2010)).

We combine sinusoidal and piecewise-constant functions in this paper, for annual and

intra-week patterns. We use *forward* data to give a clearer indication of long-run seasonal patterns. Intraday seasonality is a particular problem when modelling hourly spot prices using standard econometric methods: these patterns will not appear in the daily aggregate series, but must be taken into account when designing hourly models. Weekend prices are strongly influenced by demand surges at midday and evening mealtimes, while weekday and weekend intraday patterns usually also differ (see Wilkinson and Winsen (2002), and Bottazzi et al. (2005)): our simple hourly model deals with this problem.

Spot prices are commonly thought to be mean reverting, and this feature is described in early empirical work by Weron and Przybyłowicz (2000), Lucia and Schwartz (2002), and Simonsen (2003). Following a temporary deviation, spot prices return to some equilibrium level, which reflects economic and fundamental factors such as the marginal cost of production and seasonal weather conditions: the equilibrium need not be constant, but may be periodic, or periodic with a trend. We capture short, medium and long-term reversion by using a multivariate extension of the Ornstein-Uhlenbeck process.²

1.2 Combined spot and forward models

Use of all available market information on forwards in the model estimation ensures coherency of the spot model with observed forward dynamics. The literature describes a number of single-factor electricity spot models, usually for daily data (e.g., Huisman and Mahieu's (2003) and Weron et al.'s (2004) regime-switching models for *prices*; econometric models such those in Misiorek et al. (2006), and Koopman et al.'s (2007) seasonal Reg-ARFIMA-GARCH; and modified jump diffusions, e.g., Borovkova and Permana (2006) and Geman and Roncoroni (2006)). These models do not treat forwards explicitly, nor do they

²A number of authors have used economic theory and models based upon fundamentals to cast light upon empirical features of electricity and energy markets, including Routledge et al. (2001), Barlow (2002), Bessembinder and Lemmon (2002), Kanamura and Ōhashi (2008), and Coulon and Howison (2009).

generally give rise to manageable analytic solutions for the pricing of derivatives.

Moreover, single-factor models do not have the flexibility to match the forward volatility term structure. They are usually estimated using historical spot series, and give a volatility term structure that decreases too rapidly, as a consequence of focusing on high-frequency spot price movements, with insufficient data to accurately model lower-frequency longer-term price movements. Another strand of research has developed Heath-Jarrow-Morton type and other direct models of forward curve dynamics, which are rarely affected by short-term spikes, e.g., Borovkova (2004), and Koekebakker and Ollmar (2005). However, the resulting implied spot model is generally not Markovian, and provides a poor approximation to the complex behaviour found in electricity spot markets.

Since electricity forwards are much less volatile than spot prices, it is reasonable to use multiple, possibly unobservable, risk factors, that are separately responsible for volatile short-term spot behaviour, and the less volatile medium and long-term effects that are observed in forward prices. We follow Diko et al. (2006), who provide evidence that a three-factor diffusive model may be appropriate for a number of European spot/forward markets, including the APX (additional support is provided by Kiesel et al. (2009), who describe a two-factor model for *forward* prices). A coherent spot/forward model is also important when assessing hedging strategies, and several studies have considered this issue for energy commodities, including Schwartz and Smith (2000), Kåresen and Husby (2002), Manoliu and Tompaidis (2002) and Cortazar and Schwartz (2003). We extend these papers by combining a multi-factor mean reverting diffusion with a regime-switching spike model.

1.3 Pricing and simulation

We set our daily model in the AJD framework, which enables us to use several fundamental transform results of Duffie et al. (2000) to derive efficient closed-form solutions, up to

resolution of a system of ordinary differential equations, for the conditional characteristic function of the state variables at maturity. This approach has important implications for the efficient pricing of electricity derivatives on both spot *and* forwards, and leads to rapid and elegant pricing solutions. While Monte Carlo has been used to price very complex financial derivatives (see Boyle et al. (1997) for an introduction), pure simulation is often computationally prohibitive in electricity markets, e.g., when pricing an option on a power forward, or when computing Greeks. Whenever possible, it is useful to have available *analytic* or efficient analytic/numerical solutions. Our procedure sets the spot and forward data in state-space form. Estimation of the free parameters follows directly by the Kalman filter, and maximum likelihood. Our simple hourly spot price model then avoids some difficulties associated with econometric time-series models of hourly data (e.g., Haldrup and Nielsen (2006)), such as overparameterization, and effectively uses all available historical hourly information.

Following estimation of the daily and hourly models, we go further than reporting descriptive statistics, and assess some interesting aspects of model quality by stochastic simulation, i.e., the ability to reproduce observed market behaviour, such as spike duration, and intraday mean patterns. We extend a similar technique that was developed independently by Geman and Roncoroni (2006), by considering the simulated *distributions* of the statistics of interest, rather than one or two of their moments, which provides a more detailed picture of the model performance. We mention pricing applications of our models in Appendix A.3.

The paper is organized as follows. Section 2 develops the daily and hourly models, explains the calibration procedure, and demonstrates how the daily model can be written in state-space form using spot and forward prices. Section 3 examines the quality of the estimated model, with a numerical example using APX data. Section 4 concludes.

2 Model design and implementation

We model the daily log baseload spot price as:

$$\ln(S_t) = \theta_t + \tilde{\gamma}^\top \tilde{X}_t + \gamma^\top X_t, \quad (1)$$

in which $\theta_t := \text{YP}_t + \text{WP}_t$ contains deterministic yearly and weekly patterns, and $\tilde{\gamma}$ and γ are coefficient vectors acting on *spike* risk factors \tilde{X}_t , and *diffusive* risk factors X_t .

- The 3×1 vector \tilde{X}_t has i th element unity, and the remaining elements are zero (corresponding to active i th *spike* state), while $\tilde{\gamma}$ captures the magnitude of each spike level. We define spikes on the APX to be *spot* prices that exceed 70 euros/MWh.³ This threshold is roughly twice the mean APX baseload daily spot price. We build the empirical spike distribution: we jointly calibrate spike parameters (level 1, 2 and 3 spikes, and the probability of arrival of a level 2 spike) to match the first four central moments of the empirical spike distribution, under the constraint that level 1 and level 3 spikes arrive with equal probability. This restriction gives 4 equations and 4 unknowns, and so the system is solvable. The spike magnitude on a given date is the difference between the daily price at that date and the average of the immediate pre-spike and post-spike levels. Once a spike has been assigned to a level, we construct a 4×4 one-day transition matrix to describe movements between non-spike dates and the three spike levels. The spike regimes are observed rather than latent, and this facilitates parameter estimation. Other electricity price papers that use observable regimes for the same reason include Haldrup and Nielsen (2006) and Haldrup et al. (2010). We remove the spikes from the spot series for the rest of the calibration.

³Similar threshold spike detection can be found in, e.g., Becker et al. (2007), Weron (2008), and Christensen et al. (2009). We have tested various extensions, including nonparametric “local” spike identification, but the basic principle remains the same: we identify spikes in a “sensible” way from the spot data.

- We model the annual seasonal pattern by a parametric function:

$$YP_t = \rho_1 + \rho_2 \cos \left[\frac{2\pi(t - \rho_3)}{365.25} \right],$$

with a level term and a sinusoidal function to approximate the annual cycle. We calibrate the parameters ρ_1, ρ_2 and ρ_3 by matching available quarterly forward data.

- We use a piecewise-constant function with 7 values for the weekly pattern:

$$WP_t = \ln\{\text{Hol}_t \text{DayLevel}_{t=\text{sun}} + (1 - \text{Hol}_t) \text{DayLevel}_t\},$$

and associates each day of the week with coefficients DayLevel_t , which we calibrate using the *spot* price series, after removal of spikes and holidays. For holiday dates, the DayLevel_t is a weighted (by the fraction of population on/not on holiday on that date) average of the Sunday $\text{DayLevel}_{t=\text{sun}}$ and the $\text{DayLevel}_{t=\text{mon}, \dots, \text{sun}}$. For the APX market, $\text{Hol}_t \in \{0, 1\}$, if a given day is a public holiday or not.

- We use an additive three-factor mean reverting model, with diffusive coefficients $\gamma = (1, 1, 1)^\top$, and a vector X_t of diffusive risk factors that follows an affine diffusion:

$$dX_t^{(i)} = -\kappa_i X_t^{(i)} dt + \sigma_i dW_t^{(i)}, \quad i = 1, 2, 3,$$

with $E[dW_t^{(i)} dW_t^{(j)}] = 0$, for $i \neq j$. In matrix notation, this model is written as:

$$dX_t = K_1 X_t dt + H_0^{1/2} dW_t, \quad (2)$$

in which $K_1 = \text{diag}(-\kappa_1, -\kappa_2, -\kappa_3)$ and $H_0 = \text{diag}(\sigma_1^2, \sigma_2^2, \sigma_3^2)$. An additive three-factor model was chosen following experimentation with two-factor models, and given

the principal components results of Diko et al. (2006) and others. The risk factors revert to zero at a speed of κ_i . Following estimation, the elements of X_t can be ordered by their respective speeds of mean reversion; these are usefully interpreted as independent short, medium, and long-term risk factors, corresponding to the largest, medium, and smallest speeds of reversion. Often, the estimated volatility components σ_i , which appear in $H_0^{1/2}$, fall with $|\kappa_i|$. Model (2) can be extended to correlated risk factors, although the statistical motivation for that is still unclear in power markets.

2.1 Affine jump diffusions

Equation (2) is a special case of the class of AJD processes. Transform results on AJDs enable (near-)analytical treatment of a wide class of derivative pricing problems, and computationally tractable estimation. For further technical details, we refer the reader to the seminal paper by Duffie et al. (2000), who derive the closed form of the conditional characteristic function (CCF) of the state vector X_T at maturity T , given information at time t ; and to Dai and Singleton (2000) and Duffie et al. (2003). Knowledge of the CCF is the same as knowledge of the joint conditional density function of X_T . Duffie et al. (2000) provide two transforms that enable efficient pricing of forwards and European options. Applications of AJDs to finance include Eraker (2004) and Johannes (2004).

A process X is an affine jump diffusion if:

- It is an $n \times 1$ Markov process relative to a filtration \mathcal{F}_t (“information”), that solves the stochastic differential equation:

$$dX_t = \mu(X_t)dt + \sigma(X_t)dW_t + dZ_t, \quad (3)$$

written under the physical measure P , and driven by the $n \times 1$ \mathcal{F}_t -adapted standard

Brownian motion W_t , with $n \times 1$ and $n \times n$ parameter functions μ and σ . The process Z_t is a pure jump, with fixed jump amplitude distribution ν and arrival intensity $\lambda(X_t)$. For an extension to multiple jump components, see Duffie et al. (2000, Appendix B). We assume that μ , σ and λ are regular enough that (3) has a unique strong solution, in the technical sense of Karatzas and Shreve (1999, Section 5.2). For complex $n \times 1$ vectors c , define the “jump transform” $\zeta(c) = \int_{\mathbb{R}^n} \exp(c^\top z) d\nu(z)$, which is assumed to be known in closed form, whenever the integral is well defined.

- The drift vector μ , “instantaneous” covariance matrix $\sigma\sigma^\top$ and jump intensity λ have an affine dependence on X :⁴

$$\mu(X_t) = K_0 + K_1 X_t \quad : \quad K_0 \text{ is } n \times 1, K_1 \text{ is } n \times n, \quad (4)$$

$$\sigma(X_t)\sigma(X_t)^\top = H_0 + H_1 X_t \quad : \quad H_0 \text{ is } n \times n, H_1 \text{ is } n \times n \times n, \quad (5)$$

$$\lambda(X_t) = l_0 + l_1 X_t \quad : \quad l_0 \text{ is } n(n+1) \times 1, l_1 \text{ is } n(n+1) \times n. \quad (6)$$

Then, the CCF ψ of X_T , given current information about X at time t , and maturity T , has an exponential-affine form:

$$\psi(u, X_t, t, T) := E^\Theta[e^{u^\top X_T} | \mathcal{F}_t] = e^{\alpha_t + \beta_t^\top X_t}, \quad t \leq T. \quad (7)$$

The expectation is taken with respect to the distribution of X determined by the parameters $\Theta = (K_0, K_1, H_0, H_1, l_0, l_1)$. This is a fundamental result. Duffie et al. (2000) show

⁴Here, $H_1 X_t$ denotes the $n \times n$ matrix $(H_1)_{ijk}(X_t)_k$, and Einstein summation notation is used. We implicitly sum over all repeated indices in tensor products. Further, $\beta_t^\top H_1 \beta_t$ denotes $(\beta_t)_i (H_1)_{ijk} (\beta_t)_j$.

that α and β satisfy the complex-valued Riccati ordinary differential equations:

$$\dot{\beta}_t = -K_1^\top \beta_t - \frac{1}{2} \beta_t^\top H_1 \beta_t - l_1^\top [\zeta(\beta_t) - 1], \quad (8)$$

$$\dot{\alpha}_t = -K_0^\top \beta_t - \frac{1}{2} \beta_t^\top H_0 \beta_t - l_0^\top [\zeta(\beta_t) - 1], \quad (9)$$

with boundary conditions $\beta_T = u$ and $\alpha_T = 0$. This system of equations may either be solved analytically or, when no closed-form exists, by numerical methods such as fourth-order Runge-Kutta, or similar.⁵

2.2 Diffusive risk factors and regime-switching spikes

We treat the affine diffusion and affine jump components of (3) separately, with CCFs ψ_{AD} and ψ_{AJ} . We assume that the spikes and diffusive components observed in electricity spot prices are independent, i.e., the diffusive component has no impact upon spike occurrence, so that the affine jump diffusion CCF can be constructed simply as $\psi := \psi_{AD}\psi_{AJ}$. Separating the treatment of spikes and diffusion can be reasonable in power markets, in which *extreme* price spikes are of interest, and are relatively easy to identify. Under conditions (4)–(6), the pure affine diffusion part in diffusive state X_t , follows $dX_t = \mu(X_t)dt + \sigma(X_t)dW_t$, and has CCF given by $\psi_{AD} = e^{\alpha_t + \beta_t^\top X_t}$ from (7). We assume for simplicity that there is no stochastic volatility, so that $H_1 = 0$. Then, α_t and β_t satisfy the following equations (that simplify (8)–(9)) under appropriate boundary conditions:

$$\dot{\beta}_t = -K_1^\top \beta_t, \quad (10)$$

$$\dot{\alpha}_t = -K_0^\top \beta_t - \frac{1}{2} \beta_t^\top H_0 \beta_t. \quad (11)$$

⁵Interesting extensions to (4)–(6) include models in which parameters are permitted to be linear-quadratic functions of the state, e.g., Cheng and Scaillet’s (2007) LQJD, with linear jump component and linear-quadratic diffusion part. We show that the AJD framework already provides a good approximation to the behaviour of power markets, and do not investigate theoretical extensions here.

For K_0 and K_1 constant, and H_0 a positive-definite matrix, standard integration gives the following analytical solutions to (10)–(11):

$$\beta_i(t) = e^{(T-t)\Delta_i} \beta_i(T), \quad (12)$$

$$\begin{aligned} \alpha(t) &= \alpha(T) + \beta_i(T)(K_0)_i [e^{(T-t)\Delta_i} - 1] \Delta_i^{-1} \\ &+ (1/2) \beta_i(T) \beta_j(T) [e^{(T-t)(\Delta_i + \Delta_j)} - 1] [\Delta_i + \Delta_j]^{-1} (H_0)_{ij}. \end{aligned} \quad (13)$$

We implicitly sum over all repeated indices in (13), $\Delta_r := (K_1)_{rr}$ is the r th diagonal element of K_1 , and $i, j = 1, 2, \dots, n$. We further assume that there is no drift in the mean μ , so that $K_0 = 0$ (which gives the desired model (2)). Then, (13) simplifies to:

$$\alpha(t) = \alpha(T) + (1/2) \beta_i(T) \beta_j(T) [e^{(T-t)(\Delta_i + \Delta_j)} - 1] [\Delta_i + \Delta_j]^{-1} (H_0)_{ij}. \quad (14)$$

To summarize, the diffusive model (2) captures mean reversion of three independent risk factors, and their volatilities. The model has a CCF $\psi_{AD} = e^{\alpha_t + \beta_t^\top X_t}$, in terms of parameters α_t and β_t , that (through (12) and (14)) are themselves in terms of t , T , K_1 (i.e. $\kappa_1, \kappa_2, \kappa_3$), and H_0 (i.e. $\sigma_1, \sigma_2, \sigma_3$), which are either known, or can be estimated.

On the other hand, we model spikes using a three-state regime-switching process. This technique enables us to flexibly capture observed spike behaviour such as duration of more than one day. The *spike* model has a pure affine jump (AJ) representation, in the spike state \tilde{X} .

Proposition 1 *The conditional characteristic function ψ_{AJ} , corresponding to an m -state regime-switching process, has exponential-affine form $e^{\tilde{\alpha}(T-t, u) + \tilde{\beta}(T-t, u)^\top \tilde{X}_t}$, with \tilde{X}_t an $m \times 1$ vector with i th element unity, and the remaining elements zero (corresponding to an active i th spike state). The regime-switching process has an equivalent representation as a pure affine jump process, through appropriate mapping of parameters in (6).*

We give a complete proof in Appendix A.2. The importance of Proposition 1 is that it enables us to model spikes realistically using a regime-switching process, and then to write the spike model as an AJ (with corresponding CCF ψ_{AJ}), which can be combined with the diffusive CCF ψ_{AD} , to give the full AJD CCF $\psi = \psi_{AD}\psi_{AJ}$. In other words, we pre-calibrate and model spikes independently of the diffusive risk factors, but are then able to reintroduce the spikes when pricing derivatives.

It remains to estimate the parameters of the affine diffusive CCF ψ_{AD} , which follows by writing the log spot and log forwards in state-space form, and using the recursive Kalman filter to construct a likelihood function, which we then optimize numerically.⁶

2.3 State-space representation of diffusive factors

We first write the affine diffusion (2) in state-space form, i.e., in terms of a measurement equation, which provides a connection between observable spot and forward prices, and unobservable state components; and a transition equation, which describes the dynamic evolution of the diffusive risk factors. State-space techniques have been widely applied in econometrics, and are useful for time-varying coefficient and stochastic volatility models.⁷

Measurement equation

In an AJD, the logarithm of the price of an interval forward, $\ln(f_t(T_1, T_2))$, can be approximated by an affine function of the spike and diffusive risk factors \tilde{X}_t and X_t , where t is the pricing time, and T_1 and T_2 are the start and end of the delivery period. By definition

⁶Other methods of econometric estimation of the parameters of affine diffusions (ADs) are covered by Singleton (2001), who uses the closed-form structure of the CCF of discretely-sampled observations from an AD, with non-latent state variables and Fourier inversion, to derive (conditional) maximum likelihood estimators, and shows that these can be computationally demanding for non-scalar X . He also constructs generalized method-of-moments estimators directly from the partial derivatives of the CCF, evaluated at zero, which avoids the need for Fourier inversion.

⁷For detailed discussions see Durbin and Koopman (2004) and Harvey et al. (2004); and Coulon and Howison (2009) for an interesting application to electricity price modelling.

of the risk-neutral probability measure Q , we have:

$$\begin{aligned} \ln(f_t(T_1, T_2)) &= \ln \left(E^Q \left[(T_2 - T_1)^{-1} \sum_s S_s \middle| \mathcal{F}_t \right] \right) \\ &= \ln \left(E^Q \left[(T_2 - T_1)^{-1} \sum_s e^{(\theta_s + \tilde{\gamma}^\top \tilde{X}_s + \gamma^\top X_s)} \middle| \mathcal{F}_t \right] \right) \end{aligned} \quad (15)$$

$$= \ln \left((T_2 - T_1)^{-1} \sum_s e^{\theta_s} E^Q \left[e^{\tilde{\gamma}^\top \tilde{X}_s} \middle| \mathcal{F}_t \right] E^Q \left[e^{\gamma^\top X_s} \middle| \mathcal{F}_t \right] \right) \quad (16)$$

$$= \ln \left((T_2 - T_1)^{-1} \sum_s e^{\theta_s} e^{\tilde{\alpha}_t(s) + \tilde{\beta}_t^\top(s) \tilde{X}_t} e^{\alpha_t(s) + \beta_t^\top(s) X_t} \right) \quad (17)$$

$$\approx \ln \left((T_2 - T_1)^{-1} \sum_s e^{\theta_s + \tilde{\alpha}_t(s) + \alpha_t(s) + \beta_t(s)^\top X_t} \right) \quad (18)$$

$$:= \ln \left((T_2 - T_1)^{-1} \sum_s e^{\theta_t(s) + \beta_t(s)^\top X_t} \right), \quad (19)$$

with the summations taken over $s \in [T_1, T_2]$, and $s > t$. Equation (15) follows from the spot model (1), and (16) from independence of spike and diffusive risk factors, and the deterministic nature of the yearly and weekly patterns in θ_s . Equation (17) follows directly from the definition of the CCF (7), and (18) by setting $\tilde{\beta}_t(s) = 0$ (the justification comes from numerical observation: this term is close to zero in practice for $T_1 - t \gtrsim 20$ days). Equation (19) follows by defining $\theta_t(s) := \theta_s + \tilde{\alpha}_t(s) + \alpha_t(s)$. The coefficients $\alpha_t(s)$ and $\beta_t(s)$ are solutions of the equations (10)–(11), with dependence on s . We pre-calibrate the terms θ_s and $\tilde{\alpha}_t(s)$ following the method explained before Section 2.1.⁸

In order to write $\ln(f_t(T_1, T_2))$ as an affine function of the diffusive risk factors X_t , we

⁸We assume for simplicity in this paper that $P = Q$, i.e., the physical and risk-neutral measures are identical. Culot (2003, chapter 2) shows that an AJD under P may be written as an AJD under Q , through appropriate transformation of parameters.

approximate (19) by its first-order Taylor expansion about $X_t = 0$:

$$\ln \left((T_2 - T_1)^{-1} \sum_s e^{\theta_t(s)} \right) + \left(\frac{\sum_s \beta_t(s) e^{\theta_t(s)}}{\sum_s e^{\theta_t(s)}} \right)^\top X_t := \bar{\alpha}_t(T_1, T_2) + \bar{\beta}_t^\top(T_1, T_2) X_t. \quad (20)$$

We have written both the log spot *and* all the log forwards as affine functions of the *same* risk factors X_t . This manipulation is very convenient for estimation. Denote by $f_t^{(i)}$ the i th forward with delivery period $[T_1^{(i)}, T_2^{(i)}]$, $i = 1, 2, \dots, M$ (and M is the number of included forward products). Given expressions (1) and (20) for spot and forward prices, and with spikes removed from the spot series, we write:

$$Y_t := \begin{pmatrix} \ln(S_t) \\ \ln(f_t^{(1)}) \\ \vdots \\ \ln(f_t^{(M)}) \end{pmatrix} = \begin{pmatrix} \theta_t \\ \bar{\alpha}_t(T_1^{(1)}, T_2^{(1)}) \\ \vdots \\ \bar{\alpha}_t(T_1^{(M)}, T_2^{(M)}) \end{pmatrix} + \begin{pmatrix} \gamma^\top \\ \bar{\beta}_t^\top(T_1^{(1)}, T_2^{(1)}) \\ \vdots \\ \bar{\beta}_t^\top(T_1^{(M)}, T_2^{(M)}) \end{pmatrix} X_t + \epsilon_t := r_t + A_t X_t + \epsilon_t,$$

in which r_t is $(M + 1) \times 1$, A_t is $(M + 1) \times n$, and X_t and ϵ_t are $n \times 1$. Hence, the measurement equation is given by:

$$Y_t = r_t + A_t X_t + \epsilon_t. \quad (21)$$

The observed variables Y_t include the log spot price, and a selection of log forward prices at various liquid maturities. The choice of variables in Y_t is important, since the spot provides information on high-frequency short-term movements, while forwards contain valuable information concerning market participants' expectations of future economic conditions, and have an impact on estimation of lower-frequency, longer-term movements. In our APX example, we include month+1, month+2, quarter+1, quarter+2, quarter+3,

quarter+4, year+1 and year+2 forwards. The term $\epsilon_t \sim N(0, \Phi_t)$ in (21) can be interpreted as a “measurement error”, and is added to deal with observations that are not exactly contemporaneous, or with model shortcomings that make it impossible to reproduce all of the observed prices.

Transition equation

It is well known that an affine diffusion such as (2) can be discretized to give:

$$X_t = s_t + B_t X_{t-1} + \epsilon_t, \quad (22)$$

in which X_t and s_t are $n \times 1$, B_t is $n \times n$, and $\epsilon_t \sim N(0, \Psi_t)$. The elements of s_t and B_t solve ordinary differential equations with appropriate boundary conditions, as do the elements of the covariance matrix Ψ_t . Equation (22) models the risk factor evolution over time.

Kalman filter

The Kalman filter is used recursively to compute an estimate of the state variables at time t , given available information \mathcal{F}_t . When model innovations and initial unobserved variables are normally distributed, the Kalman filter enables convenient construction of the likelihood function. In the linear Gaussian state-space framework, the measurement and transition equations are given by (21) and (22), where Y_t is observed, r_t , A_t , s_t , B_t are coefficients (from the solution of ODEs, of known *form*, but with unknown parameters, namely κ_i and σ_i), and ϵ_t and ε_t are independent innovations. If the system matrices r_t , A_t , B_t , s_t , Φ_t , Ψ_t are known and nonstochastic, (so that they can change in a predetermined way over time, but may depend upon unknown parameters, which can be estimated), then the Kalman filter gives a minimum mean squared error (MSE) estimator of X_t conditional on \mathcal{F}_t . If the

assumption of normality is relaxed, the estimator still minimizes the MSE within the class of linear estimators. The Kalman filter is used to construct the log likelihood as follows:

$$\ln(L(Y)) = -\frac{(M+1)N}{2} \ln(2\pi) - \frac{1}{2} \sum_{t=1}^N \ln |F_{t|t-1}| - \frac{1}{2} \sum_{t=1}^N \nu_{t|t-1}^\top F_{t|t-1}^{-1} \nu_{t|t-1}, \quad (23)$$

with prediction error $\nu_{t|t-1} := Y_t - Y_{t|t-1}$; and $Y_{t|t-1}$ is the conditional forecast of Y_t ; $F_{t|t-1}$ is the conditional variance of the prediction error; and N is the sample size. A derivation of this standard result is given in Appendix A.1, for ease of reference. Estimation using (23) gives values for all free parameters *and* an estimate of the unobserved state X_t ; these are used later in both simulation and pricing.

Optimization problem

We estimate the diffusive reversion parameters and volatilities $\Pi = \{\kappa_1, \kappa_2, \kappa_3, \sigma_1, \sigma_2, \sigma_3\}$, from (23) by maximum likelihood, subject to the constraint that $\kappa_i > 0$ and $\sigma_i > 0$, i.e., $\max_{\pi \in \Pi} \ln(L(Y; \pi))$. We calibrate the variance matrix Φ_t of ϵ_t in (21) by minimizing the difference between market and model quarterly forwards. For speed, and to improve the fit of the volatility term structure, subsequent calibrations can proceed as follows: we fix $\Pi_L = \{\kappa_1, \kappa_2, \kappa_3, \sigma_1\}$, and estimate $\Pi_C = \{\sigma_2, \sigma_3\}$ by matching available at-the-money *forward option* prices, quoted on the last date of the data sample, i.e.,

$$\min_{\pi_C \in \Pi_C} \sum_{i=1}^J (\text{FO}_{\text{mod}}^i(\pi_L, \pi_C) - \text{FO}_{\text{mkt}}^i)^2.$$

Hence, we can minimize the pricing errors on the forward options, where J is the number of options to be fitted; FO_{mkt} are market prices; and FO_{mod} are the model-implied prices, computed using the pricing methods detailed briefly below. We perform the numerical optimizations using a BFGS quasi-Newton search with numerical gradient computation,

and linear backtracking to choose step length. In practice, we set $J = 2$, and fit month+1 (short-term) and year+1 (long-term) forward options.

Unconstrained maximum likelihood tends to underestimate the volatility term structure. Essentially, forward volatility is modelled only up to the level that can be explained by the volatility of the three underlying factors. Constraining the maximum likelihood to match the volatility term structure will “correct” the estimation when the number of risk factors is not sufficient to fully describe the joint dynamics of the forwards of all maturities. Attention must be given to the choice of constraints, since it may not be possible to match the entire volatility term structure due to insufficient degrees of freedom.

2.4 Hourly spot model: historical profile sampling

Given the daily model, we generate the hourly spot prices at some given future date by random sampling from a historical dataset. For a given future day with day number $d = 1, 2, \dots, 365$, we assign an hourly profile, i.e., a spot price for each of the hours 1–24, that has been selected from all previously observed hourly profiles; conditional on matching (a) the day type: weekday or weekend, and (b) the spike type: spike day or no spike day. We refer to this procedure as “historical profile sampling” (or HPS). The hourly HPS model takes as input the future daily mean spot price for day d , constructed by stochastic simulation from the estimated daily model. Then, the historical sampling dataset is constructed by:

- (WEEKDAY, NO SPIKE) For a weekday d that has daily mean spot below a threshold τ , the historical sampling dataset includes all weekdays that have daily mean spot lower than τ , and that are within ± 20 day numbers of d . For the APX example, we choose $\tau = 70$ euros/MWh. A probability is assigned to each observation from the sampling dataset using a triangular density function, which takes its maximum at d ,

and gives a non-zero probability to *all* profiles in the *sampling* dataset.

- (WEEKDAY, SPIKE) For a weekday d that has daily mean spot above the threshold τ , the sampling dataset includes all historical hourly profiles observed on weekdays for which the daily mean exceeded τ . We again assign observations a probability according to a triangular density function.
- (WEEKEND) For Saturdays (Sundays/holidays) d , the sampling dataset includes all Saturdays (Sundays/holidays) from previous years that are within ± 20 day numbers of d . We treat public holidays as Sundays.

We first normalize the historical hourly spot prices, by dividing by the daily mean. Then, $S_t[\text{hourly}] = g_t[\text{hourly}]S_t[\text{daily}]$, where g is the normalized hourly profile. The HPS is nonparametric in the sense that our only choice is that of g ; and it is coherent with the daily model, since the volatility of the daily mean is unchanged. Using this method, we are able to approximate quite complicated intraday behaviour, such as winter evening peaks, summer midday peaks, and the different patterns observed on weekdays and weekends. Subsequent work by Schneider (2010) applies this technique to generation of hourly spot trajectories for the German EEX. Various extensions are possible, including non-constant threshold τ , and alternative clustering of historical data, e.g., the split of weekday *non-spikes* into (a) weekday non-spike low day, and (b) weekday non-spike high day.

3 Empirical example

The data under study consist of the APX hourly electricity spot prices in euros/MWh (www.apx.nl). The exchange opened on 2 March 1999, and spot data is available on weekends and holidays. The early part of the series revealed a number of likely data errors,

and missing data, and we discard the period 2 March 1999 to 31 December 2000.⁹ Our illustrative sample ends on 2 June 2005, which gives a total of 1,614 daily observations. Unless otherwise indicated, “spot” refers to the baseload daily average series.

The three largest spikes occur at 660.34 euros/MWh, 368.80 euros/MWh and 637.37 euros/MWh (11, 12, and 13 August 2003, respectively), and are particularly striking. The spot series is highly volatile, positively skewed and leptokurtic, which is consistent with many empirical findings, and which reflects spike components. The standard deviation of the spot price is 33.10 euros/MWh, which is 92% of the mean value. Intraweek seasonality is evident in the autocorrelation function of the daily baseload spot series. Strong intraday patterns can also be seen in the autocorrelation function of the hourly baseload spot. The one-day log returns R_t and squared log returns R_t^2 are clearly not independent, and the first-order autocorrelations of R_t and R_t^2 are -0.203 and 0.153 , and are statistically significant. The Ljung-Box statistics for up to fourteenth-order serial correlation in R_t and R_t^2 are 807.77 and 173.60, and are highly significant when compared to the limiting chi-squared distribution.

Average peak-period prices and unconditional volatility are much higher than during the off-peak period, and are highest around midday and in the early evening; while skewness and kurtosis reflect the occurrence of spikes in peak hours, and their absence from the off-peak. Intraweek seasonal patterns are apparent in both level and volatility, and there is evidence that price behaviour is very different on weekdays and at the weekend. There is strong dependence in the baseload spot, and considerable persistence in the off-peak period prices. The autocorrelation function declines slowly for all hours 18:00–09:00. There are 230–231 observations per day type. There is little correlation between Monday–Wednesday

⁹Also, the spot series changes dramatically at 1 January 2001, due to changes in market infrastructure. Prior to 2001, three regulated tariffs were in place for end users, and all exchange bids were made at these prices.

baseload and the corresponding days in the previous 2 weeks. However, this dependence is considerably higher for Thursday–Sunday. First and second-order periodic autocorrelations show that Monday has a large impact on Tuesday and Wednesday, while Thursday and Friday, and Saturday and Sunday, are also closely related. These results strongly support explicit modelling of intraweek behaviour, and weekday/weekend levels.

We also use APX forward data, in euros/MWh, from Platts, an independent energy market data publishing company (www.platts.com). For instance, on 2 January 2001, we have quotations for baseload forwards Y2001D003, Y2001W02, Y2001M02, Y2001M03, Y2001Q2 and Y2002, i.e., day 3 (3 January 2001), week 2 (8-14 January 2001), month 2 (1-28 February 2001), month 3 (1-31 March 2001), quarter 2 (1 April-30 June 2001), and year 2002 (1 January-31 December 2002). The price quotations are the means of the bid and ask prices. A typical APX forward trade would be for 5–15 MW of power, in 5 MW increments. Finally, we use quotations of options on APX forward contracts, taken from ICAP Energy (eu.icapenergy.com). In the example, we use two such options, quoted on 26 May 2005: Option no.1 (an at-the-money put), on a Y2006 forward, with underlying forward price(=strike) 46.53 euros/MWh, maturity 17 December 2005, price 2.61 euros/MWh, and implied volatility 19.0%; and Option no.2 (an at-the-money call), on a Y2005M07 forward, with underlying forward price 43.75 euros/MWh, maturity 27 June 2005, price 3.46 euros/MWh, and implied volatility 67.2%.

3.1 Results

We now report calibrated (spikes, seasonal patterns) and state-space estimated (diffusion component) parameters. The 3 spike levels are:

$$(\text{nospike}, \text{level1}, \text{level2}, \text{level3}) = (0.66071, 1.49352, 2.79031),$$

with one-day transition matrix:

$$G_{1\text{-day}} = \begin{pmatrix} 0.966 & 0.004 & 0.026 & 0.004 \\ 0.370 & 0.397 & 0.204 & 0.029 \\ 0.370 & 0.029 & 0.572 & 0.029 \\ 0.370 & 0.029 & 0.204 & 0.397 \end{pmatrix}.$$

The exponentials of the spike levels approximately translate to multiplicative factors of an average spot level under “small”, “medium” and “large” spikes, i.e., $e^{0.66071} \approx 1.94$, $e^{1.49352} \approx 4.45$, and $e^{2.79031} \approx 16.29$. The *one-period* transition matrix $G_{1\text{-day}}$ is calculated from the instantaneous transition matrix G of Proposition 1 as $G_{1\text{-day}} = e^{G-I_n} := (p_{ij}(1))$, with $i, j = 0, 1, 2, 3$ (no spike, and spike levels 1, 2 and 3). We see that $p_{00}(1) \approx 0.966$, i.e., there is a 3.4% probability of some spike arriving in the next period given that the spot price is in a non-spike regime. However, once a spike regime has been entered, the non-negligible probabilities of remaining in some spike regime reflect the ability of the model to capture possible multiple-day spike durations. The *long-run* transition matrix $\lim_{q \rightarrow \infty} e^{(G-I_n)q} := (p_i(\infty))$ gives $(p_0(\infty), p_1(\infty), p_2(\infty), p_3(\infty)) \approx (0.917, 0.009, 0.064, 0.009)$, i.e., the overall probability of some spike arriving is 8.2%, which closely matches market observations. A rough calculation based upon the one-period transition matrix gives the expected number of periods between leaving the no-spike regime and returning to the no-spike regime after a spike, as $\sum_{q \in \mathbb{N} \setminus \{1\}} \sum_{i \neq 1} \{e^{(G-I_n)q}\}_{i1} \approx 1.7$ (days), which is again reasonable.

The yearly pattern is:

$$(\hat{\rho}_1, \hat{\rho}_2, \hat{\rho}_3) = (3.71177, 0.08064, 2.17914),$$

and corresponds to high winter and low summer prices, while the weekly pattern:

$$\begin{aligned} & (\text{Mon, Tue, Wed, Thu, Fri, Sat, Sun}) \\ & = (1.09953, 1.09809, 1.10188, 1.12976, 1.06094, 0.83001, 0.67980) \end{aligned}$$

reveals relatively high Monday–Thursday prices, that fall slightly on Friday, and are significantly lower on Saturday and Sunday, as expected. The estimated mean reversion parameters:

$$(\hat{\kappa}_1, \hat{\kappa}_2, \hat{\kappa}_3) = (0.32, 0.0068403, 0.000026268)$$

represent diffusion half-lives of 2.2 days, 101.3 days, and more than 72 years, respectively; and we see from:

$$(\hat{\sigma}_1, \hat{\sigma}_2, \hat{\sigma}_3) = (0.16943, 0.038299, 0.008129)$$

that the volatilities are $\hat{\sigma}_1 > \hat{\sigma}_2 > \hat{\sigma}_3$. In-sample forward curves are closely matched. The model volatility term structure (model vol), estimated over 27 May 2002 to 26 May 2005, slightly overstates the empirical volatility term structure (empirical vol), although the absolute difference decreases with time-to-maturity $T - t$ (days). Unsurprisingly, the largest error corresponds to the more volatile spot, although generally the error is small:

$T - t$ (days)	product	empirical vol	model vol
0	spot	0.351	0.375
30	month+1	0.0375	0.0392
60	month+2	0.0182	0.0322
91	quarter+1	0.0144	0.0280
182	quarter+2	0.0133	0.0172
273	quarter+3	0.0119	0.0124
364	year+1	0.0152	0.0125
365	quarter+4	0.0105	0.0106
729	year+2	0.00586	0.00975

3.2 Model assessment

To assess the estimated models' ability to robustly reproduce observed price behaviour, we generate multiple simulated price series, across some time period of interest (we choose calendar year 2003). We identify various model-independent statistical or business key features of the observed time series, such as spike duration, autocorrelation, or intraday seasonal patterns. Using the simulated spot series, we then build a Gaussian kernel-smoothed scenario distribution for each key feature, which is compared to that which is calculated from historical data. This method gives a more detailed picture of model performance than analysis of the moments of the scenario distribution alone.

A selection of output is plotted in Figures 1–6. We generate 1,000 daily and 250 hourly spot scenarios. Figures 1–6 illustrate the intraday behaviour of the HPS model, using hourly scenarios. In each figure, we compute the ratio of the spot mean conditional on hour, to the unconditional spot mean, given both season (summer: April–September, or winter: October–March) and day type (weekday, Saturday or Sunday). The observed

mean ratios for each hour are linked with a solid line (cubic spline), and the bands around each observed mean ratio correspond to 90% and 95% *model scenario* bands (the kernel distribution is not plotted here). We see that the model accurately captures the midday and (when appropriate) evening peaks in prices, given different seasons and day types.

4 Conclusions

We propose a practical model for daily electricity spot and forward prices, with regime-switching spikes, that incorporates various stylized features of power prices, including mean reversion and seasonal patterns. We model spike behaviour flexibly within the affine jump diffusion framework by using a regime-switching process, that enables us to replicate the short duration and extreme nature of price spikes. The model is estimated using both spot and forward market price data, in a two-step procedure, with pre-calibrated “structural” elements, and diffusive parameters that we estimate using maximum likelihood and the Kalman filter (see also Cartea and Figueroa (2005)). Spot data is appropriate for estimation of short-term shocks, spikes, and intraweek seasonality, while the coarser granularity of the forward curve is used to estimate medium/long-term shocks, and annual seasonality. The calibration procedure is motivated by the properties and limitations of power price data, and by the planned uses of the model. We also develop a simple nonparametric model for hourly spot prices, that builds upon the daily model. We illustrate the performance of the models using a simulation-based assessment methodology, which shows in particular the ability of the hourly model to sensibly reproduce complicated intraday patterns. We use several results on affine jump diffusions to give closed-form solutions for interesting power derivatives, in contrast to many “classical” power models, while remaining empirically tractable.

In short, we describe a general and flexible treatment of power (and energy) price

modelling, that covers many important characteristics of electricity prices, and that can be used efficiently for derivative pricing and hedging applications. It is straightforward to adapt the model to the specifics of a particular market, and variations on the approach presented in the paper have been used successfully in real business settings.

A number of extensions of the research in this paper are possible. We can imagine potential model modifications for a more realistic description of the observed spot series, e.g., by changing the annual pattern to account for multiple annual peaks; adapting the regime-switching process to allow for time-dependent spikes; or weakening the restrictions on the AJD coefficient matrices to enable modelling of stochastic volatility; and correlations between risk factors; as well as multivariate power/fuel models; and removal of linearity in the pricing expressions (with use of nonlinear filtering techniques). These would come at the expense of an increase in the computational burden, and would obscure the main messages of this paper.

A Appendix

A.1 Kalman filter derivation of log likelihood

The basic filter comprises prediction and updating algorithms, which estimate X_t given \mathcal{F}_{t-1} and \mathcal{F}_t , respectively. We use the following notation: $X_{t|t-1} := E[X_t|\mathcal{F}_{t-1}]$ is the conditional expectation of X_t , $P_{t|t-1} := E[(X_t - X_{t|t-1})(X_t - X_{t|t-1})^\top]$ is the conditional covariance matrix of X_t , $Y_{t|t-1} := E[Y_t|\mathcal{F}_{t-1}]$ is the conditional forecast of Y_t , $\nu_{t|t-1} := Y_t - Y_{t|t-1}$ is the prediction error, and $F_{t|t-1} := E[\nu_{t|t-1}^2]$ is the conditional variance of the prediction error.

(a) Prediction. Given \mathcal{F}_{t-1} , compute state $X_{t|t-1}$ and covariance $P_{t|t-1}$, and estimate Y_t :

$$X_{t|t-1} = s_t + B_t X_{t-1|t-1}, \quad (24)$$

$$P_{t|t-1} = B_t P_{t-1|t-1} B_t^\top + \Psi_t, \quad (25)$$

$$\nu_{t|t-1} = Y_t - A_t X_{t|t-1} - r_t, \quad (26)$$

$$F_{t|t-1} = A_t P_{t|t-1} A_t^\top + \Phi_t. \quad (27)$$

Given that X_1 and $\{\epsilon_t, \varepsilon_t\}$ are Gaussian, and $F_{t|t-1}$ is positive-definite, then the conditional distribution of Y_t is multivariate normal: $Y_t|\mathcal{F}_{t-1} \sim N_{M+1}(Y_{t|t-1}, F_{t|t-1})$, i.e.,

$$\ln(\phi(Y_t|\mathcal{F}_{t-1})) = -\frac{(M+1)}{2} \ln(2\pi) - \frac{1}{2} \ln |F_{t|t-1}| - \frac{1}{2} \nu_{t|t-1}^\top F_{t|t-1}^{-1} \nu_{t|t-1}. \quad (28)$$

(b) Updating. The inference based on information in the state variables is revised based on realization of the observed variables:

$$X_{t|t} = X_{t|t-1} + P_{t|t-1} A_t^\top F_{t|t-1}^{-1} \nu_{t|t-1}, \quad (29)$$

$$P_{t|t} = P_{t|t-1} - P_{t|t-1} A_t^\top F_{t|t-1}^{-1} A_t P_{t|t-1}, \quad (30)$$

where $P_{t|t-1} A_t^\top F_{t|t-1}^{-1}$ is the ‘‘Kalman gain’’. We assume that the inverse of $F_{t|t-1}$ always exists, i.e., positive-definite, although it could otherwise be replaced by a pseudo-inverse.

(c) Likelihood. Recursive use of (24)–(27) and (29)–(30), with (28), enables us to write the log likelihood $\ln(L(Y)) = \sum_{t=1}^N \ln(\phi(Y_t | \mathcal{F}_{t-1}))$ as:

$$\ln(L(Y)) = -\frac{(M+1)N}{2} \ln(2\pi) - \frac{1}{2} \sum_{t=1}^N \ln |F_{t|t-1}| - \frac{1}{2} \sum_{t=1}^N \nu_{t|t-1}^\top F_{t|t-1}^{-1} \nu_{t|t-1}.$$

□

A.2 Proof of Proposition 1

Part 1

The conditional characteristic function (CCF) of the m -state Markov regime-switching spike process is defined using (7) as:

$$\psi(u, \tilde{X}_t, t, T) := E[e^{u^\top \tilde{X}_T} | \mathcal{F}_t] = E[e^{u^\top \tilde{X}_T} | \tilde{X}_t = e_i],$$

in which \tilde{X}_t is an $m \times 1$ vector of spike risk factors, and e_i is an $m \times 1$ vector with i th element unity and the remaining elements zero, corresponding to an active i th spike state; and e_0 is the no-spike state. Let G be the $(m+1) \times (m+1)$ infinitesimal transition matrix with $G_{ij} dt$ the probability of moving from state $\tilde{X}_t = e_i$ to $\tilde{X}_{t+dt} = e_j$, with $i, j \in \{0, 1, \dots, m\}$.

Then, the probability of moving from spike state $\tilde{X}_t = e_i$ to $\tilde{X}_T = e_j$ is:

$$\text{Prob}(\tilde{X}_T = e_j | \tilde{X}_t = e_i) = e^{((G - I_{m+1})(T-t))_{ij}} := \bar{G}_{ij},$$

where e^{\cdot} is the matrix exponential. So,

$$\psi(u, \tilde{X}_t, t, T) = E[e^{u^\top \tilde{X}_T} | \tilde{X}_t = e_i] = \sum_{j=1}^m \bar{G}_{ij} e^{u^\top e_j} = \sum_{j=1}^m \bar{G}_{ij} e^{u_j}. \quad (31)$$

Assume that (31) can be written using the functional form of an affine jump (AJ) CCF, from (7). Then, $\sum_{j=1}^m \bar{G}_{ij} e^{u_j} = e^{\alpha + \beta^\top e_i}$, or:

$$\ln \sum_{j=1}^m \bar{G}_{ij} e^{u_j} = \alpha + \beta^\top e_i = \alpha + \beta_i. \quad (32)$$

Clearly, $i = 0$ and (32) gives:

$$\alpha = \ln \sum_{j=1}^m \bar{G}_{0j} e^{u_j}, \quad (33)$$

while $i \neq 0$ and (32) gives:

$$\beta_i = \ln \sum_{j=1}^m \bar{G}_{ij} e^{u_j} - \alpha = \ln \left(\sum_{j=1}^m \bar{G}_{ij} e^{u_j} / \sum_{j=1}^m \bar{G}_{0j} e^{u_j} \right). \quad (34)$$

The regime-switching process CCF $\psi(u, \tilde{X}_t, t, T)$ has a form that can be written as the CCF from an AJ, through (32)–(34). Furthermore, (34) satisfies the standard AJ ordinary differential equations, with $\dot{\beta}_i(t) = -\sum_{j=1}^m G_{ij}(e^{\beta_j(t) - \beta_i(t)} - 1)$ derived from either (34), or from the AJ ordinary differential equations. Practically, a regime-switching spike process can be calibrated using pre-identified spikes, and the AJ CCF then written in terms of elements of the transition matrix G .

Part 2

The regime-switching process can be written in the form (6), by appropriate choice of the jump intensity parameters l_0 and l_1 . A jump represents a move from spike state e_i to e_j . Without loss of generality, and to simplify notation, we do not consider the no-spike state e_0 here. From (3)–(6), the arrival rate λ_{ij} of the move from e_i to e_j is given by:

$$\begin{aligned}\lambda_{ij} &= l_{0,ij} + l_{1,ij}^\top e_i = \bar{G}_{ij}, \\ \lambda_{ij} &= l_{0,ij} + l_{1,ij}^\top e_{k \neq i} = 0,\end{aligned}$$

where $l_{0,ij}$ and \bar{G}_{ij} are 1×1 , and $l_{1,ij}$ and e_i are $m \times 1$. Then, $l_{1,ij}^\top e_i = \bar{G}_{ij}$, and:

$$l_{0,ij} = 0, \quad l_{1,ij} = \bar{G}_{ij} e_i. \quad (35)$$

So, from (6) and (35), for an AJ:

$$\lambda(\tilde{X}_t) = (\lambda_{ij}) = l_0 + l_1 \tilde{X}_t, \quad (36)$$

in which $l_0 = (l_{0,ij})$ is $m(m+1) \times 1$, and $l_1 = (l_{1,ij}^\top)$ is $m(m+1) \times m$. The jump from e_i to e_j has fixed size $e_j - e_i$, and so the $m \times 1$ jump amplitude distribution ν is given by:

$$\nu_{ij}(z) = 1, \quad \text{if } z = e_j - e_i, \quad (37)$$

and zero otherwise. We have shown that the elements of the AJ jump distribution in (6) can be written in terms of the transition matrix G . The regime-switching process can thus be transformed into an affine jump (see (36) and (37)), and its CCF derived, both in terms of the transition matrix G and the spike levels e_i . This step completes the proof. \square

A.3 Option pricing in the affine jump diffusion setting

We have developed pricing formulae for various power derivatives. It is always useful to have closed-form pricing solutions, since Monte Carlo methods tend to be computationally expensive, especially when computing options on forwards or option price Greeks.

Forwards are the most commonly traded financial products in power markets. In our framework, daily forwards $f_t(T)$, with maturity $T > t$, have an exponential-affine form in terms of the spike and diffusive risk factors:

$$f_t(T) = E^Q[S_T | \mathcal{F}_t] = E^Q[e^{\theta_t + \tilde{\gamma}^\top \tilde{X}_t + \gamma^\top X_t} | \mathcal{F}_t] = \phi(e^{\theta_t}, 0, (\tilde{\gamma}, \gamma)^\top, 0, (\tilde{X}_t, X_t)^\top, t, T),$$

from (1), and using Duffie et al.'s (2000) *extended* transform ϕ , which generalizes (7):

$$\phi(d_0, d_1, a, b, X_t, t, T) = E^Q[(d_0 + d_1^\top X_T) e^{(a+ib)^\top X_T} | \mathcal{F}_t] := (A_t + B_t X_t) e^{\alpha_t + \beta_t^\top X_t}.$$

Parameters A_t , B_t , α_t and β_t solve a system of complex ordinary differential equations under appropriate boundary conditions. Monthly, quarterly and yearly forwards $f_t(T_1, T_2)$, with $T_2 > T_1 > t$, are approximately exponentially-affine in the diffusive risk factors, from (20).

Duffie et al. (2000) also show that a European call option, with payoff at maturity T given by $\text{PAYOFF}(T) = (S_T - K)^+$, can be written in terms of the G -transform:

$$\begin{aligned} G(y, d_0, d_1, a, b, X_t, t, T) &:= E^Q[(d_0 + d_1^\top X_T) e^{a^\top X_T} \mathbf{1}_{b^\top X_T \leq y} | \mathcal{F}_t] \\ &= \frac{1}{2} \phi(d_0, d_1, a, 0, X_t, t, T) \\ &\quad - \frac{1}{\pi} \int_{\mathbb{R}_+} \frac{1}{v} \text{Im}[\phi(d_0, d_1, a, vb, X_t, t, T) e^{-ivy}] dv. \end{aligned} \quad (38)$$

The method of solving the above integral has been widely studied. This technique of

calculation is applied when matching available at-the-money forward option prices, quoted on the last date of the data sample, in Section 2.¹⁰

As a final word on the risk management applications of AJDs, once a contract has been valued and a position taken, it is important to be able to manage its risk by constructing a hedge against it. Hedging against a risk factor consists of taking an opposite position in proportion to the impact of the risk factor, and is often conducted through use of “Greeks”, which measure this impact. These can easily be computed in the AJD framework, e.g., G_{X_i} (delta), $G_{X_i X_i}$ (gamma) and G_t (theta), from (38).

¹⁰In industry applications, we have also experimented with fast numerical pricing of hourly forwards and call options, based upon the HPS model.

References

- [1] ATKINS, F. Some statistical properties of deregulated electricity prices in Alberta. In *Quantitative Analysis of the Alberta Electricity Market*, W. Walls, Ed. Van Horne Institute, 2006.
- [2] BARLOW, M. A diffusion model for electricity prices. *Mathematical Finance* 12 (2002), 287–298.
- [3] BECKER, R., HURN, A., AND PAVLOV, V. Modelling spikes in electricity prices. *Economic Record* 83 (2007), 371–382.
- [4] BESSEMBINDER, H., AND LEMMON, M. Equilibrium pricing and optimal hedging in electricity forward markets. *Journal of Finance* 57 (2002), 1347–1382.
- [5] BOROVKOVA, S. The forward curve dynamic and market transition forecasts. In *Modelling Prices in Competitive Electricity Markets*, D. Bunn, Ed. Chichester: John Wiley and Sons, 2004, pp. 267–284.
- [6] BOROVKOVA, S., AND PERMANA, F. Modelling electricity prices by the potential jump-diffusion. In *Stochastic Finance* (2006), M. do Rosário Grossinho, A. Shiryaev, M. Esquível, and P. Oliveira, Eds., Springer.
- [7] BOTTAZZI, G., SAPIO, S., AND SECCHI, A. Some statistical investigations on the nature and dynamics of electricity prices. *Physica A* 355 (2005), 54–61.
- [8] BOYLE, P., BROADIE, M., AND GLASSERMAN, P. Simulation methods for security pricing. *Journal of Economic Dynamics and Control* 21 (1997), 1267–1321.

- [9] CARTEA, Á., AND FIGUEROA, M. Pricing in electricity markets: A mean reverting jump diffusion model with seasonality. *Applied Mathematical Finance* 12 (2005), 313–335.
- [10] CHENG, P., AND SCAILLET, O. Linear-quadratic jump-diffusion modelling. *Mathematical Finance* 17 (2007), 575–598.
- [11] CHRISTENSEN, T., HURN, A., AND LINDSAY, K. It never rains but it pours: Modeling the persistence of spikes in electricity prices. *Energy Journal* 30 (2009), 25–48.
- [12] CORTAZAR, G., AND SCHWARTZ, E. Implementing a stochastic model for oil futures prices. *Energy Economics* 25 (2003), 215–238.
- [13] COULON, M., AND HOWISON, S. Stochastic behaviour of the electricity bid stack: From fundamental drivers to power prices. *Journal of Energy Markets* 2 (2009).
- [14] CULOT, M. An integrated affine jump diffusion framework to manage power portfolios in a deregulated market. unpublished PhD thesis, Department of Industrial Engineering, Université Catholique de Louvain, 2003.
- [15] DAI, Q., AND SINGLETON, K. Specification analysis of affine term structure models. *Journal of Finance* 55 (2000), 1943–1978.
- [16] DE JONG, C. The nature of power spikes: A regime-switch approach. *Studies in Nonlinear Dynamics and Econometrics* 10 (2006).
- [17] DIKO, P., LAWFORD, S., AND LIMPENS, V. Risk premia in electricity forward prices. *Studies in Nonlinear Dynamics and Econometrics* 10 (2006).
- [18] DUFFIE, D., FILIPOVIĆ, D., AND SCHACHERMAYER, W. Affine processes and applications in finance. *Annals of Applied Probability* 13 (2003), 984–1053.

- [19] DUFFIE, D., PAN, J., AND SINGLETON, K. Transform analysis and asset pricing for affine jump-diffusions. *Econometrica* 68 (2000), 1343–1376.
- [20] DURBIN, J., AND KOOPMAN, S. *Time Series Analysis by State Space Methods*. Oxford: Oxford University Press, 2004.
- [21] ERAKER, B. Do stock prices and volatility jump? Reconciling evidence from spot and option prices. *Journal of Finance* 59 (2004), 1367–1403.
- [22] ERLWEIN, C., BENTH, F., AND MAMON, R. HMM filtering and parameter estimation of an electricity spot model. *Energy Economics* 32 (2010), 1034–1043.
- [23] ESCRIBANO, Á., PEÑA, J., AND VILLAPLANA, P. Modelling electricity prices: International evidence. *Oxford Bulletin of Economics and Statistics* 73 (2011), 622–650.
- [24] GEMAN, H., AND RONCORONI, A. Understanding the fine structure of electricity prices. *Journal of Business* 79 (2006), 1225–1261.
- [25] HALDRUP, N., NIELSEN, F., AND NIELSEN, M. A vector autoregressive model for electricity prices subject to long memory and regime switching. *Energy Economics* 32 (2010), 1044–1058.
- [26] HALDRUP, N., AND NIELSEN, M. A regime switching long memory model for electricity prices. *Journal of Econometrics* 135 (2006), 349–376.
- [27] HARVEY, A., KOOPMAN, S., AND SHEPHARD, N., Eds. *State Space and Unobserved Component Models: Theory and Applications*. Cambridge: Cambridge University Press, 2004.
- [28] HUISMAN, R., AND MAHIEU, R. Regime jumps in electricity prices. *Energy Economics* 25 (2003), 425–434.

- [29] JOHANNES, M. The statistical and economic role of jumps in continuous-time interest rate models. *Journal of Finance* 59 (2004), 227–260.
- [30] JOSKOW, P. Restructuring, competition and regulatory reform in the U.S. electricity sector. *Journal of Economic Perspectives* 11 (1997), 119–138.
- [31] KANAMURA, T., AND ŌHASHI, K. On transition probabilities of regime switching in electricity prices. *Energy Economics* 30 (2008), 1158–1172.
- [32] KARAKATSANI, N., AND BUNN, D. Diurnal reversals of electricity forward premia. mimeo: Department of Decision Sciences, London Business School, 2005.
- [33] KARATZAS, I., AND SHREVE, S. *Brownian Motion and Stochastic Calculus*, second ed. New York: Springer, 1999.
- [34] KÅRESEN, K., AND HUSBY, E. A joint state-space model for electricity and futures prices. *Energy Power Risk Management* 7 (2002).
- [35] KIESEL, R., SCHINDLMAYR, G., AND BÖRGER, R. A two-factor model for the electricity forward market. *Quantitative Finance* 9 (2009), 279–287.
- [36] KNITTEL, C., AND ROBERTS, M. An empirical examination of restructured electricity prices. *Energy Economics* 27 (2005), 791–817.
- [37] KOEKEBAKKER, S., AND OLLMAR, F. Forward curve dynamics in the Nordic electricity market. *Managerial Finance* 31 (2005), 74–95.
- [38] KOOPMAN, S., OOMS, M., AND CARNERO, M. Periodic seasonal Reg-ARFIMA-GARCH models for daily electricity spot prices. *Journal of the American Statistical Association* 102 (2007), 16–27.

- [39] LUCIA, J., AND SCHWARTZ, E. Electricity prices and power derivatives: Evidence from the Nordic power exchange. *Review of Derivatives Research* 5 (2002), 5–50.
- [40] MANOLIU, M., AND TOMPAIDIS, S. Energy futures prices: Term structure models with Kalman filter estimation. *Applied Mathematical Finance* 9 (2002), 21–43.
- [41] MISIOREK, A., TRUECK, S., AND WERON, R. Point and interval forecasting of spot prices: Linear vs. non-linear time series models. *Studies in Nonlinear Dynamics and Econometrics* 10 (2006).
- [42] MORK, E. Emergence of financial markets for electricity: A European perspective. *Energy Policy* 29 (2001), 7–15.
- [43] PILIPOVIĆ, D. *Energy Risk: Valuing and Managing Energy Derivatives*. New York: McGraw-Hill, 1998.
- [44] ROUTLEDGE, B., SEPPI, D., AND SPATT, C. The “Spark Spread”: An equilibrium model of cross-commodity price relationships in electricity. Graduate School of Industrial Administration, Carnegie Mellon University, 2001.
- [45] SCHNEIDER, S. Power spot price models with negative prices. Munich Personal RePEc Archive, 2010.
- [46] SCHWARTZ, E., AND SMITH, J. Short-term variations and long-term dynamics in commodity prices. *Management Science* 46 (2000), 893–911.
- [47] SIMONSEN, I. Measuring anti-correlations in the Nordic electricity spot market by wavelets. *Physica A* 322 (2003), 597–606.
- [48] SINGLETON, K. Estimation of affine asset pricing models using the empirical characteristic function. *Journal of Econometrics* 102 (2001), 111–141.

- [49] WERON, R. Market price of risk implied by Asian-style electricity options and futures. *Energy Economics* 30 (2008), 1098–1115.
- [50] WERON, R., BIERBRAUER, M., AND TRÜCK, S. Modeling electricity prices: Jump diffusion and regime switching. *Physica A* 336 (2004), 39–48.
- [51] WERON, R., AND PRZYBYŁOWICZ, B. Hurst analysis of electricity price dynamics. *Physica A* 283 (2000), 462–468.
- [52] WILKINSON, L., AND WINSEN, J. What we can learn from a statistical analysis of electricity prices in New South Wales. *Electricity Journal*, April (2002), 60–69.

Daily Seasonal Coefficients (mean)
 for summer months (Apr - Sep inclusive),
 week days (Mon - Fri inclusive)
 Crosses: obs coef
 Lines: 90% CI
 Hooks: 95% CI
 Black line: fitted cubic spline on obs coef




Figure 1: Seasonal mean (summer, week).

Daily Seasonal Coefficients (mean)
 for winter months (Oct - Mar inclusive),
 week days (Mon - Fri inclusive)
 Crosses: obs coef
 Lines: 90% CI
 Hooks: 95% CI
 Black line: fitted cubic spline on obs coef




Figure 2: Seasonal mean (winter, week).

Daily Seasonal Coefficients (mean)
 for summer months (Apr - Sep inclusive), Saturday
 Crosses: obs coef
 Lines: 90% CI
 Hooks: 95% CI
 Black line: fitted cubic spline on obs coef




Figure 3: Seasonal mean (summer, Sat).

Daily Seasonal Coefficients (mean)
 for winter months (Oct - Mar inclusive), Saturday
 Crosses: obs coef
 Lines: 90% CI
 Hooks: 95% CI
 Black line: fitted cubic spline on obs coef




Figure 4: Seasonal mean (winter, Sat).

Daily Seasonal Coefficients (mean)
 for summer months (Apr - Sep inclusive), Sunday
 Crosses: obs coef
 Lines: 90% CI
 Hooks: 95% CI
 Black line: fitted cubic spline on obs coef




Figure 5: Seasonal mean (summer, Sun).

Daily Seasonal Coefficients (mean)
 for winter months (Oct - Mar inclusive), Sunday
 Crosses: obs coef
 Lines: 90% CI
 Hooks: 95% CI
 Black line: fitted cubic spline on obs coef




Figure 6: Seasonal mean (winter, Sun).