Math Meth Oper Res (2009) 69:457–473
DOI 10.1007/s00186-008-0247-4

ORIGINAL ARTICLE

Heavy-tails and regime-switching in electricity prices

Rafał Weron

Received: 10 April 2008 / Accepted: 7 May 2008 / Published online: 28 August 2008

Abstract In this paper we first analyze the stylized facts of electricity prices, in particular, the extreme volatility and price spikes which lead to heavy-tailed distributions of price changes. Then we calibrate Markov regime-switching (MRS) models with heavy-tailed components and show that they adequately address the aforementioned characteristics. Contrary to the common belief that electricity price models 'should be built on log-prices', we find evidence that modeling the prices themselves is more beneficial and methodologically sound, at least in case of MRS models.

Keywords Electricity spot price · Heavy-tails · Spikes · Markov regime-switching · Pareto distribution

1 Introduction

The recent deregulation and introduction of competitive markets has totally changed the landscape of the traditionally monopolistic and government controlled power sectors worldwide. The amount of risk borne by market participants has increased substantially, partially due to the fact that electricity is a very unique commodity. Firstly, it cannot be stored economically and requires immediate delivery, while end-user demand shows high variability and strong weather and business cycle dependence. Secondly, effects like power plant outages or transmission grid (un)reliability add complexity and randomness.

Consequently, for the valuation of electricity contracts we cannot simply rely on models developed for the financial or other commodity markets. Despite numerous

R. Weron (✉)
Hugo Steinhaus Center, Institute of Mathematics and Computer Science,
Wrocław University of Technology, Wrocław, Poland
e-mail: rafal.weron@pwr.wroc.pl

attempts (for reviews see, e.g. Benth et al. 2008; Bunn 2004; Kaminski 2004; Weron 2006), the need for realistic models of price dynamics capturing the unique characteristics of electricity and adequate derivatives pricing techniques still has not been fully satisfied. It is the aim of this paper to study electricity price processes and suggest models that can address the most pertinent characteristics, in particular, the heavy-tailed price distributions and price spikes.

The paper is structured as follows. In Sect. 2 we present the datasets and explain the deseasonalization procedures. In Sect. 3 we study the distributions of price changes. Next, in Sect. 4 we calibrate Markov regime-switching (MRS) models to deseasonalized prices and evaluate their goodness-of-fit. Finally, in Sect. 5 we summarize the results.

2 Data

In this study we use mean daily (baseload) spot prices from four major power markets: EEX (Germany), OMEL (Spain), PJM (US) and NEPOOL (US). For each market the sample totals 1,827 daily observations (or 261 full weeks) and covers the period 2 January 2001–2 January 2006, see Fig. 1.

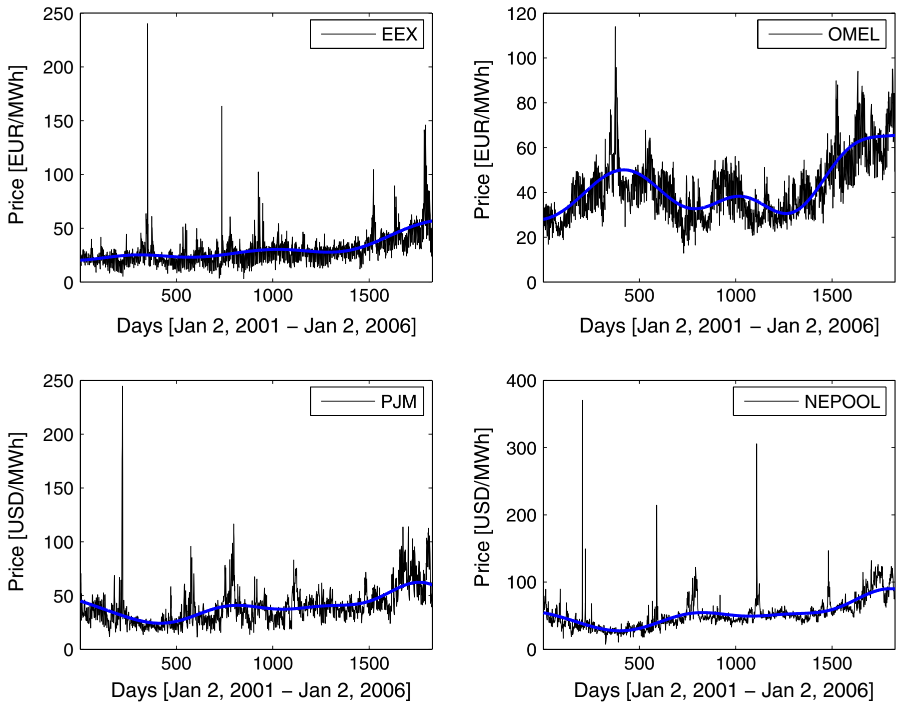

Fig. 1 Mean daily spot prices and their long-term seasonal components (*thick lines*) for EEX, OMEL, PJM and NEPOOL power markets from 2 January 2001 to 2 January 2006

It is well known that electricity demand exhibits seasonal fluctuations, which mostly arise due to changing climate conditions (temperature, number of daylight hours) and business activities (working hours versus leisure periods). Also the supply side (e.g. hydro units) shows seasonal variations in output. These fluctuations in demand and supply translate into the seasonal, mean-reverting behavior of spot electricity prices. In addition to strong seasonality on the annual, weekly and daily level, spot electricity prices exhibit very high volatility and abrupt, short-lived and generally unanticipated extreme price changes known as spikes or jumps (Park et al. 2006; Simonsen 2005; Weron 2008). This behavior can be very well observed in the German EEX market and the New England Pool. Not surprisingly, the share of hydro production in both markets is very small and 'hedging' the volume risk is difficult. In case of tight demand-supply balance—in particular due to unexpected events like plant outages or power line disconnections—there are no units available that can generate electricity instantly and at low marginal costs. Recall, that coal-fired and nuclear plants often need a few hours for start-up. Gas-fired units, on the other hand, have high marginal costs and when they are used the spot prices spike. For this reason the Spanish OMEL market, which has the largest hydro share (ca. 30%) of the four analyzed markets, exhibits different price dynamics with the lowest price volatility and least spikes.

Apart from the above mentioned characteristics, note that all spot prices depicted in Fig. 1 show a clear upward trend towards the end, starting in late 2004. Some prices almost double in just a year time due to a combination of higher fuel prices and the introduction of emission costs in Europe—EU Emission Trading scheme started in January 2005 (Benz and Trück 2006; Paolella and Taschini 2008).

The first crucial step in defining a model for electricity price dynamics consists of finding an appropriate description of the seasonal pattern. There are different suggestions in the literature for dealing with this task; for a recent review consult Trück et al. (2007). Here we follow the 'industry standard' and represent the spot price P_t by a sum of two independent parts: a predictable (seasonal) component f_t and a stochastic component X_t, i.e. $P_t = f_t + X_t$. Further, we let f_t be composed of a weekly periodic part s_t and a long-term seasonal trend T_t, which represents both the changing climate/consumption conditions throughout the year and the long-term non-periodic structural changes.

The deseasonalization is conducted in three steps. First, T_t is estimated from daily spot prices P_t using a wavelet filtering-smoothing technique (for details see, e.g. Trück et al. 2007). Recall, that any function or signal (here: P_t) can be built up as a sequence of projections onto one father wavelet and a sequence of mother wavelets: $S_J + D_J + D_{J-1} + \cdots + D_1$, where 2^J is the maximum scale sustainable by the number of observations. At the coarsest scale the signal can be estimated by S_J. At a higher level of refinement the signal can be approximated by $S_{J-1} = S_J + D_J$. At each step, by adding a mother wavelet D_j of a lower scale $j = J-1, J-2, \ldots$, we obtain a better estimate of the original signal. Here we use the S_8 approximation, which roughly corresponds to annual ($2^8 = 256$ days) smoothing, see the thick lines in Fig. 1. The price series without the long-term seasonal trend is obtained by subtracting the S_8 approximation from P_t. Next, the weekly periodicity s_t is removed by applying the moving average technique (for details see, e.g. Weron 2006) and subtracting the resulting 'mean' weekly pattern. Finally, the deseasonalized prices, i.e. $P_t - T_t - s_t$, are shifted so that the minimum of the new process is the same as the minimum of P_t (the latter alignment is required if log-prices are to be analyzed). The resulting deseasonalized time series X_t can be seen in Fig. 6.

Note that the above procedure differs from the one used by De Jong (2006) for (roughly) the same datasets: moving average technique versus weekly dummies for s_t and wavelet approximation versus a sum of a sinusoid and an exponentially weighted moving average for T_t. We believe that our approach is more robust, especially with respect to the long-term seasonal trend. Due to these differences, the qualitative results can be compared between the papers, but the quantitative rather not.

3 Distributions of electricity prices

It has been long known that financial asset returns are not normally distributed. Rather, the empirical observations exhibit excess kurtosis. This heavy-tailed character of the distribution of price changes has been repeatedly observed in various financial and commodity markets. There are also reports of heavy-tailed behavior of electricity prices. However, to our best knowledge, the studies were conducted either only for one market (Bottazzi et al. 2005; Byström 2005; Eberlein and Stahl 2003; Rachev et al. 2004; Weron 2006), one distributional class (Mugele et al. 2005), samples of relatively small size (Deng and Jiang 2005) or (log-)returns only (Chan and Gray 2006; Khindanova and Atakhanova 2002). Especially the latter two limitations can lead to qualitatively different conclusions. In particular, (log-)returns (i.e. first differences of log-prices) generally exhibit lighter tails than first differences of prices themselves.

Following Weron (2006), we fit Gaussian and three relatively popular and versatile classes of heavy-tailed distributions—hyperbolic, normal inverse Gaussian (NIG) and α-stable—to electricity price changes from the four markets. Calibration of the hyperbolic and NIG distributions is performed via maximum likelihood (ML) as their probability density functions (PDF) are given in explicit form (though using special functions; for numerical details see, e.g. Weron 2004):

$$f_{\mathrm{H}}(x) = \frac{\sqrt{\alpha^2-\beta^2}}{2\alpha\delta K_1(\delta\sqrt{\alpha^2-\beta^2})} e^{-\alpha\sqrt{\delta^2+(x-\mu)^2}+\beta(x-\mu)}, \tag{1}$$

and

$$f_{\mathrm{NIG}}(x) = \frac{\alpha\delta}{\pi} e^{\delta\sqrt{\alpha^2-\beta^2}+\beta(x-\mu)} \frac{\mathrm{K}_1(\alpha\sqrt{\delta^2+(x-\mu)^2})}{\sqrt{\delta^2+(x-\mu)^2}}, \tag{2}$$

respectively. Both laws are characterized by four parameters: steepness (or 'tail index') α, skewness β (with $0 \le |\beta| < \alpha$), scale $\delta > 0$ and location $\mu \in R$. The normalizing constant $\mathrm{K}_\lambda(t)$ is the modified Bessel function of the third kind with index λ (here $\lambda = 1$), also known as the MacDonald function. Both distributions exhibit 'semi-heavy' tails: heavier than Gaussian, lighter than power-law. The log-density of the hyperbolic law forms a hyperbola (hence the name), while the tails of the NIG law satisfy the following asymptotic relation:

$$f_{\mathrm{NIG}}(x) \approx |x|^{-1.5} e^{(\mp\alpha+\beta)x} \quad \text{for} \quad x \to \pm\infty. \tag{3}$$

Inference for the α-stable distribution is more tricky (for details see, e.g. Rachev and Mittnik 2000; Weron 2004). Here we use the regression method, which is slightly less accurate but faster than (approximate) ML. Recall, that with the exception of three special cases ($\alpha = 2, 1, 0.5$), the α-stable PDF does not have a closed form expression and, consequently, the PDF has to be numerically approximated for ML estimation. The regression method proceeds iteratively, until some prespecified convergence criterion is satisfied, by performing regressions on transformations of the α-stable characteristic function:

$$\phi(t) = \begin{cases} \exp(-\sigma^{\alpha}|t|^{\alpha}\{1 + i\beta \mathrm{sign}(t) \tan \frac{\pi\alpha}{2}[(\sigma|t|)^{1-\alpha} - 1]\} + i\mu t), & \alpha \neq 1, \\ \exp(-\sigma|t|\{1 + i\beta \mathrm{sign}(t) \frac{2}{\pi} \log(\sigma|t|)\} + i\mu t), & \alpha = 1. \end{cases} \tag{4}$$

The distribution is characterized by four parameters: tail index $\alpha \in (0, 2]$, skewness $\beta \in [-1, 1]$, scale $\sigma > 0$ and location $\mu \in R$. When $\alpha = 2$, the Gaussian distribution results. When $\alpha < 2$, the variance is infinite and the tails asymptotically decay as a power-law (hence are heavier than those of the hyperbolic and NIG laws).

Let us now return to the dataset. It does not make sense to analyze differences (or returns) of raw prices due to the spurious skewness resulting from weekly seasonality (Weron 2006). If the data is deseasonalized then the distribution of price or log-price differences is more prone to modeling. The Gaussian, hyperbolic, NIG and α-stable fits to the deseasonalized price changes (i.e. $X_t - X_{t-1}$, for $t = 2, \ldots, 1827$) are summarized in Table 1. Right tails of the corresponding cumulative distribution functions (CDF) are plotted in Fig. 2. The goodness-of-fit statistics leave no doubt that the price distributions in all markets have much heavier tails than the Gaussian law. Here we utilize the Anderson–Darling (AD) and Kolmogorov (K) statistics:

$$\mathrm{AD} = n \int_{-\infty}^{\infty} \frac{[F_n(x) - F(x)]^2}{F(x)[1 - F(x)]} dF(x), \tag{5}$$

$$\mathrm{K} = \sup_x |F_n(x) - F(x)|, \tag{6}$$

which measure the distance between the empirical CDF, $F_n(x)$, and the fitted one, $F(x)$; n is the sample size. The Anderson–Darling statistic may be treated as a weighted Kolmogorov statistic which puts more weight to the differences in the tails of the distributions. Approximate critical values for these goodness-of-fit tests can be obtained via the bootstrap technique (for implementation details see Chapter 13 in Čižek et al. 2005), in this study, though, we do not perform hypothesis testing and just compare the test values. Naturally, the lower the values the better the fit.

Apparently, the α-stable distribution yields the best fit for markets with a very spiky price behavior—EEX and NEPOOL. In particular, the tails of the empirical CDF are well approximated by a power-law, see Fig. 2. The Spanish OMEL market is at the

Table 1 Parameter estimates and goodness-of-fit statistics for Gaussian, hyperbolic, NIG and α-stable distributions fitted to the deseasonalized (with respect to the weekly period and annual seasonality) price changes: $X_t - X_{t-1}$, for $t = 2, \ldots, 1827$

Distribution	Parameters				Test values	
	α	σ, δ	β	μ	AD	K
EEX						
Gaussian		10.1818		0.0017	+INF	7.9005
Hyperbolic	0.2126	0.1414	−0.0024	0.1095	+INF	2.0055
NIG	0.0556	3.4255	−0.0011	0.0691	2.3600	1.2190
α-stable	1.5025	2.9612	−0.1572	−0.3880	**0.5185**	**0.6265**
OMEL						
Gaussian		5.3788		−0.0007	10.8666	2.4719
Hyperbolic	0.2927	2.2668	−0.0174	0.4888	**0.5000**	**0.6469**
NIG	0.1907	5.4420	−0.0131	0.3734	0.5534	0.6913
α-stable	1.7748	3.1411	−0.2554	−0.1513	1.2293	1.1421
PJM						
Gaussian		8.1489		−0.0134	+INF	4.8914
Hyperbolic	0.2099	0.0001	−0.0078	0.3418	+INF	0.9910
NIG	0.0781	4.1792	−0.0058	0.2996	**0.3981**	**0.5162**
α-stable	1.5584	3.3174	−0.1148	−0.0891	0.7273	0.7518
NEPOOL						
Gaussian		14.3490		−0.0100	+INF	9.6336
Hyperbolic	0.1898	0.0003	−0.0060	0.3215	+INF	2.2445
NIG	0.0335	3.2320	−0.0023	0.2125	2.2277	0.9246
α-stable	1.4202	2.9764	−0.1215	−0.2217	**0.5624**	**0.7343**

'+INF' denotes a very large number (infinity in computer arithmetic). The best fits for each market, in terms of the lowest statistics, are emphasized in bold. Compare with Fig. 2

opposite end—it has the lowest price volatility and tails which taper off much faster. Its price change distribution is best approximated by a hyperbolic law, with the NIG fit being only slightly worse. The PJM market is somewhere in between—the tails are lighter than power-law, but significantly heavier than hyperbolic (note the +INF value for the AD statistic in Table 1). This ordering can be also observed in terms of the tail indexes. The EEX and NEPOOL α's are the lowest, followed by those of PJM. The tail indexes for the OMEL market are significantly larger than the rest.

The relatively good fit of α-stable and NIG laws to electricity prices has been already utilized in the context of time series modeling (Mugele et al. 2005; Weron and Misiorek 2007). This line of study may be further developed by considering periodic ARMA (Nowicka-Zagrajek and Wyłomańska 2008) or fractional ARIMA (Burnecki et al. 2008) time series with α-stable (or NIG) innovations, which can address both the seasonality and heavy tails prevailing in electricity prices.

As we have mentioned earlier, (log-)returns—i.e. first differences of log-prices—generally exhibit lighter tails than first differences of prices themselves; compare

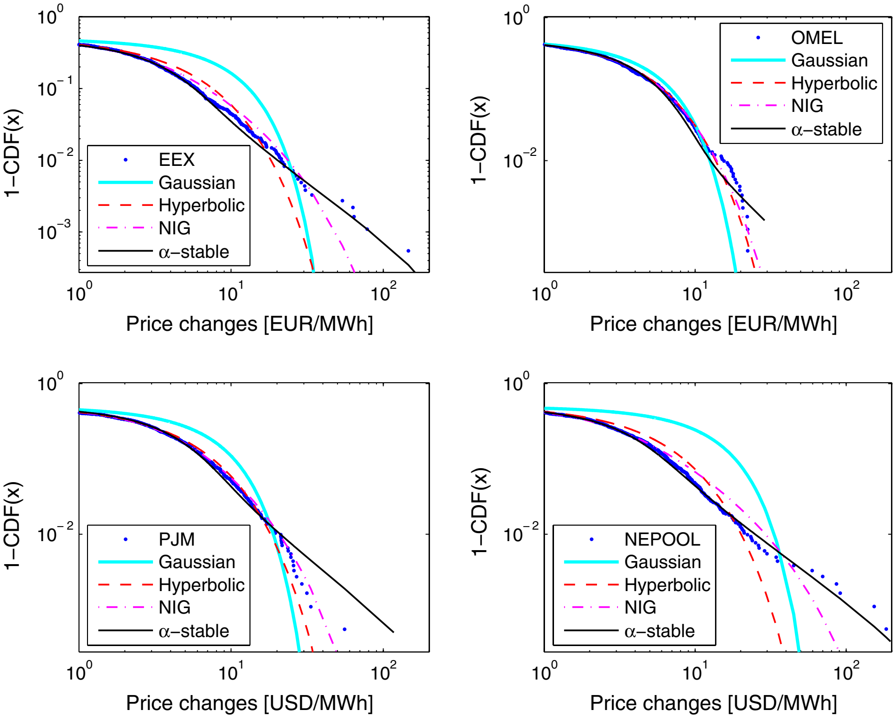

Fig. 2 Right tails of the distributions of deseasonalized price changes (first differences) and the fitted Gaussian, hyperbolic, NIG and α-stable distributions. EEX (*top left*) and NEPOOL (*bottom right*) prices have the heaviest, power-law tails, while OMEL (*top right*) the lightest, but still heavier than Gaussian. PJM (*bottom left*) prices are characterized by 'semi-heavy' tails. Compare with Table 1

Tables 1 and 2. Now, the α-stable law yields the best fit only for EEX returns; it even overestimates the most extreme observations, see Fig. 3. NEPOOL returns are best modeled by the NIG distribution and for the remaining two markets the hyperbolic and NIG laws lead to comparable fits, with the latter performing just a bit better.

Why does this happen? Well, because the logarithmic transformation dampens the spikes and hence extreme returns. It also makes the distribution of electricity price returns more symmetric as the low prices become even lower. This is confirmed by the values of sample skewness (i.e. the third central moment, divided by the cube of standard deviation) for price differences: 3.9915, −0.0848, −0.7859, −4.9609, compared with those for returns: 0.6530, −0.0295, −0.5183, −0.6864, for EEX, OMEL, PJM and NEPOOL, respectively.

This empirical exercise clearly shows that different models should be used for price changes and different for log-price changes. As we will see in the next section, contrary to the common belief that electricity price models 'should be built on log-prices', in some cases modeling prices themselves may be more beneficial and methodologically sound.

Table 2 Parameter estimates and goodness-of-fit statistics for Gaussian, hyperbolic, NIG and α-stable distributions fitted to the deseasonalized returns (or log-price changes): $Y_t - Y_{t-1}$, $Y_t = \log(X_t)$, for $t = 2, \ldots, 1827$

Distribution	Parameters				Test values	
	α	σ, δ	β	μ	AD	K
EEX						
Gaussian		21.3989		0.0053	+INF	5.0547
Hyperbolic	0.0769	1.3352	−0.0011	0.3958	3.7089	1.4347
NIG	0.0271	11.1972	−0.0003	0.1338	1.1503	0.9255
α-stable	1.5185	8.8920	−0.0135	−0.2353	**0.4103**	**0.6358**
OMEL						
Gaussian		15.3814		−0.0020	15.3015	2.7183
Hyperbolic	0.0989	4.4556	−0.0061	1.3866	0.5226	0.7813
NIG	0.0578	13.6146	−0.0044	1.0488	**0.4904**	**0.6706**
α-stable	1.7010	8.5069	−0.2009	−0.3675	1.2750	0.9243
PJM						
Gaussian		15.9973		−0.0431	17.0345	2.6161
Hyperbolic	0.0936	3.3472	−0.0045	1.0533	0.3334	**0.5096**
NIG	0.0546	13.6774	−0.0044	1.0550	**0.2760**	0.5745
α-stable	1.7020	8.7366	−0.0704	0.1688	1.0258	0.8235
NEPOOL						
Gaussian		19.4973		−0.0330	41.6959	3.9451
Hyperbolic	0.0798	0.0001	−0.0031	0.9341	1.2395	0.7431
NIG	0.0336	11.6958	−0.0025	0.8556	**0.2895**	**0.4842**
α-stable	1.5808	8.9218	−0.1613	−0.4092	0.6964	0.7025

Compare with Table 1

4 Markov regime-switching models

Price process models lie at the heart of derivatives pricing and risk management systems. If the price process chosen is inappropriate to capture the main characteristics of electricity prices, the results from the model are likely to be unreliable. On the other hand, if the model is too complex the computational burden will prevent its on-line use in trading departments. In a way, the MRS models offer the best of the two worlds; they are a trade-off between model parsimony and adequacy to capture the unique characteristics of power prices.

The underlying idea behind the MRS scheme is to model the observed stochastic behavior of a specific time series by two (or more) separate phases or regimes with different underlying processes. In other words, the parameters of the underlying process may change for a certain period of time and then fall back to their original structure. Thus, regime-switching models divide the time series into different phases that are called regimes. For each regime one can define separate and independent underlying price processes. The switching mechanism between the states is assumed to be governed by an unobserved random variable.

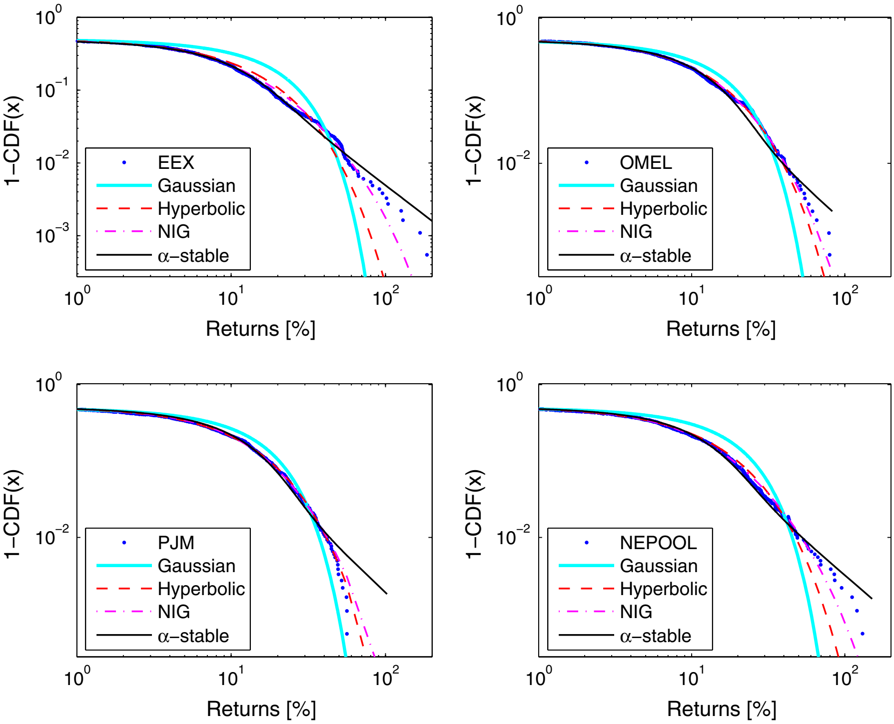

Fig. 3 Right tails of the distributions of deseasonalized log-price changes (i.e. returns) and the fitted Gaussian, hyperbolic, NIG and α-stable distributions. Compare with Table 2. Note also that the returns in all markets: EEX (*top left*), OMEL (*top right*), PJM (*bottom left*) and NEPOOL (*bottom right*) have lighter tails than the corresponding price changes in Fig. 2

For example, the spot price can be assumed to display either low or very high prices at each point in time, depending on the regime $R_t = 1$ or $R_t = 2$. Consequently, we have a probability law that governs the transition from one state to another. The price processes being linked to each of the two regimes are assumed to be independent from each other. The transition matrix $\mathbf{Q}$ contains the probabilities q_{ij} of switching from regime i at time t to regime j at time $t+1$, for $i, j = \{1, 2\}$:

$$\mathbf{Q} = (q_{ij}) = \begin{pmatrix} q_{11} & q_{12} \\ q_{21} & q_{22} \end{pmatrix} = \begin{pmatrix} q_{11} & 1 - q_{11} \\ 1 - q_{22} & q_{22} \end{pmatrix}. \tag{7}$$

Because of the Markov property the current state R_t at time t of a Markov chain depends on the past only through the most recent value R_{t-1}. Consequently the probability of being in state j at time $t+m$ starting from state i at time t is given by

$$P(R_{t+m} = j \mid R_t = i) = (\mathbf{Q}')^m \cdot e_i,$$

where $\mathbf{Q}'$ denotes the transpose of $\mathbf{Q}$ and e_i denotes the ith column of the 2×2 identity matrix.

Calibration of MRS models is not straightforward since the regime is only latent and hence not directly observable. Hamilton (1990) introduced an application of the expectation–maximization (EM) algorithm of Dempster et al. (1977) where the whole set of parameters θ is estimated by an iterative two-step procedure. In the first step the conditional probabilities $P(R_t = j|P_1, \ldots, P_T; \theta)$ for the process being in regime j at time t are calculated based on starting values $\hat{\theta}^{(0)}$ for the parameter vector θ of the underlying stochastic processes. These probabilities are referred to as smoothed inferences. Then in the second step new and more exact ML estimates $\hat{\theta}$ for all model parameters are calculated by using the smoothed inferences from step one. With each new vector $\hat{\theta}^{(n)}$ the next cycle of the algorithm is started in order to reevaluate the smoothed inferences. Every iteration the EM algorithm generates new estimates $\hat{\theta}^{(n+1)}$ as well as new estimates for the smoothed inferences. Each iteration cycle increases the log-likelihood function and the limit of this sequence of estimates reaches a (local) maximum of the log-likelihood function.

To our best knowledge, Ethier and Mount (1998) were the first to apply MRS models to electricity prices. They proposed a two state specification in which both regimes were governed by AR(1) price processes and concluded that there was strong empirical support for the existence of different means and variances in the two regimes. Huisman and Mahieu (2003) proposed a regime-switching model with three possible regimes in which the initial jump regime was immediately followed by the reversing regime and then moved back to the base regime. Consequently, their model did not allow for consecutive high prices (and hence did not offer any obvious advantage over jump-diffusion models). This restriction was efficiently relaxed by Huisman and de Jong (2003) who proposed a model with only two regimes—a stable, mean-reverting AR(1) regime and a spike regime—for the deseasonalized log-prices. The third regime was not needed to pull prices back to stable levels, because the prices were assumed to be independent from each other in the two regimes. They assumed that the dynamics of the spike regime could be modeled with a simple normal distribution whose mean and variance were higher than those of the mean-reverting base regime process. Bierbrauer et al. (2004) extended the model by allowing log-normal and Pareto distributed spike regimes to cope with the heavy-tailed nature of spike severities, while Bierbrauer et al. (2007) used a model with exponentially distributed spikes. De Jong (2006) proposed yet another modification of the basic two-regime model with autoregressive, Poisson driven spike regime dynamics and compared it to a number of MRS models.

What differentiates our study from the ones mentioned above is the fact that we not only look at some statistics of goodness-of-fit (like log-likelihoods, spike frequencies and severities) but also identify which prices or log-prices are classified as being in the spike regime. As we will see later in this section, in some models surprisingly many ‘non-spiky’ prices are wrongly classified.

We start with modeling deseasonalized log-prices, i.e. $Y_t = \log(X_t)$, following the common belief that electricity price models ‘should be built on log-prices’. We consider a two-regime specification with the base regime dynamics given by a mean reverting Ornstein–Uhlenbeck process:

$$dY_{t,1} = (c_1 - \beta Y_{t,1})dt + \sigma_1 dW_t, \tag{8}$$

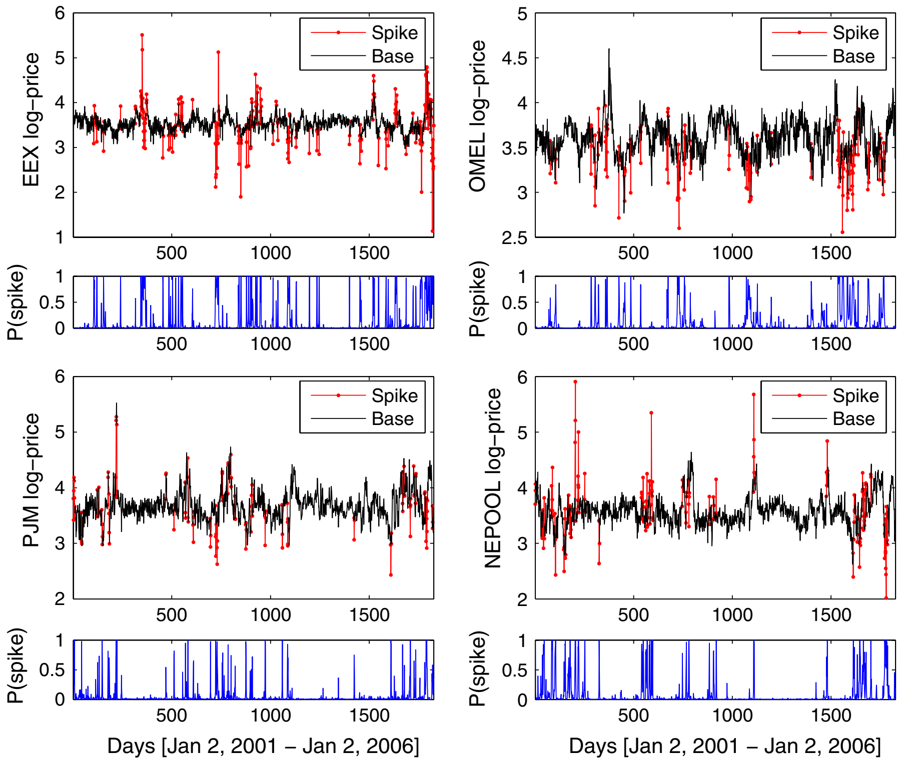

Fig. 4 Calibration results of the MRS model with log-normal spike regime to the deseasonalized log-prices from the EEX (*top left*), OMEL (*top right*), PJM (*bottom left*) and NEPOOL (*bottom right*) power markets. The corresponding lower panels display the probability $P(R = 2)$ of being in the spike regime. The log-prices classified as spikes, i.e. with $P(R = 2) > 0.5$, are additionally denoted by dots. Compare with Table 3

where W_t is Brownian motion. Note that (8) can be discretized as an autoregressive time series of order one, i.e. AR(1). The dynamics in the spike regime follow one of two qualitatively different distributions, namely log-normal:

$$\log(Y_{t,2}) \sim \mathrm{N}(c_2, \sigma_2^2), \tag{9}$$

or Pareto:

$$Y_{t,2} \sim F_{\text{Pareto}}(c_2, \sigma_2^2) = 1 - \left(\frac{\sigma_2^2}{x}\right)^{c_2}. \tag{10}$$

A specification with Gaussian spikes is left out from the analysis, because it yields similar fits to the log-normal model, at the same time being less stable (with respect to parameter estimates).

The estimation results for all four datasets are summarized in Tables 3 and 4. As expected, in both models the probability of remaining in the base regime is very high: $q_{11} \approx 0.96$ for the log-normal model and $q_{11} \approx 0.98$ for the Pareto specification. The probability of remaining in the spike regime is much lower, but still relatively

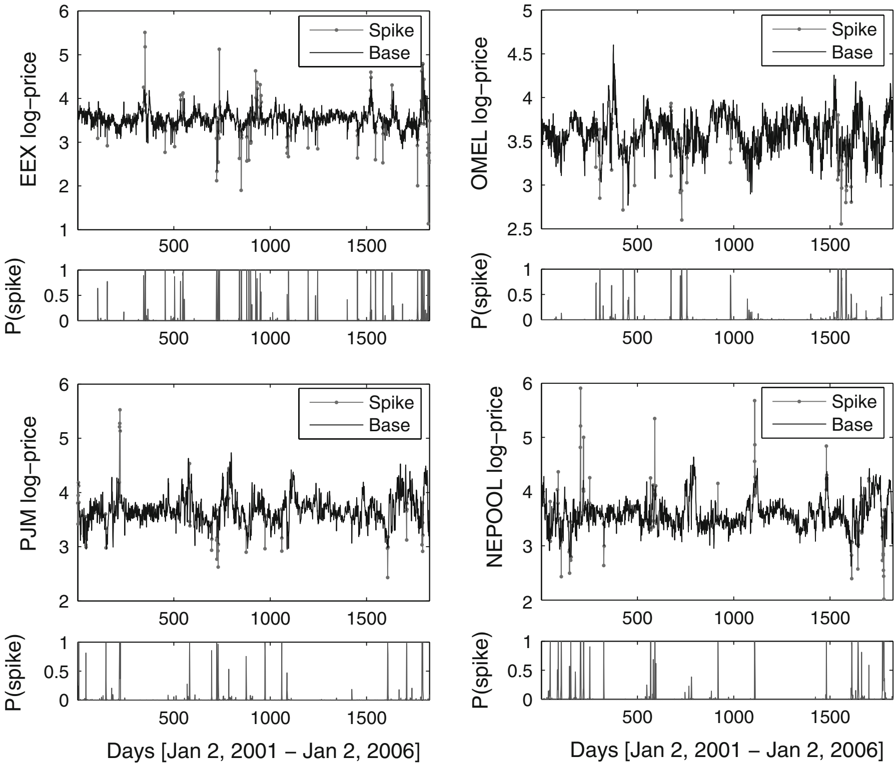

Fig. 5 Calibration results of the MRS model with Pareto spike regime to the deseasonalized log-prices from the EEX (*top left*), OMEL (*top right*), PJM (*bottom left*) and NEPOOL (*bottom right*) power markets. The corresponding *lower panels* display the probability $P(R = 2)$ of being in the spike regime. The log-prices classified as spikes, i.e. with $P(R = 2) > 0.5$, are additionally denoted by dots. Compare with Table 4

high: roughly there is a 50% chance that the log-price will stay in the spike regime for the next day. Unlike jump-diffusions, regime-switching models allow for consecutive spikes in a very natural way.

Considering the unconditional probabilities $P(R = i)$, the probability of being in the spike regime ($i = 2$) for the log-normal model is much higher than for the model with Pareto spikes: 0.5–0.12 versus 0.02–0.035. Using a heavy-tailed distribution, like the Pareto law, gives lower probabilities for being and remaining in the spike regime and a clearly higher variance. In fact, for EEX and NEPOOL log-prices the fitted Pareto spike distribution is so heavy-tailed (tail index $c_i < 2$) that the variance does not exist. This is not a problem as in all markets prices are capped. If the same price caps are imposed on the models, the model generated prices will exhibit finite variance as well.

In Figs. 4 and 5 the deseasonalized log-prices Y_t and the unconditional probabilities of being in the spike regime $P(R = 2)$ for all four markets are displayed. The log-prices classified as spikes, i.e. with $P(R = 2) > 0.5$, are additionally denoted by dots. Surprisingly many very low prices are classified as spikes. What is even more

Table 3 Calibration results of the MRS model with log-normal spike regime to the deseasonalized log-prices from the EEX, OMEL, PJM and NEPOOL power markets

Regime	Parameters			Statistics			
	β_i	c_i	σ_i^2	$E(Y_{t,i})$	$\mathrm{Var}(Y_{t,i})$	q_{ii}	$P(R=i)$
EEX							
Base	0.2941	1.0358	0.0118	3.5220	0.0235	0.9546	0.8844
Spike		1.2296	0.0317	3.4745	0.3882	0.6530	0.1156
OMEL							
Base	0.1895	0.6862	0.0132	3.6205	0.0386	0.9720	0.9041
Spike		1.2040	0.0061	3.3437	0.0682	0.7360	0.0959
PJM							
Base	0.1301	0.4821	0.0157	3.7049	0.0645	0.9679	0.9465
Spike		1.2562	0.0153	3.5391	0.1937	0.4322	0.0535
NEPOOL							
Base	0.1353	0.4852	0.0157	3.5863	0.0624	0.9689	0.9179
Spike		1.2711	0.0253	3.6101	0.3340	0.6525	0.0821

Compare with Fig. 4

disturbing, some of the spikes are not extreme enough to be classified as such. The extreme example is the least volatile OMEL market where practically no spikes (in the sense: upward jumps) are identified. This undesired behavior can be also observed in Tables 3 and 4. If we define the 'expected spike size' as the difference between the expected values in the spike and base regime, i.e. $E(Y_{t,2}) - E(Y_{t,1})$, we will see that it can be negative! Similar results were reported by De Jong (2006) for models with Gaussian spike regime, but were not considered as evidence for wrong model specification. Finally, we note that the model with Pareto spikes performed a little better in this respect than the log-normal one, but only for the spiky EEX and NEPOOL markets.

It turns out that the calibration scheme generally assigns all extreme prices to the spike regime, no matter whether they truly are spikes or only sudden drops. But these 'sudden drops' are actually not so extreme. They appear such only because of the logarithmic transformation which enhances low prices, at the same time dumping high prices. More importantly, these artificial sudden drops are not that interesting from the point of view of price modeling and derivatives valuation, because in absolute terms the price changes are small and the related price risks are negligible. Hence, when calibrating models to log-prices we needlessly try to match some of the insignificant characteristics.

Having this in mind, we fitted both MRS models to deseasonalized prices X_t. The results for the model with Pareto spikes are presented in Table 5 and Fig. 6. This time the calibration of the log-normal model failed to converge to reasonable values. Apparently the spikes were too extreme. Comparing with the results for the Pareto model for log-prices (Table 4; Fig. 5), we can observe that now practically all the spikes in all four markets are identified correctly. Moreover, the number of 'sudden

Table 4 Calibration results of the MRS model with Pareto spike regime to the deseasonalized log-prices from the EEX, OMEL, PJM and NEPOOL power markets

Regime	Parameters			Statistics			
	β_i	c_i	σ_i^2	$E(Y_{t,i})$	$\mathrm{Var}(Y_{t,i})$	q_{ii}	$P(R=i)$
EEX							
Base	0.3371	1.1883	0.0177	3.5255	0.0315	0.9769	0.9663
Spike		0.9463	1.1378	+INF	+INF	0.3388	0.0337
OMEL							
Base	0.2127	0.7655	0.0160	3.5984	0.0421	0.9891	0.9747
Spike		4.2728	2.5541	3.3345	1.1450	0.5821	0.0253
PJM							
Base	0.1582	0.5839	0.0180	3.6904	0.0618	0.9891	0.9797
Spike		2.9025	2.4292	3.7061	5.2432	0.4752	0.0203
NEPOOL							
Base	0.1701	0.6098	0.0197	3.5852	0.0634	0.9856	0.9699
Spike		1.7811	2.0149	4.5945	+INF	0.5378	0.0301

Compare with Fig. 5

Table 5 Calibration results of the MRS model with Pareto spike regime to the deseasonalized prices from the EEX, OMEL, PJM and NEPOOL power markets

Regime	Parameters			Statistics			
	β_i	c_i	σ_i^2	$E(X_{t,i})$	$\mathrm{Var}(X_{t,i})$	q_{ii}	$P(R=i)$
EEX							
Base	0.3393	11.4438	17.0811	33.7241	30.3112	0.9714	0.9500
Spike		0.3887	3.1200	+INF	+INF	0.4563	0.0500
OMEL							
Base	0.2019	7.4034	19.7234	36.6744	54.3369	0.9874	0.9730
Spike		0.9577	12.8600	+INF	+INF	0.5463	0.0270
PJM							
Base	0.1618	6.4045	26.0606	39.5883	87.6327	0.9802	0.9587
Spike		0.6195	11.3500	+INF	+INF	0.5399	0.0413
NEPOOL							
Base	0.1403	5.1072	23.9272	36.4026	91.7063	0.9777	0.9550
Spike		0.5152	7.5000	+INF	+INF	0.5272	0.0450

Compare with Fig. 6

drops' classified as spikes is much lower and, at the same time, the unconditional probabilities of being in the spike regime $P(R=2)$ are 50–100% higher (except for OMEL, but there are not too many spikes in this market anyway), which suggests that the calibration scheme does a better job of identifying the spikes. Finally, the tail indexes of the spike regime are lower indicating heavier tails.

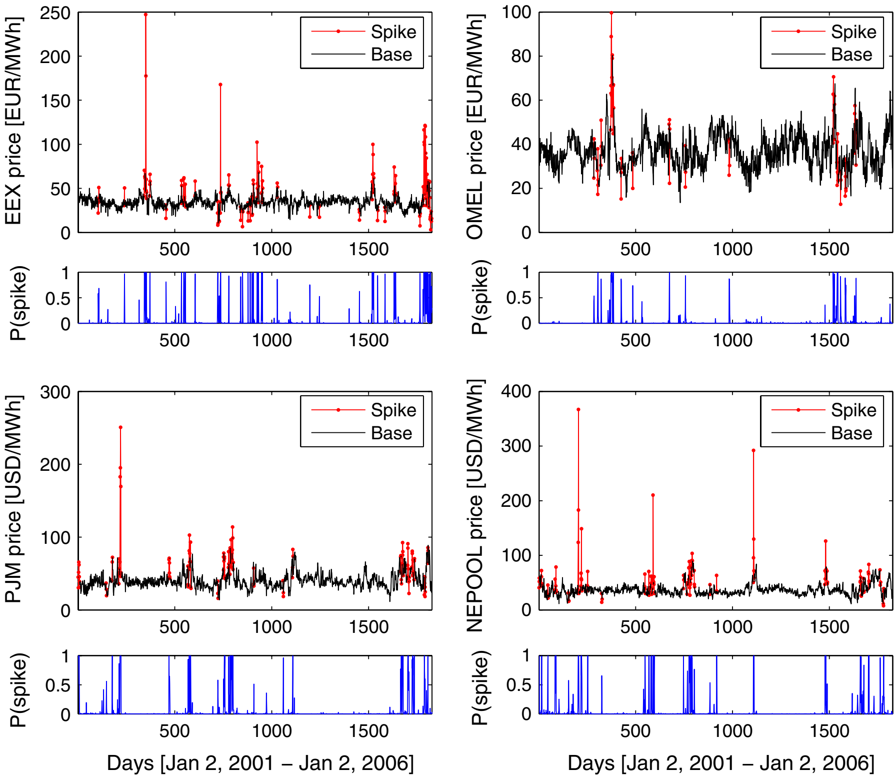

Fig. 6 Calibration results of the MRS model with Pareto spike regime to the deseasonalized prices from the EEX (*top left*), OMEL (*top right*), PJM (*bottom left*) and NEPOOL (*bottom right*) power markets. The corresponding *lower panels* display the probability $P(R = 2)$ of being in the spike regime. The prices classified as spikes, i.e. with $P(R = 2) > 0.5$, are additionally denoted by *dots*. Compare with Table 5

5 Conclusions

In this paper we have focused on two stylized facts of electricity prices: extreme volatility and price spikes, which lead to heavy-tailed distributions of price changes. The results reported in Sect. 3 show that electricity spot prices and log-prices are heavy or semi-heavy tailed. The tail behavior differs between markets: EEX and NEPOOL exhibit the most extreme behavior (power-law tails), while OMEL the least (hyperbolic tails). We attribute this fact to the availability or lack of cheap hydro generation which can be used to 'hedge' market imbalance in a matter of minutes. We also note that (log-)returns (i.e. first differences of log-prices) generally exhibit lighter tails than first differences of electricity prices themselves.

In view of this, in Sect. 4 we calibrate MRS models with semi-heavy (log-normal) and heavy-tailed (Pareto) components, both to deseasonalized prices and log-prices. Contrary to the common belief that electricity price models 'should be built on log-prices', we find evidence that modeling the prices themselves is more beneficial and methodologically sound, at least in case of MRS models. It turns out that for log-price models the calibration scheme generally assigns all extreme prices to the spike regime, no matter whether they truly are spikes or only artificial sudden drops (i.e. due to taking the logarithm of small values). This is not the case with the Pareto model calibrated to prices—now practically all the spikes in all four markets are identified correctly.

References

Benth FE, Benth JS, Koekebakker S (2008) Stochastic modeling of electricity and related markets. World Scientific, Singapore

Benz E, Trück S (2006) CO_2 emission allowances trading in Europe—specifying a new class of assets. Probl Perspect Manag 4(3):30–40

Bierbrauer M, Menn C, Rachev ST, Trück S (2007) Spot and derivative pricing in the EEX power market. J Bank Financ 31:3462–3485

Bierbrauer M, Trück S, Weron R (2004) Modeling electricity prices with regime switching models. Lect Notes Comput Sci 3039:859–867

Bottazzi G, Sapio S, Secchi A (2005) Some statistical investigations on the nature and dynamics of electricity prices. Physica A 355:54–61

Bunn DW (ed) (2004) Modelling prices in competitive electricity markets. Wiley, Chichester

Burnecki K, Klafter J, Magdziarz M, Weron A (2008) From solar flare time series to fractional dynamics. Physica A 387:1077–1087

Byström HNE (2005) Extreme value theory and extremely large electricity price changes. Int Rev Econ Financ 14:41–55

Chan KF, Gray P (2006) Using extreme value theory to measure value-at-risk for daily electricity spot prices. Int J Forecast 22:283–300

Čižek P, Härdle W, Weron R (eds) (2005) Statistical tools for finance and insurance. Springer, Berlin

De Jong C (2006) The nature of power spikes: a regime-switch approach. Stud Nonlinear Dyn Econom 10(3) (article 3)

Dempster A, Laird N, Rubin DB (1977) Maximum likelihood from incomplete data via the EM algorithm. J R Stat Soc 39:1–38

Deng S-J, Jiang W (2005) Levy process-driven mean-reverting electricity price model: the marginal distribution analysis. Decis Support Syst 40:483–494

Eberlein E, Stahl G (2003) Both sides of the fence: a statistical and regulatory view of electricity risk. Energy Power Risk Manag 8(6):34–38

Ethier R, Mount T (1998) Estimating the volatility of spot prices in restructured electricity markets and the implications for option values. PSerc Working Paper, pp 98–31

Hamilton J (1990) Analysis of time series subject to changes in regime. J Econom 45:39–70

Huisman R, de Jong C (2003) Option pricing for power prices with spikes. Energy Power Risk Manag 7.11: 12–16

Huisman R, Mahieu R (2003) Regime jumps in electricity prices. Energy Econ 25:425–434

Kaminski V (2004) Managing energy price risk: the new challenges and solutions, 3rd edn. Risk Books, London

Khindanova I, Atakhanova Z (2002) Stable modeling in energy risk management. Math Methods Oper Res 55:225–245

Mugele C, Rachev ST, Trück S (2005) Stable modeling of different European power markets. Invest Manag Financ Innov 2(3):65–85

Nowicka-Zagrajek J, Wyłomańska A (2008) Measures of dependence for stable AR(1) models with time-varying coefficients. Stoch Models 24(1):58–70

Paolella MS, Taschini L (2008) An econometric analysis of emission allowance prices. J Bank Financ (forthcoming). doi:10.1016/j.jbankfin.2007.09.024

Park H, Mjelde JW, Bessler DA (2006) Price dynamics among U.S. electricity spot markets. Energy Econ 28(1):81–101

Rachev ST, Mittnik S (2000) Stable paretian models in finance. Wiley, London

Rachev ST, Trück S, Weron R (2004) Risk management in the power markets (part III): advanced spot price models and VaR. RISKNEWS 05/2004:67–71 (in German)

Simonsen I (2005) Volatility of power markets. Physica A 355:10–20

Trück S, Weron R, Wolff R (2007) Outlier treatment and robust approaches for modeling electricity spot prices. In: Proceedings of the 56th session of the ISI, invited paper meeting IPM71, Lisbon, Portugal, 22–29 Aug 2007. Available at MPRA: http://mpra.ub.unimuenchen.de/4711/

Weron R (2004) Computationally intensive value at risk calculations. In: Gentle JE, Härdle W, Mori Y (eds) Handbook of computational statistics. Springer, Berlin pp 911–950

Weron R (2006) Modeling and forecasting electricity loads and prices: a statistical approach. Wiley, Chichester

Weron R (2008) Market price of risk implied by Asian-style electricity options and futures. Energy Econ 30:1098–1115

Weron R, Misiorek A (2007) Heavy tails and electricity prices: do time series models with non-Gaussian noise forecast better than their Gaussian counterparts? Prace Naukowe Akademii Ekonomicznej we Wrocławiu 1076:472—480. Available at MPRA: http://mpra.ub.unimuenchen.de/2292/