Regime-switching models for electricity spot prices: Introducing heteroskedastic base regime dynamics and shifted spike distributions

Janczura, Joanna and Weron, Rafal

Hugo Steinhaus Center, Wroclaw University of Technology

April 2009

Regime-switching models for electricity spot prices: Introducing heteroskedastic base regime dynamics and shifted spike distributions

Joanna Janczura
Institute of Mathematics and Computer Science
Wrocław University of Technology
50-370 Wrocław, Poland
Email: joanna.janczura@pwr.wroc.pl

Rafał Weron
Hugo Steinhaus Center
Wrocław University of Technology
50-370 Wrocław, Poland
Email: rafal.weron@pwr.wroc.pl

***Abstract*—We calibrate Markov regime-switching (MRS) models to mean daily spot prices from the EEX market. Our empirical study shows that (i) models with shifted spike regime distributions lead to more realistic models of electricity spot prices and that (ii) introducing heteroskedasticity in the base regime leads to better spike identification and goodness-of-fit than in MRS models with the standard mean-reverting, constant volatility dynamics.**

I. Introduction

Electricity is a very unique commodity. In addition to strong seasonality on the annual, weekly and daily level, spot electricity prices exhibit mean reversion, very high volatility and abrupt, short-lived and generally unanticipated extreme price changes known as spikes or jumps [3], [15], [18]. The aim of this paper is to suggest parsimonious models for electricity spot price dynamics that can address the most pertinent characteristics and, hence, be used for derivatives pricing. To this end, we test a range of Markov regime-switching (MRS) models, which by construction should be very well suited for the volatile electricity spot prices.

Motivated by recent findings [19] we focus on MRS models for the prices themselves; not the log-prices as in most other studies. Further, we introduce two novel features in the context of MRS modeling of electricity spot prices: heteroscedasticity in the base regime and shifted spike regime distributions. The rationale for the former comes from the observation that price volatility generally increases with price level, since positive price shocks increase volatility more than negative shocks (so-called 'inverse leverage effect' [13]). The CIR square-root process [4] is tested as a heteroscedastic alternative to the standard mean-reverting dynamics. Shifted spike regime distributions, on the other hand, are required for the calibration procedure to correctly separate spikes from the 'normal' price behavior. As in [19] we use two spike distributions: semi heavy-tailed lognormal and heavy-tailed Pareto.

The paper is structured as follows. In Section II we present the datasets and explain the deseasonalization procedures. In Section III we introduce the MRS models for deseasonalized prices and in Section IV we evaluate their goodness-of-fit. Finally, in Section V we summarize the results.

II. Data preprocessing

Due to space limitations in this paper we concentrate only on the German EEX market. For results of the whole study covering various European and American power markets see [11]. We use mean daily (baseload) spot prices from the period January 1, 2001 – January 3, 2009. To see how the presented methods perform under different market conditions the sample is split into two subsamples of 1463 daily observations (209 weeks each): January 1, 2001 – January 2, 2005 and January 3, 2005 – January 3, 2009, see Figure 1. Note, that starting in late 2004 the spot prices exhibit an upward trend and higher volatility, largely due to a combination of higher fuel prices and the introduction of CO_2 emission costs.

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 [16]. 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.

As in [19] 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 [16], [18]). 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: $S_J + D_J + D_{J-1} + ... + D_1$, 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 $j = J-1, J-2, ...$, 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 blue lines in Figure

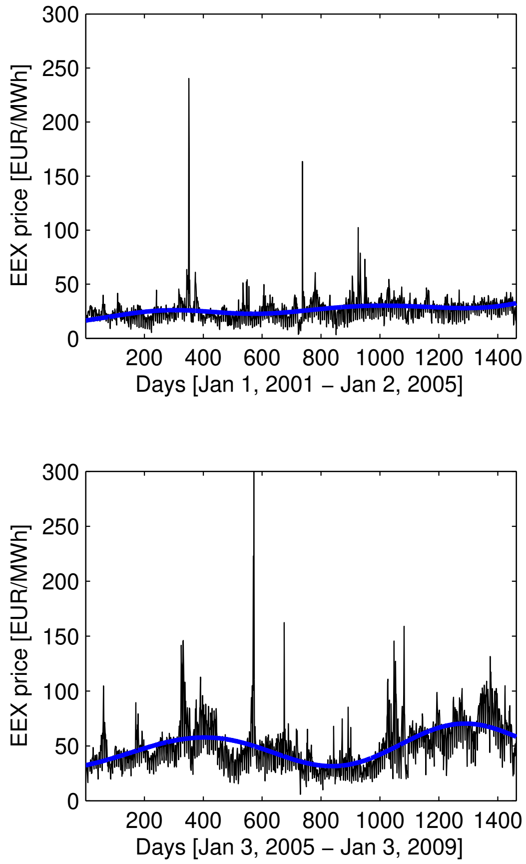

Fig. 1. Mean daily EEX spot prices and their long-term seasonal components (thick blue lines) for the two studied periods.

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 (see e.g. [18]) 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 Figure 3.

III. Markov regime-switching models

The underlying idea behind the Markov regime-switching (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. The switching mechanism between the states is Markovian and is assumed to be governed by an unobserved random variable. The underlying processes, though, do not have to be Markovian, but are (typically) independent from each other.

In this study, the spot price is assumed to display either normal (base regime $R_t = 1$) or high (spike regime $R_t = 2$) prices each day. 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 $i, j = \{1, 2\}$:

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

Because of the Markov property the current state R_t at time t 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 is the ith column of the 2×2 identity matrix.

To our best knowledge, the MRS models were first applied to electricity prices in [7]. A two state specification was proposed, in which in both regimes the log-prices were governed by autoregressive processes of order one, i.e. AR(1). Huisman and Mahieu [10] 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 (in fact of log-prices) and hence did not offer any obvious advantage over jump-diffusion models. This restriction was relaxed in [9], where a model with only two regimes – a stable, mean-reverting AR(1) regime and a spike regime – was proposed for deseasonalized log-prices. The third regime was not needed to pull prices back to stable levels, because the prices were independent from each other in the two regimes. The dynamics of the spike regime was modeled with a simple normal distribution whose mean and variance were higher than those of the mean-reverting base regime process. This simple yet versatile model was further extended by admitting lognormal, Pareto [2] and exponential [1] spike regime distributions, as well as, autoregressive Poisson driven spike regime dynamics [5].

The above mentioned models have two common features. Firstly, not the prices X_t themselves, but rather log-prices $Y_t = \log(X_t)$ are considered. Secondly, the base regime (and in some cases the spike regime as well) is driven by a mean-reverting diffusion process of the form:

$$dY_t = (\alpha - \beta Y_t)dt + \sigma dW_t, \quad (2)$$

where W_t is Brownian motion (i.e. a Wiener process), $\frac{\alpha}{\beta}$ is the long term mean reversion level, β is the speed of mean reversion and σ is the volatility. In the fixed income literature this popular process is known as the Vasicek [17] model, in mathematics as the (generalized) Ornstein-Uhlenbeck process, while in signal processing – when discretized – as an autoregressive time series of order one, i.e. AR(1).

In a recent paper, Weron [19] provided evidence that modeling the prices themselves is more beneficial and methodologically sound than the log-prices, at least in case of MRS models. 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 sudden drops. This property can be seen in Figure 2, where a MRS model with Vasicek base regime and lognormal spikes is fitted to

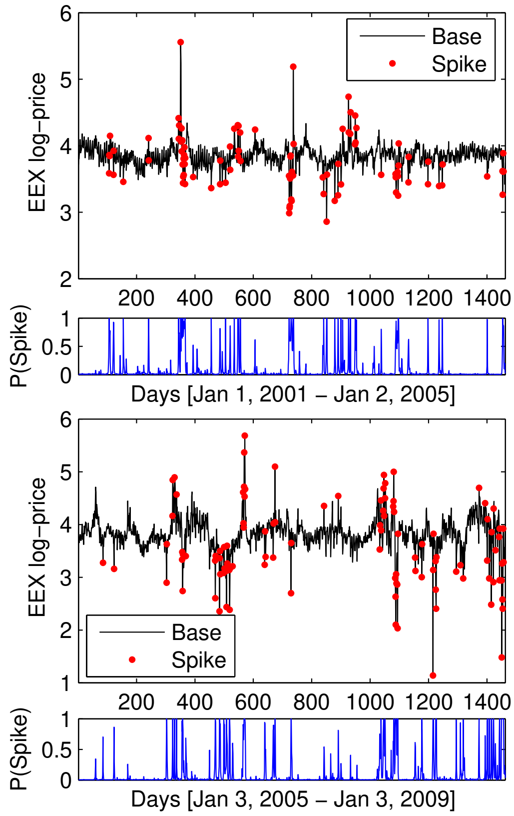

Fig. 2. Calibration results of the MRS model with Vasicek base regime and lognormal spikes fitted to the deseasonalized EEX log-prices in the two periods. The corresponding lower panels display the probability $P(R_t = 2)$ of being in the spike regime. The log-prices classified as spikes, i.e. with $P(R_t = 2) > 0.5$, are additionally denoted by dots.

deseasonalized EEX log-prices. 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, in this study we fit MRS models to deseasonalized prices X_t.

Furthermore, if we define the 'expected spike size' as the difference between the expected values in the spike and base regime, it turns out that in a significant number of cases it can be negative! For instance, such results were reported by de Jong [5] for models with a Gaussian spike regime, but were not considered as evidence for wrong model specification. As observed in [19], using price models instead of log-price models alleviates this unwanted feature to a large extent. Nevertheless, still some of the low prices are classified as being in the spike regime. To eliminate this problem, we introduce in this paper shifted spike regime distributions which assign zero probability to prices below a certain quantile (here: the median $m = \text{median}(X_t)$) of the dataset. We consider 'shifted lognormal':

$$\log(X_t - m) \sim \text{N}(\alpha, \sigma^2), \quad X_t > m, \tag{3}$$

and 'shifted Pareto' laws:

$$X_t \sim F_{\text{Pareto}}(\alpha, \lambda) = 1 - \left(\frac{\lambda}{x}\right)^{\alpha}, \quad X_t > \lambda \geq m. \tag{4}$$

Regarding the second common feature, we introduce an alternative base regime dynamics and compare it with the 'industry standard' Vasicek model; with Y_t replaced by X_t in eqn. (2). The competitor is the so-called square-root or CIR [4] process:

$$dX_t = (\alpha - \beta X_t)dt + \sigma\sqrt{X_t}dW_t. \tag{5}$$

Note, that in this model the volatility is dependent on the current price level X_t, i.e the higher the price level the larger are the price changes. Compared to Vasicek dynamics, we can expect that in the CIR model the less extreme price changes will be classified as 'normal' and not spiky.

Calibration of MRS models is not straightforward since the regime is only latent and hence not directly observable. Hamilton [8] introduced an application of the Expectation-Maximization (EM) algorithm [6] where the whole set of parameters θ is estimated by an iterative two-step procedure. The algorithm was later refined by Kim [12]. In the first step the conditional probabilities $P(R_t = j|X_1, ..., X_T; \theta)$ for the process being in regime j at time t, so-called 'smoothed inferences', are calculated based on starting values $\hat{\theta}^{(0)}$ for the parameter vector θ of the underlying stochastic processes. Then, in the second step, new and more exact maximum likelihood (ML) estimates $\hat{\theta}$ for all model parameters are calculated. Compared to standard ML estimation, where for a given probability density function f the log-likelihood function $\sum_{t=1}^{n} \log f(X_t, \theta)$ is maximized, here each component of this sum has to be weighted with the corresponding smoothed inference, since each observation X_t belongs to the jth regime exactly with probability $P(R_t = j|X_1, ..., X_T; \theta)$.

The parameters of the 'shifted lognormal' regime are obtained as the ML estimates of the standard lognormal distribution fitted to (deseasonalized) prices with subtracted median and weighted by the smoothed inferences. In the 'shifted Pareto' case the situation is more complicated. Recall, that the likelihood function of the standard Pareto distribution is monotonically increasing with λ, hence, reasonable ML estimates exist only for the shape parameter α. The scale parameter is thus typically set to $\hat{\lambda} = \min(X_t)$, the rationale coming from the fact that the distribution is defined only for $X_t > \lambda$. In order to ensure that prices lower than the median are not identified as spikes we let $\hat{\lambda}$ be the maximum of the median price and the smallest price satisfying the condition $P(R_t = j|X_1, ..., X_T; \hat{\theta}) > 0$, i.e. the smallest potential spike. Finally, the base regime parameters are estimated via ML with each price being weighted by the smoothed inferences. In every iteration the EM algorithm generates new estimates $\hat{\theta}^{(n+1)}$ as well as new estimates for the smoothed inferences.

TABLE I
CALIBRATION RESULTS, DESCRIPTIVE STATISTICS AND p-VALUES OF THE K-S GOODNESS-OF-FIT TEST

Regime		Parameters			Probabilities		Model statistics		K-S test p-values		K-S value
		β	α	σ^2, λ	q_{ii}	$P(R=i)$	Median	IDR	Regime	Model	Model
EEX (Jan. 1, 2001 - Jan. 2, 2005)											
Base	CIR	0.36	16.99	0.54	0.9911	0.9896	46.54	16.93	**0.015**	**0.015**	0.0407
Spike	Logn		3.44	0.69	0.1598	0.0104			**0.770**		
Base	Vasicek	0.36	16.62	22.99	0.9864	0.9579	46.63	16.81	0.002	0.007	0.0439
Spike	Logn		2.56	0.92	0.6906	0.0421			**0.338**		
Base	CIR	0.36	16.99	0.55	0.9913	0.9899	46.58	16.97	**0.015**	**0.016**	0.0404
Spike	Pareto		1.70	46.93	0.1492	0.0101			**0.073**		
Base	Vasicek	0.36	16.53	23.23	0.9863	0.9609	46.63	16.80	0.004	**0.013**	0.0415
Spike	Pareto		3.14	46.93	0.6635	0.0390			**0.026**		
EEX (Jan. 3, 2005 - Jan. 3, 2009)											
Base	CIR	0.30	13.81	1.65	0.9884	0.9824	45.19	31.69	<0.001	<0.001	0.0714
Spike	Logn		3.41	2.77	0.3519	0.0176			**0.038**		
Base	Vasicek	0.28	12.32	53.43	0.9751	0.8756	45.41	30.89	<0.001	<0.001	0.0890
Spike	Logn		2.79	1.05	0.8244	0.1244			**0.041**		
Base	CIR	0.30	13.82	1.65	0.9882	0.9821	45.20	31.59	<0.001	<0.001	0.0719
Spike	Pareto		1.38	45.73	0.3549	0.0179			**0.030**		
Base	Vasicek	0.26	11.29	53.97	0.9732	0.8786	45.40	30.83	<0.001	<0.001	0.0553
Spike	Pareto		2.67	45.73	0.8061	0.1214			0.002		

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.

IV. EMPIRICAL RESULTS

The deseasonalized prices X_t and the unconditional probabilities of being in the spike regime $P(R_t = 2)$ for both datasets are displayed in Figure 3. The estimation results are summarized in Table I. Additionally in this table we provide probabilities of staying in each regime q_{ii}, unconditional probabilities $P(R = i)$ of being in regime i, basic descriptive statistics and results of the Kolmogorov-Smirnov (K-S) goodness-of-fit tests.

As expected, in each model the probability of remaining in the base regime is very high: from 0.9732 in the Vasicek-Pareto model up to 0.9913 in the CIR-Pareto specification. The probability of remaining in the spike regime is much lower, but still relatively high. Obviously, unlike jump-diffusions, MRS models allow for consecutive spikes in a very natural way. Moreover, the probabilities that the price will remain in the spike regime for the next day is much higher in the Vasicek base regime models. Indeed, when fitting the Vasicek base regime models clusters of spikes can be observed, see Figure 3. Looking carefully at spike classification (Figure 4), we notice that only some of those prices are extreme enough to be regarded as spikes. This is not the case in the CIR base regime models, where – regardless of the spike distribution – spikes are identified correctly.

Considering unconditional probabilities, the probability of being in the spike regime $P(R_t = 2)$ for the CIR base regime models is lower than for the Vasicek base regime models, since there are fewer prices identified as spikes. CIR base regime dynamics leads to lower probabilities of being and remaining in the spike regime but also higher spike regime variance. In fact, the fitted Pareto spike distribution is so heavy-tailed (tail index $\alpha < 2$) that the variance does not exist. This is not a problem as the bids in the EEX market are capped [14]. If the same cap (3000 EUR/MWh) is imposed on the models, the model generated prices exhibit a finite variance as well.

In order to evaluate the goodness-of-fit, we report basic descriptive statistics and results of the K-S tests. The former include the median and the Inter-Decile Range (IDR), i.e. the difference between the ninth (90% quantile) and first (10% quantile) deciles. The quantile-based measures rather than the less robust to outliers moment-related statistics are used. The model statistics (Table I) should be compared with the corresponding values for the datasets: a median of 46.92 (45.73) and an IDR of 16.30 (30.97) in the first (second) period. The models with Vasicek base regime seem to have a slightly higher median and slightly smaller IDR than the CIR base regime models. Nevertheless, in terms of these statistics all models provide very good fits for both periods.

The K-S tests are more discriminatory. In the considered models, though, neither the prices themselves nor their differences or returns are independent and identically distributed (i.i.d.). Hence, the K-S tests cannot be applied directly to prices (or returns). Instead we have to use the following procedure. First, the data is split into two subsets: spikes (i.e. prices with probability $P(R_t = 2) > 0.5$) and the base regime. The discretization of equations (2) and (5) leads to

$$\varepsilon_t = (X_{t+1} - (1-\beta)X_t - \alpha)/\sigma, \tag{6}$$
$$\varepsilon_t = (X_{t+1} - (1-\beta)X_t - \alpha)/\sigma\sqrt{X_t}, \tag{7}$$

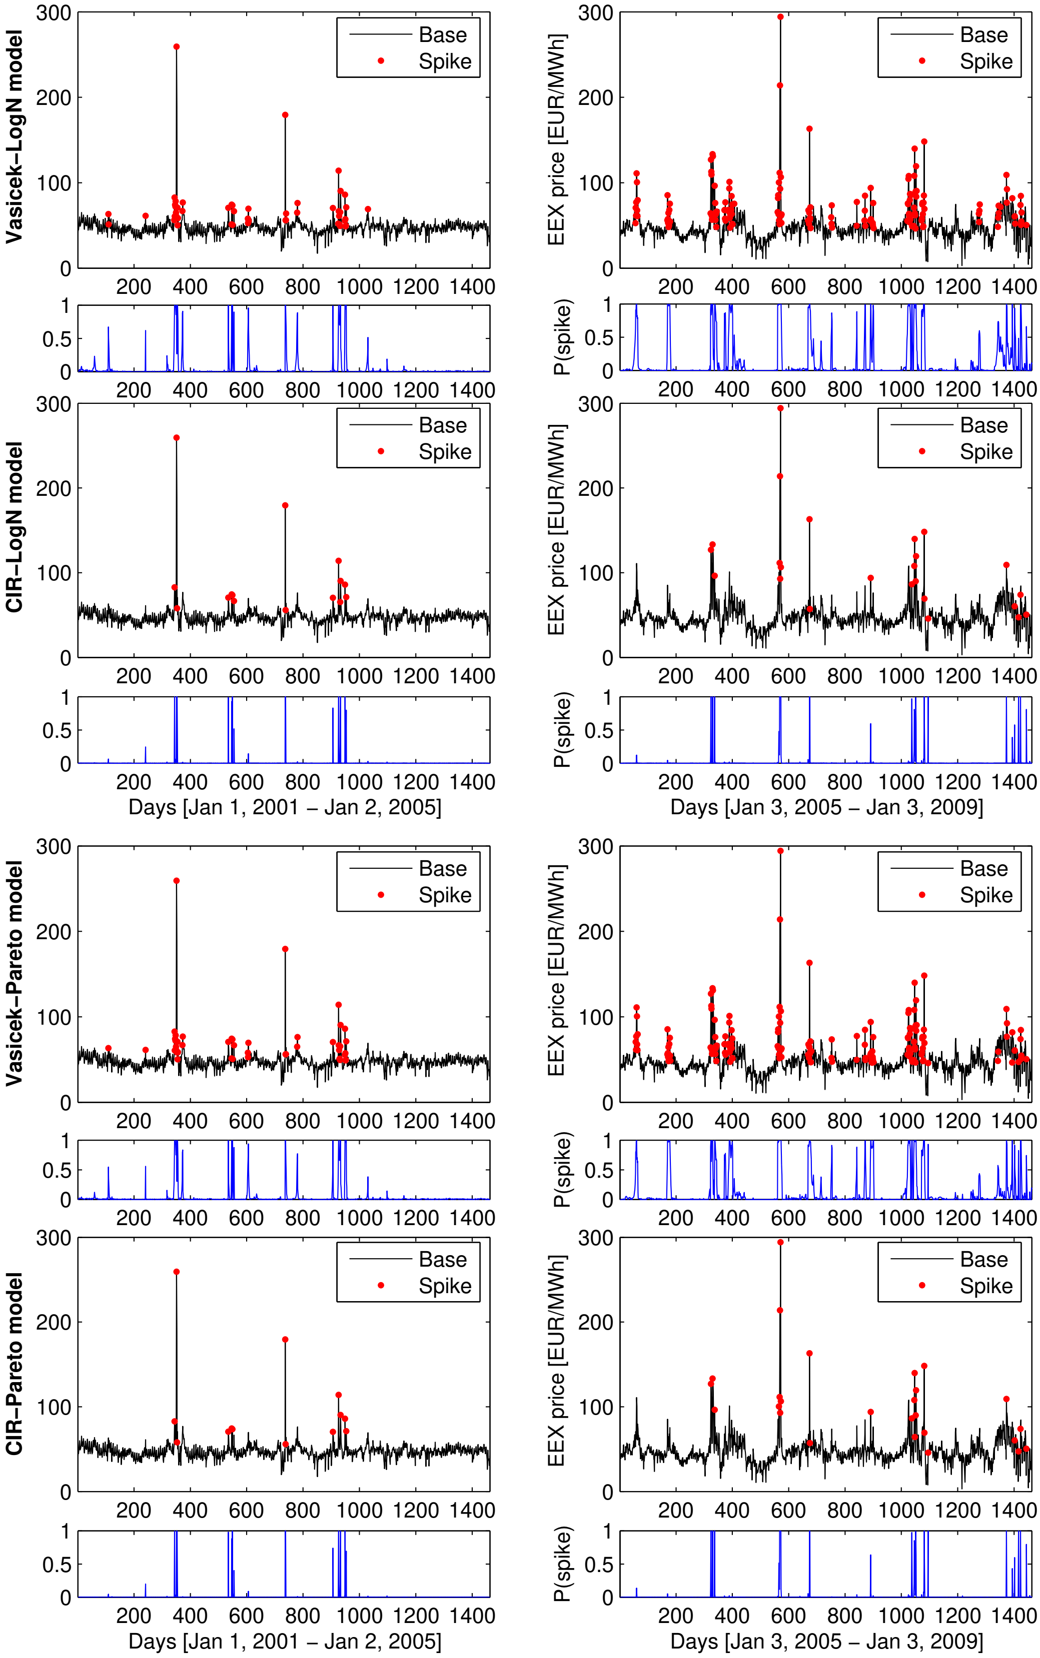

Fig. 3. Calibration results for the MRS models with shifted lognormal (*top 4 panels*) or Pareto (*bottom 4 panels*) spikes and Vasicek or CIR base regimes. The corresponding lower panels display the probability $P(R_t = 2)$ of being in the spike regime. The prices classified as spikes, i.e. with $P(R_t = 2) > 0.5$, are additionally denoted by dots.

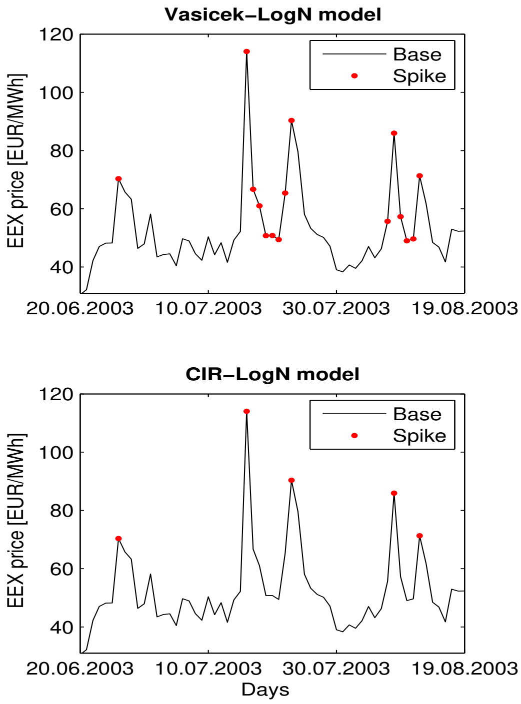

Fig. 4. Sample results for the MRS models with shifted lognormal spikes and Vasicek (*top*) or CIR (*bottom*) base regimes. Note, that only some of the prices classified as spikes in the Vasicek base regime model are extreme enough to be regarded as spikes. This is not the case in the CIR model.

for the Vasicek and CIR models, respectively. The ε_t's are now i.i.d. Gaussian random variables. Hence, we can apply one of the above transformations to the base regime data and obtain two i.i.d. samples: Gaussian and lognormal (or Pareto, depending on spike specification) distributed. Combining these two subsets yields a sample of independent variables with the distribution being a mixture of the lognormal (or Pareto) and Gaussian laws. The probability that a given price X_t comes from the spike distribution is equal to $P(R_t = 2)$, while that it comes from the Gaussian law is equal to $P(R_t = 1) = 1 - P(R_t = 2)$. We perform the K-S test for both subsets (base regime and spikes), as well as, for the whole sample.

The results of the K-S tests (Table I) indicate that good fits are obtained for prices in the first, less volatile period, since in all but one (i.e. the Vasicek-lognormal) model the p-values are larger than 0.01. Recall, that p-values larger than 0.01 indicate that we cannot reject the hypothesis about the chosen price model at the 1% significance level. Moreover, each spike distribution gives an acceptable result. The p-values obtained for the spike regime indicate that, although we cannot reject the hypothesis about the Pareto spike distribution, the lognormal distribution gives a better fit. In the second, more volatile period none of the models yields a satisfactory fit. Nevertheless, the spike distribution, which seems to be the most important feature from the risk analysis point of view, is well fitted in three out of four cases. Only Pareto spikes with the Vasicek base regime give a p-value lower than 0.01.

A comparison of the test results for the base regime provides statistical evidence that the CIR model is more suitable than the Vasicek model. In the first period, the p-values for the CIR base regime with both spike distributions are larger than 0.01, while the Vasicek p-values do not exceed 0.004. Moreover, for both periods the values of the K-S statistics are lower for the CIR base regime dynamics (with the exception of the Pareto spike model in the second period). We have to note also, that the deseasonalization method does not do a perfect job in the second, more volatile period. Perhaps the fits in this period could be improved by applying a different (more restrictive) deseasonalization method.

V. Conclusions

Our empirical study provides evidence for two facts, which have important consequences in risk management applications. Firstly, models with shifted spike regime distributions lead to more realistic descriptions of electricity spot prices. Secondly, by introducing CIR-type heteroskedasticity in the base regime – in place of the standard mean-reverting, constant volatility dynamics – we obtain better spike identification and (generally) better goodness-of-fit.

Acknowledgments

The authors would like to thank Tomasz Piesiewicz of EnergiaPro Gigawat for providing the EEX price data.

References

[1] M. Bierbrauer, C. Menn, S.T. Rachev, S. Trück (2007) Spot and derivative pricing in the EEX power market. *J. Banking & Finance* 31: 3462-3485.
[2] M. Bierbrauer, S. Trück, R. Weron (2004) Modeling electricity prices with regime switching models. *Lect. Notes in Computer Sci.* 3039: 859-867.
[3] D.W. Bunn, ed. (2004) Modelling Prices in Competitive Electricity Markets. Wiley, Chichester.
[4] J.C. Cox, J.E. Ingersoll, S.A. Ross (1985) A theory of the term structure of interest rates. *Econometrica* 53: 385-407.
[5] C. de Jong (2006) The nature of power spikes: A regime-switch approach. *Stud. Nonlinear Dynamics & Econometrics* 10(3), Article 3.
[6] A. Dempster, N. Laird, D.B. Rubin (1977) Maximum likelihood from incomplete data via the EM algorithm. *J. Royal Statistical Soc.* 39: 1-38.
[7] R. Ethier, T. Mount (1998) *Estimating the volatility of spot prices in restructured electricity markets and the implications for option values.* PSerc Working Paper 98-31.
[8] J. Hamilton (1990) Analysis of time series subject to changes in regime. *J. Econometrics* 45: 39-70.
[9] R. Huisman, C. de Jong (2003) Option pricing for power prices with spikes. *Energy Power Risk Management* 7.11: 12-16.
[10] R. Huisman, R. Mahieu (2003) Regime jumps in electricity prices. *Energy Economics* 25: 425-434.
[11] J. Janczura, R. Weron (2009) *An empirical comparison of alternate regime-switching models for electricity spot prices*. HSC Working Paper.
[12] C.-J. Kim (1994) Dynamic linear models with Markov-switching. *J. Econometrics* 60: 1-22.
[13] C.R. Knittel, M.R. Roberts (2005) An empirical examination of restructured electricity prices. *Energy Economics* 27(5): 791-817.
[14] A. Ockenfels, V. Grimm, G. Zoettl (2008) *Electricity market design.* EEX Report.
[15] I. Simonsen (2005) Volatility of power markets. *Physica A* 355: 10-20.
[16] S. Trück S, R. Weron, R. Wolff (2007) *Outlier treatment and robust approaches for modeling electricity spot prices*. Proceedings of the 56th Session of the ISI. Available at http://mpra.ub.uni-muenchen.de/4711/.
[17] O. Vasicek (1977) An equilibrium characterization of the term structure. *J. Financial Econometrics* 5: 177-188.
[18] R. Weron (2006) *Modeling and forecasting electricity loads and prices: A statistical approach*. Wiley, Chichester.
[19] R. Weron (2008) Heavy-tails and regime-switching in electricity prices, *Math. Meth. Oper. Res.*: Online First DOI 10.1007/s00186-008-0247-4.